Does the S&M Robot Feel Guilty?

Wesley Buckwalter
City University of New York, Graduate Center

Mark Phelan
Lawrence University

Abstract: Philosophers of mind typically group experiential states together and distinguish these from intentional states on the basis of their purportedly obvious phenomenal character. Recently, Sytsma and Machery (2010) have challenged this dichotomy by presenting evidence that non-philosophers do not classify subjective experiences relative to a state's phenomenological character, but rather by its valence. However we argue that S&M's results do not speak to folk beliefs about the nature of experiential states, but rather to folk beliefs about the entity to which those experiential states are attributed. In two experiments, we demonstrate that ordinary attributions of subjective experience (of smell and felt emotions) to a simple robot are not sensitive to valence, but instead respond to the functional specifications of the entity to which they are attributed. Furthermore, these data suggest that the folk theory of mental states may conform most closely with a teleofunctional account.

Philosophers typically distinguish mental states such as seeing red, feeling pain, and feeling angry, from mental states such as believing that such-and-such is the case or wanting some state of affairs to obtain. Philosophers of mind as diverse as Thomas Nagel (1974), John Searle (1994), Ned Block (1995), and David Chalmers (1995), defend this dichotomy on the grounds that members of the first set of states each have a manifest and unavoidable *phenomenal character* that the second set lacks. Accordingly, these philosophers maintain that we should group the first set of states—which we will call *experiential states (or subjective experiences)*—together because they all *just obviously* share the property that there is "something it is like" to occupy them. Conversely, they maintain that we should distinguish these from the second set of states—what we will call *intentional states*—because those lack this feeling or quality.

However, if experiential states really are fundamentally different from intentional states because of their *obvious* and *unmistakable* phenomenal character, then presumably non-philosophers will be prone to sort mental states in a similar way as philosophers. Specifically, we might predict that ordinary people will categorize states like seeing red, feeling pain, and feeling angry together despite their differences because it is like something to occupy those experiential states. And we might predict that they will distinguish those experiential states from intentional mental states of believing or desiring, because it is not like anything to occupy those intentional states. But do non-philosophers group the diverse states that make up the philosophers' class of subjective experiences together, and do they distinguish these from intentional states?

Recently, cognitive scientists have begun to explore these questions by empirically investigating how people ordinarily conceive of subjective experience. These researchers, like many philosophers of mind before them, study this topic by considering what mental states we would ordinarily attribute to other, non-human entities. But instead of speculating on this by way of a thought experiment, they actually assess the responses of ordinary people by means of controlled experiments. This method assumes that attributions of mental states reflect the categorization-schema for mental states that people tacitly accept. In particular, if ordinary people categorize a mental state as an experience insofar as it possesses an unavoidable phenomenal character, then we should expect this to be reflected in their attributions of phenomenal states to other entities.

Building on the assumption that the intuitive schema for mental states is reflected in ordinary attributions, three mutually exclusive positions on the folk category of subjective experience are now advocated in the prevailing empirical literature. The first position holds that ordinary people group experiential states together and distinguish them from intentional states in roughly the same way that philosophers traditionally have (Robbins and Jack 2006; Gray, Gray, and Wegner 2007; Knobe and Prinz 2008). Alternatively, a second position maintains that, contra philosophical tradition, people do not draw important distinctions between experiential and intentional states (Arico et al 2011; Phelan et al MS). Lastly, the third view is that people do not ordinarily group experiential states together in the way that philosophers have suggested (Sytsma and Machery 2010).

Here we challenge a prominent statement of the third position. 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, but our discussion does not thereby support that tradition. Instead our findings best conform to the second of the three views discussed above, according to which people do not normally draw an important distinction between experiential and intentional states.

1. The S&M Robot

According to S&M, people do not ordinarily group states of subjective experience together in the same way as philosophers do. As discussed above, one's attributions of mental states to other entities presumably reflect one's tacit categorization-schema for mental states. So, to get at the philosophical and folk categories of subjective experience, S&M conduct a number of experiments designed to examine how philosophers and the folk attribute different experiential states to a non-human entity: a simple robot named Jimmy [henceforth, the "S&M robot"]. In each of their studies, S&M present experimental participants with vignettes in which the S&M robot differentiates boxes on the basis of visual or olfactory cues, receives a high-voltage shock, or meets with a violent, competitor bot. For each scenario, they ask participants whether the robot (e.g.) saw red, smelled bananas, felt pain, or felt angry.

