

Figments of reality

The evolution of the curious mind



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge, CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Ian Stewart and Jack Cohen 1997

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 1997

First paperback edition 1999

Printed in the United Kingdom at the University Press, Cambridge

Typeface Swift 9/14pt. System QuarkXPress™ [SE]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Stewart, Ian, 1945–

Figments of reality : the evolution of the curious mind / Ian
Stewart and Jack Cohen.

p. cm.

Includes bibliographical references.

ISBN 0 521 57155 3 (hardcover)

1. Cognition and culture. 2. Consciousness. 3. Intellect.
4. Genetic psychology. 5. Pluralism (Social sciences) I. Cohen,
Jack. II. Title.

BF311.S679 1997

153.4'2-dc21 96-49357 CIP

ISBN 0 521 57155 3 hardback

ISBN 0 521 66383 0 paperback

Contents

	Preface	ix
	Figure Acknowledgements	xiii
	Prologue	1
1	The Origins of Life	5
2	The Reductionist Nightmare	33
3	Ant Country	63
4	Winning Ways	77
5	Universals and Parochials	109
6	Neural Nests	135
7	Features Great and Small	165
8	What is it Like to be a Human?	193
9	We Wanted to Have a Chapter on Free Will, but We Decided not to, so Here It Is	227
10	Extelligence	243
11	Simplex, Complex, Multiplex	271
	Epilogue	301
	Notes	305
	Further Reading	313
	Index	317

1 The Origins of Life

A woman scientist[♪] had been working for some time with a chimpanzee, teaching it to carry out various tasks such as opening a box and rewarding it with fruit. One day, after a session with the chimpanzee, she came into the coffee room half laughing and half crying, obviously very emotional. Her colleagues, a little alarmed, finally managed to get out of her what had happened. She had decided to leave the laboratory area temporarily, and had undone the bolt on the door – whereupon the chimpanzee had solemnly handed her a stick of celery.

Our prologue is one way to tell the story of who we are and how we got here. Such a story has several virtues: it demonstrates how utterly incomprehensible the universe in its entirety is, and how difficult it is for a newly intelligent upright ape to close the conceptual circle by encapsulating the sheer vastness of that universe inside its tiny brain case. It encourages humility. It is the cosmological story as we currently conceive it, the best guess that today's science can make about a past that we cannot revisit and distances too enormous for us to cross. It is a story so strange that we may be tempted to dismiss it as wild speculation, but that will not make the strangeness go away, because if that story is false then the true story must be even stranger.

Assuming there is such a thing as *the* true story of the origins of the universe, which is debatable.

From our own point of view, however – we mean the human race, not JC & IS – this story is impersonal and back to front. It starts with nothing, and ends with each one of us as some kind of accidental by-product of forces beyond our wildest imagination. It describes a universe that is largely alien to the one that we inhabit, which is a private universe filled with very different, human-scale things – friends, spouses, children, pets, plants, bricks and mortar. Each of us inhabits a personal universe; in a sense each of us is a personal universe – for if we are destroyed then our personal universe vanishes with us. The universe of cosmology is made of fundamental particles, such as electrons, and radiation, such as light; but our personal universes are made of very different kinds of things. We don't mean that our own universes aren't made from ordinary matter – we mean

that this matter is organised in a different manner. Most of the interesting features of our personal universes are people and their activities – friends and lovers, enemies and acquaintances from our work or our play. Because most of us live in cities the typical personal universe is urban, composed of buildings, rooms, out-of-town shopping centres ... What occupies most of our daily thoughts is *people* – their influence upon us, and ours upon them. There are babysitters to arrange, theatre tickets to book, bosses to placate, bank managers to be persuaded that a loan would be a sound business proposition ...

Sometimes the external ‘non-people’ world intrudes, but even then it normally does so by way of a human-made artefact: the car needs new tyres, the lawn needs mowing, a sudden attack of ‘flu needs medication. Changes arising outside our own small circle affect our lives in ways we do not anticipate and of which we may not approve – new machinery makes our job unnecessary, anti-pollution laws add to the cost of doing business, a new disease infests our food supply, vandals cut our telephone wires, or people from a country thousands of miles away, which we have never visited, start dropping bombs on us. When the outside world intrudes upon our personal universe we become conscious that the outside exists, but most of the time we still interpret the intrusion in personal terms. We look for a new job that suits our abilities, we hire a lawyer to help us avoid our expensive new legal obligations, we temporarily stop eating burgers, we call the telephone repair man, we build bomb shelters and sit in them cursing the enemy while the bombs fall.

But we do more than that. Many other creatures look up into the night-time sky and see the stars, but we stare at them, wonder how many there are, wonder how far away they are, wonder how they got there, wonder what they are made of, wonder – indeed – why they are there at all. We link them into simple patterns and weave stories around them to help us to rationalise their existence and to remember which pattern is which – the Hunter, the Hero, the Princess, the Bear, the Swan. Although we cannot get inside other animals’ heads, we see no evidence that any other creature looks outside its personal universe in this manner. Maybe chimpanzees and dolphins do; maybe the whale’s enigmatic and interminable song is an exercise in submarine philosophy – but maybe it’s just the whale’s way of saying ‘Hi, anybody out there? This is me.’[♪] Chimps and dolphins and whales don’t build astronomical observatories, they don’t make calendars to predict the seasons, they don’t carve symbolic versions of their thoughts on rocks. Maybe they’re wiser than we are, having fun instead of agonising about their place in the vast uncaring universe; but wiser or not, even the bright ones behave differently from us.

