

# An ODD description of Complex Ecosystem Testbed

---

*Bruce Edmonds  
Centre for Policy Modelling  
Manchester Metropolitan University*

## 1 Overview

### 1.1 Purpose

This model is designed to be a test-bed – an artificial ecology – into which agents representing humans and their society can be introduced, in order to explore how the arrival of such humans can affect the surrounding ecosystem. In reality it is not clear how this arrival would affect a particular ecosystem – in the past, at some times it seems that a balance between humans and the rest of the ecosystem has been established, but at other times the arrival of humans can only be described as catastrophic (Diamond 2004). This test-bed needs to be sufficiently complex and dynamic so that this forms a robust test of the impact of human society. In particular it should:

- Have an energy economy, so that predators (herbivores, predators etc.) dynamically impact upon each other;
- Allow the food-chain to dynamically change and develop;
- Explicitly represent space so as to allow the emergence and spread of new species (invasive and mutations);
- Not predetermine the species composition but allow this to co-emerge with other species and the environment.

When the human agents are added into this model it is important that they:

- Are embedded within and dependent upon the energy within the existent food chains, so need to predate upon plants and other entities
- Are embedded spatially so that they inhabit and exploit particular locations where they are and have access to
- Take action within and parallel to the developing ecological processes of change so that social change happens with ecological change.

Thus it is envisaged that the agents that represent humans will be similar to those representing other entities (people are biological entities after all), but have some different characteristics and additional processes not available to other entities. In particular, the fundamental characteristics of human agents will not be limited to genetic change, but be amenable to social processes (such as learning or developing new skills). The human agents are not described here, only the test bed in anticipation to them. For a brief exploration of the impact of human agents see (Edmonds 2014).

Multi-agent simulation could apply its expertise in terms of specifying and exploring the cognitive/social abilities of agents with respect to such a test bed, and start to tease out the complex and often counter-intuitive effects of such abilities. Knowledge about this could play a real part in helping us understand our own, fragile and complex, relationship with the ecologies we inhabit and exploit. It is time to show that multi-agent based simulation can deliver tangible results that can significantly contribute to the challenge of surviving on our planet and living with a variety of other species. The Challenge to all in the MABS community is to discover how human-like agents might be organised and so: (a) understand some of the possible consequences of various survival strategies and (b) help select which of the many possible strategies are worth trying out for real.

The simulation here can be seen as an ancestor of (Norling et al. 2008). It uses a variety of the interaction matrix described in (Caldarelli et al 1998). It relates to the ecological model described in (Edmonds 2013) but has only one type of basic resource.

## 1.2 Entities, State Variables, Scales

In this, entities (plants, herbivores and predators), are represented as individual objects. They inhabit one of a number of patches arranged in a 2D pattern that makes up the world. Each patch is well mixed so that interactions within that patch are random, but there is a probability that each individual can migrate to one of the four neighbouring patches each tick. The world is wrapped vertically and horizontally. There are two basic entities in this model:


1. **Patches** – these are very simple entities representing the local location and environment. Entities exist only at a patch. The patches are arranged as a 2D grid. The patches have their own characteristics in terms of a short bit string, which is used to determine which (if any) of the entities can extract energy from them. Each simulation tick each patch is provided with an amount of energy, which some of the entities (i.e. “plants”) can share. Patches do not store energy from past ticks. Some patches correspond to “deserts” (those with a zero bit string), which are spaces where entities can exist but no entities can gain energy from. Optionally some patches can be converted to “barriers” into which no entities can move or exist.
2. **Entities** – these represent plants and animals. Each entity is located on a patch, but entities are “well mixed” within each patch (one can not hide from another within a patch and interaction between entities happen at random within the patch). Entities have a number of characteristics, the most important are the its store of energy and its genome. The genome is a bit string that determines which kinds of patches it can extract energy from (plant-like) or which other kinds of entities it might eat (predator-like). For details as to how this is determined see the subsection on “Interaction” below. Energy is decremented each tick by the “life-tax”, if their energy drops to 0 it dies, if it reaches a higher value (given by the “reproduce-level” parameter) the entity can reproduce once. Offspring have similar characteristics to the parent, but (with a small probability: “mut-prob”) mutation can occur to the genome, and it has an initial energy

(given by the “**init-energy**” parameter). This initial energy is subtracted from that of the parent. Other attributes include: the trophic **level** it has reached (by eating those of a lower trophic level), **age** in simulation ticks, **species-id** (for speed), and the **num-eaten** (number of entities eaten).

