Dangerous Memes; or, What the Pandorans let loose

Susan Blackmore

In *Cosmos and Culture: Cultural Evolution in a Cosmic Context*
Ed. Steven Dick and Mark Lupisella, NASA 297-318, 2010

Note this is the version submitted to the editors in December 2007. It may have been edited before publication.

Cultural evolution is a dangerous child for any species to let loose on its world. And the parent species, whatever it is like and wherever it arises, will have no insight into what it has done until its offspring is already grown and making its way in the world. By then it is too late to take it back. So I shall call this motherly species 'Pandoran', after the mythical first Greek woman whose box released all the evils of mankind. We humans are Earth's Pandorans, and have let loose cultural evolution, but on other planets quite different creatures might be playing this role.

Opening such a box of tricks can even be lethal, and I suspect that there are several danger points. The first critical step occurs when one species becomes capable of behavioural imitation, or of some other process that makes copying with variation and selection possible. This creates a new replicator, making the evolution of culture inevitable. This is the first danger point, because the newly created culture – the spreading of copied behaviours and the competition to mix, match and make more – can get out of hand. Some of the behaviours may be so extravagant, or expensive, or dangerous, that they kill off their Pandorans and so obliterate themselves as well. This kind of waste is all part of how evolution works. Indeed natural selection might be called "design by death" because of all the billions of creations that have to die in spawning innovation and success for a few.

If this first danger point is passed, the Pandorans and their newly spawned culture may begin to adapt to each other, and coevolve towards a more symbiotic relationship, as diseases and their hosts sometimes do. If this succeeds, the result may be a stable mutualism that lasts indefinitely. Alternatively, with enough time, and under the right conditions, another step might be taken. That is, new mechanisms for copying, varying and selecting information could evolve outside of the Pandorans themselves, leading to a second danger point. For example, here on earth, humans invented printing, sound recording and photography, vast communication networks, broadcasting and the Internet. These are all methods of selective copying which means a new evolutionary step, and this creates a second danger point. As the copying increases, the thirst for innovation that's unleashed can be a drain not only on the Pandorans who started it but on their whole environment as well. This is what has happened here on earth, with the consequent overpopulation and technological explosion threatening the health and climate of the entire planet.

This danger point could also be safely passed, or it might prove fatal. We don't yet know what the outcome will be here on earth; it might go either way. However, our sample of one planet does at least allow us to think about the general picture and speculate about what might happen on other planets elsewhere in the cosmos.

I like to imagine a vast universe containing many planets which have conditions suitable for life to evolve. On some of those planets a species evolves that is capable of copying what others do, so unleashing this second evolutionary process. Among those planets, some survive the danger point and some do not, with the successful ones going on to spawn further evolutionary steps and face further danger points. On this picture, what should we expect to see around us? I would like to explore what might be out there on the basis of this memetic way of thinking about cultural evolution. I shall first explain a little about memes, meme theory and the importance of replicators, and then consider some of the possible fates of planets that give birth to multiple replicators as ours has done.

The science of memes

Memetics is rooted in universal Darwinism; that is the idea that natural selection is a general process of which earthly biology is just one example.

Working from his detailed observations of living things, Darwin saw what very few people had ever seen before, even though the process is always staring us right in the face. That is, if creatures vary, and if they have to compete for resources so that most of the variants die, and if the successful variants pass on to their offspring whatever it is that helped them survive, then the offspring must be better adapted to the environment in which all this happened than their parents were. Repeat that cycle of copying, varying and selecting, and design must appear out of nowhere.

My favourite word in that description is “must”. This “must” is what makes Darwin’s insight the most beautiful in all of science. You take a simple three step algorithm and find that the emergence of design for function is inevitable. Dan Dennett calls it “a scheme for creating Design out of Chaos without the aid of Mind” (Dennett 1995 p 50). This is “Darwin’s Dangerous Idea” that the algorithmic level is the level that best accounts for the wonders of nature; that all the fantastic and beautiful creatures in the world are produced by lots and lots of tiny steps in a mindless and mechanical algorithm.

The whole process can look like magic – like getting something for nothing – but it isn’t. It is not possible to get matter out of nowhere, but it is possible to get information, or new patterns of matter, apparently out of nowhere by making copies. If the copies vary slightly and not all the copies survive, then the survivors must have something that helped them win the competition – using Darwin’s term they are more “fit”; they make a better fit to their environment. Then they pass on this advantage to the next generation of copies. And so it goes on.

This is the fundamental idea that Richard Dawkins explained in his 1976 book The Selfish Gene. He emphasised the importance of thinking about evolution in terms of information rather than squishy living creatures, and he called the information that is copied the replicator. In fact “replicator” is not a very good name, implying that it is the thing that does the replicating rather than being the thing that is replicated (perhaps ‘replicatee’ would be better) but ‘replicator’ is what it is called and I will stick with that here; the concept is more important than the name.

For Dawkins it was a general law, “the law that all life evolves by the differential survival of replicating entities.” (1989 p 192); a “view of life …that applies to living things everywhere in the universe. The fundamental unit, the prime mover of all life, is the replicator.” (1989 p 264)”. With this view he wanted “to claim almost limitless power for slightly inaccurate self-replicating entities, once they arise anywhere in the universe.” (1989 p 322).

