GRESHAM

COLLEGE

TRAVELS WITH MY ANT

A Lecture by

PROFESSOR IAN STEWART MA PhD FIMA CMath
Gresham Professor of Geometry

29 October 1997

GRESHAM COLLEGE

Policy & Objectives

An independently funded educational institution, Gresham College exists

- to continue the free public lectures which have been given for 400 years, and to reinterpret the 'new learning' of Sir Thomas Gresham's day in contemporary terms;
- to engage in study, teaching and research, particularly in those disciplines represented by the Gresham Professors;
- to foster academic consideration of contemporary problems;
- to challenge those who live or work in the City of London to engage in intellectual debate on those subjects in which the City has a proper concern; and to provide a window on the City for learned societies, both national and international.

Gresham College, Barnard's Inn Hall, Holborn, London EC1N 2HH
Tel: 020 7831 0575 Fax: 020 7831 5208
e-mail: enquiries@gresham.ac.uk

Gresham Lecture
Travels with my Ant

Ian Stewart 29 October 1997

The role of science is to seek simplicity in a complex world. Nature exhibits many remarkable regularities and patterns, and science works by assuming that these patterns arise from the regularities and patterns of the underlying 'laws of nature' that govern them. For example the elliptical orbits of planets discovered by Kepler are more or less direct consequences of Newton's laws of gravitation and motion. This is a comfortable picture, which encourages a view of the relation between laws and their consequences — between cause and effect — that might be characterised as 'conservation of complexity'. That is, simple rules imply simple behaviour, *therefore* complicated behaviour must arise from complicated rules. This assumption, commonly tacit, has directed major movements in science. For example, it is why we think of the complexity of living things as a puzzle: where does the complexity 'come from'? Until very recently hardly anybody would dare to suggest that complexity need not 'come from' anywhere.

One problem with the conventional view is that, if simplicity is inherited directly from rules to behaviour, then it is difficult to reconcile a complex universe with the presumed simplicity of its rules. The usual answer is that the complexity arises through the interaction of large numbers of simple components. But recently the idea that complexity is conserved has received a series of mathematical challenges. One was the discovery of chaos in dynamical systems. Chaos can be characterised as complex behaviour arising from simple rules. Complexity Theory emphasizes the converse: that highly complex interactions taking place in large populations of systems can conspire to create large-scale but simple patterns. In mathematics and elsewhere, what philosophers call 'emergent phenomena' are becoming respectable again.

Langton's Ant

A simple example of emergence occurs in *Langton's Ant*, a cellular automaton invented by Chris Langton of the Santa Fe Institute. It is an excellent example of the ability of Complexity Theory to generate new concepts and reveal new types of behaviour in simple rule-based systems.

Begin with a grid of squares, which can be in one of two states: black or white. For simplicity suppose that initially they are all white. The ant starts out on the central square of the grid, heading in some selected direction — say east. It moves one square in that direction, and looks at the color of the square it lands on. If it lands on a black square then it paints it white and turns 90° to the left. If it lands on a white square it paints it black and turns 90° to the right. It keeps on following those same simple rules forever.

Those rules produce some surprisingly complex behaviour. For the first five hundred or so steps, the ant keeps returning to the central square, leaving behind it a series of rather symmetric patterns. Then, for the next ten thousand steps or so, the picture becomes very chaotic. Suddenly — almost as if the ant has finally made up its mind what to do — it builds a *highway*. It repeatedly follows a sequence of precisely 104 steps that moves it two cells southwest, and continues this indefinitely, forming a diagonal band (**Fig.1**). This relatively simple large-scale feature *emerges* from the low-level rules.

Theories of Everything

Langton's Ant is more than just a neat mathematical gadget. It opens up some deep questions about how scientists explain the universe. I'll take a look at some of these issues now, and return to the Ant later on to see how it illuminates them.