3. **Genes/characteristics** – in a way these are entities. Each patch or entity has them. They are a binary list of fixed length (patch characteristics are padded out with zeros to make them the same length though only the first few are significant). Entities pass these on to their offspring with a small probability of mutation. These binary lists determine which species can predate on which others and which species can extract energy from the environment.

### 1.3 Process Overview, Scheduling

This is a synchronous individual-based simulation. There is a basic energy economy; so that energy is injected into the world, divided equally between patches, each tick, which drives the ecology. Whether an individual can extract energy from a patch or predate upon another is determined by both of their bit-strings and a fixed random interaction matrix, described below. The world is illustrated in Figure 1.



**Figure 1. An illustration of the grid of 2D patches, in this case an 8x8 grid. Plants are small stars, herbivores and higher predators are circles (the more they have eaten the bigger they are up to a maximum size,). Different colours indicate different species but not all species are visually distinguishable.**

Each tick in the simulation:

1. **Input energy.** A fixed amount of energy is added to the model, equally divided between all the patches.
2. **Death.** A life tax is subtracted from all individuals, if their total energy is less than zero it is removed from the simulation. Their age is incremented
3. **Initial seeding.** (In the initial phase), until a viable population<sup>1</sup> is established, a random new individual is introduced with a given probability.

---

<sup>1</sup> This is when a non-zero population has lasted for a certain length of time.

4. **Energy extraction from patch.** The energy stored in a patch is divided among the individuals on that patch that have a positive score when its bit-string is evaluated, using the interaction as described above (against the patch's bit-string) in proportion to its relative fitness, at the simulation's efficiency rate.
5. **Predation.** In a random sequence, each individual is randomly paired with a number of others on the patch, given by the parameter "eating tries". If it has a positive dominance score against the other, the other is removed from the simulation and the individual gains a fixed proportion of its energy, given by the "efficiency" parameter. Individuals are not allowed to predate upon members of their own species.
6. **Maximum Store.** Individuals can only retain so much energy, so any above the maximum level set is discarded.
7. **Birth.** If an individual has a level of energy above that determined by the "reproduce-level" parameter, it gives birth to a new entity with the same bit-string as itself, with a probability of mutation. The new entity has an energy level of 1, which is subtracted from the energy of the parent.
8. **Migration.** With a probability determined by the "migration" parameter, the individual is moved to one of the neighbouring 4 patches (the world being "wrapped" top and bottom).
9. **Statistics.** Various statistics concerning the model are calculated.

The simulation ends after a given number of ticks if "**max-ticks**" is set to a non-zero number.

## 2 Design Concepts

### 2.1 Basic principles

There are no grand theories underlying this model, rather a number of common-sense principles. However these are worth listing as explicit assumptions.

- Entities are individuals existing in locations, they can interact (give birth, predate) only *within* such locations.
- Entities have a genome which determines their characteristics (including who can eat whom, and who can extract energy from which patch). This genome is passed to their offspring with possible point mutation.
- Free energy rains down into the system (as does sunlight), this is limited for any particular location and time (plants compete for it).
- All entities need energy to live, and use it up whilst living.
- Entities can gain energy in two ways: (a) from the environment in that location ("photosynthesising") or (b) predating on another entity.
- Offspring gain their initial energy from their parent.

The one, theory-like mechanism that the model uses is the method described below in 2.8 Interaction that uses a fixed matrix to determine which entity can gain energy (eat) from which other.

The following are *not* assumed, as in built into the model.