From the perspective of this general law, genes are only one example of a replicator. So we might expect them to show both characteristics that are true of all replicators everywhere in the universe, and features that depend on the idiosyncrasies of evolution here on earth. Dawkins wanted to explore the general principles as well as the specifics. Indeed he became quite frustrated with the way his colleagues tended to think about evolution as though it were inevitably and always a matter of genes. So at the end of the book he asked his now famous question “do we have to go to distant worlds to find other kinds of replicator and other, consequent, kinds of evolution?” (1989 p 192). His answer was, of course, no. Staring us in the face, “still drifting clumsily about in its primeval soup” (1989 p 192), is a new replicator; tunes, ideas, catch-phrases, fashions, and ways of making things are all spread from person to person by imitation. They vary and they are selected. These are the new replicators; the memes.

There is nothing mythical or hypothetical about memes, and this point is frequently misunderstood. Some people seem to imagine that memes are some kind of abstract entity that might, or might not, live inside brains; or that might, or might not, really exist (Aunger 2000). This is to miss the point that everything we copy in culture is a meme, by definition. So the words I am writing now, the picture on the cover of this book, the practice of making books like this, and making paper and ink and glue this way, are all memes. As Dennett (2006) points out, there is nothing mythical about words printed on a page, or even spoken words or recorded and broadcast words. Indeed Darwin referred to "The survival or preservation of certain favoured words in the struggle for existence" as a kind of natural selection (Darwin 1871, p 61). There is nothing hypothetical about kids listening to i-PODs, wearing pre-torn jeans, or putting pieces of metal through their ears and noses after seeing others do the same. There is no question about the existence of financial institutions, money, railways, bicycles, telephones, furniture, skyscrapers, holiday brochures, football, or the days of the week. They are all information; they are all encoded in some kind of matter and energy, and they can all be copied or not. So think of memes this way. Their core definition is "that which is imitated" or that which is copied. They are all around us.

There is no question, then, that memes exist, for unless you deny that anything is ever copied from person to person, they must. The real question is whether thinking about culture in terms of replicators and memes is useful or not (Laland and Odling-Smee 2000). I am convinced that it is, or at least that it can be. Already, considerable progress has been made in identifying the problems, clarifying the issues, and beginning some empirical investigations. For example, there have been constructive arguments about how to define memes (Aunger 2000, Blackmore 2001), whether they should be thought of as units or not (Sperber 2006, Wimsatt 1999) and whether they should be considered as inside or outside of human brains (Aunger 2000, Benzon 1996, Distin 2005). Particularly interesting areas of debate concern the extent to which behaviours are reconstructed rather than imitated (Jablonka and Lamb 2005, Sperber 2006) and whether memes are truly replicators or not (Blackmore 2001, Richerson and Boyd 2005). Then there are empirical investigations of, for example, replicating text (Pocklington and Best 1997) and memes in translation (Wright 2000), applications in sociology (Runciman 1998) and musicology (Jan 2007), and attempts to simulate and model memetic processes (Bull et al 2000, Higgs 2000, Kendall and Laland 2000).

These are all promising signs but it remains true to say that there is currently no thriving science of memetics. It is, at best, a tiny area with a few enthusiastic proponents. The reasons for this are not clear. Some people seem simply to be afraid of the whole idea of memes. This baffles me. When I gave my first ever lecture on memes at the London School of Economics in 1996, I mentioned that the word "meme" is often printed with scare quotes, and that I had even seen lecturers putting up their hands up round their ears in stylised scare quote fashion when daring to mouth the word "meme". More than a decade later I am sorry to find that I am still seeing this bizarre behaviour. Why? Is memetics really so scary? Possibly it is. Among the ideas that upset people are that all 'our' ideas are recombinations and adaptations of other people's, that all creativity comes from the evolutionary algorithm and not from the magic of human consciousness (Blackmore 2007a, Chater 2005), that our inner conscious selves may be memeplexes created by and for the memes (Blackmore 1999), that free will is an illusion, that modern computing technology is creating itself using us, and that the process of memetic evolution is not under our control (Blackmore 1999, Dennett 1995). These implications may be scary, but it is far from certain that they do all follow from a memetic view of evolution. In any case, being frightening is not a good test of the truth or falsity of a theory.

Among other reasons, some social scientists brand all evolutionary approaches to their subject "reductionist", and reject memetics along with sociobiology and evolutionary psychology (Bloch 2000) while some biologists want to keep Darwinism exclusively for their own field and so reject its application to culture (Rose and Rose 2000).

More interesting here are those who accept that culture evolves but still reject memetics. This may be either because they reject the whole idea of replicators as the driving force in evolution, or because they deny that memes are replicators. Wilson famously argued that “genetic natural selection operates in such a way as to keep culture on a leash.” (Lumsden and Wilson 1981 p 13). On his theory and related theories of gene-culture coevolution, the cultural variants (culturgens) are not independent evolving entities but are part of the human phenotype and are an adaptation that benefits human genes.

Branching out from purely genetic evolution, Jablonka and Lamb (2005) explore epigenetic and other evolutionary systems, but they do not accept that “the dreaded memes” (p 224) are replicators, and claim memetics is “seductively simple” but flawed (p 208). Richerson and Boyd (2005), whose theory of cultural evolution is arguably closest to memetics, do sometimes use the term ‘meme’ but they clearly state that “cultural variants are not replicators”: “Culture is on a leash, all right”, even if the dog on the end is big and clever, and “Culture is an adaptation”. In other words, culture was adaptive for human genes, it evolved for that reason, and it has persisted for that reason – in spite of including some maladaptive elements. In this respect they still illustrate Dawkins’ complaint about his 1970s colleagues that “In the last analysis they wish always to go back to ‘biological advantage’.” (Dawkins 1976, p 193).

