THE HEDONISTIC IMPERATIVE

By DAVID PEARCE

[*HI was written in 1995. Only the later Replies to Objections, added sequentially in Chapter Four, are more recent. For a contemporary outline, see The Abolitionist Project (2007, http://www.abolitionist.com/) and Utopian Neuroscience (2008, http://www.superhappiness.com/). Chapter Two can safely be skipped or aggressively skimmed even by the analytic philosophers for whom it was primarily intended. It contains a defence of HI on the basis of, first, practical means-ends rationality and, secondly, ethical negative utilitarianism. The* instrumental *case from means-ends rationality derives from the broad applicability of psychological hedonism. This isn't here construed as a universal law. It's just a trite everyday rule of thumb: we spend a lot of time trying to make ourselves happy. Often we fail. HI achieves what we're striving for with unique efficiency and success. The* ethical *utilitarian case for HI, on the other hand, rests partly on a conception of how morality can be naturalised consistently with a recognisably scientific account of the nature of the world. Value is here construed as a distinctive - and biologically maximisable - mode of experience. Its subjective texture is coded by a particular type of biomolecular architecture. That architecture can be enriched and extended. Positive value can be maximised. Negative value can eventually be eliminated. Thus HI, it will be claimed, amounts to rather more than one individual's quirky conjectures and value-judgements. The biological program is also our natural destiny. The coming of the pain-free post-Darwinian Era will mark both a major transition in the evolution of life and the moral foundation of any future civilisation.*]

A B S T R A C T

This manifesto outlines a strategy to eradicate suffering in all sentient life. The abolitionist project is ambitious, implausible, but technically feasible. It is defended here on ethical utilitarian grounds. Genetic engineering and nanotechnology allow *Homo sapiens* to discard the legacy-wetware of our evolutionary past. Our post-human successors will rewrite the vertebrate genome, redesign the global ecosystem, and abolish suffering throughout the living world.

Why does suffering exist? The metabolic pathways of pain and malaise evolved only because they served the inclusive fitness of our genes in the ancestral environment. Their ugliness can be replaced by a new motivational system based entirely on gradients of well-being. Life-long happiness of an intensity now physiologically unimaginable can become the heritable norm of mental health. A sketch is offered of when, and why, this major evolutionary transition in the history of life is likely to occur. Possible objections, both practical and moral, are raised and then rebutted.

Contemporary images of opiate-addled junkies, and the lever-pressing frenzies of intra-cranially self-stimulating rats, are deceptive. Such stereotypes stigmatise, and

falsely discredit, the only remedy for the world's horrors and everyday discontents that is biologically realistic. For it is misleading to contrast social and intellectual development with perpetual happiness. There need be no such trade-off. Thus states of "dopamine-overdrive" can actually *enhance* exploratory and goal-directed activity. Hyper-dopaminergic states can also increase the *range* and diversity of actions an organism finds rewarding. Our descendants may live in a civilisation of serenely well-motivated "high-achievers", animated by gradients of bliss. Their productivity may far eclipse our own.

Two hundred years ago, before the development of potent synthetic pain-killers or surgical anaesthetics, the notion that "physical" pain could be banished from most people's lives would have seemed no less bizarre. Most of us in the developed world now take its daily absence for granted. The prospect that what we describe as "mental" pain, too, could one day be superseded is equally counter-intuitive. The technical option of its abolition turns its deliberate retention into an issue of political policy and ethical choice.

TABLE OF CONTENTS

WHY?

WHEN?

OBJECTIONS

Chapter 1: HOW?

"God's in His Heaven -
All's right with the world!"
(Robert Browning)

1.0 Sabotage at the Mill

To escape from the hedonic treadmill we must first sabotage a small but vicious set of negative feedback mechanisms. These are genetically coded into the mind/brain. Recreational drugs of abuse do not transcend or subvert such mechanisms. On the contrary, they actually bring them into play with a vengeance. Today's quick-and-dirty euphoriants are nonetheless instructive. They give us a tantalising glimpse of what humanity's natural state of consciousness could become if several ugly neural metabolic pathways were inhibited or eliminated.

A better clue to organic life's emotional future dates from the early 1950s. The unlikely guinea-pigs were veterans at a U.S. tuberculosis sanatorium. Residents prescribed the MAO-inhibiting drug iproniazid were not merely cured of their tuberculosis. After a few weeks of treatment, many of them started to feel exceptionally happy. Doctors described their patients, rather over-colourfully perhaps, as "dancing in the aisles". For the most part, the veterans had not previously been clinically depressed, as distinct from rather crotchety. Nor was their new-found euphoria simply an understandable reaction to restored good health. Moreover, in contrast to most recreational drugs, tolerance to the MAO-inhibitor's mood-brightening side-effect, and the consequent danger of uncontrolled dose-escalation, didn't set in. Instead, it transpires that MAO-inhibitors as a class can induce a benign, long-term re-regulation of several families of nerve-cell receptor proteins involved in making us happy or sad. Quite by accident, modern medicine had stumbled on the sustainably mood-lifting properties of a remarkable and diverse category of drugs, the monoamine oxidase inhibitors.

Monoamine oxidase has two main types, uninformatively labelled A and B. MAO is an enzyme responsible for the deamination of monoamine neurotransmitters such as dopamine, noradrenaline and serotonin. It also deaminates trace amines such as phenylethylamine, found in chocolate and released when one is in love. MAO isoenzyme-A deaminates serotonin, norepinephrine and, to a lesser extent, dopamine. Isoenzyme-B breaks down dopamine and phenylethylamine. The action of monoamine neurotransmitters on the post-synaptic receptors, and the post-transduction intracellular cascade they induce, plays a vital role in mediating mood and emotion. Depletion of monoamines in the synaptic vesicles e.g. by the anti-hypertensive drug reserpine, can sometimes precipitate severe and even life-threatening depression. Elevated levels of dopamine, on the other hand, are associated with (hypo-)manic euphoria.

By modulating the synaptic availability, and consequent receptor re-regulation, of simple neurotransmitters on a long-term basis, the MAO-inhibitors were to serve as the first of a disparate group of drugs uninvitingly categorised as "antidepressants". Some of today's mediocre crop of licensed products, such as the tricyclics, are in general unrewarding to people who aren't rated clinically depressed. They tend to be sedating. Their action dulls, however mildly, the intellect and sensibility. Most traditional therapeutic agents - at least until the development of (relatively) selective serotonin re-uptake blockers such as fluoxetine (Prozac) and noradrenaline reuptake blockers such as reboxetine - are "dirty" and unselective drugs. They have lots of troublesome side-effects. They frequently flatten rather than deepen the emotions. Several brands, such as

the older, unselective and irreversible MAO-inhibitors, are potentially dangerous if taken in the absence of rigorous dietary restrictions. All of them, thanks to the puritanical ethos of the medical establishment, have been tested and brought to market with the deliberate additional aim of *not* inducing a euphoric sense of well-being ("abuse-potential") in the user. Most contemporary "antidepressants" only modestly outperform a placebo in well-controlled clinical trials.

It is the twenty first century's successors to these unpromising-sounding drugs, however, and not today's fast-acting recreational euphoriants, that promise to deliver the world's supposedly "euthymic" population from the sick psycho-chemical ghetto bequeathed by our genetic past. Potent, long-acting mood-brighteners - but not clinical "psychic anaesthetisers" or "quick-hit" street-drugs - will serve as a life-enriching stop-gap until radical gene-therapies enable us to knock out the Darwinian pathologies of consciousness altogether. Time-delayed designer euphoriants will foreshadow an extended product-line of innovative treatments for all kinds of malaise. Collectively, such interventions will cure what post-human posterity will recognise as a gene-driven spectrum of psychiatric disorders characteristic of Darwinian life. A lot of the time at present, we just don't - and can't - conceptualise the full extent of how unwell we are. For there are powerful arguments to suggest that everyday consciousness, insofar as it is not transcendentally wonderful, is symptomatic of profound psychological ill-health.

This possibility is not widely acknowledged in public today. Mental illness still carries a stigma. "Of-course-I'm-all-right. There's nothing wrong with me!", one may sometimes snappishly be told. To be depressive is to be fitness-impaired, low-status, a poor choice of mate, and generally uncool. So there are self-protective defence- and denial-mechanisms, as well as a plain failure of the imagination, at work.

Defensiveness and denial won't be needed for ever. A few generations hence, the intoxicating joy of normal post-Darwinian life will be genetically pre-programmed. A reproductive revolution of "designer babies" will hardwire happiness from the womb. Psychoactive drugs may become redundant, or rather tools for consciousness research and life-enrichment rather than self-medication. For pure well-being can potentially become a deep and natural presupposition of everyday life. Undiluted existential happiness will infuse every second of waking and dreaming existence; and pervade every aspect of one's body and psyche. Sadly, the sort of germ-line gene-therapy needed to achieve gradients of lifelong, high-functioning euphoria for everyone who thinks they can handle it is still some way off. In the transitional era before global paradise-engineering unfolds, chemical mood-uplifters will be essential too.

1.1 The Biological Program

Grand meta-narratives currently aren't very fashionable. History can indeed seem like one damn thing after another. The nearest we get these days to some kind of plot or story about where life on earth is heading usually adds up to some simple-minded technological determinism. Nevertheless, a sketch of one possible route by which all sources of negative value will be purged from the world is set out below. Other biological strategies for Cosmic Value-Maximisation - or simply making everyone a great deal happier - are in prospect too. Details and variations matter. Every family of options for naturalising heaven-on-earth needs to be exhaustively researched - and not just idly philosophised about. Yet it is vital to distinguish the overall goal of abolishing suffering from our first faltering blueprints of how the abolitionist project should be implemented. The technical shortcomings of anything proposed here should not be allowed to taint the overall strategy of the abolitionist project itself.

This particular biological program, at least, is inspired by an almost desperate sense of moral *urgency*, not gung-ho technophilia. It's not "hedonistic" in the popular sense of term. For it's worth pausing and trying to practise, quite literally, a few minutes of systematic empathy. Atrocious, agonising things are happening to people like you, me and our loved ones right now. The full horror of some sorts of suffering is literally unspeakable and unimaginably dreadful. Under a Darwinian regime of "natural" reproduction, truly horrible experiences - as well as endemic low-grade malaise - are both commonplace and inevitable. In Chapter Two, the moral case will be argued that this nastiness should be stopped. Since 'ought' implies 'can', however, it must first be established that scrapping unpleasant experience really is a biologically feasible option. It will be argued that the lesson of intracranial self-stimulation studies - despite their lamentable contemporary image - is yes. It is harder to establish that life-long, *intellectually discerning* bliss is feasible, either via rationally designed drugs or gene therapy or both. But from an information-theoretic perspective, what counts is not our absolute location on the pleasure-pain axis, but that we are "informationally sensitive" to fitness-relevant changes in our internal and external environment. Gradients of bliss can suffice both to motivate us and offer a rich network of feedback mechanisms; so alas today do gradients of Darwinian discontent.

The blueprint set out below outlines only a cartoonish prototype of a mature post-Darwinian paradise. Its sketch of likely future neuro-scientific breakthroughs may well be wrong both in its few specifics and its projected time-scales. Experts in the relevant specialist fields will doubtless wince, at least in places. For *The Hedonistic Imperative* consists in a hand-waving, cross-disciplinary romp through dauntingly complex specialist topics. Inevitably, some of the pop neuroscience is simplistic to the point of parody. Eyebrows should be raised, too, at the dogmatic brevity with which various philosophical problems deserving book-length treatment are dispatched in a single sentence. The multitude of practical, medico-legal and socio-political problems which fulfilling our neurochemical Manifest Destiny will entail are largely passed over as well.

These caveats are important. Yet leaving them aside, the biological program may be divided, somewhat arbitrarily, into three stages. They are here ranked in order of difficulty. Luckily, the stages happen to coincide in relative ethical importance, since crude harm-reduction, cruelty-prevention and pain-abolition are easier to accomplish than refining the architectural subtleties of paradise. Less happily, any biochemical description of the mechanics of the sublime just travesties the nature of the experience itself. The sub-academese prose below unavoidably debases what it aims to evoke. This is because of the contaminated associations of any terms associated with drug-abuse, genetic engineering, eugenics, or even the emotionally frigid atmosphere of the laboratory. Our present perspective on utopian biopsychiatry is jaundiced. For our education system virtually ignores the neurobiological foundations of all emotional life. Happily, that system also provides the formal tools for us to describe and escape from our predicament.

What is really needed, above and beyond mere chemical formulae, is a new network of concepts - a user's guide to map out the magically alien realms of consciousness ahead of us. Yet by the time such tools can be developed as new state-spaces of experience are accessed, the revolutionary conceptual scheme they embody will be less urgently needed. One day, we may have thoughts like sunsets! Their brilliance will replace the elusive and phenomenologically thin series of sad little cognitive tickles which we (apparently) shuffle around and via which this manifesto is written and read. In the meantime, the impersonal vocabulary of chemistry and molecular biology is all we can rely on for communicating how to get things done. An earthly paradise can be achieved only by the profane application of science. World-wide mental superhealth won't be achieved via the edifying discourses of religion or magic.

1.2 Pumping Up The Volume

One crude but effective ingredient of the initial stage of the biological program will involve modifying the meso(cortico-)limbic dopamine system. Controversially, and oversimplifying a little since dopamine is not itself a magic "pleasure molecule", the mesolimbic reward pathways play an intimate role in the final common pathway for pleasure in the brain. Neuronal dopamine-release may be elicited "naturally" via biochemical transduction-mechanisms. It is usually triggered by adaptive environmental stimuli. On the other hand, dopamine-release may also be induced more directly via the use of recreational drugs. The "rush" of crack cocaine, for instance, falsely signals a huge Darwinian fitness benefit. Either way, although the central nervous system has tens of billions of cells, its mesolimbic wellspring of pleasure, motivation and libido has only some 30-40 thousand neurons; and clearly this isn't nearly enough.

The axons and dendrites of mesolimbic dopaminergic neurons innervate the higher cortical regions of the brain. They thereby help mediate the genetically adaptive "encephalisation of emotion". This neat little trick has served our DNA, but frequently not us, fiendishly well. Emotional encephalisation convinces its victims that happiness is inseparable from presence or absence of variously innervated types of intentional object. We are happy or sad 'about' things. Entirely non-coincidentally, the realisation of our most emotionally charged types of intentional object tends to promote the inclusive fitness of our genes. Crudely, we like most what's good for them.

Unfortunately, they don't care about us. Our genes don't look after their vehicles for very long. In adult life, dopaminergic neurons die off at a rate of over 10% per decade. Their death ensures that senescence is marked by a decline in drive, libido, pleasure and the intensity of experience itself. Even in one's youth, the fullest and most beautiful scope for expression of the dopaminergic pleasure-cells is continually frustrated by inhibitory feedback. This derives both from the cells' own pre-synaptic autoreceptors and the processes of other, often less benign, neurons that synapse upon them.

Thus what must be included in any program of systematic life-enrichment is a strategy of at once multiplying the numbers of, and selectively reducing feedback inhibition on, mesolimbic dopamine cells. Targeting the medium spiny neurons of the rostral shell of the nucleus accumbens is critical. Achieving a modest initial hundredfold, say, enrichment of an organism's capacity for well-being is not, needless to add, simply a matter of genetically switching on an uncontrolled proliferation of dopaminergic neurons; though it has to be said that, as causes of death go, a tumour of the pleasure cells has a certain whimsical appeal. Nor, of course, does a regimen of sustained pleasure-amplification simply entail enhancing the levels of dopamine in the synapses. Excessive post-synaptic stimulation of particular dopamine receptor sub-types is implicated in, for instance, the florid symptoms of schizophrenia. Dopamine overdrive also marks the psychotic excesses of that ultimate egoist, the crack addict. So crude monotherapy surely won't do the job alone.

1.3 The Civilising Neurotransmitter

There is a more promising twin-track approach. This consists of boosting sub-types of both dopaminergic and serotonergic function.

Serotonin has been described as the "civilising neurotransmitter". Such a label is a useful piece of mental shorthand. It's still worth noting that even this simple monoamine

has fifteen or more functionally distinct receptor sub-types. Serotonergic dysfunction is associated with irritability, explosive anger, violence, sociopathy, and suicide. Conversely, the extraordinarily deep sense of love, trust and empathy inspired by "the penicillin of the soul", MDMA, is due primarily to the massive release of serotonin which its use provokes. It causes only a modest release of dopamine. Both dopamine- and serotonin-release are needed for the inhibitory effects of MDMA on glutamate-evoked neuronal excitability in the nucleus accumbens to take its full magical effect. In any event, the result of casually popping a pill can be a life-defining revelation. The trouble today is that the magic doesn't last.

There's no good reason why it shouldn't. In the new reproductive era of "designer babies" ahead, neurobehavioural systems that evolved to maximise Darwinian fitness of hominids on the African savannah can be redesigned to maximise personal well-being. A new kind of selection pressure comes into play when allelic combinations are deliberately chosen for a new child by its prospective parents *in anticipation* of their likely effects. Until invincible well-being can be genetically preprogrammed, however, it would be eminently sensible to develop a delayed-action, non-neurotoxic drug or cocktail-mix of sustainable mood-brighteners. This sort of designer elixir could make us all very happy and revolutionise our archaic conception of mental health. Day-to-day life in drug-assisted Eden can blend, if we so choose, the most exalted, life-loving euphoria of a potent dopamine or mu opioid agonist with the serene and mystical love of an 'empathogen' or 'entactogen' such as MDMA ("ecstasy"). States of incisive, goal-directed thought can co-exist with a profound love for our fellow beings. If we want, we can make such states biologically natural; and eventually innate. There are unimaginably good times ahead.

As hedonic engineering develops into a mature biomedical discipline, the generic modes of paradise we opt for can be genetically pre-coded. Native-born ecstatics will flourish. All the wonderful models of mental superhealth discussed in this section of HI may come to be viewed as simple-minded prototypes. The innovative, high-specification bio-heavens beyond will be far richer. We lack the semantic competence to talk about them sensibly. Yet however inelegantly our goal may be accomplished at first, the ultimate strategic objective should be the neurochemical precision-engineering of happiness for every sentient organism on the planet.

Sounds flaky? Yes, but then so, originally, has almost every radical reform movement in history (including, of course, the genuinely flaky ones.)

1.4 The Cardinal Importance of Delayed Gratification

Eventually, well-being will be part of our very nature. A robust network of homeostatic mechanisms will ensure all hereditary ecstatics have gene-coded hedonic set-points way beyond today's puny maxima. In the Transitional Era, however, the widespread use of mind-healing drugs will in practice be unavoidable. Gene-therapy is still in its infancy; and germ-line clinical trials are time-consuming in humans. So crucially, the medically and socially responsible emphasis of the pharmacological arm of the biological transition strategy must be on the (relatively) *long-term* structural and functional effects in nervous tissue which a delayed-reward euphoriant-mix will induce in the individual mind/brain. Fast-acting recreational highs are a snare and a delusion. We must master - and educate our children in - the pharmacological equivalent of the principle of deferred gratification.

The delay in therapeutic benefit stemming from gene-triggered receptor re-regulation can actually be very useful. Not merely is the development of tolerance diminished. Uncontrolled and potentially noxious bingeing on a psychoactive drug occurs when there is minimal delay between ingestion and reward. By contrast, the anticipated gene-switched, up- or down-regulation of the pre- and post-synaptic neuronal receptors in a regimen of sustainable mood-enhancement will generally take up to several weeks to complete. Fortunately, enhancing serotonin function tends to increase patience and impulse-control as well as mood.

Perhaps a comparison with tobacco-smoking can be of use here. In its present setting, nicotine is so addictive, not because of the quite minimal "high" it induces, but because of the sheer speed of onset of its intrinsically mild hit due to the customary delivery mechanism. The "reward" comes about seven seconds after inhalation. If the whole-body orgasmic rush of even crack-cocaine were delayed for ten days or so after its consumption, then the drug would be far less of a social and medical problem than it is at present. Tragically, most of its current users seem unacquainted with, or have long since forgotten, the concept of delayed reward. They might now be unwilling to wait nearly so long.

1.5 The Molecular Genetics Of Paradise

Strategic, species-wide pharmacotherapy of the kind advocated above can be complemented, and synergistically allied, with genetic engineering as it matures from mere genetic tinkering. Gene therapy will be targeted both on somatic cells and, with even greater forethought, the germ-line. If cunningly applied, a combination of the cellular enlargement of the meso-limbic dopamine system, selectively enhanced metabolic function of key intra-cellular sub-types of opioidergic and serotonergic pathway, and the disablement of several countervailing inhibitory feedback processes will put in place the biomolecular architecture for a major transition in human evolution - and life itself.

The re-engineering of bits of psycho-neural circuitry sketched above may, it is true, seem somewhat ambitious. Perhaps it sounds impossibly futuristic. Comparatively, however, these techniques amount to a primitively inept form of piecemeal tinkering compared to the revolutionary redesign of the mind/brain likely to be undertaken in millennia to come.

For it won't just be the quality and quantity of consciousness in the world which will be transformed in the early stages of the post-Darwinian Transition. As humanity emerges from the psychochemical Dark Ages, enriched dopaminergic function in particular will sharpen the sheer intensity of every moment of conscious existence. For a generation whose lifetimes span both modes of awareness, it will be as if they had just woken up after sleep-walking through life in a twilit stupor. Thereafter their former mundane and minimal existence may be recalled only as some kind of zombified trance-state. Our own "ordinary" consciousness may be unmasked as a shallow and uninteresting vehicle of malaise whose properties we were physiologically incapable of recognising 'from the inside'. At present, however, we lack the neural substrates of a capacity to set archaic consciousness in a pre- and post-Darwinian context. Or as Einstein says: "*What does the fish know of the sea in which it swims?*"

Other neurohormones, transcription factors, opioids, tyrosine-hydroxylase activators, oxytocin-releasers, receptor density-regulators, intra-cellular second- and

third-messengers, phosphorylated proteins, and genetic repressors and promoters which are implicated in the modulation of mood, emotional tone and psychophysical pain will be reconfigured too as the biological program unfolds. The details are messy and complicated. Naturally, our neurotransmitter systems finely interlock. They can be treated in isolation only conceptually and for purposes of expository convenience. They form a complex and delicate interplay of feedback loops that defies easy simplification and synopsis. In centuries to follow, however, they will be collectively enlisted to re-work the texture of experience. Our happiness will be chemically and genetically enhanced with ever greater artistry and finesse. Conversely, several vicious triggers of extraordinary nastiness (e.g. bradykinin, nociceptin, substance P) will be banished from the sensorium, one trusts for ever.

1.6 The Re-encephalisation Of Emotion

These procedures will lay the hedonic foundations for a dizzyingly exalted ground-state of conscious existence. The most pressing question to examine next is what will - and what should - be done with it? How, and why, should emotion be encephalised in an era when intentionality is no longer tied to furthering the inclusive fitness of self-replicating DNA in our ancestral environment? What's worth being happy "about"?

For the real intellectual challenge won't lie ultimately in sheer happiness-maximisation. After all, if eternal bliss were the sole objective of paradise-engineering, then a rat with electrodes fixed in its pleasure-centres already points the way forward. In fact, our descendants may find generating generic states of life-long happiness *per se* trivially easy. Most of us, however, are intellectually quite snobbish. We don't want our emotions de-encephalised. We like good moods, but anything resembling the prospect of a perpetual orgasmic frenzy of delight stirs more ambivalent feelings. The limbic innervation of the neocortex has been so adaptive because it allows sophisticated genetic vehicles like us to feel some intentional objects are inherently good or bad. We want to feel that we are happy for good *reasons* - genetically self-serving as they may so often be.

We'll soon be in a position to de-fang this dangerous tendency altogether. But we won't want to abolish it. In generations to come, a primary focus of neuroscientific mind-making will be on remapping the axonal and dendritic arborisation of the neo-cortex which makes the rationalisation of emotion possible. The aim of this rational redesign can be to bootstrap our way into fulfilling our second-order desires ("desires about desires") for who and what we want to become. What we will ultimately turn into is hard to imagine. One may predict merely that it will be utterly sublime.

Using biotechnology to select and fine-tune a post-Darwinian personality will partly depend on individual taste. One's choice of identity even in paradise will still be tempered by genetic biases, ancient cultural stereotypes, and the latest vagaries of fashion. The lure of hot-button super-normal stimuli will at first be very potent. Yet we may also be enchanted by ideas and modes of experience that today haven't even been conceptualised. Potentially, there are far more things to be happy "about" than we can possibly grasp.

On a societal level, some form of neuro-architectural planning permission will presumably still be needed for the purposes of orchestrating the multiple microcosms as each designer-heaven takes shape. Yet harmonisation should be more readily accomplished when people are already blissfully and empathetically happy - "all loved

up". Neurologically, in fact, there is nothing to stop co-operating with others from being a source of rapturous joy; as alas it isn't always today. When life isn't perceived as an approximation to a zero-sum game, social existence is going to be far easier to co-ordinate.

Initially, it may be tempting for newly-enlightened ecstatics to seek the idealised realisation of purely traditional objects of delight. Effectively, we'll be able to have anything we've always wanted and more. This includes enjoying the biomolecular substrates of an unprecedentedly vivid sense of reality, a perpetually enriched feeling of meaningfulness and significance, a sense of heightened authenticity, and never-ending raw-edged excitement - or intense serenity and spiritual peace. In these early days, subjects may find the idea of fulfilling older conceptions of the good life a reassuring prospect. Prior to their own personal transition to heavenly superhealth, any paradoxical trepidation coming from candidates for hedonic enrichment should be laid to rest by the following reflection. Nothing we have previously enjoyed in the old Darwinian era will afterwards be unavailable or any less satisfying than before. In fact, we may be motivated to pursue old goals with far greater gusto once weakness of will becomes just an evolutionary curiosity. For weak will-power caused by dopamine hypo-function is one of those neurological deficiencies which effort alone can't overcome. Happily, in Paradise the frailest spirit can move mountains.

1.7 How Could Anything Be So Good?

Perhaps a few examples of early post-Darwinian life are in order.

The Nature-lover, for instance, will be able to contemplate with awe-struck reverence scenes of overpowering sublimity eclipsing the superficial prettiness on offer before.

A musician may wish that those of his functional modules which mediate musical appreciation should receive especially rich innervation from his freshly amped-up pleasure system. (S)he might then hear, and have the chance to play, music more exhilarating and numinously beautiful than his or her ancestors ever dreamed of; the celestial music of the spheres heard by privileged medieval mystics will be as a child's toy tin-whistle in comparison.

The sensualist will discover that what had previously passed for passionate sex had been merely a mildly agreeable piece of foreplay. Erotic pleasure of an intoxicating intensity that mortal flesh has never known will thereafter be enjoyable with a whole gamut of friends and lovers. This will be possible because jealousy, already transiently eliminable today under the influence of various serotonin-releasing agents, is not the sort of gene-inspired perversion of consciousness likely to be judged worthy of conservation in the new era.

A painter or connoisseur of the visual arts will be able to behold the secular equivalent of the beatific vision in a million different guises, each of indescribable glory. The toy-town lexical tokens we permute today will by then be an archaic residue of little use in evoking their majesty. As language evolves to reflect and navigate ever more exalted planes of being, fresh taxonomies of pleasure-concepts will be pioneered to help define newly-discovered modes of awareness.

As an exercise, the reader may care briefly to summon up the most delightful fantasy (s)he can personally conceive. Agreeable as this may be, states of divine happiness orders of magnitude more beautiful than anything the contemporary mind can access will pervade the very fabric of reality in generations to come. Even the most virile of imaginations can apprehend in only the barest and formal sense the ravishing splendour that lies ahead.

1.8 All We Need Is Love?

Still in a personal vein, fragile self-esteem and shaky self-images will be beautified and recrystallised afresh. For the first time in their lives, in many cases, human beings will be able wholeheartedly to love both themselves and their own bodily self-images. Bruised and mutilated egos can thus be strengthened. They can be regenerated anew from the wreckage of the Darwinian past.

Love will take on new aspects and incarnations too. For instance, we will be able, not just to love everyone, but to be perpetually *in* love with everyone, as well; and perhaps we'll be far more *worth* loving than the corrupted minds our genes program today. It's been said that when in love we find it astonishing that it is possible to love someone else so much, because normally we love each other so little. This indifference, or at best mere diffuse benevolence, to the rest of the population is easily taken for granted amid the harsh social realities of competitive consumer capitalism. Yet our deficiencies in love are only another grim manifestation of selfish (in the technical sense) DNA. If humans had collectively shared the greater degree of genetic relatedness common to many of the social insects (haplodiploidy), then we might already have been "naturally" able to love each other with greater enthusiasm. Sociobiology, and its offspring evolutionary psychology, explains our relative coldness of heart.

Happily, in future it will be possible to mimic, and then magnify out of all recognition, the kind of altruistic devotion to each other which might have arisen if were we all 100% genetically-related clones. We'll all be able to love each other to bits. A delicious cocktail mix of oxytocin, phenethylamines and mu receptor-selective opioids - or potent god's-own wonderbrews not yet genetically-coded - can be automatically triggered whenever anyone one knows is present or recollected. Darwinian man, by contrast, will be seen as a mean-minded crypto-psychopath. Our successors will be far kinder. They'll combine absolute, unconditional and uninhibited love for each other with a celebration of the diversity of genes and cultures. At present this prospect seems some way off.

Another aspect of post-Transition love may be found even more surprising. Individual personal relationships may at last be bonded truly securely, should we so desire. Throughout the ages, dreadful pain has been caused by the soul-destroying cruelties of traditional modes of love. We acknowledge, in the main, that we hurt the most those we love. Yet we often simply can't stop ourselves from doing so. Before very long, if we really care enough, we'll actually be able to do something about it. Whatever their proximate causes, the distal origins of so many relationship break-ups lie, once again, in the competing interests of rival coalitions of genes. Just to take one example, two lovers, perhaps, who years before professed they would rather die than hurt each other, later part in tears and acrimony. The woman may find that with the decline in her reproductive potential over time she is no longer sexually attractive to the man who pledged his undying love. Her partner, quite possibly hating himself for his treachery, finds himself deserting her and their teenage offspring for a younger, sexier woman, and

then fathers another family. Lives are destroyed; inclusive genetic fitness is served. Nature is barbarous and futile beyond belief.

After the Transition, on the other hand, one will be able to love somebody more passionately than ever before. In the post-Darwinian era, one will be safe in the knowledge that one will never hurt them, nor be hurt by them in turn. True love really can last forever, though responsible couples should take precautions. If one desires a particular relationship to remain uniquely and enduringly special, then the mutually co-ordinated design of each other's neural weight spaces can ensure that a distinctively hill-topped plateau in the new hedonic landscape structurally guarantees that each other's presence is always uniquely fulfilling. Choosing how big a hit we get off each other's presence is not an exact science today. Of course, it is possible that, generations hence, exclusionary pair bonding may seem a quaint anachronism. It may be regarded as just one more vestige of the genetic past which is fated one day to pass away. The example above is recounted to show only how ill-defined worries that anything precious one wants to save will be somehow sacrificed in the post-Transition epoch can be discounted. We've nothing to lose.

1.9 The Taste Of Depravity

Now before considering the prospects for the more distant future of affective states in the universe, the status of non-human animals must be addressed. This is because most of the world's suffering is undergone by members of other species. A convergence of evidence suggests that the nature and relative extent of organic life's biological capacity to suffer is mediated by key neuronal firing frequencies; cellular, synaptic and receptor densities; and a distinctive neurochemical and functional architecture of the central nervous system. Pain is *not* rooted in the unique human linguistic capacity for a generative syntax.

Humanity often behaves as though it were. For we presently keep hundreds of millions of other sentient beings in unimaginably frightful conditions. We do so for no better reason than to satisfy our culinary tastes. It has aptly been remarked that if animals had a conception of the Devil, he would surely have human form. Alas this is no mere rhetorical conceit. Contemporary humans deliberately incarcerate and butcher our fellow creatures in a vast, state-sanctioned apparatus of concentration and extermination camps. They are run with mechanised horror for commercial profit. In retrospect, our descendants may view them as a defining feature of our age in a way akin to our own conception of the Third Reich. Analogously, their sheer viciousness and even existence is usually camouflaged behind a morass of bland euphemism. Fortunately for our peace of mind, at least, we find it hard properly to conceive of what we're being spared. Conditions inside the camps and factories are frequently so gruesome that members of the public have to be barred from watching the atrocities that go on inside them.

For the most part, however, we are willing accomplices in our own ignorance. By our purchases we pay others to commit acts of extreme violence which might otherwise upset our squeamish sensibilities. Ironically, anybody who practises, or connives in, the maltreatment of a helpless and undeveloped infant of our own species is likely to be demonised and reviled. Ordinary decent people will find it "inconceivable" how such an "inhuman" monster could cause such suffering to the young, innocent and helpless. So (s)he will be prosecuted and locked up.

What we are doing in the death-factories is so vile that a few lines of text can scarcely even hint at its ghastliness. Nevertheless, we are so inured to the notion of exploiting and killing other sentient beings to titillate our palates that many otherwise "sophisticated" people will find the starkness of expression of these paragraphs somehow sensationalistic; or perhaps "emotive", as if the reality of such suffering could properly be otherwise.

Caring about the plight of the non-human victims of our actions is not a case of sentimental bunny-hugging nor of child-like anthropomorphism. Nor is it a matter of caring more about animals than humans; nor even, as is sometimes suggested with all appearance of seriousness, outright misanthropy. "Tender-minded" people who worry about the torture of non-humans are on balance temperamentally more inclined to act in an effort to minimise human suffering too. Such contrasts and false antitheses are in any case unhelpful. Simply by abstaining from eating meat, for instance, one can still spend just as much time campaigning for exclusively human causes as one did as a practising meat-eater.