- That any particular species or set of species will exist at any trophic level.

- That species will form local food webs at all.
- That the predate relation is anti-symmetric or transitive.
- That anything at all will survive.
- That any equilibrium or “balance of nature” will be reached.
- That a local “balance” in the food web will be reached before it is disrupted by invasive entities or new mutations.

There are many ways in which the model is simplistic, notably:

- The genome does not determine many characteristics of entities, such as when it gives birth, its energy usage, when and how far it moves etc.
- The genome is a simple bit string.
- The translation from genome to phenotypic properties is simplistic.
- There is no sex in the model, only asexual reproduction.
- The locations are a 2D grid of squares.
- Entities are well-mixed within a location, thus potential prey can not hide from a predator within a location (but might avoid them probabilistically if part of a crowd of others).
- Entities do not learn or adapt at all.
- Unlike (Edmonds 2013) there is only one scarce resource – energy.

Each of these suggests possible model extensions/improvements.

## 2.2 Emergence

Notable features that emerge from this model include that:

- It can produce ecologies with plausible food webs.
- The fundamental interactions in the model, those that constitute the food chain, are emergent and can continually change in both time and space.
- The model creates endogenous shocks on its own, with new species appearing to sometimes-catastrophic effect on the existing food chains, affecting them radically in terms of their constituent species, their relative abundances and even the predation links.
- Mutation and migration happen in parallel, so that new species often appear before previous species have been completely spread over the space – states that could be interpreted as being in “equilibrium” are rarely observed, unless the ecology is non-viable or is dominated by a single species.
- The system is highly adaptive in the sense that the distribution of species and their composition evolves as the environment or other species change. However this adaption is not always successful and it is quite possible to wipe out all life, for example if a “super predator” arises.
- As you will see in the indicative results below (section 4) there are a number of different kinds of world state that can arise and persist for a while (though often no particular such state lasts forever as is shown in Figure 13).

## 2.3 Adaptation

Individual entities or patches do not adapt at all. This would be in stark contrast to human agents if added. However, the whole system adapts as an evolutionary ecology.

## 2.4 Objectives

Entities do not have objectives nor do they make any decisions. This might be in stark contrast to human agents if added.

## 2.5 Learning

Entities do not learn at all. This would be in stark contrast to human agents if added.

## 2.6 Prediction

Entities do not anticipate or plan at all. This might be in stark contrast to human agents if added.

## 2.7 Sensing

Entities can “sense” only the other entities in their patch, in the sense of randomly interacting with these (to try and eat them). Agents can not move into “barrier” locations, so they can tell if adjacent patches are barriers or not.

## 2.8 Interaction

Entities only interact in three ways: (a) extracting energy from their environment (b) interacting with another random entity with a view to eating them (c) giving birth at a location and thus creating a new entity.

Key to understanding how processes (a) and (b) are determined is the method by which domination is determined. This method is adapted from that in (Caldarelli et al. 1998). A random interaction matrix with the dimensions of the length of individuals’ bit-strings is generated at the start of a simulation. It is filled with normally distributed random floating-point numbers (mean 0, standard deviation  $1/3$ )<sup>2</sup>. This interaction matrix determines which entity can eat another entity in the following manner (see Figure 2 for an illustration):

1. The non-zero bits of the predator select the columns of the matrix, the non-zero bits of the potential prey select the rows.
2. The intersection of the selected rows and columns determine a set of numbers, these are summed.
3. If the sum is greater than zero the predator can eat the prey, in which case the prey dies and the predator gains a percentage of its energy value the rest is lost.

