TRANSHUMANISM AND THE BODY

THE WORLD RELIGIONS SPEAK

EDITED BY
CALVIN MERCER
AND
DEREK F. MAHER

2014
palgrave
macmillan

INTRODUCTION—THE TRANSHUMANIST FAQ: A GENERAL INTRODUCTION

NICK BOSTROM*

GENERAL QUESTIONS ABOUT TRANSHUMANISM

WHAT IS TRANSHUMANISM?

Transhumanism is a way of thinking about the future that is based on the premise that the human species in its current form does not represent the end of our development but rather a comparatively early phase. We formally define it as follows:

1. The intellectual and cultural movement that affirms the possibility and desirability of fundamentally improving the human condition through applied reason, especially by developing and making widely available technologies to eliminate aging and to greatly enhance human intellectual, physical, and psychological capacities.
2. The study of the ramifications, promises, and potential dangers of technologies that will enable us to overcome fundamental human limitations, and the related study of the ethical matters involved in developing and using such technologies.

Transhumanism can be viewed as an extension of humanism, from which it is partially derived. Humanists believe that humans matter and that individuals matter. We might not be perfect, but we can make things better by promoting rational thinking, freedom, tolerance, democracy, and concern for our fellow human beings. Transhumanists agree with this proposition but also emphasize what we have the potential to become. Just as we use rational means to improve the human condition and the external world, we can also use such means to improve ourselves, the human organism.

In doing so, we are not limited to traditional humanistic methods, such as education and cultural development. We can also use technological means that will eventually enable us to move beyond what some would think of as "human."

It is not our human shape or the details of our current human biology that define what is valuable about us, but rather our aspirations and ideals, our experiences, and the kinds of lives we lead. To the transhumanist, progress occurs when more people become more able to shape themselves, their lives, and the ways they relate to others, in accordance with their own deepest values. Transhumanists place a high value on autonomy: the ability and right of individuals to plan and choose their own lives. Some people may, of course, for any number of reasons, choose to forgo the opportunity to use technology to improve themselves. Transhumanists seek to create a world in which autonomous individuals may choose to remain unenhanced or choose to be enhanced and in which these choices will be respected.

Through the accelerating pace of technological development and scientific understanding, we are entering a whole new stage in the history of the human species. In the relatively near future, we may face the prospect of real artificial intelligence. New kinds of cognitive tools will be built that combine artificial intelligence with interface technology. Molecular nanotechnology has the potential to manufacture abundant resources for everybody and to give us control over the biochemical processes in our bodies, enabling us to eliminate disease and aging. Technologies such as brain-computer interfaces and neuropharmacology could amplify human intelligence, increase emotional well-being, improve our capacity for steady commitment to life projects or a loved one, and even multiply the range and richness of possible emotions. On the dark side of the spectrum, transhumanists recognize that some of these emerging technologies could potentially cause great harm to human life; even the survival of our species could be at risk. Seeking to understand the dangers and working to prevent disasters is an essential part of the transhumanist agenda.

Transhumanism is entering the mainstream culture today, as an increasing number of scientists, scientifically literate philosophers, and social thinkers are beginning to take seriously the range of possibilities that transhumanism encompasses. A rapidly expanding family of transhumanist groups, differing somewhat in flavor and focus, and a plethora of discussion groups in many countries around the world, are gathered under the umbrella of Humanity+ (formerly World Transhumanist Association), a nonprofit democratic membership organization.

WHAT IS A POSTHUMAN?

It is sometimes useful to talk about possible future beings whose basic capacities so radically exceed those of present humans as to be no longer unambiguously human by our current standards. The standard word for such beings is "posthuman." (Care must be taken to avoid misinterpretation. "Posthuman" does not denote just anything that happens to come after the human era, nor does it have anything to do with the "posthumous." In particular, it does not imply that there are no humans anymore.)

Many transhumanists wish to follow life paths that would, sooner or later, require growing into posthuman persons: they yearn to reach intellectual heights as far above any current human genius as humans are above other primates; to be resistant to disease and impervious to aging; to have unlimited youth and vigor; to exercise control over their own desires, moods, and mental states; to be able to avoid feeling tired, hateful, or irritated about petty things; to have an increased capacity for pleasure, love, artistic appreciation, and serenity; to experience novel states of consciousness that current human brains cannot access. It seems likely that the simple fact of living an indefinitely long, healthy, active life would take anyone to posthumanity if they went on accumulating memories, skills, and intelligence.

