

Comparative Psychology without Consciousness

Peter Carruthers

Abstract: The goal of this paper is to establish the truth of the following conditional: if a global workspace theory of phenomenal consciousness is correct, and is fully reductive in nature, then we should stop asking questions about consciousness in nonhuman animals — not because those questions are too hard to answer, but because there are no substantive facts to discover. The argument in support of this conditional turns on the idea that 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. Before getting to that argument, however, and in order to motivate the subsequent discussion, some of the virtues of global workspace theory are displayed.

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

1. Initial distinctions

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, that has a distinctive subjective *feel*, or that has a *qualitative character*. 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). Likewise, there seems to be an unbridgeable explanatory gap between all physical, functional, and representational facts and one's current conscious experience. 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.

It is important to emphasize that the concept of phenomenal consciousness is a first-person one. The various locutions employed (“like something to undergo”, “subjective feel”, “qualitative character”, and so on) are all intended just to draw one’s attention to one’s own phenomenally conscious experiences. Acquaintance with the latter is a necessary condition for grasping the concept, and no definition or third-personal explanation could confer understanding of the concept. (As Block 1995 remarks, adapting a comment about jazz often attributed to Louis Armstrong, “If you gotta ask what it is, you ain’t never gonna know.”) Hence philosophical zombies not only lack phenomenal consciousness itself; they must also lack the *concept* of phenomenal consciousness (Chalmers 2006).

Phenomenal consciousness is at least conceptually distinct from *access* consciousness (Block 1995, 2007). Both are forms of mental-state consciousness: it is mental states that can have phenomenal properties, and that can be accessible to enter into decision making, reasoning, and verbal report. As we have already seen, however, *phenomenal* consciousness is a first-personal notion. 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. And it is access consciousness that can be most directly explained in terms of the global neuronal workspace (Dehaene 2014).

It is controversial (to say the least) whether or not there is any real distinction between access consciousness and phenomenal consciousness. Put differently, although the *concepts* are distinct, it is controversial whether the two concepts pick out distinct properties or converge on the same property. There are three separable strands in this debate. One is about so-called “cognitive phenomenology” (Bayne & Montague 2011). On the assumption that thoughts and concepts can be access conscious, it is debated whether they make irreducible (as opposed to merely causal) contributions to people’s phenomenally conscious states. Some have argued that they do (Strawson 1994, 2011; Siewert 1998, 2011; Pitt 2004), whereas others have argued that they do not (Jackendoff 1987, 2002; Tye 2000; Tye & Wright 2011; Carruthers & Veillet 2011, 2017). For present purposes I will be assuming that it is only nonconceptual contents (when access conscious) that make a constitutive contribution to people’s phenomenology. But nothing substantive will turn on this assumption. The issues concerning

phenomenal consciousness in animals remain much the same, no matter how this particular debate plays out.

A second strand of debate about the reality of the access / phenomenal distinction is more consequential, however. This concerns the alleged richness of phenomenally conscious experience, as opposed to the relative paucity of content that can be globally broadcast and sustained in the global workspace. Block (1995, 2007, 2011) argues 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; however, they *can* report any given subset of those details when cued to do so after stimulus offset, suggesting that the full stimulus is present in consciousness while only being available for report piecemeal (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 / the global workspace. Critics have replied that people's intuitions that they see the stimulus in rich detail can be explained away, in part by appealing to visual summary statistics (Cohen & Dennett 2011; Cohen et al. 2016), and others have pointed out that people can only report around four items even when the cue for reporting coincides with stimulus offset, obviating any need for memory (Tsubomi et al. 2013).

This second strand of debate is highly relevant to the question whether global workspace theory can provide an adequate account of phenomenal consciousness, of course. But attempting to resolve it would require a paper of its own, and so cannot be attempted here. Recall that the goal of the present discussion is to establish the truth of a conditional: that *if* global workspace theory provides a complete account of phenomenal consciousness, then there is no fact of the matter whether nonhuman animals are phenomenally conscious. Establishing the truth of the antecedent of this conditional would require comparison with and rebuttal of many competing accounts of consciousness, including not just Block's iconic-memory account, but also higher-order theories (Rosenthal 2005; Graziano 2013), integrated-information theory (Tononi 2008; Tononi & Koch 2015), and more. This has to be a task for another occasion. It should be said, however, that each of these competitor theories is likely to have different

implications for the question of animal consciousness, which likewise need to be explored.¹

There is yet a third strand of debate over the reality of the distinction between access consciousness and phenomenal consciousness. For even if one thinks that the two concepts are coextensive in normal humans (as does Chalmers 1996), it is possible to claim that there are a special set of properties that are picked out first-personally (so-called “qualia”) that aren’t reducible to others, and aren’t explicable in third-person terms. This strand of debate will be addressed briefly in Section 2, where I outline how global workspace theory can provide a fully reductive account of phenomenal consciousness, thus avoiding any need to postulate such a special set of properties. Before we get to that, however, there is one other pair of distinctions that needs to be explained.

