

## FORMS OF HARM AND OUR OBLIGATIONS TO HUMANS AND OTHER ANIMALS\*\*

### I. INTRODUCTION

I will argue that there are two basic forms of harm (as distinct from damage); that there are profound differences between the nature of human language and the (non-linguistic) forms of communication that are employed by other animals; and that it turns out that language users can be harmed in both of the ways I outline whereas non-linguistic beings (which includes all nonhuman animals and some humans as well) can only be harmed in one of these ways. It might seem strange to think that the use of language might in itself have implications for the kinds of ways in which a being can be harmed, but as we will see, it turns out that without language a being cannot have a temporally structured sense of self-awareness and that without this it is not *possible* to harm a being in one of the two basic ways in which beings can be harmed.

### II. HARM VS. DAMAGE

In order to be clear about the focus of my argument I want to draw a distinction between the notions of *harm* and *damage*. This distinction turns on the fact that there is an important difference between the ways in which we can have a detrimental effect upon sentient beings on the one hand and all other kinds of nonsentient entities, structures, or abstract complexes on the other hand (such as plants, chairs, art works, or reputations). This is because sentient beings possess an inner, experiential dimension to their existence whereas nonsentient things don't. We can therefore have a negative impact upon the qualitative experiential state or capacity of a sentient being – no matter how rudimentary that capacity might be – whereas we cannot have a negative impact upon the qualitative experiential state or capacity of a nonsentient entity, structure, or abstract complex because they do not possess an inner, experiential dimension in the first place. I therefore suggest that we use the term “harm” to refer to detrimental effects to entities

\* Warwick Fox is Emeritus Professor of Philosophy, University of Central Lancashire.

\*\* © 2012 Warwick Fox ([www.warwickfox.com](http://www.warwickfox.com)).

that possess any kind of inner, experiential dimension (i.e., sentient beings) and “damage” to refer to detrimental effects to anything else (i.e., nonsentient entities, structures, or abstract complexes).<sup>1</sup>

I think that it could greatly help the clarity of discussions in this area if other writers were to observe what, with a nod to Hume’s famous discussion of the “is-ought” distinction, we might refer to as this “small attention”.<sup>2</sup> By all means let us argue questions concerning the permissibility or impermissibility of *damaging* nonsentient living things as well as other kinds of nonsentient entities, structures, and abstract complexes such as artistic and architectural works and reputations, but let us desist from reference to *harming* them since this only serves to muddy the waters by implicitly suggesting that sentient beings and nonsentient entities are alike – or at least on the same continuum – in terms of the kinds of detrimental effects we can have on them when in fact the opposite is true.

Of course, *damaging* someone’s property or reputation can lead to *harming* them because of the effect that this damage has upon the experience of the person themselves, but this doesn’t necessarily follow. For example, we can damage someone’s property or reputation without harming the person themselves if, say, the person is dead, or does not find out about this damage, or is not at all attached to that piece of property, or is so far beyond playing the reputation game – or such an experienced survivor of it – that they simply don’t care. Moreover, even in those cases in which damage is damage pure and simple – that is, damage that does not lead to harm – none of this should be taken as saying that damaging something is not in itself a bad thing to do; nor should it be taken as saying that any kind of harm is necessarily worse than any kind of damage. Rather, I am simply insisting that there is an important distinction to be made between the kinds of detrimental effects that can be visited upon sentient beings on the one hand and any other kind of non-sentient entity, structure, or abstract complex on the other, and that it would be helpful for us to observe this distinction.

### III. TWO FORMS OF HARM

My interest here, then, is in harm rather than damage, and I want to argue that there are two basic forms of harm. One of these is obvious to everyone. This is the kind of harm that I will refer to as *unnecessary pain and suffering*. It is important to employ the qualifier “unnecessary” here because although the experience of pain and suffering is self-evidently bad for any being that has it, this alone does not make it a form of harm. This is because the pain and suffering involved might be for the overall benefit of the being involved; for example, taking yourself or your child to the dentist, taking an

animal to a vet, telling someone that a loved one has died rather than hiding this from them so that they then later have to deal not only with the fact of the death but also the fact of the deception. And even if the pain and suffering involved is not for the overall benefit of the being concerned – and, thus, *is* a form of harm to that being – this alone does not make it a form of *moral harm*. This is because the pain and suffering involved might not be caused by a moral agent, in which case it is still a form of harm, but a form of *non-moral harm*, such as an accident or the kind of event that gets referred to in legal terms as an “act of God” (which is an interesting term to use given the “non-moral” categorization – and, thus, non-legally prosecutable nature – of this harm). Moreover, even if pain and suffering is caused by a moral agent and does harm certain beings, then, *depending upon the kind of ethical approach one adopts*, it might still be deemed to be the morally right thing to do – and, thus, morally necessary – if it is for the overall benefit of *other* entities that are considered to be of sufficient value as to justify this course of action. In sum, then, the upshot of these considerations is, first, that pain and suffering represents a form of *harm* when it can be shown to be *unnecessary* (the precise understanding of which is not only open to some degree of interpretation in its own right but might also vary considerably between people with different ethical orientations), and second, that unnecessary pain and suffering represents a form of *moral harm* when it is caused by a moral agent.

The final point I want to make in characterizing this first form of harm is to note that I refer to it in terms of unnecessary pain *and* suffering because people will sometimes want to distinguish between pain in the sense of an essentially bodily based sensation and suffering in the sense of mental or emotional distress, that is, as a more cognitively based phenomenon. Even so, pain and suffering clearly lie on the same kind of (affective) scale – they both feel bad – and whatever nuances people might sometimes want to read into these words, they are actually defined in terms of each other in dictionary definitions (e.g., my dictionary not only defines *suffering* in terms of “a state or an instance of enduring *pain*” but it also gives one of the meanings of *pain* as “emotional *suffering* or mental distress” [my emphases]<sup>3</sup>). Thus, while nothing much hangs on the nuances that people might wish to read into these words, I nevertheless use them both to imply that I am meaning to refer to the whole range of negative feelings, from physical pain to mental distress. For convenience, I will also at times refer to harm of the “pain and suffering” kind in more formal terms as *affective harm* (i.e., harm in respect of our sensations, feelings, and emotions).

