Complexity Science and 21st Century Issues
London School of Economics 25-26 March 2004

<u>Dr. Geoffrey West</u>, Distinguished Research Professor, Santa Fe Institute; Senior Fellow, Los Alamos National Laboratory.

"The Complexity, Unity and Simplicity of Living Systems"

One of the main things that I'm going to be talking about today is the role of scaling laws, particularly in biology, though I will say something about their possible extension to social organisations. This work which was done in very close collaboration with biologists and chemists and has evolved into questions about whether we can start to think about fundamental laws of biology. These would allow us to construct methodologies and theorems concerning the way life works, so that we can possibly have the same kind of quantitative predictions that we have, for example, in physics.

Now it's a priori truth I think, that there aren't Newton-type laws in biology. It would mean that if we took an arbitrary biological system, we would have a bunch of equations from which we would be able to derive its characteristics from first principles with any given accuracy. That was the old vision we had in classical physics, though now we have to add the caveat that Bob May yesterday gave us a very striking example of when we can't do that. Nevertheless in biology we can ask a different set of questions, which, whilst not supposing that we can get everything to a given accuracy, enables us to have what I would like to call a 'cross grain' description of biology.

Think about aging and mortality for example. We're all familiar with the symptoms of that and ten years ago I started to wonder what was actually going on inside my body. What is the mechanism? And perhaps because I came from physics, I asked a question, which is perhaps a different question from that a biologist would typically ask and that is: 'why is it that the span of my life is going to be of the order of 100 years?' Now given the ups and downs of life it may in fact be 65 or 70 years, but it's of the order of 100 years rather than a thousand or a million. So the question becomes: 'why can't I live a million years or conversely, why do I have a strong expectation to live longer than three months?' Why is it that a mouse lives only a few years and more importantly, where in molecular time and space scales is 100 years? If we read the literature on gerontology and mortality what we discover is that most of the arguments about this, though they may not be put in the same terms, are that it is genetically controlled. Which of course is not saying much, and the question remains as to where in the molecules of genes is it written that after about 100 years all human beings will be dead.

So it's that sort of question that we're going to try to answer. And it so happens that dealing with scaling laws, interestingly opens the window onto asking questions in that framework and in many cases being able to answer them. I think everybody would agree that life is the most complex and certainly the most diverse physical system in the universe and its complicatedness can be illustrated with this very complicated map of a person's metabolism. It just tells us the names of the chemicals and how they interface with each other though it tells us nothing about the dynamics. It is literally just a road map. But a person's metabolism is one kind of complexity.

Here's another kind of complexity and this is an ecosystem. I'm sure we're even more familiar with this picture because it's a tree. And the sort of questions again that we would like to ask and be able to answer are: 'How tall is this kind of tree?' 'How many leaves is it likely to have?' 'How many branches are there of a given size?' 'How far do we have to walk in a wood from this tree to find a tree of similar size?' 'How many plants are there of this size in the wood?' And so on. All are questions we might want to answer in a quantitative way. The idea being that there might be a theory that we could write as formulae and use it to answer all those questions.

However, looked at it in the right way, this extraordinary diversity and complexity, which we see exhibited from the microscopic to the macroscopic, may be shown to have an extraordinary simplicity that emerges across all of biology. We can start this approach by looking at it in turns of a scaling phenomenon; meaning that if, for example, we ask how any physiological variable, changes with size across a range of different organisms, we can find a pattern that is repeated. That's my point of departure and I will return to these questions later.

Here's a famous graph put together in the 1930s by a man named Max Twiberg (?). What is plotted here on a $\log_n/\log_n$ scale is basal metabolic rate (i.e. how much energy per unit time an organism needs to stay alive) against body mass. What we can see is a lovely straight line, and what he discovered was not only the remarkable fact that the relationship is a simple power law, but that it has a very simple slope that is very close to 3/4. Extend the data collection to all warm blooded and cold blooded life forms all the way down to unicellular organisms (which covers in size about 18 orders of magnitude) and we get the same straight line where the slope is the exponent of the power. Twiberg's original energy plot was in kilocalories per day (about 2500 for human beings), but if we plot it in watts we will notice that each of us sitting here requires less than a light bulb to stay alive. Which is an amazing statement of how extraordinarily efficient we are and how extraordinarily wasteful it is to leave the lights on, because that small amount of power could keep someone alive in Africa. I shall come back to this when talking more about ecosystems.

