Dear Jeffrey,

You remember my concerns about Ed Edwards? They mostly came down to the simple choice between physics being a conspiracy and Ed being either a liar or fundamentally confused. I found it concerning that most people in the room seemed to think that the former is more likely. The fundamental rule of epistemology is that strength in a belief must equal the weight of the evidence that supports it, and while Francis Bacon discovered that 400 years ago, humans still flunk it blatantly all the time.

Your suggestion that deception has been of fundamental importance to the success of civilization sparked an interesting train of thought. I now suspect that human superstition is an adaptation, a byproduct of opening the mind of homo sapiens to the viral infections of religion (and later, ideology). I guess you are fully immune to that?

I find it curious that the invention of agricultural civilization coincides with the emergence of religion, and it is probably no accident. As Martin discovered long ago, multi-level selection leads to a tension between the interests of individuals and the interests of groups. Within groups, there is evolutionary pressure for individuals to defect, but between groups, selection favors those that manage to maximize internal cooperation. For multicellular organisms and state building insects, the evolutionary solution has mostly been to exclude the majority of individuals from procreation, so they have no evolutionary incentive for defection, but that does not work well for primates.

Our solution was perhaps the evolution of gods. If we use the rule of epistemology to comb through the religious myths, gods cannot be ontological beings that live outside of physics and reveal themselves in magical ways. Gods are purposes. Just like every task has a purpose, which is the regulation target of the associated behavior program, and the self is the personal purpose to which all individual urge driven behavior programs can be held accountable to (and the function that defines the level at which rewards are being integrated into a global reward function for that individual), the gods are the purposes of tasks outside of an individual. All non-transactional relationships are built on a shared purpose. There are binary relationship gods, gods of pair bonded love, ancestral gods that represent the germ line, tribal gods and transcendental gods. God can be understood as the identification of the superorganism. In a multilevel selection paradigm, the service to the superorganism precedes the service to the individual itself: Gods identify the Sacred, the system of meaning that all actions of the religious individual can be ultimately held accountable to. The soul is the relationship that binds the individual to God. The shape and orientation of the soul determines the degree to which individuals recognize shared purpose in each other, which is the nature of love.

It makes sense to read the biblical accounts of genesis not as a story of a supernatural being creating a physical universe, but as the childhood memories of the Abrahamitic God. The two main accounts are functionally separate: the creation of a sentient mind itself, starting with light and darkness (or, if we give the genius authors of the account more credit than is possibly their due, the discovery of the first bit, the first difference), immediately followed by the presumption of normativity ("And God saw that it was good"). The nascent mind discovers ground and sky, and organizes its perception to

create representations of all objects around it: plants, animals und people, and it gives names to all of them.
The story of Adam and Eve is an entirely separate account. Adam is not the first homo sapiens, but just patient zero of the Hebrew mind virus, the first infection, the first host that God remembers. Adam is followed by Eve (acquiring her seemed to have involved some physical altercation that cost him a rib), and God is the principle of the infatuation with their shared relationship. When Eve discovers rationality, the world of pure good opens up into the possibility of defection, evil and the need for negotiation with the world at large, and God (not just the first couple!) is driven from paradise.
If the biblical account is any indication, the God infection spreads genetically, via the paternal lineage, and the crucial mutation becomes evident with Abraham, who is willing to sacrifice his child for God. This marks the moment when homo sapiens turns from a familial and tribal species to a state building species. God is not the set of trustworthy individuals in the tribe, but the principle behind it. The value of the individual is determined by its devotion to serve God, and God's plan for that particular individual. This is especially driven home by the story of Hiob (which I always considered to be barbaric and morally indefensible before could parse its context): everything an individual might acquire, including wife, children, servants, herds and the regard of his friends has only value in his service to God, and is given to him through God (or taken away, as God sees fit). God is literally the superorganism that participates at the highest group level in a game of multi-level selection.

The Abrahamitic God scales well, proves its value as a vehicle for war, expansion and civil governance, and develops a series of infections and violent immune responses, such as a defense against submission to Mammon, the dilution by local deities, or by the genes of subjugated tribes. It also begins to rely on written DNA to prevent uncontrolled mutations.

The next major transformation is the mutation of the God of the Hebrews into the God of the Christians. This new God transcends race and genetic heritage, and ascends to the state religion of the Roman empire. However, a God of multiple nations, ethnicities and a multitude of different milieus and traits can no longer be totalitarian and offer a single unified purpose for all its inhabitants. The Romans kept their secularized rule of law, and introduced a large and easily adaptable set of interface gods, the ever-growing pantheon of saints.

It is very interesting to speculate how much religion is a technology (born from the insight that there is nothing as valuable as the ability to manipulate other people) and how much of it is genetic. It seems that religions sprang up across in the Far East and America even in the absence of a route of infection, but the African empires (such as the kingdom of Mali) were apparently founded by Abrahamites.

