Word counts:

Abstracts: 101 + 248 = 349

Main text: 11,300

References: 1,495

Total: 13,199

Title:

Comparative Psychology without Consciousness

Author details:

Peter Carruthers

Department of Philosophy

University of Maryland

College Park

MD 20942, USA

pcarruth@umd.edu

http://faculty.philosophy.umd.edu/pcarruthers/

Short abstract:

This paper argues that while our best theory of phenomenal consciousness (namely, global workspace theory) can explain the all-or-nothing character of human conscious experience, nonhuman animals can only instantiate that theory to some degree. Our first-person concept of phenomenal consciousness, however, doesn't allow for degrees. So there is a mismatch between our best theory of human consciousness (which picks out a gradable property) and the way we conceptualize phenomenal consciousness. In response, we should give up asking questions about consciousness in animals altogether. This is not because the questions are too hard, but because there are no substantive facts to discover.

Long abstract:

Evidence continues to accumulate in support of the global workspace theory of consciousness, construed here as a theory of *phenomenal* or *subjectively felt* consciousness. I argue briefly that the global workspace account is preferable to all other competing theories, and that it has the resources to

fully explain the properties of human phenomenal consciousness. At the same time, there has been increasing interest in the question of consciousness in nonhuman animals. One might think that our best theory of human consciousness would help us to resolve that question, but it doesn't. This is because, while global broadcasting is all-or-nothing in the human mind, it is framed in terms that imply gradations across species. Yet our concept of phenomenal consciousness doesn't permit mental states to be *to some degree* conscious. So there is a mismatch between our best theory of consciousness, which identifies it with a gradable property, and our first-person conception of consciousness itself. Four responses are considered: (1) rejecting the proposed identity between phenomenal consciousness and global broadcasting; (2) revising our concept of consciousness to allow for degrees; (3) keeping the concept all-or-nothing, but defining a (vague) categorical boundary around human-like global broadcasting that will definitely include some species while definitely excluding some others; or (4) no longer raising the question of consciousness in nonhuman animals. I argue that #4 is the best option. This is not because the questions are too hard, but because there are no substantive facts about consciousness in nonhuman animals to discover.

Keywords:
Access consciousness, animal consciousness, degrees of consciousness, global workspace, hard problem, phenomenal consciousness

1. Introduction and initial distinctions

What do theories of consciousness suggest or imply about the distribution of consciousness across the rest of the animal kingdom? More specifically, what does our *best* theory of consciousness suggest? I will argue that the correct answer to this latter question presents a dilemma: either there is no fact of the matter, or few — if any — creatures besides humans are conscious. I will argue, moreover, that to the extent that there *is* a fact of the matter, the question is merely classificatory, and of no substantive interest. If I am correct on these points, then all talk of consciousness should be dropped from comparative psychology. There are many substantial facts waiting to be discovered about the minds of nonhuman animals (hereafter, "animals"); but facts about consciousness or its absence aren't among them.

Before proceeding further, however, some points of clarification are necessary. The kind of consciousness that forms our topic is so-called *phenomenal* consciousness. This is the sort of

consciousness that is *like something* to undergo, or that has a distinctive subjective *feel*. Phenomenal consciousness is a species of mental-state consciousness. It is mental states (smelling a rose, hearing a trumpet, or seeing the color of a tulip) that can be phenomenally conscious. *People* are phenomenally conscious derivatively, by virtue of undergoing phenomenally conscious states. It is phenomenal consciousness that is thought to give rise to the "hard problem" of consciousness (Chalmers 1996). For it seems one can conceive of a zombie — a creature that is like oneself in all physical, functional, and representational respects except that it lacks *this* feeling (the distinctive feeling of the smell of the rose one is sniffing). Hence many have been tempted to conclude that phenomenal consciousness involves properties (often called "qualia") that cannot be reduced to any combination of physical, functional, or representational ones.

Phenomenal consciousness is at least conceptually distinct from *access* consciousness (Block 1995, 2007). (Whether there is any real distinction between the properties picked out by these concepts is another matter, to which we return in Section 3.) Both are forms of mental-state consciousness: it is mental states that can have phenomenal properties, or that can be accessible to enter into decision making, reasoning, and verbal report. *Phenomenal* consciousness is a first-personal concept, however. One can only understand what that concept is intended to pick out by directing one's attention to some of one's own phenomenally conscious states. *Access* consciousness, in contrast, is functionally defined, and the concept could be fully understood by a zombie. A mental state is said to be access conscious if it is accessible to a wide range of other systems for further processing, specifically those involved in issuing verbal reports, in decision making, in reasoning, and in the formation of long-term memories. (To this standard list I would add that access-conscious states are involved in causing and modulating full-blown temporally-extended emotions.)

Mental-state consciousness (whether access or phenomenal) should likewise be distinguished from *creature* consciousness, or what Bayne et al. (2016) call "global-state consciousness". Creature-consciousness is a matter of being awake rather than asleep, or conscious as opposed to comatose. When the creature in question is a human person, then creature consciousness would normally implicate some or other form of mental-state consciousness. Whenever one is awake one is normally undergoing some conscious mental state or other. But the reverse need not be true. It seems that dreams are conscious mental states, even though the dreaming subject is asleep, and hence creature *un*conscious. Note that creature consciousness admits of levels, or degrees (albeit not linearly ordered

ones, perhaps; Bayne et al. 2016). Someone in a state of deep sleep is more unconscious than someone who is dreaming. And someone under deep anesthesia will be more unconscious than someone who is lightly sedated. Likewise, someone in a caffeine-induced state of high alertness will be more conscious than someone who is drowsy.

2. The need for theory

Before we can return to the bearing of theories of consciousness on the question of phenomenal consciousness in animals, there is a preliminary question to be addressed. This is whether one *needs* a theory of consciousness in order to say whether or not animals undergo conscious mental states. Tye (2017) argues not. Provided we think conscious states play a causal role in the production of behavior, we can apply Newton's Principle: similar effects have similar causes. So when we see an animal exhibiting behavior similar to that displayed by human beings — and specifically, when we see an animal showing flexible perceptually-grounded sensitivity to facts about its environment or its own body — then we can conclude that it enjoys conscious experience. (At least, this is so unless we know of some additional fact that might undermine — or *defeat* — Newton's Principle in any given case. Much of Tye's book is occupied with discussion of possible defeaters, including the question whether absence of a cerebral cortex is a defeater for conscious experience.)

Unfortunately, this approach is unsound. For it neglects a wide range of environmentally-sensitive forms of behavior in humans that are at least to some degree flexible, but are caused by *un*conscious perceptual states. Consider, for example, a baseball batter facing a fastball traveling at 96 miles per hour. The distance between the pitching plate and the batting plate is 60 feet, hence there are about 56 feet from the point at which the pitcher releases the ball to the point of contact for the batter. The fastball will travel that distance in 400 milliseconds. In that 400-millisecond interval the batter has to estimate the trajectory of the ball (will it pass through the strike-zone or not?), initiate a swing, and then plan and execute the motor instructions for a specific type of swing, aimed at the expected location of the ball. But our best estimate of the time that elapses between presentation of a stimulus (in this case, light hitting the retina from the ball as it exits the pitcher's hand) and conscious experience is somewhere around 350 milliseconds (Dehaene 2014; Marti & Dehaene 2017). It is almost certain, therefore, that the swing is initiated *before* the batter's perception of the trajectory of the ball becomes conscious, and that the execution of the swing is likewise guided by unconscious perceptual states throughout. It seems that while the batter's swing is *accompanied* by conscious experience, it isn't

caused by it. So Newton's Principle cannot apply. This means that when we observe an animal responding to its environment in this sort of swift online manner we cannot use Newton's Principle to infer that it enjoys conscious experience.

These points are vindicated and explained by the discovery of two visual systems in humans (Milner & Goodale 1995). As is now familiar to many people, visual processing proceeds in two distinct cortical streams. (Similar distinctions can be drawn with respect to other sensory systems. For the case of audition, see Rauschecker & Tian 2000; Hickok & Poeppel 2004.) The ventral stream runs forward from visual cortex through the lower portions of the temporal lobes, whereas the dorsal stream runs forward and upward through the parietal lobes. Processing within the ventral stream results in object recognition, with the relative positions of objects represented in allocentric space. Processing in the dorsal stream, in contrast, is coded in limb-centered and body-centered spatial coordinates, and is specialized for online guidance of action. Moreover, the two streams utilize distinct cellular pathways that remain distinct from one another all the way back to the retina. The ventral steam is a continuation of the parvocellular pathway, whereas the dorsal stream is a continuation of the magnocellular pathway. Notably, given the points about speed made in the previous paragraph, transmission of signals through the magnocellular (dorsal) pathway is significantly faster than through the parvocellular one.

