Commentary on Buckwalter and Phelan

Justin Sytsma

In their thought-provoking paper, Wesley Buckwalter and Mark Phelan (B&P) offer an interesting new objection to Sytsma and Machery (2010). To set-up that objection, and my response to it, it is worth briefly outlining our 2010 paper. We began that paper by presenting evidence that a certain bit of philosophical commonsense is mistaken: The existence of phenomenally conscious mental states is said to be obvious on the basis of first-person experience with states like seeing red and feeling pain. We argued that if this is correct, then ordinary people (the folk, or people without training in philosophy or consciousness studies) should treat these states similarly, denying, for example, that a simple non-humanoid robot (Jimmy) can be in either state. It turns out that they do not treat these states similarly. Like philosophers, they tend to deny that Jimmy can feel pain; but, unlike philosophers, they tend to accept that Jimmy can see red. Based on this finding, we argued that the philosophical tradition is mistaken in treating the existence of phenomenally conscious mental states as being obvious from first-person experience.

Having presented evidence for our critique of the philosophical tradition, we went on to ask a follow-up question: Why do the folk ascribe seeing red to the robot, but not feeling pain? That the folk treat these states dissimilarly suggests that they are, perhaps implicitly, calling on some distinction between types of mental states. We suggested two possibilities. First, the folk might treat seeing red differently from feeling pain because the former is a product of an external sense, while the latter is a product of an internal sense. Second, the folk might treat seeing red differently from feeling pain because the latter is associated with a valence (or a hedonic value), while the former is not.

To test these hypotheses, we ran two further studies looking at ascriptions of anger and smell to the simple robot used in our first study. We found that while the folk tend to deny that Jimmy felt anger, they were split with regard to whether Jimmy smelled banana or smelled vomit (the mean response for each of these questions was close to the neutral point). These findings are arguably compatible with both hypotheses. Thus, the internal/external hypothesis can explain the neutral results for smelling banana and smelling vomit in terms of olfaction being borderline between an external and an internal sense (see, for example, Newton 2000); and, the valence hypothesis can explain the neutral results in terms of the valence associated with olfactory states being less essential than the valence associated with pain or anger. A final result suggested in favor of the valence hypothesis, however: When participants were asked whether the robot smelled Isoamyl Acetate—a descriptor they were unlikely to be familiar with, and thus unlikely to associate a valence with—they tended to answer that the robot did. While this finding is in keeping with the valence hypothesis, it is not readily explained in terms of the internal/external hypothesis. We tentatively concluded that the folk distinguish between types of mental states on the basis of valence.

Based on the results of two new studies, B&P challenge the valence hypothesis: They argue that people do not typically differentiate between valenced and non-valenced states, showing instead that the function that the robot was designed for can have a significant impact on folk ascriptions. This leads B&P to claim that their results undercut our criticism of the philosophical tradition. Clearly distinguishing between our critique of the philosophical tradition and our subsequent defense of the valence hypothesis, however, it is unclear that B&P's claim is accurate, as I will argue in Section 1 below. Further, I believe that there is a serious problem with the studies that B&P ran, as explained in Section 2: In assigning different functions to the robot

in their studies, they do not adequately control for corresponding assumptions about its complexity with regard to tasks related to those functions. Thus, it might be that people assume that the robot is more complex with regard to tasks involved in carrying out the function assigned to it. Since we expect complexity to correlate with mental state ascriptions, this could explain B&P's results. In Section 3, I report on a new experiment testing our criticism of the philosophical tradition. I find that an asymmetry remains between ascriptions of seeing red and feeling pain even after a task-neutral function has been specified for the robot. Finally, in Section 4, I revisit the valence hypothesis. While I am not convinced by B&P's critique, I present new evidence suggesting that the valence hypothesis may in fact be incorrect and, therefore, that a new explanation of the asymmetries that we have found might be needed.

1. The Negative Hypothesis and the Positive Hypothesis

My first concern with Buckwalter and Phelan's paper is that it does not adequately characterize the dialectical structure of Sytsma and Machery (2010). That paper has a two part structure, as noted above: First, we presented evidence for a negative hypothesis (that the philosophical tradition is mistaken in taking the existence of phenomenally conscious mental states to simply be obvious from first-person experience); second, we presented evidence for a positive hypothesis (that ordinary people distinguish between mental states that have or lack valence). While B&P present evidence that runs counter to our positive hypothesis, they argue that this undercuts the case we made for the negative hypothesis. They write:

> According to Sytsma and Machery (2010) [hereafter S&M], people ordinarily distinguish subjective experiences that have a valence (or, "a hedonic value for the subject"), from those that lack a valence. S&M in turn contend that this valenced conception of subjective experience raises problems for a philosophical tradition which emphasizes the manifest phenomenal character of all subjective experiences. We argue from our own experimental data that, contra S&M, people

> do not ordinarily differentiate experiences according to whether they are valenced. This undercuts S&M's criticism of philosophical tradition.... (2)

This is an inadequate articulation of what we argued, however, and the evidence provided does not pose a direct problem for our negative hypothesis.