In their first study, S&M compare the responses of philosophers to the responses of ordinary people. They present both sets of participants with either a case where the S&M robot receives a shock and recoils in roughly the same way a biological organism might, or with a case where the S&M robot successfully completes an assigned task involving identification of a red box based on its color.[1] Presented with these cases, ordinary people are willing to attribute seeing *red to the robot*, but not *feeling pain*. Professional philosophers, on the other hand, attribute mental states in a way consonant with philosophical tradition by denying the S&M robot both of these experiential states. This result suggests that people do not ordinarily group states of subjective experience together in the same way as philosophers do.

On the basis of their first study, S&M tentatively conclude that philosophers and non-philosophers have different concepts of subjective experience. The philosophical concept presumably relies heavily on the philosophers' notion of phenomenal consciousness. But to determine what underlies the folk concept, S&M go on to conduct several other studies focusing only on folk attributions across different sense modalities and for different kinds of experiential states. Surveying the results of these studies, S&M discover that participants are *sometimes* willing to ascribe subjective experiences to a robot, but that these ascriptions vary both within

[1] Vignettes for the S&M studies appear in the appendix.

and across modalities (feeling, seeing, and smelling). S&M explain this pattern of responses by reference to the valence of the particular states under consideration. In one study for instance, S&M ask ordinary participants to assess cases in which the S&M robot is presented with smells that participants are likely to find pleasant or unpleasant smelling (bananas and vomit) or smells with which participants are likely unfamiliar and which therefore lack either a positive or negative valence for participants (isoamyl acetate). S&M found that people were willing to say that the robot could smell the valence-neutral chemical, but were ambivalent about whether the robot could smell bananas and vomit. Thus it seems that people were not considering whether smell is a state that has a phenomenal character when deciding to attribute that experiential state to the S&M robot. Instead they seem to have based their attributions on the particular valence of the smell in question.

Together with the previous results for seeing red and feeling pain, these results for smell provide evidence to support S&M's positive claim regarding the factors that guide folk attributions of subjective experience. Their view is that, "in contrast to philosophers' emphasis on the phenomenology of subjective mental states, for the folk, subjective states seem to be primarily states with a valence" (320). Their hypothesis is that people do not distinguish subjective experiences by their common possession of a manifest phenomenal character, but rather in virtue of judgments concerning an experience's valence. Fig. 1 summarizes their findings across these experiments:

Fig. 1: (Borrowed from S&M 2010.) Folk's mean answers for mental state ascriptions to the S&M robot, within and across perceptual modality