Posthumans could be completely synthetic artificial intelligences, or they could be enhanced uploads (see the section, "What is Uploading?"), or they could be the result of making many smaller but cumulatively profound augmentations to a biological human. The last alternative would probably require either the redesign of the human organism using advanced nanotechnology or its radical enhancement using some combination of technologies, such as genetic engineering, psychopharmacology, antiaging therapies, neural interfaces, advanced information management tools, memory-enhancing drugs, wearable computers, and cognitive techniques.

Some authors write as though simply by changing our self-conception, we have become or could become posthuman. This is a confusion or corruption of the original meaning of the term. The changes required to make us posthuman are too profound to be achievable by merely altering some aspect of psychological theory or the way we think about ourselves. Radical technological modifications to our brains and bodies are needed.

It is difficult for us to imagine what it would be like to be a posthuman person. Posthumans may have experiences and concerns that we cannot fathom, thoughts that cannot fit into the three-pound lumps of neural tissue that we use for thinking. Some posthumans may find it advantageous to jettison their bodies altogether and live as information patterns on vast

superfast computer networks. Their minds may be not only more powerful than ours but may also employ different cognitive architectures or include new sensory modalities that enable greater participation in their virtual reality settings. Posthuman minds might be able to share memories and experiences directly, greatly increasing the efficiency, quality, and modes in which posthumans could communicate with each other. The boundaries between posthuman minds may not be as sharply defined as those between humans.

Posthumans might shape themselves and their environment in so many new and profound ways that speculations about the detailed features of posthumans and the posthuman world are likely to fail.

WHAT IS A TRANSHUMAN?

In its contemporary usage, "transhuman" refers to an intermediary form between the human and the posthuman (see the section, "What Is a Posthuman?"). One might ask, given that our current use of, for example, medicine and information technology enable us to routinely do many things that would have astonished humans living in ancient times, whether we are not already transhuman? The question is a provocative one, but ultimately not very meaningful; the concept of the transhuman is too vague for there to be a definite answer.

A transhumanist is simply someone who advocates transhumanism (see the section, "What is Transhumanism?"). It is a common error for reporters and other writers to say that transhumanists "claim to be transhuman" or "call themselves transhuman." To adopt a philosophy that says that someday everyone ought to have the chance to grow beyond present human limits is clearly not to say that one is better or somehow currently "more advanced" than one's fellow humans.

The etymology of the term "transhuman" goes back to the futurist FM-2030 (also known as F. M. Estfandiary), who introduced it as shorthand for "transitional human." Calling transhumans the "earliest manifestation of new evolutionary beings," FM maintained that signs of transhumanity included prostheses, plastic surgery, intensive use of telecommunications, a cosmopolitan outlook and a globetrotting lifestyle, androgyny, mediated reproduction (such as in vitro fertilization), absence of religious beliefs, and a rejection of traditional family values. However, FM's diagnostics are of dubious validity. It is unclear why anybody who has a lot of plastic surgery or a nomadic lifestyle is any closer to becoming a posthuman than the rest of us. Nor, of course, are such persons necessarily more admirable or morally commendable than others. In fact, it is perfectly possible to be a

transhuman—or, for that matter, a transhumanist—and still embrace most traditional values and principles of personal conduct.

TECHNOLOGIES AND PROJECTIONS

BIOTECHNOLOGY, GENETIC ENGINEERING, STEM CELLS, AND CLONING: WHAT ARE THEY AND WHAT ARE THEY GOOD FOR?

Biotechnology is the application of techniques and methods based on the biological sciences. It encompasses such diverse enterprises as brewing, manufacture of human insulin, interferon, and human growth hormone, medical diagnostics, cell cloning and reproductive cloning, the genetic modification of crops, bioconversion of organic waste and the use of genetically altered bacteria in the cleanup of oil spills, stem cell research, and much more. Genetic engineering is the area of biotechnology concerned with the directed alteration of genetic material.

Biotechnology already has countless applications in industry, agriculture, and medicine. It is a hotbed of research. The completion of the human genome project—a "rough draft" of the entire human genome was published in the year 2000—was a scientific milestone by anyone's standards. Research is now shifting to decoding the functions and interactions of all these different genes and to developing applications based on this information.