This reveals the fundamental difference between all other theories of gene-culture coevolution and memetics: for the former the final arbiter is genetic advantage so that culture must always remain on its leash, even if the leash sometimes gets very loose; for memetics both genes and memes have replicator power and can drive change and creativity. For other theories cultural traits are an aspect of the human phenotype, but for memetics they are living things in their own right.

These include a vast range of memeplexes (co-adapted complexes of memes) that are copied, stored and propagated by their human hosts, using a wide variety of adaptations. All these, according to Humphrey, “should be regarded as living structures, not just metaphorically but technically.” (Dawkins 1989 p 192).

Some survive predominantly because they are useful to their hosts (e.g. effective financial institutions, scientific theories, or useful technologies); others depend on fulfilling human desires and preferences (e.g. the arts, music, and literature); and still others are positively harmful, tricking their hosts into propagating them. We humans are selective imitation devices (Blackmore 2001); we try to select only the useful or valuable memes but are inevitably tricked by some of the rest.

Popular examples of the tricksters are informational viruses, such as chain letters, email viruses and pyramid schemes. These reveal the classic viral structure; an instruction to copy the whole memeplex, backed up with threats and promises to ensure compliance. Dawkins (1993) pointed out that the major religions of the world have just this structure. Billions of people all over the world are infected with a religion at an early age when they have little memetic immunity, usually by their own parents whom they love and trust. They then spend the rest of their lives paying the price of adherence to false beliefs, and in turn infect others. Thus we can see the whole history of religions as an evolutionary competition for the replication of information. What matters here is not specifically whether the ideas are true, or whether believing them benefits their carriers (although both of these may play a role), but whether the religion can successfully get itself stored and replicated using humans as its meme machines. The winners are those that outdo the competition by developing adaptations such as enjoyable rituals, memorable stories, glorious art and music, explanations for life’s mysteries (whether true or not), or nasty meme tricks such as threats of hell, and death to the infidel. The religions we see surviving around us today are the few big winners in that long and mindless competition to infect human minds.

There have been ferocious debates about this way of looking at religions, with opponents arguing that God is not a virus (Bowker 1995, McGrath 2007), evidence accumulating that religions do not improve societies’ health (Paul 2005) and arguments that faith itself is harmful (Dennett 2006, Harris 2006). Sadly there has been little in the way of meme-based empirical studies of religions and belief, which is badly needed if memetics is to thrive.

The value of memetics may also be judged by its theoretical contribution to understanding human evolution here on earth, and the reasons why humans alone have acquired complex and evolving culture. In trying to explain human uniqueness many different critical turning points have been suggested, such as the acquisition of tools, the evolution of language, increasing group size, or crossing the "symbolic threshold" (Deacon 1997, Dunbar 1996, Mithen 1996). All these theories have in common the familiar assumption that genes are the final arbiter, and that "adaptive" means adaptive for genes or for the organisms carrying those genes.

By contrast memetics claims that the turning point in human evolution was the advent of imitation. Indeed it is imitation that makes us human (Blackmore 2007b). Once early hominids could imitate well enough, they let loose a new replicator and, because of this, their evolution shifted to an entirely new phase.

What does "well enough" mean? As a general principle, replicators are more effective when they are copied with high fidelity (accuracy of copying), fecundity (number of copies made) and longevity (lifetime of the copies) (Dawkins 1976). So we can imagine imitation becoming good enough to set the new evolutionary process going, and then memes competing so that those of higher fidelity, fecundity and longevity increase in the newly growing memepool. From that point on humans alone are the product of two replicators, not just one. As Dawkins put it "Once this new evolution begins, it will in no necessary sense be subservient to the old." (1989 p 193-4).

This flowering of a new replicator was what I explored in The Meme Machine, leading to a completely new way of understanding how humans came to have such unique features as their excessively large brain, true language, extensive tool use, a love of music, art and religion, and complex culture. I argued that, by a process of "memetic drive", memes changed the environment in which human genes were selected and so drove genes to produce ever larger brains that were better at imitating the currently successful memes. In this way our brains became selective imitation devices, adapted to copying some kinds of memes more easily than others (Blackmore 2001).

A good example here is the evolution of language – long a highly contentious issue with many competing theories (Dunbar 1996, Pinker 1994, 2007). On this memetic view language, like art and all of culture, is not seen as an adaptation of benefit to humans and their genes, but as a parasite turned symbiont. Indeed all of cultural evolution is seen as happening for the benefit of the memes and in spite of posing a threat to humans and their genes. The human genes did, however, survive, but the creature that was once their vehicle (i.e. the human body) gradually turned into a better and better copying machine for the new replicator, the memes. That is how we humans became such effective meme machines.

Like Dawkins and Dennett I am convinced that treating memes as replicators is the way to understand human uniqueness and our evolutionary origins. Furthermore I think it allows us to speculate about how similar replicators may have evolved on other planets and what this might mean for cultural evolution in the cosmos.

A universe of replicators

We are able to ask the question "Are we alone in the universe?" because our distant ancestors began to imitate and therefore created memes and language. We are able to contemplate communicating with other worlds because our more recent ancestors invented new copying machines and therefore let loose more replicators.

