Inventing Reality

TOWARDS A MAGICAL TECHNOLOGY

Tom Graves

Inventing Reality

Towards a Magical Technology

Tom Graves

This book is for sale at
http://leanpub.com/inventingreality

This version was published on 2014-04-02

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

Contents

Preface

Contents

- **The practical art of magic** - In which we find new ways to travel through the subjective swamp of ideas and experiences, and explore a variety of tools and techniques for merging magic and technology.
- **The joker in the pack** - In which we review a cautionary tale or two, concerning the inevitability of Murphy's Law - and the need for managing it in medical madness, mistaken magic and mixed-up muddles of many kinds.
- **What's the use?** - In which we return to the usefulness of technology, and of magic - especially when merged together into a truly magical technology which is efficient, reliable, elegant, appropriate.
- ***Further reading***

Preface to the third edition (2006)

Inventing Reality is a brief essay which explores the practical relationships between the routine realms of modern technology, and the peculiar yet equally real expression and experience of magic in the everyday world. Traditionally, these two approaches to working within the world seem to be at opposite ends of a spectrum – one subjective, the other objective. Yet in practice they have far more in common with each other than with other approaches such as science or spirituality: so much so, in fact, that they can be merged together into a truly magical technology, through which we can invent the reality – or realities –

that we need.

This new edition extends the 1998 online edition with a new chapter, 'Round The Bend', which uses the twists and turns of the traditional labyrinth to describe the process of learning new skills. Beyond that I have made only a few minor edits, to correct date references and the like, and the two quadrant diagrams in Chapter 3 have been redrawn: the original printed book, for reasons which were never explained, had contained the rough drawings rather than the finished artwork provided.

Other than these changes, the text remains essentially the same as in the original. I wrote the book in late 1985, using a brief gap in a heavy workload to put down some ideas that I'd then already been brewing for some years. Even so, it was perhaps a little premature: it took a full decade more for the ideas which were expressed in a light-hearted yet very condensed form in this book to become commonplace anywhere else, and it still seems to be experienced by many readers as a 'turning-point'. I expanded on some of its themes, in more practical detail, in the *Wyrd*[1] series; but much still remains to be done. Somewhen I'll be able to find the time to do so… somewhen… perhaps…!

Enjoy!

Tom Graves (Drummond, Victoria [Australia], May 2006)

[1] http://www.wyrdsmiths.com/booksby.htm

Acknowledgements

This essay was originally published in book form by Gateway Books[2], The Hollies, Wellow, Bath BA2 8QJ, England, first as *Towards A Magical Technology* (1986: ISBN 0-946551-30-8) and later updated and reissued as *Inventing Reality: towards a magical technology* (1990: ISBN 0-946551-64-2). I owe many thanks to Alick Bartholomew at Gateway Books, who also (in previous publishing incarnations!) successfully published two previous books of mine, *Dowsing: Techniques and Applications* (still in print as *The Diviner's Handbook*) and *Needles of Stone* (now available online at the Isle of Avalon[3] site, http://www.isleofavalon.co.uk).

Works by other authors which were quoted from or referenced in this essay are listed in the Bibliography section at the end of the book.

Regrettably, I am unable to give proper credit to the cartoonist who drew the illustrations in Chapter 3 and Chapter 6, as at present I have no record of his name.

[2] http://www.gatewaybooks.com

[3] http://www.isleofavalon.co.uk/ndlstone.html

What's in a word?

People say sometimes that technology has lost its magic – its joy, its wonder, its meaning. But if technology has lost its magic, has magic lost its technology?

Magic, technology: two awkward words. But what's in the words themselves?

And what do we mean by magic and technology anyway? How do we define those words?

It's usually magic *or* technology: two separate worlds, and never the twain shall meet. One world is real, the other false, you'd probably say: and your choice as to which is the 'true' one will say a lot about your point of view.

But the purpose of this study is to put the two words together: to look at them in an old way rather than a new one, a world in which magic and technology meet. In defining a word like 'magic' or 'technology', we limit its range of meaning; if we _un_define the words rather than define them – use them to describe rather than delimit – we open them to a wider context in which both words and their many meanings can work with each other. An exercise in thinking about thinking, and in exploring the often strange places to which this takes us.

In doing this, we can find a framework in which ideas and actions can be useful *in practice* – one in which both

magic and technology make sense in the real world we experience. The result should help us to put the magic back into technology, and the technology back into magic – moving towards a magical technology.

Is it just coincidence?

Magic is a peculiar thing. We know what we mean by the word – or do we? A dictionary defines it as 'illusion and trickery', but that's hardly what we mean when we talk of a magical occasion. To make practical use of a word, we need to have meanings from our own experience: magic as joy, as wonder, as wisdom. (Occultists talk of mages, whilst the biblical story refers to the three Magi – wise-men). Yet even these meanings are too limiting: we need words with meanings that somehow go beyond the words themselves.

Magic, many people would say, is just coincidence. Yet the word 'coincidence' is itself a prime example of the same problem. To say that magic is 'just coincidence' is to describe one known but undefinable 'something' in terms of another. For many people, each word is used to dismiss the other: neither is true, neither is valid, since, by definition, a coincidence is an event of no meaning. For others, every coincidence has meaning, part of preordained fate or, for the paranoiac, a part of a deliberate plot against them and the world. A coincidence, it seems, has either no meaning at all, or far too much meaning.

