

Transhumanist Ethics¹

Nick Bostrom

ABSTRACT

Transhumanism is a grassroots movement that advocates the voluntary use of technology to enhance human capacities and extend our health-span. Our human quest to improve ourselves has a long history; the origins of transhumanism can be traced to renaissance humanism and back to antiquity. Only recently, however, have we begun to understand the technological prerequisites for enhancing human nature and to think systematically think about the risks and opportunities. Transhumanism is emerging as the most promising alternative to conservative ethical systems that see human nature as something that cannot or should not be changed, an attitude increasingly in tension with technological possibilities and people's legitimate desire to benefit from them. The first part of this essay examines the axiological foundations of transhumanist ethics. The second part looks at human germ-line genetic engineering from a transhumanist perspective and argues that it helps us to formulate an ethically responsible stance that addresses concerns about inequalities and commodification of human life. Rather than remaining stuck in a simplistic "for or against" dilemma, we may recognize that there are many different sorts of possible modifications and consider which ones should be encouraged, which should be discouraged, and what other social policies may be required to ensure a fair and open future.

¹ For their comments I'm grateful to Mike Treder, Eugene Leitl, Jeffrey Soreff, Robert A. Freitas Jr., James Hughes, and Mark Walker.

1. What is transhumanism?

Transhumanism is a loosely defined movement that has developed gradually over the past two decades. It takes an interdisciplinary approach to understanding and evaluating the opportunities for enhancing the human condition and the human organism opened up by the advancement of technology (Bostrom et al. 1999; WTA 2002). Attention is given to both present technologies, like genetic engineering and information technology, and anticipated future ones, such as molecular nanotechnology (Drexler 1992) and artificial intelligence (Bostrom 1998; Kurzweil 1999; Moravec 1999).

The enhancement options being discussed include radical extension of human health-span, eradication of disease (Freitas Jr. 1999), elimination of unnecessary suffering (Pearce 2001), and augmentation of human intellectual, physical and emotional capacities. Other transhumanist themes include e.g. space colonization and the possibility of creating superintelligent machines, along with other possible developments that could profoundly alter the human condition. The ambit is not limited to gadgets and medicine, but encompasses also economical and social and institutional designs, cultural development, psychological skills and techniques – in general, any workable method for overcoming fundamental human limitations.

Transhumanists view human nature as a work-in-progress, a half-baked beginning that we can learn to remold in desirable ways. Current humanity need not be the endpoint of evolution. Transhumanists hope that by responsible use of science, technology and other rational means we shall eventually manage to become *posthumans*, beings with vastly greater capacities than present humans.

Some transhumanists take active steps to increase the probability that they personally will survive long enough to become posthumans, for example by choosing a healthy lifestyle or by making provisions for having themselves cryonically suspended in case of deanimation (i.e. after becoming “dead” by the current legal definition) (Ettinger 1964; Hughes 2001). In contrast to many other ethical outlooks, which in practice often have a reactive attitude to new technologies (viewing developments like human cloning mainly as moral threats), the transhumanist philosophy is guided by an evolving vision to take a more proactive approach to technology policy.

This vision, in broad strokes, is to create the opportunity (for those who so wish) to live much longer and healthier lives, to enhance our memory and other intellectual faculties, to refine our emotional experiences and increase our subjective sense of well-being, and generally to achieve a greater degree of control over our own lives. This affirmation of human potential is offered as an alternative to customary injunctions against “playing God”, “messing with Nature”, “tampering with our human essence” and displaying punishable “hubris”.

It is important to emphasize that transhumanism is not techno-optimism. While future technological capabilities carry immense potential for beneficial deployments, they also could be misused to cause enormous harm – ranging all the way to the extreme possibility of intelligent life becoming extinct. Other potential negative outcomes include widening social inequalities or a gradual erosion of the hard-to-quantify assets that we care deeply about but tend to neglect in our daily struggle for material gain (e.g.

meaningful human relationships and eco-diversity). Such risks must be taken very seriously, and thoughtful transhumanists fully acknowledge this.²

Transhumanism has roots in secular humanist thinking³, yet is more radical in that it advocates not only traditional means (such as education and cultural refinement) of improving human nature but also direct application of medicine and technology to overcome some of our basic biological limits.

Not being the creation of any one person, transhumanism does not have a single canonical formulation that defines once and for all what transhumanism is. Rather, it is a grassroots movement that has emerged gradually. Much of the discourse has taken place on mailing lists and Internet discussion forums in a rather disorganized fashion (although there are now also regular conferences and one academic journal specifically devoted to transhumanism). This paper could therefore be construed as a sympathetic interpretation, or a proposal, rather than as a straightforward exegesis. One particular difficulty is to disentangle the value component in this worldview from factual assumptions about the likely consequences of various courses of action. Yet in this paper, I shall try to mainly focus on the normative aspects.

Before moving on to discussing the transhumanist perspective on human germline genetic engineering specifically, let me first describe what in my view constitutes the core of transhumanist axiology.

² Some of the most careful analyses to date of the threats to human survival and long-term flourishing have been done by transhumanists. See e.g. (Drexler 1985, 1988; Hanson 1998; Freitas Jr. 2000; McCarthy 2000; Bostrom 2002).

³ Needless to say, the basic idea of expanding human capacities can be traced back much further, to antiquity and indeed to the very earliest preserved records of human culture – witness the quest for the herb of immortality in the Sumerian *Epic of Gilgamesh*.

