Ecology for a small planet

Inaugural Lecture, Prof RJ (Bob) Scholes (21 June 2016)

I am an ecologist. Ecologists are people who study the interactions between organisms and their environment. So, what they would see in this scene is an adult male *Loxodonta africana,* browsing on *Acacia tortillis* and grazing *Urochloa mossambicensis*. But I usually describe myself as a *Systems* Ecologist: how does that affect the way I view this scene?

Here I am, in my helicopter. From that dizzy height, I can see just the essentials, and few of the details. I see a plant eater, with the unusual attribute of being 'very large'. More importantly, I see relationships, such as between this bulky herbivore and the plants it is eating. This relationship does the herbivore good (shown by the positive sign), but harms the plant (negative sign). Since the product of the interaction signs is overall negative, this would be a self-stabilising or oscillatory relationship. If the herbivore eats too much the feedback will result in it eating less. Let's focus those binoculars a bit more sharply, to see more of the ecosystem. Now I see another set of plant eaters, eating the same and different plants, and their interactions. This new loop is overall positive, leading to either a virtuous or vicious cycle. That makes the entire system is inherently unstable: if you give it a nudge, it can trundle off in a completely unexpected direction. This is a cartoon of how a systems ecologist thinks: now let's look at an important, real-world example from early in my career.

Pastoralists in South Africa for centuries recognised that there are two types of veld: one on which their cattle grew fat all year (sweetveld) and one where they starved in winter (sourveld). The aerial photo is of one of the sweetveld patches within the sourveld at Nylsvley, where Mary and I worked after doing our postdocs. The patches are the result of human habitation seven centuries ago. Range scientists have concocted all kinds of explanation for the phenomenon of sweet- and sourveld, mostly based on the properties of the grasses: sweetveld has sweet species, and sourveld has sour species. But this is a tautology. With my students William Ellery and Andrew Blackmore, we proposed a better explanation. When you find discrete types in nature there is usually a threshold at play, causing the system to 'bifurcate' into alternate states. Ruminants like cattle and antelope have an absolute requirement for 0.5% nitrogen in their forage, below which they cannot digest the grass, no matter how abundant. By examining the sources and leaks of nitrogen in the system, in relation to its ability to capture carbon, we were able to show why sourveld occurs in high-lying, wet, infertile places, and sourveld in lower, warmer, and drier places. It was a property of the *system*, not the *species*!

I have been interested in three main things over my career: biogeochemistry, which is a fancy word for the cycling of elements such as carbon and nitrogen; Earth Observation, how we measure Earth System processes; and how ecosystems behave over time, especially went you push them to breaking point, which is what is happening all around the world in our current era. My colleagues from Europe and North America tell me that is about two-and-a-half things too much! Focus, Bob! But I see an overall connecting tread: how do you take ecological understanding to scale.

In 1978, the Wits honours class took a fieldtrip to Namibia. At the Gobabeb desert research station we were tasked with estimating how much water the narrow green ribbon of *Faidherbia albida* trees were drawing from the Kuiseb riverbed. Some bright spark had suggested defoliating the trees to free up water for the new uranium mine at Rossing. So we measure the transpiration rates of leaves, multiplied by the number of leaves on a tree, trees per river profile and the length of the river. A typical upscaling approach. The answer was so large that we ate it, lest it fall in enemy hands!

What had we done wrong? The main problem was not measurement error, but that we had ignored the interaction terms, something an ecologist should never do! At each level of the scaling hierarchy from stomata to leaf to canopy to landscape, interactions modify the flux, to such an extent that the process driving forces shift as the scale of observation changes: from vapour pressure deficit at tiny scales, through canopy coupling at patch scale to energy budget at landscape scale. Welcome to the Alice-through-the-looking-glass world of ecological scaling. Why is this critically important? As human activity comes to dominate the planet – we call the present time the Anthropocene – we are witnessing ecological crises at greater and greater scales. Historically, ecologists only had the tools to study them at the scale of a small patch. We needed to learn how to scale up, and fast!

The central question of this talk is as follows: Is big Ecology just little ecology with lots of zeroes? Here is a carbon budget I prepared for Nylsvley. The units are in g/m^2. Here is a budget for the whole Earth, the first observationally-constrained one, which I helped to prepare. The difference in area is about 15 orders of magnitude! Is the numerator just a million billion times bigger too? It is really hard not to commit serious errors when you extrapolate that far. This would be a dangerous error, when you consider that this budget underlies the whole issue of climate change.