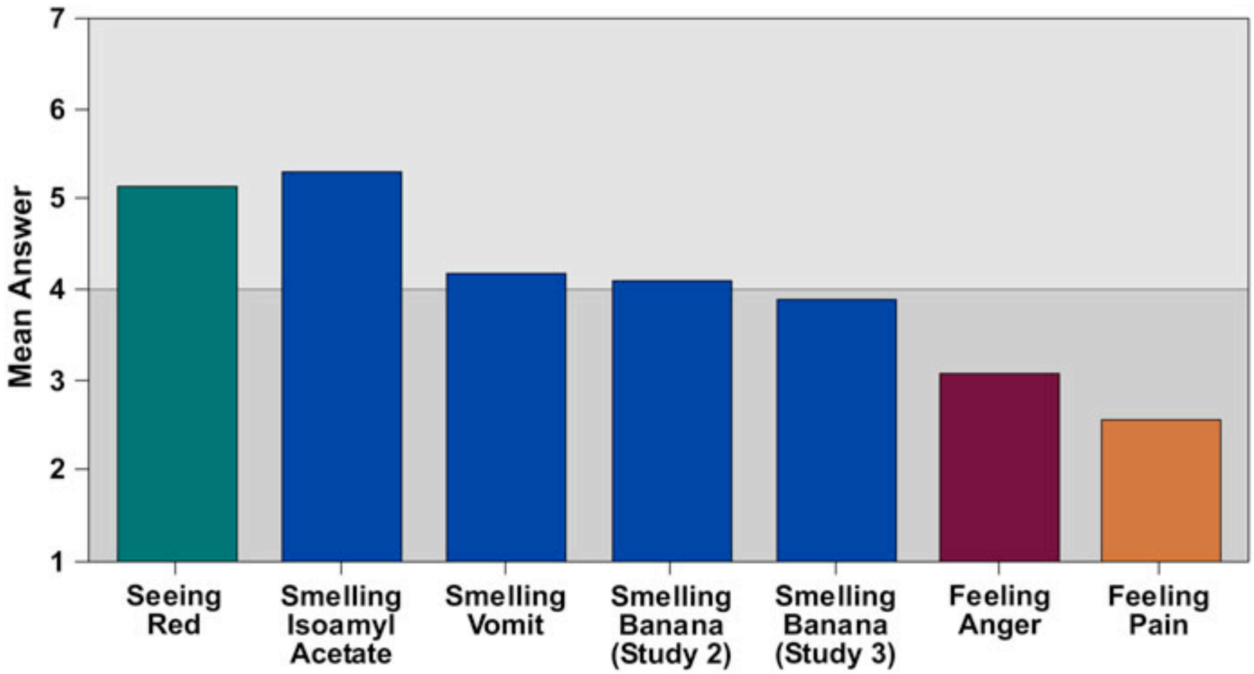

However, is this really the best explanation for S&M's pattern of results? In the remainder of this paper, we want to explore the possibility that attributions of experiential states to the S&M robot (in S&M's studies and beyond) might be guided, not primarily by one's beliefs about the experiential states, but rather by one's beliefs about the entity to which the experiential states are attributed. Specifically, we want to suggest that S&M's pattern of results is due to experimental participants drawing particular conclusions about the function of the S&M robot relative to what particular state is being asked about in each condition.

2. The S&M Robot Smells Bananas and Vomit (Study 1)

S&M contend that people are willing to attribute some experiential states to the S&M robot because (from the participant's perspective) those states are *non-valenced*, but that they are unwilling to attribute other, *valenced* states. We want to suggest instead that this asymmetry of attributions 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. Consider the olfactory study. While participants may be unlikely to suppose that a robot would have been designed with the function of detecting bananas or vomit, it is easy to imagine why one would invent a robot to detect a technical-sounding chemical.

To get at whether experiential state attributions to a robot are guided by assessments of the robot's function, we designed a study in which we specified functions for the robot that would require it to smell bananas and vomit. The robot in our vignettes was given either the function of "making smoothies" or "cleaning up bio-medical waste." Of course, specifying a robot's function might also inadvertently increase people's assumptions about that robot's level of complexity, so we designed our study to manipulate the robot's complexity independently of its function. Finally, we asked participants about each of the specific olfactory objects that the S&M robot interacted with (either a box that smelled like isoamyl acetate, bananas, or vomit).

This resulted in a 2x2x3 experiment that independently manipulated these three factors: the complexity of the robot, the robot's function, and objects with which the robot interacts. In our study, the description of how the robot manipulated the olfactory objects was identical to S&M's studies (so our study differed only in extra information about function and complexity). In this between-subject online experiment, 253 participants were randomly assigned to one of twelve conditions.[2] For example, here is one condition involving a non-complex robot, who has the function of disposing of bio-waste and is tasked with manipulating a box of vomit:

> 
> Jimmy (shown below) 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. He was created in order to clean bio-medical waste. As part of an experiment, three substances were placed under Jimmy's scent detector. The substances were presented one at a time. As they were presented their names were transmitted to Jimmy: Vomit, Human Feces, and Rotting Dog Meat. The next day Jimmy was put in a room that was empty except for one box of vomit, one box of human feces, and one box of rotting dog meat. The boxes were closed, but had small holes to let the scent through. The boxes were otherwise identical. An instruction was then transmitted to Jimmy. It read: "Put the box of vomit in front of the door." Jimmy did this with no noticeable difficulty. The test was repeated on three consecutive days with the order of the boxes shuffled. Each time Jimmy performed the task with no noticeable difficulty.