Mental-state consciousness (whether access or phenomenal) should be distinguished from *creature* consciousness, which can be either *transitive* or *intransitive* (Rosenthal 2005). Whenever a creature (whether human or animal) is aware of some object or event in its environment or body, it can be said to be (transitively) conscious of that object or event. Put differently, a creature is transitively-conscious of an object or event when it *perceives* that object or event. It is debatable whether or not transitive creature consciousness requires mental-state consciousness. For it is debatable whether the perceptual states that enable a creature to be aware of its environment must be conscious ones. At any rate, it is worth noting that there are many kinds of case where one would pre-theoretically ascribe creature consciousness to an agent — since the agent is displaying perceptual sensitivity to the environment — but where the states in virtue of which it acts as it does are *not* conscious ones. This point can be illustrated using the discovery of two visual systems in humans (Milner & Goodale 1995; Goodale 2014).

¹ Higher-order theories imply that only creatures capable of thoughts about their own mental states can be phenomenally conscious, of course, which probably constricts the set of conscious creatures quite considerably. Integrated-information theory, in contrast, implies that all creatures are to some degree conscious. Indeed, it implies something that comes close to a kind of panpsychism, as Tononi & Koch (2015) point out, since almost all complex systems contain some degree of integrated information. Thus, one’s cerebellum will be to some degree conscious, as will one’s spinal cord. Moreover, the same challenge arises for integrated-information theory as for global workspace theory: it implies that consciousness admits of degrees across species, whereas our concept of phenomenal conscious implies that consciousness is all-or-nothing. See the discussion that follows. The implications of Block’s iconic-memory account, in contrast, are much less clear, and will depend on how, specifically, one chooses to individuate iconic memory.

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 and 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).

Observing D.F.'s behavior, an uninformed person would surely say, "She reaches as she does, and orients her hands as she does, because she *sees* the orientation of the pencil and the position of the letter box." From a common-sense perspective we would explain D.F.'s successful interactions with her environment by saying that she *sees*, or is *creature conscious*, of the objects in her environment. Of course it is possible to insist that this is mistake. One could claim that D.F. lacks transitive creature consciousness of orientation and shape on the grounds that the perceptual states in question aren't conscious ones. This would be to insist that transitive creature consciousness requires appropriate forms of mental-state consciousness. I have no objection to such a stipulation. But notice that if we take this route, then we can no longer move directly from the fact that an animal displays perceptual sensitivity to its environment to the conclusion that it is aware of its environment, or is conscious of its environment. For it seems that humans display numerous forms of sensitivity to the environment that don't involve mental-state consciousness.²

² All forms of fast online skilled engagement with the environment may provide one set of examples. Consider 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,

Turning, now, to *intransitive* creature consciousness, this is what Bayne et al. (2016) call “global-state consciousness”. It 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 *unconscious*. Note that intransitive 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 state of caffeine-induced high alertness will be more conscious than someone who is drowsy.

To summarize, our target is phenomenal consciousness, which is first-personally characterized; and our question concerns the distribution of phenomenal consciousness across the animal kingdom.

Phenomenal consciousness is a property of (some) mental states, and is at least conceptually distinct

light hitting the retina from the ball as it exits the pitcher’s hand) and access-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.

Other environmentally-sensitive actions that can take place in the absence of conscious experience would include the behavioral components of emotional states like 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, and 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.

from access consciousness, which can be accounted for in third-personal terms. (Section 2 will try to motivate the claim that phenomenal consciousness can be fully reductively explained in terms of a certain sort of access consciousness.) In addition, mental-state consciousness can be distinguished from creature consciousness, whether transitive or intransitive. Note that both forms of creature consciousness admit of degrees (one can be more or less aware of the properties of a stimulus, and one can be more or less awake). It will be argued in Section 4, in contrast, that it is all-or-nothing whether or not a given mental state is phenomenally conscious.

2. Motivating global workspace theory

A global workspace account of consciousness was first developed in detail by Baars (1988, 1997, 2002), and was originally expressed in cognitive (rather than neural-network) terms. A close relative of Baars' theory (published subsequently but independently) is Tye's (1995, 2000) PANIC theory. (PANIC stands for Poised Abstract Nonconceptual Intentional Content.) The basic idea of both approaches is that some perceptual and perception-like contents (including visual and auditory imagery) are globally broadcast and received by a wide range of other systems in the mind, whereas others have more specialized or local uses. Conscious states are those that are broadcast to systems for reporting, planning, reasoning, decision making, and remembering, whereas unconscious states are those that are *not* so broadcast, although they can have a variety of other roles or effects.

