



IAN STEWART
JACK COHEN

Figments of Reality

THE
EVOLUTION
OF
THE
CURIOUS
MIND

CAMBRIDGE

more information - www.cambridge.org/0521571553

This page intentionally left blank

Is the universe around us a figment of our imagination? Or are our minds figments of reality?

In this refreshing new look at the evolution of mind and culture, bestselling authors Ian Stewart and Jack Cohen eloquently argue that our minds necessarily evolved in an inextricable link with culture and language. They go beyond conventional reductionist ideas to look at how the mind is the response of an evolving brain trying to grapple with a complex environment. Along the way they develop new and intriguing insights into the nature of evolution, science and humanity.

FIGMENTS OF REALITY

Ian Stewart and Jack Cohen

Figments of reality

The evolution of the curious mind



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
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Ian Stewart and Jack Cohen 2004

First published in printed format 1997

ISBN 0-511-03842-9 eBook (Adobe Reader)

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 |

ZARATHUSTRAN THEORY OF EVERYTHING:

$$E = 8$$

Preface

Isn't it strange that the animal we used to be developed into the creature that we now are? How – and why – did human intelligence and culture evolve? How did we evolve minds, philosophies and technologies? And now that we have them, where are they taking us?

The orthodox answer to these questions looks inside our brains to see what they are made of and how the various components operate. This leads to a story based upon DNA biochemistry, the evolution of nerve cells as pathways for sensory information, and their organisation into complex networks – brains – that can manipulate neural models of natural objects and processes. Mind is seen as a property of an unusual brain – complex enough to develop culture – but here the 'reductionist' story starts to lose its thread. Many people see mind as something that transcends ordinary matter altogether. Philosophers worry that the universe around us may be a figment of our own imagination.

In *Figments of Reality* we explore a very different, but complementary, theory: that minds and culture co-evolved within a wider context. Every step of our development is affected by our surroundings. Our minds are rooted in ordinary matter; they are complex processes – or complexes of processes – that happen in material brains. Our brains are linked to reality by their molecules; but they are also linked to reality on another level, their ability to model reality within themselves.

Those links have had important effects on the evolution of the brain and the mind. For example, even our sense organs are not totally pre-programmed: far from it. Instead, as we grow, our senses are 'tuned' to detect particular features of our surroundings. Mind is not immaterial transcendence: it is the response of an evolving brain to the need to survive in a complex environment. And with the evolution of culture, that environment has become self-modifying and self-referential, and human minds have done the same.

Evolution and tuneable senses have produced minds that can grapple with reality by operating upon *features* – high-level structures/processes in the brain that correspond to large-scale regularities in the surrounding world. For example, a goat eats leaves because they *look* like leaves, not because its nerve cells have a chemical affinity for chlorophyll. If plants had evolved differently,

using a purple chemical for photosynthesis, then goats would be looking for purple leaves instead; but otherwise they would be much the same as present-day goats. We shall investigate how the mind explores its own mental landscape and works with the features that it finds there. This leads to a new theory of the relation of individual minds to the human culture in which they reside.

This is a different view from that of current physics, which, for instance, sees a table as ‘mostly empty space’ because of atomic theory, and thereby directs our attention *away* from important human-scale features such as ‘wooden’, ‘solid’, ‘brown’. Such ‘commonsense’ features were important for evolution, and remain important for understanding many areas of science. For example, the evolution of the goat as a successful herbivore depended upon its ability to perceive leaves, not upon its understanding of biochemistry.

How can a conscious, intelligent mind evolve? Instead of giving a reductionist answer based upon internal fine structure we take an external, contextual view. We see the accumulating knowledge of generations of intelligent beings as a thing or process with its own characteristic structure and behaviour: *extelligence*. Extelligence constantly modifies and organises itself through continuing interactions with innumerable individuals. As a result, extelligence has become greater, more permanent, and far more capable than any individual intelligence. However, extelligence makes no sense without intelligences to interact with it: the two are ‘complicit’. The developing mind of each child interacts with extelligence by way of language, and the two-way flow between individuals and their surrounding culture changes both. Intelligence is fostered in the child, and extelligence is fostered in the culture. Thus the evolution and structure of the brain cannot be divorced from the evolution and structure of human society and its environment, the universe.

Our minds co-evolve with everything that influences them. Minds *are* figments of reality, processes going on inside structures made from ordinary matter whose behaviour evolved *in order to* mimic, model, and manipulate natural processes. This explains why they are ‘unreasonably effective’¹ at perceiving and reorganising their environment. The human condition is a complicit interaction between culture and individual minds, each shaping the other.

