PURPOSE IN LIFE AND SCIENCE

To see the purpose of creation, look in a mirror

Problem

As I look at my face in the bathroom mirror, I see an uneven pulp of ageing flesh around two beady eyes. My stare becomes hypnotic and the flesh seems to melt. I see a spectral zombie face shimmering behind two black holes. The holes glaze over and I get dizzy. Who am I, where is this, why does it go on? My gaze shifts and the conventional answers swim back into focus.

When we contemplate the deeper questions of life, answers are hard to find. Philosophers have argued about them since time immemorial. Organized science weighs in with new answers, slicker and more complex than before but often unsatisfying. Science fiction author Douglas Adams imagined asking a giant computer the ultimate question of the meaning of life, the universe, and everything, and getting the answer forty-two. Something is missing.

A life without purpose is no life at all. This mantra, addressed to the bathroom mirror, homes in on my unease. Science is pretty good at telling me who I am and where this is, and I have learned to accept its answers with good grace, but science is weak on meaning and purpose. Philosophers gloss this by stressing the contrast between facts and values. Yet a sense of purpose and the values that express it are essential to the life we contemplate in the mirror. The problem is that they must come from somewhere. I cannot just invent a purpose and expect it to work for me any more than I can invent facts and expect scientists to be grateful for them.

I plan in this essay to tackle the problem indirectly by looking at how science frames the ultimate question and its possible answers. Maybe forty-two is not so hopeless after all. But if finding the purpose of life is our goal here, we had better start with a preliminary definition of purpose in its most general form. To have a purpose is to be oriented toward a projected future state that one may desire or intend and that may cause one to act in order to realize that state. The definition presupposes a reality in which an active subject with some form of both imagination and free will experiences a series of states of reality that are ordered in time.

Time

The problem of purpose starts in logic itself. If events are ordered in time, any purposes attached to my actions here and now are explicable in full only by reference to future events. And this leads to a failure of epistemic closure, where in principle I can never know enough about things to make watertight plans. For example, if my purpose now is to slake my thirst, I cannot know everything in advance about the causal nexus that determines how, in a range of hypothetical circumstances, my thirst will respond to my drinking some liquid from the kitchen. So my action will be a shot in the dark, made on the basis of a presumptive model of how things will turn out. If my mental model is wrong and I drink the wrong liquid, the consequences may be dire. My model may be constructed in accordance with the best methodology and so on, but still turn out to be quite false, and only a confrontation with reality can help me improve the situation. In the longer term, we shape our purposes by trial and error. But however we do so, our purposes are abstractions, shaped now on the basis of idealized models of future states that ignore various complications and entanglements that may easily invalidate them.

The problem of timing here is inexorable. We live in the present and have knowledge of the past. We use that knowledge to build models of the future, and shape our purposes accordingly. But the models are hypothetical. If we build them using the right facts and a deep enough appreciation of the laws and regularities that bind those facts into an intelligible unity, we can plan and act with confidence. Even then, unforeseen circumstances can intervene. Given the epistemic asymmetry of past and future, and our location in a moving present with a limited view, all our plans, purposes, promises, and prophecies are hostages to fortune.

Randomness

There is an irreducible randomness about the future. Our concept of randomness turns out to be so central to the story we tell in physics and biology that it is worth spelling out first, before we dive deeper. Our paradigm of a random process is throwing dice, where each of the numbers one through six has an equal probability of appearing on the uppermost face of a die, and no available theory enables us to predict which of the six will show up in a given throw.

To analyze a dice throw, we sort all possible future states into six groups and say the probability for the future state realized by the throw to be in any one of these groups is the same as for the other five groups, and is therefore one sixth, given that the probabilities all add up to one. This is a symmetry argument. The probabilities are the same because the groups are equivalent in all but spot count on the uppermost die face. The throw breaks that symmetry. Now one face is uppermost and the spot count is fixed. This group of outcomes now has probability one, and the probability of all the other outcomes, given this fact, collapses to zero. To confirm randomness, we can throw the dice a few more times and check that all the numbers one through six come up with frequencies that become more equal the longer we keep throwing. If the individual results are unpredictable, the series is random.

