The third replicator (title to be discussed)

For *New Scientist* May 2009

By Susan Blackmore

We humans have let loose something extraordinary on our planet – a new replicator – and the consequences may be dangerous and are certainly unpredictable. Born of Earth's first replicator, genes, our species has already let out the second replicator, memes, and I believe that what we are now seeing, happening all around us in a vast technological explosion, is the birth of a third. We are Earth's Pandoran species. We are watching our home turn into an R3 planet and we are blissfully oblivious to what we have let out of its box.

This is how the cosmos looks through the eyes of Universal Darwinism – the principle that all evolutionary processes basically work the same way. Given some kind of copying machinery that makes lots of slightly different copies of the same information, and given that only a few of those copies survive to be copied again, then design must appear. This is the evolutionary algorithm that creates design out of destruction. You might call it "design by death" since clever designs thrive because of the many, many failures that don't.

This process is well understood when applied to biology, but the point of Universal Darwinism is to apply it wherever any information is copied with variation and selection. The copied information is the "replicator". So this is how we come to ask about the number and type of replicators on this or any other planet. An R1 planet has life; an R2 planet has culture; an R3 planet can reach out beyond itself.

Here on Earth our first replicator is fairly uncontroversial. It is genes, or the information coded in molecules of DNA. Or perhaps I should say that this is the first replicator we know about. We do not know precisely what came before DNA but it was probably some kind of "naked replicator", a simple self-replicating molecule that copied itself by using available resources from its environment; perhaps it was RNA or some other relative. This was eventually co-opted as part of the cellular copying machinery inside living things. The replicator itself is stored in stable molecules of DNA, is copied with exquisite precision, and used to build vehicles or interactors that carry it around, protect and propagate it. These vehicles – the gene machines, or lumbering robots, as Richard Dawkins likes to call them – are our familiar plants and animals, the wonderfully prolific and diverse creatures of the first replicator.

About four billion years later, something extraordinary happened. Members of a particularly sociable species of bipedal ape began to imitate each other. Imitation is a kind of copying, and so a new evolutionary process was born. Instead of cellular chemistry copying the order of bases on DNA, an ape was using its big brain to copy a new kind of information - gestures, sounds, and other behaviours. Each individual was stealing the products of another's learning and passing it on again. This might not have been very accurate copying, and even now there are no other creatures on earth that can imitate the way humans do, but it was enough to start a new evolutionary process. Dawkins called the new replicator "memes". A living ape, that had been a vehicle of the first replicator, was now the copying machinery for the next.

The idea of memes as a cultural analogue of genes has been much maligned, and probably the majority of biologists still reject it. Yet memetics has much to offer in explaining human nature. According to meme theory, humans are radically different from all other species because they alone can imitate; they alone are meme machines. Human intelligence is not just a bit more or a bit better than other kinds of intelligence. It is something completely different, based on a new evolutionary process and a new kind of information.

Most biologists disagree. They work on the assumption that human culture and language evolved because they were adaptive for human genes; they assume that, as sociobiologist E.O. Wilson put it, the genes must always keep culture on a leash. Memetics throws out that assumption. Yes, imitation itself must once have been adaptive for the apes who started it, but evolution has no foresight. There was no intelligent designer there to say "Eh, imitation is a kind of copying – watch out what you're letting loose".

So the memes began to proliferate. What began as an adaptation for the apes soon became an unanticipated parasite – a new evolving creature that changed them and their world forever. Their brains were redesigned by the pressure of having to copy ever more, and ever more complex, memes. The idea of memetic drive is that once memes were proliferating, individuals benefited from copying the latest and most successful memes, and then passed on any genes that helped them do so. So brains were forced to get bigger and bigger, and to become especially good at copying the most successful memes, leading to the peculiar human penchants for language, art, music, ritual and religion.

This process was dangerous. Small brains are much more efficient if you don't have to copy anything, but once memes are around you cannot survive unless you do. So brains had to get bigger, and big brains are expensive to produce, dangerous to give birth to, and expensive to run. Then there is the question of what is copied. If you start copying anything at all then you might copy dangerous memes, like throwing yourself off a cliff or using up all your resources in pointless rituals. This creates an arms race between two selfish replicators – memes benefiting from brains that copy anything and everything; genes benefiting from brains that are smaller, more efficient and highly selective.

Either of these dangers might have finished our ancestors off, but we pulled through. The result was a compromise, with human brains being just about as big as our bodies can stand, and yet selective enough to avoid copying too many lethal memes – martyrdom, cigarette smoking and BASE jumping being among the exceptions. In the same way that a disease pathogen tends to co-evolve with its host to become less lethal, so memes co-evolved with their gene-based hosts. Languages, religions, skills and fashions that began as parasites turned into symbionts. Not only do we get along with our memes now, but we could not live without them.

