

Unfinished draft. Please do not quote or cite without authors' permission.

The Folk Psychology of Consciousness*

Adam Arico[†], Brian Fiala[†], Robert F. Goldberg[‡], and Shaun Nichols[†]

[†]*Department of Philosophy, University of Arizona*

[‡]*Center for Cognitive Neuroscience, University of Pennsylvania*

Abstract: We find it so natural to think that other people and animals are conscious that it masks a deep puzzle: why do we think anything else is conscious? This paper develops the "Agency Model," according to which specific low-level cues lead people to categorize an entity as an AGENT, and that suffices to generate an inclination to attribute conscious mental states to the object. To explore whether people have this inclination, a speeded object-attribution experiment was conducted. The paper describes several models consistent with the data and discusses some of the philosophical implications of the various models.

* All authors made roughly equal contributions to the paper; the order of authors is alphabetical. Versions of this paper were presented at the University of Arizona Center for Consciousness Studies, Toward a Science of Consciousness 2008 Conference, University of Arizona's Experimental Philosophy Lab, the Experimental Philosophy Workshop at the 2008 meeting of the Society for Philosophy and Psychology, and the *Integrating Science and the Humanities* conference at the University of British Columbia. We would like to thank the audiences for their comments. In addition, we are grateful to Mike Bruno, Peter Carruthers, Yali Corea-Levy, Jane Erickson, Ian Evans, Jen Giesselmann, Josh Greene, Terry Horgan, Bryce Huebner, Tony Jack, Chris Kahn, Benji Kozuch, Uriah Kriegel, Trevor Kvaran, Theresa Lopez, John Pollock, Jesse Prinz, Philip Robbins, Deena Skolnik-Weisberg, Justin Sytsma, and especially Josh Knobe and Edouard Machery for suggestions. Arico would also like to thank the Earhart Foundation for providing financial support in the form of a H.B. Earhart Graduate Fellowship.

1. The problem of other conscious minds

We find it so natural to think that other people and animals are conscious that it is easy to overlook a deep puzzle: *why* do we think anything else is conscious? Consciousness is widely regarded to be a paradigmatically private affair, something that could never be publicly observed. So what prompts us to think that others have conscious minds at all?¹

The fact that we are inclined to attribute conscious states to other creatures has long been recognized as puzzling by philosophers (Augustine *De Trinitate* 8.6.9; Reid 1969[1785]). We can think of this puzzle as the *descriptive* problem of other conscious minds.² Philosophers have made various speculations about what leads us to think that others have conscious states. J. S. Mill develops a version of the most famous answer to the puzzle: I come to believe in other minds by drawing an analogy with my own case. He writes:

by what considerations am I led to believe, that there exist other sentient creatures; that the walking and speaking figures which I see and hear, have sensations and thoughts, or in other words, possess Minds?... I conclude that other human beings have feelings like me, because, first, they have bodies like me, which I know, in my own case, to be the antecedent condition of feelings; and because, secondly, they exhibit the acts, and other

¹ For the purpose of this paper, we focus only on the cognitive process(es) underlying token attributions of conscious experiences to different entities. We are not here concerned with the folk's metacognitive recognition that one's attribution is an attribution *of consciousness*, but with the process that disposes the folk to attribute a particular state (which the cognitive scientist or philosopher might classify as *phenomenal*) to a given entity. The question of how the folk conceptually categorize those states that they do attribute is orthogonal to the question we are attempting to answer here.

² We focus on the problem of why we are inclined to attribute conscious states to others. This can be thought of as part of a broader problem, 'the problem of other minds,' which asks why we are inclined to attribute any mental states whatsoever to other beings. But for philosophers like Augustine, Reid, and Mill, the problem of other conscious minds would have been a critical part of the broader problem of other minds.

outward signs, which in my own case I know by experience to be caused by feelings (Mill 1865).

Mill's statement suggests the following procedure, which involves two applications of analogical reasoning: First, I know that my feelings come via my body, and I appreciate the analogy between my body and the bodies of others. Second, I know that certain behaviors of mine are caused by mental states, so when I see other analogous bodies exhibit similar behaviors, I again deploy analogical reasoning to conclude that their behaviors are also caused by mental states. It's easy to see why this view has been historically attractive. Analogy provides a sensible and familiar explanation for our belief in other minds.

Although the analogical theory has historical weight and a few contemporary advocates (e.g. Hill 1984; Hyslop 1995), the currently dominant view is that we believe in other conscious minds because it's the best explanation for what we observe. This approach has been developed in recent years by a number of people. For instance, Robert Pargetter asks, 'What is the nature of the inferences that we all so commonly, and rightly, make from certain behavioural evidence to the mental lives of other people?' He suggests that 'these inferences should best be viewed as being common scientific or hypothetic inferences, or arguments to the best explanation' (Pargetter 1984, 158). The idea is that we come to believe in other conscious minds by using good inductive techniques – appealing to other conscious minds is the best explanation for the behavior we observe.³ This basic approach has been widely adopted in recent philosophy. Indeed, it's promoted in standard textbooks in the philosophy of mind (Churchland 1988, 71-72;

³ It should be noted, though, that Inference to the Best Explanation and Analogy need not be seen as mutually exclusive competitors. One could in principle take a particular argument from analogy to be the best explanation for certain kinds of behavior. In other words, one might reason along the following lines: 'well, the best explanation for that person's behavior, which is quite similar to my own, is that they are experiencing the same things as myself.'

Graham 1998, 57-63).⁴

The philosophical appeals to analogy and best explanation are based on speculation and informal observation. But over the last several decades there has been an impressive body of empirical work on how children and adults attribute mental states across a broad range of conditions (for reviews see e.g., Goldman 2006, Gopnik & Meltzoff 1997, Nichols & Stich 2003, Perner 1991, Wellman 1990). Some of the mental states that have been explored in this literature are, of course, conscious mental states (see e.g., Paul Harris on emotions 1989). Although this work has taught us a great deal, for example, about the attribution of false belief in young children, very little of this work is focused directly on the traditional question, ‘what leads us to think that particular individuals are bearers of *conscious* mental states?’. That is, there is little work aimed at explaining what leads us to categorize something as a candidate for having conscious states at all.

In addition to the intrinsic interest of the descriptive problem, solving the problem will also help us to evaluate whether the processes by which we gauge consciousness are in fact good processes. A familiar view in philosophy is that the way we in fact identify others as conscious is, in general, a rationally sound method (Mill 1865; Pargetter 1984). That is, the path by which we identify others as conscious is, roughly, a rationally appropriate path. However, it is an old and familiar philosophical worry that the basis for our thinking that others are conscious is entirely spurious (e.g. Descartes 1641/1986). It’s clear that determining whether such a worry is misplaced requires a plausible characterization of the actual basis on which we come to think of

⁴ Despite the enthusiasm for the best explanation approach, a number of philosophers have worried that it is inadequate for capturing why we attribute *consciousness* to others. The worry is that, while it is easy enough to see why the best explanation for behavior would require attributing information-processing states to others, it’s much harder to see why the best explanation for behavior would require attributing *conscious* states (e.g. Melnyk 1994).

others as conscious.

Moreover, solving the descriptive problem of other minds can also help us to assess our reactions to particular cases. For example, intuitions about whether a particular individual (or class of individuals) is conscious can inform moral judgment in both philosophy and everyday life. In the philosophy of mind, our natural intuitions about whether a particular kind of thing could be conscious have played an important role in arguments over the nature of consciousness itself (Block 1978, Searle 1980). If we determine that our intuitions are trustworthy under the relevant circumstances, then so much the better. But if we determine that our intuitions are liable to mislead in particular ways, then we might attempt to guard against these biases. The mechanisms that generate our intuitions might be too liberal – that is, they might lead us to think an entity has conscious states when it does not. The mechanisms might also be too conservative – that is, they might lead to mistaken intuitions that a certain kind of entity does not have consciousness when, in fact, it does. But before we can assess whether the mechanisms lead to defective intuitions, we need to know more about the mechanisms themselves.

2. The problem of other minds in psychology

The descriptive problem of other minds has a long history in philosophy, and questions about how people attribute mental states to others have been intensely explored in cognitive science over the last two decades (see e.g., Goldman 2006, Nichols & Stich 2003, Perner 1991, Wellman 1990). As noted above, little of this work is focused directly at our question of what leads people to attribute *conscious* states—like pain—to others. However, there is one strand of work that provides a very promising framework. This work focuses on what generates the tendency to think that a certain object has mental states *at all*.