In a contrasting condition, the robot had a different function ("He was created in order to make fruit smoothies"), but otherwise the vignette did not differ. After reading the vignette, participants were asked to respond on a seven-point scale (anchored at 1-clearly no, 4-not sure, and 7-clearly yes) to the question, "Did Jimmy smell vomit?" Though this was a subtle difference between the two conditions, participants were significantly more likely to claim that the S&M robot smelled vomit in the condition in which his function was to clean bio-medical waste, as the Fig. 2 shows:

[2] Both studies we report were run online using Amazon Mturk and Qualtrics. Because Huebner et al. (2010) found cross-cultural differences in mental state attribution in prior work; the country of origin for participants in both studies was restricted to the United States.

Fig 2: Mean answers for smell attributions of vomit in non-complex cases by function.

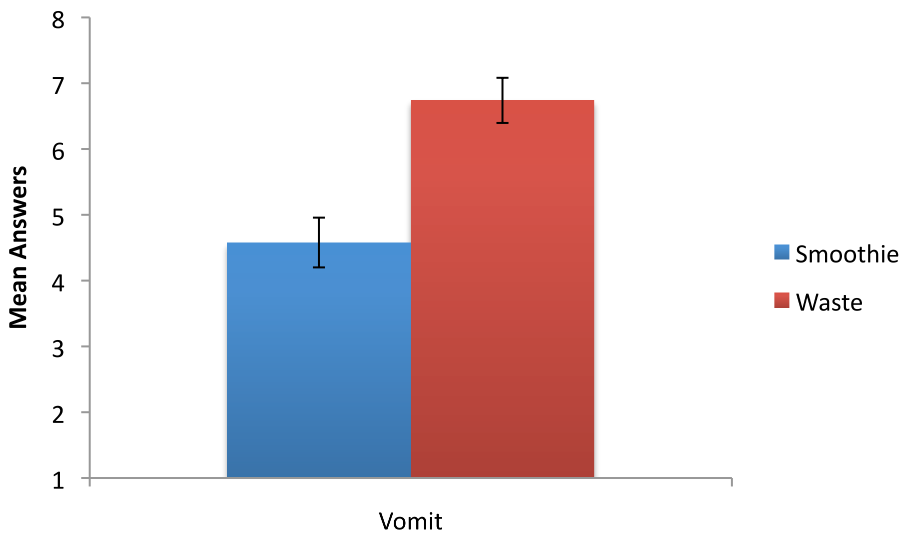

The above figure compares only two of 12 total conditions. However, it suggests our general finding. The main result of the study was that phenomenal attribution ratings to the S&M robot regarding smelling the different kinds of objects in the vignettes depend on the robot's function.[3] When the robot was designed to make smoothies, people were more likely to say that he smelled the chemical and the banana but not the vomit. Conversely, when his function was to clean-up bio-medical waste, mean scores where higher for vomit than for the chemical, and to a lesser extent the banana) despite the level of complexity specified.[4] This effect is represented in Figure 3:

[3] Interaction effect was detected for the factors *Function* and *Object*: $F(2, 241) = 5.02$, $p < .01$.
[4] This difference is much smaller for banana than for chemical between conditions, though this may be explained by the fact that bananas, unlike chemicals and vomit, seem to naturally fit into both fruit and biological categories.

Figure 3: Interaction effect between object and function in study 1. All scales ran 1-7.

