O

Optimal Designs

Holly Elmore
Harvard University, Cambridge, MA, USA

Products and Byproducts of Natural Selection

What is the proper way to interpret the appearance of design in nature? Before 1859 and the publication of Darwin's *On the Origin of Species*, the answer seemed obvious – design means a designer. In *Natural Theology,* theologian William Paley famously analogized the appearance of design in the living world to the appearance of design in a man-made object, a watch:

> . . . [S]uppose I had found a watch upon the ground, and it should be inquired how the watch happened to be in that place . . . There must have existed, at some time, and at some place or other, an artificer or artificers, who formed [the watch] for the purpose which we find it actually to answer; who comprehended its construction, and designed its use. . .. Every indication of contrivance, every manifestation of design, which existed in the watch, exists in the works of nature; with the difference, on the side of nature, of being greater or more, and that in a degree which exceeds all computation. (Paley 1809)

In his book *The Blind Watchmaker* (Dawkins 1986), Richard Dawkins answers Paley's supposition:

> The analogy . . . between watch and living organism is false. All appearances to the contrary, the only watchmaker in nature is the blind forces of physics, albeit deployed in a very special way. A true watchmaker has foresight: he designs cogs and springs, and plans their interconnections, with future purpose in his mind's eye. Natural selection, the blind, unconscious, automatic process which Darwin discovered, and which we now know is the explanation for existence and apparently purposeful form of all life, has no purpose in mind. [. . .] If it can be said to play the role of watchmaker in nature, it is the *blind* watchmaker.

Paley's observation is still true today: we instinctively recognize purpose and intention in living things the same way that we recognize design in manmade artifacts. We now know that living things evolved into their present forms, and we look at how well organisms "fit" within the challenges of their environments in order to determine what selective pressures shaped them. But how well should we expect organisms to meet those challenges? What level of craftsmanship does the blind watchmaker employ? What does it mean when we see less than optimal designs in nature?

Case Study: The Spandrels of San Marco

We may have explained away the need for Paley's intelligent watchmaker, but the question remains, how good of a watchmaker is our blind watchmaker, natural selection? Are living things *particularly* well designed? And does *everything* have to be the product of natural selection? Some

T. K. Shackelford, V. A. Weekes-Shackelford (eds.), *Encyclopedia of Evolutionary Psychological Science*,
https://doi.org/10.1007/978-3-319-16999-6_2170-1

living things have features that seem poorly adapted to their environments. It's also possible that some features that look like great designs to human minds might only appear that way in hindsight or through the lens of overly felicitous post-hoc reasoning. How many features are mere "spandrels" – byproducts of genuine adaptations? In their oft-touted 1979 paper, *The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme,* Stephen Jay Gould and Richard Lewontin compare the design of living things to that of the Cathedral of San Marco in Venice, which features beautiful mosaic portraits of the evangelists in its spandrels – the spaces between the cathedral's dome and the curved arches that support it:

> The design is so elaborate, harmonious, and purposeful that we are tempted to view it as the starting point of any analysis, as the cause in some sense of the surrounding architecture. But this would invert the proper path of analysis. The system begins with an architectural constraint: the necessary four spandrels and their tapering triangular form. They provide a space in which the mosaicists worked; they set the quadripartite symmetry of the dome above.
>
> . . .[E]volutionary biologists, in their tendency to focus exclusively on immediate adaptation to local conditions, do tend to ignore architectural constraints and perform just such an inversion of explanation. (Gould and Lewontin 1994)

The metaphorical spandrel here is what Gould would later term an exaption – "a character, previously shaped by natural selection for a particular function (an adaptation), [coopted] for a new use" (Gould and Vrba 1982). This framing of exaption can be misleading, as there are no raw evolutionary materials – every feature that is genuinely adapted for a given purpose must have begun as part of a different feature – but it is useful to consider whether the selection pressures an organism faces in a modern environment are the same as those that contributed the most to its design. "Spandrel" is still used interchangeably with "byproduct" in evolutionary biology and psychology, and the phrase "just a spandrel" has something of a pejorative connotation.