In all my previous work in memetics I have deliberately used the term 'meme' to apply to any information that is copied between people, or between people and artefacts such as books or computers. I have often wondered whether 'artificial' memes such as websites and high-tech goods deserve a different name from 'natural' human memes such as spoken words, skills, music, art and religions, but have never found any principled way of distinguishing them. So on the grounds that a false distinction is worse than none I have stuck to the term 'meme' throughout. However, being asked to write about memes in the cosmos has forced me to think about this more deeply and to conclude that there is indeed an important transition from memes copied by human brains to information copied by technology other than human brains. These 'technological memes' are riding on top of both genes and memes to form a new layer of evolution. I'd like to call them 'temes'.

The justification is this: replicators do not evolve on their own but coevolve with the machinery that replicates them. In the case of earth's first level replicator, DNA, we have only a sketchy understanding of its origins (Maynard Smith and Szathmáry 1995) but we now see an exquisitely co-adapted system of DNA and cellular copying machinery on which most living things on earth depend. These living things can be thought of as the "vehicles", or gene machines, that carry the genes around and protect them (Dawkins 1976), or as the "interactors" that interact with the environment to produce differential effects on gene replication (Hull 1988). In the case of human evolution those vehicles eventually became the copying machinery for a new replicator, memes. Could it then be that the memes will do the same, building themselves meme-vehicles that in turn become the copying machinery for a new kind of replicator, temes. I suggest that this is what is happening all around us now.

The emergence of a new replicator is probably always a complicated and messy process (as it must have been with both of the first two levels), but we may discern some important steps indicating a new transition. When humans invented writing, about five thousand years ago, this increased the longevity of memes as compared with speech, and so new written memes spread along with the skill of making them. But the process still required human hands to do the copying, and human brains to select which texts to copy and distribute. When printing was invented just a few hundred years ago this increased both the fidelity and the fecundity because many copies could be made quickly. It also shifted the copying from human hands to the printing presses, but human meme machines still did the selection. Subsequently, new technologies have increased the fecundity, fidelity and longevity in many new ways. For example, communications systems spread memes over longer distances, and broadcasting increases the number of copies sent out; sound and image recording systems raise longevity by storing previously ephemeral information; and new technologies raise fidelity enormously by making the shift from analogue to digital recording. The systems that do this are increasingly copying, storing and propagating information without human help and at some point (which may never be precisely pinned down) they deserve the new name of 'teme machines'.

I said that these technologies are copying, storing and propagating temes, but arguably the most significant shift is only now getting under way; that is when the processes of variation and selection, as well as those of copying and storage, move from human meme machines to the new teme machines. This was hardly happening at all in the mid 1970s when Dawkins invented the idea of memes, but it is happening faster and faster now, just a few decades later.

We can see this in the invention of the World Wide Web and the search engines that are essential to its use. We should remember how recent this is, and how few people (if any) realised the significance of search engines when they first appeared. Now we take them for granted, and assume that if we want to know something the whole of the web is at our finger tips. Their significance here is that they are beginning to shift the job of selection of information from human brains to teme machines. For most searches a human still chooses what to search for, but the results of that search are determined by numerous previous searches and other aspects of the stored information, and there are increasingly many ways in which software does the selection rather than people. There are also sites that can write an essay by selecting and recombining bits of old essays, or write poetry or prose. Then there are evolutionary algorithms used to solve problems in engineering, medicine or robotics.

Are these systems true teme machines? If they carry out all three of the essential processes of copying, varying and selecting information outside of human control then perhaps they are, but at the moment these processes tend not to be bound together into something like a single organism. For example, books don't carry around their own printing presses, and cars don't carry around production lines; yet biological organisms do carry their own replicating machinery around inside them.

Why? This is perhaps the same question that Dawkins (1982) pondered when, in The Extended Phenotype, he asked why life is packaged into organisms at all. If the answers from biology can be generalised we might expect future teme machines to have all three processes built into them.

This relates to a question that has caused a lot of trouble for memetics; whether memes have the equivalent of a phenotype-genotype distinction. In biological evolution genes are not directly selected. Instead they act as instructions for building bodies (their vehicles or interactors) whose properties are the phenotype. These phenotypic properties are (at least in large part) determined by genes and in turn determine whether those genes are passed on again or not. This is a highly effective system because errors in building the phenotype are not passed on in the genotype. Dennett (1995) describes it as a system for retaining the "good tricks" that selection stumbles upon, and Dawkins (1982) emphasises the design value of a 'return to the drawing board' in each generation.

Superficially, memes may not appear to work this way, and some critics have argued that this disanalogy between genes and memes is a reason for rejecting memetics (McGrath 1995, Midgley 2000, Wimsatt 1999). For example when people pick up a new saying, or hear a tune or story and pass it on, they copy what they have heard. In this case there is no obvious genotype/phenotype distinction or replicator/vehicle separation (Aunger 2000). But if someone copies the recipe for making a new kind of figgy pudding, then it is easy enough to see the written recipe as a "memotype" and the resulting pud as the "phemotype"; the analogy is quite close because if Frederick makes a nasty mess of his pudding but passes the recipe on to Frances, then Frances does not have to follow Frederick's mistakes. But what if Frances watches him pick the figs, whip up the mixture, and set the oven? Then the distinction is gone (Blackmore 1999). Some memeticists have tried to work with a memotype/phemotype distinction (Benzon 1996, Speel 1997), but this has not been widely adopted and can be very confusing.