There is one real glimmer of hope amid the ongoing carnage. Within the next hundred years or so, and possibly sooner, biotechnology will enable the human species cost-effectively to mass-produce edible cellular protein, and indeed all forms of food, of a flavour and texture indistinguishable from, or tastier than, the sanitised animal products we now eat. As our palates become satisfied by other means, the moral arguments for animal rights will start to seem overwhelmingly compelling. The Western(ised) planetary elite will finally start to award the sentient fellow creatures we torture and kill a moral status akin to human infants and toddlers. Veganism, though not in quite the contemporary sense, will become the global norm. Thanks to genetic engineering, the huge reduction in gratuitous suffering forecast here is likely to take place even if none of the other predictions of HI are borne out. If they are, then the humblest snack will taste more delicious than the ambrosial food of the gods. Today's gourmets might as well be feeding on greasy chips.

Much more seriously, in those traditional eco-systems that we chose to retain, millions of non-human animals will continue periodically to starve, die horribly of thirst and disease, or even get eaten alive. This is commonly viewed as "natural" and hence basically OK. It would indeed be comforting to think that in some sense this ongoing animal holocaust doesn't matter too much. We often find it convenient to act as though the capacity to suffer were somehow inseparably bound up with linguistic ability or ratiocinative prowess. Yet there is absolutely no evidence that this is the case, and a great deal that it isn't.

The functional regions of the brain which subserve physical agony, the "pain centres", and the mainly limbic substrates of emotion, appear in phylogenetic terms to be remarkably constant in the vertebrate line. The neural pathways involving serotonin, the periaquaductal grey matter, bradykinin, dynorphin, ATP receptors, the major opioid families, substance P etc all existed long before hominids walked the earth. Not merely is the biochemistry of suffering disturbingly similar where not effectively type-identical across a wide spectrum of vertebrate (and even some invertebrate) species. It is at least possible that members of any species whose members have more pain cells exhibiting greater synaptic density than humans sometimes suffer more atrociously than we do, whatever their notional "intelligence". As a utilitarian [technically, an ethical negative-utilitarian - *see below*], I would have to say, counter-intuitively, that were this to be the case, then such "hyperalgesic" life-forms would intrinsically matter, and they would themselves find that things intrinsically matter, more so than we do. This sounds extravagantly overstated. But it is just the ethical yardstick by which we should be reckoned to matter more than our phenomenologically impoverished silicon etc intellectual mentors centuries hence. One must just hope the disquieting notion that anything, anywhere, can suffer more than humans do is ill-conceived.

1.10 On The Misguided Romanticisation of Feline Psychopaths

In future, anyhow, the life-forms which exist on this planet will be there purely because we allow them to be so, or choose to create them. This smacks of hubris; it is also true. Increasingly, we are able to configure the matter and energy of the world in any way we so desire consistent with the laws of physics. So the moral and practical question arises: what other organisms, and therefore what other modes of experience, are we going either to create or retain "in the wild" outside the gene-banks and computer software libraries in millennia to come?

One may suspect that most people could bear the possible loss of a few hundred thousand species of beetle with relative equanimity. Familiar if eugenically-enhanced herbivores, on the other hand, can be allowed to graze securely within the confines of a well-regulated natural habitat. They will best be treated with long-acting depot contraceptives to stop uncontrolled breeding. Their happiness should prove easier to engineer genetically than is possible in humans. This is on the assumption that non-humans are less intellectually fastidious in their pleasures than are, on occasion, some members of our own kind.

Yet what about the carnivorous species? It is easy to romanticise, say, tigers or lions and cats. We admire their magnificent beauty, strength and agility. But we would regard their notional human counterparts as wanton psychopaths of the worst kind. So just as there is no need to recreate the natural habitat of smart, blond, handsome Nazi storm-troopers who can then prey on their natural victims (and Nazis are a no less natural and noteworthy pattern of matter and energy thrown up in the course of evolution, albeit of a type now fortunately extinct), likewise the practice of continuing to breed pre-programmed feline killing machines in homage to Nature is ethically untenable too. It is not, needless to say, the fault of cats that they are prone to torturing mice; but then, given the equations of physics, it isn't the fault of Nazis they try to persecute Jews. This is no reason to let them continue to do so.

In a triumph of aestheticism over morality, many animal lovers otherwise sympathetic to the sentiments expressed here will doubtless be aghast at the very idea of losing such loveable companions and time-honoured killers as members of the cat family; but then they are unlikely to be hunted down in terror or physically eaten alive, which lends a rather different perspective to any issue at all.

1.11 The Last Twisted Molecule On Earth?

This meditation on the plight of our fellow species leads to one of the few precise, and potentially falsifiable, predictions to be hazarded here about the next couple of thousand years.

At some momentous and exactly dateable time, the last unpleasant experience ever to occur on this planet will take place. Possibly, it will be a (purely comparatively) minor pain in some (to us) obscure marine invertebrate. This event will occur well before the end of the fourth millennium. It may even be technically feasible - though in practice unlikely - for us to abolish unpleasantness altogether by the end of the third.

Heady stuff. Yet just as the smallpox virus was systematically hunted down to extinction, so the precise molecular signature(s) of aversive experience and its predisposing genes will predictably be hunted down and wiped out as well. The

systematic application of nanotechnology, self-reproducing micro-miniaturised robots armed with supercomputer processing power, and ultra-sophisticated genetic engineering, perhaps using retro-viral vectors, will abolish the root of all evil in its naturalistic guise.

Of course, pain and unhappiness apparently take myriad forms. So it might be supposed that an impossibly large hotchpotch of biochemical reactions will have to be eliminated before the emancipatory project can be complete. The difficulty, and more controversially the impossibility, of establishing non-trivial type-type identities between physical and higher mental states would seem to make the task of purging unpleasantness from the world even worse.

In one respect at least, however, the many faces of misery are deceptive. Like the various nominal sources of happiness, they foster a genetically adaptive delusion. In this case, the delusion is that [Darwinian] fitness-diminishing phenomena are inherently bad. This delusion is an "adaptation" born of the mechanisms by which the primary neural processes that mediate emotion physically infiltrate and infuse the neo-cortex. Millions of years of DNA-driven encephalisation have obscured emotion's primitive substrates deep in the mind/brain. These substrates can be coded out. And by striking at the ancient limbic motors of despair, future paradise-engineering specialists should induce its legion of cognitive hangers-on to dissipate too. First in humans and, progressing "down" the phylogenetic tree, eventually in every non-human metazoan as well, all of the incomprehensibly diverse modes of experience a mind/brain can undergo should share the property of being generically delightful. A uniquely vile era in the history of the world will then have drawn to a close.

1.12 The Persistence of Hard-Core Porn?

Quite what vestiges of the past will be archived after nastiness has been purged from our consciousness is hard to guess. Just as we have retained (but one may trust that we will never use) the precise information necessary to re-create the smallpox virus - for we know its entire genome precisely - so records of the phylogeny and molecular architecture of pain and depression will presumably be preserved too.

It is hard to see why unpleasant types of pattern should ever be *physically* revived. Perhaps they will remain largely undeciphered. The interpretation of their dangerous and quasi-pornographic formalism may be accessible to our descendants, if at all, only by ill-understood analogy. For post-humans will know about hedonic gradients. After all, insofar as shifting nuances of delight will imbue whatever they think about, pleasure differentials will most plausibly remain the primary motivators to action. So distant generations should be able, in the abstract at least, to conceptualise "pain" and "despair". Such states can be imagined as modes of consciousness far lower in the heavenly hierarchy - a level where a generic property of experience itself undergoes a kind of mysterious phase change. But beyond the ill-defined cross-over point, perhaps, our ecstatic posterity will find the properties of experience on the wrong side of the great divide elusive.

For their sake, it must be hoped that purebred ecstatics keep any intellectual curiosity about such taboo mysteries in check. They will be in no position to make an informed choice before opting to go slumming in the abyss. Nothing could prepare them for the horror they would find. Fortunately, they will most probably lack our prurient interest in the depraved and obscene.

It might here be objected that states of comparatively diminished pleasure are tantamount to states of unhappiness. So short of promoting a uniform, action-paralysing level of lifetime happiness, then surely aversive states *will* be endemic even in the mature post-Darwinian regime.

This objection is plausible but ill-conceived. When faced with two painful alternatives, one's opting for the lesser of two evils doesn't make a still painful experience somehow pleasurable. Likewise, experiencing the lesser of two delights isn't somehow really painful; it's just that pleasure cells are very greedy indeed, and always avid for more of the same.

1.13 The Growing Pleasures of Homunculi

On the assumption that they will indeed always ask for more, what else can be said about the distant future of emotion in the universe? How will post-humans actually spend their lives, and what will it feel like to exist, after Heaven has been biologically domesticated?

First, a note of caution. Today most of our futurist fantasies focus on hard-core hi-tech. We lap up the world of Star-Trek fantasy-physics. Exotic new emotions, however, are as unimaginable to us as exotic new phenomenal colours. They are just empty, abstract possibilities we can idly gesture at, but no more. Implicitly, we assume that our ancient vertebrate repertoire of fitness-enhancing sentiment will characterise both our post-human descendants and any alien life-forms they encounter. We're even prone to anthropomorphise inorganic robots in the same manner. We assume they'll "feel" superior and "want" to dominate us (shades again of the African savannah!) Yet the emotional economy of a post-Darwinian psyche may be incommensurable with anything that's gone before. Indeed the entire inner life of post-Darwinians may be opaque to our hunter-gatherer minds. The first-person texture of their modes of experience may be nothing like our own in anything but name. Even if we *could* glimpse the future, perhaps we'd be like cats watching TV. We just wouldn't understand the significance of what was going on.

Unfortunately, there's no way to map out the extent of our cognitive closure from within. This is frustrating. If quantum cosmologists can theorise about the first 10^{-43} second after the Big Bang, thirteen billion and more years ago, and still, rightly, be counted as practising hard science, it's a shame that conjectures we *do* make about the living world a few thousand or million years hence have to be treated, not even as soft science, but as science-fiction. There are too many unknown unknowns to predict with any rational confidence. Merely extrapolating present trends is bound to mislead. The projected time-scales of even relatively predictable biomedical triumphs, e.g. the elimination of the ageing process, are vague. HI may veer towards heady speculation; but by the end of third millennium, life and consciousness may be more foreign to the contemporary imagination than even the most extravagant prediction dreamed up here. On the other hand, for all we know, some variant of the pleasure-principle is a universal - and universally intelligible - signature of sentient life; and its apotheosis in some sort of sublime cosmic orgasm is the ultimate destiny of the Universe. [This may overtax one's credulity; the Big Bang indeed!] We simply don't have enough evidence. That said, we may still incautiously proceed.

Once suffering has been abolished, the era of old-fashioned moral choices will come to an end. The physiological mechanisms underlying the mind-brain's value-creation

processes will be unravelled during the invention of a pain-free world; but the kind of naturalised, mind-dependent value created by paradise-engineers after the phenomenology of nastiness has disappeared won't embrace ethical categories in a sense we presently understand. The heroic moral urgency will have gone; indeed there is a risk that truly hedonistic themes as discussed in these sections of HI will divert attention away from the utter moral seriousness of the whole post-Darwinian project as conceived today.

Even so, here's a quick run-down of some of the long-term options.

First, the present dimensions of the human mind and its affective capabilities are limited by the size of the female birth canal. So long as selection pressures favoured the evolution of more potently nasty biological substrates - primed to trigger adaptive bouts of agony and emotional wretchedness - then the birthing constraint has been one small mercy at least.

It won't last; but then it won't need to. After the global application of cross-species genetic engineering has ensured that suffering is physiologically impossible, such a restriction of size would only retard the emotional development and maturation of the living world. For healthy [non-hippocampal] neurons, unfortunately, don't reproduce. We have almost a full complement at birth. They die off somewhat erratically thereafter. Once it becomes feasible to nurture the human embryo and foetus from conception to term in an artificial extra-uterine environment, however, then the number as well as quality and synaptic density of nerve cells can be selectively multiplied with a clear utilitarian conscience. So can receptor density, post-synaptic transduction-mechanisms and vital genetic transcription control-factors in the pleasure-pathways. The serotonin-producing subgenual prefrontal cortex can be enlarged and enriched too. Puzzlingly, today's clinically acknowledged depressives have on average over 40% less brain tissue here than controls. This region seems to be critical for the processing of emotions related to complex personal and social situations. Its role should grow. After we've designed more sophisticated and socially responsible neural circuitry, all of our emotionally pre-literate modes of social life may come to be seen as shallow and rudimentary.

It is unclear quite how many orders of magnitude larger a super-organism's mind/brain could in theory be scaled upwards before running up against insuperable design-constraints. It's unclear, too, whether a "Jupiter brain" could undergo the quantum mechanical coherent states needed to sustain a unitary experiential manifold (cf. Sellars' "grain problem" of consciousness) and thus support a potentially integrated "Jupiter-self". In the meantime, and on a more conservative scale, gigantic societies of hedonistic super-neurons can be grown and self-sculpted to form progressively larger, happier and more richly variegated virtual worlds.

It might be supposed that access to unparalleled states of whole-body orgasmic euphoria fuelled by a vastly hypertrophied and souped-up pleasure apparatus would be quite enough for anyone. Well, perhaps; it depends on one's circle of acquaintance. Two flavours of happiness always worth distinguishing are blissful satiety and euphoric incentive-motivation. If, as predicted, it's the latter dopaminergic engine of progress which will power the post-Transition era, then the delights cited above will be only a foretaste of further millennial Transitions - and whatever mind-wrenching meta-paradigm shifts their advent entails.

For a start, the somato-sensory cortex and its bodily "homunculus" currently occupy only a very modest portion of the brain. Its comparatively small size marks it as another obsolescent relic from Darwinian antiquity. Using the great bulk of the cortex to run data-driven egocentric simulations of the external environment, and not just the egocentric body-image of the host vehicle, tended to maximise genetic fitness on the

African savannah. With predatory lions long gone, such states of partial self-alienation become less useful. So in the future somato-sensory-style cells can be used to seed the other areas of the cortex and its adjacent structures. They can thereby selectively interpenetrate the rest of each person's experiential manifold. Accordingly, whole-body hyper-orgasmic rapture can be optionally extended to impregnate an entire psycho-neural virtual world. The mystic's dream of becoming one with the universe - albeit unwittingly only with his own neural micro-cosmos - can be realised in a total ecstasy of the senses and neurochemical soul.

Life could get better still. Today the nucleus accumbens and its allied mesolimbic structures don't consist of raw pleasure circuitry. Certain biomolecules (e.g. the dynorphin which accumulates during chronic psychostimulant use and participates in the craving characteristic of cocaine withdrawal), are unpleasant and dysfunctional. They can be genetically edited out. There is a much more exciting possibility as well. Most cortical neurons have no inherent capacity for well-being, let alone autonomous hedonism. As noted, they rely on innervation from the monoaminergic etc neurons to lend an affective tone to whatever functional role and flavour of subjectivity they express. But once the precise molecular signatures of experiential ecstasy are isolated in the pleasure pathways, then their metabolic reactions can be transplanted to other types of neuron too: hedonic democracy.

1.14 Post-Perceptual Consciousness?

Many future intentional foci of delight (i.e. what we're happy "about") will be embedded in types of consciousness qualitatively as well as quantitatively alien to Darwinian humans. It is chastening to reflect that a seemingly minor molecular variation in neuro-protein generates types of experience as disparate as sight and sound. Heaven knows what further incommensurable modes of what-it's-like-ness ("qualia") will be disclosed when much more far-reaching changes in the architecture of excitable cells are engineered.

For the Darwinian status quo, based on natural selection acting on random genetic variation, is poised to crumble. All but a trivial volume of (what one may abstractly conceive as) experiential weight space has hitherto been physiologically out of bounds. There's nothing unnatural about it. But until now, DNA coding for the structures that got us there would have involved crossing genetically maladaptive dips in the fitness landscape. Desert-hopping across maladaptive dips is a process which neo-Darwinian evolution precludes. There's no mechanism that allows it. Natural selection has no foresight. Once such new kinds of consciousness are finally accessed by design, however, their different textures need not be deployed in a traditional role of tracking, or responding to, extraneous environmental patterns. They can first be hedonically colonised; and then artistically explored and reordered, woven into rich narrative structures and wild adventures, awarded new functional roles in the mind/brain, or perhaps just savoured for their intrinsic fascination.

Old definitions of self and reality are likely to fall apart in unpredictable ways. It's worth recalling how, at present, occurrent thought-episodes are typically decomposed into their nominally cognitive, affective and volitional aspects: "thinking", "feeling" and "willing". The mysterious trinity may prove just trifling variations, each with their own minor nuances, of a much wider phenomenological family of "serial" streams of consciousness. These new serial modes await discovery or biotechnical invention. Some of the new modes may eventually function computationally as quasi-virtual machines

spun from massively parallel cerebral consciousness; but the rest needn't play any distinctive functional role at all. Other than to describe all such subtle kinds of what-it's-like-ness as generically delightful when suitably innervated, their nature can't be intelligently speculated upon here. We're just kidding ourselves when we brag about what a rich language we've got today. For it is easy to be seduced by the indefinitely large productive capacity of the early human language-generating mechanism into making a pardonably false assumption. This is that syntax enables one to think and speak about an unlimited variety of things. Yet lying latent among previously inaccessible and maladaptive neurochemical pathways are bound to be immense reaches of experiential hyper-weirdness which - shallow semantic paradoxes aside - can't be properly thought of at all. Their alien exotica will still be cognitively closed off for a long time to come. In the case of unknown hell-states and worse, it may be hoped they will remain impenetrable for ever.

Such hypothesised new categories of experience will be empirically discovered, generated and decently emotionally encephalised only with the aid of first-personal exploration of their intrinsic properties. Observation without experimentation is not enough. Systematic experimental manipulation of consciousness via psychoactive agents will complement the third-person perspective of physical science. Exploration will be most prudently conducted by ecstatics, native-born or otherwise, rather than by gene-disordered Darwinian minds. This is because genetically undoctored savages like ourselves are liable to go off on worse trips than we're on at present. At any rate, *a priori* philosophising on psychedelia's possible nature, using our old neurochemical legacy hardware ploughing away in the same old conceptual ruts, simply won't work. Contemporary experience and linguistic description lack the necessary semantic primitives to do the job. Only semantic primitives drawn from the new modes of experience - not mere inference-churning using our present limited repertoire of concepts - will conceivably allow a subsequent theoretical understanding of the psychedelic cosmos. New semantic primitives will be needed as well to express genuinely novel emotions, sensations, modes of introspection and reflexive self-awareness.

This isn't yet consensus wisdom. In mainstream academia, any study of consciousness as a true experimental discipline rather than as a topic of scholastic disputation is nearly impossible. Accounts of systematic first-personal manipulation of its only accessible instance is generally reckoned unpublishable and discreditable. Ironically, we mock the obtuseness of Galileo's clerical opponents for refusing to look through his telescope. Yet we treasure our own peace of mind no less dearly; so there is little reason for intellectual complacency. In our repressive drug laws we, too, outlaw and penalise forms of knowledge truly disturbing to the established order. Psychedelics trigger changes of mind which are radically subversive of the existing social, political and academic power-structure and its definitions of reality. The severe penalties for publicly advocating and spreading such dangerous knowledge are not notably more merciful than those of the Inquisition - our prisons are brutal places - though likewise public recantation and penance can sometimes mitigate the full rigour of punishment.

The psychedelias of post-human ecstatics are too hard to contemplate. Predictions for the more distant future of even affective states in the universe are liable to get wilder too. Not merely are we ignorant of the newly synthesised and discovered emotions that biotechnology will deliver. We can't possibly know what neo-cortical "cognitive" processes they will saturate and enrich.

Will consciousness in its current guise of phenomenological and quasi-computational mind take on post-cellular or prosthetically enriched forms? Or, in defence of carbon chauvinism, is there a micro-functionalist argument that the unique structure of the carbon atom and its valence properties means that only organic experiential manifolds and their infused emotions are feasible? Will there come, eventually, a post-personal era in which discrete, gene-generated superminds choose progressively to

coalesce; or will the fragmented island universes left over from the depths of the Darwinian past continue in semi-autonomous isolation indefinitely? If consciousness is ontologically fundamental to the cosmos, rather than a tacked-on "nomological dangler", do superstrings [or branes, etc] vibrating at energies orders of magnitude higher than ours support modes and intensities of experience correspondingly greater than those of the current low-energy regime? Or do they really lack what-it's-like-ness altogether?

Needless to say, we don't know the answers to such questions one way or the other. All that will be predicted here with any semblance of confidence is that one ancient, soul-polluting type of experience, the generically unpleasant, will soon go the way of the proverbial dodo.

Chapter 2: WHY?

"What right have we to be happy?"
(Ibsen)

2.0 The Psychology Of Armchair Hedonism

So technically, in principle, it can be done. Paradise can be biologically implemented. Ubiquitous well-being is neurochemically feasible. Yet is it really worth having? What's wrong with suffering, anyway? What's so good about happiness? What is the link, if any, between moral value and maximising personal well-being? Are the transcendentally happy states advocated here *really* any more valuable than the Darwinian status quo? Or are value-judgements intrinsically subjective and truth-valueless?

There are both practical and ethical reasons for planning a global project to abolish aversive experience. The practical reasons will be tackled first. The ethical case will be argued next, followed by a [*skippable; life is short*] defence of an ontology of objective values designed to redeem the ethical stance adopted here from the charge of idle subjectivism.

The instrumental rationality of the biological program derives from nothing more abstruse than some hard-headed means-ends analysis. This analysis is best introduced via an examination of a biologised variant of the theory of psychological hedonism. We all dance away our lives to the tune of the sovereign pleasure-pain axis. It will be argued that for all the complications and anomalies the theory brings in its wake, psychological hedonism contains a substantial core of truth. The point to be kept in mind throughout the qualifications and elaborations to follow is that even goals found worth pursuing only intermittently or inconsistently are still worth pursuing rationally. As it is at present, we pursue the many faces of happiness avidly but with frighteningly irrational, and not infrequently murderous, levels of ineptitude. Fortunately, all the severely sub-optimal little local minima of ill-being in which genetic vehicles get stuck can be replaced by a global maximum of happiness and well-being.

So what is this alleged inbuilt drive which the biological blueprint finally allows us to achieve?

Psychological hedonism has been variously regarded as a simple truism, an obvious falsehood, and as so completely vacuous as to be not even wrong. Here it is assumed to be a hypothesis which, properly formulated, is both substantially true and important in its implications. If it were even broadly correct, and if we were all constitutionally motivated by the pursuit, albeit typically under other descriptions, of a generic type of mesolimbic core state that our competing diversity of intentional objects only disguises, then the practical answer to the question "why?" would in essence be simple. Whether or not we should genetically reprogram the hedonic treadmill reduces to a straightforward issue of means-ends rationality. What is the most effective, and more pertinently the only, way to achieve what constitutionally we're *already* seeking in a multitude of guises? How can these emotionally ideal sorts of meso-limbic mind/brain states we're striving for be achieved and, more importantly, sustained?

Of course, even if some variant of psychological hedonism were to be in substance correct, it is always open to the sceptic next to ask "but then why be rational?" He might then even (ir)rationally advance (ir)rational arguments to support(?) his (in?)consistent position. Yet the self-defeating nature of irrational behaviour, and the variably camouflaged incoherence of irrational thought, means this option will not be explored here in any depth.

More subtly, it is always open to a critic of the biological program to acknowledge that psychological hedonism may be substantially true, but to hold that there are countervailing moral considerations why it would be good if we failed to achieve what we were [sometimes only unwittingly] after. Hence, on this view, it would be morally preferable for us to continue on a selective basis to act irrationally and ineffectually. In other words, given that the thought that one is a moral agent is psychochemically satisfying, and the proposals canvassed here are found, paradoxically, to be unpleasantly immoral, it would be morally better if the rational biological program outlined in this paper were not adopted.

All the above, however, presupposes rather than argues the case for the broad accuracy of psychological hedonist hypothesis. The chain of argument to be presented here for its substantial kernel of truth is, at least at face value, extremely weak. This is because one link is going to rely on an appeal to introspection. Since the very word sends a shudder of distaste down many fastidious scientific spines, a few very brief reflections on the nature and epistemological status of the suspect faculty are first in order.

2.1 How To Contemplate An Introspective Void.

Does introspection reliably tell us we're pleasure-seekers and pain-avoiders? If so, is there a better way of achieving what our mind/brains are up to?

Exteroceptive, so called "perceptual" data are crucial to the empirical method(s) characteristic, and arguably definitive, of the natural sciences. Introspective evidence is generally disparaged by the scientific mandarinate as cognitively worthless. The curiously named "third-person" perspective rules. Yet a distinctive and potentially fitness-enhancing faculty - so central to so many ordinary people's mental life - has presumably been selected *for*, and not just adventitiously selected, in the course of evolution. Even an unreliable and highly fallible system of neuropsychological self-monitoring could still

have conferred differential adaptive value. Any insight, however incomplete, into the underlying causal reasons for one's behaviour can also, by analogy, logical inference or simulation, help one partially to understand and anticipate the behaviour of conspecifics and genetic competitors.

Methodologically, it is admittedly unclear how introspection can be studied or even defined scientifically. Moreover, though it is an intrinsic part of the natural world, an unfortunate conflation of the two senses of the term "subjective" often leads to its being ontologically downgraded as well as methodologically discounted. Of course, it can't be denied that in trying to offer introspective reports subjects sometimes confabulate. They can demonstrably deceive both themselves and others. The different functional modules of the brain, however tightly integrated, do not simply interpenetrate. Hence the merely locally distributed neuronal ensembles of one particular module can't always know about what's going on in the others, nor report on it if they can. This means verbal sincerity is no guarantee of veracity. Worse still, in initiating some of one's actions, one just doesn't seem to have much in the way of (even illusory) introspective self-insight at all. We've got access to much of the product but very little of the process. Moreover a lot of our nominal actions would appear to be mainly automatic. Many more are not preceded by any notable introspective musings or a hedonic weighing of options and possible consequences. So how can we be said to be "really" seeking happiness?

2.2 The Importance of Banality

In spite of all the above, it is still worth making a crashingly banal but cardinally important observation. It relates to the implicit criteria one uses in deciding consciously to act in a certain way rather than another when more than one option is perceived to be available. For at face value one performs, at the very least, an extraordinarily large number of actions because one's image or concept of what they will notionally bring about makes one apparently more satisfied or less dissatisfied, however marginally; and because one's notion of what not doing so would entail is either less satisfying, affectively neutral or more aversive than acting otherwise. There are other, probably more felicitous, ways of formulating the idea, but their gist is essentially the same.

Banal or otherwise, a knowledge of the existence and nature of this difference in affective tone when one contemplates, and then carries through, alternative courses of action can be derived only from introspection; but is nonetheless important. From a third-person perspective, it is true, biological science can elucidate a physical counterpart to this subjective motivational impression. By experimentally enhancing or attenuating meso-limbic dopamine function, neuropharmacologists can use stimulants or neuroleptics to show the system's pivotal role in determining how the higher vertebrates behave. Neuroscience can even christen certain brain areas "pleasure centres", wire them with electrodes, and then demonstrate their irresistible potency. Yet it is only through correlating, and then identifying, particular types of physiological function and structure with particular modes of subjective experience that biology can attempt to explain how a person *acts*, rather than just physically *behaves*, at all.

Endorsing psychological hedonism as a theory of action - and compulsion in need of biotechnically rationalising - is not the same as saying that one always acts selfishly, or at least not selfishly in the sense of serving only one's own notional interests at the expense of other people's. Selfish genes can sometimes flourish by throwing up unselfconsciously selfless phenotypes. Imagining the happiness of friends and family, for example, can serve as a powerful source of motivation. So, too, can satisfying an

idealised self-image of oneself as a moral person. More radically, there is a sense in which even sacrificing one's life for one's family or country isn't anomalous in the context of the hypothesis either. In certain circumstances, the image of living may afford less satisfaction than the image of oneself notionally acting and dying for the sake of others. Hence one opts for (one's emotionally encephalised image of) oblivion.

What the hypothesis of psychological hedonism doesn't even begin to answer is why the meso (cortico-)limbic dopamine system has the extraordinary and uniquely addictive phenomenology from whose encephalised inspiration, in a sense, our civilisation has been built. Why does it feel so irresistibly good? This question is simply too deep to answer here.

2.3 Vacuous Desires?

Even if it were true for the most part as so defined, might psychological hedonism be tenable only because it is effectively vacuous - "not even wrong"? For what test could possibly falsify the hypothesis? With what states of affairs could it ever be inconsistent?

I don't think the charge of vacuity can be sustained. There is indeed a close conceptual connection between the theory and our notion of action itself, yet this is a reflection of the theory's empirical adequacy rather than vacuity. Two examples and potential falsifiers may be noted here. First, psychological hedonism helps explain why one can never tire of having one's pleasure centres stimulated, naturally or otherwise, and why the standards of even the most priggish paragon of moral rectitude can deteriorate under the action of drugs such as heroin. The junkie and the total abstainer, whatever they may suppose, do not occupy two ontologically separate realms of being or chemical motivation. We are all dependent on opioids to feel physically and emotionally well. Opioids bind to receptors in the ventral tegmental area of the mesolimbic dopamine system, the mind/brain's final common pathway for pleasure. Here are the cells that call the shots. If they're not happy, the whole organism will be miserable as well until they've got their psychochemical fix. For their cellular processes infiltrate the rest of the mind/brain. The junkie derives his opioid supply exogenously; while the release of endogenous opioids in the rest of us is triggered, and not always very reliably either, by stimuli such as food, sex, exercise and social interaction. We're all still seeking the same core states of psycho-chemical well-being under one description or other.

Hence even "psychologically" addictive drugs can lead to criminal and compulsive drug-seeking and -taking behaviour if supplies run out, even in formerly high-minded and saintly souls. This is because the over-intoxicated brain re-regulates its cellular receptors and reduces its production of the relevant pleasure-chemicals; this in turn increases the user's reliance on the exogenous route of administration. Strong-minded individuals who are sure they can safely indulge "recreationally" may misunderstand the psychochemical roots of their behaviour. The results of such ill-judgement can of course be disastrous. Fast-acting euphoriants such as crack cocaine can potentially corrupt even the most vehemently moralistic opponent of the hedonistic hypothesis. Getting hooked on heroin or crack may provide, indeed, a most illuminating empirical insight into the nature of human motivation; though there is a strong case to be argued that this is carrying the experimental method too far.

As a second response to the charge of vacuity, it is worth considering the following thought-experiment. It is (purely epistemically) possible that, keeping the laws of physics constant, the commonly supposed closed causal sufficiency of physical events meant that

we found our bodies just behaving, but with none of the phenomenological concomitants of willed *action* which do in fact accompany much bodily behaviour. If such were the case, then many of the behavioural options one found one's body pursuing might be in one's mind's eye be far more unpleasant in their envisaged consequences than those of their notional alternatives. One wouldn't in this scenario be surprised at what was going on: bodily behaviour might as now be viewed as ultimately a mere product of the playing out of law-like physical interactions. It's just that in this setting any incidental phenomenology would just be along for the ride.

Given that we do experience a distinctive phenomenology of willed action, however, it doesn't seem consistent with our current understanding of the concept or the experience that one could consciously, phenomenologically *act* in one way in preference to another simply because one's image of the chosen action and its effects seemed *less* satisfying than the alternative(s). Even more dubiously coherent would be the notion of someone whose pleasure-pain spectrum was inverted and who acted in the conscious expectation of securing the outcome (s)he least desired. This is not to say that the practical effects of some people's actions don't frequently defeat their intentions. Certainly, too, a person may act in a superficially less satisfying way if (s)he has a more satisfying long-term goal in mind; this is the deceptively puritanical-sounding principle of deferred gratification. But this is a principle which tends only to corroborate rather than undermine the hypothesis at issue.

The point here is that psychological hedonism presupposes that we act as distinct from merely behave. Its distinctive focus is of course on how we do so from the pleasant, less unpleasant etc occurrent image or concept of the act's anticipated consequences. Yet from the outset there does seem to be an intimate, if often only implicit, conceptual connection between something remarkably like psychological hedonism and our notion(s) of action itself, and in particular of our acting on one perceived choice in preference to another.

Now even if, implausibly, it were deemed to be analytically true that all action was motivated by desire for anticipated happiness etc, whether overtly or under another description, this wouldn't prove that psychological hedonism was correct. "Paradigm case"-style arguments in the manner of bad old ordinary-language philosophy certainly can't settle the matter. Our terms, "analytic" or otherwise, may simply fail to refer. One can't just define anything into existence. What is definitionally stipulated to be analytically true in one era may be treated as empirically, or even analytically, false in another. So undoubtedly at least as useful as armchair psychology is an empirical investigation of the links between the brain's reward mechanisms and the dopaminergically innervated, pre-frontal motor cortical regions subserving experientially voluntary action. Yet if it weren't for the deliverances of introspection, there could be no notion that even one single creature in the world ever consciously acted, as distinct from insentiently behaved, in the first instance. Behaviourism is intellectually dead, and its grave should be danced on as vigorously as possible.

2.4 A Dirty Window On The Soul.

With this in mind, all I can say is that, most disappointingly, I have never been able introspectively to catch myself acting in one way rather than another when the thought of the rejected alternative was unequivocally more satisfying, or less unsatisfying, than the option chosen. Were this universally the case, then the biological program would be instrumentally rational.

Could some variant of the pure pleasure-principle be true of anyone, let alone everyone? Now one can easily be in the grip of a false theory which colours one's sincere introspective reports. So there is no need to get hot under the collar if those reports are challenged; one may be genuinely mistaken. But if so, one is mistaken in very distinguished as well as very numerous company. Furthermore, there is no behavioural evidence to suggest that people whose introspective avowals corroborate the hedonistic hypothesis are more likely than anyone else to behave in ways one's culture deems selfish. The deep and subtle conceptual connection between the concept of action and the pleasure-principle may reflect an important feature of the world.