2. Human limitations

The range of thoughts, feelings, experiences, and activities accessible to human organisms presumably constitute only a tiny part of what is possible. There is no reason to think that the human mode of being is any more free of limitations imposed by our biological nature than are those of other animals. In much the same way as Chimpanzees lack the cognitive wherewithal to understand what it is like to be human – the ambitions we humans have, our philosophies, the complexities of human society, or the subtleties of our relationships with one another, so we humans may lack the capacity to form a realistic intuitive understanding of what it would be like to be a radically enhanced human (a “posthuman”) and of the thoughts, concerns, aspirations, and social relations that such humans may have.

Our own current mode of being, therefore, spans but a minute subspace of what is possible or permitted by the physical constraints of the universe (see Figure 1). It is not farfetched to suppose that there are parts of this larger space that represent extremely valuable ways of living, relating, feeling, and thinking. If only we could reach them!

The Space of Possible Modes of Being




Figure 1. “We aint seen nothin’ yet!” (not drawn to scale). The term “transhuman” denotes transitional beings, or moderately enhanced humans, whose capacities would be somewhere between those of unaugmented humans and full-blown posthumans. (A transhumanist, by contrast, is simply somebody who accepts transhumanism.)

The limitations of the human mode of being are so pervasive and familiar to us that we often fail to notice or reflect upon them. It takes an almost childlike naiveté to question them.⁴ Let us therefore pause to consider some of the more basic ones.

Lifespan. Because of the precarious conditions in which our Pleistocene ancestors lived, the human lifespan has evolved to be a paltry seven or eight decades. This is, from many perspectives, a rather short period of time. Even tortoises do better than that.

We don’t have to use geological or cosmological comparisons to highlight the meagerness of our allotted time budgets. To get a sense that we might be missing out on something important by our tendency to die early, we only have to bring to mind some of the worthwhile things that we could have done or attempted to do if we had had more time. For gardeners, educators, scholars, artists, city planners, and those who simply

⁴ Good science fiction can be helpful in this regard, as a crutch to our imagination.

relish observing and performing in the cultural or political variety shows of life, three scores and ten is often insufficient for seeing even one major project through to completion, let alone for undertaking many such projects in sequence.

Human character development is also cut short by aging and death. Imagine what might have become of a Beethoven or a Goethe if they had still been with us today. Maybe they would have developed into rigid old grumps interested exclusively in conversing about the achievements of their youth. But maybe, if they had continued to enjoy health and youthful vitality, they would have continued to grow as men and artists, to reach levels of maturity that we can barely imagine. We certainly cannot rule that out based on what we know today. Therefore, there is at least a serious *possibility* of there being something very precious outside the human sphere. This constitutes a reason to pursue the means that may let us go there and find out.

Intellectual capacity. We have all had occasions when we wished we were a little smarter. The three-pound, cheese-like thinking machine that we carry around in our skulls can do some neat tricks, but it also has significant shortcomings. Some of these – such as forgetting to buy milk or failing to attain native fluency in languages you learn as an adult – are obvious and require no elaboration. These shortcomings are inconveniences but hardly fundamental barriers to human development.

Yet it is possible, and I would argue plausible, that there is a more profound sense in which our intellectual apparatuses limit our modes of mentation. I mentioned the Chimpanzee analogy earlier: just as is the case for the great apes, our own cognitive makeup may foreclose whole strata of understanding and mental activity. The point here

is not about any “in principle” impossibility; we need not suppose that posthumans would not be Turing computable or that they would have concepts that could not be expressed by any finite sentences in our language, or anything of that sort. The impossibility that I am referring to is more like the impossibility for us current humans to visualize a 400-dimensional hypersphere or to have read, with perfect recollection and understanding, every book ever written in any human language. These things are impossible for us because, simply put, we lack the brainpower. Similarly, we may lack the practical ability to intuitively understand what being a posthuman would be like or to grasp playing field of posthuman concerns.

Further, our current intellects may cap our ability to discover philosophical and scientific truths. It seems at least conceivable that the failure of philosophical research to arrive at solid, generally accepted answers to many of the traditional big philosophical questions could be due to the fact that we aren’t smart enough to be successful in this kind of enquiry. Our intellectual limitations may be confining us in a Platonic cave, where we the best we can do is theorize about shadows (representations sufficiently oversimplified and dumbed-down to fit inside a human brain). A posthuman might see relativity theory the way Einstein might view a television sound-bite version of his own theory.

Bodily functionality. We enhance our natural immune systems by getting vaccinations, and we can imagine further enhancements to our bodies that would protect us from disease or help shape our bodies according to our desires (e.g. by letting us control our bodies’ metabolic rate). Such enhancements could improve the quality of our lives.

A more radical kind of upgrade might be possible if we suppose a computational view of the mind. It may then be possible to upload a human mind to a computer, by replicating *in silico* the detailed computational processes that would normally take place in a particular human brain (Drexler 1985). Being an upload would have many potential advantages, such as the ability to make back-up copies of oneself (favorably impacting one's life-expectancy) and the ability to transmit oneself as information at the speed of light. Uploads might live either in virtual reality or directly in physical reality by controlling a robot proxy. (Since the theoretical possibility of uploads is more controversial than the other types of enhancement that we are discussing here, I should stress that the argument in no way depends on it.)

Sensory modalities, special faculties and sensibilities. The current human sensory modalities are not the only possible ones, and they are certainly not as highly developed as they could be. Some animals have sonar, magnetic orientation, or sensors for electricity and vibration; many have a much keener sense of smell, sharper eyesight, etc. The range of possible sensory modalities is not limited to those we find in the animal kingdom. There is no fundamental block to adding say a capacity to see infrared radiation or to perceive radio signals and perhaps to add some kind of telepathic sense by augmenting our brains with suitably interfaced radio transmitters.