I suggest that a better way to understand what is happening is this: any system that copies the instructions for making a product is better than one that copies the product itself (Blackmore 1999). The gene/vehicle distinction is precisely that; a system for copying instructions (genes) for making living things (vehicles).However, such systems have to evolve over time. In biology we are seeing one that is billions of years old and has long ago achieved this distinction. But when we turn to memes, they are at most a few million years old, and we see a system still in flux. In some cases a distinction has appeared, and very effective it is too, but in other cases it has not.

Think of music. Some people do listen to others singing or playing and then copy the sounds directly, and in this case there is no replicator/vehicle distinction. But most music is now produced in recording studios and then copied in factories onto various media, ending up listened to in homes, or via phones, MP3 players and other personal music systems. The information on a CD or music file is instructions to create the pattern of sounds. When someone hears the music they don't copy it directly but go and buy it or download it from somewhere else. So this is all copy-the-instructions, and the CDs and digital players are meme vehicles.

The same applies to cars, clothes, fridges, furniture and almost all the household goods we take for granted. Almost no one sees a table they like and then goes out and cuts down a tree to make one like it from scratch. Instead they go and buy a similar, or even identical, one that has been produced from the same factory line as the one they first saw. If lots of people buy a particular model of car then the factory makes more from the production line it already had. It does not copy the ones already out on the streets.

Applying this idea we can now see that technological evolution is rapidly making the shift from copy-the-product to copy-the-instructions for making a product. This is not because there is any inherent goal for the process or any designer making it happen, but because better replicators overtake poorer ones, and better replicating machinery takes over from poorer machinery. If downloading digital files from the Internet produces more, more accurate, and more easily copied music than records, tapes, and CDs, then digital files will take over. It's as simple as that.

Another important shift is also taking place. Much of our technology is designed merely to store memes (such as books, CDs and so on) but increasingly also to copy them. If the analysis here is correct, then this is a critical step in the evolution of temes. That is, the meme vehicles we have created themselves become replication machinery for the next level of replicator. This is equivalent to the step that occurred when hominid gene machines began imitating and so became replication machinery for memes.

For now, we earthly Pandorans are still needed for various stages in the copying and selecting of temes, but teme machines are very rapidly evolving and it may not be long before there are self-repairing computer systems, self-maintaining power stations, artificial systems for choosing which new technology to build and which to ignore, and all the processes of replication will be shifted out of the hands of humans and into the world of teme machines. At this point the temes could carry on even if the Pandorans all died out. We are not at that point yet, but it is interesting to wonder whether we will get there before we wipe ourselves out, or not.

With this sketch of how replicators on earth have been doing we can now move on to consider the general case of replicators anywhere in the cosmos.

Extraterrestrial memes

Who is out there in the rest of the cosmos? Are there lots of other cultures that have evolved in completely different ways from ours? Are there lots of civilisations potentially capable of communicating with us? Or are we all alone?

These questions have been asked often enough, but in trying to answer them most people have been searching for "extraterrestrial intelligence" and so, I suggest, have got the emphasis wrong from the start. From what I have said so far it will be clear that I want to put the emphasis elsewhere, on how replicators arise, and what happens when one kind of replicating entity builds on the products of the previous one. From this perspective, intelligence is a product of replicator power, not its precursor: it is not intelligence per se that forces a great leap forward in living creatures or creates the possibility of culture, but the appearance of a second level replicator. This second replicator creates an environment in which greater intelligence is adaptive. Intelligence therefore increases, helping to provide a situation in which a third level of replicator can arise. This third level entails the creation of replicating machinery outside of any original Pandora species; a crucial step towards interplanetary communication, and so towards the possibility of others communicating with us.

In 1961, the astrophysicist, Frank Drake came up with his famous equation for estimating the number of intelligent civilisations in our galaxy capable of communication with us (Drake and Sobel 1992). He never intended his equation to provide quick answers, but rather to stimulate discussion and to direct research towards the important factors involved. In the same spirit I would like to suggest a different equation. Like Drake's it describes the number of intelligent civilisations in our galaxy capable of communicating with us, but it begins with a guess at the number of planets in our galaxy, ignores intelligence, and concentrates on the appearance of replicators. This equation is

N = n x fR1 x fR2 x fR3 x L

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

n = the number of planets in our galaxy

fR1 = the fraction of planets in n where a first level replicator survives

fR2 = the fraction of planets with R1 where a second level replicator survives

fR3 = 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

We know far too little to make sensible guesses about the values of most of these fractions. Even so, the equation provides a useful structure for thinking about the possibilities. I guess that the emergence of each new level of replicator marks a danger point at which the new replicator may simply fail, may wipe out the next level down, or may wipe out all previous levels. If this is right we can now make some further guesses about what may be happening elsewhere.

No replicators

We don't know whether "replicating entities" are necessary, sufficient, or both for life. If we assume both then all those planets which do not have the conditions in which replicators can arise must remain lifeless and therefore cultureless. Conceivably there are ways in which culture could appear that don't depend on replicators, or on individual organisms constructed by replicators, but if so I do not know how.

The First Level Replicator R1

Possibilities: R1 emerges and fails

R1 emerges and sustains life

The first level of replicator is the one that makes some form of life possible. This might have evolved from an even simpler replicating substance (e.g. a cheme (Szathmáry 1999)) or begun as a naked replicator (such as a simple self-copying molecule) before evolving to construct some kind of packaging or vehicle to protect and propagate it. I will say little about this first stage because it has been the subject of so much research and debate. Nevertheless it is clear that there are probably a large number of planets in our galaxy capable of sustaining a first level replicator that creates living things. These living things might be very different from any on earth. They might live under seas, within the solid mass of their planet, in an atmosphere of some kind, or they might roam on the surface. Conceivably they might appear on asteroids or interstellar particles. The way they develop will make it more or less likely that a second replicator will emerge.

