

PETER CARRUTHERS

SUFFERING WITHOUT SUBJECTIVITY

ABSTRACT. This paper argues that it is possible for suffering to occur in the absence of phenomenal consciousness – in the absence of a certain sort of experiential *subjectivity*, that is. (“Phenomenal” consciousness is the property that some mental states possess, when it is *like* something to undergo them, or when they have subjective *feels*, or possess *qualia*.) So even if theories of phenomenal consciousness that would withhold such consciousness from most species of non-human animal are correct, this needn’t mean that those animals don’t suffer, and aren’t appropriate objects of sympathy and concern.

1. INTRODUCTION

What sense can be made of the idea of *suffering* – pain and/or grief and disappointed desire – in the absence of phenomenal consciousness? More generally, can we understand the idea of a subjective harm whose harmfulness doesn’t consist in or involve the subject’s possession of phenomenally conscious mental states? (By a “subjective” harm I mean an event whose harmfulness in one way or another derives from its causal impact upon the subject’s on-going mental life, as opposed to such things as objective damage to health, or the merely objective frustration of desire. I shall say more about this below.)

Carruthers (1999) argues that the most basic kind of subjective harm is the first-order (non-phenomenal) frustration of desire. The fundamental form of harm, on this view, consists in the co-activation within a creature’s practical reasoning system of a first-order desire together with the first-order belief that the state of affairs that is the object of the desire doesn’t obtain. (Example: an animal that currently wants to drink, and believes that it is not presently drinking.) The argument to this conclusion deploys two lemmas: first, that the harm of unpleasant



sensations like pain consists, not in the sensation itself, but in the frustration of the subject's desire to be rid of that sensation; and second, that a creature that experienced no felt disappointment when its desires were frustrated would still be an appropriate object of sympathy and concern. Each of these two lemmas is controversial, and would be resisted by many. And the sort of harm in question is an extremely *thin* one, being confined to the bare frustration of desire, independent of any resulting emotion of disappointment.

The goal of the present paper is to find an alternative route to the same overall conclusion. Namely, that there could well be subjective harms in the absence of phenomenal consciousness, sufficient to warrant sympathy. But our focus, here, will be on harms like pain and grief, that certainly deserve to be described as kinds of *suffering*. I shall be arguing that we can make sense of the idea that suffering, too, can occur in the absence of phenomenal consciousness.

These questions are of theoretical interest in their own right. For what is at issue is the way in which subjective harms and the appropriate objects of sympathy should properly be understood and characterized. But our questions also possess some practical bite. This is because there are theories of the nature of phenomenal consciousness that might withhold phenomenal consciousness from many species of non-human animal, as we shall see in the section that follows. Before we come on to that, however, something more needs to be said to delimit our overall target: subjective harm.

There is a perfectly respectable sense in which plants can be harmed, of course. Lack of water or sunshine will generally harm them, as will the application of a herbicide to their foliage. While it can be appropriate to feel a kind of *concern* at harms resulting to a plant (if one cares about the plant's health and flourishing) it would seem that *sympathy* is out of place here. While people do sometimes express sympathy for harm done to a plant ("Oh, you poor thing, your leaves are all withered!"), this seems most plausibly interpreted as mere anthropomorphic whimsy. Only a being that is a subject of mental states is an appropriate object of sympathy, surely. I

propose to say, then, that a subjective harm is a harm that has some sort of impact upon the mental life of a subject.

(Animals and humans can *also* be harmed in the same physical/biological sense that is applicable to plants, of course. Animals, too, can be deprived of what is necessary for normal growth, health, and physical flourishing. But such harms will, almost inevitably, have an impact on the animal's mental life as well. A sick animal – in contrast with a sickly plant – is generally a suffering animal, or an animal that cannot achieve some of the things that it wants, either now or in the future. When we feel sympathy at an animal's sickness, it is really these further effects that we have in mind, I suggest.)

What *sort* of impact must something have on the mental life of a subject in order to count as a subjective harm, however? Here we can distinguish two possibilities: causal and a-causal. On the causal account, subjective harms are events that either cause a creature's goals to be subjectively frustrated (this is the thin sense of "harm" at issue in Carruthers, 1999), or that in one way or another cause a creature to *suffer*, either by causing it pain, or by causing it to undergo negative emotional states such as terror, grief, or disappointment. On the a-causal account, the mere fact that the things that an animal desires don't occur (as a matter of objective fact and independently of the animal's beliefs) is sufficient to constitute a harm. Thus many have claimed that what one might call "objective frustrations of desire" are a species of harm (Feinberg, 1977; Nagel, 1979). On this view, I am harmed if my wife cheats on me behind my back, for example, even if I never know and nothing in our relationship ever suffers as a result. For my desire that she should be faithful to me has not, as a matter of fact, been fulfilled (even if I continue to believe that it has).

This is a debate that I don't propose to enter into here, important as it is for our understanding of the appropriate objects of sympathy. One reason is that I have discussed it at some length elsewhere (Carruthers, 1992, 1999), arguing that only things that in one way or another have a negative impact on a creature's mental life (whether directly or indirectly, immediately or in the long run) should count as harms. But

more importantly, if we were to opt for the a-causal construal of harm and the appropriate objects of sympathy, then our question about the implications of an absence of phenomenal consciousness for sympathy would immediately be foreclosed. For even if a creature lacks phenomenally conscious mental states, provided that it nevertheless has desires, then those desires will be capable of being objectively frustrated. And then it would be appropriate to feel sympathy for the frustrated desires of an animal, whether or not that animal is ever phenomenally conscious.

In what follows I propose to take for granted that subjective harms to an animal require some sort of causal impact upon that animal's mental life, because making such an assumption is necessary to get my topic started. If this assumption should prove false, then that will just mean that there is yet another reason for thinking that sympathy for animals in the absence of phenomenal consciousness can be appropriate, that is all. More narrowly, I shall be assuming that harm to a subject (in the sense relevant to sympathy and moral concern) means causing the subject to *suffer*. My question is whether there can be states of suffering that aren't phenomenally conscious ones; and if there can, whether creatures that possess such states are appropriate objects of sympathy.

