The Post-Darwinian Transition[*]

DAVID PEARCE

Introduction

> *"As often as Herman had witnessed the slaughter of animals and fish, he always had the same thought: in their behaviour toward creatures, all men were Nazis"*
>
> Isaac Bashevis Singer

Most people who approach *Taking Animals Seriously* will share an unspoken presupposition. This is that animal activists take animals *too* seriously. They lack a sense of proportion. It's not that gratuitous cruelty to members of other species is morally defensible. Surely it isn't. If pressed, then all but the amoral, sociopathic or philosophically bewitched are likely to grant that wanton animal-abuse is best discouraged. Instead, the pervasive assumption is simply that animal suffering doesn't really *matter* much compared to the things that happen to human beings - to *us*. *They*, after all, are only *animals*: *objects* rather than our fellow *subjects*. Animal consciousness, insofar as it exists at all, is minimal and uninteresting.

Contrast one's likely reaction on learning that the infant or toddler next door is being abused. Let's suppose that the abuse is being inflicted for fun or profit - or, more broadly, for purposes that can be described only as frivolous. In such a case, then one's intuitions are equally clear. The suffering of the victim has to be taken very seriously. One has a duty actively to prevent it. The interests of the child *take precedence over* the wishes of the abuser. In extreme cases, the adults involved in persistent abuse may need to be legally restrained or even locked up. Indeed, it is cases of *failure* on our part to take action to prevent it - or *failure* to take action by the social services or child-protection agencies - that demand justification. To treat the suffering caused by child-abuse *lightly* would be to show a sense of disproportion when confronted with the nature of the practices involved - and our capacity to do something about them.

Yet here lies the crux.

After Darwin, a huge and accumulating convergence of physiological, behavioural, genetic and evolutionary evidence suggests - but cannot prove - an appalling possibility. This is that hundreds of millions of the non-human victims of our actions are functionally akin - intellectually, emotionally and in their capacity to *suffer* - to very young humans. In the light of what we're doing to our victims, the consequences of their also being *ethically* akin to human babies or toddlers would be awful; in fact, almost too ghastly to think about.

When we're confronted with such an emotive parallel, all sorts of psychological denial and defence-mechanisms are likely to kick in. Undoubtedly, too, animal-exploitation makes our lives so much more convenient. Not surprisingly, in view of what we're doing to them, there is a powerful incentive for us as humans to rationalise our actions.

Numerous pretexts and rationalisations aimed at legitimating animal exploitation are certainly available; most of them seek to magnify the gulf between "us" and "them". Intellectually, however, they prove on examination to be surprisingly thin.

Some of the alleged differences between "them" and "us" are entirely spurious: humans alone have souls, we are asked to believe, or enduring metaphysical egos. Other inter-species are genuine. There are the dissimilarities of gross physical appearance; the neuroanatomy of

[*] http://www.hedweb.com/animals/transit.htm

Broca and Wernicke's areas; the capacity of certain mature humans to define allegedly reciprocal notions of right and duty; or perhaps the elaborate network of social relationships in which typical human child-rearing practices are situated. Human babies and veal calves aren't literally carbon-copies of each other. Nor is the development of an individual organism just a fast-forward re-run of evolutionary history. So *pace* Haeckel, it's not simply the case that "ontogeny recapitulates phylogeny". Yet once one accepts that inflicting readily avoidable suffering *per se* is morally wrong, then it is questionable how such differences that do exist between human and (at least) advanced vertebrate non-human beings are *morally* relevant differences.

This argument isn't likely to sway the radical sceptic about animal consciousness. For in trying to appraise the sentience of other living beings - even one's adult fellow humans - it is notoriously hard to *prove* anything at all. The price of intellectual rigour, however, is a morally frivolous solipsism-of-the-here-and-now. Without merely begging the question, there's simply no logically compelling ground - just Santayana's "blind animal faith" - for believing that anything exists beyond the contents of *this* current frame of consciousness. Yet one wouldn't, for instance, let an (ostensibly) floundering toddler drown in a pond on the grounds of one's rational incapacity to penetrate beyond the veil of perception, devise a satisfactory theory of meaning, or prove the veridicality of memory, etc. Nor would one let the toddler perish because one intellectually believed that value-judgements were subjective and ethical claims truth-valueless. For when the consequences of being wrong are so terrible, then ethically one just has to play safe.

In this review essay, at least, the more radical forms of philosophical scepticism about mind - though not about ethics - will simply be set aside. Such neglect may be justified on the grounds that if one were the proverbial brain-in-a-vat etc, then no *harm* would come from acting (pseudo-)morally; albeit no good either.

Instead, rather than attempting to defeat the sceptic, a less counter-intuitive and naturalistic metaphysic will simply be assumed. Reality is indeed outlandishly weird in some of its properties. Yet there actually is a mind-independent world populated by embodied fellow subjects of experience; if there isn't, then one is harmlessly talking to oneself. Within the mind-independent world, there are fellow creatures who suffer, sometimes quite horribly. And granted merely that functionally equivalent young humans do sometimes suffer intensely, it seems overwhelmingly probable [see below] that the non-humans we treat as disposable objects of our convenience suffer horribly from what we do to them as well. If it can defensibly be argued that it's inherently morally wrong to harm and kill small children, then by parity of reasoning it is morally wrong to harm and kill functionally equivalent non-human victims too. To argue otherwise, it would be necessary either to dispute the premise, or alternatively to show that there are *morally* relevant differences between any human and any non-human which license our inconsistent attitudes and behaviour towards the two groups.

Radical scepticism again aside, one might still hope, usually on unspecified grounds, that the neurochemical substrates that mediate pain, anxiety and terror in humans may mediate a providentially different texture of experience in our fellow vertebrates - or perhaps some sort of low-grade sentience which we don't seriously have to bother about. Once again, one can't *prove* that they don't. Perhaps the astonishing evolutionary conservation of neurochemical pathways which underlie nociception ["pain-perception"] construed in a narrowly physiological sense - involving serotonin, the periaquaductal grey matter, bradykinin, ATP receptors, the major opioid families, substance P etc - may amount to a wildly misleading coincidence; or are part of a spooky conspiracy designed to mislead us. Or again, perhaps the kinds of aversive experience that non-humans undergo are still rather dull and dim - akin to some of our own aches or itches. They may be a bit unpleasant, but they're scarcely of deep moral consequence.

Yet not merely is this type of optimistic - or self-servingly sceptical - perspective radically non-Darwinian. It also violates the principle of the uniformity of Nature. The uniformity of Nature is a principle that admittedly flies under numerous variant formulations. It undoubtedly lends itself to all manner of philosophico-scientific subtleties. Yet complications aside, the existence of some sort of constancy of natural law is an assumption on which any non-sceptical account of human knowledge - or even mutually intelligible discourse - depends. So the onus of proof is on someone who seeks to *deny* some such basic uniformity - or makes an *ad hoc* exception just in the realm of the organic physiology of consciousness - to explain why the principle allegedly breaks down precisely at the most morally expedient place for *homo sapiens*.

Now the idea that our descendants might regard our treatment of the creatures we hunt, butcher and factory-farm today in the sort of light we ourselves regard the abuse of human infants is - to typical Western scientific minds at least - intuitively *absurd*. At face value, it just isn't *credible*. Animal-abusers and child-abusers occupy radically different categories in our scheme of things. Yet this hypothesised gulf rests fundamentally on intuition; not on argument. Over the millennia, it has been genetically adaptive for us to exploit other creatures. Using them as expendable objects has helped strands of human self-replicating DNA leave lots more copies of itself ("maximise its inclusive fitness"). The very "naturalness" and adaptiveness of animal-exploitation, however, serves as a reason for us to trust our moral intuitions and their verbal rationalisations less, not more. For the wells of rationality have been poisoned from the outset. Our capacity for fair judgement is biochemically corruptible and genetically corrupted. Other things being equal, genes promoting a capacity for self-serving rationalisation will tend to get differentially favoured over those promoting impartial detachment. The literally self-centred nature of our individual virtual worlds - for we each live in a self-assembled neuronal VR world grotesquely focused on one egocentric body-image - attests to the technically defined *selfish* character of DNA-driven consciousness. In consequence of this inbuilt distortion, the 'reflective equilibrium' sought after by fans of ethical common-sense neglects the systematic genetic biases coded into the mechanisms by which our intuitions are formed. Such biases leave our intuitions, and the consequences we extract from them, even less dependable than intuitive folk-physics. Ethically, we simply can't be trusted; or trust ourselves.

For if several hundred million *human* toddlers or babies were abused and killed each year - for food, fun, or scientific curiosity - then the compelling moral urgency of the animal issue would be undeniable. We'd find it hard to dispute the moral crisis - unless habit had made us so wholly desensitised to what we were doing that the mass-slaughter of human youngsters, too, had become "natural". In fact, our intermittent moral anguish over the surgical abortion of embryos/foetuses/unborn human children shows we are not always blind to the interests of the weak and defenceless; and our victims within the womb are neurologically and psychologically far less developed than the victims of our last meal. Perhaps the best hope of a revolutionary change in human attitudes to the victims of our ongoing animal holocaust is a dawning recognition on the part of many millions of people. This is that our current ethical stance to non-humans isn't just morally wrong but intellectually incoherent.

So much for the rationale for this sort of book - and this review.

David DeGrazia's treatise is an uneasy but impressive mixture of ethics, meta-ethics and scientifically-informed analytic philosophy of mind. It is a work of scholarship in the best sense. Not once, in spite of his obvious intensity of feeling and sense of the moral urgency of the issues, did I notice him slipping into overheated rhetoric or polemics.

Actually, the issue isn't that simple. Fastidious restraint in one's language is sometimes a mixed blessing, even for the purposes of intellectual comprehension rather than advocacy. This is because moral apathy, the widespread sense that one's victims *don't* need taking seriously, is always easy if one doesn't really grasp the nature of what one is talking about. Clinically

descriptive text is only more faithful to reality than its value-laden emotive counterpart if coolness of prose more accurately conveys to the reader what is really being described. And generally it doesn't. Vivisection experiments in medical journals, for example, are written up with practised, peer-sanctioned deceit. Academic philosophical treatments of animal-abuse are less Orwellian. Nonetheless, they commonly retreat into the abstruse in-house theory of rights and duties. Soon they clog up with dense layers of abstraction. This may be unavoidable; but the trouble with academic formality of language is that its remoteness from the raw immediacies of suffering *hides* how bad that suffering really is; and the desperate moral urgency of doing something to stop it.

Our own semantic competence, then, shouldn't be taken for granted. Don't assume that you straightforwardly know what you're thinking and talking about if you assume that the suffering of various categories of other beings doesn't matter. By way of analogy, we'd recognise that someone who has seen only black-and white picture postcards of, say, Van Gogh's *Sunflowers*, hasn't really grasped the nature of Van Gogh's painting. The difference between a small grey postcard and the original masterpiece is so vast that one couldn't *trust* the artistic judgement of someone who has only experienced the former to pass judgement on the latter. (S)He wouldn't know what he was saying. Yet we're far more ready to grant that someone can grasp the content of - and thus potentially pass *moral* judgement on - what is meant by, for instance, 'factory farming', 'slaughterhouse methods', or 'veal crates', even though they've merely read a piece of text (nominally) about it. The morally dangerous presumption of semantic competence is widespread and implicit even though the words of abuse themselves evoke only an inadequate mild unease or distaste. Such rarefied sentiments cannot possibly capture or evoke the felt *horror* of what takes place from the perspective of the victim.

For perspectival facts and subjective "raw feels" are an objective feature of Reality; even though we don't scientifically understand why they exist. Without them, nothing could *matter* - whether to itself or to anything else. If *we* the abusers could apprehend the horrors we perpetrate on the abused as fellow subjects rather than ill-conceived objects, then we *couldn't* be so complacent about what we're doing. But the world isn't like that. Worse, the victim's viewpoint isn't a perspective with which most of us even try to empathise - not even for a few seconds. Who cares? Get a life! Alas, the culture of abuse is just too pervasive.

To a large extent, we are in any case deliberately shielded from what we're paying for. Our willing complicity - and sometimes wilful failure of the imagination - doubtless contributes to the still prevalent sense that what we're doing to other life-forms doesn't in truth matter all that much. So it's worth quoting - however unrepresentative they are of DeGrazia's book as a whole - from the only two pages in *Taking Animals Seriously* which really begin to hint at what happens in contemporary animal husbandry.

Since World War Two, traditional family farms have largely gone out of business. They have been superseded by what's blandly known as factory-farming. Factory-farms seek to raise as many animals as possible in the smallest possible space in order to maximise profits. The single-minded pursuit of profit has the corollary that animals are nothing but meat-producing objects. They have been overwhelmingly treated as such. Here is DeGrazia talking about the fate of the 100 million mammals and 5 billion birds slaughtered annually in the USA alone:

"After hatching broiler chickens *are moved to enclosed sheds containing automatic feeders and waterers. From 10 000 to 75 000 birds are kept in a single shed, which becomes increasingly crowded as they grow at an abnormally fast rate. Crowding often leads to cannibalism and other aggressive behaviors; another occurrence is panic-driven piling on top of each other, sometimes causing suffocation. Concerns about the possibility of aggression have led many farmers to debeak their chickens, apparently through sensitive tissue. By slaughter time, chickens have as little as six tenths of a square-foot apiece. There is typically little*

ventilation, and the never-cleaned droppings produce an air thick with ammonia, dust and bacteria."