But look again at the word 'coincidence'. Literally, it's 'co-incide-ence': two events coinciding, coming together in place, in time or in any other context. And that is all: the meaning, or lack of meaning, of the co-incidence is quite

separate, to be derived from the context, the circumstances in which it occurs. So if someone asks you whether some event 'was real or a coincidence', a proper answer would be "Yes": without knowing the circumstances, the total context in which to interpret it, there is nothing more we can say.

This is far from trivial, for everything we perceive is a coincidence of one kind or another. Our senses are geared to notice change, coincidence – and have real difficulty in 'perceiving' continuity. In fact, the *only* thing we perceive – whether sight, sound, smell, taste, touch, feeling or whatever – is coincidence, the co-incidence of events of one kind or another. What we think of as 'reality', the real world, is our interpretation of those coincidences.

We tend to think in terms of 'cause *or* coincidence'; but the idea of 'cause' is itself an interpretation of coincidences, so the choice of one *or* the other is hardly valid. For example, what caused you to be reading this book? You could say that the whole of your life has been a chain of coincidences leading to *this_word, at _this* moment. And any meaning, any 'fact', you may derive from these coincidences is your choice, your reality.

> Thirty spokes share the wheel's hub;
> It is the centre hole that makes it useful.
> Shape clay into a vessel;
> It is the space within that makes it useful.
> Cut doors and windows for a room;
> It is the holes which make it useful.
> Therefore profit comes from what is there;

> Usefulness from what is not there.
> (Lao Tzu, *Tao Te Ching*)

Events are what we see, 'what is there'; the meaning is 'what is not there', the context, the part from which usefulness comes. Anything and everything you realise – literally, 'make real' – has been and still is shaped by your interpretation of events, of coincidences. In effect, what we think of as fact – literally, a 'making' – is our choice: we invent 'facts' to give meaning to those events. That's all there ever is: coincidence and its context, information and its interpretation.

What we think of as the facts of technology are better described as practical and predictable coincidences. And magic, as many people will tell you, is only coincidence. Just coincidence.

Working in the world

But to work within the world, we have to start somewhere. In this culture, it seems, most people would start by saying "Can't we explain this scientifically?" But this carries with it the assumption that everything is amenable to a scientific study – and this is itself a main component of our difficulties with working in the world, as we shall see later. Since we are dealing with *people* rather than abstract ideas, we need to start from where those people – you and me – are;

we need to start with ways of working on the world, our skills at operating within the world.

Many skills we tend to ignore: yet walking, speaking, writing, reading and the like are all learned skills. Operating many kinds of machinery – a bicycle, a telephone, a car – is so much a part of us that we tend to forget how much we have learned, how many coincidences relating to them we have learned to interpret. For example, pressing one foot hard on the floor is hardly a natural response to danger, and yet that's what you learn to do when faced with danger in a car – and you'll find yourself doing it even if you're not driving! The 'explanation' of what you do, and why you do it, says as much about you as it does about the processes in which you're involved: in any skilled work, the same results may be achieved in totally different ways, according to how each person approaches the tasks in hand.

Learning a skill is a magical process, in many ways. A skill, in effect, is how each person resolves *for themselves* the mechanics of the skill – the 'real world' – with the way in which they approach it. Technology is what we see as the outward form of skills; and magic, perhaps, is the inner form. We'll be looking at this in more detail later on; for now, let's limit ourselves to the way in which coincidence and its interpretation form a key part of skills, and thus of technology as we see it in practice.

We should remember that we need to see this *in practice*: it's all too easy to drift off into a sterile discussion of theory without any practical grounding. So it's useful to

keep some practical examples in mind. To keep up the idea of technology *or* magic for the moment, we'll use as examples, throughout this study, a skill from each side: the harsh world of computer programming, and the perhaps more dubious world of dowsing, or water-divining. What is interesting is that, by the time we've finished, you'll probably see that they turn out to be much the same: the way in which they work is surprisingly similar.

Let's start, then, with a (very) brief summary of what each skill involves.

Programming

For some people, the computer represents the ultimate de-humanising force in current technology; so it's worth remembering that a computer is only a tool. The English term 'computer' is confusing, since most computers – such as the one I'm using in writing this study – do very little computing, or 'number-crunching', at all; the French term *ordinateur* – literally, 'an orderer' – describes the process much better. A computer is, in effect, a very fast but very stupid idiot: a logic machine, following logical rules or instructions, made up in a sequence that we call a 'program'.

Each set of rules is usually built up on top of other layers of rules: a programming language, then the 'operating system', machine-language below that, and so on, right

down through the logic built into the processor, to the simple logic of switches – two states, on or off.

Within the terms of its logic, everything a computer does is 'true'. But whether it is appropriate, or useful, is quite another matter.

But let's start by looking at the basic principles. And if we keep it to the bare essentials we can break it down to just five concepts.

The first is the idea of doing things *in sequence*. (Parallel processing, which is essential for handling anything happening at speed in the real world, handles many sequences simultaneously, passing results and instructions from one strand to another). One instruction is given, then another, then another – millions of times a second.

The second key principle is the idea of naming things: 'if you don't know what it is, *give it a name*'. (This is a key principle taken from algebra: that mysterious character called 'X', for example, who was *somewhere_between 1 and 3 – or, more accurately, 'X' _holds* a number between 1 and 3.) Once named, this store or 'address' can hold some value, which might remain constant or be varied – as a 'variable' – by some other process, and can be referred to by its name. (Many programming languages, such as Pascal and C, need these names to be formally 'declared' within the program; other languages, such as most versions of BASIC, allow you to allocate names at whim; but the principle of something given an arbitrary name for practical convenience is common to all.) The

value stored in this named place is just a value: what it *means* – as a number, a letter, a place, a pointer, a reference or whatever – depends on the context. It's just information, waiting for the context to give it its meaning.