What is remarkable about this is not just that the slope is 3/4, but that each one of these organisms and each biological system within it, and every biological system that has ever evolved, has done so by Darwinian natural selection in its own environmental niche. And though in earth environments we might have thought that there could be some regularity, the extraordinary diversity would seem to rule out such a generalisation. And so the challenge was to explain how this could be so.

I will show you in a minute a whole bunch of variables, some mundane and some as profound as this, that display an almost identical sort of behaviour so that we end up with a vision that each organism (from uni-cells to huge plants and trees and mammals) gets the right amount of energy for every piece of its physiology and every part of its physiology is designed so that gets to be the size it needs to be. What this means is that our aorta, for example, is the right size for our body as is our heart and so on.

I said we need 100 (?) watts, but if we interpret this number as the amount of energy that we need to stay alive and a figure can be worked out for every other animal, then that's the energy that our body has evolved to require. But the amount of energy we now require since we have become cultured and civilised and social is 100 times 100 watts. We need that energy to enable us to live as we do and if we plot the graph from primitive people up to the modern age we hit a gigantic leap about 5000 years ago and it's interesting to ask the question; 'how big an organism do we each effectively function as?' It turns out we're bigger than a blue whale. And that's the problem.

OK, if the metabolic rate is scaling to the mass at 3/4 and doubling the mass doubles the number of cells, then the metabolic requirement decreases faster than the number of cells increases our cells require less energy than our dog's and a whale requires even less energy per cell. So there is an increase in efficiency with size; the bigger we are, the less work the cells need to do to keep us alive. Here from this elementary biology book we can see how dramatic that is; a gram of mouse requires three times that of a gram of dog and nine times that of an elephant. It shows how extraordinarily efficient we are. If we took our 10^{14} cells and cultured them in a Petri dish then the amount of energy we would need would not be 100 watts but 100 million watts. So we are phenomenal in that an extraordinary integration has led to an extraordinary efficiency and if we go down to our mitochondrial cells where our energy is created and even all the way down to the molecules where the relationship still holds we are talking about over 27 orders of magnitude.

I said that I'd show you other variables so here's our aorta. We can see from the graph of blood flow rate against our body mass has a slope that is very close to 3/8 and we can also see that its cross-sectional area against body mass also has a slope of 3/4. And the intriguing thing about this is that it's the same with tree trunks, which scale in exactly the same way. Here's another example of heart rate against size for a bunch of mammals and this line has a slope very close to -1/4. What is also interesting about this is that when we're born we're some way up the graph and we slide down it as we grow up.

What initially prompted me into this kind of analysis was the difficulty in getting control data on lifespan. Because there is a lot of fluctuation due to local environmental effects you need to dig deep into a lot of data to find generic information but what you find is that lifespan increases roughly with mass to the 1/4. And the important point here is that if heart rate decreases with mass to 1/4 then the number of beats in a lifespan is invariant. So big things live very long but have very slow heartbeats and little things that have very fast heart beats don't live very long but the number of heart bats is roughly the same. That's nothing very fundamental about hearts, but what is fundamental and true of all aerobic metabolism is that the number of times the reaction takes place in an organism during its lifespan to produce ATP (adenosine tri-phosphate), which is the currency of energy, is invariant. Again we see genome length versus cellular size shows a relationship close to 1/4 although there's some spread of data. And lastly a graph of the white (?) to grey matter of the brain again shows a beautifully straight line on the $\log_n/\log_n$ plot.

I won't show any more collected data examples, but I do want to show you how well the theory that I'm about to give you, can predict and explain most of what I have already told you. The theory is mathematical and technical and quite complicated but the conditions are simple. For example, I've shown you a picture of a tree and if we ask the question: 'what is the average distance in a forest between trees of similar size (i.e. same diameter)?' then it turns out that the theory predicts that it is a linear relationship with the diameter of the tree trunk. This wasn't realised before this work and the data from a virgin forest in Costa Rica bears this out very well. Another example along those lines is that if we ask: 'how many trees are there of a given size in a forest?', it turns out that the theory makes it an inverse square of the tree diameter and we can see from a graph that the slope is -2, which again is predicted. The thing that's interesting here is that we have two sets of data: one from

1947 and one from 1981 and we get the same slope, even though none of the trees in 1981 were there in 1947. And if we ask: ‘how many branches are there on a given size?’ we find the same rule applies. What the theory predicts is that the structure of the forest is a reflection of the structure of an individual tree, not just the topological structure but things like how much energy is flowing in each branch or how much water is going up and so on. Here’s a graph of leaf number verses size for a given tree and again it predicts that the relationship is with the square of the diameter.

