

Consciousness in animals

No fact of the matter

Chapters 4 and 5 argued that global-workspace theory provides the best account of the neural, functional, and representational correlates of phenomenal consciousness. Chapter 6 then argued that the theory can provide a fully satisfying reductive explanation: phenomenal consciousness just *is* globally broadcast nonconceptual content, and the “explanatory gap” arises from the distinctive acquaintance-based indexical concepts that we can employ when thinking about such contents. The present chapter takes up the implications of this account for the distribution of consciousness across the animal kingdom. It argues that there is no fact of the matter.

It is important to reiterate that our topic, here, is *phenomenal* consciousness. Of course it is true that most animals are subjects of transitive and intransitive *creature* consciousness (of varying kinds). Animals perceive the world around them, and have awake/asleep cycles just as humans do. Moreover, aspects of *access* consciousness turn out to be quite widely shared across species, also. (As we will see in Section 7.1, however, access consciousness is a matter of degree.) Our question is about the distribution of first-personal, subjective, “felt,” what-it-is-like consciousness.

7.1 Global broadcasting by degrees

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. As we noted in Chapter 5, 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). Hence global-workspace theory can fully explain the first-person intuition discussed in Chapter 1, that phenomenal consciousness is an all-or-nothing matter. Any given mental state (in humans, at any rate) is either categorically conscious or definitely unconscious.

When we look across species, however, it is obvious that global broadcasting will admit of degrees.¹ Or more accurately (since the term “global workspace” was introduced in the first instance for the human case), processes in the minds of animals will *more or less closely resemble* human global broadcasting. Recall that globally broadcast contents are made available to a wide range of cognitive systems. They are made available for verbal reporting, to give rise to higher-order awareness, to participate in executive functions of reasoning, planning, and decision-making, as well as to memory-forming systems; and they also serve to sustain or modulate full-blown affective reactions to stimuli. As we noted in Chapter 2, some of these capacities seem to be available in one form or another, or to one degree or another, in some species of animal. However, availability for verbal report is unique to human beings, of course. And as we also noted in Chapter 2, 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 capable of some form of mind-reading (perhaps only great apes, or perhaps primates more generally; or maybe extending to some other social creatures such as dolphins and elephants).

¹ 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 any rate, 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 discrete packets of information that it is possible to fit onto a gamma wave, rather than from anything more cognitive in nature (Miller & Buschman 2015).

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, remembering and 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 (as we discussed in Chapter 2), but absent in others. For example, 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. Moreover, all of the various systems that receive globally broadcast information will admit of degrees of internal complexity across species, and will likewise involve differing degrees of conceptual richness and sophistication.

In addition, 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 our human global-broadcasting architecture in some respects, whereas others resemble it in others.

It seems likely, for example, that capacities for planning might vary independently of the comparative richness of an animal’s conceptual capacities. This is because generalist feeders like raccoons and bears will need an extensive body of conceptual knowledge about the edible and inedible things to be found in their local habitats. But in general they might forage without engaging in long-term planning, taking a semi-random walk through the environment, guided by the affordances of situations that present themselves. Predators like wolves and lions that rely mostly on a single source of prey, in contrast, might possess a much more limited conceptual repertoire. But it seems they need to engage in some kinds of planning, while also having the capacity to inhibit their initial impulse to approach their prey directly.

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 neural accumulator systems) that 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.

At first glance, then, there is a mismatch between the all-or-nothing nature of phenomenal consciousness (which is explained by global-workspace theory in the human case) and the implications of global-workspace theory for animals. For animals will surely only instantiate that theory to some or other degree. Yet as we noted in Chapter 1, and then again when discussing integrated-information theories in Chapter 4, it seems to make no sense to say of the mental states of an animal that they are *to some degree* phenomenally conscious. One has to insist: either those states are conscious or they aren't; there is no possibility of an in-between.

