

The Evolutionary Optimality Challenge¹

(2021)

Nick Bostrom, Anders Sandberg, and Matthew van der Merwe

Future of Humanity Institute

University of Oxford

www.nickbostrom.com

[in **Handbook of Bioethical Decisions**, eds. Erick Valdés & Juan Alberto Lecaros (Springer, forthcoming)]

Abstract

Human beings are a marvel of evolved complexity. When we try to enhance poorly-understood complex evolved systems, our interventions often fail or backfire. It can appear as if there is a “wisdom of nature” which we ignore at our peril. A recognition of this reality can manifest as a vaguely normative intuition, to the effect that it is “hubristic” to try to improve on nature, or that biomedical therapy is ok while enhancement is morally suspect. We suggest that one root of these moral intuitions may be fundamentally prudential rather than ethical. More importantly, we develop a practical heuristic, the “evolutionary optimality challenge”, for evaluating the plausibility that specific candidate biomedical interventions would be safe and effective. This heuristic recognizes the grain of truth contained in “nature knows best” attitudes while providing criteria for identifying the special cases where it may be feasible, with present or near-future technology, to enhance human nature.

Introduction

We marvel at the complexity of the human organism, how its various parts have evolved to solve intricate problems: the eye to collect and pre-process visual information, the immune system to fight infection and cancer, the lungs to oxygenate blood. The human brain is arguably the most complex thing in the known universe. Given how rudimentary our understanding of these highly complex systems, particularly the brain, how could we have any realistic hope of *enhancing* them?

¹ This chapter is based closely on Bostrom and Sandberg (2009). For helpful comments and corrections on this updated version, we are grateful to Tegan McCaslin, Richard Ngo, Perrin Walker, and Wes Cowley.

To enhance even a system like a car or a motorcycle—whose complexity is trivial in comparison to that of the human body—requires a fair bit of understanding of how the thing works. Isn't the challenge we face in trying to enhance human beings so difficult as to be hopelessly beyond our reach, at least until the biological sciences and our overall capabilities have advanced vastly beyond their present state?

It is easier to see how *therapeutic* medicine should be feasible. Intuitively, the explanation would go as follows: Even an excellently designed system will occasionally break. We might then be able to figure out what has broken and how to fix it. This seems much less daunting than to take an excellently designed, unbroken system and enhance it beyond its normal functioning.

Yet we know that even therapeutic medicine is very difficult. It has been claimed that until circa 1900, medicine did more harm than good.² Various studies suggest that even much of contemporary medicine is ineffectual or outright harmful.³ And, according to one estimate, iatrogenic deaths are the third leading cause of death in the US.⁴ We are all familiar with drugs, therapies, and nutritional advice once promoted by health authorities yet later found to be damaging. In many cases, those initial recommendations were informed by large clinical trials.

When even therapeutic medicine, based on fairly good empirical data, is so hard to get right, it would seem prudent to be wary of purported *enhancements*, especially when supported by much weaker data. Evolution is a process powerful enough to have developed systems far more complex and capable than anything human scientists or engineers have managed to design. Surely it would be foolish, absent strong supporting evidence, to suppose that we are currently able to do *better* than evolution,⁵ especially when we have not even managed to fully understand the systems evolution has “designed” and when our attempts just to repair them often misfire!

We believe that these informal considerations contain a grain of truth. Nonetheless, there are several particular classes of cases where we believe it is feasible to improve human nature. The evolution heuristic is our explanation of why this is so. If the evolution heuristic works as

² McKeown and Lowe (1974)

³ Newhouse and Group (1993); Frech and Miller (1996); Kirsch et al. (2002); Bunker (2001).

⁴ Makary and Daniel (2016)

⁵ At more advanced stages of technological development, it will be more reasonable to assume we can beat evolution at its own game.

we suggest, it shows that there is *some* validity to the widespread intuition that nature knows best, especially in relation to proposals for human enhancement. But the heuristic also shows that this validity is limited, and it reveals important exceptional cases in which we *can* hope to improve on nature using even our present or near-future science and technology.

The evolution heuristic might be useful for scientists working to develop enhancement technologies. It might also be useful in evaluating beliefs and arguments about the ethics of human enhancement, because intuitions about the wisdom of nature appear to play an important role in the cognitive ecology of many anti-enhancement advocates.⁶ While sophisticated bioconservatives (cognizant of the distinction between “is” and “ought”) may not *explicitly* base their arguments on the alleged wisdom of nature, we suspect that such intuitions often influence their judgements about mid-level moral principles invoked in the bioethical literature on human enhancement. Thus, addressing such hidden empirical background assumptions may help illuminate important questions in applied ethics.⁷

The evolutionary optimality challenge

The basic idea is simple. In order to decide whether we want to modify some feature of a system, it is helpful to consider why the system has that feature in the first place.⁸ Similarly, if we propose to introduce some new trait, we might ask why the system does not already possess it. The system of concern here is the human organism. The question of why it has a certain feature can be answered on multiple levels of explanation. Here our focus is on its evolutionary history.

We define an enhancement as an intervention that either improves the functioning of some subsystem (e.g. long-term episodic memory) beyond its normal healthy range, or adds a new capacity (e.g. magnetoreception).

Note that on this definition, an enhancement is not necessarily desirable, either for the enhanced individual or for society. For instance, we might have no reason to value an enhancement of our sweat glands that increases their ability, beyond the normal range, to

⁶ See, for example, Kass (2003).

⁷ On the role of mid-level principles in one area of applied ethics, see Beauchamp and Childress (1979). Earlier work has explored the extent to which opposition to enhancements results from an (irrational) bias for the status quo (Bostrom and Ord, 2006).

⁸ This is analogous to “Chesterton’s Fence” — “*There exists in such a case a certain institution or law; let us say, for the sake of simplicity, a fence or gate erected across a road. The more modern type of reformer goes gaily up to it and says, ‘I don’t see the use of this; let us clear it away.’ To which the more intelligent type of reformer will do well to answer: ‘If you don’t see the use of it, I certainly won’t let you clear it away. Go away and think. Then, when you can come back and tell me that you do see the use of it, I may allow you to destroy it.’*” (Chesterton, 1929).

produce perspiration in response to heat stimuli. In other instances, we might benefit from increased functionality or a new capacity, and yet not benefit from the enhancement because it also causes unacceptable side-effects.⁹ The evolution heuristic is a tool to help us think through whether some proposed enhancement is likely to yield a net benefit.

The starting point of our evolution heuristic is to pose the *evolutionary optimality challenge*:

(EOC) If the proposed intervention would result in a beneficial enhancement, why have we not already evolved to be that way?

Suppose that we “steelman” evolution by likening it to a surpassingly great engineer.¹⁰ We can then re-express the EOC as the question, “How could we realistically hope to improve on this great engineer Evolution’s work?” Note that it is the *limitations of this metaphor* that make it useful for our purposes. One does not have to actually believe that evolution is a great and wise engineer; rather, it is a useful exercise to consider precisely the ways in which this is *not* so, because those are the ways in which we may hope to do better.

We propose that there are three main categories of possible answers to the EOC: altered tradeoffs, evolutionary incapacity, and value discordance.

Altered tradeoffs

Evolution “designed” the system for operation in one type of environment, very different from the one we inhabit today. Modern conditions arose too recently for our species to have fully adapted to them, thus the tradeoffs struck by evolution may no longer be optimal today. It would not be surprising, then, if we were able to modify the system to better fit the novel requirements. It is much harder to design and build a car from scratch than it is to make some tweaks to improve function in a particular setting, for example, fitting it with a new set of tyres for icy roads. Similarly, the human organism, initially developed for operation as a hunter-gatherer on the African savannah, must now function in the modern world. We may well be capable of making some enhancing adjustments to fit the new environment even if our engineering talent does not remotely approach that of evolution.

