Lucidity, science, and the arts: what we can learn from the way perception works

MICHAEL E. McINTYRE

Department of Applied Mathematics and Theoretical Physics,
Centre for Mathematical Sciences,
Wilberforce Road,
Cambridge CB3 0WA, UK
http://www.atm.damtp.cam.ac.uk/people/mem/

Invited keynote lecture to the

4th Symposium on Human Development held in Kobe, Japan
on 4 December 1999, under the auspices of the
Research Center for Human Science,
Faculty of Human Development, Kobe University.

Published version reprinted with minor corrections from

Networking of Human Intelligence:
Its Possibility and Strategy
a special millennium issue of the
Bulletin of the Faculty of Human Development,
Kobe University, vol. **7** no. 3 (2000), pp. 1–52.

Lucidity, science, and the arts: what we can learn from the way perception works[1]

Michael E. McIntyre

Department of Applied Mathematics & Theoretical Physics,
University of Cambridge, Silver Street, Cambridge, CB3 9EW, U.K.
http://www.atm.damtp.cam.ac.uk/people/mem/

The following is an edited transcript of the spoken lecture, minus material that depended on live demonstrations (only summarized here) plus material that did not have time to be included.

Introduction

Thank you, Professor Itoh, for those kind words of introduction, which I managed to follow with the help of the excellent simultaneous interpretation that we are privileged to have here. I want to thank the organizers of this very courageous symposium for asking me here, and to congratulate them on the excellent organization and on an exceptionally interesting programme. I am deeply honoured to be part of it and to be addressing an audience of such wide interdisciplinary interests on such an important set of topics.

There is no doubt, in my mind at least, that science and the arts have profound human importance and that the various points of view that they generate have much to teach us all. They prompt us to ask questions like:

- What are the arts, and why are they important? What is music, and where does it come from?
- What is science? What is mathematics?
- What is lucidity (of writing, speaking, thinking, design,...)?
- What is the Platonic world? Where does our powerful sense of "truth", "beauty", "perfection", etc., come from?

[1] *Coordinator's Note:* Some of Professor McIntyre's material is early draft material for a book in preparation. He tells me that he would welcome comments on how to improve it. The lecture was preceded by the coordinator's brief introduction, which is not included in this proceedings.

- Why are so many people instinctively "hypercredulous" — believing, or feeling an urge to believe, in some unquestionable Absolute Truth or Answer to Everything? — as with "it's *all* in the genes", "it's *all* down to culture", "market forces are *the* answer", "targets and performance-measures are *the* answer"
- Why do people instinctively want to dichotomize or polarize everything? — as with "nature *versus* nurture", "science as Mere Opinion *versus* science as Absolute Truth", "good genes *versus* bad genes", etc. (the eugenics myth — oblivious to the unimaginably complex interactions among genes and their products)
- What *are* these things called "instincts"?
- How can we better understand the power and limitations of science?
- What is "self"? What is "free will"?

I can hardly answer, in one lecture, every question on such a list, even with such a generous allocation of lecture time — especially since ideas like "instinct" have taken such a battering for so long and from so many vested interests. The whole debate about "instinct", by the way, has just been put into perspective in a beautiful new book by Patrick Bateson and Paul Martin (1999). But I think I can put us within reach of some good answers by emphasizing, first, a few points about how our human and pre-human ancestors must have evolved, according to today's best evidence, and, second, a few points about how we perceive and understand the world — evidence for which includes demonstrations like those from Professor Shimojo this morning, and others that I'll show and that you can easily check for yourself, with little or no special equipment. The two sets of points are related of course.

That is, I'll take it as reasonable to assume that our mental apparatus for what we call perception and understanding — cognition, if you prefer — must have been shaped by the way our ancestors survived over many millions of years. And the way they survived, though clearly a very "instinctive" business indeed, was not, the best evidence says, "all in the genes". Nor was it "all down to culture", not even in the last few tens of millennia.

Arguably, the most crucial point missed in the sometimes narrow-minded debates about "instinct", "nature versus nurture", "sociobiology", "evolutionary psychology", etc., is that most of what is involved in perception and cognition, and in our general functioning, takes place well beyond the reach of conscious thought. Some people find this hard to accept. Perhaps they feel offended, in a personal way, to be told that there is the slightest aspect of their existence that might, just possibly, not be under full and rigorous conscious control. But as we were reminded this morning, and as I shall remind you again with some simple demonstrations, it is very easy to show what is, after all, a priori likely, and is in any case well known: *plenty* of things in our brains take place "involuntarily", that is, entirely beyond the reach of conscious thought and conscious control.

Anyone who has driven cars or flown aircraft to any great extent will probably remember experiences in which accidents were narrowly avoided, ahead of conscious thought. The typical experience is often described as *witnessing oneself* taking, for instance, evasive action when faced with a head-on collision. It is all over by the time conscious thinking has begun. This has happened to me, and I think such experiences are quite common. Many years ago, the anthropologist–philosopher Gregory Bateson (1972) put the essential point rather well, in classic evolutionary cost–benefit terms:

> "No organism can afford to be conscious of matters with which it could deal at unconscious levels."

Gregory Bateson's point applies to us as well as to other living organisms. Why? There is a fundamental reason, combinatorial largeness. Every living organism has to deal all the time with a combinatorial tree, a combinatorially large number, of branching possibilities for future events — a point to which I shall return. Combinatorially large means exponentially large, like compound interest over millennia. Each branching of possibilities multiplies, rather than adds to, the number of possibilities. Being conscious of all these possibilities would be almost infinitely costly.

The "instinctive" avoidance of head-on collision in a car — the action taken ahead of conscious thought — was not, of course, something that came exclusively from genetic memory. Learning was involved as well. But much of that learning was itself unconscious, stretching all the way back to the (instinctive) infantile groping that discovers the world.

A journey toward self-understanding

As must be obvious by now, my approach to the foregoing questions will be essentially that of a scientist. Scientific thinking is my profession. My work is mainly on theoretical research to understand the Earth's atmosphere, as Professor Itoh has indicated, with some recent spinoffs for the Sun's interior. But long ago I almost became a musician. (Or rather, it would be more accurate to say that for short periods of my life I have been a musician.) So I know something about the artistic impulse as well. Science and the arts are both intimately bound up with the way perception works. And common to both of these human activities — whatever the popular mythologies may say to the contrary — is the urge to create, the joy of creation, the thrill of lateral thinking, and awe and wonder at the whole phenomenon of life itself and at the astonishing universe we live in.

One of the greatest of those wonders is our own adaptability. Of course we, the human species, are still in terrible trouble. We are genetically almost the same as our hunter-gatherer ancestors, tribes of people driven again and again to migration and warfare by, among other things, huge and rapid climate fluctuations. (Why else is our species — a single, genetically-compatible species with its single "human genome" — spread all around the globe?) I shall show you some of the new evidence for those climate fluctuations, gleaned from tiny gas bubbles trapped for hundreds of millennia in the Earth's great ice caps.

Our ancestors must have had not only language and lateral thinking — and I shall argue music, dance, poetry and storytelling too — but also the mechanisms of conviction politics, scapegoating and all the rest. It is now almost certain that they had language for well over a hundred millennia, at quite a sophisticated level. I shall remind you of the recent, remarkable, and practically conclusive evidence for that as well.

Even one hundred millennia is far longer than linguists used to tell us. But on today's evidence it was at least that long ago that our species started to spread around the globe from origins in Africa. The idea that language and culture started much more recently, an idea that you still hear repeated from time to time, makes little sense any more, if it ever did. What tends to be forgotten is that language and culture can be mediated purely by light waves and sound waves. This is a very convenient, an eminently *portable*, form of culture for a tribe on the move. And light waves and sound waves are such ephemeral things: together with body painting they have the annoying property of leaving no archaeological trace.

And now, in a mere flash of evolutionary time, a mere few centuries, we have shown our adaptability in ways that seem to me more astonishing than ever. We no longer burn witches. We no longer panic at the sight of a comet. Demons in the air have shrunk to a small minority of alien abductors. We dare to tolerate individual opinions. We dare to talk about an amazing new thing called human rights, and sometimes even take the idea seriously. The Pope dares to apologize for past misdeeds. We have space travel, satellite data links, and the Internet, bringing into existence a new freedom of information and disinformation and opportunities to exercise critical judgement, and to build computational systems of unprecedented power — all of it based on recent, and also amazing, achievements in software reliability (Valloppillil *et al.* 1998). We can read, though not yet understand, genetic codes. On large and small scales we are carrying out extraordinary new social experiments with labels like "free-market democracy", "children's democracy" (Neill 1970), "emancipation of women", and "Grameen Bank". I hope everyone has heard of the Grameen Bank of Bangladesh; if not, see the links on my home page. The emancipation of women is already, in Bangladesh for instance, leading to the stemming of population growth. Some of this was unthinkable as recently as fifty years ago.

This lecture reflects my own journey toward the frontiers of human self-understanding. One often hears it said that we live in an age devoid of faith. Well, I have a personal faith, a belief, a conviction, a passion, that the urge and the curiosity to increase our own self-understanding can continue to evolve us towards societies that will not, of course, be utopias, but could well be spiritually more healthy and generally more civilized than today.

When the Mahatma Gandhi visited England in 1930 he was asked by a journalist, in front of the newsreel cameras, "Mr Gandhi, what do you think of modern civilization?" The Mahatma replied, "That would be a good idea". The optimist in me hopes you agree. Part of such a civilization would not only be some kind of new covenant between science and society, including a clearer recognition of the power and limitations of science and its power to help us understand our own nature, but also a healing of the estrangement between science and the arts.

Many others, of course, have made such journeys. But in my case the journey began in perhaps a slightly unusual way. Music and the visual and literary arts were always there in the background; but the conscious journey began with a puzzle, with wondering why lucidity — in communication and in thinking — is often found so difficult to achieve, and then wondering what lucidity really is, in a functional or operational sense. I won't actually say very much about lucidity, as such, in this lecture — there is more about it in my published papers referenced below as **I**, **II**, and **III** — except to say that I want to use and understand the term in a wider and deeper sense than usual. I am not talking so much about what you can find in style manuals and in books on how to write, excellent and useful though some of them are. (Strunk and White 1979 is a gem.) What especially interested me, and still interests me, is the connection with such things as pattern perception and with the "organic-change principle" recognized, I think, by most artists.

For clear biological reasons, patterns intelligible to, and interesting to, the human brain usually involve what can reasonably be called "organic change". This means that some things change, usually by small amounts, while others stay invariant. Harmony in music is one example, as explained in Note 58 of **I** and in the expanded version on the Internet. Musical harmony is a particularly interesting case because "small" has not one but two quite different meanings, leading to the idea of "musical hyperspace". In music there are "hyperspace leaps", in the sense of going somewhere that is both nearby and far away. That is how some of the magic is done in many styles of Western music. There are many other examples of organic change and its effectiveness in the arts, in communication, and in thinking — in thinking about practically anything at all.

Fundamentally, as argued in **I**, our peculiar perceptual sensitivity, and cognitive sensitivity, to organic change seems to stem from the survival-value of being able to recognize the difference between living things and dead or inanimate things. To see a cat stalking a bird is to see organic change. The sensitivity is, I would dare to assert, very deeply instinctive. Not surprisingly from this viewpoint, the ability to recognize the difference between the living and the inanimate has been shown, experimentally, to be well developed in human infants a few months old.

Notice by the way how intimately involved, in all this, are ideas of a very *abstract* kind. There are clues here about mathematics and the Platonic.

For a start, the idea of some things changing while others stay invariant is itself very abstract. I'll point out some more examples when we come to questions about mathematics. As argued in **II**, I think it's easy to show that we all have *unconscious mathematics*, and an *unconscious power of abstraction*. This last is almost the same as saying that the brain can handle combinatorial largeness. The brain can handle — unconsciously of course — vast numbers of possibilities at once. That is fundamentally what abstraction means. The implication is that the roots of mathematics and logic — and abstract cognitive symbolism generally — lie far deeper, and are evolutionarily far more ancient, than they are usually thought to be. This in turn opens the way to understanding what the Platonic world is.

So I am talking about lucidity in a sense that cuts far deeper than, and goes far beyond, the niceties and pedantries of style manuals. But before anyone starts thinking that it's all about cloud-cuckoo-land and ivory-tower philosophy, let me remind you also of some hard practical realities. What I am talking about is relevant not only to thinking and communication but also, for instance, to the design of machinery, of technological systems, of software, and of user interfaces. Unreliable, unmaintainable software whose internal design violates lucidity principles is still almost the norm, especially in the world of commerce. (This is partly because of ignorance but partly for other reasons, of a purely commercial kind, related to what is called monopoly lock-in; see Valloppillil *et al.*, *op cit.*.) And even *user interfaces* that violate lucidity principles, making them "unfriendly", are still commonplace. The self-evident cost of all this seems to me to indicate, above all, that lucidity principles are not yet widely understood and appreciated, outside specialist communities, or perhaps that they still tend to be regarded as some kind of luxury.

Underlining this last point is the fact that one finds user-unfriendly interfaces even where they imperil safety in a big way. The Three Mile Island nuclear reactor was a particularly clear example, as I pointed out in **I**. You don't need to be a professional psychologist to appreciate the point. Before the nuclear accident for which Three Mile Island is well known, the control panels were like a system of traffic lights in which red sometimes means stop but sometimes go. (The problem is all too familiar: "User-friendliness is a luxury we can't afford; go away and read the manual: look, it *says*, on page 199, that red means 'go' at a T junction but 'stop' at a crossroads, so you've no excuse", etc.)

