

**No Problem:
Evidence that Problem Intuitions Are Not Widespread¹**
Justin Sytsma and Eyuphan Ozdemir

The meta-problem is “the problem of explaining why we think that there is a problem of consciousness” (6). This presupposes that *we* think there is a problem in the first place. We challenge the breadth of this “we” in this essay and a series of companion pieces that are underway.

It is well known that “consciousness” is used in a number of ways, most of which do not pose a relevant problem of consciousness. The problem at issue is the problem of explaining *phenomenal consciousness*. This phrase is generally used to characterize a diverse range of mental states that are thought to have a distinctive qualitative aspect, and they are typically indicated via examples. For instance, Chalmers writes that “it is widely accepted that seeing a bright red square and feeling pain are phenomenally conscious, and that one’s ordinary background beliefs... are not” (7). As this suggests, seeing red and feeling pain are prototypical examples of (supposed) phenomenal states: they are thought to have something important in common—that there is “something it is like” to be in each state.

Do *we* think that there is a problem of phenomenal consciousness? To Chalmers credit, he recognizes that the question of the *distribution of Problem Intuitions*—the general prevalence of intuitions that there is a problem of explaining phenomenal consciousness—is an empirical question. Nonetheless, he thinks that “the central intuitions are widely shared well beyond philosophy” (13) and works “under the assumption that these intuitions are widely shared, or at

¹ For or a special issue of *Journal of Consciousness Studies* on David Chalmers’s “The Meta-Problem of Consciousness” edited by François Kammerer. We want to thank the audience at the third Australasian Experimental Philosophy Conference for helpful feedback.

least that they have a widely shared basis” (15). The intent of the concluding hedge seems to be to note that while Problem Intuitions are rejected by some people (such as the present authors), “in many cases of rejection, there is an underlying intuition that is psychologically outweighed by other forces” (14).

Our focus in this paper will be on whether Problem Intuitions are widely shared among lay people. In Section 1, we argue that there is *already* sufficient empirical evidence to cast doubt on this. We then add to this body of evidence in Section 2. More specifically, we contend that lay people by and large lack the concept of phenomenal consciousness. And if this is correct, then it undercuts not only the assumption that problem intuitions are widespread, but the assumption that they have a widely shared basis. We return to this point in Section 3, where we consider whether solving the meta-problem remains important if Problem Intuitions are not widespread.

1. Existing Evidence

While Chalmers recognizes that the distribution of Problem Intuitions is an empirical question, he believes that they are widespread. This requires discounting present evidence to the contrary. Thus, Chalmers contends that the literature on folk mental state attributions is largely irrelevant to the question and asserts that the work that does bear on it is not compelling. He writes (13):

the largest body of research concerns the distribution of consciousness... for example, do people think that machines or corporations can feel pain? Some attempts have been made to connect this research to the hard problem of consciousness, but for the most part, the intuitions in question have not been the core problem intuitions.

In fact, this literature is not primarily concerned with the distribution of consciousness. One major concern is whether lay people have something like the philosophical concept of

phenomenal consciousness, and machines and corporations have been used as *comparison cases* in addressing this question.

Whether or not lay people have something like the philosophical concept is directly relevant to assessing the distribution of Problem Intuitions: if most people don't have the concept in the first place, then it would seem to follow that they don't have intuitions that phenomenal consciousness poses a problem. Sytsma and Machery (2010; SM) argue that this is the case. They note that those who assert there is a hard problem generally begin by asserting that the existence of phenomenal consciousness is pretheoretically obvious—that it is “the most central and manifest aspect of our mental lives” (Chalmers 1995, 207). They then argue that if this is the case, the distinction between phenomenal states and non-phenomenal states should be evident in how lay people categorize mental states.

SM then provide evidence that lay people do not tend to treat prototypical examples of (supposed) phenomenal states similarly. In their first study they present evidence that lay people happily attribute one (supposed) phenomenal state to a simple, non-humanoid robot (*seeing red*) while denying another (*feeling pain*).² In contrast, philosophers tend to deny both states of the robot. SM argue that this casts doubt on the standard justification offered for the reality of the hard problem. Put in terms of the present target article, they challenge the claim that Problem Intuitions are widespread, and they leverage that conclusion against the hard problem.