Switching subjects slightly, this is the growth period for rats i.e. Mass against time and we’re probably familiar with this sort of graph for baby growth. We could construct this for every single organism and the theory predicts that everybody grows at the same rate (relative to body mass?) except trees and plants. And the reason for that is because they keep their dead cells whereas almost all animals shed them. So one of the main things we get from this is that if we look at things in this very coarse grained way, everything beats to the same clock. Looked at in this way, an elephant is a blown up mouse and a mouse a blown up version of a cell of a mouse. Interestingly the theory was used in the study of cancer by people at the Harvard Medical School to understand the growth of cancer cells. The results in fact were not as they should have been because the theory is based on the assumption that cells are shed whereas a cancer forms a necrotic centre which is to say it keeps its cells, and so violates the rule. So the interesting question now is: ‘what is the right formula?’ And people may be close to finding that. If we ask the question: ‘what is the difference between growing a cancer inside our body and growing a baby?’ the difference is that babies get rid of their dead cells. I’ll tell you one more thing before we move into what we believe underlies all this. Changing subjects entirely, here’s the ratio of the mitochondrial volume relative to whole body mass and the theory predicts about 1000 mitochondria per cell for us because the mitochondria are how we get our energy.

So to summarise the rather long introduction. Life is almost certainly the most complex and diverse system in the universe, yet it exhibits these extraordinarily simple scaling or power laws and the exponent of the power is typically a simple multiple of ¼. Therefore the mystical number in the universe that we live in is 4, and this somehow plays an extraordinary role in life. And I’d also like to say that this work has been done in marvellous collaboration with Tim Brown and his colleagues at the University of Mexico and the work has now expanded to Los Alamos.

Anyway, here’s the idea that we propose underlies this and if we’re interested in social structures I think it’s worth trying to translate the theory into possible metaphors for social organisations. One of the major problems that multi-cellular organisms have, and remember we are made of 10^{14} cells, is how are all these cells going to be serviced in a democratic and efficient fashion? We find that the way natural selection has dealt with that in the free market of biology is that organisms have become multi-cellular and have evolved hierarchical branching networks. Maybe not all hierarchical but overwhelmingly hierarchical. They also take something macroscopic and distribute it to something microscopic and vice versa. So in summary, the generic universal structure of these networks means that:

1. Life at all scales is supported by networks from the microscopic to the macroscopic and some of the networks are real physical structures like the cardiovascular system inside us and some show a dynamic relationship like ecosystems in the natural world.

2. The terminal units of the network, e.g. the capillaries of the circulatory system or the petioles (leaf stalks) of trees or the mitochondria in cells are invariably in the class of networks that show scaling properties. As different species evolve within a given design, natural selection does not re-invent the basic unit where energy is being transferred.

3. Of the infinitude of possible networks that could have evolved, whether circuitry, renal, neural or respiratory, the ones that have evolved under the continuous feedback of natural selection over geologic time show optimal structure. For example the cardiovascular system that we all share shows an optimal design that minimises the cardiac output. Meaning that if I took the fourth branch of my arterial system and doubled its length my heart would have to work harder, but if I halved its length my heart would also have to work harder.

4. Generic observations from the microscopic to the macroscopic can be put into mathematical form that enables us to predict.

Let's take this notion of efficiency or optimal design. How would we model it? I was going to give you a tree as an example, but I think I'll do the circulatory system. Here in words is what we do. We take an arbitrary network of the circulatory system in three dimensional space and fix its size. Then we stick a heart on the end of the aorta as a pump. Now comes the hard bit. We have to calculate how much energy is being dissipated from the network as a function of its physical characteristics and then we have to minimise it relative to all those characteristics. The size of our aorta is such that when a blood pulse comes out of the heart almost no energy is dissipated due to frictional forces. As soon as the artery branches the radii of the branches and the frictional forces will have an effect. If the cross sectional area is arbitrary some of the pulse would probably be reflected back so we need an equation that specifies the minimum configuration for which we do not have any reflections. So the constraint on the arterial system is that the cross-sectional area of the parent tube is the sum of the cross-sectional areas of the daughter tubes. It would be exactly the same way that we would design the national grid so that electrical waves don't bounce back because it minimises the amount of energy we have to put into the system. As the tubes get smaller and smaller viscous or frictional forces become more and more important and energy starts to get dissipated as heat which does damage. The blood slows down and eventually the pulse ceases. Blood comes out of our aorta at about 600 cm^3/sec. and the flow ends up almost stopping at the end of the network, so that we get efficient diffusion to the cells.