Subsequent work has expanded the evidence-base for global workspace theory to include especially the neural networks involved. Much of this work has used minimal contrasts between conscious and unconscious conditions in combination with EEG, MEG, and/or fMRI.³ (See Dehaene 2014 for a review.) 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

³ Some have disputed whether the states in such "unconscious" conditions are genuinely unconscious, or merely ones that subjects have low confidence about (Peters & Lau 2015). This debate cannot be addressed here.

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, they 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 visual cortex 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). Note that the prefrontal cortex is thought to house the main hubs of the networks responsible for decision making, reasoning, and verbal report, and that it interacts heavily with the hippocampus and medial temporal cortex in the formation of long-term memories (Buchner et al. 1999; Preston & Eichenbaum 2013). So we appear to have a neural-network-based account of all of the main components of access consciousness, at least.

Critics have sometimes objected that the minimal-contrast methodology controls for stimulus factors but not for behavioral or downstream cognitive ones (Frässle et al. 2014; Pitts et al. 2014; Michel 2017). 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. Hence there is ongoing debate over the question whether prefrontal cortex is specifically involved in consciousness itself, as opposed to its causes and effects (Boly et al. 2017; Odegaard et al. 2017).

All sides in these debates should accept that specific regions of temporal and parietal cortices are critically involved in processing and representing particular sensory contents, of course, such as faces, houses, colors, and so on. But according to global workspace theory, these contents need to be broadcast more broadly — and to prefrontal cortex, in particular — in order to become conscious. Notice, however, that since these contents are thought to be broadcast quite widely, their representation in prefrontal cortex is likely to be diffuse and distributed. As a result, one can explain why prefrontal content-related activity sometimes doesn't show up in brain-imaging studies, as Odegaard et al. (2017) point out. Likewise, the same point can explain why the prefrontal MEG response to a briefly presented stimulus should be less predictive of conscious experience than the MEG signal

emanating from visual areas (Andersen et al. 2016). For visual processing of the stimulus is concentrated in visual cortex, of course, whereas (when conscious) the resulting contents will be broadcast to multiple cortical sites, distributed in a variety of functionally-distinct regions. Moreover, studies using cellular recording techniques together with binocular flash suppression *have* shown that the content of a stimulus is represented in neuronal activity in prefrontal cortex under passive viewing conditions, but not when that stimulus is perceptually suppressed (Panagiotaropoulos et al. 2012). Although this study was conducted with monkeys, and so cannot provide direct evidence of prefrontal involvement in *phenomenal* consciousness, it does imply that the prefrontal activity observed in instances of conscious experience in humans can't just be a matter of action-preparation.

In addition, a recent study attempted to control for both behavioral and attentional factors (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 (Weiskrantz 1986). 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 further into prefrontal cortex, and was processed much more deeply, than in blindsight cases that shared the same motor response.

Of course these debates are ongoing, and will continue. I don't pretend to have resolved them here. All I mean to establish is that global workspace theory is a *contender*, and should be taken seriously. But even if the theory were accepted, and it were thus accepted that access consciousness (in the form of global workspace theory) and phenomenal consciousness coincide, what reason would there be for thinking that phenomenal consciousness is thereby *explained*? For as we have already noted, some have accepted that global workspace theory is a correct account of access consciousness, and have accepted that access consciousness and phenomenal consciousness are always coextensive, while denying that global workspace theory can *explain* phenomenal consciousness (Chalmers 1996). This is on the grounds that, even if the two coincide, there remains an *explanatory gap* between global accessibility, on the one hand, and phenomenal consciousness, on the other. For it is possible to conceive of a creature that

has the former but not the latter, for example (a zombie).

The key reason for thinking that phenomenal consciousness doesn't just *co-occur* with global broadcasting, but that it *is* globally broadcast nonconceptual content, is that 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 indexical, acquaintance-based, concepts for the experiences in question that lack any intrinsic connections with physical, functional, or representational concepts (Carruthers 2000; Tye 2000; Carruthers & Veillet 2007; Prinz 2012).

On this account, not only can one come to recognize red, without knowing anything further about it; but one can entertain thoughts about experiences of red without the resulting thoughts 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*”, and one can think, “In other people, states like *this-R* might be caused by looking at green grass.” 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 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 intended to be fully reductive. 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 acquaintance-based indexical 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 indexical 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.