The concept of architectural constraints adds explanatory power to our story of the spandrels of San Marco. The spandrels do not exist simply to display mosaics of the evangelists, but, contra Gould and Lewontin, neither are the mosaics an afterthought entirely secondary to the architecture of the church. In *Darwin's Dangerous Idea* (1996), Dan Dennett shows that our understanding of the spandrel could benefit from an adaptationist analysis. He begins by pointing out that San Marco features only one kind of spandrel, more properly known as a "pendentive":

> Not only were the pendentives just one among many *imaginable* options; they were just one among the readily *available* options. Squinches had been a well-known solution to the problem of a dome over arches in Byzantine architecture since about the seventh century.
>
> What the actual design of the San Marco spandrels—that is, pendentives—has going for it are mainly two things. First, it is (approximately) the "minimal-energy" surface (what you would get if you stretched a soap film in a wire model of the corner), and hence it is close to the minimal surface area (and hence might well be viewed as the optimal solution if, say, the number of costly mosaic tiles was to be minimized!). Second, this smooth surface is ideal for the mounting of mosaic images—and that is why the Basilica of San Marco was built: to provide a showcase for mosaic images.
>
> The conclusion is inescapable: the spandrels of San Marco aren't spandrels even in Gould's extended sense. They are adaptations, chosen from a set of equipossible alternatives for largely aesthetic reasons. They were designed to have the shape they have precisely in order to provide suitable surfaces for the display of Christian iconography. (Dennett 1996)

Dennett advocates the intentional stance when approaching biological and, by extension, psychological phenomena, assuming that complex features are adaptations that evolved to serve a purpose. Though Gould and Lewontin rightly point out natural selection is not the direct cause of all characteristics, it *is* the only cause of complex adaptations. Even though the spandrels of San Marco are in some sense byproducts, they can also be viewed on their own terms as adaptations when compared to equally possible alternatives. It was by assuming the intentional stance that Dennett was able to find a possible adaptive purpose to the spandrels.

The debate over the appropriate perspective on possible adaptations may seem to come from the armchair – why should we care whether adaptation is more or less pervasive in shaping living beings than the genetic drift, developmental constraints, and genetic inertia that also play a role? Can't we just look at the evidence for any given case?

One reason that we must consider grand questions of causation is that the balance of the various evolutionary forces informs our null expectations. Any attempt to reverse engineer relies on the reasoning that complex, well-adapted features don't arise by chance alone, so the degree of their aptness for a given purpose is evidence that they evolved for that purpose.

More fundamentally, the intentional stance is often the only place to start when asking "why" questions about design in living beings. "Adaptationism" is a productive approach for generating hypotheses and making novel, testable predictions (Barkow et al. 1992). While any particular explanation of adaptive purpose is just a hypothesis until it is validated, it is important to have the intellectual warrant to proceed from the assumption that complexity has an adaptive explanation. In fact, the same instinctive recognition of design that led Paley to believe in an Intelligent Watchmaker sets today's evolutionary scientists on the course to investigating adaptations.

What Is Optimal?

A Note on Terminology

In common parlance, "optimal" is used interchangeably with "maximal," but for our purposes, it is useful to distinguish the two. In engineering, optimal means "satisfactory, the best solution given all the constraints," while maximal means "greatest possible," and these are the definitions that I will adopt here. It is important to note that these terms are not always distinguished in biology or psychology, and their definitions may even be reversed.