But now I'm getting ahead of myself. What, indeed, are these things called instincts, these internal forces preceding conscious thought, and what role did evolution have in forming them? What indeed do we know about human and pre-human evolution? If it isn't *all* in the genes, if it isn't just nature doing eugenics — dichotomizing genes into "good" *versus* "bad" and

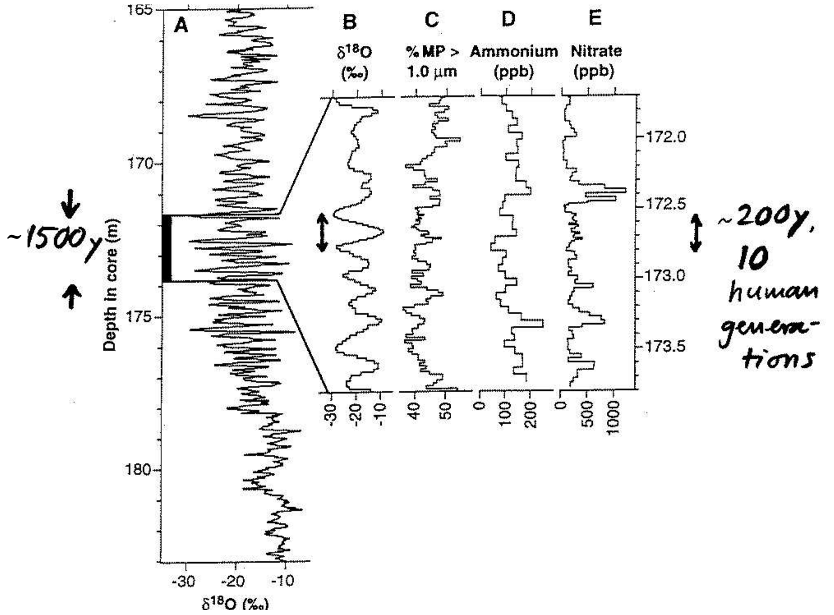

Figure 1: 1500 years of strong climate fluctuations in the Asian monsoon region, around 30 000 to 40 000 years ago, seen at high resolution (right-hand plots); δ^{18}O, plot B, measures the fluctuations in the ratio between the oxygen isotopes ^{18}O and ^{16}O. From Thompson *et al.* (1997).

wiping out the "bad" ones — then how on Earth did it all work?

Climate instability and our ancestors

Let us first remind ourselves of the timespan, or rather timespans, of climate change and human evolution. These have been much informed by recent research.

An example is shown in figure 1. This plot is from the marvellous work by Lonnie Thompson and his collaborators, who have taken heavy equipment high into the Andes and the Tibetan Plateau to drill for ice cores, then subjected them to laborious and painstaking analysis. Figure 1 is from the Guliya Ice Cap on the Qinghai–Tibetan Plateau (Thompson *et al.* 1997). From the tiny bubbles of gas trapped in the ice, analyzed with high-precision measuring instruments to look at chemical and isotopic composition, we can see the traces of past climate fluctuations. Time runs vertically upward.

The left hand plot, covering a timespan several tens of millennia ago, shows enormous fluctuations in a quantity denoted by δ^{18}O, measuring the relative amounts of two oxygen isotopes and usually taken as a temperature proxy. The expanded view on the right shows a 1500-year chunk of the record. This is from around 30 or 40 millennia ago. The first plot in the expanded view, plot B, is from the same δ^{18}O record. Notice that the typical period of a single one of the oscillations is of the order of 200 years, that is, 10 human generations if we assume that a generation takes 20 years.

But the largest amplitudes of the fluctuations in δ^{18}O are far too large for the usual proxy temperature formulae to be credible. If you were to apply such formulae, then the temperature fluctuations you would deduce would be an embarrassingly large number of degrees Celsius. All we can say is that vegetation and food supplies must have changed drastically during the lifetimes of single individuals, generation after generation. In that part of the world, such changes would very likely have been related to changes in monsoon and rainfall patterns. It is highly plausible that there would have been cycles of plenty and population growth followed by hardship and, in many cases, migration. The picture would no doubt have been complicated by the chaotic dynamics of ecological systems, but it is difficult to believe that such large changes could not have had large effects.

Let us zoom out to the larger timespan, 240 000 years, shown in figure 2. This is from the work of another research group famous for drilling and analyzing very deep ice cores from the ice cap that forms the high plateau of East Antarctica (Jouzel *et al.* 1993). Here we have coarser time resolution but a far longer timespan. Figure 2 shows two trace gases in the ice bubbles, upper and lower plots, and an estimated temperature, middle plot, from isotope ratios (δ^{18}O and deuterium). Time now runs from right to left.

You can see that there are big fluctuations over the larger timescales as well. The timespan of 1500 years that I emphasized in figure 1 has shrunk to the width of the thin vertical line near the left of figure 2; 1500 years is 75 human generations. Figure 2 has a timespan of, notionally, 12 000 human generations.

We can zoom out again and contemplate an even longer timespan (figure 3). We are now talking about the past 3.5 million years or 3 500 millennia or, notionally, 175 000 generations, perhaps more. I'm not going to show you any climate information this time — and I don't think anyone has

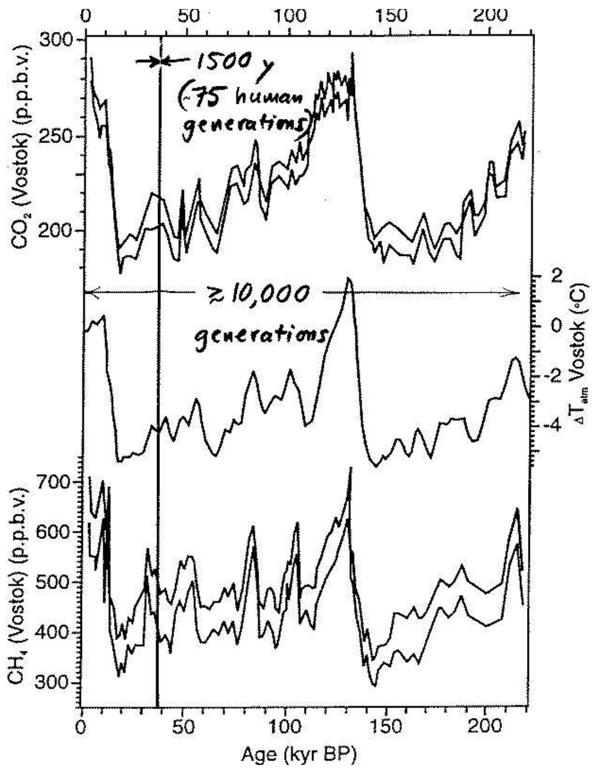

Figure 2: Climate indicators from the high Antarctic (Vostok ice cores) over the last 240 millennia, from Jouzel *et al.* (1993). The top and bottom plots show error-bar envelopes. The width of the thin vertical line roughly indicates the 1500 years or 75 human generations shown at the left of figure 1

such information at anywhere near the time resolution seen in figure 1 — but paleoclimatologists would probably think it plausible that, over most of the past half million years or so at least, the climate, being a "glacial" regime generally similar to that in figure 1, probably carried on with its huge, unstable fluctuations in much the same way as suggested by figures 1 and 2.

Figure 3 is a rather busy picture and I'm going to talk about it for a little while. Let us focus for a moment on the two-million-year mark. Two million years ago is a very significant time for human and pre-human evolution. It is the beginning, very roughly — this, we think, was not a sharp transition — of what we call the Pleistocene period, which covers something like the past two million years all the way to a time around 10 000 years ago. Then

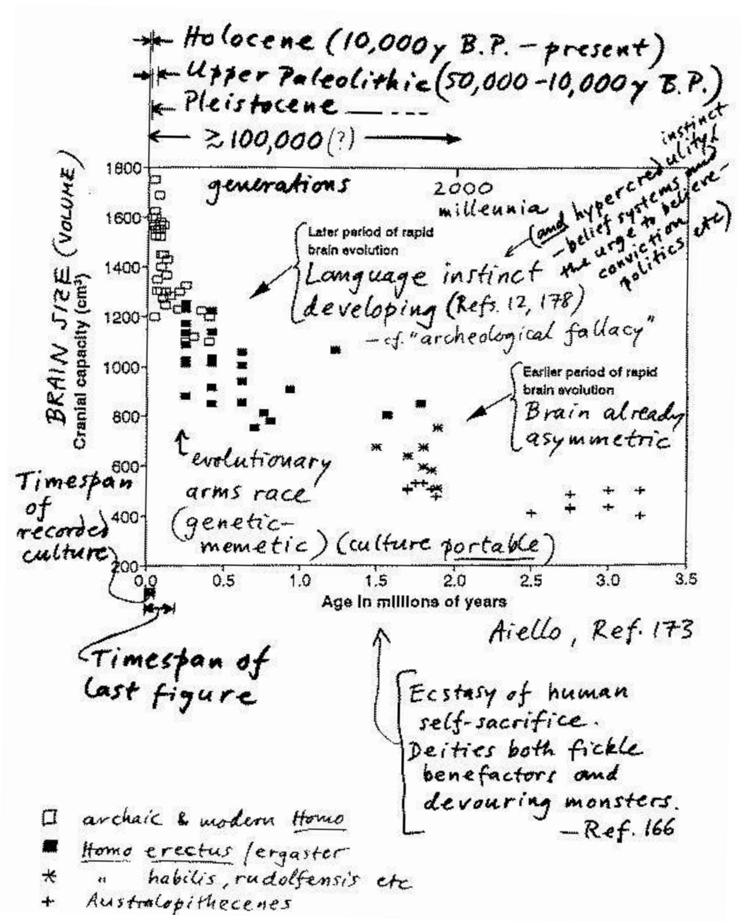

Figure 3: Human and pre-human evolution over the last 3.5 million years, from the Australopithecenes to modern *Homo*. Plot of brain size is from Aiello (1996) (ref. 173 of **III**); but the messy scribblings are mine, as used in the spoken lecture, with apologies to Professor Aiello! (Some of the brain-size measurements have since come up for revision, but the general trend is not in doubt.) Ref. 166 from **III** is Ehrenreich (1997). The climate entered the Pleistocene, an unstable glacial regime, around 2 million years ago

the Pleistocene ends relatively suddenly, and gives way to what we call the Holocene, which is the last 10 000 years up until the present. The Holocene had, and has, we still hope, a relatively warm, relatively stable climate. It was that relatively stable climate that made agriculture possible, and with it the first city-sized human communities (as well as lots of new human and animal diseases, as Jared Diamond has reminded us). Notice how tiny a sliver of time the Holocene occupies in figure 3, as indicated at top left. Even 10 000 years, 10 millennia, is a mere flash of evolutionary time.

Archeologists will tell you about the period called the Upper Paleolithic, also indicated at top left, which started very roughly 30 000 to 50 000 years ago depending on where you do your archaeology. Even this is scarcely more than a flash of evolutionary time. But the Upper Paleolithic is when things get interesting for an archaeologist. It is the timespan of recorded culture. Instead of just a rather restricted set of utilitarian stone tools, archeology suddenly sees, in these last few tens of millennia, very obvious *evidence* of human high culture — many visual artifacts, beads and bracelets, exquisitely beautiful cave paintings, musical instruments, and so forth.

What's actually plotted in figure 3 is not climate information, but rather the fossil record of the brain sizes of our ancestors, measured in millilitres or cubic centimetres. There are some dramatic trends. Over the last half million years, well into the glacial regime with its climate probably like that of figure 1, you can see that the increase in brain size, viewed on this timescale, was very rapid. Looking further back in time, you can see a less clearly delineated, but also very big, increase in brain size, starting just under 2 million years ago, and reasonably assumed by most researchers to be connected with the onset of the Pleistocene, in which the climate cooled and forests shrank in places like central Africa.

The crosses at bottom right correspond to the Australopithecenes, thought to be our early ancestors, or to be very like our early ancestors. By the way, not all these points are considered completely accurate today. After the results were published it was discovered that some of the measurements needed to be redone; but the corrections are unlikely to be big enough to affect the broad picture.

The filled squares in figure 3 correspond to *Homo erectus*, the first pre-human species fully committed to bipedality and ground-dwelling. *Homo erectus* seems to have been the first pre-human species to have migrated almost world wide, from origins in Africa. There are *Homo erectus* remains in China and in Indonesia, though as far as we know none in Australia. There is some evidence, very recently discovered, that at least one *Homo erectus* group devised, or stumbled on, some means of crossing deep-water channels between Indonesian islands, when sea levels were low enough and the destination was in line of sight. (In the Pleistocene, sea levels were sometimes 120 metres lower than today.) But they probably couldn't launch themselves on to the open sea. *Homo erectus* appears to have been pretty smart and much more adaptable than, say, today's chimpanzees, and better at toolmaking, but probably didn't have anything approaching our language ability. So boatbuilding and astronomical navigation would be a bit much to expect.

The open squares in figure 3 are ourselves and our immediate ancestors, archaic and modern *Homo*. The fossil record is too sparse to be able to say exactly where the lines of ancestry are continuous, though the transition from something like *Homo erectus* to something like our species is currently thought to have taken place in Africa. This is very much an area of ongoing research, especially as the hunt for more fossils continues and the DNA techniques increasingly come into play.