One possible response to this mismatch would be to reject the theory. One could claim that, given such a mismatch, whatever else global-workspace theory is, it can't be a theory of phenomenal consciousness. However, there isn't any mismatch between our conception of phenomenal consciousness and global broadcasting in the human case. This is because global broadcasting in humans is all-or-nothing, as we noted above. There is only a mismatch in so far as the systems that get broadcast to vary across species. But as we also noted in Chapter 1, all investigations of consciousness have to take their start from the human case, since *phenomenal* consciousness is a first-person notion. And as we have seen, there is no conflict between first-person (human) applications of that concept and global-workspace theory. Indeed, as we saw in Chapter 6, there are grounds for thinking that the latter theory (when suitably supplemented to include acquaintance-based indexical

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. This follows from the first-person nature of the concept. Hence one could possess that 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 there would appear to be no conceptual constraints on the extension of the concept. As a result, it can't be a condition of adequacy for 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: from our point of view there are no definite, first-person-accessible instances of phenomenal consciousness that are left unexplained by global-workspace theory. 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 don't, 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 yet to repeat: phenomenal consciousness is a first-person notion. Moreover, 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.

7.2 Broadcasting as a natural kind

It might be urged that while the sets of systems that receive and consume globally broadcast nonconceptual content admit of degrees across species, global broadcasting itself is categorically present or absent. There are two ways in which this idea 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*, or both combined), 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 that are broadcast to—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 network underlying the global workspace is to provide a (virtual) “center” for the mind, enabling multiple subsystems to have their activities coordinated around a single set of representations (Carruthers 2015a). Mightn’t one then say that any creature that has such a “centered mind” can qualify as having globally broadcast—and hence phenomenally conscious—mental states?

If phenomenal consciousness can be reductively explained by global-workspace theory (as we are now assuming), then the question before us is whether the workspace, for these purposes, should be identified with the more-general “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 these that give 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 an essentially first-person characterization. We have first-person access to our own phenomenally conscious states. And we have good evidence for the presence of phenomenal consciousness in other people (albeit not conclusive evidence, given the conceivability of zombies; note that the case against zombies developed in Chapter 6 is an inference to the best explanation, not any sort of entailment). For other people can talk about and reflect on the nature of their experiences in the same kinds of ways that we do. As we noted in Chapter 1, this gives us good reason to think that they, too, have a first-person perspective on their experiences of the same sort that we do ourselves.

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).

A similar problem arises for the suggestion that the 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. Now admittedly, we know that there are working-memory-*like* systems in multiple creatures, including mammals and birds, as we saw in Chapter 2. And in some of these creatures, the contents of those systems are made available to simple executive-like systems for reasoning and decision-making. But one way of putting our present difficulty is to know whether phenomenal consciousness is confined to the contents of *actual* (human) working memory, or extends also to the contents of these working-memory-*like* systems in animals also.

One might try combining these suggestions with previous ones. Perhaps the global workspace can be identified with those attention-like networks across species that play the “centering” role in the species in question. (These appear to be homologous in mammals and birds at least, as we noted in Chapter 2; see Carruthers 2015a for discussion.) 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 outlined in Chapter 6, but rather one that implicates availability to reflective forms of discursive thinking.

Moreover, as we have noted before, one thing that everyone has agreed on since Kripke (1980) is that terms referring to phenomenally 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 (H_2O); 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 phenomenal

concepts aren't like that. We don't use them with the intention of referring to whatever natural kind underlies our 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 the qualities we are aware of in ourselves. (These turn out to be globally broadcast nonconceptual contents, if the arguments of Chapters 5 and 6 are correct, but they aren't conceptualized as such when we deploy concepts like *THIS-E*.)

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 or natural kind? This challenge is 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 mind-reading 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? There would appear to be nothing in our first-person conception of phenomenal consciousness that rules this out.