The potential medical benefits are too many to list; researchers are working on every common disease, with varying degrees of success. Progress takes place not only in the development of drugs and diagnostics but also in the creation of better tools and research methodologies, which in turn accelerates progress. When considering what developments are likely over the long term, such improvements in the research process itself must be factored in. The human genome project was completed ahead of schedule, largely because the initial predictions underestimated the degree to which instrumentation technology would improve during the course of the project. At the same time, one needs to guard against the tendency to hype every latest advance. (Remember all those breakthrough cancer cures that we never heard of again?) Moreover, even in cases where the early promise is borne out, it usually takes ten years to get from proof-of-concept to successful commercialization.

Genetic therapies are of two types: somatic and germ line. In somatic gene therapy, a virus is typically used as a vector to insert genetic material into the cells of the recipient's body. The effects of such interventions do not carry over into the next generation. Germ-line genetic therapy is performed on sperm or egg cells, or on the early zygote, and can be

inheritable. (Embryo screening, in which embryos are tested for genetic defects or other traits and then selectively implanted, can also count as a kind of germ line intervention.) Human gene therapy, except for some forms of embryo screening, is still experimental. Nonetheless, it holds promise for the prevention and treatment of many diseases, as well as for uses in enhancement medicine. The potential scope of genetic medicine is vast: virtually all disease and all human traits—intelligence, extroversion, conscientiousness, physical appearance, etc.—involve genetic predispositions. Single-gene disorders, such as cystic fibrosis, sickle cell anemia, and Huntington's disease, are likely to be among the first targets for genetic intervention. Polygenic traits and disorders, ones in which more than one gene is implicated, may follow later (although even polygenic conditions can sometimes be influenced in a beneficial direction by targeting a single gene).

Stem cell research, another scientific frontier, offers great hopes for regenerative medicine. Stem cells are undifferentiated (unspecialized) cells that can renew themselves and give rise to one or more specialized cell types with specific functions in the body. By growing such cells in culture, or steering their activity in the body, it will be possible to grow replacement tissues for the treatment of degenerative disorders, including heart disease, Parkinson's, Alzheimer's, diabetes, and many others. It may also be possible to grow entire organs from stem cells for use in transplantation. Embryonic stem cells seem to be especially versatile and useful, but research is also ongoing into adult stem cells and the "reprogramming" of ordinary cells so that they can be turned back into stem cells with pluripotent capabilities.

The term "human cloning" covers both therapeutic and reproductive uses. In therapeutic cloning, a preimplantation embryo (also known as a "blastocyst," a hollow ball consisting of 30–150 undifferentiated cells) is created via cloning, and from it embryonic stem cells could be extracted and used for therapy. Because these cloned stem cells are genetically identical to the patient, the tissues or organs they would produce could be implanted without eliciting an immune response from the patient's body, thereby overcoming a major hurdle in transplant medicine. Reproductive cloning, by contrast, would mean the birth of a child who is genetically identical to the cloned parent: in effect, a younger identical twin.

Everybody recognizes the benefit to ailing patients and their families that come from curing specific diseases. Transhumanists emphasize that, in order to seriously prolong the healthy life-span, we also need to develop ways to slow aging or to replace senescent cells and tissues. Gene therapy, stem cell research, therapeutic cloning, and other areas of medicine that have the potential to deliver these benefits deserve a high priority in the allocation of research monies.

Biotechnology can be seen as a special case of the more general capabilities that nanotechnology will eventually provide [see "What Is Molecular Nanotechnology?"].

WHAT IS SUPERINTELLIGENCE?

A superintelligent intellect (a superintelligence, sometimes called "ultraintelligence") is one that has the capacity to radically outperform the best human brains in practically every field, including scientific creativity, general wisdom, and social skills.

Sometimes a distinction is made between weak and strong superintelligence. Weak superintelligence is what you would get if you could run a human intellect at an accelerated clock speed, such as by uploading it to a fast computer (see the section, "What Is Uploading?"). If the upload's clock-rate were a thousand times that of a biological brain, it would perceive reality as being slowed down by a factor of a thousand. It would think a thousand times more thoughts in a given time interval than its biological counterpart.