Climate, language, and the arts

Let me try to summarize where we have got to. Paleoclimatic records at various time resolutions, including recent ones in which very high time resolutions have been achieved — of which figure 1 is only one recent example among many — strongly suggest, and for more recent times clearly demonstrate, that the climate of the past million years or so was largely glacial, and typically very unstable, with many large and rapid fluctuations. Rapid means so rapid that there are significant changes in an individual's lifetime. Endnotes 179, 185, and 186 of **III** provide a few more keys to a burgeoning literature. In some parts of the world at least, notably the North Atlantic and West European areas, there were climatic temperature changes of the order of several degrees Celsius "within decades". In the Asian monsoon region there were the large oscillations with periods of the order of 2 centuries shown in figure 1. And again and again there were large changes in sea level, including episodes of sea-level rise at rates up to about 3 metres per century. Together with the vagaries of chaotic ecodynamics and of human population growth, the climate fluctuations must have kept tribes on the move in what became a global scale diffusion of our species, almost a "random walk" in the technical sense of the term, in an unwritten saga of growth, famine, migration, and warfare. The runaway brain evolution, as Christopher Wills (1994) has called it, whose signature is seen at the left of figure 3, is probably the signature of — quite literally — an “evolutionary arms race”.

It is strongly arguable that such an arms race was part of what drove, and was driven by, the development of language and of related cognitive abilities — the exercise of which produces, among other things, what we now call art in its various forms. I have discussed this in more detail in **III**. In brief, all these things would have felt a runaway selective pressure to develop because they enabled tribes of our ancestors to become larger and better organized. This would have advantages in the whole business of surviving in the teeth of climate fluctuations, migration, warfare, and other contingencies. How does a large group, say ten or more times the size of a group of chimpanzees, get organized for survival? They can’t, like chimps, depend on group bonding by grooming (*e.g.* Goodall 1990). They have to use light and sound as well as touch. How do you form even bigger groups if warfare or migration demands it? How do you organize scouting parties, with unknown dangers ahead? It is obvious that language, for one thing, would be an advantage.

And, speculation aside, what is now more than strongly arguable — because recent evidence shows it conclusively — is that our language ability resides in genetic memory. It does so in a very clearcut sense that I shall explain shortly. The implication of all this is that language is far more ancient than used to be supposed. If the language ability resides in genetic memory, then the development of that ability must have involved significant genetic evolution, in the sense that the genome changed significantly. In other words, language could not have been just a quick adaptation using existing genetic diversity, or pre-existing general intelligence — whatever the rather disreputable phrase “general intelligence” might mean. Rather, language must have involved a substantial, systematic change in the genome itself. The timespan for such a change, geneticists will tell you, is likely to be hundreds of millennia at least. The strong, indeed almost inescapable, suggestion from this is that our language ability must have been developing simultaneously with the increase in brain size, most spectacularly over those same last few hundred millennia — and very likely an intimate part of, and a major driving force in, the brain development itself.

And the moment you accept the likelihood of the runaway selective pressure I am talking about, you can see that there would have been a concomitant pressure to develop the essentials, the core, of something like what we call culture, as well as, of course, what we call religion. The two would not have been separate. How does a large group stay together, feel inspired to act together, survive traumas, and generally be what we might call spiritually and emotionally healthy and strong as a group? It needs ways of expressing and reaffirming group bonding, what is sometimes called "solidarity" today. That has always, until recently, in nearly all cultures known to us, involved not only what I call hypercredulity (see **III**) — in the form of unquestioning belief in the tribe's Absolute Truths, something that begins to be possible once you have language — but also an intimate involvement with what we call the arts. (Part of our own spiritual malaise, I would dare to suggest, is that today we tend to put *The Arts* on a distant pedestal divorced from real life. My wife Ruth and I had a few wonderful years, long ago, in professional chamber music, and were struck by how much more responsive the audience would be if we said a few friendly words before playing — it didn't seem to matter how trite — just establishing ourselves as ordinary humans who were *not* on a pedestal.)

So I am arguing that there must have been runaway selective pressure for the development of what amounts, in essence, to artistic activity. The musical side of this, at least, might indeed be even more ancient than the language side. The seeds of the Orpheus legend may lie very deep in the collective unconscious, as Jung called it. After all, today's ground-dwelling baboons need group bonding to survive (in fairly large groups), and seem to make use of vocalization for that purpose. For us, group musical activity — singing and dancing — can evoke feelings of emotional release and togetherness, of transcending the self, with less risk of provoking conflict than does the use of language (Cross 1999). Musical vocalization and gesture can replace, or supplement, chimpanzee-style grooming in a way that is effective in a much larger group. It has natural beginnings in the play between human mothers and infants, itself quite obviously a kind of vocal and gestural grooming.

Today, any musician who knows how to bring music to life, in any style of what is aptly called live performance, knows that, somewhere in our unconscious being, there is an "internal musician" for whom singing and dancing are essentially the same thing and to whom, in the end — after the notes have been learned and the technical obstacles cleared away — control of the performance must be handed over, in order to bring the performance to life (*e.g.* Green and Gallwey 1987). In a deeply felt sense, one must let the music itself take control of the performance. Professor Bando's wonderful lecture and musical demonstrations suggest to me that he would agree with this.

Another thing we know from the fossil record is that, as far back as

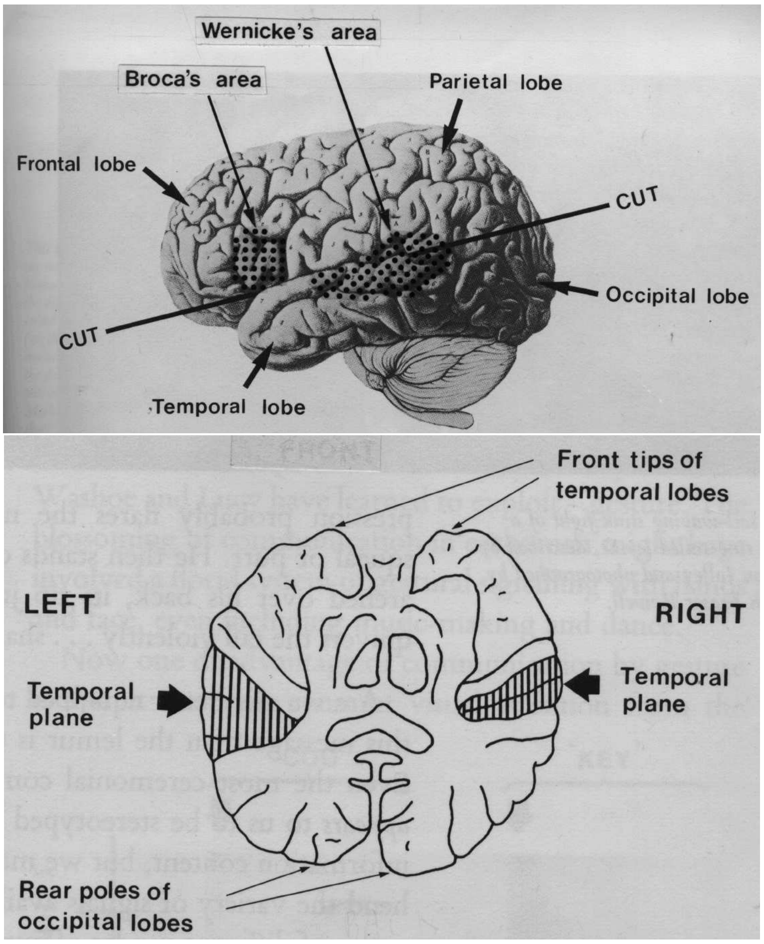

Figure 4: Side view (upper picture) and nearly-horizontal section (lower picture) through the human brain. From Blakemore (1977). Upper drawing by Louis Pierre Gratiolet (1815–1865)

2 million years ago, the brain was already highly asymmetric: it had very much the same kind of left-right asymmetry that our brain has today, related to the ability to vocalize in an elaborate way and to follow details in the sound when heard (*e.g.* Aiello 1996 & refs.). Figure 4 shows two views of the modern human brain. The upper drawing is a side view of the brain from the left. In the normal right-handed person the two regions shown coarsely shaded, known as Broca's area and Wernicke's area, are bigger on this left-

hand side. They are crucial, respectively, to our ability to speak and listen. If you imagine taking a nearly-horizontal cut through the brain and looking down from above (lower drawing) you see that Wernicke's area (hatched) is much bigger on the left than the right. The presence of such asymmetry 2 million or more years ago is consistent with the idea that our ancestors were already using elaborate vocalization as well as visual gesture, and that the selective pressure to develop such vocalization came from the value, for survival, of large groups and hence of group bonding though sound waves, vocal grooming if you will.

The Nicaraguan Big Bang

What, then, of my brazen assertion that, in any case, our language ability, however it developed, resides in genetic memory? Let me try to make the assertion a little more carefully. In the modern human brain, genetic memory contains, implicitly, the complete wherewithal — the seeds of the self-assembling, self-organizing yet input-sensitive components — from which to build the syntactic machinery, the grammatical machinery, of language.

And what evidence supports this last assertion? There are several lines of evidence. But by far the strongest, I would say conclusive, evidence is that from Nicaraguan sign language. This evidence emerged in time to make a brief appearance in Steven Pinker's excellent and powerfully-argued book *The Language Instinct*. I have studied the source material myself, as referenced in **III**; also Kegl *et al.* (1999).

In Nicaragua, before the Sandinista government came to power in 1979, there was no sign language. Deaf people were socially isolated. Following their socialist ideals the Sandinistas created two schools for the deaf, in the city of Managua. At those schools, the children began signing to each other.

This was a crisis for the school authorities. They had to find out what the children were saying or the situation would get completely out of hand. In June 1986 they called in an expert on sign language. The expert was Professor Judy Kegl of the University of Southern Maine, then at Northeastern University. She and her co-workers Ann Senghas and Marie Coppola found themselves documenting something that had never been observed scientifically before: the birth of an entirely new language. An event of such uniqueness and importance would probably be called a "big bang" these days; and it is not only unique but also very remarkable.

There were two stages. First, the older children, older than about seven years of age, created what linguists call a pidgin — syntactically feeble, incoherent, chaotic, and unstable. It lacks the usual grammatical devices — verb tenses, noun cases (nominative, possessive, etc.), adjectival clauses, conditional "if" clauses, and all the rest. Then, second, the younger children, between about three and seven years of age, added the missing grammar. They converted the pidgin into what linguists call a creole. This means a syntactically powerful, coherent, complete, and stable language. They did it with no adult input. All the children needed was the social stimulus, the need to communicate. No adult knew their language until Judy Kegl and co-workers learned it from the children!

When observing your own children learning to speak, you may well have noticed hints of their grammatical powers. In the case of English — I haven't found out what the Japanese equivalents are — small children will say things like "I keeped mouses in a box". What they have to learn more consciously is the irregular form of standard English: "I kept mice in a box". They have unconscious knowledge of verbs and nouns and their cogitive–syntactical functions; they have noticed, unconsciously, the usual inflections corresponding to those functions in the regular verbs and nouns of the language in their environment. They have noticed, unconsciously, that "keep" functions as a verb, and "mouse" as a noun; and they have applied the usual inflections. As Noam Chomsky recognized long ago, it is the *departures* from coherent syntax — the culturally evolved *irregularities* — that have to be learnt.

The archaeological fallacy

You can still find in the archaeological literature and, I think, in some linguistic circles and in other high places in academia, the idea that language and culture — and indeed, in some variants of the idea, the cognitive *ability* to develop such language and culture — suddenly blossomed at the beginning of the upper Paleolithic, at the same time as recorded visual art, far faster than the genome could have evolved. Up to a point I sympathize with the position that you mustn't believe anything until you have hard evidence, but that's not quite the same as saying that the absence of evidence implies that a thing can't possibly exist. And in the case of language, at least, we do now have very hard evidence indeed, the existence and documented origin of the Nicaraguan sign language. It no longer makes the slightest sense to say that language is "all down to culture", that it was suddenly invented at the beginning of the upper Paleolithic then passed on from parents to children.

What now seems more than likely is that, long before the Upper Paleolithic, there were highly developed cultures based on what we call dance, poetry, storytelling, rhetoric, and song, carried with them by migrating tribes and leaving no archeological trace. **III** discusses this further, together with the related question of the role of the play instinct — juvenile play as the deadly serious rehearsal for adult life. And Bateson and Martin (1999) have now given us an excellent and authoritative chapter on juvenile play and its significance, informed by much detailed behavioural observation.

Any biological competence, including the language ability, will atrophy if it is not used to the full. Poetry, storytelling, rhetoric, conviction politics, and their juvenile-play precursors seem to me to be likely aspects — natural consequences — of using the language ability to the full. Kittens chase each other's tails; children bounce sentences off each other.