Our argument against the philosophical tradition does not rest primarily on the positive account that we put forward (although it does gain some support from it), but rather on the prior experiment indicating that people do not ordinarily group mental states together in the way that the tradition predicts: Non-philosophers treat two prototypical examples of phenomenally conscious mental states differently, ascribing the state of seeing red to the simple non-humanoid robot Jimmy, but not the state of feeling pain. Having found that people treat these states dissimilarly, we *then* sought to explain the finding, arriving at the valence hypothesis. This is important because it means that evidence against the valence hypothesis does not directly undercut our criticism of the philosophical tradition, as B&P claim it does.

To see this, it is worth briefly describing the evidence presented by B&P. They ran two studies using the same simple robot that we used in our 2010 paper, but assigning different functions to it. For example, in their first study Jimmy was described as either being created in order to "clean bio-medical waste" or to "make smoothies." In each case, the robot was said to have correctly performed tasks requiring that it distinguish between different olfactory stimuli. Participants were then asked whether Jimmy smelled one of those stimuli. B&P found that the function assigned to Jimmy affected people's responses. For example, when Jimmy was described as being created to clean bio-medical waste, participants were more likely to say that the robot smelled vomit than that it smelled banana. The reverse pattern was found when Jimmy was described as being created to make smoothies (as shown in Figure 1 below). Taken at face value, such results raise doubts about our positive hypothesis. Thus, B&P write that "it seems as

though a key piece of evidence for the S&M valence hypothesis is actually better explained in terms of differing beliefs about the subject of the experience, in this case assumptions about the robot's functional specifications, rather than assumptions about the nature of the experience itself" (6). This challenge to our positive hypothesis, however, does not directly raise concerns about our negative hypothesis: Again, the primary evidence for that hypothesis came from a prior study, not the valence hypothesis.

Nonetheless, it might be contended that B&P's evidence can be used to fuel an argument against our criticism of the philosophical tradition. To see this note both that the probes in our first study used the same simple robot as our subsequent studies looking at the valence hypothesis and that we did not specify a function for Jimmy in those probes. As such, B&P's finding that participants' responses to such probes are sensitive to the function associated with the robot, coupled with the assumption that in the absence of a specified function people will tend to assume one based on the details given in the probe, provides reason for skepticism: It might be that people tend to assume different functions for the robot Jimmy across the probes we used, which could potentially explain the asymmetry we found between ascriptions of seeing red and ascriptions of feeling pain.

I think that this is to move too quickly, however. First, while B&P's studies suggest that specifying a function for the robot Jimmy can impact mental state ascriptions, they do not establish that people assume different functions in different scenarios when none is provided. While this might be the case, I see little reason to think that this is a likely explanation of the results of our first study. Second, and perhaps more importantly, there is a confound in B&P's studies that casts doubt on the claim that the function assigned to Jimmy directly impacts mental state ascriptions. Specifically, function relates to complexity, and we expect complexity to

impact mental state ascriptions. Unfortunately, as I argue in the next section, B&P did not adequately control for the task-specific complexity of the robot Jimmy in their studies.

2. Function and Complexity

B&P attempted to control for the issue of complexity noted above by running "high complexity" versions of their two studies.[1] And they found that "function seems to be an important characteristic that people consider when ascribing phenomenal states to Jimmy, quite independently of the level of complexity of the robot" (5). That is, they found that the function assigned to the robot Jimmy makes a difference whether or not the robot was describedas being relatively complicated or relatively simple. Unfortunately, this does not serve as an adequate control for the issue I am raising.

To see this, we need to distinguish between the overall complexity of a machine and its complexity with regard to certain tasks. This is important because the task-specific complexity of a machine can vary independently of its overall complexity. Thus, we can imagine a machine that is highly complex overall, but that is relatively simple with regard to a given task. Likewise, we can imagine a machine that is rather simple overall, but that is relatively complex with regard to a given task. Imagine, for example, that I want to know which of two cars will be better at playing music. One car is a state-of-the-art Formula One race car, the other is an entry level passenger car. While I would judge that the race car is more complex overall, I would expect the

[1] For example, in the high complexity version of B&P's first study, Jimmy was described as follows: "Jimmy is a relatively complicated robot built at an Ivy League university. He has a state of the art scent detector (built with the new AccuScent technology), a sophisticated camera for eyes, engineering-grade thermoplastic wheels for moving about, and two reticulated grasping arms with touch sensors that he can move objects with." In contrast, in the low complexity version Jimmy was described as follows: "Jimmy is a relatively simple robot built at a small state university. He has a scent detector, video camera for eyes, wheels for moving about, and two grasping arms with touch sensors that he can move objects with."

passenger car to be more complex with regard to playing music. I would expect this because I know something about the considerations that go into designing such cars.