Strong superintelligence refers to an intellect that is not only faster than a human brain but also smarter in a qualitative sense. No matter how much you speed up your dog's brain, you're not going to get the equivalent of a human intellect. Analogously, there might be kinds of smartness that wouldn't be accessible to even very fast human brains given their current capacities. Something as simple as increasing the size or connectivity of our neuronal networks might give us some of these capacities. Other improvements may require wholesale reorganization of our cognitive architecture or the addition of new layers of cognition on top of the old ones.

However, the distinction between weak and strong superintelligence may not be clear-cut. A sufficiently long-lived human who didn't make any errors and had a sufficient stack of scrap paper at hand could in principle compute any Turing computable function. (According to Alonzo Church's thesis, the class of Turing computable functions is identical to the class of physically computable functions.)

Many but not all transhumanists expect that superintelligence will be created within the first half of this century. Superintelligence requires two things: hardware and software. Chip manufacturers planning the next generation of microprocessors commonly rely on a well-known empirical regularity known as Moore's law. Formulated in 1965 by Intel cofounder Gordon Moore, the law originally stated that the number of components on a chip doubled every year. In contemporary use, the "law" is commonly understood as referring more generally to a doubling of computing power, or of computing power per dollar. The human brain's processing power

is difficult to determine precisely, but common estimates range from 10^{14} instructions per second (IPS) up to 10^{17} IPS or more. The lower estimate, derived by Carnegie Mellon robotics professor Hans Moravec, is based on the computing power needed to replicate the signal processing performed by the human retina and assumes a significant degree of software optimization. The 10^{17} IPS estimate is obtained by multiplying the number of neurons in a human brain (~100 billion) with the average number of synapses per neuron (~1,000) and with the average spike rate (~100 Hz), and assuming ~10 instructions to represent the effect on one action potential traversing one synapse. An even higher estimate would be obtained, for example, if one were to suppose that functionally relevant and computationally intensive processing occurs within compartments of a dendrite tree.

Most experts, Moore included, think that computing power will continue to double about every 18 months for at least another two decades. This expectation is based in part on extrapolation from the past and in part on consideration of the developments currently underway in laboratories. The fastest computer under construction is IBM's Blue Gene/L, which when it is ready in 2005 is expected to perform ~$2*10^{14}$ IPS. Thus it appears quite likely that human-equivalent hardware will have been achieved within not much more than a couple of decades.

How long it will take to solve the software problem is harder to estimate. One possibility is that progress in computational neuroscience will teach us about the computational architecture of the human brain and what learning rules it employs. We can then implement the same algorithms on a computer. In this approach, the superintelligence would not be completely specified by the programmers but would instead have to grow by learning from experience the same way a human infant does. An alternative approach would be to use genetic algorithms and methods from classical artificial intelligence (AI). This might result in a superintelligence that bears no close resemblance to a human brain. At the opposite extreme, we could seek to create a superintelligence by uploading a human intellect and then accelerating and enhancing it (see the section, "What Is Uploading?"). The outcome of this might be a superintelligence that is a radically upgraded version of one particular human mind.

The arrival of superintelligence will clearly deal a heavy blow to anthropocentric world views. Much more important than its philosophical implications, however, would be its practical effects. Creating superintelligence may be the last invention that humans will ever need to make, since superintelligences could themselves take care of further scientific and technological development. They would do so more effectively than humans. Biological humanity would no longer be the smartest life form on the block.

The prospect of superintelligence raises many big issues and concerns that we should think deeply about in advance of its actual development. The paramount question is: What can be done to maximize the chances that the arrival of superintelligence will benefit rather than harm us? The range of expertise needed to address this question extends far beyond the community of AI researchers. Neuroscientists, economists, cognitive scientists, computer scientists, philosophers, ethicists, sociologists, science-fiction writers, military strategists, politicians, legislators, and many others will have to pool their insights if we are to deal wisely with what may be the most important task our species will ever have to tackle.

Many transhumanists would like to become superintelligent themselves. This is obviously a long-term and uncertain goal, but it might be achievable either through uploading and subsequent enhancement or through the gradual augmentation of our biological brains, by means of future nootropics (cognitive enhancement drugs), cognitive techniques, IT tools (e.g. wearable computers, smart agents, information filtering systems, visualization software, etc.), neural-computer interfaces, or brain implants.

