Simulation and Analysis in Agent-Based Modelling of Land Use Change

Nicholas M. Gotts, J. Gary Polhill and William J. Adam

October 30, 2003

Abstract

The FEARLUS modelling system is being developed to enhance understanding of rural land use change, primarily at a regional scale, using spatially explicit agent-based models. Interactions between the owners or managers of neighbouring Land Parcels are central to the phenomena of interest, and to the models developed. We have concentrated particularly on the dynamics of competition between land managers selecting land uses in different ways.

The physical environment is represented in FEARLUS as a grid of cells, giving FEARLUS a resemblance to a cellular automaton. Cellular automata encourage researchers to focus on ways in which relatively simple local interactions can generate global dynamical phenomena in spatially extended systems. Computational models based on cellular automata are increasingly used in research on land use change, but standard cellular automata may not be an adequate basis for modelling entities of multiple types and spatial scales, or forms of interaction between entities which vary in spatial range over time; typical FEARLUS models differ from a standard cellular automaton in a number of ways. However, we believe that when a simulation model produces a phenomenon of interest, researchers should search for similar but simpler models that also produce this phenomenon, in order to identify the mechanisms responsible. We also believe that analytical techniques can usefully be combined with simulation. In this spirit, the paper reports analytical work and simulation experiments comparing the models which are currently at the focus of our work, with variants closer to the standard cellular automaton framework.

1 Introduction

This paper is about simple models of complex systems: specifically, spatially explicit agent-based social simulation of processes underlying land use change, in which both space and human agency are represented in rather simple and abstract ways. We argue by example that the analysis of simple formal models, and simulation studies on more complex ones, should not be regarded as alternative and even opposed approaches to the formal study of social systems, as sometimes appears to be the case (Binmore 1998, Chattoe 1996), but as complementary. Similarly, work with highly constrained models — such as those obeying all or most of the rules for standard cellular automata (CA) — and related models in which some of these rules are relaxed, can and should be combined. Nor should research necessarily begin with analysis, and resort to simulation only when analysis fails: without the stimulus of phenomena discovered through simulation experiments, it may not be obvious what it is most useful to analyse.

The work described forms part of the FEARLUS project[1], which concerns rural land use change, but some of the issues discussed are relevant across a broad range of contexts involving systems of interacting, territory-holding agents, making decisions about how to use the territory they hold — see Cioffi-Revilla and Gotts (2003).

FEARLUS is aimed at using spatially explicit agent-based simulation modeling to increase understanding of the processes underlying land use change, particularly at the regional scale and

[1]FEARLUS stands for 'Framework for Evaluation and Assessment of Regional Land Use Scenarios'.

in the medium to long term. Its main motivation is the difficulty encountered in forecasting land use change (on the basis of biophysical properties of the land and available economic returns), if land managers are regarded as unlimited in computational capabilities and driven purely by profit maximisation, as is generally the case in earlier approaches to rural land use change (Benson 1995, Parry 1996). There is a growing current of opinion within the land use research community that to model the drivers of land use change successfully:

> Simulation of decisions by and competition between multiple actors and land managers is required.

(Veldkamp and Lambin 2001, p.2). Those making land use decisions may be influenced in various ways by their neighbours (as well as wider social influences): the most obvious include imitation based on examples of the success of innovative land uses or techniques (and conversely, avoidance of innovations seen to fail). Imitation is one means of economising on computational resources, and/or compensating for an absence of knowledge, and is known to be one way in which land managers choose what to do (Pomp and Burger 1995).