There is a deeper philosophical issue that causes confusion between optimal and maximal designs – actualism vs. possibilism. The *Stanford Encyclopedia of Philosophy* explains:

> Actualism is the philosophical position that everything there is – everything that can in any sense be said to be – *exists*, or is *actual*. Put another way, actualism denies that there is any kind of being beyond actual existence; to be is to exist, and to exist is to be actual. Actualism therefore stands in stark contrast to possibilism, which [. . .] takes the things there are to include possible but non-actual objects. (Menzel 2016)

The disagreement boils down to what benchmark we are using to evaluate natural selection's designs. In our terms, adaptationists like Dawkins tend to be actualists – focused on how well natural selection solves a problem within its constraints and dismissing other possible sets of constraints as beside the point because those were not the constraints that *actually* impinged on natural selection in the given case. Those like Gould and Lewontin that focus on contingency and point out the inadequacy of nature's solutions often appeal to different possible, but not actual, circumstances under which the trait could have evolved to be better at solving the problems it now faces. Possibilists tend to see actual reality as just a subset of what could have been, whereas actualists tend to view what actually exists as all that could have been. I cannot resolve the actualism vs. possiblism debate in this chapter, but mention it in the hopes of empowering the reader to recognize sources of disagreement and misunderstanding around the concept of optimal designs.

I will henceforth use the word "optimal" to mean "best in the actualist sense" and "maximal" to mean "best in the possibilist sense." The concept of the adaptive landscape illustrates how a solution can be optimal, a local peak on the landscape, but not maximal, a global peak with respect to fitness. [Define fitness?]

Optimal Designs Are Peaks on the Adaptive Landscape

> Darwinism is not a theory of random chance. It is a theory of random mutation plus non-random cumulative natural selection. . .. Natural selection . . . is a non-random force, pushing towards improvement. . .. Every generation has its Darwinian failures but every individual is descended only from previous generations' successful minorities. . .. [T]here can be no going downhill - species can't get worse as a prelude to getting better. – Richard Dawkins, Climbing Mount Improbable

Adaptive landscapes (also known as "fitness landscapes") map genotype or phenotype to reproductive fitness. For any genotype or phenotype, all possible versions can be plotted as points on an adaptive landscape to show how fit each possible version of the phenotype or genotype is. Sewall Wright (1932) described an adaptive landscape as a *n*-dimensional space where "the entire field of possible gene combinations be graded with respect to adaptive value under a particular set of conditions." For simplicity's sake, they are often rendered in two dimensions and only with respect to one genotype or phenotype as below:

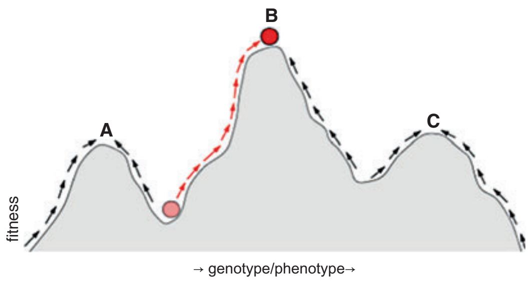

IMAGE adapted from: Public Domain, https://commons.wikimedia.org/w/index.php?curid=193687

Here we can clearly see how selection works given its constraints. Natural selection is a hill-climbing algorithm – it can only push a population up toward the nearest peak on the adaptive landscape. Theoretically, adaptive landscapes reflect all the constraints that shape them – the abiotic and biotic environment as well as the organism's other features. The constraints on a feature may be static or shift over time. When the constraints shift, the adaptive landscape reflects that shift, because the map of phenotype or genotype to fitness has changed. Even though it would be impossible to map every constraint, the metaphor of an adaptive landscape is useful because it shows us why only some paths are available for natural selection to climb. Natural selection cannot tell which peak is highest (maximal); it can only tell which direction is up.

A change in the shape of the landscape does not change natural selection itself, but it can channel the action of natural selection into a new direction by introducing new selection *pressures*. In adaptive landscape terminology, a selection pressure is the incline of the peak. The stronger the selection pressure, the steeper the incline. There is a great deal of room for chance and contingency in traversing an adaptive landscape, particularly in what genetic variation becomes available (i.e., which spots individuals could occupy on the landscape), but chance of this kind is not chance in the action of natural selection. Natural selection is always an optimization algorithm that accumulates chance events, such as beneficial mutations, in a nonrandom way.