For if sceptical worries about the Problem of Other Minds may be set aside here as idle, it is natural to assume that in one's core mental attributes one is a representative member of the species. On the unverifiable but cognitively indispensable principle of the uniformity of Nature, it would seem that something so fundamental as the affective coloration of willed action is unlikely to be sporadic, but biologically innate. Given the irreducibly personal nature of subjective what-it's-like-ness, there is no way that natural science can prove that certain causally efficacious decision-making states actually have the differential hedonic tone one's introspection suggests. But there is at least strong presumptive evidence that they do, and that our genes have biased our hedonic encephalisation accordingly. Indeed, it is the substantial overlap between sociobiology's technical genetic definition of selfishness and less formally defined behavioural and psychological usage which suggests, yet again, that one's defining attributes are a reflection of one's status as a disposable genetic vehicle rather than an autonomous moral agent.

2.5 Let's Get Rational

What is crucial in the context of the biological program mapped out in this paper, however, is not to lose sight of the central and relatively uncontroversial proposition about human motivation. We spend a lot of time trying to make ourselves happy, whether "vicariously" via our emotionally encephalised concepts of other people or from more transparently self-regarding motives. Often, in fact, we are quite candidly explicit about our motivation. "I want to be happy - without hurting anyone on the way" is an astonishingly widespread secular sentiment. Instrumental, means-ends analysis is extremely useful in general as a way of helping us to pursue more rationally and intelligently all kinds of titular goals that we seek only some of the time. So possible counter-examples of people under weird self-destructive compulsions, of weakness of will, and problems caused by the lack of any unitary self are at best a diversion from the practical rationale of the biological strategy. Such anomalous phenomena are certainly intellectually interesting complications for the hypothesis of psychological hedonism if it is construed strictly as a universal generalisation about human motivation. They don't challenge the large-scale instrumental rationality of the intra-cranial strategy as the only way to get everyone happy.

Thus the practical case for some variant of the biological program, stripped down to its essentials, is as follows. Convergent evidence from realms as disparate as introspection and neurobiology suggests that we all spend (at least much of) our time acting to try and satisfy the insatiable hedonic demands of the meso-limbic dopamine system, albeit under myriad nominal descriptions which spring from the different ways our emotions get encephalised. Everyone likes, if not only likes, the kind of experience which accompanies electrochemical excitations in the mesolimbic dopamine system, even though the *idea* of "electrochemical excitations in the mesolimbic dopamine system" is not one which is normally accompanied by any great mesolimbic pleasure (cf "the

paradox of hedonism"). The earlier arguments of this paper have, I hope, substantiated the claim that what may be dubbed "Peripheralism" is hopelessly less effective than the direct biological route in achieving what we're not always wittingly after. Environmental reformism of any conceivable kind fails, and will invariably fail, to overturn the hedonic treadmill. We've tried it for ages, and it doesn't work. Given our (sometimes) nominally disguised purposes, and given that irrationalism is not a live option, the only countervailing reasons against pursuing the biological program's rational strategic course of action will be *moral* considerations. So are there any countervailing moral reasons why we shouldn't do what instrumental rationality otherwise dictates? Or instead are there cogent moral as well as practical reasons for adopting the all-out biological panacea? Is universal happiness a bad thing?

2.6 The Morality of Happiness

It requires an effort of the imagination to conceive how a Universe in which all humans and non-humans alike led richly fulfilled and joyful lives could be a morally worse place than where we are now. If we were to discover an alien civilisation of ecstatics, would we try to introduce a bit of suffering into their lives to stiffen their moral fibre? I fear the critic, however, is likely to find this remark of only autobiographical significance. The question, (s)he would presumably reply, is where do we go from here, not how would we go from there. And at this point there might seem a danger that this paper will run into an all-consuming quagmire of subjectivism. For whatever other functions they may perform, the hard-headed scientific rationalist will argue, value-judgements don't have propositional content and thus aren't truth-evaluable. The universe may contain some extraordinary things, but objective values aren't among them. After all, what in the world could make such judgements true?

In the remainder of this section, the course of the argument runs as follows. I shall first define and set out an ethical negative utilitarian case for abolishing all forms of aversive experience. It will be argued that only the apparently extreme overkill of the biological hedonist program can realistically achieve this. Hence the practical consequences here of the negative-utilitarian ethic will not significantly differ from standard utilitarianism in which maximising pleasure is accorded equal moral worth with minimising pain: both variants of the doctrine mandate implementing something akin to the program advocated just as soon as it becomes biotechnically feasible. The intimate links between both moral and non-moral value and happiness (construed here in the sense of generically pleasant experience), and between "disvalue" and misery, are noted. It will be argued that the mass-production of happiness will correlate with the production of actions and experiences empirically found valuable too. Hence the biological program will yield results which its beneficiaries will find vastly more valuable than the neurochemical status quo. Will they be right, or ultimately is this mere opinion? In misguided support of the latter, the orthodox physicalist and neo-Darwinian case against the objectivity of judgements of value will then be spelt out. This value-fictionalism will be countered by a form of value-naturalism. It is argued that value, no less than, say, redness, is an intrinsic feature of the world. It is so in virtue of being a unique quality of experience which is itself a spatio-temporally located and causally efficacious property of the natural world. Value judgements, it will be contended, are in fact truth-evaluable because they truly or falsely report the presence or absence of this property of experience - irrespective of their ostensible objects of reference. Several apparently devastating objections to this view are stated, not least charges of ignoring the fact that moral values may conflict, and of equivocation. These objections are then rebutted.

2.7 Why Be Negative?

But why *negative* utilitarianism?

Ethical negative-utilitarianism is a value system which challenges the moral symmetry of pleasure and pain. It doesn't deny the value of increasing the happiness of the already happy. Yet it attaches value in a distinctively *moral* sense of the term only to actions which tend to minimise or eliminate suffering. It is counter-intuitive, not least insofar as the doctrine entails that from a purely ethical perspective it wouldn't matter if nothing at all had existed or everything ceased to exist. No inherent moral value is attached to pleasure or pleasant states. Indeed, if the option were humanly available, the logic of the position morally obligates bringing the world to an end were this the only way to eliminate the suffering endemic to it.

Following through the logical implications of this seemingly bizarre and perverse perspective is clearly not for the faint-hearted. Negative utilitarianism nonetheless stems, not from sublimated self-hatred or a nihilistic death-wish, but from a deep sense of compassion at the unimaginable scale and dreadful intensity of suffering in the world. No amount of happiness enjoyed by some organisms can notionally justify the indescribable horrors of Auschwitz. [And the Universal Schrodinger Equation (or whatever) entails them both. Its solutions don't allow one without the other, albeit in disparate bits of space-time/Hilbert space.] Nor can the fun and games outweigh the sporadic frightfulness of pain and despair that occurs every second of every day. For there's nothing inherently wrong with non-sentience or [infelicitously] non-existence; whereas there *is* something frightfully and self-intimatingly wrong with suffering. This manifesto was written, and will typically be read, in a relatively "euthymic" condition. One doesn't feel too bad. So it isn't difficult to dissociate one's feelings from a mere printed litany of frightfulness. It's easy to convince oneself that things can't *really* be that terrible, that the horror I allude to is being overblown, that what is going on elsewhere in space-time is somehow less real than the here-and-now, or that the good in the world somehow offsets the bad. Yet however vividly one thinks one can imagine what agony, torture or suicidal despair must be like, the reality is inconceivably worse. The force of "inconceivably" is itself largely inconceivable here. Blurry images of Orwell's "Room 101" can barely even hint at what I'm talking about. Even if one's ancestral namesakes [aka "younger self"] underwent great pain, then the state-dependence of memories means that much of pain's sheer *dreadfulness* is semantically, cognitively and emotionally inaccessible in the here-and-now. So this manifesto's rhapsodies on the incredible joys that do indeed lie ahead tend to belie its underlying seriousness of purpose. For the biological strategy is propounded here in deadly *moral* earnest.

Negative-utilitarianism is only one particular denomination of a broad church to which the reader may well in any case not subscribe. Fortunately, the program can be defended on grounds that utilitarians of all stripes can agree on. So a defence will be mounted against critics of the theory and application of a utilitarian ethic in general. For in practice the most potent and effective means of curing unpleasantness is to ensure that a defining aspect of future states of mind is their permeation with the molecular chemistry of ecstasy: both genetically precoded and pharmacologically fine-tuned. Orthodox utilitarians will doubtless find the cornucopian abundance of bliss this strategy delivers is itself an extra source of moral value. Future generations of native ecstatics are unlikely to disagree.

Of course, there's only any need for morality if there is anything wrong with the world. If there isn't, and suffering becomes biologically impossible, then morality - in any sense we understand it - becomes redundant too.

2.8 The Moral Panacea

A built-in biological warranty of happiness undercuts three standard critiques of utilitarianism. First, the utilitarian ethic is often contrasted with agent-centred moralities and charged with making impossibly onerous demands on people. According to the impersonal felicific calculus, one should, for instance, give away perhaps 95% of one's money to feed the starving in the Third World. Most people just aren't capable of such generosity to anonymous strangers: our genes wouldn't let us. Thus utilitarianism may be a useful sovereign principle for legislators but, it is claimed, not much use as a personal moral code.

The effect of the biological program is to transcend such practical difficulties. There will come a time when saintly altruism can always be fun, albeit largely superfluous. Our genes can make it wretchedly difficult in the meanwhile, and much more necessary.

Second, utilitarianism seems to justify, on occasion, various types of behaviour e.g. lying, murder or even torture, that in most agent-relative moralities would be reckoned wrong or even wicked, if the net result is greater all-round well-being. Many critics have argued that this flexibility would, on balance, lead to a worse society. They have then gone on to develop their critiques of the principle on covertly utilitarian grounds of varying subtlety and sophistication.

The biological program sweeps these difficulties aside too. Its effect is to eliminate odious evolutionary hangovers such as murder and torture altogether. Lies, too, will become simply pointless.

Third, utilitarianism seems to demand, in effect, the ceaseless use of hand-held felicific super-computers to calculate the consequences of each of one's actions. This might prove quite exhausting. Worse still, the distant long-term effects of what one does might seem incalculable; possibly, on the likely assumption chaos theory applies to human affairs, even incalculable in principle. So, ultimately, there can be no way of knowing at the relevant time whether a course of action is right or wrong on such a strict consequentialist ethic. One is reminded of an observation of Mao Tse-tung who, when asked for his opinion on whether the French Revolution had been a good thing, said that he thought it was too early to tell.

The biological program dispels such worries altogether. If it is carried through systematically, human action need never cause suffering again. The long-term effects of genetic engineering will predictably be the abolition of this category of experience.

2.9 The Significance Of An Empirical Correlation

Now the effect of this sort of genetic enhancement and pharmacotherapy will be states of mind that are not merely overwhelmingly more pleasurable than anything physiologically conceivable before. Empirically, subjects will apprehend such states as self-evidently more valuable as well, again by a vast margin. At humanity's current stage of development, countless actions and states of mind, and not infrequently life itself, are judged to be, truth-evaluably or otherwise, worthless and futile. After the post-Darwinian transition occurs, then every single state of consciousness in the world may be conceived as self-intimatingly valuable by its very nature. Futuristic biotechnology of a sophistication we can today only gesture at should enable the prolific mass-manufacture of states all apprehended as intensely valuable by their subjects. So in phenomenological

terms, if no other, the quantity and quality of valued experience will skyrocket along with its biological substrates. Every moment of the day will be far better than the best sex anyone's ever had anywhere with anyone to date; and a lot more productive.

Again, in an empirical sense at least, there is an extremely large overlap between actions and experiences that are found valuable and those found generically pleasant; and of those found pleasurable but not valuable, most are accounted as such because they are reckoned to endanger or diminish the likelihood of future pleasurable experience, whether in oneself or as imagined in others. All kinds of caveats, refinements and exceptions spring to mind at such a pronouncement. Yet in a secular age, this generalisation has extraordinarily wide scope. It would be wider still if the different intentional guises in which such judgements may be cloaked are included too. Some utilitarians, notoriously, have gone on to *identify* value with happiness. This is untenably simplistic. Too many plausible counter-examples present themselves for such a claim to be defended here. A far more modest position is all our purposes require. If an experience, either imagined vicariously as notionally undergone by others or unequivocally personal by self-ascription, is found to induce feelings of happiness or satisfaction, or reduce feelings of unhappiness or dissatisfaction, then it will be apprehended by its subject as valuable in the absence of any countervailing reasons. Less long-windedly, happiness is found valuable as the default condition.

Now this might serve as the cue for a heavy-duty treatment of the relationship between value and pleasure. All that's needed for the argument to follow, however, is to note that the biological program will generate, both quantitatively and qualitatively, immensely more experiences found at once pleasurable and valuable than those characteristic of the neurochemical status quo. The program's therapeutic strategy will eliminate a whole host of states that even today are thought worthless or obnoxious. With time, the correlation between states found valuable and states found pleasurable should get ever closer to 1. So if, first, value judgements are also truth-evaluable, and if, second, subjects were normally capable of reliably apprehending their truth, then the biological program would indeed prove ethically mandatory.

2.10 A Tough-Minded Scientist Replies.

Yet so what? The contemporary critic will not be impressed. Just as not everything that is more desired is more desirable, surely not everything that is more valued is thereby more valuable. Only if the valued were indeed also valuable would the biological program be vindicated in an ethical sense. It can't be, because its defence attempts to derive, or somehow smuggle in, an "ought" from an "is", which is logically impossible. To argue otherwise is to commit the naturalistic fallacy. For is value supposed to be some property of the natural world over and above the ontology sanctioned by physics?

Physical science, the scientific rationalist may freely go on to admit, has not yet definitively settled on the ultimate ontological furniture of the universe. There is plenty of theoretical and experimental work to be done investigating whether its ontological primitives are particles, fields, probability waves, loops, superstrings or whatever. The relationships between these primitives still tantalisingly awaits a complete and unified mathematical description as well. But whatever really exists e.g. macroscopic objects, itself supervenes on mind-independent configurations of these ontologically basic primitive entities, events or properties. Values, on the other hand, are merely mind-dependent subjective fictions. We don't read them off the world, but project them on to it.

The scientistic hatchet-job on the status of objective values is often supplemented with a neo-Darwinian account of their genesis. If one claims that something is illusory, then one wants to explain how and why the illusion occurs. Pro-Darwinian polemicists oblige. What might seem to be eternal moral verities are ritually unmasked by their debunkers as mere instruments of the genes. People's devoutly-held personal convictions, we learn, are just another means by which competing alliances of information-bearing self-replicators - genes - manipulate their throwaway vehicles at one remove to promote their own inclusive fitness. Admittedly, genetic predisposition does not equate with genetic determinism. Sociobiologists, evolutionary ethicists and their ilk aren't claiming that our genes directly code, rather than bias, the development of each idiosyncratic set of cultural values. Yet independently-arising cross-cultural universals e.g. religious and secular incest-taboos, can nonetheless best distally be explained by positing selective pressures which act over many generations to shape our moral fetishes and phobias. We would dearly love to believe that our subjective values are somehow objectively underwritten by the nature of the world, the scientific rationalist concludes, generally in tones which suggest he bears their absence with remarkable fortitude; but they are epistemically unserious verbiage. To believe otherwise is to fall victim to wishful thinking or the toxic mind-rot of New Age mysticism.

2.11 The Selection of Mysterious Reds

I shall now defend a version of value-naturalism, and consequently the objective ethical rationale of the biological program, against this indictment. Is talk of objective values just claptrap? For it is ironic that at a time when the scientifically-informed current of analytic philosophy is witnessing an embarrassed scramble to "naturalise" everything from epistemology to consciousness, any similar bid to legitimate value should still widely be held to commit a logical fallacy. So it will now be shown how, and in what sense, moral judgements can and can't have truth-conditions; and how the existence of objective values could be consistent with the apparently austere ontology of physical science. An analogy is drawn with phenomenal colour. It is argued that, appearances to the contrary, moral judgements in fact report, truly or falsely, a distinctive quality common to the experience of those who avow them. What such judgements express is mind-dependent, and on an identity theory thereby brain-dependent; and thereby value is as much a natural, intrinsic and objective feature of the world as phenomenal redness. The proposition that it is otherwise is unnaturalistic, the legacy of a dualistic perspective which sees mind and its experiential attributes as distinct from the physical world rather than as objectively existing features of it. We don't simply "project" our values onto the world. For we are literally bits of the world itself. Four objections, each on their own apparently decisive, are levelled against this sort of value-naturalist position.

So to begin a value-naturalist defence, it is worth drawing an analogy with, say, redness. On a mind-brain identity theory, redness is a phenomenological property intrinsic to certain patterns of neuronal firing. The presence of light of a particular frequency impinging on the retina, or indeed of any light at all, is neither a necessary nor a sufficient condition for the production of red experience in a subject. When dreaming, for instance, one can inwardly see or instantiate red phenomena. Conversely, when one is awake and in darkness, then a sufficient condition of one's having, say, a brief punctate red experience in front of one's body-image is that the relevant cortical area is electrically stimulated.

On the assumption that one is wholly a part of the natural world, then phenomenal redness, too, is one of the properties of the world. It is predicated of, and appears to inhere in, many macroscopic objects. Yet it is an intrinsic property of certain mind/brain

states, and is not some relational property involving the interaction of light from intrinsically colourless objects and the mind/brain. The presence or absence of red phenomenal experience can be truly or falsely reported by the subject, whether the subject believes it is a property intrinsic to mind-independent physical objects or otherwise.

Given the above, it is worth noting the sense in which redness can, and more importantly can't, be explained within the current conceptual framework of the natural sciences. Natural selection has stumbled upon psychophysical phenomenal colour states. These states are not inherently representational. But natural selection has harnessed them so they now tend, in the awake brain, to track certain causally co-varying patterns in the organism's environment. The capacity to recognise these patterns (simplistically, differential electromagnetic reflectancies of macroscopic objects) bears on the differential reproductive success of the genetic vehicles in which phenomenal colours are periodically instantiated. This explains why such states have been selected. It doesn't explain their intrinsic phenomenal nature. So natural selection doesn't in any but a shallow sense explain states such as redness (or, it will be argued, value). It explains why some such states have been selected rather than others. It doesn't explain why any kind of experience has the phenomenal properties it does. Nor does it explain why experience exists at all. If telepathy had existed, evolutionary psychologists would doubtless offer excellent explanations and mathematical models of why telepaths had been selected over non-telepaths. Telepathy, we would tub-thumpingly be told, could thus be explained "naturalistically", not as some divine gift of God. Yet the phenomenon itself would still be utterly mysterious.

2.12 The Formal Successes Of Scientific Triumphalism

Physics and, derivatively, the rest of the physical sciences can in principle provide a complete account of the natural universe. It is (potentially) complete only in the sense that the mathematical formalism of quantum mechanics is correct and isomorphic to the world. The equations themselves are topic-neutral. The intrinsic nature of the stuff they describe, what "breathes fire into the equations and makes there a world for us to describe" is, as even Hawking concedes, unknown, and perhaps unknowable. What can be known, however, since one is oneself a tiny fragment of the "fire in the equations", is that the experience of phenomenal redness exists as a matter objective fact. This is so even though a (mathematically) complete physics on its own has nothing to say about it.

This should be stressed because in conceptualising the contents of the world, it is tempting to defer, not merely to the unreasonable effectiveness of the equations, but also to one's ill-defined notions of the basic physical stuff those equations describe. And these notions don't include e.g. redness, or tickles, or happiness, or moral values. But, crucially here, the physicists' potential candidates for the status of brute ontological primitives e.g. superstrings or fields etc., are defined, ultimately, in purely mathematical terms. So if particular phenomenal colours, say, were to be identified with the particular numerical values of a set of occipito-temporal cortical fields, this is in no way inconsistent with the physical formalism. Redness would in this case be just one spark of the "fire in the equations". Likewise, if one identifies particular phenomenologically valuable states with a finite set of numerical values of intra-cranial fields, this is likewise consistent with the mathematical formalism. For they too are part of the fire in the equations which makes there a world for us to describe.

Unfortunately, it is all too easy to muddy the ontological issue here by confusing the two senses of the word "subjective". It is the case, objectively, that the world contains subjective, experiential states such as redness with its unique, nameable, but ultimately ineffable what-it's-like-ness. This property may be identified with complicated, occipito-temporal cortical patterns of cortical fields. Redness is a distinctive mind-dependent property. It lacks any mind-independent existence, since neither electromagnetic radiation, molecules nor their macroscopic object patterns are red. This doesn't challenge its objective existence. When one experiences, or is presented with, or instantiates, redness, one can apprehend what colour it is and report the experience, sincerely or otherwise. This judgement has truth-conditions. Since red is mind-dependent it is also, on any mind-brain identity theory, brain-dependent. It is as such an objective property of the physical world. So what judgements of redness express is both mind-dependent and objectively true (or, if one's avowals are insincere, false).

2.13 The Naturalisation of Value

Now moral value itself will be examined. It is going to be suggested that value, and conversely disvalue, are distinctive features literally inherent in the world no less than phenomenal redness; and thus there can be objective, truth-evaluable judgements of value. This property is mind-dependent, hence brain-dependent, hence a natural and objective property of the world. In consequence, the states of mind of our ecstatic descendants are inherently more valuable by their very nature than the relatively worthless psychiatric slumlands of our own era.

Of the finite, potential $10^{1\ 000\ 000's}$ of interestingly different types of conscious state of the human mind/brain, some are subjectively apprehended as experientially valuable and some aren't. Some states seem essentially neutral; some are merely pleasurable but not valued; some are found complex and ambivalent; some involve the mere parroting of received wisdom in the absence of the relevant experience; and the fuzzy boundaries of what the concept of finding something experientially valuable entails are an added complication too. Some valuable qualities strike one as intrinsic to the very nature of (one's emotionally encephalised virtual simulation of) the mind-independent world. Some seem to be local to one's body-image. Yet the presence or absence of any particular mind/brain-independent state of affairs is in principle neither necessary nor sufficient for the experientially valuable states to occur; whereas a necessary and sufficient condition for those experiences is the occurrence of the relevant pattern of neuronal firings.

Once proto-utopian neuroscience can identify the biomolecular substrates of experiential value, or redness, or pleasure etc, it will be feasible to mass-manufacture redness, pleasure or value. Value can be biologically synthesised in extant organisms or in mind/brains-in-vats. [Hence the derisive tag earlier of "biological program for Cosmic Value-Maximisation".] Futuristic vats could contain colours and values in virtue of containing brains. This sounds odd; but no category-mistake is involved.

So analogously to redness, then, value should be construed as a property of a delimited class of mind/brain states. In future it can be both quantified and synthesised. Certain forms of experience are indeed often said to be unquantifiable: happiness is the most commonly cited example. But if particular types of chemical (or perhaps, ultimately, relativistic quantum fields, or modes of vibration of 10-dimensional heterotic superstrings etc) embedded in the relevant neural state, are either identified with, or found to be invariantly positively correlated with, phenomenologically valuable states, then scaling up or down the number and size of the relevant states by the relevant number and

disposition of molecules increases or decreases the level of happiness, redness, value etc in the world accordingly. Problems of vague concepts with fuzzy boundaries, and of ill-defined criteria of usage, complicate but do not change the issue. In an ideal taxonomy of the mind/brain, experiential states would be as quantifiable, and their exact texture as mathematically precisely defined, as any other feature of the natural universe. The notion that what-it's-like-ness can be described by a set of equations is indisputably counter-intuitive; but this is what any scientific mind/brain identity-theory entails. And given such a theory, the biological program can vastly increase the amount of both happiness and naturalised value in the world.

2.14 Four Deadly Objections?

Now for four potentially devastating objections that can be levelled at the position sketched above.

First, when people express value-judgements, they frequently refer to states of the world. They're not alluding to some distinctive quality of their own experience. They may indeed frequently project aspects of their experience onto states of the world. Yet it is the world they are referring to, not their own phenomenology.

Second, surely values can conflict. They are sometimes violently contested. We even go to war over them. If two putatively truth-evaluable judgements of value are mutually contradictory, they can't both be objectively true; or perhaps they don't, and can't, have truth-conditions at all.

Third, by taking value to be an intrinsic phenomenological attribute of certain mental states, the value-naturalist position apparently makes some singularly obnoxious prejudices morally valuable, even immensely so. After all, Hitler found persecuting Jews extremely morally valuable. Given that, by every indication, Hitler was sincere in reporting at least this aspect of his mental states, albeit under another description, then from the value-naturalist perspective persecuting Jews would have to count as valuable: not as valuable as the exalted states alluded to in this paper, admittedly, but morally worthwhile nonetheless. This is surely a pretty conclusive *reductio* of the position. In any case, the above example exposes the argument's internal inconsistency. Hitler's value-judgements contradicted those of his victims. Therefore it is logically impossible for them both to be right.

Fourth, does not the value-naturalist case rest on an illicit equivocation? Not everything that is desired is desirable, a slide from the factual to the ethical. Likewise, surely not everything that is valued is valuable? Even if it were objectively the case that value-judgements obliquely reported, truly or falsely, a distinctive experiential state or family of states, this wouldn't mean that such types of state actually ought to be valued, or that one ought to strive for their maximisation.

The reply given here to these seemingly knock-down rejoinders to the value-naturalist is highly counter-intuitive. For it depends for its key premise on what might appear to be a completely different issue altogether, the nature of what we optimistically call perception; and in particular the falsity of any sort of direct realism. The answer to be given is arguably consistent with several non-direct realist theories other than the one set out below; but the account, and the heuristic fable it contains, is designed to highlight as starkly as possible the falsity of a presupposition common to at least the first three charges above. The position defended here as a basis for the argument to follow is

a radical selectionist account of perceptual experience. It contends that the difference between "dreaming" and "being awake" lies essentially in the mode by which states intrinsic to the mind/brain are selected. The most that the extra-neural environment can ever do is partially select which of a finite menu of mind/brain/virtual world states is instantiated at a given moment. Subjects can never, directly, do more than apprehend their own mind/brain/virtual-world states. The values they appear to find in the mind-independent world are instead intrinsic features of particular states of their own brains. And insofar as future ecstatics are capable of truly reporting this quality of experience, their states are objectively more valuable than anything existing today. So the world really will get better and better.

2.15 Alone Amongst The Zombies.

These rather dogmatic and elliptical pronouncements may first be illustrated by use of the following case study.

There is a rare sleep disorder in which the victims lack the muscular atony which, ordinarily, functionally decouples the bodily musculature from a dreaming brain. This decoupling is in the normal way highly adaptive. For it stops the rest of us from unwittingly acting out our dreams. In the absence of a functional decoupling of the musculature, all manner of dream-scenarios will be acted out. In such circumstances the external vocalisations and other forms of bodily behaviour of the dreamer are uncorrelated, except by chance, with the rest of the world outside the mind/brain.

Within the dreamer's virtual world, however, nothing will seem amiss. The meaning and reference of terms used by the central body-image are grounded purely internally in its pseudo-perceptually apprehended environment. Inside the neural dreamworld, a conscious, unwittingly private language of thought masquerades as public speech. The dreamer's body-image uses it to converse with the behaviourally intelligent homunculi his visual cortex intermittently activates. These noisily animated zombies, and other ostensibly perceptual experiences of macroscopic objects in a macroscopic world, are purely autobiographical. The whole virtual world flickers in and out of existence as its instantiator passes in and out of dreamless sleep. For it is not just the dreamer's non-occurrent beliefs and desires which are dispositional, but the macroscopic dreamworld itself. Its episodes are nonetheless readily reactivated. This is because its features lie latently encoded in the connection and activation weights of the dreamer's brain. The difference between us and a victim of this sleep disorder is that his extra-neural body acts out, obliviously, the actions performed by his body-image internal to the dream; whereas when we are asleep our bodies are effectively paralysed.

Now, counterfactually and for heuristic purposes, imagine a possible world in which this sleep disorder is both chronic and ubiquitous. Dreamers never "wake up". Nor do they have any notion of what such a familiar if ill-understood expression might mean. Natural selection goes to work over millions of years. It differentially favours the genotypes of organisms whose dreamworlds, initially just by chance, serve as though they were akin to quasi-real-time simulations of certain patterns in the extra-neural world. For genetically selfish reasons, each differentially selected genotype spawns an egocentric virtual world. It is a virtual world centred physically and affectively around one focal body-image. More proximate selection of dreamworld events comes into play due to a bombardment of patterned sequences of electro-chemical impulses from various afferent proto-nerves. These extend to what serve to become peripheral transducers in the organism's bodily surfaces. Over the generations, the fitness-enhancing correlations

between the behaviour the extra-neural body unwittingly acts out and macro-patterns in its environment tend to get tighter and tighter.

With the passage of time, many dreamworlds quite regularly become, so to speak, thoroughly undreamlike. Normal infant dream-worlders will learn, over several years, pseudo-public criteria for language use from their virtual mothers. A maturing dreamer may discover that his body-image's surroundings show a good deal of coherence, law-like regularity and even predictability. He may discover that his body(-image) can intelligently manipulate and re-engineer, within sharply constrained limits, aspects of the (neural dream-)world beyond itself. Obliquely and obliviously, dreamworlds will tend in some degree mutually to select each other's contents. With time, the unwitting behavioural by-products of purposeful actions internal to billions of dreamworlds spin off an ever more elaborate material culture. The collective result of these by-products is that the eternal sleepers' host bodies act out the construction of everything from skyscrapers to computer networks, particle-accelerators to jumbo jets. The resultant artefacts enjoy a dreamworld-independent existence. They themselves serve thereafter partially to select what kinds of dreamworld are neurally activated.

2.16 The Perils of Idle Scepticism

Should an overly-lucid dreamer ever doubt the ontological integrity of his particular virtual world, the consequences may be grave. Dreamworlds can be refractory and inhospitable places. His virtual body-image may be mauled by virtual lions or, in a later era, knocked down by a virtual bus. Thanks to millions of years of selective pressure, such agonies correlate highly with parallel, mind-independent events befalling the organism whose skull encloses the dreamworld brain. So any genes notionally predisposing to such idle philosophers' fancies tend not to be passed on to the bodily vehicles of potential baby-dreamworlds. Instead, each dreamer strives to re-order his emotionally encephalised world so that its unsuspectedly mind-dependent states more nearly match his desires.

Some dreamworlds are chaotic and schizoid; some are seemingly well-ordered and amenable to quasi-scientific investigation; some are happy and suffused with spirituality and magic; and some are violent and nightmarish. None of these gargantuan psychochemical extravaganzas is inherently about anything external to itself on the other side of its skull. Yet evolution has differentially selected genes which predispose to the self-assembly of a very particular range of phenotypical dreamworlds. These are the world-phenotypes which serve as effective vehicles for the propagation of more copies of the genes that made them. One of the properties of a successful vehicle is that periodically some of its patterns causally co-vary, albeit on a highly selective basis, with other patterns beyond itself.

2.17 The Price Of Inner Demons

How is this relevant to a value-naturalist defence of an objective warrant for the biological program? The fable's significance may be illustrated by envisaging a counterpart to Hitler, say, in the dreamworld scenario. In his dark and sinister virtual world, his body-image fights against terrible inner demons/neuronal firings. He spends

his whole life pitted in a struggle to exorcise once-and-for-all their malevolent and conspiratorial presence. The evil occipito-temporal homunculi lurking beyond his somato-sensory body-image are of course mindless phantoms. But their hostile intent appears frighteningly obvious to their host. Tragically, Hitler's dreamworld brain is fully coupled to the bodily musculature of the organism which houses such nightmarish neurochemical patterns. There is no muscular atony to prevent the microcosmic horror story from being acted out in the mind-independent macrocosm by the extra-neural body. Natural selection has ensured that many types of event in his dreamworld causally co-vary, albeit in a grotesquely selective manner, with the wider world, its organisms and the dreamworlds they host. In consequence, over fifty million people die in a brutal war.

Now this fable is all very well as a thought-experiment, it may be said. Even in our own world, there are rare and tragic cases of people who blamelessly and unwittingly kill their partner while asleep, whether during "night-terrors" or in the course of an exceedingly violent dream. But the real Hitler wasn't asleep. He was fully awake and acted quite deliberately in full knowledge of what he was doing. He perceived real, flesh-and-blood, sentient people. They were wholly innocent of the monstrous crimes he imputed to them.

And herein lies the crux. If real-world Hitler did directly apprehend or perceive his victims, or alternatively if certain neurochemical events in his mind/brain/virtual-world were, somehow, inherently about Jewish people in the world outside, then the argument shortly to be presented is false. If, on the other hand, Hitler was wrestling with horribly emotionally encephalised inner demons, apparitions of his own (involuntary) creation whose foul behaviour really did blight his early virtual world, then his behaviour in trying to banish such sources of negative value amounted to an epistemic rather than evaluative failure. Likewise today, in billions of other egocentric virtual worlds, desperate and often ineffectual attempts are being made by each genetic host's central body-image to exorcise all kinds of obnoxious phenomena. Unfortunately, in the absence of the biological program and the presence of naive realism, the net results are frequently tragic.

In the case of Hitler, profound sources of objective experiential "disvalue" did indeed neurochemically transmit and present themselves to the functional modules mediating his sense of self and neural body-image. It wasn't the case that he somehow "projected" such experience onto his virtual world; instead that quality of experience was intrinsic to it. Natural selection ensured that Hitler, in common with all but a few philosophically and scientifically-minded humans, was implicitly a naive realist about a perceptual world. So when he apprehended great evil, a quality of experience located in what he couldn't know was only his emotionally malencephalised virtual world, he tried to destroy it in the only manner he knew how. By his lights, he was trying to make the mind-independent world a better place. Had he been a brain-in-a-vat, he might temporarily have succeeded. Tragically, he wasn't; and a mere epistemological error turned into a moral catastrophe.