Culture depends upon communication, which we achieve with language. Language, the first step towards extelligence, co-evolved with brains and made minds, complicit with hands and technology, and the discovery of patterns

1 The symbol ¹ (‘note’) indicates that there is a note at the back of the book which takes the appropriate topic further, or provides a reference. Various technicalities are relegated to the Notes, as are discussions of interesting distractions.

and laws. Mind can only think about mind once language equips it with a recursive (that is, self-referential) feature-detection system. Once it has this, *self-awareness* is an immediate, essentially trivial property, because 'self' is a feature too. The existence of features makes it possible to employ a mental map instead of the real territory.

The greatest single step in organic evolution was the aggregation of different bacteria to make the nucleated cell. Similarly, the greatest step so far in our cultural evolution has been the aggregation of different cultures to make multicultures. There are many kinds of multicultural, ranging from multinational corporations to major cities like New York. But the self-complication of human culture will not stop there, because it is a self-propelled process. Today's multicultures are like the creatures of a colony, coexisting as more or less isolated 'ghettos'. Tomorrow's multicultures will be more like genuine multicellular organisms, in which extelligence is specialised like the different tissues of a complex animal. Our new communication technologies are beginning to knit all of the different multicultures into a new entity, a superculture: Humanity.

This will be our story.

And here is the place to thank everybody who has contributed to it. JC is grateful for the hospitality of the University of Warwick, which provided him with a room and a phone. IS had a room and a phone too, but then, he works there. The manuscript of *Figments of Reality* was subjected to critical reading by a variety of people. We are grateful to them for their sterling efforts, which improved the book considerably. Naturally we take full responsibility for any remaining errors or infelicities (unless they were the other author's fault). Our editors at Cambridge University Press, Tim Benton and Barnaby Willitts, deserve special mention since they were exposed to more than one version of the manuscript. In alphabetical order, the other readers were: Daniel Goldenberg, Steve Gould, Mac Hanson, Rabbi Dr Margaret Jacobi, Mike Leci, Mal Leicester, Christine McNulty, Alan Moore, Alan Newell, David Poyser, Terry Pratchett, Helen and Gareth Rees, Lena Sarah, Paulo Sousa, Heather Spears, Colin Touchin, and Elizabeth Viau.

A word about the brief stories that head each chapter. They are there because they illuminate, perhaps indirectly, the main themes of the chapter concerned. All of them are, to the best of our knowledge, true. (Except one, which we invented, for very good reasons – only to find that it contained more truth than we had thought. We confess the falsehood early in the appropriate chapter, and explain the unexpected germ of truth shortly afterwards.) Some of our readers complained that one or two of these stories were not 'politically correct'.

However, we left them unchanged, because we feel that political correctness should be secondary to truth. We emphasise that the stories are not intended to be derogatory or offensive, and if you think any of them is, then you're reading things into it that we didn't intend. For example, we mention a woman scientist who becomes emotional. We cast her as a woman because, as it happened, she was. The emotion, to us, is a positive and necessary feature of the story, and if the same thing had happened to a male scientist, he'd have been just as emotional, and we would have told it that way. Several stories rest upon human frailties or idiosyncracies, but we are not holding anyone up to ridicule. The aim of those stories is to illustrate what a strange but wonderful animal we humans are.

Figments was written in a variety of places – lots of trains between Coventry and London Euston, benches in St. James's Park in London, benches on the Thames Embankment – even sometimes at a desk. Many of the places were aeroplanes – Ryanair from Birmingham to Dublin, American Airlines from Chicago to San Francisco, Delta Airlines from Salt Lake City to Cincinnati ... One of the ironies of the complex world of the late twentieth century is that one of the best places to find the solitude to write is at 35,000 feet travelling at 550 mph.

Not for long, we fear: already telephones are proliferating across the backs of aircraft seats.

A tropical island would be more comfortable, but comfort is not conducive to wordage. Stuck on a tropical island, it is too easy to consume coconuts and rum punch all day long without ever having the stimulus to put pen to paper. Stuck in a Boeing 767 on an overnight flight, with eight hours to pass and about ten cubic feet of space to pass it in, one's only companion a heap of gin miniatures and a can of tonic water, the attraction of taking refuge in the world of the imagination is far stronger. All you need is a legal pad and a pen, or for technofreaks, a laptop.

We used both. We're not fussy and we type fast.