Although unnecessary pain and suffering is the most obvious and

ubiquitous kind of harm we can think of, it does not exhaust the category of harm because it is entirely possible for a being to die or be killed without experiencing any unnecessary pain and suffering. For example, if you undergo a general anesthetic, then the only thing you will feel – if you even feel this – is a slight prick in your arm when the needle is inserted for injecting the anesthetic. You will not experience this as a harm, however, because you believe that what is being done is for your benefit. But what if – for the sake of the argument and not for the sake of worrying anyone! – the anesthetist is sufficiently deranged to intend your death or else sufficiently inexperienced or neglectful to cause it by accident? Or what if something else goes wrong during the procedure and you die without regaining consciousness? If so, then your very being ceases to exist without you having experienced any form of affective harm being done to you. And yet the greatest possible harm has been done to you; your very being has been eliminated. These considerations mean that pain and suffering on the one hand and death on the other hand are dissociable: we can experience pain and suffering without dying and we can die without experiencing pain and suffering. Yet both are clearly forms of harm; they both clearly represent detrimental effects upon the qualitative experiential state or capacity of our being. Moreover, even if we were to experience some minor degree or relatively short period of pain and suffering in dying or being killed, the central issue at stake in this situation would not be the pain and suffering that we had experienced so much as the fact that we had died or been killed. It is therefore important to distinguish two basic forms of harm: affective harm – that is, harm of the unnecessary pain and suffering kind – and this second form of harm concerning death per se (note that I will sometimes use the term “death per se”, as here, to refer to death in the absence of considerations regarding pain and suffering).

#### IV. THE SECOND FORM OF HARM: WHAT KIND OF DEATH?

But what kind of death are we referring to when we refer to death per se as a harm? After all, a member of our own species, to take our own case, can exist on a merely biological level; on a biological plus merely sentient level; on a biological, sentient plus merely self-aware-in-the-moment level; and, in the normal case of people beyond the age of about four, on a biological, sentient, self-aware-in-the-moment, and self-aware-over-time level. To exist on a merely biological level would be to exist in a “persistent vegetative state” of a kind that lacked any sense of feeling whatsoever; to exist on a biological plus merely sentient level would be to be a sentient being and no more; to exist on a biological, sentient plus merely self-aware-in-the-

moment level would be to possess what we could refer to as a *temporally isolated sense of self-awareness*; and to exist on a biological, sentient, self-aware-in-the-moment, and self-aware-over-time level would be to possess, in the normal case, what we could refer to as an *enduring temporally structured sense of self-awareness*, *autobiographically structured awareness*, or an *autobiographical sense of self*.

Let us consider these cases in turn. Imagine, first, if you were stripped of all these layers of existence except the biological one; for example, you might be in a deep coma or otherwise sedated to the point where you were still alive but unable to feel anything and, therefore, had no self-awareness either in the moment or over time. Suppose also that there was no possibility of you ever being in anything other than this nonsentient vegetative state. In this case you could not be caused harm of the pain and suffering kind because you are no longer sentient. But neither could your biological death harm *you* in the only sense that can matter to you, since you have no sense of self in the first place. *You* – your own sense of self – would have already died at the moment that you lost your self-awareness forever.

Now consider a situation in which you were both biologically alive and sentient but, perhaps through some catastrophic neurological damage, had been stripped of any form of self-awareness. Although you would no longer be in a plant-like nonsentient vegetative state, you would nevertheless be in a similar cognitive situation to the one we suppose many nonhuman animals to be in: you would experience things at a basic, first-order level, but you would not be self-aware, that is, you would not be aware at a second – or higher – order level *of* your awareness. To put it another way, you would experience things as they occurred but you would not have any sense of self that was, to borrow an apt phrase employed by the developmental psychologist Katherine Nelson in a related context, “distinct from ongoing experience”<sup>4</sup>. Now in this case, you can certainly be caused harm of the pain and suffering kind because you are sentient. But as was the case in regard to biological death per se, the death of your sentience per se could not harm *you* in the only sense that can matter to you, since you have no sense of self in the first place. In this case, too, *you* – your own sense of self – would have died at the moment that you lost your self-awareness forever.

Now consider the more complex situation in which you are biologically alive, sentient, and self-aware-in-the-moment, but have been stripped of your sense of being self-aware-over-time, that is, stripped of your temporally structured sense of self-awareness and, thus, your autobiographical sense of self. Unfortunately, it turns out that some people do exist in this condition: there are rare examples of people who have suffered brain dam-

age – primarily to the hippocampus, which is required to transfer memories from short-term (or working) memory to long-term memory storage – such that they live in a temporal window of *seconds*. They are self-aware in the moment, but lack any kind of temporally structured sense of self-awareness. Clive Wearing, one of the most profoundly affected patients with this condition, is continually reporting that he has “just become conscious” or “just woken up for the first time”, since he has no memory of having been conscious a minute ago.<sup>5</sup> Formerly an accomplished musician, he can still talk, play the piano, and recognize his wife every time she comes into the room (since repetition of these practices prior to his illness has left him with what memory researchers refer to as *implicit memory* and, specifically, *procedural memory* for these things, which operates at an unconscious level), but he cannot forge any new memories or consciously link the pieces of his life together (i.e., he lacks *explicit memory* – also called *declarative memory* – and, especially, that form of explicit memory referred to as *autobiographical memory*). In referring to what it is like to be in this only ever self-aware-in-the-moment condition, he repeatedly says an astonishing thing: “It’s like being dead” or “It’s like death!”<sup>6</sup> And, indeed, his autobiographical self is dead; he keeps dying to himself, as it were, in this moment, and the next moment, and the next; no trace of ever having experienced those moments remains.

In what ways can a person in this condition be harmed? Obviously such a person can be caused harm of the pain and suffering kind in any given moment because they are sentient. But can biological death per se or sentience death per se (by which I mean the death of any capacity to feel anything at all) be a harm to them? If we took Clive Wearing at his word that it is already “like being dead” to be in his condition, then we might conclude that biological or sentience death per se cannot harm him any further since he has already died in the only sense that matters to him. But even if we were to set this dramatic remark to one side as (contrary to my previous note) some kind of rhetorical flourish, we can still inquire in a more dispassionate way into the question of what biological or sentience death would mean to a person in this condition.

Consider, then, the hypothetical case of someone who is self-aware in any given moment but has a temporal window of next to nothing: they have a sense of existing on a moment-by-moment basis, but that’s it; they have no sense of having just *had* a sense of existing or that they *will* have a sense of existing. Now the death of someone in this condition could not possibly mean that they are cut off from their memory claims upon the past, their dreams, plans, and projects for the future, and their self-aware location of

the present in that autobiographical context, because they have already lost any sense of these things. They cannot, in other words, be cut off from anything of any autobiographical interest to them – they cannot be cut off from their own story, as it were – because they have no autobiographical capacities; their own story stopped at the internal, experiential level for them at the moment they lost their temporally structured sense of self, their sense of their own existence through time. Since then they have lived in *experientially isolated moments* and so can only die in an *experientially isolated moment*. Thus, for them, death cannot even mean, as we casually say, “end of story”, since their inner, experiential story ended at some previous time. The upshot is that it is difficult to see how the death of such a temporally isolated form of self-awareness can constitute any kind of harm to someone who exists in this condition (i.e., independently of how others around them might feel about this and independently of questions regarding pain and suffering). It would simply mean that the person happens to die in this moment rather than some other moment – and that’s it. But this is, in effect, what is happening to them at the experiential level in each moment anyway. How, then, can the death of their temporally isolated self-awareness harm them any further?