All of the points just mentioned are still hotly debated, of course. Many have argued that the phenomenal concept strategy for responding to zombie-arguments and the explanatory gap cannot

succeed (Chalmers 2006; Goff 2011). Others have denied that there are such things as phenomenal concepts in any case (Ball 2009), and Tye himself has changed his mind on the topic (Tye 2009). It would require a separate paper to resolve these issues, also. That is not my purpose. All I have attempted to do here is to sketch out how global workspace theory supplemented with the phenomenal concept strategy should be highly attractive, at least, to anyone with physicalist and naturalistic commitments. For it holds out the promise of providing a fully reductive explanatory theory of phenomenal consciousness, without needing to appeal to any special properties of “qualia” (Chalmers 2006) or “mental paint” (Block 2010) that might otherwise need to be postulated to exist. As a result, it is worth considering the implications of global workspace theory for the question of phenomenal consciousness in animals. That will be the topic of the remainder of this paper.

3. 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 can 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; van Vugt et al. 2018).

When we look across species, however, it seems obvious that global broadcasting will admit of degrees.⁴ 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

⁴ This is not to say that *all* aspects of global broadcasting admit of degrees, of course. And in particular, it is possible that the raw limits on broadcasting quantity are the same across primates, at least, amounting to a limit of roughly two items per hemisphere (Tsubomi et al. 2013; Luria et al. 2016). It may be that these limits derive from the number of packets of information that it is possible to fit onto a gamma wave, rather than from anything more cognitive in nature (Miller & Buschman 2015).

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 such as 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.

Moreover, given the numbers, complexity, and complex relationships among the set of systems to which phenomenally conscious states are broadcast in humans, it seems almost inevitable that similarities and differences in those systems across species will be complex, multifaceted, and cross-cutting. Indeed, it seems quite 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 plainly 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.

In addition, 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. In addition, 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.

It might be replied that while the sets of systems that receive and consume globally broadcast representations admit of degrees across species, global broadcasting itself is categorically present or absent. There are two ways in which such a reply might be advanced. One would be to claim that global broadcasting is individuated by its *function*, which remains constant across species, rather than by the set of systems broadcast to. The other would be to identify global broadcasting with a core component of the network involved in humans (most likely *attention*, or perhaps *working memory*), and to claim that this component is a natural psychological kind, remaining the same, and being largely homologous, across multiple species and animal groups. Let us take these ideas in turn.

Functions can either be *current* functions ("what role does this component *actually play* in the operations of the overall system?") or they can be *evolutionary* functions ("what role was this component *selected to play*?"). Taken in the first way, global broadcasting can't be separated from the set of systems broadcast too — for all will enter into a description of what global broadcasting actually does in the human mind. But taken in the second way, it might well be the case that the adaptive function of global broadcasting can be specified without mention of the human-specific set of consumer systems. Indeed, I have argued elsewhere that the evolutionary function of the global broadcasting / global workspace network is to provide a (virtual) "center" for the mind, enabling multiple subsystems to have their activities coordinated around a single set of representations (Carruthers 2015). Mightn't one then say that any creature having such a "centered mind" can qualify as having globally broadcast — and hence phenomenally conscious — mental states?

If phenomenal consciousness can be reductively explained by the global workspace (as we are now

assuming), then the question before us is whether the workspace, for these purposes, should be identified with the abstract “centered mind” property, or rather with the human-specific set of consumer systems that we have experimental evidence of in the human case (systems for reporting, reasoning, decision making, and remembering). To answer this question, we need to examine what does the actual explanatory work in reductively explaining phenomenal consciousness. And it seems obvious that no explanatory role is played by the claim that conscious states play a “centering” role in the mental lives of people. Rather, what does the explanatory work is the fact that the contents of the global workspace are made available to systems for reflective higher-order awareness and discursive thought. For it is this that gives rise to the “hard problem” of consciousness, the conceivability of zombies, and so on.

Of course, if we had independent evidence of phenomenal consciousness in a wide range of creatures, then that might require us to focus on the centering function of global broadcasting for our explanation. But we have no such evidence — the distribution of consciousness across species is precisely what is at stake in the current discussion. The only species for which we have direct evidence of consciousness is the human species. This is hardly surprising, when we recall that phenomenal consciousness has a first-personal characterization. We have first-person access to our own phenomenally conscious states. And we have good evidence (albeit not conclusive — remember zombies!) for the presence of phenomenal consciousness in other people, given that they can talk about and reflect on the nature of their experiences.

What, then, of the suggestion that the global workspace might be identified with one of its key components (specifically attention), rather than with availability to the full suite of receiving systems? One problem for such a suggestion is 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 idea seems to become almost a notational variant of the global workspace account we have been discussing thus far. 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).⁵

One might try combining this suggestion with the previous one. Perhaps the global workspace can be identified with those attention-like networks across species (which appear to be homologous in mammals and birds at least; see Carruthers 2015 for discussion) which play the “centering” role in the species in question. This would be to identify the global workspace with a biologically-individuated natural kind. But a similar problem arises here as before: it is not *this* notion that carries the explanatory burden in the reductive explanation of phenomenal consciousness in humans sketched earlier, but rather one that implicates availability to reflective forms of discursive thinking.