Humans also enjoy a variety of special faculties, such as appreciation of music and a sense of humor, and sensibilities such as the capacity to be sexually aroused by erotic stimuli. Again, there is no reason to think that what we have exhausts the range of

the possible, and it is certainly possible to imagine higher levels of sensitivity and responsiveness.

Mood, energy, and self-control. Despite our best efforts, we often fail to feel as happy as we would like. Our chronic levels of subjective well-being seem to be largely genetically determined. Life-events have little long-term impact; the crests and troughs of fortune push us up and bring us down, but there is little long-term effect on self-reported well-being. Lasting joy remains elusive except for those of us who are lucky enough to have been born with a temperament that plays in a major key.

In addition to being at the mercy of a genetically determined setpoint for our levels of well-being, we are limited in regard to energy, will-power, and ability to shape our own character in accordance with our ideals. Even such “simple” goals as losing weight or quitting smoking prove unattainable to many.

Some subset of these kinds of problems might be necessary rather than contingent upon our current nature. For example, we cannot both have the ability easily to break any habit and the ability to form stable, hard-to-break habits. (In this regard, the best one can hope for may be the ability to easily get rid of habits we didn’t deliberately choose for ourselves in the first place, and perhaps a more versatile habit-formation system that would let us choose with more precision when to acquire a habit and how much effort it should cost to break it.)

3. The core transhumanist value: exploring the posthuman realm

Given that our current human mode of being spans such a tiny subspace of what is possible, it is not farfetched to suppose that there are parts of this larger space that represent extremely valuable ways of living, feeling, thinking, and relating.

What may some of these be? We can conceive, in the abstract at least, of crackling sensual pleasures more blissful and thrilling than any in human history; aesthetic contemplation more rapturously sublime or more perfectly Apollonian; nonpareil levels of personal development and maturity allowing for the first time that precious inner core of each one of us to suffuse and fill out our whole personas; a vastly richer understanding of the human condition, derived from having savored life more fully and reflected more deeply; a keener intelligence and a quicker wit, grasping the whole of science better than any current expert understands her own specialty; philosophical thinking more profound and disillusioned; and love so passionate, ever-fresh, and secure that its reality surpasses our maddest moonstruck longings. We can also conceive of some of the secondary effects of such capacities – wonderful new art forms, truer science, more enlightened philosophy, and closer unions between lovers.

In addition to the values that we can conceive of, there may be others that our current limited perspective hides from our view. Our intuitions about values might be constrained by the narrowness of our experience and the limitations of our powers of imagination. We must leave room in our thinking for the possibility that as we develop greater capacities, we shall come to discover values that will strike us as being of far higher order than those we can realize as un-enhanced biological humans. Our vision of posthuman goods may be as myopic as a monkey troop's fantasy of the boons of human

existence presumably is. What we see in our wildest fancies could be the equivalent of a big bunch of bananas.

The conjecture that there are greater values than we can currently fathom does not imply that values are not defined in terms of our current dispositions. Take, for example, a dispositional theory of value such as the one described by David Lewis (Lewis 1989). According to this theory, *X* is a value for you if and only if you would want to want *X* if you were perfectly acquainted with *X* and you were thinking and deliberating as clearly as possible about *X*. (*X* could be said to be a value *simpliciter* – i.e. not merely a value *for somebody* – if and only if everybody are disposed to want to want *X* under ideal conditions.) On this view, there may be values *X* that we don't currently want, and that we don't even currently want to want, because we may not be perfectly acquainted with *X* or because we are not ideal deliberators. Some values pertaining to certain forms of posthuman existence may well be of this sort; they may be values for us now (and be so in virtue of our current dispositions) that we may not be able to fully appreciate with our current (less than ideal) deliberative capacities and our lack of the receptive faculties required for full acquaintance with them.

This point is important because it shows that the transhumanist view that we ought to explore the realm of posthuman values *does not* entail that we should forego our current values. The posthuman values can *be* our current values, albeit ones that we haven't yet clearly comprehended. Transhumanism doesn't say that we should favor posthumans over humans, but rather that the right way of favoring humans is by enabling us to realize our ideals better and that some of our ideals may well be located outside the space of modes of being that are accessible with our current biological constitution.

When Leon Kass wrote that “No friend of humanity cheers for a post-human future” (Kass 2001), he made a false statement.

While we can overcome many of our biological limitations, it is possible that there are some limitations that are impossible for us to transcend, not only because of technological difficulties but also on metaphysical grounds. Depending on what our views are about what constitutes personal identity, it could be that certain modes of being, while possible, are not possible *for us*, because any being of that kind would be so different from us that they could not be us. Concerns of this kind are familiar from theological discussions of the afterlife. Christian eschatology holds that some souls will be allowed by God to go to heaven after their time as corporal creatures is over. Before being admitted to heaven, these souls would undergo a purification in which they would lose many of their previous bodily attributes. Skeptics may doubt that the resulting spirits would be sufficiently similar to our current minds for it to be possible for them to be parts of the same person. A similar predicament arises for transhumanism: if the mode of being of a posthuman is radically different from that of a human, then we may doubt that a posthuman could be the same person as a human, even if the posthuman originated from the human.

An investigation of the concept of personal identity is beyond the scope of this essay, but a couple of points can still be made at this juncture.