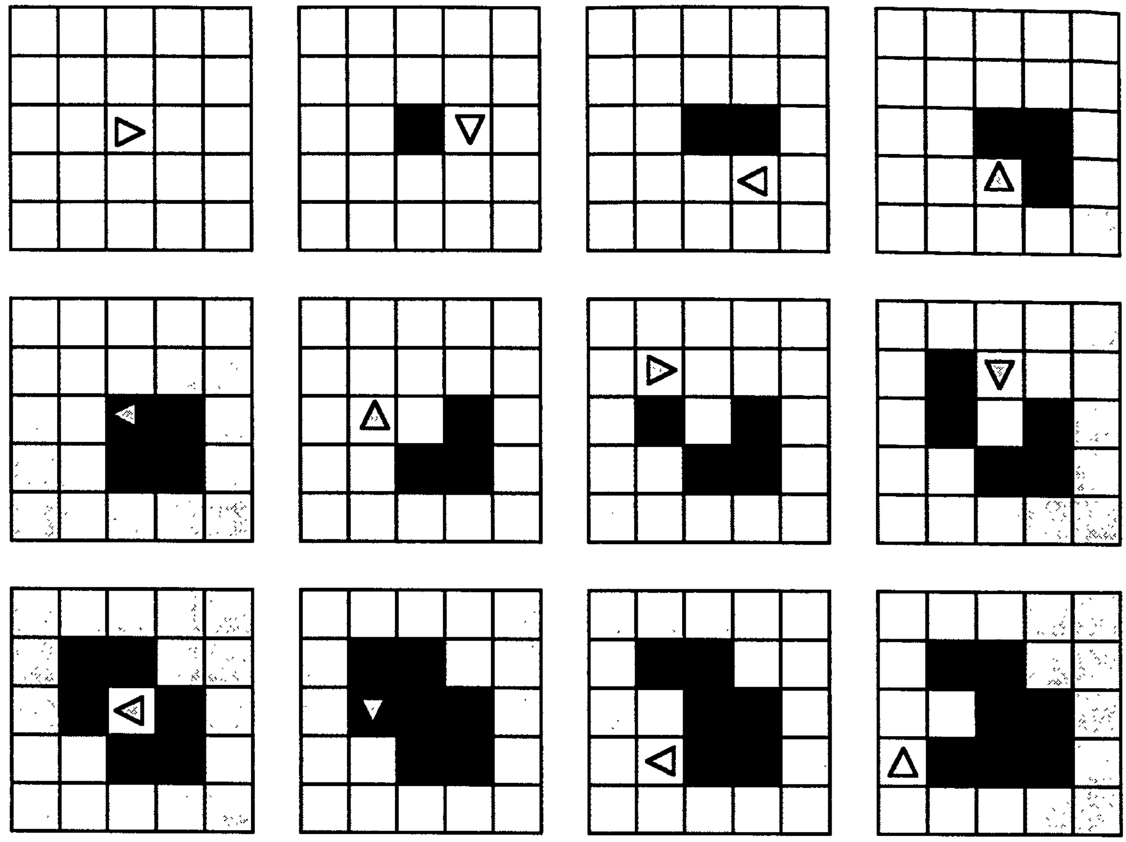

Fig.1 The first twelve steps followed by Langton's Ant (grey arrow). (For clarity, squares not yet visited are shown light grey: these should be treated as 'white' when applying the rules.)

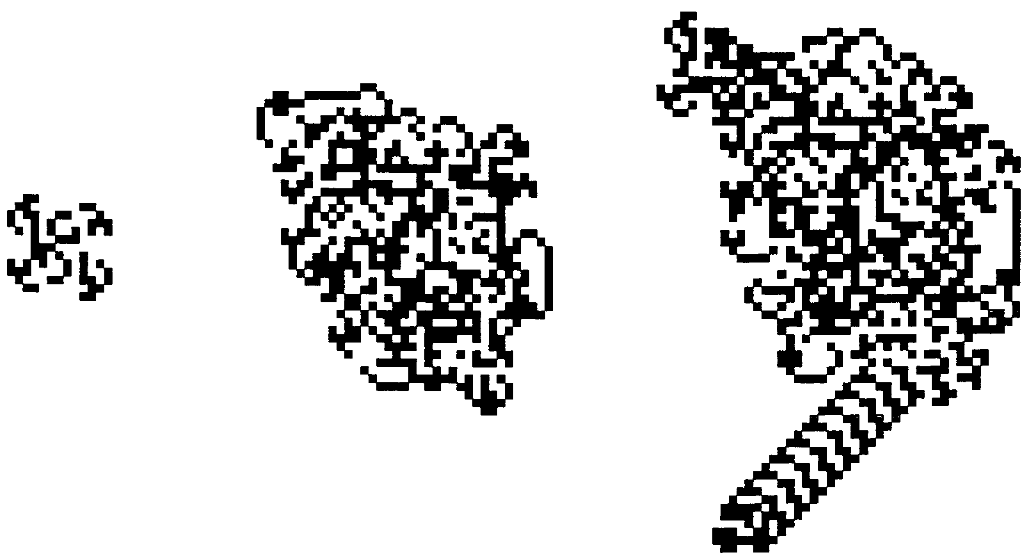

Fig.2 Three stages in the infinite journey of Langton's Ant: symmetric, chaotic, and highway-building.

In the traditional view, a scientific explanation of some phenomenon consists of a deduction, from natural laws, of that phenomenon. We look down a 'mental funnel' from a natural phenomenon, and 'see' the underlying 'laws' or rules (**Fig.3**). Notice that the arrow of explanation runs upwards, from rules to phenomena — whereas the arrow of discovery runs in the opposite direction.

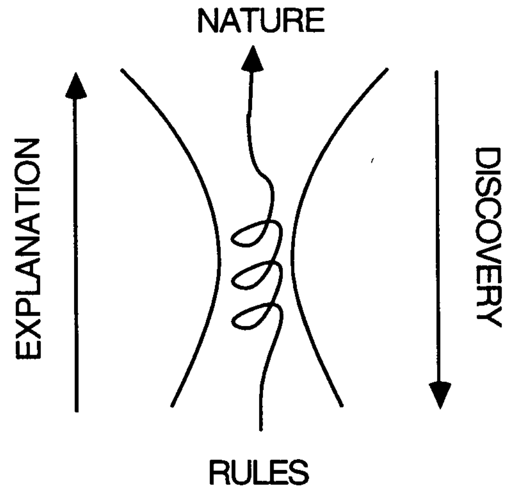

Fig.3 Anatomy of an explanation: conventional view.