2. HIGHER-ORDER THEORIES OF PHENOMENAL CONSCIOUSNESS

Representationalist theories claim that phenomenal consciousness can and should be reductively explained in terms of some or other form of *access* consciousness. (An intentional state is access-conscious when it is available to, or is having the right kinds of impact upon, other mental states of the subject of the required sort; see Block, 1995.) *First-order* theories of the kind defended by Dretske (1995) and Tye (1995, 2000) claim that phenomenal consciousness consists in a certain sort of intentional content (*analog*, or *non-conceptual*)¹ being available to the belief-forming and decision-making systems of the

creature in question. Some higher-order theories claim, in contrast, that phenomenal consciousness consists in the targeting of analog first-order perceptual states by *higher-order* states – either by higher-order thoughts (Rosenthal, 1993), or by higher-order experiences (Lycan, 1996).

In the version of higher-order account developed and defended by Carruthers (2000), perceptual states (and more generally, all states with analog content such as mental images, bodily sensations, and emotions) become phenomenally conscious when they are available to a faculty of higher-order thought which is capable of framing beliefs about those very states. Because of such availability, and in virtue of the truth of some or other version of “consumer semantics”, the states in question acquire a dual analog content. (All forms of consumer semantics claim that the intentional content of a state depends, in part, on what the “down-stream” systems that consume the state are capable of doing with it or inferring from it. Teleosemantics is one form of consumer semantics; see Millikan, 1984, 1989; Papineau, 1987, 1993. Functional or inferential role semantics is another; see Loar, 1981; Block, 1986; McGinn, 1989; Peacocke, 1992.) Thus one and the same percept is *both* an analog representation with the content *red*, say, *and* an analog representation of *seeming red* or *experience of red*.²

Carruthers (2000) claims that such an account is warranted, because it can successfully explain all of the various supposedly puzzling features of phenomenal consciousness. Briefly (and by way of motivation for the discussion that follows), the account can explain how phenomenally conscious experiences have a *subjective aspect* to them; this is their higher-order analog content. These higher-order contents are the *mode of presentation of* our own experiences to us, rendering the latter subjective in something like the way that worldly properties acquire a subjective dimension via the modes of presentation inherent in our first-order perceptual contents.

The account can also explain how phenomenally conscious experiences can be available to us for introspective recognition, since their higher-order contents will present them to us in much the same sort of way that our first-order perceptions

present to us a world of colors and shapes. And it can explain how we can come to have *purely* recognitional concepts of our phenomenally conscious experiences. These will be recognitional concepts grounded in the higher-order analog contents of our experiences, similar to the way that first-order recognitional concepts of color, say, are grounded in the first-order analog content of perception, only without any of the surrounding beliefs about the mechanisms by means of which such recognitions are effected. Moreover, the account can explain why phenomenally conscious properties should seem especially *ineffable*. For when we deploy purely recognitional concepts of them they cannot be further described, and – unlike public colors and textures – they cannot be exhibited to others. And it also explains why there should seem to be a pervasive *explanatory gap* between the higher-order account itself and the qualities of our phenomenal experiences – again, the blame falls squarely on our purely recognitional concepts. (See Carruthers, 2000, 2004a, b for extended discussion.)

If some or other first-order account of phenomenal consciousness is correct, then this kind of consciousness will be widespread in the animal kingdom, and will perhaps even be present in insects (Tye, 1997). If the above form of higher-order account is acceptable, in contrast, then it seems unlikely that many animals besides ourselves will count as undergoing phenomenally conscious experiences. It is still hotly debated whether chimpanzees are capable of higher-order thought, for example, and powerful arguments can be adduced to the contrary (Povinelli, 2000). And if chimpanzees *aren't* so capable, then probably *no* animals besides ourselves are phenomenally conscious, if a higher-order thought account of phenomenal consciousness is correct. I should stress, however, that I have no axe to grind here. If chimpanzees or other animals should turn out to have higher-order thoughts, then all well and good; they will turn out to be phenomenally conscious also. The point is just that for a higher-order thought theorist, the capacity for phenomenal consciousness is conditional on a capacity for higher-order thinking; and the latter capacity is unlikely to be widespread amongst non-human animals.

This is the background theoretical framework that gives the practical bite to our question whether suffering is possible without phenomenal consciousness. For on the answer to this question might turn the appropriateness of sympathy for non-human animals. I should emphasize that my goal in the present paper isn't to *defend* a higher-order account of phenomenal consciousness, however. It is rather to explore what room might be left by such an account for forms of suffering that aren't phenomenally conscious in character.

3. EXPLAINING COLOR EXPERIENCE

My goal over the next four sections will be to approach the question whether pains that aren't phenomenally conscious might nevertheless be appropriate objects of sympathy and concern. In this discussion I shall set to one side the suggestion (made in Carruthers, 1999) that what really makes pain bad is the subject's frustrated desire to be rid of the pain. So we may as well assume that pains are somehow intrinsically aversive, from the perspective of the subject. I need to approach my goal in stages, however, beginning with discussion of a more neutral case: that of color perception.

To a first approximation, color percepts are analog first-order representations of the surface properties of distal objects. Such perceptual states are analog in the sense of being fine-grained – our perceptions of color are characteristically much finer grained, admitting of many more distinctions amongst hues, than any concepts that we can form and recall. This is widely, if not universally, agreed upon.³ But quite how the intentional content of our color perceptions should be characterized is very much in dispute. On the sort of externalist view championed by Tye (1995), color percepts are analog representations of the physical reflectance-determining properties of surfaces. However, it is also possible to defend an internalist account, according to which the content of the percept is to be characterized in abstraction from its normal worldly causes, as well as one that is non-reductive, where we would characterize

the content of a percept of red, say, by saying that it represents *this* property (exhibiting a red surface).⁴

These disputes needn't detain us here. For present purposes, the point is just that there is nothing in the notion of a color percept, thus characterized, that requires such perceptions to be intrinsically phenomenally conscious. This is all to the good, since there is widespread evidence of non-conscious visual perceptions in general, and of non-conscious color perceptions in particular (Milner and Goodale, 1995; Weiskrantz, 1997). And of course a first-order theorist, too, will want to allow for the possibility of color percepts that aren't phenomenally conscious, where these would be representations with the right sort of analog content, but that aren't available to the required kinds of belief-forming and decision-making processes. Similarly, a higher-order theorist such as myself will claim that color percepts are only phenomenally conscious when they are non-inferentially available to a faculty of higher-order thought, and hence acquire at the same time a higher-order analog content of *seeming color* (Carruthers, 2000).