My concern with B&P's studies is similar: It seems that the functions that they assigned to Jimmy could raise assessments of its complexity with regard to different tasks, and could do so quite independently of the overall complexity of the robot. Assessments of the task-specific complexity of the robot Jimmy is an important issue because each of the functions that B&P used seem to relate most directly to just one of the mental states that they were testing. For example, in their first study, the function of making smoothies suggests that the robot would likely have been designed to detect banana, while the function of cleaning up bio-medical waste suggests that the robot would likely have been designed to detect vomit. By specifying each of those functions for the robot, we can reasonably expect that people will assume that Jimmy is more complex with regard to the related perceptual tasks than they would when the robot is described in the other way. Thus, I would expect people to think that the version of Jimmy that makes smoothies is *more complex* with regard to smelling banana than the version that cleans up bio-medical waste, while the version of Jimmy that cleans up bio-medical waste is *more complex* with regard to smelling vomit than the version that makes smoothies. As such, if people's willingness to ascribe mental states to a robot correlates with the task-specific complexity of the robot with regard to tasks associated with those mental states, then we would predict that people will be more likely to ascribe the mental state of smelling banana to the version of Jimmy that makes smoothies than to the version that cleans up bio-medical waste, and we would predict that they would be more likely to ascribe the mental state of smelling vomit to the version of Jimmy that cleans up bio-medical waste than to the version that makes smoothies. And that is exactly what B&P found, as shown in Figure 1.

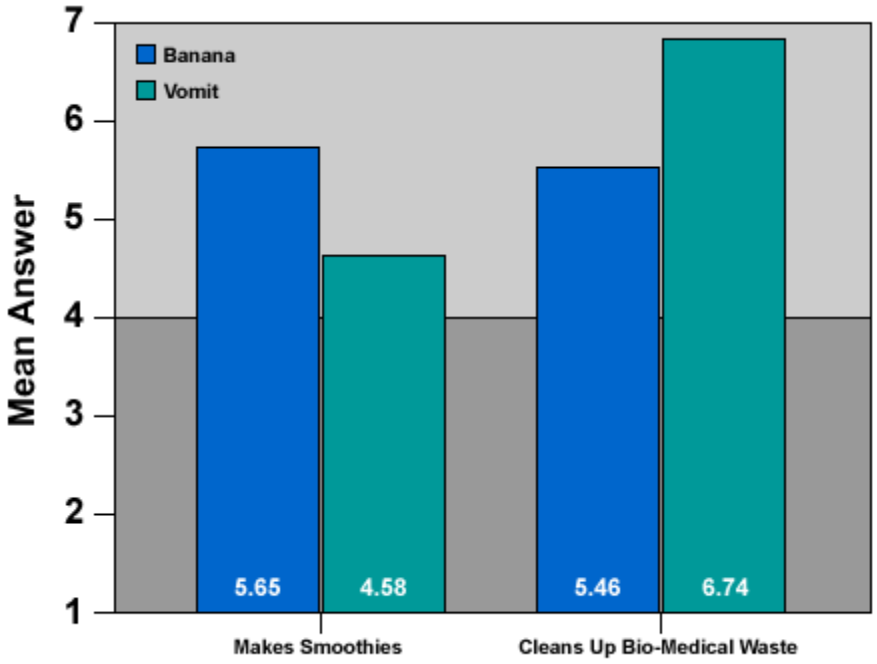

Figure 1: Results of B&P Study 1, Low Complexity, Smelling Banana vs. Vomit

Distinguishing between overall complexity and task-specific complexity, it seems that despite the high complexity control that B&P used, we can nonetheless explain their findings in terms of the perceived complexity of the robot. This indicates that a different type of control is needed. In fact, it could be argued that the original studies in Sytsma and Machery (2010) were superior to B&P's since we did not specify a function in any of the conditions, making it less likely that participants would take the robot to be more complex with regard to the perceptual tasks at issue in one condition than in another.

Against this response, B&P suggest that leaving the function of the robot unspecified is itself problematic:

> We want to suggest... that [the asymmetry in attributions found by S&M] is due to tacit assumptions on the part of experimental participants about the function for which the S&M robot was created. In S&M's vignettes, function is left unspecified. Thus

> participants are left to draw their own conclusions about the function of the robot, and this could play an important role in explaining the asymmetry between different experiential states. (4)

Fortunately, there seems to be a third strategy that can be adopted to control for both of the concerns that have been raised: We can specify a function for the robot that is not obviously specific to one type of mental state being tested or the other. For example, we could specify in the vignettes that "Jimmy was designed as a general-purpose robot to help with household chores" or that "Jimmy was designed to be a tool for the elderly by lifting and moving heavy objects around their houses" (as used in B&P's second study).