Which leads us to the third principle: an instruction can imply some *operation*, some context, using these stores by name. (An instruction is itself something that has been given a name; a programming 'language' is a convention describing a list of named instructions and their use.) Most of these instructions are exceedingly trivial – it's only the sheer speed of processing that gives the computer program any semblance of intelligence. A typical operation might be to copy a value from one place to another, or to compare two values – almost the limit at the base-level of many processors – or, in a 'higher' level of logic, to show some value on the screen.

The next key point is the idea of *'perhaps'*: a logic decision, given in most programming languages as an 'IF … THEN … ELSE …' structure – 'if so-and-so is true, *then* I'll do this, *else* I'll do that'. The context is built up logically, layer upon layer, with each decision made logically upon the context already defined by other instructions, other operations; these decisions thus become the apparent 'intelligence' of the program.

Finally, we have the idea of *'do it again'*: we can break the sequence – usually after a 'perhaps', a logic decision – and start the sequence from another place, repeating instructions and skipping over others.

And really, that's all there is to it: layer upon layer of logic based on these five concepts. For example, here's a short program using terms from the programming language 'BASIC':

```
10 LET X = 1
20 PRINT SPACE$(X) "Hello there!"
30 LET X = X + 1
40 IF X < 10 THEN GOTO 20
50 PRINT "That's a very basic piece of BASIC"
60 END
```

It doesn't matter if you haven't come across BASIC before: all that matters is that if you type this sequence of instructions on a suitable computer, and then type RUN, it produces this result on the screen:

```
 Hello there!
  Hello there!
   Hello there!
    Hello there!
     Hello there!
      Hello there!
       Hello there!
        Hello there!
         Hello there!

That's a very basic piece of BASIC
```

And yes, it *is* a trivial program. (I'll admit it's also a very old dialect of BASIC: but a more modern version that could do without line-numbers or the much-unloved GOTOs wouldn't illustrate the points here so well.)

You can see that we have a sequence: in this case, the line numbers, from 10 to 60. Each line consists of one or more named instructions.

Line 10 puts a number into a store which we label `X`.

In line 20 we use the current number in that store to tell the `SPACE$()` instruction how many spaces to print, followed by the words `Hello there!`.

In line 30 we add 1 to the number currently stored in `X`, and put the result back in `X` – so the number increases on each pass.

In line 40 we use that stored value to decide – depending on whether it is still less than 10 – to 'do it again' by going back to line 20 or to move on to line 50. Note that `Hello there!` is only shown nine times – the question in the instruction is 'Is X *less than* 10?'.

And in line 60 we say that the sequence has ended, so that another layer of logic – the 'operating system' level – can take over.

Note how the meaning of the value stored in the place labelled X has different meanings at different points in the program: its meaning, its *use*, depends on the context in which it is used.

This example is trivial: it doesn't do anything useful. But as with all computer programs, it's a set of made-up rules; the skill of the programmer is in making up sets of rules that *do* do something useful. As we'll see with dowsing, the question is not whether a program is 'true' – which it must be, by definition, since a digital computer of this type can only make true/false decisions of logic – but whether the rules it is given actually relate to the real world we experience. As with dowsing, the rules need to relate to the world in a way that is efficient, reliable, elegant, and, above all, appropriate: a way that relates to *people*. Otherwise, there's no point – an exercise in expensive and sometimes dangerous triviality.

Dowsing

You've possibly seen dowsers in action, or played with it yourself at some point in the past. In principle, you wander around with some kind of instrument – traditionally a hazel twig, but nowadays more often a couple of bent wires or a bob on a thread – looking for something or other. When you stand over what you're looking for, the instrument reacts: the hazel rod bends up or down, the bent wires cross over, the plumb bob or 'pendulum' describes a circle hanging on its thread. In other words, the response marks the coincidence of what you're looking for and where you are.

The reaction of the instrument is in fact its reaction to your

hands moving; and your hands move because you tell them to. Not consciously, but as a reflex response to something or other – *un_defined. If you like, your hands move in response to a set of rules which you invent, which state that they_should* move when that coincidence occurs. It's much the same as with riding a bicycle: you *direct* the process rather than control it. Indeed, if you do try to control it by deliberate action, you're more likely to make a mess of it than if you leave it to 'work itself'.

In essence, that's all there is to it; which is why many people think that there is nothing to it. They can't understand how anything so simple could work, *therefore* it can't possibly work.

Which is to miss the whole point. Perhaps the shortest summary of the skill is to say that *dowsing is entirely coincidence and mostly imaginary*: and the catch with that, as we've seen, is that it depends on how you interpret those two words. Literally, the dowsing reaction marks the coincidence of what you're looking for and where you are; and both of those things – what you're looking for, and where you are – are defined by images, by descriptions, by imagination.

You decide what you are looking for, by describing it as an image. Traditionally, this was done by holding beside the rod a sample of the sought-for material – a sample of water, or coal, perhaps, or some other minera; but unless you insist that dowsing is the sensation of some as-yet-undefined 'radiations' (for which there is, strangely enough, no sci-

entific evidence), there is no *structural* difference between holding a physical sample, and holding a written label, or a drawing, or even just a description of it in your mind's eye. Conceptually, they're all images.