One of the earliest investigations on the topic is Heider & Simmel's study on how adults describe an animation involving geometric objects (1944). They found overwhelmingly that adults describe the animation by adverting to mental states. For example, at one point in the animation, the big triangle repeatedly bumps up against the inside edge of a rectangle, and nearly all subjects say that the triangle wants to get out of the box. More recent work shows that children respond in much the same way to these sorts of stimuli. Like the adults in Heider & Simmel's study, when children are asked to describe what they saw, they advert to the goals, beliefs, and intentions of the triangles in a 2D animation (e.g. Bowler & Thommen 2000; Abell et al. 2000; Dasser 1989). If you've watched one of these animations, the results will come as no surprise. It's extremely natural to see these objects as having mental states because the motion trajectories of the triangles 'push the right buttons' to trigger mind attribution.⁵ This becomes evident when one contrasts Heider-Simmel style animations with an animation of triangles moving about the surface in straight lines at constant speeds. In that case, there is no inclination to start attributing mental states to the triangles. Motion alone is not sufficient. But it remains possible that relatively simple motion cues suffice for agency attributions. For instance, change in speed plus change in direction might be sufficient to generate an attribution of a mind, even if nothing can be discerned about the goals or thoughts of that mind (Scholl & Tremoulet 2000, 305). Of course, as adults, we don't cave to our first-blush intuitions of mentality here – we know, on slight reflection, that the images don't have minds. Nonetheless, there clearly is a mechanism that generates these powerful, if overridable, inclinations to attribute mental states,

⁵ It's possible that people regard the Heider-Simmel stimuli as fictional. But the point can also be made by adverting to natural reactions to the emerging class of robots – Roombas and Rodney Brooks' Mobots – that behave in goal-like ways. Watching these seems to trigger an attribution of mental states. Thanks to Bernie Kobes for these points.

and this mechanism likely plays an important role in everyday attributions of mental states.

Susan Johnson and colleagues use very different techniques to discern the mechanisms underlying the attribution of a mind to an individual. There are several ways that a baby might reveal that she thinks an object has a mind: she might follow the 'gaze' of the object, try to communicate with the object, imitate the behaviors of the object, or attribute goals to the object. By exploiting this variety of indicators, Johnson provides strong evidence that infants attribute minds as a result of particular kinds of relatively simple features that serve as cues. In one experiment, 12 month old infants were shown a fuzzy brown object under a variety of different conditions (Johnson, Slaughter & Carey 1998). In one condition, the fuzzy brown object (with no facial features) interacted contingently with the infant by beeping and flashing when the infant babbled or moved; in another condition, the fuzzy brown object exhibited an equivalent amount of flashing and beeping, but in this condition the activity was not contingent on the infant's behavior. In both conditions, children's looking behavior was measured when the fuzzy brown object 'gazed' at one of two objects by making a smooth, 45 degree turn towards the object and remaining in this orientation for several seconds. What Johnson and colleagues found was that infants would follow the 'gaze' of the fuzzy brown object, but only when its beeping and flashing were contingent. In another set of conditions the fuzzy brown object did not flash or beep, but Johnson and colleagues found that babies were more likely to follow the 'gaze' of the fuzzy brown object when it had eyes than when it did not.

In other experiments, babies were shown a stuffed orangutan that had a face and exhibited contingent interaction. Babies imitated the behavior of the stuffed animal and made communicative gestures toward it, indicating that the babies coded the object as having a mind (Johnson et al. 2001).

In more recent experiments, Johnson and colleagues devised a new object, the blob, a bright green object about the shape of an adult shoe that had no facial features but could beep and move around on its own. In one condition, a confederate engaged the blob in ‘small talk’, and the blob beeped contingently with the confederate; in the other condition the blob beeped randomly. Again they found that babies would follow the ‘gaze’ of the blob, but only in the contingent interaction condition.

Finally, and perhaps most impressively, the blob design has recently been coupled with Amanda Woodward’s (1998) goal attribution experiment. Woodward (1998) showed babies an arm moving towards one of two locations, each containing one of two objects. Then the locations of the objects were switched. Babies looked longer when the arm reached to the same location that now held a different object, suggesting that the babies expected the arm to reach for the same goal-object. Shimizu & Johnson (2004) found something similar with the blob – babies looked longer when the blob moved in the same direction but towards a different object. But this *only* occurred when the blob had behaved contingently with the confederate. What’s particularly striking about this case is that it shows that once the child regards the blob as an agent the child then makes *predictions* about the blob’s behavior.

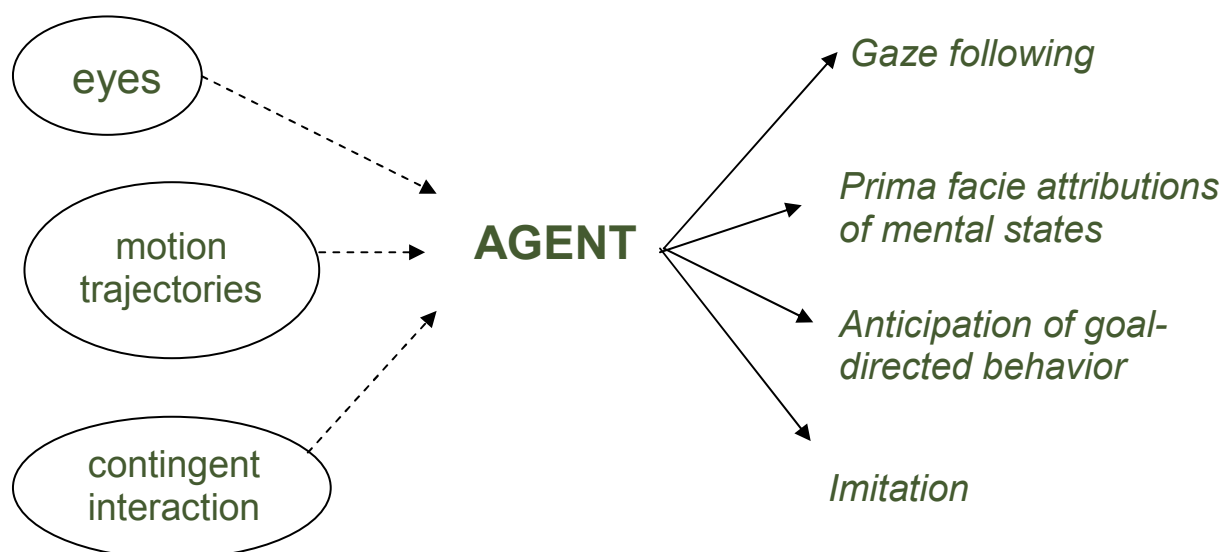


Figure 1

We will use AGENT as a technical category, central features of which are represented in figure 1.⁶ If an entity exhibits one of the relatively simple features mentioned above, this will trigger the inclination to categorize the entity as an *AGENT*. Once an individual identifies an entity as an AGENT, this sets up certain dispositions in the person – e.g., the disposition to attribute goals, to imitate, to anticipate goal directed behavior, and prima facie attributions of mental states (as in the Heider-Simmel studies).⁷

Our proposal is that, when it comes to the tendency to attribute mental states to an individual, there is nothing special about *conscious* mental states. Rather, once a person categorizes an individual as an AGENT, she will be inclined to attribute conscious mental states to the individual.⁸ This account, which we'll call the 'AGENCY model', can be represented as a minor adjustment to the account of the AGENCY concept, yielding figure 2. The AGENCY model builds naturally on extant work to explain what leads us to think that others have conscious states. But there is a developing literature in philosophy and psychology that seems to challenge it.

⁶ In keeping with the idea that AGENT is a technical category, we do not mean to suggest that the AGENT concept corresponds to any English word. We do, of course, maintain that it is an internally represented concept, though.

⁷ As we will elaborate later, this is not an exhaustive characterization of the AGENCY category.

⁸ For present purposes, we are neutral on whether the AGENT concept is innate. If the concept is innate, then the view begins to look similar to Thomas Reid's solution to the problem of other minds (1969 [1785]). Our point is that, regardless of how the AGENT concept is acquired, once an entity is categorized as an AGENT, an inclination to attribute conscious states to the entity will be present.