Moreover, 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 concepts for substances like water *are* natural-kind ones. Even before we knew anything about chemistry, we used the concept *water* to refer to the underlying nature or essence of the recognizable stuff that fills our lakes and rivers (H₂O); 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 problem. Given that phenomenal consciousness is successfully explained in our own case by full-blown global broadcasting, what reason is there to identify it with the operations of some much-more-minimally described system? This challenge will be especially acute, since many of the remaining components of human-like global broadcasting 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?

⁵ A similar problem arises for the suggestion that presence of working memory might be sufficient for global broadcasting across species, abstracting somehow from the actual human consumer systems for the contents of working memory. For there are lots of reverberating short-term memory systems in humans and other animals that have nothing to do with consciousness.

In conclusion of this section, then, the form of global broadcasting that actually figures in the reductive explanation of the only instances of phenomenal consciousness that we have first-person knowledge of (our own) is complex, multifaceted, and multi-component. And when we look across the minds of other species of animal we will find systems that more or less resemble our own in multiple cross-cutting ways. In short: global broadcasting admits of degrees across species.

4. Phenomenal consciousness is all-or-none

In contrast to the global workspace, it seems that applications of the concept *phenomenally conscious* are all-or-none. Or so I shall now argue. (For a different but converging argument, see Simon 2017.)

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 (whether intransitive or transitive) 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, transitive creature consciousness can be more or less detailed, as well as more or less rich and 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, even if one is aware of many fewer properties of it. 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 mental 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.” (In effect, this is about degrees of transitive creature consciousness once again.) But still the inchoate experience in question was fully phenomenally conscious, it seems.

A number of recent studies of conscious experience have employed a graded perceptual awareness scale, and have used it to suggest that consciousness might admit of degrees, after all (Ramsoy & Overgaard 2004; Overgaard et al. 2006; Andersen et al. 2016; Tagliabue et al. 2016). But the scale in question arguably measures degrees of detail and clarity in perceptual *content*, rather than degrees of

mental-state consciousness. It is used in conjunction with reports of the content of experience (for example, *square* or *diamond*), and comprises the four options (1) not seen at all; merely guessing; (2) merely glimpsed; a feeling there was something there; (3) almost clear experience; partial content seen; and (4) clear experience; unambiguous content. But there is still something that it is (unequivocally) *like* to have a mere glimpse of something, even if the content of the glimpse is left wholly ambiguous (beyond a mere impression of shape of *some* sort, perhaps). Degrees of content are one thing (and are undeniably real), whereas degrees of phenomenal consciousness are quite another (and are arguably *not* real).⁶

Rosenthal (2018) argues that in addition to degrees of detail and vividness in the *content* of a perceptual state, there can be degrees of *awareness of* the occurrence of that perceptual state. One might be unsure whether one saw anything on a given trial in a backward-masking experiment, for example. If a higher-order theory of consciousness were correct, then degrees of awareness of a state would amount to degrees of conscious status for that state. But we are presently assuming a first-order account (global workspace theory). If there is no perceptual state, of course, then that state cannot be phenomenally conscious. But it is nevertheless determinately *like something* to be unsure whether or not one saw anything on the screen at the moment in question.⁷

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 three possible ways to respond: (1) give up on the identity between phenomenal consciousness and global broadcasting; (2) allow that consciousness is 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 (3) stop even raising the question of consciousness in nonhuman animals, because there is no fact of the matter.

⁶ A similar point can be made with respect to summary scene statistics. Even if a mere glimpse gives one just the impression of a peopled scene of some kind, or buildings of some sort, it is still determinately *like* something to have a visual impression of a bunch of people, even if one is aware of no further details.

⁷ In any case, even if we agreed that a fleeting indeterminate visual experience were to some lesser degree phenomenally conscious, this wouldn't help us make sense of degrees of consciousness across species. For in the latter case we may be dealing with a fully-detailed temporally-extended perceptual state, but one that is available to fewer or less sophisticated consuming systems for the broadcast information.

I will argue that #3 is the best option.

5. Give up on workspace theory?

We have a mismatch between our *concept* of phenomenal consciousness (which is all-or-nothing) and arguably 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 could 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 3. 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 2, there are grounds for thinking that the latter theory (when suitably supplemented to include acquaintance-based indexical 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 philosophers 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

phenomenal 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. Draw a (vague) boundary?