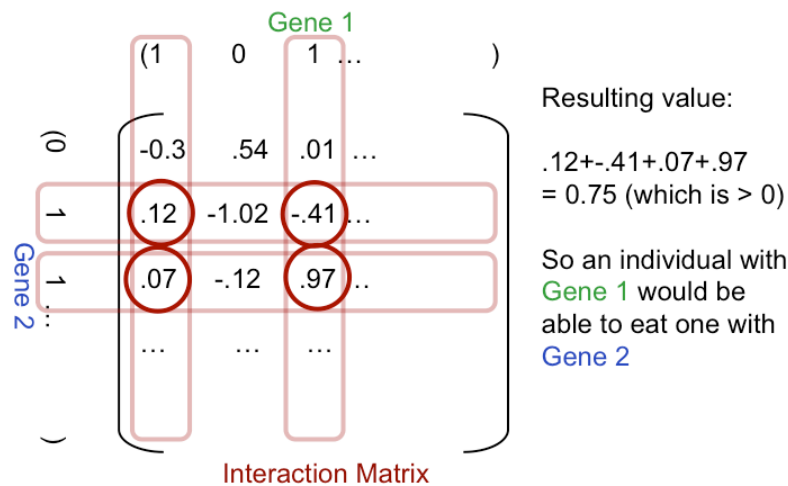
Essentially the same process is used to determine which entities can extract energy directly from the environment, except that the part of the prey is taken by

---

<sup>2</sup> This could be forced to be anti-symmetric but in the runs here it is not. Forcing the matrix to be anti-symmetric slightly simplifies the possible interactions but (in all the variants I have observed) does not substantially change them.



the patch with its bit string (padded with zeros to reach the appropriate length). In this case only those with scores greater than zero get any of the patch's energy. The patch's energy is divided between all qualifying individuals in proportion to their score against the patch. This scheme has the consequence that no individuals can extract energy from a patch with a bit-string of all zeros. Thus all the simulations reported below will have some patches that act as "deserts", that is patches where individuals cannot extract any energy from the environment (although they may pass through the patch using previously stored energy or predate upon other individuals there).



**Figure 2. The use of the interaction matrix to determine predation as well as energy extraction from a patch to give its relative fitness.**

This interaction scheme allows complex food webs to be evolved, for example via a genetic "arms-race" between predator species and prey species, since it allows for adaption with respect to another specific species. In other words fitness is not an absolute number but relative to the environment, if it extracts energy from this, or another species. (Caldarelli et al. 1998) showed that this kind of scheme can be used to evolve complex ecologies with plausible characteristics including food-webs with similar network characteristics to observed food-webs.

At the start of the simulation, the random interaction matrix is generated. Each patch is allocated a random bit-string with the given number of bits, padded out with zeros to make it the same length as individuals' bit-strings. The environmental complexity (**num-env-chars**) is the number of significant characteristics that patches have – the number of bits in its bit-string. Bit strings of length 2 allow for 4 types of patch, of length 3 8 types etc.<sup>3</sup>.

## 2.9 Stochasticity

There are various sources of stochasticity in the model.

- The initial characteristics and positioning of patches are random.
- The genome of initial "seed" entities is random.

<sup>3</sup> Due to the fact that a patch with a zero-length bit string would be equivalent to having a bit-string of all zeros and hence a desert, only environmental complexities starting at length 1 or above are explored.

- Which individuals are paired with which others within a patch for possible predation is random.
- When mutation happens at birth and which bits are affected is random.
- When movement occurs is probabilistic and where they move to is random.

## 2.10 Collectives

There are no real collectives in the model. However there are two ways in which entities could be considered to form collectives. Firstly, when located on the same patch, and hence could predate upon each other or compete for environmental energy there. Secondly, as member of the same “species” – that is entities with the same genome. However entities have no way of knowing which others are the same as them, and may predate upon them (unless this is prevented by turning off the “**allow-cannibals?**” switch).

## 2.11 Observation

The world is divided into patches the background colouring indicates its characteristics. Individuals are small stars (if never eaten another individual) or circles (getting bigger for each other they eat) if they have. Colouring indicates species. The position of entities within patches is not significant and randomised.

The following graphs are displayed.