These considerations lead us to the heart of the question of what it is, exactly, that makes death a harm in and of itself. The only plausible answer to this question surely lies in the fact that it cuts us off forever from our self-evidently valuable awareness of our own existence over time. To put it in the personal form, my death cuts me off forever from my memory claims upon *my* past, *my* dreams, plans, and projects for the future, and *my* self-aware location of the present in *this* autobiographical context, and I value these capacities (if not always their contents) and do not want to lose them. But if this is what is wrong about any form of death that cuts us off from these things, then it means that the kind of death we mean when we refer to death as a harm must be the death of our temporally structured sense of self-awareness; the death, in other words, of our autobiographical sense of self. Conversely, it follows that biological death, sentience death, or the death of temporally isolated self-awareness cannot harm us any further if we have already lost our temporally structured sense of self-awareness (and, thus, our autobiographical sense of self) since we would have already died to ourselves, as it were, in the only sense that can really matter to us.

These considerations suggest that the kind of death that we mean – that we must mean – when we refer to death as a harm is the death of our sense of our own existence through time, the death of our autobiographical sense of self. However, it is important to realize that autobiographical

death is not necessarily an all-or-nothing affair: whereas we tend to think of biological death in something approaching an either/or categorization (at least until the final period of the demise of a being we tend to pronounce it as being either “alive” or “dead”), autobiographical death can come in degrees, that is, we can lose some aspects of our autobiographical capacities without losing all of them. We can, for example, begin to lose aspects of the “brainware” we need in order to sustain our autobiographical selves. Dementia, which is unquestionably a form of harm to its victims, is a well known example of this process. But even if our brainware remains in good working order, it is also possible that we can lose whole sections of the autobiographical memories that serve to sustain our autobiographical selves for other reasons. These could conceivably vary from psychological trauma to more exotic causes such as being on the receiving end of a “memory zapping device” (if such devices exist or came to exist, then they would deserve to be called “weapons of autobiographical destruction”). Anything, in short, that diminishes the *capacities* that someone has to sustain their autobiographical self represents a form of harm to them. This means that the kind of harm that we mean – that we must mean – when we refer to death as a harm is not biological death per se, sentience death per se, or even the death of temporally isolated self-awareness, but rather the *death or diminishment of our autobiographical capacities*. I therefore suggest that we refer to this kind of harm as *autobiographical capacity harm*.

Well, I suggest this with one proviso. Just as we previously needed to distinguish between causing pain and suffering per se and causing *unnecessary* pain and suffering in order to identify only the *unnecessary* form of pain and suffering as a form of harm, so we need to distinguish in this context between causing the death or diminishment of autobiographical capacities per se and causing the *unwanted* death or diminishment of autobiographical capacities in order to identify only the latter as a form of harm. This is because there will be certain circumstances in which a person might, say, wish to have certain traumatic memories excised (if that could be done) or have a genuine and justifiable wish to die. Now if it is entirely unreasonable to think that these wishes will be reversed or if it is the case that they cannot be reversed (e.g., they might have previously expressed a considered wish – perhaps in a “living will” – for their life to be terminated in certain kinds of circumstances and now be in an irreversible non-conscious state), then agents who actively terminate or actively assist in terminating the person’s memories or life, either painlessly or in the most painless possible way under the circumstances, are clearly helping the person rather than harming them since they are giving the person both what they really

want and what is in their interests in terms of cessation from further pain and suffering (including mental distress). Obviously strict tests need to be met (e.g., concerning soundness of mind, degree of deliberation, and medical circumstances) in order to satisfy the criteria associated with a genuine and justifiable wish for these procedures to be carried out. But whatever the legal situation might be in regard to these matters, we are not harming a person *from the perspective of the person concerned* if we assist them in those cases in which these tests have been met. In contrast, we generally think that we are harming them in the most intrusive or extreme way possible if we cause the death of either some of their memories or their entire temporally structured sense of self-awareness (i.e., the death of their autobiographical sense of self) when they do not want these things; if, in other words, we cause their *unwanted* autobiographical death or diminishment.

In conclusion, then, there are two basic forms of harm: *unnecessary pain and suffering*, which we can also refer to as *affective harm*, and *unwanted autobiographical death or diminishment*, which we can also refer to as *autobiographical capacity harm*.

## V. WHICH KINDS OF BEINGS CAN BE HARMED IN WHICH WAYS?

Once we have identified and clarified what the basic forms of harm are, it becomes important to know which kinds of beings can be harmed in which ways. Consider the question of sentience first. There is inevitably a grey area in terms of precisely where sentience begins in the animal kingdom, but the evidence we have for the connection between the degree and complexity of central nervous system organization that an entity has and its capacity for some degree of sentience is overwhelming. Accordingly, Peter Singer reasonably suggested in the first, 1975 edition of his influential book *Animal Liberation* that, on the basis of what we know about animal behavior and physiology, it seemed reasonable to draw the sentience cut-off line somewhere between the more developed nervous systems of crustaceans, such as lobsters, crabs, prawns, and shrimps, and the less developed nervous systems of mollusks, such as oysters, scallops, and mussels. In Singer's view, this meant that it was not permissible for us to eat shrimps, for example, but it was permissible for us to eat oysters. However, by the time of the second, 1990 edition of his book, Singer had decided that he couldn't be sure that mollusks didn't feel pain and that since it was easy to avoid eating them, it was better to do so. If we do this, then according to Singer, "This takes us to the end of the evolutionary scale [sic; it's a tree or a bush, not a scale], so far as creatures we normally eat are concerned; essentially, we are left with a vegetarian diet."<sup>7</sup> For the philosopher who has done the most to

champion what we might in this context refer to as the “ethics of affective harm” – that is, ethics focused on questions of pain and suffering – the cut-off point for sentience in the animal kingdom may therefore exclude some living things that are classified as animals (e.g., sponges and corals are animals too), but it does include pretty much any animal we might want to eat.

As far as causing death per se is concerned, Singer notes in the opening chapter of his book that the issue of the “wrongness of killing a being is more complicated” than that of “inflicting suffering on animals” and that he has therefore “kept, and shall continue to keep, the question of killing *in the background*” (my emphasis) so as to focus on the issue of animal pain and suffering.<sup>8</sup> Singer then sidelines the issue of causing death per se until devoting three pages to it in the final chapter of his book, at which point he reiterates that “I have kept [this issue] in the background up to this point ... because it is so much more complicated than the wrongness of inflicting suffering.”<sup>9</sup> But is the fact that a moral issue is *difficult* a good enough reason for Singer not to address it throughout the main body of his book when, as I have argued here, death (or diminishment) of a particular kind represents *the* other basic form of harm? Instead of seriously grappling throughout his book with the admittedly more complicated question of whether or not it is permissible to eat other sentient beings if they are decently reared and then killed with a minimum of pain – a point on which he is ultimately equivocal – the fact that Singer sidelines the issue of causing death per se throughout the main body of his book allows him to jump to vegetarian conclusions all too easily along the way.