*"*Laying hens *live their lives in "battery" cages made entirely of wire. Cages are so crowded that hens can seldom fully stretch their wings; debeaking is common practice to limit the damage of the hens' pecking cagemates. For hours before laying an egg, a hen, deprived of any nest, paces anxiously amid the mob; at egg laying time, she must stand on a sloped, uncomfortable wire floor that precludes the instinctual behaviors of scratching, dust bathing, and pecking for food. Unnatural conditions, lack of normal exercise and demands for high egg production cause bone weakness. Some hens undergo forced molting, stimulated by up to twelve days without food. When considered spent, hens are stuffed into crates and transported in uncovered trucks for slaughter; during handling and transport, many (over two thirds in one study) incur broken bones. Laying hens and broiler chickens have the same fate; They are shackled upside down, fully conscious, on conveyor belts before their throats are cut by an automated knife. (Hens' brothers have short lives due to their commercial uselessness. After hatching, they are dumped into plastic sacks and left to suffocate, or ground up while still alive to make feed for their sisters.)"*

*"*Hogs*, a highly intelligent and social species, have virtually nothing to do in factory farms except stand up, lie down, eat and sleep. Usually deprived of straw and other sources of amusement, and separated from each other by iron bars in small crates, hogs appear to suffer greatly from boredom. Sometimes they amuse themselves by biting a tail in the next crate. Industry's increasingly common response is to cut off their tails - a procedure that, like castration of males, is usually done without anesthesia. Hogs stand on either wire mesh, slatted floors, or concrete floors - all highly unnatural footings. Poor ventilation and accumulated waste products cause powerful fumes. Hogs are often abused at the loading and unloading stage of transport, particularly at the slaughterhouse. Rough handling sometimes includes the use of whips and electrical 'hot shots'."*

*"*Veal calves *are probably worse off than other farm animals. Shortly after birth, they are taken from their mothers and transported considerable distances - often with rough handling, exposure to the elements, and no food or rest. At the veal barn, they are confined in solitary crates too small to allow them to turn round or even sleep in a natural position. Denied solid food and water, they are given a liquid milk replacer deficient in iron (in order to produce the gourmet white flesh), resulting in anemia. Because it is drunk from buckets, rather than suckled, the liquid food often enters the rumen rather than the true stomach, causing diarrhea and indigestion. The combination of deprivations sometimes results in such neurotic behaviors as sucking the boards of crates and stereotyped tongue-rolling."*

"Like their veal-calf siblings, many dairy cows, as calves, never receive colostrum - the milk produced by their mothers which helps to fight diseases. More and more they are confined either indoors or in overcrowded drylots (which have no grass). Unanesthetised tail docking is increasingly performed. In order to produce some twenty times the amount of milk a calf would need, dairy cows are fed a diet heavy in grain - as distinct from the roughages for which their digestive tracts are suited - creating health problems that include painful lameness and metabolic disorders, which are exacerbated by confinement. About half U.S. dairy cows at any one time have mastitis, a painful udder. Many cows today are given daily injections of Bovine Growth Hormone to stimulate additional growth and increase milk production (despite a surplus of dairy products). Although their natural life span is about twenty to twenty-five years, at about age four, dairy cows become unable to maintain production levels and are transported for slaughter. Most processed beef comes from them."

*"*Cattle raised specifically for beef *are, on the whole, better off than the other farm animals already described. Many of the cattle get to roam in the outdoors for about six months. Then they are transported long distances to feedlots, where they are fattened up on grain rather*

than grass. Craving roughage, the cattle often lick their own and other cattle's coats; the hair that enters the rumen sometimes causes abscesses. Most feedlots do not confine intensively. Their major sources of distress are the boredom likely to result from a barren environment, unrelieved exposure to the elements, dehorning (which cuts through arteries and other tissue), branding, the cutting of ears into special shapes for identification purposes, and unanesthetized castration (which involves pinning the animal, cutting his scrotum, and ripping out each testicle)."

"Transporting hogs and cattle for slaughter - which can entail up to three days without food, water, or rest - typically results in conspicuous weight loss and other signs of deprivation. The slaughtering process itself is likely to cause fear. The animals are transported on a conveyor belt or goaded up a ramp in the stench of their fellows' blood. In the best of circumstances, animals are rendered unconscious by a captive-bolt gun or electric shock before their throats are slit."

This horrible suffering occurs, one has to remind oneself, primarily because we enjoy the taste of meat; and because our appetites are financially profitable.

Worlds That Matter

> *"The day may come when the rest of animal creation may acquire those rights which never could have been withheld from them but by the hand of tyranny...a full-grown horse or dog is beyond comparison a more rational, as well as a more conversable animal, than an infant of a day, or a week, or even a month old. But suppose the case were otherwise, what would it avail? The question is not, can they reason? Nor, can they talk? But* can they suffer? *Why should the law refuse its protection to any sensitive being? The time will come when humanity will extend its mantle over everything that breathes..."*
>
> Jeremy Bentham
> PRINCIPLES OF MORALS AND LEGISLATION

Chapter One of *Taking Animals Seriously* contains a concise and extremely useful view of the recent scholarly literature. Perhaps it would be a good idea if future editions included a potted cross-cultural and historical context too. This background would be helpful lest the unwary student suppose that the moral status of animals was the discovery of a far-sighted bunch of Anglophone analytic philosophers twenty-five years ago.

DeGrazia divides recent scholarly output into two generations. This schema is not entirely convincing, but it's still convenient:

- 'First generation' work on animal ethics was written by utilitarians, most notably Peter Singer (*Animal Liberation* 1975 rev. ed. 1995); and animal rights theorists, most notably Tom Regan (*The Case for Animal Rights*; Berkeley: University of California Press; 1983).
- 'Second generation' scholarship, characteristic of authors such as Mary Midgley (*Animals and Why They Matter* Athens, GA: University of Georgia Press, 1983) and S.F.Sapontzis (*Morals, Reasons and Animals* Philadelphia: Temple University Press, 1987) the latter abandon system-building and previous efforts to ground ethics in reason-derived ahistorical norms. Also forming part of this 'second generation' scholarship are Rosemary Rodd's scientifically sophisticated contribution (*Biology, Ethics and Animals*: Oxford, Clarendon 1990); and, rather incongruously, philosopher Peter Carruthers' *The Animals*

Issue. Carruthers advances the thesis that the mental states of animals are all non-conscious.

Taking Animals Seriously itself seeks to transcend the old utility-versus-rights debate. It aims throughout to explore the mental life and moral status of animals in an empirically-informed manner. DeGrazia deploys the non-foundationalist "coherence-based" methodology of ethical justification that he develops in Chapter Two to argue that many kinds of animal do - and many don't - have moral status. He takes great pains to explain what the crucial but disastrously ill-named principle of equal consideration of interests for animals actually means; and, no less relevantly, what it doesn't. DeGrazia spells out (p 37) that "giving as much moral weight to human interests as we give to relevantly similar human interests does *not* entail:

- identical rights for humans and animals
- a moral requirement to treat humans equally
- the absence of any morally interesting differences between animals and humans "

Drawing on a wide range of ethological research, DeGrazia sets out the principled grounds on which *morally* relevant similarities and differences can be identified in potential bearers of moral status. He argues, convincingly, that a very diverse range of animals have feelings, desires and beliefs. Intriguing and disconcerting evidence is presented that a whole repertoire of mental properties, and even language, are not, as many non-Darwinian-minded philosophers have claimed - all-or-none properties peculiar to humans.

DeGrazia also offers a good discussion of the contemporary academic literature in animal physiology and ethology. With plenty of complications and some exceptions, modern research suggests that the distinction between vertebrates and non-vertebrates - by itself, under such a description, an ethically trivial distinction - may in fact serve as a rough-and-ready marker for much more profound and morally important differences altogether.

Admittedly, *all* non-humans might have been *radically* unlike humans and still commanded moral status. Or rather the issue would still need to be argued, not just presupposed. After all, one day insentient silicon robotic isomorphs of organic life may have *ostensible* functional analogues not just to pain, but to morals, meta-ethics, and even to traditional religious casuistry; though perhaps silicon theologians overtax our imagination. So it still needs to be spelt out why insentient objects, artefacts or life-forms - whether amenable to functional description or otherwise - *don't* merit genuine moral status; and why they have don't have any *interests* which need to be taken into account. [This sweeping statement disguises a contentious assumption: that only organic systems have unitary experiential manifolds as distinct from discrete and fleeting specks of consciousness. We are such manifolds because only organic minds have a functional architecture based on the extraordinary and unique valence properties of the carbon atom. Carbon's functionally unique attributes are indispensable to the formation of the "warm" quantum coherent states which hypothetically mediate unitary fields of experience. This QM-invoking solution to Sellars' notorious "grain problem" of consciousness is discussed in my review of Chalmers. It's worth noting that organic functionalist arguments for moral carbon chauvinism are scarcely received wisdom; and must rank as speculative]

As it happens, however, the neo-Darwinian synthesis confirms the fact that human and non-human vertebrates are similar where not type-identical in the category that matters most. This is the category that grounds, and gives rise to, our very notion of *mattering* in the first instance - the pleasure-pain axis. A universe without any kind of feelings in its ontology would be a universe in which nothing *mattered* or had any *importance*; and in the realm of phenomenology, appearance and reality are one-and-the-same. Things which are felt intensely matter more. If your pain - or an animal's or extra-terrestrial's pain - doesn't matter to me, that

it is because I have failed to apprehend it. For it is a different, anaemic experience or spuriously objectified third-person fact which I have in mind instead.

If we *could* apprehend the real first-person agonies of a member of another species, or even acknowledge that such agonies are part of the real ontology of the world, then we might be less callous in our treatment of non-humans. Unfortunately - doting pet-owners apart - we find cross-species empathy very hard; and for the sake of our victims, if not always perhaps good ethological method, it might be better if we actually "anthropomorphised" more, not less. Although not logically sound, the best way to promote the desperately needed revolution in our treatment of other life-forms may well be to convince people that *in the relevant respects* non-humans are just like "us" - or possibly reshape our notions of just who *we* are. This mode of persuasion is more likely to be effective simply because it consists in forcing us to think through the full implications of what we already believe. It doesn't ask us to revise our basic values and presuppositions. This would be a far harder task altogether. Precepts such as "act so as to minimise needless suffering" are, for sure, infuriatingly imprecise. Yet their unexceptionable woolliness helps to command assent and lays out a minimum of common ground needed to take the argument forward.

Ethical utilitarians explicitly focus on our shared capacity for pain and pleasure: the sovereign nice-nasty axis construed in the broadest sense. Unfortunately, the DNA-driven "encephalisation of emotion" makes many of us grant greater moral weight to a high nominal IQ than emotional well-being. IQ is an ill-defined and ideologically-disputed notion bound up with the vaunted human capacity to churn out logical inferences. As traditionally conceived, the notion ignores our shared capacity for *feeling* and "emotional IQ" altogether. And it is the quality and intensity of feeling which determines whether - and how much - those logical inferences, or anything else, actually matters to anyone at all. To some extent, I fear, DeGrazia's otherwise admirable and extensive discussion of animal cognition encourages this tendency to focus on intentional objects [this further "essentially contested" term is philosophy-speak for what we think etc "about"] rather than why these objects do - or don't - have any significance. This preoccupation with the "cognitive" is both a shame and a danger; for to focus on animal minds interpreted in a narrow intellectualist sense is to focus on an area where in many respects they are demonstrably inferior to least most mature humans; whereas in the case of their ethically crucial capacity to suffer, the evidence is at best mixed.

So next, granted some minimal principle of the uniformity of nature, it's worth briefly exploring the biochemical substrates of two particularly distressing modes of aversive experience. Where are they found, and where are they absent, within the phylogenetic tree? How should their absence or prevalence lead us to re-examine our traditional ideas of the moral status of members of other species; and, crucially, to the way that we behave towards them?

Fear and Anxiety

> *"It is not obvious what is obvious"*
>
> Daniel Dennett

With a few exceptions, nearly all the anxiety-mediating agents (e.g. the beta-carbolines) found to date have as their site of action the benzodiazepine receptors (*TAS*; p 121). Beta-carboline ligands which bind to the benzodiazepine receptors induce in humans

"...intense inner strain and excitation, increased blood pressure and pulse, restlessness, increased cortisol and catecholamine release, and stereotyped rocking motions. The administration of anxiety-producing beta-carbolines to primates caused piloerection (hair

raising) and struggling in the restraint chair, increased blood pressure and pulse, increased cortisol and catecholamine release, and increased vocalization and urination." (TAS p121)

Again with a handful of exceptions, the *anti*-anxiety properties of alcohol, the barbiturates and the benzodiazepines (the 'minor tranquillisers': Valium etc) can be tied to a large, single, multifunctional receptor complex. This single neurochemical substrate includes a barbiturate- and ethanol-binding site, a chloride ion channel, and a binding site for neurotransmission. It has been shown that there are high-affinity saturable and specific receptors for the benzodiazepines in the vertebrate central nervous system. Following the landmark study of Nielsen, Braestrup and Squires (*Evidence for a late Evolutionary Appearance of a Brain Specific Benzodiazepine Receptor*, Brain Research 141 (1978) 342-466), persuasive evidence has accumulated that all vertebrates - including the bony fishes - have these receptors. Such receptors were found to be absent in all the invertebrate species tested (originally the woodlouse, earthworm, locust, lobster and squid); and also from the cartilaginous fishes.