WHAT IS UPLOADING?

Uploading (sometimes called "downloading," "mind uploading," or "brain reconstruction") is the process of transferring an intellect from a biological brain to a computer.

One way of doing this might be by first scanning the synaptic structure of a particular brain and then implementing the same computations in an electronic medium. A brain scan of sufficient resolution could be produced by disassembling the brain atom for atom by means of nanotechnology. Other approaches, such as analyzing pieces of the brain slice by slice in an electron microscope with automatic image processing have also been proposed. In addition to mapping the connection pattern among the 100 billion-or-so neurons, the scan would probably also have to register some of the functional properties of each of the synaptic interconnections, such as the efficacy of the connection and how stable it is over time (e.g., whether it is short-term or long-term potentiated). Nonlocal modulators such as neurotransmitter concentrations and hormone balances may also need to be represented, although such parameters likely contain much less data than the neuronal network itself.

In addition to a good three-dimensional map of a brain, uploading will require progress in neuroscience to develop functional models of each species of neuron (how they map input stimuli to outgoing action potentials, and how their properties change in response to activity in learning). It will also require a powerful computer to run the upload and some way for the upload

to interact with the external world or with a virtual reality. (Providing input/output or a virtual reality for the upload appears easy in comparison to the other challenges.)

An alternative hypothetical uploading method would proceed more gradually: One neuron could be replaced by an implant or by a simulation in a computer outside of the body. Then another neuron, and so on, until eventually the whole cortex has been replaced and the person's thinking is implemented on entirely artificial hardware. (To do this for the whole brain would almost certainly require nanotechnology.)

A distinction is sometimes made between destructive uploading, in which the original brain is destroyed in the process, and nondestructive uploading, in which the original brain is preserved intact alongside the uploaded copy. It is a matter of debate under what conditions personal identity would be preserved in destructive uploading. Many philosophers who have studied the problem think that at least under some conditions, an upload of your brain would be you. A widely accepted position is that you survive so long as certain information patterns are conserved, such as your memories, values, attitudes, and emotional dispositions, and so long as there is causal continuity so that earlier stages of yourself help determine later stages of yourself. Views differ on the relative importance of these two criteria, but they can both be satisfied in the case of uploading. For the continuation of personhood, on this view, it matters little whether you are implemented on a silicon chip inside a computer or in that gray, cheesy lump inside your skull, assuming both implementations are conscious.

Tricky cases arise, however, if we imagine that several similar copies are made of your uploaded mind. Which one of them is you? Are they all you, or are none of them you? Who owns your property? Who is married to your spouse? Philosophical, legal, and ethical challenges abound. Maybe these will become hotly debated political issues later in this century.

A common misunderstanding about uploads is that they would necessarily be "disembodied" and that this would mean that their experiences would be impoverished. Uploading, according to this view, would be the ultimate escapism, one that only neurotic body-loathers could possibly feel tempted by. But an upload's experience could in principle be identical to that of a biological human. An upload could have a virtual (simulated) body giving the same sensations and the same possibilities for interaction as a nonsimulated body. With advanced virtual reality, uploads could enjoy food and drink, and upload sex could be as gloriously messy as one could wish. And uploads wouldn't have to be confined to virtual reality: They could interact with people on the outside and even rent robot bodies in order to work in or explore physical reality.

Personal inclinations regarding uploading differ. Many transhumanists have a pragmatic attitude: whether they would like to upload or not depends on the precise conditions in which they would live as uploads and what the alternatives are. (Some transhumanists may also doubt whether uploading will be possible.) Advantages of being an upload would include:

- Uploads would not be subject to biological senescence.
- Back-up copies of uploads could be created regularly so that you could be rebooted if something bad happened. (Thus your life-span would potentially be as long as the universe's.)
- You could potentially live much more economically as an upload since you wouldn't need physical food, housing, transportation, etc.
- If you were running on a fast computer, you would think faster than in a biological implementation. For instance, if you were running on a computer a thousand times more powerful than a human brain, then you would think a thousand times faster (and the external world would appear to you as if it were slowed down by a factor of a thousand). You would thus get to experience more subjective time, and live more, during any given day.
- You could travel at the speed of light as an information pattern, which could be convenient in a future age of large-scale space settlements.
- Radical cognitive enhancements would likely be easier to implement in an upload than in an organic brain.