Multiple lines of evidence suggest that, while the two visual streams interact with one another in their early stages, their outputs are distinct; and while the outputs of the ventral stream can be conscious, the outputs of the dorsal stream are inaccessible to consciousness. (For a recent review, see Goodale 2014.) For example, damage to the dorsal stream can result in optic ataxia, in which visually guided action becomes halting and inaccurate while conscious experience is fully intact; whereas damage to the ventral stream can result in agnosia (an inability to consciously identify objects), while leaving visually-guided motor control intact.

Consider, in particular, D.F., a patient with bilateral temporal-lobe damage studied extensively by Milner & Goodale (1995). D.F. suffers from complete visual-form agnosia. While she can still experience colors and textures (and hence might be able to guess at the identity of a banana from its distinctive yellow color and mottled texture), she can no longer experience the shapes or orientations of objects. For example, she cannot recognize a banana from a line drawing of one, she is at chance when judging the orientation of a pencil (whether upright or horizontal), and she is at chance when judging whether a

block of wood is square or oblong. But her reaching-and-grasping behavior is indistinguishable from normal. She will orient her hand appropriately when grasping a horizontally-held or vertically-held pencil, using a normal finger grip (but without being able to report the pencil's orientation in advance). And she is just as accurate as neurotypical people when posting a letter through a letter box arranged at various angles (while remaining at chance when consciously judging those angles).

It is also thought to be the dorsal visual stream that underlies the well-known phenomenon of *blindsight* (Weiskrantz 1986), in which people with damage to primary visual cortex become blind (that is to say: incapable of conscious experience) in a portion of their visual field, while being well above chance in identifying simple letters and patterns of movement presented in their blind field. For there are projections from the retina to parietal cortex via the superior colliculus (bypassing primary visual cortex), whereas there are no such subcortical projections direct to temporal cortex. Moreover, we know that the monkey Helen, who had the whole of primary visual cortex surgically removed, could nevertheless move around and pick up objects normally, while being unable to identify objects except by touch or taste (Humphrey 1974). (It may well be that blindsighted people don't generally act like this because they *believe* themselves to be blind. It should be noted, however, that what initially suggested that blindsight in humans might be a possibility was the observation that a blindsight patient could walk across a cluttered room much more easily than one might have expected; see Weiskrantz 1986.)

It appears to be the dorsal stream, too, that underlies habitual sensorimotor action. There is good reason to think that habitual behavior can occur without being initiated or guided by conscious experience (Lisman & Sternberg 2013), and we know that there are projections from the dorsal visual stream to dorsolateral striatum (Saint-Cyr et al. 1990), which is the locus of sensorimotor control in habitual responding (Wood & Rünger 2015). Indeed, as control of behavior shifts during learning from goal-oriented to habitual, one can trace a corresponding shift of control from ventral to dorsolateral striatum (Redgrave et al. 2010; Burton et al. 2015). This comports well with one's everyday experience. One can engage in routine activities (like driving a car) without conscious awareness of the stimuli involved in the activity, and while one's conscious mind is wholly engaged with other matters.

It seems likely that sleepwalking implicates a similar dorsal route. While there are a number of different kinds of motor-involving sleep disorder, sleepwalking is the sort that takes place during the deepest form of sleep, in the absence of any of the rapid eye-movements (REMs) that would normally indicate

the presence of dreaming. Indeed, sleepwalking (and in some documented instances, even sleep *driving*!) takes place with prefrontal cortex and most other regions of cortex fully suppressed (as in non-REM sleep), while there is local activation of sensory cortices, the amygdala, and motor regions of cingulate cortex (Terzaghi et al. 2012; Januszko et al. 2015).

Finally, in this catalog of environmentally-sensitive actions that can take place in the absence of conscious experience, one should include the behavioral components of emotional states such as fear and anger. Although such emotions generally give rise to distinctive forms of conscious experience (at least in humans), the behaviors in question are initiated directly and swiftly and in the absence of conscious experience by subcortical circuits that include especially the amygdala (LeDoux 1996, 2017; Panksepp 1998). It is these circuits that are responsible for the fight-or-flight response, setting in train a variety of physiological changes (increased heart-rate and breathing-rate, and so on), activating emotion-expressive facial expressions (the fear-face, the anger-face, and so on), and initiating adaptive forms of behavior guided by the affordances of the environment (fleeing, freezing, or fighting). Such emotion-expressive actions will automatically run through to completion unless inhibited by executive commands. Emotionally arousing stimuli will attract one's attention, of course; hence one will soon undergo conscious experiences of the relevant events. No doubt such experiences play a role in sustaining and/or modulating one's on-going emotions. But they aren't necessary (indeed, they aren't present) at the outset.

There are thus a wide variety of actions in the human case that display perceptual sensitivity to the environment, that may sometimes be accompanied by conscious experience, but that are actually caused by perceptual states that are deeply *un*conscious. This has implications for any attempt to use Newton's Principle ("similar effects have similar causes") to determine whether animals are phenomenally conscious. For with respect to any given kind of animal behavior, we would then have to ask whether the behavior in question is more similar to human consciously-caused behavior than it is to human *un*consciously-caused behavior. And there is no way to answer this question from the behavior alone. Indeed, in order to provide a systematic, more than merely-intuitive, answer to the question one would need to develop and appeal to a theory of the difference between conscious and unconscious forms of perception. Armed with such a theory, one could then raise (and investigate) the question whether both forms of perception are manifest in animals. This is where we go next.

3. Theories of consciousness

I will focus here on the class of theories that have the resources to explain the distinction between conscious and unconscious mental states (setting aside theories such as panpsychism [Strawson 2006], as well as those that draw on subatomic quantum indeterminacies [Hameroff & Penrose 1996], which seem incapable of doing so). The best-known such theory is the global workspace account put forward by Baars (1988, 1997, 2002), which is a close relative of Tye's (1995, 2000) PANIC theory. (PANIC stands for Poised Abstract Nonconceptual Intentional Content.) The basic idea is that some perceptual and perception-like contents (including visual and auditory imagery) are globally broadcast to a wide range of other systems in the mind/brain, whereas others have more specialized or local uses. Conscious states are those that are broadcast to systems for reporting, planning, reasoning, and remembering, whereas unconscious states are those that are *not* so broadcast, although they can have a variety of other roles or effects.

The basic idea behind global broadcasting theory has been heavily investigated by cognitive scientists in recent decades, especially using minimal contrasts between conscious and unconscious conditions in combination with EEG, MEG, and/or fMRI. (Much of this work is summarized and reviewed in Dehaene 2014.) Minimal contrasts are ones where the stimuli remain fixed but where consciousness is sometimes present, sometimes absent. One such case is binocular rivalry, where distinct images are presented to the two eyes, only one of which is conscious at any one time (Tong et al. 1998). A close relative of binocular rivalry is continuous flash suppression, where an image (of a face, say) can be presented continuously to one eye while the other eye is presented with a flashing Mondrian-type stimulus (Wilke et al. 2002; Tsuchiya & Koch 2005). The latter dominates initially, while the other stimulus emerges into consciousness from the noise. Briefly presented stimuli can also be backward masked to render them unconscious, while nevertheless being processed quite deeply, extracting high-level content that can prime subsequent behavior (Breitmeyer & Ogmen 2000). If the intensity and timing parameters are set correctly, it can be arranged so that a backward-masked stimulus of a given intensity and duration is only visible on fifty percent of the trials. (The so-called *attentional blink* can be arranged similarly; Raymond et al. 1992; Marti et al. 2012.) The general finding in such experiments is that unconscious stimuli give rise to local reverberations in sensory and higher-level association areas of temporal cortex, whereas stimuli that are consciously experienced give rise to much more wide-spread coordinated activity, involving both prefrontal and parietal cortices as well (Dehaene 2014).