Individuals in a population occupy one point on the landscape at any given time, and the population evolves when its position on the landscape changes. Occupying a peak on the fitness landscape means the population has attained a local maximum in fitness, or an optimal design.

Natural Selection Cannot Cross a Fitness Valley

Natural selection increases the proportion of alleles in the population that lead to greater reproductive success. This principle is axiomatic. Therefore, natural selection cannot push a population to cross a valley in the fitness landscape. Populations can get stuck on a local maximum and rendered unable to access higher peaks, because doing so would require passing through generations in which individuals of lower fitness had higher reproductive success than those of higher fitness – which natural selection will not do by definition.

Just because natural selection cannot push a population across a valley doesn't mean it is impossible for populations to move from peak to peak. There are other evolutionary forces, particularly genetic drift, that can move a population across the landscape regardless of the fitness consequences. Wright (1932) referred to this phenomenon as "shifting balance." In effect, drift allows a population to sample a larger area of the fitness landscape. Usually, this simply means lower mean fitness in that population, but sometimes a subset of individuals strays far enough away from one peak to "shift the balance" so that natural selection begins to drive it up another peak.

Genetic Drift Gets in the Way of Natural Selection

When genetic drift is strong enough to overcome selection, you get stuck with features that were not chosen for any particular reason but also are within a range of acceptable harm or sub-optimality. Because genetic drift is essentially sampling error, the balance of selection and drift depends on the size of the population. In small populations, drift can be more influential relative to selection because the sampling of alleles with each new generation is less representative of the previous generation than it would be in a larger population. Response to selection is measured in generations, so the longer generation times of large animals like humans put a speed limit on natural selection while also increasing the role of genetic drift. This means that evolution happens very slowly for humans unless a selection pressure is very strong.

Natural selection is the only process that can lead to complex adaptations. Drift can sometimes lead to better designs, but only by accident. Drift cannot preferentially discover high ground on the adaptive landscape the way natural selection can. And since there are many more ways to fail to solve a problem than to solve it, drift is usually a drag on fitness.

What "Optimal" Means

In adaptive landscape terminology, optimal means a local peak. A local peak may also happen to be a global peak, or maximum. When speaking of natural selection, an optimal design can only mean one that achieves a local maximum on the fitness landscape with regard to selection pressures felt probabilistically throughout a lineage's past. There are no absolutely optimal designs. A design can only be "optimal" with respect to its constraints. If there were no constraints on a design, there would be nothing to judge it by.

What "Optimal" Does Not Mean

Optimal designs are the best available, not the best possible. (Even designs that appear to occupy global maxima could still be made better if not for the physical constraints of our universe.) Nor does optimal mean inevitable or predetermined. How well a design feature performs a given task depends on the solutions available in design space, the genetic variation available to select on, and response to selection, which is not always strong or fast.

Also note that optimal designs may not be recognizably elegant – sometimes a "kluge," a clumsy but effective design (Marcus 2008), is the optimal solution to a problem given the energy budget or developmental constraints.

Optimal for What Purpose? And for Whom?

> The mind is what the brain does. – Marvin Minksy
>
> The whole purpose of our search for a 'unit of selection' is to discover a suitable actor to play the leading role in our metaphors of purpose. – Richard Dawkins, The Extended Phenotype

The topic of optimal designs is contentious and confused enough when we're talking about architecture and spandrels. It's even more confusing to talk about the design of the mind because the mind doesn't just follow whatever implicit goal natural selection moves it toward. The mind adopts its own goals. Failure to distinguish natural selection's "goals" from the mind's volition can make discussing the evolution of the mind nearly impossible.