A couple of other points about uploading:

- Uploading should work for cryonics patients provided their brains are preserved in a sufficiently intact state.
- Uploads could reproduce extremely quickly (simply by making copies of themselves). This implies that resources could very quickly become scarce unless reproduction is regulated.

TRANSHUMANISM AND NATURE

WHY DO TRANSHUMANISTS WANT TO LIVE LONGER?

This is a personal matter, a matter of the heart. Have you ever been so happy that you felt like melting into tears? Has there been a moment in your life of such depth and sublimity that the rest of existence seemed like dull, gray slumber from which you had only just woken up?

It is so easy to forget how good things can be when they are at their best. But on those occasions when we do remember—whether it comes from

the total fulfillment of being immersed in creative work or from the tender ecstasy of reciprocated love—then we realize just how valuable every single minute of existence can be, when it is this good. And you might have thought to yourself, "It ought to be like this always. Why can't this last forever?"

Well, maybe—just maybe—it could.

When transhumanists seek to extend human life, they are not trying to add a couple of extra years at a care home spent drooling at one's shoes. The goal is more years that are healthy, happy, and productive. Ideally, everybody should have the right to choose when and how to die—or not to die. Transhumanists want to live longer because they want to do, learn, and experience more; have more fun and spend more time with loved ones; continue to grow and mature beyond the paltry eight decades allotted to us by our evolutionary past; and in order to get to see for themselves what wonders the future might hold. As the sales pitch for one cryonics organization goes:

> The conduct of life and the wisdom of the heart are based upon time; in the last quartets of Beethoven, the last words and works of "old men" like Sophocles and Russell and Shaw, we see glimpses of a maturity and substance, an experience and understanding, a grace and a humanity, that isn't present in children or in teenagers. They attained it because they lived long; because they had time to experience and develop and reflect; time that we might all have. Imagine such individuals—a Benjamin Franklin, a Lincoln, a Newton, a Shakespeare, a Goethe, an Einstein (and a Gandhi)—enriching our world not for a few decades but for centuries. Imagine a world made of such individuals. It would truly be what Arthur C. Clarke called "Childhood's End"—the beginning of the adulthood of humanity. (Cryonics Institute. http://www.cryonics.org/Accessed June 30, 2014.)

ISN'T THIS TAMPERING WITH NATURE?

Absolutely, and it is nothing to be ashamed of. It is often right to tamper with nature. One could say that manipulating nature is an important part of what civilization and human intelligence is all about; we have been doing it since the invention of the wheel. Alternatively, one could say that since we are part of nature, everything we do and create is in a sense natural too. In any case, there is no moral reason why we shouldn't intervene in nature and improve it if we can, whether by eradicating diseases, improving agricultural yields to feed a growing world population, putting communication satellites into orbit to provide homes with news and entertainment, or inserting contact lenses in our eyes so we can see better. Changing nature for the better is a noble and glorious thing for humans to do. (On the other hand, to "pave

paradise to put up a parking lot" would not be glorious; the qualification "for the better" is essential.)

In many particular cases, of course, there are sound practical reasons for relying on "natural" processes. The point is that we cannot decide whether something is good or bad simply by asking whether it is natural or not. Some natural things are bad, such as starvation, polio, and being eaten alive by intestinal parasites. Some artificial things are bad, such as DDT-poisoning, car accidents, and nuclear war.

To pick a topical example, consider the debate about human cloning. Some argue that cloning humans is not unnatural because human clones are essentially just identical twins. They were right in this, of course, although one could also correctly remark that it is not natural for identical twins to be of different ages. But the more fundamental point is that it doesn't matter whether human clones are natural or not. When thinking about whether to permit human reproductive cloning, we have to compare the various possible desirable consequences with the various possible undesirable consequences. We then have to try to estimate the likelihood of each of these consequences. This kind of deliberation is much harder than simply dismissing cloning as unnatural, but it is also more likely to result in good decisions.

These remarks hopefully should seem trivial. Yet it is astonishing how often polemicists can still get away with arguments that are basically (thinly disguised) ways of saying, "It is good because it's the way it has always been!" or "It is good because that's the way Nature made it!"

