

True Complexity and its Associated Ontology

Contribution to John Wheeler's 90th Birthday celebration:

George F R Ellis,
Mathematics Department, University of Cape Town

Section 1: True complexity and the natures of existence

My concern in this article is true complexity and its relation to physics. This is to be distinguished from what is covered by statistical physics, catastrophe theory, study of sand piles, the reaction diffusion equation, cellular automata such as The Game of Life, and chaos theory. Examples of truly complex systems are molecular biology, animal and human brains, language and symbolic systems, individual human behaviour, social and economic systems, digital computer systems, and the biosphere. This complexity is made possible by the existence of molecular structures that allow complex bio-molecules such as RNA, DNA, and proteins with their folding properties and lock-and-key recognition mechanisms, in turn underlying membranes, cells (including neurons), and indeed the entire bodily fabric and nervous system.

True complexity involves vast quantities of stored information and hierarchically organised structures that process information in a purposeful manner, particularly through implementation of goal-seeking feedback loops. Through this structure they appear purposeful in their behaviour ('teleonomic'). This is what we must look at when we start to extend physical thought to the boundaries, and particularly when we try to draw philosophical conclusions – for example, as regards the nature of existence - from our understanding of the way physics underlies reality. Given this complex structuring, one can ask, What is real?, that is, What actually exists?, and What kinds of causality can occur in these structures?

This article aims to look at these issues. It contains the following further sections: Section 2: The nature of true complexity; Section 3: The natures of existence; Section 4: The nature of causality; Section 5: The relation to fundamental physics, including some comments on the relation to ultimate reality.

2: The nature of true complexity

In broad outline, emergence of complex systems has the following features:

1. Emergence of complexity occurs in terms of (a) function (simple structures causally underlying functioning of more complex structures), (b) development (a single initial cell growing to a complex interlocking set of 10^{13} cells), and (c) evolution (a universe region containing no complex systems evolving to one containing billions of them), each occurring with very different timescales.

Natural selection (see e.g. Campbell 1991) is seen as the mechanism that allows this all to come into being through an evolutionary process.

2. Complex systems are characterised by (a) hierarchical structures delineating both complexity and causality with (b) different levels of order and descriptive languages plus (c) a relational hierarchy at each level of the structural hierarchy

This is summarised in the hierarchy of structure/causation relating to human beings (Figure 1).

Sociology/Economics/Politics
Psychology
Physiology
Cell biology
Biochemistry
Chemistry
Physics
Particle physics

Figure 1: A hierarchy of structure and causation. Each lower level underlies what happens at each higher level, in terms of physical causation. For a much more detailed exploration of such hierarchies, see the web pages (Ellis 2002).

As expressed by Campbell (1991, pp.2-3), "With each upward step in the hierarchy of biological order, novel properties emerge that were not present at the simpler levels of organisation. These emergent properties arise from interactions between the components.... Unique properties of organized matter arise from how the parts are arranged and interact .. [consequently] we cannot fully explain a higher level of organisation by breaking it down to its parts". One can't even describe the higher levels in terms of lower level language.

Furthermore, one can't comprehend the vast variety of objects at each higher level without using a characterisation structured in hierarchical terms, e.g.

"animal - mammal – domestic animal – dog – guard dog - Doberman - Fred",

"machine – transport vehicle - automobile – sedan – Toyota – CA687-455".

The hierarchical characterisation used may be based on (i) appearance, (ii) structure, (iii) function, (iv) geographic location and/or history (e.g. evolutionary history), or (v) an assigned labelling (e.g. alphabetic or numeric, each themselves hierarchical in nature). Note that in the end these categorisations go from the generic to the individual/specific.

3. These hierarchical structures are modular – made up by structural combinations of simpler (lower level) components with their own state variables, incorporating encapsulation and inheritance, enabling re-use and modification

In general many lower level states correspond to a single higher level state, because a higher level state description is arrived at by averaging over lower level states and throwing away a vast amount of lower level information ('coarse-graining'). The number of lower level states corresponding to a single higher level state determines the entropy of that state. This is lower level information hidden in the higher level view.

4. Complex emergence is enabled by (a) bottom up and (b) top down action, the latter occurring by coordinating lower level actions according to the system structure and boundary conditions.

Higher level structures and boundary conditions can structure lower level interactions in a coordinated way; this is top-down action in the hierarchy. This affects the nature of causality significantly, particularly because inter-level feedback loops become possible (Figure 2). Reliable higher level behavioural laws occur if the variety of lower level states corresponding to a particular higher level initial state all lead to the same higher level final state, thus enabling same-level action.




Figure 2a: Bottom-up only.




Figure 2b: Bottom-up and Top-down.

Figure 2: Bottom-up and Top-down action. The fundamental importance of top-down action is that it changes the causal relation between upper and lower levels in the hierarchy of structure and organisation, cf. the difference between Fig 2a and Fig 2b.

Causality in coherent complex systems has all these dimensions (bottom up, same level, top down): there are explanations of each kind that are simultaneously applicable.

5. Living systems involve purposeful use of information to control physical functions in accord with higher level goals. They are structured as (a) feedback control systems that (b) can learn by (c) capturing, storing, recalling, and analysing information which is used to set the system goals; this involves (d) pattern recognition and (e) utilisation of simplified predictive models.

It is these capacities that make the difference between complicated and complex systems. They enable strongly emergent phenomena such as the existence of the rules of chess (as well as the resulting strategies of chess players) and of money and exchange rates. There is no implication here as to how the information is stored (it might be encoded in particular system energy levels, sequences of building block molecules, synaptic

connection patterns for example). Also there is no implication here that the amount of useful information is described by the Shannon formulae.

This adds up to **Organised Complexity** (Sellars 1932, Simon 1962, Churchman 1968, Simon 1982, Flood and Carson 1990, Bar-Yam 2000). Here we see non-material features such as information and goals having causal effects in the material world of forces and particles, which means they have an ontological reality. We now examine aspects of these various features in more detail.

2.1 Complexity and Structure

The essential point of systems theory is that the value added by the system must come from the relationships between the parts, not from the parts per se (Emery 1972, von Bertalanffy 1973). True complexity, with the emergence of higher levels of order and meaning, occurs in *modular hierarchical structures*, because these form the only viable ways of building up and utilising real complexity. The principles of hierarchy and modularity have been investigated usefully in the context of computing, and particularly in the discussion of *object-oriented programming*, see (Booch, 1994), and it is helpful to see how these principles are embodied in physical and biological structures.

A: Modularity (Booch 1994, pp.12-13, 54-59). The technique of mastering complexity in computer systems and in life is divide and rule – decompose the problem into smaller and smaller parts, each of which we may then refine independently (Booch, p.16). By organising the problem into smaller parts, we break the informational bottleneck on the amount of information that has to be received, processed, and remembered at each step; and this also allows specialisation of operation. This implies creation of a set of specialised modules to handle the smaller problems that together comprise the whole. In building complex systems from simple ones, or improving an already complex system, you can re-use the same modular components in new combinations, or substitute new more efficient components, with the same functionality, for old ones. Thus we can benefit from a library of tried and trusted components. However the issue then arises as to how we structure the relationships between the modules, and what functional capacity we give them. Complex structures are made of modular units with *abstraction*, *encapsulation*, and *inheritance*; this enables the modification of modules and re-use for other purposes.

1: Abstraction and Labelling (Booch 1994, pp. 20, 41-48): Unable to master the entirety of a complex object, we choose to ignore its inessential details, dealing instead with a generalised idealised model of the object. *An abstraction* denotes the essential characteristics of an object that distinguishes it from all other kinds of objects. An abstraction focuses on the outside view of the object, and so serves to separate its essential behaviour from its implementation; it emphasises some of the system's details or properties, while suppressing others. A key feature is that *compound objects can be named and treated as units*. This leads to the power of abstract symbolism and symbolic computation.