All of which brings us to the question of which kinds of beings can be harmed by the kind of death that actually constitutes a form of harm. As I have argued, this means the question of which kinds of beings possess a sense of their own existence through time, that is, an enduring temporally structured sense of self-awareness or autobiographical sense of self. This issue seems to have a much clearer answer in terms of the empirical evidence that bears on it than that of precisely where sentience begins and ends in the animal kingdom. The answer is essentially this: the only kinds of beings who possess an autobiographical sense of self are linguistically-enabled beings. By “linguistically-enabled beings” here I am referring not to beings that merely have the potential to learn language, but rather to beings that are actually enabled by – that are, if you like, actually “running” – language in their mental operations. And by “language” here I mean language as that term is formally understood in linguistics, that is, as involving the use of symbols in the context of a generative grammar (i.e., a set of shared rules that determine the ways in which these symbols can be used in order to

be meaningful to the rest of the group who share that particular form of symbolic communication). Now it is certainly the case that there are rare examples of people who possess some or even a high degree of language function but who suffer from some kind of neurological or psychiatric condition that prevents them from possessing an autobiographical sense of self. However, my claim here is not that the possession of language guarantees the possession of an autobiographical sense of self in every single case but rather that it enables the development of an autobiographical sense of self in all normally developed individuals who are not suffering from some intervening neurological or psychiatric condition. Conversely, my claim here is that there is no possibility of developing an autobiographical sense of self in the absence of possessing language (i.e., language might not be sufficient for autobiographical awareness, but it is necessary). This is a strong claim; what is the evidence for it?

#### VI. WHAT IS IT LIKE TO BE AN OTHERWISE INTELLIGENT HUMAN BEING WHO LACKS LANGUAGE?

Perhaps the starkest evidence for the claim I have just made comes not from comparative studies that attempt to get at the question of whether our closest evolutionary relatives possess a sense of self that persists through time but rather from looking, more directly, at the question: What is it like to be an otherwise intelligent human who lacks language? In my view this is a profoundly important and much neglected question. I have therefore spent some time in recent years trying to locate first-person accounts by those extremely rare examples of people who:

- (i) were born deaf or became deaf as infants and were not exposed to sign language until relatively late in the course of normal linguistic development but who were then able to learn language to the point where they were later able to tell us what it was like to have been without language;
- (ii) had learned language in the normal course of their development but then lost it completely (global aphasia due to stroke) before subsequently recovering to the point where they too could tell us what it was like to have been without language.

I have also been interested in the second-person accounts of people without language that have been provided by their language teachers or other sensitive observers such as neuropsychologically trained clinicians.

Now I would dearly like to quote from all of these accounts at considerably greater length than I am going to – they are “mind-blowing” – but given that I will be exceeding my word limit allocation as things stand, that these accounts are available in the public domain (albeit in some little known locations), and that I have quoted from and discussed some of them at greater length elsewhere<sup>10</sup>, I can only provide the briefest overview of these accounts in this context in order to devote as much of this paper as possible to my primary line of argument.

Suffice to say, then, that first-person accounts by people who did not gain language until relatively late in their development suggest that until they learned language they lacked both self-awareness and temporal awareness. In other words, they had no sense of any autobiographical awareness. The most famous first-person account – and, indeed, one of the very few first-person accounts we have – is undoubtedly that of Helen Keller (1880-1968) who only began to learn a tactile form of sign language (since she was both deaf and blind) at the age of seven. In an extremely insightful essay, which bears the revealing title “Before the Soul Dawn” (and which has been out of print for many years until its republication in 2003), Keller tells us that:

Before my teacher came to me I did not know that I am. I lived in a world that was a no-world... My inner life, then, was a blank without past, present, or future... When I learned the meaning of “I” and “me” and found that I was something, I began to think. Then consciousness [presumably she means consciousness in the sense of self-awareness here] first existed for me... Thought made me conscious of love, joy, and all the emotions... and the blind impetus, which had before driven me hither and thither at the dictates of my sensations, vanished forever.<sup>11</sup>

The contemporary deaf French actress Emmanuelle Laborit (born 1971), who likewise only began to learn sign language at the age of seven, says that:

I believe there was no sense whatsoever of time progression in my mind during that period. Past, future, everything was on the same time-space line... I was completely helpless, completely unaware of time passing. There was daylight and the darkness of night, and that was it.

I still can't assign dates to things during the period from my

birth to age seven, or arrange what I did in chronological order. Time was in a holding pattern. I just experienced things as they happened. ...I lived each [event] as an isolated experience, in the present.<sup>12</sup>

It was only in the context of learning language that Laborit was released from her isolated state of non-self-aware existence in the present and realized that she existed, that she “was somebody”: “I was seven years old. I had just been born and come of age in one fell swoop.”<sup>13</sup>

This lack of any temporal sense, in particular, fits with the impressions that sensitive neuropsychiatric observers such as Oliver Sacks have of people in this situation. For example, Sacks provides an account of Joseph, an apparently intelligent and inquisitive deaf boy, who had not been diagnosed as deaf until the age of four and had not been exposed to sign language until his entry to the school in which Sacks met him at the age of eleven. Sacks, who notes that Joseph was only “just beginning to pick up a little Sign”, comments as follows:

It was not only language that was missing: there was not, it was evident, a clear sense of the past, of “a day ago” as distinct from “a year ago”. There was a strange lack of historical sense, the feeling of a life that lacked autobiographical and historical dimension, the feeling of a life that only existed in the moment, in the present.<sup>14</sup>

These first-person accounts and clinical impressions fit with the impressions that the teachers of such people have. Susan Schaller provides a remarkable account of teaching sign language for the first time to an otherwise intelligent twenty-seven year old deaf Mexican student named Ildefonso. (Although Ildefonso does manage to learn sign language, he never manages to tell Schaller what it was like to have been without language at the experiential level, only how hard it was for him before that time.) Schaller notes that Ildefonso “had no concept of time as we learn it”; that for a long time he “could not understand any lesson on time”; and that nothing remained more difficult throughout the course of teaching him than trying to get him to understand temporal concepts.<sup>15</sup> In reviewing other contemporary cases of late language learners, Schaller found the same problem and notes in regard to another case that “The most difficult task, *as usual*, was schedules and time. The student’s only time was the present” (my emphasis).<sup>16</sup>

Feral, or wild, children – children who have survived from a very young age in the company of animals – never successfully manage to learn language, which is not surprising given that they have experienced a fundamentally different form of socialization to languageless children who have been brought up in otherwise normal human company. But the impressions that close observers glean from these languageless children nevertheless serve to reinforce the general picture that has emerged from the accounts we have just considered by and about languageless people who have been otherwise normally socialized. The science writer John McCrone provides the following composite picture of feral children based on his examination of a range of accounts by those who have had direct contact with them:

[T]hey seemed somehow to lack memory and self-awareness... [their thoughts were] limited to the world of the here and now... They could make simple associations and learn to recognize familiar people and situations. But they seemed unable to reflect on the past or the future, or to have any insight into their own plight.<sup>17</sup>