WILL TRANSHUMAN TECHNOLOGIES MAKE US INHUMAN?

The important thing is not to be human but to be humane. Though we might wish to believe that Hitler was an inhuman monster, he was, in fact, a human monster; and Gandhi is noted not for being remarkably human but for being remarkably humane.

The attributes of our species are not exempt from ethical examination in virtue of being "natural" or "human." Some human attributes, such as empathy and a sense of fairness, are positive; others, such as tendencies toward tribalism or groupishness, have left deep scars on human history. If there is value in being human, it does not comes from being "normal" or "natural," but from having within us the raw material for being humane: compassion, a sense of humor, curiosity, the wish to be a better person. Trying to preserve "humanness," rather than cultivating humaneness, would idolize the bad along with the good. One might say that if "human" is what we are, then "humane" is what we, as humans, wish we were. Human nature is not a bad place to start that journey, but we can't fulfill that potential if we reject any progress past the starting point.

ISN'T DEATH PART OF THE NATURAL ORDER OF THINGS?

Transhumanists insist that whether something is natural or not is irrelevant to whether it is good or desirable [see also "Isn't This Tampering with Nature?" and "Why Do Transhumanists Want to Live Longer?"].

Average human life-span hovered between 20 and 30 years for most of our species' history. Most people today are thus living highly unnaturally long lives. Because of the high incidence of infectious disease, accidents, starvation, and violent death among our ancestors, very few of them lived much beyond 60 or 70. There was therefore little selection pressure to evolve the cellular repair mechanisms (and pay their metabolic costs) that would be required to keep us going beyond our meager three scores and ten. As a result of these circumstances in the distant past, we now suffer the inevitable decline of old age: damage accumulates at a faster pace than it can be repaired; tissues and organs begin to malfunction; and then we keel over and die.

The quest for immortality is one of the most ancient and deep-rooted of human aspirations. It has been an important theme in human literature from the very earliest preserved written story, *The Epic of Gilgamesh*, and in innumerable narratives and myths ever since. It underlies the teachings of world religions about spiritual immortality and the hope of an afterlife. If death is part of the natural order, so too is the human desire to overcome death.

Before transhumanism, the only hope of evading death was through reincarnation or otherworldly resurrection. Those who viewed such religious doctrines as figments of our own imagination had no alternative but to accept death as an inevitable fact of our existence. Secular world views, including traditional humanism, would typically include some sort of explanation of why death was not such a bad thing after all. Some existentialists even went so far as to maintain that death was necessary to give life meaning.

That people should make excuses for death is understandable. Until recently there was absolutely nothing anybody could do about it, and it made some degree of sense then to create comforting philosophies according to which dying of old age is a fine thing ("deathism"). If such beliefs were once relatively harmless, and perhaps even provided some therapeutic benefit, they have now outlived their purpose. Today, we can foresee the possibility of eventually abolishing aging, and we have the option of taking active measures to stay alive until then, through life extension techniques and, as a last resort, cryonics. This makes the illusions of deathist philosophies dangerous, indeed fatal, since they teach helplessness and encourage passivity.

Espousing a deathist viewpoint tends to go with a certain element of hypocrisy. It is to be hoped and expected that a good many of death's

apologists, if they were one day presented with the concrete choice between (A) getting sick, old, and dying, and (B) being given a new shot of life to stay healthy, vigorous, and to remain in the company of friends and loved ones to participate in the unfolding of the future, would, when push came to shove, choose this latter alternative.

If some people would still choose death, that's a choice that is of course to be regretted, but nevertheless this choice must be respected. The transhumanist position on the ethics of death is crystal clear: death should be voluntary. This means that everybody should be free to extend their lives and to arrange for cryonic suspension of their deanimated bodies. It also means that voluntary euthanasia, under conditions of informed consent, is a basic human right.

It may turn out to be impossible to live forever, strictly speaking, even for those who are lucky enough to survive to such a time when technology has been perfected, and even under ideal conditions. The amount of matter and energy that our civilization can lay its hands on before they recede forever beyond our reach (due to the universe's expansion) is finite in the current most favored cosmological models. The heat death of the universe is thus a matter of some personal concern to optimistic transhumanists.