Many mammals and birds are capable of color perception, of course. But according to higher-order accounts of phenomenal consciousness, their percepts of color are unlikely to be phenomenally conscious ones. For if animals lack a capacity for higher-order thought, then their color percepts won't (of course) be *available to* higher-order thought. And consequently those percepts won't possess a dual analog content (e.g. both analog *red* and analog *seeming red*) in the way that our human color percepts do. What is it *like* to undergo a non-conscious perception of red? In the relevant sense, it isn't *like* anything. And nor, of course, can we reflectively *imagine* such a color percept. For any image that we form in this way will itself be phenomenally conscious, and hence will be quite inappropriate for representing the content of the animal's experience.⁵

Color percepts that aren't phenomenally conscious will nevertheless have many of the same causal roles as color percepts that are, on a higher-order account. Thus a fruit-eating monkey might reach out for a particular piece of fruit because she *sees* that the fruit is red, *believes* that red fruit of that type is

ripe, and *desires* to eat ripe fruit. Nothing here requires that the percept of red in question should be a phenomenally conscious one. Percepts of color that are merely first-order and non-phenomenal can perfectly well figure in many normal processes of belief formation and practical reasoning. And indeed, while a human in the same situation *would* enjoy a phenomenally conscious experience of red, it won't normally be *because* it is phenomenally conscious that the person acts as she does – that is, it won't be the higher-order analog content *seeming red*, but rather the first-order analog content *red*, that figures in the causation of her action.⁶

The monkey chooses as she does because she has learned that redness is a sign of ripeness, and because she values ripeness. But there may well be species that are “hard-wired” to prefer, or to shun, certain colors – like the male sticklebacks who are inflamed to rage (in the breeding season) by the sight of anything red. So there might be creatures for whom redness is intrinsically attractive, or intrinsically repellant. And of course this, too, can all take place at a purely first-order level, independent of phenomenal consciousness. There are morals here to be drawn for the case of pain, as we shall see shortly.

4. PAIN AS A PERCEIVED SECONDARY QUALITY OF THE BODY

Tye (1995) argues that feelings of pain are best understood on the model of color perception, but in this case as percepts of secondary qualities of one's own body (see also Martin, 1995; Crane, 1998). A number of considerations lend support to this sort of view. One is that pains are experienced as intrinsically located in space, in something like the way that colors are. Colors are generally perceived as located in particular positions in space, on the surfaces of surrounding objects. (Sometimes they can be perceived as “free floating”, independent of any perceived object, as when one has a red after-image. But even in these cases colors are experienced as being external to oneself, even if one knows that they aren't.) In a similar way, pains are perceived as located in one's own body – either within it, or on a region of its surface. Just as color percepts seem to have

something like the content, “*That* [analog] property covers *that* region of that object”, so pain percepts seem to have a content along the lines of, “*That* [analog] property is located in/on *that* region of my body.”

Another supporting consideration is that the sub-personal processes that generate pain sensations have just the same sort of constructive, interpretative, character as one finds in connection with perception generally (Ramachandran and Blakeslee, 1998). This is demonstrated by, among other things, the fact that phantom-limb pains can be made to disappear by providing people with appropriate visual feedback. For example, a patient whose phantom hand was locked into a fist, with the nails felt as digging painfully into the palm, was temporarily cured by inserting his good hand – initially also closed in a fist – into a box divided by a mirror, arranging himself so that his phantom was *felt as* positioned in the place represented by the mirror-image of his good hand. When he opened his real hand, and so by looking at the reflection watched “both” his hands open, the patient also *felt* his phantom hand opening, and the pain disappeared (Ramachandran et al., 1995).

In claiming that pains *should* be understood as secondary qualities of the body, I am not saying that they are understood in this way by our common-sense psychology. One significant difference between pain and vision, for example, is that we ordinary folk have a proto-theory of the mechanisms that mediate vision. (We know that what you see depends upon the direction of your gaze, and that seeing requires light, for example.) In contrast, we have very little idea of how experiences of pain are caused. Partly for this reason, and partly because pains aren’t inter-subjectively available in the way that colors are, we don’t have much use for the idea of an unfelt pain, in the way that we are perfectly comfortable with the idea of unperceived colors. Likewise we speak of “having” or “feeling” our pains, rather than perceiving them; whereas we talk about perceiving, rather than “having”, colors. And so on. None of this alters the fact that pain experiences are representational, however, nor the fact that such experiences are very similar in their structure and their type of content to color

experiences. Both involve analog representations of a certain sort of secondary quality, distributed over a particular region of space.

Given such a construal of pain experiences (as perceptions of secondary qualities of the body) then essentially the same range of options open up regarding their phenomenally conscious status. A first-order theorist such as Tye (1995) will say that all pain perceptions are phenomenally conscious, provided that they are available to the belief-forming and decision-making processes of the creature in question. My sort of higher-order theorist will say, in contrast, that pain perceptions are only phenomenally conscious when they are available to a higher-order thought faculty, and hence acquire a dual analog content (Carruthers, 2000). On this latter view, a feeling of pain is only phenomenally conscious when it acquires the higher-order analog content *feeling of pain* (in addition to the first-order analog content *pain*), in just the same way that a perception of red is only phenomenally conscious when it acquires the higher-order analog content *experience of red* in addition to the first-order analog content *red*. And it is the higher-order analog content that confers on the states in question their distinctive subjective aspect, and gives them the kinds of qualia or *feel* required for phenomenal consciousness.⁷

5. THE AVERSIVENESS OF PAIN

A great many kinds of non-human animal will be capable of feeling pain, of course, in the sense of perceiving the relevant quality as located in some specific region of their body. But on a higher-order account, it could well be the case that no animals except human beings undergo pains that are phenomenally conscious. Since it is possible that humans are unique in possessing a faculty of higher-order thought, it may be that we are the only creatures to undergo perceptual states with dual analog content, whether those states are percepts of color or of pain. But none of this seems relevant to the aversiveness – the awfulness – of pain. For animals plainly find their pains

aversive; and it may be that such aversiveness is intrinsic to the perception of those states. In the way that some animals might be wired up to find certain colors intrinsically aversive or attractive (when perceived), so it may be that all animals are wired up to find the pains that they feel intrinsically aversive.