2: Encapsulation and Information Hiding (Booch 1994, pp.49-54): Consumers of services only specify what is to be done, leaving it to the object to decide how to do it: “No part of any complex system should depend on the internal details of any other part”. *Encapsulation* is when the internal workings are hidden from the outside, so its procedures can be treated as black-box abstractions. To embody this, each class of object must have two parts: an *interface* (its outside view, encompassing an abstraction of the common behaviour of all instances of the class of objects) and an *implementation* (the internal representations and mechanisms that achieve the desired behaviour). Efficiency and usability introduce the aim of reducing the number of variables and names that are visible at the interface. This involves *information hiding*, corresponding to coarse-graining in physics; the accompanying loss of detailed information is the essential source of entropy.

3: Inheritance (Booch, p. 59-62) is the most important feature of a classification hierarchy: it allows an object class – such as a set of modules - to inherit all the properties of its superclass, and to add further properties to them (it is a ‘is a’ hierarchy). This allows similarities to be described in one central place and then applied to all the objects in the class and in subclasses. It makes explicit the nature of the hierarchy of objects and classes in a system, and implements generalisation/ specialisation of features (the superclass represents generalised abstractions, and subclasses represent specializations in which variables and behaviours are added, modified, or even hidden). It enables us to understand something as a modification of something already familiar, and saves unnecessary repetition of descriptions or properties.

B: Hierarchy (Booch 1994, pp. 59-65). Hierarchical structuring is a particularly helpful way of organising the relationship between the parts, because it enables building up of higher level abstractions and permits relating them by inheritance. A hierarchy represents a decomposition of the problem into constituent parts and processes to handle those constituent parts, each requiring less data and processing, and more restricted operations than the problem as a whole. The success of hierarchical structuring depends on

- (a) implementation of modules which handle these lower-level processes,
- (b) integration of these modules into a higher-level structure.

The basic features of how hierarchies handle complexity are as follows:

“Frequently, complexity takes the form of a *hierarchy*, whereby a complex system is composed of inter-related subsystems that have in turn their own subsystems, and so on, until some lowest level of component is reached” (Courtois 1985) “The fact that many complex systems have a nearly decomposable hierarchic structure is a major facilitating factor enabling us to understand, describe, and even ‘see’ such systems and their parts” (Simon 1982). Not only are complex systems hierarchic, but the levels of this hierarchy represent *different levels of abstraction*, each built upon the other, and each understandable by itself (and each characterised by a different phenomenology). This is the phenomenon of *emergent order*. All parts at the same level of abstraction interact in a well-defined way (which is why they have a reality at their own level, each represented in a *different language* describing and characterising the causal patterns at work at that level). “We find separate parts that act as independent agents, each of which exhibit some fairly complex behaviour, and each of which contributes to many higher level functions.

Only through the mutual co-operation of meaningful collections of these agents do we see the higher-level functionality of an organism. This is *emergent behaviour* – the behaviour of the whole is greater than the sum of its parts, and cannot even be described in terms of the language that applies to the parts. Intra-component linkages are generally stronger than inter-component linkages. This fact has the effect of separating the *high-frequency dynamics* of the components – involving their internal structure – from the low-frequency dynamics –involving interactions amongst components” (Simon 1982) (and this is why we can sensibly identify the components). In a hierarchy, through encapsulation, objects at one level of abstraction are shielded from implementation details of lower levels of abstraction.

2.2 Bottom-up and Top-down action.

A: Bottom-up action. A fundamental feature of the structural hierarchy in the physical world is *bottom-up action*: what happens at each higher level is based on causal functioning at the level below, hence what happens at the highest level is based on what happens at the bottom-most level. This is the profound basis for reductionist world views. The successive levels of order entail chemistry being based on physics, material science on both physics and chemistry, geology on material science, and so on.

B: Top-down action. The complementary feature to bottom-up action is *top-down action*, which occurs when the higher levels of the hierarchy direct what happens at the lower levels in a coordinated way (Campbell 1974, Peacocke 1993). For example depressing a light switch leads to numerous electrons systematically flowing in specific wires leading from a power source to a light bulb, and consequent illumination of a room.

In general many lower level states correspond to a single higher level state, because a higher level state description is arrived at by averaging over lower level states and throwing away a vast amount of lower level information (coarse graining). Hence, specification of a higher level state determines a family of lower level states, any one of which may be implemented to obtain the higher level state (the light switch being on, for example, corresponds to many billions of alternative detailed electron configurations). The specification of structure may be loose (attainable in a very large number of ways, e.g. the state of a gas) or tight (defining a very precise structure, e.g. particular electrons flows in the wiring of a VLSI chip in a computer). In the latter case both description and implementation require far more information than in the former.

The dynamics acts on a lower level state L_i to produce a new lower level states L'_i in a way that depends on the boundary conditions and structure of the system. Thus specifying the upper state H_i (for example by pressing a computer key) results in some lower level state L_i that realises H_i and then consequent lower level dynamic change. The lower level action would be different if the higher level state were different. It is both convenient and causally illuminating to call this *top-down action*, and to explicitly represent it as an aspect of physical causation; this emphasizes how the lower level changes are constrained and guided by structures that are only meaningful in terms of a higher level description.

The question now is whether all the states L_i corresponding to a single initial upper state H_1 produce new lower level states that correspond to the same higher level

state H'_1 . Two major cases arise. Different lower level realisations of the same higher level initial state may result, through microphysical action, in different higher level final states (Figure 3). Here there is no coherent higher level action generated by the lower level actions. On the other hand, top-down action generating same-level action occurs when each lower level state corresponding to a specific higher level state results in the same final upper level state, so that every lower level implementation of the initial higher level state gives the same higher level outcome (Figure 4).




Figure 3: Low level action that does not result in coordinated high level action.




Figure 4: Same-level action results from coordinated low-level action that results in coordinated high level action.

In this case, consistent behaviour occurs at the higher level, regarded as a causal system in its own right - there is now effective higher level autonomy of *same level action*, which we can consider in its own right independent of the lower levels:

$$H_1 \longrightarrow H'_1$$

This higher level action is effective by coordinating actions at the lower levels, resulting in coherent higher level action. Whether this happens or not may depend on the particular coarse-graining (i.e. higher level description) chosen.

Multiple top-down action enables various higher levels to co-ordinate action at lower levels in a complex system in a coherent way, and so gives them their causal effectiveness. It is prevalent in the real physical world and in biology, because no real

physical or biological system is isolated. Boundary effects as well as structural relations effect top-down action. I will illustrate this with a series of examples.

1: Gas in a cylinder with piston. The cylinder walls together with the piston position determine the gas pressure and temperature. Both are macro-concepts which make no sense at the micro level.

2: Interaction potentials. Potentials in the Schroedinger equation, or in the action for the system, represent the summed effects of other particles and forces, and hence are the way that the nature of both simple and complex structures can be implemented (from a particle in a box to a computer or a set of brain connections). These potentials describe the summed interactions between micro-states, enabling internal top-down effects. Additionally one may have external potentials imposed in the chosen representation, representing top-down effects from the environment on the system.

3: Nucleosynthesis in the early universe. The rates of nuclear interactions depend on the density and temperature of the interaction medium. The microscopic reactions that take place in the early universe, and hence the elements produced, thus depend on the rate of expansion of the universe (determined by the Friedmann equation).

4: Quantum measurement. *Top-down action* occurs in the quantum measurement process - the collapse of the wave function to an eigenstate of a chosen measurement system (Isham 1997, Penrose 1989). The experimenter chooses the details of the measurement apparatus - e.g. aligning the axes of polarisation measurement equipment - and that decides what set of microstates can result from a measurement process, and so crucially influences the possible micro-state outcomes of the interactions that happen. Thus the quantum measurement process is partially a top-down action controlled by the observer, determining what set of eigenstates are available to the system during the measurement process. The choice of Hilbert space and the associated operators and functions is made to reflect the experimenter's choice of measurement process and apparatus, thus reflecting this top-down action.