The brain is the hardware from which the mind arises. Like a computer, the hardware was designed so that the machine could be applied to certain general purpose computations (Pinker 1997). Unlike a computer, the brain evolved by natural selection. When we use our on-board computer, we can adopt instrumental, shorter term goals than natural selection. Mind and brain evolved so that we could adopt those short term goals, but natural selection cannot dictate what those instrumental goals should be. Is a mind optimal if it follows natural selection's "goal," or is it optimal if it does what it, the mind, wants?

Richard Dawkins refers to goals and purposes adopted by minds as "neo-purposes" to distinguish them from the "archeo-purpose" of natural selection (ref: https://www.youtube.com/watch?time_continue=1776&v=mT4EWCRfdUg).

> Brains have evolved with various capacities that assist the survival of the genes that made them. Among these evolved capacities is the ability to set up goals, or purposes. [. . .] The brain is a kind of on-board computer used to control the body's behavior in ways that are beneficial to the genes that built it. It can perceive the outside world. It remembers things. It learns the consequences of its actions—good and bad. It sets up simulated models in imagination [. . .] It sets up purposes or goals in the sense of neo-purpose. The capacity to have a mental goal, or neo-purpose, is an adaption with a survival value, or archeo-purpose.

The mind can easily adopt neo-purposes that are at odds with the archeo-purpose of the genes and the blind processes that created it, particularly as it finds itself in environments the genes haven't "seen" yet through natural selection. Using contraception, for example, allows our brains to satisfy the desire for sex without ever producing the children sexual urges evolved to promote. The misgivings about children that lead us to use contraception have a long history as well – but the recent advent of highly effective birth control methods has changed the adaptive landscape by changing the consequences of these drives (Dawkins 2009). Philosopher Alan Gibbard (1990) points out that the implicit goals of the genes shouldn't dictate our goals as their vehicles:

> It is crucial to distinguish human goals from the Darwinian surrogate of purpose in the "design" of human beings [. . .] The Darwinian evolutionary surrogate for divine purpose is now seen to be the reproduction of one's genes. That is not, as far as I know, been anyone's goal, but the biological world looks as if someone quite resourceful had designed each living thing for that purpose [. . .] A like conclusion would hold if I knew that I was created by a deity for some purpose of his: his goal need not be mine.

Our brains are much more optimal for an environment that has disappeared, one in which the desire for sex was enough to ensure the creation of offspring. John Tooby and Leda Cosmides argue that the most productive approach to evolutionary psychology considers the ancestral environment of modern humans, particularly the Paleolithic era during which anatomically modern humans evolved (Barkow et al. 1992). So what is being optimized when there are multiple competing interests at stake?

Game Theory and Evolutionary Stable Strategies

When multiple agents are in conflict, it no longer makes sense to speak of optimal designs because there is no one goal with respect to which the design *can* be optimal. What we get are Nash equilibria or evolutionarily stable strategies (ESS), strategies which are unbeatable by any newly introduced strategy (Maynard Smith 1972). Essentially, a stable design is reached when there is a stalemate among all the conflicting goals. These situations are the province of game theory.

Symbiosis and Parasitism

The same human individual can host a number of different organisms with different goals. Human gut microbes may have a range of psychological effects, including the ability to affect mood (De Palma et al. 2014). Humans may also host parasites that increase their own fitness by manipulating our bodies, such as when the influenza virus causes us to sneeze and spread viral particles. Parasites can also influence their own fitness by manipulating a host's mind, such as when *Toxoplasma gondii* makes mice more aggressive and willing to confront cats, who are the parasite's desired next host. *Toxoplasma gondii* cannot successfully complete its life cycle in humans, but infection with *T. gondii* can nonetheless affect human behavior (Flegr 2007). Here, a parasite employing very well designed host manipulation causes very suboptimal behavior for both intended and unintended hosts.

Intragenomic Conflict

Human genes can only successfully reproduce by making new humans, but that does not mean that all the genes in the human genome share the same goals. The genes within the genome have largely overlapping goals, but the fiercest battles are often fought over scraps. For instance, alleles that are marked as either coming from the mother or the father (imprinted) promote behavior in offspring that benefits that parent-of-origin's reproductive success (Moore and Haig 1991). Where mom and dad's genes disagree, there are evolutionary games.