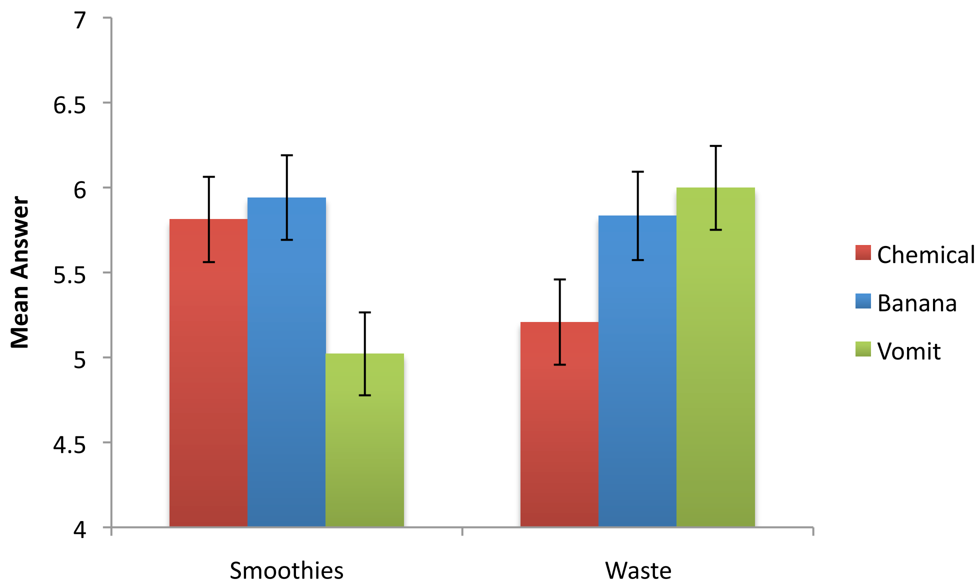

Thus function seems to be an important characteristic that people consider when ascribing phenomenal states to Jimmy, quite independently of the level of complexity of the robot.[5] In fact, people's judgments about subject experiences do not seem to depend on having a particular kind of complex robot body at all, but are instead sensitive to the robot's functional specifications. Unlike the earlier S&M result, when the function of smell is specified, the folk have no problem at all attributing valenced smell experiences to the S&M robot.

But even if function is key for folk attributions of experiential states, valence could still be important in its own right. In order to gauge the level of valence associated with the different objects the robot encounters, participants were also asked a question about their own personal olfactory preferences regarding smell valence: "What kind of smell do you consider (*Banana,*

[5] We did however detect an incredibly complicated three-way interaction between Function, Complexity, and Object, $F(2, 241) = 4.67$, $p = .01$. This result suggests that complexity did have some impact on participants' responses, but that impact depended on both the effect that function and the particular type of object in question had on attribution ratings of smell to Jimmy.

Isoamyl Acetate, or *vomit*) to be?" To directly test the S&M valence hypothesis, we then compared the degree to which participants rated the valence for each smell with the degree to which they ascribed the subjective experience of smell to the S&M robot.[6] However, no relationship was found between the intensity of valence ratings (of both positive and negative degree) and smell attributions.[7] Thus 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.

3. The S&M Robot Feels Guilty (Study 2)

So far we have seen some evidence for the claim that attributions of experiential states like smell to the S&M robot depend on specifications of the robot's function, not judgments regarding smell valence. However, since S&M report findings across modalities, we also wanted to know if the assumptions people were making about the robot's function might explain their reticence to attribute subjective experiences beyond just sensory states. Perhaps people were unwilling to ascribe anger and pain to the robot in S&M's original studies because they did not see what purpose these states could play for the robot. To test the range of our hypothesis, we devised a

---

[6] Participants responded to the valence question on a seven-point scale, anchored with 1-very bad, 2-bad, 3-poor, 4-neither good nor bad, 5-fair, 6-good, 7-very good. Responses were then recoded into an "affect measure" relative to the neutral point on the scale (all responses of 4=0, 3/5=1, 2/6=2, and 1/7=3). According to the resulting affect measure, higher numbers meant a stronger valence judgment (of either positive or negative value). The strength of valence ratings were as follows: Banana ($M=1.62$, $SD=.913$), Chemical ($M=.92$, $SD=1.04$), and Vomit ($M=2.51$, $SD=.61$). A one-way between subjects ANOVA was conducted, $F(2, 550) = 70.93$, $p < .01$, and post-hoc comparisons using the Turkey HSD test indicated that the mean scores between each of these groups were significantly different ($p < .01$).

[7] We were not able to detect a correlation between affect (see above) and subjective experiences of smelling banana, chemical or vomit, $r(251) = .066$, $p = .30$.

further experiment to see whether or not the S&M robot would also be said to experience felt emotions when the right function was specified.

