Philos. Technol. (2012) 25:439–458
DOI 10.1007/s13347-012-0091-6

INTRODUCTION

Evolution, Genetic Engineering, and Human Enhancement

Russell Powell · Guy Kahane · Julian Savulescu

Received: 23 September 2012 / Accepted: 24 October 2012 / Published online: 9 November 2012

Abstract There are many ways that biological theory can inform ethical discussions of genetic engineering and biomedical enhancement. In this essay, we highlight some of these potential contributions, and along the way provide a synthetic overview of the papers that comprise this special issue. We begin by comparing and contrasting genetic engineering with programs of selective breeding that led to the domestication of plants and animals, and we consider how genetic engineering differs from other contemporary biotechnologies such as embryo selection. We go on to consider the implications of genetic engineering for human nature, human evolution, and persistence of the human species. Finally, we question whether genetic interventions warrant the extraordinary ethical scrutiny they are often given, and we show how the misleading "genetic blueprint" metaphor has imposed a faulty structure on the enhancement debate. We conclude by considering the nature of biological development and the sobering limits it places on what genetic engineering can reasonably hope to achieve

Keywords Enhancement · Evolution · Genetic engineering · Human nature · Selection

There is growing evidence, including the successful genetic modification of plants and animals, that humans will eventually be able to alter many of their biological and psychological characteristics by directly modifying their genes. Traits that could be

R. Powell (✉)
Boston University, 745 Commonwealth Avenue, Boston, MA 02215, USA
e-mail: powell@bu.edu

G. Kahane · J. Savulescu
Oxford Uehiro Centre for Practical Ethics, Oxford Centre for Neuroethics,
Suite 8, Littlegate House, St Ebbes Street, Oxford OX1 1PT, UK

G. Kahane
e-mail: guy.kahane@philosophy.ox.ac.uk

J. Savulescu
e-mail: julian.savulescu@philosophy.ox.ac.uk

targeted for genetic intervention range from susceptibilities to disease to complex cognitive abilities and psychological dispositions, including those implicated in our identities, even in human nature itself. Anticipation of the coming biomedical enhancement revolution has given rise to a sizable and diverse body of literature examining the ethical dimensions of emerging biotechnologies and their implications for human flourishing and wellbeing (for recent samplings, see Savulescu and Bostrom 2009; Savulescu et al. 2011).

Although ethical discussions of biomedical enhancement have forged links with contemporary problems in the philosophy of mind, normative ethics, and political philosophy, they have so far rarely benefited from a substantive engagement with the philosophy of biological science. It is this gap that the present special issue is designed to fill. Broadly speaking, its aim is to consider how biological theory can advance ethical debates and policy discussions surrounding genetic engineering and human enhancement. This involves more than simply ensuring *consistency* with current empirical work in the life sciences—it means appreciating the relevance of conceptual and methodological problems in biology for ethical debates that arise in connection with the new biosciences.

Although there are many forms of biomedical enhancement, in this special issue we emphasize genetic engineering for substantive reasons beyond mere manageability. Few other medical interventions trigger such powerful moral intuitions, and are treated with such ethical and regulatory caution, as modifications of the human genome. Moreover, and relatedly, genes are thought to play a special causal role in biological evolution and the development of organisms from embryo to adult (hereinafter "development"). Whether the actual roles of genes in evolution and development justify the exceptional attitudes taken toward genetic engineering, and whether they provide reasons for restraining the development and use of certain biotechnologies, are questions to be explored throughout this issue and in future work.

Moral philosophical positions on genetic engineering and other emerging reproductive biotechnologies, such as embryo selection and cloning, often repose on particular metaphysical presuppositions about the causal-historical structure of the living world. In some cases, these biological "framing assumptions" are blatantly at odds with the theoretical and empirical foundations of the life sciences. This is true, for example, for ethical theories that attempt to derive normative principles from essentialistic conceptions of species that are inconsistent with the "population-thinking" that forms the core of the modern evolutionary biological worldview. In other cases, as we shall see, the tension between certain bioethical positions vis-à-vis genetic enhancement and contemporary biological theory is a bit subtler. Although crude forms of genetic determinism and reductionism are widely rebuked in the literature on enhancement, more sophisticated incarnations of these theses continue to structure bioethical discussions of genetic modification technologies.

There are many ways that biological theory can inform or otherwise contribute to ethical discussions of genetic engineering and biomedical enhancement. In this introductory essay, we highlight some of these potential contributions, and along the way provide a synthetic overview of the papers that comprise this special issue. We begin by comparing and contrasting genetic engineering with the ancient programs of selective breeding that led to the domestication of plants and animals, and

we consider how genetic engineering differs from other contemporary biotechnologies such as embryo selection. We go on to consider the implications of genetic engineering, and biomedical enhancement more broadly, for human nature, human evolution, and persistence of the human species. Next, we explore evolutionary perspectives on transgenic and synthetic organisms, and whether biomedical science will be able to improve on natural biological design. Finally, we query whether genetic interventions warrant the special ethical scrutiny they are often given, and we show how the related and misleading "genetic blueprint" metaphor has imposed a faulty structure on the enhancement debate. We conclude by considering the nature of biological development and the limits it places on what genetic engineering can reasonably hope to achieve.

1 From Selective Breeding to Genetic Engineering

In a sense, the deliberate genetic modification of organisms is not new. Humans have been intentionally transforming the heritable underpinnings of plant and animal phenotypes ever since the advent of selective breeding programs in the Agricultural Revolution, which had a profound impact on human society and even triggered bouts of human evolution. The domestication of plants and animals led to a nutrient surplus that enabled the transition from small hunter-gatherer societies to much larger sedentary populations with specialized divisions of labor, which in turn allowed for the development of sophisticated material and social technologies and ultimately the modern world. These novel modes of subsistence also generated new selection pressures on human populations, which evolved metabolic capacities to make use of agricultural food resources (e.g., lactose digestion) and developed a range of gene-based immunities to agriculture-related pathogens. The ultimate social, ecological, and environmental ramifications of the Neolithic Revolution were so extraordinary that many geologists believe they warrant the recognition of a new geological epoch: the aptly named *Anthropocene*. Perhaps the biotechnology revolution will lead to equally profound changes in the human condition.

As noted above, throughout the Holocene humans were engaged in a systematic attempt to bring about desired trait distributions in crops and animals through programs of selective breeding. The term "selective breeding" is preferable to "artificial selection" because the latter connotes something "unnatural" or "unreal", when in fact the practices surrounding domestication merely altered the fitness values of competing phenotypes in populations of the domesticated lineages, much as abiotic and non-human biotic interactions do "in the wild". The evolutionary interaction between humans and domesticated animals is simply an instance of the more general and quite "natural" phenomenon of biological mutualism (Sterrett 2002). Though many feats of domestication were remarkable (just compare maize with its teosinte ancestor), authors have long noted the limited ability of selective breeding to effect major transformations in domesticated lineages. These limitations were touted, in fact, by evolutionary skeptics in Darwin's time as evidence of the immutability of species and the causal impotence of natural selection (Richards 1997).