Critics have sometimes objected that the minimal-contrast methodology controls for stimulus factors but not for behavioral ones (Frässle et al., 2014; Pitts et al. 2014). For in order to report that a stimulus is visible on a given trial, participants need to say so, or press a button to indicate so, or respond in some other way that differs from how they respond when the stimulus remains unconscious. But a recent study attempted to control for behavioral factors also (Salti et al. 2015). Participants were first trained to use all eight fingers when indicating the spatial position of a stimulus on a clock-face, with a unique finger used for each position. They were then presented with the same stimuli under conditions of backward masking that rendered the stimulus invisible on about fifty per cent of the trials, but they were nevertheless required to guess at a location on *every* trial. Since these "guesses" on invisible trials were accurate at levels well above chance, one can infer that something like blindsight was involved. The investigators then trained pattern classifiers on the resulting EEG and MEG data to determine where in the brain the information about the specific location of the stimulus was represented. The finding was that in consciously-seen trials the spatial content reached much deeper into prefrontal cortex, and was processed much more deeply, than in blindsight cases that shared the same motor response. (Recall that one would expect blindsight information to reach motor cortex, at least, since it can serve to guide simple forms of action.)

Notice that if global workspace theory is correct, then phenomenal consciousness and access consciousness coincide with one another. This is denied by Block (2007, 2011a), who thinks that the contents of phenomenal consciousness are richer than those accessible for planning or reporting. The main evidence provided, is that people claim to see more details in a briefly presented stimulus than they can thereafter report (Sperling 1960; Landman et al. 2003; Sligte et al. 2008). As a result, Block thinks that phenomenal consciousness should be identified with the contents of iconic memory, whereas access consciousness comprises the contents of working memory.

Critics have replied that people's intuition that they see the stimulus in rich detail can be explained away, in part by appealing to background scene statistics (Cohen & Dennett 2011; Cohen et al. 2016). In addition, there is evidence that people can only report about four items even when the cue for reporting coincides with stimulus offset, obviating the need for working memory (Tsubomi et al. 2013). People can be shown an array of colored squares for a full second, whereupon just one of the squares remains illuminated. But that square is now divided into two colored halves, one of which is the original color of the square, the other of which is drawn from elsewhere in the display. The task is just to indicate which

is which. Participants still show the standard four-item limit (two in each hemisphere), which is also present in their neural EEG signatures that flatten off at that four-item limit. These findings seem plainly inconsistent with Block's view. If one's perception of a colored square were conscious, one would think that one would immediately know which side of the square had changed when one half of it is replaced by a differently-colored rectangle.

There is a great deal more that could be said on these topics, of course. But I will be assuming henceforward that phenomenally conscious states are also access consciousness. This doesn't necessarily mean, however, that all access-conscious contents are also phenomenally conscious. For it is arguable that conceptual as well as nonconceptual contents can be globally broadcast (as when one sees a figure *as* a duck or *as* a rabbit, for example). But it can also be argued that globally broadcast concepts don't make a constitutive contribution to the phenomenal properties of the globally broadcast perceptual states in which they participate (Carruthers & Veillet 2011, 2017). This is because only nonconceptual content gives rise to the sorts of "hard problem" thought experiments that typify phenomenal consciousness. In the discussion that follows, however, nothing of significance will turn on this point.

It may be worth comparing global workspace theory with some other nearby theories that are informed by cognitive science and have attracted significant attention. One is Prinz's (2012) AIR theory of consciousness. ("AIR" stand for Attended Intermediate-level Representations.) The claim is that attentional signals directed at nonconceptual representations of color, shape, and so on are both necessary and sufficient for conscious experience. I agree that attention is necessary for consciousness. One can quibble over whether it is fully sufficient (Carruthers 2015). But the larger point is that attention to nonconceptual content is the *cause* of, rather than being *constitutive* of, consciousness. Attention results in the global broadcast of the attended representations. If there were nothing to broadcast *to*, there would be no consciousness. Attention would be like an engine spinning its wheels with nowhere to go.

Another way to approach the same point is to notice that attention can either be individuated in network terms, as the mechanism that normally issues in global broadcasting of attended representations (this is the standard notion employed in cognitive science); or it can be individuated computationally, in terms of its function in boosting the targeted neural activity while suppressing

competing populations (Mole 2011; Wu 2014). Understood in the first way, the AIR theory becomes almost a notational variant of the global workspace account (except for its emphasis on nonconceptual as opposed to higher-level content; as just noted, this is an aspect of the theory I agree with). But understood in the second way, it would plainly be inadequate as a theory of consciousness. For there are multiple boosting-while-suppressing mechanisms in the human brain that remain unconscious. For instance, the mechanism that selects between competing motor plans is thought to operate in just this manner (Wu 2014).

A second proposal worth contrasting with the global workspace account is Tononi's integrated information theory (Tononi 2008; Tononi & Koch 2015). On this view, consciousness can be measured by the extent to which information is integrated via complex reverberating feedback loops. The main problem with the theory is that it is unclear how to turn it into an account of mental-state consciousness, except by piggy-backing off the global workspace view. Since information integration is a graded notion, and is present at all different levels of cognition, the theory seems to imply a sort of panpsychism. For even paradigmatically unconscious blindsight-like states will contain some degree of informational integration. But this is panpsychism that is unmotivated, and does no explanatory work. And the first-personally available difference between conscious and unconscious states will reduce to the distinction between states that are *highly* informationally integrated (because globally broadcast) and those that are much less so (because they have only local or specialized effects). But it is the global broadcasting, not the informational integration, that explains why conscious states are available to be remembered and reported, as well as entering into planning, whereas unconscious states are not.

Finally, let me mention the class of higher-order theories of consciousness. These come in two basic varieties. One claims that it is the actual presence of a higher-order thought about an experience that renders that state phenomenally conscious (Rosenthal 2005). This view faces notorious difficulties. One is that it is hard to know what one should say when the content of experience and the content of the higher-order thought don't align with one another (Mandik 2009; Weisberg 2011). What happens to one's phenomenally conscious experience if one undergoes a perception of red but *believes* that one is perceiving orange? Another difficulty is to explain how higher-order thoughts (which are fully conceptual, note) can give rise to the distinctive fine-grained character of phenomenally conscious experience (Block 2011b). And yet another problem is to explain why we should be entertaining multiple higher-order thoughts about our perceptual states at every moment of our waking lives (Carruthers

2000). On a global workspace account, it is easy to explain why we should always be subject to some or other phenomenally conscious state. This is because it is the central workspace that serves to focus the entire organism on the most relevant stimuli or the most relevant activities whenever it is awake (Carruthers 2015). But what could possibly be the point of continually entertaining higher-order thoughts about (some subset of) our current perceptual states? And notice that the answer to this question cannot be: to provide flexibility in responding. For it is the first-order contents that guide planning and decision making, not generally the higher-order knowledge that one has such contents.

The other main variety of higher-order approach to consciousness is dual-content theory (Carruthers 2000, 2005). On this view, every globally broadcast perceptual or percept-like state is both a first-order representation of the world or one's own body *and* a higher-order nonconceptual representation of the first-order percept. Each globally broadcast percept with the content *red*, for example, will at the same time have the content *seeming red* or *experience of red*. This is said to be a byproduct of the fact that globally broadcast contents are available to a higher-order thought system (or "mindreading" faculty) that is *capable* of entertaining thoughts about one's own mental states, combined with the truth of some or other form of consumer semantics. (According to the latter, the content of a given representation depends not just on its cause but also on what the consumer systems for that representation are apt to do with it.) Dual-content theory avoids all the difficulties that attend its actualist cousin. But it does require us to buy into a specific theory of content-determination; and there is little or no direct evidence to favor it over a first-order global workspace theory. Moreover, the arguments that might be thought to support dual-content theory do not, in reality, do so (Carruthers 2017). Hence it is rational to prefer the simpler form of first-order account.