The grid of locally interacting *Land Parcels*[2] within FEARLUS models make them strongly reminiscent of CA, which are increasingly used in research on land use change (White, Engelen and Uljee 1997). A standard CA consists of cells organised in some regular arrangement such as an n-dimensional grid. Each cell may be in one of a finite number of states, and the state of a cell at time $t+1$ depends only on the states of a fixed and finite set of neighbour cells at time t. The transition rule governing changes in cell state may be deterministic or stochastic, but is the same for all cells (with the possible exception of cells on the edge of the grid) and all times. The constraints embodied in this definition encourage a focus on ways that relatively simple local interactions can generate global dynamical phenomena in spatially extended systems, but these same constraints can make factors operating at different spatial scales, and entities which change spatial extent over time, hard to represent. Typical FEARLUS models differ from a standard CA in several ways (although there are parameter settings which produce CA). However, when a simulation model gives rise to a phenomenon of interest, we believe researchers should seek similar but simpler models that also produce this phenomenon, to identify the mechanisms responsible. In this spirit, the paper will report analytical work and simulation experiments comparing the models currently at the focus of our work with simpler variants, many of them closer to the standard CA framework.

Our approach to simulation makes considerable use of simulation *experiments*. Simulations may quite legitimately be used simply to demonstrate that a model system *can* demonstrate a particular form of behaviour. However, if the model has any stochastic elements (including the selection of initial parameters), as most social simulation models do, it is desirable to go beyond this by using experimental and statistical techniques to discover how it *usually* behaves. Without this, we often cannot be sure that an observed behaviour is not a fluke. Moreover, the ability to *compare* the behaviour of a simulation model under different parameter settings is central to understanding its behaviour, and this demands the ability to test whether apparent differences are real.

A decision to build and run simulations should not mean that analysis is permanently sidelined: simulation studies frequently produce unanticipated outcomes — either unexpected answers to the questions investigated, or phenomena which do not directly bear on those questions, but raise issues not previously considered. In such cases, simulation can suggest new goals for analytical work, and analyses of simplified versions of the simulation system may illuminate its behaviour. This paper illustrates this point with two examples drawn from FEARLUS.

[2] A number of terms will be used to refer to elements of FEARLUS models, some of which could also refer to real-world entities. In FEARLUS model elements' names, each word begins with an upper-case letter (e.g, 'Land Parcel'). Each such term is italicised when first used.

2 Method

2.1 General Description of FEARLUS Models

A FEARLUS model consists of a set of *Land Managers* (representing households, not individuals), and their *Environment*, which includes a grid of square, hexagonal or triangular *Land Parcels*, and a set of possible *Land Uses*. Every *Year*, Land Managers select a Land Use for each Land Parcel they own. The parameters of a FEARLUS model also specify how to determine the *External Conditions*, representing a combination of economic and climatic factors, and encoded as a bitstring, the length of which is a model parameter. The bitstring can vary from Year to Year but applies across the whole grid. External Conditions for any number of Years may be stored in a file, but generally the initial bitstring is determined randomly, and each subsequent bitstring is produced from its predecessor by applying a stochastic process which applies a predetermined *Flip Probability* (f) to each bit independently: if $f = 0$ the initial bitstring will be retained throughout; if $f = \frac{1}{2}$, each Year's bitstring is independent of its predecessors and the External conditions are *temporally uncorrelated*. If $0 < f < \frac{1}{2}$, the External Conditions change, but are temporally correlated. Each Land Parcel has a set of *Biophysical Characteristics*, encoded as a bitstring and fixed for the direction of a simulation run (again, the length of these bitstrings is a model parameter; it is the same for all Land Parcels). There are also two numerical parameters which do not vary over space *or* time: a *Break Even Threshold* (BET), specifying how much *Yield* must be gained from a Land Parcel to break even, and the *Land Parcel Price* (LPP).

In *Year Zero*, Land Parcels are assigned to Land Managers, and there is a random allocation of Land Uses to Land Parcels. Land Managers have an *Account*, initially set to zero (the Year Zero Yield does not affect this, but is available as information in Year 1). The rest of the run repeats the following annual cycle:

1. *Selection of Land Uses.* The Land Use for each Land Parcel is selected by its Land Manager, using the latter's *Land Use Selection Algorithm.*

2. *Calculation of External Conditions.*

3. *Calculation of Yields.* Yield from a Land Parcel is determined by matching the concatenated bitstrings for the Parcel's Biophysical Characteristics and the current External Conditions, against one representing the requirements of the current Land Use: the Yield is simply the number of matching bits.