It is curious that almost all life on earth depends on DNA and uses the same (or small variations on the same) code for translating stored information into instructions for protein synthesis. Whether it is possible to have multiple replication systems operating side by side, or whether one inevitably defeats all competitors, I do not know.

Second Level Replicator

Possibilities: R2 emerges in Pandoran species but fails

R2 emerges and kills Pandorans

R2 emerges, coevolves with Pandorans, and sustains culture

A second level replicator, R2, is one that emerges from a living thing created by a first level replicator. That is, R1 builds vehicles, and those vehicles become the copying machinery for R2, copying a different kind of information with variation and selection.

This brings us to where I began this chapter, with the idea that acquiring a second replicator is a dangerous step. This is because the new replicator (precisely because it is a replicator) will multiply selfishly. Since there is so much misunderstanding over the term “selfish replicator” I will just explain that this does not mean that genes or memes or temes have plans, desires, or intentions; it means only that their ability to multiply whenever conditions are right means that they exert replicator power regardless of the effect this has on anything else. This is all they can do, for they are just information mindlessly undergoing replication.

The danger is much like that of a parasite infecting a host. Conventional parasites are built on the same replicator as their hosts, while R2 parasites use a new replicator, but the principles may be similar, as may the possible outcomes.

First, R2 may emerge but never really take hold. This might happen if, for example, its effects on individual Pandorans are harmful so that individuals that do not acquire the copying ability die out, leaving ordinary Pandorans to prevail. Or it might happen if the conditions on the planet do not provide a sufficient advantage to R2 for it to thrive . On earth it seems that if environmental conditions change slowly then new skills are best dealt with by R1 (assimilated into the genome), if they change fast each individual is better to learn new skills for itself (individual learning). Only when individual learning is too costly and the environment neither too variable nor too stable is imitation worthwhile (Richerson and Boyd 2005). Once copying is underway it must be of sufficiently high fidelity, and there must be enough individuals capable of it for R2 to take off. If these principles generally hold true then there may be planets where R2 appears but never thrives. There may even be planets where this happens many times either in series or in parallel.

Second, R2 may begin to proliferate but then prove to be lethal. It could be that the new ability to replicate information is so energy intensive that the Pandorans' resources are exhausted. It could be that the things copied (behaviours, signals, chemicals or whatever) themselves become lethal. Or it could be that R2 begins to transform the Pandorans into better R2 machines and this transformation proves lethal.

Here on earth the dangers of R2 included the high energy requirements of building and sustaining a very large brain, and the danger of death during childbirth for a species that walks upright and therefore needs a narrow pelvis. Conventional theories attribute the increase in human brain size to tool use, the acquisition of language, or other adaptations of benefit to human genes, but I have argued that it was meme driven and therefore potentially dangerous (Blackmore 1999). In other words, the acquisition of memes might have killed us off.

Third, if this danger point is passed, R2 can become established on its planet. To do this R2 must adapt to R1 and vice versa. In other words the two replicators must coevolve. This may be a kind of arms race with the R1 Pandorans striving to survive and multiply copies of R1, while R2 strives to transform the Pandorans into better machinery for multiplying R2. If the poor Pandorans are not wiped out by this process then they become simultaneously R1 machines and R2 machines (more specifically, they remain R1 vehicles while becoming R2 replicating machines). The information they copy one to another itself evolves and this is what we call culture.

This whole process is analogous to coevolution between parasites and their hosts. So what determines whether a parasite will kill its host or will coevolve to become symbiotic? There is evidence from disease pathogens that when transmission is horizontal the parasite is more dangerous than when transmission is vertical (down the generations) (Maynard Smith and Szathmáry 1995). This is interesting with respect to a second level replicator such as memes. For example, it seems likely that early memetic transmission was mostly vertical, within families, and probably remained largely so during most of human evolution, which may be relevant to the fact that we survived. But now transmission is increasingly horizontal between peers rather than from parents to their children, and may be correspondingly more dangerous.

Another interesting question is whether it is possible for different kinds of R2 to appear at once on the same planet, resulting in multiple kinds of culture. Here on earth human brains are the only ones to have been turned into meme machines, and humans have the only significant culture. There is simple cultural transmission in song birds, cetaceans and chimpanzees (Whiten et al 1999) but not the kind of cumulative culture that marks a fully evolving new replicator. This is reminiscent of the fact that DNA has close to a monopoly as earth's R1, and prompts speculation as to whether there might be a general principle here. It seems plausible that if two second level replicators appear at once they will compete and the winner take all. This might even be the reason for the curious fact that we humans are the only remaining species in the hominid line. Perhaps Neanderthals and other related species also embarked on an R2 transformation but either their R2 proved lethal (perhaps through driving their heads to become too large for safe childbirth or through copying dangerous traditions) or they competed directly with early humans and lost. This is no more than speculation but might just prove interesting by prompting new hypotheses about what happened in hominid evolution, and what might happen elsewhere in the cosmos.