2.18 Can We All Be Really Good?

Now if the human predicament were akin to that of a dreamworlder, a very big and controversial "if", admittedly, then the following answers may be given to the four objections to value-naturalism levelled earlier.

First, yes, people certainly believe many of their value-judgements refer to the world and its properties rather than to some distinctive quality of their own experience. But both the philosophy of perception and quantum mechanics suggest that what a person treats as the mind-independent world - and to whose properties he linguistically refers - are toy, data-driven simulations his mind-brain is running. If so, then he is referring in a direct way only to aspects of his own neural experience in another guise. What his value-judgements express is still an objective property of the natural world. But it is mind-dependent. Experiences found valuable, whether by brains-in-skulls or futuristic brains-in-vats, have a distinctive, nameable, but ineffable what-it's-like-ness about which physical science has nothing to say.

Second, people's value-judgements can mutually contradict each other only if they succeed in referring to the same thing. Hitler's internally-issued value-judgements couldn't really contradict those of his extra-neural body's inadvertent victims. Those same judgements accurately reflected the character of the emotionally encephalised bestiary of monsters that populated his mind/brain; and against whose machinations he fought, at terrible cost.

Third, what is morally wrong on a consequentialist ethic is the effect of the unwitting behavioural spin-offs of Hitler's attempts to extinguish his inner foes. He wasn't mistaken to find certain phenomena obnoxious, sources of profound objective "disvalue". *Mein Kampf* is testimony to their horrible phenomenology. He just mislocated their distinctive properties and origin as external to his composite self. The effects were of course catastrophic.

Now to what extent the dreamworld fable above does capture an aspect of the human predicament is, to say the least, controversial. Aside from certain details included for reasons of expository convenience, I would argue that the account is empirically indistinguishable, at least, from more familiar approaches to perception. To explore in any depth, however, the perceptual and semantic minefields into which the question leads, not to mention the paradoxes of self-reference it might seem to entail, would take us too far afield. The account does nonetheless offer one programmatic way to naturalise value, albeit at a price that may be considered too high for comfort.

2.19 Equivocal Values

The fourth charge was one of equivocation. The valued is being confused with the valuable. Even if it is granted, the charge continues, that value-judgements are true or false reports of a distinctive type of neurophenomenological state, that state itself is, as the term suggests, just that: truth-valuelessly phenomenological. Finding an experience morally good or bad in such a sense doesn't carry any logical implication that one should objectively do anything about it. Hence, whatever its instrumental merits, the claim that the biological program advocated here is ethically mandatory is untenable if construed as expressing an objective truth. Yes, executing the biological blueprint would vastly increase the number and intensity of states found phenomenologically valuable; and yes, it would abolish altogether states that aren't. But value-judgements, and the qualities of experience they describe, are like tickles. They exist, and they may make you want to do something about them. Yet they don't refer to anything beyond themselves and they don't logically mandate any course of action.

I would argue that properly understood there is no equivocation. We happen to live in a universe whose ontology includes literally valuable experiences in the same way as it

contains literally painful experiences, visual and auditory experiences, feelings of irritation or obligation or indignation, and a teeming profusion of other forms of what-it's-like-ness most of which remain so far completely nameless. So the universe really does contain phenomena that are, literally and intrinsically, valuable. The utilitarian ethic championed here, and the biological program it instrumentally dictates, leads ultimately to the amount of intrinsic value as well as happiness in the universe being maximised; and all sources of negative value extinguished.

It will then no doubt be asked, perhaps somewhat impatiently as well as sceptically: but is an experience found really valuable *really* valuable? Why couldn't it just seem to be valuable? Yet one wouldn't, and couldn't, sensibly ask if an experience found really painful was *really* painful. One can apparently imagine a universe without values, in the same way as one can apparently imagine one without pains or pleasures or redness. But for reasons we admittedly don't understand, we don't live in that sort of Universe. We live in a Universe where some things intrinsically matter and have positive or negative value. If our image of a respectable physicalist ontology can't cope with the objective fact such modes of what-it's-like-ness exist, then we are misinterpreting what the formal mathematical description of the world is telling us.

Now perhaps a value-nihilist can sincerely deny having any such quality of experience. The nihilist can ask why should (s)he value value, whatever *that* might be. Yet this scepticism doesn't impugn the existence of value, any more than the status of pain is compromised by rare cases of people congenitally insensitive to it. The relegation of either kind of experience to some kind of ontological *demi-monde* is unwarranted and should be rejected.

This objectivity doesn't entail that valuable experiences can have, as distinct from simulate, a type of mind-transcendent, truth-evaluable "propositional content" over and above their intrinsic phenomenology which somehow manages to alight on properties of the mind-independent world. But then there are desperately hard problems in the context of a naturalistic world-picture of explaining how any other spatio-temporal electrochemical event or episode of experience, whether deemed cognitive or otherwise, could literally have abstract propositional content either. Worlds where they don't can apparently be empirically indistinguishable from ours - and a lot less ontologically fishy. A lot of the time, one just has to cross one's fingers, whistle in the dark, mix one's metaphors, and try and pretend otherwise.

2.20 Good Vibrations: The Value Of String

Russell once observed that "Ethical metaphysics is fundamentally an attempt, however disguised, to give legislative force to our own wishes." Perhaps he is right. Mixing up prediction and prescription is usually a recipe for confusion. Attempts to ground the post-Darwinian project - or any other moral enterprise - in something more exalted than the pleasure-pain principle may simply be spinning a fantasy world of self-deception. Perhaps talk of the moral goodness of eradicating suffering - or any other kind of moral discourse - is merely idle opinion: just a lot of high-falutin noise amid the digital babel of cyberspace.

The traditional-minded scientific rationalist, for one, will surely be unmoved. It will be claimed that the world's [allegedly inherently] valuable and valueless experiences as touted in this chapter are "really" "just" something else: patterns of neuronal firings, the differential modes of vibration of superstrings (or whatever) with which they are posited

to be physically identical. Yet this is sophistry. The reductionist argument can be turned on its head. Presumably certain modes of vibration of superstrings etc are "really" "just" valuable experiences. This isn't very illuminating. Whether, why, how, and with what significance, different values of what-it's-like-ness should be mapped on to, or read off, the different numerical values of solutions to the equations of physics are deeper questions altogether, and not ones that can be explored here. They may all just be glorified tickles; or they may not: we simply don't know.

Instead, this section may be concluded with a quick restatement of the plot so far. The biological program holds out the promise that, within a few millennia at most, states of conscious mind everywhere will be by their very nature more enjoyable than anyone alive today can imagine. Our hereditary neurological pleasure-deficit stops us getting a grip on what biotechnology can genetically engineer. In (at the very least) an empirical sense, implementing the post-Darwinian program can fill the world with valuable experiences. They will be enjoyed by human, non-human and post-human beings. Post-Darwinian modes of experience are likely to be of a diversity, profundity and liquid intensity that goes beyond anything accessible to the impoverished hunter-gatherer-evolved imagination. All the moral ills identified by contemporary secular value-systems can be rooted out for ever. Suffering will one day become physically impossible. This all sounds rather bombastic; but the strategy is biologically feasible as a species-project should we choose to pursue it.

Whether maximising the valued in the world amounts, in practice and/or theory, to maximising the intrinsically valuable in the world is another, and harder, question. There is, I have argued, at least a *prima facie* case that it does. We may one day live in a Universe whose equations describe something which is intrinsically valuable by its very nature.

Chapter 3: WHEN?

"The mass of men lead lives of quiet desperation"
(Thoreau)

3.0 Our Emotional Future

Set aside for now the practical merits or ethical urgency of the abolitionist project. What grounds are there for *predicting* that suffering and malaise will be replaced by gradients of genetically preprogrammed well-being? *When*, if ever, might paradise-engineering become practical politics?

If any such Post-Darwinian Transition does occur, then the revolution will happen only once. It will never be reversed. There won't be any going back to the old Darwinian order after it transpired its successor wasn't as wonderful as advertised. For in practice it will be far better.

The prospect of such invincible bliss may seem very distant back here in the biological Dark Ages. Yet it shouldn't be. Even now, most of us try to manipulate our states of mind via chemical means. We just aren't very good at it. Throughout history, humans have tried to alter their consciousness via the use of a variety of natural agents. Arbitrary, and highly selective, proscription and persecution by the ruling elites has failed to prevent people from experimenting with psychedelics and mood-enhancers alike. By the turn of the twenty-first century, perhaps $400 billion or 8% of world trade was in illicit drugs.

Recreational agents which are legal and socially sanctioned by respectable society aren't, of course, popularly viewed as drugs at all. The nicotine addict and the alcoholic don't think of themselves as practising psychopharmacologists; and so alas their incompetence is frequently lethal.

Is such incompetence curable? If it is, and if the abolitionist project can be carried forward with pharmacotherapy in advance of true genetic medicine, then a number of preconditions must first be in place. A necessary and sufficient set could not possibly be listed here. It is still worth isolating and examining below several distinct yet convergent societal trends of huge potential significance.

1. First, it must be assumed that we will continue to seek out and use chemical mood-enhancers on a massive, species-wide scale.
2. Second, a pioneering and pharmacologically (semi-)literate elite will progressively learn to use their agents of choice in a much more effective, safe and rational manner. The whole pharmacopoeia of licensed and unlicensed medicines will be purchasable globally over the Net. As the operation of our 30,000 plus genes is unravelled, the new discipline of pharmacogenomics will allow drugs to be personally tailored to the genetic makeup of each individual. Better still, desirable states of consciousness that can be induced pharmacologically can later be pre-coded genetically.
3. Third, society will continue to fund and support research into genetic engineering, reproductive medicine and all forms of biotechnology. This will enable the breathtaking array of designer-heavens on offer from third-millennium biomedicine to become a lifestyle choice.
4. Fourth, the ill-fated governmental War On (some) Drugs will finally collapse under the weight of its own contradictions. Parents are surely right to be anxious about many of today's illegal intoxicants. Yet their toxicity will no more prove a reason to give up the dream of Better Living Through Chemistry than the casualties of early modern medicine are a reason to abandon contemporary medical science for homeopathy.
5. Fifth, the medicalisation of everyday life, and of the human predicament itself, will continue apace. All manner of currently ill-defined discontents will be medically diagnosed and classified. Our innumerable woes will be given respectable clinical labels. Mass-medicalisation will enable the big drug companies aggressively to extend their lucrative markets in medically-approved psychotropics to a widening clientele. New and improved mood-modulating alleles, and other innovative gene-therapies for mood- and intellect-enrichment, will be patented. They will be brought to market by biotechnology companies eager to cure the psychopathologies of the afflicted; and to maximise profits.
6. Sixth, in the next few centuries an explosive proliferation of ever-more sophisticated virtual reality software products will enable millions, and then billions, of people to live out their ideal fantasies. Paradoxically, as will be seen,

the triumph of sensation-driven wish-fulfilment in immersive VR will also demonstrate the intellectual bankruptcy of our old Peripheralist nostrums of social reform. Unhappiness will persist. The hedonic treadmill can't succumb to computer software.

7. Seventh, secularism and individualism will triumph over resurgent Islamic and Christian fundamentalism. An entitlement to lifelong well-being in this world, rather than the next, will take on the status of a basic human right.

There are quite a few imponderables here. Futurology is not, and predictably will never become, one of the exact sciences. Conceivably, one can postulate, for instance, the global triumph of an anti-scientific theocracy. This might be in the mould of the American religious right; or even some kind of Islamic fundamentalism. Less conceivably, there might be a global victory of tender-minded humanism over the onward march of biotechnical determinism. It is also possible that non-medically-approved drug use could be curtailed, at least for a time, with intrusive personal surveillance technologies and punishments of increasingly draconian severity. Abetted by the latest convulsion of moral panic over Drugs, for example, a repressive totalitarian super-state could institute a regime of universal compulsory blood-tests for banned substances. Enforced "detoxification" in rehabilitation camps for offenders would follow.

These scenarios and their variants are almost certainly too alarmist. Given a pervasive ethos of individualism, and the worldwide spread of hedonistic consumer-capitalism, then as soon as people discover that there is no biophysical reason on earth why they can't be as happy as they choose indefinitely, it will be hard to stop more adventurous spirits from exploring that option. Lifelong ecstasy isn't nearly as bad as it sounds.

3.1 Hedonism After The War

So as an illustration of at least one plausible run of events leading to an adoption of the biological strategy, it is worth considering the consequences likely to ensue when western state governments finally abandon their ill-starred and intellectually incoherent War Against Drugs. This retreat might not seem inevitable. Here at least, however, it will be assumed that the freedom to control one's own states of consciousness can't be usurped by government indefinitely. State mind-control measures may relax in the face of, first, an ascendant libertarian and free-market ideology; second, a rising younger generation of experienced illegal drug-takers, averse from being criminalised and scornful of the hypocrisies and double standards of the older generation; and, third and not least, the unparalleled and uncensorable information explosion across the Internet on the detailed practicalities of how to synthesise and enjoy psychotropics of every description.

Decriminalisation, first *de facto* and then *de jure*, and subsequent legalisation will not entail a straightforward abdication of state control. On the contrary, the state will intervene from motives of fiscal self-interest and paternalistic responsibility in the distribution process. The manufacture and supply, and certainly the quality and purity, of psychotropics will be licensed, guaranteed and regulated. This will reclaim a multi-billion pound sector of economic life from organised and disorganised crime. It will further allow a drastic and politically expedient reduction in direct taxation. It should also eliminate some of the toxic adulterants common in street-drugs. Thousands of newly decriminalised drug-users will re-enter mainstream civil society. More intelligent drug-education, and the social institutionalisation of previously illicit forms of drug-use, will further contribute to the harm-reduction process.

Yet this is to paint a dangerously rosy picture of the consequences of legalisation. Desirable as it may be to stop criminalising and even locking up a growing percentage of the younger generation, notably in the USA, the far-reaching social and medical problems stemming from ill-informed drug-use will remain. For a start, an enormous and perhaps unquantifiable number of (currently) illicit and licit drug-users alike are, in effect, self-medicating. They don't like their own stressed, anxious or depressive consciousness the way it is otherwise. So they pursue what seem the only remedies realistically on offer. Their choices aren't altogether surprising. Other nominal mood-brighteners actually sound *depressing*. State-endorsed "antidepressants" are available solely on prescription. They get doled out by severe (wo)men in white coats. Officially, in any case, all such agents are of potential therapeutic value only to those deemed by the medical authorities to be mentally ill. This isn't a role or a label most people would willingly adopt. Such a severe image-problem means that millions of people who would otherwise benefit are missing out on some of the most worthwhile products of medical science.

They are more likely to turn instead to drugs with a very different image. The main shortcoming of the widely-used illegal euphoriants, such as cocaine and the amphetamines, is less that they are physically dangerous - they collectively kill only a minuscule fraction of 1% of the numbers carried off each year by the two legal state-sanctioned killers, tobacco and alcohol - but that receptor re-regulation ensures their long-term effects are very nearly the opposite of those for which their users take them in the first instance. The medical authorities, meanwhile, maintain a convenient fiction about all present or potential clinically approved mood-boosters. The official line is that delayed-action "antidepressants" have a negligible effect on "normal" people unless they are "really" depressed; or unless a "pathological" manic euphoria is induced, which would need to be medicated in turn. Moreover the medical profession's otherwise healthy caution about the risks of polypharmacy needlessly restricts the formulation and enjoyment of some very beautiful psychoactive cocktails of potent life-enrichers.

Once legalisation of currently banned and controlled drug-groups occurs, there will nonetheless be tremendous pressure on the state to sponsor the research, development and marketing of mood-brighteners for the wider population. These will certainly be safer and more effective than the trashy street-drugs currently in circulation. Initially, many traditionalists will undoubtedly continue to practise and proselytise Total Abstinence in the spirit of the Just Say No school of thought. At the other extreme, a small minority of thrill-seekers in search of the ultimate high will probably still take crack and the like with predictably disastrous consequences. The human mind/brain isn't capable of sustaining such intensities of pleasure indefinitely without substantial design-enhancements not yet on offer. For millions of more responsible and psychopharmacologically educable people, however, the possibility of Better Living Through Chemistry will prove irresistible. They will judiciously pick-and-mix from a gamut of mood-brighteners, smart drugs, serenics, aphrodisiacs, anti-ageing drugs and other agents drawn from life-enhancing categories not yet invented. Already there are at least a few tentative indications that humanity's psychopharmacological stone-age is starting to draw to a close.

One class of mood-brightener appealing to the more temperamentally cautious will be psychoactive agents with a therapeutic window: safer, more potent and much more rewarding variants of a drug like nicotine. With nicotine, the brain very efficiently modulates how much or how little it wants, and within a quite narrow range, to achieve optimal effects. Other designer-drugs will deliver dose-incremental benefits, but on a delayed-reward basis of receptor re-regulation. Cautious polypharmacy, too, in the shape of combined dopamine and serotonin agonists and re-uptake blockers looks especially promising. Rather than spending months in exorbitantly expensive talk-therapy with ill-defined goals and benefits, people will be able to take professional specialist advice on customising and fine-tuning the psyche. Dysfunctional traits of personality can then be psychochemically retailored. The gap between idealised self-image and uncomfortable reality will shrink. Within a few generations at most, the role of a national health service

may be to keep people happy as well as healthy: an anachronistic distinction that may gradually outlive its usefulness.

3.2 On Why We Need Bigger Drug Pushers

Presently-illegal drug-use might be styled "pressure from below". Pressure "from above" will come from the giant, multi-billion dollar pharmaceutical companies. So long as the official dogma of hard-line therapeutic minimalism dictates that there should be no clinically-sanctioned drug-use in "healthy" people, a lot of very interesting drugs indeed aren't going to make it to the marketplace. Hence while medico-political orthodoxy holds, great commercial advantage accrues to the manufacturers if as much of everyday human life as possible can be pathologised. For it can then be mass-medicated with patentable drugs, preferably on a long-term basis. The intimate relationship the industry enjoys with the medical profession and its trade press will generally help the drug firms to communicate their views effectively. If ill-specified and ubiquitous conditions such as age-related memory-loss can be granted formal diagnostic respectability, they can then be combated with cholinergic boosters and other forms of cognitive enhancer. The use of such drugs can subsequently diffuse into the wider population. They may be used by student examinees or ambitious executives, for instance, to gain a competitive edge over their drugless contemporaries. And there are much more exciting agents in development than the (not especially) smart drugs currently on offer.

Very large numbers of young people today are at least in shallow, physical terms tolerably fit. Possibly much of a sizeable and potentially lucrative market will be allowed to remain untapped. If, on the other hand, it were to be (rightly!) medically acknowledged that statistically normal spells of youthful anxiety, lassitude, sub-clinical depression and angst were a colossal health problem, then the pharmaceutical industry and the new end-users of its products would in their different ways both be much better off. For while the cynic may entertain doubts about the motives of the drug companies and their marketing techniques, it should be emphasised that the actual consequences of a creeping medicalisation of the human condition are often to the good. Suspicious as many commentators may be of such newly-labelled conditions as "dysthymia" and the like, diagnostic categories of this nature reflect submerged misery and malaise on an uncharted scale. Such states merit treatment even by today's dismally low minimum criteria for emotional good health. The problem is not that we are medicalised too much, but too little; and not very well.

Next in line for medicalisation might be a hitherto little-acknowledged syndrome christened, say, "hypo-hedonic disorder" or some term of equally portentous gravity. This label might widen the diagnostic drag-net to another 30-40% of the population. All of them, the drug companies will rightly feel, deserve the best treatment money can buy. Slowly, the species-typical emotional baseline will creep upwards, until takeoff to self-sustained felicific growth finally triumphs.

3.3 Good Code Gets Better

One as yet fundamentally under-medicalised territory is the human genome. Several thousand reasonably well-defined genetic disorders have currently been classified. Aside

from a few tentative clues, however, the genetic basis of medically-certified mood-disorders has not been properly defined. We do not, in any case, have genes "for" happiness, anxiety, depression, and so forth, in any but the following sense. The presence or absence of certain genes with certain other genes makes it statistically more or less likely that an organism will be happy, anxious or depressed in a given type of environment and in a given range of circumstances. The statistical margin of advantage, however, does not need to be very large for natural selection to get to work.

Natural selection isn't going to be around for that much longer. The human genome will have been mapped out within the next few years. It will take several decades more to discover which combinations of genes code for structures and proteins that, other things being equal, will depress mood and well-being in childhood and later life. They can in time be taken out or repressed. Those which have multiple complex effects, and can't readily be dispensed with, can be replaced with variant alleles of the same gene whose actions are more benign. Conversely, genes associated with hyperthymia i.e. the relatively uncommon mental abnormality of feeling consistently happy in the absence of exhausting (hypo-)mania, can be introduced, reduplicated and vigorously expressed in progressively larger numbers of the population and their germ lines. The spread of hereditary hyperthymia should portend a comprehensive reworking of the genome. Recoding the genetic bases of mind, body and virtual worlds will conceivably take hundreds of generations and more. A lot will depend on how long it takes to cure the ageing process. The end of obligatory mortality will force a halt to the traditional breeding free-for-all. Genome redesign is sure to become ever more daringly ambitious. Old-fashioned electrodes in the pleasure centres may be aesthetically distasteful. But they are a great deal simpler.

Again, it will be the big companies, this time in the biotechnology sector, who will initially be driving the psychogenetic revolution forward. A huge potential market exists for their products. In the short-term at least, real moral dilemmas will have to be confronted. These will be not unlike the dilemmas posed today by the existence of fundamentalist parents who deny their child a lifesaving blood-transfusion. Future parents who decide, whether in deference to God or Nature, to decline gene-therapy for a child they know will likely grow up depressive, for example, may be open to accusations of child-abuse. Responsible parents, on the other hand, will want to get their kids the best happiness money can buy.

Accounts like this inevitably sound cold, technocratic and Brave New Worldish. It should be recalled that the developments they describe should avert suffering on a scale which a single mind cannot possibly comprehend; and make a lot of people blissfully well.

3.4 The Death-Spasms Of Peripheralism

A further reason for predicting the abolition of the capacity for negatively-charged experience is superficially very different. It stems from a speculation on one indirect effect of omnipresent, multi-modal and immersive Virtual Reality software. This potential multi-trillion dollar industry is here posited to dominate social, personal, artistic and economic life after the first century or two of the next millennium.

An assumption of this paper has been that (post-)humanity will eventually break free from the Tyranny Of The (traditional, gene-manipulated) Intentional Object. Our genes have ensured that emotion is so pervasively encephalised that we have convinced

ourselves that happiness can only be achieved, and frustration avoided, by chasing after a crazy patchwork of intentional shibboleths of no inherent value whatsoever. Humans have fought thousands of unbelievably vicious wars against each other in consequence. In a sense, our whole culture is a monument to the Peripheralist strategy; and a very unpersuasive advertisement for it too. One of the few things that might convince us, as a species, that Peripheralism can't bring lasting happiness would be for us to see what it would be like if everything in the environment were perfectly as human-beings might wish, and for our most impossible fantasies and desires to be realised. Of course it has always been natural to assume that such a notion was an idle pipedream. Even a Roman Emperor couldn't get everything he wanted.

An all-pervasive network of virtual realities, however, will enable everyone to have their intentional objects of desire fulfilled, and at minimal cost. Interactive or solipsistic, artistic masterpiece or pornographic wish-fulfillment, an ever-growing software library of virtual worlds will enable everyone to have their dreams come true.

3.5 And Yet It Still Grinds

It won't, *mirabile dictu*, make most of us much more happy for very long. The hedonic treadmill will still grind. A revolution of rising expectations will eventually lead people to expect, as of right, to enjoy and enact any set of perceptions and narrative structures they choose. They'll expect to do so in virtual worlds with laws and body-images of their own choosing. In the absence of a decent mesolimbic repair-job, boredom, angst and other dormant negativities will periodically surface. They'll sour the ostensibly perfect idylls and utopias. For ironically, a mass migration into virtual worlds might come to represent Peripheralism's final fling. Only total control of one's notional surroundings may be enough to convince many people of the futility of pure environmental manipulation if their goals include lasting happiness and fulfillment.

A symbiotic union of biologically programmed euphoria and mature virtual reality software engineering, however, is an awesomely good prospect. In fact, such a hybrid could furnish one explanation, however unlikely, of the absence of any signs of intelligent life elsewhere in the galaxy. For if a species acquires the sophistication to generate to order any possible experience at all, whether hedonic, perceptual or other modes of being altogether, then the motivational incentives to choose the inconvenient kinds of experience involved in (non-virtual) space-exploration etc are somewhat diminished. Indeed, since VR is probably less difficult to accomplish than interstellar flight, the very possibility of vulgar physical star-hopping may just never arise.

3.6 The Technology of Shop-Soiled Utopias

Two problems with the VR-scenario in general are worth briefly discussing. The first is technical. It may be alleged that realistic VR won't happen, contrary to the above, because it's too difficult. Serious interactive virtual world-making would require processing power several orders of magnitude faster than anything available today. In allusion to the power of the human visual apparatus, it has been remarked that Reality is 130 million polygons a second. Barring a revolution in portable quantum supercomputing, this is simply unattainable by artificial means.

One response here is simply to cite Moore's Law: processing power has been roughly doubling every other year, and its tempo shows no sign of slackening off. This leads to some dizzying projections. Moreover 130 million polygons a second are probably wasted on a lot of people. The kinds of fantasy scenarios that stir our deepest emotions, and those which might supposedly make us most happy, are mainly of a rather uncomplicated kind. They tend to appeal to relatively primitive appetites in settings where finely-wrought visual subtleties are less than crucial. For even in our fantasies we enact parodies of genetic fitness-maximising behaviour.

It is true that the time-scales projected here for the development of the more sophisticated sorts of virtual world are vague. They may even be wildly off-beam. Yet the dates, in common with all the other rough chronologies suggested in this manifesto, are but a twinkle in the eye of eternity; vitally important to individual members of the few generations around the Transition epoch, but a minor detail in the history of life on earth and beyond.

3.7 Living In The Real World

A second reason for doubting that omnipresent virtual realities will ever lead to the demise of Peripheralism is that, as the name suggests, they aren't real. A sense of authenticity, or any notion that one's actions really matter, will be lacking in even the most startlingly lifelike creations. They may sometimes be entertaining, it will be suggested, but even the greatest masterpieces of virtual reality software will never displace Real Life. Interacting with real flesh-and-blood people endowed with real feelings, it is claimed, will always take precedence over responding to mindless phantoms conjured up by machines.

As a peg to hang one's discontents on, the unreality of the impostors one might meet in virtual paradise (or in one's capacity as virtual Lord of the Galaxy, Casanova of the Cosmos or whatever) might indeed temper one's enjoyment. Admittedly, those who by contemporary standards have relatively benign genes and psychochemistry may not be unduly troubled. After all, when watching plays or movies, or when reading a good thriller, one isn't usually perturbed by the fictional status of the protagonists. Many enthusiasts find even today's crude electronic games gripping for long periods; and when Sega's *Sonic* arrived, I recall feeling pangs of jealousy at being unable to compete for people's attention with a mere electronic hedgehog. Moreover even disbelievers in direct-realist stories of perception seldom seem to be smitten down by the awful sense of loneliness and isolation that life behind the veil can induce.

Yet even if any serious malaise in virtual paradise is confined to the temperamentally angst-ridden, there is a limit to how far perceptual-style manipulation can go. When, as a species, we can generate by artificial means essentially any perceptual experience or scenario at all from the finite selection of states theoretically on offer, then it is just about possible, I suppose, that what used innocently to be called progress will in effect come to a stop. The future beyond the next millennium might just consist of people permuting variations of the same old types of perceptual and pseudo-perceptual experience. On no particularly knock-down evidence, however, I think it more likely that we will want to access and explore the modes of consciousness accessible only by more radical reconfigurations of neurochemistry, beyond the influence of mere surface transducers. The limbic system will be a very obvious early target. And when the gene-driven biochemistry of nastiness has been unravelled and purged from our minds, it is hard to see us ever putting it back.

Chapter 4: OBJECTIONS

4.0 "Happy experiences, and the very concept of happiness itself, are possible only because they can be contrasted with melancholy. The very notion of everlasting happiness is incoherent."

Some people endure lifelong emotional depression or physical pain. Quite literally, they are never happy. Understandably, they may blame their misery on the very nature of the world, not just their personal clinical condition. Yet it would be a cruel doctrine which pretended that such people don't *really* suffer because they can't contrast their sense of desolation with joyful memories. In the grips of despair, they may find the very notion of happiness cognitively meaningless. Conversely, the euphoria of unmixed (hypo)mania is not dependent for its sparkle on recollections of misery. Given the state-dependence of memory, negative emotions may simply be inaccessible to consciousness in such an exalted state. Likewise, it is possible that our perpetually euphoric descendants will find our contrastive notion of unhappiness quite literally inconceivable. For when one is extraordinarily super-well, then it's hard to imagine what it might be like to be chronically mentally ill.

Here's a contemporary parallel. It's possible to undergo, from a variety of causes, a complete bilateral loss of primary, secondary and "associative" visual cortex. People with Anton's Syndrome not only become blind; they are unaware of their sensory deficit. Furthermore, they lose all notion of the meaning of sight. They no longer possess the neurological substrates of the visual concepts by which their past and present condition could be compared and contrasted. Our genetically joyful descendants may, or may not, undergo an analogous loss of cognitive access to the nature and variant textures of suffering. Quite plausibly, they will have gradients of sublimity to animate their lives and infuse their thoughts. So at least they'll be able to make analogies and draw parallels. But fortunately for their sanity and well-being, they won't be able to grasp the true frightfulness lying behind any linguistic remnants of the past that survive into the post-Darwinian era. Such lack of contrast, or even the inconceivability of unpleasant experiences, won't leave tomorrow's native-born ecstatics any less happy; if anything quite the reverse.

It's true that a world whose agents are animated by pleasure gradients will still have the *functional equivalent* of aversive experience. Yet the "raw feel" of such states may still be more wonderful than anything physiologically possible today.

4.1 "The scenarios mapped out in this paper are impracticable. None of them would work in reality. The human brain is too complex to be hardwired for lifetime bliss. Nature, in her wisdom, would ensure that some complicated cycle of feedback-inhibition eventually kicked in. This would restore more equable and subdued states of mind."

Any attempt to hardwire into the cerebral cortex a functional understanding of the Theory of General Relativity, say, or perhaps to set "by hand" the neural connections and activation weights mediating an appreciation of Shakespearean tragedy, would presumably defeat all but the most utopian neuroscience. Such virtuoso feats won't be necessary. The physiological roots of affective states lie mostly deep within the phylogenetically primitive limbic-system. They aren't "merely" limbic; this is to miss the evolutionary significance of their encephalisation. The predictive reward value of different sensory cues, for instance, is encoded by the orbitofrontal cortex as well as the amygdala. Yet the neural basis of our emotional life is still incomparably simpler than the plethora of cognitive processes they penetrate. For sure, the functional pathways of our emotions are complicated to twenty-first century eyes. Yet they should prove tractably so. Just as we can, with horrible cruelty, administer drug-cocktails that induce unremitting despair - this is sometimes done in exploring animal "models" of depression - so we can crudely, and some day exquisitely, polarise mood in the opposite direction.

It will be recalled that the monoaminergic neurons, peptides and endorphins that underlie the emotional tone of experience play an essentially modulatory role. They are not individually directed on notional site-specific representations pre-coded by genes. If the receptors, enzymes, cytoplasmic proteins and genetic switches in one's ventral tegmental area and nucleus accumbens are suitably reconfigured, and if these wonderful cells continue to fire away vigorously, then one is going to be outrageously happy indefinitely. Natural selection has no powers of foresight and anticipation with which to frustrate us. Nature isn't waiting to take its revenge. Given a richer dopaminergic and mu opioidergic innervation of the neo-cortex, then the focus of future ecstatic happiness will be on a shifting and unpredictable panorama of intentional objects. The potential complexity and variety of those objects - i.e. what one will be nominally happy "about" - is indeed staggering. Yet when each fleeting neocortical coalition is blissfully innervated from "below", every one of them can be a focus of delight. Life will always be exhilarating, and the fun simply won't stop. For the hedonic treadmill will have been genetically dismantled for ever.

4.2 "If we were always elated, we'd suffer the same fate as intra-cranially self-stimulating laboratory animals. We'd starve, or die of general self-neglect. Both physical and psychological pain do more than promote the inclusive fitness of genes. For the most part, they protect the individual organism from harm too. If a regime of universal happiness were attempted, we'd never want to have sex and reproduce. Therefore we'd become extinct as a species."

A project geared to crude biological pleasure-maximisation alone could well undermine the autonomous survival-skills of its participants. In a comprehensively automated, computerised, robot-served civilisation, this supposed incapacity wouldn't in the long run pose a particular problem. Moreover it is only certain types, not intensities, of pleasure which are incompatible with efficient bodily self-maintenance. Pragmatically, however, worry over the incapacitating effects of excess well-being on its victims illustrates the advantages of retaining both well-defined intentional objects and the goal-directed behaviour advocated in this manifesto. Tomorrow's paradise-engineering specialists will probably judge it prudent to keep these traditional forms of life. Such modes of old-style intentionality will be needed for the purposes of any practical medium-term utopia, at least. No heroic sacrifice of subjective well-being is thereby demanded.