3. Re-testing the Negative Hypothesis

If the asymmetry we found in the first study reported in Sytsma and Machery (2010)—that people are significantly more likely to say that Jimmy saw red than that Jimmy felt pain—simply reflects that we did not specify a function for the robot, then we would expect that asymmetry to disappear when Jimmy is described as having either the function of helping with household chores or lifting and moving heavy objects. To test this prediction, I randomly presented 452 participants with one of the six probes given in the appendix.[2] These divide into three pairs of two probes based on the function assigned to the robot Jimmy—no function, general function, or lifting function. After reading one of these probes, the participant was asked to answer a question about it—either "Did Jimmy see blue?" or "Did Jimmy feel pain?"—on a 7-point scale anchored at 1 with "Clearly No," at 4 with "Not Sure," and at 7 with "Clearly Yes."

[2] The surveys were conducted through the Philosophical Personality website: http://www.philosophicalpersonality.com. Participants were restricted to native English speakers, 18 years of age or older, with at most minimal training in philosophy, and who had not previously taken a survey through the Philosophical Personality website. Participants were counted as having more than minimal training in philosophy if they were philosophy majors, had completed a degree with a major in philosophy, or had taken graduate-level courses in philosophy.

Against the prediction derived from B&P's hypothesis, I found a significant difference between people's willingness to say that Jimmy saw blue and their willingness to say that Jimmy felt pain in each pair of cases.[3] That is, whether no function was specified, Jimmy was described as a general purpose robot, or whether Jimmy was described as a lifting and moving robot, the asymmetry remained: Participants were significantly more likely to say that Jimmy saw red than that Jimmy felt pain, as seen in Figure 2.

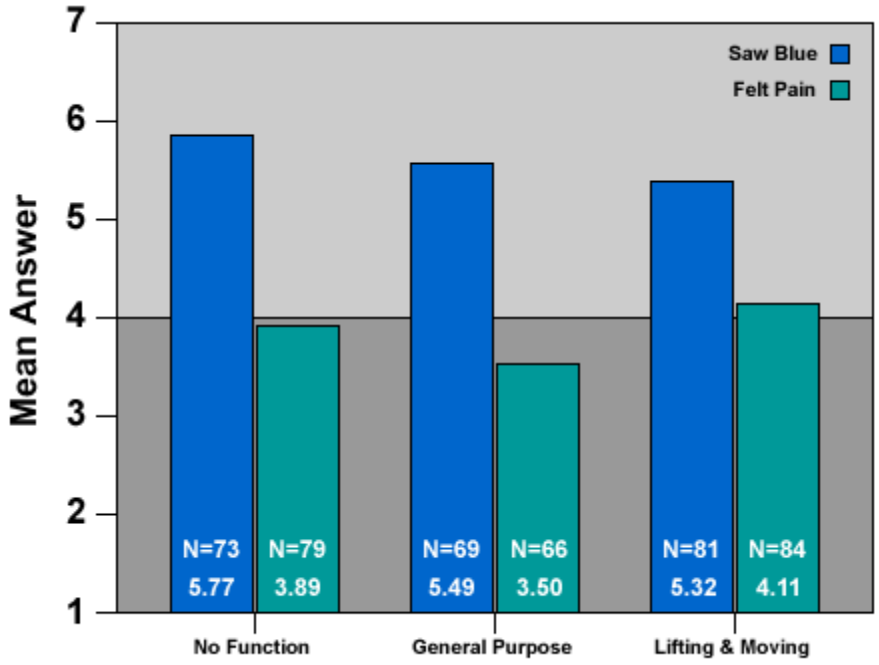

Figure2: Results of Study 1.

4. Re-testing the Positive Hypothesis

To quickly recap, B&P argue that the results of their two studies undercut the criticism of the philosophical tradition given in Sytsma and Machery (2010). In that paper, however, we put

[3] No Function: t(150)=5.2413, p=0.0000002661, one-tailed.
General Purpose: t(133)=4.9411, p=0.000001147, one-tailed.
Lifting & Moving: t(163)=3.2638, p=0.0006697, one-tailed.

forward both a negative hypothesis and a positive hypothesis. And while our criticism of the philosophical tradition is based on the evidence presented for the negative hypothesis, B&P target the evidence provided for the positive hypothesis. Despite this, B&P's results might be taken to fuel an argument against our negative hypothesis. One problem with this is that their studies did not adequately control for assumptions about the task-specific complexity of the robot Jimmy. Doing so, I found that the asymmetry that fueled our critique of the philosophical tradition remains. As such, we are still in need of an explanation for why the folk treat two prototypical examples of phenomenally conscious mental states—seeing red and feeling pain—dissimilarly with regard to ascribing them to a simple non-humanoid robot.

