the present time. Even if START II process is continued and START III also is completed both USA and Russia will still have more than 2000 NW in their arsenal. This level will be reached even optimistically only after about ten years. According to the projection by the US National Academy, this stage will have to be followed by a stage where US and Russia will cut their weapons to a level of about a thousand. This will set the stage for China, France and the UK to enter into disarmament negotiation. Countries like India, Israel and Pakistan will enter the picture only at a much later stage.

We should therefore concentrate on efforts that will halt clandestine transfer of materials and technology relating to nuclear weapons and missiles. At present, the United States is the sole arbitrator in these matters and its action is governed merely by its self-interests rather than true facts as have been established sometimes even by the US agencies themselves. We should try to find a mechanism by which, an arm of the UN, for example, can seriously investigate complaints from member states whose security is threatened by theft of nuclear materials by terrorist groups or clandestine transfer of weapons/missiles-related materials and technology and enable the Security Council to take remedial steps. (For example, punitive measures or at least restoration of status quo ante.)

It will only be a matter of time, before let us say a cargo ship registered in country A, with goods meant for country B from country C, is surreptitiously used by some terrorist groups to explode a crude nuclear device in some port of call of this ship. The concept of nuclear deterrence is such an unstable idea that, a suicidal group can easily trigger a nuclear war and destroy all human civilization. It is a sad commentary on human affairs that we have to depend on nuclear weapons to prevent large scale conventional wars between major powers.

Let me summarize:

1. Nuclear devices or weapons already have been in existence in South Asia for some years now and cannot be wished away.
2. If, in the opinion of our scientists, we need to do additional tests, we should carry them out at an appropriate time.
3. We must cap our nuclear goals – we must limit the role of nuclear weapons only for a second strike in retaliation for a nuclear strike against India only. We must maintain our integrity in not transferring technology in clandestine fashion or offer nuclear protection to any other country even if such a situation arises in future.
4. The CTBT and FMCT should be discussed in parliamentary committees with inputs from knowledgeable citizens, before we sign them.
5. We will be justified in demanding a better mechanism for preventing clandestine transfer of materials and technology related to NW/missiles as a precondition to our signing the CTBT and FMCT. Given our excellent track record in honouring international agreements, we are in a strong position to demand this. This goal makes more sense than abolition of all nuclear weapons which is definitely not going to happen in the next couple of decades.

ACKNOWLEDGEMENT. In writing this I have received help from friends both in Government and outside it. I have also used information available from websites. However, I am solely responsible for any factual errors or misunderstanding. The views that I have expressed bear no relation to those of my employer, the Indian Institute of Science.

J. Pasupathy is in the Centre for Theoretical Studies, Indian Institute of Science, Bangalore 560 012, India.

COMMENTARY

Emergence, reductionism and the seamless web: When and why is science right*

P. W. Anderson

In modern science 'what you see' is very seldom 'what there is'. Every object or fact is explained in terms of something very far from direct human perception. From the days of Copernicus, Galileo and Newton, when the mythical charioteers driving the sun, moon, and stars across the heavens were replaced by a structure of abstract natural laws, the tendency has been to analyse more and more deeply, reducing more and more of the complex world around us to the consequences of ever simpler, but more abstract and microscopic laws. In the 20th century this process accelerated as bodies were reduced to atoms and the atoms, at first immutable, then were reduced mostly to electrons and empty space. The nucleus then itself became doubly composite, with the finest particles being objects which can never appear in isolation.

Physicists – and scientists in general – love to do two things: to take

*This text is based on the Pagels lecture delivered at Aspen, Colorado, August 1999.

apart, to analyse into smaller and simpler components; and to mystify, to say it is not really THIS, it's THAT. They like to take upon themselves the role of the shaman or the mullah. Everything comes from a First Cause – the First Equation – and only the appropriate scientist can investigate this with his very expensive equipment, and understand it with his abstruse theories. This attitude is typified by Steven Weinberg's book *Dreams of a Final Theory* where he repeats his claims of many years ago that the only really fundamental science is the search for the ultimate constituents of matter. He classifies all of science by following 'little white arrows' from the everyday object, inevitably leading back through analysis after analysis to the 'fundamental' problem: what are the particles, which may even not be particles at all but strings in more dimensions than you can believe, made of?