5: The arrow of time. *Top-down action* occurs in the determination of the arrow of time. One cannot tell how a macrosystem will behave in the future on the basis of the laws of physics and the properties of the particles that make up the system alone, because time-reversible micro-physics allows two solutions - one the time reverse of the other - but the macro-system allows only one of those solutions (Davies 1974, Zeh 1992). The prohibition of one of the allowed microsolutions is mathematically put in by hand to correspond to the real physical situation, where only entropy-increasing solutions in one direction of time occur at the macrolevel; this does not follow from the microphysics. For example, Boltzmann brilliantly proved the H-theorem (increase of macroscopic entropy in a gas on the basis of microscopic molecular collisions); but Loschmidt pointed out that this theorem works equally well with both directions of the arrow of time. Physically, the only known solution to this arrow of time problem seems to be that there is top-down action by the universe as a whole, perhaps expressed as boundary conditions at beginning of space-time (Penrose, 1989), that allows the one solution and disallows the other. This is related to the quantum measurement issue raised above: collapse of the wave function takes place with a preferred direction of time, which is not determined if we look at the microlevel of the system alone.

6: Evolution. Top-down action is central to two main themes of molecular biology: First, the *development of DNA codings* (the particular sequence of bases in the

DNA) occurs through an evolutionary process which results in adaptation of an organism to its ecological niche (Campbell 1991). This is a classical case of top-down action from the environment to detailed biological microstructure - through the process of adaptation, the environment (along with other causal factors) fixes the specific DNA coding. There is no way you could ever predict this coding on the basis of biochemistry or microphysics alone (Campbell 1974).

7: Biological development. Second, the *Reading of DNA codings*. A second central theme of molecular biology is the reading of DNA by an organism in the developmental process (Gilbert 1991, Wolpert, 1998). This is not a mechanistic process, but is context dependent all the way down, with what happens before having everything to do with what happens next. The central process of developmental biology, whereby positional information determines which genes get switched on and which do not in each cell, so determining their developmental fate, is a top-down process from the developing organism to the cell, based on the existence of gradients of positional indicators in the body. Without this feature organism development in a structured way would not be possible. Thus the functioning of the crucial cellular mechanism determining the type of each cell is controlled in an explicitly top-down way. At a more macro level, recent research on genes and various hereditary diseases shows that existence of the gene for such diseases in the organism is not a sufficient cause for the disease to in fact occur: outcomes depend on the nature of the gene and on the rest of the genome and on the environment. The macro situations determine what happens, not specific micro features by themselves, which do work mechanistically but in a context of larger meaning that largely determines the outcome. The macro environment includes the result of conscious decisions (the patient will or will not seek medical treatment for a hereditary condition, for example), so these too are a significant causal factor.

8: Mind on body. *Top-down action* occurs from the mind to the body and thence into the physical world. The brain controls the functioning of the parts of the body through a hierarchically structured feedback control system, which incorporates the idea of decentralised control to spread the computational and communication load and increase local response capacity (Beer 1981). It is a highly specific system in that dedicated communication links convey information from specific areas of the brain to specific areas of the body, enabling brain impulses to activate specific muscles (by coordinated control of electrons in myosin filaments in the bundles of myofibrils that constitute skeletal muscles), in order to carry out consciously formulated intentions. Through this process there is top-down action by the mind on the body, and indeed on the mind itself, both in the short term (immediate causation through the structural relations embodied in the brain and body) and in the long term (structural determination through imposition of repetitive patterns). An example of the latter is how repeated stimulation of the same muscles or neurons encourages growth of those muscles and neurons. This is the underlying basis of both athletic training and of learning by rote. Additionally, an area of importance that is only now beginning to be investigated by Western medicine is the effect of the mind on health (Moyers 1993), for example through interaction with the immune system (Sternberg 2000).

9. Mind on the world. When a human being has a plan in mind (say a proposal for a bridge being built) and this is implemented, then enormous numbers of micro-particles (comprising the protons, neutrons, and electrons in the sand, concrete, bricks,

etc that become the bridge) are moved around as a consequence of this plan and in conformity with it. Thus in the real world, the detailed micro-configurations of many objects (which electrons and protons go where) is to a major degree determined by the macro-plans that humans have for what will happen, and the way they implement them. Some specific examples of top down action involving goal-choice are:

(i) *Chess rules* are socially embodied and are causally effective. Imagine a computer or alien analysing a large set of chess games and deducing the rules of chess (i.e. what physical moves of the pieces are allowed and what not). It would know these are inviolable rules governing these moves but have no concept of their origins, i.e. whether they were implied by modification of Newton's laws, some potential fields that constrain the motion of the chess pieces, or a social agreement that restricts their movement and can be embodied in computer algorithms. Note that the chess rules exist independent of any particular mind, and indeed may survive the demise of the society that developed them (they can for example be written in a book that becomes available to other societies).

(ii) *The Internet*. This embodies local action in response to information requests, causing electrons to flow in meaningful patterns in a computer's silicon chips and memory, mirroring patterns thousands of miles away, when one reads web pages. This is a structured influence at a distance due to channelled causal propagation and resulting local physical action.

(iii) *Money and associated exchange rates*; the money is a physical embodiment of the economic order in a society, while the exchange rates are socially embodied but are also embodied for example in ink on newspaper pages, and in computer programmes utilised by banks.

(iv) *Global Warming*. The effect of human actions on the earth's atmosphere, through the combined effect of human causation moving very large numbers of micro-particles (specifically, CFC's) around, thereby affecting the global climate. The macro-processes at the planetary level cannot be understood without explicitly accounting for human activity (Schellnhuber 2000).

(v) *Hiroshima*. The dropping of the nuclear bomb at Hiroshima was a dramatic macro-event realised through numerous micro-events (fissions of uranium nuclei) occurring through a human-based process of planning and implementation of those plans.




Figure 5. Branching Hierarchy of causal relations. The hierarchy of physical relations (Figure 1) extended to a branching hierarchy of causal relations. The left hand side involves only (unconscious) natural systems; the right hand side involves conscious choices, which are causally effective. In particular, the highest level of intention (ethics) is causally effective

Because of this, the structural hierarchy, now interpreted as a causal hierarchy, bifurcates (see Figure 5). The left hand side, representing causation in the natural world, does not involve goal choices: all proceeds mechanically. The right hand side, representing causation involving humans, is to do with choice of goals that lead to actions. Ethics is the high-level subject dealing with the choice of appropriate goals. Because this determines the lower level goals chosen, and thence the resulting actions, ethics is causally effective in the real physical world. This is of course obvious as it follows from the causation chain, but to make the point absolutely clear: a prison may have present in its premises the physical apparatus of an execution chamber, or maybe not; whether this is so or not depends on the ethics of the country in which the prison is situated.

2.3 Information and Goal-seeking

A key feature is the use of current and past information to set goals that are then implemented in feedback control systems. This is the context in which stored information is the effective core of complex history-dependent behaviour.

A: Feedback control . The central feature of organised action is *feedback control*, whereby setting of goals results in specific actions taking place that aim to achieve those goals (Ashby 1958, Beer 1966, 1972). A comparator compares the system state with the goals, and sends an error message to the system controller if needed to correct the state by making it a better approximation to the goals (Figure 6).