Modern genetic engineering constitutes a qualitative break from these ancient efforts in at least two major respects. First, unlike selective breeding, genetic

engineering (and the related discipline of synthetic biology) bypasses the sexual recombination of genomes in order to carry out precision genetic modifications and even full genomic customizations, promising exponentially greater control over the development of organisms and their properties. Genetic engineering can produce combinations of genes that are unlikely to ever be achieved via the "natural" sexual recombination of genomes. Much as selective breeding regimes engender fitness conditions that open up regions of phenotypic space that are unlikely to ever be explored by "natural" evolutionary processes, genetic engineering opens up regions of phenotypic space that are effectively off limits to selective breeding.

Rather than altering genes indirectly through the manipulation of breeding behavior (or artificial insemination/IVF), genetic engineering directly modifies the genomes of gametes or early embryos in order to produce a desired phenotype that will be transmitted to the next generation through ordinary reproduction. In this respect, genetic engineering acts like "directed mutation", or mutational pressures that are not random with respect to their fitness consequences—where fitness values are determined by the landscape of human desires and intentions.

Embryo selection, which involves the combination of IVF and pre-implantation genetic diagnosis, is currently being used to ensure that serious genetic disorders are not transmitted to offspring. Embryo selection may be particularly effective at eliminating monogenic diseases or selecting for a few genes associated with positive traits. However, when scores of genetic variants are being targeted, relying on random sexual recombination to produce an embryo with the desired overall combination of genes is totally infeasible. Having said this, we should not assume that embryo selection will be limited to the current handful of embryos that can be generated per cycle of in vitro fertilization. Stem cells could be used to create thousands of eggs from the same woman, and standard fertilization procedures could create nearly as many embryos with different genetic profiles. Nevertheless, when hundreds of genes and gene combinations are at stake, the process of random sexual recombination becomes incredibly inefficient. In this respect, a successful program of genetic engineering would constitute a radical improvement over embryo selection.

On the other hand, genetic engineering may encounter ethical problems that embryo selection avoids. For instance, one might argue that germline genetic engineering modifies the traits of an individual who will come to exist at some time in the future and who thus might be harmed by such interventions. In contrast, embryo selections decide *which* individual will come to exist at some time in the future, and thus the individual who ultimately develops from the selected embryo cannot be harmed by the aforementioned selection so long as her life is worth living, since absent the selection she would not have existed—some other person resulting from a different embryo would have existed. In this way, defenders of embryo selection can fall back on the "non-identity" problem whereas the proponents of genetic engineering cannot (Savulescu 2001; Savulescu and Kahane 2009).

Embryo selection has already begun (albeit modestly) to change the distributions of gene-based disorders in human populations. The same is true of genetic testing, which can influence human reproductive behavior by enabling people to choose genetically compatible partners or to forego reproduction altogether in order to avoid passing on deleterious mutations to their offspring. As more information is gleaned from genomes, and the genes implicated in positive traits such as intelligence,

memory, athletic ability and so forth are identified, genetic technologies will play an increasingly important role in human reproductive decisions. Genetic engineering is a more direct and (in principle) effective way of addressing the reproductive challenges to which genetic testing and embryo selection are currently providing solutions.

A second way in which genetic engineering differs from the selective breeding programs of the Neolithic Revolution is that the latter efforts were confined to non-human animal populations. The only deliberate, systematic attempt to modify the biological trait distributions of human populations culminated in the spectacular scientific and moral failure that was worldwide eugenics movement of the twentieth century, which resulted in some of the worst atrocities of recent human history. Much of the opposition to human genetic engineering stems from understandable fears that state-based eugenics programs could be recapitulated, and that liberal approaches toward genetic engineering technologies could lead down a slippery slope to human rights violating eugenics programs of the past. Much ink has been spilled comparing and contrasting the motivational and institutional factors that drove the old coercive eugenics programs from those that underpin the contemporary liberal project of genetic engineering (for discussions, see Harris 1993; Buchanan et al. 2001, ch. 2; Agar 2004). Even if human genetic engineering is limited to individual choice in the private sphere, the cumulative effect of such choices could ultimately have evolutionary and social implications that are even more profound than those associated with the Neolithic Revolution.

In addition to raising important ethical issues concerning justice, autonomy, harm, and other potentially conflicting values (for the seminal philosophical treatment of such conflicts, see Buchanan et al. 2001), the widespread use of genetic engineering raises a host of evolution-specific concerns that have rarely been touched upon. Science-savvy moral philosophers, and ethically savvy philosophers of science, have a unique (and complementary) role to play in delineating the intricate relationship between evolutionary facts and normative conclusions as they relate to emerging biotechnologies. Setting the agenda for this interdisciplinary role is a goal of this special issue.

2 The Implications of Genetic Enhancement for Human Evolution, Human Nature, and Persistence of the Human Species

Biological evolution refers to a genetically heritable change in the mean gene or trait distribution in a population over generational time. Gene and trait distributions have been changing under the mechanistic influences of evolutionary processes (selection, drift, mutation, migration, etc.) ever since the origins of life on Earth some 3.8 billion years ago. It is only in the last few decades that humans have developed the ability to induce site-specific genetic modifications in organisms, including humans, without manipulating patterns of reproduction. If these genetic modifications target the germline (gametes or early embryos) rather than somatic cell lines, they will be heritable across generations. Assuming human germ line modifications are widely undertaken, this could result in a change in the distribution of genetically heritable traits in human populations over time, thus impacting on human evolution.

Just what that evolutionary impact is likely to be, and the ethical implications that it might have, are questions ripe for biologically informed philosophical discussion.

Predicting long-term evolutionary outcomes is notoriously difficult for reasons that relate to the stochastic nature of evolutionary processes. Drift and mutation are inherently stochastic processes: they are epistemically random because they involve a complex configuration of innumerable causes, and they may be metaphysically random due to their causal dependency on quantum effects (a debate for another venue). Although natural selection admits of probabilistic predictions regarding future evolutionary trajectories in spatiotemporally restricted intervals (e.g., metal-tolerant plant varieties tend to out-reproduce metal-intolerant ones in contaminated soils), the fitness differences that lead to evolution by natural selection supervene on a vast array of geological, climatological, and ecological variables which are themselves stochastic and hence unpredictable. Darwin put the point this way: "natural selection will always act according to the nature of the places which are either unoccupied or not perfectly occupied by other beings; and this will depend on infinitely complex relations" (1859, 119).