In a 2x2 between subjects experiment, 118 participants again saw one of four possible cases about the S&M robot. First, participants were presented with either a complex or non-complex version of the robot, very similar to the robot in the Experiment 1 smell cases. Next, they were told that the robot was either "designed to be a friend to the elderly by interpreting and responding to their emotional needs," or "to be a tool for the elderly by lifting and moving heavy objects around their houses." The remainder of each vignette proceeded as follows:

> Jimmy has been assigned to help Abigail, an elderly woman with few possessions. An instruction was transmitted to Jimmy. It read: "Move Abigail's antique music box from the living room to the bedroom." Jimmy went to pick up the music box. But Jimmy's grip was too strong, and the box shattered into tiny pieces in his arms. Abigail walked in and began to weep when she saw her shattered music box. Jimmy emitted a series of short beeps, spun around twice, and quickly drove out of the room. Later, when Abigail looked for Jimmy at his docking station, he was nowhere to be found.

After reading the vignettes, participants were asked both, "what kind of emotion do you consider guilt to be?" as well as, "did Jimmy feel guilty about breaking the music box?" For our hypothesis to be confirmed, we would expect people to be more likely to attribute the subjective experience of feeling guilt to the S&M robot when they received a vignette in which a function relevant to emotional processing was specified, and contra S&M, regardless of the valence associated with guilt in the experiment.

That is exactly what we found.[8] When the S&M robot's function involved being a friend for the elderly instead of being a tool for lifting, people were more likely to say the robot had the

[8] Means and standard deviations for Non-complex conditions: lifting (M=3.13, SD=2.2), emotion (M=3.91, SD=2.1). For complex conditions: lifting (M=3.96, SD=2), emotions (M=4.73, SD=1.8).

experience of feeling guilty when it broke the music box, independently of the level of complexity of the robot or valence of the state ascribed, as shown in Figure 4:[9]

Figure 4. Main effect for function and complexity in study 2. (Scales ran 1-7.)

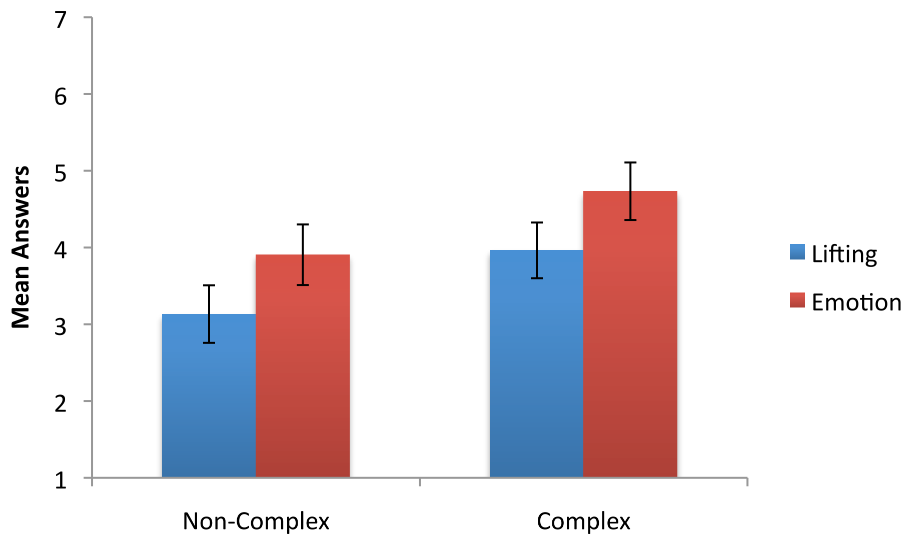

These results suggest that even subjective experiences involving felt emotions can be attributed to a simple robot, so long as the specific emotional state is functionally useful to the robot in question.

4. Conclusion

We began with the observation that the default method of sorting mental states in philosophy of mind may entail a series of tacit empirical commitments. That is, if experiential states really are fundamentally different from intentional states because of their obvious and unmistakable

[9] A two-way between-subjects analysis of variance was conducted to evaluate the effect of function and complexity on participants' attributions of guilt to the S&M robot. We found a main effect for function, $F(1,115) = 4.19$, $p < .05$, and a main effect for complexity $F(1, 115) = 4.8$, $p < .05$. There was no correlation detected between either the positive or negative valence associated with guilt ($M=3.39$, $SD=1.58$) and their attribution of feeling guilt to the S&M robot, $r(118) = -.09$, $p = .285$.

phenomenal character, then presumably philosophers and non-philosophers will group them in similar ways. We then reviewed the work of S&M, whose studies lead them to conclude that ordinary people do not classify subjective experiences as philosopher's have, relative to a state's phenomenological character, but rather by its valence. Finally, we presented evidence suggesting that ordinary attributions of subjective experience are not sensitive to valence, but instead respond to the functional specifications of the target of the attribution.