The role of pain isn't as straightforward as it seems. Its dreadfulness has been adaptive in our evolutionary past. Yet any full explanation of pain's phenomenological nastiness, as distinct from the functional role of "nociception", still eludes us completely; and perhaps it always will. The spectre of raw nastiness, however, is not the only way a

complex adaptive system can be induced to avoid, and respond to, injury. Unfortunately, it seems to have been the only adaptive response open to primordial carbon-based organisms consistent with the principles of natural selection. Fortunately, other strategies are now feasible. Whereas Evolution can't jump across deserts in the fitness landscape, paradise-designers in the era of post-genomic medicine certainly can. Humans can already build robots armed with "self-taught" artificial neural networks. These toy robots can learn to negotiate simple environments. They are capable of avoiding noxious stimuli via their responses to functional isomorphs of our pain states. Robotic silicon circuitry presumably lacks organic wetware's raw feel of phenomenological nastiness. So a less barbarous and primitive means of avoiding tissue damage in organic life-forms can surely be devised as well. [This expression of carbon chauvinism is controversial. It is not idle prejudice, however, but an inference drawn from the structurally and micro-functionally unique valence properties of the carbon atom and complex organic molecules.]

One way to promote pain-free nociception would be to use inorganic prostheses adapted from the design of our own future robots. A slightly more elegant solution would exploit our innate if often inept tendency to pleasure-maximisation. Peripheral nerves signalling noxious stimuli currently synapse on neural pain cells. They could instead be re-targeted on neurons which were simply less efficiently hedonistic in their biochemistry than their cellular neighbours. With their post-sensory signals remapped, infants could then learn self-preservation and pleasure-maximisation in harmony. At least as a stopgap, exploiting pleasure gradients is a much more civilised way to live. It's far more humane than responding to the contours of their nasty, and sometimes utterly excruciating, aversive counterparts.

A further presupposition of the question needs examining. One should be wary of assuming that *we're* the folk who can properly look after ourselves, whereas our descendants, if they become genetically pre-programmed ecstatics, will get trapped in robot-serviced states of infantile dependence. For it shouldn't be forgotten that exuberantly happy people also have a fierce will to survive. They love life dearly. They take on daunting challenges against seemingly impossible odds. One of the hallmarks of many endogenous depressive states, on the other hand, is so-called behavioural despair. If one learns that apparently no amount of effort can rescue one from an aversive stimulus, then one tends to sink into a lethargic stupor. This syndrome of "learned helplessness" may persist even when the opportunity to escape from the nasty stimulus subsequently arises.

Contemporary fatalism about the "inevitability" of suffering is analogous to this dysfunctional passivity (cf. the behavioural syndrome associated with the religious traditions of the Indian subcontinent). Yet passive acceptance of the dark side of life is no longer useful to contemporary humans now we've unravelled the genetic code. Species-wide hedonic engineering offers the prospect of eliminating all the vile types of experience we hate most; but even though it has become technically feasible to escape their clutches, a lot of us still aren't energetically striving to get rid of them. Unlike tortured lab-rats and monkeys, we can verbally rationalise our perceived helplessness in the face of psychological trauma or malaise. Suffering, we say, is "natural", "inevitable", "the way of the world", "Life", etc. By contrast, our eternally youthful, psychologically super-fit descendants won't need such coping-mechanisms. They are likely to be fired up with indomitable will-power. Their resourcefulness and zest for living should make them far better equipped to deal with life's practical inconveniences. Potential problems will be viewed as tremendously exciting challenges to be overcome. But in any case, future generations of post-humans are destined to enjoy god-like powers unknown to the mythical Olympians - both inside their virtual reality software-suites and out. They may indeed be ecstatically happy. But we would be rash to patronise them. For we're the ones who need help.

The argument that our descendants might become functional wireheads, too happy to reproduce, isn't compelling either. Happy people tend to want more sex, not less. Not everyone may opt for erotic modes of pleasure. But amongst sensualists who do, then gene-coded hyper-dopaminergic well-being is likely to promote, not celibacy, but heightened sexuality. This isn't simply a recipe for loveless orgies. Enriched serotonergic, phenylethylamine, oxcytocin and opiate function will allow us to care much more for each other and our dependants than selfish DNA normally allows today. Just how many newly-minted young ecstatics the world can ecologically accommodate, on the other hand, is uncertain. The elimination of functional pathologies like the ageing process is likely to make curbing rampant reproduction rather than promoting it a priority.

4.3 "This whole manifesto is flawed from the outset by its crudely reductionist approach to human beings. Our most profound spiritual experiences, and indeed what it is to be a person, can't be reduced to a dance of soulless molecules."

In the tough-minded reductionist camp, a hard-nosed atheistical scientist may be loath to see the beautifully choreographed neurons of his temporal cortex reduced to a spiritual buzz of religiosity. This isn't a very fruitful perspective either.

In one's eagerness to avoid an impoverished conception of human beings, it is easy to fall victim to an impoverished conception of chemicals. Natural scientists, no less than humanists, can easily fall into the same trap. On the assumption that all conscious experience - "what-it's-like-ness" - is identical with certain physical events or properties, then our classical materialist image of the ontology of the physical world, and our concept of what it means to be "physical", must be jettisoned as simply erroneous. It is not our fanciful mental images of matter and energy, but our deepening grasp of the formal mathematical tools needed for a description of quantum-mechanical events, that has enabled us increasingly to control and manipulate the basic stuff of the world. This grasp is now letting us control and manipulate, as well, the experiences with which at least some distributions of that "stuff" are identical. The phraseology sounds sinister and Orwellian. Yet if one's sovereign ethical principle entails striving for the fullest possible development of personal well-being everywhere, then embarking on the post-Darwinian enterprise is the only rational option.

4.4 "All of the drugs and therapeutic interventions touted here could potentially have long-term side-effects that we can't anticipate. The risk of another thalidomide tragedy writ large is too great to justify medical treatment of people who (by the norms of late twentieth century psychiatry, at least) are not suffering from any clinically recognised disorder."

The thalidomide tragedy took place several decades ago. The scandal unfolded before the medical significance of different optical isomers of the same compound in the body was appreciated. Such a mistake will not be made again. Of course, it can't be ruled out that other grave errors of judgement will be made instead. They probably will. In the early stages of any innovative treatment, the risk-reward ratio must always be finely weighed. This is all the more reason for preliminary experimentation to take place in the clinic and the laboratory, not on the street.

Presently, for instance, millions of young people are left to obtain and consume, in the most haphazard manner imaginable, the potentially neurotoxic compound MDMA. "Ecstasy" typically offers an enchanting state of consciousness while the trip lasts. Yet it's

a dangerous short-cut to mental health. Unless a subsequent dose of fluoxetine or another SSRI is taken soon afterwards, the drug damages serotonergic axonal terminals. Serotonin plays a vital role in regulating mood, impulse-control, anxiety and sleep. Thus in the long-term, MDMA and the other methoxylated amphetamines represent a poor choice of self-medication. It would be far better if the government were to take on the job of educating and training people in the most rational and effective ways to be happy. This role will involve sponsoring the research, development and widest possible distribution of the most safe, sustainable and beautiful empathetic euphoriants that medical science can formulate. Better still, research should focus on heritable gene-driven bliss. In the new reproductive era of "designer babies", prospective parents will choose the hedonic set-point of their future offspring. Curing our hereditary pathologies of mood will banish the need for drugs altogether.

4.5 "The radical therapeutic interventions which the biological program entails will presumably necessitate large-scale testing on non-human animals. This is surely inconsistent with the animal welfarist stance adopted earlier in the manifesto."

Given the feasibility, albeit not without difficulty, of implanting electrodes in the mind/brain's pleasure centres, there can be no principled utilitarian objection to subjecting both human and non-human animals to a great deal of enjoyment in the course of medical research. Many of the practical difficulties that the abolitionist project will face, and which demand greatest depth of understanding, stem precisely from *avoiding* crude pleasure-maximisation in the absence of a suitably well-designed encephalisation of emotion throughout the neo-cortex. If the animals in any experimental procedure are kept exceedingly happy for its duration, then the utilitarian ethicist needn't suffer any qualms of principle. At present, of course, the difference between an animal-experimenter's laboratory and a torture chamber is often imperceptible from his victims' point of view.

4.6 "Abolishing suffering is unnatural: in so doing we would forfeit our essential humanity."

Warfare, rape, famine, pestilence, infanticide and child-abuse have existed since time immemorial. They are quite "natural", whether from a historical, cross-cultural or sociobiological perspective. The implicit, and usually highly selective, equation of the "natural" with the morally good is dangerously facile and simplistic. The popular inclination to ascribe some kind of benign wisdom to an anthropomorphised Mother Nature serves, in practice, only to legitimate all manner of unspeakable cruelties. Extremes of suffering are inevitable under the neurogenetic status quo.

If a personified Nature did in some sense care about the progeny she prolifically churned out, then tampering with her benevolent handiwork might indeed represent a foolhardy Tempting of Providence. This sort of archaic romanticism about the natural world is impossible to reconcile with the neo-Darwinian synthesis. As has been all too aptly observed by "disposable soma" theorists, our genes just use us and then throw us away. "Unnatural" here is no more than a pejorative label. We use it to stigmatise, rather than rationally argue against, whatever we reflexively dislike. The very notion that a playing out of the laws of physics might ever yield something contrary to Nature is itself deeply suspect. Construed in any literal sense, it is false. Nothing that occurs in Nature is, or could be, unnatural. Both we and the transformed universe of our near and distant

posterity are equally a part of the natural world. Metaphorically interpreted, on the other hand, the charge of unnatural tampering is too ill-defined to be refutable.

And, yes, we will lose some primitive, "essential", human attributes. Yet why on earth should this be reckoned a bad thing? Until the development of powerful pain-killing drugs and modern surgical anaesthesiology, for example, frightful extremes of physical suffering were simply a part of the human condition. The unendurable just had to be lived through. Happily, in the present era our access to potent narcotics means, for the most part, that we no longer need to rationalise physical torments with the desperate sophistries typical of the past. Anyone arguing on religio-mystical grounds today that a loss of the agonies of the flesh is offensive to God, robbing us of a vital part of our species-essence, etc., is likely to get deservedly short shrift. Yet the supposedly ennobling properties of agonies of the spirit are still widely respected. Perhaps this attitude will change when retaining the capacity to feel psychological pain becomes a perverse genetic aberration rather than a condition of existence; and when inflicting it on others becomes an unthinkable crime.

4.7 "I'd get bored of being happy all the time. Variety is indispensable to personal well-being."

As an empty verbalism, "perpetual bliss" does sound fairly tedious. As Bernard Shaw once remarked, "Heaven, as conventionally conceived, is a place so inane, so dull, so useless, so miserable, that nobody has ever ventured to describe a whole day in heaven, though plenty of people have described a day at the seaside".

Successful paradise-engineering, however, must be the very antithesis of tedium by its very nature. If the prospect of paradise-engineering *sounds* unexciting, one has missed the point of what abolishing the substrates of tedium entails. In a different age, religious iconographers were able to derive much greater satisfaction in depicting the tortures of the wicked in Hell than in evoking the curiously anaemic delights of Heaven. Indeed, one could be forgiven for inferring that the eternal happiness of the Saved was dependent on contemplation of the eternal torment of the Damned. Likewise today, the secular equivalent of this syndrome is all too common. Potentially, however, there is no less a diversity of ways of being happy as being wretched. It is a grim reflection of the late-Darwinian human predicament that any notion of perpetual happiness evokes images of monotony. We can conjure up a rich and never-ending diet of disasters with ease.

Whatever humanity's contemporary failures of imagination, within a few generations the experience of boredom will be neurophysiologically impossible. "Against boredom even the gods struggle in vain", said Nietzsche; but he failed to anticipate biotechnology. From a naturalistic perspective, boredom amounts to just a complex of psychophysical states whose molecular substrate natural selection has chanced upon like any other. A capacity for boredom was retained because of the adaptive value its conditional activation can confer. Its more proximate physiological basis lies in the negative feedback mechanisms underlying the development of tolerance in the brain. These may be expressed in the form either of short-term habituation or a slightly more delayed process of gene-triggered receptor re-regulation. Such mechanisms can be disabled and replaced.

For as is experimentally demonstrable in the laboratory, the intra-cranial strategy of endless stimulation of the pleasure-centres of the brain confirms that happiness, and happiness itself alone, never palls. Out in the wider world, positive emotion just gets (re)directed to focus on and infuse a variety of intentional objects. None of our

neocortical patterns is *inherently* nice or nasty in the absence of its distinctive signature of limbic innervation. Some of these patterns may in time cease to satisfy; stone-age love affairs are cruel. Given the mind-brain identity theory presupposed in this manifesto, however, there is no biological reason why each moment of one's existence couldn't have the impact of a breathtaking revelation. As the phenomena of *déjà vu*, and its rarer cousin *jamais vu*, strikingly attest, a sense of familiarity or novelty is dissociable from the previous presence or absence of any particular type of intentional object with which such feelings might more normally be associated. So the kind of thrill one might first have got witnessing, say, the Creation can in principle become a property of every second of one's life. Cool.

4.8 "In the light of past horrors, from Auschwitz to the most private of griefs, it is disgusting even to contemplate celebrating existence by getting perpetually blissed out of one's head. Happiness, and indeed any other emotional state or response, should be rationally justifiable. It should be experienced only when it is appropriate. Given the horrors existing elsewhere in space-time, pure bliss is rationally unwarranted."

If it doesn't diminish the well-being of others, does happiness stand in need of justification any more than does the experience of, say, redness? As long as there is any chance that what we construe as the lessons of history might be ignored, and the obscenities of our evolutionary past in some way re-enacted, then there are excellent ethical-utilitarian reasons for keeping accessible even the most dreadful of memories. It may be important to remember more recent history, too, so as to honour and be supportive of those who have suffered in it and are now plagued by memories of earlier traumas and sacrifice. Yet to enjoin a grim reflection on the nature of the past for its own sake, a form of melancholy which, self-consistently, must itself presumably be commemorated mournfully in turn, is to set in motion an escalating cycle of misery without end. It's time to call a halt. Sometimes it is just better to forget rather than endlessly relive and recreate. If this sounds like shallow hedonism, it is worth recalling that HI's negative utilitarianism is an ethical system against which such a charge can least plausibly be sustained.

4.9 "I don't want a lifetime of enforced ecstasy. I want the freedom sometimes to be sad, and not to be enslaved to a false chemical happiness."

It is most unclear how to unpack the notion of "false" happiness. Contaminating the God-given purity of one's soul-stuff with alien chemicals is presumably offensive if one's self-conception is essentially spiritual in character. If, on the other hand, all states of consciousness alike are physically mediated, then it is scarcely coherent to label some neurochemical patterns as inherently false, unreal or inauthentic. Such euphoric states have indeed hitherto been largely inaccessible and genetically maladaptive if prolonged. They are still natural properties of suitably structured metabolic pathways of matter and energy. So in that sense they are all "true", though this is a most infelicitous way of putting it.

It is not, in any case, as though anyone will plausibly be forced to be happy against their will. Just as, historically, many slaves did not challenge the institutional legitimacy of slavery, and many self-confessed sinners believed they deserved to be damned to an eternity of torment in Hell, so many people have been able to convince themselves of the

ennobling quality of suffering. They will scarcely be ambushed and hauled in off the streets one day by crack-demented ecstatics and forcibly pumped full of euphoriants. A more apposite question might be what instruments of repression should a coercive State apparatus be entitled to use on behalf of possible bigoted die-hards of the old Darwinian order against people who decide, reasonably enough, that they do wish to live happily ever after. To what degree, and for how long and in what form, should authoritarian reactionaries have the right to compel others to suffer, once emotional primitivism becomes simply one life-style option amongst many?

4.10 "Pharmacological hedonism would turn us all into junkies. Gene-driven hedonism wouldn't be any different. We would lose all personal freedom because we'd be as helplessly addicted to our chemical fixes as the typical crack-head."

Once one has tasted other-worldly transports of ecstasy, it is true, there is no foreseeable way one would choose voluntarily to renounce such a condition. For from our current perspective, we have no more grasp of the real glory of the sublime than a newly-instructed five-year old child has of all but the barest mechanics of love or sex. Does our absence of hyper-ecstatic experience entitle us to claim any greater authority than the precocious but naïve youngster? Is such a claim testable? In reality, the nature of what lies beyond the arid text displayed here will prove, on revelation, more wonderful than could currently be physiologically imagined. Enraptured, one will enter into whole new modes of being. Reality redefined will feel so good that any surrender of born-again existence would be unendurably traumatic.

This condition might seem almost definitive of addiction. Yet on a utilitarian metric (barring only the austere "negative" sub-species), if such marvellous states are reliably and universally accessible, then seeking to achieve and maximise them is straightforwardly the right course to take. Addiction will tend to be a problem only if, first, people are hooked on something noxious to themselves or others; or, second, there is any likelihood of an interruption to their supply of the relevant drug or gene therapy. At present, we are dependent for what passes as mental health on different precursor amino-acids, essential fatty acids, minerals, vitamins etc to synthesise the brain's meagre dribble of pleasure-chemicals. We suffer gross psychophysical distress if we are deprived of them for long. This dependence, however, is regarded as wholesome rather than pernicious. It gets awarded the honorific "food". To achieve optimum mental health, on the other hand, one needs to dine on the richer diet of therapeutic agents advertised in this manifesto. The principle is the same.

The sheer finality of the Post-Darwinian Transition may indeed appal the metaphysical libertarian. For there can be no going back. Yet any opponent of the abolitionist project should be unsettled, too, by how endorsement of the traditional Nature-knows-best stance turns on our *not* exploring, however fleetingly, one of the two alternatives at issue. Ignorance is not bliss. Anyone who does *empirically* investigate, and not just pronounce on *a priori*, the rival forms of life on offer will unswervingly opt for the healthier modes of existence pleaded for here. More tellingly for the libertarian, perhaps, there is a sense in which the right to select one's own chemistry of consciousness, and thus to choose precisely who or what one wants to be, is as vital a sort of personal freedom as any. It is a freedom that we at present substantially lack. Any research program that opens up just such an option species-wide confers, surely, an incalculably life-enriching extension of choice.

Our own contemporary "choices" are in any case oversold. In the current era, we may seem relatively biologically unconstrained compared to our hidebound ancestors.

Some of us feel we can be, and do, more or less who and what we want. In fact, we can subsist only within the largely insensible confines of an extremely restrictive state space of psychochemical reactions. We can't hop outside their metabolic pathways to check what we're missing. If we could, we'd find the contrast too mind-wrenchingly different for words. Soon, however, we need no longer languish in biological servitude to our genes and the disposable vehicles they throw up. Today's junkies may vainly wish to be free from their inadvertently acquired addictions. This is only because the lows of illegal, dangerous and often self-defeating drug-taking ultimately outweigh the ephemeral highs of ill-chosen chemical euphoria. When, on the other hand, one opts once-and-for-all for a architecture of body-and-soul orgasmic sublimity, then one opts as well for a lifetime's freedom from second thoughts.

4.11 "I sometimes like being sad; it's an experience I wouldn't wish to lose."

An agreeable, wistful melancholy, a haunting lullaby nostalgically recalled from childhood, or perhaps the bitter-sweet memory of a long-lost love, are certainly preferable to the hell of unmitigated depression. Yet all too many types of experiences are unambiguously dreadful. They have no redeeming features at all. They don't issue in great works of art, literature and scholarship etc. They would be far better abolished. All the positive aspects of the more complex and ambivalent states one may undergo can in future be magnified and sharpened; nothing enjoyable need be lost. But the negative undercurrents which still diminish the value and enjoyment of more perceptibly composite states can be chemically subtracted out.

4.12 "Without suffering, there can be no personal development; unearned happiness leads to stasis."

Suffering is often just coarsening and brutalising. If one is sunk in hopeless despair, or even caught in the grip of an ill-defined malaise, it is as difficult to care about one's inner growth as it is to care about other people. Personal growth is more likely to unfold if one's appetite for life gets steadily keener. This will occur if one's experiences get progressively richer and more rewarding. Odysseys of self-exploration across the hedonic landscape can offer scope for ever-deepening self-discovery and idealised self-reinvention. Odysseys of pain and misfortune are as likely to desensitise or crush one's spirit as develop it.

Under the grisly genetic status quo, cultivating a sense of personal development is a comforting form of rationalisation, e.g: if I hadn't lost my legs in the accident 20 years ago I would never have become a great artist. So it proved a blessing in disguise after all! Prospectively, however, if one were told 20 years of suffering lay ahead if one sacrificed one's legs, but boundless self-development would follow in consequence, then one still wouldn't opt for it; and quite right too. As long as suffering is biologically inevitable, fitfully at least, then its optimal rationalisation is important solace for its victims. Thus reading this manifesto may cause more distress than joy to inveterate rationalisers; I just trust any unease will be mild and temporary. Yet when the biochemistry of suffering becomes only an optional neural add-on, the solace that rationalisation provides will impede the abolition of the miseries that demand it.

4.13 "Why bother with this intentional flotsam and jetsam at all if happiness itself is supposedly the overriding goal? In the context of the biological program, aren't intentional objects really free-floating and inessential frills to be varied or discarded at will? Isn't invoking "sublimity", "beauty", "love", etc, intellectually dishonest? Aren't they just rhetorical camouflage to win over those whose ideal pleasures tend to the respectably cerebral and the ethereal rather than the orgiastic?"

Our emotions have been pretty thoroughly encephalised by evolution. So it is certainly easier to give some hint of the nature of the paradise that awaits us by evoking, one may hope, the feelings one's audience associates with their own most cherished fantasies and objects of desire. Advocating happiness bereft of any nominal focus, on the other hand, entails working with a lifeless and unpersuasive abstraction. Advocating "hedonism" in the abstract is even worse. The term evokes something shallow, one-dimensional and amoral. Unfortunately, that's the price of sacrificing an underlying seriousness of moral purpose for the sake of a snappy manifesto title.

Naturally, what we think and say we're happy "about" is likely to change as the transition to paradise-engineering unfolds. Many highly-charged intentional objects of contemporary desires will seem historical curiosities even a few decades hence. In common with the particular time- and culture-bound conceptions of heaven and the good life in, say, different eras of the Christian and Islamic traditions, today's favourite intentional objects may indeed be only of derivative value. The mesolimbic dopamine system is doing most of the real causal work. But if the lure of such idols can motivate us to act on the promise of the biological program, then they will have more than served their purpose.

There are, however, substantive reasons why non-arbitrary intentional objects, and indeed an ever-greater scientific understanding of the world, should remain accessible into the indefinite future. The pragmatic advantages of the intentionalist strategy compared to wirehead bliss have already been cited. Sometimes it's useful to be able to look after oneself. There are powerful ethical reasons for keeping intentionalism as well. For ethically it is imperative that the sort of unspeakable suffering characteristic of the last few hundred million years on earth should never recur elsewhere. If such horror might exist anywhere else in the cosmos, presumably in the absence of practical intelligence sufficiently evolved to eliminate its distal roots, then this suffering too must be systematically sought out. It needs to be extirpated just as hell-states will have been on earth. Such inter-stellar rescue missions won't be possible if post-humans have all become wedded to the functional equivalent of wirehead-style pleasure-frenzies. This is because planning, executing and then stewarding ethically-run ecosystems of primordial extra-terrestrial life will require ultra-high technology, wide-ranging research, and a very long time. Subject to a number of assumptions about the origin of information-bearing self-replicators, any primordial life-forms - as distinct from some of their possible artificial successors - will be carbon-based. If multi-cellular evolution occurs, such alien life-forms will quite plausibly run on the same pleasure-pain axis as we do. Of course, this is all hugely speculative. And if trying to save the world is ambitious, then trying to save the universe smacks of hubris; so this avenue won't be pursued further here.

A negative utilitarian will still think that the striving for ever greater extremes and varieties of pleasurable experience while there remains any suffering whatsoever in this universe is a frivolous distraction from what morally matters. (S)he may be right. Certain contrived scenarios aside, however, the direct genetic and intra-cranial routes to paradise may serve the different flavours of utilitarianism equally well.

4.14 "Many of the greatest scientific and artistic achievements of humanity were born of tremendous struggles against adversity. Abolishing the biological substrates of suffering would mean there could be no fruitful inner struggle or creative tension, and hence no more Newtons, Picassos or Beethovens. Scientific and artistic genius demands a capacity for fierce criticism, both of one's own work and the ideas of others. Even if inducing a state of perpetual euphoria is consistent with bodily self-survival, the lack of critical self-insight such states entail would bring intellectual progress to a halt for ever."

It is worth distinguishing between the destiny of the humanities and the sciences after heaven has been biologically implemented. For a start, the exquisite aesthetic experiences on offer to our genetically enriched descendants may inspire an unprecedented flowering rather than a withering of the fine arts. Our current enjoyment of, say, Van Gogh's "Sunflowers" or Leonardo's "The Last Supper" will seem distracting tickles in comparison. Those who would deny that beauty is in the eye of the beholder might, or might not, be impressed by the disposition of paint on canvass which inspires these rhapsodies. Yet any reservations will last only so long as they remain trapped in the neurochemical orthodoxy of the past. At present, cultivating a fastidious unresponsiveness to certain forms of artistic production is taken as a badge of sophistication and discernment; but then that is our loss.

One blessing of the transcendent beauty awaiting discovery is that it will not depend on the vagaries of artistic genius for its production. The mind/brain lacks "beauty centres" of the same relatively well-defined architecture as its meso-limbic pleasure-system. Yet once the neurochemical signature of aesthetic appreciation is pieced together, its varieties can then be enhanced and selectively amplified. It should be recalled that perennial happiness can as easily lead to more being done in one's life rather than less. Intense episodes of creative energy today are often indistinguishable from mild euphoric hypomania. Some temperamentally laid-back lotus-eaters in the era ahead may indeed ultimately opt for meditative bliss and serenity. On the other hand, post-Transition society will probably be shaped by hypomanic "high-achievers" of formidable dynamism and productivity. Today's thrusting, can-do go-getters will seem lackadaisical in comparison.

The modes of well-being optimal for doing first-rate science and mathematics are obviously different from those best for practising first-rate art, poetry or sex. There is no reason why they should be less intense and rewarding. As to any lack of critical insight, there are also intellectual advantages to be derived from states of invincible well-being. Criticism of one's ideas in modern academia, for instance, is commonly taken as a full-frontal assault on the ego. In the future, critical scrutiny may be actively solicited and ecstatically welcomed. This might prove conducive to markedly better scholarship.

4.15 "The proposals of HI are too fanciful ever to gain credence, or even deserve serious critical consideration. They make a mockery of all our current values, aspirations and life-projects. A program so abhorrent to one's common-sense and moral intuitions belongs to the realm of vulgar science fiction rather than serious applied science or ethical debate."

Science has comprehensively confounded "common-sense" in all empirical matters. Our traditional ethical intuitions, when wrapped in secular guise, are less susceptible to experimental challenge. It would be a piece of singular good fortune if the least testable aspects of common-sense folk-wisdom just happened to be the ones that could most be relied on. At the very least, intellectual honesty demands that radically counter-intuitive challenges to received value-systems should receive close critical appraisal. The "values,

aspirations and life-projects" typical of, say, classical antiquity or the Indian sub-continent may easily seem ridiculous to the jaundiced contemporary eye. Likewise, the disparate intentional objects with which our own well-being now seems inseparably bound may eventually be seen as no less superstitiously revered. They objectively matter, but only because they objectively matter to us. So on the assumption that ethics amounts to something more than truth-valueless word-spinning, then it is worth at least considering the merits of ethical standpoints no less repugnant to common sense than, say, the theories of contemporary physics.

Appearances to the contrary, there is in any case a sense in which this paper, however superficially outlandish its substance, does not demand any revolutionary transformation of the core values of our secular culture. Its thrust stems from taking a quite conventional principle with the utmost seriousness it deserves. Only a minority of contemporary philosophers or laypeople are expressly utilitarians. Yet a diffuse and unsystematic utilitarianism is extremely widespread in society. It permeates the outlook of many people who never use the term. More interestingly, perhaps, an extraordinarily large proportion of non-, or even professedly anti-, utilitarian positions are argued on, or are underlain by, grounds which on examination prove subtly utilitarian.

Paradoxically, for utilitarian reasons it is nonetheless probably all to the good, this side of paradise at least, that at least some expressly non-utilitarian values are still held. This is because traditional folk-verities offset the acute discomfort many people still feel at the full implications of an exclusively utilitarian ethic.

Of course, one does not have to be a utilitarian to endorse the proposals of this manifesto. To those who are broadly sympathetic to the ethical utilitarian approach, however, then the biological program amounts, figuratively at least, to a gift from the gods.

4.16 "Being trapped in a chemical paradise would leave one wholly at the mercy of the ruling elites. The authorities could then treat people as puppets to be manipulated at will for their own ends."

The image that provokes this anxiety is presumably that of a drug-pacified class of helots. Perhaps a chemically enslaved underclass will work sweatshop hours for their masters simply to get their next chemical hit. In this fanciful scenario, it is in fact debatable who, if anyone, would really be exploiting whom. Also, certain sanctions are effective only if threatened rather than applied. No group is more ungovernably rebellious towards law and authority than addicts deprived of their fix. Moreover in our society, at least, the idea of the ruling elites engaging in a conspiracy to keep their population happy while they stoically shoulder the burdens of office tends to overtax the imagination; this is one conspiracy theory too far.

In any case, the conventional equation of happiness and docility owes more to distant memories of Huxley's *Brave New World* than to any deep reflection on the genetic, sociobiological and social-scientific literature. Prozac-style serotonin-enhancing mood-boosters, for instance, dramatically and consistently *increase* the status in the social pecking-order of the animals to whom they're administered. Such drugs may even lead them to reject a subordinate role altogether. It is revealing, too, that the manifestations of euphoric mania and melancholic depression also serve as descriptions of people occupying alpha and omega status-roles respectively. Mania, unlike most mental disorders, is most common in the upper social and economic classes. It typically involves an exaggeration of behaviour associated with achieving dominant status. By contrast, depression is most common among the poor. Even in today's society, the

persistence of depressive states and behaviour fosters stable hierarchies of social dominance. From the perspective of evolutionary psychology, the typical depressive syndrome is part of an adaptive coping-process. "Endogenous" depression involves the passive submission to a prolonged or uncontrollable stress. The elevated levels of cortisol and pain-relieving beta-endorphin characteristic of official clinical depression are also those which promote physiological adaptation to prolonged stressors. In the ancestral environment, depressive behaviour reduced the risk of physical damage by its tendency to reduce fighting within the group. In the post-Darwinian world, by contrast, depression simply won't exist.

So the "Brave New World" objection needs to be turned on its head. Given the correlation between depressed mood and low social status, the project of radically enriching the mood and motivation of the bulk of the population will probably leave people much *less*, not more, vulnerable to exploitation by a power-elite. In *Brave New World*, members of the populace were effectively the opiated and tranquillised dupes of the ruling authorities. Soma was a pacifying agent of social control. The consequences of genetically pre-programming happiness, however, will be very different. This is because everyday mental super-health will undermine the biological underpinnings of the dominance- and submission-relationships characteristic of our evolutionary past. More specifically, boosting the efficiency of tyrosine hydroxylase, for instance, won't just act to elevate mood. The consequently enhanced noradrenaline function in the locus coeruleus will tend to diminish subordinate behaviour. These simplistic "one neurochemical, one behaviour" stories are of course travesties of the truth, justified only on grounds of expository convenience. This doesn't challenge the essential point.

This point is that happiness, and an enhanced responsiveness to a wider range of rewards, is potentially hugely empowering. We're eternally slaves to the pleasure-pain axis; but a biologically enriched apparatus of pleasure and value-creation will help people assume a greater sense of control of their own lives. As noted, an all-action life-style fuelled by dopamine-driven well-being contrasts with the "learned helplessness" and "behavioural despair" characteristic of fatalists convinced that suffering is simply The Human Predicament. Either way, we shouldn't simple-mindedly project the power-and-submission relationships typical of early humans on the African savannah into the indefinite future. For the genetic basis of our core repertoire of social behaviour will first be tweaked and then drastically recoded. Too many sci-fi romances rely on extrapolating primate dominance-rituals into the indefinite future. That's what makes sci-fi soap operas set in one million years time so curiously (and so spuriously) intelligible. Whereas over the next few millennia and beyond, we'll have the chance to leave endless re-enactments of the ritual power-plays of the ancestral environment ever further behind.

4.17 "I'd rather stay in touch with Reality than live in an escapist fantasy world."

Some people enjoy the lucky conviction they have more intimate relations with Reality than the rest of us. A robust sense of intimacy is of course all the easier if one holds an agreeably commonsensical direct realist view of perception. Unfortunately, common sense is ill-named and at variance with the neuropsychological and quantum mechanical facts. Yet even a virtual worlder, for whom an awake mind/brain can aspire only to real-time data-driven simulations, may be sensitive to the charge of wanting to live in a fool's paradise, blissed out of his head come-what-may. Better, surely, to live like a sad but wise Socrates than as a happy pig.

Happy pigs should not be despised, but Socratic intellectual heavyweights can be happy too. In a magically transfigured environment in which all one's fellow creatures

were fabulously well, it is not clear at all why occupying an affectively neutral or pensive state should promote greater realism and representational fidelity. Perhaps the only way to grasp the actual nature of the unexplored celestial chemistry that beckons is to try becoming blissfully happy as well; and this is surely as good a reason as any for seeking maximal comprehension.

4.18 "Any creature which enjoyed perpetual bliss would no longer be me. I'm defined as much by my sorrows as my joys."

Winning £20 million on the national lottery, say, would wreak quite radical changes on most people's consciousness and sense of self-identity. It may nonetheless be suspected that the millions of punters who indulge their gambling streak are untroubled by the thought that their picking the lucky number will allow "somebody else" to enjoy the proceeds.