Chalmers notes the evidence provided by SM, but then quickly dismisses it: “this result is predicted by Chalmers (1996, p. 18), which observes that ordinary mental terms like this have both a functional reading and a phenomenal reading, with sensational terms such as ‘pain’ more likely to suggest a phenomenal reading than perceptual terms such as ‘see’” (fn7). We find it

² These findings have been replicated a number of times (Sytsma and Machery 2012, Sytsma 2012, Cova et al. 2019).

unclear why Chalmers contrasts “see” with “pain” rather than “feel.” Nonetheless, taking this objection at face value, it should be noted that there is a large body of evidence indicating that lay people do *not* typically treat “pain” as a sensational term (Sytsma 2010a; Reuter 2011; Sytsma 2012; Reuter, Phillips, and Sytsma 2014; Kim et al. 2016; Sytsma and Reuter 2017; Reuter and Sytsma 2018; Reuter, Sienhold, and Sytsma 2018). Instead, the evidence suggests that people distinguish *pains* from *feelings of pain*, treating the former as objects of sensations, not as sensations themselves.

More charitably, Chalmers might be read as objecting that lay people are more likely to read “see” in a functional sense in SM’s study while philosophers are more likely to read it in a phenomenal sense. This objection was not only addressed in SM’s original paper, but in a good bit of follow-up work (Sytsma 2009, 2010a, 2012, 2013, 2016, ms). Rather than rehash this work here, however, we’ll focus on extending it in the next section.

Chalmers also notes that he thinks that “many or most subjects have concepts of specific phenomenal states such as feeling pain or experiencing colour,” but that he is “neutral on whether they also have a unifying concept of phenomenal consciousness” (fn7). SM’s conclusion need not be phrased in terms of a unifying concept, however. Rather the point is that people don’t treat the two *prototypical* examples of (supposed) phenomenal states that Chalmers notes similarly, suggesting that they don’t recognize the phenomenality of these states. And if this is correct, then it undermines the claim that they treat these states as being phenomenal.

Before turning to our new study, it is worth noting that more recent work in experimental philosophy and psychology also supports the contention that problem intuitions are not widespread. Although Peressini (2014) takes issue with the results of SM (see Sytsma ms for a response), on the basis of a new study he too contends that lay views do not support there being a

hard problem. He writes that the hard problem depends on “an understanding of subjective experience that includes a metaphysically robust qualitative component, one that is not present in the folk conception” (886). And more recently, Weisman, Dweck, and Markman (2017) have found that perceptual states and bodily sensations tend to be treated as falling into different clusters across several large-scale studies of US adults, which again suggests that people do not tend to see these states as having something important in common.

Finally, we want to note a further set of studies reported by Gottlieb and Lombrozo (2018). Chalmers briefly discusses this article, writing that they “find that people judge that phenomena tied to subjective experience and to privileged access are relatively hard to explain” (14). This gives the erroneous impression that this work indicates that people find these phenomena to be hard to explain. But, if anything, their work suggests the exact opposite! Gottlieb and Lombrozo look at judgments about whether “science could one day fully explain” a range of phenomena, finding among other things that these explanatory judgments were inversely correlated with judgments that the phenomenon had an “introspectively accessible phenomenology” (122). In their first study, this was measured using a question that involved language philosophers would recognize as signaling phenomenal consciousness: “[the phenomenon] involves a subjective experience (a feeling of what it is like) that only the individual experiencing it can know.” This was further broken down in a second study, which included a question about “subjective experience”: “[the phenomenon] has a subjective experience associated with it—a ‘feeling’ of what it is like.”

Although we doubt that lay people tend to read this language as indicating phenomenal consciousness, if they do then Gottlieb and Lombrozo’s studies provide evidence that lay people generally do *not* harbor Problem Intuitions. In their first study, while 32 out of 46 items had

mean ratings above the neutral point for “introspectively accessible phenomenology,” only two of those 32 items had mean ratings below the neutral point for explanatory judgments—“believing in God” and “being able to feel and identify different types of textures with the sense of touch”—and both of these ratings are readily explained on skeptical grounds. In their second study, 39 out of 46 items had mean ratings above the neutral point for “subjective experience,” but again only two of those 39 items had mean ratings below the neutral point for explanation judgments (the same two as the first study).