This paradigm of 'explanation' has emerged over a period of centuries, and began its spectacular rise to prominence in the work of Isaac Newton and his contemporaries — although it can be traced back through Kepler, Galileo, Aristotle... . As scientists looked down more and more mental funnels they started to find common rules. For example the rules known as 'quantum mechanics' can be found down the funnel leading from chemistry, where they explain chemical bonding. They can also be found down the funnel from cosmology, where they explain the origin of the universe in the Big Bang (**Fig.4**).

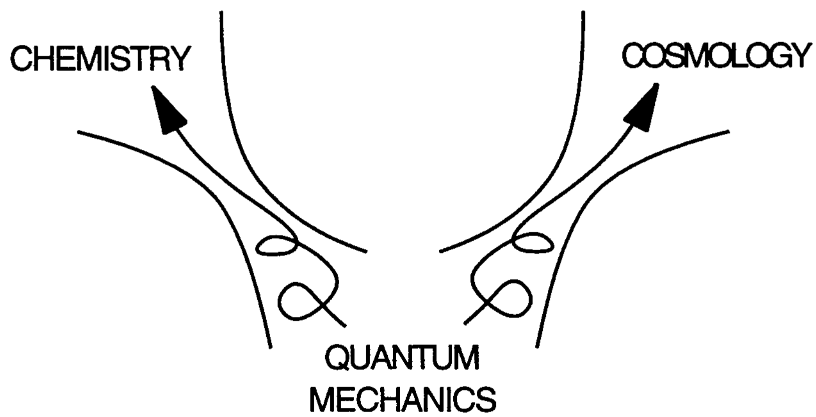

Fig.4 Funnel diagram for quantum mechanics.

It is very impressive to find the same rules down two such disparate funnels, explaining such different things, and this led to a general feeling that rules that are common to many funnels must be somehow more 'fundamental' than those that are not. Most of the funnels in modern physics ultimately end up in two systems of rules: quantum mechanics and general relativity. Unfortunately these two systems are mutually contradictory. Quantum mechanics is indeterminate and considers matter to be ultimately indivisible, whereas general relativity is a determinate theory of continuous space and time. The contradiction is philosophical rather than operational, in that either one viewpoint, or the other, is appropriate to most questions. Nonetheless, it shows that neither theory can be considered truly fundamental. One way out of this impasse would be to find a 'deeper' set of rules that explained both quantum mechanics and general relativity. This long-sought synthesis has been dubbed the Theory of Everything, on the grounds that — in our language — it lies at the bottom of *all* funnels (**Fig.5**). If we can find it, so the rhetoric goes, then there will be some mathematical equation, simple enough to wear on a tee-shirt, that will proclaim the end of physics.