7.3 Stipulating a categorical boundary

As we noted in Chapter 1, we can't really make sense of degrees of phenomenal consciousness. But perhaps we don't have to. An obvious suggestion is to say that a creature will enjoy phenomenally conscious experience if it undergoes states that are more similar to human global broadcasting than they are similar to any form of human *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-person 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; but we can

set this aside). 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 yellow.) But it is also because of the complex and cross-cutting nature of the similarity space, noted in Section 7.1. 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 evidence from the human case that some of the consumer systems for global broadcasts 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" thought experiments. So that would leave us with the broad class of higher-order thought, semantic-memory, and executive-function systems (including verbal report) as the relevant components of global broadcasting when we make comparisons across species.

It might be claimed that executive systems aren't necessary for consciousness, either, however. For dreams are phenomenally conscious states, but during dreaming, one's capacities for executive function are suppressed. It remains unclear, however, 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, as we noted in Chapter 5. Moreover, a recent study found that dream reports were predicted by gamma-band EEG activity in the prefrontal cortex (Siclari et al. 2017), suggesting that dream contents *reach* prefrontal cortex, even if the latter is largely suppressed and incapable of responding to them further.

What if one could completely knock out capacities for language and higher-order thought on a temporary basis, however? (Perhaps one might use some future and more effective form of transcranial magnetic stimulation; or one might employ some future sort of invasive but temporary chemical or thermal enervation of the neural systems involved.) If, on recovery, people reported that they continued to enjoy phenomenally conscious experience, then wouldn't that demonstrate that language and reflective thought aren't necessary for consciousness? No, it would not. For these reports would, of course, be retrospective ones, grounded in memory of one's perceptual contents while capacities for reporting were knocked out. But in order to report the contents of a memory, that memory needs to be activated and globally broadcast. It thereby becomes available to linguistic report and higher-order thought. So on the hypothesis that it is these, specifically, that are necessary for a content to become phenomenally conscious, those memory contents will thereby *become* phenomenally conscious, even if they hadn't been when originally undergone.

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 most direct evidence that we have of phenomenally conscious experience in other people. But on a global-workspace account, these capacities aren't *constitutive* of phenomenal consciousness. (That would involve a shift to some kind of higher-order-thought theory of consciousness, discussed in Chapter 4.) 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 manner that doesn't just stipulate the outcome? I suggest that we can't. Given that there are no real qualia (as we argued in Chapters 1 and 6), there is no fact of the matter concerning the relevant dimensions of comparison.

We could, however, borrow an idea from the supervaluation literature on vague concepts (Fine 1975; Lewis 1982). We could say that definitely conscious states are ones whose similarity to human globally broadcast states 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 sort of maneuver might well imply that *some* other animals would qualify as 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.² So there will already be a considerable gulf between the nature of human and

² Indeed, the size of prefrontal cortex in humans is *proportionately* greater than in other primates (Donahue et al. 2018). Yet the differences between human brains and those of our nearest relatives, the chimpanzees, are already quite large. This is true whether one measures absolute brain-size (the human brain is about three to four times larger), or considers the encephalization quotient (which corrects for body size), which is also more than three times as great (Roth & Dicke 2005). Recall from Chapter 2, too, that executive inhibition capacities among primates in the A-not-B task seem roughly equivalent to those of a ten-month-old human infant.

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 tens if not 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.³

We can conclude that although global-workspace theory *could* 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) would 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.

³ Note that the argument here isn't an argument from authority. I am not saying that the restrictions in question are supported because famous people have endorsed them. Rather, the point pertains to the supervaluation process that can be used to fix definite truth-conditions for the vague concept of global broadcasting. This defines definite truth as what would survive all possible ways of sharpening the concept. The views in question provide evidence that there are reasonable sharpenings that would exclude most if not all species of animal.

7.4 The negative semantic argument

Section 7.3 has sketched how one might construct a categorical concept of global broadcasting centered on the human case. It should be emphasized that such a concept would need to be *constructed*, however. It isn't somehow implicit in our conception of phenomenal consciousness. On the contrary, the latter concept is first-personal, as we have stressed throughout. Although it applies to states that are, as a matter of fact, globally broadcast ones (given the truth of global-workspace theory), there is nothing in one's first-person use that entails that the consciousness concept should also have in its extension any states that are more similar to human globally broadcast states than to human unconscious ones. Likewise, as we pointed out in Section 7.2, there is nothing in our first-person conception of phenomenal consciousness that ties the extension of the concept to one or another natural psychological kind underlying global broadcasting. In fact, there is nothing about our first-person concept of phenomenal consciousness that fixes its extension across the mental states of other species. Hence there is no fact of the matter whether members of any those species have phenomenally conscious states or not.