The portable cultures in question are of course nothing other than what we already call "oral traditions". Long before the Upper Paleolithic, their storytelling would have probably included, in various forms, the usual sorts of saga about heroic journeys and exploits. This makes sense biologically. For one thing, what more efficient way, other than a powerful oral tradition, could there be to maintain tribal memory of past famines, migrations, surrounding geography, or whatever was important to survival over the centuries? It seems to me to be no accident that we feel more than a touch of magic today when we hear a well-told tale, not least a tale of an epic journey, preferably with an element of exaggeration or fantasy (Neill 1970). Look at the massive success of the film *Star Wars*, and the power of tales like Ursula Le Guin's *The Left Hand of Darkness*. Everyone has favourite examples. Australian aborigines "sing" their actual and dream worlds, and tell us that their songs are (among other things) maps of the beautiful land in which they live, the land they love. Their art contains fantasy but also reflects the real world, just as does the greatest art of our own Western high culture. It seems to me to be no accident that Gustav Mahler felt that a symphony must contain the world, and Carl Nielsen that music is life. And in the harsh and treacherous conditions of the Pleistocene, portability would have been useful. Lugging material works of art around during exploration, migration, and warfare might just not have been our ancestors' highest priority.

Of course it's possible that maverick individuals might nevertheless, from time to time, have played with visual artefacts — not hard to think of once they thought of painting their bodies — long before the Upper Paleolithic explosion of the recorded arts when the idea of making visual artefacts became, we must suppose, suddenly fashionable and widely accepted. We can imagine the likely reaction of the tribal elders long before that change in fashion: "If the gods had meant us to fool with beads and bracelets, they wouldn't have given us poetry, song, and dance." Or, more perhaps more realistically, in thunderous tones: "Violate not the sacred shapes of our ancestral tools. Your transgression will dissipate their power and doom us all." One can imagine an even sharper reaction to experiments with the beginnings of musical instruments. Such fascinating sounds must surely attract evil spirits! "Foolish one, do not play with fire! It is a mockery of the sacred songs. A terrible punishment will be visited on us." Priorities would have been decided unconsciously, then as today.

Fickle benefactors and devouring monsters

What, then, of religion? To be sure, the unquestioning belief in the tribe's Absolute Truths, plus the dichotomization Us *versus* Them,[2] would have been powerful in tribal bonding, including bonding in warfare, and powerfully selected for. But is that all there is to it? I suspect not. Why for instance do so many tribal–religious mythologies seem to be full of deities that are both fickle benefactors and devouring monsters? Where does the enormous "high" — the ecstasy — of war and self-sacrifice come from? What evokes martyrdom and the transformation of ordinary humans, male or female, into killers or willing sacrificial offerings, or both? As Ehrenreich points out (1997 and **III** ref. 166), these things, too, probably have pre-linguistic roots.

Think again of the circumstances around 2 million years ago. Our ancestors were abandoning the shrinking forests and becoming committed to a ground-dwelling existence, to being fully bipedal. That meant, of course, increased vulnerability to the big cats and other large predators. The price of versatility and geographic mobility was to live very dangerously, even if groups were already large.

[2] I shan't have time to say much about the dichotomization instinct. There is a more careful discussion in **II** and **III**. Us or Them merely extends primordial dichotomies such as edible or inedible, fight or flight, male or female. The dichotomization instinct imposes this pattern on everything, puts everything into extremist, black-and-white terms, whether appropriately or not. It is an aspect of hypercredulity.

Lacking, at first, the most highly developed hunting skills and the most sophisticated weapons, those early ancestors of ours would not only have been regular victims of such predators but would also, very probably, Ehrenreich argues, have depended on scavenging from carcasses of other large animals killed by the same predators — contributing to the high quality diet both demanded by, and made possible by, increasing brain size and the commitment to live dangerously and get smart quickly, in all kinds of ways including the ways of social intelligence.

Here is another runaway selective pressure. It may well have been what drove the first episode of brain expansion seen in figure 3, the expansion beginning around 2 million years ago. And yes, the large predators were indeed, quite literally, fickle benefactors and devouring monsters, terrifying and awesome from our ancestors' viewpoint, just as they would be, at close quarters, from ours. The self-assembling, self-organizing yet input-sensitive components of the corresponding mental imagery may well have been established in genetic memory even earlier, over many millions of years.

Such a scenario, followed later by the transition to being predators ourselves, well explains, it seems to me, why our psyches and mythologies are so full of these deities that combine the seemingly illogical roles of benefactor and monster, deities to whom some individuals will gladly, even ecstatically, become martyrs for the sake of the group or tribe, and with whom individuals must also, from time to time, do battle — what is called fighting with dragons or wrestling with angels.

Images like these seem to permeate all mythologies; they are another very deep part of our collective unconscious. You see hints of them from the way small children, and their parents, love to "play monsters." As the great composer Michael Tippett (1995) has reminded us, these awesome, terrifying and sometimes beautiful deities, these "tygers burning bright" and their protean variants, their organically-changing proliferations, part human, part beast, endlessly multifarious, divine and demonic, are part of what underlies some of the greatest achievements of artists of all kinds.

Darwinian and sociobiological fallacies

I have been talking rather a lot about evolution, in a Darwinian sense, and I want to make an important cautionary remark. There is a popular mistake, or myth — you read it again and again in the newspapers, you even hear it from biologists and respected scientists who should know better — which says that Darwinian evolution is *hypercompetitive*, that it is all about individuals fighting each other, about survival of the fittest in a naïve, comic-strip sense, about *nothing but* "nature red in tooth and claw".

This myth, or the adherence to it, is sometimes labelled "brutal biological determinism", or, when applied to humans, "sociobiology" (quite contrary to the original meaning of the word, I heard Edward O. Wilson, who coined it, say on the radio recently). Sometimes the myth is compounded with other myths and given labels like the "dominance of the selfish gene", understood in a very narrow sense. "Wicked people like Wilson are saying that it's all in the genes," we are told (again contrary to what Wilson actually wrote and said); "such people say that human beings are just like simple computers, programmed like automatons to follow some rigid program."

Well, of course, it ought to be obvious that that is absurd, and that Darwinian processes, Darwinian evolutionary ideas, don't really say that. The most essential idea is that whatever helps survival is selected for. Such selection must have gone on even when the detailed circumstances and mechanisms for selection are highly complex, and not easily understood by us. And the processes I have been talking about are nothing if not highly complex.

Take love and altruism, for instance. In practically all of this lecture so far I have, in effect, been arguing that such things as love and altruism are, among other things, for a species like ours, a crucial part of what makes for survival. So I am saying that Darwinian ideas explain the existence of human love and altruism just as well as they explain the existence of wings on birds. For survival, birds depend on wings — among other things — and humans depend on love and altruism.

This point seems to have escaped many of today's politicians, and others who seem to have embraced a very dangerous, and I would say hypercredulous, belief that More Competition is Always Better (and who will even quote Darwin to "prove" it). But for survival in the real world one can put it thus, following a recent fashion: *co-op-etition* is usually better than hypercompetition. To say this is to recognize an aspect of hard practical reality, long recognized in the scientific and more recently in the commercial world.

Co-op-etition means a balance between cooperation and competition, and it is very interesting that just now the world of big business — including some multinational corporations — are recognizing, more openly than before, that co-op-etition is often better than hypercompetition. ("Co-op-etition" is their neologism, by the way, not mine. They just recently invented the word.) In ordinary commercial hypercompetition, by contrast, you just think of competition as warfare and try to annihilate or swallow up your competitors, as documented for instance in Valloppillil *et al.* (1998) and elsewhere in the literature, in works by Richard d'Aveni for instance. I am told that d'Aveni coined the word hypercompetition. But people are realizing that, like ordinary war, hypercompetition is getting too expensive. Businesses, in other words, now seem to be recognizing more openly — in the West at least; I don't pretend to know anything about what's happening here in Japan — that co-op-etition is much more efficient and cost effective.

Now, on the longer timespans of human evolution, it's pretty obvious that co-op-etition is much better for gregarious species like our own and, I've said this already, social intelligence and all these things that keep a tribe together — including ordinary friendship, love, altruism, play, including children's play, rehearsals for real life (we call it play even though it is deadly serious) — are important for survival, hence selected for.

The genetic–memetic dynamic, the ozone hole, and instincts

So let's put the myth that "Darwin equals hypercompetition" behind us. Immediately there is an important implication: that evolution is more complex than you read about in the textbooks. For human evolution in particular, "memes" — this is Richard Dawkins' word for an infectious idea, a cultural virus, something that captures people's imagination, that replicates itself within the cognitive brain and propagates through a culture — memes as well as genes are involved, and so I prefer to speak of genetic–memetic evolution. I understand there is no exact Japanese equivalent, but "genetic" is an adjective made from "gene" and "memetic" an adjective made from "meme". Genes and memes act together. They interact in a non-trivial way, and that's completely crucial to the evolution of species like ours.

Only by recognizing some kind of genetic–memetic interaction or dynamic can we begin to explain, it seems to me, such things as the existence of a linguistic syntax-builder in genetic memory, hence the existence of Nicaraguan sign language. We are talking about an intimate interplay of cultural and genetic factors, inextricably bound together in a complex mutual interaction — an intricate dynamic operating over multiple timescales.

But wait a minute. How can this be? How can there be such an *interplay*? How can such a fast thing as cultural or memetic evolution *interact* with such a slow thing as genetic evolution?

Let us digress for a moment. Dynamically evolving systems with multiple timescales are nothing new. They are commonplace, not only in the biosphere but elsewhere in the real world. Some of them have been well studied by mathematicians and physicists. One that I happen to have studied myself, in my own professional research, is the Earth's atmosphere. Even ignoring the birds and the butterflies, we have here a system describable in almost the same way as above. It is a system in which a large number of physical and chemical processes are inextricably bound together in a complex mutual interaction — an intricate dynamic operating over multiple timescales!

For the atmosphere, the significant timescales range from femtoseconds (thousand-million-millionths of seconds) all the way to centuries. And in the case of the atmosphere, a case that is relatively well understood, there is not the slightest doubt that processes on one timescale crucially affect those on another.

For instance, if you want to understand why the deepest ozone hole is in the south (despite the pollutants causing it being emitted in the north), and why the ozone hole will take centuries or longer to disappear, then you will have to consider a subtle, and systematic, interplay of photochemical and fluid-dynamical processes over the whole range of shorter timescales. Photochemistry, the interplay of colliding atoms, molecules, and ultraviolet photons, is the fastest significant process, with femtosecond timescales. Its products, such as ozone, affect the heating and cooling of the atmosphere and the air motions that take man-made pollutants on epic journeys across the globe. Many of the air motions themselves, such as the turbulence you sometimes encounter in passenger aircraft, seem on first acquaintance to be mere random fluctuations. But some of them can, and do, have systematic effects accumulating over long times. The chief ways in which the faster processes systematically interact with the slower ones are now fairly well understood. Certain kinds of random-seeming fluctuations can have a ratchet-like character that makes their average long-term effect very significant, even over centuries (*e.g.* McIntyre 1998b).

Returning, now, to human evolution, I am suggesting that *its* multiple timescales need not surprise us even though, as individuals, we experience only the fastest fluctuations of cultural whim and fashion — many of which seem on first acquaintance to be mere random fluctuations! And the kinds of long-term selective processes I've been emphasizing demand that the genome, language ability, and cultural ability all evolved together, along with the survival-related uses of language and culture that I have mentioned already — storytelling, courtship, rhetoric, advocacy, invective, mimicry, dance, music. The rapid and seemingly ephemeral memetic processes become a subtle, ratchet-like part of the whole complex environment of selective pressures on the genome.

Every fluctuation in linguistic repertoire, for instance, has the potential to spawn new opportunities for individuals with particular gene variants to be rhetorically powerful, influential, and reproductively successful in novel ways. As Bateson and Martin (1999), and Wills (1994), and others thinking about these questions increasingly emphasize, the "environment" seen by the genome can be as much an internal and social environment as an external one.

So intimate, then, is the interplay between genes and memes, between predisposition to learn and what is learnt, that when we witness ourselves "instinctively" avoiding a head-on collision on the roadway, for instance, we should not be surprised to find ourselves talking about something that is neither completely innate nor completely learned, even though it plainly acts ahead of conscious thought. Such is almost certainly the nature of the urges and actions that we experience as "instinctive". I think it hardly ever makes sense to say that they are entirely innate or entirely learned.

We could perhaps begin to define "instincts" as potential or realized behaviour patterns that are not *consciously* learned, and not *consciously* evoked. That is, they influence our actions in some way that, at the time, comes from beyond conscious reach. And so much unconscious learning goes on — most of all in infancy, but also in later life — that it becomes a futile and meaningless question to ask whether behaviour comes from nurture or nature, as if there were a clearcut distinction. Does a lock work entirely by its key or entirely by its tumblers? That sort of lock-and-key question makes no sense at all.

In the remainder of this lecture I want to look at some aspects of our mental powers and propensities, which evolved in such ways, and under such pressures, and under such conditions. To understand something about these mental powers and propensities of ours is to make progress toward understanding ourselves. One of the most basic things to understand is the role of unconscious assumptions, and how crucially we depend on such assumptions without the slightest awareness that we are making them.

The starting point is to recall that everything we think and do depends on perception. Our survival depends, and always depended, on how we make use of sensory data from the outside world. Professor Shimojo has reminded us that significant clues about this can be found in easily-demonstrable perceptual phenomena. And because perception and cognition overlap, we learn about cognition too, by implication. If there can be perceptual illusions, then so too can there be cognitive illusions — something to remember whenever you're being sold something (Piattelli-Palmarini 1994). If I stick to some of the simpler cases, I can remind you, for a start, that there exist many highly relevant, instructive, and easily reproducible perceptual phenomena that any observant person can check for themselves in simple experiments.

The model-fitting hypothesis: self and consciousness

The most important key to understanding perception, and, by extension, cognition — including such things as where mathematics comes from — is the idea, or hypothesis, that perception works by model-fitting.