First, one can envision some very considerable enhancements that would not make it impossible for the post-transformation you to be the same person as the pre-transformation you. Surely, you could obtain quite a bit of increased life expectancy, intelligence, health, memory, and emotional sensitivity, without your original person

necessarily ceasing to exist in the process. A person's intellectual life can be transformed quite radically by getting an education, a person's life expectancy can be extended substantially (say by being unexpectedly cured from a lethal disease), and yet these developments are not viewed as spelling the end of the original person. In particular, it seems that modifications that *add* to one's person (especially if they are self-chosen?) can be very substantial, more so, perhaps, than modifications that detract (such as serious brain damage). If most of what you currently are – including your memories, the activities open to you, the range of feelings you can experience – is preserved, then adding extra capacities on top of that would not easily cause you to cease to exist.⁵

Second, preservation of personal identity, especially if this notion is given a narrow definition, is not everything. One can value other things than oneself. Or one might regard it as satisfactory if some parts or aspects of one's person survive and flourish, even if it entails giving up some parts of oneself such that one no longer counts as being the same person. Which parts of oneself that one would be willing to relinquish for other goods may not become clear until we are better acquainted with the options.

To sum up, transhumanism affirms the quest to develop further so that we can explore hitherto inaccessible realms of value. Technological enhancement of the human organism is a means to this end. There are limits to how much can be achieved by low-tech means such as education, philosophical contemplation, or moral self-scrutiny (the methods proposed by classical philosophers with perfectionist leanings, for example Plato, Aristotle, and Nietzsche), or by means of creating a fairer and better society (as social reformists such as Marx or Martin Luther King envisioned). This is not in any way

⁵ Although in some cases, if the new opportunities that are opened up to you make you exercise the old capabilities less often, there could perhaps be a contingent loss of personal identity.

to denigrate what we can do with the tools we have today. Yet we have good reason to suppose that there might well be extremely worthwhile and valuable modes of being that are foreclosed to us as long as we retain our current biological limitations. This is so even if we adopt a theory of value that defines value in terms of actual human preferences. It is highly likely, however that we will eventually develop technologies that will enable us to overcome some of our limitations. The prospect of there being great values outside of the human sphere constitutes us a strong reason, which we could call the “transhumanist imperative”, to seek to develop such technologies and use them to expand human capacities so that we can begin to explore the wider realm of modes of being.

4. The transhumanist perspective on germ-line genetic engineering

Having examined the normative basis of the transhumanist movement, we shall turn to consider how it applies to a concrete issue, human germ-line genetic engineering.

From the foregoing description it should be obvious that certain types of objections against germ-line modifications are not accorded much weight by a transhumanist interlocutor. For instance, objections that are based on (variously disguised versions of) the idea that there is something inherently wrong or morally suspect in using science to manipulate human nature are regarded by transhumanists as wrongheaded.

Moreover, transhumanists emphasize that particular concerns about negative aspects of genetic enhancements⁶, even when these concerns are legitimate, must be

⁶ The distinction between enhancement and treatment can be notoriously hard to draw (Parens 1998). The transhumanist position is at an advantage in this regard, because it denies any fundamental ethical importance to this distinction. This is not to say that attempting to draw such a distinction, even roughly and with some element of metaphysical arbitrariness, cannot be useful. For instance, when funds are

judged against the potentially enormous benefits that could come from genetic technology successfully employed.

For example, many commentators worry about the psychological effects of the use of germ-line engineering. The ability to select one's children's genes and create "designer babies" will, it is claimed, corrupt parents, who will come to view their children as mere products (Kass 2001). We will then begin to evaluate our offspring according to standards of quality control, and this will undermine the ethical ideal of unconditional acceptance of children, no matter what their abilities and traits. Are we really prepared, asks the objector, to sacrifice on the altar of consumerism even those deep values that are embodied in traditional relationships between child and parents? Is the quest for "perfection" worth this cultural and moral cost?

Faced with such criticism, how should a transhumanist respond? Certainly not by dismissing the concerns as irrelevant. Transhumanists no less than other people recognize that the depicted outcome would be bad. We don't want parents to love and respect their children less. We don't want social prejudice against people with disabilities to get worse. The psychological and cultural effects of "commodifying" human nature are potentially important.

The first thing to emphasize, however, is that these dystopian scenarios are speculations. There is, at least to my knowledge, no solid ground for believing that these alleged consequences would actually happen. What relevant evidence we have, for instance regarding the treatment of children who have been conceived through the use of

limited, public health insurance must define which interventions are covered, and one way of doing that might be by paying for treatment but not enhancement. As I argue below, however, there are cases where this simplistic rule would be inadequate; some enhancement-options may become available that it might be appropriate for the state to subsidize.

in vitro fertilization or embryo screening, suggests that the pessimistic prognosis is both alarmist and wrong. Parents will in fact love and respect their children even when artificial means and conscious choice play a part in procreation.

We can push this thought farther by offering a counter-speculation. Maybe germ-line enhancements will lead to *more* love and parental dedication. Maybe some mothers and fathers will actually find it *easier* to love a child who, thanks to enhancements, is bright, beautiful, and healthy. Maybe the practice of germ-line enhancement will lead to better treatment of people with disabilities, both because a general demystification of the genetic contributions to human traits could make it clearer that people with disabilities are not to *blame* for their disabilities (as some prejudiced people seem to intuitively believe) and also because a decreased incidence of some disabilities could lead to more assistance being available for the remaining affected people to enable them to live full, unrestricted lives through various technological and social supports.

Speculating about possible psychological or cultural effects of germ-line engineering can therefore cut both ways; good consequences no less than bad ones are possible. In the absence of sound arguments for the view that the negative consequences would predominate, such speculations are no reason against moving forward with this technology.