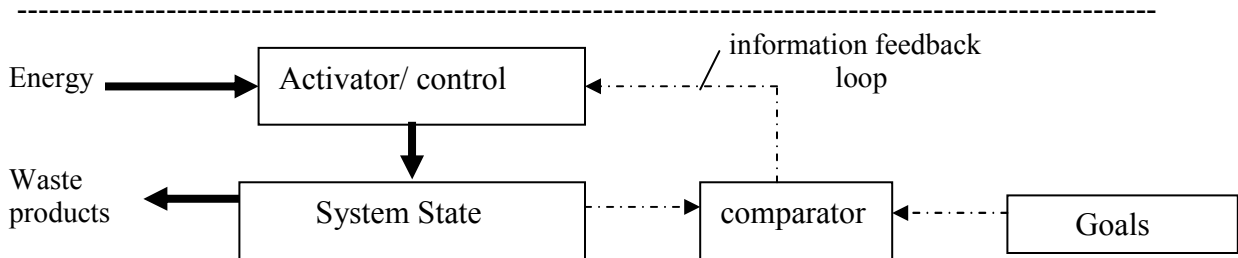


Figure 6: The basic feedback control process. The second law of thermodynamics requires energy input and heat output in active processes, which must occur then in an open system.

The linkages to the comparator and thence to the controller are *information linkages* rather than power and/or material linkages like that from the activator to the system (the information flow will use a little power but only that required to get the message to where it is needed). Examples are controlling the heat of a shower, the direction of an automobile, or the speed of an engine. Thus it is here that the *key role of information* is seen: *it is the basis of goal choice in living systems* (and artifacts that embody feedback control). The crucial issue now is what determines the goals: where do they come from? Two major cases need to be distinguished.

Homeostasis: In-built goals. There are numerous systems in all living cells, plants, and animals that automatically (i.e. without conscious guidance) maintain homeostasis - they keep the structures in equilibrium through multiple feedback systems that fight intruders (the immune system), and control energy and material flows, breathing and the function of the heart, body temperature and pressure, etc. They are effected through numerous enzymes, anti-bodies, regulatory circuits of all kinds (for example those that maintain body temperature and blood pressure) (Milsum 1966). The inbuilt goals that guide these activities are implicit rather than explicit, for example the temperature of the human body is maintained at 98.4F with great accuracy without that figure being explicitly pre-set in some control apparatus, but certainly these goals are identifiable and very efficiently attained. They have developed in the course of time through the processes of evolution, and so are historically determined in particular environmental context, and are unaffected by individual history. In manufactured artifacts, the goal may be explicitly stated and controllable (e.g. the temperature setting of a thermostat). *Not only are the feedback control systems themselves emergent systems, but also the implied goals are emergent properties that guide numerous physical, chemical, and biochemical interactions in a teleological way.* They represent distilled information about the behaviour of the environment in relation to the needs of life, and so they represent implicit information processing by the organism. At the higher levels they include the instinctive behaviour of animals. These feedback control loops are hierarchically structured with maximum decentralisation of control from the higher to the lower levels, as is required both in order to handle requisite variety (Ashby 1958) and the associated information loads (Beer 1981), and for maximal local efficiency (ability to respond to local conditions).

Goal seeking: Chosen goals. However at higher levels in animals, important new features come into play: there are now explicit behavioural goals, that are either learnt or are consciously chosen. *It is in the choice of these goals that explicit information processing comes into play.* Information arrives from the senses and is sorted and either discarded or stored in long term and short term memory. Conscious and unconscious processing of this information sets up the goal hierarchy (with ethics the topmost level) which then controls purposeful action.

B Information origin and use: Responsive behaviour depends on purposeful use of information: capture, storage, transmission, recall, and assessment to control physical functions in accord with higher level goals. The computations are based on stored variables and structured information flows, so hidden internal variables affect external behaviour. Current information is filtered against a relevance pattern, the irrelevant being discarded, the moderately significant being averaged over and stored in compressed form, the important being selectively amplified and used in association with current expectations to assess and revise immediate goals. The relevance pattern is determined by basic emotional responses such as those delineated by Panksepp (1998), which provide the evaluation function used in a process of neural Darwinianism (Edelman 1990) that determines the specifics of neural connections in the brain. In this way emotional responses underlie the development of rationality. Expectations are based on causal models based on past experience ('frames', e.g. how to behave in a restaurant), constantly revised on the basis of newer experience and information.

Thus feedback control systems based on sophisticated interpretation of present and past data enable purposeful (teleological) behaviour. *Memory* allows both the long term past and the immediate environmental context to be taken into account in choosing goals, providing historical information used to shape these goals in conjunction with present data. Long term memory allows a non-local (in time) kind of causation that enables present and future behaviour to be based on interpretations of long past events (e.g. remembering that an individual let one down in important ways years ago). *Learning* allows particular responses to develop into an automatic skill, in particular allowing some responses to become inbuilt and so able to be rapidly deployed (e.g. driving a car, many sports moves, and so on).

Symbolic Representation: At the highest level, the process of analysis and understanding is driven by the power of symbolic abstraction (Deacon 1997), codified into language embodying both syntax and semantics. This underpins other social creations such as specialised roles in society and the monetary system on the one hand, and higher level abstractions such as mathematics, physical models, and philosophy on the other – all encoded in symbolic systems.

Information guides all this, and is manifestly real in that it has a commercial value that underlies development of a major part of the international economy (the information technology sector). The meta-question of how context influences behaviour is guided and constrained by a system of ethics based on an overall world-view associated with meaning. This will also be encoded in language and symbols.

These are all strongly emergent phenomena that are causally effective. They exist as non-material effective entities, created and maintained through social interaction and teaching, and are codified in books and sometimes in legislation. Thus while they may be represented and understood in individual brains, their existence is not contained in any individual brain and they certainly are not equivalent to brain states. Rather the latter serve as just one of many possible forms of embodiment of these features.

2.5 Mathematical and Physical description

How to model all this? There are two approaches to quantitative modelling of hierarchical systems and emergent properties that may be useful. On the one hand, network mathematics and related network thermodynamics tackle the problem directly, see e.g. Peacocke (1989) and Holland (1998), and on the other some of these issues are tackled in the studies of neural networks and genetic algorithms (Carpenter and Grossberg 1991, Bishop 2000) and of control systems (e.g. Milsum 1966). What are needed are computer hierarchical models plus heuristic understanding of interplay of components, together with mathematical models of specific sub-systems and networks and physical models of molecular structure and interactions (needing mathematical models of 3-dimensional structure) that allow this to come into existence in complex systems.

Section 3: The natures of existence

In this section I propose a holistic view of ontology, building on the previous proposals by Popper and Eccles (1977) and Penrose (1997). I clearly distinguish between ontology (existence) and epistemology (what we can know about what exists). They should not be confused: Whatever exists may or may not interact with our senses and measuring instruments in such a way as to clearly demonstrate its existence to us.

3.1 A Holistic view of Ontology.

I take as given the reality of the everyday world – tables and chairs, and the people who perceive them – and then assign a reality additionally to each kind of entity that can have a demonstrable causal effect on that every day reality. The problem then is to characterise the various kinds of independent reality which may exist in this sense. Taking into account the causal efficacy of all the entities discussed above, I suggest as a possible completion of the proposals by Popper and Eccles and Penrose that the four worlds indicated in Figure 7 are ontologically real. These are not different causal levels within the same kind of existence, rather they are quite different kinds of existence, but related to each other through causal links. The challenge is to show firstly that each is indeed ontologically real, and secondly that each is sufficiently and clearly different from the others that it should be considered as separate from them. I now discuss them in turn.

-
- **World 1: Matter and Forces**
 - **World 2: Consciousness**
 - **World 3: Physical and biological possibilities**
 - **World 4: Mathematical reality**

Figure 7: The different kinds of reality implied by causal relationships can be characterised in terms of four Worlds, each representing a different kind of existence.

3.2 Matter and Forces

World 1: *The Physical World of Energy and Matter*, hierarchically structured to form lower and higher causal levels whose entities are all ontologically real.