It is too early to tell whether our days are necessarily numbered. Cosmology and fundamental physics are still incomplete and in theoretical flux; theoretical possibilities for infinite information processing (which might enable an upload to live an infinite life) seem to open and close every few years. We have to live with this uncertainty, along with the much greater uncertainty about whether any of us will manage to avoid dying prematurely, before technology has become mature.

ACKNOWLEDGEMENTS AND DOCUMENT HISTORY

The Transhumanist FAQ was conceived as an attempt to develop a broadly based consensus articulation of the basics of responsible transhumanism. The aim was a text that could serve both as a guide to those new to the field and as a reference work for more seasoned participants.

Close to a hundred people have contributed in some way in the making of this document. The current version is a thorough revision of the version of 1999. Considerable new material has been added and many old sections have been substantially reworked. In preparing version 2.0, the following people have been especially helpful:

Eliezer Yudkowsky, who provided editorial assistance with the entire document and many comments on particular issues of substance; Dale Carrico, who proofread the first half of the text and eliminated a whole army of typos and infelicities; and Michael LaTorra, who did the same for the second half; and "Reason," who then went over the whole document again, as

did Frank Forman and Sarah Banks Forman. Useful comments of either substance or form have also been contributed by (in alphabetical order): Michael Anissimov, Samantha Atkins, Milan Cirkovic, José Luis Cordeiro, George Dvorsky, James Hughes, G.E. Jordan, Vasso Kambourelli, Michael LaTorra, Eugen Leitl, Juan Meridalva, Harvey Newstrom, Emlyn O'Reagan, Christine Peterson, Giulio Prisco, Reason, Rafal Smigrodzki, Simon Smith, Mike Treder, and Mark Walker. Many others have over the years offered questions or reflections that have in some way helped shape this document, and even though it is not possible to name you all, your contributions are warmly appreciated. All remaining errors are my own fault.

Since this new version builds on its predecessor, the contributors to the latter are also indirect contributors to the present document. The people who assisted with the first version included especially Anders Sandberg (several of the original sections were based heavily on his material and the section on transhumanism's historical precedents still very distinctly bears his mark), Kathryn Aegis, and Natasha Vita-More (who also both provided substantial chunks of text). The presentation in the cryonics section was, and still is, directly inspired by an article by Ralph Merkle. Several people contributed to the definition of transhumanism, in particular Kathryn Aegis and Max More. Greg Burch, David Pearce, Kathryn Aegis, and Anders Sandberg kindly offered extensive editorial comments. Ideas, criticisms, questions, phrases, and sentences to the original version were contributed by (in alphabetical order; some contributors are known only by their email addresses): Kathryn Aegis, Alex (intech@intsar.com), Brent Allsop, Brian Atkins, Scott Badger, Doug Bailey, Harmony Baldwin, Damien Broderick, Greg Burch, David Cary, John K. Clark, Dan Clemmensen, Damon Davis, Jeff Dee, Jean-Michel Delhotel, Dylan Evans, EvMick@aol.com, Daniel Fabulich, Frank Forman, Robin Hanson, Andrew Hennessey, Tony Hollick, Joe Jenkins, William John, Michelle Jones, Arjen Kamphius, Henri Kluytmans, Eugene Leitl, Michael Lorrey, mark@unicorn.com, Peter C. McCluskey, Erik Moeller, J. R. Molloy, Max More, Bryan Moss, Harvey Newstrom, Michael Nielsen, John S. Novak III, Dalibor van den Otter, David Pearce, pilgrim@cyberdude.com, Thom Quinn, Anders Sandberg, Wesley R. Schwein, Shakehip@aol.com, Allen Smith, Geoff Smith, Randy Smith, Dennis Stevens, Derek Strong, Remi Sussan, Natasha Vita-More, Michael Wiik, Eliezer Yudkowsky, and zebo@pro-ns.net

I would like to thank you all for helping creating this FAQ and for making transhumanism possible.

NICK BOSTROM
Oxford, October 2003

ON the bank at the end
Of what was there before us
Gazing over to the other side
On what we can become
Veiled in the mist of naïve speculation
We are busy here preparing
Rafts to carry us across
Before the light goes out leaving us
In the eternal night of could-have-beens

NOTE

* *Editor's note*: This piece is an extract from a longer article. For the complete article and more of Bostrom's work on transhumanism, see www.nickbostrom.com