Inevitably, the full story is messier. As DeGrazia notes, the discovery of peripheral benzodiazepine-receptors with a presumably non-anxiety role [and also the development of 5HT1a mixed agonists such as buspirone with anti-anxiety properties], means the intricacies of the evolutionary story are vastly more complicated than any lightning sketch can show. Yet overall, there is strong evidence that all vertebrates, and some invertebrates, suffer anxiety and fear.

Pain

> *"Every particle of factual evidence supports the factual contention that the higher mammalian vertebrates experience pain sensations at least as acute as our own. To say that they feel pain less because they are lower animals is an absurdity; it can easily be shown that many of their senses are far more acute than ours - visual acuity in certain birds, hearing in most wild animals, and touch in others; these animals depend more than we do today on the sharpest possible awareness of a hostile environment. Apart from the complexity of the cerebral cortex (which does not directly perceive pain) their nervous systems are almost identical to ours and their reaction to pain remarkably similar, though lacking (so far as we know) the philosophical and moral overtones. The emotional element is all too evident, mainly in the form of fear and anger."*
>
> Richard Serjeant

Pain is, if anything, even more basic in a phylogenetic sense than fear or anxiety. The opiates which (at least) functionally mediate analgesia have been found in earthworms. It's disturbing to realise that the most "primitive" experience one can undergo, at least if one's own life is anything to go by, is also the most intense. Abstract, more-or-less serial thought, by contrast, tends to be at best faint, elusive and ethereal in its phenomenal properties.

This contrast needs stressing. Consciousness is sometimes claimed to be the prerogative of the higher vertebrates, or even of humans alone in view of our comparatively superior cognitive prowess. Yet - quite incongruously from such an anthropocentric perspective - our most abstruse and distinctive cognitive skills are usually those most minimally penetrable to introspective access; while some of our most primitive feelings are also the most intrusive and subjectively *important*.

DeGrazia carefully distinguishes between our concepts of phenomenological pain and physical nociception. He notes the problems had by people with a congenital inability to feel pain or leprosy. In a fascinating speculation, he argues that (p 111)

"Pain seems to be a development of consciousness in creatures endowed with a highly developed response system known as nociception. Consciousness may have developed as a free-rider on certain inherited gene groups that included relatively complex information processing; or it may have evolved as a way of focusing an organism's attention to those areas of information processing that are most valuable at a given time. Either way, pain was apparently the new conscious companion of responses to potentially harmful situations (in these creatures, nociception) in the animals in which consciousness emerged."

Just as relevantly, DeGrazia sets out how insects lack the extensive CNS processing-mechanisms implicated in pain-perception among vertebrates. The locust, for instance, keeps on eating while being devoured by a mantis. It's hard to imagine a vertebrate animal retaining any semblance of equanimity while meeting such a fate. DeGrazia suggests that whereas the startle-reflex would confer survival advantage similar to acute pain, insects with short life-spans and modest learning needs would derive negligible advantage from it. There would be little or no selection-pressure favouring a neural capacity for any such experience.

This issue is actually more problematic than it sounds. The difference between 'little' or 'no' selection-pressure is huge from an evolutionary perspective. Even a 1% reproductive advantage conferred by a capacity to experience phenomenological pain - *if* it were functionally significant and causally efficacious *qua* phenomenological pain - would allow natural selection to get to work over millions of generations. Nonetheless, in qualified support of DeGrazia, it seems unlikely that organisms without a single *central* nervous system possess a unitary experiential manifold - let alone a unitary sense of self to which moral status could readily be attached. Even if the multiple ganglia of a locust each feel an extremely rudimentary kind of aversive experience - as IMO is quite likely - the mantis-devoured locust's feeding head doesn't participate in it - whereas a "toothache", for instance, seems to penetrate to the very heart of our whole existence. [The encephalisation of pain and emotion is tremendously adaptive; and computationally hard to match by our silicon robots. From an information-theoretic perspective, the saturation of our (neo-cortical) cognitive processes and organic virtual worlds by (mainly limbic) feelings in complex experiential manifolds may offer computational advantages over a "classical" computational architecture]

Fascinating as they are, DeGrazia's speculations on the origins of consciousness face serious difficulties. His account of conscious pain doesn't offer a solution to the "zombie problem" famously highlighted by David Chalmers. If zombie-nociception would do the same functional job that's allegedly performed by phenomenological pain, then it's hard to understand why selection-pressure didn't favour mere nociception. Phenomenological pain, unlike zombie nociception, doesn't logically "supervene" on an [apparently] exhaustive specification of the microphysical facts. For it to do so, those putative microphysical facts would have to be heroically reconstrued; and a primitive "what-it's-likeness" posited as the stuff which the quantum mechanical formalism describes instead: naturalistic panpsychism. Moreover, even if phenomenological pain really were functionally advantageous to genetic vehicles *in virtue of* its horrific subjective texture, the evolutionary story wouldn't have explained, in any deep sense, why and how that uniquely awful texture of nastiness occurs. The story would simply explain why it was differentially selected over phenomenological states. In sum, we still don't understand why the laws of physics didn't generate a world the constituents, configuration and behaviour of which - and the neurophysiology of its organisms - was type-identical to our own, but where consciousness was absent. [One proposed solution is found in Cosmic Consciousness For Tough Minds]

Such philosophical argument over the (non-)existence of consciousness wouldn't matter ethically if it weren't for an insidious muddle over the two radically different senses of "objective". The confusion allows the third-person ontology favoured by orthodox natural science and its *lumpen*-academic cheerleaders to get privileged over the first-person perspective. For in reality, what it's subjectively like to be a desperately distressed veal calf, for example, is an objective fact about the world. The fact that such horror may also be notionally captured by a set of observer-independent equations is morally irrelevant. The facts about subjective mind-dependent states are objectively true. Physics gives us a formal mathematical description of the world. It says nothing about the insentience or otherwise of what "breathes fire in the equations and makes there a world for us to describe."

Speculative metaphysics aside, DeGrazia's profound conclusion is that:

"...affective beings (who have feelings), conative beings (who have desires) and cognitive beings seem to be co-extensive on our planet with the vertebrates, give or take a few species..."

such as the cephalopods. All vertebrates are endowed with the limbic and autonomic systems which contain the basic biological substrates of pain, anxiety and fear.

It should be stressed that this conclusion doesn't, as it stands, mean that morally speaking we can do anything we like to invertebrates. If DeGrazia is broadly correct in his dichotomy, then a (quasi-)Kantian indirect duty view - the idea that the only reason we should avoid cruelty to animals is that such practices corrupt the character of agents and make them more likely to behave badly toward humans - might still be adapted and enlarged so that "we" is taken more broadly than it does now. Working within this sort of framework, the frivolous killing of invertebrates, such as stamping on a fly for the sake of it or through mere irritation, might still be discouraged. It should be deplored on the grounds that the attitude of mind underlying such actions promotes cruelty to morally important vertebrates too. Yet the conclusion that - simplistically - vertebrates are special enables us non-arbitrarily to avoid treating a fly or a worm with the same consideration we should accord a fellow vertebrate. It's a dreadfully crude division; but it's a very useful start.

DeGrazia's account is still problematic in other ways. Even granted his vastly more generous conception of mentality than hominid chauvinists, *Taking Animals Seriously* is too ready, I think, to link - without further argument - moral status to intelligence and complexity. It would be better instead if such attributes were treated as markers for the property that generates and defines mattering itself. This involves the capacity to suffer, or rather the capacity to undergo experience imbued with significance and located on a broadly-defined pleasure-pain axis.

Again, there are a lot of complications to research into the biological basis of mattering. Even utilitarians, who stress the moral primacy of the pleasure-pain axis, are liable to assume that degrees of sentience are somehow inevitably bound up with intelligence and the ability to process information. Our lack of introspective access to the workings of the distinctively human language modules ought to alert us to the pitfalls of intellectualism here. No substantive argument is presented in DeGrazia or elsewhere for believing that our unusual adaptation of an extraordinarily hypertrophied intellect has been accompanied by a matching hypertrophied capacity for suffering relative to other less cognitively sophisticated vertebrate species. The reason is that no supporting evidence exists for such a notion.

Non-humans demonstrably possess greater acuity in many of the "special senses", notably olfaction, hearing and vision. What grounds have we for supposing that no such heightened sensitivity to *pain* isn't found elsewhere in the animal kingdom? One must hope that it isn't; pain is vile enough to "one of us" as it is. We simply don't know enough about the pain-centers

of a whale or an elephant, for instance, to establish whether approximate equality of biological propensity to anguish is really the case. Greater encephalisation of emotion most likely *does* extend the nominal range and nuances of things one can be unhappy 'about'; though in the case of vertebrates with acute special senses, it may well be humans who are comparatively obtuse in our lack of discriminative power, perhaps fortunately so. Yet it's not clear that encephalisation by itself can intensify aversive experience in the absence of limbic structures to mediate any such additional nastiness. The assumed role of intelligence is a link that too many accounts of possible candidates for moral status presuppose.

In any event, if suffering really is the selfish-DNA-driven, out-of-control evolutionary nightmare that the evidence suggests, with no higher purpose to dignify it, then there's no indication of any mechanism by which it could ever be checked simply because it felt unspeakably bad. Perhaps the most that can be hoped is that the substrates of a pain so all-consumingly bad that it sapped the capacity for thought and (genetically) adaptive behavioral responses would - other things being equal - get selected against. Less optimistically, it is generally assumed - ignoring the philosophers' zombie problem - that pain's adaptive motivating force is in some degree proportionate to its intensity. The worse the pain, the greater the incentive to escape it. This perspective has grimmer and more sinister implications altogether.

So just how bad can pain be? In view of the great weight here being placed on the parallel between small children and non-human animals, it's worth asking if children suffer as adults and to the same degree. At least when cortical myelination is complete, then (once again, given certain assumptions) young children may well suffer as intensely as adults. Indeed, it's not perverse to raise the possibility that youngsters sometimes suffer more. This might sound implausible. Yet on the crudest level, children literally have more (irreplaceable) brain cells of the kind that mediate emotional experience, albeit with a different dendritic arborisation etc. Over the years, 'neural Darwinism' (an admittedly somewhat misleading term) also acts to winnow out many dysfunctional and non-functional inter-neuronal connections. This may enhance a person's intellectual performance but diminish the raw amount and intensity of consciousness. It all depends on what gets winnowed where. Moreover, efficient brains use less energy and do things more "automatically" - and less consciously. Further, as one ages, the mind/brain progressively loses nerve cells - even though their loss may elicit a compensatory sprouting to repair any functional deficits, and even though physical cellular shrinkage rather than cell-death may account for much well-attested cerebral weight-loss. Certainly, many older adults report that they feel things less intensely than they did in their callow but emotionally tempestuous youth.

The evidence of a direct causal connection between intellectual prowess and intensity of feeling, then, is still to be found; and perhaps never will. Furthermore, as pain gets worse, one's capacity for abstract thought, and capacity to exhibit one's vaunted intelligence however it's defined, diminishes. The suffering one undergoes doesn't thereby matter less. On the contrary, it can become all that matters. I do not know what it is normally like to be a whale or a pig. But I suspect that *in extremis* it is very similar to what it is like for me to be in terror or extreme pain: simply horrific. The same type of post-synaptic metabolic cascades get triggered. And *contra* Wittgenstein, if a lion could talk, we might understand it rather well: for we have in common a core biological repertoire of raw appetites and emotions, not to mention genes, metabolic pathways and brain structures to match.

DeGrazia probably wouldn't go this far. In his discussion of beliefs, desires and language, he concentrates once again on grading their relative sophistication, connectivity and systematicity rather than the felt texture of the sentiments/limbic processes that infuse their individual episodes of cortical activation. Yet why should creatures whose adults are more intelligent inherently matter more? If intelligence could be used as a marker for the intensity of emotion and the biochemical creation of significance, then IQ might at least serve as a useful yardstick for something that inherently mattered. If it can't be so used, then one might

as well argue that a *Pentium Pro* is morally superior to a *Intel 386*. The right answer is surely that processing power and moral status are simply incommensurable categories. For the most phylogenetically primitive sorts of consciousness - most intrusively pain - appear to be the most intense; whereas the most recent kinds of consciousness in evolutionary terms, notably those implicated in linguistic processing, have a subtle and introspectively opaque texture so elusive that certain philosophers have even doubted its existence.

Possibly, even more rarefied modes of consciousness are feasible. If so, it is unclear why hyper-intelligent transhuman beings who might undergo such novel cognitive processes should matter more than we do. This is unless post-humans feel things more deeply for other neurological reasons altogether; for they may have designed themselves hyper-emotional psychochemical states too.

Even here, we must be careful with our terminology. The term "intelligence" itself is too riddled with covert value-judgements about what does and doesn't rank as even *cognitively* important to be very useful. Its shifting usage reflects shifting power-relationships; not the carving of Nature at the conceptual joints. Yet *if* some value-neutral sense of intelligence is salvaged, and *if* the argument that relative IQ is morally relevant is taken seriously, then we would also have to accept that ultra-smart Mensa masterminds matter more in ethical terms than less intellectually agile members of our own species. It's not clear why this should really be the case. Perhaps high-powered intellects might still potentially matter more, in a merely instrumental sense, if they were more creative of socially useful inventions - though such comparisons are usually invidious and probably best avoided. Moreover, to add another complication, acknowledged genius does seem to have some kind of limited positive correlation with a tendency to manic-depression. This tendency might be morally relevant because people with "bipolar disorder" do tend to feel things more intensely, and its soft-bipolar forms are linked to unusually high creativity. So if one is trying to press the issue, then I suppose one could even make some sort of case that manic-depressives do *inherently* matter more because things matter more to them - for only in the naïve third-person ontology of scientism do things that matter have to be observer-independent any more than tickles have to be observer-independent. Yet the argument gets pretty tortuous.