The idea is that the brain fits internal models to sensory data. In other words, ordinary perception is a sort of unconscious "science in miniature"; and science, whose job is to fit models to data, is an extension of ordinary perception. Notice that I am keeping Absolute Truth out of this; as argued in **III**, it is much more useful to talk about goodness-of-fit! Neurophysiological and neuropsychological studies indicate that the model-fitting done by the brain is a very active process, even though we are normally quite unaware of it. It is done, of course, almost entirely at unconscious levels.

What do I mean by "sensory data"? I do not mean the "sense data" of some philosophical traditions. Those are merely the subjective percepts themselves. I mean the photons of light hitting the back of your eyes, or the sound waves vibrating the drums of your ears or the molecules entering your nose. I mean the actual raw information that hits you from the outside world. I am suggesting, by the way, that we take the existence of that world for granted (by hypothesis, if you will — though the hypothesis that there exists a single outside world is a pretty good one, having stood up to a great deal more consistency-checking than most hypotheses I can think of).

The model-fitting hypothesis accounts for very many otherwise inexplicable perceptual phenomena, as will be illustrated shortly. It even gives us a simple and sensible approach to the questions about self and free will — consciousness and so on, if you prefer — as I have argued very carefully in **II**, with due attention to the positivist and other philosophical pitfalls.

In brief, even though our actual brains are committee-like, with many distinguishable "divisions" or "modules" with various evolutionary origins, the internal models to which we fit sensory data must include a single "self-model". They must do so because, to fit the data in a way that is useful for survival, my brain needs an internal model of the world and the things in it *including* the particular thing in it that is my body. Otherwise, I would become disoriented. When I avoided that head-on collision, my brain had to have a pretty good model of me in my moving vehicle, as well as models of the roadway and its side margins and the oncoming vehicle, all of which had to be put together coherently in model spacetime and fitted to the sensory data. Otherwise, I'd have had a survival problem!

This point relates to the very interesting remarks made by Professors Yoshida and Bando this morning, on what I would call the *versatility* and *expandability* of the self-model: how the boundary of self can expand for instance in the psychotherapy room, or in the concert hall — sometimes to infinity! It has been pointed out that the driver of a vehicle tends to feel the vehicle as part of himself or herself, like the clothes we wear. I once heard it said of a famously skilful Cessna pilot, "Oh, he *wears* the aeroplane."

It is the brain's internal models that give rise to subjective percepts. Mulifarious and multifaceted though the actual brain may be, the single self-model gives rise to the subjective percept of a single, conscious self, just as, multifarious and multifaceted though visual cortical activation may be (fine-grain maps, coarse-grain maps, colour maps, motion maps, etc), the internal model of the oncoming vehicle gave rise to the percept of a single oncoming vehicle.

Perception researchers sometimes refer to something called "cross-modal transfer" — a good example is given in Bateson and Martin, and more are referred to in **II** — in which an object is perceived through one sensory mode or data type, such as touch, and later recognized through another, such as vision. Monkeys, for instance, have been trained to demonstrate this. The model-fitting hypothesis easily makes sense of such phenomena. If a single internal model is being fitted to whatever incoming data are available, then the recognition of the same object via alternative sensory modes is just what one would expect. Indeed, the word "transfer" becomes superfluous: with a single model, no "transfer" is required. All along, the commonest naturally-occurring situation has been for more than one mode of sensory data to be used simultaneously in the model-fitting. There is no biological reason why the brain should not make use of all the available data. A classic example is the "McGurk" or lip-reading effect, in which the perceived sound of speech is changed by visual data.

But once again I am getting ahead of myself. What, conceptually speaking, are these things called models?

Models, computers, and mathematics

By model I simply mean a partial and approximate representation of reality. That's the sense in which the word is normally used in science. Notice that it is also not far from the sense used in connection with children's model boats, cars, houses, or what have you.

These latter are models, in the same sense, as well as being real objects. More precisely, they can be regarded as models in this sense because a model house can be regarded as a partial and approximate representation of a real house. The essential points are that it represents some aspects of a real house but not others, and that everything is approximate.

When I was a small child I used to be fascinated by model houses that had electric lighting. But I never saw one that had a roof that could keep the rain out in the same way as a real house. And if you wanted to represent that correctly, the fluid dynamics wouldn't scale, so you'd have a problem. So that's the first point. Model houses, boats, or whatever, are exactly this: *partial and approximate* representations of reality, if you mean by reality the real full-size boats, cars, houses, or whatever.

Now these days you can have practically the same things in what is charmingly called a "virtual reality" system. You can run a computer programme and have some kind of video device or headphones or goggles or gloves, and you can see and perhaps even touch the model houses etc. in the virtual reality system. As a matter of fact, in their present state of development, such model houses often look rather similar to, or even simpler than, the wooden toys that children often have. Virtual-reality houses are a lot simpler than the real thing, so they are another illustration of exactly the same idea. They are partial and approximate representations of reality.

But remember now that these models are made of computer code. They are not made of wood or plastic. And notice also the following point: if you inspected that computer code, it would not directly tell you that it represented a model house. It wouldn't be at all obvious. You would see a bunch of symbols and strange characters and arcane-looking words. The point I'm making is that the *same model* can have *different representations*, and that it need not be obvious that the different representations are equivalent to each other.

Likewise, models can equally well be made of neural patterns and allosteric protein circuitry. Protein molecules are often allosteric, which means that they behave like logic elements or miniature transistors.

By the way, textbooks about the brain tend to give the impression that brain function is all in the neurons and synapses, interconnected in a network in which synapses are thought of as terribly simple things, a bit like transistor inputs. But we now know that this is only the tip of a hypermassively parallel computational iceberg. There are enormous numbers of allosteric protein molecules inside even one synapse. A single synapse is beginning to look more and more like a big parallel computer in its own right, as people look in more and more detail with ever-improving experimental techniques. It seems practically certain that in the next few decades the picture we have of how the brain works is going to get immensely more complicated than it is now. (This point seems to have escaped researchers who claim, on the basis of the simple network picture, that artificial intelligence will soon outstrip the human brain.)

Anyway, that's one way to make models, in the relevant sense. Here's another way. Models can also be made of mathematical equations.

That's almost obvious from what was just said about virtual reality. Mathematical equations are, for this purpose, not too different from computer code. Mathematical equations can be regarded as instructions to compute. Of course, there are non-trivial differences. Mathematics is a very beautiful thing that allows you to handle infinite numbers of computations. So in a very significant way it goes further than computer code. But for the purposes of this argument it's not really all that different. Whether one thinks of a finite or an infinite number of computations is for our purposes only a detail.

So it shouldn't surprise us that models can be made of mathematical equations. There is an essay, that every physicist knows of, I think, by the famous physicist Eugene Wigner who talked about the "unreasonable effectiveness of mathematics" in describing the real world. But what's surprising is not the fact that mathematics comes in; what's surprising is the fact that very simple mathematics comes in when you are building models of sub-atomic Nature. It's not the mathematics that's unreasonable; it's the simplicity. So I think Wigner should have talked about the "unreasonable simplicity of sub-atomic Nature". It just happens, and I don't think anyone knows why, that at the level of electrons, protons, and neutrons, the sub-atomic level, things *are* very simple. That's just the way nature seems to be at that level. And of course that means that the mathematics is simple too.

After all, if I were talking about a house that happened to be perfectly spherical, and perfectly smooth — I admit that it wouldn't be much use as a house — then the mathematics to describe it would also be simple. All you would have to say is the house has a certain radius. The rest could be expressed by arguing from symmetry. You could say the spherical house has a certain radius, and that it looks the same from all angles. That's enough — you've then described the whole thing. Well, an electron, say, is a bit like that, with one slight difference: an electron can look just *two* different ways, no matter what direction you look at it from. This is what's called "quantum-mechanical spin one-half". It is strange but nevertheless simple, nearly as simple as a perfectly spherical house.

Let me continue this discussion of what models are, for just half a minute more, because the next example is very remarkable. Figure 5 shows what we can do with these same simple models of electrons and other similarly simple entities, including the entities called photons. This is the famous graph of what's called the Planck function. The Planck function has a very simple equation, which you can find in any physics textbook. It is simple because it comes from a model built from simple entities. This model is called the quantum theory of blackbody radiation, if you want the technical term. It predicts the intensity of the radiation as a function of spatial frequency or inverse wavelength, that is, the number of waves per centimetre. In the case shown in figure 5 we are talking about photons of millimetre wavelength.

Our present models of the the whole universe, the so-called big-bang

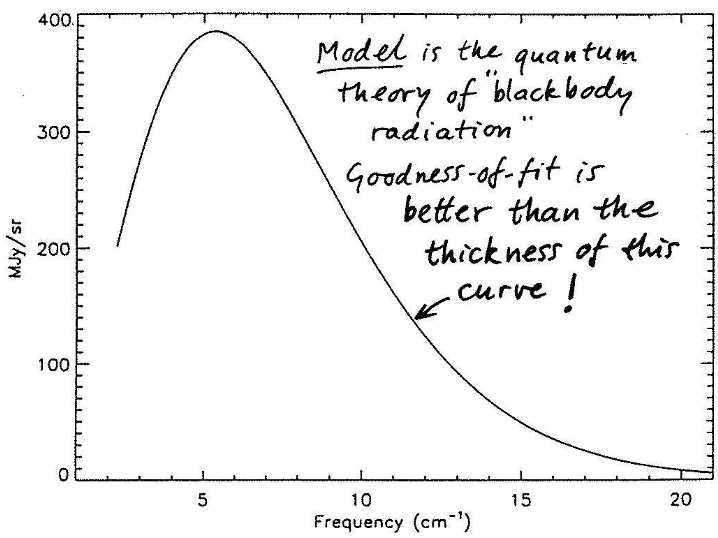

Figure 5: Model and invisible data: the Planck function plus invisibly small data error-bars. The data are measurements of the cosmic microwave background radiation from the famous COBE satellite. The agreement is much better than the thickness of the plotted curve! From Fixsen *et al.* 1996. To convert the values at left from megajansky per steradian to SI units, W $m^{-2}(m^{-1})^{-1}sr^{-1}$, multiply the values by $2.99792458 \times 10^{-12}$

cosmologies, predict that the Planck function should accurately describe the cosmic microwave background radiation, the sea of low-energy photons that isotropically fill the universe. Now there was a remarkable project to launch a satellite and measure this radiation extremely accurately. The satellite was named COBE, not to be confused with the city of Kobe here — COBE for "COsmic microwave Background Explorer". If you want to read about the project, there's a a beautifully written popular account by Mather and Boslough (1996). The results of the measurement gave us a very wonderful example of goodness-of-fit of a model to data. The Planck function fits the data to a much better accuracy than the thickness of the plotted curve in figure 5! There are no experimental points visible: you can say that there *are* experimental points in the figure, but hidden well inside the curve. I find this awe-inspiring. We really do know something about the universe, and how it works physically, from models like this — even though we cannot, if we are behaving responsibly, pretend to know the Absolute Truth.

“Walking lights” and organic change

Now let’s come back to human perception, and the model-fitting hypothesis for perception. Figure 6 shows twelve dots. When these are animated (a video of this was shown in the lecture; the animated version can also be seen on my home page, http://www.atm.damtp.cam.ac.uk/people/mem/) anyone with normal vision will immediately see what looks like a person walking.

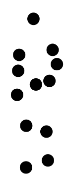

Figure 6: Cf. animation on http://www.atm.damtp.cam.ac.uk/people/mem/

This “walking lights phenomenon” is a classic in experimental psychology, though not well known outside the experimental-psychology community. It was discovered by Professor Gunnar Johansson and his colleagues, one of whom, Professor James Maas of Cornell University, kindly supplied the data whose digitized form was used in the animation. The phenomenon has been thoroughly studied, along with a host of related perceptual phenomena under the general heading of “structure from motion”. It is a compelling example of something that can hardly be made sense of without the model-fitting hypothesis — that is, can hardly made sense of without the supposition that the brain is unconsciously fitting a model to the data.

What are the data? It takes a conscious effort to recognize what a tiny amount of information we are talking about. The significant data are just twelve moving points on a two-dimensional surface. The surface is the screen in front of you, or equivalently the retina at the back of your eye.

What is the model? It is non-trivial. It involves unconscious assumptions of a very specific kind. It assumes a particular kind of motion — linked, piecewise-rigid motion — in three dimensions. The brain fits this model — or rather, a member selected, somehow, from this enormous *class* of models — to the data, to the twelve moving points. It does so ahead of conscious thought: the impression of a person walking is there in your mind before you have time to think. It actually takes a time of the order of a few tenths of a second, according to the experimental studies that have been done, even though the class of models to be selected from is combinatorially large.

It's a very robust and clearcut example. For anyone with normal vision, the model fitting takes place whether you like it or not. It's completely "instinctive"; and the percept of a person walking is completely unambiguous.

Notice incidentally that this example also illustrates what I called organic change, and its relevance to detecting lifelike behaviour. Some things are changing, continuously, while others stay invariant. Organic change is a property both of the sensory data and of the model being fitted. For instance the model has an invariant number of links and rigid members joining them, but the positions and angles are changing continuously.