Nevertheless, there are certain evolutionary concerns that we can identify and assess in connection with genetic engineering. One such concern relates to the implications of human genetic engineering for the evolvability of human populations: is human genetic engineering likely to create a biological monoculture that reduces our resistance to disease or renders human populations less flexible in the face of novel environmental challenges? If so, how might we mitigate these risks while securing the benefits of genetic modifications technologies? (for a recent exploration of these issues, see Powell 2012). A related issue concerns the relationship between human genetic engineering on the one hand, and persistence of the human species and its value on the other. Some view genetic technologies as a crucial tool for ensuring human survival and wellbeing during the period of environmental, ecological, and social turbulence that characterizes the modern epoch (Savulescu and Persson 2008; Buchanan 2011; Persson and Savulescu 2012). Others see the human genetic engineering project as poised to destroy human nature and drive the human biological species toward extinction. Francis Fukuyama expresses this worry as follows:

> While it is legitimate to worry about unintended consequences…the deepest fear that people express about technology is not a utilitarian one at all. It is rather a fear that, in the end, biotechnology will cause us in some way to lose our humanity—that is, some essential quality that has always underpinned our sense of who we are and where we are going, despite all of the evident changes that have taken place in the human condition through the course of history (2002, 101).

Fukuyama and other like-minded bioconservatives, such as Leon Kass (2003) and Michael Sandel (2007), advocate anti-enhancement legislation not because altering traits associated with human nature could lead to particular social or psychological harms, but out of deference to their deep-seated intuition that the integrity of human nature *as such* has inherent moral value. For these authors, as with many members of the public, the prospect of altering human nature is received as intrinsically repugnant on a gut-level (Kass 2000). They seek to establish the philosophical legitimacy of this intuition, which is difficult to justify in reference to canonical principles of bioethics, such as harm, autonomy or beneficence. Bioconservatives recognize that the principle-based framework of contemporary moral and political philosophy does

not provide reasons that would support this moral judgment. But the same is true, they retort, for other ubiquitous moral intuitions, such as those concerning bestiality or consensual incest, and thus they argue that our inability to articulate or defend our repugnance in some principled way should not prevent intuitions of repugnance from influencing the formation of public policy (for a discussion, see Clarke and Roache 2009).

From the standpoint of moral epistemology, it is unclear how instances of "legitimate" moral repugnance are to be distinguished from the misinformed or bigoted instances of moral repugnance that pepper the checkered moral history of recent civilizations, and the moral intuitions of tribal and hunter-gatherer societies before that (see Edgerton 1992). Kass may be right that there are instances of justified moral repugnance. But if the visceral disgust one experiences at the thought of modifying human genomes or transferring genes between species can be traced to a metaphysical confusion about human nature or the ontological status of species, then this gives us strong reason not to rely on the repugnance intuition as grounding for our considered moral judgments in these arenas. Since bioethicists contemplating the ethics of enhancement have appealed to the value of "human nature" as a reason for prohibiting human genetic modification, it is important that we consider the status of this concept in contemporary biological theory.

Does "human nature" remain a viable concept in the light of modern biology and, if so, does it possess normative significance? This two-part question is addressed by *Tim Lewens (this issue)* in his characteristically perspicacious contribution to this special issue. Lewens considers the extent to which bioethical debates over the enhancement of human nature are premised on pre-Darwinian conceptions of species, from which accounts of species natures and their purported ethical implications are derived. The "anti-essentialist consensus" in biology and the philosophy thereof is that species are not like atomic elements: they are not token members (like gold atoms) of a class (like gold) by virtue of possessing certain intrinsic properties that supervene on their microstructural constitution (like configurations of sub-atomic particles). Rather, species are individual lineages that originate, evolve, and go extinct in space and time; their members are united by virtue of the genealogical relations they stand in to one another, not by virtue of their possessing similar intrinsic or even ecological properties. Organisms may differ with respect to their intrinsic properties and nonetheless be individuals of the same species; conversely, organisms may possess similar intrinsic properties and yet be members of different species if the similarities they exhibit evolved independently. With the death knell of essentialism in biology having long since tolled, what room is left for species natures, including human nature?

Lewens considers a more plausible account of human nature that has been engineered for consistency with modern evolutionary theory. According to this "nomological" account, due to Machery (2008), traits are part of human nature if and only if they are widely shared by members of the species and result from an evolutionary process. Lewens argues that Machery's account is problematic mainly because (1) it ignores evolved variation within species (polymorphism) without a sufficient programmatic justification for doing so, (2) cultural outcomes may themselves result from evolutionary processes, (3) many "biological" traits associated with human nature rely partially on cultural resources for their development, and (4) many

"cultural" traits rely on biological resources for their acquisition and transmission. Lewens argues that any account of human nature that accommodates these issues will be forced to acknowledge that certain culturally contingent traits, such as playing rugby or practicing Catholicism, must be incorporated into any description of human nature—which amounts to a methodological gutting if not an outright *reductio* of the concept. At best, we can use human nature to refer to statistically reliable dispositions that characterize human populations. Such a permissive account of human nature offers little metaphysical grounding on which to build an ethical account of species natures.

In his reply to Lewens, *Edouard Machery (this issue)* offers pragmatic justifications for focusing on universally distributed traits, and he sketches a conceptual framework that might underpin this project. He contends that for a trait to be part of human nature it must be the proper subject of evolutionary *explanation*, which he believes excludes many of the "purely cultural" traits that Lewens believes cannot on Machery's own logic be excluded. In a separate commentary on Lewens, *Grant Ramsey (this issue)* offers an alternative account of human nature that he calls the "life-history trait cluster" account, which he believes avoids the evolutionary theoretical criticisms that dog other nomological conceptions of human nature. Ramsey argues that Machery's account provides little relevant information to would-be enhancement engineers, whereas the life-history trait cluster account of human nature can inform the ethics of human enhancement by (e.g.) helping us to identify traits that are amenable to alteration or augmentation without producing deleterious developmental sequelae. *Russell Powell (this issue)* channels Lewens's criticisms of the concept of human nature into a critique of prominent bioconservative positions on the ethics of biomedical enhancement. Powell argues that essentialist motivations creep into even biologically sophisticated discussions of human enhancement, and he argues that the interaction between biology and culture in the development of complex human traits undermines prominent bioethical views that presuppose the existence of an evolved biological potential that can be distorted or "disrespected" by agentic forces striving for mastery over our "given" evolved biology.

Even if it is the case that contemporary biological theory leaves little room for a substantive conception of human nature, it is still reasonable to ask whether biomedical enhancement is likely to undermine the persistence of the human biological *species* and any values associated with its continuance. Nick Agar (2010) has recently argued that "radical enhancements" of human individuals, or biomedical modifications that alter human capacities beyond the existing normal range, are likely to drive the human biological species toward extinction by replacing it with an "enhanced" and reproductively incompatible species, and he offers a variety of normative reasons for resisting this outcome.