When we look outside our personal universe, we find that the external world is organised in its own characteristic way. It has gravity, ecology, dinosaurs, $E = mc^2$, angles of a triangle adding up to 180° , and so on. It is impersonal: while it is perfectly reasonable to argue with your bank manager that she should increase your overdraft above £180, it is fruitless to argue with a triangle in the hope of increasing the sum of its angles above 180° . On the other hand, the external universe links into our personal world in many ways: calories in food, digital music on CDs, passenger jets, television. All these technologies depend on science, and science is our most successful way to dig into the structure of that external, impersonal universe. Television strengthens the connection between the personal and impersonal worlds by providing science programmes on how the world began or how it will end, and natural history programmes – like our pets and aquariums, house plants and gardens – provide tenuous links with the rest of living nature. All this notwithstanding, we are much more concerned about how we fit into our personal circle of friends than about how we all fit into the complex ecology of our own planet.

Those of us who are scientists behave in exactly the same way, but we tend to be more bothered by it, because we have real trouble understanding why we're doing it. Our scientific instincts tell us that the real universe out there is actually far more important, on any serious scale of events, than whether Mary told her mother she was dieting ... but somehow questions on the level of Mary's diet take up much more of the scientist's time than the whys and wherefores of galactic superclusters – even when the scientist is a cosmologist.

We lead a dual existence – *in* nature but not *of* it, perpetually reacting to our estimate of what the world will be rather than what it is right now. We mirror the world outside us with another in our heads: our perceptions of that world. It's a distorting mirror, an imperfect representation, but to us it seems *real*. In a funny self-centred way we see ourselves as existing slightly to one side of the rest of the universe. We are in control of our world, we can make choices, we have *minds* that we can make up or change. Everything else is just following the inexorable impulses of nature. When we think of an amoeba, a fox, an oak tree, or a dinosaur, we think of them as a part of nature. The amoeba fiddles about putting out pseudopods and ingesting food particles, and that's about it. The fox runs through the bushes chasing a rabbit for dinner, and when it encounters the occasional bunch of subhumans on horseback it's too busy running from the dogs to debate the morality of blood sports. The oak tree is just sitting there synthesising, drawing in water from its roots and carbon dioxide from the air, and if it's worrying about anything it's about the impending winter and dropping its

leaves – not whether the neighbouring oak tree thinks it's a cad for fertilising too many of its acorns. We see dinosaurs as eating, breathing, multiplying, and dying out against the great backdrop of natural forces, like the K/T meteorite that hit the Earth 65 million years ago and caused mayhem all over the planet. Gary Larson's 'Far Side' cartoons often work by imputing human-type motivation to animals, and they are funny because we know that most animals *don't* worry about their circle of friends.

All very well. But how much of our belief that we are special is grounded in fact, and how much is just a comfortable illusion of superiority? The belief that we are superior to other animals is a human value judgement, and as such is likely to be biased in our own favour, but there can be little doubt that we are *different* – in important ways – from the other animals on our planet. These differences must be explained. Their explanation is made more difficult, but also much more interesting, by the fact that human beings have not always been as they are now. Few of us doubt that we evolved from creatures that, like most animals, related directly to the natural world and thereby avoided all of the social problems that occupy our every waking minute and even assail us in our dreams.

How did that happen?

This question is the central issue that will shape our narrative. What was it about this particular lump of rock, in this particular spiral arm of this not terribly special galaxy, that made us the way we are? How is it *possible* for inanimate matter to turn into complex creatures like us with their own inner worlds of mind and imagination? Given that it is possible, why did it happen? Why us?

Some will ascribe it to God and be satisfied: we have nothing to say to them.

Some will ascribe it to inexorable consequences of the fundamental laws of physics, and be satisfied: we have nothing to say to them either.

We *do* have something to say, however, to those who find either answer incomplete, people who think that our presence on this planet and our curious mental abilities deserve to be explained rather than explained away. In *Figments of Reality* (henceforth abbreviated to *Figments*) we attempt to explain the evolution of human beings from a new point of view – one that differs considerably from the usual scientific story, although it retains many points of contact with it. More accurately, we shall look at the questions of mind and culture from *two* disparate viewpoints, which complement rather than contradict each other. One is the conventional scientific viewpoint: take the system to bits – in a conceptual sense – and see how those bits fit together. The other, less conventional but in our opin-

ion equally important, is to look at *context*, and see how the system is shaped by what lies around it.


Along the way we shall be forced to reassess the orthodox scientific stories about how things work, many of which are little better than myths. We don't think that such reassessment makes the orthodox stories any less 'true' (we'll air some of our prejudices about truth later), and we certainly don't think that it makes them any less 'scientific'. The point is that if you approach the questions from different directions you may find yourself wanting different kinds of answer, just as 'God' may satisfy a priest in search of virtuous living but not a programmer in search of virtual reality. We think that such changes of perspective help to make many problems of human evolution and cultural development seem less puzzling. In particular they will help us to tell the story of human mind and culture in a more accessible way – one that explains, rather than just asserts, the scientific bases of our world and of ourselves.