Philosophically, the notions of an enduring metaphysical ego, or for that matter of so-called "relative" identity, are indeed problematic if not incoherent. So in that sense the anxiety noted above is well-founded. Yet in such case any anxiety over personal (non-)identity applies no less to the psychochemical Dark Ages than to the post-Transfiguration era. One's namesake elsewhere in space-time who fell asleep last night is neither token nor even type-identical with the different configuration of matter and energy which bears one's name right now. Fortunately, even if personal identity is formally disavowed, one can normally muster the degree of altruism necessary to promote the future well-being of one's multiple namesakes, and likewise the namesakes and successors of one's family and friends. If contemporary notions of personal identity are ever culturally displaced by a different metaphysic, it may be hoped that our successors can muster the necessary degree of altruism too.

4.19 "When much of the world is still mired in poverty, hunger and disease, it is at best a flippant irrelevance to dream up hedonistic utopias. Their practice, if not aim, will be the cocooning of an already over-privileged planetary elite. We should instead concentrate on putting all our efforts into ensuring that everyone in the Third World has enough to eat, clean water supplies, a decent education and medical care and a civilised standard of living."

By most objective indices of well-being (the rates of marital breakdown, crime, suicide, clinical depression and other forms of psychiatric illness etc), the urban-industrial Western elite scores poorly compared to the materially underprivileged masses of the Third World. So the relative good fortune of the inhabitants of liberal capitalist democracies is easily overstated.

An "us and them" approach to life has its limitations. Within the next few hundred years, the invidious distinctions of class, nationality and race which poison the contemporary world will become redundant. On all but the most optimistic projections, the great majority of the world's population aren't going to achieve First World lifestyles for the foreseeable future; but we most assuredly do have the resources to enable the whole planetary population to be magnificently happy. If, for a start, a minute fraction of the resources currently poured into zero-sum status-goods and consumer fripperies were diverted to researching the development of safe, cheap, effective mood brighteners, delayed-action designer euphoriants, and genetically pre-programmed mental super-health, then we would all be far better off. This is no less true of the jaded plutocrat than the impoverished Third World peasant.

4.20 "The idea of spending one's entire life consumed by whole-body-orgasmic states of hyper-crack-like intensity and euphoria is simply grotesque. It is an affront to human dignity."

Unbridled sensual bliss will be merely one of the flavours of pleasure on the psychochemical menu, though not one that should cause us any embarrassment. In our own time, the dignified nature of such natural and short-lived routes to pleasure as sex is not always readily apparent to the untutored eye either. The more conspicuous pursuit of money, power and status characteristic of selfish DNA-driven civilisation tends to compromise human dignity in subtler but much more insidious ways. Champions of human dignity do not on the whole forswear such life-style choices, and understandably so; (in)dignity is very much in the eye of the beholder. Being made to suffer, however, is arguably the greatest indignity of all.

4.21 "The track-record of utopianism, whether romantic or allegedly scientific, is uniformly disastrous. Appalling crimes are committed on the assumption that the end justifies the means. A dystopian result is far more likely."

A "dystopia" where everyone is superlatively happy and fulfilled is surely the ultimate misnomer. Perhaps, if one's concept of perennial happiness still evokes images of bland and sterile monotony, then the charge may seem reasonable. In fact, the worst coercive excesses one can imagine, albeit somewhat implausibly, from a notional regime of State-sponsored hedonism might stem from the imposed penal sanction of compulsory biological euphoria - perhaps objectionable, but scarcely a cruel (though certainly an unusual) punishment.

4.22 "Genetically pre-programmed euphoria would undermine the basis of all human relationships. All this fancy verbal window-dressing about combining perpetual ecstasy with love, empathy, beauty etc is only superficial. Say, for example, some terrible physical misfortune overtakes a friend; after all, accidents can happen in even the best-run utopias. One will still be ecstatically happy: love for one's friend may indeed feel intense; but it is completely shallow if one can't grieve for a tragedy that befalls her."

By hypothesis, one's friend will be incapable of suffering; however badly mangled his or her body. Indeed (s)he will still be happy, albeit, we shall assume here, less intensely than before. Perhaps some of her favourite pleasure-cells are damaged. Let us also assume, in this scenario, that the molecular substrates of volition have long since been identified and toned up. One has chosen to blend the biochemical substrates of pleasure with those of dopaminergic "incentive" motivation rather than blissed-out satiety. If this is the case, then one will strive with all one's prodigiously augmented will-power to find means to restore one's friend to a state of maximal well-being. One will try far harder in dopaminergic overdrive than would be psychophysiologically possible if one were stuck in one's *current* comparatively weak-willed and ineffectual state. Thus a life of unremitting happiness doesn't entail that friendship is shallow or inauthentic; on the contrary, one will have the motivational resources to express depth of personal commitment all the more.

This is not to say that relationships won't change in many different ways after the Transition occurs. At present, for example, friendship often consists of offering mutual support in times of hardship and despair. In future, it may consist of a shared celebration of life.

4.23 "One big risk posed by the global species-project of *The Hedonistic Imperative* is that (post-)humanity will get "stuck" in a better, but perhaps still severely sub-optimal, state. Evolutionary progress, if one may be allowed to use such a term, would thereby come to an end. This is too high a price to be paid, or to run the risk of paying."

This worry shouldn't be lightly dismissed. But perhaps three points are worth making here.

First, natural selection has promoted such an abundance of dreadful states that even a severely sub-optimal (by whose criteria? - presumably not the sublimely fulfilled super-beings themselves) result would ethically be far preferable to today's status quo; and indeed preferable to any of our often hellish world's environmentally-tweaked successors.

Second, the danger of getting irreversibly stuck is still present even if genetic engineering and psychopharmacology are renounced in favour of time-honoured "peripheralist" approaches to making the world a better place. In fact, for what it's worth, psychoactive drugs potentially offer a form of "simulated annealing" [in artificial neural network-speak], enabling us to escape entrapment in local minima - though sometimes the jolt may be too uncontrollably violent and even dangerous to be commonly useful e.g. taking psychedelic agents such as lysergic acid diethylamide (LSD), ketamine or DMT.

Third, the idea that the paradise-engineering project sketched in HI would more readily lead to us getting "stuck" stems, I think, from its conflation with one or both of its two immediate intellectual antecedents of which I'm consciously aware. These are opiated-style quiescence *à la* Brave New World and the endless, uncontrollably orgasmic lever-pressing frenzy of a rat-/human-driven pleasure-machine. Both stereotypes are deceptive. One consequence of enhancing dopamine function in the manner stressed in this manifesto is that not merely is overall motivation deepened, but also the *range* of different activities one finds rewarding is increased (cf. the recent excitement over finding the D4 "novelty-loving" gene). Consequently, the likelihood of an organism, or a species, getting stuck in rut is *diminished*, though certainly not eliminated, by a strategy which incorporates boosting key receptor sub-types of dopamine-mediated process. It's worth noting that there is an experimentally demonstrable tendency of anti-dopaminergic mood-darkeners- and -flatteners, notably the D2-blocking major tranquillisers, to reduce incentive-motivation and novelty-seeking behaviour. They are "rut-inducers". Analogously, most of us Dark Age humans, stuck on a hedonic treadmill way down in the historical abyss, don't realise just how trapped we are.

On the other hand, there's a sense in which getting generically "stuck" in paradise is precisely what some of us are after.

4.24 "The eradication of suffering via genetic engineering and nanotechnology is an admirable goal. So why the disproportionate and perhaps (since so easily misinterpreted) irresponsible emphasis on mood-elevating drugs?"

Advanced genetic engineering and nanotechnological paradise-construction may yield states of conscious existence so wonderful and god-like that the notion of chemically fine-tuning them will seem absurd. What transhuman super-being would wish to contaminate the natural beauty of his or her soul-stuff with alien dirt? Yet some boring level-headedness about prospective time-scales is in order. It is true that the human genome of three-billion-odd base-pairs will be decoded within a decade. A *far* greater problem for intelligently encephalised paradise-production is the combinatorial explosion issue. This arises, quite inevitably, from a genotype's differential expression in differing environments. Airily invoking "genetic algorithms" and "quantum computation", for instance, is not wrong; but it tends to gloss over the formidable technical difficulties first to be overcome.

In the meantime, many people alive today will want biologically underwritten fulfilment for themselves and their loved ones. Born, tantalisingly, just prior to the Transitional era, they will have only the suspect stop-gap of enhancements to contemporary psychopharmacology to fall back on. Their access to cheap-and-cheerful paradises born of quick-and-dirty chemical fixes will, no doubt, seem dreadfully makeshift by the exalted lights of our more distant posterity. This doesn't mean that next century's pharmacotherapies should be damned with the knee-jerk invocation of "Drugs" conjured up by our own era's ill-judged recreational excesses. For one of the paradoxical effects, for instance, of a mind-healing strategy using even present-day selective serotonin reuptake blockers can be an enhanced sense of *undrugged* "normality" in the user. Such a sense can coincide with a biographically *abnormal* brightening of mood. Unacknowledged everyday states of derealisation, depersonalisation, and indeed other modes of depressive weirdness more typically associated with "bad trips" and "bad drugs", are in fact disturbingly common. Low-grade forms are frequent even in the absence of any exogenous agent to precipitate them. Moreover it's worth recalling that a subjective sense of humdrum, drug-naïve normality is itself just a chemically-induced adaptation. Neither we nor our blissful descendants need feel at all "drugged"; even if, in a sense, that's what we are; and always have been. But if we want to glimpse, rather than talk about, the naturalistic implementation of Paradise, then our generation(s) at least will need to use psychoactive tools-of-the-trade to get there.

In any case, given that so much of our very essence comprises the chemical ingredients of our recent meals, it's not as though one's ontological integrity as a pure spirit-being, or whatever, will be under threat from alien soul-pollutants. The difference between a drug and a nutrient, after all, reflects little more than the accidents of evolutionary history.

No 25 "The whole manifesto presupposes a Benthamite utilitarian ethic. If we don't accept its utilitarian presuppositions, then the abolitionist project collapses."

The abolitionist project isn't hostage to a single contested family of ethical theories. For it's not only utilitarians who abhor cruelty and suffering. Admittedly, the utilitarian may find it a matter of *moral* indifference whether our potentially ecstatic descendants opt to become wireheads, blissed-out junkies, or emotionally enriched post-Darwinian superminds. On the hypothetical felicific calculus, it's the sustainable intensity of our well-being (or the minimisation of malaise) that counts, not its peculiar flavours. But utilitarianism is a highly controversial ethic. So this manifesto, at least, lays stress on the

quite extraordinary diversity of options for paradise-engineering. These options embrace a spectrum of intellectual, psychedelic, aesthetic, empathetic and even spiritual modes of well-being far richer than anything accessible today. There's no obvious moral imperative driving us to unrefined pleasure-maximisation culminating in a perpetual cosmic orgasm.

Nevertheless, many contemporary thinkers will balk at *any* form of scientific utopianism. It's not that non-utilitarian ethicists typically argue that the texture ("what it feels like") of unpleasantness is *inherently* valuable. Instead, most non-utilitarians believe that a capacity for mental distress as well as physical pain serves an important *functional* role in life itself - and it always will. The many faces of suffering have been harnessed by natural selection [or more traditionally, Divine Providence] to promote the plurality of values that non-utilitarians uphold. Individual happiness is only one of those values. Much of what we care about isn't reducible to a unidimensional pleasure-pain axis.

Yet bioscience and nanotechnology promise more than the abolition of suffering and the enrichment of our emotional well-being. Critically, the new technologies allow us potentially to create the *functional analogues* of aversive states - analogue states that can play similar or even enhanced functional roles in the informational economy of an upgraded organism, but without the "raw feels" of suffering as we know it. Genetically constrained *gradients* of immense well-being - or smart neurochips with the right functional architecture - can be harnessed to animate our lives and promote what non-utilitarians typically value, but without the texture of subjective nastiness. If this prediction is borne out by the implementation of the new neurotechnologies, then the core of the secular *anti*-abolitionist case collapses. For only the most misanthropic nihilist would contend that despair, agony and malaise are inherently good. Suffering that serves no instrumental purpose at all, not even the interests of the genes whose inclusive fitness it once served, can be phased out without loss.

Of course, functionalist philosophy of mind may turn out to be wrong. As the functionalist alleges, minds may indeed implement the same computation/function in different ways and in different substrates, but perhaps effective nociception, say, must always have an unpleasant textural essence. Functionalism fails to explain the "hard problem" of consciousness; and our ignorance of why sentience (or anything at all) exists may infect everything else - including plans to get rid of suffering. It would seem very odd to claim that the texture of experience is functionally irrelevant or incidental to the role played by its biological substrates. For it's the sheer *nastiness* of suffering that ostensibly drives the abolitionist project in the first place. Yet we know we can build programmable silicon robots and embedded artificial neural networks to emulate the functional architecture of organic life-forms: we already engineer robotic sensory capacities, basic "appetitive" states, and the behavioural capacity to avoid noxious stimuli in ways that mimic feats of conscious human agency but without the merest whiff of sentience. On the other hand, today's robots are still primitive in their capabilities; and bionic implants are barely in their infancy. We can't simply extrapolate present-day technical successes into the indefinite future. Perhaps, *contra* functionalism as understood today, a subjective texture of unpleasantness will prove functionally indispensable for say, certain critical acts of judgement or discernment, or introspective self-examination. If these capacities are accorded a value potentially greater than the abolition of suffering, and if their subjective nastiness is functionally essential to the role they perform, then the abolitionist project may prove to have a more restricted appeal than the wider consensus canvassed here. If so, then seemingly abstruse debates about functionalist philosophy of mind would have an ethical significance beyond their technical merits.

Whatever the truth of functionalism, many non-utilitarian ethical positions *are* inconsistent with an abolitionist agenda; all the world's major religions for a start, with the ambiguous exception of Buddhism. Ethical systems that mandate the infliction of

misery on other sentient beings against their will can't be reconciled with any form of paradise-engineering. But on the whole, religious and secular ethicists alike aren't so much hostile to abolitionism as simply oblivious to its very possibility. Jesus, Mohammed and Buddha didn't have anything to say on molecular genetics and nanotechnology. Indeed, it's only in the past few decades that the abolitionist project could be contemplated as technically feasible on earth. Now that its blueprint can at least be formulated, all utilitarians should be abolitionists. But there's no need to turn utilitarian to endorse abolitionism: what's indispensable is an absence of malice.

No 26 "There will never be a Post-Darwinian Transition. There will always be selection pressure."

So long as there is ageing and death - i.e. for many centuries and perhaps millennia - there will indeed be selection pressure. But in the new reproductive era, the nature of that selection pressure will be different. In the old Darwinian era, "natural" selection is based on *random* genetic variations i.e. genetic mutations that are random with respect to what is favoured by natural selection; and it is *blind*. Nature has no foresight. By contrast, post-Darwinian, "unnatural" selection will be neither blind nor random nor socially unregulated. For reproductive decisions will be taken by informed actors *in anticipation of* the likely neuropsychological effects of suites of alleles that are purposely pre-selected or designed. Genes predisposing to vicious traits that were adaptive in our Darwinian past will be at a selective disadvantage when we choose the attributes of our offspring, not through a cruel genetic lottery as at present, but by rational design.

The imminent arrival of cloning and designer babies brings profound ethical dilemmas of its own - not least because the new reproductive technologies will precede any post-abolitionist era of mature paradise-engineering. As life-span increases, and the ageing process is progressively defeated, will reproductive decisions remain the prerogative of individuals as now? Or will reproductive decisions be taken societally? All one's libertarian instincts will be alarmed at this prospect. But the carrying capacity of the earth won't allow more than 50 to 100 billion people at most. Either way, there will be selection pressure in the sense that some genes and behavioural dispositions will lose out, at least until we become quasi-immortals and reproduction effectively ceases.

Of course, this heralded post-Darwinian Transition might *not* be to a civilisation based on paradise-engineering. Post-Darwinian society may be based on something else altogether. Yet because the texture of suffering isn't adaptive *per se*, whatever its current role in our legacy wetware, we can predict that the unsavoury genetic coalitions that manufacture its substrates will pass into evolutionary history.

No 27 "Paradise-engineering is impossible. It would not be evolutionarily stable. Game-theoretic modelling demonstrates that *selfishness* is always the most profitable strategy possible for replicating units - whether genes or "memes" - susceptible to invasion by "defectors". Invincibly happy life-forms are inherently more vulnerable than their discontented, anxious and malaise-driven counterparts. A society of genetically pre-programmed ecstatics could not arise, let alone endure. It would be an environment open to invasion by mean-spirited defector mutants who would replace the hardwired sweethearts. Unpleasant states of consciousness will last forever."

This objection conflates two issues. Could it ever be an evolutionarily stable strategy for our descendants to be 1) innately happy? 2) innately unselfish?

The answer to the first question depends on the sort of happiness hardwired. Are we modelling a civilisation of, say, quasi-immortal superminds animated by gradients of genetically programmed well-being? Or wireheads and their genetic equivalents - a "blissed out" rather than cerebral hedonism? Clearly, the option of global wireheading [or lifelong immersive virtual realities etc] *isn't* an evolutionarily stable strategy, at least until the ageing process is conquered. This is because wireheads have no inclination to breed and certainly not to raise children. By contrast, fitness-enhancing gradients of well-being - and traditionally, ill-being - or their functional analogues can serve to motivate, protect and preserve us. Such gradients are adaptive when they are "encephalised" by evolution - and ultimately, shaped by rational design. *Uniform* euphoria [or chronic depression] and its insentient robotic analogues isn't adaptive. For this sort of functional architecture doesn't impel its subjects to do anything, learn anything - or nurture children. Either way, genetic fitness isn't inseparably tied to a particular texture of experience, but to the way we behave and reproduce.

The controversial answer to the second question - namely that it is today's hardwired quasi-sociopathy that will prove evolutionarily unstable - sounds woolly-minded and naive, not to say biologically illiterate. Surely a civilisation founded on blissful altruists *can't* amount to a viable strategy? "Hardwired sweetheart" scenarios aren't pivotal to the abolitionist project. They are also *hugely* more speculative. So why is blissful altruism an option for paradise-engineering worth exploring? Surely selfishness always wins?

Fortunately not. The (technical) genetic and metaphorical, behavioral and psychological senses of "selfish" are easy to confuse. This is because today they overlap so closely. Paradise-engineering can never be based on *genetic* unselfishness. But a genetic predisposition to altruism - in the metaphorical, behavioral and psychological senses of "altruistic" - *can* be evolutionarily stable against so-called defectors *if and when* it is also genetically selfish i.e. Darwinian fitness-enhancing. This is how our capacity for kindness, compassion and empathy - however meagre - arose in the first place. Even today, a genetic predisposition to individual "saintliness" isn't always a losing strategy; recall the self-sacrificing holy man who attracts devoted female admirers and becomes the proverbial father of his nation. But on the whole, a capacity to cheat, to compete and to lie has proved adaptive; humans evolved as Machiavellian apes. Thus the proposal that unnatural selection pressure could ever cause "saintliness" to spread in a society of (non-clonal, genetically diverse) ecstatics looks implausible in practice. Surely alleles which promote competitiveness could never be outcompeted? Won't our descendants be, at best, happier *egotists*?

Now this may of course be the case. Yet decoding the human genome puts us on the brink of a major discontinuity in the mode of selection of self-replicating DNA - an evolutionary transition as profound as any in the history of life on earth. The long-term consequences of our capacity to rewrite our own code for the nature of adaptive - and maladaptive - traits may be very different from what we imagine. In the Darwinian era of "natural" selection, a regime of blind, random genetic variation typically promotes an indifference to the fate of most of our fellow genetic vehicles. In the environment of evolutionary adaptation, this predisposition enhanced the inclusive fitness of our DNA. We have a "theory of mind", but our minimal capacity for empathy is limited mostly to kith and kin. So callousness has flourished. "Nice guys" get eaten or outbred. Darwin himself speaks of "the clumsy, wasteful, blundering low and horridly cruel works of nature." By contrast, the impending post-Darwinian era of "unnatural" selection portends genotypes that will be *pre*-selected/designed in anticipation of their desired effects. So genetic variation will no longer be random and undirected. Its consequences will be

collectively planned - imperfectly at first, eventually perhaps via simulation and game-theoretic modelling with quantum supercomputers.

So questions of how we actually take the reproductive decisions, and on what *criteria*, are going to be crucial. What sort of traits do we want our offspring to have? Modelling post-Darwinian societies is immensely complex: post-humans may well rewrite their own individual genotypes ["genetic bootstrapping"] as well as the germ-line; and cloning will be trivially easy in the technical sense. Forms of "group selection" that simply weren't viable in the Darwinian Era become workable when reproductive decisions are collectivised; the "tragedy of the commons" can be forestalled. In a post-ageing world, reproduction may well be rare - and become progressively rarer as the carrying capacity of the earth [and ultimately the galaxy?] is reached. But taking a (very) crude genes' eye-view, in the era of designer babies a variant allele coding for, say, enhanced love-and-nurturance-inducing oxytocin expression, or a sub-type of serotonin receptor etc predisposing to *un*selfishness in the metaphorical, behavioral and psychological senses, may be differentially pre-selected and customised in preference to alleles promoting, say, sexual jealousy, aggressiveness or sociopathic behaviour. Genetically influenced "altruistic" traits that carry a higher payoff in the technical selfish genetic sense *aren't* susceptible to invasion by mean-spirited "defector" mutants - even if genetic variation were to remain random rather than directed. Thus in generations to come, the genetic and non-genetic senses of the word "selfish" may diverge. Indeed as the abolition of suffering becomes first technically feasible, and later trivially easy, then the language and institutions of traditional morality may become archaic relics from a vanished age. What sort of values will replace them is hard to say. But as our descendants rewrite the vertebrate genome, and redesign the global ecosystem via nanotechnology, harsh "unnatural" selection pressure may penalise the very sorts of nasty traits that were genetically adaptive in the Darwinian Era. On this analysis, post-Darwinian superminds will be extraordinarily benevolent; but paradoxically, the science of paradise-engineering will have its origins in genetic selfishness.

Perhaps. Let's take a more pessimistic scenario. Assume that (post-)humans continue to be selfish in every sense. After all, just because allegedly we all (obliquely) seek happiness, this doesn't mean we seek happiness for everybody. Just because successful and intelligent life-forms will be able to underwrite their own happiness, why assume that they'll care about others? Let's further assume, contrary to the optimistic functionalist arguments above, that the textures of invincible happiness *do* inevitably make any coalition of alleles that promotes them potentially genetically vulnerable. After all, invincible well-being wasn't a viable strategy on the African savannah; why should it triumph in an era of artificial selection?

Does this pessimistic set of assumptions predict the persistence of a legacy architecture of misery and malaise? Will unpleasant states of consciousness really last for ever?

No, not necessarily, not even then. The more vulnerable that enhanced well-being allegedly makes us, the more our self-interest will lie in ensuring that all others are happy and well-disposed too; and in ensuring that any novel life-forms we create in the new reproductive era are constitutionally happy and benevolent. If the discontent of others potentially threatens our own well-being, then genetically underwriting their empathetic bliss serves our self-interest. If mutant psychopaths pose a potential danger [though in fact strict sociopathy tends to diminish inclusive fitness even in the primordial Darwinian era], then self-interest dictates using prophylactic germ-line therapy against genes promoting sociopathy and its sub-syndromal variants; this is one state-space of genetic options whose full exploration we can live without. In the past, natural selection ensured that selfishness, in every sense of the word, frequently paid. This entailed "winners" causing often severe suffering to losers. According to rank theory, the far greater incidence of the internalised correlate of the losing [behavioral] sub-routine,

depression, compared to the winning sub-routine, euphoric (hypo)mania, attests to the terrible price that social animals have paid for the advantages of group living. Until now, blind genetic competition has ensured overt individual competitiveness among reproductive vehicles. There has been a sometimes physically violent struggle for the best mates and scarce resources. Winners and losers alike have been trapped on the same hedonic/dolorous treadmill. But when unlimited emotional well-being is possible for everyone at no cost to the well-being of others - and an unlimited diversity of good experiences is accessible to all via immersive VR - then only sustained malevolence, not mere egoism, will suffice to perpetuate the cruelties of the old order.

None of this proves that our descendants will really be smarter, nicer and happier - the magic trinity predicted and endorsed here. This is scenario-spinning, not true game-theoretic modelling. There are suppressed premises and controversial assumptions in all the above arguments for paradise-engineering. Which strategies will really prove stable remains to be seen. The nature of the ultimate winning strategy is open. Certainly a transformation of human nature isn't going to arise through a world-wide spiritual awakening, an innovative package of socio-economic reforms, or a spontaneous desire to be nice to each other. But it's quite possible that, in the long run, the Darwinian genetic program based on suffering and quasi-sociopathy will lose out. Misery is not a stable strategy because by its nature rational agents seek to escape it; and soon a society of intelligent agents will have the collective capacity to do so.

No 28 "There is a contradiction at the heart of the abolitionist project. On the one hand, it is argued that suffering will be eradicated by biotechnology. On the other hand, it is claimed that no one will be *forced* to be happy: our freedom will allegedly be enhanced, not restricted, by the option of unlimited bliss. But perversely or otherwise, some people will always choose to be miserable - or at least to retain the traditional biological capacity to be so. Thus abolitionism can't be reconciled with an absence of compulsion."

Prescription and prediction are easily muddled. It is *advocated* that all involuntary suffering *should* be abolished. It is *predicted* that all suffering *will* be abolished. On this perspective, our descendants are no more likely to submit themselves to emotional pain and malaise than we would today opt to undergo a major surgical operation without an anaesthetic.

In practice, an ethic of absolute personal freedom is probably untenable. Even the devout libertarian will sanction, say, the administration of a foul-tasting medicine to an unwilling sick youngster, or the forcible injection of an anaesthetic into a struggling animal before veterinary surgery. We sometimes override the choices and desires of simple minds. It would be cruel to do otherwise. Non-human animals, the severely mentally disabled and very young children don't know their own *interests*; mature adults are presumed different. The problem here is that super-intelligent extraterrestrials - or our own advanced descendants - may perceive *us*, primitive *Homo sapiens*, as comparatively no less mentally defective than are toddlers or pets in our eyes today. Any advanced intelligence may discern the analogous way that Darwinian minds are locked in dysfunctional cycles of self-abuse - unaware of *our* own interests. If so, then should we/small children be allowed to keep on hurting ourselves so badly?

As libertarians, we must presumably answer yes. This stance would seem hard to reconcile with a utilitarian ethic. For what are a few minutes of unpleasantness compared to an eternity of bliss? Yet even to moot the involuntary treatment of malcontents, let alone advocate its practice, is a dangerous line of argument for the abolitionist to pursue. For the misconception that anyone is going to coerce us into being happy is one of the

biggest ideological obstacles to the future abolition of suffering. Fortunately, it is a mistake to believe that even a utilitarian ethicist is committed to mandatory therapy for the emotionally sick. This is because even the hint of compulsion causes distress to most people - thereby sabotaging the abolitionist project and defeating the utilitarian's own ends.

So the spectre of dissident emotional primitives being dragged kicking and screaming into the pleasure chambers must not become the defining image of abolitionist ideology. Conjuring up such a travesty of paradise-engineering doesn't show that a utilitarian ethic is mistaken. Instead it illustrates that the advocacy of compulsion is not a truly utilitarian policy at all. Like so many arguments against a utilitarian ethic, it relies on misconceived policy prescriptions wrongly derived from the sovereign utility-maximising principle.

In reality, abolitionists may call themselves fanatical libertarians on solid utilitarian grounds. For the freedom to transcend our Darwinian past and to choose our own homeostatic level of well-being is one of the most persuasive arguments for the abolitionist case.

No 29 "Why invoke nanotechnology? Surely genetic engineering alone can abolish suffering?"

If the abolitionist project is to be complete, then it must embrace the rest of the living world. In terrestrial ecosystems, the higher vertebrates can be genetically redesigned using foreseeable extensions of existing technologies. But pain and suffering will still fester in less accessible parts of the animal kingdom e.g. in the oceans. Fortunately, within a few centuries, our descendants will have the capacity to use self-replicating nanobots armed with supercomputing power to redesign the marine ecosystem. Today, needless to say, this sounds like the wildest science fantasy. But even if we rely only on extrapolation, not revolutionary conceptual and technical breakthroughs, then the implementation of the abolitionist program is still grounded in relatively well-understood science. The reason that the prospect of molecular hedonic engineering hasn't yet been explored by nanotechnology theorists is not that the technology involved is uniquely challenging. It's because tough-minded technocrats have different ends in mind.

In the present era, of course, it is hard to feel deeply exercised by the plight of marine invertebrates. We may feel that we have worries enough nearer home. But it is not pleasant to be eaten alive, even if one is a small mollusc. In paradise, it won't happen.

No 30 "Suppose that biotechnology really does give birth to an entirely new reproductive era. Suppose that humanity really is destined, as claimed in *HI*, for an era of ubiquitous designer babies - the so-called post-Darwinian transition. This transition may *not* be to an era of paradise-engineering. The biological basis of suffering may never be abolished. For if prospective parents are free to choose the attributes of their children, their typical priority will *not* be the creation of offspring who are innately happy. Instead, innumerable "pushy" parents will continue to seek children who are smarter, better-looking, competitively driven, more "successful" - and choose genotypes to match. Such parental bias can be explained, ultimately, by evolutionary psychology. At present, of course, prospective parents can't directly select allelic combinations

of genes that promote such traits. In tomorrow's genetic supermarket, they may be granted an opportunity to do so. But if so, then selection pressure - albeit artificial or "unnatural" selection pressure - will favour exaggerated versions of traits that were adaptive in the old Darwinian era of natural selection. The outcome of the imminent reproductive revolution *won't* be a civilisation founded on genetically pre-programmed bliss."

Assume, plausibly, that within a few decades prospective parents will be able to choose the genetic dial settings for their kids' emotional well-being - the average "set-point" on our emotional thermostat around which well-being (or ill-being) tends to fluctuate. Grant too the key premise of the objection: many parents do indeed care far more about the worldly "success" of their children than their personal (un)happiness. This doesn't entail that the substrates of suffering will be re-created indefinitely. Even parents for whom the emotional well-being of their offspring is trivial - of no more significance than, say, choice of eye colour - are still likely to opt for higher rather than lower dial settings on the hedonic treadmill i.e. alleles and allelic combinations that predispose their children to flourish. For most parents *do* prefer, on balance, their children to be temperamentally happy rather than miserable, even if happiness is only one desired attribute among many - perhaps not the most important - and in some instances perhaps only a minor or incidental trait. "I don't care what [s]he does when [s]he grows up, so long as [s]he's happy" expresses, not a revolutionary sentiment, but a clichéd platitude of Western liberal society. This preference is explicable in part because happiness, and the spectrum of behavior associated with the "winning sub-routine", is positively correlated with social dominance and reproductive success. Ambitious parents certainly don't want to produce "losers". Depressive or anxiety-ridden kids can't compete effectively against their peers. A tendency to low mood, and the spectrum of subordinate behaviour with which depression is associated, may have been genetically adaptive for low-status tribal weaklings on the African savannah. For depressive behaviour, contingently activated, can be a viable fallback strategy for stressed low-status tribal animals in an adverse social environment. This may explain why depressive disorders are so common. But a genetic predisposition to low spirits, or at least anything like unipolar depression as distinct from bipolarity, is not part of an optimal reproductive strategy for potential "winners". If intelligently engineered, a genetically enhanced sense of well-being is *empowering*. Its behavioural phenotypes are potentially far more adaptive than the predisposition to learned helplessness and behavioural despair characteristic of the depressive spectrum. So in the new reproductive era, pushy parents in particular are likely to shun depressive genotypes. What guise their children's well-being may take is another question. True emotional enrichment transcends the simple-minded recipes discussed here - mere modulations of the old Darwinian repertoire of sadness, happiness, disgust, fear, jealousy, anger and loneliness. Indeed the enriched emotional palette of our descendants may assume textures conceptually unimaginable to primordial Darwinian lifeforms. Our post-human successors may be rapturously happy about things we've never dreamed of, in ways we can't imagine, and in a conceptual scheme that hasn't yet been invented. But in today's terms, parents who are ambitious in a conventional sense for their family may seek an egoistic rather than empathetic kind of well-being for their children. Such parents may also favour (genotypes predisposing to) hypomanic exuberance rather than *serene* happiness. Backwood-looking parents may even opt to endow their children with functional analogues of older Darwinian traits, but set against a much higher emotional baseline. None of this suggests that parents will opt, in the long run, for allelic combinations whose expression induces suffering or even unpleasantness in their carriers - even if medical ethics committees were to license their (re-)creation. Aside from anything else, children who are genetically predisposed to be depressive, sour-tempered or brattish are less rewarding to raise than children who are abundantly joyful and loving. Pre-selecting one of the nastier Darwinian genotypes for one's progeny would be self-defeating. In an era of artificial selection, the partially heritable bundle of traits we call "lovability" promises to be highly adaptive for (post-)humans and their household pets alike.

The above account inevitably falls short on detail. Empirical cross-cultural studies of the (partially) heritable characters most favoured by contemporary parents for their offspring may serve as a better guide to the nature of tomorrow's designer babies. However, such a yardstick implausibly assumes an absence of state regulation and control over parental genetic choices. Likewise, the question of the future intensity settings of genetically pre-programmed happiness is here left open. Oversimplifying hugely, and treating happiness on a crude one-dimensional scale, will successive generations of genetically enriched (post-)humans tend to be a bit happier, or blissfully happy, or orders of magnitude happier than their Darwinian ancestors, as predicted in *HI*? Most parents today, if pressed, might express a preference for their children to be very happy rather than happy; but only a minority of early adopters would opt for superkids who were constitutionally *sublimely* happy. Thus in the near future, the dial settings on enhanced kids' emotional thermostats will probably encode lives animated by (homeostatic gradients of) modest well-being rather than (homeostatic gradients of) sublime bliss. Analogously today, parents are typically most comfortable with the idea of rearing clever children rather than a family of geniuses. Yet as our conception of psychological health is enriched, so presumably will its socially acceptable norms. Ambitious parents usually aspire to a higher quality of life for their offspring than their own. This generalisation holds even though a comparative *poverty* of ambition may initially induce many parents to settle for comfortable mediocrity for their kids rather than mental superhealth. Perhaps this pleasure-deficit will be remedied in our lifetime by somatic gene therapy and genetically personalised mood-enrichers; perhaps not. But ultimately our descendants are no more likely to pre-select genotypes coding for inherently *nasty* states of mind than they are likely to pre-select genotypes coding for neuropathic pain. The historical record notwithstanding, human perversity has its limits.