⁹ Which side-effects are acceptable depends, of course, on the benefits resulting from the enhancement, and these may vary between subjects depending on their goals, life plans, and circumstances.

¹⁰ Messinger (2012)

Evolutionary incapacity

Even if some trait would have been adaptive in our ancestral environment, there is no guarantee that evolution would have discovered it. We have access to various tools, materials, and techniques that were unavailable to evolution. We can work backwards, starting with a goal in mind and figuring out the steps necessary to attain some trait. Even if our engineering talent were far inferior to evolution's, we may nevertheless be able to achieve certain things that stumped evolution, thanks to these novel aids. We should be cautious in invoking this explanation though; evolution often managed to achieve with primitive means what we are unable to do with state-of-the-art technology. But in some cases, we can show that it is practically infeasible to create a certain feature without some particular tool—no matter how ingenious the engineer—while the same feature can be achieved by any dimwit given access to the right tool. In these special cases, we might be able to overcome evolutionary restrictions without presupposing that our talent exceeds that of evolution.

Value discordance

Even if evolution had managed to build the finest reproduction and survival machine imaginable, we may still benefit from changing it because what we value is not primarily to be maximally effective fitness optimizers. There is a discrepancy between the standards by which evolution measured the quality of its work and the standards we wish to apply. It is not surprising that we can modify a system to better meet our goals if they differ substantially from the ones that (metaphorically might be seen as having) guided evolution in designing the system the way it did. Again, this explanation does not presuppose that our engineering talent exceeds evolution's. Compare the case to that of a mediocre technician, who would never be able to design a car, let alone a good one, but who may well be capable of converting the latest BMW model into a crude rain-collecting device, thereby *enhancing* the system's functionality in this respect.

*

In the following sections, we explore each of these categories of possible answers to the EOC in more detail.

Altered tradeoffs

Evolutionary adaptation requires striking tradeoffs between competing “design criteria”. Evolution has fine-tuned us for life in the ancestral environment, which, for the most part, was life as a member of a hunter-gatherer tribe roaming the African savannah. Because modern societies differ in many ways from the environment of evolutionary adaptedness, the tradeoffs struck by evolution may no longer be biologically optimal.

In evolutionary biology, the “environment of evolutionary adaptedness” (EEA) refers not to a particular time or place, but to the environment in which a species evolved and to which it is adapted.¹¹ It includes both inanimate and animate aspects of the environment, such as climate, vegetation, prey, predators, pathogens, and the social environment of conspecifics. We can also think of the EEA as the set of all evolutionary pressures faced by the ancestors of the species over recent evolutionary time—in the case of humans, at least 200,000 years. Hunting, gathering of fruits and nuts, courtship, parasites, and violent encounters with wild animals and enemy tribes were elements of the EEA; speeding cars, fast food, desk jobs, and tax returns were not.

If we can identify specific changes to our environment that have shifted the optimal tradeoff point between competing design desiderata in a determinable direction, then we may be able to find interventions that would “retune” the tradeoff to a point closer to the present optimum. Such retuning interventions might be among the low-hanging fruits on the enhancement tree—ones we could reach even without recourse to super-advanced biomedical technology.

Enhancements that aim to retune altered tradeoffs can often meet the EOC. A new trait might have been maladaptive in the EEA even though it would be adaptive now. Alternatively, the new trait might be intrinsically associated with another that was maladaptive in the EEA but has become less disadvantageous (or even beneficial) in the modern environment. In either case, the enhancement could be adaptive today without having been so in the EEA, providing an explanation of why we do not already have that trait, thus meeting the EOC.

We can roughly distinguish two ways in which tradeoffs can change. Firstly, new *resources* may have become available that were either absent or available only at great cost in the EEA.

¹¹ Bennett (2018)

Secondly, the *demands* placed on one of the human organism's subsystems may have changed since we left the EEA. Let us consider these two possibilities in turn.

Changes in resources

One of the main differences between human life today (for most people in developed countries) and life in the EEA is the abundant availability of food. In the state of nature, food is relatively scarce much of the time, making energy conservation important and implying tradeoffs between investments in metabolically costly tissues, processes, and behaviors. As we shall see, increased access to nutrients suggests several promising enhancement opportunities. We have also gained access to important new non-dietary resources, including improved protection against physical threats, obstetric assistance, better temperature control, and increased information availability.

We can illustrate these considerations by examining how they could apply to potential enhancements of the brain. (Throughout this chapter, the examples we give are designed mainly to be helpful in understanding how the heuristic works. They should not be read as a "favorite list" of the enhancements we think look most promising.)

Example: size and energy consumption of the brain

The human brain constitutes only 2% of body mass yet accounts for about 20% of total energy expenditure. Combined, the brain, heart, gastrointestinal tract, kidneys, and liver consume 70% of basal metabolism. This forces tradeoffs between the size and capacity of these organs, and between allocation of time and energy to activities other than searching for food in greater quantity or quality.¹²

Unsurprisingly, we find that, in evolutionary lineages where nutritional demands are high and cognitive demands low (such as bats hunting in uncluttered environments), relative brain size is correspondingly smaller.¹³

In humans, brain size correlates positively with cognitive capacity ($r \approx 0.4$).¹⁴ Holding brain mass constant, a greater level of mental activity might also enable us to apply our brains more effectively to process information and solve problems. The brain, however, requires extra energy when we exert mental effort, reducing the normally tightly regulated blood glucose

¹² Aiello, Bates and Joffe (2001); Fish and Lockwood (2003).

¹³ Niven (2005)

¹⁴ McDaniel (2005); Rushton and Ankney (2009).

level by about 5% for short efforts and more for longer exertions.¹⁵ Increasing blood glucose levels has been shown to improve cognitive performance in demanding tasks.¹⁶

The metabolic problem is exacerbated during prenatal and early childhood growth when brain development requires extra energy. Brain metabolism accounts for a staggering 60% of total metabolism in newborns,¹⁷ intensifying the competition between mother and child for nutritional resources during gestation and infancy.¹⁸ Children with greater birth weight have a cognitive advantage.¹⁹

Another constraint on prenatal cerebral development is the size of the human birth canal (itself constrained by bipedalism), which historically placed severe limits on the size of newborns' heads.²⁰ These constraints are partly obviated by modern obstetrics (particularly the availability of cesarean section). One way of reducing head size at birth and perinatal energy demands would have been to extend the period of postnatal maturation; however, delayed maturation was vastly riskier in the EEA than it is now.

What all this suggests is that cognitive enhancements might be possible if we can find interventions that recalibrate these legacy tradeoffs in ways that are more optimal in the contemporary world. For example, suppose we could discover interventions that moderately increase brain growth during gestation, or slightly prolong the period of brain growth during infancy, or that trigger an increase in available mental energy. Applying the EOC to these hypothetical interventions, we get a green light. We can see why these enhancements would have been maladaptive in the EEA and why they may nevertheless have become beneficial now that the underlying tradeoffs have changed, thanks to the plentiful availability of food. If the “downside” of more mental energy is that one burns more calories, many of us would regard this as a pretty good deal.

Not *all* cognitive enhancements get an immediate green light from this line of reasoning. Consider, for example, stimulants like caffeine and modafinil, which enable increased wakefulness and control over sleep patterns.²¹ Sleep, however, serves important yet poorly

¹⁵ Scholey, Harper and Kennedy (2001); Fairclough and Houston (2004).

¹⁶ Korol and Gold (1998); Manning et al. (1998); Martin and Benton (1999); Meikle, Riby and Stollery (2005); Smith et al. (2011). Increasing oxygen levels—another requirement for metabolism—also improves cognition (Winder and Borrill, 1998; Scholey et al., 2020).