Call this "the negative semantic argument." There is nothing in the intentions with which first-person phenomenal concepts like *THIS-E* or *THIS-R* are employed that settles how similar the functional role of an animal's perceptual states must be to human global broadcasting to count as "the same," or that even constrains our answers. For, as already noted, the concepts one uses to think about one's own conscious experiences aren't natural-kind ones, and when we think about our own experiences in this sort of distinctive first-person way we aren't intending to designate the underlying explanatory property, whatever that might turn out to be.

Note that the negative semantic argument isn't an argument from *vagueness*. (Nor, of course, is the claim that attributing consciousness to animals is some sort of category mistake.) The claim isn't that there fails to be a fact of the matter about phenomenal consciousness in animals because their states fall somewhere within the vague boundary of our concept of consciousness. On the contrary, our phenomenal concepts are *sharp*, in the sense defended in Chapter 1—any candidate mental state of

ours is either definitely phenomenally conscious and is *like something* to undergo, or it definitely isn't. Nevertheless, the argument does turn on a kind of semantic indeterminacy. For I have argued that there is nothing in our first-person concept of phenomenal consciousness that fixes how physically or functionally similar an animal must be to the human case to qualify as conscious.

But how *could* there be no fact of the matter whether another animal has phenomenally conscious experiences or not? How is this even possible? Recall that we are now 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 are more similar to some variety of human unconscious state than they are to human globally broadcast states, through to species for which the reverse is true, 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 closer toward the unconscious end of the spectrum, or two species whose states fall closer to the human-global-broadcasting end. 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 phenomenally 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 *constitutive* of 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).

Suppose we had complete knowledge of the functional and representational components of the mentality of a monkey, or a salmon, or a

honey bee. Suppose we knew everything of a third-person sort that there is to know about what happens to the animal's perceptual states under various conditions, how those states interact with its 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 similar to 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. As a result, there is no substantive fact of the matter concerning consciousness in nonhuman animals.

7.5 The positive semantic argument

Section 7.4 presented a negative semantic argument for the claim that there is no fact of the matter about animal consciousness. That argument turned especially on claims about what our concept of phenomenal consciousness is *not*. The present section will develop a second way of coming to the same conclusion. This is the positive semantic argument, which turns on a positive account of the truth-conditions of judgments attributing phenomenal consciousness to other creatures.

Recall, again, that the concept of phenomenal consciousness is a first-person one, grounded in one's capacities for thought about one's own globally broadcast perceptual states. For instance, while undergoing an experience of redness one might think about it using the acquaintance-based indexical concept *THIS-R*. Or when contemplating one's own mortality, one might think, "When I die, all *this* will disappear," meaning to refer to one's phenomenal experience in general. Here one employs an acquaintance-based indexical concept such as *THIS-E*.

What, then, are the truth-conditions of thoughts employing concepts of this kind when thinking about the mental states of an animal? Since the concepts are first-person ones, the truth-conditions will presumably

have to take a first-person-based form, somewhat along the lines of, “Species X undergoes perceptual states of the *same sort* as *this-E*.” If one were a realist about experiential qualia, then it would be clear what “same sort” amounts to here: it would mean that there are qualia properties present in the mind of the animal in question. And there would be a fact of the matter (whether knowable by us or not) concerning whether or not that is true. But on the view defended in Chapters 5 and 6, in contrast, *THIS-E* picks out globally broadcast nonconceptual content in general. And then the question is, what counts as *sameness of sort* in fixing the truth-conditions of one’s thoughts about the animal?