In this kind of discussion there is repeated reference to a hypothetical 'perfect computer' which could, one is told, in principle follow out all the consequences of these 'Theories of Everything', making them the only input that matters. One of these particle physicists even wrote a book called *The God Particle*.

Thirty years ago, as a cry of anguish at this arrogance of the particle physicists, I wrote an article called 'More is different', the idea being to look at things from the opposite direction. The reductionist point of view says that all this complexity in the actual world we experience is 'just' the consequence of the simpler first causes, so let us look only for them. My approach was to accept the same material facts – and to accept that indeed everything in the world did indeed obey the fundamental natural laws – but the hooker is in the world 'obey'. Does that mean anything? What is missing is that perfect, omniscient computer that is going to tell you how the complex world follows from those excessively simple equations. Complex consequences can arise from simple causes, especially when you allow the system which you are looking at to be at a different scale from the elementary entities of which it is made. The process already had a name, although I was not aware of it at the time; 'emergence', the term used by the 19th century biologists to describe the emergence of increasingly complicated forms of life on a previously inanimate planet.

The message is that the world in which we live is the consequence not only – or even primarily – of some incredibly simply but hidden god equation or 'god particle' but of the 'god principle' of emergence at every scale and every level. (Please realize that I use the deity's name only in self-defence.)

I first became aware of this fact in my own field of condensed matter physics, where the underlying laws have been known to a very high degree of accuracy at least since 1932, when P. A. M. Dirac concluded his paper on the theory of the electron with the words 'we have now laid down the principles underlying much of physics and all of chemistry' – which has been paraphrased as 'the rest is (mere) chemistry'. He implied correctly that the quantum theory of the electron (and of the electromagnetic field in a simple approximation) as well as the immutability of the nucleus could be shown to provide the substrate of laws on which all of atomic, molecular and condensed matter physics is based. Within a few years we added enough knowledge of the nucleus to underpin geology and most of astronomy; and in the forties and fifties biology was 'reduced' to chemistry.

Why then are people still making discoveries in condensed matter physics? The example of superconductivity is instructive. This phenomenon of the disappearance of resistance at low temperature in most metals was discovered in 1911, and intensively studied for decades thereafter. But the phenomenon received an explanation in terms of those fundamental laws only in 1957, and this was generally accepted by the early sixties, 50 years after the discovery and 30 years after all the 'fundamental' laws behind it were known. If the strong principle of reductionism was true, and we could just compute all consequences of the laws of motion, how could this be? In 1956, I heard probably the greatest of all calculators, Dick Feynman, declare publicly that he had calculated the best he knew how and could not find any behaviour resembling superconductivity – and he was only the best of the many brilliant physicists including Einstein and Heisenberg who had tried.

It is important to realize that even the hypothetically perfect computer would have done no better: if it had simply simulated the observed experimental behaviour of a superconductor by following all the electrons and nuclei exactly, it would have told us nothing new, unless it told us WHY. Unfortunately, it would not know what that question means. What it does mean is that there are certain concepts and constructs which allow enormous compression of the brute force calculational algorithm down to a set of ideas which the human mind can grasp as a whole.

The basic concept here which made it possible was what is called 'broken symmetry': the underlying laws have some symmetries which are obligatory for small systems like nuclei or atoms but which are not manifest in their consequences for large systems of atoms. A single atom of lead, for instance, obeys quantum mechanics and is a rather featureless ball; put a lot of them together and they form a crystalline solid (the first broken symmetry) which is a greyish, shiny, soft metal (a second one) and becomes superconducting at 8 K (still a third). Experiment had taught us the concepts of stacking which caused the first, and of quasifree electrons which cause the second; but it took a stroke of brilliance to find the third, new kind of broken symmetry, and more of the same to understand why and what caused it.

All of these concepts are only precisely definable in what we call the 'thermodynamic' or $N \rightarrow \infty$ limit; rather than building the solid up atom by atom, we conceptualize it as essentially infinite, and then back off to the finite but large size we actually are. This, in physics, is why the computer fails and has to be replaced by the much more chancy process of staring at the experiments for long enough to figure out a concept, or build a simplified model, or whatever it takes to answer that hard question WHY? The computer can never deal with an infinite system. It is particularly hard because the computer's difficulty increases exponentially, not proportionally to N.