2. New Evidence

There is already reason to doubt the assumption that Problem Intuitions are widespread. That said, we are delighted to read that Chalmers would be “delighted to see empirical research” that tests this assumption (15). And we aim to delight. We begin, here, by expanding on the work of SM.

As noted, the most prominent line of response to SM is to attempt to explain away the asymmetry they found. This is typically done by arguing that while lay people recognize that seeing red is phenomenal, the phrase “seeing red” in SM’s question didn’t trigger this reading, but instead triggered a purely informational reading akin to “detecting red.” As such, people are happy to attribute this state to a simple robot. In contrast, it is claimed that the phrase “feeling pain” is read in a phenomenal way, explaining why lay people are generally disinclined to attribute this state to the robot.

While there are already a number of responses to this objection in the literature, in the present study we take another tack, exchanging the supposedly problematic term “see” for one

generally assumed to mark phenomenal consciousness—“experience.” This manipulation was suggested by Tony Jack, as noted by SM. They write:

A critic could argue that our hypothesis predicts that if we ask ordinary people whether Jimmy *experiences red*, they would answer affirmatively. If they would answer negatively (as our critic suggests they would), then we would have to conclude that the folk conceive of subjective experience as philosophers do. (fn11)

SM then raise doubts about whether “experience” is actually used in this way by lay people (see also Sytsma 2010b). Nonetheless, if lay people *do* have the concept of phenomenal consciousness, then this seems like the most likely term in English to express it. As such, if we were to replicate SM’s results using this terminology, it would be compelling evidence for SM’s central thesis.

In addition, we had a second goal for this study—to extend the results from SM beyond American participants. Toward this, the study was carried out online via advertising on Google. Ads were displayed globally, in English, attracting a diverse pool of 1,465 participants hailing from 124 different countries. Extended details about materials, demographics, and statistical analysis can be found in the supplemental materials.

We used the same vignettes as in Study 1 in SM. That study used a 2x2 between-participants design that varied the *agent* and the *state* described: each participant read a short story about either a normal undergraduate (Timmy) or a simple robot (Jimmy) participating in an experiment in which the agent gave behavioral cues related to either seeing red or feeling pain. We extended this to a 2x2x2 between-participant design, also varying the *term* used to describe the state: participants were either given the original question from SM (“Did [agent] see red?” or “Did [agent] feel pain when he was shocked?”) or a revised question in which “see” / “feel” was replaced with “experience.” Participants answered using a 7-point scale anchored at 1 with “Clearly No,” at 4 with “Not Sure,” and at 7 with “Clearly Yes.”