We'll give you the bare bones of the story now, to act as a 'road map' for the rest of the book. First, we look at the origins of life and its evolution – both on Earth, the story of how we came into being, and elsewhere, the story of what might have happened instead and what might be happening right now on a planet of some distant sun. We describe the evolution of senses – in particular sight, hearing, and smell – showing how they have influenced the evolution of networks of nerve cells, leading to that most flexible and enigmatic of all organs, the brain. We demonstrate that, far from being mere passive observers of reality, our senses are fine-tuned during development to emphasise those features in which our brains have an especial interest. By manipulating these mental features we construct 'conceptual maps' of the reality around us, which enable us to make up our minds (take decisions) and change our minds (modify our choices in response to the consequences of those decisions). We do not so much *observe* reality as put together our personal representation of it and drape that back on to our perceptions of the external world. This facility is moderated by intelligence – the ability to reason, to solve problems – which is not merely a structural feature of large brains with intricate networks of nerves. Intelligence arose in intimate association with a marvellous non-genetic trick used by parents to provide their offspring with a head start in life, a trick that we call 'privilege'. Privilege begins with yolk and nests, and culminates – so far – in culture. We further claim that it is not intelligence alone, or culture alone, that leads to mind, but both – interacting 'complicity'.

A feature of our minds that is often singled out as *the* thing that makes us uniquely human is language. Some scientists think that language is a

necessary prerequisite for intelligence, and others that intelligence is a necessary prerequisite for language. We think that both are right – and so both are wrong, for each thinks the other mistaken and both are mistaken about ‘prerequisite’. Language and intelligence evolved *together*, both being inextricably linked to culture.

Finally, we tell of the rise of human culture, the techniques that cultures employ to survive in a changing world, and the effect of cultural differences on displaced ethnic groups, leading to multicultural societies in which individuals grapple with changes in their cultural identity. We tell of the growth of global communications that lock the multiculture in place, so that we cannot go back even if we wish to. We take a brief look at the future of human multiculture. And we wrap the entire package up and tie it with a neat bow, by means of a unifying concept – extelligence – that is the contextual and cultural analogue of internal, personal intelligence.

To kick the whole story off, we now ask a ‘warm-up’ question: how did inanimate matter give rise to life? In the Prologue we described the current view of the origins of the universe, the ‘Big Bang’ theory as it is called. Space, time, and matter arose from nothing; then the simple kinds of primal matter that existed at the prevailing high temperatures began to combine to make all of the different chemical elements – hydrogen, helium, lithium, beryllium, boron, carbon, nitrogen, oxygen ... These different atoms then combined to form chemical molecules – two hydrogens plus an oxygen to make water, one carbon and two oxygens to make carbon dioxide. The bodies of living creatures are made from millions of different molecules, all of which trace back to the nuclear reactions in the cores of stars. Literally, ‘we are stardust’, as Joni Mitchell sang about Woodstock. 

Particles building into atoms, and atoms into molecules – these we can comprehend, they’re just like bricks building into a house. But houses don’t develop a will of their own, get up, and walk away. Living creatures did, and that’s a real puzzle. How did inanimate, inorganic chemistry somehow generate the rich flexibility of life? Not all at once, that’s for sure. There was no wondrous, special moment, pregnant with significance, at which life suddenly appeared on the planet. Instead, life emerged gradually from non-life. In this respect the origins of life are a bit like the origins of a person’s life. There was a time when Maureen didn’t exist. At what time did the egg, embryo, fetus, child, become Maureen? At what time did it become human? Surely there was not a specific *moment* of becoming Maureen – though people who don’t know about it do talk of ‘the moment of fertilisation’ – except in a legal sense, at her naming cere-

mony. A person is like a painting or a novel: it progressively comes into being. Maureen started as not-Maureen and gradually became Maureen. So it was with the origins of life.

We can't go back and see what actually happened, but we can infer the kind of molecular game that must have been played out upon the primal Earth. In particular, we can understand that life could reasonably have come into being gradually and spontaneously as a consequence of perfectly reasonable chemistry. Four billion years ago, the Earth was a very different place. Its surface was barren rock, sandy desert, bubbling tar-pit, smoking sulphur-hole. Its oceans were a watery layer of chemicals dissolved out of the rocks and injected into the ocean depths by underwater volcanoes. All of the diversity of chemical elements that we find today was already present then – for apart from a continual infall of meteoric dust and a slow leakage of the lighter gases, the atoms that make up today's world are the same ones that were present four billennia ago. The difference between that ancient Earth and the one we inhabit today lies not in its atoms, but in its molecules. They are much more diverse now, and – absolutely crucially – they are organised in much more complicated ways.

Textbooks tell you that a molecule is a system of atoms connected together by interatomic forces – 'bonds'. This is true – as much as any human statement about nature is true – but it is not the whole truth. Another part of the story is that unlike atoms, molecules can become more complex. Atoms, left to themselves, do not produce types of atom that have never existed before – although some atoms can change by way of nuclear reactions, with uranium turning into lead, for example. But atoms can rather easily produce entirely new types of molecule by combining in new ways, and those molecules can also go on to produce new molecules – a process that continues to this day. If the only thing you knew about the Earth was a catalogue of its molecules, you would be able to see a distinct difference between today's catalogue and that of four billion years ago. Today's catalogue would include many enormous molecules, such as proteins and DNA, that would be missing from the early version.

So over the billennia, molecules have become more complex. However, that is by no means the whole story, because there is much more going on than mere complexity. That four billion year-old catalogue of molecules would include some amazingly complicated ones too, for instance innumerable weird conglomerations formed in the tar-pits. Similarly today's catalogue would be littered with molecules like toffee, a disordered mass of one-off constructions whose greatest similarity to each other is that every single one of them is totally boring. No, the molecules that are of greatest interest are not just *complicated* –

they are organised. They are, in fact, machines – the first machines that appeared on Earth. To be sure, they don't look much like the machines with which we are familiar – lawnmowers, cars, aeroplanes – but they have a basic property in common with these human-made devices. They can perform functions, a fancy way to say that they do things. A function is an operation which, when presented with certain inputs, produces various outputs in a reliable manner. The most obvious function of a lawnmower, for example, is to mow a lawn: here the input is a lot of straggly grass and the output is a neat, tidy swathe of green. A lawnmower can perform other functions too: propping open the door of the garden shed or holding down a pile of plastic sacks when a breeze is blowing.