Imagination is a powerful tool: we often forget how powerful. For example:

> Imagine that there's an orange on the table in front of you. (You choose what kind of orange it is: it's your choice, you're making up the rules here).
>
> You can see, in imagination, its colour and texture; see the way the light shows the dimples on its surface.
>
> Reach out and touch it; take some time to feel its surface, let the texture and weight describe itself to you.
>
> Now dig your fingernails in; feel and smell and sense the orange as the zest in the pith bursts out.
>
> Remove the skin, slowly, carefully; break the orange into its segments.
>
> Now put a segment into your mouth; feel the texture, then bite into it; taste the juice as it breaks through.
>
> Even though it's entirely imaginary, your senses will be at least part-convinced that it's real… very real… So is it real, or imaginary? You'll find that the only accurate answer here is "Yes"…

Now do the same with an imaginary sewer-pipe… you'll see just how powerful these images can be! It's all imaginary, and real at the same time – in an imaginary sense. And you can use that sense of reality to match what you're looking for; to note the coincidence between this image and the 'real' world.

You're also using images to describe where you are, the current position of 'here'. In basic dowsing, your 'here' is marked by where you stand: the rod (or whatever) moves when you stand over what you're looking for. Yet that's not all that practical when you're looking for something in a wall; so you change the rules, and say instead that 'here' is where you are pointing to. Then you can change the rules again, and say that, having marked the point above what you're looking for, you'll now get another reaction at the same distance away from that point as the object is down: 'distance out equals distance down', sometimes referred to in traditional dowsing as the 'Bishop's Rule'. Yet it's a made-up rule: it's a convenient image to describe something, rather than a 'fact' as such.

Follow the operational logic of that, and you'll see that there's no structural difference between someone out in a field dowsing with a hazel rod, and someone else looking for the same thing using a pendulum and pointing to various places on a map. A difference in degree, it's true, and probably a considerable difference in reliability; but no *structural* difference.

In both cases, people are following rules that they have

invented, to pre-limit the meaning of the coincidence that the rod or pendulum marks. It's a significant point that dowsing instruments have a very limited range of responses: it's easier to pre-limit the range of possible 'answers' – coincidences – in that way. In effect, you *declare* in dowsing that a coincidence shall have a particular meaning, and then set up conditions under which that coincidence can occur: if you like, you *program* the circumstances to have that particular meaning.

The catch comes in how well you can set up the conditions, so that a given coincidence *does* have the specific meaning you've declared. The point is not whether the reaction is 'true' – which it is, by definition, if you think about it – but whether it's *useful*: the key questions in dowsing are about whether the method being used is efficient, reliable, elegant, and, perhaps most important, whether it is an appropriate tool for the job in hand. And the answers to these questions depend on the *people* involved, not on the outer form of the technology.

In both dowsing and programming, the rules we use are not pre-defined: we make them up. We *choose* to follow certain rules, for convenience and for useful convention to be able to discuss our results with others: but that is a choice, not a requirement. If you like, we use those rules to explain to others how we work. Yet within reasonable limits, *anything goes*.

But what are those 'reasonable' limits? And what do we mean by 'reason' anyway? It all depends, I suppose, on

your point of view: it all depends on how you choose to explain things.

Can't we explain this scientifically?

Where once religion was the sole arbiter of Truth, science is seen by the main part of this culture as its leading light, the main source of its descriptions of the world. Science was (and still is, of course) a quest for knowledge in an abstract sense: but its public image, as presented in schools and colleges and in the media, is that of _the_way to analyse reality, using logic to define and delimit what is real and true (and what is not) in ever more minute detail.

In this context, technology is 'applied science'; religion is seen as an anachronism, poetry and the arts an irrelevance, whilst magic is nowhere to be seen – an aberration of the mind now finally eradicated by the ever-increasing progress of science in its explanations of the world and reality.

Our image of science is that it *explains* things for us. In fact, science is credited with all change, all advance, all progress. Which, as we shall see, is a little unfair: in reality, science has very little to do with it – and it has very little to do with science as practised. It's all a matter of your point of view.

What we call a point of view is better described as a *filter*, selecting out of the mass of information, of coincidences, those items that we consider to be valid, to be 'signal'

events: everything else is just 'noise'. And yet we choose that filter, that definition of what is real and what is not. In the same way, we let it choose us: we can see only what the filter will let us see; a pair of rose-tinted glasses only shows us a rose-tinted world. If we cannot change the filter, we cannot change the way we see the world.

Yet we face a fundamental paradox. As Stan Gooch put it in a letter to the *New Scientist* magazine:

> Things have not only to be seen to be believed – but also have to be believed to be seen.

Each point of view seems 'right'; the mistake, perhaps, is in assuming that if I'm right, then by definition everyone else must be wrong. As Edward de Bono put it in his book *Practical Thinking*, 'everyone is always right, but no-one is ever right': they may be right from their own point of view, but they simply do not have the information available to be right in a total sense, to be truly 'objective'. And you can't really be objective, as most public descriptions of science claim to be, if all you have is a point of view. de Bono went on to describe two 'laws of thinking':

> [1] An idea can never make the best use of available information, because that information trickles in over a period of time.
>
> [2] Proof is often no more than a lack of imagination, in failing to see an alternative hypothesis that would equally fit the facts.

Or, as he put it in a rather less gentle form, 'certainty comes only from a feeble imagination.'

There are plenty of certainties around, pre-packaged ways of viewing the world: the catch is that the real world we experience never quite seems to fit these world-views. A political ideal like the beautiful communard slogan 'from each according to ability, to each according to their need', seems to come out in practice as 'from each according to facility, to each according to their greed'. And religious world-views, for example, have a sad habit of assuming that people are 'perfect' – conforming to some arbitrary ideal – when they just aren't that way inclined.