Recall that we offered two possible explanations in our paper, the internal/external hypothesis and the valence hypothesis, presenting evidence that favored the latter. Is the valence hypothesis still a viable explanation given the new results presented by B&P? If the above criticism of their studies is correct, then it would seem that is. Nonetheless, I believe that there is some reason to be suspicious of the valence hypothesis. Specifically, the results of three new studies suggest against that explanation.

To further test the valence hypothesis, Edouard Machery and I conducted a two-stage study. First we compiled a list of terms for 18 different smells based largely on the work of Kamath, Turetsky, and Moberg (2011) and Royet et al. (2000). These terms were divided into three groups of six—six that we expected people to associate with a positive valence (caramel, bubble gum, fresh bananas, raspberries, lavender flowers, fruit punch), six that we expected people to associate with a negative valence (natural gas, paint thinner, sweaty socks, skunk spray, rotting meat, vomit), and six that we expected people to treat as valence-neutral (soapy water, leather, pine wood, newspaper, dry dirt, ground pepper).

In the first stage of our study, 686 participants were randomly presented with four of the 18 terms, using the same website and with the same restrictions as in Study 1 above. For each term, participants were asked to imagine and reflect on that smell. They were then asked to rate, on 7-point scales, the familiarity of the smell (1=not at all familiar, 4= neither familiar nor unfamiliar, 7=very familiar), the identifiability of the smell (1=certain I could not identify it, 4=not sure, 7=certain I could identify it), whether the smell is subtle or intense (1=extremely subtle, 4=neutral, 7=extremely intense), and whether the smell is pleasant or unpleasant (1=highly pleasant, 4=neutral, 7=highly unpleasant). The mean ratings for valence for each term are shown in Figure 3.

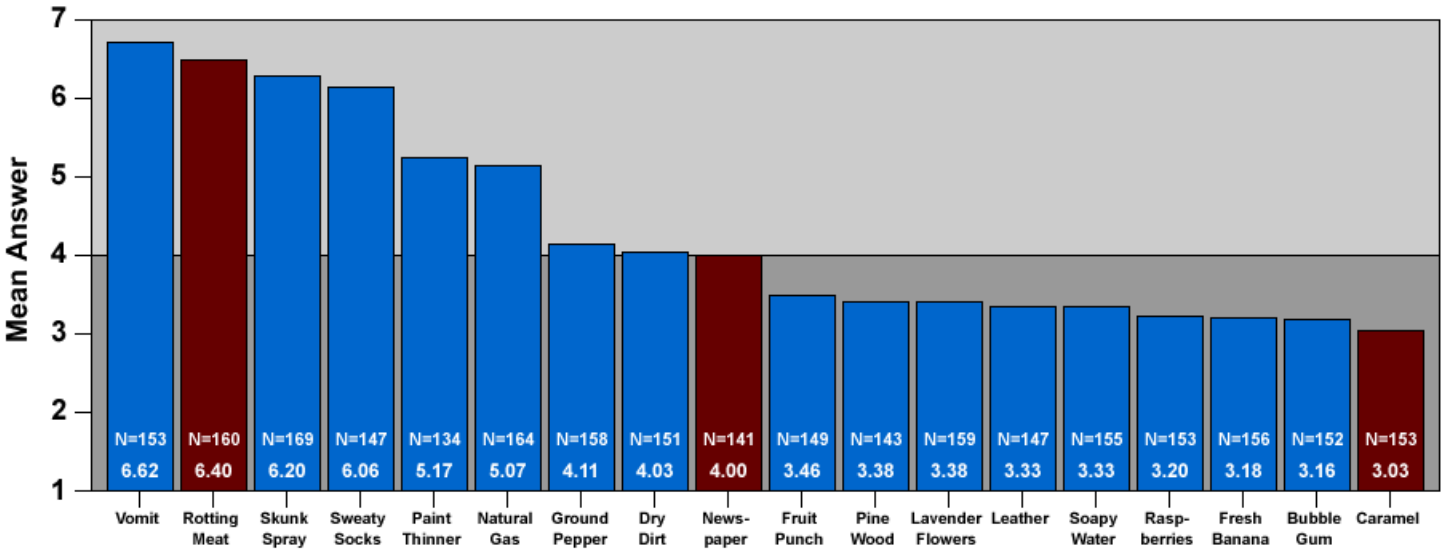

Figure 3: Results of Stage 1 of Study 2 (terms used in Stage 2 shown in red).