Whereas most authors engaged in the enhancement debate have, like Agar, focused on the enhancement of human individuals, *Chris Gyngell (this issue)* considers whether genetic engineering technologies could be used in ways that enhance (or reduce) the evolutionary resilience of the human species. Against a backdrop of species concepts, Gyngell discusses several ways that genetic modification technologies could enhance the persistence of the human species. From an evolutionary perspective, population lineages must strike a balance

between the exploration of variation space and the exploitation of ecologically propitious variation. Genetic engineering technologies could increase the exploratory potential of human populations, Gyngell explains, by generating novel beneficial variations, and they could enhance exploitation capabilities insofar as they allow for the rapid delivery and spread of beneficial genetic variants outside of reproduction. However, Gyngell also points to ways that genetic engineering technologies could be used so as to decrease the overall probability of species persistence, and he suggests that the goals of individuals pursuing genetic enhancement may be in tension with the adaptability and survivability of the human species taken as a whole, raising the specter of normative conflict and a problem of collective goods.

3 Evolutionary Perspectives on Transgenic Organisms

Transgenesis, or the transfer of genes or genetic elements "horizontally" between reproductively isolated lineages, has been carried out in single-celled organisms, plants and complex animals, revolutionizing domesticated crop and animal yield and leading to hundreds of successful pharmaceutical treatments. At present, there is no evidential basis to conclude that genetically modified organisms are more likely than "naturally" genetically modified foods (i.e., those produced via selective breeding) to cause cancer, toxicity, infection, or ecological disruption. A common concern, albeit one with limited empirical support, is that transgenesis threatens to upset the delicate balance of biotic communities and ecosystems that have evolved for millions of years in the absence human interference (for a critical discussion of this view, see Powell 2010). Many, however, find such interventions ethically disturbing even if we set aside worries about their potentially harmful consequences. The prospect of transferring genes or living tissues between species that have been on distinct evolutionary paths for tens or hundreds of millions of years, strikes many people as hubristic and disrespectful, constituting a moral affront to nature itself. The prospect of transgenesis in non-human organisms is often met with attitudes that appear to be motivated by the same "yuck factor" that underpins bioconservative positions on human genetic enhancement and human cloning.

The assumption that inter-species gene exchange disrupts the natural biological order frames much of the ethical thinking on transgenesis, whether it focuses on the potential harm of trangensis or on its supposed intrinsic wrongness. However, a better understanding of the mechanisms and modes that characterize the evolution of life on earth can affect our moral perceptions of genetic engineering technologies and temper the intuitions of disgust they often invoke. *Da Fonseca et al. (this issue)* draw upon recent work in molecular evolutionary biology to argue that human nature is fundamentally transgenic. Once we realize that transgenesis between disparate animal groups is in fact a naturally occurring and evolutionarily important phenomenon, we will begin to view genetic engineering as a way to harness this natural mechanism in the service of human wellbeing, rather than as an intrinsic disruption of the natural order.

Although horizontal gene transfer is pervasive in bacteria, the exchange of genes outside of reproduction was long thought non-existent in complex animals where the

germ line is sequestered from environmental influences. As it turns out, however, much of the genome of animals, including that of humans, is comprised of "transposable elements" many of which may have been introduced laterally by viral vectors acting between evolutionarily distant lineages. Viruses subvert host immune systems, infect germ line cells, and integrate novel genetic elements into the host genome that may be transmitted down the generations. The relative frequency and importance of lateral gene transfer in animals is currently unknown, but Stancioli et al. offer reasons to suppose that it is both commonplace and evolutionarily significant, generating adaptive variation that lineages have been able to exploit. By showing that lateral gene transfer occurs among animal lineages in the wild, the authors hope to neutralize the "yuck" factor, and to persuade bioethicists that appealing to human genetic integrity as a reason to oppose transgenesis is a relic of a pre-Darwinian biological world view.

4 Artifactual Organisms

A further question that arises in relation to transgenic or synthetic organisms is whether their artifactual nature bears on their moral status. Biocentrists claim that all living organisms are morally considerable qua goal-directed or teleologically organized entities. One problem with this account is that many human artifacts, such as thermostats, computers and certain medical equipment, are also goal-directed entities, and thus they too would have moral status on the biocentrist account—a conclusion that looms perilously close to a *reductio* of the biocentrist position.

One biocentrist response to this objection, due to Varner (1998), is to ground the interests of non-sentient biological organisms in biological function, where the function of a trait is determined in relation to a history of natural selection for a given effect. On this view, insofar as teleological artifacts have not evolved via a process of natural selection, they will not have interests and thus cannot be morally considerable. The problem with this account, however, is that not all uncontroversial organisms have a history of natural selection that can ground interests. This is the challenge raised by *Sune Holm (this issue)*, who considers the imminent prospect of synthetic organisms, or living systems that are rationally designed from nucleic acid building blocks and which possess the same teleological properties as naturally evolved organisms. Holm argues that synthetic organisms pose a problem for biocentrist accounts that rely on the above etiological account of function to distinguish organisms from artifacts. After all, why should a naturally evolved organism with a certain set of intrinsic properties have a different moral status from a synthetically created organism with the same set of intrinsic properties? In place of Varner's view, Holm defends an "organizational" account of function, which grounds interests in capacities for self-maintenance rather than in a history of selection. He contends that the organizational account better accommodates our intuitions about the interests of synthetic organisms and the lack thereof in existing teleologically organized but non-self-maintaining artifacts.

In his commentary on Holm, *John Basl (this issue)* argues that a selection-based account of interest can and should be preserved even if *natural* selection-based etiologies are not necessary to ground interests. Basl argues that although both

organizational and selection-based accounts of biological interests may have to bite the bullet when it comes to certain canonical artifacts, the organizational account implausibly attributes interests to non-biological and non-artifactual entities, such as thermodynamic systems that exhibit properties of self-maintenance (e.g., candles, tornados etc.). Even if we accept that a history of selection, or a tendency toward self-maintenance, furnishes an entity with a certain class of interests, it remains an open question whether these interests are intrinsically morally valuable—that is, it is unclear whether they generate moral reasons to treat entities bearing those interests in particular ways.

5 Enhancing the Products of Evolution

Even if one does not believe there is anything intrinsically morally problematic about modifying genomes, human or otherwise, many fear the unintended and unforeseen consequences that might ensue from such "deep" interventions, which many see as proposing to tinker with the fundamental building blocks of life. Prominent bioethicists (e.g. Fukuyama 2003; Parens 1995) and even government-commissioned bodies have invoked the wisdom of nature in support of anti-genetic enhancement legislation, finding said wisdom in the evolutionary process. For example, the U.S. President's Council on Bioethics appeals to evolutionary considerations to justify its skepticism regarding the ability of genetic engineering to improve human wellbeing, stating that "the human body and mind, highly complex and delicately balanced as the result of eons of gradual and exacting evolution, are almost certainly at risk from any ill-considered attempt at 'improvement'...." (2003, 287–288). The assumption is that natural selection will gradually improve organismic function over long spans of evolutionary time, resulting in biological design that is optimally suited for its place in the stable economy of nature and unlikely to be improved upon by human genetic engineers.