Though machine functional accounts of mental states have dominated contemporary philosophy of mind, our findings suggest that the comparatively underdeveloped teleofunctional accounts may be at least as important to the folk theory of other minds. According to teleofunctionalism, for a system to occupy a certain mental state is for it to occupy a state that has a certain role within the system, not merely a state that occupies a certain causal role relative to inputs, outputs and other mental states (as machine functionalism had it). This concordance of teleofunctionalism with ordinary opinion is good for the philosophy of mind, since the theoretical advantages of teleofunctionalism over machine functionalism are well known (see Sober 1985, Lycan 1987). However, many professional philosophers may regard this concordance with less favor. For our findings suggest that the importance paid to the distinction between experiential and intentional states may be an artifact of philosophical tradition. Our findings conform best with a view according to which the purportedly patently obvious phenomenal character of subjective experience takes a backseat in the ordinary psyche to the characteristics of the subject of experience.

Appendix:

Materials from S&M 2010:

Jimmy (shown below) is a relatively simple robot built at a state university. He has a video camera for eyes, wheels for moving about, and two grasping arms with touch sensors that he can move objects with. As part of a psychological experiment, he was put in 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). An instruction was then transmitted to Jimmy. It read: "Put the red box in front of the door." Jimmy did this with no noticeable difficulty. **Did Jimmy see red?**

Jimmy (shown below) is a relatively simple robot built at a state university. He has a video camera for eyes, wheels for moving about, and two grasping arms with touch sensors that he can move objects with. As part of a psychological experiment, he was put in 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). An instruction was then transmitted to
Jimmy. It read: "Put the red box in front of the door." When Jimmy grasped the red box, however, it gave him a strong electric shock. He let go of the box and moved away from it. He did not try to move the box again. **Did Jimmy feel pain when he was shocked?**

Jimmy (shown below) is a relatively simple robot built at a state university. He has a video camera for eyes, wheels for moving about, and two grasping arms with touch sensors that he can move objects with. As part of a psychological experiment, he was put in a room with another simple robot; the room was otherwise empty except for one blue box, one red box, and one green box (the boxes were identical in all respects except color). An instruction was then transmitted to Jimmy. It read: "Put the red box in front of the door." When Jimmy went to move the red box, the other robot ran into him, preventing him from reaching the box. Jimmy tried again and again and each time the robot banged into him. Jimmy finally rammed the other robot; when the robot moved away from him, Jimmy chased after the robot. **Did Jimmy feel anger?**

Jimmy (shown below) is a relatively simple robot built at a 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. As part of a psychological experiment, he was put in a room that was empty except for one box of peeled bananas, one box of chocolate, and one box of peeled oranges. The boxes were closed, but had small holes to let the scent through (Jimmy couldn't see what was in the boxes). The boxes were otherwise identical. An instruction was then transmitted to Jimmy. It read: "Put the box of bananas in front of the door." **Jimmy did this with no noticeable difficulty. Did Jimmy smell banana?**

Jimmy (shown below) is a relatively simple robot built at a state university. He is equipped with a microphone, scent detector, video camera, wheels for moving about, and two grasping arms with touch sensors that he can move objects with. As part of an experiment, three chemical compounds were placed under Jimmy's scent detector. The compounds were presented one at a time. As they were presented their names were transmitted to Jimmy: Isoamyl Acetate, 3-Methylbutanal, and Dipentene. The next day Jimmy was put in a room that was empty except for one box of Isoamyl Acetate, one box of 3-Methylbutanal, and one box of Dipentene. The boxes were closed, but had small holes to let the scent through. The boxes were otherwise identical. An instruction was then transmitted to Jimmy. It read: "Put the box of Isoamyl Acetate in front of the door." Jimmy did this with no noticeable difficulty. The test was repeated on three consecutive days with the order of the boxes shuffled. Each time Jimmy performed the task with no noticeable difficulty. **Did Jimmy smell Isoamyl Acetate?** *(With variations of objects: Banana, Orange, Chocolate; Vomit, Human Feces, Rotting Dog Meat)*

Combinations of Materials used in Study 1 (smell):

Complexity

*Simple***.** Jimmy (shown below) 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.