There was also a cost to the rest of life on earth. Wherever they went humans took new techniques with them, spreading agriculture and changing the landscape, obliterating Earth's largest mammals, domesticating others, and changing whole ecosystems. Then much more recently they began to build radically new kinds of technology and the changes they effected dwarfed anything that had gone before. Was this just more of the same or something new?

In all my previous work in memetics I have used the term ‘meme’ to apply to any information that is copied between people, including stories in books, ideas embodied in new technology or websites stored in computers and called up by other computers half a world away. The reason I called them all “memes” was not just that other people did so, but that there seemed no principled way of distinguishing between what we might call ‘artificial’ memes such as websites and high-tech goods, and ‘natural’ human memes such as spoken words, habits, fashions, art and religions. So on the grounds that a false distinction is worse than none I stuck to the term ‘meme’. Yet an email encrypted in digital code, broken into tiny packets, and beamed around the planet does seem different from someone shaking hands and saying “Hi”. Could there be a fundamental principle lurking here? If we ask what make memes count as a different replicator from genes, would that help us decide what would make something else count as a different replicator from memes?

Put that way, the answers are easier to see. Memes are a new kind of information (behaviours rather than DNA), copied by a new kind of machinery (brains rather than chemicals inside cells). This is a new evolutionary process because all of the three critical stages – copying, varying and selection – are done by those brains. So now we can turn to all that new technology and ask whether the same applies.

Electronically processed binary information is of a different kind from imitated behaviours, skills, stories or fashions: computers and servers are quite different in how they operate from living brains. There really is a new evolutionary process because all three critical stages are carried out by machinery other than human brains. The answer is “yes”.

Think of Google. This familiar helper is a virtual machine running on digital technology and constructed out of numerous programs written in binary code. Google uses masses of stored digital information, searches through it, selects what it needs and then displays the selections in new combinations. Although originally designed by people to help other people search the web, it increasingly operates without human intervention, and the results of its searches go not only to computer screens for people to look at, but to other programs, commercial applications and even viral software.

Or think of programs that write original poetry or cobble together new student essays, or programs that store information about your shopping preferences and suggest books or clothes you might like next. Or think about Second Life, a virtual world run on multiple computers in which avatars live virtual lives in virtual homes with virtual money. These, and many other recent inventions, behave as Darwinian machines in their own right.

So we seem to have a third replicator in our midst, but how did it creep up unnoticed? Looking back through the history of technology we now know what to look out for – meme vehicles that are not just passive carriers or storage spaces but that copy, vary or select information as well. Cave paintings or modern art can be copied but they don’t do any copying themselves. Nor do houses, tools, bridges, money, scientific theories or the Tower of London. Trains, cars and planes are complex machines and they spread memes around by carrying people, but they don’t replicate themselves or anything else. Writing was a great step forward because words could be stored safely and copied, but it was still human hands that did the copying – until the advent of the printing press. Here was a machine that actually did the copying itself. All right, it required people to run it, and to select what to print and how many copies, but it was a first step in the process that ended up converting written words into binary code. After

that everything became more flexible. Texts could be moved around, copied, varied and selected, by telegraph, phone, fax, and ultimately by computer. Then came more linking up, wider networks, the Internet, and finally – unpredicted and unexpected – the World Wide Web and the massive opportunities for further evolution that this provided.

All this makes perfect sense through the eyes of Universal Darwinism, but to see it this way we have to throw off the self-centred idea that we designed all this stuff for our own benefit. Instead we have to step back and see it as yet another evolutionary process playing out its inevitable game of design by death. The accelerating expansion, the increasing complexity, and the improving interconnectivity of all three replicators are signs of the same fundamental design process driving them all. Road networks look like vascular systems, and they look like computer networks because interconnected systems do better than isolated systems. The Web today connects billions of computers in trillions of ways just as a human brain connects billions of neurons in trillions of ways. The former is limited by the size of a human skull, while the latter is limited only by the resources of a whole planet, but their uncanny resemblance is because they are doing a similar job.