Against such appeals to the wisdom of nature, Powell and Buchanan (2011) have argued that comparing evolution to a master engineer is a deeply flawed analogy that distorts rational assessments of risk in connection with new biotechnologies. In contrast to the skepticism of the above authors, Powell and Buchanan are optimistic that in many cases deliberate genetic modification could improve upon the products of evolution, given the severe constraints on mechanistic selection with respect to producing outcomes consistent with human good, and the ways in which forward-looking genetic engineering can overcome many of these constraints. One of the things they point out is that natural selection is not a process that acts for the good of the species or even for the good of the organism. Natural selection acts on human traits not because they enhance human *wellbeing*, but because they increase biological *fitness*: a morally neutral, technical term in biology that refers to an organism's expected reproductive success. People in developed nations who have the highest subjective and objective levels of wellbeing also generally have lower expected reproductive success, due mainly to family planning in the service of education, career, and other aspects of human flourishing.

A poignant illustration of how human wellbeing and biological fitness come apart relates to the phenomenon of senescence, or the ageing of the body and its

subsystems. The disease states associated with ageing are surely not in the interest of organisms. Why then do organisms age and why do different species age at different rates? In his seminal paper on the evolution of senescence, G.C. Williams (1957) argued that ageing is the result of a tradeoff between two opposing selective forces. The first is the investment in genes/traits that increase the vigor of youth, but which have incidental detrimental consequences for the ageing organism. Every animal has a reproductive probability distribution that rises at reproductive age and declines thereafter, even in the absence of senescence. This is because an organism's reproductive probability at a given time is a function of it surviving to that age, and an organism always has a lower chance of surviving to time $T+1$ than it does to T. Williams thus hypothesized, and offered empirical data in support of the claim, that genes which increase fitness in earlier stages in the life cycle will be disproportionately selected as compared to those which enhance fitness at later stages in the life cycle. Selection models suggest that even genes that confer a minor reproductive benefit early in life can become universally distributed in a population even though they come with catastrophic costs that manifest in later phases of the life cycle and cause senescence-related deterioration (Williams and Nesse 1991).

The tendency to invest disproportionately in youth is opposed by a second selective force that invests in cell maintenance and repair, and which to some extent, but never completely, mitigates the incidental detrimental effects of the former. Different species have struck the balance between these two opposing forces in different ways, depending on predation pressures, life histories, mating regimes, and so forth. The fact that different species senesce at different rates would seem to suggest that senescence is to an important degree under genetic control and thus amenable to genetic manipulation. Even if we thought it appropriate to defer to the "wisdom" of the evolutionary process, senescence-related traits would still be legitimate targets for enhancement insofar as they result from evolutionary "neglect" rather than selective attention.

Nevertheless, one might advance an evolutionary argument in support of the claim that attempting to enhance senescence-related traits will tend to produce undesirable unintended consequences. *Bennett Foddy (this issue)* explicates and critiques one such "argument from evolution", which he takes to be the strongest of the various ethical objections to using biotechnology to slow or eliminate the processes of ageing in humans. Noting the biological consensus that ageing is the result of evolutionary neglect rather than active optimization, Foddy considers a modified version of the argument from evolution: namely, that ageing is a necessary byproduct of an underlying optimal biological state. He argues, however, that there are two key reasons to be skeptical of this claim. First, the modern problem of senescence and its associated disease states arose in response to agriculture, medicine, public health, and other trappings of "civilization". Senescence is a recent socio-economic problem that our ancestors, who lived relatively fast and died relatively young, simply did not face, and thus its costs were not factored into the evolutionary trade-off calculus that putatively solidified in the environment of evolutionary adaptation. Second, the current science of ageing indicates that senescence is heavily affected by environmental variables, such as total caloric intake or exposure to free radicals (which increase oxidative stress)—factors that have increased markedly since the Agricultural Revolution. Foddy claims that manipulating these novel environmental variables, or engaging

in biomedical interventions to counteract their effects, would be consistent with acknowledging an evolutionarily optimized system of senescence.

6 Evolutionary Mismatch and the Prospects of Biomedical Enhancement

Some authors have advocated the use of biomedical enhancement as a means by which to ameliorate the constraints that natural selection has imposed on human behavior and cognition in the modern environment (Persson and Savulescu 2012). Drawing attention to the mismatch between traits that have contributed to biological fitness in the Paleolithic past and those that are currently conducive to human flourishing is an approach to the ethics of enhancement taken by *Earp, Sandberg and Savulescu (this issue)* in their defense of a moral obligation to use neuroenhancements to maintain and reinforce human pair bonds. They suggest that the same brain regions and biochemical pathways that underwrite monogamy in certain non-human mammalian species could be tweaked or coopted to achieve similar results in humans, bringing human behavior in better alignment with human moral values.

Earp et al.'s approach to the neuroenhancement of monogamy is informed by evolutionary theory in two crucial respects. The first is explanatory: understanding the environment of evolutionary adaptation can go some way toward explaining the relationship difficulties we face in the modern social environment. Moral values and commitments often conflict with deeply entrenched mating dispositions, leading to high rates of adultery, divorce and other forms of social strife with significant psychological fallout. The second way in which their account is informed by evolutionary theory is methodological in nature: evolutionary hypotheses regarding adaptation and common ancestry can help enhancement researchers to identify the proximate pathways implicated in human pair bonding, thus revealing potential targets for biomedical intervention.

The authors accept that the evolved psychological components of human nature might serve as a default guide to ethics and public policy, but they argue that biology-consistent norms should be overridden when they involve unjustified harm to others or when they conflict with basic moral principles. Because adultery leads to significant harms to partners and children, and because it is inconsistent with our expressed values and obligations, it represents an instance where our biological drives should to the extent possible be overridden by our moral values. And neuroenhancement, they claim, is a morally preferable alternative to social, legal and other environmental modifications that have been attempted in the past with limited success, in part because it addresses the problem of human pair bonding at its proximate biological "roots".

An extensive literature on biases and heuristics has identified numerous cognitive biases that may have been adaptive in the ancestral evolutionary environment, but which prevent human beings from acting rationally in a wide range of practical contemporary circumstances. These include cross-culturally robust phenomena such as confirmation bias, positive illusions, hedonic adaptation, hyperbolic discounting, probabilistic reasoning, and other failures of prudential rationality. *Neil Levy (this issue)* considers the mismatch between our evolved cognitive capacities and our ability to pursue personal conceptions of the good. He argues that manipulating social environments rather than "onboard capacities" is a preferable avenue by which

to achieve the legitimate goals of cognitive enhancement. This is because the cognitive sub-systems responsible for pervasive cognitive biases also produce substantial cognitive benefits in the modern environment, and so the alteration of onboard capacities is likely to come at a significant functional price. Even if the mind is thoroughly modular, developmental interconnections within cognitive modules can substantially limit the ability of neuroenhancement to manipulate cognitive traits in isolation from others. As we shall see, similar causal complexities may present a major obstacle not only to neuroenhancement, but also to the genetic enhancement project writ large.