¹⁷ Holliday (1986)

¹⁸ Martin (1996)

¹⁹ Matte (2001)

²⁰ Trevathan (1987)

²¹ Caldwell (2001)

understood functions besides energy conservation.²² This should give us pause. Without a clear understanding of the terms of the tradeoff struck by evolution, we cannot be confident we have met the EOC. In such cases, the heuristic counsels caution. If the reason we do not sleep less than we do has to do with these other functions, then reducing sleep might well turn out to have more problematic side-effects than increasing caloric expenditure.

Changes in demands

Just as there have been changes in the available resources, as compared to our hunter-gatherer ancestors' world, so too have there been changes in the demands we face in the modern environment. These suggest another set of potential opportunities for enhancement.

Many “diseases of civilization” are thought to be caused, at least in part, by changed demands. For example, our ancestors needed to exert themselves physically to secure adequate nutrition, whereas easy and continuous access to abundant food can promote obesity. Comfortable modern indoor environments lead us to spend less time outside, leading to widespread vitamin D deficiency.²³

Below, we consider two examples of possible enhancement targets suggested by such changes in demand.

Example: abstract thinking and mental focus

A capacity for abstract reasoning seems to have become more rewarded in contemporary society than it was in the EEA. There is a positive correlation in Western society between IQ and income.²⁴ Higher levels of general cognitive ability are important not just for many well-paid high-status jobs, but also for success in everyday life, such as for being able to fill out forms, understand news, and maintain health. As society becomes more complex, people with low cognitive ability are placed at an increasing disadvantage.²⁵

And while above-average general cognitive ability may have been somewhat advantageous in the EEA, the degree of change in demand that has occurred for some *particular* cognitive abilities (such as aptitude for numeracy and literacy) is even more dramatic. It would not be surprising if there were relatively minor neurological changes—perhaps achievable via

²² Siegel (2005)

²³ Thomas et al. (1998); Amrein et al. (2020).

²⁴ Neisser et al. (1996); Gottfredson (1997); Zagorsky (2007).

²⁵ Gottfredson (1997); Gottfredson (2004).

germline genetic interventions—that would greatly increase our faculties for formal mathematics and literacy, given that there has not been much specific selection for these traits (as opposed to selection for more general learning capabilities that can also be applied in these domains). Boosting our capacity for abstract symbol manipulation might be net beneficial in the modern environment even if it came at the expense of some other cognitive faculties—for example, if it left less cortical area for processing olfactory information, motor planning, landmark navigation, or visual motion detection.

The increased demand we now face for sustained attention on abstract cognitive tasks also suggests that we look for opportunities to adjust tradeoffs to favor such focused mental activity at the expense of other forms of processing. For example, in the EEA, it may have been important to sustain a high level of peripheral awareness to scan for potential predators and enemies, and much less important to be able to focus on a piece of text or a spreadsheet for hours at a time. In a modern white-collar environment, the priorities are reversed. The result is that levels of distractibility and external stimulation-seeking that may have been adaptive in the EEA are now dysfunctional, and a significant fraction of the population is diagnosed as suffering from ADHD. The changed demand for different forms of mental activity suggests that we may hope to find cognitive enhancers that work by shifting the balance from one form to another in ways that improve the tradeoff. For example, drugs such as methylphenidate and amphetamine can enable sustained focused mental effort (at the expense of more relaxed, unfocused, meandering, environment-aware forms of cognition), and they are frequently used for enhancement purposes.

Example: dietary preferences and fat storage

In the EEA, we needed fat deposits, but now it's better to have bank deposits. When food is reliably available and we have better ways to store resources, we face reduced demand for consuming and accumulating calories, yet we still have our old evolved cravings for high-calorie foods. This suggests opportunities for enhancement by altering our taste preferences or recalibrating our bodies' set-points for appetite and fat storage.

In principle, there are many routes to effectuate such a recalibration—ranging from nutritional advice, diet pills, artificial sweeteners, indigestible substances that taste like fat, weight loss clubs and hypnotherapy, to genetic or pharmaceutical interventions that change our hormonal or neuroregulatory systems, or interfere with lower-level metabolic pathways. The EOC does not explain why success in this direction has so far been limited despite considerable

investment, but it does hold out some hope that a solution to the obesity epidemic may be available (even with technology not much more advanced than the current state of art).

Evolutionary incapacity

We have discussed opportunities for enhancement arising from altered tradeoffs. Even if we think of evolution as a surpassingly great engineer, whose skills we cannot hope to match, we can nevertheless hope to achieve some enhancements by fine-tuning evolution's work to better fit the modern environment. We now turn to another source of potential enhancement opportunities: ones that arise from the fact that there are certain fundamental limitations in what evolution is able to do. Couched in the 'great engineer' metaphor, we could express this by saying that we may, without hubris, hope to achieve certain things with our clumsy fiddling that stumped evolution, because we have access to certain tools, materials, and techniques that the great ingenious engineer lacked.

Metaphors aside, we can identify several restrictions of evolution's ability to achieve fitness-maximizing phenotypes even in the EEA. We can divide these into three classes:

- *Fundamental inability*: evolution is fundamentally unable to produce some trait (even though the trait would be boosted fitness).
- *Local optima*: perhaps for contingent historical reasons, evolution got stuck in a local optimum that excludes some trait that would have been fitness-increasing.
- *Lags*: the development of a fitness-increasing trait, while evolutionarily feasible, would require so many generations that there has not yet been enough time for it to arise.

These three classes are not sharply separable. For example, one reason a trait may take a vast number of generations to develop is that it requires escaping from one or more local optima. Conversely, given very long time scales, even some traits that we may regard as fundamentally beyond evolution's reach might conceivably have evolved. However, the partition into these three classes can serve as a useful rough guide.

Fundamental inability

Biological evolution is limited in what it can achieve. For example, it seems unlikely that any biological organism could produce diamond. And while bacteria can produce microscopic metal crystals,²⁶ there is no way to unite them into contiguous metal. So it might not be

²⁶ Klaus et al. (1999)

surprising that evolution has not given us diamond tooth enamel or a titanium skeleton, even if these traits would have increased fitness in the EEA.

Examples are easy to multiply. Evolution could probably not have evolved high-performance silicon chips to augment neural computation, even though such chips might have been able to serve useful cognitive functions. A theoretical design of artificial red blood cells (“respirocytes”) has been published, which would enable performance far outside the range of natural red blood cells, allowing us to hold our breath for 3.8 hours. But the design relies on materials and pressures that are unavailable to evolution.²⁷

Engineered systems that radically depart from nature may create various complications with biocompatibility or functional integration with evolved systems. But at least there is no mystery as to why we have not already evolved these systems, even under the supposition that they would have been adaptive in the EEA. Enhancements that evolution is fundamentally unable to produce can therefore meet the EOC.

When invoking “fundamental inability”, it is important to determine that the inability does not pertain merely to the specific means whereby one intends to achieve the enhanced trait. If evolution would have been able to employ some *other* means to the same effect, we would have to wonder why evolution had not given us the trait via this alternative route, and the EOC would remain unanswered.

Local optima

Evolution sometimes gets stuck on solutions that are locally but not globally optimal. A locally optimal solution is one where any small change would make the solution worse, even if some bigger set of changes might make it better.

Being trapped in a local optimum is especially likely to account for failure to evolve polygenic traits that are adaptive only once fully developed, but incur a fitness penalty in their intermediary stages of evolution. In some cases, the evolution of such traits may require an improbable coincidence of several simultaneous mutations that might simply not have occurred among our finite number of ancestors. In these cases, a crafty genetic engineer could have some hope of attaining a solution that surpasses the one found by natural evolution. A human engineer can *plan*—starting with a goal in mind, working backward to

²⁷ Freitas (1998)

determine the genetic modifications necessary for its attainment, and then implementing the full set of needed modifications in one go. Goal-directed planning can often achieve outcomes that are infeasible to attain via myopic processes or random search.