It should be emphasized that terms like “aversive” and “awful”, here, can be given a fully cognitive (in the sense of “mentalistic”) characterization. To say that animals find their pain-properties awful, isn’t just to say that they tend to *behave* in such a way as to avoid or ameliorate those properties. Rather, it is to say that animals very much *want* their pains to go away. Only a theory like that of Searle (1992) – which ties the very notions of “belief” and “desire” to a capacity for phenomenal consciousness – would need to deny this (on the assumption that animals lack states that are phenomenally conscious). And there are, in any case, many good reasons for thinking Searlean theories to be false, which we needn’t pause to detail here. (In brief: non-phenomenal uses of “belief” and “desire” are *rife* throughout cognitive science.) So there are good reasons to allow that many non-human animals can be subjects of belief and desire, even if higher-order thought accounts of phenomenal consciousness are correct. Indeed, there is a case for thinking that beliefs and desires are *very* widely distributed throughout the animal kingdom, being possessed even by ants, bees, and other navigating insects (Tye, 1997; Carruthers, 2004c).

It should also be stressed that there is nothing tacitly higher-order involved when we say that animals *find their pains to be awful*. Granted, *to find X to be Y* might involve coming to *believe* that X is Y, in general. And so if an animal finds its own pains awful, we can allow that it must *believe* that its pains are awful. So it has a belief *about* its pain. But this isn’t a higher-order belief, any more than beliefs about colors are higher-order. What the animal finds to be awful is the property that its perceptual state represents as being present in some region of its body (the pain), not its own perception of that property. And the animal’s corresponding belief concerns that represented property, rather than its own state of *representing* that prop-

erty. In short, the intentional object of the *awfulness* attribution is pain itself (when perceived), not the perception of pain.

So there seems nothing to prevent animals from finding their pains awful, even if their pain experiences aren't phenomenally conscious ones. Just as importantly, the awfulness of pain for human beings is unlikely to have anything much to do with its phenomenally conscious properties, on this account. What causes us to think and behave as we do when we are in pain will be the first-order perception of a secondary quality of the body, just as happens in other animals. The difference is merely that, because of the availability of the states in question to a faculty of higher-order thought, those states will *also* have acquired a higher-order analog content, paralleling and representing their first-order analog contents. But it is the first-order content that will carry the intrinsic awfulness of pain, if intrinsically awful it is.

One way of seeing this point is to reflect that the very same intentional law, or nomic generalization, will apply to both humans and other animals, even given the truth of a higher-order account of phenomenal consciousness. The generalization will be something along the lines of this: *when perceiving a state of that sort (pain) as occurring in its body, the subject is strongly motivated to do things that might remove or minimize the property perceived.* This law-like generalization is entirely first-order in character, and can remain true and applicable even when the perceptual state in question (the feeling of pain) *also* possesses a higher-order analog content. Another way of seeing the same point is to notice that, from the perspective of a higher-order theory of phenomenal consciousness, it is likely that a subject who is undergoing a phenomenally conscious pain would *still* have behaved in many of the same ways even if the state in question *hadn't* been phenomenally conscious (e.g. because the subject had failed to possess, or had lost, their capacity for higher-order thought).

Reflecting on the phenomenology of our own pain perceptions is yet another way of appreciating that it is the first-order content of pain-states that carry their intrinsic awfulness. Suppose that you have just been stung by a bee while walking

barefoot across the grass, and that you are feeling an intense sharp pain between two of your toes. How would you characterize what it is that you want to cease, or to go away, in such a case? A naïve subject would surely say, “I want *that* to go away [gesturing towards their foot, meaning to indicate the pain]”. What is wanted is that *that property* – the one that you feel between your toes – should cease. The object of your desire, and the focus of your attention, is the property that is *represented* as being present in your foot, not the mental state of *representing* that property. And this would have been the very same even if your pain-state had lacked a higher-order analog content, and hence hadn’t been a phenomenally conscious one (for a higher-order theorist).

Of course, a more sophisticated subject who knows about pain receptors, and who knows that analgesics work by blocking the signals generated by those receptors from reaching the brain, might say, “I don’t care whether you get rid of *that* [pointing to her foot and meaning to indicate the pain perceived], what I want is that you should get rid of the *feeling* of pain”. And this would indeed be a second-order judgment, comparable to saying that I want you to remove, not the redness of an object, but my *perception* of the redness. But this surely isn’t the basic case. In the basic case, the object of aversion is the pain represented. And such aversions can be entirely first-order in character.

The awfulness of pain on this account, then, has nothing much to do with its phenomenally conscious properties. What is found awful is that the relevant secondary quality (pain) is perceived as occurring in one’s body; and this can be a purely first-order (non-phenomenally conscious) affair. Feelings of pain are perceptual states that we share with many other animals, who seem to find the properties perceived just as aversive as we do. Of course, here as in the case of color, we have no idea how to *imagine* what a pain might be like that isn’t phenomenally conscious. Any pain that we reflectively imagine is *ipso facto* going to be a phenomenally conscious one. But given the right theoretical background, this needn’t prevent us from recognizing that non-phenomenally conscious pains exist (in-

deed, are extremely common, if the pains of all non-human animals have this status). Nor should it prevent us from recognizing that non-phenomenally conscious pains are experienced as awful in essentially the same sort of way that we experience our own pains as awful.