The brain appears to be a major arena of conflict, judging by the fact that imprinted genes are disproportionately expressed in the brain compared to other tissues (Davies et al. 2005). The contested resource in this case may be how much support offspring can draw from relatives that are asymmetrically related to them at the causal locus. For example, babies with the imprinting disorder Angelmann's syndrome have only the father's imprinting. Angelmann's babies only sleep an hour or two a day and they want a lot of milk. This is very exhausting for the mother, who wants time to rest and recover. In the inverse disorder, Prader-Willi syndrome, which results from having only the mother's imprinting, babies sleep far more than usual and don't often want to breastfeed. There may be no one optimal amount of sleep for the baby – the amount of sleep we need may simply be a compromise between conflicting goals of the mother's genes and the father's genes (Haig 2014).

Examples of Optimal Designs and Suboptimal Equilibria in Psychology

The Regions of the Brain: Layers of a Kluge

> Evolution often proceeds by piling new systems on top of old ones. The neuroscientist John Allman has captured this idea nicely with an analogy to a power plant he once visited, where at least three layers of technology were in simultaneous use, stacked on top of one another. The recent computer technology operated not directly, but rather by controlling vacuum tubes (perhaps from the 1940s), which in turn controlled still older pneumatic mechanisms that relied on pressurized gases. If the power plant's engineers could afford the luxury of taking the whole system offline, they would no doubt prefer to start over, getting rid of the older systems altogether. But the continuous need for power precludes such an ambitious redesign. (Marcus 2008)

As Gary Marcus explains in his book *Kluge: the Haphazard Construction of the Human Mind*, a kluge is "a clumsy or inelegant – yet surprisingly effective – solution to a problem." Marcus compares the construction of the human brain to the power plant control systems described above. The hindbrain is the most ancient part of the brain, shared by all mammals and reptiles, and controls our most basic reflexive functions such as breathing and alertness. The midbrain is involved in visual and auditory reflexes, and the forebrain is responsible for cognition, language, and judgment. Not only are higher functions like language dependent on more basic functions like breathing, higher regions of the brain are inextricably linked to the midbrain and the hindbrain (Bear et al. 2007).

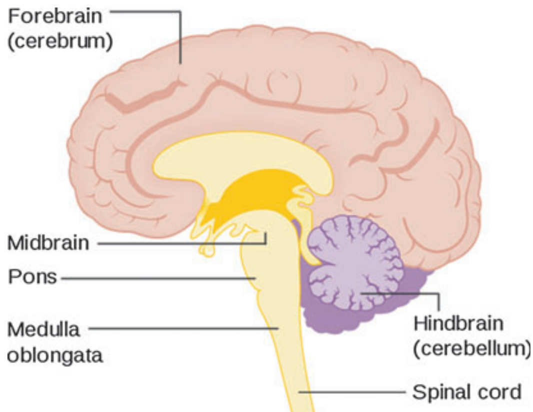

The forebrain sits on top of the midbrain which sits on top of the hindbrain. The structures are arranged in descending order of evolutionary age. https://upload.wikimedia.org/wikipedia/commons/thumb/8/83/Diagram_showing_the_brain_stem_which_includes_the_medulla_oblongata%2C_the_pons_and_the_midbrain_%282%29_CRUK_294.svg/572px-Diagram_showing_the_brain_stem_which_includes_the_medulla_oblongata%2C_the_pons_and_the_midbrain_%282%29_CRUK_294.svg.png

We have no reason to believe our brains are at a global maximum on the adaptive landscape. Ascending layers of updated systems in our brains indicate a history of updates on top of updates while the machine was still running. Evolution was constrained, in this case, by the existing brain's crucial job in survival and reproduction.