7 Genetic Exceptionalism and the Genetic Blueprint Metaphor

Our knowledge of the developmental processes that lead from the embryo to the mature organism is still quite hazy, but one thing that is becoming clear is that development is so thoroughly interactive that it is difficult to accord genes a privileged causal role as compared to the myriad other factors that shape developmental outcomes (see Oyama 2000; Griffiths 2001; Robert 2004; Griffiths and Gray 2005). This emerging picture of development presents a serious challenge to gene-centric views of development that underpin a range of positions on the ethics of genetic enhancement.

Nuclear DNA is only one part of a staggeringly complex developmental apparatus involving an interaction between genetic, regulatory and environmental factors, exhibiting ubiquitous norms of reaction, gene-by-gene and gene-by-environment interactions, and so on. DNA is causally powerless without a legion of other biomolecular agents acting within the sub-cellular and extra-organismic milieu. Rather than claiming that DNA is causally determinative of the phenotype, proponents of gene-centric views of development instead suggest that DNA is "causally privileged" because it represents a "blueprint" that contains the "encoded instructions" or "programs" for "building" the phenotype. Biologists who use these engineering, computing, and information theoretic metaphors are usually well aware of their limitations (Kitcher 2001), but these limitations are rarely communicated to or acknowledged by non-specialists (including bioethicists and science journalists) who are quick to assimilate these catchy metaphors.

More sophisticated versions of the blueprint metaphor make use of the concept of genetic information: they contend that the genome contains the "information" necessary to build the organism and that this information is "read" in development. On this view, genes play a central "managerial" role, orchestrating the construction of the organism with the aid of non-genetic developmental mechanisms. Despite its intuitive appeal, no biologists or philosopher of biology has thus far been able to articulate and defend a substantive account of genetic information that (1) extends to the phenotype proper (as opposed to merely coding for the amino acids that make up the proteins that will ultimately comprise the phenotype) and (2) accords genes a privileged causal role in biological development (Godfrey-Smith 2000).

Why is it so problematic to say that genes play a special informational role in development? Information in the simplest *causal* sense refers to reliable co-variation between a sender and a receiver under specified channel conditions. Genes carry

information about phenotypes insofar as they reliably co-vary with phenotypes. However, because of the many-to-many relations of genes and phenotypes, genes often fail to correlate strongly with particular phenotypes, especially when it comes to complex traits (monogenic diseases are an exception). Moreover, on the causal account of information, genes are not informationally privileged developmental units. This is because, as Griffiths and others have argued (Griffiths 2001; Sterelny and Griffiths 1999), non-genetic factors will also co-vary with phenotypic outcome, affecting development in substantial and particular ways.

Some authors have argued that genes contain a type of semantic information due to the unique role they play in trans-generational inheritance and the evolution of adaptation (e.g., Shea 2012). But even these authors do not defend the thesis that genes play a special causal role in *ontogeny* (development)—and it is the casual role in ontogeny that is most relevant in the context of genetic engineering, which seeks to achieve certain phenotypic outcomes through genetic modifications that tend to *correlate* with those outcomes.

In contrast, semantic information entails an intentional imperative that confers normativity on developmental outcomes *regardless of developmental correlation*. Gene G for trait T involves an intentional imperative in relation to T, even when T goes unrealized in environmental contexts that alter the causal power of G. This allows us to say that G's "message" was "misread". Because semantic conceptions of genetic information are in Griffiths (2006) words "intrinsically unsuited to express the causal link between genes and complex phenotypes", they tell us little about the causal relationship between a gene and a complex phenotype, which is precisely what the genetic engineer needs to know. On the other hand, correlational genetic information does provide these crucial details, but to fall back on a correlational account of genetic information is to accept that genes do not contain information that other aspects of development (such as temperature, nutrition, socialization etc.) do not also contain.

Despite the conceptual drubbing it has taken from philosophers of science (e.g. Sarkar 1996; Griffiths 2001; Pigliucci 2010), the notion that the genome serves as a "blueprint" of the organism now runs so deep in the popular psyche that it arguably deserves to be labeled a fallacy. The genetic blueprint metaphor motivates a wide range of ethical, political and regulatory positions on genetic technologies. For instance, UNESCO's *Universal Declaration on the Human Genome and Human Rights* ties human nature and human dignity explicitly to the human genome. In medical law and ethics, genetic information is widely regarded as a special class of medical information that requires correspondingly unique levels of regulation, clinical attention, and ethical oversight. Few other medical technologies are treated with such regulatory caution and ethical trepidation.

This attitude of "genetic exceptionalism" has percolated into moral philosophical discussions of human genetic enhancement. Skeptics of the biomedical enhancement enterprise view genetic intervention as a qualitatively different sort of developmental intervention, both with respect to its alleged negative impact on human autonomy (e.g., Habermas 2003), and its ability to tinker with the essential building blocks of normal species design (e.g., Fukuyama 2002; Kass 2003; Sandel 2007). Defenders of the biomedical enhancement project have been quick to point out the fallacious reasoning of the aforementioned critics, such as their implicit or explicit commitment

to some form of genetic determinism—but they (enhancement enthusiasts) have done little to undercut the misleading blueprint metaphor. To the contrary, they have only reinforced the notion that a "map" of the human genome presents would-be enhancement engineers with a schematic of the human organism ready for rational redesign in the service of human wellbeing, if only we could figure out how to decipher "the code". Both parties to the enhancement debate, therefore, have a tendency to sensationalize the causal properties of the gene, and hence the power of genetic technologies to impact on human wellbeing.

None of this is to say that the role of genes in development is unimportant—only that their causal role may be more subtle, more complexly configured, and more contingent on context than many participants in the enhancement debate have recognized. This has important implications for the ethics of enhancement. The recognition that genes are causally comparable to other developmental factors should help to disabuse an apprehensive academy and public of the notion that genetic interventions are qualitatively different than other biomedical or traditional "environmental" enhancements, such as improvements in public health or early childhood education. And since genetic interventions will very often require paying close attention to environmental factors, the common worry that a focus on genetics will lead to the neglect of environmental factors also loses much of its force.

It is actually a common mistake of both critics and proponents of enhancement to forget the more general point that most interventions count as genuine enhancements only when considered in light of a certain context or environment (Savulescu and Kahane 2009). This is because the targets of enhancement are often only of instrumental value—valuable only in virtue of their effects. For example, a decrease in aggressive tendencies might be beneficial in the context of a peaceful, well-ordered society but could prove deadly in a violent and chaotic world. In this way, what counts as an enhancement is often a contextual matter: whether certain interventions that increase functional capacities are also likely to increase human wellbeing hinges in part on our ability to anticipate the future developmental context in which the intervention and its consequences will unfold. This, by the way, is as true for traditional environmental enhancements as it is for genetic interventions.