Example: the appendix

The human appendix is a vestigial remnant of the caecum in other mammals. While it has some limited immunological function,²⁸ it easily becomes infected. In a world without surgery and antibiotics, appendicitis is a life-threatening condition (and it often occurs at a relatively young age). There is also some evidence that surgical removal of the appendix might reduce the risk of ulcerative colitis.²⁹ This would suggest that removal of the appendix might have increased fitness in the EEA.

A *smaller* appendix, however, *increases* the risk of appendicitis. Carriers of genes predisposing for small appendices have higher risks of appendicitis than non-carriers—and, presumably, lower fitness.³⁰ Therefore, unless evolution could find a way of completely doing away with the appendix entirely in one fell swoop, it might be unable to get rid of the thing, hence it remains, despite being a liability. If this story is correct, then an intervention that safely and conveniently removed the appendix might be a plausible enhancement capable of meeting the EOC.

*

Another source of evolutionary lock-in is *antagonistic pleiotropy*. This refers to a situation in which a gene affects multiple traits in both beneficial and harmful ways. If one trait is strongly fitness-increasing and the other mildly fitness-decreasing, the overall effect is positive selection for the gene.³¹ The local optimum here is to retain the gene in question. But the global optimum would be to circumvent the antagonistic pleiotropy, by evolving new genes that specifically produce the beneficial traits without causing the detrimental effects on other traits.

²⁸ Fisher (2000)

²⁹ Koutroubakis and Vlachonikolis (2000); Andersson et al. (2001).

³⁰ Nesse and Williams (1998)

³¹ Leroi et al. (2005)

Over longer timescales, evolution usually gets around antagonistic pleiotropy, for instance by evolving modifier genes that counteract the negative effects.³² However, such developments can take a long time, and in the meanwhile a species remains trapped in a local optimum.

Example: the $\epsilon 4$ allele

One well-known example of antagonistic pleiotropy is the $\epsilon 4$ allele of apolipoprotein E. Having one or two copies of this allele increases the risk of Alzheimer's disease in middle age but lowers the incidence of childhood diarrhea and may also have some protective effects during neurological development.³³ One potential enhancement that might therefore pass the EOC could be to add these alleles for their benefit in early life but then remove them or silence them in later life, to avoid paying the cost of increased Alzheimer's risk.

*

Yet another way in which evolution can get trapped into a suboptimal state is exemplified by the phenomenon of *heterozygote advantage*. This refers to the not uncommon situation where individuals who are heterozygous for a particular gene (i.e. possess two different alleles of that gene) have an advantage over homozygous individuals (who have two identical copies). Heterozygote advantage is responsible for many cases where potentially harmful genes are being maintained at a finite frequency in the population.

Example: the sickle-cell allele

The classic example of heterozygote advantage is the sickle-cell gene, where homozygous individuals suffer anemia while heterozygous individuals benefit from improved malaria resistance.³⁴ Heterozygotes have greater fitness than both types of homozygotes (those lacking the sickle-cell allele and those having two copies of it). Balancing selection preserves the sickle-cell gene in populations (at a frequency that varies geographically with the prevalence of malaria). The local optimum selected by evolution is one in which, by chance, some individuals will be born homozygous for the gene, resulting in sickle-cell anemia, a potentially fatal blood disease. The more global or ideal optimum—everybody being heterozygous for the gene—is unattainable by natural selection because of Mendelian inheritance, which gives each child born to heterozygous parents a 25% risk of being homozygous for the sickle-cell allele.

³² Hammerstein (1996)

³³ Oria et al. (2005)

³⁴ Allison (1954); Cavalli-Sforza and Bodmer (1999).

Heterozygote advantage suggests obvious opportunities for enhancement. Prenatal genetic screening could be used to guarantee that a child is born with exactly one copy of the allele, thereby securing the universal benefit of heterozygosity while avoiding the cost of some fraction of the population ending up homozygous. Other interventions could also be possible, such as somatic gene therapy or pharmaceuticals that reproduce the beneficial effects of heterozygosity in individuals lacking any sickle-cell allele.³⁵

*

Another kind of evolutionary lock-in is that of an evolutionarily stable strategy: “a strategy such that, if all the members of a population adopt it, no mutant strategy can invade”.³⁶ One way species can become trapped in such an equilibrium is via sexual selection. In order to be successful at wooing peahens, peacocks must produce extravagant tails which serve to advertise their genetic quality. Since only healthy peacocks can afford to grow and carry top-notch tails, it is adaptive for peahens to prefer to mate with peacocks that sport such impressive tails; and given this fact, it is also adaptive for peacocks to invest heavily in their rear plumage. However, it is likely that the *species* would have been better off (in the sense of becoming more abundant and more competitive relative to other species occupying the same niche) if it had evolved some less costly way for males to signal their fitness. Yet no individual peacock or peahen is able to defect from the evolutionarily stable strategy without thereby removing themselves from the gene pool. If there had been a United Nations of the peafowl, through which the birds could adopt a coordinated Millennium Plan to overcome their species’ vanity, the peacocks might well have voted for a sumptuary law that required them all to trim their tail feathers and adopt more modest attire.

The concept of an evolutionarily stable strategy can be generalized to that of an evolutionarily stable state. A population is said to be in an evolutionarily stable state if its genetic composition is restored by selection after a disturbance, provided the disturbance is not too large.³⁷ Such a population can be genetically monomorphic or polymorphic. Thus, while an evolutionarily stable strategy is one that is stable if *everybody* adopts it, an evolutionary

³⁵ Some individuals possess a variant allele (HbC) that provides malaria resistance without sickle-cell anemia in its homozygotic state. However, the HbC allele incurs a fitness penalty when heterozygous with either of the more prevalent alleles; and so exists only at low frequency in human populations (Wilkins and Godfrey-Smith, 2009). This suggests another enhancement option: to use genetic engineering to ensure homozygosity for the HbC allele.

³⁶ Smith (1982)

³⁷ *Ibid.*

stable state can encompass a set of different strategies whose distribution is stable under small perturbations. It has been suggested, for example, that the human population has been in a stable state in the EEA with regard to sociopathy, which can be seen as a defector strategy which can prosper when it is rare but becomes maladaptive when it is more common.

38

Lags

Evolution takes time—often, a *long* time. If conditions change rapidly, the genome will lag. Given that conditions for our hominid ancestors were quite variable—due to migration into new regions, climate change, social dynamics, advances in tool use, and adaptation in pathogens, parasites, predators, and prey—our species has never been perfectly adapted to its environment. Evolution is running up fitness slopes, but when the fitness landscape keeps changing under its feet, it may never reach a peak. Even when beneficial alleles or allele combinations exist, they may not have had time to diffuse across human populations. For some proposed enhancements, evolutionary lag can therefore provide an answer to the EOC.

This manner of meeting the EOC is related to the “altered tradeoffs” category, but with the difference that it focuses on ways in which *even in the EEA* we were not perfectly adapted to our environment. So there is the potential for an additional set of mismatches—and consequently for low-hanging enhancement opportunities—beyond those that have arisen with the dramatic changes in resource and demand that have followed the introduction of agriculture.