There is a complex morass of issues here. It's not worth getting bogged down in them. This is because, as will be seen in section five, the imminence of the post-Darwinian Transition will ensure that traditional ethical dilemmas get swept away into evolutionary history. Traditional casuistry, and moral league tables in the Great Chain of Being, are likely to become obsolete. For as has been remarked, an angel in heaven is no one in particular.

Losing Our Minds

> *"If the experimenter would not be prepared to use a human infant, then his readiness to use non-human animals reveals an unjustified form of discrimination on the basis of species, since adult apes, monkeys, dogs, cats, rats and other mammals are more aware of what is happening to them, more self-directing and, so far as we can tell, at least as sensitive to pain as a human infant."*
>
> Peter Singer
> ANIMAL LIBERATION

What are the ingredients of conscious mentality? Do unconscious minds matter? Why is consciousness ethically important?

In common with many philosophical treatments of the nature of mind, the neglect of a well worked out theory of perception leads DeGrazia to omit the greater part of organic life's

mental furniture altogether. For the default option of classical perceptual realism doesn't just amount to a false theory of the world. It also leads to an impoverished conception of mind. By way of contrast - and in defiance of our ingrained direct realist intuitions - if the existence of the mind-independent world can only ever be *inferred* from our individual virtual world models, and not directly apprehended or "perceived", then the properties which one normally ascribes to classical macroscopic objects are inherently mental and covertly autobiographical. This is so even when one is "awake". The "awakened" condition is a mysterious mode of consciousness in which minds/brains are popularly supposed to attain a state of direct self-transcendence that they lack while dreamfully asleep.

More soberly, "out there" and "in here" don't mean what they seem. Natural selection has indeed harnessed fields of neurally organised forms of consciousness and worked them into simulating something else. It has modelled the fields of macroscopic, "medium-sized dry objects" whose patterns occupy quasi-classical branches of the quantum mechanical Multiverse; and whose (partly) causally covarying simulations stretch beyond the somato-sensory homunculus into the wider reaches of the neocortex as a whole. Yet each simulation is still ineradicably mental and indexical to its creator. So it's still an adaptive con-job. If classical realism is indeed false, and quantum mechanics is the sovereign theory of the whole cosmos, then mental life is presumably biologically ancient; even pre-Cambrian. This is on the assumption that there's an unreflective "perceptual" mentality which runs (at least) a long way "down" the phylogenetic tree. Most notably, in creatures with central nervous systems, this mentality takes the guise of species-specific virtual worlds. These are vivid "experiential manifolds" centred around a somato-sensory body-image.

Naturally, the billions of noisy, colourful, refractory, virtual worlds churned out by evolution through self-replicating DNA are not recognised and categorised by their host vehicles as [more-or-less] functionally organised modes of consciousness. For the most part, recognition of their mind-dependence would be functionally irrelevant to the organism. The cognitive skills which recognition entails demand a form of meta-representational capacity which would be neither energy-efficient nor cost-effective in genetic terms. For the functional role of most aspects of animal (and human) mind is to do duty for the *non-mental* as efficiently as possible. Natural selection, blindly as ever, spawns the machinery for generating dynamic simulations of the local environment. Most intimately, simulating this local environment involves running toy egocentric models of throwaway DNA vehicles (aka living organisms) themselves. Neurally-active genes code (in vertebrates) for egocentric "somato-sensory" simulations of the host organism's body. And if any intelligent organism is ever tempted to wonder why the whole world seems to be centred on itself, then it's worth asking where else, if anywhere, in a world-model would it be more advantageous for purely selfish genes to travesty one's significance in the great scheme of things.

The virtual world of a gazelle or a chimpanzee may not be as complex as that of a mature adult human. Yet it is still subjectively immense and vastly complex. Awake or dreaming, the greater part of each experiential manifold in which such simulations consist will often seem harshly or indifferently mind-independent. Each virtual world dwarfs the egocentric body-image at its centre. This sort of self-alienation is hugely adaptive for the host vehicle whose mind/neural network is running the simulations. It doesn't make the modelling process anything less than a genetically predisposed charade. Perhaps it will take the routine accessibility and long-term habitability of "artificial" virtual worlds - immersive, multi-modal and generated by (non-organic) VR-optimised computers - for the nature of this hard-wired perceptual realist illusion to sink in. Or perhaps, for many of us, its illusoriness never will; inferential realism is a philosophical position to be contemplated, not a form of life to be lived.

DeGrazia discusses animal minds minus their virtual worlds at some length and considerable depth. With his "coherence-based" approach to ethics, he appraises the moral status of animals, not in terms of the straightforward foundations of a pleasure-pain calculus,

but in the light of the surprising richness and unsuspected diversity of mental life as defined even within the confines of an orthodox perceptual realist conception of mentality. Again, though fascinating and informative, DeGrazia's discussion is - I think - fundamentally skewed. This is because of his implicit reliance on an untenable realist theory of perception and his consequent diminished conception of the realm of consciousness. The more mental properties which get palmed off onto the rest of the world, the less there is to occupy one's mind - or non-human animal minds either.

Classical realism, however, is hopelessly at odds with what quantum mechanics, quite aside from *a priori* philosophical argument, tells us about the mind-transcendent world. Hilbert space and common-sense folk-physics are inconsistent; and ultimately it's common sense which has to be sacrificed if any unitary world picture is to be salvaged from the epistemological wreckage. Our quasi-hard-wired theatre of classical macroscopic objects leads most of its subjects to locate the contents of their visual fields in the spuriously accessible Outer World rather than in their world-*selected* organic minds. It is this mistaken theory of perception which gives rise to the intractable mind-body conundrum in the first instance.

At least we recognise there's a problem here. Philosophers in the grip of the ancient World-Knot ask: How can a cheesy grey mass of warm porridge - exquisite functional organisation notwithstanding - possibly give rise to consciousness, to *this* thought or *that* sensation? And the simple answer is that it can't. Classical brains - as apparently disclosed by the deliverances of naive realist perception augmented by hi-tech microscopy - are a mind-dependent artefact of particular sorts of QM-coherent quantum minds. Acknowledging this dependence isn't a disguised plea for Idealism or Scepticism. It's just a call for sophisticated realism-by-inference-to-the-best-explanation. [Philosopher Bryan Magee describes his thought in the school chapel that, by blinking, he effaced his entire perceptual world, as an "indescribably awful" realisation. He then succumbs to transcendental idealism. But surely it's better to treat the inferred mind-independent environment which presumably gives rise to each microcosm as a theoretical posit in good standing. It's a hypothesis with a whole lot of explanatory and predictive power.] Within the hugely more vast Multiverse, tiny mind-dependent classical worlds of "medium-sized objects" are themselves a highly adaptive facet of primordial-DNA-driven psychology. They get neurally activated in living organisms. In common with a schizophrenic's voices, a classical world as a whole is apprehended by its host's pre-frontal cortical module(s) as "out there" and not inside one's [somato-sensory-cortical] head. Unlike schizophrenic voices, however, these shifting classical constructs tend causally to co-vary pretty closely with certain gross macro-patterns in local regions of the the Multiverse as a whole. This is why classical mind-worlds elsewhere - *not* a chimerical Classical World - have flourished.

There are natural mechanisms, not daily miracles, at work here. Classical mental macro-worlds are predisposed - but are not, strictly, genetically pre-programmed - to self-assemble if their constituent neurons get their weights and connections trained by appropriate sequences of stimuli from surface transducers. Formally, virtual mind-worlds can be described with the mathematical tools of artificial neural nets. Their behaviour mimics the implementation of powerful learning algorithms. Organic virtual worlds aren't classically programmed; their nets get "trained up" by peripheral input. And naturally they have a different ontology from that which the cognitive modules they interface with normally suppose; for we can't hop outside our models.

Such loose talk will appal anyone with a horror of "Cartesian materialist" homunculi. Are they all watching an infinite regress of mini-TV screens? There certainly *is* a real mystery here. But anyone who doubts the existence of little men in the head [irrespective of their theory of perception] should try a spell of lucid dreaming. They sure ain't anywhere else.

Virtual world prototypes stretch a long way back into the evolutionary past. Certainly, they extend far further than fully-fledged *second*-order representations, such as beliefs and desires, which function as though they were "about" the worlds on which they focus. These occurrent beliefs and desires, however, tend to serve as our foremost exemplars of mental states. It's these relatively late arrivals which act as simulated vehicles for "propositional content". Most mental life involves more mundane features than anything so exotic; although sometimes its features can be gruesome. When someone wantonly kills a mouse, for instance, the killer extinguishes an entire virtual world too, albeit a murine macro-world rather than its humanoid counterpart.

One of the reasons we have an impoverished conception of the mental life of animals, then, is that our individual egocentric visual worlds aren't construed as part of our *mental* life at all in everyday existence. The non-consensual and more environmentally-autonomous virtual worlds of the schizophrenic or well-frazzled acid-head are the exception, not the rule here. We readily grant that the voices etc in a schizophrenic's "external" environment are autobiographical features of the troubled individual's mental life. But the consensus-hallucinations of more typical types of virtual world enjoy no such recognised status; because much of us is constituted by these very hallucinations. If they were to let slip their true colours, the inference to their mind-dependence might interfere with their functional role in the informational economy of the organism. "Indescribably awful" realisations are unwholesome.

Perhaps the covertly mental status of what are often thought of as paradigmatically "physical" properties is most readily disclosed in lucid dreaming. Within a given dream, the huge rock-face, say, which one is climbing is the virtual rock-face of a virtual mountain. One measures kilometres; one is dealing in cubic centimeters. How the metrics of phenomenal spaces can arise embedded in crumpled-up neural minds defies our present understanding. But the mentality of putative "physical" properties is just as real if, to take a more savage example, one is all too awake on the African savannah and being mauled by a lion while a corresponding virtual lion is tearing bits off one's body-image. For if one is awake and being chased by a virtual lion [or a QM superposition of virtual lions etc], then its very probable causal covariation with [a QM superposition of] a real-world hungry predator(s) means it is highly adaptive to treat one's simulation as a mind-independent reality. Think about the difference for a second or so; and you're lunch.

A cruel but striking experimental procedure is instructive here. Vivisectors sometimes surgically abolish an animal's capacity for muscular atony. This state of effective paralysis normally stops our bodies acting out our dreams. Permanent surgical ablation of the region responsible for the functional decoupling of the bodily musculature from its neural command centers during dreams, on the other hand, ensures that the dreaming organism unwittingly enacts its inner psychodramas as it sleeps; just as, controversially, we all do quasi-veridically when awake. Thus the cat doesn't just have simple beliefs and desires about the virtual mouse it chases. Its virtual body-image image chases the virtual mouse within the vast virtual spaces of a feline dreamworld. Its beliefs have both 'narrow' mental content and narrow so-called 'perceptual' content too. Sweeping aside lots of complications, they are beliefs about first-order representations expressed under another description; though since the cat's fleeting murine representations are neither "transparent" or "projectible", post-classical AI sometimes drops the "representational" tag altogether. Common-sense distinguishes, within each experiential manifold/virtual world, between experience and the object of experience. This supposed object of experience is something non-experiential to which we fancy we've got shared direct access and with which we are mysteriously 'presented'. In fact, our solid, refractory chairs and tables, sticks and stones - and squishy grey brains - are themselves distinctive modes of experience. They can be neuronally fired up by electrodes, psychedelic drugs, dreams; or selected from the psychoneural menu while one is awake by peripheral input

from the mind-independent environment. Yet that environment is only *inferred* as the best possible explanation of the experiential evidence.

Hence a cat does not have simple beliefs and desires about a mind-independent mouse. Instead, a crude feline mouse-simulation is taking place; together with a relatively undeveloped non-verbal system of second-order representations. First and second-order representations interact and partially interpenetrate as each simulation dynamically evolves in feline neural nets. When the cat is awake, key features of its world-simulation tend to causally covary with a vastly more complicated creature - the living mouse. Yet mystical feats of feline self-transcendence are no more feasible than human clairvoyance. It can be known from the evanescence of dreams that Nature can conjure up and destroy whole CNS immensities of virtual worlds in milliseconds. These feats of destruction and creative world-making happen whenever we are awake and blink. Animal minds are no less gappy but equally real.

So what happens to the virtual world of the dreaming cat which chases phantom mice? Does it disappear and get replaced by the *real* world when the cat wakes up? No. But the virtual world now gets tightly sculpted, and its shifting contents neurally selected, by peripheral input. A catworld is a quite simple toy world compared to human virtual worlds. But it is still an intensely sentient mental microcosm in its own right; and, tragically, it is a killer-world with intensely sentient victims.

Why does this matter ethically?