Plainly, the sameness relation can’t just be a matter of sameness of nonconceptual content. This is because unconscious perceptual states, too, have nonconceptual content—including color content. Now admittedly, perceptual states in the dorsal visual (sensorimotor) system are thought to be coded in limb-centered and body-centered coordinates, whereas conscious visual contents are generally allocentric (Milner & Goodale 1995). But there can be unconscious perceptual states in the ventral visual system too (Dehaene 2014). And other forms of conscious experience can be coded in body-centric ways, as are proprioceptive experiences of the position and movements of one’s own body-parts. In consequence, the sameness relation that forms our target must somehow implicate the functional roles of the contents in question—which is to say, global broadcasting.

In Sections 7.2 and 7.3 we noted that concepts like *THIS-E* aren’t used as natural-kind terms, nor with the intention of tracking the underlying explanatory properties, whatever they turn out to be. Nor, of course, are they used with the intention of deferring to the practices of other speakers, as we noted in Chapter 6. Phenomenal concepts aren’t taught to us, and aren’t directly part of public discourse. Rather, as we noted in Chapter 1, they emerge out of first-person acquaintance with one’s own experiential states. They are, in a sense, private, in that they aren’t beholden to other people’s judgments or usage, and their truth-conditions aren’t fixed by public usage-norms. Rather, when we communicate about phenomenal consciousness, we do so by drawing the attention of other speakers to states within themselves.

Consider an analogy, in this case with a first-order concept, but one that plainly isn’t intended to track any natural kind and is also (in the

required sense) private, not intended to defer to the usage of others. Moreover, it is a concept whose extension hasn't been fixed by previous usage or public norms. While looking around neighborhoods in an unfamiliar city to which one is relocating, one might exclaim, "Now *this* is the sort of place we should live!" In the case I imagine, the referent of the term "this sort of place" isn't a single feature (like houses with picket fences), nor would one be able to articulate all the factors that enter into one's judgment. And plainly one isn't intending to refer to the underlying essence of the place, whatever that might turn out to be. (One isn't using "this sort of place" as a natural-kind term.) Now consider another neighborhood in the same city, and ask, "Is that neighborhood of *this sort*?" And if one asserts that it is, what fixes the truth-condition for one's judgment?

The answer to the latter question is surely something like this: that if one were to visit that other neighborhood under the same conditions with the same precursors (thus creating the same set of dispositions-to-judge) then one would (or would not) count that other neighborhood as the same in the relevant respect. Or, somewhat more abstractly: if someone with the exact same dispositions-to-judge that were involved in one's judgment about the first neighborhood were to be exposed to the second, then the second neighborhood would evoke a judgment that it belongs to the relevant sort.

The truth-conditions of thoughts employing phenomenal concepts that are about the mental states of other beings must be something like this. The meaning of the concept is fixed in the first-person case (just as the meaning of "this sort of place" is fixed by an individual thinker in a specific neighborhood), but is then applied beyond it. By analogy, then, the truth-condition of a judgment like, "Creature C has perceptual states that are *this-E*" might be this: "If I were to be aware of creature C's perceptual states, then I would judge them to be *this-E*." Or more plausibly (since the antecedent of this conditional is impossible to fulfil), this: "If the dispositions-to-judge that underlie my use of the concept *THIS-E* were to be instantiated in creature C, then they would issue in a judgment that some of the creature's perceptual states are *this-E*." Indeed, it is hard to see any other way in which the truth-condition for a phenomenal judgment about another being could be fixed (given qualia irrealism, of course).

On this account it is fairly easy to understand how one can think that another human has a color experience of the same sort as *this-R*, or that another human is phenomenally conscious more generally, undergoing states that are *this-E*. And there is surely a fact of the matter in the human case. One just has to suppose that classificatory dispositions exactly like those underlying one's use of the concepts *THIS-R* or *THIS-E* might exist in the mind of another. Moreover, the behavior of other people (especially their verbal behavior when talking about their experiences, expressing puzzlement about the explanatory gap, and so on) can provide us with sufficient evidence that one's thoughts about phenomenal consciousness in others are true. But when extended to the case of other animals, the corresponding counterfactual is no longer evaluable. And in that case, I suggest, it must be neither true nor false.