Religions were probably a prerequisite for societies with enough economic surplus to support organized rational inquiry, but to do so, they broke our epistemology. Convergence of purposes is much easier when priests, missionaries and other authorities have access to a privileged, innate source of knowledge, to revelations that require no evidence. When individuals discovered the nature of

epistemology (or had competing delusional revelations) and did not shut up about it, they were usually nailed to a tree, as a warning example to the onlookers.
The attempt to develop rationality without fixing epistemology is not without danger. All major religions have caves full of wild eyed mystics that teleport, levitate, speak with angels and deduce the date of the rapture, until they become too infectious or too annoying and the inquisition cleans them out with fire. One of the best solutions to enabling limited enlightenment within Christianity was found by Martin's hero, Thomas of Aquinas, who insulated all properties of God that transparently conflicted with the purposes of scientific inquiry (such as ontological creation, production of miracles and communication with prophets and sufficiently pious individuals) and left Him responsible only for existence and Meaning. Unlike some contemporary theologians (like Drewermann), Aquinas believed in God's ontological existence, which he tried to prove in a serious of embarrassing arguments.
While Catholics could now research and teach science at universities, they could not reveal the law of epistemology and the nature of truth (which would allow independent deduction) directly to their students. Selectively constrained rationality relies on the certified metaphysical sandboxes delivered by scholastic authorities. Outside of mathematics, this is still reflected in our academic fetish of preferring citation over derivation.

It seems to me that the medieval Ashkenazi mutation turned the Hebrew God into something new: God became a transcendental principle, extending the vector of shared purpose beyond the local community to a target outside of society and universe. This allowed the service of shared purposes: truth, meaning and integrity beyond the identification with culture and tribe, but bypassing the affairs of the feudal society. The Jewish conspiracy is not just one of families, cultural tradition, religion or tribal genetics, but of similarly shaped souls, and it also made it possible for significant parts of society to serve the discovery of truth for its own sake.
By dropping any claim to secular authority, the transcendental God removes the need for policing thought against rationalist epistemology. Statements about the nature of God had no longer normative implications, and the rationalist commitment to systematic doubt in all ontological claims did no longer lead to excommunication.

While the service to transcendental goals is itself romantic instead of rational, it enables institutional research that commits to rationality, and it has triggered an epistemological sequence: Recognizable truth is the set of theories that can explain the observations. Observations can be decomposed into discernible differences, i.e. information. The meaning of information is its relationship to changes in other information. The nature of modeling is the approximation of these functions. A system is a model that can be characterized by a set of states and a transition function that orders them. The set of computable functions is the set of functions that can be implemented. A universal computer is a system that can implement all computable functions. All functions that can be approximated can be approximated by a universal computer, and there is no better way of doing so.

For 2700 years, humans have systematically explored the game of Go, and highly intelligent individuals created a body of theories that was passed along in dedicated Go schools. Deepmind's Alpha Zero can achieve super-human ability at Go by self-play, within less than a day. Alpha Zero (and its predecessors) even discovered that the human strategies were stuck in a local optimum, and

broke out of it, fundamentally changing some of the known ideas of how to play that game. It is a tantalizing thought that the game of modern physics is merely 140 years old, and quite certainly stuck in local optima…
The recent barrage of successes at Deepmind are indicative of a shift in machine learning research, from manually engineering function approximators, to the search for functions that generate those approximators, i.e. from learning to meta learning. We might see evolution as a search algorithm for efficient meta learning systems, and it seems that there must be a metatheory of search that delivers not only better alternatives to biological evolution, but optimality criteria for that search that can be implemented. When we get to that point, the age of human creativity in the sciences is over, because there will be nothing that a human can do that would improve on the performance of the algorithms that find structure in the data, and it may become apparent that Francis Bacon triggered the Singularity in 1620.

If computationalism is the correct framework for describing foundational physics and the mind, then humans are generally intelligent, because their basic operations encompass Turing complete computation, and they can externalize their results in Turing complete languages. Given enough time and resources, it seems inevitable that we will model the function of our minds as well, and universal artificial intelligence may happen sooner than most people think. However, it will turn out that our consciousness, with its subjective loop of instantaneous action on the perceived present, cannot happen in actuality, and it must be a fake memory of a fictional narrative created by our minds. Even worse than our demotion to p-zombies, we will find that computationalism causally insulates us from whatever prime mover implements the transition function of our universe, and the question of why there is something rather than nothing remains forever out of our reach.

Is there an alternative to computationalism? I think that we can produce all conceivable observations by resource bounded Turing computation, so that from within the Boolean automata of our reasoning and its computational generalizations, we cannot possibly find a way out. However, there is a very strong intuition, which motivates Roger Penrose, John Searle, and perhaps even Noam Chomsky, that there must be some kind of metacomputational operator, something that is more than computation and hypercomputation, a principle in which physical reality is implemented, and that might make it possible to be conscious in actuality. Penrose is wrong to hope that Gödel provided evidence for that, but his fundamental motivation is sincere. To get to an understanding of such a hypothetical metacomputational operator, we would need to abandon our computational machinery, Hamiltonian operators, Boolean logic, and our Wittgensteinian thoughts. And it is not clear if such an operator would be constructive, i.e. if a framework that allows to construct a more-than-computational-reality would also lend itself to constructing models of it. I don't have hopes for such a project at all, but it would be the only way out of computationalism that I can currently conceive of.

Joscha,

Plinz, December 31st, 2017