This is the basic world of matter and interactions between matter, based at the micro level on elementary particles and fundamental forces, and providing the ground of physical existence. It comprises three major parts:

World 1a: Inanimate objects (both naturally occurring and manufactured).

World 1b: Living things, apart from humans (amoeba, plants, insects, animals, etc).

World 1c: Human beings, with the unique property of being self-conscious.

All these objects are made of the same physical stuff, but the structure and behaviour of inanimate and living things (described respectively by physics and inorganic chemistry, and by biochemistry and biology) are so different that they require separate recognition, particularly when self-consciousness and purposive activity (described by psychology

and sociology) occurs. The hierarchical structure in matter is a real physical structuration, and is additional to the physical constituents that make up the system themselves. It provides the basis for higher levels of order and phenomenology, and hence of ontology.

There is ontological reality at each level of the hierarchy. Thus we explicitly recognise as being real, quarks, electrons, neutrinos, rocks, tables, chairs, apples, humans, the World, stars, galaxies, and so on. The fact that each is comprised of lower level entities does not undermine its status as existing in its own right (Sellars 1932). We can attain and confirm high representational accuracy and predictive ability for quantities and relations at higher levels, independent of our level of knowledge of interactions at lower levels, giving well-validated and reliable descriptions at higher levels accurately describing the various levels of emergent non-reducible properties and meanings. An example is digital computers, with their hierarchical logical structure expressed in a hierarchy of computer languages that underlie the top-level user programmes. The computer has a reality of existence at each level that enables one to meaningfully deal with it as an entity at that level (Tannenbaum 1990). The user does not need to know machine code, and indeed the top-level behaviour is independent of which particular hardware and software underlie it at the machine level. Another example is that a motor mechanic does not have to study particle physics in order to ply her trade.

3.3 Consciousness

World 2: *The World of Individual and Communal Consciousness:* ideas, emotions, and social constructions. This again is ontologically real (it is clear that these all exist), and causally effective.

This world of human consciousness can be regarded as comprising three major parts:

World 2a: Human Information, Thoughts, Theories, and Ideas.

World 2b: Human Goals, Intentions, Sensations, and Emotions.

World 2c: Explicit Social Constructions

These worlds are different from the world of material things, and are realised through the human mind and society. They are not brain states, although they can be represented as such, for they do not reside exclusively in any particular individual mind. They are not identical to each other: World 2a is the world of rationality, World 2b is the world of intention and emotion, and so comprehends non-propositional knowing, while World 2c is the world of consciously constructed social legislation and convention. Although each individually and socially constructed in a complex interaction between culture and learning, these are indeed each capable of causally changing what happens in the physical world, and each has an effect on the others. In more detail:

World 2a: The world of Human Information, Thoughts, Theories, and Ideas. This world of rationality is hierarchically structured, with many different components. It includes words, sentences, paragraphs, analogies, metaphors, hypotheses, theories, and indeed the entire bodies of science and literature, and refers both to abstract entities and to specific objects and events. It is necessarily socially constructed on the basis of varying degrees of experimental and observational interaction with World 1, which it then represents with varying degrees of success. World 2a is represented by symbols, particularly language

and mathematics, which are arbitrarily assigned and which can themselves be represented in various ways (sound, on paper, on computer screens, in digital coding, etc).

Thus each concept can be expressed in many different ways, and is an entity in its own right independent of which particular way it is coded or expressed. These concepts sometimes give a good correspondence to entities in the other worlds, but the claim of ontological reality of entities existing in World 2a makes no claim that the objects or concepts they refer to are real. Thus this world equally contains concepts of rabbits and fairies, galaxies and UFOs, science and magic, electrons and aether, unicorns and apples; the point being that all of these certainly exist *as concepts*. That statement is neutral about whether these concepts correspond to objects or entities that exist in the real universe (specifically, whether there is or is not some corresponding entity in World 1) or whether the theories in this world are correct (that is, whether they give a good representation of World 1 or not).

All the ideas and theories in this world are ontologically real in that they are able to cause events and patterns of structures in the physical world. Firstly, they may all occur as descriptive entries in an encyclopaedia or dictionary. Thus each idea has causal efficacy as shown by existence of the resulting specific patterns of marks on paper (these constellations of micro-particles would not be there if the idea did not exist, as an idea). Secondly, in many cases they have further causal power as shown by the examples of the construction of the Jumbo Jet and the destruction of Dresden. Each required both an initial idea, and resulting detailed plan and an intention to carry it out. Hence such ideas are indubitably real in the sense that they must be included in any complete causal scheme for the real world. You can if you want to deny the reality of this feature - and you will end up with a causal scheme lacking many causal features of the real world (you will have to say that the Jumbo Jet came into existence without a cause, for example!) .

World 2b: The world of Human Goals, Intentions, Sensations, and Emotions. This world of motivation and senses is also ontologically real, for it is clear that they do indeed exist in themselves, for example they may all be described in novels, magazines, books, etc, thus being causally effective in terms of being physically represented in such writings. Additionally many of them cause events to happen in the physical world – for example the emotion of hate can cause major destruction both of property and lives, as in Northern Ireland and Israel and many other places. In World 2b, we find the goals and intentions that cause the intellectual ideas of World 2a to have physical effect in the real world.

World 2c: The World of Explicit Social Constructions. This is the world of language, customs, roles, laws, etc, which shapes and enables human social interaction. It is developed by society historically and through conscious legislative and governmental processes. It gives the background for ordinary life, enabling Worlds 2a and 2b to function, particularly by determining the means of social communication (language is explicitly a social construction). It is also directly causally effective, for example speed laws and exhaust emission laws influence the design both of automobiles and road signs, and so get embodied in the physical shapes of designed structures in World 1; the rules of chess determine the space of possibilities for movements of chess pieces on a chess board. It is socially realised and embodied in legislation, roles, customs, etc.

3.4 Physical and biological possibilities

World 3: *The world of Aristotelian Possibilities.* This characterizes the set of all physical possibilities, from which the specific instances of what actually happens in World 1 are drawn.

This world of possibilities is ontologically real because of its rigorous prescription of the boundaries of what is possible - it provides the framework within which World 1 exists and operates, and in that sense is causally effective. It can be considered to comprise two major parts:

World 3a: The world of physical possibilities, delineating possible physical behaviour .

World 3b: The world of biological possibilities, delineating possible biological organisation.

These worlds are different from the world of material things, for they provide the background within which that world exists. In a sense they are more real than that world, because of the rigidity of the structure they impose on World 1. There is no element of chance or contingency in them, and they certainly are not socially constructed (although our understanding of them is so constructed). They rigidly constrain what can happen in the physical world, and are different from each other because of the great difference between what is possible for life and for inanimate objects. In more detail:

World 3a: The world of physical possibilities, delineating possible physical behaviour (it is a description of all possible motions and physical histories of objects). Thus it describes what can actually occur in a way compatible with the nature of matter and its interactions; only some of these configurations are realised through the historical evolutionary process in the expanding universe. We do not know if laws of behaviour of matter as understood by physics are prescriptive or descriptive, but we do know that they rigorously describe the constraints on what is possible (you cannot move in a way that violates energy conservation; you cannot create machines that violate causality restrictions; you cannot avoid the second law of thermodynamics; and so on). This world delineates all physically possible actions (different ways particles, planets, footballs, automobiles, aircraft can move, for example); from these possibilities, what actually happens is determined by initial conditions in the universe, in the case of interactions between inanimate objects, and by the conscious choices made, when living beings exercise volition.