4. *Harvest.* The Account of each Land Manager is adjusted. For each Land Parcel owned, the Yield for that Parcel is added, and the BET subtracted.

5. *Selection of Land Parcels for sale, and retirement of insolvent Land Managers.* Each Land Manager whose Account is in deficit puts up for sale (at the LPP) as many of their worst-performing Land Parcels as necessary to clear the deficit. A Land Manager unable to do this while retaining at least one Parcel, leaves the simulation.

6. *Land Sales.* The selected Land Parcels are sold in random order. One ticket in a lottery is issued for each *Grid Neighbour* of the Parcel belonging to a Land Manager with at least the LPP in their Account (a Land Parcel's Grid Neighbours are the eight Parcels orthogonally or diagonally adjacent to it) — so Land Managers owning multiple Grid Neighbours get multiple tickets. One ticket is assigned to a potential new Land Manager. A Land Manager must buy the Land Parcel if selected. A new Land Manager starts with an Account of 0 *after* buying the Land Parcel.

For a fuller description of FEARLUS models, see Polhill, Gotts and Law (2001).

2.2 FEARLUS Models and Cellular Automata

The aspects of most FEARLUS models which clash with the letter or spirit of the CA definition given in the Introduction are as follows:

1. A Land Manager can acquire an *Estate* including multiple Land Parcels. These can be arbitrarily distant from each other, and can grow and shrink from Year to Year, yet have an effect on each other (via the Land Manager) from one Year to the next. Furthermore, an Estate and its Land Manager (and by extension, each Parcel in the Estate) has a *Social Neighbourhood*, consisting of all the Estates that include Land Parcels which are Grid Neighbours of any of its own Parcels. Any Parcel in the Social Neighbourhood can affect the choice of Land Use. Hence, it is not generally possible to define a fixed set of neighbours for cells in FEARLUS models - unless one counts all Parcels as Grid Neighbours of each other, which is strictly in line with the definition (for a finite grid), but undermines the notion that a CA is structured into overlapping but non-coincident neighbourhoods.

2. Even without multi-Parcel Estates, there is an aspect of the tie between Land Parcel and Land Manager that violates the CA definition: no limit is set on the amount that a FEARLUS Land Manager's Account can contain, so a cell would have an infinite set of possible states, defined by the possible wealth levels of its Land Manager.

3. A FEARLUS model cell has three kinds of attribute which can distinguish it from other cells (other than its grid location): the attributes of its Land Manager, its current Land Use, and its Biophysical Characteristics. The last of these cannot change. The CA definition above states that the rule table is the same for each cell, but cells with different Biophysical Characteristics have, in effect, different rules determining the other aspects of a cell's state.

4. The CA definition also specifies that the transition rule is the same at all times, but unless the External Conditions bitstrings are of length zero, or the flip probability is zero, this will not hold in FEARLUS models. This is the most fundamental way in which FEARLUS models depart from the CA paradigm: the whole point of the External Conditions is to represent factors which act across the whole Environment, but change over time.

2.3 Simulation Environments

In all models considered here, the Land Parcels are square, arranged in a 7×7 grid, with opposite sides joined to produce a toroidal topology. (The small size permits experiments involving relatively large numbers of runs, while the toroidal topology avoids edge effects; unpublished data indicate that for the kind of experiments reported here, neither larger grids nor different topologies greatly affect results.) The bitstrings defining Land Uses' preferred conditions always contain 16 bits. The bitstrings defining the Land Parcel's Biophysical Characteristics and those defining External Conditions must therefore total 16 bits, but Environments differ in the division of these bits between Biophysical Characteristics (variable across space, but fixed over time) and External Conditions (uniform across space but variable over time). External Conditions may be either correlated or uncorrelated from Year to Year: in the former case, the Flip probability is $\frac{1}{8}$; in the latter, $\frac{1}{2}$. Similarly, the Biophysical Characteristics of Land Parcels may be either clumped or unclumped. In either case, each bit is initially set to 0 or 1 with equal probability and independently, for every Land Parcel. In the 'clumping' process used here, which is carried out on each bit-position in turn during initialisation, adjacent Land Parcels are selected at random to swap non-matching bit-values, for as long as there is a swap which will increase the number of neighbouring Land Parcels pairs that have the same value. This process maintains the number of Parcels having each bit-value, unlike the clumping procedure described in (Polhill et al. 2001).