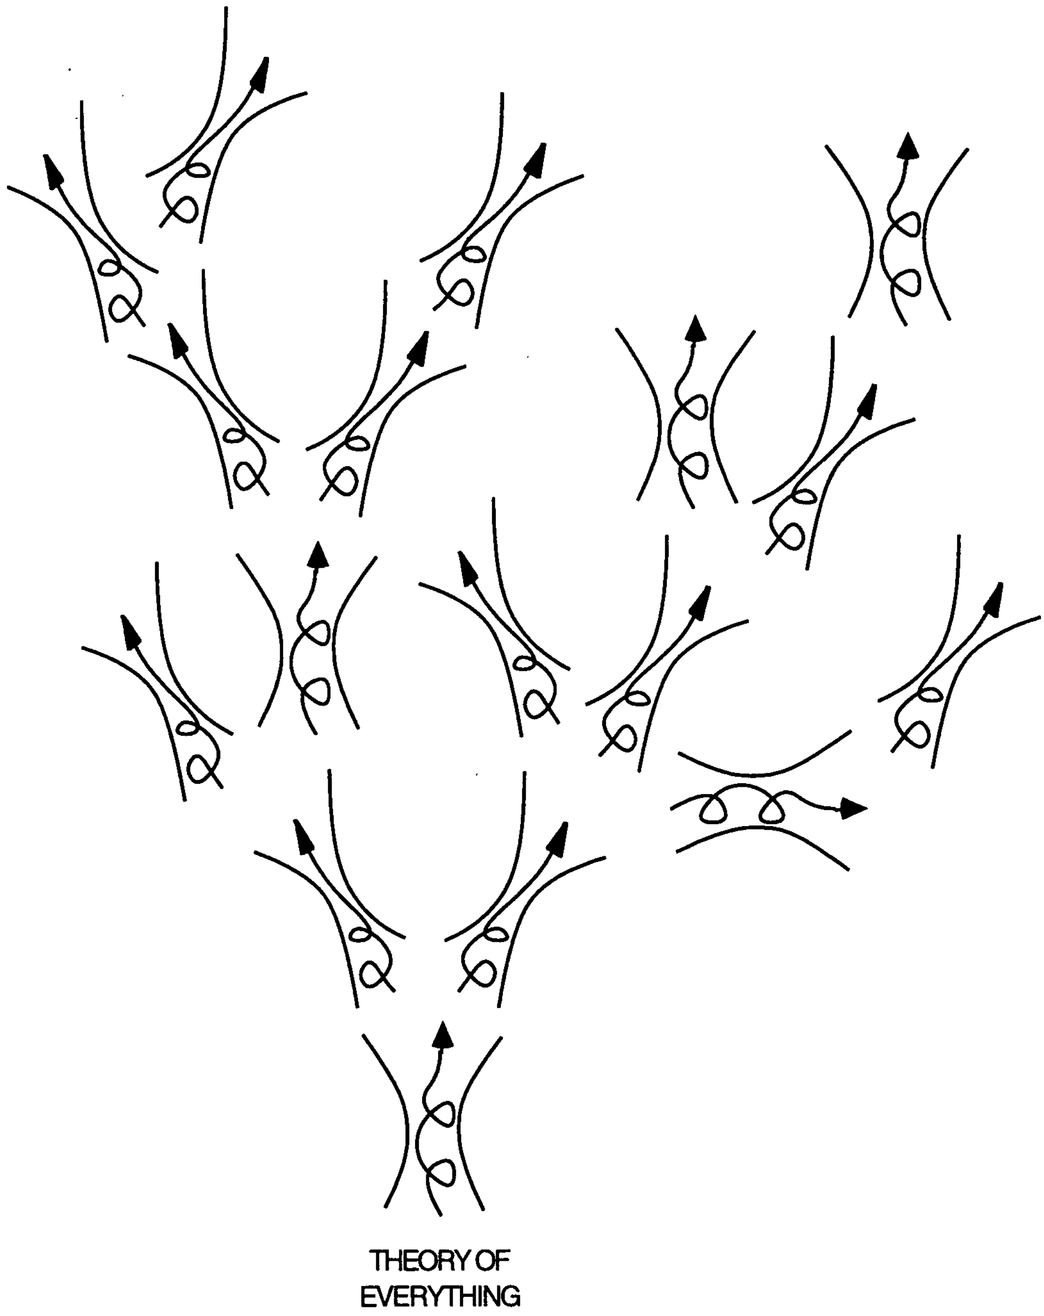

Fig.5 Is there a single Theory of Everything at the bottom of every funnel?

The philosophy involved in such a concept of explanation is called *reductionism*, though the tee-shirt version is admittedly an extreme case. Reductionism leads to a model of science as a hierarchy of rules, each valid on an appropriate level of explanation. The rules on a given level are all consequences — at least to some useful degree of approximation — of those on lower levels. Deeper levels are more fundamental, and if a Theory of Everything exists, it lies at the deepest level of all. In an extreme but common view, it is not just a useful approximation to the truth: it *is* the truth.

The Goats in the Machine

To what extent does the reductionist vision of science correspond to how the universe actually operates?

The complexities of the universe arise in an enormous variety of ways. Take a deep breath: that may seem simple enough, but the oxygen you are using was separated from water by the sun's energy, probably by chlorophyll in a leaf; and the molecules from which it is comprised were assembled in leaves all over the globe. It takes ten pages of complicated chemical and energetic equations to explain the biology of photosynthesis in leaves; and science is quite literally incapable of predicting the weather-patterns that brought the resulting oxygen molecules to your lungs. And once it has arrived there, things become no simpler. The diffusion of oxygen into your blood depends upon the loading rules for haemoglobin as it passes through your lungs. Haemoglobin and the photosynthetic pigment chlorophyll are enormously complicated chemical machines, whose changes of shape as they function can barely be modelled on our most complex computers.

Now watch a goat eye a rosebush. Think of all the cones in her eye, all the connections to her brain, all her muscles as she walks; all the image-processing algorithms carried out by her visual cortex, all the control signals sent from brain to muscles. Watch her amble over and take a mouthful of leaves. The way her chewing muscles grind the leaves is barely understood by specialists, and as for what happens to the leaf pulp when it meets the bacteria in her special stomach...

Yet there are large-scale simplicities too. There are simple ecological models that explain how goats eating leaves have turned the Sahara from a fertile plain, providing ancient Rome with much of its food, into a sandy waste. They warn that the Greek goat/olive-tree economy is heading in the same direction — but try to explain that to a Greek olive-grower.

We can scale the same kind of story up to cosmological dimensions. The basic ideas of Einstein's General Relativity are apparently simple, but our best computers can model only very simple systems according to that scheme. Basically, we are restricted to one-body systems, either one star, one black hole, or one universe. A binary star is beyond the current reach of General Relativity — not just as regards solving the equations, but as regards writing them down to begin with. The same goes for an accurate analysis of the solar system. There is more than a hint of hubris in a claim that we understand the origins of the universe, when the mathematics involved is incapable of explaining binary stars or the solar system.

Everywhere you look there are things like molecules, solar systems, leaves, and goats, which are much too complicated even to begin to understand; and processes that can be followed only in very simplified versions by experts. Instead of seeing simplicities down the reductionist funnels, we are trapped in the *reductionist nightmare* (**Fig.6**), in which the funnels keep branching forever without hitting bottom.