As we noted in Section 4, we can't really make sense of degrees of phenomenal consciousness. But perhaps we don't have to. An obvious suggestion is 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 *unconscious* 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 *unconscious* 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 aren't 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 3. 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 *unconscious* perceptual states. The states

involved in online and habitual forms of motor control are one thing (Lisman & Sternberg 2013; Goodale 2014; Wood & R nger 2015), whereas the perceptual states that issue in direct affect-expressive behavior are another (Panksepp 1998; LeDoux 2017); and both are quite different again from unattended states in the ventral visual and auditory systems that have a variety of priming effects independently of global broadcasting (Dehaene 2014). 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. But 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. Moreover, a recent study found that dream reports were predicted by gamma band EEG activity in prefrontal cortex (Siclari et al. 2017), suggesting that dream contents *reach* prefrontal cortex, even if the latter is suppressed and incapable of responding to them further.

Some components of global broadcasting are *epistemically* more important than 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; see Rosenthal 2005.) 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 along 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 *unconscious*. How could we decide between these possibilities in a way that doesn't just stipulate the outcome? I suggest that we can't. There is no fact of the matter concerning the relevant dimensions of comparison.

We could, however, borrow an idea from the supervenience literature on vague concepts (Fine 1975; Lewis 1982). We could say that definitely-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 well-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 their states are more similar to human *unconscious* 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 *unconscious* 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.

We can conclude that although global workspace theory can provide the basis for 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. In connection with the vast majority of species, the phenomenally conscious status of their perceptual states would have to be something that we stipulate rather than discover.

7. Doing without consciousness

It seems there may be 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 (if any) have states that are close enough to being globally broadcast in the human manner that any reasonable precisification 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 aren't, 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). And more important, "lighting up" is fully explained (we are supposing) by functional differences between the two types of state, not by the appearance of any new kind of property (qualia).

But how *could* there be no fact of the matter whether a given species of animal enjoys phenomenally conscious states or not? Granted it might be difficult — if not impossible — for us to know (Block 2002). But surely, it might be said, consciousness itself must be present or absent. In reply, recall that the concept in question is a first-person one, grounded in one's acquaintance with one's own conscious states. Then the truth-conditions for a statement of the form, "Animal X has conscious states" must be something along the lines of, "Animal X has states that are relevantly like *this-R*." If it were true that the concept *this-R* picked out a distinctive sort of property (a quale), then there would indeed be a fact of the matter. The relevant similarity would be whether or not the mind of animal X contains that sort of property (hard as this might be for us to know). But on the global workspace account outlined in Section 2, there is no such property. Rather, what first-person concepts like *this-R* pick out is globally broadcast nonconceptual content. Hence the property that we are actually asking about, when we ask whether an animal is conscious, is globally broadcast nonconceptual content. And as we have seen in Section 6, there seems no way to answer this question that doesn't involve some degree of stipulation on our part, imposed on a multi-dimensional similarity-space.

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 the animal's 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 globally broadcast ones in humans (on all reasonable ways of precisifying the

concept of global broadcasting, perhaps) than they are like any of the various kinds of unconscious state that humans undergo, 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. This, too, suggests that there is no substantial fact of the matter concerning consciousness in nonhuman animals.

8. Of Martians, mice, and men

Someone might object as follows to argument outlined above. 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 and distinguish among 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 acquaintance-based indexical concepts for 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 fails to 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 are 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 acquaintance-based indexical 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.

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 may be misconceived. Supposing that global workspace theory is our best theory of what consciousness is, 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.

9. Moral matters?

It might be said that even if the question of animal consciousness has no significance for comparative psychology, it matters greatly for the question of our ethical treatment of animals.

One way of developing this line of thought is to focus on empathy, since empathy requires first-personal identification with the feelings of the subject empathized with. To empathize with someone means imagining, in a first-person way, what that person is feeling. Since any state that we consciously imagine in this way will perforce be phenomenally conscious, it might be felt to be important to know whether it is *appropriate* to feel empathy for any given species of animal. If that animal is capable of phenomenally conscious experience, then it would be; but if it isn’t, then it wouldn’t.

Empathy can be morally problematic even in the human case, however (Bloom 2016). And it is especially problematic in connection with non-human animals (even setting aside questions about consciousness). For imagination is likely to be a highly unreliable guide to the mental life of an animal. This is because 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 empathetic 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 one’s own feelings onto it.

Indeed, it seems plain that sympathy can be independent of questions of consciousness (Carruthers

1999, 2004). To be sympathetic towards the plight of a creature is to be motivated to help it, grounded in a third-person appreciation of its needs and mental states. Then provided one cares about a given animal or type of animal, one can still feel sympathy for it, even if one is convinced that the mental states of the animal aren't phenomenally conscious ones.