Nowadays we have a new kind of superconductivity which occurs at much higher temperatures and has been frustrating us for a dozen years now by clearly requiring a whole extra concept or set of concepts which the community cannot agree upon. It certainly has the 1957 kind of broken symmetry, and another feature (called for short d-wave) which we invented back in 1960; but there is still a missing link, which I am sure I understand but unfortunately so do several other people and their ideas at least appear not to be the same as mine. But what is clear is that it is inspiration, not perspiration, which is required.

It is not explicitly the thermodynamic limit, but the more general idea that large systems can exhibit behaviour which is generically different from that of the units from which they are made, which is the lesson to be carried on to more general contexts. There is an intellectual autonomy which the higher or more complex level has from the simpler substrate on which it is built: organisms have properties which are meaningless for cells, as for example differentiation into specialized tissues. Societies do things which individual people do not. For example, every advanced society seems to reinvent the concept of money, and so on. One can wholly understand the laws of motion of the substrate in which a given phenomenon appears, yet have no handle on how it arises or what it consists of. In condensed matter physics this is almost the rule: again and again, one finds the experimentalists calling to the attention of the bewildered theorist something bizarre, new, and unexpected. Nature is better at lateral thinking than the human mind, and the human mind is better at it than any other calculator yet conceived.

It seems to me now, as it did 30 years ago, that the important and exciting problems in science cluster around these contacts between the different levels of scale and of complexity. The perennial and puzzling problem of the origin of life, and the even more difficult problem of the origin of Monod's 'teleonomy' – independent purposive action – are of this kind, as is the recently fashionable question of consciousness, surely an emergent property if any is.

From this point of view the social sciences are not distinct in principle, only in their distance from the substrate of fundamental certainty – the number of 'little white arrows' which are missing. Yet the existence of parallelisms across cultures, economies and societies tempts us into conjecturing that there must be some kind of general understanding of how social hierarchies arose from a simpler substrate. This seems a rather less inflammatory way to present what E. O. Wilson means by 'consilience'.

Let me not elaborate these ideas any further but shift gears and talk about why reduction may be worth the trouble, if it is not for the purpose of predicting behaviour from the underlying fundamental laws – which, I have tried to persuade you, is basically futile. If we are interested in, say, geology, we do not take the concept of plate tectonics which underpins geology with the physics of the earth's interior, and calculate the geological history of the earth *a priori*. Whether the answer would even be unique is not at all guaranteed: in Steve Gould's phrase, if we play the tape twice, in the case of historical sciences, it can come out quite differently. Instead, we use the concept to relate one kind of geological observation to another. It makes what appeared to be a hodgepodge of unrelated events and histories into a connected, coherent whole. This is a good example of what I mean by the 'seamless web' of science: mostly by using the kind of understanding which you can achieve by reductionist explanations of the phenomena at the more complex level in terms of the simpler one, science becomes an integrated whole rather than a congeries of independent facts. In chemistry, for instance, the quantum theory of the chemical bond changed the periodic table of the elements from a useful empirical generalization full of unexpected exceptions and provisos, to a logical outcome of the symmetries of quantum mechanics.

The argument is that we carry out the process of reduction not for the futile goal of attempting to ascertain the predictions of some First Cause, making all of the sciences simply 'applied cosmo-particle physics', but for the very real goal of relating the different fields into a seamless web of knowledge, where the validity of one part is buoyed up by its inextricable tie-up with another. Evolution, for instance, is no longer simply the province of the taxonomic biologist and palaeontologist, but of the geologist and molecular biologist, with important relationships to nuclear chemistry and to important parts of astronomy; and great swatches of these and other fields would have to be completely reconstructed if one had to abandon evolution. The Kansas state school board seems indeed to have realized that they had to abolish the Big Bang along with evolution: their logic, if not their science or their good sense, is sound.