If we now turn to those similarly rare first-person accounts by people who had language but then lost it completely (global aphasia due to stroke) before subsequently recovering to the point where they too could tell us what it was like to have been without language, then we again find they lack any temporally structured sense of self-awareness. Reflecting on the loss of language that followed her stroke, the neuroanatomist Jill Bolte Taylor says:

Instead of a continuous flow of experience that could be divided into past, present, and future, every moment seemed to exist in perfect isolation... I stopped thinking in language and shifted to taking new pictures of what was going on in the present moment. I was not capable of deliberating about past or future-related ideas because those cells were incapacitated.<sup>18</sup>

Reflecting on the loss of language that followed his stroke, the psychologist Scott Moss wrote:

It was as if the stroke had benumbed any emotional investment in the future and I simply shrugged at my perception of my imminent demise... If I had lost the ability to converse with others, I had also lost the ability even to engage in self-talk [i.e., he

“could not use words even internally” in thinking to himself]. In other words, I did not have the ability to think about the future – to worry, to anticipate or perceive it – at least not with words. Thus, for the first five or six weeks after hospitalization I simply existed... It was as if without words I could not be concerned about tomorrow...

It was also fascinating to me how completely and totally fixed I was on the “here and now”... So both the past and the future had faded for me, and I existed almost exclusively in the present... I was unable to generate a gestalt of either my previous life or the future, and therefore life beyond the immediate situation was meaningless.<sup>19</sup>

These rare reports by and about people who lack language suggest that to be without language is to be *time blind*. But why should this be the case? What is it about the nature of language that enables its users to have a sense of themselves dwelling in time and those who lack it to have no temporally structured sense of self-awareness, no sense of autobiographical time?

## VII. THE POWER OF LANGUAGE: WHY DOES LANGUAGELESSNESS EQUATE TO TIMEBLINDNESS? WELCOME TO THE WORLD OF INDICES, ICONS, AND SYMBOLS

Following a set of distinctions first set out by the American philosopher C. S. Peirce (1839-1914), those who study communication commonly accept that there are three basic ways in which we can refer to anything, namely, by means of *indices*, *icons*, and *symbols*. An index is anything that refers to something by in some way directly indicating, or pointing to, its presence in the immediate environment. (The term *index* derives from the Latin for pointer, from *indicāre* to disclose, show; hence, we refer to our “pointing finger” as our “index finger”.) This can, of course, be achieved by means of pointing with an index finger, but it can also be achieved by means of making a sound – say, a particular type of grunt – that vocally indicates the immediate presence of something in a one-sound-one-meaning kind of way in the immediate presence of the stimulus with which it is associated (such as a snake or a bird of prey overhead).

In contrast, an icon does not achieve its referential aim by directly indicating something in the immediate environment but rather by means of displaying a likeness or resemblance to it. (The term *icon* derives from the Greek *eikōn*, image, from *eikenai*, to be like.) Iconic forms of reference

can be quite simplified and schematic, such as the icons on the computer I am using as I write this or the icons on the doors that indicate women's and men's toilets (indeed, sometimes these icons can be too schematic!), but they can also be very detailed, such as a painted portrait or a photograph. The important point, however, is that they achieve their referential aim by means of a likeness or resemblance.

In the absence of any kind of linguistic support, an index only makes sense when it directly indicates something that is "here, now". To attempt to point to or directly indicate – either manually or vocally – something that is not here or not happening now is entirely meaningless. If we cannot observe what is being "indicated", then we have no "other end of the line", as it were, to which we can tie the indicator. Thus, pointing to a horse that isn't here or to "two days ago" is utterly meaningless. An icon, on the other hand, can resemble something that isn't "here, now", but because it relies on resemblance to achieve its referential aim, it can still only refer – again, in the absence of linguistic support – to something that has at some point been observed. To make an icon – a resemblance – of something that has never been observed or that cannot be observed is meaningless because, again, we have no "other end of the line" to which we can tie the would-be "resemblance".

In contrast, symbols are fundamentally different from either indices or icons because the essence of a symbol is that it has an *arbitrary* relationship with what it stands for; that is, it bears no necessary connection with that to which it refers. As the very root of the term suggests, a symbol and what it refers to are, as it were, just "thrown together" (*symbol* derives from the Greek, *symbolon*, sign, from *sumballein* to throw together, from *syn-*, with or together + *ballein*, to throw). But if this is the case, then how does a symbol work; how does it achieve its referential aim? We can see that an index achieves its referential aim by being linked to the presence of something "here, now" and an icon achieves its referential aim by being linked to the thing it resembles, but how on earth is a symbol supposed to achieve its referential aim if it bears no particular "here, now" or resemblance connection to whatever it is supposed to refer to? The answer is that it does so by virtue of the fact that a group of people implicitly or explicitly *agree* – or, given the age at which language is typically acquired, perhaps we should say *come to act in agreement* – that a particular otherwise arbitrary sound or shape either stands for something else or else serves to make certain kinds of connections between the sounds or shapes that stand for other things (in which case the sound or shape being employed has a grammatical function). This group then preserves these agreements and corrects those

who violate them. The nature of these agreements – these referential understandings – can change slowly over time of course, but the group concerned remain the speakers of a common language so long as they keep “travelling together” in terms of their collective implicit and explicit referential agreements.<sup>20</sup>

On first consideration, it might seem quite counterintuitive to employ a form of reference that has no obvious connection with that to which it refers. But it is precisely the fact that symbols lack any obvious connection to what they refer to that explains the power of symbolic communication. Whereas indices rely on the “here, now” presence of what is being indicated and icons rely on the fact that they in some way physically resemble their target, *symbols can refer to anything, anywhere, anytime, whether it is observable in principle or not, simply because its users agree that this is what the symbol will mean.* (Thus, we have a roughly shared understanding of what a *unicorn* is – and would certainly know one if we saw one – even though unicorns are not only unobservable but do not even exist.) In view of this, those who study comparative cognition and communication – that is, the reasoning and communicative abilities of different species – refer to indexical and iconic forms of reference as being “stimulus-bound” or “context-bound” because their use (in the case of indices) or their nature (in the case of icons) is tightly bound to the presence or nature of the stimulus being referred to. In contrast, symbols float free of the presence or nature of what is being referred to; we could substitute the sound and spelling we currently use for the word “unicorn” tomorrow and carry on just as happily, providing we all adopted this change.