Examples like this of apparently random processes suggest that the best definition of a random series is that no exact description of such a series can be given that is shorter than simply listing its successive terms. In this sense, the series is algorithmically incompressible. By contrast, the digits in the decimal expansion of pi, for example, can be given by various simple mathematical formulas, so that series is algorithmically compressible and not random.

This definition appeals to a presumed impossibility that we may not always find plausible. Depending on the example, we may suspect that algorithmic compression is possible in principle and merely difficult or tiresome in practice. So we may allow a process to be practically random even if a deeper analysis could conceivably unearth algorithmic simplicity.

Entropy

Physics abounds with random processes and erects huge theoretical edifices on them. Any typical physical system, such as a glass of water, involves vast numbers of tiny parts, such as water molecules, whose combination in physically interesting configurations generates spaces of possible system states with even vaster numbers of inhabitants.

The dynamical laws that govern these huge configurations of particles are either deterministic or probabilistic. If the laws that govern the evolution of a system from one state at an earlier time to another state at a later time are deterministic, then from an exact specification of the earlier state we can compute an exact specification of the later state, and if we are fast enough we can predict system states. If the laws that govern the evolution of a system are probabilistic, the best we can do is to predict probabilities for the various possible outcomes as a system evolves from an earlier to a later state.

Classical mechanics is deterministic in principle, though not always in practice, since inaccurate measurements in systems that admit chaotic behavior may limit predictions very severely, so that such systems show effectively random behavior. If I spill some water, the splash is random for all practical purposes.

Quantum mechanics is deterministic in principle, assuming that systems can be in superpositions of several states at once, though hardly ever deterministic in practice, since any measurements or interactions at all disturb a system and cause it to change its state randomly. If I measure the thermal motion of an atom in a water molecule, whether it jiggles this way or that is random. The difference between classical and quantum mechanics is that what seems like a mere practical constraint in classical systems is a constraint of principle, with no hint of a way out, in quantum systems.

The second law of thermodynamics says that as any closed system evolves in time, its entropy tends to increase. That is, if the overall configuration of the state of a system has a certain probability, relative to a set of states including both that state and any states into which it may evolve, then after a lapse of time, the later state of the system has an overall configuration with at least as high a probability. What this means is that a system tends to evolve along a path through a given state space in the direction of increasingly probable configurations. This law is well confirmed, fundamental, and universal. A dropped glass shatters, ice melts in warm water, and we all die in the end.

The law also works in reverse, for retrodiction as well as prediction. Given a state of a system at some later time, the temporal symmetry of the fundamental dynamical laws dictates that retrodicting an earlier state of the system in the state space should also find an increase in entropy. So given a partially melted ice cube in a drink, we can retrodict that it was even more melted a few minutes earlier. Failing additional assumptions about the earlier state of the drink, this is a valid retrodiction. In general, given only a first state and some dynamical laws, any second state is probably going to be at least as probable as the first. Temporal asymmetry appears only when we locate special events in the past, such as the big bang or putting ice in the glass.

Evolution

Entropy applies to closed systems, but living systems are open. A living system is a process that feeds on inputs with low entropy and produces outputs with high entropy. In more concrete terms, living processes convert things with rather special or improbable forms, such as foodstuffs, into a few things with even more improbable forms, such as fresh eggs, and a lot of things with rather ordinary forms, such as waste. The more ordinary or probable a thing, the more randomness is involved in the assembly of its ingredients.

Randomness is minimal in genetic material. Information is negentropy and genes are compressed packages of information. Genes are coded recipes for constructing convoluted protein molecules. These molecules work together as nanomachines and ultimately as organisms. So genes carry instructions for replicating organisms. But despite the most elaborate copying and quality control mechanisms, randomness infects genes too, to create mutations. Some mutations are preserved by natural selection and the result, over sufficiently many generations, is the evolution of new species. The evolution of species may be seen as the exploration of ever more complicated and indirect, yet effective, ways for genes to project copies of themselves into the future. In the process, information is concentrated and randomness is wrung out more and more thoroughly.