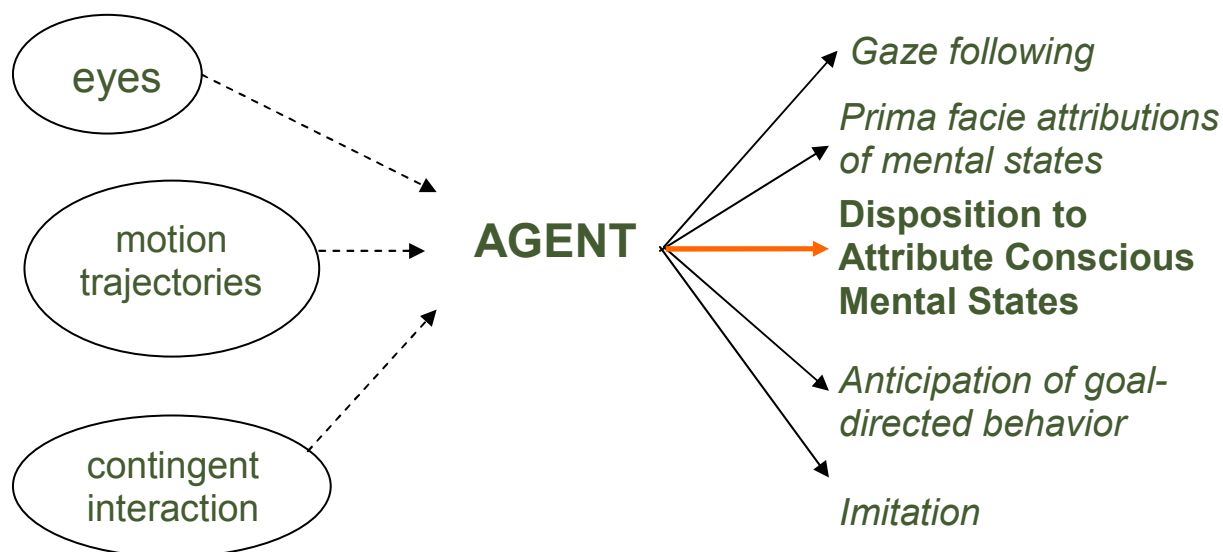


Figure 2: The AGENCY Model

3. Agency and Experience: A Dissociation?

Although it is *prima facie* plausible that the attribution of agency would suffice to dispose one to attribute conscious states, several researchers have recently argued that agency and experience are independent dimensions of mental attributions (Gray, Gray, and Wegner 2007, Knobe & Prinz 2008, Robbins & Jack 2006).

Heather Gray, Kurt Gray and Daniel Wegner collected data from over 2000 respondents to an online survey. Subjects were presented with 78 pairings of 13 different characters, and then asked to rate (on a 5-point Likert scale) which character was had a greater capacity for possessing a certain particular mental trait (e.g., ‘Which character is more capable of experiencing joy?’). The researchers maintain that the responses reveal two quite different dimensions of mind perception: Agency and Experience (Gray et al. 2007, 619). Moreover, Gray and colleagues conclude from the data that the two features, Agency and Experience, can be dissociated. God, for instance, rates at the top of the Agency scale, but at the absolute bottom of Experience; frogs and fetuses, conversely, rate high on the Experience scale while lacking any

degree of Agency.

Although the work of Gray et al. might be thought to challenge our AGENT-based approach, there is a natural reply. The capacities that fall under the category of Agency in their study are fairly sophisticated: self-control, morality, memory, emotion recognition, planning, communication, and thought. Paradigmatic entities that would trigger AGENT in our sense, like frogs, are judged to lack many of those sophisticated capacities. Thus, while frogs are extremely low on ‘agency’ by the standards of Gray and colleagues, frogs certainly qualify as AGENTS on our notion. Furthermore, the traits that fall under Gray and colleague’s category of ‘experience’ also often involve sophisticated capacities, e.g., personality, pride, embarrassment, and joy. So, Gray and colleagues’ results leave open the possibility that a minimal notion of agent of the sort we have in mind would produce a tendency to attribute minimal kinds of experience, like pain. The AGENCY model aims to capture the most basic features that generate the disposition to attribute the most basic conscious states.

Joshua Knobe and Jesse Prinz present a different kind of evidential argument for dissociating agency and experience, suggesting that people attribute agency while positively *resisting* consciousness attributions. Their studies focus on attributions of mental states to groups. In one study, Knobe and Prinz presented subjects with two sets of sentences; one set attributed non-conscious mental states to a group entity, while another attributed conscious mental states to that same entity. The non-conscious state attributions consisted of sentences like ‘Acme Corp intends to release a new product this January’ and ‘Acme Corp wants to change its corporate image.’ The conscious state attributions included such sentences as ‘Acme Corp is now experiencing great joy’ and ‘Acme Corp is feeling excruciating pain.’ Knobe and Prinz found that subjects rated sentences about Acme Corp.’s non-conscious mental states as sounding

‘natural’, while they rated sentences about Acme Corp.’s conscious mental states as sounding ‘weird’. The moral of these experiments, they say, is that ‘people are unwilling to ascribe to group agents states that require phenomenal consciousness’ (Knobe & Prinz 2008).⁹

Knobe & Prinz’s work indicates that people regard groups as agents in *some* sense, but this notion of agency might differ from our more primitive notion of AGENT. Group agents might fail to activate the primitive notion of AGENCY.¹⁰ Indeed, although this has not been studied yet, we suspect that there is a developmental pattern concerning the attribution of mental states to groups such that young children will be less likely than adults to attribute mental states to groups. If this turns out to be right, it would suggest that children are not identifying groups as

⁹ One possible deflationary interpretation of Knobe & Prinz’s results is that participants were speaking figuratively when attributing mental states to groups. We conducted pilot studies to investigate this possibility. To screen out those who fail to distinguish between literal and figurative language, we gave subjects a series of sentences, some clearly figurative (e.g., ‘Einstein was an egghead’) and some clearly literal (e.g., ‘Carpenters build houses’), and asked them to rate those sentences on a 7-point Likert scale of literalness (1= ‘Figuratively True’ 7= ‘Literally True’). We also provided those subjects with sentences attributing different mental states to individuals (e.g., ‘Some millionaires want tax cuts’) and to groups (e.g., ‘Some corporations want tax cuts’), and compared their judgments regarding different types of mental states. We found that when subjects rated group attributions, they tended to rate non-conscious state attributions as ‘literally true’ and gave significantly higher ratings for non-conscious than conscious mental state attributions ($t(66) = 7.735, p < .001$). For further discussion of Knobe & Prinz’s experiments, see Arico (forthcoming) and Sytsma & Machery (forthcoming a).

¹⁰ Arico (forthcoming) argues that groups can be ‘snuck in through the side door’ of AGENT categorization by situating the attribution in a context traditionally associated with Agents. He found that subjects rate attributions of conscious states to groups as sounding significantly ‘weirder’ than those same attributions to individuals. He also found, however, that this reaction is mitigated when the attribution includes a contextual setting. Arico explains this by saying that while groups do not display the low-level features that immediately incline us to categorize them as AGENTs, given sufficient description of the context of attribution, cognitive associations could lead one to categorize groups as AGENTs.

AGENTS in our sense.¹¹

It is important to note that neither of these accounts presents a positive proposal about the descriptive question with which we started – what features trigger our attributions of conscious states? This is not surprising given that a central purpose of these accounts was to show the independence of *experience* and *agency*. But it does mean that the accounts do not yet give a full answer to the question, ‘why do we believe that others have conscious mental states?’. And of course that’s exactly what the AGENCY model attempts to do. In what follows we will explain the model and its empirical commitments in more detail.

4. The AGENCY Model

We suspect that a more fundamental notion of ‘agency’ does provide a sufficient basis for attributing conscious states. We have focused on the relatively simple features that serve as cues, which trigger the AGENT concept. But of course these sorts of cues are not the *only* way

¹¹ In an unpublished paper, Bryce Huebner (forthcoming) has suggested that folk psychology functions differently in attributions of beliefs (Agency, loosely speaking) than in attributions of pains and emotions (Experience). Huebner found that people attribute beliefs but not conscious states to robots. Furthermore, other research has found that people are reluctant to attribute pain and emotion to robots even when the robots exhibited appropriate behavior (Sytsma & Machery forthcoming b). Once again, this suggests that there is an attribution of AGENCY with no concomitant inclination to attribute conscious states. But alternative explanations are available. One is that, as we suggested for Knobe & Prinz, the AGENCY category is not really triggered by the robots. However, given that Huebner’s stimuli included a photo of a robot that possessed one of the relatively simple features that serve as cues mentioned above (eyes), this is not a satisfactory explanation. Another possible explanation is that the AGENT category *is* triggered and *does* generate the disposition to attribute conscious states, but that disposition is subsequently suppressed. Perhaps the disposition is suppressed by some generally-held social schema about robots and sensations/emotions; or perhaps it is suppressed by higher-order deliberations about the robot’s physiological makeup. Whatever the reason for suppressing the disposition, these results are only problematic for our theory if it can be shown that subjects were not initially disposed to attribute phenomenology to the robots.

to activate the AGENT concept. For example, merely thinking about an acquaintance (who is not physically present) may trigger the AGENT concept. Relatedly, if a trusted source of testimony *tells* me that there is a person in the other room, then this may typically lead me to think that there is an AGENT in the other room. For our purposes, the role of the relatively simple features that serve as cues is especially important, since we are particularly interested in identifying a set of minimal sufficient conditions for attributing conscious states. We leave it as an open question whether testimony, thought, and related processes like inferential reasoning are able to trigger the AGENT category *directly* (without mediation by simple featural cues), or only *indirectly* by presenting the relevant featural cues in mental imagery, by conceptual or semantic association to the relevant features, and so forth. The key idea is that relatively simple features (e.g., motion pattern, facial features, contingent interaction) will suffice to trigger the attribution of AGENCY, and this will in turn produce an inclination to attribute conscious states to the individual. Hence, the core of the AGENCY model is a causal sufficiency hypothesis:

Hypothesis 1. Typically, if an entity is categorized as an AGENT, then there will be an inclination for attributing conscious states to that entity.