Returning to the possibilities for an R2 planet, the direction of cultural evolution will depend on the kind of species in which that culture first emerged, how much scope there is for transforming that species into an efficient R2 machine, whether the species has limbs suitable for constructing rich material culture, and what material resources there are available for such construction. The general principle will apply throughout that replicators of high fidelity, fecundity and longevity will prevail. This is likely to result in such innovations as communication systems, behavioural traditions, education, symbolic culture, tools and buildings. As the culture goes on evolving it may produce analogues of writing (because it increases longevity), long range communications systems (because they increase fecundity) and digitisation of information (because it increases fidelity). Gradually these may, depending on conditions, create signals detectable far from the home planet itself, but this only becomes a serious possibility once the way is already paved for the emergence of a third level replicator.

Third Level Replicator

A third level replicator, R3, is one that emerges from vehicles built by R1 and R2. That is, level 2 replicators build vehicles to protect and propagate themselves, and those vehicles become the copying machinery for a third level of replicator. On earth this is occurring as memes (R2), use humans (R1 vehicles and R2 replicating machines), to build R2 vehicles and these become machinery for copying temes (R3).

The possibilities with three replicators multiple rapidly, and we have little to go on when looking at our single planet for clues. So I will not speculate much further but confine myself to sketching out some of the more obvious possibilities. A few such possibilities if R3 takes hold are listed in Table 1.

	Outcome	**Remaining replicators**

1	R3 wipes out all of R1, R2 and R3	0
2	R3 kills R2 and its Pandorans; the rest of R1 survives	R1
3	R3 wipes out all of R1 and R2, but R3 survives	R3
4	R3 emerges and coevolves along with R1 and R2	R1, R2, R3
5	R3 machines merge with R2 machines	R1 R2/3
6	R4 emerges from R3	?

The first (disaster) scenario is one that may yet happen here on earth: that is, R3 gets beyond the control of its Pandorans, irreparably damages its planet, and so obliterates all life as well as itself. This could happen through overuse of resources, by causing drastic climate change, or by nuclear or other accidents. A second, slightly less drastic, outcome is that a planet's Pandorans are all killed but other life forms survive and can then begin the evolutionary process all over again. This is perhaps a more likely outcome here on earth given the widespread presence, variety and resilience of DNA based life, especially bacteria and insects.

A third possibility is that R3 coevolves with earlier replicators for long enough to get to the point where it does not need them any more, for example by creating self-repairing and self-replicating R3 machines. It might then alter its planet's atmosphere, climate or resources in such a way that all the rest of life dies out but the R3 systems survive.

If disasters like this are avoided, the earlier replicators might survive in several ways. As a fourth example, a stable symbiosis might be reached with all three replicators coexisting. We have all three on earth now but this is very recent and already serious dangers are apparent. So it is impossible to judge how stable such a system is likely to be.

Even if a three-replicator system can survive I suspect that there would be pressures for change. One change might be the merging of R2 with R3 machines. This is already happening here on earth with the invention of neural implants and other prostheses, and with the beginnings of the technology to create entirely artificial creatures based on DNA. This might be one way in which the descendents of earlier replicators might carry on in new forms.

Finally there is the possibility of yet another level of replicator emerging. What would R4 be like? It is hard enough to speculate about R3, but my best guess is this. An R3 level culture might develop the machinery to copy itself and so to seed new variants of that culture on different planets. This would lead to competition between variant cultures, and the evolution across the cosmos of R4 level civilisations. If this happened, planets separated by large distances would play a role analogous to islands here on earth; creating relatively isolated conditions in which cultures would evolve in different directions.

Note that the vehicles of earlier replicators, such as biological organisms, would not need to travel themselves as long as adequate artificial vehicles had been built and could cope with the widely different conditions on different planets. If it occurred, this kind of colonisation of the cosmos ought to be visible by its effects, and we have seen no evidence of it so far. Perhaps this is because, like all the previous transitions, the creation of a new level of replicator is dangerous.

Perhaps all the transitions I have discussed are incredibly dangerous and we earthlings are lucky to have got through to the third replicator stage: perhaps the dangers are small, and the rest of the cosmos is teeming with systems we have failed to detect. We do not know. Indeed we have no idea how to assess the values of any of the variables f_{R1}, f_{R2}, f_{R3}, or L but perhaps by thinking about evolution this way we Pandorans may see a little more clearly what sort of child we have let loose.

References

Aunger, R.A. (2000) Introduction to Aunger, R.A. (Ed) *Darwinizing Culture: The Status of Memetics as a Science*, Oxford University Press

Benzon,W. (1996) Culture as an evolutionary arena. *Journal of Social and Evolutionary Systems,* **19**, 321-362

Blackmore,S.J. (1999) *The Meme Machine,* Oxford, Oxford University Press.

Blackmore,S. 2001 Evolution and memes: The human brain as a selective imitation device. *Cybernetics and Systems*, **32**, 225-255.

Blackmore, S. (2007a) Memes, minds and imagination. In *Imaginative Minds*. Ed. Ilona Roth, London, British Academy

Blackmore, S.J. (2007b) Imitation makes us human. In *What makes us human?* Ed. Charles Pasternak, Oxford, Oneworld, pp 1-16.

Bloch, M. 2000. A well-disposed social anthropologist's problems with memes. In R. Aunger (Ed.), *Darwinizing Culture: the status of memetics as a science.* Oxford; New York: Oxford University Press.