6. FIRST-ORDER PAIN-STATES AS APPROPRIATE OBJECTS OF SYMPATHY

Given that pains that aren't phenomenally conscious are nevertheless found awful by those who possess them, it is *prima facie* plausible that the subjects of such pains are appropriate objects of sympathy and concern, despite the absence of phenomenal consciousness from their lives.⁸ But is it really so obvious that pain-states that aren't phenomenally conscious are appropriate objects of sympathy, however? And how is such an intuition to be grounded? For doesn't Carruthers (1989) rely on the contrary intuition in arguing that non-conscious pains *aren't* appropriate objects of sympathy? Is there really anything more, here, than a clash of equally unwarranted intuitions?⁹

Carruthers (1989) imagines the case of Mary, who is unusual in that she only ever has pains in her legs that are non-conscious. These pains cause her to rub the offended part, and in severe cases to scream and cry out, but all without any *conscious* awareness of pain on her part. It is suggested that these pains aren't appropriate objects of sympathy and concern, largely because Mary herself isn't bothered by them – she may find it inconvenient when she can't walk properly on a broken ankle (and we may feel sympathy for her inconvenience), but she isn't consciously distressed by her pains, and she doesn't seek our assistance (other than physical assistance) or sympathy.

The example of Mary goes much further than is appropriate as a test of the question whether pains that aren't phenomenally conscious are appropriate objects of sympathy, however. This is because Mary's reactions to her broken ankle don't appear to be cognitively mediated ones – i.e. they aren't mediated by her

beliefs and desires, but are in some sense purely behavioral. (She appears not even to know that she *has* a pain in her ankle, except by inference from her own behavior.) Or at any rate, if her reactions *are* cognitively mediated, we can say that the cognitions in question aren't available in anything like the normal way to inform speech and to guide practical reasoning.

We can set Carruthers (1989) a dilemma, in fact. *Either* Mary's behavior isn't cognitively mediated; in which case, whatever it is that is causing her behavior, it isn't a state that intrinsically gives rise to *aversion* (i.e. to a desire that the represented property should cease). But in that case what causes her behavior isn't *pain* in the sense that we are discussing. *Or* Mary's behavior *is* cognitively mediated by a representation of a pain-property in her ankle together with an aversion to the presence of that property. But these states are unusual in not being available to inform her speech (e.g. leading her to say, "There is a pain in my ankle" – note, *not* a second-order statement), nor to guide her practical reasoning (e.g. leading her to seek help in satisfying her desire for the pain to cease). But on this horn of the dilemma, we have no reason to trust Mary's own assessment of her situation. For if Mary *does* have a strong desire for her pain to cease, but something is blocking the expression of this desire in speech and action, then she surely *is* an appropriate object of sympathy.

The real point (and what really grounds the claim that pain-states that aren't phenomenally conscious can be appropriate objects of sympathy) is that there is good reason to think that *in our own case* it is the first-order property represented (the pain) that is the object of aversion, not the mental state of representing that property (the perception of the pain). The aversiveness of pain in our own case is primarily a first-order affair. So if sympathy is appropriate for our own pains, it is equally appropriate in the case of pain-states that aren't accompanied by higher-order analog contents representing that the pain in question is being perceived. Hence sympathy is equally appropriate in connection with states that aren't phenomenally conscious, given a higher-order thought account of phenomenal consciousness.

I have been arguing that there can perfectly well be aversive pain-states that aren't phenomenally conscious, and that such states are nevertheless appropriate objects of sympathy. But some philosophers will be inclined to object that the very *idea* of a pain-state that isn't phenomenally conscious makes no sense. They will insist that the idea of *pain* is the idea of a state with a certain distinctive introspectible *feel* to it. So a state that lacked *feel* (that wasn't phenomenally conscious) just couldn't be a pain. But equally, of course, philosophers were once inclined to object just as strongly to the idea of a non-conscious visual percept. Yet the last 20 years of psychological research has now convinced most of us that the idea of non-conscious visual experience does make sense after all.

The more general point is that there can be states that are similar to phenomenally conscious visual experiences, or similar to phenomenally conscious pains, in almost all respects *except* that of possessing an introspectible *feel*. In particular, there can be states that possess just the same sorts of (first-order) intentional contents, and that have many of the same functional roles, as their phenomenally conscious counter-parts.¹⁰ In which case it is surely well-motivated to treat them as states of the same mental kind. Anyone is free to stipulate that the term "pain", or the term "visual percept", should be reserved for states that are phenomenally conscious, of course. But in that case all of our discussion could be couched in terms of states that are in other respects *similar to* pains, or *similar to* visual percepts. And then just the same issues would arise, concerning whether these *pain-like* states are appropriate objects of sympathy and concern.

7. EMOTION AS INVOLVING SOMASENSORY PERCEPTION

I have argued that, in the case of bodily pain at least, there can be suffering without subjectivity. But what of the forms of suffering characteristic of negative emotional states like grief? And can there be feelings of frustration without phenomenal consciousness? For example, can a creature experience

disappointment at a lost opportunity for mating or predated without undergoing mental states that are phenomenally conscious ones? And can a creature feel anything like grief at the loss of a partner or offspring in the absence of such consciousness? I propose now (in this section and the one following) to sketch reasons for positive answers to these questions. Since most of the groundwork has already been laid, our discussion can be comparatively brisk.

Before we begin, however, we need to recall that our task is to investigate whether there can be forms of *suffering* that are worthy of sympathy in the absence of phenomenal consciousness, not whether sympathy might be warranted on other grounds. Carruthers (1999) argues that creatures whose desires are frustrated in a thinner sense than this (not involving anything naturally described as “suffering”) could be appropriate objects of sympathy – e.g. a creature that *wants* to drink but *believes* that isn’t presently drinking. Our present task is to see whether we can make sense of the idea that a creature can suffer emotionally in the absence of phenomenal consciousness, and whether such a creature would be an appropriate object of sympathy on that ground alone.

The main claim underlying a positive answer to these questions is that feelings of frustration, like feelings of pain and emotional states more generally, are best understood in terms of perceptions of secondary qualities of the body. As is now well known, emotional states like frustration, anger, and grief give rise to a variety of kinds of distinctive physical and physiological effects, such as changes in heart-beat, breathing, muscle tension, and bodily posture, as well as changes within the chemistry of the blood, such as its levels of adrenaline (Damasio, 1994; Rolls, 1999). When we are in such a state, we will be aware (via somasensory perception) of a complex bodily *gestalt*. And it is the content of this percept that gives rise to the distinctive phenomenology of emotion.