Although these and related debates will no doubt continue, in what follows I propose to assume a global workspace account of phenomenal consciousness. Indeed, I propose to assume that phenomenal consciousness doesn't just *co-occur* with global broadcasting, but that it *is* globally broadcast nonconceptual content. For the resulting theory can then be extended to explain, in a scientifically acceptable way, why people find phenomenal consciousness so puzzling — indeed, it can explain why they should so frequently come to believe that consciousness *cannot* be scientifically explained. This is because one can form purely recognitional concepts for the experiences in question that lack any intrinsic connections with physical, functional, or representational concepts (Carruthers 2000; Carruthers & Veillet 2007). Not only can one come to recognize red, without knowing anything further

about it; but one can come to recognize *experiences of red*, without the resulting concept embedding either the concept *experience* or the concept *red* (Carruthers 2017). Let us designate such a concept by *this-R*. Then one can explain how we come to be able to entertain the sorts of thought-experiments that give rise to the "hard problem" of consciousness (Chalmers 1996). One can think, "There could be a being exactly like me in all physical, functional, and representational respects who nevertheless lacks *this-R*." Similarly, when Mary emerges from life in her black-and-white room (Jackson 1986), she will come to know, for the first time, a proposition like, *red tomatoes give rise to experiences such as this-R*, despite (by hypothesis) having previously known everything there is to know of a physical and functional sort about color vision.

It is important to emphasize that the resulting theory is a fully reductive one. There are no special properties attaching to phenomenally conscious states — there are no qualia. What there are, are globally broadcast nonconceptual contents. But such states come to be philosophically puzzling because of the purely recognitional concepts we can form for them. Yet the existence of these puzzles can in turn be reductively explained. Provided that the notions of nonconceptual content, global broadcasting, and recognitional concepts are in good scientific standing (as they surely are), what we have is the outline of a fully explanatory scientific account of phenomenal consciousness. This is what I will assume going forward, as we ask what bearing such a theory has on the question of animal consciousness.

(Those who are unconvinced of the theory can of course read what follows in the form of a conditional: *if* phenomenal consciousness is nothing other than globally broadcast nonconceptual content, then what should be said about phenomenal consciousness in nonhuman animals?)

4. Degrees of broadcasting across species

Dennett (2001) once defended a theory very much like the global workspace account, except that it allowed for degrees. He said that consciousness is like fame in the brain. And just as someone can be more or less famous, or can become just a little bit more famous, so a mental state can be more or less conscious, or get just a little bit more conscious. But it seems he was mistaken. Global broadcasting in humans appears to be an all-or-nothing phenomenon. There is a step-function underlying global broadcasting. Either activation levels in the neural populations in question remain below threshold, in which case there is no global broadcasting (albeit some additional local or specialized effects); or those activation levels hit threshold, and full global broadcasting results (Dehaene & Naccache 2001; Del Cul et

al. 2007; Marti & Dehaene 2017).

When we look across species, however, it is obvious that global broadcasting will admit of degrees. Or more accurately (since the term "global workspace" was introduced in the first instance for the human case), processes in the minds of animals will *more or less closely resemble* human global broadcasting. Recall that globally broadcast contents are made available to a wide range of cognitive systems. They are made available for verbal reporting, to give rise to higher-order awareness, to participate in executive functions of reasoning, planning, and decision making, as well as to memory-forming systems; and they also serve to sustain or modulate full-blown affective reactions to stimuli. Availability for verbal report is unique to human beings, of course. And it seems likely that there can only be higher-order awareness of one's own mental states in a relatively small class of nonhuman creatures (perhaps only great apes, or perhaps primates more generally; or maybe extending to some other social creatures like dolphins and elephants). Moreover, what we refer to collectively as "executive function" is really a set of different capacities that are at least partly independent of one another (Miyake et al. 2000; Diamond 2013). These functions include: selecting from among competing action-schemata, mentally rehearsing actions, inhibiting actions, forming intentions for the future, implementing intentions, switching between tasks, directing attention, searching memory, and modulating emotion. It seems quite likely that some of these capacities will be present in some creatures, but absent in others. In addition, all of the receiver systems for globally broadcast information in humans will admit of degrees of internal complexity across species, and will likewise involve differing degrees of conceptual richness and sophistication.

Recall that the concept *phenomenally conscious* is a first-person one: it can only be grasped through acquaintance with one's own phenomenally conscious mental states. We have good reason to think that the property picked out by the concept is that of globally broadcast nonconceptual content. And as we have just seen, this property admits of degrees across species. However, it seems that the concept, *phenomenally conscious*, is all-or-none. Or so I shall now argue. (For a different but converging argument, see Simon 2017.) In which case we have a mismatch: our first-person concept of consciousness is all-or-none, but the all-or-none property picked out by that concept in the human (first-person) case admits of degrees across species. Later sections will consider what our response to this mismatch should be.

As many people have noted, it is hard to conceive of a case of experience that is partly *like something* to undergo, partly not. Put differently: it is hard to conceive of a perceptual state that is partly phenomenally conscious, partly not. Of course *creature* consciousness admits of degrees, as we noted in Section 1. One can be semi-conscious, or only partly awake. But some of the states one is in when semi-conscious are fully (unequivocally) phenomenally conscious nonetheless. Moreover, the contents of consciousness can be more or less rich or vivid. Contrast looking at an object in full light versus looking at the same object when so dimly lit that one can barely make it out. But still, it is fully (unequivocally) *like something* to be looking at a dimly lit object. Similarly, although someone participating in a backward-masking experiment might report, "I'm not sure whether I saw it or not", the state reported (uncertainty about a brief experience) is fully phenomenally conscious. Finally, although people sometimes use the language of partial consciousness, this seems to refer, not to a state with partly conscious status, but rather to partial *understanding* or partial *interpretation*. Thus someone who witnessed a mugging on a train platform might report, "It all happened so fast, I was only partly conscious of what was happening." But still the inchoate experience in question was fully phenomenally conscious, it seems.

So we have a mismatch: the concept of phenomenal consciousness is all-or-none; the property picked out by that concept is that of globally broadcast nonconceptual content, which is all-or-none in humans; yet global broadcasting admits of degrees across species. There appear to be just four possible ways to respond: (1) give up on the identity between phenomenal consciousness and global broadcasting; (2) change the concept of consciousness to allow for degrees, enabling us to say that other creatures are conscious, not fully, but to some degree; (3) keep the concept all-or-nothing, but define a (vague) categorical boundary around human-like global broadcasting that will definitely include some species while definitely excluding others; or (4) stop even raising the question of consciousness in nonhuman animals. I will argue that #4 is the best option.

5. Give up on the theory?

We have a mismatch between our *concept* of phenomenal consciousness (which is all-or-nothing) and our best theory of the property that the concept picks out (which admits of degrees across species). One might respond to this mismatch by rejecting the theory. One might claim that, given such a mismatch, whatever else the theory of global broadcasting is, it can't be a theory of phenomenal consciousness.

There is no mismatch between the concept of phenomenal consciousness and global broadcasting in the human case, however. This is because global broadcasting in humans is all-or-nothing, as we noted in Section 4. There is only a mismatch in so far as the systems that get broadcast to vary across species. So there is no conflict between first-person (human) applications of the concept *phenomenally conscious* and the theory of global broadcasting. Indeed, as we saw in Section 3, it seems that the latter theory (when suitably supplemented to include purely recognitional phenomenal concepts) can explain everything that needs to be explained about human phenomenal consciousness.

Moreover, although many people have intuitions about which species of animal are likely to enjoy conscious experiences, such intuitions aren't among the possession-conditions for the concept of phenomenal consciousness. One could possess the latter concept while denying that *any* animals are phenomenally conscious, or while claiming that they *all* are. (Indeed, some even claim that every single physical particle in the universe possess a little bit of phenomenal consciousness; Strawson 2006.) So it can't be a constraint on a theory of human consciousness that it should be able to accommodate our intuitions about animals — especially since those intuitions vary quite widely across people.

Another way to put the point is this: there are no definite, first-person-accessible, instances of phenomenal consciousness that are left unexplained by the theory of global broadcasting. The fact that the theory is of little help to us in settling the question of phenomenal consciousness in animals isn't a legitimate reason for rejecting it. While we might have *hoped* for such help from a theory of consciousness, consciousness in animals isn't among the first-person facts that an adequate theory of human consciousness is required to explain.

If there were facts about which animal species enjoy phenomenally conscious experiences and which do not, and if we knew of those facts, then we could use our knowledge of them to triangulate an improved, more precise and more general, theory of consciousness. We could look for just those aspects of global broadcasting that are common to all creatures capable of phenomenal consciousness, which would presumably be a subset of those involved in human consciousness. But we have no direct, theory-independent, access to phenomenal consciousness in animals. And I will suggest in due course that it is a mistake to think that there exists some further fact of the matter regarding animal consciousness or its lack, over and above what we can in principle discover about their cognitive organization.