Humans have developed a new way to accelerate biological evolution by genetic engineering. The human genome is rather less than a gigabyte in size and replicates humans about once or twice per gigasecond. By contrast, file transfer over a broadband line can replicate gigabyte files instantly, and digital technology ensures arbitrarily high copying fidelity. Using such technologies, we are learning in this century to project organized information into the future with unprecedented speed, volume, and efficiency.

As to why this process of concentrating and replicating information should persist and intensify, we can see that the more efficient replicators crowd out the competition. Organisms are gene machines, and the genes replicate as freely as nature allows. By analogy, human brains are meme machines, where memes, roughly speaking, are ideas that prompt us to copy them. Memes replicate as freely as human nature allows. As Internet pioneers used to say, information wants to be free. This process goes on all the time, within us as well as around us. Purpose is still a puzzle.

Free Will

In the brain, the electrical firing of neurons conforms to physical laws. Changing concentrations of neuromodulator and neurotransmitter molecules cause neurons to send tiny electrical signals to each other. Over the cortex, billions of neurons create a vibrating field of electrical energy that carries the symphonic states we experience in consciousness. Cerebral neurodynamics are more complex than anything we can yet model accurately, but we have no reason to doubt that known physical laws can in principle explain everything of functional significance that goes on in our brains.

A cerebral neuronet is a physical structure that realizes a succession of states in time. In principle, we can use information about a particular state plus physical laws to predict its evolution in time. Mention of evolution here is quite fitting. Each state of the system propagates in a lawlike way to a successor state, with occasional random variations. And neural groups in the brain compete with each other for resources and for opportunities to fire and control their neighbors. Organized patterns of neural activity arise through a process of natural selection much like that governing simple organisms in a nutrient medium.

Seen from within, the activity in human brains may appear to be regulated and even dominated by plans and purposes, but from a laboratory perspective our thoughts form and reform like clouds in the sky to give at best a self-generated illusion of purpose. As neural groups in our brains interact and get organized, stable cognitive structures evolve that computer scientists call demons. A ceaseless pandemonium in our heads creates a flickering background of thoughts and images and ideas, from which we make up our own minds, get ourselves together, and exercise what we like to regard as free will, in an astonishingly complex process of self-organization. Almost all of this action in our brains is unconscious, and it is only from within that we can even begin to imagine that our conscious thoughts and choices somehow govern the rest. Scientists who have studied this agree that there is a good deal of illusion in our thoughts and opinions, both about the ongoing drama of our own selves and about our presence and influence in the external world, and they disagree only about the depth and incorrigibility of the illusion. So opinions differ on whether this author is free to go and fetch a glass of water.

All this seems to marginalize the concept of purpose. We can easily enough imagine a purpose behind the process of life in general, just as we can describe our own thoughts in terms of purpose. We can even enjoy a striking sense of explanatory power when we relate purposes to human passions and the tumult of our everyday lives. Nevertheless, the concept appears to be superfluous in science. We expect in principle to be able to translate without loss any references to purpose in our lives and thoughts into neutral scientific terms. These terms conform to the simple logic that past states plus universal laws are enough to predict in principle and thus far explain future states. The question is whether purpose is now redundant.

Purpose

If purpose in life or mind is no more real than cosmic purpose, then perhaps cosmic purpose is no less real than my purpose in fetching a glass of water. We can see the dynamic of reality as a whole as being to propagate ever more efficiently packaged and organized information into future states of reality. From the big bang to the present, the entire evolution of the universe has been in the direction of generating increasingly concentrated and convoluted packages of information that persist in what looks like an increasingly randomized environment. We can indeed imagine a purpose of sorts here, but it is abstract and empty, with no hint of a mechanism for foresight or control. It certainly seems a far cry from human purposes.