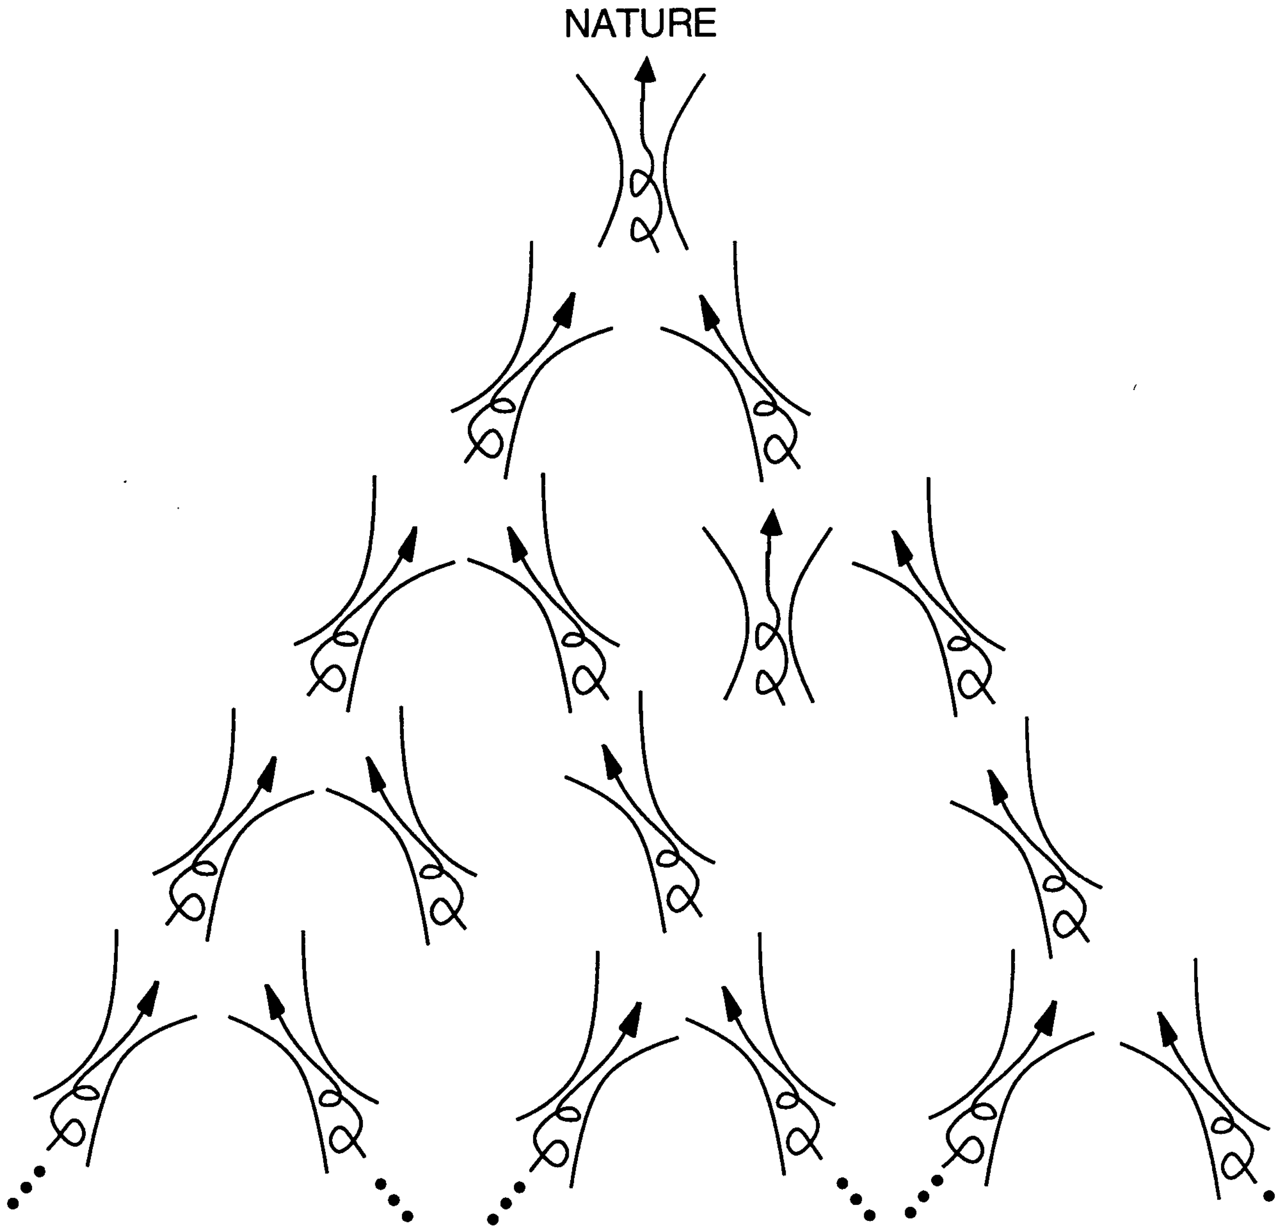

Fig.6 The reductionist nightmare.

Back to the Ant

Remember where wer were. Langton's Ant, governed only by a few very simple rules, displays surprisingly complex behaviour. First order, then chaos, then order emerging from the chaos — the famous highway.

The only rigorous way that is currently available to deduce this feature of Ant dynamics is to write down the ten thousand or so steps that lead to the 104-step cycle. Then an analysis of why the cycle repeats implies the formation of the highway structure. Now, Langton's rules are the 'Theory of Everything' for the universe that his ant inhabits — an anty-matter universe, so to speak. So here we have a feature whose existence can currently be demonstrated rigorously *only* by starting from the Ant's Theory of Everything, and following their consequences one step at a time.