Biocomplexity

You might well think it impudent of me to talk about this model-fitting process so confidently, when no-one has any detailed mechanistic understanding of how it is actually done by all the tens of billions of neurons and the enormously greater numbers of synapses and protein molecules (and intracellular organelles, microtubules, etc. etc.). On a quasi-mechanistic level we are talking about something so complicated that it is practically impossible to begin to imagine even the degree of complexity, let alone claim to understand anything much about how it works.

I want to digress for a moment to try to underline this point about complexity, which often seems not to be sufficiently appreciated or acknowledged — as starkly shown by all the irresponsible talk about genes for this, genes for that, and, most horrifying of all, "designer babies". Even what we call relatively "simple" life forms are still far, far too complicated to understand in detail. They are not simple machines; still less is the human child's mind a simple blank slate, as some education auditors think. Take the famous bacterium *Escherichia coli*, less than a thousandth the mass of a single human neuron. We now know the complete genome, to the extent of having more or less complete samples of *E. coli*'s DNA sequence in a few cases. But all this tells us is that even a single *E. coli* bacterium is a dauntingly complex dynamical system, with a combinatorially large number of internal states, fluctuating, reconforming and evolving under thermal agitation on picosecond timescales – vastly too complex for us yet to simulate in detail, even with today's most powerful supercomputers, let alone to understand in detail. We can't yet even reliably predict the function, or even the shape or conformation, of a single protein molecule knowing only the DNA sequence of the gene that produces it. And a single *E. coli* bacterium contains millions of protein molecules of up to thousands of types, many of them acting as computational elements — as allosteric circuit components — in some kind of massively parallel information-processing operation.

So, contemplating an *E. coli* in our present state of knowledge is a bit like contemplating the Internet when you have just found out how wires and transistors work and programs stored but know very little about programming languages, and nothing at all about multi-scale computational architectures and interacting layers of software. One example among many that underline the point about "programming" and "software" is a phenomenon called PTGS, post-transcriptional gene silencing. We know it happens and that it is important for survival, but we have no idea how it works. It can nullify the effect of a gene even if that gene is not switched off.

Model-fitting, prior probabilities, and repertoires of components

Despite such considerations of biocomplexity, the walking lights phenomenon and many other examples of perceptual phenomena (some of them discussed in **II**) allow us to make some reasonable general statements about the kind of thing the model-fitting process must doing — what some of its characteristics must be. These are evident from the nature of the perceptual phenomena without knowing the mechanistic details. I have discussed this more carefully and at greater length in **I** and **II**.

First of all, the model-fitting process must not only be capable of efficiently optimizing goodness of fit, but must also, somehow, be able to cope with a combinatorially large tree of possibilities. This is a very remarkable computational feat. No computer scientist has any idea, I believe, how it is done at the kinds of speeds relevant to survival. (That is part of why computers don't yet drive taxis.)

Let me remind you again what combinatorial largeness means. Here's an example: if you make even a structure as simple as a linear one-dimensional chain with ten different links, the number of ways to do that is more than three and a half million. If you make a chain with a hundred links, the number of ways to do that $\sim 10^{158}$, which means 10 with 158 zeros after it:

100,000,000,000,000,000,000,000,000,000,000,000,000,000,000,
000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,
000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,
000,000,000,000,000,000,000,000.

Now it is one thing to write such a number down, but quite another to imagine, intuitively, how large it is. I would defy anyone to do so. That's the point about combinatorial largeness: you can't begin to imagine the actual magnitudes, and you can't enumerate the possibilities one after the other, even at supercomputer speeds. More realistically complicated model structures that can somehow represent, for instance, what you commonly see and make sense of in a three-dimensional visual scene involve far greater numbers of possibilities still — so great that it becomes nontrivial even to write such numbers down, let alone to imagine their magnitude.

So the brain must have ways of drastically pruning the combinatorially large tree of possibilities. It has to solve the model-fitting problem fast enough for survival purposes, so it needs to be very fast at rejecting huge numbers of unlikely cases. It also needs to find, whenever possible, a unique solution to the model-fitting problem. In principle, many models will fit one set of data (what a mathematician might call the "nonuniqueness of inverse theory"). The tree must be pruned, if only to get uniqueness. We may think of the pruning as equivalent to the assignment of Bayesian prior probabilities — call it prior prejudice if you will. The walking lights demonstration nicely illustrates that point, showing among other things that piecewise-rigid motion in a three-dimensional space is given a high prior probability.

Another way to say the same thing, and in a way that might be closer to the actual biological reality, would be to say that the brain must, in some rather abstract sense, have a *repertoire of pre-existing model components* (whose seeds, at least, are in genetic memory). Something to represent the class of piecewise rigid, flexibly-linked objects would be an example, including within itself a smaller class of "bio-motions" or strongly lifelike piecewise-rigid motions, to which especially high prior probabilities may tend to be assigned. That may be part of why the walking lights demonstration is so reliable and robust. Conversely, models that can't be made from any of the available components have, in effect, been given zero prior probability. Such models simply aren't available to be fitted to the data.

Of course the prior probabilities might be wrong in a given case. The twelve moving points need not have been a person walking: they could have been a bunch of fireflies. They could have been any old moving points that had nothing to do with anyone walking. But the brain made an unconcious assumption that it was looking at piecewise-rigid motion in three-dimensional space, and chose to fit that kind of model to the data — all done ahead of conscious thought — and so you had the percept of the person walking.

One can imagine other cases of moving points in which piecewise-rigid motion does not fit the data. Then the brain has to try something else, such as fireflies. If no model from the available repertoire succeeds in fitting the data, then the brain has to give up. In the latter case you experience a feeling of disorientation or visual confusion. (See for instance Roger Shepard's "impossible elephant" picture in **I**.)

The model-fitting process can't afford to waste time and effort with models that are incapable of fitting things in the outside world. So the model components and the models into which they are assembled must be coherent and self-consistent, like a good scientific theory, and must take account of general properties of the outside world, such as its local Euclidean geometry. Thus the model-fitting brain must have, among other things, an unconscious knowledge of Euclidean geometry, built into its repertoire of model components. This knowledge of Euclidean geometry is one example of what I meant earlier by "unconscious mathematics".

(How much of this unconscious mathematics is innate and how much learned, or calibrated, in infancy, is another of those lock-and-key questions that makes no sense. It would probably make more sense to say that genetic seeding and genetic propensity-to-learn closely interact with incoming visual and tactile data as the infant gropes around. We do know that unconscious learning is a significant *part* of what is involved, from cases in which children are born with cataracts or other visual obstructions removed later in life. Such people do not have normal vision; indeed their experience suggests that they lack most of the repertoire of three-dimensional model components.)

Another example of unconscious mathematics is calculus, the mathematics of rates of change. This might surprise you; but I can easily show it through another classic demonstration (figure 7). If you look at the inner edges of the black marks, you see a smooth curve, a sort of white or white-ish edge with a perfectly smooth curvature. Again, this seems to be a robust and reproducible perceptual phenomenon: anyone with normal vision sees the smooth curve. By the way, we were reminded this morning that people

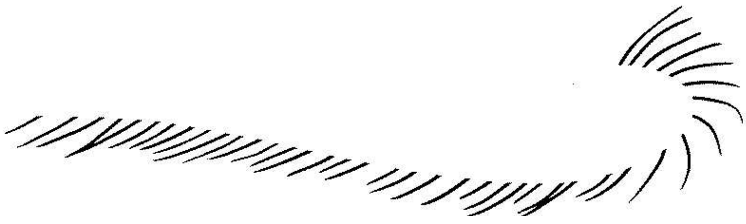

Figure 7: Illusory contour (the inner curve).

do vary in some perceptual functions; but this, like the walking lights, is one of the very robust perceptual phenomena that seem to vary very little. I once met someone who saw a yellowish curved edge rather than a white one, but that's an irrelevant detail for present purposes.

Psychologists call the smooth curve an illusory contour, or subjective contour. The word illusory is perfectly fair, in that there isn't actually a curve there. The curve is put there by your visual system. If you are a mathematician, you will recognise that the visual system is fitting a smooth curve to the ends of those black segments. It's solving — ahead of conscious thought — a problem in the so-called calculus of variations. It is here that we begin to glimpse the Platonic, that timeless world beyond the everyday, a world of perfect forms and ideas, a world of truth and beauty, where perfect circles are truly perfect and where infinitely greater wonders await discovery.

My last example in this series is an auditory one, again illustrating the point about repertoires of model components and also giving us, as a by-product in this case, another clue to the question of what music is and why it exists. The example is the so-called harmonic series of a periodic vibration. The harmonic series or equivalent time-domain information is another entity, or pattern, that the brain must have in its repertoire of model components. This is simply because our ancestors had to survive in what you might call a jungleful of animal sounds. It was important to be able to recognize the yowls, hoots, and whistles made by birds and animals, members of your own and other species. There is survival value in being able to tell who is emitting what sound. So our brains must have models of sound *sources* and sound transmission that can be fitted to the auditory data and hence yield information about the sources.

One important class of sound sources is characterized by approximately periodic acoustic vibrations, such as can be produced by organs like my vocal chords, or a bird's syrinx, which are fundamentally similar for this purpose — oscillators that vibrate almost periodically, *i.e.* with a regularly repeating vibration cycle. Here's a children's toy bird that does the same thing. It produces just such a sound through the almost-periodic motion of a vibrating reed. (The toy bird was sounded in the lecture, making a loud squawking sound.)

That's in case anybody was asleep! But you see the point: we are able to recognise such a sound as probably coming from a single source. It might be a real bird; it might be a toy like this.

So the brain's repertoire of model components for this purpose has to include information about the harmonic series or the equivalent time-domain information or both. The harmonic series is a particular pattern of frequencies (demonstrated on the piano during the lecture). This pattern is present, implicitly, in any sound made by a almost-periodic vibrator, like the reed in my toy bird. It contains important information about the mechanics of the repeating reed vibrations.

So it is no surprise to find that the pattern of frequencies represented by the harmonic series is a building block in music, too. It is present in the music of all cultures that I have ever heard — very clearly, for instance, in Japanese koto music. This is further discussed in Note 58 of **I** on musical harmony, and its extended Internet version. Most of what's in the musical harmony textbooks, and quite a lot that isn't in them, can be constructed or reconstructed from the organic-change principle plus recognition of the harmonic series' special significance for the ear-brain system.[3]

Once again, we need not debate whether the ability to recognize squawks, yowls, hoots, and whistles is fully formed at or before birth, or whether unconscious learning from auditory input has a role (before or after birth), which probably it does. One way or another, survival compels the ear–brain system to perform feats of pattern perception in which, inescapably, the harmonic series or equivalent information, internally represented in one way or another, has to be one of the most crucial patterns the system must deal with. And where, then, does musical pleasure come from? I think part of it

[3]Strangely, all the music-acoustics textbooks I've seen, and even Helmholtz' great treatise on the subject, completely miss this basic point — as is clear from their discussions of what is misnamed "just tuning", which gives *bad* tunings of frequently-used chords.

lies in the relation between survival and play. What we call play is a deadly serious rehearsal for real life. To make music is, in part, to *play* with these very same patterns. It is for survival's sake that the brain takes pleasure in playing with them. That, I suspect, goes some way toward explaining why music exists at all.

Mathematics, music, and the unconscious power of abstraction

As already emphasized, the model-fitting hypothesis implies that perceptual processing handles, somehow, a combinatorially large number of possibilities. Even after drastic pruning of the tree of possibilities, such processing must still have ways handling vast numbers of possibilities at once. This requires, or one could say *amounts to*, an unconscious power of abstraction.

Abstraction means no more and no less than the ability to handle many possibilities at once. Abstraction has to be a property of the repertoires of model components.

We have already had various hints of this. Take for instance the phenomenon of illusory contours illustrated in figure 7. As noted in **II**, "their geometrical properties, including their remarkable smoothness, are abstract properties — general enough to include vast numbers of special cases of smooth object outlines, from the edge of the moon and the edge of the sea to the curve of a perfect banana or elephant's tusk, the arch of a leaping dolphin, or the outline of a hanging grape or drop of water. Here is a Platonic seed [part of the unconscious, primeval origin] of the consciously abstract ideas that mathematicians call curvature, continuously-turning tangent, rate of change, and extremum principle, part of what is called differential geometry and the calculus of variations in space and, by extension, in spacetime.

"Indeed, we directly experience the extension to spacetime. We feel it kinaesthetically, through our sense of continuously evolving motion, for instance through athletic grace in humans and in other creatures and through the subtlety and elasticity of dance and music. In figure 7 the shape and clarity of the illusory contour is sensitive to the precise locations of the ends of the black segments. In a musical performance the kinaesthetic feel, the motion and continuity, is similarly sensitive to the precise timings of the onsets of sounds. Even so gifted a musician as Mozart had to earn, by hard work at the keyboard, his legendary ability to make fortepiano passages sound like 'flowing oil'."

Here is another example, one of the most basic and pervasive aspects of perceptual phenomena, carrying important hints as to how the combinatorial tree is pruned. This is the phenomenon that psychologists call perceptual grouping. Figure 8 shows the case described in **I**: "It consists of thirteen identical dots, which are automatically seen as four main groups. The first

Figure 8: Visual grouping by proximity

group, on the left, is seen as two pairs, and the second group as a pair plus a single. The last group is seen as three singles... The third group has some ambiguity... the brain is using spatial proximity as its cue."