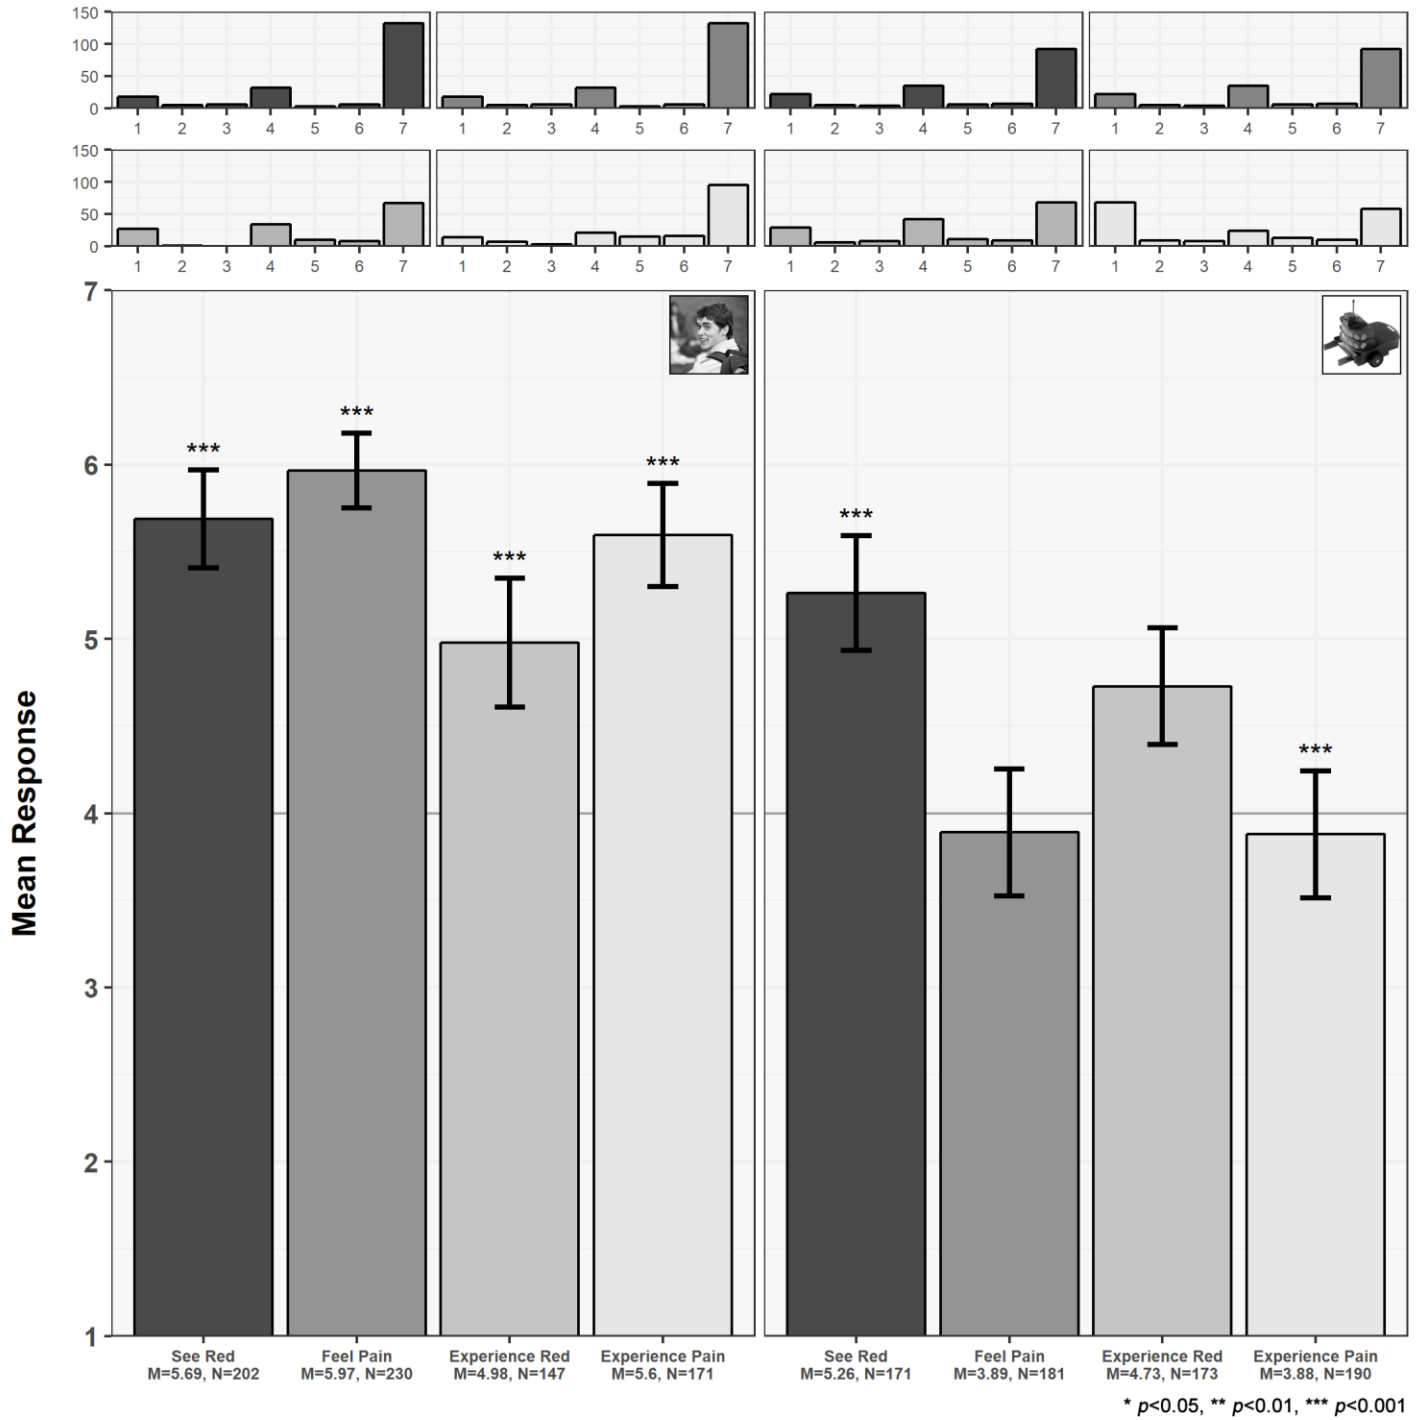


Figure 1: Results with 95% CIs; histograms above respective graphs.

The results are shown in Figure 1. Focusing first on the replication of SM, planned analyses showed that the mean responses for Timmy for both seeing red and feeling pain were

significantly above the neutral point. In contrast, while the mean response for Jimmy seeing red was significantly above the neutral point, the mean response for Jimmy feeling pain was not significantly different than the neutral point. This stands in contrast to the results from SM, where the mean response for Jimmy feeling pain was significantly below the neutral point ($M=2.54$ in the original study compared to $M=3.89$ in the present study). We return to this difference below.

Despite this difference in responses for Jimmy feeling pain, the key pattern of results from SM's study is replicated in our study. First, while the difference between the mean ratings for the two states for Timmy was negligible (0.28), the difference for Jimmy was notably larger (1.37). Second, while the difference between the mean ratings for Jimmy and Timmy for seeing red was small (0.42 points), the difference for feeling pain was large (2.08 points). In other words, as in SM, participants treated the two states quite differently for the robot, but not for the human. And, again, setting concerns about the interpretation of "seeing red" aside, this suggests that lay people do not categorize these mental states in the way we would expect if they recognize both states as being phenomenal.

Turning to the conditions where we replaced "see" / "feel" with "experience," planned analyses showed the same basic pattern: mean responses for Timmy for both experiencing red and experiencing pain were significantly above the neutral point. In contrast, while the mean response for Jimmy experiencing red was significantly above the neutral point, the mean response for Jimmy experiencing pain was not significantly different than the neutral point. Again, we find that the difference between the mean ratings for the two states for Jimmy (0.85) was greater than for Timmy (0.62). Interestingly, we find that the mean rating for each agent for "experience red" is significantly lower than the mean rating for "see red." If we found this result