For all of these reasons, philosophers considering the ethics of enhancement should shift their focus from genetic enhancement to the *enhancement of development*, which involves manipulating any variables that bias the probabilistic unfolding of the phenotype toward desirable outcome space in light of the larger environmental context. The notion that genetic modifications alter organisms at the "most fundamental" or "deepest" level is yet another unhelpful metaphor that reinforces the genetic essentialism, gene-centricism and blueprint thinking that have tinted the philosophical lenses of opposing camps in the enhancement debate—in one case the tint is rosy, in the other it is something darker.

8 The Limits of Genetic Engineering

The interactive complexity of biological development could suggest that current approaches to genetic engineering may not entail the radical control of human design that skeptics and enthusiasts of human genetic engineering have often envisioned.

The current state of medical genetics is arguably consistent with this conclusion. More than 1,800 genes have been identified as causally associated with hereditary disorders, and genetic engineering is likely to be effective at eliminating these rare monogenic diseases (Porteus et al. 2006). However, it is becoming increasingly clear in the light of genome-wide association studies that the overwhelming majority of human diseases (as well as normal complex traits) involve the operation of many genes that are individually of minor effect, and which interact with the genetic and environmental backgrounds in complex ways. Because of the interactions between genes, their protein products, and various non-genetic factors in the cellular, prenatal and postnatal environment, multi-genic traits do not follow simple Mendelian patterns of inheritance (van Heyningen and Yeyati 2004). Our ability to personalize medicine in the light of individualized genomic data continues to be highly limited as a result.

One might argue that our current limitations in this regard are "merely epistemic", in the sense that they are likely to be overcome as the science of molecular developmental biology matures. Nevertheless, recent large-sample studies of monozygotic twins suggest that the predictive value of whole-genome profiles for an etiologically diverse range of diseases, such as cancer, autoimmune, cardiopulmonary, gastroenterological, and neurodegenerative disorders, may *never* reach the threshold for clinical utility (Roberts et al. 2012). We will briefly describe this study and its implications for the enhancement project.

It is well known that not all twins of a monozygotic pair die from the same disease. But if one monozygotic twin has a disease, and if the disease has a genetic basis (i.e., genetic variation explains some proportion of variation in the disease), then the second twin, which has a nearly identical genotype, will be more likely to develop the disease than the general population. Because standard heritability studies assess the genetic contribution to disease by comparing monozygotic with dizygotic twins, they reveal only the *average* genetic contribution across all monozygotic twins for a given disease. The problem that arises in using this average to gauge the ultimate clinical utility of using genome sequencing on the general population in order to test for a given disease is that monozygotic twins are not a homogeneous reference class in the relevant regard: some twin genomes carry a much higher probability of developing a given disease than other twin genomes, thus raising the average genetic contribution to disease as measured in monozygotic twin studies.

In order to ascertain the structure of this heterogeneous probability distribution, and from that infer the degree of "actionable" medical information that could be gleaned from whole genome sequencing of the general population, Roberts et al. introduced the concept of "genometypes": genomes with identical risks of developing a given disease. They then use sophisticated statistical modeling on a large epidemiological data set of monozygotic twins to assess frequency (proportion of twin pairs that might have a given genometype) and risk combinations for 20 different genometypes. Since the observed distribution of disease among monozygotic twins is consistent with many combinations of genometype frequency and risk, the researchers were able to estimate the upper and lower bounds on the clinical utility of whole genome data for the general population, without having to identify any of the implicated genes. They conclude that at the lower boundary, most sequenced patients would not gain any useful information, since their risk would be similar to that of the

general population. In the best-case scenario, the majority of patients might be alerted to one or a few disease risks.

This result is consistent with the general finding that the vast majority of common diseases appear to involve many genes of small effect. Unfortunately, the efficacy of genetic testing and (ultimately) genetic enhancement, at least for the foreseeable future, rely on there being fewer genes of larger effect implicated in a given trait. If this pattern holds true for diseases states, it is likely to hold true as well for cognitive and behavioral traits with highly complex developmental trajectories and convoluted reaction norms, such as psychological dispositions that have been found to have a genetic basis (see Bouchard and Loehlin 2001 for a review). This not only has unfortunate implications for the prospects of personalized medicine. It also casts serious doubt on the above "epistemic-limitations" thesis. Of course, things would be different if we possessed total causal information about trait development, but such postulations take us well beyond any reasonable discussion of enhancement possibilities.

Having said this, it would be too quick to conclude that the genetic manipulation of complex phenotypes such as memory, attention, intelligence, moral emotions, athletic ability and so on will be impossible for the foreseeable future. Work with animal models does indeed suggest that the gene-based alteration of human cognitive and behavioral functions is entirely possible. For instance, "Doogie mice" are transgenic mice which, as a result of genetic intervention, have been shown to exhibit enhanced memory, learning and intelligence, as demonstrated in a series of tests involving the navigation of a water maze, the recognition of new objects, and the time to extinction of a conditioned fear response. It is not surprising that such complex cognitive traits are to some degree under "genetic control". Less often remarked upon, however, is the fact that the Doogies' non-genetically enhanced counterparts (control mice) can achieve comparable levels of cognitive acuity simply by being exposed to toys, tunnels and other basic "educational enhancements" (Cooke and Bliss 2003)—reminding us that "ecological engineering", as Levy (this issue) puts it, can be as or more effective than "internal" engineering when it comes to realizing the goals of human enhancement.

On the other hand, tremendous evolutionary feats have been accomplished merely through selective breeding, and earlier we offered reasons to think that genetic engineering would far surpass standard "artificial selection". The remarkable range of morphological, cognitive and behavioral variation exhibited by the domestic dog is a quintessential example of selective breeding success, and what could explain this success if not the genetic profiles that were indirectly selected for by filtering their associated phenotypes? Whether human morphological, cognitive, and behavioral traits will be as genetically malleable as those of *Canis lupus familiaris*, or as recalcitrant to variability and manipulation as other mammal lineages, we simply do not know.

It is far too early to predict the extent to which genetic engineering will transform human biology and society. It would be a mistake, though, to assume that genetic engineering will be the most effective and/or radical form of biomedical enhancement, or even that it will be more effective and/or radical than traditional environmental enhancements. While genetic engineering represents a quantum leap over selective breeding and goes well beyond embryo selection in its ability to produce desirable combinations of genes, our ability to fundamentally reshape normal human

design may have to await a transformation in our understanding of biological development and the etiology of complex traits.

Nevertheless, the traits of human and non-human animals, and the environments in which they develop, are increasingly being determined by the intentional choices of forward-looking human agents with an eye toward human good, rather than by undirected evolutionary processes. Eventually, the human species will be in a position to assume substantial and deliberate control over its own evolutionary biological destiny, taking decisions that could affect the fate of human nature, the human species, and the future of life on Earth. This raises profound questions at the interface of moral philosophy, biological theory and biotechnology. We hope this special issue will take steps toward mapping out this uncharted philosophical terrain, and that it will inspire others to explore the metaphysical, epistemic and ethical dimensions of the new biosciences.