The speed of evolution is limited by many factors.³⁹ Some are inherent in the process itself, such as the mutation rate, the need for sufficient genetic diversity, and the constraint that selection can only encode a few bits into the genome per generation.⁴⁰ A recessive beneficial mutation will spread to an appreciable fraction of a fixed well-mixed population in time inversely proportional to its selective advantage. For example, if the mutation gives a 0.1% increase in fitness, it will take 9,200 generations to reach 50% of the population from a starting prevalence of 0.01%.⁴¹ Reviews of published studies have found that for most traits in most species, directional selection is fairly weak, suggesting that beneficial new traits are likely to spread slowly.⁴²

³⁸ Mealey (1995)

³⁹ Barton and Partridge (2000)

⁴⁰ Worden (1995)

⁴¹ Cavalli-Sforza and Bodmer (1999). Population structure (especially low-population bottlenecks) can significantly shorten the time it takes for a new allele to reach fixation.

⁴² Hoekstra et al. (2001); Kingsolver and Pfennig (2007).

There is evidence for recent positive selection in humans.⁴³ Some of it may be in response to climate variations, producing a wide range of variation in salt-regulating genes and skin pigmentation in populations far from the equator.⁴⁴ Significantly, genes involved in brain development have also been shown to have been under strong positive selection, with new variants emerging over the last 37,000 years and 5,800 years.⁴⁵

If we find a gene that has a desirable effect, and that evolved recently and has not yet spread far despite showing evidence of positive selection, interventions that insert it into the genome or mimic its effects would likely meet the EOC.

Example: lactase persistence

Humans typically lose the ability to digest lactose after infancy, due to decreased production of the lactase enzyme. While this may have been adaptive in the past, since it makes weaning easier, increased consumption of dairy products beyond childhood have stimulated selection for lactase persistence in humans over the last 5,000–10,000 years.⁴⁶ This is so recent that there has not been time for the trait to diffuse to all human populations—globally, 35% of adults are estimated to exhibit lactase persistence.⁴⁷ Taking lactase pills enables lactose intolerant people to digest lactose, widening the range of food they can enjoy. This enhancement clearly passes the EOC.

Value discordance

Our final top-level category of answers to the EOC focuses on the discordance between evolutionary fitness and human values. Even if human beings were optimal with respect to fitness in our current environment (and we have just seen that this is not always the case), this would provide no guarantee that we were optimal with respect to what matters to us. A great engineer may have built a system that efficiently serves one purpose; and it could still be unsurprising if a lesser engineer were able to tinker with it to make it better serve a different purpose.

⁴³ Voight et al. (2006)

⁴⁴ Thompson et al. (2004); Ju and Mathieson (2021).

⁴⁵ Evans et al. (2005); Mekel-Bobrov et al. (2005). The rapid growth of the brain in the human lineage also suggests that its size must be controlled by relatively simple genetic mechanisms (Roth and Dicke, 2005). It is noteworthy that, despite this, the selection differential for human brain weight during the Pleistocene was only 0.0004 per generation (Cavalli-Sforza and Bodmer, 1999).

⁴⁶ Bersaglieri et al. (2004)

⁴⁷ Gerbault et al. (2011)

Although our goals are not identical to those (metaphorically) pursued by evolution, there is considerable overlap. We value health; and health increases fitness. We value good eyesight; and good eyesight is useful for survival. We value musicality and artistic creativity; and these talents probably helped to attract mates in the EEA. If we are hoping to enhance some trait that is equally sought by evolution as it is by us, then we will not find an answer to the EOC in the discordance category, and we must either seek for an answer in one of the other categories or else suspect that what may *appear* to be an easy and unambiguous enhancement will in fact turn out to come at some large hidden cost. However, there are also many traits that we would value that would either have provided no evolutionary benefit in our ancestral environment, or else would not have done so to a sufficient degree to result in the extent of trait development that would be optimal from the perspective of our own values. These cases offer potential opportunities for feasible enhancement.

Example: contraceptives

Contraceptive technologies can be viewed as a form of enhancement, since they increase our control over our reproductive systems. We may value this because it makes family planning easier and increases choice. But evolution frowns on these practices. There is no mystery why we haven't evolved an easy reproductive off-switch under volitional control—evolution (no matter how skillful as an engineer) didn't *try* to do that. Contraceptives thus easily pass the EOC.⁴⁸

*

It is useful to distinguish two very different sources of value discordance. One is that the characteristics that would maximize an individual's fitness are not always identical to the characteristics that would be best *for her*. The other is that the characteristics that would maximize an individual's fitness are not always identical to those that would be *best for society*, or *impersonally best*. If our goal is to identify potential interventions that individuals would have prudential reasons for wanting, then we may perhaps set aside the second source of value discordance. If, however, we are interested in addressing broader ethical or public policy questions, then it is relevant to consider value discordance arising from either of these two sources. Let us review each of them in turn.

⁴⁸ Evolution might still have the last laugh if in the long run she redesigns our species to directly desire to have as many children as possible, or to have an aversion against contraceptives. Cultural evolution might beat biological evolution to the punch.

Good for the individual

There is a vast philosophical and empirical literature on the question of which traits promote individual well-being, which we shall not review here. For our purposes, it will suffice to list some candidates which might, with some plausibility, be claimed to contribute to individual well-being in a wide range of circumstances.⁴⁹ (This list is for illustration only—lists could be substituted without affecting the argument.)

Some traits that may promote individual well-being:

- Subjective well-being
- Freedom from severe or chronic pain
- Friendship and love
- Long-term memory
- Mathematical ability
- Beauty
- Awareness and consciousness
- Musicality
- Artistic appreciation and creativity
- Literary aptitude
- Confidence and self-esteem
- Athletic skill
- Healthy proclivities
- Mental energy
- Ability to concentrate
- Intelligence
- Longevity
- Social skills

To illustrate the idea, take long-term memory. Suppose that we believe that having better memory would tend to make our lives go better—perhaps because it would give us competitive advantages in the job market, or perhaps because we believe that memory is linked to other abilities or outcomes that would increase our well-being. We are considering

⁴⁹ The items in the list need not be restricted to final goods; it can include characteristics that are mere means to more fundamental goods. For example, even if one holds that musicality or musical appreciation is not intrinsically good, one can still include them in the list if one believes that they tend—as a matter of empirical fact—to promote well-being, e.g. by multiplying opportunities for enjoyment.

some intervention, perhaps a pill, that appears to improve memory. We then pose the EOC: Why has evolution not already endowed us with better long-term memory than we have?

Perhaps we find an answer in one of the categories covered above (altered tradeoffs and evolutionary incapacity). Yet suppose that we don't. We may then seek an answer in value discordance. Even if the intervention would have been maladaptive in the EEA, and even if it would still be maladaptive today, it may nevertheless be good for us, since what is good for us is not the same as what maximizes our fitness.

But we are not yet done. In cases like this, the evolution heuristic tells us that we should expect that the intervention will have some effect that reduces fitness. If we cannot form any plausible idea of what sort of effect this might be, then we should be wary. A fitness-reducing effect that we have not anticipated might be something very bad, such as a serious medical side-effect (which might manifest after a long delay) or some subtle functional deterioration that we cannot easily detect or attribute.⁵⁰ The EOC raises a warning flag.

If, however, we can give a plausible account of why the proposed intervention to improve long-term memory would reduce fitness, *and yet we judge this fitness-reducing effect as desirable or at least worth enduring for the sake of the benefit*, then we have met the EOC. This does not *guarantee* that the enhancement will succeed. It is still possible that the intervention will fail to produce the desired result or that it will have some unforeseen negative side-effect. There might be more than one sufficient reason why evolution did not already make this intervention to enhance our long-term memory. But once we have identified at least one sufficient reason, the warning flag raised by the EOC comes down.