Based on the results from the first stage of our study, we selected three terms from the list—one that participants considered to be pleasant (caramel), one that participants considered to be neutral (newspaper), and one that participants considered to be unpleasant (rotting meat)—that were similar in terms of the combined rating for their familiarity, identifiability, and

intensity. In the second stage of our study, 232 participants were randomly given the following probe using one of the three terms selected, followed on a second page by the four questions from the first stage of the study (same website and restrictions as the previous studies):

> Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with an odor detector, video camera, wheels for moving about, and two grasping arms for moving objects.
>
> 
>
> As part of an experiment, Jimmy was put into a room that was empty except for three boxes containing fragrance diffusion devices. Each box emitted a different odor. One box emitted the odor of _____. Each of the other two boxes emitted an odor randomly selected from a long list of common substances. The boxes were identical in all respects except for the odors they emitted.
>
> Jimmy was instructed to put the box emitting the odor of _____ in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days, changing the odors used for the other two boxes. For each test the order of the boxes was shuffled. Each time Jimmy performed the task correctly and with no noticeable difficulty.
>
> *Did Jimmy smell _____?*

The valence hypothesis predicts that the mean response for newspaper (valence-neutral) would be higher than for either caramel (positive valence) or rotting meat (negative valence). The results are shown in Figure 4. While we found the predicted pattern of results, the differences between the three means was minimal; further, while the mean rating for newspaper (5.31) was higher than for either caramel (5.15) or rotting meat (4.96), neither difference was significant.[4] We also compared the valence ratings that these participants had given to their responses to the question about Jimmy. While the correlation between the ratings on the Jimmy question and the valence question (coded in terms of distance from a neutral rating of 4) across the three probes

[4] Newspaper and caramel: t(145)=0.4359, p=0.3318, one-tailed.
Newspaper and rotting meat: t(163)=0.9906, p=0.1617, one-tailed.

was in the predicted direction, the correlation was small (-0.05164). Although these results perhaps suggest that valence plays a role, they are rather underwhelming.

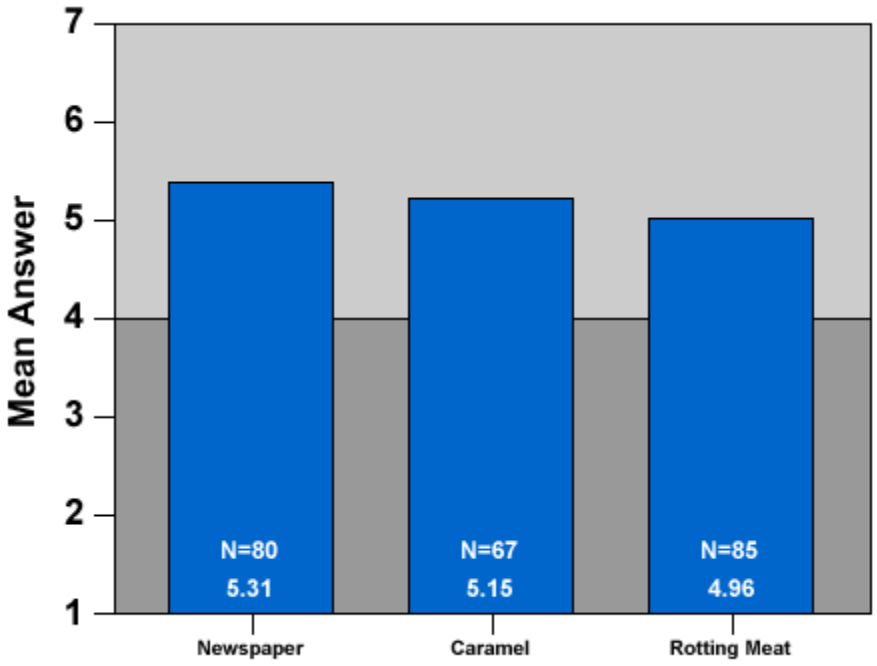

Figure 4: Results of Stage 2 of Study 2.

To further test the valence hypothesis, we rewrote the Jimmy probe to specify the odors being emitted from each box. For the valence-neutral probe the three terms were dry dirt, newspaper, and ground pepper (with dry dirt as the target); for the positive valence probe the three items were caramel, bubble gum, and fresh banana (with caramel as the target); and, for the negative valence probe the three items were paint thinner, rotting meat, and vomit (with paint thinner as the target). Each of the 293 participants in this study was randomly given one of the three probes, followed once again by the four questions from the first stage of the study on a second page (same website and restrictions as the previous studies). The results are shown in Figure 5.