Aside from the parameters specifying the amount and distribution of spatial and temporal variation in conditions affecting Yield, and the Selection Algorithms followed by Land Managers (discussed below), the only model parameters varying over the experiments reported here are the BET and LPP. In most cases, the BET is set at 8 and the LPP at 16, and in these cases, the Environment used for an experiment will be described using the following syntax:

P$< p >$[c|u]-E$< e >$[c|u]

where p is replaced by the number of bits in the Land Parcel Characteristics bitstrings, the first 'c' or 'u' (absent if $p = 0$) indicates whether these Characteristics are clumped or unclumped, e

is replaced by the number of bits in the External Conditions bitstrings, and the second ‘c’ or ‘u’ (absent if $e = 0$) indicates whether these are correlated or uncorrelated from Year to Year. Thus:

P12u-E4c

indicates 12 unclumped Land Parcel Characteristic bits and 4 correlated External Conditions bits. These characteristics of an Environment are sometimes referred to as its its *Spatio-Temporal Heterogeneity Type* (STHT).

When an Environment has a BET other than 8, and/or an LPP other than 16, this will be indicated by adding a suffix, e.g.:

P0-E16u-BET10-LPP20

indicates an environment with no Land Parcel Characteristic bits, 16 uncorrelated External Conditions bits, a BET of 10 and an LPP of 20.

FEARLUS experiments consist of a number of simulation runs. Those discussed here pit two Subpopulations against each other in *Contests* to assess their success in Environments with particular amounts and types of spatio-temporal heterogeneity, and BET and LPP values. In such a Contest, Land Managers are equally likely to belong to either Subpopulation; all members of a Subpopulation use the same Selection Algorithm. At the start of each run, each Land Parcel is assigned to a different Land Manager. At the end (after 200 Years), Subpopulation success is assessed by counting the Parcels assigned to members of each.

2.4 Aspiration Threshold Selection Algorithms

Most of the Selection Algorithms discussed here use an *Aspiration Threshold*. The Land Manager using an Aspiration Threshold checks whether the most recent Yield from a Land Parcel met the Threshold and if so, leaves the Land Use for that Parcel unchanged. The concept of aspiration thresholds comes from economic psychology: Simon (1955) uses the term *aspiration level* for the minimum price a seller will accept. The concept is linked to that of ‘satisficing’ (Simon 1957), used of agents which cease their search for a problem solution once they find one that is *good enough*, rather than persisting in the search for an optimum solution. There are reasons to expect this kind of behaviour from human economic agents: switching solutions may involve costs, as will the search for a new solution itself. Moreover, it may not be known whether a new solution will indeed turn out better than the current one. All these factors may apply to real-world land managers’ problems, although at present FEARLUS models only deal with the last.

If the Aspiration Threshold is not met, a Land Use must be selected anew. In the Selection Algorithms we focus on here, this involves either *Random Experimentation* (a random choice between the possible Land Uses, all having equal likelihood of being selected), or *Imitation*, in which the Land Use is selected from among those used in the Parcel’s Social Neighbourhood in the preceding Year. Two different forms of Imitation are used, although in both more successful Land Uses have a greater chance of selection. In the original form, called *Yield-weighted Imitation* or YI, a stochastic choice is made between all Land Uses employed within the Social Neighbourhood of the Parcel in the preceding Year, weighted by the total Yields for each Land Use across that Social Neighbourhood. In a revised form, *Stochastic Best-mean Imitation* or SBI, the choice is restricted to those Land Uses with the highest mean Yield across the Social Neighbourhood in the preceding Year, and each of this subset has an equal chance of being selected.