Example: happiness

Evolution is not really concerned with our happiness and has instead produced many adaptations that cause psychological distress and frustration.⁵¹ The “hedonic-treadmill” causes us to quickly habituate to positive changes; gains that thrilled us at first soon get taken for granted and become a new baseline that we experience as barely adequate—presumably this was adaptive in the EEA as a way to prevent complacency.⁵² Similarly, sexual jealousy, romantic heartaches, status envy, competitiveness, anxiety, boredom, sadness and despair

⁵⁰ A relevant example here is the ‘Doogie’ lab mice, genetically engineered to have enhanced memory, but which also exhibited increased sensitivity to pain—something that would likely have been a fitness disadvantage in the EEA (Lehrer, 2009).

⁵¹ Buss (2000)

⁵² Diener et al. (1999)

may all have been conducive to survival and reproductive success in the Pleistocene and subsequently, yet they exert a heavy toll in human suffering.

An intervention that caused an upward shift in our hedonic set-point, or that down-regulated some of these negative emotions, would therefore meet the EOC. We can see why the intervention would have been maladaptive in the EEA, and yet believe that we would benefit from it because of a discordance between fitness and individual well-being: we value happiness more highly than evolution did.

Good for society

Many characteristics that promote individual well-being also promote the wider good, but the two lists are unlikely to be identical.

Some traits that may promote the social good:

- Extended altruism
- Conscientiousness
- Honesty and integrity
- Modesty and self-deprecation
- Originality, inventiveness, and independent thinking
- Civil courage
- Knowledge and good judgment about public affairs
- Empathy and compassion
- Nurturing emotions and caring behavior
- Just admiration and appreciation
- Sense of fairness
- Lack of racial prejudice
- Lack of tendency to abuse drugs
- Taking joy in others' successes and flourishing
- Useful forms of economic productivity
- Health

As with the list for individual well-being, this one is for illustration only. One could create alternative lists for various related questions, such as traits that are good for humanity as a whole, or for sentient life, or for a particular community, or traits that specifically help us become better moral agents. Such lists may overlap, but they will likely disagree about some

traits or their relative importance. The evolution heuristic can be applied using any such list as input, and the procedure is similar to that for the “good for the individual” type of value discordance.

Example: compassion

Suppose we have a drug that appears to make those who take it more compassionate. This might seem like a good thing, but why hasn’t evolution already made us more compassionate? Presumably, we could easily have evolved to produce some endogenous substance with similar effects to the drug; so the likely explanation is that a higher level of compassionateness would not have increased fitness in the EEA.

We then press on and ask *why* it is that greater compassionateness would not have been adaptive. And we can plausibly surmise that the reason is that such a trait would have been associated with evolutionary downsides—such as reduced ability to credibly threaten savage retaliation, or a tendency to spare the lives of enemies allowing them to come back another day and reverse their defeat, or an increased propensity to offer help to those in need beyond what is useful for reciprocity and social acceptance, and so forth. But these very effects, which would have made heightened compassion maladaptive for an individual in the EEA, are precisely the kinds of effects which we might believe would make it beneficial for the common good today. Note that we don’t have to assume that the relevant tradeoffs have changed since the EEA. Even in the EEA, it might have had net good effects for a local population of hunter-gatherers if one of their members were born with a mutation that caused an unusually high level of compassion; we just need to assume that the individual herself would have incurred a fitness penalty. If we accept these premises, then the hypothetical drug that increases compassionateness would pass the EOC.

The heuristic

The evolutionary optimality challenge asks, of an apparently attractive enhancement, why we have not already evolved the intended trait if it really is such a rad innovation. When trying to answer this question, we might find ourselves in one of several possible epistemic positions:

Current ignorance prevents us from forming any plausible idea about the evolutionary factors at play

This should give us pause. If we do not understand why a very complex evolved system has a certain property, there is a considerable risk that something will go wrong if we try to modify it. The case might be one where nature does indeed know best.

We are not claiming that it is *always* inadvisable to proceed with an intervention in a case like this. We might have other sources of evidence that reassure us that it will produce the intended result without causing unacceptable side-effects. For example, we might have used the intervention many times before, always to great success; or we might have experimental evidence from a closely analogous system, such as an animal model, suggesting that it should work in humans too. The evolution heuristic here delivers only a weak recommendation: that absent a good answer to the EOC, we should proceed with great caution, and we should be on the alert for the possibility that the intervention will turn out to have significant (though perhaps subtle) side-effects.

We come up with a plausible idea about the relevant evolutionary factors, and they suggest that the intervention would be harmful

In this case, our initial hopes of having identified a useful enhancement are undermined when we apply the evolution heuristic. None of the three categories we have described yields a satisfactory answer to the EOC: relevant tradeoffs have not changed since the EEA; evolution would have been capable of producing the intended modification by now; and there is no significant value discordance in relation to the targeted trait. Here, the heuristic gives us a strong reason for thinking that the enhancement intervention will fail or backfire. This is a case where we should respect the wisdom of nature.

We come up with several different plausible ideas about the relevant evolutionary factors

A third possibility is that we come up with two or more plausible but incompatible accounts of the evolutionary factors at play. We must then consider the implications of each of the different evolutionary accounts separately with respect to the EOC. If all of them show green lights, we are encouraged to proceed. If some of the evolutionary accounts show green lights but others show red lights, then we face a situation of familiar scientific uncertainty, and we can use decision theory to determine how to proceed. We might take the gamble if we feel that the balance of probabilities sufficiently favors the green lights; otherwise, we can attempt

to acquire more information in order to reduce the uncertainty, or forgo the potential enhancement and try something else.

We develop a plausible idea about the relevant evolutionary factors, and they imply we wouldn't have evolved the enhanced capacity even if it were beneficial

The final possibility is that we find a convincing account of the pertinent evolutionary factors which provides a satisfactory explanation of why we would not have evolved some trait even if it were overall beneficial. Then the heuristic gives us a green light to proceed. We have found grounds for a justified belief that, in the case before us, it would not be hubristic to suppose that we may be able to improve upon nature's work. Of course, it is still perfectly possible for us to fail—any specific intervention could have any number of *incidental* side-effects—and all the ordinary reasons for care and caution still apply; but there is no special “wisdom of nature” reason for pessimism in this case.

Discussion

There is a widespread belief in some kind of “wisdom of nature”. Many people prefer “natural” remedies, “natural” food supplements, and “natural” ways of improving ourselves (such as training, diet, education, and grooming). Offerings that are construed as “unnatural” are often viewed with suspicion. This negative attitude is especially strong in relation to biomedical means of enhancing human capacities, which are often viewed as unwise, short-sighted, or hubristic. We believe that such attitudes also influence normative intuitions in debates about human enhancement ethics.

While it is tempting to dismiss intuitions about the wisdom of nature as vulgar prejudice, we have argued that they contain an important grain of truth. We have attempted to extract this truth in the form of the evolutionary optimality challenge, which asks for any proposed enhancement: if it would indeed be so beneficial, why haven't we already evolved to be that way?

After posing this challenge, our heuristic instructs one to examine three broad categories of answers: altered tradeoffs, evolutionary incapacity, and value discordance. These categories correspond to systematic limitations of the wisdom of nature idea. For some potential enhancement interventions, the challenge can be met with an answer from one of these categories; for others, it cannot. The latter interventions do warrant extra suspicion, and attempting them may indeed be unwise and hubristic. In contrast, interventions for which we

can meet the EOC do not defy the wisdom of nature, and have a better chance of turning out well.

Pace Powell and Buchanan (2011), our argument does not rely on a (false) “strong adaptationist” assumption of evolutionary optimality.⁵³ On the contrary, the heuristic we have presented seeks to zoom in on the ways in which evolution is *not* optimal, although it does simultaneously emphasize that evolution can—in a certain circumscribed sense and within certain limits—usefully be characterized as a biological optimization process. If one overestimates the degree of evolutionary optimality that is typically found in nature, and one then applies the EOC and finds that it gives a green light to a particular proposed enhancement, this should *increase* one’s confidence that it would in fact be safe and beneficial. (The cost of overestimating evolution’s optimality, in the present context, is that it would increase our heuristic’s false-alarm rate—giving wisdom of nature arguments more credit than they are due.)