Optical Illusions: The Pitfalls of Solving an Impossible Problem

In *How the Mind Works* (Pinker 1997), Steven Pinker explains that from the perspective of light entering our eyes, there are no objects. There is just light of differing wavelengths. The mind processes that information to define objects, "the movable hunks of matter that we count, classify, and label with nouns" in our field of view. But defining objects is a difficult philosophical problem. Where does the nose end and the face begin? When the does head give way to the neck? Many objects lack clear boundaries, but the brain tries to find them all the same. The famous vase-face illusion shows that our brain can flip between entertaining one object (the vase interpretation) to another (the face interpretation) with the same visual information.

The face-vase illusion. https://commons.wikimedia.org/wiki/File:Rubin2.jpg

Optical illusions reveal the biases in our sensory processing. When our minds detect an object, that interpretation pops out and excludes others. In the case of the above vase-face illusion, the vase pops out and suddenly the faces are gone until the faces pop out and suddenly the vase is gone. As this illusion demonstrates, determining what it is we are seeing based on light alone is sometimes a hopelessly ambiguous problem, but the fact that optical illusions are reserved as parlor tricks shows a very robust and perhaps optimal mental design for perceiving objects.

Cognitive Biases: Optimized Shortcuts

Human cognition is subject to a long list of biases (https://en.wikipedia.org/wiki/List_of_cognitive_biases). These biases seem like suboptimal thinking because they don't always tell us the truth and occasionally fail to give accurate information. But when we view the purpose of human cognition more holistically, we see that time and processing power are limited, so in many cases shortcuts are the only way to get to an answer at all. Sometimes heuristics are better performers than slow and cumbersome solutions to those problems that aren't important to solve perfectly (Kahneman 2011).

To give a classic example, when the risks of ignoring a predator in the bushes are far greater than the risks of incorrectly interpreting the sound of a breeze as a predator in the bushes, over-interpreting noises can be an optimal solution. When our ancestors heard a sound, their minds jumped to the most emotionally salient possibility: a predator. This is called the availability heuristic – we judge the likelihood of an event by how easily it comes to mind, and emotionally salient things come to mind more easily. This is bad for accuracy, but very good for survival. A few minutes of panic and wasted time are much less costly than failing to avoid a real attack. With more time and resources, our predator detection could be even sharper, but when a heuristic works well enough, spending more resources on improving it would be gilding the lily. Those resources could be better spent elsewhere if the cost of overreacting is small.

Summary

There is an on-going debate in evolutionary sciences about whether natural selection consistently provides optimal designs. The very idea of optimal designs hits a fault-line in biology and psychology between actualists and possibilists. The definition of optimal itself is disputed and a matter of perspective.

Natural selection can only climb hills on an adaptive landscape, and that landscape represents all the constraints on a feature's design. Genetic drift usually causes features to be suboptimal by overcoming the work of selection, but occasionally opens up the path for natural selection to climb a higher peak.

Optimal means the best design available, not the best design period. There are no absolutely optimal designs – they could always be better if some feature of the environment were better. You could almost say that the most perfect design is nothing, because then there's no problem!

Sometimes designs are not optimal with respect to their function because they are satisfying different and conflicting goals. Sometimes, as with freestanding arches and vision, conflict at one level is itself an optimal design at a higher level of function. Sometimes, as in the case of cognitive biases, a design is optimal even though it sometimes fails to perform its job because it's better to err in the right direction. Sometimes a design is suboptimal because other agents are working to exploit design weaknesses, as in the case of sensory bias.

Finally, it's important to remember that no matter how perfect or imperfect a feature appears to be, it always helps to consider what purpose it might have evolved to serve. What Stephen Jay Gould would call adaptationist thinking is a highly effective source of hypotheses, and sometimes hypotheses become part of evolutionary theory.

References

Barkow, J. H., Cosmides, L., & Tooby, J. (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. New York: Oxford University Press.