Consider how we use the concept of purpose in everyday life. I envisage a desired future state of myself and the world and I constrain my actions accordingly with the purpose of realizing that goal. This involves a series of functional mechanisms at the neural, organismic, and societal levels that are explicable within the frame we have reviewed so far. Such purposes find a place within the scientific vision of a strictly mathematical universe in which states unfold in time with an impersonal majesty that dwarfs our fleeting lives.

A problem with the majestic scientific vision is that we who conceived it are humans. People with pandemonium in their heads somehow tamed their demons sufficiently to develop mathematics and science far enough to enable them to see the vision. There is a bootstrap process here that demands a closer look. We need to understand how naturally evolved survival machines built up from protein macromolecules can grasp the fundamental principles that purportedly govern the entirety of natural creation.

The bootstrap process not only gave rise to science but is recapitulated in ontogeny. A baby has a brain that sees phenomenal surfaces vividly but understands very little. Soon mother and family emerge, and the story of purpose gathers momentum. As the child grows to adulthood, ever more of the world and its history makes sense in a cosmic drama with the self in the starring role. Then, with advancing age, the cold majesty of the big vision sinks in, and the passions of life, as well as its purpose, begin to seem as spectral as the zombie in the bathroom mirror. The sense of purpose first waxes, then wanes. As time goes by, state after state of reality unfolds, and each state replaces its predecessor as the objective setting for the ongoing realization of the unifying subject.

An anthropic principle holds sway here. On both the personal and societal levels, we say that creation unfolded as it did because we, here and now, are the ones who are looking back and reconstructing it. As a matter of logic, we need to explain our own existence. So purpose exists because we see it, and we are part of creation.

Solution

We have a story in which we are free to say that purpose permeates the whole of creation. But to stand at the summit of such a story, we need to be worthy of it. The idea is absurd that I, your humble author, about to go and fetch a glass of water, sit at the summit and pinnacle of cosmic destiny.

We can do better. Our civilization has begun to project a generic image of our collective self. The scientific consensus on our shared universe, our shared biological ancestry, our shared genome, our shared neurodynamics, our collective immersion in the same meme bath, and our common planetary destiny, adds up to a portrait of an envelope self. Each of us, until we die, pushes out the envelope in some tiny way and thus enriches this growing self. Everything in our universe leads up to this potentially infinite living being and finds there a vision of its ongoing purpose.

Long ago, the royal we was invoked to denote the universal ambition of the enlightened monarch. Now, in science, we can assert something similar. We are the subject of cosmic history, the purpose of creation. The innermost self of each human being is bathed in the light from the radiant sun of science. Purpose is a concept we must deploy to suit our passion, and if our passion is to explain the universe, then our purpose is realized when the universe makes sense. As our passion goes out to loved ones, they share our purpose in the epic of creation. Each of us has a cosmic soul. We glimpse this soul when we sense the sense of creation. Purpose is immanent in creation, and is realized when we find the harmony of all the states and parts that make it up.

If this sounds overblown, consider my purpose in fetching a drink. To realize that this is indeed my purpose, I need only consider the harmonious symphony of actions I succeed in undertaking to fetch a glass, pick up the pitcher, pour the water, and carry the glass to my desk. There, I did it! This miraculous interplay of neurons and sinews, not to mention the synchrony of expectation and reality as the glass and water behave as my mental model predicts, is as glorious as the entire cosmic drama, just smaller. Even that zombie in the mirror has a purpose, to boot the brain into action. As I see it, purpose is what makes life worth living.

Andy Ross
May 2004

Andy Ross lives in Germany and works as a developer for the international software company SAP. Earlier he worked as a physics and computer science editor at Springer-Verlag in Heidelberg, Germany.

Andy is British and his roots are in southern England. Before moving to Germany, he worked as a tutor in logic and philosophy at the University of Oxford, as a teacher of English in Japan, and as a teacher of mathematics and physics in London. He holds four degrees in philosophy, three from the University of Oxford and one from the University of London.

Andy has authored two books: a 1996 science-fiction novel entitled *Lifeball* and a 2004 collection of scientific papers on consciousness entitled *Mindworlds*. The full texts of both are available at his homepage www.andyross.net