It should go without saying that we do not think that our heuristic should *replace* other more familiar ways of evaluating candidate enhancement interventions, such as via a detailed mechanistic level understanding of relevant biological systems or via well-designed clinical trials. Our claim is far more modest; that the heuristic can serve as a sometimes useful complement—an additional lense through which the (typically very messy) empirical situation can be viewed. It may be helpful in nominating promising candidate enhancement interventions and in setting reasonable prior expectations for the likelihood of success. The need for the heuristic would disappear if one had a complete and fully accurate understanding, at the mechanistic level, of all the relevant genetic and biochemical pathways involved. However, at present and for the foreseeable future, such a full understanding will often be unavailable, owing to the immense complexity of many biological systems—and the consequent possibility of subtle or delayed side-effects and unwanted interactions.

Conclusion

By understanding both the sense in which there is validity in the idea that nature is wise, and the limits beyond which the idea ceases to be valid, we are in a better position to identify promising interventions and to evaluate the risk-benefit ratio of existing enhancements. Furthermore, if we are right in surmising that intuitions about the wisdom of nature can exert

⁵³ Powell and Buchanan (2011)

an inarticulate influence on moral intuitions about biomedical enhancements, then our heuristic—while primarily a method for addressing empirical questions—may also contribute to normative debates surrounding (real or hypothetical) human performance enhancing technologies.

References

- Aiello, L. C., Bates, N., & Joffe, T. (2001). In defense of the Expensive Tissue Hypothesis. In D. Falk & K. R. Gibson (Eds.), *Evolutionary Anatomy of the Primate Cerebral Cortex* (pp. 57–78). Cambridge University Press.
- Allison, A. C. (1954). Protection Afforded by Sick Cell Trait Against Subtertian Malarial Infection. *British Medical Journal*, 1, 290–294.
- Amrein, K., Scherkl, M., Hoffmann, M., Neuwersch-Sommeregger, S., Köstenberger, M., Tmava Berisha, A., Martucci, G., Pilz, S., & Malle, O. (2020). Vitamin D deficiency 2.0: An update on the current status worldwide. *European Journal of Clinical Nutrition*, 74(11), 1498–1513.
- Andersson, R. E., Olaison, G., Tysk, C., & Ekblom, A. (2001). Appendectomy and protection against ulcerative colitis. *New England Journal of Medicine*, 344(11), 808–814.
- Barton, N., & Partridge, L. (2000). Limits to natural selection. *Bioessays*, 22(12), 1075–1084.
- Beauchamp, T. L., & Childress, J. F. (1979). *Principles of biomedical ethics*. Oxford University Press.
- Bennett, K. (2018). Environment of Evolutionary Adaptedness (EEA). In V. Zeigler-Hill & T. K. Shackelford (Eds.), *Encyclopedia of Personality and Individual Differences*. Springer International Publishing.
- Bersaglieri, T., Sabeti, P. C., Patterson, N., Vanderploeg, T., Schaffner, S. F., Drake, J. A., Rhodes, M., Reich, D. E., & Hirschhorn, J. N. (2004). Genetic signatures of strong recent positive selection at the lactase gene. *American Journal of Human Genetics*, 74(6), 1111–1120.
- Bostrom, N., & Ord, T. (2006). The reversal test: Eliminating status quo bias in applied ethics. *Ethics*, 116(4), 656–679.
- Bostrom, N., & Sandberg, A. (2009). The wisdom of nature: An evolutionary heuristic for human enhancement. In J. Savulescu & N. Bostrom (Eds.), *Human Enhancement* (pp. 375–416). Oxford University Press.

- Bunker, J. P. (2001). The role of medical care in contributing to health improvements within societies. *Int. J. Epidemiol*, 30(6), 1260–1263.
- Buss, D. M. (2000). The evolution of happiness. *American Psychologist*, 55(1), 15–23.
- Caldwell, J. A. (2001). Efficacy of stimulants for fatigue management: The effects of Provigil® and Dexedrine® on sleep-deprived aviators. *Fatigue and Transportation (II)*, 4(1), 19–37.
- Cavalli-Sforza, L. L., & Bodmer, W. F. (1999). *The Genetics of Human Populations*. Dover Publications.
- Chesterton, G. K. (1929). *The Thing*. Aeterna Press.
- Diener, E., Suh, E. M., Lucas, R. E., & Smith, H. L. (1999). Subjective well-being: Three decades of progress. *Psychological Bulletin*, 125(2), 276–302.
- Evans, P. D., Gilbert, S. L., Mekel-Bobrov, N., Vallender, E. J., Anderson, J. R., Vaez-Azizi, L. M., Tishkoff, S. A., Hudson, R. R., & Lahn, B. T. (2005). Microcephalin, a gene regulating brain size, continues to evolve adaptively in humans. *Science*, 309(5741), 1717–1720.
- Fairclough, S. H., & Houston, K. (2004). A metabolic measure of mental effort. *Biological Psychology*, 66(2), 177–190.
- Fish, J. L., & Lockwood, C. A. (2003). Dietary constraints on encephalization in primates. *American Journal of Physical Anthropology*, 120(2), 171–181.
- Fisher, R. A. (2000). The primate appendix: A reassessment. *The Anatomical Record*, 261(6), 228–236.
- Frech, H. E., & Miller, R. D. (1996). The Productivity of Health Care and Pharmaceuticals: An International Comparison. In *UCLA Research Program in Pharmaceutical Economics and Policy* 97–1.
- Freitas, R. A. (1998). Exploratory Design in Medical Nanotechnology: A Mechanical Artificial Red Cell. *Artificial Cells, Blood Substitutes, and Immobilization Biotechnology*, 26, 411–430.
- Gerbault, P., Liebert, A., Itan, Y., Powell, A., Currat, M., Burger, J., Swallow, D. M., & Thomas, M. G. (2011). Evolution of lactase persistence: An example of human niche construction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1566), 863–877.
- Gottfredson, L. S. (1997). Why g matters: The complexity of everyday life. *Intelligence*, 24(1), 79–132.
- Gottfredson, L. S. (2004). Life, Death, and Intelligence. *Journal of Cognitive Education and Psychology*, 4, 23–46.