We will take this claim to be the central claim of the *basic* AGENCY model. But a more ambitious model would include the additional claim that categorizing an object as an AGENT is causally *necessary* for a certain sort of inclination to attribute conscious states:

Hypothesis 2. Typically, there will be quick, automatic inclination for attributing conscious states to an entity *only if* that entity is categorized as an AGENT.

We'll call the combination of these hypotheses the *radical* AGENCY model.

While the radical Agency model claims that categorizing an entity as an AGENT is typically both sufficient and necessary for being inclined to attribute conscious states to it, the

model also allows for judgments about consciousness that do not involve these quick, automatic inclinations. The model aligns with various dual-process pictures, according to which, in addition to quick, automatic, low-level cognitive processes there are slow, deliberative, controlled, high-level reasoning processes capable of operating on the same domain (see, e.g., Chaiken & Trope 1999, Sloman 1996, Stanovich & West 2002). The theses laid out above obtain only for the low-level process of consciousness attribution.¹²

Hypothesis 1 suggests that if a person categorizes an object as an AGENT, then she will be disposed to attribute conscious states to that object. This produces the interesting prediction that even relatively simple features will generate an inclination to attribute conscious states to an object. Hypothesis 2 predicts that if a person rejects the categorization of AGENT for a given object, then she will typically not have the automatic inclination to attribute conscious states to the object.

Insects provide an intriguing real-world test case here. Insects are widely regarded among neuroscientists as incapable of experiences like pain, since they lack the relevant neural structures. Nonetheless, insects exhibit all of the simple features reviewed above – eyes, distinctive motion trajectories, and contingent interaction—that serve as cues for categorizing a thing as an AGENT. As a result, the AGENCY model predicts that insects should be categorized as AGENTS and, as a result, people should have a greater tendency to attribute conscious states to insects than to objects like clouds and rivers that presumably do not get categorized as AGENTS.

¹² Alternatively, high-level reasoning can independently produce the judgment that an entity is or is not capable of having conscious experiences. One possibility is that such high-level reasonings nonetheless involve a categorization of the entity as an AGENT. But another possibility is that higher-level judgments provide an entirely independent path to attributions of consciousness. The AGENCY model we offer here remains neutral on these questions.

5. Experiment

In order to test the AGENCY Model, we ran a reaction-time study in which subjects were presented with a sequence of Object/Attribution pairs. In the present study, we were interested to learn about the processes underlying folk attributions of states that psychologists and philosophers typically consider conscious or phenomenal, rather than the folk's metacognitive categorizations of certain states as 'conscious' or 'phenomenal'. As such, the attributions we examined specified states or properties that we take to be paradigmatic of phenomenally conscious experience, such as 'Feels Happy' and 'Feels Pain'.¹³ Objects were drawn from categories such as Vehicles, Insects, and Plants.

Because the model we are proposing describes a low-level cognitive process that, we suggest, reacts automatically to simple features, we wanted to measure the mental chronometry for attributions of phenomenal states to various objects. Thus, for each Object/Attribution pair, subjects were asked to respond as quickly as possible (Yes or No) whether the object has the attribute. The AGENCY model makes predictions both about overt responses and about reaction times. For overt responses, the core of the AGENCY model, hypothesis 1, predicts that subjects should exhibit a tendency to attribute conscious states to objects that are likely to trigger AGENT categorization. Since insects exhibit the simple features that trigger the AGENT concept, and so are likely to be categorized as AGENTS, they are of particular interest. And what will be

¹³ We follow philosophical tradition in taking pain to be a paradigmatic example of a state that has phenomenal consciousness. For an alternative view, see Carruthers 2004. Also, Sytsma & Machery (forthcoming b) maintain that the folk conception of subjective experience is divided into two separate conceptions. Still, they recognize that a state like 'feeling pain' is traditionally taken to be a phenomenally conscious state. Whether the folk conceptualize it as such (either tacitly or explicitly) has no bearing on the discussion at hand (See footnote 1).

especially telling will be to contrast the responses to insects with responses to objects that seem unlikely to trigger the AGENT concept, including individuals in the categories Vehicles, Natural Moving Objects, and Plants. Hypothesis 2 (for the radical AGENCY Model) generates the further prediction that under speeded conditions, participants should be *disinclined* to make overt attributions of conscious states to anything that does *not* get categorized as AGENT.

In a reaction time paradigm, faster response times suggest that the responses are dictated (largely, if not exclusively) by the low-level, automatic process, while longer reaction times suggest that the responses are influenced by higher-level considerations, deliberations, associations, etc. Given this (standard) interpretation of response times, hypothesis 1 predicts slower reaction times when participants *deny* conscious state attributions to objects that are typically classified as AGENTS (as compared broadly to non-AGENTS). The idea is that even if someone were to overtly express the belief that insects don't feel pain (e.g. because they lack appropriate neural structures), she would still have an automatic inclination to think that they do feel pain; she thus has to overcome that automatic inclination to get out her answer of 'no', and this will take some extra time.¹⁴ Or, to put it more colorfully, when I am asked whether ants feel pain, there is a little guy inside of me saying 'yes!' And I need to repress that guy before I can get out my 'no' response.¹⁵ By contrast, there's no little guy inside of me saying that trucks or

¹⁴ A bit more fully, the prediction goes as follows. The presence of the cues biases the subject toward categorizing the individual as an AGENT with the consequent inclination to attribute conscious states to the individual; but competing processes defy these attributions. This creates an uncertainty that takes time to resolve, driving up RT times as a result. Since the whole issue here is about processing time, the way to investigate these matters is by testing performance under speeded conditions.

¹⁵ This tendency may be so strong as to persist even after extensive instruction in the domain. Goldberg & Thompson-Schill (forthcoming) found that even biology professors made predictable mistakes and slower reaction times on a speeded classification task of living/nonliving. Likewise, although neuroscientists might deny that insects feel pain, our bet is that they

rivers feel pain, so my denial in those cases should not be delayed in this way. Finally, hypothesis 2 predicts that, for *anything* that is not counted as an AGENT, there will be no inclination to attribute conscious states, so there should be no hesitation to respond ‘no’ for anything that does not trigger the AGENT concept.¹⁶

Method

Participants

Thirty-four participants (14 female, mean age = 19.2) from the University of Pittsburgh volunteered for this study to fulfill course requirements.

Materials

Subjects performed a timed property-attribution task, in which they were asked to respond positively or negatively to a series of questions attributing different sorts of properties to different sorts of entities. Properties included ‘Feels Anger’, ‘Feels Happy’, ‘Hunts’, ‘Made of Metal’, ‘Feels Pain’, ‘Feels Pride’, ‘Is A Living Thing’, and ‘Is Colored White’. Entities included fifteen word items for each of eight categories: Mammals, Birds, Insects, Plants, Artifacts, Vehicles, Inanimate Natural Objects (e.g., stone, mountain), and Moving Natural Objects (e.g., cloud, blizzard). All category items were matched for letter length, number of syllables, and lexical familiarity.

too would be significantly slower in those denials and be more likely to make mistakes (by their own lights) in speeded conditions.

¹⁶ The differences between the RT predictions flowing from hypotheses 1 and 2 is that hypothesis 1 only says that we should expect faster denials (of conscious states) for at least some individuals that don’t get categorized as AGENTs. Hypothesis 2 makes the stronger claim that we should typically expect fast denials for *any* individual that doesn’t get categorized as AGENT.