Agreed, the Theory of Everything does indeed predict exactly what the ant will do.

So could an ant reasonably wear 'Ant Rules' on a tee-shirt and proclaim the end of ant physics? That's not so clear, especially if we ask slightly deeper questions. For example, computer experiments show that the ant *always* seems to end up building a highway, even if you scatter finitely many black squares around the grid before it starts — in any pattern whatsoever. But nobody has ever been able to prove that the ant always builds a highway. It certainly can't be done by reductionist rhetoric: there are infinitely many different initial conditions to consider. So here we have a high-level 'feature' of Ant dynamics that seems to be universal, but which cannot currently be deduced from the Theory of Everything for the system. Some general deductions are possible. For instance Kong and Cohen have proved — by exploiting certain features of the system — that the ant's trajectory is necessarily unbounded, that is, it escapes from any finite region. Here's a proof.

The Cohen-Kong Theorem: ant trajectories are unbounded

It is easy to check that the Theory of Everything for Langton's Ant is time-reversible. That is, the current pattern and heading determines the *past* uniquely as well as the future. Any bounded trajectory must eventually repeat the same pattern, position, and heading; and by reversibility such a trajectory must be periodic, repeating the same motions indefinitely. Thus every cell that is visited must be visited infinitely often.

The ant's motion is alternately horizontal and vertical, because its direction changes by ±90° at each step. Call a cell an H-cell if it is entered horizontally, and a V-cell if it is entered vertically. The H- and V-cells tile the grid like the black and white squares of a checkerboard.

Select a square M that is visited by the ant, and is as far up and to the right as possible, in the sense that the cells immediately above and to the right of it have never been visited. Suppose this is an H-cell. Then M must have been entered from the left and exited downward, and hence must have been white. But M now turns black, so that on the next visit the ant exits upwards, thereby visiting a square that has never been visited. A similar problem arises if M is a V-cell. This contradiction proves that no bounded trajectory exists.

Maybe somebody will solve the highway problem with a similar 'high-level' proof. However, it is striking that even for such a simple system as Langton's ant — where we *know* the Theory of Everything because we set it up — nobody can answer one simple question: starting from an arbitrary 'environment' of finitely many black cells, does the ant always build a highway? So here the Theory of Everything lacks explanatory power. It predicts everything but explains nothing.

In contrast, the Cohen-Kong theorem — the derivation of another feature, unboundedness — *explains* that feature. It is true that the Cohen-Kong Theorem is a *consequence* of the Theory of Everything. But it is making that consequence *explicit* that explains the unbounded trajectories. Appeals to the uniqueness of the consequences of the Theory of Everything are not, of themselves, explanations. They are closer to expressions of faith.

Only slightly more complex rules lead to systems such as Conway's game of Life. Conway proved that in Life there are configurations that form universal Turing machines — programmable computers. Turing proved that the long-term behaviour of a Turing machine is undecidable — for example, it is impossible to work out in advance whether or not the program will terminate. Translated into Life terms, that implies that the question 'does this configuration grow unboundedly?' is formally undecidable. So here's a case where we *know* the Theory of Everything, *and* we know a simple question that it is provably impossible to answer on the basis of that Theory. Since 'grow unboundedly' and its negation 'stay bounded' are high-level features, here is a case where we can prove the existence of a high-level feature that *cannot* be deduced from a real, known, very simple Theory of Everything.

So why do we think that a real Theory of Everything, for *our* universe, can in any meaningful sense be an Ultimate Answer?

VariAnts

The maths of Ants goes much further. For instance, Greg Turk, L.A.Bunimovich, and S.E.Troubetzkoy have investigated generalized ants, defined by a *rule-string*.

Suppose that instead of just black and white the squares have n colors, labelled 0, 1, 2, ..., n-1. The rule-string is a sequence of n symbols **0** or **1**. When the ant leaves a cell of color k it changes it to color k+1 (wrapping n = (n-1)+1 round to 0). It turns right if the kth symbol is **1**, and left if it is **0**. It moves one square on and repeats.

Langton's original rules are summed up in the rule-string **10**. Some rule-strings give trivial ant-dynamics — for example an ant with rule-string **1** (or even **111...1**) travels forever round a 2x2 square. But any rule-string that contains both a **0** and a **1** must lead to unbounded trajectories, by the Cohen-Kong idea.

Suppose for simplicity you start with a 'clean' grid — all cells in color 0. Ant **100** creates patterns that start out looking rather like those of Langton's ant — at first symmetric, then chaotic. After 150 million steps, however, it is still behaving chaotically. Does it ever build a highway? Nobody knows. Ant **110** does build a highway, and it takes only 150 steps to do so. Moreover, it needs a cycle of only 18 steps to create the highway, instead of the 104 used by Langton's Ant. Ant **1000** is relentlessly chaotic. Ant **1101** begins chaotically, but goes into highway-construction after 250,000 steps, using a cycle of length 388. Ant **1100** keeps building ever-more-complex patterns that, infinitely often, are bilaterally symmetric. (See **Fig.7**.) So it *can't* build any kind of highway in the usual sense.

I defy anyone to give a brief, simple classification of the behaviours of all of these generalized ants, or to predict from their rule-string just what their long-term behaviour will be — even if they all start on a clean grid.