Bowker,J. (1995) *Is God a Virus?* London, SPCK

Bull,L., Holland,O. and Blackmore,S. (2000) On meme-gene coevolution. *Artificial Life,* **6,** 227-235

Chater, N. (2005) Mendelian and Darwinian views of memes. In: Hurley, S. and Chater, N. (Eds) *Perspectives on Imitation: From Neuroscience to Social Science*. 2 volumes, MIT Press Vol 2, 355-62

Darwin,C. (1871) The Descent of Man and Selection in Relation to Sex. London, John Murray

Dawkins,R. (1976) *The Selfish Gene* Oxford, Oxford University Press

Dawkins,R. (1982) *The Extended Phenotype*. Oxford, Oxford University Press

Dawkins,R. (1989) *The Selfish Gene* Oxford, Oxford University Press (new edition with additional material)

Dawkins,R. (1993) Viruses of the mind. In B.Dahlbohm (ed) *Dennett and his Critics: Demystifying Mind.* Oxford, Blackwell

Deacon,T. (1997) *The Symbolic Species: The Co-evolution of Language and the Human Brain.* London, Penguin

Dennett, D. (1995) *Darwin's Dangerous Idea,* London, Penguin

Dennett, D. (2006) *Breaking the Spell: Religion as a natural phenomenon*

Distin, K. (2005) *The Selfish Meme: A Critical Reassessment,* Cambridge, Cambridge University Press

Drake, F. and Sobel, D. (1992) *Is anyone out there? the scientific search for extraterrestrial intelligence* Delacourt.

Dunbar,R. (1996) *Grooming, Gossip and the Evolution of Language*. London, Faber & Faber

Harris, S. (2006) *The end of faith: Religion, terror, and the future of reason*. Free Press

Higgs, P.G. (2000) The mimetic transition: a simulation study of the evolution of learning by imitation. *Proceedings of the Royal Society London B.,* **267,** 1355-1361

Hull,D.L. (1988) Interactors versus vehicles. In H.C.Plotkin (Ed) *The Role of Behaviour in Evolution.* Cambridge, Mass, MIT Press

Jablonka, E. and Lamb, M.J. (2005) *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral and Symbolic Variation in the History of Life*. Bradford Books

Jan, S. (2007) *The Memetics of Music: A neo-Darwinian view of musical structure and culture.* Aldershot, UK, Ashgate

Kendal, J.R. and Laland, K.N. (2000) Mathematical models for memetics. *Journal of Memetics* 4(1) http://jom-emit.cfpm.org/2000/vol4/kendal_jr&laland_kn.html

Laland, K and Odling-Smee, J. (2000) The evolution of the meme. In. R.A.Aunger (Ed) *Darwinizing Culture: The Status of Memetics as a Science*, Oxford University Press, 121-141

Lumsden,C.J. and Wilson,E.O. (1981) *Genes, Mind and Culture.* Cambridge, Mass., Harvard University Press.

Maynard-Smith,J. and Szathmáry,E (1995) *The Major Transitions in Evolution.* Oxford, Freeman

McGrath, A. (2007) *The Dawkins Delusion?: Atheist Fundamentalism and the Denial of the Divine*, London, SPCK

Midgley,M. (2000) Why memes? In. H. Rose and S. Rose (Eds) *Alas, Poor Darwin*, London, Cape, 67-84

Mithen,S. (1996) *The Prehistory of the Mind.* London, Thames and Hudson

Paul, G.S. (2005) Cross-National Correlations of Quantifiable Societal Health with Popular Religiosity and Secularism in the Prosperous Democracies: A First Look. *Journal of Religion and Society*, **7**, 1-17 (not sure about page nos)

Pinker, S. (1994) *The Language Instinct* New York, Morrow

Pinker, S. (2007) *The stuff of thought: Language as a window into human nature*. Viking

Pocklington,R and Best,M.L. (1997) Cultural evolution and units of selection in replicating text. *Journal of Theoretical Biology,* **188,** 79-87

Richerson, P.J. and Boyd, R. (2005) *Not by Genes Alone: How culture transformed human evolution,* Chicago, University of Chicago Press

Rose, H and Rose, S. (2000) (Eds) *Alas, Poor Darwin*, London, Jonathan Cape.

Runciman,W.G. (1998) Greek hoplites, warrior culture, and indirect bias. *Journal of the Royal Anthropological Institute,* **4**, 731-751

Speel, H.C.A.M. (1997) A short comment from a biologist on William Benzon's essay ' Culture as an Evolutionary Arena'. *Journal of Social and Evolutionary System,s* **20** (3), 309-322

Sperber, D. (2006) Why a deep understanding of cultural evolution is incompatible with shallow psychology. In Nick Enfield and Stephen Levinson (eds.) *Roots of Human Sociality*, 431-449

Szathmáry,E. (1999) Chemes, Genes, memes: a revised classification of replicators. In. C.L.Nehaniv (Ed) *Mathematical and Computational Biology: Computational Morphogenesis, Hierarchical Complexity, and Digital Evolution. Lectures on Mathematics in the Life Sciences,* **26,** 1-10

Whiten,A., Goodall,J., McGrew,W.C., Nishida,T., Reynolds,V., Sugiyama,Y., Tutin,C.E.G., Wrangham,R.W. and Boesch,C. (1999) Cultures in chimpanzees. *Nature,* **399**, 682-685

Wimsatt, W. (1999) Genes, Memes and Cultural Heredity, *Biology and Philosophy* **14:** 279–310

Wright, D. (2000) Translating Science: The Transmission of Western Chemistry into late Imperial China, 1840-1900 Brill, Leiden