Procedure

The experiment used a within-subjects design. All stimuli were presented using E-Prime software (Psychology Software Tools, Pittsburgh PA), which measured participants' response time for each trial. Participants were required to make a speeded response within two seconds before the next item was automatically presented and with a response the next item was presented. For each category (e.g., Insect, Plant), all subjects responded to 120 stimulus items (8 property attributions [e.g., feels pain, feels pride] x 15 entities [e.g., for Insect: bee, wasp]). Since our primary interest was in the attribution of simple conscious states, we collapsed responses to 'feels anger', 'feels happy' and 'feels pain' for our analyses. We preplanned comparisons for overt responses to attributions of these simple conscious states between the Insect category and the categories of Vehicles, Moving Natural Objects, and Plants. We also preplanned comparisons of reaction times of negative responses to simple conscious states between the category of insect and the categories of Vehicle, Moving Natural Objects, and Plants.

Results

Participants were significantly more likely (all p 's $<.001$) to attribute simple conscious states (pain, happy, anger) to insects (70% of trials) than to Plants (10%) or items that exhibit motion, including Vehicles (7%) and Natural Moving Objects (6 %). (See chart 1). Participants were also significantly *slower* to reject the attributions of simple conscious states to insects (670 ms) than to natural moving objects (610 ms) ($t(33) = 2.17, p <.05$) or vehicles (616 ms) ($t(33) = 2.39, p <.05$). No significant difference in response times were found between insects and plants (651

ms).

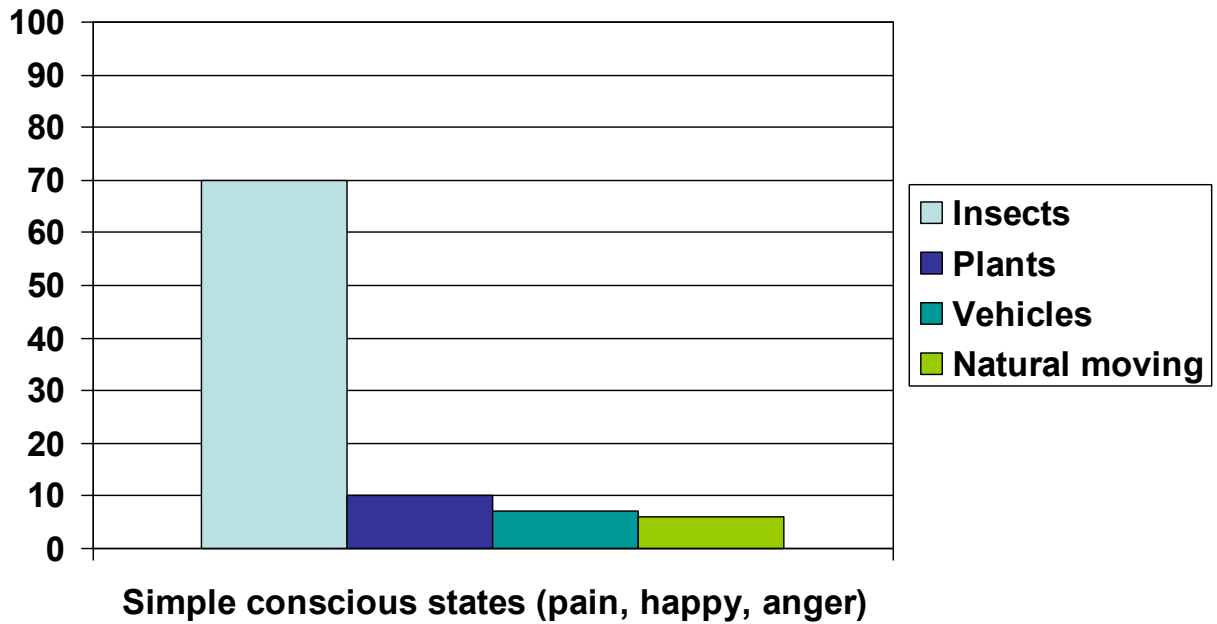


Chart 1

In light of the failure to find a significant difference in RTs between insects and plants, we conducted additional analyses comparing Plants to Vehicles and Natural Moving Objects. To our surprise, we found that participants were also more likely to ascribe simple conscious states to Plants than to Vehicles ($t(33) = 3.28, p < .01$) and Natural Moving Objects ($t(33) = 4.79, p < .001$). Participants also demonstrated significantly slower RTs for denying simple conscious states to plants (651 ms) as compared to natural moving objects ($t(33) = 4.26, p < .001$) and to vehicles ($t(33) = 4.73, p < .001$). (See chart 2).

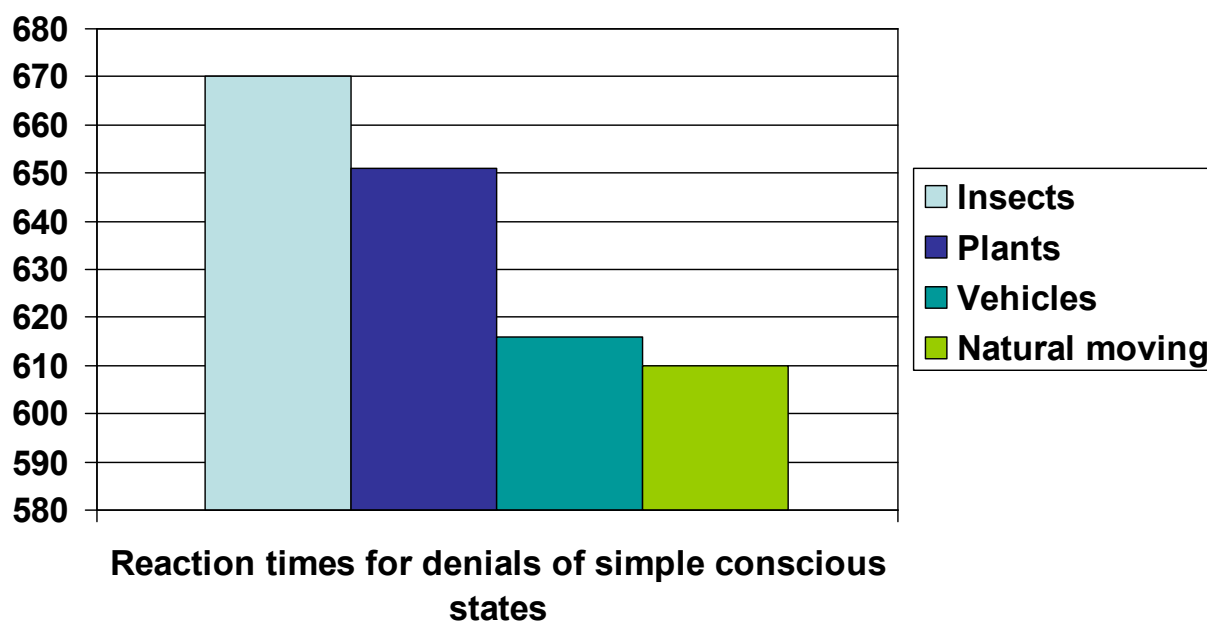


Chart 2

To verify the generalizability of these effects across the items selected, item analyses (averaging across subjects, treating each item as a random variable in the analysis) replicated each of these effects (p 's < .001). In addition, no significant differences in response time were found between insects, plants, vehicles, and natural moving objects when participants denied a superficial property (e.g., colored white).

6. Discussion

The experiment supports the central hypothesis of the *basic* AGENCY model: that categorizing an entity as AGENT is typically sufficient for generating an inclination to attribute conscious states. The overt responses suggest that being identified as an AGENT generates strong inclination to attribute simple conscious states. In particular, insects, but not vehicles or clouds or rivers, were judged to have conscious states.¹⁷ The reaction time data tell a similar story.