This contrast between the stimulus-dependent and stimulus-independent nature of indices and icons on the one hand and symbols on the other is so important that those who study comparative communication categorize the former two as *signals* as distinct from *symbols*. And what linguists mean by *language* in a formal sense is *symbolic communication*, that is, communication that takes the form of employing symbols in the context of a generative grammar (i.e., a set of shared rules that determine the ways in which these symbols can be used in order to be meaningful to the rest of the group who share that particular form of symbolic communication). Now although there is no question that nonhuman animals communicate with each other in various ways in terms of their behaviour, linguists are in widespread agreement that language per se – the use of symbols in the context of a generative grammar – is essentially unique to humans.<sup>21</sup>

The fact that language can refer to unobservable features of the world whereas, in the absence of any kind of linguistic support, signals can't (be-

cause they are tied to observable features of the world such as things that can be pointed to or things that can be resembled), is momentous. On the basis of these considerations alone, we should expect to find that beings that lack language would not be able to refer to and therefore not be able to reason about things that are central to our understanding of the world but that are inherently unobservable. Prime examples here include the mental states of others and physical causes: we cannot refer to the mental states of others or physical causes (even when they are present or occurring here and now) purely by pointing to them or trying to draw a likeness to them in some way precisely because there is nothing we can observe that we can connect up with what is being pointed to or imaged. Rather, all we can observe directly is the behaviour of other beings rather than their mental states per se or the conjunction of certain kinds of events rather than causation per se. The conclusion that follows from all this in regard to nonhuman animals is that, contrary to what many people might prefer to think (not least because we project our own forms of interpreting the world onto other, nonhuman beings), we should not expect nonhuman animals to be able to reason about inherently unobservable phenomena such as the mental states of others and physical causes. Moreover, this is exactly what carefully controlled experimental research with chimpanzees, our closest evolutionary relatives, reveals.<sup>22</sup>

But what goes for other inherently unobservable phenomena like mental states and physical causes must also go for *time* since, like them, time can only be conceived, not perceived. As the distinguished developmental psychologist Katherine Nelson puts it in her significant book *Language in Cognitive Development*:

To recapitulate, the child alone cannot discover time, because (unlike concrete objects) it is not an entity that exists to be discovered. Rather, conceptions of process and change have led different societies to conceptualize time in different ways, and those ways are conveyed to children through language forms.<sup>23</sup>

On the basis of these formal considerations concerning the nature of language, we should therefore not expect any being that lacks language to be able to reason about time because they lack the kind of referential toolkit – symbols employed in the context of an agreed set of combination rules – that would enable them to do so. And this necessarily means that a being that lacks language is not able to develop an enduring temporally structured sense of self-awareness or autobiographical sense of self. In turn, it

follows from my previous discussion regarding the two basic forms of harm that non-linguistically enabled beings cannot be caused autobiographical capacity harm because, in the absence of language (not to mention, as I have previously noted, the range of cognitive capacities that are necessary to underpin the development of language), they have no such autobiographical capacities.

Indeed, as Nelson argues, it is not even clear that a being that lacks language and, thus, a temporally structured sense of self-awareness, can have any genuine sense of self-awareness at all:

[U]ntil the various uses of language make it possible to imagine a past and future self, and to imagine that other people have different pasts and futures, as well as different presents, one cannot speak of a fully determined self *distinct from ongoing experience*... Language uniquely enables contemplating a self that is different from present experience, and imagining a self that will grow older as well as a self that was once a little baby [my emphasis].<sup>24</sup>

### VIII. FORMS OF HARM AND OUR OBLIGATIONS TO HUMANS AND OTHER ANIMALS

I have so far argued essentially to two ends. First, there are two basic forms of harm: unnecessary pain and suffering, which we can also refer to as affective harm, and unwanted autobiographical death or diminishment, which we can also refer to as autobiographical capacity harm (sections 1-4). Second, non-linguistically enabled beings can only be harmed in terms of being caused unnecessary pain and suffering whereas linguistically enabled beings can be harmed in both ways (sections 5-7). This difference arises because only linguistically enabled beings are able to develop the kind of enduring temporally structured sense of self-awareness or autobiographical sense of self that *can* be caused autobiographical capacity harm.

If we accept that we should avoid causing harm to other beings – and I have laid out the basic logic for this (I hope noncontroversial) argument elsewhere<sup>25</sup> – then the primary upshot of the foregoing argument is as follows: it is not permissible to cause unnecessary pain and suffering to sentient beings in general or to cause the unwanted death or diminishment of the autobiographical capacities of those beings that have these capacities, which is to say normal linguistically-enabled beings. Conversely, considered in the abstract, it is permissible to cause the death per se of non-linguistically-enabled beings. However, the fact that the death of non-linguistically enabled

beings can cause unnecessary pain and suffering to other sentient beings who are psychologically invested in or emotionally close to them means that, in reality, this “considered in the abstract” clause obviously needs to be understood in conjunction with its potential to cause affective harm to those other beings (whether human or nonhuman) that are psychologically invested in or emotionally close to the beings that are killed. This means that we therefore need to understand the conclusion that it is “permissible to cause the death per se of non-linguistically-enabled beings” as subject to the following constraints.

First, we *incorporate* all kinds of non-linguistically or not yet linguistically enabled sentient beings – such as companion animals and infants, respectively – into the daily texture of our personal lives on an individual basis and treat them – and expect others to treat them – *as if* their death would be a harm to them, even if this death were painless. (I will refer to these beings as *incorporated sentient beings*.) There are straightforward reasons for this, but they lie with the fact that the deaths of these beings would cause unnecessary pain and suffering to their guardians and others who are emotionally close to them rather than the fact that their deaths would – or could – cause any autobiographical harm to the beings that die. Even so, it makes sense for us to treat companion animals and human infants *as if* their deaths would be a direct harm to them – when in fact they constitute an indirect harm to their guardians and others who are emotionally close to them – since this is, as it were, the shortest route home in terms of achieving the morally desirable outcome of avoiding causing unnecessary pain and suffering to their guardians and others who are emotionally close to them. (This “as if” status is, of course, massively reinforced in the case of human infants because they normally have the potential to develop the highly valued end-state of an autobiographical sense of self, whereas nonhuman beings do not possess this potential.) Thus, we accept – and should accept – that, under normal circumstances, it is *not* permissible to cause the death per se of what I have just referred to as *incorporated sentient beings* because of the unnecessary pain and suffering that this would cause to others.

If we were to extend this line of thinking to the social worlds of *unincorporated sentient beings* (i.e., sentient beings that we have *not* incorporated into the daily texture of our personal lives on an individual basis), then it might be prudent to avoid causing even the painless death of other highly social animals<sup>26</sup> such as great apes, cetaceans, and elephants since killing these animals might increase the possibility that we are causing distress to members of their group who are used to interacting with them. (There may of course be other kinds of reasons for not killing these animals, such as

preserving biodiversity, but I am concerned in this context with the kinds of reasons that proceed from our understanding of the ways in which these sentient beings can themselves be harmed.) By the same token, if there are circumstances in which it is deemed that some of these animals should be killed – for example, culled for environmentally related reasons – then it makes sense to cull family groups rather than similar numbers of unrelated (or not as closely related) animals, and this practice is indeed followed in some situations.