If we are not the original designers we like to think we are, what is our true role and where do we go from here? We humans were vehicles for the first replicator, copying machinery for the second, and what for the third? We seem to have handed over most of the storage and copying to our new machines, but we still do much of the selection which is why the web is so full of sex, drugs, food, music and entertainment. We also run the power stations, build the factories that make the computers and repair things when they go wrong. Maybe we will be gradually incorporated into vast networks as their energy suppliers, much as free-living bacteria are thought to have become incorporated into living cells as energy-providing mitochondria. Maybe we will merge with the new machines, having brain implants and super-fast embedded connectors so that all our brains become nodes in an even larger machine which we won't be able to live without. Maybe we won't survive at all.

I said that the birth of a new replicator means a dangerous tipping point. Early life nearly destroyed itself by producing the poisonous gas, oxygen, but it pulled through and now cannot live without it. Our ancestors might have killed themselves off with their over-large brains and dangerous memes, but they pulled through and now cannot live without language and culture. This time the danger is to the whole planet. Gadgets such as mobile phones and personal computers are already using 15% of household power and rising (NS 23 May p 17); the web is using over 5% of the world's entire power and rising. We blame ourselves for climate change and depletion of Earth's resources, but perhaps we should blame this new evolutionary process that is greedy, selfish, and utterly blind to the consequences of its own expansion. We at least have the advantage that we can understand what is happening and that must be the first step towards working out what, if anything, to do about it.

BOXES

Updating Drake

We are able to ask the question "Are we alone in the universe?" because our ancestors created memes, turning Earth into an R2 planet, rich in language and culture. We are

able to contemplate communicating with other worlds because Earth is fast becoming an R3 planet, rich in digital technology that passes information around at the speed of light and could potentially send it out far into the galaxy. How many other planets have taken a similar course? And why haven't we heard from them yet?

Traditionally the search for extraterrestrial intelligence, or SETI, has looked for intelligence. In 1961 Frank Drake proposed his famous equation for estimating the number of intelligent civilisations capable of communicating with us in our own galaxy. It includes the rate of star formation, the fraction of stars with planets, the fraction of planets that can sustain life and the fraction that get intelligent life and then technology.

But perhaps intelligence and civilisation are not what we should be concentrating on. My analysis based on Universal Darwinism suggests that instead we should be looking for R3 planets. The number of those in our galaxy will depend on the probability of a planet getting a stable first replicator, then a second, and then a third. Maybe each step is hard, or maybe each is easy but dangerous. This new and simpler equation won't tell us the answers but by posing new questions it may help us understand why – so far – we have not heard from anyone else out there.

This box does not need the equations but if you want them Drake's is well known. Mine is

$N = n \times f_{R1} \times f_{R2} \times f_{R3} \times L$

where N = the number of intelligent civilisations in our galaxy capable of communicating with us

n = the number of planets in our galaxy

f_{R1} = the fraction of planets in n where a first level replicator survives

f_{R2} = the fraction of planets with R1 where a second level replicator survives

f_{R3} = the fraction of planets with R2 where a third level replicator survives

L = the fraction of a planet's life for which a third level replicator persists

Help me find a name for the third replicator.

There are genes, memes and now what? If there truly is a third replicator already spreading greedily through our computer networks and cyberworlds, what should we call it?

In a lecture I gave to TED2008 I talked about “techno-memes”, but since I was claiming these are not just another kind of meme but are something new, I wanted a name without “meme” in. So I called them “temes”.

“Genes, memes and temes” trips off the tongue easily enough, and I like the resonance, but it’s too easy to confuse “temes” with “teams” or even “teems”. Indeed the very next day a Wired article said I’d been talking about “teams”.

Can you come up with a better name? I put out a call on my website and have had about 50 emails with a total of 30 different suggestions. Two Spanish speakers pointed out that in Spanish “teme” means “afraid” or “have fear” which they seemed to find thoroughly appropriate.

Here are just some of the suggestions people sent me. Which do you think work best, or have you any better ideas?

Syntheme (for synthetic memes)

Macheme or Machime or just Cheme or Chene (for machine memes)

Softweme (for software memes)

Treme (for tertiary ene or eme) or Thirdeme

Systeme (because of the large systems involved)

Geme (for God-like meme)

Sele (to rhyme with feel, and relate to silicon)

Technemes (a longer abbreviation from technological memes), along with Techmes, Tmemes and Techmemes

Sheens (from machines, and it’s shiny)

T-reps (inspiring the trepidation of T-rex)

Esemes (for “ex silico memes)

Bytemes (combining byte and teme)

Byne or Bine (from binary, bits and bytes)

Tecplicator (from technical replicator)

Arteme or just eme or eeme (for artifical meme)

Goleme (from the artificial creature, the Golem)

Autome (because they are becoming autonomous)

Iteme (from IT)

C/meme (for computational memes)