- Tropic Level Distribution is a histogram of the trophic levels. This is of individuals not species, so if an individual has eaten another that has eaten a 'herbivore' then it has a trophic level of at least 2.
- A graph of the number of entities at each trophic level, shown on a logarithmic scale.
- Graphs of the population size and number of different genomes (roughly species) in existence.
- A line histogram of relative species abundance, with species sorted from most abundant (left) to least (right), on a logarithmic phase.
- A histogram of different abundances of species, with the log of abundance determining the bars and how many species are at that abundance on the vertical.
- A scatter gram (only calculated on request) of the log of the number of species found against the log of the area species are sampled from. This is a sample, whose size is given by **hist-sample-size**, and calculated when the button is pressed.

An number of statistics are calculated every simulation tick. These include the following.

- **Eats** – how many are eaten
- **Deaths** – total systematic deaths (eaten, starve, got too old)
- **Births** – number of new individuals born due to parent reaching required energy
- **Pop** – number of live individuals
- **Species** – how many distinct genomes there are in living individuals



- **Starve** – how many died of lack of energy
- **Old** – how many died due to reaching max-age
- **Migr** – how many migrated to another patch
- **Rand Die** – how many died randomly (due to rand-death-prob)
- **secs/tick** – how long current simulation tick took to compute (in seconds)
- **L0** – proportion of individuals that have not eaten anyone ever (herbivores)
- **L1** – proportion of individuals that have ever eaten a herbivore
- **L2** – proportion of individuals that have ever eaten a carnivore

The following are only calculated when press Calc button or at end of simulation. All are about diversity measured as the average hamming distance between all pairs of distinct individuals, as in (de Aguiar et al 2009). To make this more feasible this estimated using random sample of possible pairs of size given by parameter “**sample-size**”). These include the following:

- **pi-t** – the overall diversity
- **pi-d** – diversity within patches
- **F-st** =  $(\text{pi-t} - \text{pi-d}) / \text{pi-t}$  – how much diversity is inter-patch
- **pi-env** – diversity within all patches with same characteristics
- **F-env** =  $(\text{pi-t} - \text{pi-env}) / \text{pi-t}$  – how much diversity is of an inter-patch type
- **F-denv** =  $(\text{pi-env} - \text{pi-d}) / \text{pi-env}$  – how much within patch diversity is due to variation between patches with same characteristics

### 3 Details

#### 3.1 Initialization

A random matrix with the dimensions of the length of individuals’ bit-strings is generated at the start of a simulation. It is filled with normally distributed random floating-point numbers (mean 0, standard deviation  $1/3$ )<sup>4</sup>. This is then fixed for the duration of the simulation.

The simulation starts with no individuals, but only empty patches. The patches are each given a random set of characteristics (limited by the number of bits determined by the “**num-env-chars**” parameter). Optionally the patches are then swapped around to implement a given level of average proportion of like patches (given by “**target-prop-n**”). Then each patch might be turned into a “barrier” patch (one where entities can not move to) with a probability given by parameter “**prop-barriers**”.

Until a viable population of entities is established single random “seed” entities are introduced into a random patch (with a probability of “**init-new-species-prob**” each tick). A viable population is deemed when there are some entities alive for a consecutive period of time greater than a minimum.

---

<sup>4</sup> This could be forced to be anti-symmetric but in the runs here it is not. Forcing the matrix to be anti-symmetric slightly simplifies the possible interactions but (in all the variants I have observed) does not substantially change them.

## 3.2 Input Data

The simulation does not use any input data.

## 3.3 Submodels

### 3.3.1 Dominance Calculation

Although this is described in the section above (2.8 Interaction) it is formally described here. If the potential predator bit-string is represented by vector  $\mathbf{p} = (p_1, p_2, \dots, p_n)$ , the potential prey (or patch) bit-string is represented by vector  $\mathbf{y} = (y_1, y_2, \dots, y_n)$  and the interaction matrix by  $\mathbf{M}$  with entries  $m_{i,j}$  then  $s = \sum_{i:p_i=1, j:y_j=1} m_{i,j}$  is calculated. In the case of a potential predator-prey interaction predation (between individuals that are randomly paired within a patch) occurs if  $s > 0$  and in the case of individual patch will receive a share in the energy of the patch in proportion to  $s$  if  $s > 0$  (along side all other such individuals).