Molecules, too, can perform functions, because they interact with other molecules. And because molecules have definite shapes, these interactions are different for different molecules. For example, molecule A may have some kind of dent in its surface, just the right shape to fit a bump in molecule B. If so – and if the interatomic forces are suitable – then you would expect to find many molecules that are made from A and B fitted together. This sort of 'plug and play' construction of molecules is going on all the time. It is to some extent counter-balanced by the tendency of molecules to fall apart for various reasons, so we don't just get the whole of terrestrial existence locked together in a single super-molecule.

Molecules can also have moving parts. The bonds that join their atoms together can bend and twist, to a limited extent, and sometimes atoms can even revolve on their bonds like propellers on a spindle. This flexibility provides a lot of scope for making chemical machines with interesting functions. Some molecules can make other molecules fit together, or pull them apart. After performing their function they remain unchanged, and are ready to carry it out again and again. Such molecules are called 'catalysts'. Catalytic molecules act like a production line: provided they are supplied with the right 'raw materials' they can go on turning out copy after copy of their favoured molecule, indefinitely.

Carrying out a function is quite different from having a purpose. Molecular machines do not carry out functions because they want to do so: they carry them out because this is how they are made. Indeed it is impossible for them *not* to carry out their functions. In the same way, a rock carries out the function of rolling down a hill because it is suitably rounded and has significant enough mass for gravity to latch on to. But it does not have that rounded shape for the *purpose* of rolling down a hill. We mention this because human beings seem to have an innate tendency to confuse functions with purposes – so that, for example, 'the sun keeps us warm' becomes 'the sun was placed in the sky *in order*

to keep us warm'. This kind of purpose-centred thinking can easily lead to people worshipping the sun-god, not realising that the sun can perform the function of keeping them warm without either wishing to do so, or requiring worship to continue doing it.

At any rate, four billion years ago there were pretty much the same atoms around as there are now, but not in the same combinations, and not organised like they are now. The complex molecules that occur in living organisms and in pseudo-living entities such as viruses are known as 'organic' molecules. The atom that makes all organic molecules possible is carbon: carbon atoms have the ability to stick together and form huge, stable skeletons, to which other atoms can attach. Even carbon can perform this task only within a narrow range of temperatures, and other atoms can't do it at all, with the possible exception of silicon. This is not to say that carbon is essential for life; just that it is essential for *our* kind of life, which is the only kind we know about, and it generally looks like rather good stuff to make life from. However, the kind of organisation that we call 'life' might in principle arise in other ways – silicon-based molecules, interacting trains of electrons in metallic crystals, colliding plasma vortices in the corona of a star ... The possibility of complex molecules is important because *some* complex molecules can perform more sophisticated tasks than simple ones. Upon these more sophisticated tasks does the peculiar form of matter that we call 'life' depend. Living organisms are much more than just formless bowls of molecular soup: the manner in which their molecules are arranged is at least as important as what those molecules are. But without the potential complexity that carbon provides, molecules complicated enough to get themselves organised into organisms like us would not exist.

Life seems very different from inorganic matter – it can move of its own volition, reproduce itself, consume other substances, respond to its environment. It is therefore hardly surprising that some people think that living material is simply a different *kind* of stuff from non-living matter. This belief is known as vitalism. Its greatest defect is that there is no evidence in its favour: none of this different kind of stuff has ever been isolated. If you take a living organism to bits, right down to the molecular level, all you find is ordinary matter. We humans are made from the same atoms as the rocks, water, and air around us. The inevitable conclusion is that it is not the ingredients that differ: it is how they are organised. A living creature can be killed by bashing its head with a rock: it is hard to see how such a crass act can devitalise its esoteric immaterial substance, but easy to see how it can wreck its organisation.

In the same manner a car is made from the same atoms as the sheets of

metal, sacks of aluminium powder, and cans of polymer from which it is assembled. Its ability to move does not arise because it is made from a different kind of matter: it is merely a consequence of how that matter acts when it is put together in a particular manner. An automotive engineer would be able to explain, in more than enough detail to send any partygoer in search of the drinks tray, what is involved in this organisation. But nobody ever made a car by going out and looking for a new kind of matter that has the ability to move when petrol is poured into it.

There is a danger with the 'car' analogy if it is pushed too far. To some people, organisation implies the existence of an organiser, as the existence of a watch implies that of a watchmaker. This is a seductive line of argument, but there is no compelling reason to accept it. One of the most remarkable features of organic matter – and, we now realise, inorganic matter too under suitable circumstances – is its ability to organise *itself*. So in some ways a better analogy than a car would be a whirlpool, a tornado, or a flame: an organised structure that comes into being without conscious intervention. Our intuition is upset by self-organisation, probably because we seldom experience such behaviour directly: in our everyday world the only way to produce organisation is to work pretty damned hard to make it come about. Nevertheless, we are surrounded by and made from matter that is highly organised, and it must have got that way by some route. Either it has been organised by an organism-maker, or it has organised itself.