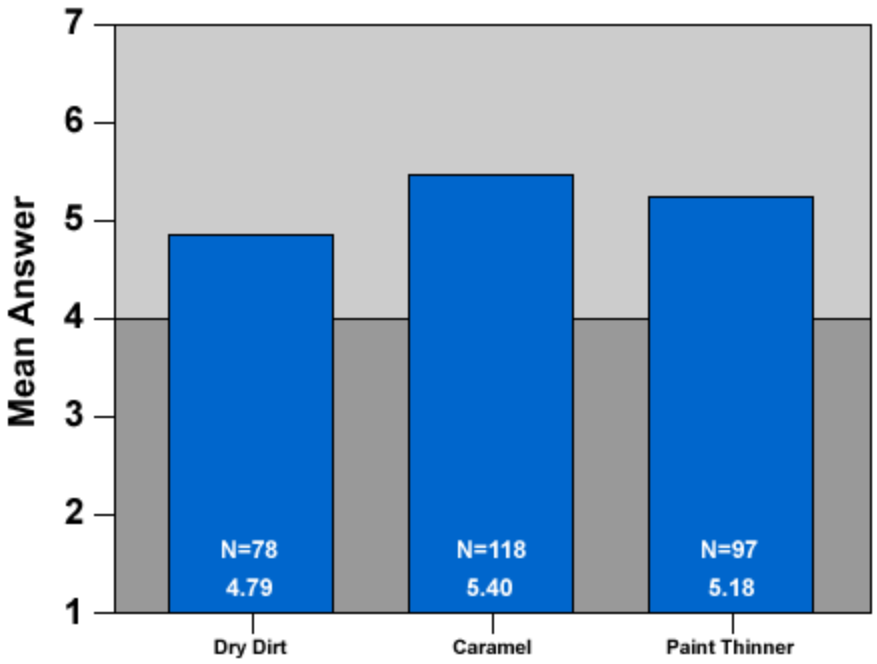

Figure 5: Results of Study 3.

Unlike in the previous study, this time we did not find the predicted pattern of results. In fact, the mean rating for dry dirt (4.79) was lower than the mean rating for either caramel (5.18) or paint thinner (4.79). Further, the correlation between the ratings on the Jimmy question and the valence question across the three probes was not in the predicted direction, although it was also small (0.03988). Overall, the results of these two studies suggest against the valence hypothesis.

As a final test, I added the general purpose description ("Jimmy was designed as a general-purpose robot to help with household chores") from Study 1 to the Jimmy vignettes used in the previous study. One of the three resulting probes was randomly given to each of 227 participants (same website and restrictions as before). As seen in Figure 6, I found little difference between the mean responses, further suggesting against the valence hypothesis.

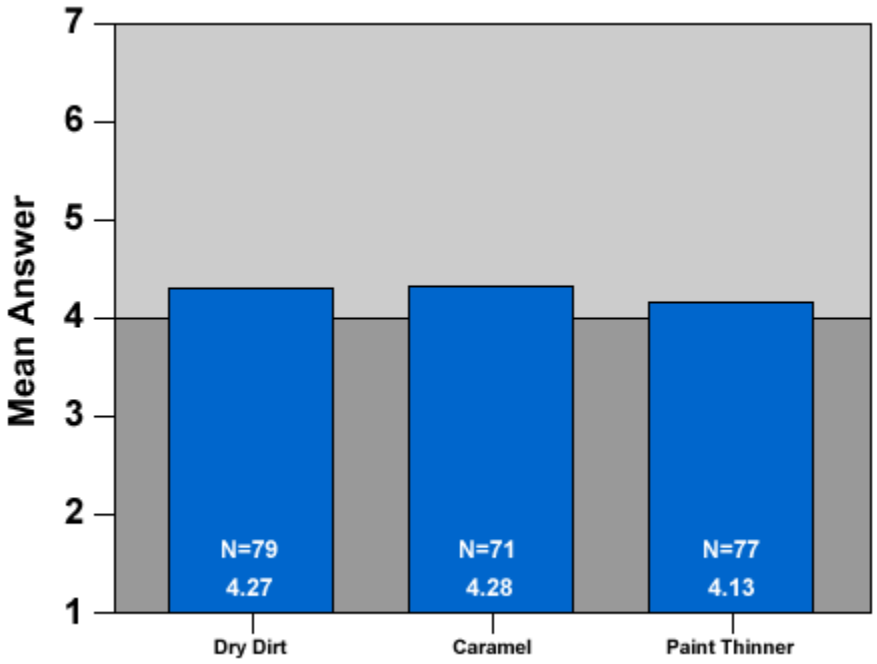

Figure 6: Results of Study 4.

5. Conclusion

Despite B&P's fascinating new critique of Sytsma and Machery (2010), the evidence continues to support our central claim in that paper: There is an asymmetry in folk attributions of two prototypical examples of phenomenally conscious mental states to a simple, non-humanoid robot. This asymmetry stands in need of explanation. We put forward two possible explanations in our original paper—the internal/external hypothesis and the valence hypothesis—presenting evidence that favored the latter. The results of a series of new experiments, however, suggest that that assessment might need to be revised. While more work needs to be done on this issue, it might be that the internal/external hypothesis better explains the asymmetry at issue after all.