This does not imply that these hypothetical ruminations of moral philosophers and other concerned parties are of no use. On the contrary, they serve to make us aware of things that could go wrong so that we can be on the lookout for untoward developments. By being aware of the perils in advance, we will be in a better position to take countermeasures to stop them from coming about. For instance, if we think that some

people would fail to realize that a human clone would be a unique person deserving just as much respect and dignity as any other human being, then we can counter this threat by working harder on educating the public about why genetic determinism is false. The theoretical contributions of well-informed and reasonable critics of germ-line enhancement may therefore indirectly, perhaps with a touch of paradox, *add* to our justification for proceeding with germ-line engineering. For to the extent that the critics have done their job, they will have alerted us to many of the potential untoward consequences of germ-line engineering and thereby contributed to our ability to take precautions, thus improving the odds that the balance of effects will be positive.

There will in all likelihood be some negative consequences of human germ-line engineering that we cannot or will not forestall. Needless to say, the mere existence of negative effects is no reason not to proceed. Any major technology (as well as any major social or political reform) has *some* negative consequences, including some *unforeseen* negative consequences. And so does, for that matter, the choice to preserve status quo. It is only after a fair comparison of the risks with the likely positive consequences that any conclusion based on a cost-benefit analysis can be reached.

In the case of germ-line enhancements, the potential gains are *enormous*. Only rarely, however, are these potential gains elaborated on at any length, perhaps because they are too obvious to be of much theoretical interest. (By contrast, uncovering subtle and non-trivial ways in which manipulating our genome could undermine deep values is philosophically a lot more challenging.) But if we think about it, we recognize that the promise of genetic enhancements is anything but insignificant. Set aside, for a moment, any worries you might have about possible downsides. Surely, being free from severe

genetic diseases is a really good thing. And having a mind that can learn more quickly, or having a robust immune system, is wonderful too. A healthier, wittier, happier humanity may be able to reach new levels culturally as well as in terms of individual lives. To achieve a significant enhancement of human capacities would be to take a step on the transhuman journey to explore some of those modes of being that are not accessible to humans as currently constituted, and perchance to discover and realize important new values. On an even more basic level, there is a great potential for alleviating unnecessary human suffering. Every day that the introduction of effective human genetic enhancement is delayed is a day of lost individual and cultural potential, and a day of torment for millions of sufferers who are being ravaged by diseases that could have been prevented.

Seen in this light, proponents of a ban or a moratorium on human genetic modification must take on a heavy burden of proof indeed if they want the balance of reason to tilt in their favor. Transhumanists conclude that this challenge has not been met.

5. Should human reproduction be regulated?

Supposing, then, that we agree with the transhumanist view that we ought to go forward with genetic enhancement of humans in some manner, how should we go about it? This question leads to a host of interesting and constructive issues about how best to proceed.

One option is the *laissez-faire* approach of permitting almost everything, leaving all choices to parents, and denying any role of the state in setting or enforcing genetic policy. While I think this attitude is consistent with transhumanism, I don't think it is the best transhumanist approach. Yet it does have something to be said for it.

One thing that can be said for adopting a libertarian stance in regard to human reproduction is the sorry track record of socially planned attempts to improve the human gene pool. The list of historical examples of state intervention in this realm ranges from the genocidal horrors of the Nazi regime, to the incomparably milder but still disgraceful semi-coercive sterilization programs of “retards” favored by many well-meaning social progressives in the past century, to the controversial but perhaps understandable program of the current Chinese government to limit population growth. In each of these cases, state policies interfered with the reproductive choices of individuals. If parents had been left to make these choices for themselves, the worst transgressions of the eugenics movement would not have occurred. Bearing this hard-earned lesson from history in mind, we ought to think twice before supporting proposals that would have the state regulate what sort of children people are allowed to have and the methods they may use to conceive them.

On the other hand, there is also an argument for thinking that the libertarian approach is *less* appropriate in the realm of reproduction than it is in other domains. For in reproduction, the most important interests at stake are those of the child-to-be-born, who cannot, of course, give its advance consent or freely enter into any kind of contract. Moreover, as a society we currently approve of many measures that limit parents’ freedoms with respect to their children. We have laws against child abuse and against child neglect. We have obligatory schooling. In some cases, we can force needed medical treatment on a child, even against its parents’ wishes.

Nonetheless, there is an asymmetry between these types of social intervention and interventions aimed at genetic enhancements. For while there is a consensus that nobody

should be subjected to child abuse and that all children should have at least a basic education and should receive necessary medical care, it is unlikely that we will reach an agreement on proposals for genetic enhancements any time soon. Many parents will resist such proposals on principled grounds, such as religious convictions.

The best policy for the foreseeable future may therefore be not to legally require any genetic enhancements, except perhaps in extreme cases such as when a severe genetic disease, for which there is no alternative treatment, can be directly targeted through a safe and reliable genetic therapy (and even there, it is dubious that the social climate is ready for any mandatory interventions.)

The scope for ethics and public policy, however, extend far beyond the passing of laws requiring or banning specific interventions. Even if a given enhancement option is neither outlawed nor legally required, society may still seek to discourage or encourage its use in a variety of ways. Through subsidies and taxes, research-funding policies, genetic counseling practices and guidelines, laws regulating genetic information and genetic discrimination, provision of health care services, regulation of the insurance industry, patent law, education, and through the allocation of social approbation and disapproval, we as a society have many gentler means of influencing the direction in which particular technologies are applied.⁷ So we may ask, with regard to genetic enhancement technologies, which types of applications we ought to promote or discourage.

⁷ See also (Glover 1984; Stock 2002) for arguments supporting wide parental discretion in selecting which genetic enhancement options to use or not to use for their children. (Buchanan et al. 2000) argue that such freedom is compatible with a potentially large role for the state in genetic policy, e.g. in making sure that there is wide access to genetic therapies, in providing information and counseling, and in ensuring that genetic practices fit into an overall just social structure that is maximally inclusive of people with disabilities.