6. Alter the concept?

Could we alter the concept of phenomenal consciousness to allow for degrees of consciousness across species? It is hard to see how this could be done. For recall that it is essential to the concept that it be introduced first-personally. It is by undergoing phenomenally conscious states, and having one's attention drawn to those states, that one comes to grasp the concept in the first place. But these states in humans are either determinately present, or determinately absent, as we saw in Section 4. We certainly cannot imagine what a partly-phenomenally-conscious experience would be like. For any experience we can imagine will be *definitely* phenomenally conscious.

Even if we can't conceive of a partly-phenomenally-conscious state directly, it might be said that we can at least form a conception of the degree to which the global workspace architecture in another species resembles our own. So we can at least get a handle on the scalar property that our concept of phenomenal consciousness picks out. But even this is doubtful. For recall the complexity of the set of systems to which phenomenally conscious states are broadcast in humans. It seems inevitable that similarities and differences in those systems across species will be complex, multifaceted, and cross-cutting. It is unlikely that there is an objective linearly-ordered similarity-space waiting to be discovered. On the contrary, there will be a multidimensional similarity *network*, with the minds of some species resembling a global broadcasting architecture in some respects, whereas others resemble it in others.

Some of the systems implicated in global broadcasting might be linearly ordered across species, of course. Thus the evolution of capacities for verbal reporting obviously presuppose prior capacities for action-selection. And likewise, one might think that if one lacked a capacity to select among competing actions, then one would also be incapable of forming intentions for the future. For to form an intention is, in effect, to select an action for future implementation. But there might be creatures that have no capacity for executive selection among actions (with choices among actions being determined through bottom-up competition among accumulator systems) which nevertheless have the capacity to *inhibit* an action. And vice versa, there might be creatures that can select between actions but lack any capacity to inhibit habitual or prepotent ones.

Moreover, capacities for planning seem likely to vary independently of the richness of one's conceptual repertoire. Generalist feeders like bears and raccoons will have rich set of concepts and categories, but might forage by employing a semi-random probability-informed walk through the environment with

little or no advance planning, guided by perceived affordances. In contrast, predators like lions and wolves that rely mostly on a single source of prey might have a much more limited conceptual repertoire, while engaging in at least medium-term planning. Moreover, it seems that capacities for top-down modulation of emotion are likely to vary independently of other components of executive function, as will long-term memory capacities.

A scale of degrees of global broadcasting, then, would be something that we would have to *stipulate*, not discover. Put differently: degrees of global broadcasting don't constitute a discoverable natural kind. This will prove to be important when we consider whether we can introduce a categorical notion to apply across species, which is where we go next.

7. Draw a (vague) boundary?

As we noted in Section 6, we can't really make sense of degrees of phenomenal consciousness. But perhaps we don't have to. An obvious suggestion is to say that a creature will enjoy phenomenally conscious experience if it undergoes states that are more similar to human global broadcasting than they are similar to any form of human *un*conscious mental state. We can set up a categorical concept that will reach beyond the human case, while staying true to the all-or-nothing first-personal nature of the concept, by fixing its extension in terms of relative closeness to the two poles drawn from the human case: the global broadcasting architecture that underlies human phenomenally conscious experience, on the one hand, and the sorts of content-availability that underlie *un*conscious forms of perception in humans, on the other.

The resulting concept will be vague, of course (which arguably our concept of phenomenal consciousness isn't; see Simon 2017). There are bound to be cases that are not determinately more similar to global broadcasting than they are to dorsal-network sensorimotor guidance, for example. In part this is because any concept defined in terms of two poles in a continually varying domain will leave cases that are indeterminate between the two. (Think of defining shades of color as *red* provided they are more similar to scarlet than they are to some central instance of bright yellow.) But it is also because of the complex and cross-cutting nature of the similarity space, noted in Section 6. Some creatures will enjoy states that are more similar to globally broadcast ones along one dimension, but more similar to sensorimotor states along another. Moreover, not only are the receiver systems for global broadcasting in humans multifaceted, but so, too, are the causal roles of *un*conscious perceptual states. The states

involved in online and habitual forms of motor control are one thing, whereas the perceptual states that issue in direct affect-expressive behavior are another; and both are quite different again from unattended states in the ventral visual and auditory systems that have a variety of downstream effects independently of global broadcasting. So both poles of the proposed category contain significant (and potentially independently varying) internal complexity.

Are some aspects of global broadcasting more important, or more relevant to the question of consciousness, than others? We have direct evidence from the human case that some of the consumer systems for global broadcast aren't *necessary* for someone to have phenomenally conscious experience. Thus a person with complete amnesia, or someone with completely flattened affect, can nevertheless describe their experiences to us and engage in "hard problem"-type thought experiments. So that would leave us with the broad class of executive function and semantic memory systems (including verbal report) as the relevant components of global broadcasting when we make comparisons across species. However, it might be claimed that these aren't necessary for consciousness, either. For dreams are phenomenally conscious states, but during dreaming one's capacities for executive function are suppressed. Yet it remains unclear whether a creature that *had no* executive function capacities (rather than suppressed ones) could have phenomenally conscious dreams. For our only direct evidence of the phenomenally conscious status of dreaming derives from the reports and reflections that we can make about our dreams when we recall them on waking.

Some components of global broadcasting are *epistemically* more important that others, of course. In particular, capacities for verbal report combined with capacities for higher-order forms of reflective thinking (of the sort that are required for one to become puzzled by the "hard problem" thought-experiments) provide the only *direct* evidence that we have of phenomenally conscious experience in other people. But on a global broadcasting account, these capacities aren't supposed to be *constitutive* of phenomenal consciousness. (That would involve a shift to some kind of higher-order thought theory of consciousness, of the sort discussed briefly in Section 3.) So we can't conclude that these capacities are metaphysically necessary for consciousness to occur.

Given the multicomponent and multifaceted nature of global broadcasting, what would it take for there to be a fact of the matter about whether or not a given mental state in an animal is more similar to a human globally broadcast one than to a human unconscious one? There will inevitably be greater

similarity along some dimensions of comparison than there are on others. So on some ways of weighting the importance of the various dimensions the state in question will come out as conscious, whereas on other ways of weighting those dimensions it will be classified as *un*conscious. How could we decide between these possibilities in any way that doesn't just stipulate the outcome? I suggest that we can't. There is no fact of the matter as to the relevant dimensions of comparison.

We could, however, borrow an idea from the supervaluation literature on vague concepts (Fine 1975; Lewis 1982). We could say that conscious states are the ones whose similarity to human global broadcasting are greater on *any* reasonable way of weighting the different components and dimensions. If phenomenal consciousness in the mind of a monkey, or a chicken, or a honey bee is not to be something that we *stipulate* but rather *discover*, then greater similarity to human global broadcasting would have to survive any reasonable precisification of the concept of global broadcasting. This may well mean that *some* other animals are determinately phenomenally conscious. But the class of such animals is likely to be quite small.

One reason for this is that a multifaceted set of executive functions is at the heart of global broadcasting in humans, and yet it is widely agreed that human executive function capacities are uniquely developed in comparison even to the other great apes, as well as containing components such as verbal report that are wholly unique to humans. So there will already be a considerable gulf between the nature of human and animal global broadcasting. In addition, the conceptual repertoire of an ordinary human will be orders of magnitude richer than even the most sophisticated of animals. (Human concepts are thought to number in the hundreds of thousands; Bloom 2002.) Hence the vast majority of creatures in the animal kingdom will either determinately lack consciousness (because they are more similar to human *un*conscious states on all ways of stipulating the weights in the similarity space), or there will be no fact of the matter whether or not they are phenomenally conscious (because their states are more similar to global broadcasting given some ways of stipulating the weights in the similarity space, but are more similar to human *un*conscious states given others, and thus fail to be classified as conscious under all reasonable weightings of the dimensions).