Many of those who are interested in the morality of our treatment of other animals 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 may well be no determinate answer to the question of phenomenal consciousness in animals for the vast majority of animal species, as we noted in Section 6. 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. I suggest that a better approach is that of Dawkins (2012, 2017), 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 both tractable and substantive.

10. Conclusion

If the global workspace theory is correct, then it provides at least the outline of a full scientific explanation of phenomenal consciousness. For it can explain, not only the distinction between conscious and unconscious mental states, but also why we should find our own conscious states to be so

scientifically puzzling. This is because we can form acquaintance-based indexical concepts for them, ones that lack any conceptual connections to the underlying scientific facts. As a result, we can conceive of zombies, and we are led 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 entertain indexical thoughts about 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 one embraces 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 supposing we knew of those facts, asking in any given case whether there is *enough* similarity for an animal's mental states to qualify as phenomenally conscious could add nothing further to what we would already know. So this isn't a question that comparative psychologists should spend any time on. On this account, 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 Sidel, Julius Schönherr, and Moonyoung Song for comments on some of the ideas presented in this article, as well as to three anonymous referees for their criticisms and advice.

References

- Andersen, L., Pedersen, M., Sandberg, K., & Overgaard, M. (2016). Occipital MEG activity in the early time range (<300 ms) predicts graded changes in perceptual consciousness. *Cerebral Cortex*, 26, 2677-2688.
- 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.
- Ball, D. (2009). There are no phenomenal concepts. *Mind*, 118, 935-972.
- 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. & Montague, M. (eds.) (2011). *Cognitive Phenomenology*. Oxford: Oxford University Press.
- 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: Oxford University Press.
- Block, N. (1995). A confusion about the function of consciousness. *Behavioral and Brain Sciences*, 18, 227-247.
- Block, N. (2002). The harder problem of consciousness. *Journal of Philosophy*, 99, 391-425.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481-499.
- Block, N. (2010). Attention and mental paint. *Philosophical Issues*, 20, 23-63.
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Science*, 12, 567-575.
- Bloom, P. (2002). *How Children Learn the Meaning of Words*. Cambridge, MA: MIT Press.
- Bloom, P. (2016). *Against Empathy: The case for rational compassion*. New York: Ecco Press.
- Boly, M., Massimini, M., Tsuchiya, N., Postle, B., Koch, C., & Tononi, G. (2017). Are the neural correlates of consciousness in the front or in the back of the cerebral cortex? Clinical and neuroimaging evidence. *Journal of Neuroscience*, 37, 9603-9613.
- Breitmeyer, B. & Ogmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics*, 62, 1572-1595.

- Buckner, R., Kelley, W., & Petersen, S. (1999). Frontal cortex contributes to human memory formation. *Nature Neuroscience*, 2, 311-314.
- 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 & M. Montague (eds.), *Cognitive Phenomenology*, Oxford: Oxford University Press, 35-56.
- Carruthers, P. & Veillet, B. (2017). Consciousness operationalized, a debate realigned. *Consciousness and Cognition*, 55, 79-90.
- Carruthers, P. (1999). Sympathy and subjectivity. *Australasian Journal of Philosophy*, 77, 465-482.
- Carruthers, P. (2000). *Phenomenal Consciousness*. Cambridge: Cambridge University Press.
- Carruthers, P. (2004). Suffering without subjectivity. *Philosophical Studies*, 121, 99-125.
- Carruthers, P. (2015). *The Centered Mind: What the science of working memory shows us about the nature of human thought*. Oxford: 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.
- Chalmers, D. (2006). Phenomenal concepts and the explanatory gap. In T. Alter & S. Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge*, Oxford: 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.
- Dawkins, M. (2017). Animal welfare with and without consciousness. *Journal of Zoology*, 301, 1-10.
- Dehaene, S. & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79, 1-37.
- Dehaene, S. (2014). *Consciousness and the Brain*. Viking Press.
- 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*, Brighton: 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*. New York: Farrar, Straus, and Giroux.
- Goff, P. (2011). *A posteriori* physicalists get our phenomenal concepts wrong. *Australasian Journal of Philosophy*, 89, 191-209.
- 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: Oxford University Press.
- Jackendoff, R. (1987). *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press.
- Jackendoff, R. (2012). *A User's Guide to Thought and Meaning*. Oxford: Oxford University Press.
- Kripke, S. (1980). *Naming and Necessity*. Oxford: Blackwell.
- Kurzban, R. (2012). *Why Everyone (Else) is a Hypocrite: Evolution and the modular mind*. Princeton: 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.
- Lisman, J. & Sternberg, E. (2013). Habit and nonhabit systems for unconscious and conscious behavior: Implications for multitasking. *Journal of Cognitive Neuroscience*, 25, 273-283.
- LeDoux, J. (1996). *The Emotional Brain*. New York: 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.
- Luria, R., Balaban, H., Awh, H., & Vogel, E. (2016). The contralateral delay activity as a measure of visual working memory. *Neuroscience and Biobehavioral Reviews*, 62, 100-108.
- 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.
- 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.