for just Jimmy, we might take it to indicate that “experience” was in fact triggering a phenomenal reading in *some* participants where “see” did not. But we also find this result—and with a larger effect size—for Timmy, who participants should presume is phenomenally conscious if they have that concept. So, there is no reason to interpret the dip in means for Jimmy as providing evidence of a phenomenal reading.

Most importantly, we find that even after switching terms, the difference between the mean ratings for Jimmy and Timmy for “experience red” was negligible (0.25), while the difference for “experience pain” was much larger (1.72). That is, we yet again found that participants treated the two states quite differently. If lay people tend to have something like the philosophical concept of phenomenal consciousness and that concept is indicated by the term “experience,” then we would have expected a very different pattern of results. As such, the results suggest either that lay people don’t have the concept or that the word “experience” doesn’t suggest the concept (or both). But, if neither “see” nor “experience” suggest a phenomenal reading, then what term would? If phenomenal consciousness is indeed as central and manifest as advocates assert, we would not expect it to be so difficult to draw out evidence of this distinction from lay participants. That participants continue to treat two prototypical examples of (supposed) phenomenal states dissimilarly, thus provides reason to doubt that they have something like the philosophical concept in the first place. And absent that, there is no reason to think that Problem Intuitions are widespread.

Given the diversity of the sample, it is worth checking for demographic effects, including native language. Exploratory multiple regressions suggest significant demographic effects for gender and location of birth, a borderline significant effect for religiosity, and some interaction

effects including between gender and location of birth. We don't find notable effects for native language, education, or training in philosophy, psychology, or the sciences.