Continuing as in **I**, "Perceptual processing is flexible... If the pattern is animated and the leftmost three dots start to move, performing a rigid rotation or translation, or both, then those three dots are immediately regrouped as a triplet despite their uneven spacing. There is an English phrase to describe the phenomenon: the three dots 'move as one'. Such grouping phenomena are basic to structure-from-motion perception. They seem to be part of how the brain uses [Bayesian] prior probabilities, drawing on genetic and non-genetic memory, to prune the enormous combinatorial tree of possible internal models while trying to maintain consistency with incoming sensory data."

In the lecture, an animated version was demonstrated, beginning with the static figure 8 as above but then showing the leftmost three dots moving as one, in vertical translational motion, causing them to be immediately regrouped as a triplet.

Perceptual grouping is yet another manifestation of the unconscious power of abstraction. As noted in **II**, "Relative spatial proximity, 'motion as one', and 'belonging to a group' are abstract properties. They are abstract in the sense of being general enough to include or exclude vast numbers of special cases, hence vast numbers of model-fitting possibilities..."

"Perceptual grouping is recognizably a Platonic seed — part of the unconscious, primeval origin — of the consciously abstract ideas of set theory and integer used in mathematics. These ideas include for instance a recognition that the sequence of positive integers $1, 2, 3, ...$ has no end. One can think of adding 1 to any positive integer, arbitrarily chosen, or equivalently of adding one object to any group of objects arbitrarily chosen. Though simple, the example is typical of mathematical thinking."

Notice the implications for the origins of "abstract cognitive symbolism", sometimes thought to have suddenly arisen at the onset of the Upper Paleolithic (recall top left of figure 3). The foregoing arguments suggest not only where mathematics comes from, but suggest also that the roots of mathematics and logic, of the Platonic world, the world of perfect forms and ideas "already there", as it is sometimes described — and abstract cognitive symbolism generally — lie far deeper, and are evolutionarily far more ancient, than they are usually thought to be.

II gives some further discussion of these matters, including comments on the proof of Fermat's Last Theorem and on Professor Penrose's ideas about mathematical truth.

Some further demonstrations

At around this point in the lecture, I showed several more examples from the standard repertoire of visual perceptual phenomena: (*a*) the barber's pole illusion, a case of robustly and persistently fitting the "wrong" model, axial rather than rotatory motion; (*b*) the rotating Necker cube, illustrating the active flipping between two metastable models with comparable prior probabilities that both fit the data; (*c*) Roger Shepard's "impossible elephant", a variation on the theme of Escher's "impossible triangle", to which the brain will fit no stable or metastable model whatever, and (*d*) a standard demonstration of so-called "gap filling" (figure 9). There is an ongoing philosophical debate as to whether this should be called gap filling or not; but the debate is — I would dare to suggest — neatly circumvented by the model-fitting hypothesis. The data may have gaps, but the model fitted does not.

This last demonstration (*d*) was also by way of a joke to demonstrate wrong prior probabilities. Figure 10 shows what is revealed when the fence in figure 9 is removed.

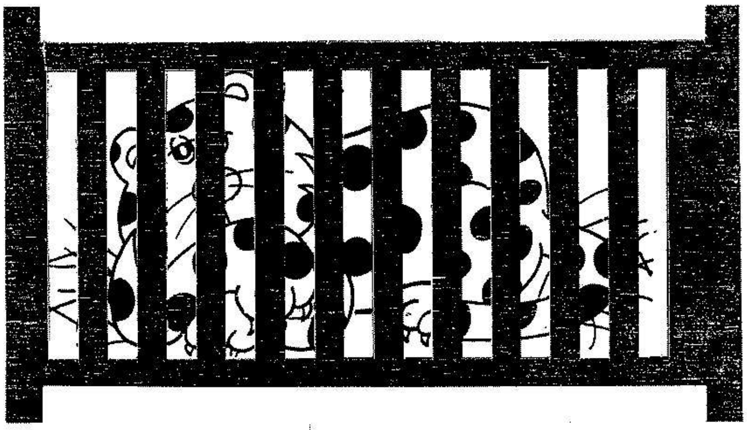

Figure 9: Leopard behind bars?

Acausality illusions

Acausality illusions, or "illusions of acausality", are perceptual phenomena in which causes seem to occur after their effects. These have sometimes been regarded as mysterious; but the model-fitting hypothesis dispels the mystery.

The phenomena make perfectly good sense in terms of the model-fitting hypothesis and standard physics because perceived time in the brain is — and for reasons of coordination *has* to be — an internal model property, and not a physical variable. (Think once again about avoiding a high-speed, head-on collision.) A careful discussion, with literature references, is given in **II**. Here I will be brief, simply presenting two outstandingly clear and simple examples of musical acausality illusions, consisting of a quotation from Mozart and a slight variant of it, shown in the top and bottom halves of figure 11. Audio clips are available via my home page in the form of .wav files (and hopefully MP3 soon, as well), of the piano solos and the orchestral accompaniments referred to below, all of which were demonstrated in the lecture.

As in **II**, "consider... the change of harmony at the third bar, shown by an arrow, in the Mozart example [reproduced in the top half of figure 11]. Even at the fairly slow tempo shown by the time markers, the harmony change is perceived to occur at the time of the arrow, even though the information defining the change is entirely contained in the following two notes, the first

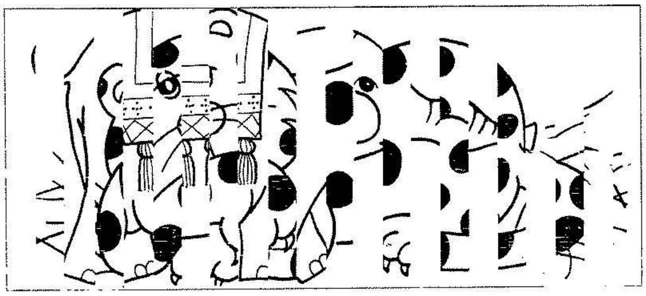

Figure 10: Leopard–elephant?

of which does not sound until one-third of a second later. This acausality illusion depends only on the listener's general familiarity with the 'language' of Western tonal music. It does not depend on the listener's having previously heard the particular piece of music, any more than understanding a sentence of speech depends on having previously heard the particular sentence. The perceived time of the harmony change is especially clear to a musically trained listener because, as with sports training, musical training cultivates an acute awareness of perceived times through the need to monitor and control timings in performance...

"The perceived timing of the harmony change at, not after, the time of the arrow in [the top half of figure 11] is confirmed by the way in which a composer or arranger skilled in Western tonal music would devise an or-

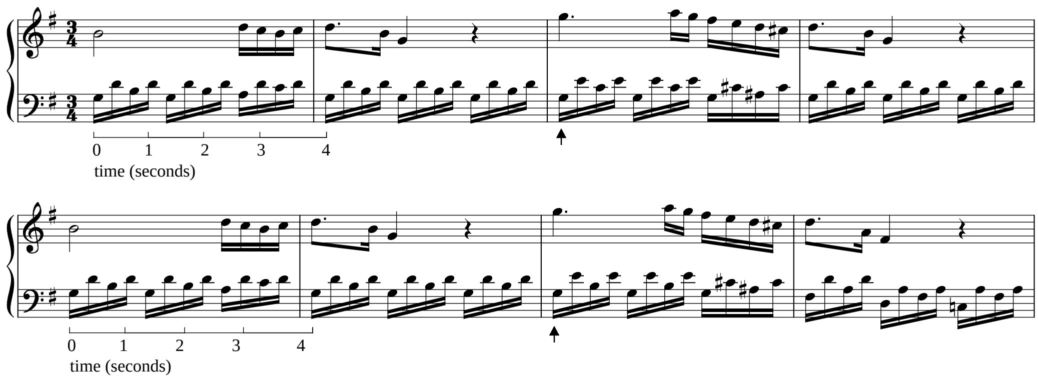

Figure 11: Mozart K 545 and variant

chestral accompaniment for this example, if intending to preserve its musical sense. Such an accompaniment would change harmony at, not after, the time of the arrow...

"If you happen to be familiar with the particular Mozart piano solo quoted in [the top half of figure 11], and think that such familiarity might be influencing how you hear things at the time of the arrow, then try listening instead to the piano solo shown in [the bottom half of figure 11]... There is a new harmony change, strongly audible and very different from Mozart's. But its perceived timing still corresponds to the time of the arrow, even though the example is identical to Mozart's for a full two-thirds of a second beyond the arrow. A suitable orchestral accompaniment... would again change harmony at, not after, the time of the arrow."

Epilogue: public understanding

I want to return briefly, now, to the public understanding of science and related matters. This is one of my concerns as a professional scientist, and it must surely concern anyone who cares about our future and our children's future. It is one of the reasons for my own journey toward human self-understanding. And self-understanding should in turn give us, among other things, a better understanding of the power and limitations of science.

As I wrote in **II**, a good metaphor for our time, especially for those demanding Absolute Certainty and Absolute Truth from scientists, is "science as our eyes and ears on an uncertain future". In this lecture, and more extensively in the published papers **I**, **II**, and **III**, I have given reasons for thinking that such metaphors are far from superficial. Not only can science be viewed as an extension of ordinary perception — with, of course, suitable precautions and cross-checks — but such an extension is fundamentally what science is. This is because ordinary perception, as well as science, works by fitting models to data from the outside world.

Like ordinary perception, science can make certain aspects of the outside world vividly clear to us, and can do so with impressive, indeed awesome, accuracy in some cases. Figure 5 is one example. Conversely there are, and always will be, aspects that are unclear, illusory, or out of sight altogether, beyond the research frontier.

What is it like to be a research scientist? With science as an extension of ordinary perception, one can put it like this. Trying to make sense of things near the research frontier will always be — fundamentally — like driving in the fog. It is, and will be, like driving a vehicle in swirling fog on an unfamiliar, unmapped, twisty road with many branches, and with plenty of oncoming traffic and potential collisions. One has to live with uncertainty, one has to keep one's eyes and ears open, and one has to expect surprises.

Does it make sense to harass, criticise, or sack the driver who admits to uncertainty about what lies ahead? It might be a better idea, in the real world, "Star Wars" notwithstanding, to sack the driver who shuts his eyes and blocks his ears and claims infallible prior knowledge of what's ahead, who sets rigid targets and objectives and sticks to them regardless of what emerges from the fog. This is accurately relevant to public issues involving science.

It is relevant for instance to the ongoing fight against new diseases, and — to take an example where I have some specialist knowledge of my own — it is relevant to the problem of monitoring and predicting environmental change, for instance whether the sea level will rise, and if so how quickly, and by how much, over the next century or two, say by 3 cm or 300 cm, or whether deep ocean circulations could flip chaotically and devastate local climates and economies. For well understood reasons, there is great uncertainty over questions like these last two: we are indeed, collectively, driving in the fog, and the "vehicle" at risk — strong in some ways and fragile in others — is our planetary life support system. We had better keep our eyes open.

A precondition for doing so will be the survival of open, independent science alongside commerce and commercial science. Open science is needed not only because of hazards we already know about, but also, even more, because of the combinatorially large numbers of unforeseen and unforeseeable ways in which things can and will go wrong. The number of such possibilities will increase very steeply as technological systems become ever more complex. One reason I have been quoting Valloppillil *et al.* (1998) — the well known Internet document called *Halloween I* — is its unique importance as a testimony to the superiority of the methods and ethics of open, independent science when dealing with a combinatorially large unknown.

Halloween I is actually about computer software. But there are clear implications for the safety and reliability not only of today's vast software systems, many of them critical to business and finance, but also of tomorrow's systems of genetic engineering, and of other complex systems. Understanding these issues of safety and reliability is very much tied up with understanding the power and limitations of science.

The survival of open science in the teeth of commercial, legal, and political pressures cannot be taken for granted. But the need for it is increasingly well understood, within as well as outside the world of commerce. Openness and independence are preconditions not only to scientific competence in general, but also — as *Halloween I* testifies with rare cogency — preconditions to the best possible engineering, to maximizing the safety, reliability, and security of the complex technological systems on which we depend, whether connected with computer science and electronic commerce, with quantum physics, with molecular genetics, or with any other field.

All this may sound paradoxical to the man in the street. What has the survival of open science got to do with commercial success and reliable technology? Isn't open science some kind of cultural luxury? Why shouldn't the most powerful software be developed in conditions of commercial secrecy, rewarding those who create it? Why shouldn't the same principle apply to life-saving genetic medicine? How can secrecy not be the best form of security? Why should open, independent, dangerous, expensive scientific thinking have anything to do with such practicalities, and with hard commercial realities? Again, the reason is combinatorial largeness, something that the man in the street does not yet understand.

It is *despite* combinatorial largeness — against all the odds — that there are such things as reliable scientific knowledge, and complex yet reliable computer software. Such reliability has always depended on the collective model-fitting, the "massively parallel problem-solving", as *Halloween I* calls it, made possible by the existence of open international scientific communities. These are cross-cultural groups of people driven more by the urge to understand or improve something than by commercial pressures or political agendas or tribal loyalties, and whose interest in a problem has been ignited somehow and whose reward is the prestige of helping to solve the problem. The chances of noticing the unforeseeable are then enormously improved.

That lesson was first learnt during the Renaissance, after centuries of scant progress with alchemy and the like. It is now being learnt again in the world of commerce, as the problems to be solved become ever more complex and difficult. Open communities with their interest ignited have problem-solving abilities of an order described by *Halloween I* as "simply amazing", when seen from the viewpoint of an organization powered by huge financial incentives but constrained by commercial pressures and commercial secrecy. That such an organization should have recognized, contrary to its own in-house culture, that openness and independence can be more powerful than financial incentive may yet be seen as a turning point in the development of advanced human societies — though it will be a close-run thing. One need hardly add that open, independent communities are the only source of credible scientific advice when things go wrong.