¹⁷ Insects are familiar objects, and while this lends ecological plausibility to such cases, one would also like to have evidence

Again, our prediction was that there would be a tendency to attribute simple conscious states like pain to insects, so denying that insects feel pain would require one to override an initial natural inclination.¹⁸ As predicted, we found that denials of simple conscious states to insects were indeed much slower than denials of such states to vehicles and natural moving objects.¹⁹

Hypothesis 2, that AGENCY is necessary for the inclination to attribute conscious states, also looks to get some support from the overt responses. Overall, items that lacked AGENCY cues were unlikely to be afforded the simple conscious states. Insects were overwhelmingly more likely to be granted simple conscious states than were the items that did not have paradigmatic AGENCY features. But the situation with RTs is more complex. Participants showed slower response times when denying simple conscious states to plants as compared with similar responses to vehicles and natural moving objects, and this seems to run contrary to the *radical*

using novel objects. That, of course, is a central virtue of the work by Heider & Simmel and Johnson. To further explore the AGENCY model, it will be important to follow up that work by looking directly at attributions of conscious states. So, for instance, in the standard method developed by Johnson, infants will come to regard the blob (a very un-AGENT like entity) as an AGENT if it beeps contingently with a person's speech. This raises an obvious question about attribution of conscious states. Would infants in the contingent interaction condition be more likely to regard the blob as suffering pain if it is visibly damaged (as compared to a non-contingently interacting condition)? A variety of behavioral measures would be possible, like monitoring for facial distress or concern if the blob is, say, poked with a pin. In addition, people show physiological reactions to suffering in others, and it would be good to test for such physiological responses as well. If infants do show such behavioral indicators of pain attribution, then this would provide additional support for the idea that identifying something as an AGENT suffices to incline people to attribute conscious states to it.

¹⁸ There is a further question about why we attribute the particular states that we do (in this case, pain, anger, and happiness). Then there is the broader question of exactly what's included within the range of states that we are sometimes willing to attribute to AGENTS. We leave this as an important area for future research.

¹⁹ Because participants only had two seconds to respond to each item, it's unlikely that differences were caused by extended deliberations on each item. Rather, these response time differences are likely to reflect more immediate decision processes as shown in seminal and contemporary accounts of semantic processing (e.g. McCloskey & Glucksberg 1979; Smith et al. 1974).

Agency model's prediction that *only* items categorized as AGENTS will generate the automatic inclination to attribute conscious states to the item.

As noted, the RT evidence on plants contravened our predictions, and this demands further consideration of how to model the role of AGENCY in the attribution of conscious states. There are a number of different models that can accommodate the data. We will not try to be exhaustive. Instead, we will review three models that we find especially interesting. First, though, we would like to note that in this experiment we also collected responses for each item on whether or not the object 'is a living thing'. In contrast to vehicles and natural moving objects, participants overwhelmingly judged insects and plants to be living things. So, throughout our discussion we will rely on the idea that plants (and insects) are coded as ALIVE.²⁰ The relevance of this fact will become more salient in the following discussion of the three explanatory models.

One possibility is to retain the *radical* AGENCY model and explain away the unexpected reaction times for plants. For instance, since many things coded as ALIVE are also AGENTS, ALIVE might have high cue validity – it might provide a good heuristic for thinking that an object has mental states. That is, if something is coded as ALIVE then that's a good cue that it's an AGENT, even though it remains possible for something to be coded as ALIVE but not as AGENT.²¹ As a result, it takes some time to process the fact that plants fall into the class of

²⁰ Still, as noted above for the AGENT concept, we want to allow for the possibility that the ALIVE concept is not a perfect match for any word in English.

²¹ As we discuss in the following section, there is some evidence in the developmental psychology literature for thinking that categorizing plants as AGENTS is, in fact, a prerequisite for categorizing plants as ALIVE. That is, rather than merely seeing ALIVE as a highly-valid cue for AGENCY, we might see ALIVE categorizations as being routed through AGENT categorizations.

things that are ALIVE but not AGENTS (this proposal is illustrated in figure 3). The key point of this model is that people do not actually have a disposition to go all the way with attribution of conscious states to plants. Rather, they just need to do a bit more processing in order to exclude plants than they do for vehicles. Note that this would also explain why there were comparatively few overt responses in favor of plants having simple conscious states.

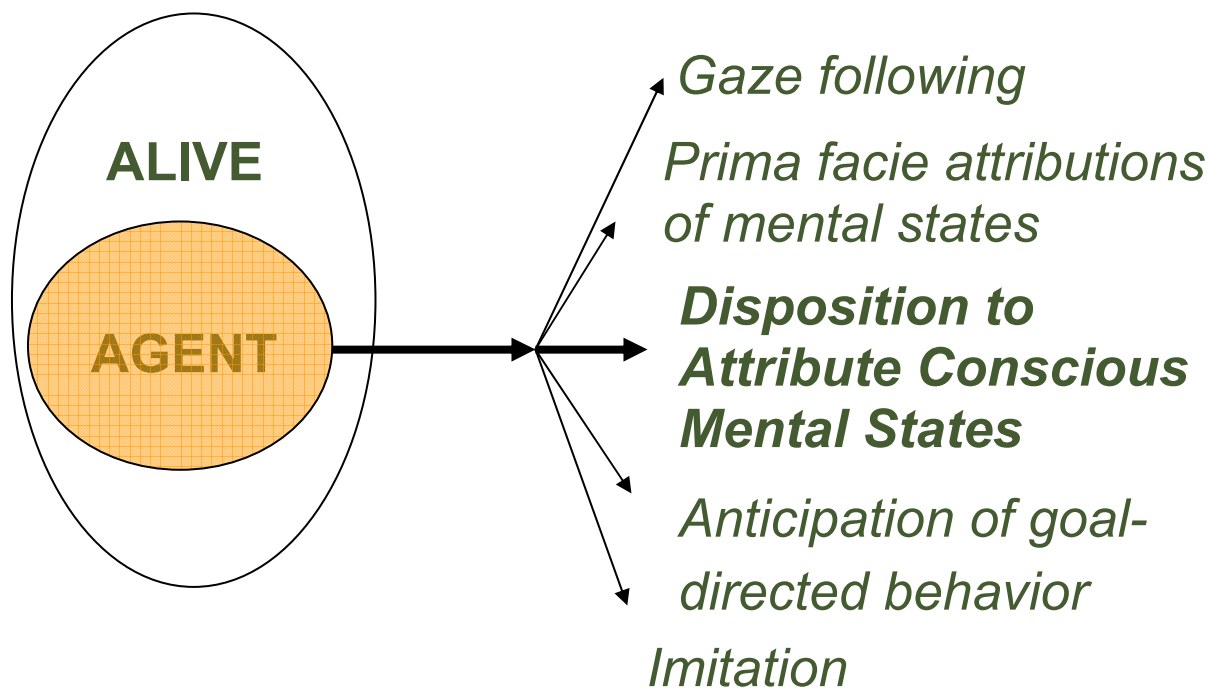


Figure 3: The Agency Model, with ALIVE as cue

A second possible model holds that categorizing an object as ALIVE is actually a (previously unnoticed) cue, which then triggers the individual to categorize that object as AGENT. That is, perhaps activating the notion that a thing is ALIVE would also activate the notion of AGENT.²² We'll call this the LIFE-to-AGENCY model (see figure 4). On such an

²² It is also, of course, possible that categorizing an individual as an AGENT would bring with it the categorization of the individual as ALIVE. For instance, it might be that once the child identifies Johnson's blob as an AGENT, the child would also think of the blob as ALIVE. We set this issue aside for present purposes.

account, the presumed tendency to attribute conscious states to plants is mediated by the AGENT attribution.




Figure 4: Life-to-Agency Model

A third and much more dramatic possibility is that AGENCY is not the important category at all. Perhaps what really matters for attributing conscious states is categorizing an object as ALIVE (as depicted in figure 5). This explanatory model provides a new way to defend the view that agency and experience are independent dimensions. An object, according to this explanation, can be regarded as something that is capable of having experiences even if it is not regarded as an AGENT. This approach would suggest that the deep conceptual link is not between AGENCY and consciousness, but rather between LIFE and consciousness. As such, we call this the LIFE model.



Figure 5: Life Model

7. Developmental questions

The evidence from our experiment does not decide between the three models we've just sketched. Nonetheless, we can adjudicate between them by considering data from related areas of inquiry. One critical source of evidence for deciding between the models is developmental literature. We know that in our experiment, adults tended to overtly deny that plants have conscious states. But given their reaction times, the question is whether there is an automatic inclination in favor of attributing conscious states to plants which then gets suppressed. If there is an automatic inclination to make such attributions, then we might find fully overt attributions of conscious states to plants if we look earlier in development, before children have acquired whatever it is that suppresses overt attributions in adults.

Both the LIFE-to-AGENCY and the LIFE models share the idea that attribution of LIFE is sufficient for the inclination to attribute conscious states. Such hypotheses might get confirmation from developmental evidence of the attribution of conscious mental states to plants. Indeed, in a pair of classic studies, Inagaki & Hatano (1987, 1991) found that Japanese children attributed conscious states to tulips. Strikingly, they found that most children (72%) in

their study said that tulips feel pain, but most (80%) deny that stones feel pain. The children also overwhelmingly (98%) denied that tulips *think* (1987, 1016). This is a significant data point in favor of the idea that attributions of LIFE are sufficient to generate the tendency to attribute simple conscious states. However, one (deflationary) explanation for the Inagaki & Hatano results is that the results are culturally specific to Japanese children.