If one believes that physical laws are prescriptive rather than descriptive, one can view this world of all physical possibilities as being equivalent to a complete description of the set of physical laws (for these determine the set of all possible physical behaviours, through the complete set of their solutions). The formulation given here is preferable, in that it avoids making debatable assumptions about the nature of physical laws, but still incorporates their essential effect on the physical world. Whatever their ontology, what is possible is described by physical laws such as the Second Law of Thermodynamics:

$$dS > 0 ,$$

Maxwell's laws of Electromagnetism:

$$\mathbf{F}_{[ab;c]} = 0, \quad \mathbf{F}^a{}_b{}_{;b} = \mathbf{J}^a, \quad \mathbf{J}^a{}_{;a} = 0 ,$$

and Einstein's Law of Gravitation:

$$\mathbf{R}_{ab} - \frac{1}{2} \mathbf{R} g_{ab} = \mathbf{k} T_{ab}, \quad T_{ab};b = 0.$$

These formulations emphasize the still mysterious extraordinary power of mathematics in terms of describing the way matter can behave, and each partially describe the space of physical possibilities.

World 3b: The world of biological possibilities, delineating all possible living organisms. This defines the set of potentialities in biology, by giving rigid boundaries to what is achievable in biological processes. Thus it constrains the set of possibilities from which the actual evolutionary process can choose - it rigorously delineates the set of organisms that can arise from any evolutionary history whatever. This ‘possibility landscape’ for living beings underlies evolutionary theory, for any mutation that attempts to embody a structure that lies outside its boundaries will necessarily fail. Thus even though it is an abstract space in the sense of not being embodied in specific physical form, it strictly determines the boundaries of all possible evolutionary histories. In this sense it is highly effective causally.

Only some of the organisms that can potentially exist are realised in World 1 through the historical evolutionary process; thus only part of this possibility space is explored by evolution on any particular world. When this happens, the information is coded in the hierarchical structure of matter in World 1, and particularly in the genetic coding embodied in DNA, and so is stored via ordered relationships in matter; it then gets transformed into various other forms until it is realised in the structure of an animal or plant. In doing so it encodes both a historical evolutionary sequence, and structural and functional relationships that emerge in the phenotype and enable its functioning, once the genotype is read. This is the way that directed feedback systems and the idea of purpose can enter the biological world, and so distinguishes the animate from the inanimate world. The structures occurring in the non-biological world can be complex, but they do not incorporate ‘purpose’ or order in the same sense. Just as World 3a can be thought of as encoded in the laws of physics, World 3b can be thought of as encoded in terms of biological information, a core concept in biology (Kuppers, 1990, Pickover 1995, Rashidi and Buehler 2000) distinguishing the world of biology from the inanimate world.

3.5 Abstract (Platonic) reality

World 4: *The Platonic world of (abstract) realities that are discovered by human investigation but are independent of human existence. They are not embodied in physical form but can have causal effects in the physical world.*

World 4a: Mathematical Forms. The existence of a Platonic world of mathematical objects is strongly argued by Penrose, the point being that major parts of mathematics are discovered rather than invented (rational numbers, zero, irrational numbers, and the Mandelbrot set being classic examples). They are not determined by physical experiment, but are rather arrived at by mathematical investigation. They have an abstract rather than embodied character; the same abstract quantity can be represented and embodied in many symbolic and physical ways. They are independent of the existence and culture of human beings, for the same features will be discovered by intelligent beings in the Andromeda galaxy as here, once their mathematical understanding is advanced enough (which is why

they are advocated as the basis for inter-stellar communication). This world is to some degree discovered by humans, and represented by our mathematical theories in World 2; that representation is a cultural construct, but the underlying mathematical features they represent are not - indeed like physical laws, they are often unwillingly discovered, for example irrational numbers and the number zero (Seife 2000). This world is causally efficacious in terms of the process of discovery and description (one can for example print out the values of irrational numbers or graphic versions of the Mandelbrot Set in a book, resulting in a physical embodiment in the ink printed on the page).

A key question is what if any part of logic, probability theory, and physics should be included here. In some as yet unexplained sense, the world of mathematics underlies the world of physics. Many physicists at least implicitly assume the existence of

World 4b: Physical laws, underlying the nature of physical possibilities (World 3a).

Quantum field theory applied to the standard model of particle physics is immensely complex (Peskin and Schroeder 1995). It conceptually involves *inter alia*,

- Hilbert spaces, operators, commutators, symmetry groups, higher dimensional spaces;
- particles/waves/wave packets, spinors, quantum states/wave functions;
- parallel transport/connections/metrics;
- the Dirac equation and interaction potentials, Lagrangians and Hamiltonians;
- variational principles that seem to be logically and/or causally prior to all the rest.

Derived (effective) theories, including classical (non-quantum) theories of physics, equally have complex abstract structures underlying their use: force laws, interaction potentials, metrics, and so on.

There is an underlying issue of significance: *What is the ontology/nature of existence of all this quantum apparatus, and of higher level (effective) descriptions?* We seem to have two options.

(A) *These are simply our own mathematical and physical constructs* that happen to characterise reasonably accurately the physical nature of physical quantities.

(B) *They represent a more fundamental reality* as Platonic quantities that have the power to control the behaviour of physical quantities (and can be represented accurately by our descriptions of them).

On the first supposition, the ‘unreasonable power of mathematics’ to describe the nature of the particles is a major problem – if matter is endowed with its properties in some way we are unable to specify, but not determined specifically in mathematical terms, and its behaviour happens to be accurately described by equations of the kind encountered in present-day mathematical physics, then that is truly weird! Why should it then be possible that *any* mathematical construct whatever gives an accurate description of this reality, let alone ones of such complexity as in the standard theory of particle physics? Additionally, it is not clear on this basis why all matter has the same properties – why are electrons here identical to those at the other side of the universe? On the second supposition, this is no longer a mystery – the world is indeed constructed on a mathematical basis, and all matter everywhere is identical in its properties. But then the question is how did that come about? How are these mathematical laws imposed on physical matter? And which of the various alternative forms (Schroedinger, Heisenberg, Feynman; Hamiltonian, Lagrangian) is the ‘ultimate’ one? What is the reason for variational principles of any kind?

Another proposal is

World 4c: Platonic aesthetic forms, providing a foundation for our sense of beauty.

In this paper, those further possibilities will not be pursued. It is sufficient for my purpose to note that the existence of a World 4a of mathematical forms, which I strongly support, establishes that this category of world indeed exists and has causal influence.

3.5 Existence and Epistemology.

The overall family of worlds. The major proposal of this section is that *all these worlds exist – Worlds 1 to 4 are ontologically real and are distinct from each other*, as argued above.

These claims are justified in terms of the effectiveness of each kind of reality in influencing the physical world. What then of epistemology? Given the existence of the various worlds mentioned above, the proposal here is that,

Epistemology is the study of the relation between World 2 and Worlds 1, 3, and 4. It attempts to obtain accurate correspondences to quantities in all the worlds by means of entities in World 2a.

This exercise implicitly or explicitly divides World 2a theories and statements into (i) true/accurate representations, (ii) partially true/misleading representations, (iii) false/poor/misleading representations, and (iv) ones where we don't know the situation. These assessments range from statements such as 'It is true her hair is red' or 'There is no cow in the room' to 'electrons exist', 'Newtonian theory is a very good description of medium scale physical systems at low speeds and in weak gravitational fields' and 'the evidence for UFO's is poor'. This raises interesting issues about the relation between reality and appearance, e.g. everyday life gives a quite different appearance to reality than microscopic physics – as Eddington pointed out (Eddington, 1928) a table is actually mostly empty space between the atoms that make up its material substance, but in our experience is a real hard object. As long as one is aware of this, it can be adequately handled.

There is a widespread tendency to equate Epistemology and Ontology. This is an error, and a variety of examples can be given where it seriously misleads. This is related to a confusion between World 2 and the other Worlds discussed here which seems to underlie much of what has happened in the so-called 'Science Wars' and the Sokal affair (Sokal 2000). The proposal here strongly asserts the existence of independent domains of reality (Worlds 1, 3, and 4) that are not socially constructed, and implies that we do not know all about them and indeed cannot expect to ever understand them fully. That ignorance does not undermine their claim to exist, quite independent of human understanding. The explicit or implicit claim that they depend on human knowledge means we are equating epistemology and ontology – just another example of human hubris.