Minimising the energy relative to the network's characteristics and putting it as a mathematical equation means that we have an analytic description of our arterial network and one of the things we notice about the minimisation is that the volume of blood flow is linear with the mass of the organism and this gives us a theory for any organism with a beating heart. So if we want to know the characteristics blood flow in the eighth branch of a cat's circulatory system there is a formula. It may not be exactly our cat but averaged across a million cats this would be the right answer.

Now how does this energy requirement reflect back to the metabolic rate which is the amount of energy we need to keep us alive? We are kept alive by taking oxygen into our lungs which is transferred to the blood which then transfers it to the cells. So the blood flow rate in our aorta is actually a proxy for our metabolic rate. Since the dissolvability of oxygen in blood (per blood cell) is the same for everybody we can calculate the blood flow rate scaled with body mass and again we get a gradient of 3/4. I'll come back to the more generic question of the origin of the 4 but that's how we do the calculation.

I previously mentioned the word efficiency and we should note that the cells of our body are actually decreasing their metabolic rate as a function of our size, and that all mammals are functioning in much the same way. So let's talk about efficiency. How does that come out of this? Well if we calculate the total resistance of the circulatory system network as a function of size what we discover is that as body size increases the total resistance is decreasing. It decreases as mass to the 3/4 and it decreases because the number of outlets is increasing faster. So our total resistance is decreasing as mass to the 3/4 but the flow rate through the system is increasing as mass to the 3/4. So this is like Ohms law in which volts = current x resistance. So we can determine what our voltage or blood pressure is, by multiplying the metabolic rate by the total resistance. It is actually invariant, meaning that the blood pressure of a whale is the same as ours, and the same as a shrew. Now that's surprising because our aorta is such, a whale's is much bigger and a shrew's I cannot see because it's a fraction of a millimetre. And the theory predicts this.

So it's the decrease in resistance that drives increasing efficiency and there is a decrease in metabolic rate of a cell with size. And so we can see it is the network, the total macroscopic network that is somehow controlling the rate at which a cell has to behave. So instead of starting our explanation with the microscopic and build up to saying how this gives rise to the properties of the macroscopic we're saying something the other way round. It is the macroscopic network that is controlling what is going at the microscopic level. The production of energy is being controlled by the network or delivery system. Remove the network and instead of having cells whose energy output is decreasing with body size (mass to the ¼) we can predict that if we remove the cells and culture them in vitro they would end up being the same even though they've come from different organisms across eight orders of magnitude.

So the mathematical theory tells us not only what the slope is on a $\log_n/\log_n$ plot but it also tells us the value all warm blooded animals should evolve to in vivo. Contrast that with the value in vitro and this is very close to the prediction. We can see from the graph where the lines cross. If we look at what mass the smallest mammal should be we find it's a gram and that is what a shrew is. There are two ways of looking at this; one is that down at the end where the two cross, whatever mammal that is, its cells are working as hard as any cell possibly could. We can't get any faster than that and therefore we can't get any mammal smaller than that. So it's no wonder that a shrew doesn't live very long because its cells are working like crazy and literally wearing it out. This is the origin of a theory of aging and mortality.

There are a couple of other things I want to mention. One of the things this theory predicts is that all rates decrease as mass to the ¼, therefore time increases as mass to the ¼. There's also something else I didn't refer to and that is the role of temperature because for almost all of life, except for us mammals, temperature is a very important controlling variable and it is indirectly for us, because it controls the environment we live in. Boltzmann's formula tells us how the metabolic rate is related to mass and absolute temperature (Kelvin) and this is called the activation energy of the biochemical processes that underlie whatever is going on. So roughly:

Metabolic rate is proportional to $M^{-1/4}e^{-E/KT}$ where M = mass, E = energy, T = thermodynamic temperature and K = Boltzmann's constant.

And time is proportional to $M^{1/4}e^{E/KT}$

So these are the generic forms that all rates and times have to obey and one of the things that comes out of this is if we look at metabolic rate, growth rate, mortality rate, even the rate of evolution, all of them obey this . So if we correct for mass in terms of networks and temperature from the fact that we have statistical mechanics, then everything lives at the same rate, including evolving at the same rate even though the absolute times are extraordinarily different, by as much as 20 orders of magnitude. One of the interesting questions if we start talking about social organisations is: 'are there serious analogues to mass and temperature?' To which the answer is a cautious 'yes'.