It wouldn't do so at all, if it weren't for the fact that virtual worlds and the extended cortical minds they embody have been "emotionally encephalised" thanks to natural selection. The limbic system insinuates its processes into the furthest reaches of cortical mind. What happens in virtual worlds *inherently* matters because they're shot through with limbic-driven emotional meaning and significance. The encephalisation of emotion has extended not just to the cortical regions playing host to second-order representations typified by the occurrent belief-episodes of folk psychology. Our limbic processes, most notably those of the monoaminergic neurons, also infiltrate each egocentric virtual world and its vast cortical arrays too - in man and mouse alike. This infiltration accounts for the circumstance that neither we nor other animals merely "project" our feelings and values onto the [virtual] world. For many of each world's most striking features don't just *seem* inherently terrifying, delightful, beautiful, desirable, nasty, etc. They *are* inherently terrifying, delightful, beautiful, desirable nasty, etc. Thanks to the outgrowths of of our limbic emotional powerhouse, that girl [warthog/hippopotamus,gazelle etc], for instance, does not just *seem* sexy. She possesses the inherent property of sexiness as part of her very essence [or, more precisely, as part of her fleeting psychochemical excitation]. This identification is possible only because virtual worlds are strictly mental: in the realm of phenomenology the difference between seeming and reality dissolves. Just so long as the relevant causal covariation with the mind-independent world is retained, the emotional saturation of a [virtual] world tends also to be highly adaptive. Our genes have outrageously biased what matters to their neural creations - us - so as to differentially further their own reproductive prospects. None of this would be possible if classical perceptual realism were true; but then its intellectual sell-by date has now passed, even though the hard-wired illusion remains.

The philosophical and scientific story of mind-making is much more complicated than the simplistic outline offered here. Yet it's abundantly clear that natural selection has ensured that many organisms have horrible minds, live in horrible virtual worlds, and suffer horrible deaths. Is this an immutable law of Nature? Probably not.

The Post-Darwinian Transition

"There are a thousand hacking at the branches of evil to one who is striking at the root"

Henry David Thoreau

Within the next few centuries, a quite startling option will become technically feasible. Nanotechnology and genetic-engineering will allow us to abolish the biological substrates of suffering in all sentient life. The unpleasant forms of consciousness are set to pass into evolutionary history. Potentially, unhappiness and all its vestiges will become biologically obsolete.

Bombastic fantasy? Like most predictions of events further than a few years ahead, prophecy of this nature all sounds rather fanciful. In reality, how likely are we ever to implement a biological blueprint for universal bliss?

The Hedonistic Imperative outlines why and how tomorrow's biotechnologists will be equipped to practise systematic paradise-engineering. Its consequences will be of a beauty and a grandeur that we can scarcely begin to comprehend. Here, a more narrow issue will be discussed. If we *do* decide biologically to naturalise the sublime, will we consider it ethical to sustain in other species the barbarities of Darwinian regime we've chosen to abandon ourselves? Is it likely that any notion of ethical progress will end when we've liberated members only of our own particular race from the gene-driven malaise of the past? Or instead, mercifully, will a serotonergically enriched capacity for empathy incalculably deepen our compassion for the sufferings of others, while a genetically amped-up "dopaminergic overdrive" propels post-humanity's cross-species biological rescue-job into the Post-Darwinian Era? In the latter scenario, our descendants in the transitional phase are likely to be vastly more moral as well as happier than us. This is because they will be capable of both a greater empathy and a greater capacity to act upon it. Such parameters are genetically tunable; and can be drastically enhanced. So what are the odds of this happening? Are we stuck with the selfishness of a Machiavellian ape from the African savannah until the end of time; or is something else better in store?

First, consider how we might react if we discovered an extraterrestrial civilisation of organic creatures, let's call them Ecstatics. Their geneticists long ago banished the biological substrates of extreme anguish and everyday malaise alike. Ecstatic life is utterly wonderful. Gradients of well-being animate how they think and act; unpleasantness has simply been written out of the script. Their whole existence is endlessly exciting and profoundly fulfilling, day and night. Ecstatics think of physical and mental pain as bestial hangovers from evolutionary history. More commonly, they find it hard to conceptualise such severe mental illness at all.

Let us suppose that these angelically happy super-beings aren't electrode-studded wireheads hooked up to pleasure-machines. So they don't spend their lives like lever-pressing laboratory rats, frantically practising intra-cranial self-stimulation. Nor are they dull-witted opiated dupes of a ruling elite *a la* Huxley's *Brave New World*. These are the two simple-minded scenarios typically evoked by the prospect of getting rid of life's nastiness; and they tend to exhaust our normal range of imaginative possibilities if asked to evaluate what eternal happiness would amount to on earth. Instead, Ecstatics are genetically pre-programmed to enjoy rapturous states of consciousness throughout every day of their lives. Joy is a background precondition of daily existence. Their everyday textures of awareness have a diversity, intensity and sublimity which our own human legacy wetware cannot normally glimpse, still less sustain. The rock-bottom baseline of Ecstatic mental health still ensures that each moment of their lives comes as an exhilarating revelation. Boredom is neurologically impossible. Moreover, Ecstatics aren't merely happier than DNA-driven emotional primitives from the Darwinian era. They enjoy biologically enriched neural substrates of motivation too. Thus

Ecstatics are driven by a willpower far stronger than anything of which contemporary humans are physiologically capable. So they don't sit around all day in a contented zombified stupor. On the contrary, their raw dynamism and irrepressible appetite for life far exceeds our own.

Now what sort of arguments might we try and use to convince Ecstatics that they should restore, or create, a taste of the suffering and everyday discontents that pervade our own late-DNA world? How might we explain and justify any potentially ennobling and life-enriching properties which [we sometimes tell ourselves] unpleasant modes of consciousness often possess; and which Ecstatics are in danger of forgetting? Would we try and *compel* them to rewire their minds for genetically predisposed suffering - for their own good, naturally. Or is the very idea itself monstrous?

Ecstatics themselves, we may suppose, regard experiential nastiness of any kind as coarsening, brutalising, and pornographic. Their ancestors abandoned such obscenities a long time ago. So how might we persuade them that their intuitions of unnatural obscenity and immorality are wrong? What valuable but nasty properties precisely might we identify within our own mode of existence that richly fulfilled Ecstatic lives were lacking? We may suppose that, for their part, Ecstatics treat our reluctance to share irresistible happiness as part of a hereditary, mood-congruent thought-disorder. Have they committed a terrible collective mistake? If we argue the case for traditional life's fitful mayhem and misery over the new gene-driven paradise, could we be rationally confident we were acting as anything nobler than vehicles - and byzantine mouthpieces - for selfish DNA?

Perhaps; but it's not easy to show how.

Ecstatic aliens are science-fiction, to the best our knowledge at any rate. By contrast, the impending need to justify suffering - as and when opting to retain its neural mechanisms becomes a life-style choice rather than brute biological fate - isn't fictional at all. Of course, the idea that something as apparently *inevitable* as suffering will ever require ideological justification may seem a cruel joke today. The textures of unpleasantness are integral to our lives and even our loves. Yet as we understand and progressively manipulate the substrates of mood and emotion, we will need, sooner or later, to defend the deliberate infliction or conservation of their nastier modes of operation in others. For it will be late-/post-humans who decide when, where and how other life-forms will suffer. And if we *aren't* prepared to tolerate such tampering with our own or anyone else's DNA-driven psychophysiology, then we will need to think hard about what laws or other punitive sanctions to use against people who *do* want eternal happiness. A cry of "Just say no!" probably won't prove discouragement enough to stop them.

Realistically, the use of systematic coercion to enforce legacy-Darwinism is unlikely to work indefinitely. Ethics aside, that's one pragmatic reason why it shouldn't be tried. Yet if life-long super-ecstasy is genetically codeable [as, of course, would be life-long tortured hellishness or depression; for there are solutions to the (generalised) Universal Schrodinger Equation which make Auschwitz look like a fun-filled utopia], should the unprecedented well-being it delivers become the hereditary birthright of only a single trans-human super-species? Or should it be zoologically universalised? In centuries ahead, should we intervene in the rest of the living world to rescue its entrapped life-forms from the "natural" horrors of which they are the helpless and blameless victims? Or should we just leave them to it? Is there anything morally wrong with applying nanotechnology and genetic engineering systematically to reorder the natural world so that it's a fabulous place to live for the whole lot of us? Or is this sheer hubris, since selfish DNA makes a morally better world than anything conscious mind can engineer by design? What are the arguments for and against creating a naturalised heaven-on-earth for all our fellow creatures, and all our states of consciousness?

Such questions today have a pronounced air of unreality. Doesn't this guy have a job? If getting rid of human suffering sounds wild, scrapping *animal* suffering, too, sounds positively flaky. Actually, the moral argument for abolishing non-human animal suffering as the technical obstacles come to seem less mountainous is stronger than for humans. For at least sophistical arguments can be concocted to justify the need for obligate human malaise. One will be told how it builds the character, ennobles the spirit, and leads to great works of art and literature etc, though if one listened to some critics of the prospect of universal happiness, one might be forgiven for supposing that humanity's consuming passion was producing great literary classics, not pursuing money, power, drugs, and sex. Yet *animal* suffering is not character-building, nor does it lead to life-affirming works of art etc. It is just nasty and pointless. The nearest it gets to mimicking any kind of Meaning is the way it serves *as though* it had the purpose of helping self-replicating DNA to leave more copies of itself. But that's as far as it goes. It's not good for anything but some twisted bits of DNA. So the case for abolishing unpleasantness in animals is *at least* as compelling as it is for humans once mature biotechnology turns its abolition into an implementation-problem rather than a harebrained philosopher's fantasy.

But won't a world without traditional predatory carnivores in all their bloody and savage glory be less diverse and therefore more boring? Aren't cats cool?

The boredom issue is a gigantic red herring. If we wanted to, future neuroscientists could make a lifetime spent watching grass grow into a nailbiting psychological cliffhanger; although, as it happens, no such contrivances will be needed. Stripped of its predisposing genes and neuronal substrates, boredom will become physiologically inaccessible to anyone. Its peculiar vapid texture was just a phase certain forms of early DNA life went through. The particular kinds of neural negative feedback mechanism which boredom reflects will become obsolescent too. By contrast, everything in the post-Darwinian world will be much more vividly intense than today's "normal" life. Moreover, as a bonus, naturalised biological paradise will be a far more richly differentiated place too. For we've scarcely even begun to explore the galaxy of wonderful experiences it's possible to savour and delight in. These won't be only the sparkling deliverances of newly-engineered senses. They'll include new modes of introspection and meditative consciousness extending way beyond the shallow reveries of anything neurochemically accessible today, even by the deepest-dyed mystic. Natural selection previously stopped us accessing these outlandish modes of experience. This is because coding for their substrates would have involved either occupying, or crossing, maladaptive gaps in the genetic fitness landscape. For now, however, we're stuck unwittingly ringing the changes in our own mediocre repertoire.

It's true that the post-Darwinian world won't be *maximally* diverse. There won't be any suicidal despair, jealousy, bubonic plague or child-abuse. Moreover if we *did* want to maximise ecological diversity, we could breed creatures that naturally prey on humans. For if we arrange matter and energy in the right way, it's feasible to design obligate predators who can thrive only on human flesh. But who cares? What's the point? The absence both of vileness and the mundane deformations of consciousness which we presently take for granted would be morally bad only if diversity were inherent good. But it's only good today insofar as it stops hyper-dopaminergic novelty-seekers from getting bored. When boredom is impossible, and bliss biologically ubiquitous, then why adulterate perfection with ugliness?

For his part, DeGrazia is right, I think, to argue for the intellectual incoherence of many of our traditional intuitions. He is also right to argue that we must radically change our attitude to non-humans. Yet then - understandably perhaps - his intellectual nerve fails. He falls back on a conventional conservatism when contemplating the fate of victims of the primeval Darwinian order.

Until recently, it's true, the only appropriate response after absolving oneself of any direct *personal* complicity in the suffering of other life-forms has indeed been been to leave

things to Nature. The trouble is that this approach amounts to a far less benign solution than its soothing verbal formulation suggests. Urban-dwelling animal activists are, on the whole, far too romantic about the natural world. With our cloistered, media-filtered conception of the Great Outdoors, we implicitly rely on a filter of sanitised wildlife programmes to tell us what the animal kingdom is supposedly all about.

In fact, Nature documentaries are mostly travesties of real life. They entertain and edify us with evocative mood-music and travelogue-style voice-overs. They impose significance and narrative structure on life's messiness. Wildlife shows have their sad moments, for sure. Yet suffering never lasts very long. It is always offset by homely platitudes about the balance of Nature, the good of the herd, and a sort of poor-man's secular theodicy on behalf of Mother Nature which reassures us that it's not so bad after all.

That's a convenient lie. If you had just gone through the horror of seeing your loved one eaten alive by a predator, or die slowly of thirst, you would find such clichés empty. Yet in Nature this kind of thing happens all the time. It's completely endemic to the prevailing red-in-tooth-and-claw Darwinian regime. Lions kill their targets primarily by suffocation; which will last minutes. The wolf pack may start eating their prey while the victim is still conscious, though hamstrung. Sharks and the orca basically eat their prey alive; but in sections for the larger prey, notably seals. An analogous scenario in which intelligent extraterrestrial naturalists turned the stylised portrayal of *our* death-agonies into a lyrical spectacle for popular home entertainment is repugnant. Yet as long as we revel in the production of animal snuff-movies in the guise of wildlife documentaries, that is often the role we play in the tragic lives of photogenic members of other species here on earth.

There is, of course, a danger in harping on about the terrible extent of suffering indigenous to Nature. One runs the risk that such accounts may be used by hunters and non-obligate meat-eaters as a license for our massively *adding* to the savageries which already exist. It's simply the way of the world, we are told. There's so much suffering around already that increasing it a bit won't make much difference.

This sort of cavalier attitude to the fate of others is morally catastrophic. Any reversion to the traditional cruelties of a primordial selfish-DNA regime after abolishing its ghastly late-industrial culmination would amount to a calculated act of barbarism - possessing all the ecological naturalism of a heritage-industry theme-park without any of its redeeming folksiness.