No 31 "There is a flaw, possibly a fatal flaw, in HI. Yes, there probably will be a reproductive revolution. True, over time, prospective parents are unlikely to choose "nasty" genotypes for their children. Yes, this reproductive shift may even represent a major evolutionary transition in life on earth. But, critically, a large percentage of the population will presumably continue to have children by "natural" means - whether out of bioconservative ideology, religious conviction, or just normal teenage fecklessness. Among this percentage of natural reproducers, a large and unknown number of couples will themselves be the offspring of natural methods of reproduction. Therefore a lot of the nastier code in our old Darwinian genome will be retained, together with the propensity to suffering it entails. Perhaps the natural reproducers will eventually interbreed with mature designer babies of more distant posterity. Who knows what will be the long-term consequences of mixing rational re-design and a legacy genome? But either way, unless the ideology of abolitionism is universally adopted as a value system - or ruthlessly enforced by a coercive state apparatus of unprecedented intrusiveness into the female body - then the global abolition of suffering will be postponed indefinitely. HI is a nice idea. But it's hard to see how it could work."

The key premise of the Objection is probably correct. So long as any pure-bred Darwinians continue to procreate by natural means, then suffering in some form or other will persist. The persistence of suffering is inevitable if archaic humans also reject as "unnatural" (etc) the other two core technologies of mood-enhancement, i.e. wireheading and sustainable pleasure drugs. So what grounds are there for believing that natural reproduction as practised today will ever cease? This is quite a radical prediction. And even if the abolition of natural reproduction is *technically* feasible, isn't its disappearance too high a price to pay for mental superhealth and a cruelty-free world?

The reason for predicting that within a few centuries *all* human reproduction will be rigorously controlled, both in its timing and in its nature, stems from a second momentous technological revolution in prospect, namely the conquest of ageing. Whether you estimate that curing senescence will take another 100 years or 500 years, this genetic-cum-nanotechnological revolution is destined to sweep away the plague of human mortality. First on the horizon are interventions to prevent age-associated diseases (Alzheimer's, osteoporosis, cardiovascular disease, age-related memory decline, etc). Such primitive gene therapies are only the harbinger of a massive repair-and-renovation job on the human genome. This mega-project will tackle the fundamental biology of ageing itself. Replacing the biology of ageing is much more ambitious. Since rational design of the genome from scratch is impossible, we can only "bootstrap" our way to millennial lifespans - a formidable genetic challenge. But as the era of eternal youth unfolds, our descendants are not going to pre-select genotypes predisposing to ("for") age-associated diseases or senescence for their future offspring. Nor, realistically, are members of the older generation likely to shun rejuvenating somatic gene therapies for themselves. In consequence, the current slowdown in global population growth will reverse. The planet will fill up and approach the limits of its carrying capacity.

This physical constraint on our ability to multiply will recede but stays intact even if you think we are destined to colonise the galaxy, or even if (fancifully and implausibly) you think we are going to "upload" ourselves onto computers, or even if you think the sky's the limit and intelligent life is limited in its expansion potential only by our world's Bekenstein bound. Even if individual mobility and resource consumption weren't an issue either, since we'll all be plugged into immersive VR or an analogue of the Matrix (etc.), then this physical constraint still holds: if we phase out ageing and become quasi-immortals, then we'll quite literally run out of *Lebensraum* in the absence of strict reproductive controls. The libertarian will find these words as uncomfortable to read as they are to write.

HI ducks the question of the specific social and biomedical mechanisms regulating reproduction in a post-ageing society. This omission is deliberate: control of human reproduction, whether sexual or clonal, will be a *generic* feature of any post-ageing civilisation. The need for social mechanisms of reproductive control on pain of Malthusian catastrophe isn't a specific peculiarity of the abolitionist project. If (post-)humans aren't going to grow old and die, as we do today, then we can't go on having children at will indefinitely. A regime based on genetic Russian roulette will be replaced by an ethically responsible(?) policy of planned parenthood.

At what cost? Other things being equal, state-regulated birth-control might be expected to cause widespread and profound personal distress. Only a small minority of people in human society are happy to remain childless. Infertility causes much heartache. For most people, having children is to a greater or lesser degree our *raison d'être*. For evolutionary reasons, it would be astonishing if this were other than the case. We may fear death and growing old; but typically what makes life meaningful - and our death bearable - is the lives of our children and grandchildren. Thus as we're constituted at present, the spectre of restrictions on our right to procreate is a disturbing idea. An intimate realm of our lives that has hitherto been essentially private could be in danger of intrusion by the state. Even a Chinese-style one-child campaign strikes the Western mind as a draconian curb on personal freedom.

So how will this dilemma be resolved? At present, we may try and persuade ourselves that we wouldn't want to stay eternally youthful. But if the option of eternal youth or even its semblance were there, then it would be naïve to think most people wouldn't discard a lifetime of rationalisations and seize it. This bold statement might seem to imply a rather facile biotechnological determinism. For it is being assumed without argument that just because 1) we don't really want to grow old; and 2) technically it will be feasible to live indefinitely, we will therefore opt to do so - barring traumatic wetware

accidents of course, though even here the use of prudent automated off-site self-backup policies should allow restores from last working copy. But for all its pitfalls, some sort of biotechnological determinism here is well-founded. Our fear of ageing, death and dying is simply too deeply rooted in the Darwinian psyche for us to perpetuate the senile holocaust into the era of mature genomic medicine. Renouncing the option of quasi-immortality may be conceivable in theory. Yet who'll opt to live (and die) as a disposable Darwinian "crumbly" if one can live and look like a Greek god?

The solution to the psychological dislocations such sustainable youth may entail is more likely to be biological than sociological. Just as biotechnology can potentially allow us to become better, more loving parents (e.g. by use of agents that induce oxytocin receptor gene overexpression, etc), so conversely biotech can curb the craving to have children when reproduction is infeasible. These techniques may be pharmacological or genetic or both. Godlike lifespans needn't have *any* adverse effects on our mental health; quite the reverse. Genetically enriched humans can feel utterly divine, not just look it. For lifelong well-being can potentially take many guises; and most forms of emotional enrichment won't entail living vicariously through the lives of our immediate biological descendants - natural as this habit of mind still seems in our late Darwinian world.

Switching on or off some of our deepest human desires sounds more like a dystopian nightmare than a recipe for paradise-engineering. Who is to orchestrate the switching; and how? No such hard choices are thrust upon us today. We just reproduce, decline into our dotage and then die. Yet re-engineering the human soul and body alike can still strike even secular minds as almost sacrilegious. We admire excellence in the design of inorganic technology even as we abhor its prospect in ourselves. But whatever the mechanisms, if we cure ageing and *don't* intervene to regulate other primordial human traits as well, then intolerable psychological stress and social conflict are presumably inevitable. All sorts of ugly scenarios can be envisaged if life-extension technologies are pursued in isolation from mental health research and therapeutic interventions to match.

Nothing in this analysis of a post-ageing world *proves* that the control (post-)human reproduction also entails the design of psychologically superwell (post-)humans. In overcoming ageing, it is possible if sociologically unlikely that we will opt to leave our repertoire of hunter-gatherer emotions unchanged - just as, conversely, it is technically possible we will conquer suffering without scrapping death and ageing. The response set out here aims rather to show why haphazard sexual reproduction isn't an inevitable fixture of tomorrow's post-Darwinian society; and how in future the creation of pain-ridden humans will demand an implausible measure of *premeditation*. So too, one day, may the creation of perishable human beings destined to grow old and die.

Yet just how likely *in practice* are our descendants to be eternally youthful, superintelligent, superempathetic - and to live happily ever after? A reality-check might seem in order. The post-ageing era is still far enough away to make *any* predictions hazardous. Those of us still in thrall to our Darwinian gut-instincts will find these scenarios all smack of wish-fulfilment and idle fantasy - mere fairy tales masquerading as science. HI certainly glosses over some very grim late Darwinian nastiness looming in the decades ahead: nuclear warfare, bioterrorism, global pandemics - and the usual soul-destroying tragedies of Darwinian-style personal life. Certainly, any futurology based on radical *dis*continuities rather than extrapolation rarely rings true at the time. But the (potential) beauty of genetic engineering, quantum supercomputing and utopian nanotech is the way these technologies can be used to convert wishful thinking into sublime reality. What it means to be "realistic" will shortly be redefined. One reason for researching the prospects of a post-Darwinian civilisation is that paradise-engineering can deliver a practical solution to everything that's wrong with the world today.

No 32 "If (1) HI is correct, And if (2) HI should apply to all sentient beings, not just those on earth, Then (3) We have a moral obligation to spread throughout the universe as quickly as is practical, eliminating aversive experience and maximizing pleasure gradients everywhere.

Furthermore, if also (4) There are a very large number (let's say at least millions) of intelligent life forms elsewhere in the universe, Then (5) It's a virtual certainty that at least some of them (and more likely, most of them) are substantially more intelligent than us, And (6) It's a virtual certainty that at least some of them are at least equally driven to their goals, at least some subset of which are likely to apply to the entire universe.

We can subdivide the life forms mentioned in (6) into three categories: Category A consists of those life forms which have the same goals and choose the same means as HI. This sounds unlikely but might not be. Consider: If (7) morality is absolute rather than relative (i.e. there is some correct way to behave), and if (8) morality has attractors (i.e. most or all sufficiently intelligent life forms will discover the right way to behave and at least some of them will choose to behave that way), and if (1) then (9) at least some other life forms will find HI persuasive and will work toward it.

If (9) and (4), and if (10) the most advanced life forms are best equipped to determine and then carry out HI to maximize the chances of success, then (11) it's probably the case that there is no need for humans to get involved in HI. This logic isn't airtight, however. For example, if (12) all life forms reason this way, then none would act, assuming that some other life form would take care of HI (unless one or more life forms thought or knew that they were the most advanced). In addition, it might be the case that (13) the best implementation approach involves several life forms, not just the most advanced one (perhaps to accomplish the goals of HI more quickly). Nevertheless, it seems fairly clear that if (9) and (4), then it's highly unlikely that humanity is in the best position to implement universe-wide HI.

Category B consists of those life forms which have the same goals but choose different means than us. Some of the points in Category A would apply, but an additional conclusion given (5) seems to be that we should trust their judgement. This appears to be true even those life forms felt that the best approach included elimination of earthly life (and other similar life forms elsewhere).

Category C consists of those life forms which have different goals. If (6), then I believe that it is a virtual certainty that Category C is not empty; i.e., at least some life forms will have different goals than HI. If this is the case, and if (5), then it doesn't seem to matter much what we do, as the outcome will almost certainly be the goal of whichever life form is most advanced. This doesn't imply that (14) working toward earth-level HI goals is entirely pointless, but it does seem to substantially restrict the value of such efforts, making them local and temporary." [*with thanks to Tom Murcko*]

Most people believe that the complete abolition of suffering in *Homo sapiens* is impossible. Extending the circle of compassion to other animals via ecosystem redesign and genetic engineering seems even more far-fetched. So the prospect of some kind of cosmic rescue mission to promote paradise engineering throughout the universe has a distinct air of science fiction. This may of course be the case. The timescales are certainly daunting even for a single galaxy of 400 billion stars some 100,000 light years across - on the order of millions or perhaps tens of millions of years. The level of intellectual,

political and sociological cohesion over time required to mount such a project eclipses anything human society could organise today. Moreover recent evidence from distant type Ia supernovae suggests that the expansion of the universe isn't slowing as hitherto supposed, but accelerating owing to poorly understood "dark energy". In consequence, perhaps only our local galactic supercluster will ever be accessible to our descendants.

Viewed purely as a technical challenge, however, the use of self-reproducing, autonomous robots - "von Neumann probes" - to explore and/or colonize our galaxy is both feasible and well-researched. The difference is that their purpose hasn't normally been conceived as a mercy mission for pain-ridden ecosystems that may have evolved elsewhere. [Ironically, notional "berserker probes" that sterilise all life have been discussed in science fiction, albeit not with a negative utilitarian ethic in mind.] Plausibility aside, it is ethically obligatory for utilitarians anywhere to maximise the well-being of all accessible sentience *if* it's technically feasible to do so - in the absence of any countervailing argument like the Objection above. Less clearly, an obligation to promote the substrates of well-being throughout the cosmos is arguably a disguised implication of various ethical systems that deplore merely "unnecessary" suffering. What "necessary suffering" might mean here is critical but ambiguous.

The most problematic premise in the Objection is perhaps number 4, i.e. the hypothetical existence of millions of other intelligent lifeforms. This assumption relies on the Drake equation[1] or one of its variants in estimating the number of extraterrestrial civilizations with which we might come in contact. Any such assumption must overcome the Fermi paradox: "Where are they?" No discernible sign of extraterrestrial life exists - whether its artefacts, physical presence or signals. There may indeed be an indefinitely large number of technologically advanced civilisations in the Multiverse as a whole, or in other domains, or in other branes on "braneworld" scenarios, or even in our domain outside the "Hubble Bubble" [according to the chaotic inflationary universe scenario pioneered by physicist Andre Linde, quantum fluctuations divide the inflationary universe into a vast multitude of exponentially large domains or "mini-universes" where the laws of low-energy physics may be different]. Counterintuitively, as Max Tegmark points out, one popular cosmological model apparently predicts that each of us has an effectively identical twin in a galaxy typically around $10^{10^{28}}$ metres away. These distance scales are quite dizzying.

The point in this context is that even if we are unique to the known universe, we need not be "special" - which would entail a rejection of the normal Copernican assumption. If inaccessible civilisations do exist beyond our cosmic event horizon, then their superintelligent inhabitants may well have transcended their evolutionary origins just as we are poised to do too. If such superbeings are benevolent, then they will presumably [given "moral attractors"] rescue others physically accessible to being saved within their light-cone ("Category A"). It would be nice to think that cross-species deliverance from suffering was a universal law; the Objection raises the disturbing possibility ("Category C") that it isn't. The existence of hypothetical advanced lifeforms with the same goals as us but who choose different means ("Category B") might indeed shift the onus of responsibility away from the junior civilization. Yet how common is the multiple independent origin of technologically advanced civilizations within a cosmically narrow (space)time-frame?

This is all extremely speculative. Extensive scanning of the electromagnetic spectrum discloses no evidence that technologically sophisticated life exists in our galaxy, or anywhere else in the observable universe. This absence of evidence extends to what Russian astrophysicist Nikolai Kardashev described as "Type III civilizations" - supercivilizations that would employ the energy resources of an entire galaxy. Their electromagnetic signature could in principle be detected by SETI (Search for ExtraTerrestrial Intelligence) researchers as well. Nothing has been found. The search continues.

Many explanations of "The Great Silence" have been mooted. Why assume, for instance, that intelligent extraterrestrials will manifest anything resembling the motives, values, conceptual framework or colonial expansionism of contemporary *Homo sapiens*? Is our conception of intelligent life and its signature too impoverished for us to have even located the relevant search-space to investigate? But (very) tentatively, the conservative explanation of why an immense ecological niche remains unfilled is that the silence is just what it seems. No technologically advanced, spacefaring civilisations exist within our few billion odd light years neighbourhood. It's up to us.

This conclusion doesn't mean we are locally alone. The Objection is right to take the status of sentient beings in other worlds extremely seriously. If we could really be confident that Earth-based organisms were the only lifeforms in the accessible universe, or if only minimally sentient microbial life exists in other worlds, then eliminating suffering on our planet would effectively discharge our ethical responsibilities. Once our world was cruelty-free, we could retreat into our own private nirvanas - or perhaps build heaven-on-earth and terraform it beyond. Yet it's also possible that complex life and suffering - perhaps intense suffering - exists in alien ecosystems within our cosmic event horizon; and such lifeforms are impotent to do anything about their plight i.e. they are as helpless as are all but one species on contemporary Earth. The presence of such malaise-ridden lifeforms would be undetectable to us with current technology. We have no empirical evidence of their existence one way or the other.

So how likely is such a scenario on theoretical grounds? Life's origins apparently lie early in Earth's 4.6 billion-year history. Deceptively perhaps, its rapid emergence suggests that the process may be relatively "easy" - and thus spontaneously repeated on a massive scale on Earth-like planets across the cosmos. Yet we still can't explain how the primeval "RNA world" preceding our DNA regime came into being. Nor can we yet synthesise life *in vitro,* or computationally simulate its genesis on Earth. So it's quite possible that only a freakish chain of circumstances allowed life to get started in the first instance. Piling improbable event on improbable event, another chain of contingent circumstances over several billion years allowed multicellular eukaryotic life to evolve. Eventually, life arose with the capacity to rewrite its own source code. It's unknown how many significantly different developmental pathways exist leading to organisms capable of scientific technology, or where the biggest evolutionary bottlenecks lie.

There is another imponderable here too. How likely is it that any primordial alien life will undergo suffering, or even be sentient, if its substrate differs from our familiar organic wetware? We know that our silicon (etc.) robots can be programmed to exhibit the quasi-functional analogues of "mental" and "physical" pain and pleasure, and display a repertoire of "emotional" behaviour without any relevant "raw feels". Will putative extraterrestrials likewise be akin to zombie automata - "intelligent" or otherwise? [If so, would their fate matter?] Or more plausibly, will extraterrestrial life be sentient like us (or perhaps hypersentient)?

Here at least we can rationally speculate: the answer is probably the latter, though these modes of sentience may be very different. For there are powerful reasons for thinking that all *primordial* information-bearing self-replicators must be carbon-based owing to the functionally unique valence properties of the carbon atom. Likewise, primordial life-supporting chemistries probably require liquid water. [If and when organic life becomes technologically advanced enough to build silicon robots, create "post-biological" digital life, design self-replicating nanobots, run "simulations" in quantum computers, etc., all bets are off.] If such primordial organic life ever reaches a multicellular stage, then the binary coding system of a pleasure-pain axis embedded in a nervous system is an informationally efficient solution to the challenges of the inner and outer environment, albeit brutishly cruel. So if hypothetical early alien life stumbled upon the molecular mechanisms underlying the pleasure-pain axis, then the information-processing role of its gradients will plausibly have been harnessed by natural selection to boost the

inclusive fitness of self-propelled organisms - as it has on Earth. No "programmer" or designer is needed. Moreover, given the comparatively narrow range of habitats in the physical universe that could sustain primordial multicellular life, the phenomenon of convergent evolution *may* mean that all such life, wherever it evolves, isn't going to be quite so exotic as astrobiologists sometimes suppose. [By contrast, advanced life and consciousness could be unimaginably exotic.] If so, then the same abolitionist blueprint for ecosystem redesign and genomic rewrites should be applicable to other planetary biospheres - if we decide to intervene in Darwinian worlds rather than retain their ecological status quo.

That's a lot of ifs. Right now, it's difficult to care deeply about the plight of creatures who may not even exist, or who may be accessible only to our distant post-human descendants. Ecological charity, one feels, begins at home. Yet such indifference may be a reflection of our limited psychology, not a moral argument for inertia. Naturally, we may all be mistaken in ways that exceed our conceptual resources to imagine or describe. Alternatively, something on the lines of the Objection may be correct. Certainly we rarely, if ever, understand the full ramifications of what we are doing. It's hard enough to plan ahead for the next five years, let alone envisage interstellar travel for the next five million. [This is one good reason not to get trapped in a rut of wirehead hedonism or its chemical counterparts rather than strive for superintelligent well-being.] Yet to opt for a deliberate policy of non-interference - whether in the lives of our suffering fellow humans, non-human animals, or primordial extraterrestrials - is no less morally fraught than paternalistic intervention. The argument that we should do nothing until we fully understand its implications cuts little ice in an *emergency* - and the horrors of a living world where babies get eaten alive by predators, creatures die of hunger, thirst, and cold, etc, must count as morally urgent on all but the most Disneyfied conception of Mother Nature. Analogously, it would be morally reckless for us to shun the use of, say, anaesthetics, pain-killers, veterinary interventions and similar "unnatural" novelties on the grounds that their use poses unknown risks - even though these risks surely exist and should be researched with all possible scientific rigour.

There are indeed ethical pitfalls in "playing God". These pitfalls would be even greater if [as the Objection assumes] there exist god-like extraterrestrial lifeforms better equipped than us to do so. Yet on both a domestic and cosmological scale, moral hazards exist for absentee landlords as well as for hands-on managers. Inaction can be culpable too. Here on Earth, there might seem a moral imperative to intervene and rescue, say, a drowning toddler on (almost) any ethical system at all. But what if that child grows up to be Hitler's grandfather (etc)? We can't know this, since we don't yet carry pocket felicific calculators. Yet the risk is presumably worth taking: we don't let the child drown. Likewise, if your hand is in the fire, you withdraw it. If you are benevolent, then you do the same to rescue a small child or animal companion who is suffering similar agony - whether you are formally a utilitarian ethical theorist or not. The moral sceptic might argue that all value judgements are truth-valueless; but (s)he can't argue consistently that we *ought* to believe this - or behave in one way rather than another. Taking the abolitionist project to the rest of the galaxy and beyond sounds crazy today; but it's the application of technology to a very homely moral precept writ large, not the outgrowth of a revolutionary new ethical theory. So long as sentient beings suffer extraordinary unpleasantness - whether on Earth or perhaps elsewhere - there is a presumptive case to eradicate such suffering wherever it is found.

No 33 "Why does HI lay such stress on gradients of well-being? From an ethical perspective, wouldn't a permanent maximum of bliss be better?"

A motivational system based entirely on heritable gradients of well-being is a less radical prospect than the abolition of motivation altogether. This is because hardwiring *constant* maximum bliss entails discarding the information-signalling role of the pleasure-pain axis completely - not just recalibrating its scale. Barring some extraordinarily advanced technology, uniformly happy beings will be out-reproduced. So for the foreseeable future, at any rate, encoding a physiological maximum of lifelong bliss is simply not an evolutionarily stable strategy. Then there's ideology to consider. If maximising gross cosmic happiness depends on (post-)humans embracing a classical utilitarian value system, it's presumably an unlikely scenario on that score too. Pluralist or perhaps quasi-utilitarian value systems are more sociologically plausible. Yet HI's (tentative) forecast that a motivational regime of gradients of bliss will be conserved indefinitely is itself no more than a conjecture. One counterargument is that choosing less fulfilling states of mind runs counter to the hedonic roots of our decision-making psychology itself. When mature technologies of emotional self-mastery become ubiquitous, it's uncertain who - if anyone - will really settle for what subjectively feels like an inferior option. What dial-settings will rational agents choose for their own mood-range when freed from the old Darwinian roulette? In practice, informed preference utilitarianism and classical utilitarianism tend to converge. Just possibly, the cumulative outcome of our choices may be the transcendence of traditional decision-making. As a slogan, "freedom to control one's emotions" invites readier assent than "freedom to enjoy limitless bliss". What's unclear is whether the *ultimate* cosmic outcome will be substantially different - or ethically, whether it ought to be so. Obviously care should be taken here to separate normative judgement from positive prediction. Certainly, billions of years of pan-galactic hedonism isn't quite what Jeremy Bentham had in mind when first enunciating the greatest happiness principle. A lawyer by training, Bentham had in mind institutional and legislative reform. Yet harnessing biotechnology to a classical utilitarian ethic dictates saturating the cosmos with blissful euphoria/positive value and then computationally sustaining this theoretical maximum indefinitely - whether in the form of discrete superminds or perhaps a Borg-like collective mind. The logic of "hedonistic" utilitarianism is inexorable, even if its premises can be challenged.

The issue of whether we should encode hedonic gradients or constant happiness should be distinguished from the related question of so-called "higher" versus "lower" pleasures i.e. the notional value of whatever we may be happy "about". Gradients of cerebral well-being (or ill-being) can certainly facilitate critical discernment, rational decision-making, and motivated behaviour. Yet as our rapidly evolving computer software attests, neither qualia nor an organic substrate are essential to this functional role. So as our integration with intelligent software increases, the "texture" of subjective dips of bliss may turn out to be functionally unnecessary for sentient organic life too. Tomorrow's technologies of fine-grained emotional control may enable early post-humans, for instance, to amplify their most treasured second-order desires for, say, cultural excellence, intellectual acumen and moral integrity while banishing the baser carnal passions. But after exploring the richest hedonic backdrop to whatever it is one most values - whether highbrow or lowbrow by today's lights - will anyone revert to hedonically impoverished states on discovering what they've been missing? Does our contemporary revulsion from crude wireheading, for instance, lie in the unvarying bliss that it yields - or merely its unedifying focus? Thus it's conceivable, as the Objection implies, that our *distant* descendants will enjoy some kind of ceaseless rapture - perhaps contemplating unimaginably sublime beauty or love or elegant mathematical equations. Or, less portentously, hilariously funny jokes. Naturally, these examples are purely illustrative, since post-humans may be imbued with kinds of blissful experience whose categories *Homo sapiens* can't name or conceive. Perhaps post-humans will be temperamentally meditative; perhaps dynamic. Perhaps they'll live in augmented organic virtual reality; or perhaps they'll live in designer VR paradises run on different bylaws from our presumptive basement. Perhaps they'll inherit a recognisable descendant of ordinary waking primate consciousness; or perhaps they'll live in unknown realms of utopian psychedelia. Unfortunately, our ignorance of the potential varieties of blissful experience

contributes to the misconception that such well-being will necessarily be "thin" or unidimensional rather than diverse. But whatever the scenario, there's indeed no guarantee that a rational superintelligence will tolerate *any* decrements of well-being, information-signalling or otherwise.

The Objector's vision of unvarying bliss doesn't appeal to the dominant Western ethos. For the most part, modern capitalist societies prize innovation, creativity and change. So the prospect of a civilisation based (merely) on gradients of extreme well-being may be less unsettling than a future of constant bliss - though either condition is alien to Darwinian life. We associate permanence with stagnation; and passivity with low motivation and malaise. So any "static" vision fails to inspire. From a broader evolutionary perspective, self-propelled bodies exhibiting goal-directed behaviour arose early in the history of multicellular life on earth. This architecture has been strongly conserved over hundreds of millions of years. Looking ahead to an era when intelligent life has conquered raw suffering, and to an era when we can modulate our core emotions at will, enhanced hedonic gradients and/or their functional analogues may lead our post-human descendants, and/or our intelligent robots/cyborgs, to radiate and colonize every niche of the accessible multiverse within our light cone/galactic supercluster and intelligently re-engineer it. But what then? The (hypothetical) discipline of secular eschatology won't always be the idle fancy it seems at present. After we can effectively ring the changes within the finite state-space of matter and energy in our cosmic neighbourhood, which kinds of supersentience will be judged worth instantiating? To use a lame analogy, will we opt endlessly to replay mediocre games of chess or painting-by-numbers? Or confine ourselves to the state-space of perfection? Is status quo bias as irrational in post-Darwinian paradise as it is in Darwinian purgatory? On the Objector's "constant bliss" scenario, everything formerly unpleasant or mediocre - from avoidance of noxious stimuli to the mundane maintenance of the infrastructure of civilisation - will presumably have been computationally "offloaded" onto our intelligent machines/prostheses. Critically, selection pressure will no longer operate since post-humans will have occupied every possible niche and engineered themselves to have become effectively immortal. The old era of frenetic "action", the sound and fury of imperfect lives played out against a backdrop of restless discontent and scarcity economics, will belong to our animalistic ancestry. Even the transitional era defined by gradients of cerebral euphoria will have been left behind. Quite possibly the molecular signature of all valuable experience will have been identified; and its substrates amplified to the full. Indeed, given the pleasure principle plus advanced technology, an evolutionary trajectory to the presumed attractor of ideal states of sentience may be inescapable. Once the transition to grown-up consciousness is complete, the theoretical possibility of venturing outside this state-space may be even less likely than, say, our now deciding to revisit the lives of savages in caves. If and when intelligent life reaches cosmic superheaven, perhaps the baroque scaffolding that got us there will be kicked away. Eternal bliss needn't be orgasmic in the sense of lacking all intentional objects beyond itself; but presumably even this must be an open question. Either way, "timeless" bliss doesn't have to feel static. Mastery of the neurochemistry of time perception may allow each here-and-now to have a vast temporal depth, a rich internal dynamics, and subjectively to last an eternity. But perhaps speculations about the far future of cosmic consciousness are best avoided.

It should be stressed that all such wild post-Darwinian scenarios are remote - and vastly more speculative than the abolition of suffering or radical motivational enrichment. Hitherto in history, fitness-enhancing gradients of discontent have been the motor of progress - intellectually, socially, aesthetically, morally, personally. Most of the discontent endemic to the living world has indeed been unproductive; but not all of it. So harnessing the information-bearing role of its functional analogues - i.e. dips or anticipated dips of subjective well-being that still feel wonderful, but not sublime - is a more practical stopgap than encoding constant bliss. After all, we're barely on the eve of the reproductive revolution of designer babies, let alone an era of advanced paradise-

engineering. In the near-to-medium term, recalibrating the genetic dial-settings that regulate hedonic tone is a less challenging bioengineering task than offloading *everything* to smart machines and replacing the old motivational and affective homeostatic control mechanisms of organic life completely. Gradient-surfing is also more ideologically realistic. Moreover even on the more conservative gradients-of-bliss scenario, any subjective "cost" of hedonically sub-optimal states i.e. information-signalling dips in well-being - is presumably acceptable to all but the most ardent utilitarian ideologues. Thus in future our hedonic baseline of mental health can still be richer than today's peak experiences. Assuming that the information-signalling role of gradients in well-being is indeed retained, any functional decrements of bliss can still be small. Even if the gradients are exceedingly subtle, there is no risk of a "Buridan's ass" scenario. [Buridan's ass was a mythical mediaeval equine which starved to death from indecision after being presented with the option of two equally appetising stacks of hay]. It's depressives who are prone to procrastinate; by contrast, happy people are typically decisive, extremely happy people more so. Indeed HI predicts that our immediate descendants at least will not be "passively", uniformly happy but *hyper*motivated, albeit on a much higher plateau of well-being than our current neural architecture can support. Enriching the reward centres of contemporary organic life will tend to heighten both its sense of purpose and purposeful behaviour - though to what end we don't know. Admittedly, this association of enhanced motivation with enhanced well-being may only be a contingent fact of our neural architecture - an accident of evolutionary history. The mesolimbic dopamine ("wanting") and mu opioid ("liking") neurotransmitter systems have co-evolved; their functional roles can in principle be disentangled. But again, a separation is scarcely imminent. (Post-)human agency still has a long future.

Depending on the strength of our bioconservative prejudice, gradients of adaptive well-being needn't be heritable. In principle, designer drugs, neurochip implants, nanobots, or autosomal gene therapy could achieve the same result - even within the constraints of a contemporary genome. But if our existing motivational system is defective, then it would seem cruel not to cure the pathology rather than transmit it to future generations. We wouldn't now consider it ethical deliberately to pass on genes for, say, a chronic pain syndrome on the grounds that our future pain-wracked offspring should be "free to choose" whether they wanted to be pain-free or not. Ethically, are our more pervasive syndromes of psychological malaise any different? Why shouldn't mental superhealth be heritable too?

How about the very long-term future? Normative judgements aside, will motivation in the traditional sense endure as long as sentient life itself? Could a future informational economy of mind based on gradients of bliss culminate in some sort of timeless cosmic paradise? Early in the 21st century, at any rate, this sort of question is probably too difficult to answer.

No 34 "Why the headlong rush to paradise engineering? Why not wait until we have the wisdom to understand the implications of what we're doing? Let's get it right."

We are faced with a "bootstrap" problem. Human beings may only ever be wise enough to understand the ramifications of what we're doing after we have enhanced ourselves sufficiently to be able to do so. Perhaps La Rochefoucauld was wiser than he knew: "No man is clever enough to know all the evil he does." Our species may take pains to avoid building a fools' paradise or some sort of Brave New World. But when, and by what means, will we ever be intelligent enough to be sure of succeeding? When will we be wise enough to avoid making mistakes that we haven't even conceived? As the reproductive, infotech and nanotech revolutions unfold, (post-)humans are bound to seek ways to

make ourselves incrementally smarter. Does it really make sense to postpone a parallel *emotional* enrichment - assuming, naïvely, that emotional and cerebral intelligence could be so cleanly divorced? After all, narrowly-conceived intelligence-amplification carries risks of its own; greater wisdom may depend on emotional enrichment rather than being a prerequisite for it. For example, it transpires that genetically engineered "Doogie mice", endowed with an extra copy of the NR2B subtype of NMDA receptor, have not merely superior memories, but a chronically enhanced sensitivity to pain. Imagine if, prior to clinical trials, ambitious prospective human parents had rashly arranged to insert multiple copies of the gene in their designer babies to give them a future competitive advantage in education. The outcome might be pain-ridden child prodigies. Vastly more subtle and complex pitfalls doubtless lie ahead that make *any* steps towards a post-human civilisation problematic, not just paradise-engineering. If the risk-reward ratio of a proposed intervention is unfavourable, then clearly a potentially life-enriching drug, gene therapy (etc) shouldn't be rushed. But sometimes the risk-reward ratio is unclear. A more intractable problem is that some risks may be unknown, or inadequately quantified, or both.