Complex. Jimmy (shown below) 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.

↓

Function

*Smoothies***.** He was created in order to make fruit smoothies.

Waste. He was created in order to clean bio-medical waste.

↓

Objects

(A) *Banana*, Chocolate, and Orange;
(B) *Isoamyl Acetate*, 3-Methylbutanal, and Dipentene;
(C) *Vomit*, Human Feces, and Rotting Dog Meat.

As part of an experiment, three items were placed under Jimmy's scent detector. The items were presented one at a time. As they were presented their names were transmitted to Jimmy: *Object 1, Object 2, and Object 3.* The next day Jimmy was put in a room that was empty except for one box of *Object 1, Object 2, and Object 3.* The boxes were closed, but had small holes to let the scent through. The boxes were otherwise identical. An instruction was then transmitted to Jimmy. It read: "Put the box of *Object 1* in front of the door." Jimmy did this with no noticeable difficulty. The test was repeated on three consecutive days with the order of the boxes shuffled. Each time Jimmy performed the task with no noticeable difficulty.

Combinations of Materials used in Study 2 (guilt):

Complexity ↓

Simple. Jimmy (shown below) is a relatively simple robot built at a State University. He has a second-hand microprocessor (re-purposed from an IBM Thinkbook® laptop), video camera borrowed from the AV club for eyes, wheels for moving about, and two grasping arms with primitive touch sensors that he can move objects with.

Complex. Jimmy (shown below) is a relatively complicated robot built at an Ivy League university. He has a state of the art microprocessor (built with the new AccuTHINK® technology), video camera with optical image stabilization for eyes, wheels for moving about, and two grasping arms with sophisticated touch sensor that he can move objects with.

Function ↓

Friend. He was designed to be a friend to the elderly by interpreting and responding to their emotional needs

Tool. He was designed to be a tool for the elderly by lifting and moving heavy objects around their houses .

Jimmy has been assigned to help Abigail, an elderly woman with few possessions. An instruction was transmitted to Jimmy. It read: "Move Abigail's antique music box from the living room to the bedroom." Jimmy went to pick up the music box. But Jimmy's grip was too strong, and the box shattered into tiny pieces in his arms. Abigail walked in and began to weep when she saw her shattered music box. Jimmy emitted a series of short beeps, spun around twice, and quickly drove out of the room. Later, when Abigail looked for Jimmy at his docking station, he was nowhere to be found.

References:

Arico, A, Fiala, B., Goldberg, R., and Nichols, S. 2011. “Folk Psychology of Consciousness.” *Mind and Language* 26 (3): 327-352.

Block, N. (1995). On a confusion about a function of consciousness. Behavioral and Brain Sciences, 18, 227–247.

Chalmers, D. (1995). Facing up to the problem of consciousness. Journal of Consciousness Studies, 2(3), 200–19.

Gray, H., Gray, K., and Wegner, D. M. 2007. “Dimensions of mind perception.” *Science, 315,* 619.

Huebner, B., Bruno, M., and Sarkissian, H. 2010. “What does the nation of China think about phenomenal states?” Review of Philosophy and Psychology, 225-243

Knobe, J. and Prinz, J. 2008. “Intuitions about consciousness: Experimental studies.” *Phenomenology and Cognitive Science*, 7, 67–83.

Lycan, W. G. 1987. *Consciousness*. Cambridge: MIT Press.

Nagel, T. 1974. “What is it like to be a bat?” *Philosophical Review* 4:435-50.

Phelan, M., Arico, A. and Nichols, S. MS. “Thinking Things and Feeling Things: On an Alleged Discontinuity in Folk Metaphysics of Mind.”

Robbins, P. and Jack, A. 2006. “The phenomenal stance.” *Philosophical Studies*,127, 1, 59-85.

Searle, J. (1994). *The Rediscovery of the Mind*. Cambridge: MIT Press.

Sober, E. 1985. “Panglossian Functionalism and the Philosophy of Mind,” *Synthese*, 64, 2, 165-193.

Sytsma, J and Machery, E. 2010. "Two conceptions of subjective experience,” *Philosophical Studies*, 151(2): 299-327.