As my remarks about *E. coli* may suggest to you, genetic engineering is going to be a far more subtle and complex matter than computer programming. Safety and reliability will be of public concern for the foreseeable future; and, especially if present trends continue — in patent law for instance — today's unreliable commercial software could well be followed by tomorrow's unreliable genetic engineering. We might even have unreliable Earth systems engineering, either by design or, as at present, by accident. For these reasons alone the survival of open, independent science as a distinct entity, alongside commerce but independent of commerce, can no longer be described as some kind of luxury. The survival of open science is a basic and urgent necessity, as *Halloween I* has inadvertently illustrated. That, above all, is why the public understanding of science is important; and if this Symposium and this lecture contribute to such understanding, directly or indirectly, in the slightest way, then they will not have been held in vain.

In my youth I nearly became a full-time professional musician, as hinted earlier. One reason I chose science instead was not only because I had some talent for it but also because of the vicious and destructive hypercompetitiveness found in parts of the music business: the idea of competition as warfare, the dirty tricks and the winner-take-all culture, the shameless exploitation of young artists, now documented in the book *When the Music Stops* by Norman Lebrecht. I cannot vouch for all the details in that book, but the nature of the beast comes through recognizably. Science as a profession attracted me and many others because, by contrast, it had a remarkable "gift culture" one of whose key attributes is the scientific ethic — a kind of chivalry or code of honour between competing colleagues, aspired to if not always attained, and able to keep our natural competitiveness, and alchemical secretiveness, sufficiently within bounds to permit a degree of openness, cooperation, and massively parallel problem-solving.

Good, credible, reliable science has always depended on the existence of that same culture and ethic, whose effectiveness first became apparent during the Renaissance. Over decades and centuries, the effectiveness has been demonstrated in countless ways, including the ignition of interest and the recruitment, again and again, of prodigiously talented people for modest remuneration. The building of complex yet reliable computer software is only one illustration among many.

But the culture, the ethic, and the recruitment are now in great peril. This in turn imperils our future problem-solving abilities and the availability of independent thinking and advice on matters of public concern. Parts of the science enterprise are becoming more and more like the professional music business writ large, with commercial and political forces spurring hypercompetitive behaviour.

An example, if you need one, is gene patenting, or DNA sequence patenting, the attempt to claim ownership of naturally-occurring genetic information and all its future uses — logically no different from patenting the Earth. It is not the patenting of inventions, but the patenting of naturally-occurring patterns. This is not only an unprecedented kind of patenting, but also a new hypercompetitive weapon of mass destruction — destruction, that is, of all competition and independent research — a weapon now going into its intensive testing phase in the law courts and elsewhere. I'm among those who hope the weapon won't work.

The politically influential book by Terence Kealey (1997) has gone so far as to urge that all science and engineering should be done under commercial pressure, giving openness and independence little chance of survival. This bears thinking about each time your computer system crashes and your customers are turned away, or a nuclear power plant goes out of control, or another disease agent jumps a species barrier. The belief, largely unconscious, that More Competition is Always Better is arguably, as I suggested earlier, one of the most dangerous hypercredulous beliefs of our time.

Conversely, those advanced societies that find new and effective ways to value and protect the scientific ethic and other, similar professional ethics — notwithstanding their pricelessness, their invisibility to the audit culture, their unquantifiability as "performance measures" — will have a powerful long-term advantage through balancing competition with cooperation.

This then is the central paradox of the crisis in science and of the larger crises in democratization, commercialization, globalization, and auditing. It is a paradox increasingly recognized, as already noted, within the commercial world itself, reflected in the current talk about "co-op-etition", "work satisfaction", "sustainability", "partnerships", the "third way", and "alliances not takeovers". By limiting hypercompetition, a society can become *more* competitive.

I would dare to suggest that it can become more civilized as well, and spiritually healthy. Perhaps it could even aspire to being what the Mahatma Gandhi would call civilized. The optimist in me hopes you agree! Perhaps we shall have our new covenant, and a degree of healing of the division between scientific and artistic creativity.

I don't see why not. It seems to me that scientific understanding can only increase, not decrease, our awe and wonder at the nature of things, at the miracle of knowing as much as we do know, and at the promise of yet deeper knowing and understanding. Knowing that poetry and music can be found in acoustic time series, and in binary numbers etched into an optical disc, must surely enhance, not diminish, their power to move us. Knowing that science and the arts both have a deep biological significance, that they have to do with genes and our ancestors' survival, with juvenile play and with education that works, and with the marvellous and multifarious developments that we call culture — knowing that nature and nurture intimately and subtly work together, that nurture is part of nature, in infancy and beyond — surely all this should enrich, not impoverish, human life. Admitting that the known laws of physics are only approximate, though exquisitely accurate, and that they point to deeper mysteries not yet fathomed, must surely increase, not decrease, our sense of the grandeur of the universe.

I think we shall dare to understand, even better than today, the nature and origin of our instinctive powers. It could begin with schoolchildren puzzling over the walking lights phenomenon, and similar perceptual phenomena, and asking whether seeing is believing. It could begin with a few examples of combinatorial largeness. This is a journey that has barely begun. Thank you for listening.

REFERENCES

Aiello, L. C., 1996: Terrestriality, bipedalism and the origin of language. *Proc. Brit. Acad.*, **88**, 269–289. Reprinted in: W. G. Runciman, J. Maynard Smith, and R.

I. M. Dunbar (eds.), 1996: *Evolution of social behaviour patterns in primates and man*, Oxford, University Press, and London, British Academy, 297 pp.

Bateson, G. L., 1972: Style, grace and information in primitive art. In: *Steps to an Ecology of Mind: Collected Essays on Anthropology, Psychiatry, Evolution and Epistemology*, pp. 101–125. San Francisco, Chandler; Aylesbury, Intertext; Northvale, New Jersey, Jason Aronson, 510 pp. Gregory Bateson, son of the genetics and embryology pioneer William Bateson (1861-1926), began his career as an anthropologist. His writings contain much wisdom.

Bateson, P., Martin, P., 1999: *Design for a Life: How Behaviour Develops.* London, Jonathan Cape, Random House, 280 pp. Also to be published in the USA in 2000 by Simon and Schuster. This lucid and authoritative new book, from acknowledged experts on the evidence from animal and human behaviour, shows how great a misconception it is to suppose that "it's all in the genes" or "it's all down to culture", or even to speak simplistically of "how much nature and how much nurture".

Blakemore, C., 1977: *Mechanics of the Mind.* Cambridge, University Press, 208 pp.

Cross, I., 1999: Is music the most important thing we ever did? Music, development and evolution. In: *Music, Mind and Science* ed. Suk Won Yi, Seoul: Seoul, National University Press, http://www.mus.cam.ac.uk/~cross/MMS/ and snubook@plaza.snu.ac.kr

Ehrenreich, B., 1997: *Blood Rites: Origins and History of the Passions Of War.* London, Virago/Little, Brown, 292 pp.

Fixsen, D. J., Cheng, E. S., Gales, J. M., Mather, J. C., Shafer, R. A., and Wright, E. L., 1996: The cosmic microwave background spectrum from the full COBE FIRAS data set. *Astrophys. J.*, **473**, 576–587. COBE stands for "cosmic background explorer", and FIRAS for "far infrared absolute spectrophotometer", where "absolute" has its technical experimental meaning, namely that measured values contain no undetermined additive constant: there is no freedom to shift the points plotted in Figure 5 up or down relative to the scale at the left.

Goodall, J., 1990: *Through a Window: My Thirty Years with the Chimpanzees of Gombe.* London, Weidenfeld & Nicolson; also Boston, Houghton Mifflin, 1991, 268 pp.

Green, B., Gallwey, W. T., 1987: *The inner game of music.* London, Pan Books, 248p.

Jouzel, J., Barkov, N. I., Barnola, J. M., Bender, M., Chappellaz, J., Genthon, C., Kotlyakov, V. M., Lipenkov, V., Lorius, C., Petit, J. R., Raynaud, D., Raisbeck, G., Ritz, C., Sowers, T., Stievenard, M., Yiou, F., Yiou, P., 1993: Extending the Vostok ice-core record of palaeoclimate to the penultimate glacial period. *Nature*, **364**, 407–412.

Kealey, T., 1997: *The economic laws of scientific research*. New York, St. Martin's, 382 pp. See
http://www.atm.damtp.cam.ac.uk/people/mem/papers/LHCE/padavid.html

Kegl, J., Senghas, A., Coppola, M., 1999: Creation through contact: sign language emergence and sign language change in Nicaragua. In: *Language Creation and Language Change: Creolization, Diachrony, and Development*, ed. Michel DeGraff. Cambridge, Massachusetts, USA, MIT Press, 573 pp., 179–237.

Mather, J. C., Boslough, J., 1996: *The Very First Light: The True Inside Story of the Scientific Journey Back to the Dawn of the Universe*. London, BasicBooks, 328pp. This is a popular account of COBE, the "cosmic background explorer" satellite, including the courage, and hard labour, involved in rescuing the project, involving 1500 or more participants, from the aftermath of the Space Shuttle disaster. Dr John Mather was the project leader. Reviewed by Professor Joseph Silk in *Nature*, **385**, 691, under the heading "Genius triumphs over jealousy". It is difficult to imagine the vast effort involved in conceiving, designing, building, using, and crosschecking the exquisitely precise yet robust instrumentation that made these astonishingly accurate measurements possible.

I McIntyre, M. E., 1997a: Lucidity and science I: Writing skills and the pattern-perception hypothesis. *Interdisciplinary Science Reviews*, **22**, 199–216. Updates and material supplementing the text of this paper are available on the Internet at http://www.atm.damtp.cam.ac.uk/people/mem/

II McIntyre, M. E., 1997b: Lucidity and science II: From acausality illusions and free will to final theories, mathematics, and music. *Interdisciplinary Science Reviews*, **22**, 285–303. Updates and supplementary material are available on the Internet at http://www.atm.damtp.cam.ac.uk/people/mem/

III McIntyre, M. E., 1998a: Lucidity and science III: Hypercredulity, quantum mechanics, and scientific truth. *Interdisciplinary Science Reviews*, **23**, 29–70, 316. Updates and supplements, including an important **CORRIGENDUM**, are available on the Internet at http://www.atm.damtp.cam.ac.uk/people/mem/ (or google Lucidity Charlesworth)

McIntyre, M. E., 1998b: Breaking waves and global-scale chemical transport in the Earth's atmosphere, with spinoffs for the Sun's interior. *Prog. of Theor. Phys. (Kyoto)*, Supplem. No. **130**, 137–166. See also an important *Corrigendum*, 1999: *Prog. of Theor. Phys. (Kyoto)*, **101**, 189.

Neill, A. S., 1970: *The Last Man Alive: A Story for Children from the Age of Seven to Seventy.* London, Gollancz, 192 pp. First published London, Herbert Jenkins, 1938. Alexander Sutherland Neill is famous for his insight into human children and human nature and for founding — successfully though a century or two ahead of its time — the first children's democracy, the legendary Summerhill School in Leiston, Suffolk, UK. See also http://members.tripod.com/thelastmanalive/home.html and ...//www.atm.damtp.cam.ac.uk/people/mem/papers/LHCE/dilnot-analysis.html and links.

Piattelli-Palmarini, M., 1994: *Inevitable Illusions: How Mistakes of Reason Rule Our Minds.* New York, Wiley, 242 pp.

Strunk, W., White, E. B., 1979: *The Elements of Style*, 3rd edn. New York, Macmillan, 92 pp. A miniature gem on basic writing technique — an elegant pocket book, perfect for bus queues, showing great insight and compressing the essentials into a few dozen small pages.

Thompson, L. G., Yao, T., Davis, M. E., Henderson, K. A., Mosley-Thompson, E., Lin, P.-N., Beer, J., Synal, H.-A., Cole-Dai, J., Bolzan, J. F., 1997: Tropical climate instability: the last glacial cycle from a Qinghai–Tibetan ice core. *Science*, **276**, 1821–1825.

Tippett, M., 1995: *Tippett on music*, ed. M. Bowen. Oxford, University Press, 318 pp.

Valloppillil, V., and co-authors, 1998: *Halloween I.* On the Internet, with commentary by E. Raymond, at http://www.opensource.org/halloween/halloween1.html. Among other things this is an inadvertent but important testimony, from the Microsoft Corporation, to the quality of open-source software such as GNU and Linux and its remarkable reliability and "robustness...at every level", by comparison with Microsoft's own software products. The functioning of today's Internet depends on that reliability and robustness, which in turn depends, as *Halloween I* makes very clear, on the existence of an open, independent computer-science community. *Halloween I* is also about hypercompetitive aims and tactics. In tomorrow's world we may expect there to be fundamentally similar, though far more complex, issues concerning genetic engineering.

Wills, C., 1994: *The Runaway Brain.* London, HarperCollins, 358 pp. This beautiful book presents abundant and cogent evidence — see for instance the discussions of genotype–culture feedbacks and the 'Baldwin effect' — for a most intimate and subtle interplay between nature and nurture, between biological and cultural evolution. Further comments in Note 13 of **I**.