That said, there are real questions about the extent to which we *can* extend this line of thinking to other, unincorporated sentient beings (hence, I noted only that “it might be prudent” to act this way). This is because any loss that these non-linguistically enabled beings experience must be of a far more restricted order than that experienced by autobiographically imbued beings such as ourselves. The death of a non-linguistically enabled being cannot represent a significant autobiographical marker in the lives of any members of its group because these beings lack autobiographical capacities in the first place, nor can any loss that these animals experience be based on some kind of sorrow for the loss of the (unobservable) inner world of the being that has died because they have no sense of this either.<sup>27</sup> Rather, any loss that a highly social nonhuman animal experiences would have to be based purely on a short-term (not autobiographically remembered) sense of the loss of a particular, familiar form of enjoyable interaction. But even this cannot apply – or apply to anything like the same extent – in regard to the kinds of animals that we usually keep or treat as stock animals (e.g., cattle, pigs, goats, sheep, poultry, and fish) since these effectively constitute collections of individuals rather than anything approaching the degree of social organization and enduring individualized (and, thus, not readily substitutable) forms of interaction that are to be found in highly social animals.

Taken together, these considerations lead to the more qualified conclusion that it is permissible to cause the painless death of *unincorporated* sentient beings that *do not belong to highly social species*. Thus, so far as the question of death per se goes, it is permissible to cause the painless death of the kinds of animals that we usually keep or treat as stock animals as well as many kinds of wild mammals, birds, fish, amphibians, and reptiles. In reality, however, the deaths that we do cause these animals often do involve pain, sometimes considerable pain. Does this then mean that although it may be permissible in principle to cause the painless death of unincorporated sentient beings that do not belong to highly social species, we should not cause these deaths in practice because doing so typically involves causing pain? This is the point at which all thinkers concerned with this issue need to get

out of their analytical armchairs, as it were, and “get real”.

I have separated out the issues of there being two basic forms of harm in order to be as clear as possible about these forms of harm and which kinds of beings can be caused which forms of harm. But once we have done this we are then obliged to “get real” in the sense of thinking about these matters in “real world” terms. When we do this, then the first thing we need to realize is that it is just a blunt and inescapable fact of ecological life that the kinds of deaths that unincorporated sentient beings typically experience in nature are rarely painless, to put it mildly. As Mark Sagoff argues, animals typically die violently in nature through predation, starvation, disease, parasitism, and cold; most do not live to maturity and very few die of old age; and many might “reasonably prefer to be raised on a farm, where the chances of survival for a year or more would be good, and to escape the wild, where they are negligible.”<sup>28</sup> The next thing we need to realize is that, as Jeremy Bentham – Singer’s own inspiration for his “animal liberation” approach to animal ethics – put it: “The death they [i.e., that nonhuman animals] suffer in our hands commonly is, and always may be, a speedier, and by that means a less painful one, than that which would await them in the inevitable course of nature.”<sup>29</sup> Thus, animals will typically experience less pain in the course of being killed humanely than they would in the course of dying or being killed “naturally”.

When we put these real world considerations together with the argument presented here to the effect that death per se cannot be a harm to non-linguistically-enabled beings, then the general conclusion follows that although we should in all cases avoid causing unnecessary pain and suffering to sentient beings, it is nevertheless permissible – because it will typically cause less pain and suffering than would otherwise be the case – to cause the deaths of unincorporated sentient beings that do not belong to highly social species so long as we seriously seek to minimize the pain and suffering we cause in doing this.

In conclusion, I note two final contextualizing points in regard to the argument I have presented here and the question of our obligations towards nonhuman animals in particular. First, there are many possible grounds for arguing the case for vegetarianism or veganism. These can range from dietary considerations regarding human health, to wider humanitarian (“feeding the human population more efficiently”) kinds of arguments, to ecologically (including biodiversity) based kinds of arguments, to straightforward moral arguments regarding our obligations to nonhuman animals regardless of any of the foregoing arguments. The focus of this paper is only relevant to this last, “moral” category of arguments and, as I hope I have

made clear, the implication of my argument in regard to the question of “moral vegetarianism” or “moral veganism” is that it is permissible to cause the swift death of certain kinds of nonhuman animals, but that we have a serious obligation to see that they have been treated decently in their lives (i.e., that they have not been caused unnecessary pain and suffering) prior to their deaths. That said, I have made no claims one way or the other about the success or otherwise of the other dietary, humanitarian, and ecological kinds of arguments I have just mentioned, each of which is complex in its own right. The second point to note here is that – as indicated by the title of my paper – I have only been concerned in this paper with the question of harming, not the question of helping; thus, I have only been concerned with our *negative* obligations to humans and other animals (i.e., with where the boundaries begin and end in regard to what it is not permissible to do to them) rather than with any *positive* obligations that we might also have to help humans and other animals (i.e., with where the boundaries begin and end in regard to what we are positively obliged to do for them). The question of our positive obligations to humans and other animals is another, similarly complex issue, which I have discussed at some length elsewhere<sup>30</sup>, but that discussion does not alter the conclusions reached here in regard to our negative obligations.

## NOTES

1. For a related discussion concerning the importance of distinguishing what has variously been referred to as a first-person, personal, subjective, or internal perspective from a third-person, impersonal, objective, or external perspective in the context of discussing the value of different kinds of entities, see my “Foundations of a General Ethics: Selves, Sentient Beings, and Other Responsively Cohesive Structures”, in *Philosophy and the Environment*, Royal Institute of Philosophy Supplement: 69, edited by Anthony O’Hear (Cambridge: Cambridge University Press, 2011), 47-66; also available from the “sample online papers” section of my website at [www.warwickfox.com](http://www.warwickfox.com).
2. David Hume, *A Treatise of Human Nature* (1739), Book III, Part I, Section I (any edition).
3. *Collins English Dictionary: Complete and Unabridged*, 6<sup>th</sup> ed. (Glasgow: HarperCollins, 2003).
4. Katherine Nelson, *Language in Cognitive Development: The Emergence of the Mediated Mind* (Cambridge: Cambridge University Press, 1998), 344.
5. Clive Wearing’s wife, Deborah Wearing, provides an intimate account of her husband’s condition in *Forever Today: A Memoir of Love and Amnesia* (London: Corgi Books, 2005); Oliver Sacks offers a clinical neurologist’s perspective on Wearing’s condition in “The Abyss”, *The New Yorker*, September 24, 2007, 100-112 (also available online).
6. Deborah Wearing notes her husband saying this four months after the onset of his illness in 1985 (*Forever Today*, 181); Sacks notes Wearing saying this to him twenty years later in 2005; and Wearing says the same thing in a television documentary I have seen on him. Given that repetition is a post-illness feature of Wearing’s speech – not that Wearing