Of course one doesn’t need to know anything about hearts, or about breathing, or about bodily postures in order to experience these states of the body. One can still have analog perceptual representations of the various changes involved.

And just as the contents of our color percepts take the form, “*This* surface has *this* [analog] property”, and the contents of feelings of pain take the form, “*This* region of my body has *this* [analog] property”, so too the contents of the somasensory percepts distinctive of emotional states take the form, “*These* regions of my body are undergoing *these* [analog] changes”. And just as pain properties are experienced as intrinsically aversive, so it may be that many of these somasensory *gestalts* are experienced as aversive or attractive.

It is important to realize that there is a sense in which the phenomenology of frustration is peripheral to the main motivational roles of desire and frustration of desire, however. When I am hungry, for example, I am in a state that is focused outwards on gaining access to and ingesting food, rather than inwards on the removal of an aversive sensation. I might sometimes feel such sensations, of course – such as a tightening of the stomach and feelings of bodily weakness – and I may believe that eating will remove them, thus providing myself with an additional motive to eat. But it is possible to desire food without ever experiencing such feelings. (Rats, for example, will work for food even if they have never experienced hunger. See Dickinson and Balleine, 2000.) And such feelings plainly aren’t, in any case, the *primary* motivation to eat – the *desire to eat* itself provides that.

Similarly (to take a different example), when I drive to the airport to collect a loved one, my goal is to be in the presence of the loved person once again, rather than to remove some unpleasant feeling or sensation. And the disappointment that I feel when her flight is cancelled (the sinking in my stomach, the tightening in my chest, and so forth), consists in a set of states *caused by* the frustration of my desire, rather than being constitutive of it.

What, then, is the role of somasensory feeling in our lives, if these feelings aren’t part-and-parcel of desiring as such? According to the model of human practical reasoning proposed and elaborated by Damasio (1994), we continually monitor our bodily states via somasensory perception while contemplating envisaged options and outcomes. The thought of a fearful event

makes us a bit afraid, and thinking about the loss of a loved one makes us somewhat bereft. We sense the resulting changes in ourselves, and use these as cues to the desirability or otherwise of the envisaged state of affairs. In effect, the functional role of these somasensory percepts (whether positive or aversive) lies in the mental rehearsal of action, rather than primarily in action itself.¹¹

Given the correctness of some such account as this, it seems unlikely that the aversiveness of disappointment is as “primitive” and widespread within the animal kingdom as is the aversion to pain. For as we noted above, the most basic kinds of desire-based and emotion-based attraction and aversion are directed outwards to objects and states of affairs in the world, rather than towards properties of our own bodies. A creature that is afraid of a nearby predator, say, has an aversion to remaining in the presence of the fearful thing. And although the creature’s state of fear will give rise to a suite of physiological changes, one wouldn’t necessarily expect the animal’s perceptions of these changes to be aversive in their own right. For why would this be needed, given that the animal *already* has an aversion to the fearful object? Somasensory properties probably only began to be aversive or attractive in and of themselves when animals started to engage in mental rehearsal of various sorts. For they would then have needed some way of determining the desirability or otherwise of an envisaged scenario. And it would appear, from the human case at least, that the monitoring of somasensory responses is what emerged as the solution to this problem.

While the aversiveness of some emotional states is probably by no means universal amongst animals, it may still be quite widespread amongst mammals, at least. For there is good reason to think that mental rehearsal isn’t unique to humans. For example, much of the evidence that has been cited (controversially) in support of the higher-order thought abilities of chimpanzees can also be used (much *less* controversially) to support the claim that this species of great ape, at least, engages in mental rehearsal. For example: a subordinate ape knows the location of some food hidden within an enclosure, and from

previous experience expects to be followed by a dominant who will then take the food. So the subordinate heads off in another direction and begins to dig. When the dominant pushes her aside and takes over the spot, she doubles back and retrieves and quickly eats the food.

Such examples are generally discussed as providing evidence that chimps can engage in genuine (higher-order thought involving) *deception* – that is, as showing that the subordinate chimp is intending to induce a false belief in the mind of another (Byrne and Whiten, 1988; Byrne, 1995) – whereas critics have responded that chimpanzees may just be very smart behaviorists (Smith, 1996; Povinelli, 2000). But either way, it seems that the chimp must engage in mental rehearsal, predicting the effects of walking in the wrong direction and beginning to dig (the dominant will follow and take over the digging), and discerning the opportunities for hunger-satisfaction that will then be afforded.

The suffering that arises from desire-frustration, then, and that is involved in other similar sorts of negative emotional state such as grief and terror, lies in the somasensory perception of a distinctive bodily *gestalt*, which is caused by the state in question, and which is experienced as aversive or intrinsically unwelcome. The question, now, is whether such suffering can exist in the absence of phenomenal consciousness.

8. THE HARM OF FELT FRUSTRATION WITHOUT PHENOMENAL CONSCIOUSNESS

Notice that the results of somasensory monitoring to produce bodily *gestalts*, just like percepts of pain, are first-order analog representations of secondary qualities of the body. So on any higher-order approach to phenomenal consciousness, such states won't be intrinsically phenomenally conscious ones. On the contrary, according to my own account, it will only be in cases where those states are available to a higher-order thought faculty capable of entertaining thoughts about those very states, that the percepts in question will acquire a dual analog

content (both first-order and higher-order), and will hence come to have the kind of subjectivity that is distinctive of phenomenal consciousness.