Bear, M. F., Connor, B. W., & Paradiso, M. A. (2007). *Neuroscience: Exploring the brain*. Philadelphia: Lippincott Williams & Wilkins.

Buller, D. J. (2005). *Adapting minds: Evolutionary psychology and the persistent quest for human nature*. Cambridge, MA: MIT Press.

Christian, B., & Griffiths, T. (2016). *Algorithms to live by* (1st ed.). New York: Henry Holt & Co.

Davies, W., Isles, A., & Wilkinson, L. (2005). Imprinted gene expression in the brain. *Neuroscience & Biobehavioral Reviews, 29*(3). https://www.sciencedirect.com/science/article/pii/S0149763404001678

Dawkins, R. (1986). *The blind watchmaker*. New York: Norton.

Dawkins, R. (1996). *Climbing mount improbable*. New York: Norton.

Dawkins, R. (1982). *The extended phenotype: The gene as the unit of selection*. Oxford: Freeman.

Dawkins, R. (2009). The Purpose of Purpose. https://www.youtube.com/watch?time_continue=1776&v=mT4EWCRfdUg.

Dennett, D. C. (1996). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.

De Palma, G., Collins, S. M., Bercik, P., & Verdu, E. F. (2014). The microbiota–gut–brain axis in gastrointestinal disorders: Stressed bugs, stressed brain or both? *The Journal of Physiology, 592*(Pt 14). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4214655/.

Flegr, J. (2007). *Effects of* toxoplasma *on Human Behavior. Schizophrenia Bulletin., 33*(3), 757–760. https://doi.org/10.1093/schbul/sbl074 https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2526142/.

Gibbard, A. (1990). *Wise choices, apt feelings: A theory of normative judgment*. Cambridge, MA: Harvard University Press.

Gould, S. J., & Vrba, E. S. (1982). Exaptation – A missing term in the science of form. *Paleobiology, 8*(1), 4–15.

Gould, S. J., & Lewontin, R. C. (1994). The spandrels of san Marco and the Panglossian paradigm: A critique of the Adaptationist Programme. In E. Sober (Ed.), *Conceptual issues in evolutionary Biology_* (pp. 73–90). Cambridge: MIT Press.

Haig, D., & Wharton, R. (2003). Prader-Willi syndrome and the evolution of human childhood. *American Journal of Human Biology, 15*(3), 320–329 https://www.ncbi.nlm.nih.gov/pubmed/12704708.

Haig, D. (2014). Troubled sleep: Night waking, breastfeeding and parent-offspring conflict. *Evolution, Medicine, and Public Health., 1*, 32–39.

Kahneman, D. (2011). *Thinking fast and slow* (1st ed.). New York: Farrar, Straus, & Giroux.

Marcus, G. (2008). *Kluge: The haphazard evolution of the human mind* (pp. 11–12). Houghton Mifflin Harcourt Kindle Edition.

Maynard Smith, J. (1972). *"Game theory and the evolution of fighting". On evolution*. Edinburgh: Edinburgh University Press. ISBN 0-85224-223-9.

Menzel, C. (2016). In E. N. Zalta (Ed.), *Actualism. The Stanford encyclopedia of philosophy (winter 2016 edition)* https://plato.stanford.edu/archives/win2016/entries/actualism/.

Moore, T., & Haig, D. (1991). Genomic imprinting in mammalian development: A parental tug-of-war. *Trends in Genetics, 7*(2), 45–49.

Paley, W. (1809) Natural Theology. 12th edition.

Pinker, S. (1997). *How the mind works*. New York: Norton 2009 edition.

Queller, D., & Strassman, J. (2009). Beyond society: The evolution of organismality. *Philosophical Transactions of the Royal Society B, 364*, 3143–3155.

Wikipedia. List of cognitive biases. Accessed 2 Sep 2017. https://en.wikipedia.org/wiki/List_of_cognitive_biases.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the 6th International Congress of Genetics, 1*, 356–366.