These three Strategies for choosing a Land Use are combined in different ways to give rise to five families of Aspiration Threshold Selection Algorithms used in the experiments reported here. Members of a family differ only in the level of their Threshold:

1. HR: Land Managers employing a *Habit/Random Selection Algorithm* always use Random Experimentation if their Aspiration Threshold is not met. This Algorithm was first used inadvertently, due to an error in setting the probability of using different strategies. We were surprised to find that in some Environments it outperformed the next two Algorithms in this list, and in many others did about as well.

2. HYI: Those employing a *Habit/Yield-weighted-Imitation Selection Algorithm* always use Yield-weighted Imitation if their Aspiration Threshold is not met.

3. HRYI: Those employing a *Habit/Random/Yield-weighted-Imitation Selection Algorithm* choose stochastically whether to use Random Experimentation (with probability 1/16) or Yield-weighted Imitation if their Aspiration Threshold is not met. (The precise value of 1/16 is arbitrary, but that it is small is not. Earlier work (Polhill, Gotts and Law 2001) suggested that a small admixture of Random Experimentation could make a big difference to HYI in some circumstances.)

4. HSBI: Land Managers employing a *Habit/Stochastic-Best-mean-Imitation Selection Algorithm* always use Stochastic Best-mean Imitation if their Aspiration Threshold is not met.

5. HRSBI: This Selection Algorithm relates to HSBI precisely as HRYI does to HYI, with Land Managers using Random Experimentation with probability 1/16 and Stochastic Best-mean Imitation with probability 15/16 if their Aspiration Threshold is not met.

For each of these families of Selection Algorithm, there are two extreme possibilities: the Aspiration Threshold may be set too high ever to be attained (above the maximum possible Yield), or too low ever to be missed (at or below 0). In the former case, the below-threshold method of selecting a Land Use will always be employed, and the resulting Selection Algorithm can be symbolised by dropping the initial H from the acronyms given above: R (but in keeping with earlier papers we will in fact call this 'RS' for 'Random Selection'), YI, RYI, SBI, RSBI. In the latter case, the result is a Selection Algorithm that always leaves the Land Use unchanged; this is symbolised H.

In addition to these families of Selection Algorithm, some others are described, and used for comparative purposes, in section 4.

Since we do not wish to make unsupported assumptions about the statistical properties of the populations of simulation runs from which we draw our experimental samples, we currently use non-parametric statistical tests only. As in Polhill et al. (2001), we use three types of experiment here:

- A type 1 experiment, the most extensively used, consists of a number of simulation runs involving the same two Selection Algorithms, in a type of Environment defined by a fixed set of parameters. These runs differ from each other only because a new seed is generated at the start of each run, for use in the pseudo-random processes employed in the model. The binomial test is used to determine whether one Selection Algorithm has finished significantly more runs holding a majority of the Land Parcels than the other.

- A type 2 experiment compares the performance of two Selection Algorithms against a third 'comparison Algorithm' in a given type of Environment, using a *paired replicates* approach. The Environments for the two members of a matched pair of runs have the same Land Uses, Land Parcel Characteristics, and External Conditions. The sign test is used to determine whether one of the two Selection Algorithms being compared performs better against the comparison algorithm in significantly more of the matched pairs of runs than the other.

- A type 3 experiment uses a similar approach to compare contests between two Selection Algorithms in *two types of Environment.* The sign test is used to determine whether Selection Algorithm A performs better against Selection Algorithm B in Environment type I or II.

For many of the experiments reported, we made predictions of the results (based on the results of exploratory experiments, and our hypotheses about the mechanisms producing them). We report these predictions along with the results. The latter are taken to be significant when they reach the .01 level (1 tailed); we report achievement of the .001 and .0001 levels of significance when this occurs.