6. Which modifications should be promoted and which discouraged?

The concept of externalities is crucial in addressing this question. An externality is a cost or a benefit of an action that is not carried by the decision-maker. An example of a negative externality is that of a firm that lowers its production costs by polluting the environment – the firm enjoys most of the benefits while escaping the costs (in terms of environmental degradation) which may instead be paid by people living nearby.

Externalities can also be positive, such as when somebody is putting time and effort into creating a beautiful garden outside their house – the effects are enjoyed not exclusively by the gardener but spill over to passersby. As a rule of thumb, sound social policy and social norms seek to internalize many externalities so that the incentives of the producer more closely match the social value of production. We may levy a pollution tax on the firm and give our praise and approbation to the home gardener who beautifies the neighborhood.

Genetic enhancements aimed at the obtainment of *status goods* (goods that are desirable only in so far as they provide a competitive advantage) tend to have negative externalities. Take for example stature. There is evidence that being tall is statistically advantageous, at least for men in our society. Taller men earn more money, wield greater social influence, and are viewed as sexually more desirable. Parents wanting to give their child the best possible start in life may therefore rationally choose a genetic enhancement that may add an inch or two to the expected length of their offspring. Yet for society as a whole, there seems to be no advantage whatsoever in people being taller. If everybody

grew two inches, nobody would be better off than they were before. Money spent on a status good like length has no net effect on social welfare and is therefore, from society's point of view, wasted.

Health is a very different type of good. It has intrinsic benefits. If you become healthier, you are personally better off and others are not any worse off. There may even be a positive externality of enhancing one's own health: if you are less likely to contract a contagious disease, others benefit by being less likely to get infected by you. Being healthier, you may also contribute more to society and consume less of publicly funded healthcare.

If we were living in a simple world where people were perfectly rational self-interested economic agents and where social policies had no implementation costs or unintended effects, then the basic policy prescription regarding genetic enhancements would be relatively straightforward: internalize the externalities of genetic enhancements by taxing enhancements that have negative externalities and subsidizing enhancements that have positive externalities. Unfortunately (or perhaps fortunately), our world is not that simple. Crafting policies that work well in practice is *much* more complicated.

Even just determining the net size of the externalities of a particular genetic enhancement can be highly non-trivial. Take for example an enhancement that increases memory or intelligence. There is clearly an intrinsic value to this in as much as most of us would probably like to be a bit smarter even it didn't have the slightest effect on our standing in relation to others. It's simply nice to be able to understand better. But additionally, there would be important externalities, both positive and negative. On the negative side, others would suffer some disadvantage from your increased brainpower in

that their own competitive situation would be worsened. Being more intelligent, you may be more likely to attain high-status positions in society, positions that would otherwise have been enjoyed by a competitor. On the positive side, others may benefit from enjoying witty conversations with you, from trading with you, from marrying you, from the taxes you pay, from the novels you write, from the scientific discoveries you make, and in myriad other ways.

If in the case of intelligence enhancement, the positive externalities outweigh the negative ones, then a prima facie case exists not only for permitting genetic enhancements aimed at increasing intellectual ability, but for encouraging and maybe subsidizing them too. Whether such policies remain a good idea when all practicalities of implementation and political realities are taken into account, is a different question and beyond the scope of this paper.

But at least we can conclude that an enhancement that has both intrinsic benefits for the enhanced individual and significant net positive externalities for the rest of society – health and intelligence enhancements, for instance, seem to belong to this category – should not be banned; the question, rather, is how strongly they should be encouraged, and by what means. By contrast, enhancements that achieve only positional advantages (such as augmentation of stature or physical attractiveness) should not be encouraged by society, and one might even make a case for social policies aimed at reducing expenditure on such goods, e.g. a progressive tax on consumption (Frank 1999).

7. The issue of equality

An important kind of externality in germ-line enhancements is their effects on social equality. This has been a focus for many opponents of germ-line genetic engineering, who worry that it will widen the gap between haves and have-nots. Today, children from wealthy homes enjoy many environmental privileges, including access to better schools and social networks. Arguably, this constitutes an inequity against children from poor homes. We can imagine scenarios where these inequities grow much larger thanks to genetic interventions that only the rich can afford, adding to the environmental advantages already benefiting privileged children. Pushing such a trend to its extreme, one can even speculate about the members of the privileged stratum of society eventually enhancing themselves and their offspring to a point where the human species, for many practical purposes, splits into two or more species that have little in common except a shared evolutionary history (Silver 1997). The genetically privileged may be ageless, healthy, super-geniuses of flawless physical beauty, who are graced with a sparkling wit and a disarmingly self-deprecating sense of humor, radiating warmth, empathy and relaxed confidence. The non-privileged would remain at today's level but perhaps deprived of some their self-respect and suffering occasional bouts of envy. The mobility between the lower and the upper classes might be reduced practically to zero so that a kid born to poor parents, lacking genetic enhancements, would have no chance whatsoever of successfully competing against the super-kids of the rich. Even if we stipulate that no discrimination or exploitation of the lower class occurs in this hypothetical scenario (let's say the members of the elite get all the best jobs because they truly *are* the best candidates, and that they always act with integrity, obey the law, and pay their taxes,

etc.), there is still something disturbing about the prospect of a society with such extreme inequalities.

How should a transhumanist respond to this worry? In my view, the best response contains many parts. The first is to acknowledge that widening inequalities is a legitimate cause for concern.