The problems with the 'organism-maker' hypothesis have been rehearsed by philosophers and theologians for as long as anyone cares to remember. Its obvious advantages (it 'solves' the problem to many people's satisfaction) are countered by its equally obvious defects. For instance, who or what organised the organiser? And where is the organiser? The 'self-organisation' hypothesis has far more to offer to those who share the scientist's wish to understand nature and not just postulate it. It is a daring hypothesis, which does not solve the problem unless we can explain *how* and *why* living matter self-organises. It is becoming clear that there is nothing inherently self-contradictory in the idea that organisation sometimes comes 'for free', and it is also becoming clear that limited laboratory-scale systems and computer simulations indulge in self-organised behaviour far more often than we might have anticipated. *Why*, we are still unsure, but we know that it is so. Perhaps our universe is special in being like that; perhaps all universes must be. Which, we don't know.

The self-organising ability of life becomes clear only over long time-scales: compare an organism such as a mouse, today, to a lump of rock four

billion years ago. One of the most obvious ‘unusual’ features of life, however, can be seen on far shorter timescales: its ability to reproduce. Life makes new life – and pretty much the *same* life. People make new people, cats make new cats, nematode worms make new nematode worms, and amoebas make new amoebas. This is an amazing ability, and it certainly looks very different from ordinary chemistry.

However, we tend to underestimate what ‘unaided’ chemistry is capable of, and that distorts our assessment of how amazing or unlikely life is. Thirty years ago, biology was thought to be very complex and chemistry relatively simple. The chemical story of the origins of life seemed to require the construction of a conceptual pyramid of ever-complicating processes, rising from the lowly plains of test-tube chemistry to the lofty heights of biology. Nowadays we understand that this picture is wrong. ‘Unaided’ chemistry – chemistry that does not require a living organism to make it happen – goes all the way up. Even simple unaided chemistry is a lot more complicated than the textbooks would have us believe. For example, if a mixture of two parts hydrogen to one part oxygen is ignited, then it explodes, giving water. The old textbooks see this as a single chemical reaction: $2\text{H}_2 + \text{O}_2 \rightarrow 2\text{H}_2\text{O}$. (We don’t write this in the apparently simpler form $\text{H}_2 + \text{O} \rightarrow \text{H}_2\text{O}$, by the way, because reactions are about *molecules*, and a molecule of oxygen is O_2 , not O.) Newer textbooks will tell you that there are at least ten other molecules involved as intermediaries, and the more closely you look, the more of them you will find. The old textbooks tell you what to start with and what it ends up as, but not what happens in between. When reactions as basic as this turn out to be so complex, it is not surprising that more sophisticated kinds of chemistry are *far* more complex. Moreover, as our understanding of the complexity of chemistry grew, we also came to recognise that biochemistry is a lot closer to ‘unaided’ chemistry than we used to think. In fact modern industrial processes, which make extensive use of catalysts, sit right at the junction of ‘unaided’ chemistry and very similar biochemistry.

Another reason why we are so puzzled by life arising from ‘mere’ chemistry is that it is very difficult to find, on the Earth, now, the kind of chemistry that long ago gave rise to life. This is because life has invaded all of the possible habitats for such chemistry, from the deep oceans, tens of miles deep in granite cracks, to high in the atmosphere – so their chemistry has been changed out of all recognition. Rusting would be a good example, except that on Earth it is nearly always ‘assisted’ by bacteria, who take a tithe of the energy. So let’s imagine iron rusting on the surface of a lifeless planet. Recall the concept of catalysis: a molecule is a catalyst if it assists in the production of another molecule, or molecules,

without itself being used up in the process. Sterile rusting proceeds by auto-catalysis – given a bit of rust on iron it catalyses more of *itself*. Such a process is recursive, it pulls itself up by its own bootstraps, so you need a bit of the product to get it started. (Stop worrying: we never said that that initial bit of product was produced by the *same* recursive process. See later.)

Many recursive systems are known in real chemistry and technology, but they are largely missing from school or college chemistry because they don't fit the simplified theories being taught there. The catalytic convertor in a car oxidises pollutants using just such a system. The catalytic surface does its work in a series of expanding rings, just like the very best example of this kind of chemistry, the Belousov-Zhabotinskii (BZ) reaction of figure 1. This is an extremely photogenic instance of recursive chemistry, with expanding rings of blue in a rusty red solution. For forty years after such systems were first described, most chemists did not believe they could work: they seemed to be contrary to that most famous – and misunderstood – of scientific laws, the Second Law of Thermodynamics.

They are not, and neither is life.

Thanks to the epic researches of Maurice Wilkes, Rosalind Franklin, Francis Crick, and James Watson in the 1950s, we know that one remarkable molecule – more properly, a family of very similar molecules – underlies almost all terrestrial life. That molecule is DNA, whose initials stand for 'deoxyribose




Figure 1 Typical 'target' patterns in the Belousov-Zhabotinskii reaction. As time passes, the rings expand.

nucleic acid' (or 'deoxyribonucleic acid' according to taste). DNA forms the genetic material of almost all organisms. A few viruses use RNA, 'ribose nucleic acid' or 'ribonucleic acid', but DNA and RNA come from the same molecular stable. DNA has a simple but clever molecular structure in which twin strands spiral like a staircase. The treads are made from four types of molecule called 'bases', held together by a framework of sugars and phosphates. This structure allows DNA to do two important things: encode information, and replicate. The information is represented by the sequence of bases, and includes such things as the structure of key proteins without which organisms cannot be built, and sequences that determine *when* they will be built. DNA replicates by separating the two strands, in which the bases are complementary to each other, and re-creating a matching strand for each, thereby producing two copies of the genetic information from one original. (This description, though standard, is an oversimplification, but it is sufficiently accurate for our present purposes.) Throughout *Figments* we shall distinguish replication, the creation of exact or nominally exact copies, from reproduction, the creation of *similar* copies – in particular, similar enough that they too can reproduce. Normally DNA replicates, but when the occasional inevitable copying error – the technical term is *mutation* – creeps in, then the molecule is better thought of as reproducing.