Another example, the concept of the 'supernatural', is a crucial one to our study. The concept itself is another product of this limited thinking: the assumption that a scientific reductionism can explain everything also implies that the logical structure, the world-view, that it builds also defines and delimits the real world, the natural world. Anything outside of the structure is 'supernatural' – literally, above the natural – and, by a tortuous piece of circular reasoning, cannot exist: except, perhaps, as a product of the imagination.

Yet this distinction is quite artificial: if something is perceived, it's natural. It exists, even if 'only' in imagination – so it's natural. If everything is coincidence, and our perception of events is limited only by what we choose to perceive, then there can be no distinction between real *or* imaginary, natural *or* supernatural.

To work on the world *as it is experienced*, to work on it in different ways, we need to be able to explain the world in a variety of ways: we need to be able to change our point of view.

Points of view

In the last chapter we saw that coincidence and meaning are quite separate – we can't really say that any thing or any co-incidence can be said to be caused by any other event, point of view or whatever. Anything goes.

The trouble with that concept is that if anything goes, we are left with nothing with which to make sense of the world. Without some way of handling it, we have no way of predicting anything at all, and everything and nothing happens at the same time – a certain recipe for instant insanity. So we need some kind of model which allows us the flexibility to allow anything to happen, yet still operate in something resembling a controlled manner.

One approach which is useful is to separate the information we experience from its interpretation, and describe this content and context each as an axis in a simple two-dimensional diagram. In one direction, we have a spectrum of information, ranging from outer tangible or sensory data, to information we derive from within ourselves, as feelings and intuitions. The other axis describes a range of methods of interpretation, from indefinable value judgements through to the strict true/false analyses of logic:

a spectrum of interpretation between value and truth – whatever either of those might mean. The model looks something like this:

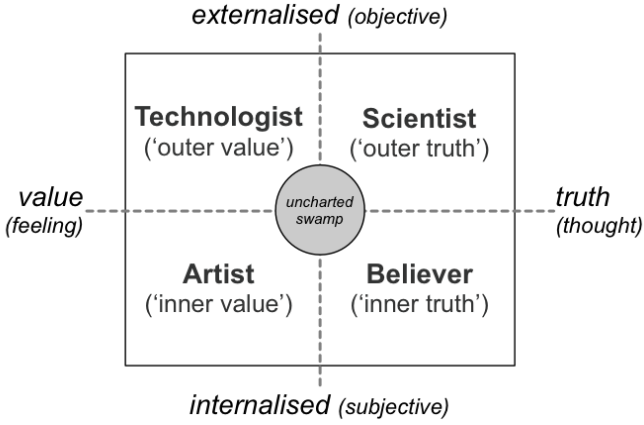

This gives us four quadrants, or modes, in which we both collect information and interpret it: inner value, inner truth, outer truth, outer value. Each of these quadrants is only a way of working on the world, a mode to describe reality as we experience it through using that way of interpreting the world. Reality is, if you like, the sum of everything that could be experienced through these four very different ways of working on the world.

This kind of model of reality can be found in Carl Jung's work, for example, or Ken Wilber's; but there is a particularly interesting variant on the theme in *SSOTBME*, a delightfully bizarre book on 'thinking about thinking'. The book was published anonymously, so for convenience I'll

refer to its author as Leo (a simpler name than his fictitious character 'Lemuel Johnstone').

Leo builds this model by describing the whole of reality as a swamp. Not a featureless swamp – every point of view, every experience, every possible coincidence of events is included. There are also endless opportunities to wallow in the mires of confusion, and to disappear beneath the surface without trace… It sometimes seems that the safest move would be not to move at all, to stay still. But even that isn't certain: the surface seems to quake with the tide of events, so that even the safest-seeming point of view will seem doubtful after a while. Nothing stays the same for long: indeed, the only real constant is change itself.

As Leo puts it, there are four main ways in which to exist within this kind of reality. Each one coincides with a quadrant of the model above: inner value, inner truth, outer truth, outer value. Each is best described as a *mode* of operation, in which certain possibilities – such as movement, in this sense of moving from one point of view or one experience to another – exist solely because others – such as stasis, developing one particular point of view – do not occur.

The first way of working on this world is to skim the surface of the swamp, travelling in a hydroplane at high speed. The whole point is the speed, and the variety of ideas and experiences that come from just travelling about with no particular place to go.

This is a mode of *inner value*, which we could call the

artistic mode.