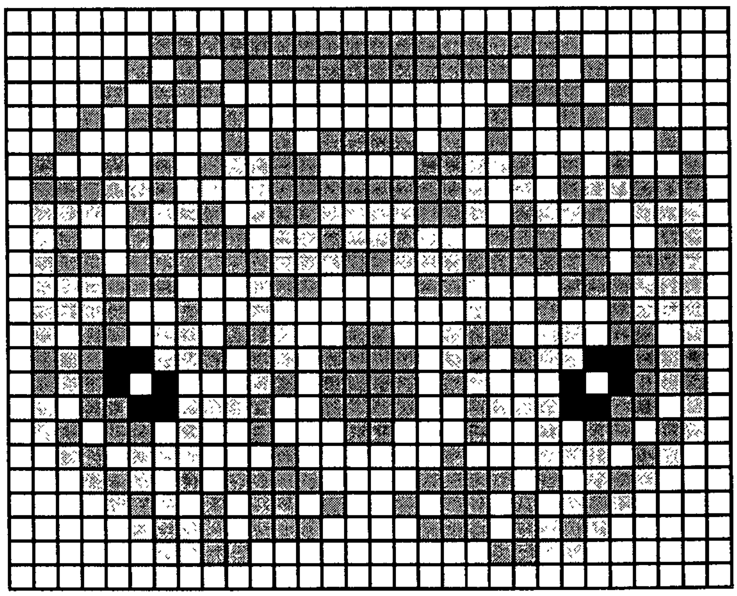

Fig.7 A symmetric pattern produced at step 16,464 of ant **1100**.

Features

Let us refer to the universal simplicities that can be extracted from reductionist theories — such as the Ant's highway, or the culinary preferences of goats — as 'features'. The use of features as a means to understanding nature underlies nearly all of science. The theories of physics are derived from simple 'toy' systems — a single electron in a potential well, two point masses orbiting under inverse square law gravity, and so on. The regularities uncovered by analysing such systems, or performing experiments on them, are codified as mathematical 'laws', which are then generalised to *all* systems. Then the consequences of those generalisations are investigated, but only for marginally more complex systems. For example Newton's inverse square law of gravitation was derived