There is another reason for thinking that the class of animals capable of consciousness will be quite small, if this is the set of creatures that would satisfy all reasonable precisifications of the notion of global broadcasting. This is because both philosophers (Dennett 1978) and psychologists (Kurzban 2012)

have claimed — seemingly reasonably — that conscious mental states are all and only those that can be verbally reported. Likewise, both philosophers (Rosenthal 2005) and psychologists (Graziano 2013) have claimed — again, not unreasonably — that consciousness requires capacities for higher-order thought. On the former approach, consciousness will be restricted to humans; on the latter, it will likely be restricted to a small set of highly social and intelligent species.

Is there any way to cut through the multidimensionality of the space of similarities and differences? Could we say that attention, for example, is the critical factor? Can we say that any creature with a recognizable attentional network can qualify as conscious? We have already noted in Section 3, however, that the attention network is normally individuated in terms of its relationship with global broadcasting. The attention network is the system whose directed signals result in global broadcasting of the representations targeted by those signals. Hence in order to individuate attention, we would first have to individuate global broadcasting. And as we also noted earlier, it won't do to individuate attention in terms of boosting-while-suppressing functions. In part this is because such functions are known to occur unconsciously (for instance, in motor selection) in the human case. And in part it is because there are systems with such properties in animals that look much more like sensorimotor ones than global broadcasting ones. For instance, there are neurons in dragonflies that "lock on" to one among a number of moving targets, boosting the representation of the target and suppressing representations of competitor targets, which enables to dragonfly to pursue one item of prey at a time (Wiederman & O'Carroll 2013).

Exactly similar points can be made if the suggestion is that working memory is sufficient for phenomenal consciousness. For working memory, too, is individuated in terms of its role in global broadcasting. There are lots of reverberating short-term memory systems in humans and other animals that have nothing to do with consciousness.

Perhaps it might turn out, however, that working memory is a natural psychological kind, present in multiple species, and individuated, in part, by its central role in coordinating the activity of the whole organism (Carruthers 2015). But this would be of little help to us. For why should we identify phenomenal consciousness with such a kind? One thing that everyone has agreed on, at least since Kripke (1980), is that terms referring to conscious mental states aren't used as natural-kind terms. In contrast, it is generally agreed that our concept *water* is a natural-kind concept. Even before we knew

anything about chemistry, we used it to refer to the underlying nature or essence of the recognizable stuff that fills our lakes and rivers; and it turned out that it was that very same stuff that presents as ice in some circumstances (frozen water) and as mist in others (evaporated water). But our concepts of the felt qualities of our conscious experiences aren't like that. We don't use them with the intention of referring to whatever natural kind underlies those experiences, whatever that might turn out to be, and however that kind might be presented in other creatures. On the contrary, we mean to refer just to those felt qualities themselves.

So we still face the same dilemma. Given that phenomenal qualities are successfully explained in our own case by full-blown global broadcasting, what reason is there to identify them with the operations of some much-more-minimally described working memory system? This challenge will be especially acute, since many of the remaining components of human-like global broadcast might *also* be natural kinds. This might be true of the language faculty, for example, or the mindreading network. So why shouldn't we say that it is availability of nonconceptual content to one of *these* kinds that fixes the extension of consciousness?

We can conclude that although global workspace theory can give us a categorical concept to employ across species, it seems likely that only a small subset of species (perhaps a set restricted to humans) will qualify as determinately phenomenally conscious, on this approach; and equally, only a small subset will qualify as definitely *lacking* phenomenally conscious states, likewise. For the vast majority of species, the phenomenally conscious status of their perceptual states would have to be something that we stipulate rather than discover.

8. Doing without consciousness

It seems that there are a great many species of animal for whom there is no fact of the matter as to whether they have phenomenally conscious states or not. Is this something that should disturb us? And should we now try to settle which species have states that are close enough to being globally broadcast in the human manner that any reasonable precicisification of the concept of global broadcasting would show them to be conscious? I think both questions should be answered negatively.

Recall that we are supposing that globally broadcast nonconceptual content is what phenomenal consciousness *is*. Hence there is no special property that gets added when a perceptual state is globally

broadcast — there are no qualia. As we transition from species whose states aren't close enough to human global broadcasting to qualify as phenomenally conscious to species whose states *are* sufficiently close, nothing lights up, and nothing magical appears. There is just nonconceptual content that is available to a greater range of systems, or to systems with greater internal complexity or conceptual sophistication. There will be functional differences between the two sets of contents of course. But then there will likewise be functional differences among the mental states of two species that fall on the unconscious side of the divide, or two species that each fall on the consciousness side. The differences in question don't differ in kind in any deep way.

In the human case there is a big introspective difference between states that are conscious and states that are not, of course, and the difference is akin to "lighting up". For globally broadcast states are ones that we are immediately aware of having, whereas un-broadcast states are ones that we aren't aware of having at all (except through third-person-type interpretative inferences). But introspective availability is just one facet of global broadcasting in humans. It isn't supposed to *constitute* phenomenal consciousness (that would turn the global workspace account into a higher-order theory, rather than a first-order one). Still, one might think that if a state is conscious then there is something that we can *imagine* when we imagine the mental life of the animal, whereas if a state is unconscious there isn't. And this seems like a big difference. Indeed, for many people this is what they want a theory of consciousness in animals *for* — they want to know in what cases it is appropriate to take the perspective of the animal, and to imagine what it is like to be the animal.

Let us take these suggestions in turn. First, is phenomenal consciousness necessary for an animal to have a perspective on the world? I suggest not. All perceptual states of any complexity are perspectival, whether they are conscious or not. They represent the world as seen, or heard, or touched, or smelled from the spatial location of the agent or its parts. Hence they are subjective in the sense that they represent some properties of the world but not others. All visual contents will represent some of the properties in the line of sight of the subject, for instance, but not those that are occluded by other objects or are on the non-visible surface of an object.

Turn, now, to the suggestion that the consciousness question needs to be settled so that we can imagine what it is like to be the animal, since this depends on the conscious status of its mental states. But in fact, imagination is a highly unreliable guide to the mental life of an animal in any case. For anything we

imagine, and any set of images we form, will be globally broadcast to our own set of consumer systems, which differ quite significantly from those of the target animal. So the result will be a distinctively-human conscious state, rather than a strictly animal one. For the same reason, we should not use imagination to ground our sympathy directed at the plight of an animal. For what we imagine will be broadcast to our own affective and valuational systems, not the animal's. So the resulting empathic affective response is likely to be quite wide of its target. If one thinks that it might be important to be sympathetic towards the situation of an animal, one should seek an accurate third-person understanding of its needs and affective states, not project our own feelings onto it.

Suppose we had complete knowledge of the functional and representational aspects of the mentality of a monkey, or a salmon, or a honey bee. Suppose we knew everything of a third-personal sort that there is to know about what happens to its perceptual states under various conditions, how those states interact with valuational ones, and how the resulting behavior is determined. Would it add anything to our knowledge to have done the comparative work needed to know whether states in these animals are more like global broadcast ones in humans than they are like any of the various kinds of unconscious state, or vice versa? I suggest not. That comparison, and the resulting classification of the state as conscious or unconscious, wouldn't add anything to what we already knew. There is no extra property of the mind of the animal that accrues to it by virtue of its similarity or dissimilarity to the human global workspace.

There has been a flurry of interest in consciousness in animals lately, including books by Godfrey-Smith (2016) and Tye (2017), as well as the inauguration in 2016 of a new scientific journal, *Animal Sentience*, which is devoted to the study of the topic. I suggest that this interest is misconceived. Given our best theory of what consciousness is (globally broadcast nonconceptual content), the question of animal consciousness is essentially just a classificatory one. All of the substantive questions are of the sort traditionally addressed by comparative psychologists without needing to raise the question of consciousness — How does the mind of a given type of animal actually work? What features of the environment is it sensitive to, and in what ways? Among these questions, of course, will be ones that relate to components of global broadcasting — Is it capable of higher-order thought? Is it capable of planning? Is it capable of mental rehearsal of action? To what extent does it have a system with the psychometric properties of human working memory? And so on. These are important questions, and represent the true hard problem of animal minds. But once the answers are known, there is no point in

asking the further question whether the systems in question are similar enough to human global broadcasting (on all reasonable ways of weighting the various cross-cutting dimensions of similarity) to receive a classification of "conscious". This latter question will tell us nothing we didn't already know, and an answer to it would serve no useful purpose.