- would be aware that he is repeating himself – one assumes that he has said this in various other contexts as well and that it represents a sustained – albeit continually forgotten – view of his condition “from the inside” and not a one-off rhetorical flourish.
7. Peter Singer, *Animal Liberation*, 2<sup>nd</sup> ed. (London: Jonathan Cape, 1990), 174-175.
  8. *Ibid.*, 17.
  9. *Ibid.*, 228; his full discussion of this issue runs from 228-230.
  10. Warwick Fox, *A Theory of General Ethics: Human Relationships, Nature, and the Built Environment* (Cambridge, MA: The MIT Press, 2006), see esp. chapters 6-8.
  11. Helen Keller, *The World I Live In* (New York: New York Review Books, 2003; first pub. 1908), 72, 74, 76. I quote considerably more of Keller’s account in *A Theory of General Ethics* and show how it maps onto research findings that speak to the question of what it is like to be a nonhuman primate.
  12. Emmanuelle Laborit, *The Cry of the Gull* (Washington, D.C.: Gallaudet University Press, 1998), 7.
  13. *Ibid.*, 2.
  14. Oliver Sacks, *Seeing Voices: A Journey into the World of the Deaf* (London: Picador, 1991), 39-40.
  15. Susan Schaller, *A Man Without Words* (London: Ebury Press, 1992), 116 and 194-95; see also chapter 11 and Afterword.
  16. *Ibid.*, 197.
  17. John McCrone, *The Myth of Irrationality: The Science of the Mind from Plato to Star Trek* (New York: Carroll & Graf Publishers, 1993), 104.
  18. Jill Bolte Taylor, *My Stroke of Insight: A Brain Scientist’s Personal Journey* (London: Hodder and Stoughton, 2008), 48, 68.
  19. C. Scott Moss, *Recovery with Aphasia: The Aftermath of my Stroke* (Urbana, IL: University of Illinois Press, 1972), 4-5 and 10. Extensive extracts from this book (which include the very brief extracts I have just given) are also available in *Injured Brains of Medical Minds*, edited by Narinder Kapur (Oxford: Oxford University Press, 1997). Kapur’s book also contains an additional overview article by Scott Moss, written seven years after his stroke, in which he reiterates that (even up to five months after his stroke) “my entire existence was focused exclusively on the here-and-now. The stroke [which, we must not forget, resulted in complete loss of language function] had abolished my memory of the past and projection of the future had no meaning for me” (p. 78).
  20. Although this “social agreement” solution to the problem of symbolic forms of reference might sound like a simple and straightforward one, it isn’t. The range of cognitive capacities that are necessary to underpin this level of social understanding – such as “theory of mind” (or understanding the intentions and perspectives of others) and “joint”, “shared”, or “triadic” forms of attention – are formidable and serve to explain why only normally developed humans are capable of symbolic communication. I discuss these cognitive achievements, and the research findings that bear on them, at some length in *A Theory of General Ethics*, chapters 6-8.
  21. For more on all the matters discussed in this section – the nature of indices, icons, and symbols; the differences between signalling and language; and the uniqueness of human language – see Terrence Deacon, “Biological Aspects of Language”, in *The Cambridge Encyclopedia of Human Evolution*, edited by Steve Jones, Robert Martin and David Pilbeam (Cambridge: Cambridge University Press, 1994), 128-133; Terrence Deacon, *The Symbolic Species: The Co-evolution of Language and the Human Brain* (London: Allen Lane, Penguin, 1997), esp. chapter 3, appropriately entitled “Symbols aren’t Simple”; Peter Gärdenfors, *How Homo Became Sapiens: On the Evolution of Thinking* (Oxford: Oxford University Press, 2003), esp. chapter 6; and my *A Theory of General Ethics*, chapters 6-8.

22. For the best short overview introduction to these research findings for a general audience see Daniel Povinelli, "Behind the Ape's Appearance: Escaping Anthropocentrism in the Study of other Minds", *Daedalus: Journal of the American Academy of Arts and Sciences*, Winter 2004, 29-41. Povinelli's team has been at the forefront of this research; thus, see also: Daniel Povinelli and Timothy Eddy, *What Young Chimpanzees Know about Seeing*, *Monographs of the Society for Research in Child Development* (Chicago: The University of Chicago Press, 1996); Daniel Povinelli, "Can Animals Empathize?", *Scientific American Presents: Exploring Intelligence* 9.4 (1998): 67, 72-75; Daniel Povinelli, *Folk Physics for Apes: The Chimpanzee's Theory of How the World Works*, reprint edition (Oxford: Oxford University Press, 2003); Daniel Povinelli and Jennifer Vonk, "Chimpanzee Minds: Suspiciously Human?", *Trends in Cognitive Sciences* 7.4 (2003): 157-160; Jennifer Vonk and Daniel Povinelli, "Similarity and Difference in the Conceptual Systems of Primates: The Unobservability Hypothesis", in *Comparative Cognition: Experimental Explorations of Animal Intelligence*, edited by Edward Wasserman and Thomas Zentall, 363-387 (Oxford: Oxford University Press, 2006); Derek Penn and Daniel Povinelli, "On the Lack of Evidence that Chimpanzees Possess Anything Remotely Resembling a 'Theory of Mind'", *Philosophical Transactions of the Royal Society*, B, 362 (2007): 731-744; Derek Penn, Keith Holyoak, and Daniel Povinelli, "Darwin's Mistake: Explaining the Discontinuity between Human and Nonhuman Minds", *Behavioral and Brain Sciences* 31 (2008): 109-130; Daniel Povinelli, *World without Weight: Perspectives on an Alien Mind* (Oxford: Oxford University Press, 2012). See also Michael Tomasello's important book *The Cultural Origins of Human Cognition* (Cambridge, MA: Harvard University Press, 2000). Again, I refer to and discuss much of this research in *A Theory of General Ethics*, chapters 6-8.
23. Katherine Nelson, *Language in Cognitive Development: The Emergence of the Mediated Mind* (Cambridge: Cambridge University Press, 1998), 288. See also Daniel Povinelli, "The Self: Elevated in Consciousness and Extended in Time", in *The Self in Time: Developmental Processes*, edited by Chris Moore and Karen Lemmon, 75-95 (Mahwah, NJ: Lawrence Erlbaum, 2001).
24. Nelson, *ibid.*, 344.
25. Warwick Fox, "Foundations of a General Ethics: Selves, Sentient Beings, and Other Responsively Cohesive Structures" (details in n. 1, above; see pp. 50-52).
26. I use the term "social animals" here in the formal sense to refer to a level of social organization that goes beyond mother-offspring bonding to include "permanent groups of adults living together, and relationships between individuals that endure from one encounter to another". Although there can be reasons to be wary of information on *Wikipedia* on some occasions, I take this definition of "social animal" from the *Wikipedia* article on this topic on this occasion since it seems to me to be both succinct and reliable.
27. See Gärdenfors, *How Homo Became Sapiens*, 131-132, on the apparent lack of awareness of death in chimpanzees and baboons.
28. Mark Sagoff, "Animal Liberation and Environmental Ethics: Bad Marriage, Quick Divorce", reprinted in *Environmental Philosophy: From Animal Rights to Radical Ecology*, 3<sup>rd</sup> ed., edited by Michael Zimmerman, 87-96 (Upper Saddle River, NJ: Prentice Hall, 2001), 92.
29. Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation* (Mineola, NY: Dover Publications, 2007; first pub. 1789), 311.
30. Fox, *A Theory of General Ethics*, see esp. chapters 5-8.