This passive abdication of responsibility - directed at humans it would be called culpable neglect - is still a popular option among animal advocates. It is encouraged even though Nature is so often frightfully cruel - in its effects though not through some purposive malevolence. Nature is nasty not because most creatures have the sophisticated theory of mind and higher-order intentionality required to encompass human-style sadism. The reason is merely that natural selection has placed no check at all on how bad suffering can be wherever its existence - or any behavioural capacity associated therewith - has let some gene coalitions leave more copies of themselves than others. Encephalising horrific modes of experience so they get conditionally activated is a very effective way of spurring living vehicles to behave in ways likely to maximise the inclusive fitness of their DNA. It's utterly vicious and compelling. But life doesn't have to be like that.

It is only quite recently that a strategy for genetically engineering the complete abolition of aversive experience in humans has even been mooted - let alone a strategy extending the rescue-mission to non-human animals. But then it is only quite recently that earthly blueprints for its biological implementation could be devised.

Happily, a completely unprecedented revolution is in the offing. First, the option of worldwide genetically preprogrammed sublimity is no longer technically inconceivable. Scrapping the root of all evil in its biochemical manifestation is winning acknowledgement as at least a *theoretical* possibility; even though the prospect is typically regarded as wild and eccentric. Certainly, the mass-use use of long-acting depot contraceptives, cross-species retroviral gene therapy, self-reproducing micro-miniaturised nano-robots with supercomputer processing power and therapeutic bioengineering capabilities, etc, all sounds outrageously sci-fi - certainly not practical politics today. Yet this credibility-gap may close quite abruptly. The call for worldwide paradise-engineering isn't an empty plea for new physics, biological wonder-tissue or superluminal warp-drives. At first, of course, the family of ideas underlying the whole post-Darwinian enterprise will seriously occupy the minds of only the [currently-defined] scientific and political fringe. The technical challenges posed by abolishing all of what's wrong with the world are far less formidable than the alteration in mind-set needed to plan the post-Darwinian biological project in the first instance. Yet as meta-paradigm-shifts go, getting rid of aversive experience isn't *conceptually* difficult. Paradoxically, it's the tender-minded people who care most about animals who are also the folk likely to be most appalled at the hard-headedness required to implement the indefinitely sustainable psychological and physical well-being in prospect for life on earth. Gung-ho testosterone-driven technophiles, on the other hand, are less likely to care about the suffering of lesser creatures which their own personal technical expertise makes preventable. The conservatism of the tender-minded is understandable; but profoundly reactionary. If triumphant, its living victims will continue to be sacrificed on the altar of a Mother Nature whose existence in rose-tinted guise has no place outside the romantic imaginations of its creators. For if our cars and computers should be precision-engineered, then why shouldn't the biomolecular architecture of consciousness? Well-being is too important to be left to selfish DNA.

DeGrazia, in discussing our alleged lack of positive obligations to animals, uses the pejorative phrase "meddle with nature" (p277) Certainly, if trying to subvert the biological *status quo* meant starving the lion to save the gazelle, it is mostly futile. Preventing suffering for the one is effectively causing suffering to the other. (The situation of obligate predators such as cats is very different from omnivorous humans. Indeed, as DeGrazia aptly notes, "while many steak-loving humans like to regard themselves as part of this vast chain of carnivorousness, they neglect the fact that omnivores do not need meat to survive with good health. Indeed, overall, meat may do us more harm than good")

Yet a much more revolutionary imagination is needed. DeGrazia simply doesn't entertain the possibility that genetic-engineering might enable us to abolish all aversive experience and replace the monotony of the hedonic treadmill with a fabulous diversity of enjoyable states. His implicit conservatism is perhaps understandable: the notion of a race of beings who are animated wholly by pleasure gradients is - for now at any rate - a figment of otherworldly dreamers, in spite of its biotechnical feasibility, and - as millennial global species-projects go - technical simplicity. It's worth recalling that physicists and visionary AI buffs routinely discuss proposals far more exotic. Huxley's static and dystopian vision of *Brave New World* - where the chemically-tranquillised masses were all sedated and opiated dupes of the power elite - has had the unfortunate effect of asphyxiating professional scholarly thought and political action on the immense range of paradise-engineering options in prospect.

A host of down-to-earth practical objections to making paradise happen do of course spring to mind. Suffering might seem too widespread and diverse in the animal kingdom ever to be eliminated altogether. Inside the ghetto of malaise, its modes can sometimes seem infinitely varied. We can readily understand how a fellow creature can be always in some way unhappy or in pain - most of us know someone whose life is spent in such a state. The idea that the reverse condition is ubiquitously feasible - that each of us could always feel happy and gloriously well - initially strikes us as bizarre.

Yet our intuitions are utterly misplaced. We elevate one generic yet parochial feature of consciousness, albeit a feature which has pervasively innervated the mammalian neocortex thanks to evolution - to the status of timeless feature of the world. Moreover the idea that the only way to ensure perpetual happiness would be to turn us into incapacitated wireheads is no more realistic than the correlative notion that the only way to make someone perpetually miserable would be to implant electrodes into their pain centres. Alas, this simply isn't the case.

Valuing Consciousness

> *"The dissolution of commercial animal farming as we know it obviously requires more than our individual commitment to vegetarianism. To refuse on principle to buy products of the meat industry is to do what is right, but it is not to do enough. To recognise the rights of animals is to recognise the related duty to defend them against those who violate their rights, and to discharge this duty requires more than our individual abstention. It requires acting to bring about those changes that are necessary if the rights of these animals are not to be violated. Fundamentally, then, it requires a revolution in our culture's thought about, and its accepted treatment of, farm animals... But prejudices die hard, all the more so when they are insulated by widespread secular customs and religious beliefs, sustained by large and powerful economic interests, and protected by the common law. To overcome the collective entropy of those forces against change will not be easy. The animal rights movement is not for the faint heart"*
>
> Richard Ryder

Hard-nosed scientists and traditional analytic philosophers are likely to feel that much of this review essay is idle opinion. OK, this guy gets worked up about getting rid of cruelty and suffering; but so what? Moral seriousness here is implicitly taken by philosophical sophisticates to be *intellectually* frivolous. The language of morals is basically verbiage because value-judgements aren't reckoned to have truth-conditions. They amount instead to an obliquely autobiographical commentary on the state of mind of their author. By common scientific consent, they don't consistently pick out objective features in the mind-independent universe. In issuing value-judgements, all we are doing - assuming adherence to a pre-QM classical realist fantasy of perception - is unwittingly projecting our feelings onto the world "out there" as [supposedly] disclosed by the senses. We think we're "reading off". In fact, we're "reading in". Admittedly, we don't have any wholly satisfactory theory of truth - any more than meta-ethicists have any wholly satisfactory theory of the meaning of value-judgements. Yet only some sort of correspondence theory of truth is going to be viable; and unlike statements of fact, what value-judgements express clearly *doesn't* correspond to anything which could potentially make them true or false. They are therefore, it's alleged, just (in)convenient fictions: DNA-driven adaptations, explicable perhaps in terms of human evolutionary psychology, but still inventions of the human mind. In this context, the very title *Taking Animals Seriously* gives the game away. Why not take animals lightly? Surely, there's simply no fact of the matter either way?

In defiance of the ill-named naturalistic fallacy, I'm going to argue there *is* a fact of the matter. A post-Darwinian world where suffering has been replaced by states which *seem* self-intimatingly valuable really *is* more valuable, no less than it will be more blissful. Furthermore, all-pervasive well-being is not just more valuable than the endemic miseries of the status quo. It's better than DeGrazia's well-intentioned but ultimately cosmetic reformism, which still leaves the bloody and pain-infested Darwinian legacy biologically entrenched both in and around us.

How come? Doesn't this sort of claim mix up prediction with prescription? Surely value isn't like, say, pain? Value-judgements at least purport to have propositional content; and thus are potentially true or false. Whether they're merely privately entertained or verbally expressed, they serve as vehicles for expressing something over-and-above the ill-defined phenomenological properties of particular spatio-temporally located episodes occurring in the mind/brains of people who physically make them. Pain doesn't have this sort of content. It's just painful. It's not "about" anything. Pain is self-intimating. Value isn't. In the case of pain, for sure, the "seeming" and the reality genuinely are indistinguishable. At least in its rawest and purest form, the experience of pain isn't shot through with theoretical assumptions about its nature above-and-beyond its self-disclosing nastiness; or often even with [fallible] attributions of its cause. You can't be mistaken about being in agony. So at the risk of succumbing to a naive semantic empiricism, we can be depressingly confident that no revolutionary scientific discovery could ever reveal that pain didn't exist; the distinction between Appearance and Reality vanishes when the reality at issue is appearance itself. Yet to claim the same about value, and to claim that value is self-authenticating simply because [currently] peak experiences appear self-intimatingly valuable, is mere tricksy verbal manoeuvre. Mapping out the world's ontology needs hard experimental work; not an exercise in inward-looking contemplation.

So what's going on? Values may seem to be about something external to experience if one retains a classical realist theory of perception. The rival picture of billions and billions of autobiographical virtual worlds, each chattering with mentalese masquerading within as public speech, is scarcely conventional wisdom. And yet if values "really" existed outside the distinctive quality they lend to certain forms of experience - dressed up in virtual world furniture or otherwise - they would be weird and cognitively inaccessible objects, wholly out of place in a naturalistic world-picture. So there's simply no need to posit such ontological extravagances at all. But if they can't intelligibly be treated as platonic objects, they are even less plausible as candidates for natural properties of the [inferred] mind-independent world. There simply isn't room for goodness and badness in the scientific world-picture as revealed by physics. So why not banish values altogether? Surely the "naturalistic fallacy" was debunked a long time ago?

Yet the naturalistic fallacy is only a fallacy if value is propositionalised and treated as something external to experience itself; and since causally inert, non-spatio-temporal abstract propositions are scientifically unnaturalisable, this exile is probably ill-advised. [How a natural world can *simulate* a world where truth-evaluable "propositional content" exists is another story. Ours does, quite uncannily; natural selection is quite superb at simulating the miraculous, though it's unaccountably silent concerning the date on which the first semantic miracle allegedly occurred]. Many modes of experience that are apprehended as valuable - typically bound up today with encephalised feelings involving neural representations of genetically advantageous qualities - aren't normally conceptualised by their subjects as particular modes of experience. Their valuable aspect is known under other descriptions entirely. Value appears instead to inhere in the properties of the particular (and typically, reproduction-enhancing) objects, properties, people or behaviour that excite such judgements. Yet to identify value as dwelling beyond our psychoneural virtual worlds - as a spooky, ill-located sort of ontological furniture - is not simply incoherent. It is to presuppose one has non-inferential access to the mind-transcendent universe that we simply do not, and could not, possess.

But then what is the ontological status of value? What is it *really*?

Well, what is the ontological status of phenomenal blueness? It's a mode of experience. Awake or dreaming, it's mind-dependent, albeit often selected from the mind/brain's finite menu of states by patterns of peripheral stimuli. It can't be defined in terms of anything else external to itself. Formally, its occurrence in one's mental world may indeed be field-

theoretically encoded by the equations of QM; and natural selection has ensured that awake mind/brains normally undergo it when optic nerves are triggered by electromagnetic radiation differentially reflected from macro-patterns in the local environment. Yet this doesn't make blueness any less unique and irreducible, or "really" something else. If we ever understand why the quantum mechanical formalism codifies the structures and interrelationships of different kinds of consciousness to yield the exact textures it does - or indeed any texture at all - we'll be able to manipulate and create blueness [or manufacture and maximise valuable experiences] in a more effective and more quantitative sense than we can today. Yet this won't denature its phenomenology; or it would be something else, a property of a different kind in a different kind of world.

The story of value is more complicated on account of of its accumulated ideological baggage; but phenomenologically at least, maximally valuable and maximally happy worlds would seem to be coextensive. A world without misery is a world without moral dilemmas. A maximally happy post-Darwinian cosmos is unsurpassable in both its quantity and quality of apprehended value. It's impossible to feel blissfully fulfilled and find blissful fulfilment valueless. Conversely, a world or Everett hell-branch which was literally full of suicidal despair and pain would not just *seem* utterly valueless. It would quite literally lack any positive value at all. If the predictions of HI are borne out, on the other hand, the world of our descendants will be biologically supercharged with value to a degree exceeding our present notational resources. Images of hyper-intelligent but jaded alien civilisations of sophisticates, bored beyond measure with their meaningless lives, are misconceived. They owe too great a debt to watching repeat-episodes of Star Trek; and not enough to contemporary biomedical research.

The likelihood of an ultimate total reconciliation of the phenomenology of well-being and value is not obvious. Traditional and unbiologically-inspired utilitarianism poses various ethical dilemmas, or at least uncomfortable consequences. Superficially, one can imagine possible worlds which were bliss-ridden in a "baser" and more debauched fashion than merely moderately happy but "edifying" worlds. But the comparison is deceptive. "Empty" or "base" happiness [the sort of happiness most commonly associated with taking dirty street-drugs or the furtively pursued pleasures of the flesh], insofar as it is indeed apprehended as "empty" or "base", is sullied happiness and ill-deserving of the name. Such happiness certainly won't be maximal; so the dilemma of possible tradeoffs doesn't arise.