Moreover (and again like percepts of pain) it seems likely that the causal roles of bodily *gestalts* (and particularly their aversiveness) are attached to their first-order, rather than to their higher-order, contents. It is the perception of increased heart-rate, dry mouth, and so forth distinctive of fear that is aversive in the first instance. Such states would very likely have been aversive prior to the evolution of a higher-order thought faculty if (as seems plausible) mental rehearsal was also available earlier. And they seem likely to remain aversive in the absence of such a mechanism (e.g. in severe autism).¹²

Thus there may well be animals that are capable of experiencing grief, for example, in the absence of phenomenal consciousness. These animals will perceive in themselves a complex bodily *gestalt* (physical lassitude and so forth), caused by the knowledge that a mate or infant is dead or lost, say. And they will experience this bodily state as intrinsically aversive. But all of this can be purely first-order in character. The feelings, and the aversion to their objects, can exist in the absence of a capacity for higher-order thought, and so without the states in question having the dual-level analog content distinctive of phenomenal consciousness. And insofar as we realize that this is so, it seems perfectly appropriate to feel sympathy for the animal's state, even if we cannot imagine what such a state might be like.¹³

Of course, some philosophers will wish to object (as they did in the case of pain) that they can make no sense of the idea of states of disappointment, or of grief, that aren't phenomenally conscious ones. But the reply is essentially the same, too. People can give whatever stipulative definitions they want. But then the same issues can be raised concerning states that are in many respects *similar to* the targets in question (disappointment, grief). And it seems undeniable, moreover, that such states *can* exist, and very probably *do* exist given the correctness of some or other higher-order account of phenomenal consciousness.

9. CONCLUSION: SUFFERING WITHOUT SUBJECTIVITY

The conclusion of this paper is the one that I advertised at the beginning: even if the pains and disappointments experienced by non-human animals are never phenomenally conscious, they can still be appropriate objects of sympathy and concern. In the relevant sense, there can be suffering without subjectivity. I want to emphasize, however, that the conclusion isn't (yet) that concern for the sufferings of non-human animals is morally required of us. That is a distinctively moral question, to be answered via considerations of moral theory (see Carruthers, 1992, Chapters 1–7). All that has been established here, is that sufferings that aren't phenomenally conscious can be *appropriate* objects of sympathy, not that such sympathy is morally mandatory. I also want to emphasize that the assumptions made (and sometimes sketchily defended) in the course of this paper (e.g. a higher-order thought account of phenomenal consciousness, and secondary-quality accounts of pain and disappointment) are for the most part just that: assumptions. My goal has only been to demonstrate that suffering without subjectivity is *possible*, or *makes sense*; not to show that it is ever actual.

ACKNOWLEDGEMENTS

I am grateful to all those who participated in discussions of this and related material, presented at a University of Maryland (Baltimore County) philosophy colloquium, and also at a symposium devoted to animal consciousness at the Association for the Scientific Study of Consciousness meeting in Memphis (both held in Spring 2003); and especially to Daniel Dennett, Susan Dwyer, Robert Lurz, Jessica Pfeifer, Paul Pietroski, and Carol Voeller. I am also grateful for the insightful comments of an anonymous referee for this journal.

NOTES

¹ *Analog* content is intentional content that is more fine-grained than any concepts we could possess (think of the fine shades of color that we can

distinguish, but not remember), even if it is nevertheless imbued with, or shaped by, concepts. *Non-conceptual* content is intentional content that is not only fine-grained but independent of, and prior to, concept-deployment. See Carruthers (2000, Chapter 5) for discussion.

² Note the awkwardness that we face, here and throughout this paper, in having to describe intentional contents that are analog and/or non-conceptual in character using non-analog (or “digital”) concepts like *red* or *experience of red*. Although an analog content can’t be *expressed* in language, however (in the way that a propositional/conceptual content can), it can still be talked about and referred to.

³ For the competing views, see on the one side Bermúdez (1995), Tye (1995, 2000), Carruthers (2000), and Kelly (2001), and on the other side McDowell (1994).

⁴ For discussion of these and other options, see Carruthers (2000, Chapter 5).

⁵ In my view, this is one of the sources of our temptation to think that animals *must*, somehow, be phenomenally conscious. For we have no idea how to imagine an experience that *isn’t* phenomenally conscious. See Carruthers (2000, 2005). But of course imagination is a notoriously fallible guide to possibility. And in the present case we have to hand an explanation of *why* reflective/conscious imagination should be of no use in representing the perceptual states of non-human animals, given that those states are non-conscious ones. (Note that I can allow that non-conscious images are possible, and that such images would be of the right sort for representing non-conscious experiences. But we can’t form and direct the content of such images at will, in such a way as to provide ourselves with an answer to the question of what the non-phenomenally conscious experiences of an animal might be like.)

⁶ See Carruthers (2005), for extended defense of these claims. All that we need for present purposes is that they *make sense*, and that they are *permitted by* a higher-order thought theory of phenomenal consciousness. For recall that our goal is just to explore what room *might* be left by such higher-order theories for non-phenomenal forms of suffering.

⁷ Pain-states that aren’t phenomenally conscious are still *felt*, of course. This is just to say that the pain (the secondary quality of the body) is perceived. But pains that are felt can nevertheless lack *feel* in the quasi-technical philosophical sense of possessing the distinctive sort of subjectivity that constitutes a mental state as a phenomenally conscious one.

⁸ Whether we are *required* to feel sympathy for (or to otherwise to take account of) the pains of non-human animals, however, is another, and in my view distinctively moral, question. It is a question that needs to be answered by considerations of moral theory rather than philosophical psychology. See Carruthers (1992, Chapters 2–7).

⁹ I am grateful to an anonymous referee for pressing these objections on me.

¹⁰ In the case of visual experiences, what we have actually discovered is that there are *no* states that play all of the roles that common sense assigns to

visual experiences; namely, leading to the formation of belief, feeding into planning in relation to the perceived environment (“I’ll go *that way*”), and guiding the details of our bodily movements on-line, as well as possessing a distinctive *feel*. Rather, it turns out that there is one visual system realized in the temporal lobes that is concerned to provide inputs to conceptual thought and planning (and whose outputs are conscious), and another visual system realized in the parietal lobes that provides on-line guidance of movement, but whose outputs aren’t conscious (Milner and Goodale, 1995; Clark, 2002). It seems better to say that *both* of these kinds of state are types of vision, rather than to say that neither is.

¹¹ Of course, many of the bodily changes that we perceive will themselves have functional roles relevant to action, such as the collection of changes constitutive of the “flight-or-fight” response.