- Marti, S., Sigman, M., & Dehaene, S. (2012). A shared cortical bottleneck underlying attentional blink and psychological refractory period. *NeuroImage*, 59, 2883-2898.
- Michel, M. (2017). Methodological artefacts in consciousness science. *Journal of Consciousness Studies*, 24: 11-12, 94-117.
- Miller, E. & Buschman, T. (2015). Working memory capacity: Limits on the bandwidth of cognition. *Daedalus*, 144, 112-122.
- Milner, M. & Goodale, M. (1995). *The Visual Brain in Action*. Oxford: 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: Oxford University Press.
- Odegaard, B., Knight, R., & Lau, H. (2017). Should a few null findings falsify prefrontal theories of conscious perception? *Journal of Neuroscience*, 37, 9593-9602.
- Overgaard, M., Rote, J., Mouridsen, K., & Ramsøy, T. (2006). Is conscious perception graded or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and Cognition*, 15, 700-708.
- Panagiotaropoulos, T., Deco, G., Kapoor, V., & Logothetis, N. (2012). Neuronal discharges and gamma oscillations explicitly reflect visual consciousness in the lateral prefrontal cortex. *Neuron*, 74, 924-935.
- Panksepp, J. (1998). *Affective Neuroscience*. Oxford: Oxford University Press
- Peters, M. & Lau, H. (2015). Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLIFE*, doi: 10.7554/eLife.09651.
- Pitt, D. (2004). The phenomenology of cognition or What is it like to think that P?, *Philosophy and Phenomenological Research*, 69, 1-36.
- 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.
- Preston, A. & Eichenbaum, H. (2013). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, 23, R764-R773.
- Prinz, J. (2012). *The Conscious Brain*. Oxford: Oxford University Press.
- Ramsøy, T. & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, 3, 1-23.
- 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.
- Rosenthal, D. (2005). *Consciousness and Mind*. Oxford: Oxford University Press.
- Rosenthal, D. (2018). Consciousness and confidence. *Neuropsychologia*.
doi.org/10.1016/j.neuropsychologia.2018.01.018
- 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.
- Siclari, F., Baird, B., Perogamvros, L., Bernardi, G., LaRocque, J., Riedner, B., Boly, M., Postle, B., Tononi, G. (2017). The neural correlates of dreaming. *Nature Neuroscience*, 20, 872-878.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton, NJ.: Princeton University Press.
- Siewert, C. (2011). Phenomenal thought. In T. Bayne & M. Montague (eds.), *Cognitive Phenomenology*, NY: Oxford University Press, 236-267.
- Simon, J. (2017). Vagueness and zombies: why “phenomenally conscious” has no borderline cases. *Philosophical Studies*, 174, 2105-2123.
- 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. (1994). *Mental Reality*. Cambridge, MA: MIT Press.
- Strawson, G. (2006). Realistic monism – why physicalism entails panpsychism. *Journal of Consciousness Studies*, 13: 10-11, 3-31.
- Strawson, G. (2011). Cognitive phenomenology: Real life. In T. Bayne & M. Montague (eds.), *Cognitive phenomenology*, Oxford: Oxford University Press, 285-325.
- Tagliabue, C., Mazzi, C., Bagattini, C., & Savazzi, S. (2016). Early local activity in temporal areas reflects graded content of visual perception. *Frontiers in Psychology*, doi: 10.3389/fpsyg.2016.00572
- 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. & Wright, B. (2011). Is there a phenomenology of thought? In T. Bayne & M. Montague (eds.), *Cognitive phenomenology*, Oxford: Oxford University Press, 326-344.
- Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Tye, M. (2009). *Consciousness Revisited: Materialism without phenomenal concepts*. Cambridge, MA: MIT Press.
- Tye, M. (2017). *Tense Bees and Shell-Shocked Crabs: are animals conscious?* Oxford: Oxford University Press.
- van Vugt, B., Dagnino, B., Vartak, D., Safaai, H., Panzeri, S., Dehaene, S., & Roelfsema, R. (2018). The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science*, 360: 6388, 537-542.
- Weiskrantz, L. (1986). *Blindsight*. Oxford: Oxford University Press.
- 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. & Runger D. (2015). Psychology of habit. *Annual Review of Psychology*, 67, 289-314.
- Wu, W. (2014). *Attention*. London: Routledge.