### 3.3.2 Parameters

The parameters of the basic ecological model are as follows, with their default values in square brackets.

- **gene-size:** [100] the number of bits in an individual's gene
- **num-env-chars:** [3] the number of, effective, bits in the characteristics of a patch, functions similar to an individual's gene,
- **migration-prob:** [0.01] the probability that any individual will move to another patch each time click
- **mut-prob:** [0.01] the probability that a newly born individual will have its gene mutated
- **food-rate:** [500] how much energy is put into the world each time click, evenly divided among patches,
- **efficiency:** [0.9] what proportion of energy of something eaten goes to predator, or to herbivore from patch,
- **reproduce-level:** [3] if an individual's energy gets to this point it gives birth once, new-born's energy being then subtracted from it,
- **eating-tries:** [2] each tick each individual tries to eat this number of others on the same patch, but this only happens if they dominate them via the interaction matrix
- **max-age:** [80] if  $> 0$  individuals die when they reach this age in simulation ticks, otherwise no age ceiling,
- **max-store:** [20] if  $> 0$  this is the upper bound on what energy individuals can accumulate, rest is lost to system,
- **max-time:** [1000] if  $> 0$  the time at which the simulation is halted

The following are included for completeness but are not often varied.

- **life-tax:** [0.25] how much energy subtracted from each individual each time click, dies if energy is 0 or below,
- **init-energy:** [1] the energy of a new born, this is subtracted from the parent at birth,

- **init-new-species-prob:** [0.01] probability of a new individual with a random genome being introduced each time click
- **stop-new-species-once-established?:** [true] if set stops new individuals with new genomes being introduced into the simulation once a viable population is established
- **initial-species-variety:** [0] 0/1: if 0 simulation starts with a single individual, if this is not successful relies on a new individual being introduced via init-new-species-prob, , if 1 starts with a full population with random genes
- **rand-death-prob:** [0.02] the probability an individual randomly dies each tick
- **anti-sym-mat?:** [false] forces the interaction matrix generated at start to be anti-symmetric
- **migrate-near?:** [true] if true then individuals migrate, if they do, to a neighbouring patch if false to a random other patch
- **allow-cannibals?:** [false] if true individuals of the same species can eat each other, if not can only eat those of another species
- **neutral?:** [false] if true then all individuals on a patch get the same amount of energy from the patch regardless of their bit-string, otherwise they have to dominate the patch via the interaction matrix to extract energy.

## 4 Simulation Outcomes

A brief summary of some of the outcomes/behaviour of this model are described.


### 4.1 Archetypical resultant ecological states

There are four different ecological kinds of outcome observed in this model: 1, a non-viable outcome where nothing thrives or reproduces, defined as being fewer than 10 individuals in the whole space, , 2, a situation where one, or two, plant species dominate, 3, a plant ecology, not case 1 or 2, with no herbivores or higher predators and, 4, a mixed ecology like case 3 but with herbivores and higher predators. In practice if there are fewer than 10 individuals there are usually one or no individuals within a few simulation ticks, and if there are either one or two species or many. Thus although the division is somewhat arbitrary, it very clearly distinguishes four cases between observed simulation trajectories. Furthermore these four kinds tend to persist for many simulation ticks so that each can be meaningfully identified. These are each described with outcomes from a typical run below. Many of the later results will be in terms of the occurrences of each of these four types.

Each description is accompanied with three figures (left) is a visualisation of the patches and individuals, the colours of the background patches indicates its bit-string, plants are indicated by a small star, individuals higher up the food-chain are indicated by a circle whose size is related to how many other individuals they have eaten; (centre) is a graph of the number of species over time; and (right) is a graph showing the number of individuals of each trophic level on a shifted log scale.

#### 4.1.1 Non-Viable Ecology


Here species do not manage to extract any energy from the environment, so any introduced species quickly starve with no reproduction. There is only ever one individual since when this one dies a new random one is introduced into the simulation.