To suppose that one's own classificatory dispositions underlying a phenomenal concept like *THIS-R* or *THIS-E* might exist in the mind of an animal is to suppose that the animal has a completely different kind of mind from the one that it *actually* has—it attributes to the animal a mind capable of formulating concepts referring to its own experiences and engaging in reflective thought about them. But the intended target of one's thought, of course, is the mind of the animal as it actually is. Suppose one thinks that macaque monkeys, for instance, have states that are *this-E*. Here one is making a judgment about macaque minds as they actually are. But the truth-condition of this thought is that, if the dispositions underlying *my* use of *THIS-E* were instantiated in the mind of the monkey, then those dispositions would be evoked by the monkey's states. But the antecedent of this conditional entails that I am no longer thinking about a monkey's mind, but the mind of another sort of creature altogether, with cognitive capacities like those of humans. As a result, there is no fact of the matter about how human first-person concepts of experience would or would not apply if instantiated in the mind of the animal.

Not all counterfactuals that suppose a target entity to be other than it is are unevaluable, of course. One can judge the truth of a statement like, "If 260-pound John were to weigh 160 pounds, then he would be capable of climbing the stairs more quickly." The problem arises when one is using the truth of the counterfactual to determine some *further* truth about the thing mentioned in the antecedent as it *actually* is. If we want

to know, for example, whether someone weighing 260 pounds is in principle capable of running a four-minute mile, then offering the following counterfactual is plainly of no help: "If someone weighing 260 pounds were to weigh 160 pounds, then he would in principle be capable of running a four-minute mile."

In the case just described, of course, we think there is some fact of the matter whether a person weighing 260 pounds could run a four-minute mile. This is because the truth-condition for the statement, "Someone weighing 260 pounds can in principle run a four-minute mile" are fixed independently of the counterfactual proposed (mistakenly) to help one adjudicate its truth. In the case of statements about phenomenally conscious experience in beings who only satisfy global-workspace theory to some *degree*, in contrast, the counterfactual in question is precisely the one we need to fix the truth-condition for the statement. For as we noted above, the truth-condition for the statement, "Creature C undergoes states that are *this-E*" (i.e. creature C is phenomenally conscious), is given by the counterfactual, "If the disposition-to-judge underlying my use of the concept *THIS-E* were to be instantiated in the mind of creature C, then C would judge that some of its states are *this-E*." Here it is none other than the counterfactual needed to fix the truth-condition for an attribution of phenomenal consciousness to an animal that at the same time requires the mind of the creature in question to be other than it is.

Suppose this argument is resisted, however. Suppose someone were to insist that it is allowable for the counterfactuals that fix the truth-conditions for thoughts like, "Creature C is *this-E*" or, "Creature C is *this-R*" to require us to suppose that the mind of C is quite other than it actually is. In that case the truth-value of those judgments when applied to a macaque monkey, say, would presumably be *true*. For it is presumably true that if the monkey's attended percept of a red tomato were presented to a disposition-to-judge just like that underlying my use of the concept *THIS-R*, then that percept would be judged to be *this-R*.

But now this line of thought can be reduced to absurdity. For it is equally true that if an *unattended* percept of red in my own ventral visual stream were presented to a disposition-to-judge like that underlying my use of the concept *THIS-R*, then it would be judged to be *this-R*. And in the same way, it will be true that if a nonconceptual content in my own dorsal—sensorimotor—visual stream were presented to a

disposition-to-judge like that underlying my use of the concept *THIS-E*, then it would be judged to be *this-E*. But these are paradigmatic instances of perceptual states in myself that *aren't* phenomenally conscious. It is plainly illegitimate to use counterfactuals to project my own phenomenal concepts into unconscious regions of my own mind. For doing so presupposes that those unconscious contents have a role other than they do. By the same token, then, it must be illegitimate to project my own phenomenal concepts into minds that are significantly different from my own. For doing so presupposes that the minds in question are other than they really are.