There's still an obvious problem here. The sort of naturalistic analysis advanced here conceptually entails that values are real but mind-dependent. Alarm-bells start ringing here. "Mind-dependent" makes them sound ontologically second-rate. The tension between the two categories arises, however, only if one thinks of the mind as somehow outside the world "looking in". *If* mind/subjectivity weren't a natural feature of the cosmos, then the mind-dependence of value might indeed impugn its status. Yet mind is as much a part of the natural world as are atoms and molecules. The "subjectivity" of value no more threatens its reality than the "subjectivity" of pain makes surgical anaesthetics redundant. For sure, it's all in the mind. But minds are all in the world. What's morally pernicious is third-person metaphysic of contemporary science. Such a metaphysic muddles the two senses of "objective" and thereby sows confusion. The behaviour of the stuff of the world is indeed amenable to description by a mathematical formalism. This formalism encodes how, and to what degree, it matters to a creature that (s)he's undergoing extreme pain. It is objectively true that the world contains such phenomena that matter desperately. The suffering of someone I have never met and don't know of doesn't matter to *me* - and preventing it can't be desperately important to me - but this doesn't mean it is actually any less desperately important. All first-person facts are created equal. Simply because it doesn't matter to [misleadingly called] other 'observers' (who are, on a perceptual irrealist perspective, actually nothing of the kind) then it's easy to suppose that it's more "objective" to discount the plight of others. Yet moral apathy in the guise of observer-independence reflects a morally harmful fallacy of equivocation, not a scientific fact.

Isn't this value-naturalist position self-subverting? The critic who finds the status quo, or DeGrazia's tidied-up state-of-Nature, preferable to ubiquitous happiness may charge that the immense value self-ascribed by our ecstatic descendants to their lives and consciousness is delusive. Phenomenologically, of course, their biological wonderworld is more valuable, simply in virtue of the sheer number, intensity and variety of experiences apprehended as worthwhile. Yet the ubiquity of their value-steeped phenomenology no more means earthly paradise is really more valuable than the medieval penchant for finding witches means that the medieval world was really more witch-ridden. So, surely, I am disagreeing with the critic's position. Yet if value is, as claimed here, a distinctive texture of experience rather than something expressed by propositional content, then I can't self-consistently say that the critic is wrong. For this would be to propositionalise value rather than treat it as a distinctive quality of consciousness. Propositionalising value is the very practice being argued against. And yet if one can't contest judgements of value, then what's the point of this review?; straight studies in the neuroethology of mind might as well be left to the academic journals. Surely this rejoinder counts as a *reductio* of the whole neo-naturalist argument against the value-sceptic?

But this is all far too quick. What's happening here is that the appalled critic is himself occupying a valueless, malaise-infected state of consciousness. Within that malaise-infected state, (s)he is no more capable of semantically capturing the nature of mature post-Darwinian life than my current state of mind can capture the excruciating agonies of an Ottoman torture-chamber. The critic's frame of mind testifies to the value-starved and mean-spirited nature typical of psychopathologies bred by the *present* DNA-regime; not the illusory imperfection of our outrageously wonderful future.

Darwinism With A Human Face?

> *"Because one species is more clever than another, does it give it the right to imprison or torture the less clever species? Does one exceptionally clever individual have a right to exploit the less clever individuals of his own species? To say that he does is to say with the Fascists that the strong have a right to abuse and exploit the weak - might is right, and the strong and ruthless shall inherit the earth.*
> Richard Ryder

DeGrazia draws together his discussion by extracting the principles set out in the list below. Broadly, they represent an extension to other species of the "principle of nonmaleficence": basically, don't cause unnecessary harm. Stated baldly and in the absence of DeGrazia's detailed reasoning behind each of them, they might seem arbitrary in number. Why not list fourteen or sixteen? As short-term stopgaps, they would nonetheless seem good working precepts - number fifteen excepted [see below]. Yet piecemeal tinkering is not enough. Fundamentally, the principles they embody still amount to an endorsement of "Darwinism with a human face". Even their thorough and *complete* state-sanctioned implementation - currently a world away - would leave atrocious "natural" suffering set to continue for millions of years indefinitely. The self-replicating biological machinery which manufactures the world's pain would continue to churn out its living vehicles for as long as Earth is capable of supporting organic life. DeGrazia's guidelines - admirable as is indeed the humane mind which formulated them - are still implicitly conservative of the old DNA regime which threw up the architecture of unpleasantness in the first instance. Only a blueprint for scrapping the generative mechanisms responsible for the mass-production of nastiness in the living world - essentially, bad base-pairs of self-replicating DNA - gets to the root of what needs to be done. Here, however, is DeGrazia's makeshift recipe. It's a start.

WHAT IS TO BE DONE?

1. Don't cause unnecessary harm.

2. Make every reasonable effort not to provide support for institutions that cause or support unnecessary harm.
3. Don't cause significant suffering for the sake of your or others' enjoyment.
4. Apply equally any standards allowing the causing of suffering.
5. Don't kill sentient animals unnecessarily.
6. The presumption against killing humans, Great Apes and dolphins is virtually absolute.
7. For a large class of sentient animals - at least fish, herpetofauna (amphibians and reptiles) - and birds - the presumption against killing these animals is ordinarily weaker than that against killing humans, Great Apes and dolphins.
8. Don't confine sentient animals unnecessarily (where confinement is understood as the imposition of external constraints on movement that significantly interfere with one's ability to lead a good life).
9. There is a strong presumption against confining nondangerous sentient animals.
10. The presumption against confining innocent humans, Great Apes and dolphins is virtually absolute.
11. To the extent that we can separate out freedom interests in practice, for a large class of sentient animals - at least fish, herpetofauna and birds - the presumption against confining them is ordinarily weaker than that against confining humans, Great Apes and dolphins.
12. The conditions of any justified confinement must be responsive to the animal's needs.
13. There is a presumption against disabling sentient animals (that is, damaging their ability to function in ways that significantly interfere with their ability to live a good life) and if they are nondangerous, the presumption is virtually absolute
14. Provide for the basic physical and psychological needs of your pet, and ensure that she has a comparably good life to what she would likely have if she were not a pet
15. If (hypothetically) there appears to be a genuine conflict between benefiting an animal and respecting her autonomy, unless the expected benefit is very great and the apparent infringement of autonomy very marginal, respect autonomy.

NUMBER FIFTEEN

Why number fifteen? Well, the stress on "autonomy" will surely play well with a lot of readers, not least the domestic audience of rugged individualists resident in DeGrazia's American academic home-turf. In the ideological aftermath of America's triumph in the Cold War, anything which smacks of statism, welfarism, and socialistic paternalism is still taboo; and DeGrazia extends such distaste to the rest of the animal kingdom too. Yet even the most admirably libertarian paediatrician, for instance, would hesitate to apply an anti-interventionist philosophy to human infants and toddlers; their cognitive limitations are too severe, even though their juvenile feelings are intense. Intellectually, few animals are any brighter - as distinct from more self-sufficient - than human toddlers; and if, as argued throughout this review-essay, their moral status is similar, a non-interventionist philosophy toward (at least) vertebrate animals is equally misplaced. They need our help, even though they don't know it.

Set in much wider perspective, our faith in individual autonomy, humanoid or otherwise, is fundamentally misguided because there is nothing truly autonomous about constituting a throwaway neurochemical robot [built as though it were] designed by Evolution to leave more

copies of one's genes. Many of the things we're genetically predisposed to want, feel and do, alas, are profoundly psychologically damaging to the emotional well-being of each of us. Yet they fester and multiply because they serve the reproductive advantage of the DNA which spawned vehicles like us. The superficially anomalous way in which certain organisms are endowed with a [highly constrained] measure of notional *choices* for action - in our mind's eye, we can run toy simulations of alternate scenarios which plausibly ensue from initiating different behavioural options - is an immensely useful adaptation, given a depressingly compelling gene's-eye view of the world. In humans, this faculty even gets philosophically dignified by the name of Free Will. Yet no life-form gets to choose the laws of physics and chemistry that determine which distributions of matter and energy are instantiated when and where - and that includes humans and their nominal choices. Neither the world's Master [Wheeler-DeWitt etc] Equation, nor any of its solutions, are dictated by mere mortals. Raising the phenomenology of voluntary action and anticipated consequences into some sort of metaphysical principle or separate ethical ideal is thus rationally ungrounded, to say the least. For the neural substrates of volition, willed action, and our sense of Freedom itself are as biologically manipulable as any other chemical reaction; and have been pressed into service by selfish DNA. Autonomy only matters to the extent that perceived restraint feels aversive; it's just another state of mind in the great cosmic mathematical dance.

In an extremely limited context, DeGrazia is probably right to endorse a hand's-off approach. Misguided attempts to subvert any organism's (pseudo-)autonomy which result in more animals suffering rather than less are evidently best avoided. Yet this is an argument for greater understanding of the neurobiological substrates of what promotes their well-being, not inertia. In practice, what count as sins of commission and sins of omission are time- and culture-bound. A lot of the future ways of helping animals will require massive and systematic *intervention*: genetic engineering, long-acting depot contraception, ecosystem re-design, and eventually nanotechnology-based reworking of the whole molecular architecture of the vertebrate nervous system. This might seem to involve radical discontinuities in the evolution of life and consciousness; but such transitions have occurred before, and - extrapolating - there may be many more. Serious futurology is not a game of trend-spotting.

If You Think It's Murder, Act Like It

> *"Animal Liberation will require greater altruism on the part of human beings than any other liberation movement. The animals themselves are incapable of demanding their own liberation, or of protesting against their condition with votes, demonstrations or bombs. Human beings have the power to continue to oppress other species forever, or until we make the planet unsuitable for living beings. Will our tyranny continue, proving that we really are the selfish tyrants that the most cynical of poets and philosophers have always said we are? Or will we rise to the challenge and prove our capacity for genuine altruism by ending the ruthless oppression of species in our power, not because we are forced to do so by rebels or terrorists, but because we recognise our position is morally indefensible? The way in which we answer this question depends on the way each one of us, individually, answers it."*
>
> Peter Singer
> ANIMAL LIBERATION

An old philosophical tradition consists in simply expounding the truth as one sees it. One then just waits until sheer force of argument allows one's conclusions to become generally known and luminously self-evident. Perhaps this sort of dispassionate engagement with the issues will indeed prove enough to rescue billions of presently unborn victims of human inhumanity to non-humans in decades to come.

Unfortunately, disembodied rationality, even if it existed, would be causally impotent; and very little good to anyone. Certainly, placing one's faith in the dawning light of reason isn't always a recipe for success where immensely powerful and hostile vested interests are at stake. And the vested interests defending animal-exploitation are very powerful indeed.

So how can the revolution in our treatment of our fellow subjects be brought about?

There was a time not so long time not so long ago when the idea of philosophers in the dominant Western analytical tradition actually *doing* anything to promote ethical conduct would be regarded as incongruous to the point of being laughable. Within the academic profession itself, making any first-order ethical judgements - as distinct from practising meta-ethical analysis of the linguistic meaning of value-judgements - was regarded as faintly disreputable. *Being* right in one's judgements counted for more than *doing* right. Indeed, even after the demise of Oxford-style linguistic philosophy, students in ethics classes were - and still are - graded purely and simply on the basis of what they say - and not for what they do outside the seminar-room. Admittedly, the question of whether those who presume to mark and grade students of Ethics would be better qualified to sit in judgement than their charges is at the very least open to debate. But either way, if morality of any kind is more than empty vapouring, then surely deeds rather than empty pieties are what matter. This demands a strategy of action.

As it stands,*Taking Animals Seriously* is a well-written, well-researched piece of analytic philosophy. If it convinces any sceptical or under-motivated readers of the intellectual underpinnings of the case against animal-abuse - as I think it should - then it will serve a valuable purpose - *if* it then bestirs them to take action. What the work as a whole doesn't do is offer policy prescriptions in a political sense. It amounts to an ethical treatise-cum-philosophical study of non-human minds, not a political tract. DeGrazia advances closely-reasoned arguments on what ought to be the case. He sets out, in general terms, how we ought ethically to behave. Yet *Taking Animals Seriously* ducks any investigation of the specific mechanisms by which even its own *comparatively* tame proposals can be made to happen. Certainly, it doesn't devise any sort of political strategy, or lay out an organisational framework, by which a revolution in the treatment of our victims can be brought about.

So what should be done - or, much better, what should we ourselves *do* - to try and stop the holocaust?

First, here's a schematic review of some of the options.

Acts of *violence* against the abusers rarely help the abused. They are part of the very Darwinian heritage one is trying to transcend. This isn't a call for sanctimoniousness - as distinct from clarity - in condemning the actions of the minuscule handful of activists tempted to pursue this sort of radical activism. Passive, turn-the-other-cheek acceptance of unprovoked violence directed against oneself may be admirable. Acquiescence in its institutionalised and unprovoked infliction on others demands less stoicism and no great heroics.

None of these caveats changes the fact that violent action against other persons is - in general - disastrously ill-conceived. Most of us are ourselves, in any case, so implicated by our consumerist lifestyles in the exploitation of others that the singling out of some abusers rather than others is often a matter of punishing visibility rather than objective consequence. Typically, it is the ramifications of our acts that are frightful, not the inherent character of the agents who commit them. Thus condemnation of animal-abusers - or those who in vain try physically to stop them - should at best be seen instead as a rhetorical tool of purely instrumental value. It's not a matter of some metaphysical assignment of guilt. Under the eye of eternity, even a Hitler or a Pol Pot is no more guilty - or innocent - than a smallpox virus. For we are all animals. Our behaviour is exhaustively described by a set of natural laws which

we didn't choose and of whose playing out we are all a part. Fortunately, it transpires that a non-obvious consequence of these laws is the development of a species blessed with a capacity to overthrow the Darwinian regime to which those same laws gave rise. Suffering, it transpires, has temporal boundaries as well as spatial ones.