Although it is often described as such, DNA is not a *self*-replicating molecule: leave a mass of DNA in a beaker and you won't get more of it. It replicates only with the aid of many other molecules, known by names like transfer RNA, messenger RNA, and enzymes. We mention these merely to drive home that DNA needs an entire 'support team' in order to replicate: it no more makes copies of itself than a document in a photocopier makes copies of itself. Moreover, the fact that DNA contains 'information' is far less important than the physical (that is, chemical!) form that the information takes. All molecules 'contain' information – the positions of their atoms, for example, are a kind of information, as you will quickly discover if you build molecular models. The information in DNA is useful *not* because it is information, but because it is information stored in a form that other chemical machines can manipulate. As an analogy, the positions of the wood fibres that make up this page encode a huge amount of information, but when you read the page the only *useful* information – for you – comes from the letters printed on it.

The process that allows DNA to replicate is another autocatalytic recursive cycle, only here it is a *collection* of molecules that catalyses itself. The DNA contains the defining information for the molecules in the support team. The support team helps DNA to replicate, and the DNA helps to replicate its own

support team. Recursion often feels disturbing, but how *else* could a replicative process work? What makes recursive processes disturbing is the feeling that they can never get started – the ‘chicken and egg’ problem. Actually that’s not a serious problem at all, just a case of sloppy thinking caused by incorrectly extrapolating the process backwards. It’s relatively easy to get a replicative process *started*. What you can’t do – without destroying the process – is *stop* it. The way to start a chicken-and-egg process is to create a suitable start-up configuration, one that is part of the process only the first time round. For example a non-chicken might be persuaded to lay an egg that grows into a chicken, whose eggs also grow into chickens, and so on forever. Clearly you can’t play this trick if you start with a perfectly replicating non-chicken and absolutely nothing untoward happens to its egg; but if it is a reproducing non-chicken, subject to variations that do not affect the reproductive abilities of its offspring, there’s no conceptual problem at all – just a technical one of actually making the trick work. The answer to the hoary philosophical teaser then becomes no more than a question of definition. Is a chicken egg one that was laid by a chicken, or is it an egg that grows into a chicken? In the former case, the chicken came first (from a non-chicken egg); in the latter case, the egg came first (laid by a non-chicken).

There are other ways to get a replicative or reproductive system started. One is for it to ‘piggyback’[♪] on a pre-existing replicative or reproductive system. This is how documents replicate: they piggyback on photocopiers, which are replicated by humans working in factories. The photocopiers in turn piggyback on human reproduction. Of course it’s not possible for every replicative/reproductive process to piggyback on a previous one, or else there is a genuine chicken-and-egg problem, so at least one process has to get started some other way (and act as a start-up configuration for everything that subsequently piggybacks on it). That other way is best described as ‘scaffolding’: *before* the replicative loop closes up, the process is assisted by something else, which drops out of the loop permanently *after* it is closed. Once a system acquires the ability to replicate, it spreads rapidly and takes over any disorganised substrate.

Although the loop formed by DNA and its support team is in principle replicative, in practice it is ‘only’ reproductive. The procedure is so complex that it seldom takes place without errors. Moreover, in sexually reproducing organisms, the reproductive procedure introduces ‘mix-and-match’ modifications. This should not be thought of as a defect. Reproductive systems are much more interesting than mere replicative ones, precisely because they can change. Replication is just the same thing repeated forever. Reproduction has room for flexibility – it can produce a chicken from a non-chicken’s egg.

That possibility leads to evolution, which in various ways forms the subject of the next three chapters. Before tackling such a subtle subject we shall deal with a more down-to-earth question: how did DNA replication get started? The process looks too complex to have arisen from raw scaffolding: most probably it piggybacked. There are hints of possible precursors in the DNA replication process itself. Over the years, many different proposals have been made, and we mention them here to show that there are *several* plausible solutions to the problem of how life got started on its reproductive path.

One is the 'RNA world'; a second, due to Graham Cairns-Smith, is clay; and a third is Stuart Kauffman's concept of an autocatalytic network of molecules. The RNA world is a hypothetical period of evolution when DNA did not yet play a role in the replication of proto-living forms: instead, the simpler molecule RNA held centre stage and reproduced without help from DNA's band of molecular assistants. Back in the 1950s Stanley Miller, a student of Harold Urey, performed experiments showing how amino acids – the building blocks for proteins – arose spontaneously in a simulation of the Earth's primal chemistry. Variations on this system have provided all the raw materials for life, either DNA-based or RNA-based. The possibility of an RNA world, predating today's DNA/RNA combination, first became apparent in the 1980s when Tom Cech and Sydney Altman¹ discovered special RNA molecules now called ribozymes. These acted as a catalyst in a reaction that snipped out parts of themselves – one element of the recursive process needed for replication. Jack Szostak then employed a laboratory version of molecular evolution to produce more efficient ribozymes which could copy long RNA sequences. In 1996 David Bartel found some that are as effective as some modern protein enzymes. RNA 'self'-replication – employing molecular assistants, but not DNA – has not yet been achieved, but it looks far more plausible.