It is true, of course, that we have vast inequalities in the world today. While we regard some of these as unfair, we also accept a wide range of inequalities because we think that they are deserved, or because they have social benefits (such as economic productivity), or because they are an unavoidable concomitant to free individuals making their own (sometimes foolish) choices about how to live their lives. Some of these justifications can also be used to exonerate some inequalities that could result from germline engineering.⁸

Moreover, although it's a simple and boring point, it must be stressed that a technology leading to an increase in unjust inequalities is not a sufficient reason for discouraging the development and use of that technology; we also have to consider its benefits. The benefits include not only positive externalities but also intrinsic values that reside in, e.g., the enjoyment of good health, a soaring mind, and emotional well-being. As argued above, the benefits of such enhancements could be enormous.

Further, society need not be a passive bystander. We can move to counteract some of the inequality-increasing tendencies of enhancement technology with social policies. One way of doing so would be by widening access to the technology by subsidizing it or providing it for free to children of impecunious parents. In cases where the enhancement

⁸ Although probably not the justification that appeals to desert – it would be hard to argue that an enhanced child *deserved* its extra talents more than other kids did. (Hard, but not impossible. Some defend the view that we deserve our talents and the competitive advantages they give us, see e.g. (Nozick 1974).)

has considerable positive externalities, such a policy may actually benefit everybody, not just the recipients of the subsidy. In other cases, one could attempt to support the policy on the basis of social justice and solidarity.

Even if all genetic enhancements were made available to everybody for free, however, this might still not completely allay the concern about inequity. Some parents may choose not to give their children any enhancements, and these children would then suffer diminished opportunities through no fault of their own. How should society respond to such a dilemma?

Not, it seems, by banning genetic enhancements. For the problem stems from the fact that we give parents discretion over how they reproduce – that is what would prevent a society that is prepared to make genetic enhancements available for free from providing them to all children. It would then be peculiar to argue that society should respond to the problem by limiting the reproductive freedom of those parents who wish to use genetic enhancements. If we are willing to limit reproductive freedom and give society a mandate to regulate in this domain for the sake of reducing inequities, then society may as well use its mandate by making some enhancements obligatory for all children rather than by banning enhancements. For by having everybody enhancing to the same degree, we not only prevent the increase in inequalities but also reap the intrinsic benefits and the positive externalities that would come from the universal application of enhancement technology. If, on the other hand, reproductive freedom is regarded as too precious to be curtailed, then neither requiring nor banning the use of reproductive enhancement technology is an option. In that case, we would either have to tolerate inequities as a price worth paying for reproductive freedom or seek to remedy the inequities in ways that

don't infringe on reproductive freedom, such as by giving financial or social support to children who have been born without enhancements.

Finally, it should be emphasized that the inequality-objection against germ-line engineering is based on the hypothesis that germ-line engineering would in fact increase inequalities if left unregulated and no countermeasures were taken. This hypothesis might be false. It may well turn out to be technologically much easier to cure gross genetic defects than to enhance an already healthy genetic constitution. (We currently know much more about many specific inheritable diseases, some of which are due to single gene defects, than we do about the genetic basis of talents and desirable qualities such as intelligence and longevity, which depend on complex constellations of large numbers of genes.) If this turns out to be the case, then the trajectory of human genetic enhancement may be one in which the first thing to happen is that the lot of the genetically worst-off class is radically improved, through the elimination of diseases such as Tay Sachs, Lesch-Nyhan, Downs Syndrome, early-onset Alzheimer's disease. As the lowest-hanging fruits get picked, the focus would gradually shift to curing more marginal disadvantages. The effect would be a major *leveling* of inequalities, not primarily in the monetary sense but with respect to the even more fundamental parameters of basic opportunities and quality of life.

8. The changing technological and moral context

This last point illustrates the need for critically examining our assumptions about what the technological progress may have in store. Many assumptions, upon reflection, turn

out to be naïve and simplistic. Bad policy could easily result from basing our decisions on ill-conceived judgments about the likelihood of future scenarios.

While this paper has focused mostly on the normative dimension, transhumanist thinking is at least as concerned with trying to get a better grasp of the positive issues, such as the specifics about which types of technological capabilities will become available at which stages, how various technologies may interact and how they will impact society, and the likely consequences of different possible courses of action.

Transhumanist thinking emphasizes that it is an error to consider particular technological changes in isolation from other technological and social changes that will occur over the coming decades. To decide whether human germ-line engineering is a good thing, we should not imagine a world like the one we are currently living in only with germ-line engineering added. For by the time germ-line engineering has been developed, widely adopted, the enhanced children have matured, and they have had time to have a significant influence on the world (it will take more than half a century for this to happen), much else will have changed in the meantime. It would be foolish to ignore the changing technological context.

Take the objection that germ-line genetic engineering is uniquely hazardous because the changes it would bring are irreversible; they will affect all generations to come. This objection rests on a misunderstanding. Germ-line interventions could reverse the effects of earlier germ-line interventions. Moreover, technological and medical progress will hardly end with the invention of germ-line genetic engineering. Eventually, it is likely that medicine will progress to the point where germ-line modifications can be

reversed even in adults, using e.g. somatic gene therapy.⁹ Considering the lag time intrinsic to human germ-line enhancements (because of our slow maturation), it is likely that by the time it could have any large-scale impact on society, effective means of reversing germ-line modifications will already have been developed.

It is possible that all coming generations would choose to retain the modifications so that while they would not be irreversible, they might nonetheless never be reversed. This possibility, however, is not unique for germ-line interventions. The abolition of slavery and the introduction of general suffrage might never be reversed; indeed, we hope they won't be. Yet this was no reason for resisting these reforms. Likewise, the potential for everlasting consequences is no reason to oppose genetic modifications. If immunity against horrible diseases and enhancements that expand the opportunities for human growth are passed on to subsequent generations forever, so be it.