References

Agar, N. (2004). *Liberal eugenics: In defence of human enhancement*. Blackwell.

Agar, N. (2010). *Humanity's end: Why we should reject radical enhancement*. Cambridge, MA: MIT Press.

Bouchard, T. J., & Loehlin, J. C. (2001). Genes, evolution, and personality. *Behavior Genetics, 31*(3), 243–273.

Buchanan, A. (2011). *Beyond humanity.* Oxford: Oxford University Press.

Buchanan, A., Brock, D., Daniels, N., & Wikler, D. (2001). *From chance to choice: Genetics and justice*. Cambridge University Press.

Clarke, S., & Roache, R. (2009). Bioconservatism, bioliberalism and the wisdom of reflecting on repugnance. *Monash Bioethics Review, 28*, 4.1–4.21.

Cooke, S. F., & Bliss, T. V. P. (2003). The genetic enhancement of memory. *Cellular and Molecular Life Sciences, 60*, 1–5.

Darwin, C. (1859/1964). *On the origin of species by means of natural selection*. Harvard University Press

Edgerton R. (1992). *Sick Societies: Challenging the Myth of Primitive Harmony.* The Free Press.

Fukuyama, F. (2002). *Our posthuman future*. New York: Farrar, Straus and Giroux.

Fukuyama, F. (2003). *Our posthuman future*. New York: Farrar, Straus and Giroux.

Godfrey-Smith, P. (2000). Explanatory symmetries, preformation, and developmental systems theory. *Philosophy of Science, 67*, S322–S331.

Griffiths, P. E. (2001). Genetic information: a metaphor in search of a theory. *Philosophy of Science, 68*(3), 394–412.

Griffiths, P. E. (2006). The fearless vampire conservator: Philip Kitcher, genetic determinism and the informational gene. In: C. Rehmann-Sutter, & E. M. Neumann-Held (Eds.) *Genes in development: Rethinking the molecular paradigm* (pp. 175–198). Duke University Press.

Griffiths, P. E., & Gray, R. D. (2005). Three ways to misunderstand developmental systems theory. *Biology and Philosophy, 20*, 417–425.

Habermas, J. (2003). *The Future of Human Nature*. Polity.

Harris, J. (1993). Is gene therapy a form of eugenics? *Bioethics, 7*(1–2), 178–187.

Kass, L. (2000). The wisdom of repugnance: Why we should ban the cloning of humans. In: reprinted in G. McGee (Ed.) *The human cloning debate*. Berkeley Hills Press.

Kass, L. (2003). Ageless bodies, happy souls: biotechnology and the pursuit of perfection. *The New Atlantis* 9–28.

Kitcher, P. S. (2001). Battling the undead: How (and how not) to resist genetic determinism. In R. Singh, C. Krimbas, D. Paul, & J. Beatty (Eds.), *Thinking about evolution. Historical, philosophical and political perspectives*. Cambridge: Cambridge University Press.

Machery, E. (2008). A plea for human nature. *Philosophical Psychology, 21*, 321–329.

Oyama, S. (2000 [1985]). *The ontogeny of information*. Cambridge: Cambridge University Press.

Parens, E. (1995). Should we hold the (germ) line. *The Journal of Law, Medicine & Ethics, 173*, 173–176.

Persson, I., & Savulescu, J. (2012). Moral enhancement, freedom, and the god machine. *The Monist, 95*(3), 399–421.

Pigliucci, M. (2010). Genotype-phenotype mapping and the end of the "genes as blueprint" metaphor. *Philosophical Transactions Royal Society B, 365*, 557–566.
Porteus, M. H., Connelly, J. P., & Pruett, S. M. (2006). A look to future directions in gene therapy research for monogenic diseases. *PLoS Genetics, 2*(9), 1285–1292.
Powell, R. (2010). What's the harm? An evolutionary theoretical critique of the precautionary principle. *Kennedy Institute of Ethics Journal, 20*(2), 181–206.
Powell, R. (2012). The evolutionary biological implications of human genetic engineering. *The Journal of Medicine and Philosophy, 37*(3), 204–226.
Powell, R., & Buchanan, A. (2011). Breaking evolution's chains: the prospect of deliberate genetic modification in humans. *The Journal of Medicine and Philosophy, 36*(1), 6–27.
Richards, R. A. (1997). Darwin and the inefficacy of artificial selection. *Studies in History and Philosophy of Science C, 28*(1), 75–97.
Robert, J. S. (2004). *Embryology, epigenesis, and evolution*. Cambridge: Cambridge University Press.
Roberts, N. J., Vogelstein, J. T., Parmigiani, G., Kinzler, K. W., Vogelstein, B., & Velculescu, V. E. (2012). The predictive capacity of personal genome sequencing. *Science Translational Medicine*. doi:10.1126/scitranslmed.3003380.
Sandel, M. (2007). *The Case Against Perfection: Ethics in the Age of Genetic Engineering*. Harvard University Press.
Sarkar, S. (1996). Biological information: A sceptical look at some central dogmas of molecular biology. In S. Sarkar (Ed.), *The philosophy and history of molecular biology: New perspectives* (pp. 187–232). Dordrecht: Kluwer.
Savulescu, J. (2001). Procreative beneficence: why we should select the best children. *Bioethics, 15*, 413–426.
Savulescu, J., & Bostrom, N. (Eds.). (2009). *Human enhancement*. Oxford: Oxford University Press.
Savulescu, J., & Kahane, G. (2009). The moral obligation to create children with the best chance of the best life. *Bioethics, 23*(5), 274–290.
Savulescu, J., & Persson, I. (2008). The perils of cognitive enhancement and the urgent imperative to enhance the moral character of humanity. *Journal of Applied Philosophy, 25*(3), 162–177.
Savulescu, J., ter Meulen, R., & Kahane, G. (Eds.) (2011). *Enhancing human capabilities*. Oxford: Wiley-Blackwell.
Shea, N. (2012). Representation in the genome and in other inheritance systems. *Biology and Philosophy.*
Sterelny, K., & Griffiths, P. E. (1999). *Sex and death: an introduction to the philosophy of biology.* Chicago: University of Chicago.
Sterrett, S. G. (2002). Darwin's analogy between artificial and natural selection: how does it go? *Studies in the History and Philosophy of Science C, 33*, 151–168.
van Heyningen, V., & Yeyati, P. L. (2004). Mechanisms of non-Mendelian inheritance in genetic disease. *Human Molecular Genetics, 13*(2), R225–R233.
Varner, G. (1998). *In nature's interests? Interests, animal rights and environmental ethics*. Oxford: Oxford University.
Williams, G. C. (1957). Pleiotropy, natural selection, and the evolution of senescence. *Evolution, 11*(4), 398–411.
Williams, G. C., & Nesse, R. M. (1991). The dawn of Darwinian medicine. *Quarterly Review of Biology, 66*(1), 1–22.