In May 1996 the chemist Jim Ferris discovered a way in which long RNA strands (10–15 bases in length) might have formed in the primal environment. If he added montmorillonite – a kind of clay – to the chemical mix, then long RNA chains formed on the surface of the clay. This was especially interesting in view of Cairns-Smith's earlier speculations that clay might provide a replicative structure upon which RNA could piggyback, and we will briefly describe what he had in mind. Clay is a complex combination of aluminium, silicon, oxygen, magnesium, calcium, iron, and many other elements. Clays can dissolve in water and precipitate out again. Their crystalline forms employ rarer elements to structure themselves into exotic shapes: scrolls, curlicues, spirals. Like most crystals, these shapes can act as templates to produce more shapes of the same kind, building

up on top. When some external event causes the crystal to break, each piece can act as a template for further growth, so these clay forms can replicate – indeed, reproduce (figure 2). They can even compete with each other, because some shapes are better at extracting particular substances from solution. Clays are probably the nearest thing on Earth to a silicon-based life-form, a replicating system upon which others can piggyback. As Cairns-Smith realised, carbon compounds naturally stick to the surfaces of clay crystals, and they catalyse organic reactions. In particular they catalyse processes of polymerisation, in which molecules of the same kind are added to each other, forming long chains and other structures. By this process amino acids could become proteins, and simple bases could link up to form RNA, DNA, or – mostly – other nucleic acids. As Ferris showed, Cairns-Smith's chemical intuition was justified, which adds weight to his view that the origin of our kind of life was subsequent to a much more primitive kind of clay life, a story that he calls 'genetic takeover'. It is a story of a smooth transition from inorganic chemistry to our kind of life, eventually resulting in creatures about as organised as a bacterium, without a nucleus – what biologists call a prokaryote. We can even have our cake and eat it too: perhaps DNA piggybacked on RNA and RNA on clay.

The autocatalytic network idea is rather different: it presents a set of circumstances in which 'scaffolding' is almost inevitable, rather than being just a convenient coincidence. A replicating molecule would be one that catalyses itself, but that's just a bit *too* convenient and seems not to happen naturally. (Rust doesn't really count: it needs iron, water, and oxygen too, not to mention bacteria.) However, it's *much* easier to come up with a 'support team' of molecules in




Figure 2 Mutation and replication in stacks of clay platelets.

which each member catalyses other members. Such a team ‘closes up’ into an autocatalytic network if *every* member of the team is catalysed by some other combination of members. Then the entire team acquires the ability to replicate.♪

Autocatalytic networks, in particular, illustrate just how close ‘unaided’ chemistry can get to genuine life. Now: suppose that we could provide such a network with one more feature, its own ‘identity’, so that it could exist – and replicate – as a well-defined entity, instead of just dispersing into the vast chemical ocean. Then we would have a rudimentary chemical ‘life-form’. Here’s a possible scenario, proposed some years ago by Alexander Oparin: actual events were probably more complex and possibly quite different. There are molecules known as lipids – fats – which look rather like tadpoles. Their heads are attracted to water molecules but their tails are repelled. Moreover, lipids quite like to stick together. So lipids at the surface of a watery environment – such as the primal ocean – naturally align themselves in sheets, with all the heads on one side and all the tails on the other. These sheets are biological membranes, and their key property is to separate regions of space from each other. In our previous terminology, they are naturally occurring chemical ‘machines’ whose function is to separate the watery ocean into distinct regions. (In real biological membranes there is a *double* wall, its molecules aligned tail-to-tail with heads on both ‘external’ surfaces, but the general function is the same.) Doron Lancet suggested that such membranes might also close up into tiny capsules, and other molecules could diffuse in or out. Now, said Lancet: suppose that, by chance, the molecules trapped inside such a capsule when it first forms happen to make up an autocatalytic network. Then – fuelled by raw materials that diffuse in from the diverse but disorganised primal ocean – the network will replicate. Suppose further that the lipid whose molecules create the capsule is also part of that autocatalytic network. Then the capsule will swell up as its contents replicate, and eventually it will grow to such a size that it becomes unstable, in the sense that a large capsule will tend to break up into smaller ones – each containing the chemical support team needed to make the process continue. There is no obvious end to the process: what we have is in effect a prototype cell, which sucks in nutrients, grows, and divides into cells of the same type. So, given a fair sea and a following wind, autocatalytic networks can in effect organise their own spatial geometry to form replicating organisms – or, at least, proto-organisms. More complex replicating molecular systems now have something to piggyback on.

All the above is speculative: its purpose is to show you how *easily* life might get its act together as a result of natural combinations of ordinary physical

and chemical features of the inorganic world. What actually happened? In Chapter 4 we shall take a look at how and why this kind of piggybacking eventually led to life as we know it – with the driving process, of course, being evolution. For the moment we shall set the scene by leaving out the evolutionary element, describing what seems to have happened without asking why.

Lipid capsules filled with autocatalytic networks of chemicals are close enough to genuine organisms of a bacterial grade of complexity – prokaryotes – that it is easy to imagine how prokaryotes might have come about, although the most plausible theories of their origins are less simple – as we see shortly. At any rate, whichever route it was that produced prokaryotes, we know that they appeared, multiplied, evolved, and took over the surface of the Earth soon after there was a liquid sea. We know that it happened very quickly, though we can't be sure whether it happened as soon as there was liquid water, or whether it required a few million years after that. On geological timescales the difference is immaterial; the point is that it happened so fast that the process must have been chemically and physically 'easy' – for an entire planet, cooling down from a bombardment of meteorites. In fact some kinds of meteorite contain organic compounds, so, for all we know, Molecules from Outer Space may have been the initial scaffolding, as Fred Hoyle and Chandra Wickramasinghe suggested many years ago.¹

Before more complex forms evolved, prokaryote life dominated the seas for three billion years. During that time many different kinds invented photosynthesis, a way to power their recursive chemistry by extracting energy from sunlight. In so doing they excreted a highly toxic waste product – oxygen. At that time few, if any, organisms made use of oxygen, a highly reactive chemical: it still causes problems today, because it lets things catch fire. The build-up of oxygen changed our atmosphere completely, to the extent that it is a long way from chemical equilibrium – that is, without life the level of oxygen would decrease considerably as it reacted with minerals, oxidising them. Life doesn't stop those reactions happening, but it puts the oxygen in faster than unaided chemistry can take it out again.