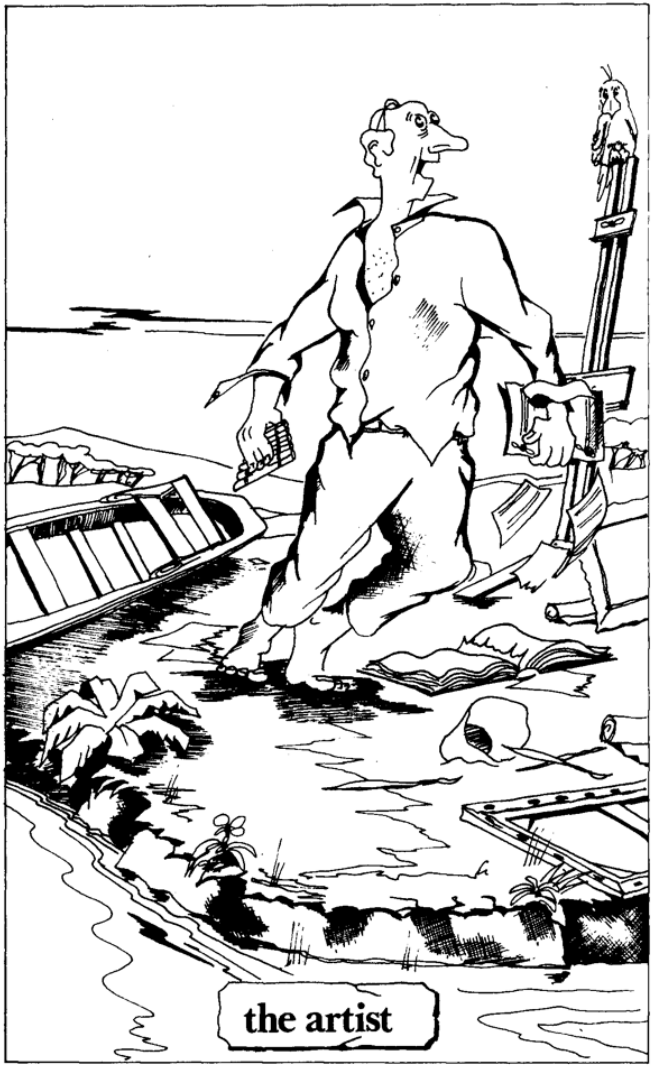
the artist

Playing with this description a little further, we can see that this is hardly a safe way of operating within a swamp: it's all too easy to crash into some unexpected experience, to run out of fuel (inspiration?), or to decide to settle down in some uncharted spot with no hope of future supplies or common experience shared with anyone else. But it's certainly the quickest way of scanning a wide range of experiences and points of view – although, at that speed, it's not going to be too easy to make sense of anything other than that they were, indeed, experiences and points of view.

Another *modus operandi* seems quite the opposite of the artistic one: to develop one point of view as far as it will go, right out into another dimension. You state that that point of view is true – inviolably and absolutely true – and build on it, like a pole in the swamp.

This is a mode of conviction, of faith, of *inner truth* – the mystical or religious mode.


the mystic

Again, playing with this image a little further, the higher you climb up the pole, the more of the swamp will come within your view: the more you climb, the more true will seem the point of view. In the distance you can see other poles, other points of view – some of them way out in the distance indeed – but you can hear that experiences from those poles, especially from further up each pole, seems much the same as your own. The mystics, those people who are well and truly up the pole and with their heads in the clouds, can see and share a vast range of vision – even though most of it seems like cloudy thinking to us.

The only trouble with this mode is that you can't actually *experience* anything else, since, by its definition, you have to stay with that one point of view; and it seems a sad fact that each pole has to be counterbalanced by a vast morass of struggling bodies, each of whom has grasped the pole and disappeared beneath the surface, screaming "I have the truth" as they did so.

A friend and I were once intercepted at the station by an evangelist wanting to give us 'coffee and the word of God'. It was a predictable set-up, and the coffee was poor, too.

After an hour of circular 'discussion', my friend yelled "This is a f***ing waste of time", and left – which was sensible, as we were about to miss the

last train. I said "I'd better go" and made for the door.

Our evangelist said "God Bless You" (I could hear the

capitals); I stammered "Oh, G-god b-bless you". The reply came as I left the door: "But he already has…"

The third mode in this model is to build a solid platform, a safe predictable area in which everything is true and interrelated in logic. Everything is patently obvious, there are no surprises on the platform itself – although around the edges things may not be quite so predictable as they seem.

This is a world of *outer truth*, a scientific world.


OCT
NOV
DEC
the scientist

To many people on the platform, the platform itself *defines* reality, and encompasses the whole of truth. To this point of view, which we could call public-science or 'scientism', anything beyond the platform is unreal; their duty is to build higher and higher walls around the platform, to protect the good citizens from the ignorance and superstition beyond. In fact this has very little to do with science as practised – we could suggest that these are same people who would have screamed "I have the truth" around the poles of religion, except that the solidity of the platform prevents them from decently disappearing beneath the surface as would have happened elsewhere.

The platform is woven between a group of poles, more often called the ivory towers of academia; their mystics are the 'pure scientists' whose breadth of vision is matched only by the impenetrability of their thought. And at the edge of the platform are practical scientists, researchers working at the limits of the known world – having discovered, by some means or other, some unintended hole in the fortress walls of scientism.

Scientists, says Leo, are like people in wheelchairs – they need firm level ground to move about on. To move, they must extend the platform, extending the boundaries of science, cutting down shady dogmas and filling in soggy hypotheses (to use Leo's graphic image). But when they arrive at some new place, it is just as boring and predictable as anywhere else on the platform – hobgoblins and foul fiends having rushed away at the sound of the first myth

being exploded. For the problem with this mode of working is that, by the time it has finished, what it seeks has ceased to be the swamp, has ceased to be reality as it is (or was) – it is just an artificial platform, an 'objective' world with no room for personal experience at all.

But working away at the edges of the platform are another group, commonly but quite wrongly called 'applied scientists'. At one edge you'll find the psychiatrists, not bothering too much about which theory is absolutely true, but using ideas from Jung, Adler, Freud, Laing and anyone else's work they can lay their hands on. And at another edge there'll be electrical engineers switching between wave and particle theories of light and energy, blithely unconcerned about their mutual incompatibility in logic.

At first this does look like science, but only because of the safety-line of 'if it doesn't work, go back to theory' – in other words go back to 'outer truth'. But in fact this is a quite different mode, in which you carry the platform with you, spreading your weight on swampshoes to allow you to move with relative freedom from place to place, idea to idea, to find a point of view which is _useful_at that time, rather than supposedly 'true' in any absolute sense.