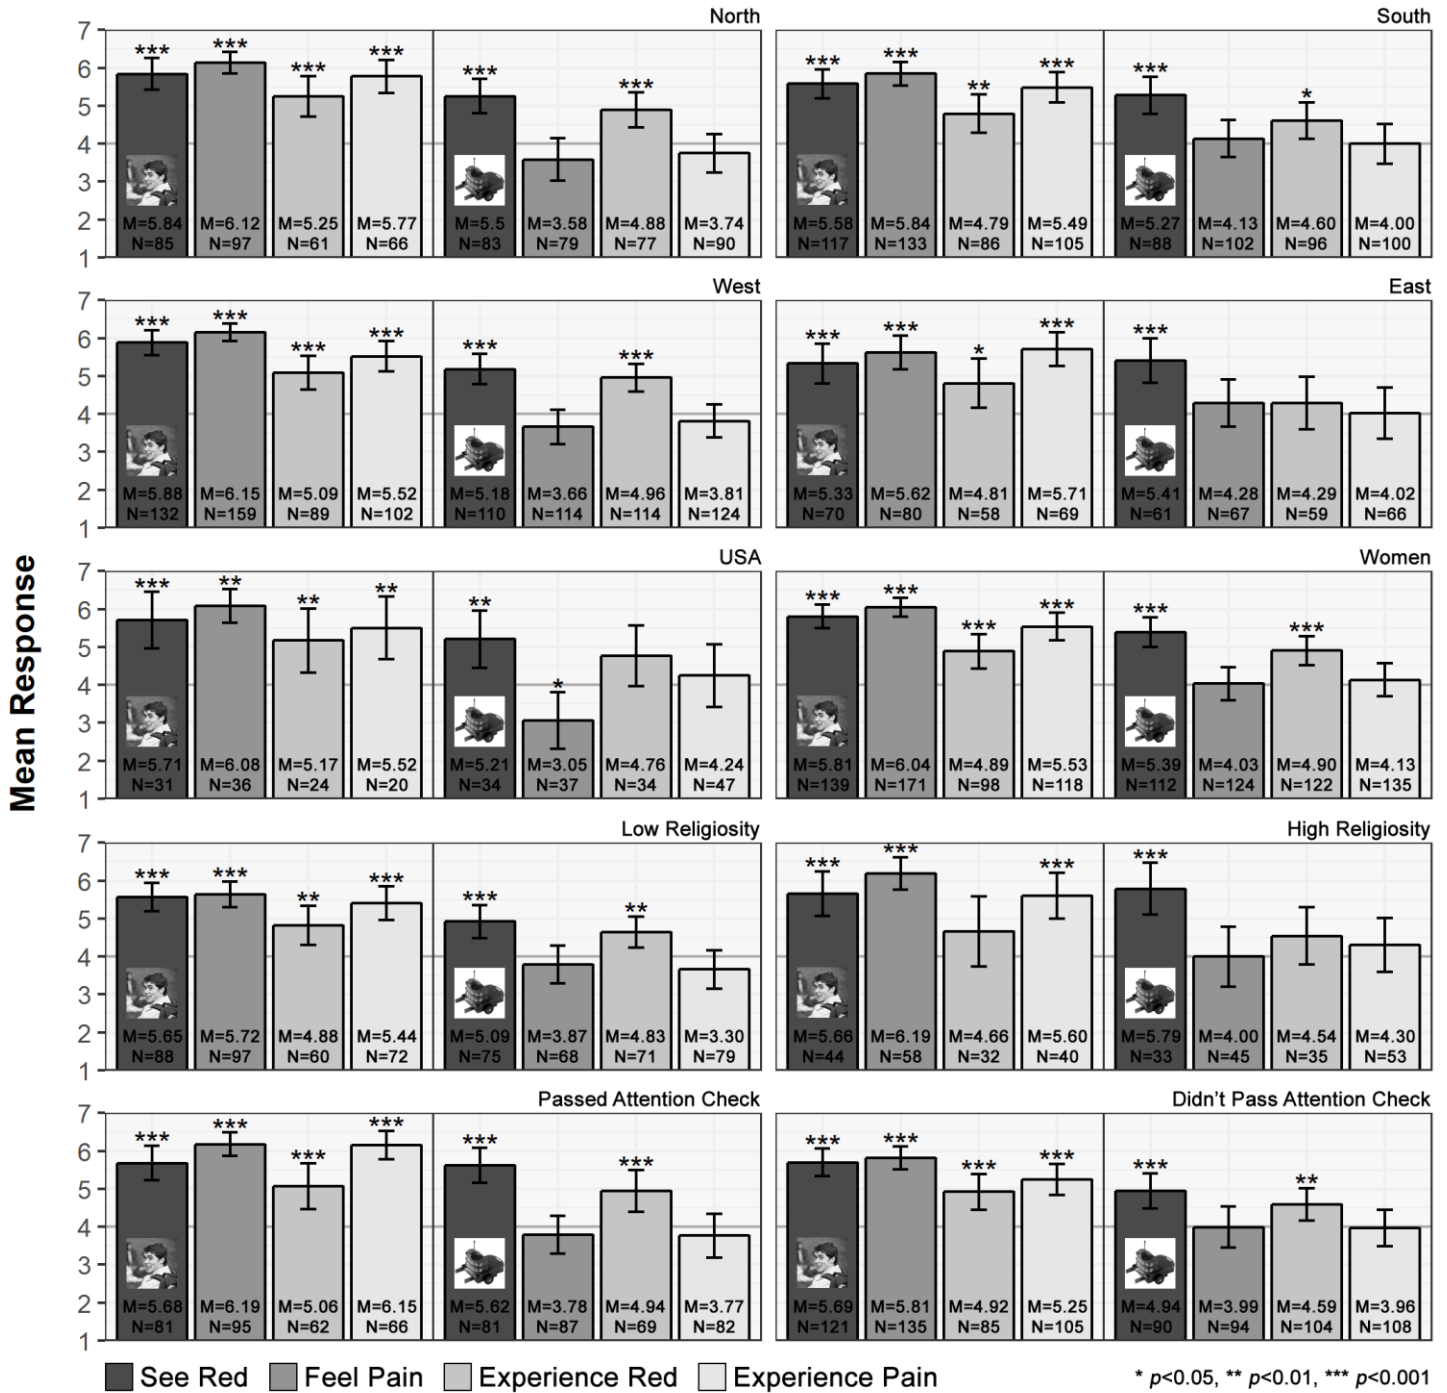


Figure 1: Results by demographics with 95% CIs; low religiosity included scores of 1 or 2 on 5-point scale, high religiosity scores of 4 or 5.

Looking more closely at location, the effect is found while breaking down by region, as well as when splitting the sample based on Global North/South or East/West, with the latter showing the strongest effects.³ We also see an effect for those born in the USA versus the rest of the world. While a number of things could be said here, the key thing to note is that none of the groupings show a general skepticism toward ascribing mental states to the robot, including that none tend to deny that the robot experiences red; in fact, with the exception of participants from the East and non-women, the mean rating for Jimmy experiences red is significantly or borderline significantly above the neutral point in each case. Splits are shown in Figure 2.

One worry is that participants might not have been suitably attentive to the task, perhaps reflecting the recruitment strategy we used. As detailed in the supplemental materials, however, a stringent attention check was included and responses did not vary importantly with attention.

Finally, we noted above that the mean rating for Jimmy feeling pain was notably higher than that reported by SM. The demographic breakdowns offer a potential explanation of this difference. Looking at the ratings for participants from the USA ($N=37$, $M=3.05$, $SD=2.25$), the mean rating is much closer to that reported by SM ($M=2.54$).

3. Is the meta-problem still important?

Chalmers writes that “variable factors (e.g. cultural, linguistic, and theoretical background, and other factors that vary with historical period and individual psychology)... may sometimes overwhelm the contribution of more universal factors,” such as “mental states and introspective mechanisms that most humans share” (15). Nonetheless, his “working assumption is that there are also near-universal factors that play a central underlying role in explaining problem intuitions

³ See supplemental materials for details.

where they are present” (15). We hold that this assumption is false: if the occurrence of phenomenal mental states and mechanisms for introspecting them is near-universal, then we would expect people to generally treat prototypical examples of (supposed) phenomenal states similarly, but they don’t. While it might be that the claimed near-universal factors are typically overwhelmed in lay people today, we see no reason to expect that to be the case.