¹² Autism admits of a very wide spectrum of severity, of course, and not all autistic people are wholly lacking in a capacity for higher-order thought.

¹³ Here, as in the case of pain, any attempt at reflective imagination aimed at answering our own question will inevitably involve images that are phenomenally conscious, and hence inappropriate to represent the mental state of the animal.

REFERENCES

- Bermúdez, J. (1995): ‘Non-conceptual Content’, *Mind and Language* 10, 333–369.
- Block, N. (1986): ‘Advertisement for a Semantics for Psychology’, *Midwest Studies in Philosophy* 10, 615–678.
- Block, N. (1995): ‘A Confusion About a Function of Consciousness’, *Behavioral and Brain Sciences* 18, 227–261.
- Byrne, R. (1995): *The Thinking Ape*, Oxford: Oxford University Press.
- Byrne, R. and Whiten, A. (1988): *Machiavellian Intelligence*, Oxford: Oxford University Press.
- Carruthers, P. (1989): ‘Brute Experience’, *Journal of Philosophy* 86, 258–269.
- Carruthers, P. (1992): *The Animals Issue*, Cambridge: Cambridge University Press.
- Carruthers, P. (1999): ‘Sympathy and Subjectivity’, *Australasian Journal of Philosophy* 77, 465–482.
- Carruthers, P. (2000): *Phenomenal Consciousness: A Naturalistic Theory*, Cambridge: Cambridge University Press.
- Carruthers, P. (2004a): ‘Phenomenal Concepts and Higher-order Experiences’, *Philosophy and Phenomenological Research* 67, 316–336.
- Carruthers, P. (2004b): ‘HOP over FOR, HOT Theory’, in R. Gennaro (ed.), *Higher Order Theories of Consciousness*, Philadelphia: John Benjamins.
- Carruthers, P. (2004c): ‘On Being Simple Minded’, *American Philosophical Quarterly* 41, 205–220.

- Carruthers, P. (2005): 'Why the Question of Animal Consciousness Might Not Matter Very Much', *Philosophical Psychology* 18.
- Clark, A. (2002): 'Visual Experience and Motor Action: Are the Bonds too Tight?', *Philosophical Review* 110, 495–520.
- Crane, T. (1998): 'Intentionality as the Mark of the Mental', in A. O'Hear (ed.), *Current Issues in Philosophy of Mind*, Cambridge: Cambridge University Press, pp. 229–252.
- Damasio, A. (1994): *Descartes' Error: Emotion, Reason and the Human Brain*, New York: Putnam's Sons.
- Dickinson, A. and Balleine, B. (2000): 'Causal Cognition and Goal-directed Action', in C. Heyes and L. Huber (eds.), *The Evolution of Cognition*, Cambridge, MA: MIT Press.
- Dretske, F. (1995): *Naturalizing the Mind*, Cambridge, MA: MIT Press.
- Feinberg, J. (1977): 'Harm and Self-interest', in P. Hacker and J. Raz (eds.), *Law, Morality and Society*, Oxford: Oxford University Press, pp. 285–308.
- Kelly, S. (2001): 'Demonstrative Concepts and Experience', *Philosophical Review* 110, 397–420.
- Loar, B. (1981): *Mind and Meaning*, Cambridge: Cambridge University Press.
- Lycan, W. (1996): *Consciousness and Experience*, Cambridge, MA: MIT Press.
- Martin, M. (1995): 'Bodily Awareness: A Sense of Ownership', in J. Bermúdez, N. Eilan and A. Marcel (eds.), *The Body and the Self*, Cambridge, MA: MIT Press.
- McDowell, J. (1994): *Mind and World*, Cambridge, MA: Harvard University Press.
- McGinn, C. (1989): *Mental Content*, Oxford: Blackwell.
- Millikan, R. (1984): *Language, Thought, and Other Biological Categories*, Cambridge, MA: MIT Press.
- Millikan, R. (1989): 'Biosemantics', *Journal of Philosophy* 86, 281–297.
- Milner, D. and Goodale, M. (1995): *The Visual Brain in Action*, Oxford: Oxford University Press.
- Nagel, T. (1979): 'Death', in *Mortal Questions*, Cambridge: Cambridge University Press, pp. 1–10.
- Papineau, D. (1987): *Reality and Representation*, Oxford: Blackwell.
- Papineau, D. (1993): *Philosophical Naturalism*, Oxford: Blackwell.
- Peacocke, C. (1992): *A Study of Concepts*, Cambridge, MA: MIT Press.
- Povinelli, D. (2000): *Folk Physics for Apes*, Oxford: Oxford University Press.
- Ramachandran, V. and Blakeslee, S. (1998): *Phantoms in the Brain*, New York: Fourth Estate.
- Ramachandran, V., Rogers-Ramachandran, D. and Cobb, S. (1995): 'Touching the Phantom Limb', *Nature* 377, 489–490.
- Rolls, E. (1999): *Emotion and the Brain*, Oxford: Oxford University Press.

- Rosenthal, D. (1993): 'Thinking That One Thinks', in M. Davies and G. Humphreys (eds.), *Consciousness*, Oxford: Blackwell, pp. 197–223.
- Searle, J. (1992): *The Rediscovery of the Mind*, Cambridge, MA: MIT Press.
- Smith, P. (1996): 'Language and the Evolution of Mind-reading', in P. Carruthers and P. Smith (eds.), *Theories of Theories of Mind*, Cambridge: Cambridge University Press.
- Tye, M. (1995): *Ten Problems of Consciousness*, Cambridge, MA: MIT Press.
- Tye, M. (1997): 'The Problem of Simple Minds', *Philosophical Studies* 88, 289–317.
- Tye, M. (2000): *Consciousness, Color and Content*, Cambridge, MA: MIT Press.
- Weiskrantz, L. (1997): *Consciousness Lost and Found*, Oxford: Oxford University Press.

*Department of Philosophy
University of Maryland
College Park, MD 20742
USA
E-mail: pcarruth@umd.edu*

Copyright of Philosophical Studies is the property of Kluwer Academic Publishing / Academic and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.