About 1.5 billion years ago new forms of life, with much more complicated recursive chemistry, arose to exploit the new reactions made possible by oxygen. These were the eukaryotes, and their most important feature was the possession of a nucleus. People often talk of bacteria as being 'unicellular' and creatures like us as being 'multicellular', as if you can evolve a human being by sticking a lot of bacteria together, but this is quite the wrong image. Bacteria are *not* single-celled organisms, because they are not cells. They have a few features in

common with cells, and they seem to have evolved into cells, but even a single cell is considerably more sophisticated than a bacterium.

Eukaryotes can be single-celled – a well-known example is *Amoeba* – but they can also be many-celled. The eukaryote cell differs in significant ways from a bacterium. It is larger – typically about 10,000 times as large by volume. Even in a single-celled eukaryote, the cell possesses a range of ‘organelles’, component sub-units with some special function, such as the nucleus (which contains most of the cell’s DNA) and mitochondria (which protect the cell against oxygen and provide much of its energy). The currently accepted theory,¹ which goes back at least a century and was revived in 1967 by Lynn Margulis, is that the cell arose from independent bacteria of various kinds by a process of symbiosis, which may have started out as parasitism. A simple, but misleading, way to say this is that various bacteria ‘got together’ to produce a cell. A more accurate way to say it is that cells *emerged* from the coevolution of bacteria. We don’t just mean ‘appeared’; we are using the word in the sense of ‘emergent phenomenon’. This term comes from philosophy, and is used when the behaviour of a system appears to transcend anything that can be found in its components – where the whole seems ‘greater than the sum of its parts’. Here the point is that if you put a lot of bacteria together and wait long enough, then the overall system will home in on the cell as a viable way to organise its business.

We shall have plenty more to say later about emergence.

In a similar manner, multicellular creatures mostly arose *not* by combining separate cells together into a colony, but by starting with a single cell and letting it *divide* repeatedly – ‘multiplication by division’. In this manner a single large cell became an aggregate of sub-cells *with the same DNA*, a useful degree of genetic coherence that made it possible for the entire system to co-evolve simply and naturally. But now each sub-unit was free to specialise, if the result helped to keep the organisms’ reproductive cycle going, so eukaryotes evolved different types of cells, with different capabilities. Just as molecules added entirely new dimensions of complexity to atoms, eukaryotes added entirely new dimensions of complexity to organisms. The new atmosphere opened the way to oxygen-breathing organisms with a faster lifestyle; life embarked upon a wild romp of self-complication.

Sometimes an apparently minor change had major implications on a global scale. At some point some varieties of marine organism stopped excreting their wastes in liquid or semi-liquid form, and instead produced them in solid form. Probably the first such animal was a swimming worm, but it might conceivably have been a trilobite – the timing is right, but the evidence is slim.

This minor change in water content made a huge difference, because the solid wastes sank to the bottom of the shallow seas, forming an anaerobic layer for soft-bodied organisms to graze. One animal's waste became another's resource – just as had happened earlier for the toxic oxygen wastes of the bacteria.

In 1909 evidence for one of the more curious stages in the evolutionary process came to light in Yoho National Park, high in the Canadian Rocky Mountains. Charles Walcott, secretary of the Smithsonian Institute and America's leading palaeontologist, discovered a large number of unusual fossils in a rock formation known as the Burgess Shale. The story has been grippingly told by Stephen Jay Gould in *Wonderful Life*. The fossils were unusual because they were formed from soft-bodied creatures. Normally conditions are unsuitable for soft parts to fossilise, but in this case something like a mud-slide had overwhelmed the pool in which they were living. Walcott took a cursory look, assigned them to various known groups of organisms, filed them away in drawers and forgot about them. In 1971 Harry Whittington of Cambridge University recognised that the Burgess Shale organisms are far more interesting than Walcott had supposed. They represent an early explosion of multicellular life: the anatomical diversity in that one small pool was much greater than that over the entire global ecosystem today. Not in terms of the number of species, but the number of phyla. A phylum is one of the largest units into which organisms are classified. For example today's many-jointed 'arthropods' – members of the three phyla of crustaceans (shrimps and the like), chelicerates (spiders, scorpions and their kin), and uniramians (insects and more) – *all* evolved from three groups present in the Burgess Shale. However, more than twenty other radically different arthropod designs are found in the Burgess Shale creatures too, only one of which – the now-extinct trilobites – went on to establish itself as a major player. In recent years several palaeontologists have suggested that the diversity of the Burgess Shale organisms is not *quite* as great as was first supposed, but they are without doubt a weird and varied bunch: for instance figure 3 illustrates *Opabinia*, which has a nozzle at the front, a claw at the back, five eyes, gills on top of its body, and a three-segment tail.

Just *one* of the Burgess Shale creatures was part of the evolutionary lineage that led to humanity. Since most of the Burgess Shale creatures quickly died out, for no obvious structural reasons, Gould deduces that it was largely a matter of luck which of them survived – and he inferred that our presence on this planet owes much to Dame Fortune and not much to Good Design. But is this really so? We give our answer in Chapter 5. But whatever the interpretations placed on it, the Burgess Shale fossils show that around 570 million years ago, at