IS & JC

Coventry, September 28, 1996

Figure Acknowledgements

Figure 3

From *Wonderful Life: The Burgess Shale and the Nature of History*, by Stephen Jay Gould. Copyright © 1989 by Stephen Jay Gould. Reprinted by permission of W. W. Norton & Company, Inc.

Figure 21

Jack Cohen, *Reproduction*, Butterworths, London 1977, p.179, Figure 10.1.

Figure 22

Jack Cohen, *Reproduction*, Butterworths, London 1977, p.186, Figure 10.7.

Figure 24

Reprinted by kind permission of Alice O'Toole. From *Scientific American* April 1996, p.24.

Prologue

Fifteen thousand million years ago the universe was no bigger than the dot at the end of this sentence.

A tiny, tiny, *tiny* fraction of a second before that – but there was no fraction of a second before that. There was no time before the universe began, and without time, there can be no ‘before’. (As well to ask what lies north of the North Pole.) There was no space, no time, and no matter. But when the space that was coextensive with the universe had grown to the size of a dot, time had already begun to tick. The temperature within the dot was far too high for matter to exist, but there was plenty of what was required to create matter: radiation. The primal dot seethed with radiant energy.

During time’s first duodecillionth (10^{-39}) of a second of existence, the universe was a ‘false vacuum’, a state of negative pressure in which every fragment of space repelled every other fragment. Space exploded exponentially, and in that near-infinitesimal instant the universe inflated from a tiny dot to a ball many light-years across as its negative pressure literally blew it apart. As the temperature dropped the false vacuum gave way to a true vacuum, a state of zero pressure, and the era of inflation ceased. The universe, now large enough to be interesting, continued to expand under its own momentum – but more sedately, at a rate of a few thousand kilometres per second.

When time was one ten thousandth of a second old, the temperature of the universe dropped to a trillion degrees. Pairs of particles, one of matter, one of antimatter, were winking into existence and out again, born in and dying as fluctuations of radiant energy. Matter and radiation were in perfect balance. However, the balance between matter and antimatter was imperfect. For every 999,999,999 antiprotons there were 1,000,000,000 protons. From that imbalance came everything that we know.

When time attained the grand old age of one second, the temperature of the nascent universe had fallen to a mere ten billion degrees. Electrons and antielectrons, colliding in pairs, filled the universe with bursts of neutrinos and antineutrinos. Neutrons, no longer stable, decayed into protons and electrons.

Two minutes after time began (some say one and a half minutes, others three) the universe had cooled to one billion degrees, and matter as we know it began to assemble. Neutrons paired incestuously with their proton

offspring to form creation's first atoms – heavy hydrogen, otherwise known as deuterium. Deuterium fused into helium and matter began to diversify.

After half an hour the universe changed: now it was three quarters hydrogen, one quarter helium. The pace of change slowed. It took seven hundred thousand years before the universe cooled enough to become transparent to light. By then, matter had formed itself into almost a hundred different elements. It took a hundred million years for that matter to clump itself into galaxies, and for the first stars to shine.

Ripples in the early fabric of spacetime, amplified by the inexorable tug of gravity, folded in on themselves, collapsing under their own mass, leaving huge voids hundreds of millions of light-years across, bubbles of emptiness filling the universe like foam. On the surfaces of the bubbles, matter condensed into vast sheets and filaments. One such structure – let us call it the Distant Superattractor – made itself felt a billion light-years away, as its gravitational attraction sucked matter inwards towards its centre. There was nothing that greatly distinguished it from trillions of equally enormous clusters of matter.

Smaller – but still many thousands of light-years across – was a clump of matter known as the Great Attractor. Like all of the matter in a region of space billions of light-years across, it streamed towards the Distant Superattractor. Within and around the Great Attractor, matter arrayed itself into a hierarchy of ever-smaller clumps, which were sucked towards the Great Attractor even as it made its way towards the Distant Superattractor. One such was the Local Supercluster, a group of tens of thousands of galaxies that surged collectively towards the constellation Virgo at 700 kilometres per hour. The Local Supercluster was composed of more than a hundred galactic clusters, none differing significantly from the rest – the M101 cluster, the M81 cluster, the Virgo cluster, the Local Group ... A typical cluster is several million light-years across, and is composed of hundreds of galaxies; an individual galaxy is some hundred thousand light-years in diameter, a vast swirl of matter that rotates once every quarter of a billion years.