The issue of *nonviolent* direct action to prevent the institutionalised atrocities of the death-camps and factory-farms is more complicated; and ultimately less clear-cut. In the end, however, one's conclusions - I think - have to be broadly similar. Brute force usually doesn't work. Once again, however, such a plea for legalism shouldn't be used as pretext for the usual knee-jerk sanctimony. Unwarranted self-righteousness tends to get directed against those with the courage to show more than an easy but vicarious stoicism at the suffering of the oppressed. For there is an incongruity to our hand-wringing over *Why Didn't We Bomb the Death Camps?* for instance, and our condemnation of the Nazi experiments on humans as the ultimate abomination - while at the same time we collaborate with regimes guilty of sanctioning the very same acts of killing and cruelty against highly sentient non-humans (who are invariably described as *only* animals, as though their vulnerability and helplessness meant they mattered less then the reigning *Herrenvolk*). We may be baffled how Eichmann and Mengele could be decent family men and yet do such terrible things to children. Yet their attitude to their helpless victims was not radically different from ours to "inferior beings". Only a few of us actively enjoy causing suffering to those we exploit and kill. For the most part, we are simply oblivious to it. Or we (mis-)conceive it as too trivially insignificant to worry about. In our case, our victims are, after all, *subhuman* - not even *untermenschen* - and our use of the very word "animals" coveys a sense of superiority and disdain. For the most part, it is simply a matter of convenience to treat non-humans in the way we do. Abusing animals for money, taste or curiosity and even *fun* has for long simply been a part of the way the world works.

Of course, in the wake of the growth of the animal-rights movement, there has recently arisen a hitherto unfelt need to demonise and demean our non-human victims - and those who try to help them - now that our previously well-nigh unquestioned right to kill and exploit them is being challenged. Bloodsports enthusiasts, for instance, currently spend a lot of time cataloguing the alleged depredations of our *victims* on the environment. Recreational animal-killers go to extraordinarily lengths to avoid admitting that they themselves enjoy hunting and killing other creatures for fun. But then until a few years ago such rationalisations seemed scarcely called for. Selfish DNA had honed our intuitions so that the most agonising bloodshed seemed simply "natural".

Given the enormity of what we're perpetrating, then why the qualms one may feel about going beyond restrictive legalism - itself a tendentious term, since all governments in the world today rest on some original act of illegality? Aren't all methods of bringing the old order to an end justified? The dilemma for the ethical utilitarian, typically of a legalistic and profoundly pacifistic bent, is that history doesn't show that the quiet conscience, or even full-throated protest, is more successful than physical intervention. The most effective tactic historically has been a combination of both. Thus from a utilitarian perspective, it's hard to know whether sabotaging the economic and technological infrastructure of the death factories, and physically destroying the machinery of killing, vivisection and factory-farming, isn't morally justified if at all feasible. Perhaps it is. This is because threats to the apparatus of oppression steeply push up the costs of abuse. [Threatening the oppress*or*s raises darker issues altogether] The ghastly if sometimes high-minded atrocities committed in university 'research-labs', for example, have been scaled down - though they still continue - not in the main through Damascene conversions, conscience-stricken crises of faith in the still hours of the night, or the genteel promptings of ethics committees. They've been restricted because of the high price of security measures needed to safeguard those laboratories where the horrors have been taking place. It would be nice - and extraordinarily convenient for the liberal bourgeois sensibility, whether wired for the digital age or otherwise - if history recorded that oppressor groups did succumb

to polite and dignified protest. Unfortunately, not many instances of such spontaneous acts of collective goodwill spring to mind. We're far too good at rationalising base self-interest.

Tactically, on the other hand, there are certainly strong arguments in favour of legalism. Despite the strict pacifism of most ALF activists, it's inevitable that the profiteers and the bureaucrats in charge of the non-human killing-apparatus, and the beneficiaries of the whole economic empire of ancillary services on which it depends, will talk - invariably without irony - of the violence and terrorism of their *opponents*. In the Orwellian lexicon of the killers and their apologists, the destruction of 'private property' - i.e. the instruments of mass-killing - is invariably dubbed 'violent' and "terroristic." The institutionalised physical abuse and killing of non-humans, on the other hand, is bizarrely categorised as law-abiding and peaceful(!). Transposing the respect due to sentient beings as *subjects* to physical *objects* is simply one of the more grotesque examples of the ideology of animal-abuse. Whatever the grisly ironies, the fact remains that the power of the modern state is always likely to snuff out direct action. The only possible exception is the coordination of mass civil-disobedience which follows breakthroughs to a critical mass of public support; after which it should be unnecessary. Moreover direct *physical* action against the material infrastructure of abuse also distracts attention from the arena where the decisive battle will actually be lost or won. The battle for the "hearts and minds" of the human population is a phrase lamed by overuse; but it's as relevant as ever. To believe otherwise is to fall victim to a utopian romanticism which misreads the realities of political power in the modern state.

On balance, then, the slaughter and abuse of our victims will probably be preventable only when a majority of the population in mainstream human society can be induced to accept that our present-day systematic abuse and killing of non-humans is morally wrong. For all its manifold failings, liberal capitalist democracy does offer the mechanisms to enforce majority-decisions when consent is obtained through the ballot box and its impending digital successors. True, much of our nominal democracy today is indeed a sham. Yet it is not a complete sham. If enough of the population oppose a ruling government, then the regime in question can be peacefully ousted. Likewise, if enough of the population come to recognise that institutionalised killing and abuse of non-humans is morally wrong, then such killing and abuse can be curtailed; and subsequently abolished in law. The full resources of the state can then be deployed to enforce that abolition.

Unlikely? Over the past hundred and fifty years, the state has steadily extended its quasi-monopoly of coercive acts in *human* society to an extent that would have once been unimaginable. For sure, violence today as practised by e.g. teachers on school-students, husbands on wives, and citizens against each other, still occurs. Yet it's vastly less common than it was in the past. It is increasingly taboo. Legal sanctions against interpersonal violence, and enforcement-mechanisms to prevent it, have steadily grown in depth, scope and effectiveness. Endorsing the liberal-democratic state's quasi-monopoly on violence, and calling for it to be *extended* rather than challenged, might sound a wildly paradoxical plea. It sounds even odder from a tender-minded radical who advocates a wholesale and nonviolent*revolution* in our behaviour to non-humans. Yet the machinery of the animal holocaust - organised as now in the final flourish of the Late-Darwinian Era on a scale and systematicity that dwarfs anything practised by our ancestors - is likely to be dismantled by essentially peaceful and legal means. Full-blown revolutions (as distinct from political coups) are rare even when the victims are human and can potentially fight back. When the oppressed are mute and helpless, the preconditions for an insurrection of the oppressed do not exist; and they never will.

So there is clearly a daunting struggle ahead. Life-stylism by itself is not remotely enough. Simply refusing to pay others to commit acts of violence on one's behalf is indeed important; for boycotting meat-products and their producers diminishes the financial incentives for killing and abuse. Yet the cultivation of personal purity - though commendable - can become a disastrous distraction. Such distraction occurs when an otherwise admirable desire to banish all

trace of *personal* complicity in animal-abuse eclipses the struggle to promote collective action against the institutional system of animal-exploitation as a whole. Animal-abuse itself needs to become a criminal offence. It can't be left as a matter of consumer choice or personal taste. When it gets phased out, it will be abolished "from above" as much as "from below". This can only happen, however, if most people - if necessary a bare plurality - can be persuaded that it is morally unacceptable for *anyone* to do it.

This transformation depends on inducing a fundamental shift in the beliefs and values of a majority of politically active adult humans. Or, much more optimistically phrased, it depends on extracting and making starkly explicit the full consequences of beliefs and values we already hold: namely the extension of the kind of love and privileges given to, *and genuinely deserved by*, Rover, the adored family pet, to the similar creatures we are paying to have abused and butchered. A whole range of life-forms typically treated as *objects* must come to be treated as fellow *subjects*.

Just how likely is this shift to occur? And by what means? Can present trends to vegetarianism, and increasingly veganism, be extrapolated deep into the next millennium?

To some extent, the sea-change in prospect is likely to be demographic and generational rather than the product of mid-life conversion experiences. The defensibility of animal-abuse, even under its innumerable euphemisms, tends to seem less "obvious" to younger people. The struggle that will be waged is both ideological and scientific. In the battle to win converts, the Net offers an immense opportunity to subvert the ideology of oppression. As global Net-use and digital convergence accelerate dizzyingly, and web-enabled devices promise to proliferate all over the globe, it's becoming clear that here is where the long-term ideological battle will be won or lost.

This rallying-cry might seem a naïve piece of Net evangelism. And trusting that the weak and the vulnerable might ever be protected by the powerful might seem naïve in the extreme. Yet once the incentive of self-interest has been stripped away, and genetic-engineering allows us to produce whatever food-products we like without causing death or suffering, then our argumentative blind-spot is likely to disappear. Invoking genetic-engineering as a solution to today's biggest source of systematic animal-abuse, namely the factory-farming of live animals for the purpose of eating their flesh, is likely to make most progressive radicals queasy. For unquestionably there is enormous scope for biotechnology to be perverted for purposes which have nothing to do with the global welfare of either humans or animals. Yet unless genetically-engineered test-tube meat delicacies can be mass-produced cost-effectively, we will have to rely exclusively on moral arguments against animal-abuse. This will mean a far longer delay before a liberation of the oppressed and immeasurably more bloodshed.

Isn't this plea for a futuristic cruelty-free diet a cop-out? Doesn't gesturing in the direction of future food technologies just enable their armchair advocates to live comfortable lives of genteel digital radicalism in the meantime, while our victims live miserable lives followed by gruesome deaths?

Possibly such advocacy *is* self-serving. One should never underestimate the human capacity for self-deception. Yet biotechnology offers the most effective *long-term* global strategy for success in purging the world of cruelty and pain. Simple cost-considerations are likely to make genetically-engineered single-cell protein food cheaper and healthier (no pesticide-, hormonal- and antibiotic- residues etc). Money and morals fused together make a potent combination.

Of course, most of us aren't genetic engineers. We can't grow mouthwatering steaks-in-vats ourselves, or synthesise drugs evoking the illicit tastes and textures of the depraved appetites of the past. So what can we do instead? Surely not just wait until some ill-defined

(bio-)technological determinism sweeps the old regime aside. Organising systematic ideological warfare on behalf of our victims using the new electronic media is going to be vital. Sound-and-video-footage of the kind that simply wouldn't be allowed on traditional TV must be smuggled out from the factory-farms and death-factories. It must be disseminated over imminent Web-TV to the widest possible audience. Admittedly, the publication of such horror-footage will provide morbid titillation to corrupted minds; but it will upset most people. It may even shock some of them into action - or abstinence. For if one had to watch the life and death of the creature the remains of whose body was sitting on one's plate, then one almost certainly would be too revolted to eat it. After all, at present that piece of meat seems so *innocuous*. Bad things often do.

Trying to counter the billions of dollars worth of propaganda currently pumping out the opposite message, namely the terrible myth that non-humans are merely objects to be used, produced and eaten, might seem an unequal struggle. It is. Yet puncturing the tissue of deceptions on which the reigning speciesist ideology rests is potentially feasible. It can be done if the tender-minded activists who most strongly oppose animal-abuse can conquer their visceral technophobia.

One final plea can be entered here. It concerns the grey area where life-stylism gets converted into something more powerful than the force of a good example. It's important that meat-eating should start to become *socially unacceptable*. Only after this is it likely to be criminalised if practised on the corpses of once-sentient animals rather than on tasty, genetically-engineered vat-proteins.

Is widespread social stigmatisation of eating dead animals really a serious prospect within the foreseeable future? Perhaps surprisingly, yes. Couched *in the abstract*, the infliction of needless suffering on other beings is acknowledged by most people to be morally wrong. We need merely to make the connection between what we're doing and the suffering our actions cause at several removes: supermarkets today are cunningly designed to evoke warmth and friendliness, not sinister graveyards. By way of context, over the past twenty years or so *overt* racism has become socially taboo within more and more parts of society. So have the more virulent forms of, say, sexism and homophobia. Violence against children, too, a habit universally recognised by child-abuse experts as the cause of potentially long-lasting psychological damage, is heading in the same direction; though likewise in practice there is a fearful way to go. One may predict - as well as advocate - that some time over the next few decades, a similar growth of stigmatisation will attach to eating traditional meat-products derived from "livestock". This is a process that has already halfheartedly begun in progressive circles. Even on a relatively timid extrapolation of this trend, our descendants may view meat-eating with the revulsion and incomprehension we reserve for cannibalism or genocide against humans; and in particular the Nazi Holocaust. So if there is any sense at all to the notion of moral progress, it would be useful to try and imagine why posterity might see us in such an ugly light. DeGrazia gives us some telling clues. *Taking Animals Seriously* is an admirable, original and important book. Yet in the end, DeGrazia is too much an ideological prisoner of the old DNA regime to contemplate its total wipeout.

It would be nice to end on an uplifting note. Such uplift would also be misleading and facile. Right now as you read these words, mass-killings and systematic animal-abuse continue at unimaginable levels. We are quite literally paying its perpetrators to kill their victims on our behalf. A sense of guilt and horror, not complacency, is needed to stop us. At the very least, if one is looking for a postscript to the call to take animals seriously, then it might well be: "If you think it's murder, act like it."