The simplest way of accumulating ecological quantities to larger scales I call 'paint-by-numbers'. [I need to explain to young colleagues that when I was a child, a common gift was a paint set which came with a numbered drawing. In order to be a Rembandt, all you had to do was paint the numbered tube into the shape with the same number.] You classify a region into discrete ecosystems, assign each a mean value, and sum them up by weighting the values by the areas they represent. Here is the ecosystem map of South Africa, and there is the tiny bit I contributed.

The problem is that ecosystems are not internally homogeneous and seldom have nice discrete boundaries. You can sidestep these problems by representing the ecosystem processes as a spatially-continuous rather than discrete field. That is what I did to recalculate the greenhouse gas emissions resulting from veld fires in Africa, during the SAFARI '94 campaign. Previous estimates, which treated all savannas as one homogeneous class, suggested that about 4 million tonnes of methane annually were emitted from wildfires per year. Taking account of the spatial variation in fuel loads and fire types resulted in a new estimate, eight times lower.

Here is a more recent example of this approach, a calculation I made of the aboveground carbon stock of the whole of South Africa, for purposes of reporting to the UN Framework Convention on Climate Change. This time, the fundamental resolution of the calculation is 1 km^2, made possible by recent improvements in vegetation observation from space. But although we moving to finer and finer resolution, there is a limit to how much better our estimates can be: they still do not take into account the interactions!

The best way to do so is to measure the process at the actual scale at which it occurs. This is a Stommel diagram, a way of showing what characteristic domains in space and time various processes occupy. For instance, the process of photosynthesis can be measured at the leaf scale using a small chamber which you clamp over the leaf [compare this to the potometer we were using just five years previously]. Fifteen years later, we could measure photosynthesis over a whole landscape, using eddy covariance. This is the flux tower at Skukuza, which I established in 2000 – making it the longest-running flux tower in Africa, or in savannas anywhere in the world. Fifteen years later again, we can estimate photosynthesis at the scale of an entire continent, by measuring the changing atmospheric CO_2 concentrations with great precision, such as is done at the Global Atmosphere Watch site near Cape Point, and then applying model inversion to work out where they must have come from.

In 2000, before such approaches became routine, we had the clever idea of using the nearly-closed atmospheric gyre over southern Africa during our dry winter months to measure the accumulation of gases and smoke at regional scale: 'Africa in a test tube'. The SAFARI 2000 campaign revealed many surprises: the emissions from Highveld coal burning did not harmlessly exit over the Indian Ocean, but recirculated over the continent; and much of the soupy pollution during winter was very widespread, not just over industrialised Gauteng, and originated from veld fires in Angola.

The one thing that the World Summit on Sustainable Development in Johannesburg in 2002 agreed on was that we could not manage the planet as an interconnected system if we had only fragmentary information. An international organisation called the Group on Earth Observation was set up with South Africa as co-chair, tasked with developing a 'system of systems' to share global information in nine domains, and fill the gaps where they existed. Because of my experience on the Global Climate Observing System and as chair of the Global Terrestrial Observing system, I was appointed to the four-person design team. Today GEOSS is in its second ten-year phase, and has over a hundred member countries and a similar number of member organisations. Once it was established, I was appointed the chair of its activities related to biodiversity: the Group on Earth Observation Biodiversity Observation Network.

Biological diversity is in rapid decline worldwide: we are at the start of the 'sixth extinction crisis'. The fifth was the collision between Earth and a large comment 56 million years ago, which wiped out the dinosaurs; and there were 4 cataclysms before that. This one seems to be due to us, but is surprisingly hard to provide the global picture. One problem is that 'extinction' is really hard to prove until it is way too late. Another is that biodiversity data is notoriously local, patchy and tricky to scale up. Researchers have proposed literally hundreds of different approaches to measuring biodiversity loss.

A masters student, Oonsie Biggs, and I developed a way of summarising biodiversity loss data which avoided many of these pitfalls. It is called the Biodiversity Intactness Index, and incorporates richness (R), area (A) and land use impacts (I). Normalising by area makes it possible to apply at a huge range of scales, without serious distortions. The concept is now in wide use, in several forms, around the world.

But you can't capture a multidimensional concept like 'biodiversity' with one number. On the other hand, hundreds of contradictory indicators is a recipe for confusion and inaction. The current work is on what is the minimal, robust set of information for multiple purposes: the concept of 'essential variables', highly informed by system analysis and decision support theory.

Having the evidence is not enough to ensure that sound and timely decisions are made. Can you trust the data? How do you deal with contrarian voices, whether they be on the link between HIV and AIDS or between human actions and climate change? You need a better-functioning interface between science and policymaking.