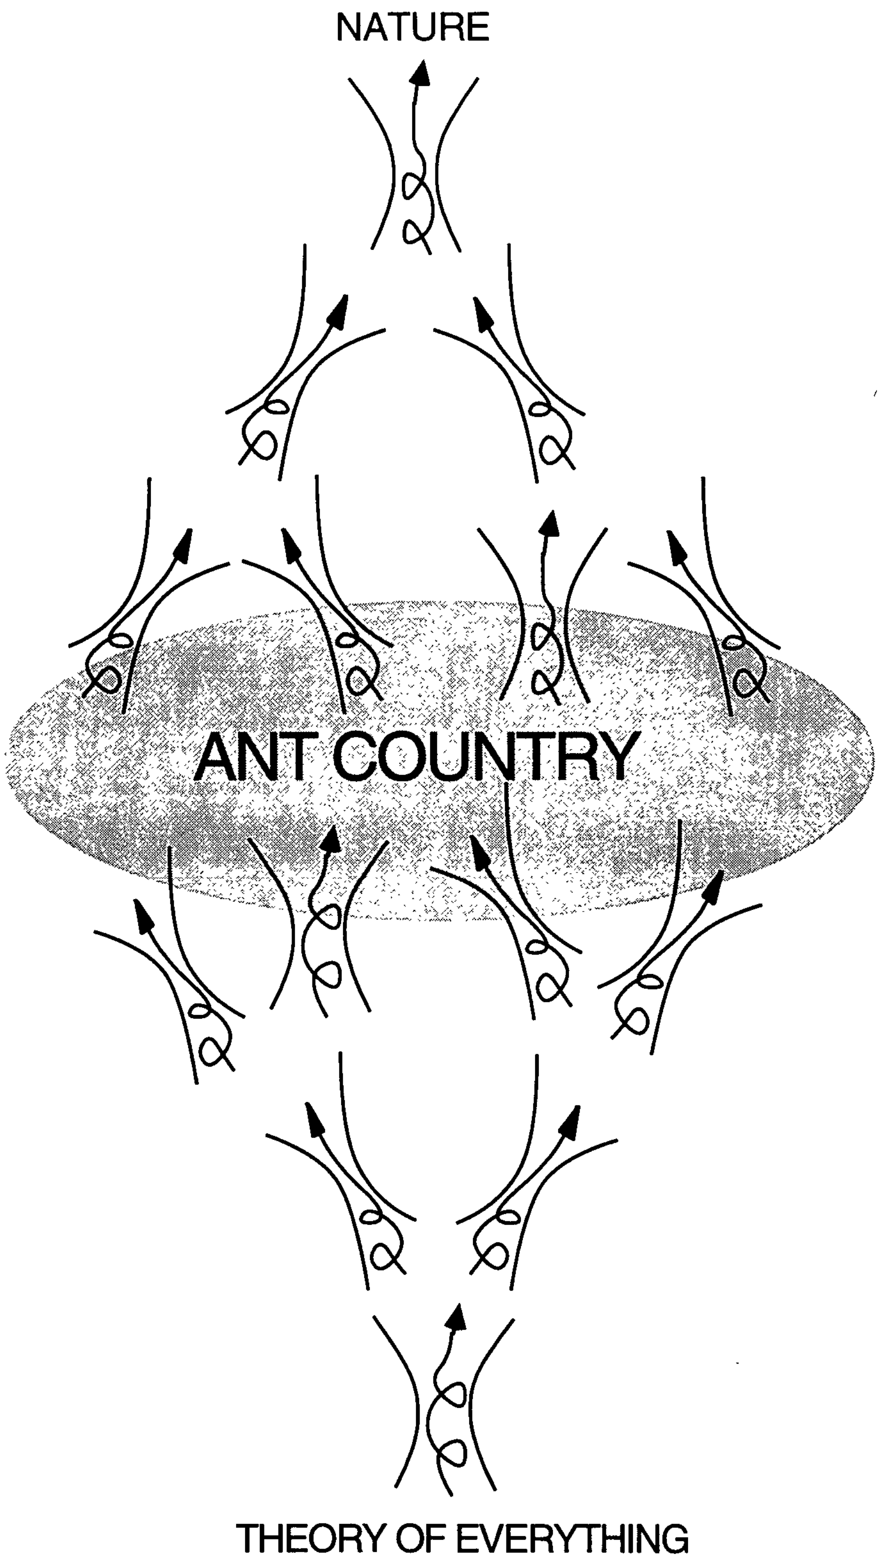

Fig.8 Ant Country.

from the elliptical orbit of the planet Mars, assumed to hold good for all particles in the universe, and applied — with great success — to systems such as Earth/Moon/Sun. We don't know that Newton's law really applies, in all its reductionist glory, to every particle in a galaxy. (It may well do, but we can't possibly *check* that.)

In Newton's derivation of elliptical orbits, planets are modelled as uniform spheres. However, real planets are not uniform spheres. The very universality claimed in Newton's law — *every* particle attracts *every* other particle — requires us to re-examine the derivation. Mars is not a uniform sphere, but a highly complex system of atoms — each one of which exerts its tiny contribution to the gravitational field of the whole. The reductionist rhetoric forces us to contemplate a very complex system of interacting forces.

Fortunately in this case it is not necessary actually to carry out the computation of such a gravitational field — which is impossible. The inverse square law gravitation has a useful feature, in our formal sense. By performing appropriate analytical estimates, based upon reasonable assumptions about the density of atoms in Mars's interior, it is in principle possible to prove that the combined field is a very close approximation to what it would be if the mass of Mars were concentrated at its centre. When we match the orbit of the real Mars to the predicted ellipse, we are exploting this feature: we are not pursuing the reductionist rhetoric. All of our scientific theories work with features: they are how we organize the structure of the theories in those terms. And the 'laws' that we prize are regularities in those features, not fundamental rules for the universe as a whole.

Ant Country

We take Langton's ant — and even more so, its more elusive generalizations — as a symbol for the gap between the top-down reductionism of the reductionist nightmare, and the bottom-up reductionism of the Theory of Everything. Bottom-up analysis proceeds from the putative Theory of Everything and ascends levels of description by deducing logical consequences of those laws in a hierarchical manner. Top-down analysis proceeds from nature and looks down mental funnels to see what lies inside. Thanks to the reductionist nightmare, the top and the bottom do not meet. Instead they both diverge into deductions too lengthy for the human mind to comprehend them. This 'nomansland' between top and bottom is *Ant Country* (**Fig.8**).

What of science's claim that its bottom-up rules explain the top-down behaviour of nature? How does the reductionist chain of logic traverse Ant Country? The answer is that it does not — it just claims to. The link between bottom and top is achieved through the intermediary of models. A uniform sphere (bottom) and a planet (top) are identified, for conceptual purposes: the planet is modelled by the sphere. The rules explain the sphere's gravitational field; this explanation is transferred — by analogy, not logic — to the planet. I'm not disputing that this process often *works*. But it breaks the alleged reductionist chain. It gives the illusion that the simplicity of the rules leads *directly* to the simplicity of elliptical orbits for real planets.

Instead, the explanatory story must enter the uncharted territory of Ant Country. And that is where the emergent phenomena live, it is where they come from. Ant Country is where complexity is created from nothing, where systems organize themselves into more complex systems without anything equally complex telling them how to do it. But very few scientists are even aware that Ant Country is there, let alone have any intention of exploring it.

REFERENCES

Elwyn R. Berlekamp, John H. Conway, and Richard K.Guy, *Winning Ways* (vol.2), Academic Press, New York 1982.

Jack Cohen and Ian Stewart, *The Collapse of Chaos*, Viking, New York 1994.

Richard Dawkins, *The Blind Watchmaker*, Longman, London 1986.

A.K.Dewdney, Computer Recreations, *Scientific American* (September 1989) 180-183; *Scientific American* (March 1990) 121.

David Gale, *Mathematical Entertainments, Mathematical Intelligencer* **15** No. 2 (1993) 54-55.

Jim Propp, Further Ant-ics, *Mathematical Intelligencer* **16** No. 1 (1994) 37-42.

Ian Stewart and Jack Cohen, *Figments of Reality*, Cambridge University Press, Cambridge 1997.

Steven Weinberg, *Dreams of a Final Theory: the Search for the Fundamental Laws of Nature*, Hutchinson 1993.