In the Local Group were two dozen galaxies: Andromeda, M33, and one – not greatly different from any other – known simply as 'The Galaxy'. Like most galaxies it was spiral in form, although unusually it had two smaller close companions known as the Greater and Lesser Magellanic Clouds. Its spiral arms – like those of most other galaxies – were density waves, places where its component atoms piled up against each other. Along the crests of those waves the pressure became so intense that it sparked nuclear reactions, and stars came into being.

In this one galaxy there were more than a hundred billion stars.

One such star – not especially different from its companions – had spectral class G2, meaning that its surface temperature was about average (six thousand degrees) and the light that it emitted (at a level of brightness that was also close to the average) showed a prominent trace of calcium. Like many stars, it was enclosed in a cloud of cosmic debris – stardust blown across the intergalactic space in shockwaves generated by explosions in the galactic core. All of the different chemical elements born in stars’ nuclear furnaces were present – some in abundance, others the merest traces. Among them, fused into existence by a coincidental resonance of nuclear vibrations, was the element carbon.

As the universe grew older, and colder, and larger, this particular cloud of stardust – like many others – began to condense, the grains sticking to each other, to form irregular lumps of methane ice, dense clouds of gas, fragments of rock. As it condensed, it also collapsed into a flattish disc, spinning on its axis, a swirl of cooling matter that collided, bounced, broke, stuck, aggregated. As time passed, a mere instant on cosmic scales, the clumps became fewer, but bigger. Crushed under their own gravity, they formed flattened spheres – planets. The G2 star acquired a solar system.

None of this was especially unusual.

Each planet, forming in its own particular place, found itself in possession of the features that its mode of formation would naturally create – a rocky core, a methane-hydrogen atmosphere, a surface flowing with molten metal or dotted with lakes of acid, encircling companions ... Each planet acquired its own identity. This in particular was true of the third planet, counting outwards from the central star. Much of its surface was covered by a thin layer of water. It had an atmosphere, mostly nitrogen. And its surface temperature was within the range at which water remained liquid. Although no other planet in this particular solar system resembled the third in these respects, it was probably much the same as many other planets around many other stars in many other galaxies.

Everywhere, even in the depths of intergalactic space, atoms bumped against each other and stuck to form molecules. On the third planet this happened more often than in the vacuum between the stars, because there were more atoms to bump into. The individual features of the third planet constrained the kind of molecule that occurred, producing structures that would not have occurred on a methane world or an ice giant. One day there arose a collection of molecules that could make copies of itself – a replicating system. Perhaps it came together accidentally in the primal soup of the oceans, perhaps it was given a helping hand by the receptive surfaces of rocks or clays. However it happened,

PROLOGUE

the replicator did what replicators do – it replicated. Over and over again. After a fairly short time the planet became distinctly unusual, its chemistry subverted and reorganised by the voracious replicator. The replicator made the occasional mistake, but some mistakes could also replicate, and soon a kind of long-term War of the Replicators was under way, as ever more sophisticated molecular collectives did battle for the right to continue replicating.

It all got rather complicated.

For instance: one group of replicators acquired the knack of converting starlight into food.

For instance: an early success, the bacterium, attained such numbers that one of its metabolic by-products, the corrosive gas oxygen, came to occupy a substantial portion of the planet's atmosphere.

For instance: other groups of replicators evolved the ability to leave the solid ground and soar upon the gases of the atmosphere.

For instance: sixty-five million years ago an especially successful type of replicator was exterminated, planetwide, by the impact of a large rock. Other tiny hairy warm-blooded replicators suddenly found that their main competition had vanished from the face of the third planet, and their rapidly diversifying successors exploded across continents and oceans.

For instance: today, two of the descendants of those tiny creatures are busy delineating their own limited version of the entire story in strange, angular geometric symbols, impressed in contrasting pigment upon sheets of compressed white vegetable matter, in the hope that other creatures of similar kind will scan the sheets with their light-detecting sensors – and in some inexplicable manner imbibe meaning and significance and make them part of themselves. Typically for these replicators we find a tiny portion of the ungraspable universe making a glorious, hopeless attempt to encapsulate that awe-inspiring whole inside its own tiny form, improbably employing weak electrical impulses that scuttle along a network of a trillion tiny fibres – vibrant, alive, and even more ungraspable than the universe that it is attempting to grasp.

A circle closes.

A mystery opens.

1 The Origins of Life

A woman scientist¹ 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.’¹ 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 'complicitly'.

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 ‘pre-requisite’. 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 autocatalysis – 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