Regarding accelerating technological change, my own opinion is that human germ-line genetic enhancements will likely be obsolete before they have been implemented on a wide scale. Nanomedicine, therapies and enhancements based on advanced molecular nanotechnology (Freitas Jr. 1999), will one day make plain old genetic engineering seem crude and ineffective by comparison. Quite possibly, some even more capable technology, such as superintelligence and uploading, will likewise make nanomedicine obsolete before *it* gets off the ground. (Still, many of the basic ethical issues involving the creation of new persons will remain whatever the specifics of

⁹ Some ways of doing germ-line engineering would include reversibility in the original design, e.g. by making the expression of an inserted gene dependent on continued stimulation with a drug that the patient could stop taking at any point; but this doesn't help, of course, in the case of modifications that shape early development.

implementation (Glover 1984; Parfit 1984); our discussion of the ethics of germ-line engineering pertains to a wider set of potential technological options.)

It is not only our technological context that will change; our moral context will change too. By this I mean that many of our current moral norms and ethical intuitions are (often implicitly) based on factual assumptions about human nature and about the likely consequences of various attitudes and ways of behaving.¹⁰ When human nature changes and the circumstances of our lives undergo profound transformation, as could happen over the coming decades owing to accelerating technological progress, then the factual underpinnings of many of our current value judgments will need to be revisited.

One can hold up a picture of a future and say, “Doesn’t this look frightening! Would you really want to live in such an alien world?” When evaluating the long-term consequences of some proposed policy option, however, one must remember that it is not we-as-we-are-now that would inhabit the hypothetical world; rather, the concerned parties are either our descendants or ourselves-at-a-much-later-stage-of-development. Our attitudes, lifestyles, and ways of relating to each other may have adjusted in the meantime. When the anticipated consequences eventually occur, they may fit naturally into the kind of lives that are being lived at that time. Instead of placing a possible future world next to our current world and asking, “Which would you prefer to live in?”, it may be better to envision the choice as being between different routes we could take, different journeys that we could get to experience. To transhumanists, the journey of unlimited human growth and expansion, of exploration of the transhuman and eventually the

¹⁰ In this I agree with Francis Fukuyama’s view in (Fukuyama 2002). But while Fukuyama thinks that the appropriate response is to limit the development of technology so as to protect our moral system from having to deal with upsetting new possibilities, I believe that it is better to improve and update our ethics so that it can provide guidance under the new circumstances.

posthuman realm, appears infinitely more attractive than remaining at home forever in our biologically defined cottage.

References

Bostrom, N. (1998). "How Long Before Superintelligence?" *International Journal of Futures Studies* 2. <http://www.nickbostrom.com/superintelligence.html>

Bostrom, N. (2002). "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards." *Journal of Evolution and Technology* 9. <http://www.nickbostrom.com/existential/risks.html>

Bostrom, N., et al. (1999). *The Transhumanist FAQ*. <http://www.transhumanist.org>

Buchanan, A., et al. (2000). *From Chance to Choice: Genetics & Justice*. Cambridge, Cambridge University Press.

Drexler, K. E. (1985). *Engines of Creation: The Coming Era of Nanotechnology*. London, Forth Estate.

Drexler, K. E. (1988). "A Dialog on Dangers." *Foresight Background 2, Rev. 1*. <http://www.foresight.org/Updates/Background3.html>

Drexler, K. E. (1992). *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. New York, John Wiley & Sons, Inc.

Ettinger, R. (1964). *The Prospect of Immortality*. New York, Doubleday.

Frank, R. (1999). *Luxury Fever: Why Money Fails to Satisfy in an Era of Excess*. New York, Free Press.

Freitas Jr., R. A. (1999). *Nanomedicine, Volume 1: Basic Capabilities*. Georgetown, TX, Landes Bioscience.

Freitas Jr., R. A. (2000). "Some Limits to Global Ecophagy by Biovorous Nanoreplicators, with Public Policy Recommendations." *Zyvex preprint April 2000*.

Fukuyama, F. (2002). *Our Posthuman Future: Consequences of the Biotechnology Revolution*. New York, Farrar, Strauss and Giroux.

Glover, J. (1984). *What Sort of People Should There Be?* New York, Penguin.

- Hanson, R. (1998). "The Great Filter - Are We Almost Past It?" *Working paper*.
- Hughes, J. (2001). "The Future of Death: Cryonics and the Telos of Liberal Individualism." *Journal of Evolution and Technology* **6**.
- Kass, L. (2001). "Preventing a Brave New World." *The New Republic* **05.21.01**.
- Kurzweil, R. (1999). *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York, Viking.
- Lewis, D. (1989). "Dispositional Theories of Value." *Proceedings of the Aristotelian Society Supp.* **63**: 113-37.
- McCarthy, T. (2000). "Molecular Nanotechnology and the World System." <http://www.mccarthy.cx/WorldSystem/intro.htm>
- Moravec, H. (1999). *Robot: Mere Machine to Transcendent Mind*. New York, Oxford University Press.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. New York, Basic Books.
- Parens, E., Ed. (1998). *Enhancing Human Traits: Ethical and Social Implications*. Washington, D. C, Georgetown University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford, Clarendon Press.
- Pearce, D. (2001). *The Hedonistic Imperative*. <http://www.hedweb.com/hedab.htm>
- Silver, L. (1997). *Remaking Eden: How Genetic Engineering and Cloning will Transform the American Family*. New York, Avon Books.
- Stock, G. (2002). *Redesigning Humans: Our Inevitable Genetic Future*. New York, Houghton Mifflin.
- WTA (2002). *The Transhumanist Declaration*, World Transhumanist Association. <http://www.transhumanism.com/declaration.htm>