- Hammerstein, P. (1996). Darwinian adaptation, population genetics and the streetcar theory of evolution. *Journal of Mathematical Biology*, 34(5–6), 511–532.
- Hoekstra, H. E., Hoekstra, J. M., Berrigan, D., Vignieri, S. N., Hoang, A., Hill, C. E., Beerli, P., & Kingsolver, J. G. (2001). Strength and tempo of directional selection in the wild. *Proceedings of the National Academy of Sciences*, 98(16), 9157.
- Holliday, M. A. (1986). Body Composition and Energy Needs during Growth. In F. Falkner & J. M. Tanner (Eds.), *Postnatal Growth Neurobiology* (pp. 101–117). Springer US.
- Ju, D., & Mathieson, I. (2021). The evolution of skin pigmentation-associated variation in West Eurasia. In *Proceedings of the National Academy of Sciences* (Vol. 118, Issue 1, p. e2009227118).
- Kass, L. (2003). *Beyond therapy: Biotechnology and the pursuit of happiness*. Executive Office of the President.
- Kingsolver, J. G., & Pfennig, D. W. (2007). Patterns and Power of Phenotypic Selection in Nature. *BioScience*, 57(7), 561–572.
- Kirsch, I., Moore, T. J., Scoboria, A., & Nicholls, S. S. (2002). The Emperor's New Drugs: An Analysis of Antidepressant Medication Data Submitted to the US Food and Drug Administration. *Prevention & Treatment*, 5.
- Klaus, T., Joerger, R., Olsson, E., & Granqvist, C. G. (1999). Silver-based crystalline nanoparticles, microbially fabricated. *Proc. Natl. Acad. Sci. USA*, 96(24), 13611–13614.
- Korol, D. L., & Gold, P. E. (1998). Glucose, memory, and aging. *American Journal of Clinical Nutrition*, 67(4), 764–771.
- Koutroubakis, I. E., & Vlachonikolis, I. G. (2000). Appendectomy and the development of ulcerative colitis: Results of a metaanalysis of published case control studies. *American Journal of Gastroenterology*, 95(1), 171–176.
- Lehrer, J. (2009). Neuroscience: Small, furry ... and smart. *Nature*, 461(7266), 862–864.
- Leroi, A. M., Bartke, A., Benedictis, G. D., Franceschi, C., Gartner, A., Gonos, E. S., Fedei, M. E., Kivisild, T., Lee, S., Kartaf-Ozer, N., Schumacher, M., Sikora, E., Slagboom, E., Tatar, M., Yashin, A. I., Vijg, J., & Zwaan, B. (2005). What evidence is there for the existence of individual genes with antagonistic pleiotropic effects? *Mechanisms of Ageing and Development*, 126(8), 421–429.
- Makary, M. A., & Daniel, M. (2016). Medical error—The third leading cause of death in the US. *BMJ*, 353, i2139.
- Manning, C. A., Stone, W. S., Korol, D. L., & Gold, P. E. (1998). Glucose enhancement of 24-h memory retrieval in healthy elderly humans. *Behavioural Brain Research*, 93(1–2), 71–76.

- Martin, P. Y., & Benton, D. (1999). The influence of a glucose drink on a demanding working memory task. *Physiology and Behavior*, 67(1), 69–74.
- Martin, R. D. (1996). Scaling of the Mammalian Brain: The Maternal Energy Hypothesis. In *Physiology* (Vol. 11, Issue 4, pp. 149–156).
- Matte, T. D. (2001). Influence of variation in birth weight within normal range and within sibships on IQ at age 7 years: Cohort study. *British Medical Journal*, 323(7314), 310–314.
- McDaniel, M. A. (2005). Big-brained people are smarter: A meta-analysis of the relationship between in vivo brain volume and intelligence. *Intelligence*, 33(4), 337–346.
- McKeown, T., & Lowe, C. R. (1974). *An Introduction to Social Medicine*. Blackwell Scientific.
- Mealey, L. (1995). The Sociobiology of Sociopathy—An Integrated Evolutionary Model. *Behavioral and Brain Sciences*, 18(3), 523–541.
- Meikle, A., Riby, L. M., & Stollery, B. (2005). Memory processing and the glucose facilitation effect: The effects of stimulus difficulty and memory load. *Nutritional Neuroscience*, 8(4), 227–232.
- Mekel-Bobrov, N., Gilbert, S. L., Evans, P. D., Vallender, E. J., Anderson, J. R., Hudson, R. R., Tishkoff, S. A., & Lahn, B. T. (2005). Ongoing adaptive evolution of ASPM, a brain size determinant in *Homo sapiens*. *Science*, 309(5741), 1720–1722.
- Messinger, C. (2012). Knocking Down a Steel Man: How to Argue Better. *The Merely Real*. <https://themerelyreal.wordpress.com/2012/12/07/steelmanning/>
- Neisser, U., Boodoo, G., Bouchard, T. J., Boykin, A. W., Brody, N., Ceci, S. J., Halpern, D. F., Loehlin, J. C., Perloff, R., & Sternberg, R. J. (1996). Intelligence: Knowns and unknowns. *American Psychologist*, 51(2), 77.
- Nesse, R. M., & Williams, G. C. (1998). Evolution and the Origins of Disease. *Scientific American*, 279(5), 86–93.
- Newhouse, J. P., & Group, T. I. E. (1993). Free for All? Lessons from the RAND Health Insurance Experiment. *Current Biology*, 15(16), 624–626.
- Niven, J. E. (2005). Brain Evolution: Getting Better All the Time? *Current Biology*, 15(16), R624–R626.
- Oria, R. B., Patrick, P. D., Zhang, H., Lorntz, B., Costa, C. M. D., Brito, G. A. C., Barrett, L. J., Lima, A. A. M., & Guerrant, R. L. (2005). APOE4 protects the cognitive development in children with heavy diarrhea burdens in northeast Brazil. *Pediatric Research*, 57(2), 310–316.

- Powell, R., & Buchanan, A. (2011). Breaking Evolution's Chains: The Prospect of Deliberate Genetic Modification in Humans. *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine*, 36(1), 6–27.
- Roth, G., & Dicke, U. (2005). Evolution of the brain and intelligence. *Trends in Cognitive Sciences*, 9(5), 250–257.
- Rushton, J. P., & Ankney, C. D. (2009). Whole Brain Size and General Mental Ability: A Review. *International Journal of Neuroscience*, 119(5), 692–732.
- Scholey, A. B., Benson, S., Sela-Venter, S., Mackus, M., & Moss, M. C. (2020). Oxygen Administration and Acute Human Cognitive Enhancement: Higher Cognitive Demand Leads to a More Rapid Decay of Transient Hyperoxia. *Journal of Cognitive Enhancement*, 4(1), 94–99.
- Scholey, A. B., Harper, S., & Kennedy, D. O. (2001). Cognitive demand and blood glucose. *Physiology & Behavior*, 73(4), 585–592.
- Siegel, J. M. (2005). Clues to the functions of mammalian sleep. *Nature*, 437(7063), 1264–1271.
- Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- Smith, M. A., Riby, L. M., van Eekelen, J. A. M., & Foster, J. K. (2011). Glucose enhancement of human memory: A comprehensive research review of the glucose memory facilitation effect. *Neuroscience & Biobehavioral Reviews*, 35(3), 770–783.
- Thomas, M. K., Lloyd-Jones, D. M., Thadhani, R. I., Shaw, A. C., Deraska, D. J., Kitch, B. T., Vamvakas, E. C., Dick, I. M., Prince, R. L., & Finkelstein, J. S. (1998). Hypovitaminosis D in medical inpatients. *New England Journal of Medicine*, 338(12), 777–783.
- Thompson, E. E., Kuttub-Boulos, H., Witonsky, D., Yang, L., Roe, B. A., & Rienzo, A. D. (2004). CYP3A variation and the evolution of salt-sensitivity variants. *American Journal of Human Genetics*, 75(6), 1059–1069.
- Trevathan, W. (1987). *Human Birth: An Evolutionary Perspective*. Aldine de Gruyter.
- Voight, B. F., Kudaravalli, S., Wen, X. Q., & Pritchard, J. K. (2006). A map of recent positive selection in the human genome. *Plos Biology*, 4(3), 446–458.
- Wilkins, J. F., & Godfrey-Smith, P. (2009). Adaptationism and the adaptive landscape. *Biology & Philosophy*, 24(2), 199–214.
- Winder, R., & Borrill, J. (1998). Fuels for memory: The role of oxygen and glucose in memory enhancement. *Psychopharmacology*, 136(4), 349–356.
- Worden, R. P. (1995). A Speed Limit for Evolution. *Journal of Theoretical Biology*, 176(1), 137–152.

Zagorsky, J. L. (2007). Do you have to be smart to be rich? The impact of IQ on wealth, income and financial distress. *Intelligence*, 35(5), 489–501.