There are, unfortunately, areas which call themselves science which do not have – or welcome – the interlevel contacts, the attempt at a more microscopic understanding which characterizes the web of consilience. One thinks of the kind of giant dietary 'study' which seems to be a favourite of journalists, of the effect, say, of broccoli on the incidence of piles, without any attempt to establish a plausible mechanism; one thinks of the electrochemical community unwilling to accept physicists' rock solid arguments against cold fusion because of their intellectual isolation; one thinks above all of the two schools of psychology, Freudianism and behaviourism, glaring at each other for decades but neither willing to investigate how the brain really works or to incorporate neurophysiological information. And the social sciences are just beginning the process of integration, but here we know we are at the cutting edge and should give them the benefit of the doubt.

What I would like to do now is to discuss this question of when and how we know science is right (to use the technical term, the epistemology of science) and actually represents something real outside of ourselves. As many of you probably know, there is a considerable school of scholars who claim that science, in particular, is 'socially constructed' as a cultural phenomenon among the 'dead white males' – or those who might as well be – who constitute the historical majority of us. Let me say immediately that I do not condemn the study of science as a social phenomenon at all. As an active participant for 50 years, at the cutting edge all of the irrational, social, emotional and personal motivations and behaviours one could possibly imagine play a very large role in what goes on, and there is hardly a major discovery I know of about which one cannot relate a longer or shorter tale of prejudice, venality, rejection, and just plain blockheaded opposition.

But now I am not talking of the latest news from the front but of the end result of this process of wrangling over the answer, which in many cases can take decades, the end result which has become incorporated in my 'seamless web'. What I want to say is simply that there is no qualitative difference between scientific knowledge and common or garden knowledge like the identify of your mother-in-law or that a lot of the objects in this room are chairs. The philosophers spent many centuries discussing what you see when you perceive a chair, and in the end the answer of Wittgenstein that such definitions are what they call 'circular' won out: the chair has a large number of properties, and when the thing you see has enough of them that it cannot be anything else, then it is a chair.

This kind of rough-and-ready definition has been confirmed and sharpened by the dreadful failures of computers in sorting out sense data compared to our own minds, and by modern experiments on perception. The way we recognize an object is very enlightening. According to experiments which follow eye motions precisely while we do everyday tasks, our eyes do not look at the whole of it – we glance at an edge, then the other, then a corner, and so on, (these are called 'saccades') and then we form a hypothesis as to what it is; with a few more glances we confirm that guess and from then on it is there in our consciousness without our looking at it in detail. Similarly, with a face: every cartoonist knows that it takes only a few features to make a face completely recognizable. Thanks to some recent work in computer science – in the study of the so-called satisfiability problem – we are beginning to have some idea of how the mathematical structure of this kind of thing works. What the group of people around Scott Kirkpatrick have been doing is studying the properties of large collections of logical propositions, chosen at random: things like 'x or y or not-z are true' involving a large number of variables N and a large number of propositions M, all of which one wants to satisfy simultaneously. This is the kind of thing you see in sense data – you notice that the brown surface is in front of the green one but has an edge to the left, and so on. What they find is that if there is little enough data it is easy to satisfy all of it, several ways; but there is a rather sharp phase transition as we increase M, to having no possibility that random data could be logically consistent. What our brains are doing is getting us past that phase transition, so that we can only find a solution because of the correlations in the data which come from the real structure of the world around us. The real world is enormously highly correlated, in ways most of which our brains learned about in early childhood: it is three dimensional, it is full of objects and beings, and so on. The trouble with solipsism is that it is so much harder to see how we could construct such a correlated world for ourselves than to suppose it is really there!

Well, I do not think that I have to draw the parallel to the web of observations and measurements which constitutes modern science. Again, we start out with individual observations – clauses – which can be fitted together in all kinds of ways, and as we learn more we find that they connect together in an increasingly rigid structure. The only new thing I am suggesting, perhaps, is the idea of a rather sudden phase transition between an era of pseudoscience, when the logical connections between different observations are so few and so tenuous that many different underlying structures are consistent with most of the data – and a possibly rather sudden transition to a regime where there is only one consistent framework in which it can be understood – and then you can even pick out the (hopefully, few, but regrettably not always) wrong observations or dishonest experimenters. It is at this stage that I believe it is as impossible to 'socially construct' science as it is to invent Alyosha Abrikosov or your mother-in-law.

P. W. Anderson is in the Department of Physics, Princeton University, Princeton NJ 08540, USA.