This is a mode in which truth is defined in terms of whether it has practical value, *outer value.*


the magician

We could describe this as a technological mode. But it has an older and easier label – a magical mode. Whilst there is a real structural difference between technology and science, *there is no structural difference between technology and magic.*

It's entirely true that, as Arthur C. Clarke put it, "any sufficiently advanced technology is indistinguishable from magic". But it's also just as true that any sufficiently advanced magic is indistinguishable from technology. They may *seem* different – sometimes very different – but the underlying approach to reality is exactly the same.

The only difference between magic and technology, in practice, is that magicians tend to be a little way out in the distance – they may be seen wandering to astrology, alchemy and other forgotten, part-ruined areas of the old platforms of science, to rest by some religious pole, or travel as fools where angels fear to tread.

This mode is hardly *safe*, as the platform of science may seem to be, but at least it works on the swamp as it is; and the point of the exploring is not to find out how true a belief, a point of view, may be, but to put it to *use.*

Technology is not 'applied science'

If we use this model, we can see clearly that technology is not 'applied science'. In fact, it's quite the other way round: in our culture, science is a codified summary of the

practical experience of technology, simplified to describe what would happen under nonexistent 'perfect' conditions. Reality is always a little different from the niceties of theory: for example, the pressure-temperature-volume equations, the 'gas laws' we learnt at school, only apply to a perfect gas that doesn't actually exist anywhere in the real world.

One of the biggest problems in technology is not the neat precision of theory, but of getting things to work in the real world. Leo, the writer of *SSOTBME*, said that one of his problems working as a mathematician in aircraft design was to make his maths sufficiently _im_precise to be useful. An important recent advance in computing practice has been the development of 'fuzzy logic' – the mathematics of being precisely imprecise – which puts the inherent uncertainties of chaos theory to practical use.

Varying tolerances in materials can make a mockery of any design, too. Low-cost versions of electronic components, for example, may be up to twenty percent away from nominal values: multiply that a few times and things can go a long way from the neat predictions of the textbook. One of the real arts of electronic design is in building circuits that cancel out any problems from loose tolerances and lower-cost components. There's a popular truism, too, that defines an engineer – a technologist – as 'someone who can make something well for a pound that any fool can make badly for ten'. In that sense, technology is an art-form in its own right.

Historically, too, most key developments in technology have come *before* the science, not after: intercontinental radio and heavier-than-air aircraft were considered impossible until theory caught up with practice. The idea of scientific research actually driving progress is mostly a myth: far more often it's been someone chasing an idea, or looking for commercial advantage. Technology does the work, but technologists call themselves 'applied scientists' to improve their credibility rating in the academic sweepstakes; science gets the credit.

In effect, science is presented as the controller of a linear progress from idea to hypothesis to theory to law:

But we can also see that this is simply another way of looking at the same quadrant diagram:

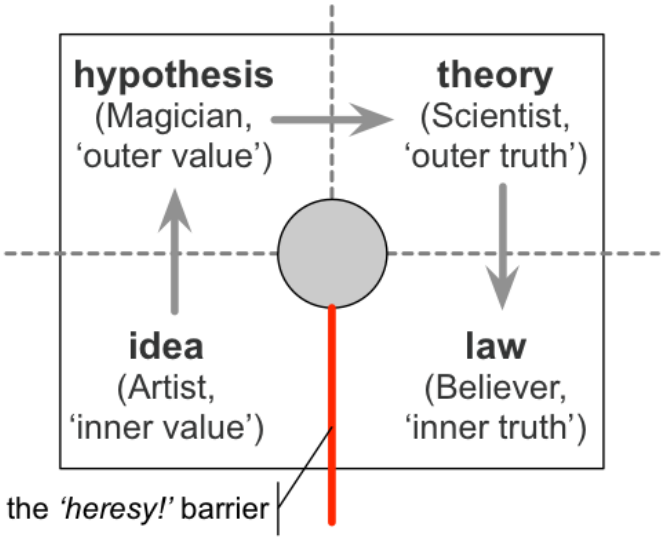

…except that a rigid wall (the 'heresy barrier') has been interposed between the finality of 'scientific law' – 'inner truth' – and the chance of any new ideas. Historically, science has progressed as a sequence of 'scientific revolutions', to use Thomas Kuhn's phrase: the Copernican revolution, the Newtonian revolution, relativistic mechanics and so on. An idea is put into practice as technology, is codified as theory, and frozen as scientific law: and the rigidity and finality of law prevents us from seeing the *context* in which that 'law' is true. The 'law' is how reality *ought* to be if the world was perfect: the law tells us God's intentions, as the religious person would put it. Or perhaps it's us telling some God how the world ought to have been made … ?

Explanations, explanations

What it's really about is *explanations*. An explanation is a model, a prediction of how something will work. Taken in a wider context, a group of explanations – a 'paradigm' – attempts to define how the world 'really works': and all too often, if something doesn't fit the explanation, *it doesn't exist*.

The explanation is seen as being more important than reality, the totality of the swamp. Because the explanation *predicts*, it makes the world *safe*. Or seem safe, at any rate. Everything is known, everything is predictable, there is no supernatural, 'I am the Way, the Truth and the Life, there is no way to the Father but by Me', 'there is no God but God, and Mohammed is his prophet…' – which can lead to a rigid state that Edward de Bono described as 'nothing-but-ness'.