Recent work by Jane Erickson & Frank Keil (2008), however, suggests that the attribution of mental states to plants is not isolated to Japanese children. Erickson & Keil offered American children different kinds of explanations for a variety of entities, including humans, plants, and non-living natural kinds. In one experiment, the child's task was to evaluate the quality of a psychological explanation for a behavior. When the behavior was produced by a plant – namely, growing toward the sun – kindergarten children regarded a psychological explanation for the behavior as a good explanation. In the case of the plant growing toward the sun, the experimenter offered the following psychological explanation: 'The flower grew towards the sun because it felt tired and hungry and knew that the sunlight would give it energy to make food.' Although the children regarded such a psychological explanation as apt for plants, they rejected parallel psychological explanations for artifacts and non-living natural kinds. In addition, Erickson & Keil found a developmental progression such that older children were less likely to allow psychological explanations for plants (though they continued to regard plants as exhibiting goal-directed behavior). All of this seems to fit well with the models on which an attribution of LIFE suffices to generate an inclination to attribute conscious states. That is, the data seem to suggest that attributions of LIFE are sufficient to trigger inclinations to attribute conscious states, which would show that the radical AGENCY model is mistaken. Further, the data are consistent with the even more dramatic possibility that the *basic* AGENCY

model, too, should be rejected in favor of the LIFE model. For the data leaves open the possibility that LIFE is the operative categorization behind dispositions to attribute conscious mental states.

Thus, combined with our own evidence with respect to reaction times for plants, one might construe the evidence from Inagaki & Hatano and Erickson & Keil as thoroughly undercutting the AGENCY models. After all, one might take these studies to show that children are willing to attribute psychological states to plants based solely on the fact that plants are alive; and if LIFE suffices to generate the inclination to attribute conscious states to an individual, then there seems to be no need to appeal to AGENCY to explain the inclination to attribute conscious states to people and insects. However, a rather different interpretation of the results is suggested by some earlier work by John Opfer and colleagues.

Opfer (2002) showed subjects—both children and adults—a series of videos in which irregularly-shaped blobs moved along various paths. Among other things, the task required subjects to determine what the object was. In one condition, the videos included goal objects that corresponded to the blob's movement, while in the other condition, the videos did not include goal objects; they then presented subjects with a battery of attributions that included both psychological traits and biological traits. According to these studies, attributions of LIFE themselves depend on the individual seeing the entity as a goal-directed agent. Even children as young as 5 years old, they say, relied on the blob's motion being goal-directed in order to attribute to it either biological states (like being alive) or mental states (like being happy or feeling pain). Opfer observes, 'the goal-directedness appeared to be the decisive factor in these displays: identical blobs that moved identically but towards no goal failed to convince children and adults that they were living things.'(116) In a separate study, Opfer & Siegel (2004) found

that 5-year-olds did not initially take plants to behave in goal-directed manners. Yet when told that plants move towards sunlight and water, 5-year-olds inferred that these motions were purposeful, goal-directed. Importantly, Opfer & Siegel found that ‘once children conclude that plants can act in goal-directed, self-sustaining ways, they also conclude that plants are alive.’(329) Thus, though one might initially be inclined to interpret Inagaki & Hatano and Erickson & Keil as supporting the LIFE and LIFE-to-AGENCY models, this interpretation seems less plausible in light of the fact that the attribution of life itself seems to depend on the recognition of an entity *as a goal-directed agent*.

Moreover, considering the work by Opfer and colleagues, it now seems natural to interpret the Erickson & Keil data as providing some support for the *basic* AGENCY model. Given that Opfer and colleagues find children's attributions of life to depend upon seeing the entity as the sort of thing that acts towards goals, and given that Erickson & Keil's attributions included language that seems to activate a goal-directed understanding of plants, we might see the Erickson & Keil data as a confirmation of an implicit prediction of the *basic* AGENCY model. According to the AGENCY model, AGENTS are, inter alia, the sorts of things that we are inclined to attribute goals to. That children approve of psychological explanations of goal-directed behavior attributed to plants just seems to show that, once children see plants as the sorts of entities that have goals and behave according to those goals, they are inclined to say that plants also can have psychological states. In other words, Erickson & Keil's data seem to provide evidence that, once they are primed to categorize plants as AGENTS, children are inclined to attribute psychological states to plants. And that just is what the *basic* AGENCY

model predicts.²³

8. Philosophical Implications

Our proposed account has some relatively clear implications regarding the descriptive problem of other conscious minds. The picture on offer suggests that our cognitive systems are differentially sensitive to very specific features – eyes, distinctive motion trajectories and contingent interaction – the detection of which is normally causally sufficient to bias the subject toward attribution of conscious states. Importantly, the cognitive processes that operate on these low-level cues appear to operate even when global considerations militate against the idea that the target has mental states. For example, my belief that a triangle from Heider and Simmel’s animation is not a mentalistic agent does not stop me from seeing the triangle as ‘wanting to get out of the box’. Similarly, my belief that insects are not conscious does not prevent me from having the inclination to attribute pain and anger to them. The simple features suffice to trigger the inclination to attribute conscious states. It is natural to think of the relevant cognitive process as exhibiting two central features of modularity: domain specificity and informational encapsulation (Fodor 1983). The process is domain specific to the extent that a limited number of simple features are able to trigger that process. In the normal case, things like clouds, blizzards, cars and mountains do not exhibit the right sorts of features, and hence fail to trigger the process.²⁴ Since the relevant process fails to incorporate important information from the subject’s

²³ If categorizing plants as ALIVE requires categorizing them as AGENTS, then the *radical* Agency Model is also preserved: the delayed responses for plants are no longer an obvious counter-example, since the automatic inclination to attribute simple conscious states (which is eventually overridden) can plausibly be explained by attributions of LIFE tacitly depending upon attributions of AGENT.

²⁴ Of course, such objects may *appear* to exhibit some of these features under abnormal circumstances.

broader set of beliefs (for example, the belief that that triangle is *literally* a two-dimensional figure), the process also seems, to that extent, informationally encapsulated. These points apply equally to the various models under consideration (whether they give a central role to AGENCY, LIFE, or both), for all models give special weight to a small set of relatively simple featural cues.

Consider again the two traditional answers to the descriptive problem of other conscious minds: analogy and inference to the best explanation. Both analogical reasoning and inference to the best explanation are thought to be paradigm cases of non-modular processing (Fodor 1983, 88-89, 107). First, neither process is domain specific: analogies can in principle relate anything to anything, and explanations can in principle target arbitrary explanatory domains. And second, neither analogical reasoning nor inference-to-the-best-explanation is informationally encapsulated; both are highly sensitive to information from all over the cognitive system (Fodor 1983, 104-107). If this assessment of the processes is correct, then there is a tension between the traditional answers to the descriptive problem of other conscious minds and the view that we are promoting. For unlike traditional accounts, on our account it is natural to view some of the mental processes responsible for the attribution of conscious states as modular.

The modularity hypothesis provides a familiar framework for articulating the difference between the AGENCY model and the traditional analogy and best-explanation models. But the important differences can still be drawn without adopting the modularity framework. The important point is that on the AGENCY model, the way we attribute consciousness to other things often involves simple cues (like contingent behavior) that quickly, automatically, and effortlessly trigger the AGENT concept, which in turn leads directly to an inclination to attribute conscious states. These operations differ radically from traditional intellectualist processes like analogical reasoning and inference to the best explanation. This point again resonates with

various dual-process models that describe two distinct pathways for reaching the attributions.

Kahneman provides a nice summary of dual-process architecture:

The operations of System 1 are typically fast, automatic, effortless, associative, implicit (not available to introspection), and often emotionally charged; they are also governed by habit and are therefore difficult to control or modify. The operations of System 2 are slower, serial, effortful, more likely to be consciously monitored and deliberately controlled (Kahneman 2003, 698).