Let me go back to the aorta and the smallest mammal. Blood is pumped down the tubes until they become so small that energy becomes dissipated in the system and blood slows down. Now as the organism gets smaller and smaller, the aorta is getting smaller and smaller as a relationship of 3/8, but eventually the aorta would become so small that it could not sustain a pulse. Also as it gets smaller and smaller the viscous effect gets bigger and bigger and we'd end up with a mammal with a beating heart but no pulse. That is very inefficient because it means that energy is being dissipated in every branch of the network and if we calculate the metabolic rate of such a system it no longer scales as mass at 3/4. It scales linearly so there's no advantage in terms of scaling or increase of efficiency. And the pigmy shrew is just about on that limit.

The energy that is being dissipated is inevitably damaging. We know that systems that are self sustaining but they are inevitably killing us because entropy is increasing and we're producing huge amounts of dissipated energy in the form of heat and things like free oxygen radicals which cause structural damage. Now we have a quantitative mathematical theorem which means that on the average we can calculate all these quantities so we can calculate how much damage we're doing to ourselves. But as we well know we also repair ourselves. Without repair we'd calculate that we'd only live a few weeks, but if we put in repair which is not given by this theory we can get an estimate of how long we should live. And thus we get an answer of the order of 100 years. We get this by the multiplication of one extraordinarily large quantity (the number of cells we have in our body) and the extraordinarily small quantity which is the analogue of E to the KT which is the energy associated with half an electron volt. I forget what it is in Joules but this gives the possibility of making a quantitative theory of aging and why it is we only live 100years.

One of the things we do know about repair is that we have a repair mechanism that has evolved. But in any calculation we have to allow for the wear and tear of life. In the past our great great grandparents didn't live much beyond 35-40, In the US in 1885 the average lifespan of a man was 48 though it is now 75. But the test of this theory is not to predict the statistics but to predict the maximum conceivable life span and no person ever has lived beyond 127 years. In other words we know if we look at mortality curves lifespan is increasing but there's an endpoint. The only way to change it is by changing repair mechanisms. That's the only kind of gene tampering that would work.

We are repairing ourselves and you are sitting here repairing yourselves now, but there's one piece of you that's hard to repair except if you doze off and that's your brain. I noticed coming into London at midnight that there was a huge traffic back-up because they were repairing the roads when the city was asleep and that's the analogue of what sleep is in this speculation. You shut down or repress in order to repair. So we can calculate from that how much sleep we need and what we get is an interesting scaling law of the ratio of sleep time to wake time and that is predicted on a scale of mass to the 1/6. The 1/6 comes about because the brain is roughly speaking the only organ in your body that does not scale as a minimum and in fact scales as mass to the 3/4. And if you take that into account you get 1/6.

When we did our calculations, on the blackboard so to speak, I thought: 'What would it be for an elephant?' So we got this formula and we stuck in numbers and it came to three hours and I said that can't be right because I thought: 'What does an elephant do for the rest of the night?' But it turns out it's true – an elephant sleeps for roughly three hours.

So let me finish on this note and say that scaling laws are a wonderful way to get into a problem because if you start to see regularity, simple power laws, it is a window into something universal. So in the biology case, either you don't believe in science and you say this is just a bunch of coincidences or you say, the fact that everything you look at has a simple power law that has an exponent that is a multiple of ¼ means that there must be something that is unifying all the data and underlies it, and that's the route that we took. Our claim is that it is networks which scale, but what is nice about scaling is that it has led to a theory of many different things. The question then becomes: 'To what extent can this kind of paradigm be extended to social organisations, in particular corporate structures and urban systems?' And I shall leave it there.

Questioner: If you plot energy versus mass for birds in flight you get a perfect line for a $\log_n/\log_n$ plot and if you plot planes you get a perfect line. Every plane that has ever flown is on the same line except one and that's the one that flew all the way around the world. Have you considered this?

Answer: Yes, I have thought about it. That's how I got into a lot of this work. The logarithm of horsepower verses size and it's a pretty good line and it goes over six orders of magnitude and it's linear and the heart rate is 2/3. One of the great things which I didn't talk about is that life has found a way of using the fractality inherent in hierarchical networks to optimise the amount of energy that it takes in and the 4 is actually 3+1, it's actually three dimensions of Euclidean space and the 1 is coming from the spatial aspects of fractal geometry.

My initial prejudice was only that maybe this would lead to something. The thing that we've learnt from this is, driven by the data and driven by the theory, is that evolution has optimised the infinitude of variables that have controlled how life has evolved. What is extraordinary is that energy is the controlling variable and the great challenge now is the question of how we integrate information such as there is in the genes or the brain with energy. That I see as the major challenge in much of modern biology and contrary to my own prejudices it seems that information follows energy (and I include resources etc in 'energy') rather than information determining energy. Human beings were determined by natural selection and information has followed on behind.