In the 1980s the world faced a nasty shock – the emergence of a hole in the ozone layer over Antarctica, which as letting in harmful UV radiation. An international process of evaluating the technical information was developed, which led to the Montreal Protocol on Ozone-Depleting Substances, and a quarter of a century later, the hole is on its way to being fixed. Climate change is a similarly important, contentious and complex problem; so the Intergovernmental Panel on Climate Change was set up to periodically evaluate the evidence in a comprehensive, rigorous and transparent way. I have served on the third, fourth and fifth IPCC assessments, in various roles. The IPCC, which won a Nobel prize in 2007, has been credited for moving the climate change political agreements forward at a time when other environmental treaties seem to be stuck; so in the year 2000 a group of ecologists launched the Millennium Ecosystem Assessment, broadly patterned on the IPCC. I was co-Chair of the global status and trends working group. The Intergovernmental Platform on Biodiversity and Ecosystem Services arose from that legacy, and I co-chair its current assessment of land degradation and restoration. Inspired by the power of the assessment approach, I convinced the South African Department of Environmental Affairs to apply it to the highly emotive issue of elephant culling; and the process led to a much better policy outcome. Most recently, we applying assessment techniques to the equally polarised question of fracking for shale gas in the Karoo. The global observation systems have in many ways paralleled these developments.

Which brings me to my interest in the dynamics of complex systems. Savannas, an ecosystem type which covers a third of South Africa and two-thirds of Africa, have been a key test case for ecosystem stability theory. Most major global ecosystems are dominated by a single plant life-form: forests by tree, grasslands by grasses. But savannas have both, apparently persistently, in defiance of competition theory. They occupy about a sixth of the world, so this clearly is not some minor anomaly! Savannas are subject to abrupt shifts to higher tree density, phenomenon called bush encroachment. The classical explanation for the coexistence of trees and grass, alas still alive, is the Walter hypothesis, which states that grasses use shallow water and trees use deep water, so they are able to coexist. These isolines represent the conditions for which $dT/dt=0$ amd $dG/dt=0$. At the point where they interact, the system as a whole is at equilibrium. My PhD, under Prof Brian Walker, showed this hypothesis to be both based on false assumptions and insufficient to explain coexistence. A conceptual model based on asymmetric competition between trees and grass, and the effect of fire on trees has two stable states: a wobbly few-tree state, and a robust many tree state. A seminal paper by myself and Steve Archer in 1997 showed that a 'minimal' savanna model had to take into account at least these things, and triggered a cottage industry of publications on this topic, some of it my own work. Putting together even a relatively simple set of governing equations for such savannas leads to an exotic range of behaviors, from unstable to stable to bistable. To briefly summarise a big topic: even moderately complex ecosystems have multiple stable states. Finding hypotheses for coexistence is not hard, but identifying the true mechanism is! Equilibrium ideas and disturbance ideas, like trees and grass, are not mutually exclusive!

Recently I have carried this kind of theory to the existential problem of the seaworthiness of the global spaceship with 10 billion people on board. Here is an example, sketching the 'safe space' between the Scilla of food security and the Charybdis of climate change. Because these two defining challenges of the 21st century are not independent of one another, the 'frontiers' defining the tolerable limits of one as a function of the other, curve. The size of the safe space tells us something about resilience. We are currently outside the safe space (nearly a billion people go hungry). By mid-century, unless we shift the curves, we are still unable to feed the hungry, but now we are perilously close to the climate tolerance edge. We can drop this line of how much food we need to grow per person, by being less wasteful. We can shift this curve of the impact of agriculture on the climate by making developing technologies which grow more food for less greenhouse emissions, and we can shift this curve of the effect of climate on agriculture adapting our agriculture to changing circumstances. Some of my current work is in this area.

What strikes me in reviewing my thirty year career, is how it mirrors what we know about change in complex systems. There are many feasible outcomes, separated by quite unpredictable forks in the road. I had chance encounters with people and problems along the way, which bumped me onto different, unintended, less-travelled paths. To young people, my advice is to be born to a supportive family, and then make the best of whatever opportunities arise, guided by your curiosity and instinct for the right thing to do. I have been inspired by the example of Basil Schonland, one of the most distinguished professors ever to grace this university, and who founded the CSIR, where I spent a productive quarter of a century before returning to Wits to teach. As a brilliant graduate of the world's top physics research lab in Cambridge, he wondered how to keep his edge in far-away Africa. Gazing out of his office window, he saw a Highveld summer lightning storm, and realised that it provided an opportunity his Europe-based peers would seldom witness. The demands of working in South Africa are not necessarily an impediment, they can be an opportunity. We are forced to be multi-skilled. The same international colleagues who berate me for not specialising sufficiently are amazed by how much I know about so many things. I have made a niche out of seeing the connections between things; and that has made all the difference.