Regarding attributions of consciousness, we are essentially postulating *both* a System 1 process or ‘low road’ involving automatic processing of relatively simple featural cues, and a System 2 process or ‘high road’ involving relatively slow, controlled and deliberative reasoning.²⁵

In addition to the implications for the traditional descriptive problem of other minds, the models we’ve sketched here might also have import for the status of particular intuitions about the consciousness (or not) of other entities. A number of prominent philosophers have built explicit theories of mind partly on the basis of our intuitions about what is conscious and what is not. The account we have offered might be used to explain intuitions associated with some famous thought experiments deployed in philosophy of mind. One such case is Ned Block’s famous example in which we are to imagine that all the residents of China are rigged up with radio transmitters so as to functionally mimic a living brain (Block 1980). The intuitive reaction to this case is presumably that the nation of China (as a whole) does not have any mental states, and surely not any conscious states! This is then supposed to support the conclusion that merely getting a system to functionally mimic a brain does not make that system conscious. For

²⁵ Another possible model that does not invoke modularity involves a fairly radical conception of analogical reasoning and inference to the best explanation, on which these kinds of inferences can take place at a ‘low’ level (i.e. exhibit speed and automaticity), and are sensitive to simple feature cues.

intuitively, it seems clear that functionally organizing the residents of China to behave like a brain would not make the collective nation have conscious states.

Notice, however, that if our proposal is correct, there is a potential explanation for these intuitions that does not involve the denial that the nation of China enjoys conscious states. Instead, it may be that the example tends to provoke these intuitions because the sorts of featural cues that typically incline a subject toward attributions of consciousness are not salient with respect to the nation of China.²⁶ Each Chinese person, on the other hand, does possess the relevant cues (which are readily imaginable when considering Block's scenario). Hence, we have an inclination to attribute conscious states to individual Chinese people, but not to the nation of China. This is just one example, but it illustrates how the answer to the descriptive problem might influence the way the relevant intuitions are used in philosophical debate. For depending on what one thinks about the epistemic status of the relevant psychological processes, one might be led either to dismiss the intuitions, or to give them special weight.

An answer to the descriptive problem might also bear on the idea that people intuitively embrace a 'folk dualism' (Bloom 2004), according to which the mind is radically different than the body.²⁷ It's plausible that one aspect of such a dualism is the apparent gulf between consciousness and physical objects. For instance, when we think about a brain as a massive collection of neurons that has various chemical and physical characteristics, it is not at all intuitive that this mass has consciousness. Something similar can, of course, be said for other bodily organs. Even after we are told that the brain is the part of the body responsible for

²⁶ This is not to say that the *only* path to attributions of consciousness is via the simple cues. Rather, the point is that if an entity fails to manifest the cues that trigger the AGENCY category, then it will be significantly less natural to attribute conscious states to that entity. Instead, attributions of consciousness in such cases are likely to be the result of deliberate high-level reasoning.

²⁷ See Chalmers (2002) for a philosophical review of the gulf between consciousness and physical objects.

consciousness, this does not render it *intuitive* that the brain is where conscious experience occurs. We suggest that part of the reason for this is that when we consider brains as hunks of physical stuff, we are considering descriptions that exclude the sort of cues that tend to activate the low-level processes that generate the intuitive sense that an entity is conscious.²⁸ Hence, it's not surprising if we find some initial resistance to the idea that the physical brain is conscious. In this light, it's somewhat ironic that, while people have difficulty thinking of the brain as conscious, they have no trouble at all thinking that *ants* are conscious. On the contrary, our experiments indicate that people have trouble thinking that ants are *not* conscious.

†*Department of Philosophy*

University of Arizona

‡*Center for Cognitive Neuroscience*

University of Pennsylvania

²⁸ Though, as noted above, there might be other pathways that generate the judgment that an entity has conscious states.

References:

- Abell, F., Happé, F. and Frith, U. 2000: Do triangles play tricks? attribution of mental states to animated shapes in normal and abnormal development. *Journal of Cognitive Development*, 15, 1-20.
- Arico, A.J. Forthcoming: Folk psychology, consciousness, and context effects.
- Block, N. 1978: Troubles with functionalism. In W. Savage (ed.), *Perception and Cognition: Issues in the Foundations of Psychology*. Minneapolis: University of Minnesota Press.
- Bloom, P. 2004: *Descartes' Baby*. New York, NY: Basic Books.
- Bowler, D. and Thommen, E. 2000: Attribution of mechanical and social causality to animated displays by children with autism. *Autism*, 4, 147–71.
- Carruthers, P. 2004: Suffering without subjectivity. *Philosophical Studies*, 121, 99-125.
- Chaiken, S. and Trope, Y. (eds). 1999: *Dual Process Theories in Social Psychology*. New York, NY: Guilford.
- Chisholm, R. 1977: *Theory of Knowledge*. 2nd Edition. Englewood Cliffs: Prentice-Hall.
- Cohen, S. 1984: Justification and truth. *Philosophical Studies* 46, 279-295.
- Dasser, V., Ulbaek, I. and Premack, D. 1989: The perception of intention. *Science*, 243, 365-7.
- Descartes, R. 1641/1986: *Meditations on First Philosophy*. Translated by John Cottingham. Cambridge: Cambridge University Press.
- Erickson, J. and Keil, F. 2008: The role of intention in teleological explanations. Poster at 2008 meeting of Society for Philosophy and Psychology.
- Fodor, J. 1983: *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Goldberg, R. and Thompson-Schill, S. Forthcoming: Developmental 'roots' in mature biological knowledge. *Psychological Science*.

- Goldman, A. 1979: What is justified belief? In G. Pappas (ed) *Justification and Knowledge*.
Cambridge, MA: MIT Press.
- Goldman A. 2006: *Simulating Minds*. Oxford: Oxford University Press.
- Gopnik, A. and Meltzoff, A. 1997: *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.
- Gray, H., Gray, K. and Wegner, D. 2007: Dimensions of mind perception. *Science*, 315, 619.
- Harris, P. 1989: *Children and Emotion*. Oxford: Blackwell.
- Heider, F. and Simmel, M. 1944: An experimental study of apparent behavior. *American Journal of Psychology*, 57, 243-259.
- Huebner, B. Forthcoming: Commonsense concepts of phenomenal consciousness: Does anyone care about functional zombies?
- Hyslop, A. 2005: Other minds. *The Stanford Encyclopedia of Philosophy (Winter 2005 Edition)*, ed. E. Zalta. From <http://plato.stanford.edu/archives/win2005/entries/other-minds/>.
- Inagaki, K. and Hatano, G. 1987: Young children's spontaneous personification as analogy. *Child Development*, 58, 1013-1020.
- Inagaki, K. and Hatano, G. 1991: Constrained person analogy in young children's biological inference. *Cognitive development*, 6, 219-31.
- Johnson, S., Booth, A. and O'Hearn, K. 2001: Inferring the goals of non-human agents. *Cognitive Development*, 16, 637-656.
- Johnson, S., Slaughter, V. and Carey, S. 1998: Whose gaze will infants follow? features that elicit gaze-following in 12-month-olds. *Developmental Science*, 1, 233-238.
- Kahneman, D. 2003: A perspective on judgment and choice: mapping bounded rationality. *American Psychologist*, 58, 697-720.
- Knobe, J. and Prinz, J. 2008: Intuitions about consciousness: experimental studies.

- Phenomenology and the Cognitive Sciences*, 7(1), 67-83.
- McCloskey, M. and Glucksberg, S. 1979: Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology*, 11, 1-37.
- Mill, J. 1865: *An Examination of Sir William Hamilton's Philosophy*. London: Longmans.
- Nichols, S. and Stich, S. 2003: *Mindreading*. Oxford: Oxford University Press.
- Perner, J. 1991: *Understanding the Representational Mind*. Cambridge, MA: Bradford Books/MIT Press.
- Reid, T. 1969/1785: *Essays on the Intellectual Powers of Man*. Cambridge, MA: MIT Press.
- Robbins, P. and Jack, A. 2006: The phenomenal stance. *Philosophical Studies*, 127, 1, 59-85.
- Scholl, B. and Tremoulet, P. 2000: Perceptual causality and animacy. *Trends in Cognitive Sciences*, 4, 299-309.
- Searle, J. 1980: Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417-424.
- Shimizu, Y. and Johnson, S. 2004: Infants' attribution of a goal to a morphologically unfamiliar agent. *Developmental Science*, 7, 425-30.
- Slooman, S. 1996: The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.
- Smith, E.E., Shoben, E.J., and Rips, L.J. 1974: Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, 81, 214-241.
- Stanovich, K. and West, R. 2000: Individual differences in reasoning: Implications for the rationality debate. *Behavioural & Brain Sciences*, 23, 645-726.
- Sytsma, J and Machery, E. Forthcoming a: How to study folk intuitions about consciousness. *Philosophical Psychology*.

Sytsma, J. and Machery, E. Forthcoming b: Two conceptions of subjective experience.

Wellman, H. 1990: *The Child's Theory of Mind*. Cambridge, MA: Bradford Books/MIT Press.

Woodward, A. 1998: Infants selectively encode the goal object of an actor's reach. *Cognition*,
69, 1-34.