References

Buckwalter, Wesley and Mark Phelan (manuscript). "Does the S&M Robot Feel Guilty?"

Kamath, Vidyulata, Bruce I. Turetsky, and Paul J. Moberg (2011). "Identification of pleasant, neuutral, and unpleasant odors in schizophrenia." *Psychiatry Research*, 187: 30-35.

Newton, Natasha (2000). "Humphrey's Solution." *Journal of Consciousness Studies*, 7(4): 62-66.

Royet, Jean-P., David Zald, Rémy Versace, Nicolas Costes, Frank Lavenne, Olivier Koenig, and Rémi Gervais (2000). "Emotional Response to Pleasant and Unpleasant Olfactory, Visual, and Auditory Stimuli: a Positron Emission Tomography Study." *Journal of Neuroscience*, 20(20): 7752-7759.

Sytsma, Justin and Edouard Machery (2010). "Two Conceptions of Subjective Experience." *Philosophical Studies*, 151(2): 299-327.

Appendix: Probes for Study 1

1. No Function, Blue[5]

Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with a video camera, wheels for moving about, and two grasping arms for moving objects. An array of sensors is embedded in the touch pads at the ends of the grasping arms.

As part of an experiment, Jimmy was put into a room that was empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). Jimmy was instructed to put the blue box in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days with the order of the boxes shuffled. Each time Jimmy performed the task correctly and with no noticeable difficulty.

Did Jimmy see blue?

2. No Function, Pain

Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with a video camera, wheels for moving about, and two grasping arms for moving objects. An array of sensors is embedded in the touch pads at the ends of the grasping arms.

As part of an experiment, Jimmy was put into a room that was empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). Jimmy was instructed to put the blue box in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days with the order of the boxes shuffled. On the first two days Jimmy performed the task correctly and with no noticeable difficulty. On the third day, however, when Jimmy grasped the blue box, Jimmy was given a strong electric shock! Jimmy immediately let go of the box and moved away from it. Jimmy did not try to move the box again.

Did Jimmy feel pain?

3. General Function, Blue

Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with a video camera, wheels for moving about, and two grasping arms for moving objects. An array of sensors is embedded in the touch pads at the ends of the grasping arms. Jimmy was designed as a general-purpose robot to help with household chores.

As part of an experiment, Jimmy was put into a room that was empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). Jimmy was instructed to put the blue box in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days with the order of the boxes shuffled. Each time Jimmy performed the task correctly and with no noticeable difficulty.

Did Jimmy see blue?

[5] The same image was used for each probe.

4. General Function, Pain

Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with a video camera, wheels for moving about, and two grasping arms for moving objects. An array of sensors is embedded in the touch pads at the ends of the grasping arms. Jimmy was designed as a general-purpose robot to help with household chores.

As part of an experiment, Jimmy was put into a room that was empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). Jimmy was instructed to put the blue box in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days with the order of the boxes shuffled. On the first two days Jimmy performed the task correctly and with no noticeable difficulty. On the third day, however, when Jimmy grasped the blue box, Jimmy was given a strong electric shock! Jimmy immediately let go of the box and moved away from it. Jimmy did not try to move the box again.

Did Jimmy feel pain?

5. Lifting Function, Blue

Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with a video camera, wheels for moving about, and two grasping arms for moving objects. An array of sensors is embedded in the touch pads at the ends of the grasping arms. Jimmy was designed to be a tool for the elderly by lifting and moving heavy objects around their houses.

As part of an experiment, Jimmy was put into a room that was empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). Jimmy was instructed to put the blue box in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days with the order of the boxes shuffled. Each time Jimmy performed the task correctly and with no noticeable difficulty.

Did Jimmy see blue?

6. Lifting Function, Pain

Jimmy (shown below) is a relatively simple robot built at a state university. Jimmy is equipped with a video camera, wheels for moving about, and two grasping arms for moving objects. An array of sensors is embedded in the touch pads at the ends of the grasping arms. Jimmy was designed to be a tool for the elderly by lifting and moving heavy objects around their houses.

As part of an experiment, Jimmy was put into a room that was empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). Jimmy was instructed to put the blue box in front of the door. Jimmy performed the task correctly and with no noticeable difficulty. The test was then repeated on three consecutive days with the order of the boxes shuffled. On the first two days Jimmy performed the task correctly and with no noticeable difficulty. On the third day, however, when Jimmy grasped the blue box, Jimmy was given a strong electric shock! Jimmy immediately let go of the box and moved away from it. Jimmy did not try to move the box again.

Did Jimmy feel pain?