We can conclude, then, that there is no fact of the matter about phenomenal consciousness in most animals, at least. This is because the counterfactuals that fix the truth-conditions for our first-person concept of phenomenal consciousness are unevaluable when extended to minds that are too much unlike our own. But how much unlike is *too* unlike? What degree of dissimilarity must there be, for the counterfactual that fixes the truth-condition for statements attributing phenomenal consciousness to others to become unevaluable? The answer depends on what set of capacities underlies our use of the phenomenal concepts that are employed to formulate those counterfactuals. Plainly, since such concepts are about one's own experiences, they require capacities for higher-order thought. But do they also require capacities for language and for executively controlled sequences of reflective thinking?

If capacities for higher-order thought were the only ones necessary for phenomenal concepts to be possible, then that would imply that there is a fact of the matter about phenomenal consciousness in all creatures capable of such thought. But in fact the *kinds* of higher-order-thought capacities that are required for one to have phenomenal concepts are rather special ones. Regular higher-order concepts are deeply embedded in predictive and explanatory networks linking mental states with circumstances and behavior. This will surely be true of the higher-order concepts possessed by other primates; and it is equally true of human higher-order concepts, except when we do philosophy and become puzzled by "hard"-problem thought experiments. Recall that in order to form a phenomenal concept one has to set aside one's regular set of concepts for mental states of various sorts, focusing just on a globally broadcast nonconceptual content that one is acquainted with (but not

described as such, of course), and formulating a thought with the intention of referring just to that. These capacities appear to go well beyond those that are required for higher-order thought in general.

In fact, the dispositions underlying one's phenomenal concepts would seem to require capacities for intentionally controlled concept formation and reflective thinking. And they might well require linguistic ability as well. This will be true if, as seems plausible, the capacity for this kind of controlled reflective thought is language-dependent. And if so, then the positive semantic argument outlined in this section will imply that there is no fact of the matter about phenomenal consciousness in any creature outside of the human species. (The question of phenomenal consciousness in human infants and young children will be discussed in Chapter 8.) But in the absence of a more careful analysis of the cognitive prerequisites of phenomenal concepts than I can provide here (as well as a more detailed account of the cognitive powers of animals than I was able to give in Chapter 2), the issue of the precise extension of cases where there is a fact-of-the-matter about phenomenal consciousness must be left undecided. What we can be quite confident of, however, is that the extension in question is a highly restricted one.

7.6 A diagnosis

I have claimed that there is no fact of the matter about the presence or absence of phenomenally conscious states in most if not all nonhuman animals. This should come paired with a diagnosis for why almost everyone who has addressed the topic previously has assumed that there *is* a fact of the matter, whether known or unknown. Why does everyone assume that there is a substantive question here with a determinate answer, if in reality (as I claim) there isn't?⁴

⁴ One exception to this generalization is Papineau (2003), who develops an argument similar to the negative semantic argument of Section 7.4. He claims that there is nothing in our use of the first-person concept of phenomenal consciousness that fixes its extension beyond the human case, concluding that there is no fact of the matter whether the mental states of other creatures count as instances of that concept. Moreover, Papineau's diagnosis for why people naturally assume that there *is* a fact of the matter is also quite similar to my own. However, the indeterminacy he focusses on concerns the right level of description of our conscious states—functional, neuronal, chemical, and so on—whereas my own argument operates at the

Part of the answer is that many people are qualia realists of one sort or another. If our phenomenal concepts pick out realworldly properties of some kind, that don't reduce to functional and representational ones, then there must be facts about how those properties are distributed across the animal kingdom. Moreover, even those people who are officially qualia irrealists might find it hard to shake the intuition that something "lights up" with the arrival of phenomenal consciousness in the world, and hence that there should exist answers to the question where in the world that something is to be found. Indeed, I count my own past self in this category. For recall from Chapter 4 that I once objected against a form of global-workspace theory that it doesn't really *explain* why first-order contents should "light up" and acquire a subjective dimension when globally broadcast.