Many of those who are interested in the morality of our treatment of different animal species plainly think that the question of animal consciousness is important. Indeed, some of them seem to think that this is the "magic bullet" that will settle the question of moral standing for each animal species. This is certainly true of Singer (1981, 1993) and other utilitarians (Bentham 1789). But there are two reasons to be suspicious of this approach. (Actually three, since utilitarianism is arguably a false moral theory. But we can set this aside for present purposes.) One is that there is likely to be no determinate answer to the question of phenomenal consciousness in animals for the vast majority of animal species, as we noted in Section 7. Yet one might think that there are still real questions about how these animals should be treated. And the second reason for suspicion is that it is far from clear why only phenomenally conscious states should matter in any case. There can certainly be negatively valenced states in the absence of consciousness (Winkielman et al. 2005; Barrett & Bar 2009; Lebrecht et al. 2012). So there can be states that are in some respects pain-like or suffering-like in creatures whose cognitive organization is too far distant from a human global workspace architecture to qualify them for consciousness. *Perhaps* these states should matter less than those that are conscious. But it isn't obvious that this is so, and one might think that it would be a task for moral theory to consider whether and why it should be so.

An exclusive focus on *conscious* pain and *conscious* suffering distracts from the very real work that needs to be done, teasing out the moral relevance of the various sorts of cognitive organization that we discover in animals. A better approach is that of Dawkins (2012), I suggest, who brackets questions of consciousness in our treatment of animals, focusing, rather, on questions of animal welfare, health, and flourishing. These are third-personal notions, rendering the questions in which they are embedded more tractable and substantive.

9. Of Martians, mice, and men

Suppose that there are Martians who are vastly more intelligent than us. Specifically, suppose that their executive function capacities outstrip ours by orders of magnitude, adding many new abilities while also

greatly magnifying and expanding some of the same capacities we possess. Indeed, suppose that their cognitive abilities outstrip ours by as much as our minds outstrip the mind of a mouse. Martian philosophers might then wonder whether creatures as cognitively primitive as humans could really enjoy phenomenally conscious mental states. (Suppose that Martians, like humans, can introspectively recognize the qualities of their own perceptual states.) Some of them might even argue that there is no fact of the matter. But they would be wrong.

Not only would the Martian philosophers be making a mistake in claiming that there is no fact of the matter about human consciousness, but we could surely (given sufficient time and a means of communication) convince them of this fact. For we could talk to them about our experiences, they could learn that we have become obsessed with what we call the "hard problem" of the relationship between the introspectively accessible properties of our experiences and their representational contents and functions, and they could infer that we possess purely recognitional concepts for the qualities of our globally broadcast perceptual states. It is reasonable to assume that mice are incapable of providing us with such evidence. Moreover, they are so incapable, not just because they can't talk, but because they lack the underlying abilities. I presume that mice don't bother their heads over the "hard problem" of consciousness; indeed, they are surely incapable of doing so.

Although the Martian example doesn't provide a counterexample to the claim that there is no fact of the matter about consciousness in most species of animal, it does illustrate how consciousness can be present in two creatures, one of whose executive capacities is a subset of those possessed by the other. The global workspace in a Martian encompasses our own, and then some. So not all of those Martian capacities are necessary for consciousness to exist; in particular, the enhancements that are uniquely present in them aren't. In which case, might not something similar be true in respect of the mouse? Mightn't the subset of capacities we share with the mouse be sufficient for phenomenal consciousness in the latter? And couldn't there be a fact of the matter, even though "global broadcasting" in the two species is otherwise very different?

This line of thinking presupposes that phenomenal consciousness is "a thing", however, over-and-above the sorts of availability-relations in question. Global broadcasting in humans occurs within a framework of capacities that enables us to form purely recognitional concepts for our experiences, get puzzled about zombies and the "hard problem" and so on. This warrants us in introducing a first-personal

concept of phenomenal consciousness in addition the notion of access consciousness, thereafter inquiring about the relationships between them. But given that the upshot of that inquiry is that phenomenal consciousness *is* globally broadcast nonconceptual content, we can see that there is no extra property possessed by our conscious experiences beyond their globally broadcast status. We lack the same first-person warrant for ascribing phenomenal consciousness to mice, however (whereas the Martians *are* so warranted in ascribing it to us, as we noted earlier). But more important, there is no extra property that the mice might either have or lack. They have whatever global-broadcasting-like arrangement they have. There is nothing else worth saying. It *is* worth charting the similarities and differences between the mouse's cognitive architecture and our own, of course. This is just regular comparative psychology. But the question whether the similarities are sufficient for consciousness is merely a classificatory one, and tracks no further substantial matter of fact.

10. Conclusion

The science of consciousness has made huge strides in recent years. We are now at a point where we might have the outlines of a full scientific explanation. In particular, global workspace theory is supported by multiple lines of evidence, and can be extended to explain why we should find our own phenomenally conscious states to be so scientifically puzzling. This is because we can form purely recognitional concepts for those states that lack any conceptual connections to the underlying scientific facts, thus enabling us to conceive of zombies, and leading us to think that there must be an explanatory gap of some sort between the scientific facts and the facts of consciousness. By postulating that phenomenal consciousness is globally broadcast nonconceptual content, and by supposing that we can form purely recognitional concepts for those contents, we can explain everything that stands in need of an explanation — including the appearance of an explanatory gap.

When we turn from the human case to inquire about the implications of global workspace theory for the question of consciousness in animals, however, matters are not so straightforward. For our first-person conception of phenomenal consciousness is an all-or-nothing one, whereas global broadcasting will admit of degrees across species, of a multifaceted and cross-cutting sort. We could respond by saying that creatures have phenomenal consciousness whose states more closely resemble human global broadcasting than they resemble human forms of unconscious perception. But the degrees of resemblance appealed to will either have to be stipulated rather than discovered, or they will have to be restricted to what would survive all reasonable stipulations. On the first alternative, there are no facts of

the matter about animal consciousness. On the latter, there might be; but these will likely be limited to a small class of cases, leaving a large range of indeterminacy.

More directly, I have argued that if we embrace the fully-reductive nature of global workspace theory, then these questions no longer have any deep significance. There are many facts to be discovered about the minds of animals; and among these facts will be dimensions of similarity and difference from the human global workspace. But whether in any given case there is *enough* similarity for an animal to qualify as conscious would add nothing further to what we would already know. So this isn't a question that comparative psychologists should spend any time on. The nature and functioning of animal minds should be studied; but phenomenal consciousness in animals doesn't deserve to be.

Acknowledgments

I am grateful to Chris Masciari, Joe Millum, Aida Roige, Eric Saidel, Julius Schönherr, and Moonyoung Song for comments on some of the ideas presented in this article.

References

Baars, B. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.

Baars, B. (1997). *In the Theater of Consciousness*. Oxford University Press.

Baars, B. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6, 47-52.

Barrett, L.F. & Bar, M. (2009). See it with feeling: affective predictions during object perception. *Philosophical Transactions of the Royal Society B*, 364, 1325-1334.

Bayne, T., Hohwy, J., & Owen, A.M. (2016). Are there levels of consciousness? *Trends in Cognitive Sciences*, 20, 405-413.

Bentham, J. (1789). *An Introduction to the Principles of Morals and Legislation*. Oxford University Press.

Block, N. (1995). A confusion about the function of consciousness. *Behavioral and Brain Sciences*, 18, 227-247.

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481-499.

Block, N. (2011a). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Science*, 12, 567-575.

Block, N. (2011b). The higher-order approach to consciousness is defunct. *Analysis*, 71, 419-431.

Bloom, P. (2002). *How Children Learn the Meaning of Words*. MIT Press.

Breitmeyer, B. & Ogmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics*, 62, 1572-1595.

Burton, A., Nakamura, K., & Roesch, M. (2015). From ventral-medial to dorsal-lateral striatum: Neural correlates of reward-guided decision-making. *Neurobiology of Learning and Memory*, 117, 51-59.

Carruthers, P. & Veillet, B. (2007). The phenomenal concept strategy. *Journal of Consciousness Studies*, 14, 9-10, 212-236.

Carruthers, P. & Veillet, B. (2011). The case against cognitive phenomenology. In T. Bayne and M. Montague (eds.), *Cognitive Phenomenology*. Oxford University Press, 35-56.

Carruthers, P. & Veillet, B. (2017). Consciousness operationalized, a debate realigned. *Consciousness and Cognition*, 55, 79-90.