Section 4: The nature of causality

The key point about causality in this context is that simultaneous multiple causality (inter-level, as well as within each level) is always in operation in complex systems. Any attempt to characterise any partial cause as the whole (as characterised by the phrase 'nothing but') is a fundamentally misleading position (indeed this is the essence of fundamentalism). This is important in regard to claims that any of physics, evolutionary biology, sociology, psychology, or whatever are able to give *total* explanations of any specific properties of the mind. Rather they each provide partial and incomplete explanations. There are always multiple levels of explanation that all hold at the same time: no single explanation is complete, so one can have a top-down system explanation as well as a bottom-up explanation, *both being simultaneously applicable*.

Analysis 'explains' the properties of the machine by analysing its behaviour in terms of the functioning of its component parts (the lower levels of structure). Systems thinking (Churchman 1968) 'explains' the behaviour or properties of an entity by determining its role or function within the higher levels of structure (Ackoff 1999). For example, the question: 'Why does an aircraft fly?' can be answered,

- in *bottom up terms*: it flies because air molecules impinge against the wing with slower moving molecules below creating a higher pressure as against that due to faster moving molecules above, leading to a pressure difference described by Bernoulli's law, etc. etc.,
- in terms of *same-level explanation*: it flies because the pilot is flying it, and she is doing so because the airline's timetable dictates that there will be a flight today at 16h35 from London to Berlin,
- in terms of *top-down explanation*: it flies because it is designed to fly! This was done by a team of engineers working in a historical context of the development of metallurgy, combustion, lubrication, aeronautics, machine tools, computer aided design, etc. etc., all needed to make this possible, and in an economic context of a society with a transportation need and complex industrial organisations able to mobilise all the necessary resources for design and manufacture.

These are all simultaneously true explanations. The higher level explanations rely on the existence of the lower level explanations in order that they can succeed, but are clearly of a quite different nature than the lower level ones, and certainly not reducible to them nor dependent on their specific nature. They are also in a sense deeper explanations than the lower level ones.

The point is fundamental. The analytic approach ignores the environment and takes the existence of the machine for granted; from that standpoint, it enquires as to how the machine functions. This enables one to understand its reliable replicable behaviour. But it completely fails to answer why an entity exists with that specific behaviour. Systems analysis in terms of purpose within the higher-level structure, where it is one of many interacting components, provides that answer – giving another equally valid, and in some ways more profound, explanation of why it has the properties it has. This approach determines the rationale, the *raison d'être* of the entity; given that purpose, it can usually be fulfilled in a variety of different ways in terms of structure at the micro level.

Finally, it is not just matter or information that has physical effect. It is also thoughts and emotions, and so intentions. Although physicists don't usually recognise all of these realities, their causal models of the real world will be incomplete unless they include them. Human thoughts can cause real physical effects; this is a top-down action from the mind to the physical world. At present there is no way to express this interaction in the language of physics, even though our causal schemes are manifestly incomplete if this is not taken into account. The minimum requirement to do so is to include the relevant variables in the space of variables considered. That then allows them and their effects to become a part of physical theories - perhaps even of fundamental physics

Section 5: The relation to fundamental physics

Fundamental physics underlies this complexity by determining the nature and interactions of matter. The basic question for physicists is what are the aspects of fundamental physics that allow and enable this extraordinarily complex modular hierarchical structure to exist, where the higher levels are quite different than the lower levels and have their own ontology; and what are the features that allow it to come into being (i.e. that allow its historical development through a process of spontaneous structure formation)?

5.1: A 'Theory of Everything'

The physical reasons allowing this independence of higher level properties from the nature of lower level constituents have been discussed inter alia by Anderson (1972), Schweber (1993), and Kadanoff (1993), focusing particularly on the renormalisation group; however more than that is needed to create fundamentally different higher order structures than occur at lower levels. This is a fundamentally important property of physics, underlying our everyday lives and their reality. Its source is the nature of quantum field theory applied to the micro-properties of matter as summarised in the standard model of particle physics. It would be helpful to have more detailed studies of which features of quantum field theory on the one hand, and of the standard particle/field model with all its particular symmetry groups, families of particles, and interaction potentials on the other, are the keys to this fundamental feature.

What is 'fundamental physics' in the sense: *what feature of physics is the key to existence of truly complex structures?* What for example allows modular separation of sub-nuclear, nuclear, atomic, and molecular properties from each other in such a way as to allow the development and functioning of DNA, RNA, proteins, and living cells? Whatever it is, this must claim to be the 'truly fundamental' feature of physics - it is the foundation of the complexity we see. Is the key to such a 'theory of everything'

- the general nature of *quantum theory* (e.g. superposition, entanglement, decoherence) and its classical limit?
- the specific nature of *quantum field theory* and *quantum statistics* (certainly required for the stability of matter) and/or of *Yang-Mills gauge theory* ?
- the specific *potentials* and *interactions* of the standard particle physics model and its associated *symmetry groups* ?
- the *basic particle properties* (existence of three families of quarks, leptons, and neutrinos, for example)?

- the *basic properties of fundamental forces* (effective existence of four fundamental forces; their unification properties)?
- the specific *masses* and *force strengths* involved?
- the value of specific constants such as the fine structure constant?
- or is it not any one of these, but rather the combination of all of them?

The latter seems most plausible: but if so, then why do they all work together so cunningly as to allow the high level emergence of structure discussed in this article?

Whatever the conclusion here, we ultimately face the fundamental metaphysical issue, emphasized by John Wheeler (19xx): *What chooses this set of laws/behaviours, and holds it in existence/in operation?* And on either view the anthropic issue remains: *Why does this specific chosen set of laws that has come into being allow intelligent life to exist by allowing this marvellous hierarchical structuring?* These issues arise specifically in terms of current scientific speculations on the origin of the universe by numerous workers (Tryon, Hartle, Hawking, Gott, Linde, Turok, Gasperini, et al, see Ellis 2000 for a brief summary). Most of these proposals either envisage a creation process starting from some very different previous state, which then itself needs explanation, and is based in the validity of the present laws of physics (which are therefore invariant to some major change in the status of the universe, such as a change of space-time signature from Euclidean to hyperbolic); or else represent what is called ‘creation from nothing’, but in fact envisages some kind of process based on all the apparatus of quantum field theory mentioned above – which is far from nothing! In both cases the laws of physics in some sense pre-exist the origin of the present expansion phase of the universe. In the case of ever-existing universes, the same essential issue arises: some process has chosen amongst alternative possibilities. Whatever these processes are, they do not obviously by themselves imply the possibility of the structuring discussed here.

A further interesting question is how views on all of this would change if indeed a successful physical ‘Theory of Everything’ as usually envisaged by particle physicists – perhaps based on ‘M-theory’ or superstring theory – were to be developed. Physically speaking, this would have a logically and causally superior status to the rest of physics in the sense of underlying all of physics at a fundamental level. The puzzle regarding complexity would not be solved by existence of such a theory, it would be reinforced, for such a theory would in essence have the image of humanity built into it – and why that should be so is far from obvious, indeed it would border on the miraculous if a logically unique theory of fundamental physics were also a TOE in the sense envisaged above. That would be a coincidence of the most extraordinary sort.

5.2: The relation to ultimate reality

If we ask the question, *What is “ultimate reality”?*, we find a delightful ambiguity:

- Is it the fundamental physics that allows all this to happen? - its physical causal foundations?
- Is it the highest level of structure and complexity it achieves? – which is the ultimate in emergent structure and behaviour?
- Is it the ethical basis that ultimately determines the outcome of human actions and hence of social life, and whatever may underlie this basis?