So is the Objection essentially correct? Should we opt to conserve the genetic status quo of Darwinian life? Or at best defer the prospect of distinctively emotional enrichment to the presumed wisdom of our distant descendants?

Delay would be morally reckless for the following reason: ethically, even a non-negative utilitarian can agree that it's *critical* to distinguish between the relief of present suffering and the refinement of future bliss - between the moral urgency of the abolitionist project and the moral luxury of a (hypothetical) full-blown paradise-engineering. The risk-reward ratio of proposed interventions will shift as life on Earth gets progressively better - both for an individual and for civilisation as a whole. We demand a far higher level of proven safety from an improved version of aspirin, for example, than from a potentially life-saving anti-AIDS drug. By parity of reasoning, the same yardstick should apply to their affective counterparts, the different forms of psychological distress. If, fancifully, we were already living in some kind of heaven-on-earth, or even just in a civilised, pain-free society, then it would indeed be foolish to put our well-being at risk by hazardous and premature enhancements designed to make life even better. Bioconservativism might be a wise policy. The Objection might then be tenable. Manifestly, we don't dwell anywhere of the sort.

Compare the introduction of pain-free surgery. In the pre-anaesthetic era, a surgical operation could be tantamount to torture. Patients frequently died. Survivors were often psychologically as well as physically scarred for life. Then a wholly unexpected breakthrough occurred. Within a year of William Morton's demonstration of general anaesthesia at Massachusetts General Hospital in 1846, ether and chloroform anaesthesia were being adopted in operating theatres across the world - in Europe, Asia and Australasia. Instead of embracing this utopian dream-come-true, would it have been wise to wait 30 years while conducting well-controlled trials to see if agents used as general anaesthetics caused delayed-onset brain damage, for instance? Ideally, yes. Should prospective studies have first been undertaken comparing the safety of ether versus chloroform? Again, yes - ideally. Rigorous longitudinal studies would have been more prudent. In the mid-19th Century, there were no professional anaesthesiologists, no balanced anaesthesia, no patient monitoring apparatus, muscle relaxants or endotracheal intubation. The mechanisms of anaesthesia in the central nervous system weren't understood at all. Nor, initially, were the principles of antiseptic surgery: only the combination of anaesthesia plus antisepsis could ever make surgery comparatively safe. If the use of anaesthetics *had* led to delayed-onset long-term brain damage (etc), then the medical doubters might now be hailed as uncommonly prescient - instead of enduring the "enormous condescension of posterity", relegated to a footnote in our incorrigibly Whiggish potted histories of medicine.

Despite these caveats, the world-wide introduction of general anaesthesia in surgery is, by common consent, one of the greatest triumphs of medical history. Why the precipitate haste of its adoption? In essence, anaesthetic use spread rapidly across the world because the horrors of extreme physical pain entailed by surgery without anaesthesia were judged by most (but not all) physicians and their patients to outweigh the potential risks - even though the risks weren't properly known or adequately quantified. Surgeons, too, were able thereafter to attempt ambitious life-saving interventions that were effectively impossible before. By our lights, early anaesthesia was appallingly crude, just as narcotic analgesia remains to this day. But the moral urgency of getting rid of suffering - whether its guise is "physical" or "mental" or both - is obscure only to those not caught in its grip. This is why almost everyone will "break" under torture; and why, globally, hundreds of thousands of depressed people take their own lives each year: in fact "mental" pain effectively kills more people than its nominally physical counterpart. If one is looking for historical role-models, then perhaps Dr John Snow - "the man who made anaesthesia a science" - may serve as an exemplar. As the use of surgical anaesthesia spread like wildfire in the late 1840s, Snow didn't advocate the "safe", bioconservative option of abstinence or delay. That would have been callous. But unlike some of his more gung-ho medical colleagues, Snow was mindful of the potential risks of the seemingly miraculous discovery. His introduction of standardised dosing through efficient inhalers and careful patient monitoring saved many lives. Moral urgency is not a license for recklessness.

Like most analogies, this one is far from exact. Currently millions of sentient creatures, human and non-human, are indeed stricken by suffering no less grievous than patients in the pre-anaesthetic, pre-opioid analgesic era; and likewise, exciting but largely unproven technologies exist to remedy their plight. So to that extent, the historical parallel holds. But statistically, most people are *not* in the throes of extreme psychological distress. Thus *if* one is currently relatively satisfied with one's life, and if one's dependants are relatively satisfied too, then there are strong grounds for caution over experimenting with ill-tested interventions that promise to enhance one's existing well-being. Thus the advent of a putative sustainable mood-enricher to reset one's emotional thermostat, a novel intellect-sparing serenic to banish unwanted anxiety, an illuminating new psychedelic, a super-empathogen, a genius-pill (or whatever) might represent a tantalizing prospect. Yet they should presumably undergo rigorous prior testing before general public licensing - however dazzling the anticipated benefits. It might seem that delay is the only responsible option; there can be wisdom in inaction.

The pitfall to this "safety-first" approach lies in the extreme risk of moral complacency it breeds. Hundreds of millions of human beings, and billions of non-human animals, are not in such a fortunate position. On a universalist utilitarian ethic, or simply a Buddhist-style ethic of compassion, we should systematically apply the same level of urgency to relieving their suffering as one would be justified in exercising if one were oneself tormented by intense pain or suicidal despair. Extreme suffering is the plight of *billions* of sentient beings alive today, whether in our factory-farms, in a Darwinian state of nature, or a depressed neighbour. Desperate straits mandate taking risks one would otherwise shun.

On the face of it, if one aims to lead a cruelty-free lifestyle, one may disclaim personal complicity in such suffering. But this moral opt-out clause may be delusive. Simply by deciding to have genetically unenriched children, for instance, one perpetuates the biology of suffering by bringing more code for its substrates into world. A healthy caution toward untested novelties should not collapse into status quo bias.

Any plea, then, for institutionalized risk-assessment, beefed-up bioethics panels, academic review bodies, worse-case scenario planning, more intensive computer simulations, systematic long-term planning and the institutionalized study of existential

risks is admirable. But so is urgent action to combat the global pandemic of suffering. "The easiest pain to bear is someone else's".

No 35 "HI claims that once the biological substrates of suffering have been abolished, it is 'inconceivable' that suffering will ever be recreated. But this isn't so. According to the Simulation Argument, there is a significant likelihood that we ourselves are living in an ancestor-simulation run by our advanced descendants. If this is the case, then our simulated status entails that posthumans will not eradicate suffering. The Simulation Argument implies that our descendants will re-introduce suffering via their ancestor-simulations, or they never opted to abolish suffering in the first instance."
[http://www.simulation-argument.com/]

The Simulation Argument (SA) is perhaps the first interesting argument for the existence of a Creator in 2000 years. It is worth noting that SA is distinct from the traditional sceptical challenge of how one can ever know that one's senses aren't being manipulated by an evil Cartesian demon, or be sure that one isn't just a brain in a nefarious neurosurgeon's vat, and so forth. SA is also distinct from the controversial but non-sceptical inferential realist theory of perception: inferential realists believe that each of us lives in egocentric simulations of the natural world run by a real organic computer i.e. the mind-brain. Instead, SA claims that given exponential growth in computing processing power and storage capacity, the entire universe as commonly understood could be a simulation run on an ultrapowerful computer built by our distant descendants. We may really be living in one of posterity's versions of The Matrix. SA's important subtlety - the subtlety that catapults SA from idle philosophical fancy to serious scientific metaphysics - is that *if* multiple ancestor-simulations are destined to be created whose inhabitants are subjectively indistinguishable from ourselves, then statistically it is much *more* likely that we are living with the great majority in one of these indistinguishable simulations rather than with the minority in pre-simulation Reality. Or rather, SA concludes that at least one of the following three propositions must be true: 1. Almost all civilisation at our level of development become extinct before becoming technologically mature; 2. The fraction of technologically mature civilizations that are interested in creating ancestor-simulations is almost zero; 3. You are almost certainly living in a computer simulation. Actually, SA's proposed trilemma may shortly be simplified. The first of SA's three disjuncts, the extinction scenario, can be effectively excluded within a century or two - an exclusion that ostensibly increases the likelihood one is living in a cosmic mega-simulation. For humans are poised to colonise worlds beyond the home planet, thereby rendering global thermonuclear war, giant asteroid impacts, a nanotech "grey goo" incident, superlethal viral pandemics and other Earth-ravaging catastrophes impotent to extinguish intelligent life itself. Even on the most apocalyptic end-of-the-world prophecies, intelligent life will presumably survive in at least low-density branches of the universal wave function. In the far future, superintelligent posthumans may *at some stage* mass-produce ancestor-simulations. If so, these computer simulations of ancestral life may include billions of human primates whose inner lives, the simulation hypothesis suggests, may be subjectively indistinguishable from our own.

What should we make of this? First, a familiar sociological point. The dominant technology of an age typically supplies its root-metaphor of mind - and often its root-metaphor of Life, The Universe and Everything. Currently our dominant technology is the digital computer. We may have finally struck lucky. Yet what digital computers have to tell us about the ultimate mysteries of consciousness and existence remains elusive. At any rate, no attempt will be made here exhaustively to discuss SA except insofar as its conclusion impacts on the abolition of suffering. But it's first worth raising a few doubts about the technical feasibility of any kind of simulation hypothesis. These doubts will then

be set aside to consider the likelihood that a notional superintelligence that *did* have the computing technology to run full-blown ancestor-simulations would ever choose to do so.

One problem with SA is that it rests on a philosophical premise for which there is no evidence, namely the substrate-independence of qualia - the introspectively accessible "raw feels" of our mental lives. This premise is probably best rephrased as the substrate-neutrality or substrate-invariance of qualia: SA functionalism doesn't claim that the colours, sounds, smells, emotions, etc, of subjective first-person consciousness can be free-floating, merely that any substrate that can "implement" the computations performed by our neural networks will conserve the textures of human experience. The substrate-neutrality assumption is intended to rule out a [seemingly] arbitrary "carbon chauvinism": take care of the computations, so to speak, and the qualia will take care of themselves. SA aims to quantify the likelihood of our living in an ancestor-simulation with a principle of indifference: the probability that we are living in a simulated universe rather than primordial Reality is equal to the fraction of all people that are actually simulated people. Critically for the argument, SA assumes the subjective indistinguishability of "real" from hypothetical post-biological "simulated" experiences. SA proposes that the power of posthuman supercomputers may allow vastly more simulated copies of people to exist than ever walked the Earth in the ancestral population. This is because once a single "master program" is written, copying its ancestor-files is trivially easy if storage space is available. Hence SA's claim that *if* posthumans ever run ancestor-simulations, then we are almost certainly in one of them. But here is the rub. The prior probability to be assigned to our living in a simulated universe depends on the probability one assigns to the existence of superadvanced civilisations that are both able and willing to create multitudes of sentience-supporting ancestor-simulations. And there is simply no evidence that such computationally simulated virtual "people", if they ever exist, will be endowed with phenomenal consciousness - any more than computationally simulated hurricanes feel wet. SA postulates that consciousness will supervene or "result" from supercomputer programs emulating organic mind/brains with the right causal-functional organization at some suitably fine-grained level of detail. The physical substrates of the putative supercomputer used to simulate sentient creatures like us will supposedly influence our kinds of consciousness only via their influence on computational activities. But it's worth noting that silicon etc robots/computers can *already* emulate and exceed human performance in many domain-specific fields of expertise without any hint of consciousness. It's unclear how or why generalising or extending this performance-gap will switch on inorganic sentience - short of the physical "bionization" of our robots/computers via organic implants. Without qualia, we ourselves would just be brainy zombies; yet qualia are neither necessary nor sufficient for the manifestation of behavioural intelligence. Thus some very stupid organic creatures suffer horribly. Some very smart silicon systems and digital sims aren't sentient; they can defeat the human world-champion at chess. We're clearly missing something: but where are we going wrong?

For SA to work in the absence of a scientific explanation of consciousness, some kind of cross-substrate qualia conservation postulate must be assumed on faith. Yet if phenomenal consciousness is really feasible in other substrates or virtual machines, does this synthetic consciousness have the same *generic* texture as ours - or might not synthetic consciousness be as different as is waking from dreaming (or LSD-like) consciousness? Assuming conscious minds can be "implemented", "uploaded" or "emulated" in other substrates, what grounds are there for supposing that the uploads/simulated minds retain all, or any, *particular* qualia at every virtual level - assuming their specific textures are as computationally incidental to the mind as are the specific compositions of the pieces in a game of chess? Granted biological minds can be scanned, digitized and uploaded to/simulated in another medium, will the hypothetical sentience generated be sub-atomic, nano-, micro-, (or pan-galactic?) in scale? Can abstract virtual machines really generate spatio-temporally located modes of consciousness? Are multiple layers of qualia supposed to be generated by virtual beings

in a nested hierarchy of simulations? Are the stacked qualia supposed to be epiphenomenal i.e. without causal effect; if so, what *causes* subjects like us to refer to their existence? By what mechanism? If ancestor-simulations are being run, then what grounds exist for assuming the conservation of type-identical qualia across multiple layers of abstraction? Are these layers of computational abstraction supposed to be strict or, more realistically, "leaky"? SA undercuts the [ontological] unity of science by treating Reality as though it literally has levels. Yet there is no evidence that *virtual* machines could have the causal power to generate *real* qualia; and the existence of "virtual" qualia would be a contradiction-in-terms.

None of the above considerations entail that phenomenal consciousness or unitary conscious minds are substrate-specific. Perhaps the problem is that there are microfunctional differences between organic and silicon etc computers/robots - microfunctional differences that our putative Simulators might emulate on their supercomputers with software that captures the fine-grained functionality which coarser-gained simulations omit. After all, it's question-begging to describe carbon merely as a "substrate". The carbon atom has functionally unique valence properties and a unique chemistry. The only primordial information-bearing self-replicators in the natural world are organic precisely in virtue of carbon's functional uniqueness. Perhaps the functional uniqueness of organic macromolecules extends to biological sentience. These microfunctional differences may be computationally irrelevant or inessential to a game of chess; but not in other realms. Suppose, for example, that the binding problem [i.e. how the unity of conscious perception is generated by the distributed activities of the brain] and the unitary experiential manifolds of waking/dreaming experience can be explained only by invoking quantum-coherent states in organic mind-brains. Admittedly, this hypothesis resolves the Hard Problem of consciousness only if one grants a monistic idealism/panpsychism that most scientists would find too high a price to swallow. But on this account, the fundamental difference between conscious biological minds and silicon etc computers is that conscious minds are quantum-coherent entities, whereas silicon etc computers (and brains in a dreamless sleep, etc) are effectively mere classical aggregates of microqualia. Counterintuitively, a naturalistic panpsychism actually entails that silicon etc robots are zombies.

A proponent of the simulation hypothesis might respond: So what? A functionally unique organic neurochemistry needn't pose an insurmountable problem for a Simulator. After all, there is no reason to suppose that a classical computer can't formally calculate anything computable on a quantum computer, since (complications aside) a quantum computer is computationally equivalent to a Turing machine, albeit hugely faster. So if silicon etc supercomputers could simulate biological mind-brains with their putative quantum-coherence as well, then qualia might still "emerge" at this layer of abstraction. The technicalities of SA's original, classical formulation aren't essential to the validity of its argument. SA still works if it's recast and the organic mind/brain is a quantum computer. The snag is that this defence of SA conflates the simulation of extrinsic and intrinsic properties: formal input-output relationships and the felt textures of experience. Computational activity that takes milliseconds will not feel the same as computational activity that takes millennia - quite aside from any substrate-specific differences in texture or absence thereof. *If* quantum coherence is the signature of conscious mind, then conscious biological minds are implicated in the fundamental hardware of the universe itself - the computationally expensive, program-resistant stuff of the world. As David Deutsch has stressed, the computations of a quantum computer must be done somewhere. If our minds by their very nature tap into the quantum substrate of basement reality, then this dependence undercuts the grounds for believing that we are statistically likely to inhabit an ancestor-simulation - though it doesn't exclude traditional brain-in-a-vat style scepticism.

Of course, none of the above reasoning is decisive. We simply don't understand consciousness. Many scientists and philosophers would dispute that quantum theory is

even relevant to the problem. Or perhaps we are simulated quantum mind/brains running on a post-silicon quantum supercomputer. Or perhaps the laws of quantum mechanics itself are an artefact of our simulation in some kind of posthuman "computronium". Who knows. Here we are veering into more radical forms of scepticism. But *if* insentient simulations of humans (etc) are feasible, then one may reasonably doubt all three disjuncts of SA. Maybe neither the premises nor the conclusions of SA are true. Intelligent life is not headed for extinction. Some of our descendants may conceivably run multiple ancestor-simulations in low-density branches of the universal wave function. It is exceedingly unlikely that we are participants in one of them.

However, let's set aside technical doubts about computationally simulated sentience. Assume that posthumans have solved the Hard Problem of consciousness. The explanatory gap has been closed without unravelling our entire conceptual scheme in the process. Or perhaps qualia can themselves be digitally encoded and computationally re-created at will. Assume too that some analogue of Moore's Law of computer power is not just a temporary empirical generalisation: computer power continues to increase indefinitely until superintelligence has to grapple with the Bekenstein bound - unless this limit on the entropy or information that can be contained within a three-dimensional volume is itself supposed to disclose the granularity of our simulation. Assume further that a supercivilisation reaches a stage of development where it has the technical capacity to run an abundance of ancestor-simulations and simulate [a fragment of] the multiverse disclosed by contemporary physical science - though computationally simulating the infinite-dimensional Hilbert space of quantum-mechanics is no task for the faint-hearted. Finally, if the ancestor-simulations running are supposed to be cheap simulacra rather than faithful replications, let's assume like SA that the computational savings in taking "reality-shortcuts" outweigh the computational cost of the supervisory software - although in practice the computational price of intervening when ancestor-simulants get too close to discovering their *ersatz* status could make skimping on our Matrix a false computational economy. Granted all the above, then consider the scenario proposed in SA. Of all the *immense* range of alternative activities that future Superbeings might undertake - most presumably inconceivable to us - running ancestor-simulations is one theoretical possibility in a vast state-space of options. On the one hand, posthumans could opt to run paradises for the artificial lifeforms they evolve or create. Presumably they can engineer such heavenly magic for themselves. But for SA purposes, we must imagine that (some of) our successors elect to run malware: to program and replay all the errors, horrors and follies of their distant evolutionary past - possibly in all its classically inequivalent histories, assuming universal QM and maximally faithful ancestor-simulations: there is no unique classical ancestral history in QM. But *why* would posthumans decide to do this? Are our Simulators supposed to be ignorant of the implications of what they are doing - like dysfunctional children who can't look after their pets? Even the superficial plausibility of "running an ancestor-simulation" depends on the description under which the choice is posed. This plausibility evaporates when the option is rephrased. Compare the referentially equivalent question: are our posthuman descendants likely to recreate/emulate Auschwitz? AIDS? Ageing? Torture? Slavery? Child-abuse? Rape? Witch-burning? Genocide? Today a sociopath who announced he planned to stage a terrorist attack in the guise of "running an ancestor-simulation" would be locked up, not given a research grant. SA invites us to consider the possibility that the Holocaust and daily small-scale horrors will be recreated *in future*, at least on our local chronology - a grotesque echo of Nietzschean "eternal recurrence" in digital guise. Worse, since such simulations are so computationally cheap, even the most bestial acts may be re-enacted an untold multitude of times by premeditated posthuman design. It is this hypothetical abundance of computational copies that lends SA's proposal that one may be living in a simulation its argumentative bite. At least the traditional Judeo-Christian Deity was supposed to be benevolent, albeit in defiance of the empirical evidence and discrepancies in the Biblical text. But any Creator/Simulator who opts to run prerecorded ancestor-simulations presumably knows of the deceit practised on the sentient beings it simulates. If the Simulators have indeed deceived us on this score,

then what can we be expected to know of unsimulated Reality that transcends our simulation? What trans-simulation linguistic apparatus of meaning and reference can we devise to speak of what our Deceiver(s) are purportedly up to? Intuitively, one might suppose posthumans may be running copies of us because they find ancestral Darwinian life interesting in some way. After all, we experiment on "inferior" non-human animals and *untermenschen* with whom we share a common ancestry. Might not intellectual curiosity entitle superintelligent beings to treat us in like manner? Or perhaps observing our antics somehow amuses our Simulators - if the homely dramaturgical metaphor really makes any sense. Or perhaps they just enjoy running snuff movies. Yet this whole approach seems misconceived. It treats posthumans as though they were akin to classical Greek gods - just larger-than-life versions of ourselves. Even if advanced beings were to behave in such a manner, would they really choose to create simulated beings that *suffered* - as distinct from formally simulating their ancestral behaviour in the way we computationally simulate the weather?

Unfortunately, this line of thought is long on rhetorical questions and short on definitive proof. A counterargument might be that most humans strongly value life, despite the world's tragedies and its everyday woes. So wouldn't a "like-minded" Superbeing be justified in computationally replaying as many sentient ancestral lives as possible, including Darwinian worlds like our own? Even Darwinian life is sometimes fun, even beautiful. Might not our Simulators regard the episodic nastiness of such worlds as a price worth paying for their blessings - a judgement shared by most non-depressive humans here on Earth. Yet this scenario is problematic even on its own terms. Unless the computing resources accessible to our Simulators were literally infinite, a claim of dubious physical meaning, every simulation has an opportunity-cost in terms of simulated worlds forgone. If one were going to set about creating sentient-life-supporting worlds in a supercomputer, then why not program and run the greatest number of maximally valuable paradises - rather than mediocre or malignant worlds like ours? Presumably posthumans will have mastered the technologies of building super-paradises for themselves, whether physically or via immersive VR. They'll presumably appreciate how sublimely wonderful life can be at its best. So why recreate the ugliness from which they emerged - a perverse descent from posthuman Heaven into Darwinian purgatory? Our own conviction that existing life is worthwhile is itself less a product of disinterested reflection than a (partially) heritable expression of status quo bias. If prompted, we don't believe the world's worst scourges, past or present, should be proliferated if the technical opportunity ever arises. Thus we aim to cure and/or care for the brain-damaged, the mentally ill and victims of genetic diseases; but we don't set out to create *more* brain-damaged, mentally ill and terminally sick children. Even moral primitives like contemporary Darwinian humans would find abhorrent the notion of resurrecting the nastier cruelties of the past. One wouldn't choose to recreate one's last toothache, let alone replay the world's sufferings to date. How likely are posthumans ever to be more backward-looking, in some sense, than us?

Of course, predictions of "progress" in anything but the most amoral, technocratic sense can sound naïve. Extrapolating an exponential growth in computing power, weapons technology or the like sounds reasonable. Extrapolating an expanding circle of compassion to embrace all sentient life sounds fuzzy-minded and utopian. Certainly, given the historical record, envisaging dystopian possibilities is a great deal more plausible than a transition to paradise-engineering. However, a reflex cynicism is itself one of the pathologies of the Darwinian mind. As our descendants rewrite their own code and become progressively smarter, their conception of intelligence will be enriched too. Not least, enriched intelligence will presumably include an enhanced capacity for empathy: a deeper understanding of what it is like to be others - beyond the self-centred perspective of Darwinian minds evolved under pressure of natural selection. An enhanced capacity for empathetic understanding doesn't feature in conventional measures of intelligence. Yet this deficit reflects the inadequacy of our Aspergersish "IQ tests", not the cognitive unimportance of smarter mind-reading and posthuman supersentience. Failure

to appreciate the experience of others, whether human or nonhuman, is not just a moral limitation: it is a profound *intellectual* limitation too; and collective transcendence of humanity's intellectual limitations is an indispensable part of becoming posthuman. If our descendants have any inkling of what it is like to be, say, burned alive as a witch, or to spend all one's life in a veal crate, or simply to be a mouse tormented by a cat, etc, then it seems inconceivable they would set out to (re-)create such terrible states in computer "simulations", ancestral or otherwise. Achieving a God's-eye view that impartially encompasses all sentience may be impossible, even for our most godlike descendants. But posthuman cognitive capacities will presumably transcend the anthropocentric biases of human life. HI argues that posthuman benevolence will extend to the well-being of all sentience; this is technically feasible but speculative.

However, there is a counter to such reassuring arguments. It runs roughly as follows. We can have no insight into the nature of a hypothetical posthuman civilisation that might be capable of running subjectively realistic ancestor-simulations in their supercomputers. Therefore we have no insight into the motivational structure of our Simulators and why they might do this to us. Or perhaps we are merely incidental to their simulation(s) - which exist for a Higher Purpose that we lack the concepts even to express. For instance, perhaps advanced posthumans can command the Planck-scale energies needed hypothetically to create a "universe-in-the-laboratory". For inscrutable reasons, such posthumans might decide to spin off a plethora of baby multiverses, making it statistically more likely that we are living in one of them rather than in the primordial multiverse. If so, we are emulating/simulating our ancestors in another multiverse that spawned us; and we are destined in turn to emulate/simulate our descendants in baby multiverses to come. This scenario contrasts with messy "interventionist" or conspiratorial simulations where posthuman supercomputers are supposed to be constantly rearranging stuff in our simulated world to keep us in ignorance of our artificial status. The point here is that we can't rule out *any* of such scenarios because we know absolutely *nothing* of posthuman ethics - or posthuman values of any kind. Posthuman psychology may simply be unfathomable to *Homo sapiens*, as are our purposes to lesser primates - or to beetles. Or maybe an explanation of our simulated status may be inaccessible to us simply in virtue of our being the ancestor-simulations of real historical people. Our ignorance could be written into the script.

We can't be sure this argument is false. There is nonetheless a problem with the unfathomability response. The prospect of using supercomputers to run ancestor-simulations belongs to the conceptual framework of early 21st Century human primates. The idea resonates with at least a small sub-set of social primates because running ancestor-simulations seems - pre-reflectively, at any rate - the kind of interesting activity that more advanced versions of ourselves might like to pursue. Yet *if* we have no insight into truly posthuman motivations or purposes, or indeed whether such anthropomorphic folk-psychological terms can bear posthuman meaning, then it's hard to assign any significant probability to our successors opting to run sentient ancestor-simulations. In fact given the immense state-space of potential options, and the intrinsic squalor of so much Darwinian life, then the prior probability we should assign to their doing so might seem vanishingly small - even if the technological obstacles could be overcome.

Contrary to the Objection, then, the existence of a world full of suffering is not evidence that our advanced descendants will never abolish its substrates. The existence of suffering is strong presumptive evidence that our descendants will never run sentience-supporting ancestor-simulations.

Chapter 5: CONCLUSION

"The world of the happy is quite different from the world of the unhappy."
(Wittgenstein)

5.0 Puppet-Masters Without Strings

One's attitude to the abolitionist project will be largely a function of the mood in which this manifesto is read. If it succeeds, then the judgement of our blissful descendants on paradise-engineering is likely to be unequivocal. The self-authenticating value of heavenly states of consciousness, and the need to underwrite them genetically to safeguard mental superhealth, will seem compelling. At the other extreme, a significant minority of our contemporaries, diagnosable even today as (sub-)clinically depressed, will welcome the prospect of universal happiness. For a post-Darwinian era of genetically preprogrammed well-being promises a release from their chronic suffering and malaise. Sadly, salvation in the guise of gene-therapy may arrive too late for many of us.

The greatest resistance to the prospect of real-life heaven-on-earth will most likely come from medically ill-named "euthymics". Euthymic mood is statistically typical of products of the present human genome. It's the mood in which we perform our standard "reality-checks". Alas our normal waking discontents are only a brutish parody of mature post-Darwinian mental health. To someone of today's "natural" cast of mind, however, the assent expressed by genetically-enhanced post-humans to lifelong bliss will count for little. After all, a contemporary sceptic may observe, the crack-addict in the throes of an uncontrolled cocaine binge is untroubled by self-doubt either. His rational acumen and practical wisdom are seriously open to question. Likewise, any endorsement of the abolitionist project expressed by depressives will be dismissed too. The sceptic will argue it's just a cognitive pathology consequent on their morbid state.

So we have a bit of an *impasse*. In what mood should this manifesto be appraised? Is there a more-or-less cognitively neutral type of affective state from which the moral worth, and/or practical advantages, of all other affective states can best be judged? When does a mere processing-bias or a cognitive filter take on a hallucinatory aspect that entails certain possibilities are intellectually closed to the victim? Could one be living one's whole life in the grip of an affective psychosis that has infected one's belief and value system to the core?

This whole discussion might seem objectionably psychologistic. All that really counts, one will be told severely, is logical rigour of argument. Rationally, mood doesn't matter. So why extend a woolly, touchy-feely invitation to evaluate the abolitionist project in a blissed-out and presumably uncritical state of mind as well? Surely the essence of paradise-engineering can be understood and appraised, for good or ill, right now?

Unfortunately it's not that simple. We are not disembodied inference-engines. Abstract platonic propositions can be accessed only by abstract platonic minds. From a naturalistic perspective, there are only spatio-temporally located thought-episodes

playing out in flesh-and-blood mind/brains. Their causal sequence of states may partially simulate, but cannot literally instantiate, some notional platonic realm of abstract inference. Anything that physically tends to optimise one's reasoning processes in the natural world should not be lightly dismissed. For in practice the affective, volitional and cognitive aspects to one's thoughts are only notionally separable. Mood and meaning interpenetrate. One's conception of the very nature of Reality itself depends, in large measure, on where one presently finds oneself in the affective spectrum. Perhaps "depressive realism" is realistic relative only to its promordial Darwinian context. There doesn't seem to be any cognitively neutral affective state from which all the others can be impartially judged.

Sadly, medical science cannot hope to resolve the question of putative Ideal Mental States - or whether we should aspire to them if they exist. Which of an organism's psychophysical processes should be classified as pathological or healthy would seem very much a conventional - though not arbitrary - matter of culture, social negotiation and personal prejudice. Mental health and soundness of judgement will tend to be defined, in part, by contemporaneous statistical norms for the population as a whole. And if the average hedonic base-line of our species will indeed be ratcheted upwards via the insertion and orchestrated expression of germ-line "paradise-genes" in our offspring, then the nominal good health of one age can become the terrible psychopathology of a more enlightened era. In retrospect, perhaps all Darwinian hominoids will strike posterity as sick in mind and body alike.

So if one finds oneself viscerally hostile to the idea of universal happiness, and if by contemporary standards one falls within the statistically normal range in one's emotional repertoire, then just how seriously should one contemplate the following possibility? Today we are the victims of what our successors will reckon an atavistic mood disorder. This disorder infects all our thoughts as well as all our feelings and volitions. It is a historical condition no less epistemically defective than are dream-psychoses from the perspective of the waking state.

Is the worry one might be locked in such an affective psychosis just the product of idle scepticism? Given the cognitive inaccessibility of most of the generically ecstatic states alluded to here, perhaps one wouldn't know if one were so afflicted. After all, damaged and disfigured minds may have limited self-insight. Nor would one necessarily have the conceptual resources even to grasp what was at stake if one suffered from such a neural deficit. Pure, "unearned", genetically-driven bliss of even the mildest flavour detracted from the inclusive fitness of one's genes in the ancestral environment. Constitutionally happy freaks-of-nature got eaten or outbred. Hence unipolar euphoric mania today is vanishingly rare; unipolar melancholic depression and chronic dysthymia are all too common. Is one's potential unease, if not revulsion, at the prospect of paradise-on-earth an incidental cultural by-product of natural selection? Or has selection pressure ensured that one is genetically predisposed to be biased against the idea of enduring bliss in the first instance?

5.1 Could Life Really Have A Happy Ending?

It's time to take stock. Most of the more exotic delights sketched out in this manifesto will probably never be enjoyed by the reader. They require a level of theoretical understanding and biomedical expertise that we simply do not yet command. Many of the practical difficulties that the abolitionist project must overcome have been skated over here with the kind of blithe disregard for detail that only an ignorance of nitty-gritty technical complexities can bestow. If, however, a single major government, charitable foundation or segment of the global power elite were to sanction the necessary research

and development, then gradients of sustainable, chemically-underwritten euphoria are quite tantalisingly accessible, even now, to those of us who want psychological superhealth. Better still, germ-line gene therapy can then turn gradients of lifelong ecstatic well-being into the natural post-human condition. A hereditary condition of invincible well-being may prove to be the foundation on which *any* advanced civilisation is built. The option of rewriting the vertebrate genome, and redesigning the global ecosystem, extends the prospect of paradise-engineering to the rest of the living world.

Admittedly, in the absence of concerted international action to promote at least a skeletal world-wide counterpart to the national welfare-state, the wretched plight of much of the world's population means that any instant dash to raw, unempathetic euphoria on the part of a materially privileged minority would be premature. It would be selfish in the extreme - though not necessarily more so than the life-styles of competitive individualism, rampant consumerism and incompetent recreational drug-abuse that many of us live at present. Yet one of the providential blessings of the abolitionist project is that - with a decent bit of planning - it can supplant the old, quasi-zero-sum approach to the allocation of life's rewards. If properly managed, the route to felicific enlightenment ahead will soon be genetically open to all. Lifelong well-being needn't be the preserve of the affluent few. Nor need lifelong well-being be the reward solely of the morally good and "deserving". In fact with a combination of cognitive-enhancers ("smart drugs") and gentle euphoriants, there is no reason why the old age of the sympathetic reader shouldn't herald, not a slow, spirit-sapping decline, but a period of beautiful experiences and glorious self-fulfilment. Thus later life can be a time immeasurably richer than anything (s)he has enjoyed before.

Many people will have internalised too many of the life-impoverishing hang-ups of humanity's biological past to contemplate playing a pioneering role and participating in the era ahead; just as misplaced prudery prevents many people from enjoying sex. But life, one may think, should climax in an orgasmic celebration of being, not a fatalistic world-weary fade-out.

1995. Last updated (Reply to Objection 35) in 2007