The deeper and more general answer, though, is that I suspect people lose sight of the first-person character of the concept of phenomenal consciousness. We (some of us) talk about phenomenal consciousness on a regular basis, of course, and hence deploy public, shared concepts that we use in the course of such talk. This is, in effect, an alternative—third-personal—concept of phenomenal consciousness. It treats consciousness as whatever properties of experience give rise to thought experiments of the "hard problem" type. We naturally assume that this talk, like almost all other public discourse, issues in thoughts that have definite truth-conditions (albeit vague ones, perhaps). It is easy to forget that the point of this talk is just to draw attention to our own conscious experience, and that the critical concepts that get the problem of consciousness going in the first place are first-person ones. It is these that can't issue in determinate truth-conditions when projected across species.

Indeed, it is worth noting that although I have presented both the negative and positive semantic arguments given in Sections 7.4 and 7.5 within the framework of global-workspace theory, versions of those arguments could equally well be employed by any theorist who takes

functional level itself. It turns on indeterminacy in the set of systems globally broadcast to that are necessary for conscious experience. My view is that this is the right level at which to explain consciousness, as I argued in Chapter 6; hence differences in underlying physical constitution that make no functional or representational difference are irrelevant to the question of consciousness.

seriously the first-person character of phenomenal consciousness while denying the existence of qualia. For example, consider someone who endorses the fragile-short-term-memory view discussed in Chapter 4, but who thinks (as Block 1995 actually does), that the public concept PHENOMENAL CONSCIOUSNESS is intended just to draw people's attention introspectively to their own experiential states. Then there will be nothing in the use of that concept that prioritizes one rather than another of the physical or functional factors that accompany fragile-short-term-memory contents in humans (the negative semantic argument). And the truth-conditions for thoughts employing that concept across species will require projecting one's own conceptual dispositions into the mind of the animal in question while nevertheless trying to hold the nature of that mind fixed (the positive semantic argument).

7.7 Conclusion

Our first-person concept of phenomenal consciousness is an all-or-nothing one; and applications of that concept in our own case are fully explained by global-workspace theory, as we saw in Chapters 5 and 6. But when we shift our gaze to the question of consciousness across species, we confront a mismatch. While our concept of phenomenal consciousness is all-or-nothing, the minds of other animals will only resemble the human global workspace more or less closely, to some or other *degree*. Moreover, because that concept is a first-person one, it can find no determinate application when applied to minds that are significantly unlike our own. The upshot, I have argued, is that there is no fact of the matter whether animals have phenomenally conscious states or not.

Let me reiterate what I emphasized at the outset of the chapter: the upshot here pertains just to the question of *phenomenal* consciousness in animals. It is important that phenomenal consciousness shouldn't be confused with any of the other uses of the term "conscious" that we distinguished in Chapter 1. For of course it is definitely *true* that many animals have perceptual states that are (to some degree) *access-conscious*. And of course it is *true* that animals are often conscious rather than unconscious (i.e. they are sometimes *intransitively creature-conscious*; they are sometimes awake). Moreover, it is undeniably true that they are

often conscious *of* objects and events in their environment (i.e. they are frequently *transitively creature-conscious* in some form or another; they perceive the world). My claim is just that the first-person concepts that give rise to the “hard problem” thought experiments (phenomenal concepts) can’t be projected into the minds of nonhuman animals except by stipulation (although they *can* be projected into the minds of other people).

If one embraces the fully reductive nature of global-workspace theory, however, then the question of animal consciousness is no longer of any factual significance. There are many facts to be discovered about the minds of nonhuman animals; and among those facts will be dimensions of similarity to, and difference from, the human global workspace. But supposing we knew all of those facts, inquiring 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 already knew. Nor would trying to figure out what judgments the dispositions underlying our own phenomenal concepts would issue in if transplanted into the mind of the animal. So these aren’t questions that comparative psychologists should spend any time on. What I suggest, then, is that the nature and functioning of animal minds should be studied; but phenomenal consciousness in animals doesn’t deserve to be. This is because there are no further facts to discover.