- Is it the meta-physical underpinning of the fundamental physics, on the one hand, and of cosmology on the other - whatever it is that 'makes these physical laws fly' (John Wheeler 19xx), rather than any others?

From a physical viewpoint, it can be suggested it is the latter: the causal hierarchy rests in metaphysical ultimate reality as indicated in Figure 8.

<i>Metaphysics</i>	<i>Ethics</i>
<i>Cosmology</i>	<i>Sociology</i>
<i>Astronomy</i>	<i>Psychology</i>
<i>Geology</i>	<i>Physiology</i>
<i>Materials</i>	<i>Biochemistry</i>
<i>Chemistry</i>	
<i>Physics</i>	
<i>Particle Physics</i>	
<i>'Theory of Everything'</i>	
<i>Metaphysics</i>	

Figure 8: The Hierarchy of causal relations (Figure 5) extended to show the metaphysical issues that arise both at the foundations of the physics and in terms of causation in cosmology. In both cases issues arise that are beyond investigation by the scientific process of experimental testing; hence they are meta-physical rather than physical.

Here the unknown metaphysical issues that underlie both the choice of specific laws of physics on the one hand, and specific initial conditions for cosmology on the other, are explicitly recognised. It is possible that information is the key in both cases. Others at this meeting discuss the 'it from bit' idea in the context of physics; in this paper I have emphasized it in the context of complex systems, where information plays a central role in their emergence. The loop would close if it figured in a fundamental way in both arenas.

It is a pleasure to dedicate this paper to John Wheeler, who has done so much to explore the nature of fundamental issues underlying physics.

References

- Aberle, D F et al (1950): 'The Functional Pre-requisites of a Society'. *Ethics*, Vol **60** .
- Ackoff, R (1999): *Ackoff's Best: His Classic Writings in Management* (Wiley and Sons).
- Anderson, P W (1972): 'More is Different'. *Science* **177**, 377. Reprinted in P W Anderson: *A Career in Theoretical Physics*. (World Scientific, Singapore. 1994).
- Ashby, R (1958): *An Introduction to Cybernetics*. (Chapman and Hall, London).
- Bar-Yam, Y (Ed) (2000): *Unifying Themes in Complex Systems* (Perseus Books, Cambridge, Mass).
- Beer, S (1966): *Decision and Control*. (Wiley, New York).
- Beer, S (1972): *Brain of the Firm*. (Wiley, New York).
- Bishop, C (2000): *Neural Networks for Pattern Recognition* (Oxford University Press, Oxford).
- Booch, G (1994): *Object Oriented Analysis and Design with Applications* (Addison Wesley, Menlo Park).
- Campbell, D T (1974): 'Downward Causation'. in *Studies in the Philosophy of Biology: Reduction and Related Problems* , Ayala, F J, and Dobzhansky, T (eds) (University of California Press, Berkeley).

- Campbell, N A (1991): *Biology* (Benjamin Cummings).
- Carpenter, G A and Grossberg, S (1991): *Pattern Recognition by Self-organising Neural Networks* (MIT Press).
- Churchman, C W (1968): *The Systems Approach* (Delacorte Press, New York).
- Courtois, P (1985): 'On time and space decomposition of complex structures' *Commun ACM* **28**(6), 596.
- Davies, P C W (1974): *The Physics of Time Asymmetry*. (Surrey University Press, London.).
- Deacon, T (1997): *The Symbolic Species: The co-evolution of language and the human brain*. (Penguin, London)
- Eddington, A S (1928): *The Nature of The Physical World* (Cambridge University Press, Cambridge).
- Edelmann, G (1990): *Neural Darwinianism* (Oxford University Press, Oxford).
- Ellis, G F R (2000): 'Before the Beginning: Emerging Questions and Uncertainties'. In *Toward a New Millenium in Galaxy Morphology*, Ed. D Block, I Puerari, A Stockton and D Ferreira (Kluwer, Dordrecht).
- Ellis, G F R (2002): 'The universe around us: an integrative view of science and cosmology', see <http://www.mth.uct.ac.za/~webpages/ellis/cos0.html>.
- Emery, F (Ed) (1972): *Systems Thinking*. (Penguin, London).
- Flood, R L and C, E R (1990): *Dealing with Complexity: An Introduction to the Theory and Application of Systems Science* (Plenum Press: London)
- Gilbert, S F (1991): *Developmental Biology* (Sinauer, Sunderland, Mass).
- Holland, J H (1998): *Emergence: from Chaos to Order*. (Perseus Books).
- Isham, C J (1997): *Lectures on Quantum Theory* (Imperial College Press, London).
- Kadanoff, L (1993): *From Order to Chaos. Essays: Critical, Chaotic, and Otherwise*. (World Scientific, Singapore).
- Koppers, B O (1990); *Information and the Origin of Life*. (MIT Press, Cambridge, Mass. 1990).
- Milsum, ?? (1966): *Biological Control Systems Analysis* McGraw Hill).
- Mishra, R K et al (Eds): *On Self Organisation. An Interdisciplinary Search for a Unifying Principle*. (Springer).
- Moyers, B (1993): *Healing and the Mind*. (Doubleday, New York).
- Panksepp, J (1998): *Affective Neuroscience: The Foundation of Human and Animal Emotions* (Oxford University Press, Oxford).
- Peacocke, A R (1989): *An Introduction to the Physical Chemistry of Biological Organization*. (Oxford University Press, Oxford).
- Peacocke, A R (1993): *Theology for a Scientific Age: Being and Becoming* (Fortress Press, Minneapolis).
- Penrose, R (1989): *The Emperor's New Mind*. (Oxford University Press, Oxford).
- Penrose, R (1997): *The Large, The Small, and the Human Mind* (Cambridge University Press, Cambridge).
- Peskin, M E and Schroeder, D V (1995): *An Introduction to Quantum Field Theory* (Perseus Books, Reading, Mass).
- Pickover, C A (1995): *Visualizing Biological Information*. (World Scientific, Singapore).
- Popper, K and Eccles, J (1977): *The Self and its Brain: An Argument for Interactionism* (Springer, Berlin.).
- Rashidi, H H and Buehler, L K (2000): *Bioinformatics Basics: Applications in Biological Science and Medicine* (CRC Press, Boca Raton).
- Schellnhuber, H J (1999): 'Earth System analysis and the second Copernican revolution'. *Nature* issue 6761, 2 December 1999 (supplement to vol **402**), C19.
- Schweber, S (1993): 'Physics, Community, and the Crisis in Physical Theory'. *Physics Today*, November, 34.
- Seife, J (2000): *Zero: The Biography of a Dangerous Idea* (Penguin, London).
- Sellars, R W (1932): *The Philosophy of physical realism* (Russell and Russell, New York).
- Simon, H A (1962): "The Architecture of Complexity". *Proc Am Phil Soc* **106**.
- Simon, H A (1982): *The Sciences of the Artificial* (MIT Press, Cambridge, Mass).
- Sokal, A (2000): see <http://www.physics.nyu.edu/faculty/sokal/index.html> and <http://zakuro.math.tohoku.ac.jp/~kuroki/Sokal/>.
- Sternberg, E (2000): *The Balance Within: The Science Connecting Health and Emotions*. (W H Freeman, New York).
- von Bertalanffy, L (1973): *General Systems Theory*. (Penguin, London).
- Tannebaum, A S, (1990): *Structured Computer Organisation* (Prentice Hall, Englewood Cliffs).

Wheeler, J A (19xx): what makes the laws of physics fly?
Wolpert, L (1998): *Principles of Development* (Oxford University Press, Oxford).
Zeh, H D (1992): *The Physical Basis of the Direction of Time*. (Springer-Verlag, Berlin.).

George Ellis, Cape Town
Version 2002-09-19