Nonetheless, Chalmers holds that even if his assumption is false, “solving the meta-problem will remain an important project” (15). We are not so sure. We believe that this depends in part on just how restricted of a minority shows Problem Intuitions. At the extreme, if Problem Intuitions are largely restricted to philosophers, related academics, and those who have studied their work, then we contend that the importance of solving the meta-problem would be largely dissolved.

Chalmers claims that “it will still be crucial” for those who hold that Problem Intuitions are illusory to explain them “in order to make the case that they are illusory” (15). But if Problem Intuitions are not widespread, then the burden of proof would seem to go the other way around: it would be crucial for the minority to make the case that their intuitions are *not* illusory. We believe this is a tall order. Accepting that Problem Intuitions are not widespread is in tension with two principles that Chalmers seems to embrace: (1) that phenomenally conscious mental states and mechanisms for introspecting them are near-universal; and, (2) that Problem Intuitions reliably occur in the presence of these factors. Accepting that Problem Intuitions are not widespread, to justify that those Problem Intuitions that occur should nonetheless be taken seriously, we believe that one either needs to give up on (1) or (2).

Abandoning (1) would amount to the claim that all but a select group are philosophical zombies! This seems far-fetched and rhetorically weak. Abandoning (2) is perhaps more

promising. Accepting (1) and that Problem Intuitions are not widespread, it amounts to the claim that Problem Intuitions don't typically arise in the presence of phenomenal consciousness. Perhaps one could argue that this is because recognizing that one is phenomenally conscious requires not just the right sort of mechanisms for introspection, but the right sort of introspective training. This would amount to abandoning the claim that phenomenal consciousness is "central and manifest," however, which is the standard justification for believing in the phenomenon in the first place. Perhaps there are other options for rejecting (2). But any such rejection would cast doubt on the connection between being phenomenally conscious and having Problem Intuitions, which would feed into the debunking argument and the coincidence argument that Chalmers is at pains to dismiss.

References

- Chalmers, David (1995). "Facing up to the problem of consciousness." *Journal of Consciousness Studies*, 2(3): 200–219.
- Cova, Florian et al. (2019). "Estimating the Reproducibility of Experimental Philosophy." *Review of Philosophy and Psychology*.
- Gottlieb, Sara and Tania Lombrozo (2018). "Can Science Explain the Human Mind? Intuitive Judgments About the Limits of Science." *Psychological Science*, 29(1): 121–130.
- 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.
- Peressini, Anthony (2014). "Blurring two conceptions of subjective experience: Folk versus philosophical phenomenality." *Philosophical Psychology*, 27(6): 862–889.
- 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 (2018). "Putting Pain in its Proper Place." *Analysis*.

- Reuter, Kevin and Justin Sytsma (2018). “Unfelt Pain.” *Synthese*.
- Sytsma, Justin (2009). “Phenomenological Obviousness and the New Science of Consciousness.” *Philosophy of Science*, 76(5): 958–969.
- Sytsma, Justin (2010a). “Dennett’s Theory of the Folk Theory of Consciousness.” *Journal of Consciousness Studies*, 17(3-4): 107–130.
- Sytsma, Justin (2010b). *Phenomenal consciousness as scientific phenomenon? A critical investigation of the new science of consciousness*. PhD Dissertation, University of Pittsburgh.
- Sytsma, Justin (2012). “Revisiting the Valence Account.” *Philosophical Topics*, 40(2): 179–198.
- 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 (2016). “Attributions of Consciousness.” In J. Sytsma and W. Buckwalter (eds), *A Companion to Experimental Philosophy*, Blackwell.
- Sytsma, Justin (ms). “Experiencers and the Ambiguity Objection.”
- 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 Kevin Reuter (2017). “Experimental Philosophy of Pain.” *Journal of Indian Council of Philosophical Research*, 34(3): 611–628.
- Weisman, Kara, Carol Dweck, and Ellen Markman (2017). “Rethinking people's conceptions of mental life.” *PNAS*: 114(43): 11374-11379.