**Figure 3. Typical Non-Viable Ecology (left) the world state (left) Number of Species (right) Log, 1 + Number of Individuals at each trophic level,**

#### 4.1.2 Dominant Species Ecology


Here one, or a few, species dominate. The dominant species is both a plant and a predator, eating any new other species that appear. Thus, occasionally individuals are classified as belonging to a higher order trophic level, although no other species manages to achieve a long-term survival. Very occasionally two or three dominant species occur, each destroying the others that wander into the patches they dominate.



**Figure 4. Typical Dominant Species Ecology (left) the world state (left) Number of Species (right) Log, 1 + Number of Individuals at each trophic level,**

#### 4.1.3 Rich Plant Ecology


In this case a rich plant ecology develops where many different species compete as to their efficiency in extracting energy from the different kinds of patch, and are resistant to potential herbivores who, if introduced, simply starve. In terms of the number of individuals this state often produces the greatest number of species and the highest population (in terms of number of individuals). Species only gradually replace older ones as they marginally out-compete them in terms of energy extraction.



**Figure 5. Typical Herbivore Ecology (left) the world state (left) Number of Species (right) Log, 1 + Number of Individuals at each trophic level,**

#### 4.1.4 Mixed Ecology


In the last case, successful herbivores and higher predators evolve to produce a highly dynamic ecology. There is a continual “arms race” both in terms of bit-string evolution as well as over the space of patches. There are typically far fewer species than in the rich plant ecologies since many plant species are wiped out. This typically results in a power law in numbers of individuals at each trophic level with an order of magnitude between the prevalence of each layer. Here you get a more constant replacement of older species as found in (Drossel et al 2001).



**Figure 6. Typical Mixed Ecology. (left) the world state (left) Number of Species (right) Log, 1 + Number of Individuals at each trophic level,**

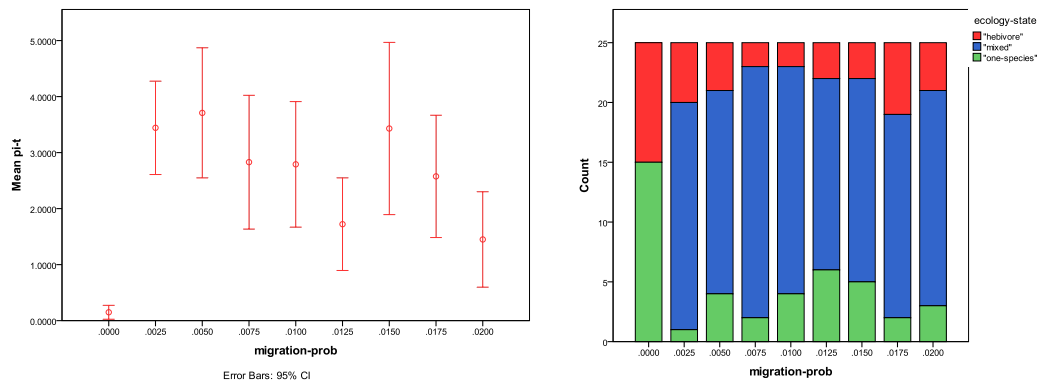
#### 4.2 Sensitivity Analysis of Basic Ecological Model

Each individual run of this simulation can give very different results starting from the same parameter values. Partly this depends on the interaction matrix generated and partly on the happenstance of mutation and movement within the world. Summaries of such runs for different parameter values may thus be misleading as averaging may give a false picture of the collection of trajectories. Thus for each parameter I show both the average effect on diversity over all the runs, the charts on the left below (with a 95% spread) but also a count of how many runs end up in each of the four states above (the bar-chart on the right below). For each of these there were 25 runs of 1000 time ticks for each parameter value.




**Figure 7. Effect of Energy Input on (left) diversity and (right) ecological type**  
**red=plant, blue=mixed, purple=single species, green=non-viable**

The more energy that is put into the system then, generally, the greater the diversity that results (Figure 7), however the response to more energy is non-linear as larger populations support more predators, which has the effect of reducing the diversity. Generally, the higher the input energy, the less frequent does a pure plant ecology result.