Carruthers, P. (2000). *Phenomenal Consciousness*. Cambridge University Press.

Carruthers, P. (2005). *Consciousness: Essays from a higher-order perspective*. Oxford University Press.

Carruthers, P. (2015). *The Centered Mind: what the science of working memory shows us about the nature of human thought.* Oxford University Press.

Carruthers, P. (2017). In defense of first-order representationalism. *Journal of Consciousness Studies*, 24, 5-6, 74-87.

Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press.

Cohen, M.A. & Dennett, D. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15, 358-364.

Cohen, M.A., Dennett, D., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences*, 20, 324-335.

Dawkins, M. (2012). *Why Animals Matter.* Oxford University Press.

Dehaene, S. (2014). *Consciousness and the Brain*. Viking Press.

Dehaene, S. & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79, 1-37.

Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLOS Biology*, 5, e260.

Dennett, D. (1978). Toward a cognitive theory of consciousness. In his *Brainstorms*, Harvester Press.

Dennett, D. (2001). Are we explaining consciousness yet? *Cognition*, 79, 221-37.

Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, 64, 135-168.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 54, 235-259.

Frässle, S., Sommer, J., Jansen, A., Naber, M., & Einhäuser, W. (2014). Binocular rivalry: frontal activity relates to introspection and action but not to perception. *The Journal of Neuroscience*, 34, 1738-1747.

Godfrey-Smith, P. (2016). *Other Minds*. Farrar, Straus, and Giroux.

Goodale, M. (2014). How (and why) the visual control of action differs from visual perception. *Proceedings of the Royal Society B*, 281, 20140337.

Graziano, M. (2013). *Consciousness and the Social Brain*. Oxford University Press.

Hameroff, S. & Penrose, R. (1996). Orchestrated reduction of quantum coherence in brain microtubules: A model of consciousness. *Mathematics and Computers in Simulation*, 40, 453-480.

Hickok, G. & Poeppel, D. (2004). Dorsal and ventral streams: a framework for understanding aspects of the functional anatomy of language. *Cognition*, 92, 67-99.

Humphrey, N. (1974). Vision in a monkey without striate cortex: A case study. *Perception*, 3, 241-255.

Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy*, 83, 291-295.

Januszko, P., Niemcewicz, S., Gajda, T., Wolynczyk-Gmaj, D., Justyna, P., Gmaj, B., Piotrowski, T., & Szelenberger, W. (2016). Sleepwalking episodes are preceded by arousal-related activation in the cingulate motor area: EEG current density imaging. *Clinical Neurophysiology*, 127, 530-536.

Kripke, S. (1980). *Naming and Necessity*. Blackwell.

Kurzban, R. (2012). *Why Everyone (Else) is a Hypocrite: Evolution and the modular mind*. Princeton University Press.

Landman, R., Spekreijse, H., & Lamme, V. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research,* 43, 149-164.

Lebrecht, S., Bar, M., Barrett, L.F., & Tarr, M. (2012). Micro-valences: perceiving affective valence in everyday objects. *Frontiers in Psychology*, 3, doi: 10.3389/fpsyg.2012.00107.

LeDoux, J. (1996). *The Emotional Brain*. Simon and Schuster.

LeDoux, J. (2017). Semantics, surplus meaning, and the science of fear. *Trends in Cognitive Sciences*, 21, 303-306.

Lewis, D. (1982). Logic for equivocators. *Noûs*, 16, 431-441.

Lisman, J. & Sternberg, E. (2013). Habit and nonhabit systems for unconscious and conscious behavior: Implications for multitasking. *Journal of Cognitive Neuroscience*, 25, 273-283.

Mandik, P. (2009). Beware of the unicorn: consciousness as being represented and other things that don't exist. *Journal of Consciousness Studies*, 161, 5-36.

Marti, S., Sigman, M., & Dehaene, S. (2012). A shared cortical bottleneck underlying attentional blink and psychological refractory period. *NeuroImage*, 59, 2883-2898.

Marti, S. & Dehaene, S. (2017). Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature Communications*, 8:1955. DOI: 10.1038/s41467-017-02079-x.

Milner & Goodale (1995). *The Visual Brain in Action*. Oxford University Press.

Miyake, A., Friedman, N., Emerson, M., Witzki, A., Howerter, A., & Wager, T. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, 41, 49-100.

Mole, C. (2011). *Attention is Cognitive Unison*. Oxford University Press.

Panksepp, J. (1998). *Affective Neuroscience*. Oxford University Press

Pitts, M., Metzler, S., & Hillyard, S. (2014). Isolating neural correlates of conscious perception from neural correlates of reporting one's perception. *Frontiers in Psychology*, 5:1078.

Prinz, J. (2012). *The Conscious Brain*. Oxford University Press.

Rauschecker, J. & Tian, B. (2000). Mechanisms and streams for processing of "what" and "where" in auditory cortex. *Proceedings of the National Academy of Sciences*, 97, 11800-11806.

Raymond, J., Shapiro, K., & Arnell, K. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18, 849-860.

Redgrave, P., Rodriguez, M., Smith, Y., Rodriguez-Oroz, M., Lehericy, S., Bergman, H., Agid, Y., DeLong, M., & Obeso, J. (2010). Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease. *Nature Reviews Neuroscience*, 11, 760-772.

Rosenthal, D. (2005). *Consciousness and Mind*. Oxford University Press.

Saint-Cyr, J., Ungerleider, L., & Desimone, R. (1990). Organization of visual cortical inputs to the striatum and subsequent outputs to the pallido-nigral complex in the monkey. *Journal of Comparative Neurology*, 298, 129-156.

Salti, M., Monto, S., Charles, L., King, J-R., Parkkonen, L., & Dehaene, S. (2015). Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. *eLife*, 4:e05652.

Simon, J. (2017). Vagueness and zombies: why "phenomenally conscious" has no borderline cases. *Philosophical Studies*

Singer, P. (1981). *The Expanding Circle*. Oxford University Press.

Singer, P. (1993). *Practical Ethics* (2nd Edition). Cambridge University Press.

Sligte, I., Scholte, H.S., & Lamme, V. (2008). Are there multiple visual short-term memory stores? *PLoS*

ONE, 3, e1699.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74, 1-29.

Strawson, G. (2006). Realistic monism – why physicalism entails panpsychism. *Journal of Consciousness Studies*, 13 (10-11), 3-31.

Terzaghi, M., Sartori, I., Tassi, L., Rustoni, V., Proserpio, P., Lorusso, G., Manni, R., & Nobili, L. (2012). Dissociated local arousal states underlying essential clinical features of non-rapid eye movement arousal parasomnia: an intracerebral stero-electroencephalographic study. *Journal of Sleep Research*, 21, 502-506.

Tong, F., Nakayama, K., Vaughan, J.T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, 21, 753-759.

Tononi, G. & Koch, C. (2015). Consciousness: here, there, and everywhere? *Philosophical Transactions of the Royal Society B*, 370, 20140167.

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *Biological Bulletin*, 215, 216-242.

Tsubomi, H., Fukuda, K., Watanabe, K., & Vogel, E. (2013). Neural limits to representing objects still within view. *The Journal of Neuroscience*, 33, 8257-8263.

Tsuchiya, N. & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8, 1096-1101.

Tye, M. (2017). *Tense Bees and Shell-Shocked Crabs: are animals conscious?* Oxford University Press.

Tye, M. (1995). *Ten Problems of Consciousness*. MIT Press.

Tye, M. (2000). *Consciousness, Color, and Content.* MIT Press.

Weisberg, J. (2011). Misrepresenting consciousness. *Philosophical Studies*, 154, 409-433.

Weiskrantz, L. (1986). *Blindsight*. Oxford University Press.

Wiederman, S. & O'Carroll, D. (2013). Selective attention in an insect visual neuron. *Current Biology*, 23, 156-161.

Wilke, M., Logothetis, N., & Leopold, D. (2003). Generalized flash suppression of salient visual targets. *Neuron*, 39, 1043-1052.

Winkielman, P., Berridge, K., & Wilbarger, J. (2005). Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Personality and Social Psychology Bulletin*, 31, 121-135.

Wood, W. & Rünger D. (2015). Psychology of habit. *Annual Review of Psychology*, 67, 289-314.

Wu, W. (2014). *Attention*. Routledge.