A central facet of this kind of closed explanation is the ascription of one unknown to another – a tautology – to prevent any questioning of the belief-system. "There are only four forces in the universe – weak nuclear, strong nuclear, electromagnetic and gravitational; and everything is causally connected", wrote scientist John Taylor, neatly sidestepping the fact that we do not know what any of these forces are, and have no causal explanations even for some key concepts such as radioactive decay.

An explanation is used to 'explain away' anything which doesn't fit. Armchair Freudians use a system so neatly closed that if you say you don't need psycho-analytic

treatment, that statement is proof that you need it! And I've had many a battle with the 'religious left' who believe so strongly in democracy that, as the only true representatives of 'the people', only people who agree with them may be allowed to speak or vote – anyone else is, by definition, 'an enemy of the people'.

A few decades ago, ten young Germans became convinced that their country was a rigid police state. The group, led by Andreas Baader and Ulricke Meinhof, conducted bombings and 'executions in the name of the people', to fight the cruelty and injustice of the state as they saw it.

The state's response was that all police were mobilised, road blocks and 'search-on-suspicion' orders were imposed throughout the country for years, and civil liberties and legitimate forms of protest were severely curtailed: the characteristics of a police state.

Germany hadn't actually been much of a police state when the Baader-Meinhof group started; but it certainly was when they had finished.

Their belief system hunted out the 'facts' it needed to 'prove' that its assumptions were true at the end; if the facts didn't exist, it *created* them. It had no means of testing whether its assumptions were valid in the first place, or whether the actions dictated by those beliefs were appropriate at all.

When explanations are used in this way, only those things predicted by the belief-system can be seen to happen: which is, of course, the whole point of using an explanation in this way. Remember that 'things have not only to be seen to be believed, but also have to be believed to be seen': so the assumptions about what will happen are also the cause of what does actually happen.

The danger is that while theorists and politicians play at 'what I tell you three times is true' (to quote Lewis Carroll), the real world isn't quite so accommodating. Whilst change is not always for the better, a rigid world-view that cannot cope with change at all will, sooner or later, sink into the swamp, often dragging others with it into insanity beneath the surface.

> The writer Dr Johnson was once standing in a street where the houses came closer and closer as they reached towards the sky. He saw two women hurling abuse and vegetables at each other, almost within frying-pan range from the topmost floor of two houses on opposite sides of the street.
>
> "They can never agree", said the good doctor, "for they are arguing from different premises."

An explanation is a tool, not a 'fact'; it's a way of describing what was perceived. Like technology itself – 'the use of

tools and techniques' – we can either use them or, through ignorance, be used by them. It's our choice.

Where people believe that their safety and sanity depend on a belief structure, an explanation-system, that structure is more important than themselves. They no longer direct it: *it_controls _them.*

The logic, 'the Truth', that they perceive tying the structure together is like a single thread holding together that platform of science – 'the real world' – linking everything in one unified structure which defines their world-view, their definition of what is 'sane' (literally, 'healthy') and what is not. Break the thread in just one place, though, and the whole complex edifice collapses… and with it their sanity. So it's no wonder people will die for a belief, will go to war over beliefs: structures like this are terrifying fragile.

The world-view becomes rigid, and the social definition of sanity is itself insane: a mad world in which security can only be achieved by the threat of 'mutual assured destruction'. To protect the strand of logic, 'reality' can only be seen in terms of the structure, is in fact *defined* and delimited in terms of the structure. And we have a tautology: the structure is 'proved' by logic to be true in terms of itself, from its own viewpoint, in order to prove that the logic was, indeed, logical.

But we knew that already: what is the use?

'Use' comes from the context, from 'what is not there'. No doubt you remember, from your school-days, that the sum

of angles within a straight-sided triangle is 180°. It's part of the logic of Euclidean geometry. But this is only true for a flat surface, not a spherical one, as many sailors have found to their cost: on a sphere, 'straight' lines are curved into another dimension, and the sum of the enclosed angles may be anything up to 540°. In other geometries the values are even more bizarre: so the truth, the 'right' answer, changes according to the context.

Yet how do you tell, sailing on the sea, whether the world is flat or curved? And if curved, whether it is spherical, or cylindrical, or sausage-shaped, or anything else? Common sense, after all, insists that it is flat; you need a context greater than your immediate perception to understand the idea that the world we stand or sail on is curved. And even that is an assumption – a *useful* point of view. In any technology that deals with the world as it is, rather than as some theory pre-defines it, anything may be 'true': what matters is its *usefulness* – whether it is efficient, reliable, elegant and apt.

> "What is matter? And does it?"
> "Now this", he said, raising his glass, "is a very different matter. And it does!"
>
> (advert for *Guinness* stout, c.1970)

Where beliefs and points of view are seen as tools rather than final Truths, a world-view is held together by many strands: some may break from changing circumstances, or

simply from old age, but the structure as a whole is flexible enough to withstand it. A tatty collection of old ropes of thought may not look so neat as a so-refined single strand of logic, but it least it isn't fragile; the structure may sway a little in the winds of change, but it's not so likely to collapse without warning in the minor earthquake of some 'scientific revolution'.

So, to work on the world as it is, rather than on how some belief defines it to be, we need some way of seeing the context wide enough to select an appropriate point of view. We need to move from explanations that *define*, to explanations that *describe*; we need to move from point to point within the swamp, yet have some way of knowing where we are.

We call this process 'learning'. And yet, since to move from belief to belief is unstable, dangerous, we also call it insane. Learning a skill, learning to work with the world in some new way, could also be called 'the subtle art of insanity' – and that's what we'll look at next.