**Figure 8. Effect of Migration Rate on (left) diversity and (right) ecological type**  
**red=plant, blue=mixed, purple=single species, green=non-viable**

A zero rate of migration means that each patch is isolated, so this severely restricts the diversity and ensures small populations. Above that the higher the migration rate, the lower the diversity since patches act less like semi-isolate demes and more like a total well-mixed population, with all species competing against all (Figure 8).




**Figure 9. Effect of Mutation Rate on (left) diversity and (right) ecological type**  
**red=plant, blue=mixed, purple=single species, green=non-viable**

A zero mutation rate means that nothing can evolve, so that there is, at most, one species. Above that, a higher mutation rate implies a higher diversity and fewer cases of a “single species” ecology (Figure 9).



**Figure 10. Effect of Gene Size on (left) diversity and (right) ecological type**  
**red**=plant, **blue**=mixed, **purple**=single species, **green**=non-viable


Longer bit-strings enable a greater diversity to develop, however the space of possibilities is so great for sizes above 40 that this is nowhere near explored and hence does not limit the growth of complexity at these scales of space and time, Figure 10.



**Figure 11. Effect of Environmental Diversity on (left) diversity and (right) ecological type**  
**red**=plant, **blue**=mixed, **purple**=single species, **green**=non-viable

Greater environmental diversity seems to decrease the occurrence of rich plant ecologies, Figure 11.





**Figure 12. Effect of Efficiency on (left) diversity and (right) ecological type**  
**red=plant, blue=mixed, purple=single species, green=non-viable**

Higher efficiencies mean that ecologies with higher trophic levels occur more often, since the higher levels can access more energy, Figure 12.

**Figure 13. An illustration of the number of species and the types over longer runs,**  
**red=plant, blue=mixed**

During longer runs, the simulation is often observed to change state between the types of ecology listed above. Figure 13 shows 5 independent runs of the model over 10,000 simulation ticks with a lower mutation rate, 0.001. The first two runs are dominated by plant ecologies with occasional periods of higher trophic levels appearing, the second two are the reverse and remain mixed ecologies for the majority of the time, though mixed ecologies are not as stable as plant ecologies. The last is dominated by a single species and remains so with very short-lived appearance of herbivores. These longer runs illustrate the long-term robustness of the four types of ecology in this model. In the top two runs illustrated in Figure 13 a rich plant ecology dominates, sustaining large numbers of individuals, with occasional short-lived intrusions of herbivores. In the next two runs a mixed ecology dominates with occasional periods when herbivores disappear. The bottom run shows the stability of the situation when one species dominates, effectively preventing any others appearing.

## 5 References

- Caldarelli, G., Higgs, P.G. and McKane, A., 1998. Modelling Coevolution in Multispecies Communities. *Journal of Theoretical Biology*, 193, 345-358.
- de Aguiar, M.A.M., Baranger, M., Baptestini, E.M., Kaufman, L. and Bar-Yam, Y., 2009. Global patterns of speciation and diversity, *Nature*, 460, 384-387.
- Diamond, J., 2004. *Collapse: How Societies Choose to Fail or Succeed*, Viking.
- Edmonds, B. (2013) Multi-Patch Cooperative Specialists With Tags Can Resist Strong Cheaters. In Rekdalsbakken, W., Bye, R.T. and Zhang, H. (eds), *Proceedings of the 27th European Conference on Modelling and Simulation (ECMS 2013)*, May 2013, Alesund, Norway. European Council for Modelling and Simulation, 900-906.

Norling, E., Powell, C and Edmonds, B., 2008. Cross-Disciplinary Views on Modelling Complex Systems. In. David, N. & Sichman, J.S., Eds., Multi-Agent-Based Simulation IX, Springer, Lecture Notes in Artificial Intelligence, 5269, 183-194.