3 The Influence of Aspiration Thresholds

3.1 Experimental Findings

Experiments on HR Algorithms, reported in detail elsewhere (Gotts, Polhill and Law 2002), indicate that the Aspiration Threshold makes a considerable difference to the performance of Subpopulations using HR Algorithms, across a wide range of Environments with Land Use bitstrings of length 16, and an LPP of twice the BET. The overall picture is complex, but can be summarised as follows:

- A Threshold equal to the BET is generally at least as good as any alternative, and the optimal Threshold appears never to exceed whichever is the greater of the the BET, and 8 — the Yield that would be expected from a random choice of Land Use, henceforth *Random Choice Expected Yield*, which is half the maximum Yield).
- Greater BET and greater predictability of Yields from one Year to the next both tend to raise the optimal Threshold.
- For a BET of 8, an Aspiration Threshold at the BET was clearly better than lower Thresholds in Environments where more than half the Land Use bitstring was matched against the Land Parcel's Biophysical Characteristics (and less than half against the External Conditions).
- However, in Environments with a lot of uncorrelated temporal variation in External Conditions, a somewhat lower Threshold appeared to do better.
- Experiments with Environments with higher and lower BETs were also carried out on HR; with a BET much lower or much higher than 8, there was a pronounced 'flattening' of differences between different Selection Algorithms. In the former case, most Land Managers would attain a Yield at least equal to the BET in most Years, reducing the possibility of Land Parcels changing hands; with a very high BET, most Land Managers will be replaced every Year.
- However, in relatively predictable Environments with a BET of 6, an Aspiration Threshold of 8 tended to do better than one of 6, despite the flattening.
- Experiments using a BET of 10 produced the most complex pattern, but in no case did an Aspiration Threshold above the BET show a clear advantage over a Threshold of 10.

We also investigated the families of Aspiration Threshold Selection Algorithm involving Yield-weighted Imitation — HYI and HRYI — using a range of 17 Environments with a BET of 8, and at least as many External conditions bits as Land Parcel bits. Results for nine of these are given in table 1.

Note the following points from the table:

- The Threshold 8 Algorithms lose their advantage over their Threshold 10 counterparts in Environments with correlated temporal variation and little or no spatial variation, where imitation of neighbours might be expected to permit higher Yields than elsewhere. The two families (HYI and HRYI) differ, however, in Environments with low spatial variation and uncorrelated temporal variation: H8RYI has a clear advantage over H10RYI, while H8YI has much less of an advantage over H10YI. This flattening effect may due to the fact that in simulations involving only HYI Land Managers, Land Uses which fall into disuse cannot be reintroduced: where there is little choice, it matters less how that choice is made (corresponding HR results resembled those for HRYI).
- This flattening of differences between HYI Algorithms compared to HRYI Algorithms also appears to extend to some extent to the contests between Algorithms with Thresholds of 6 and 8.

Table 1: Aspiration Thresholds 6, 8 and 10 for Yield-weighted Imitation Selection Algorithms. In this and most subsequent tables, columns are headed by the names of two Selection Algorithms, separated by a forward stroke '/', while rows correspond to Environments. The number of 'wins' for the two Selection Algorithms in various types of Environment ('wins' are runs in which the Subpopulation using it ended up with more Land Parcels) are given in the cells of the column. If one of the Algorithms was predicted to do better than the other in a particular cell, the figure recording its wins is italicised. Figures sufficient to confirm such a prediction at significance levels of .01, .001 or .0001 (one tailed) are given one, two or three asterisks respectively, whether or not such a prediction was actually made. Note the differences between Environments with correlated and uncorrelated temporal variation, and the tendency to more one-sided results in the left half of the table.

STHT	H6RYI/H8RYI		H8RYI/H10RYI		H6YI/H8YI		H8YI/H10YI	
P0-E16c	0	*60****	24	36	18	*42**	23	37
P1u-E15u	*42**	18	47***	13	34	26	32	28
P1u-E15c	1	*59****	35	25	11	*49****	31	29
P2u-E14u	*39*	21	*53****	7	*39*	21	*38*	22
P2u-E14c	2	*58****	28	32	6	*54****	31	29
P4u-E12u	36	24	*52****	8	31	29	*44***	16
P4u-E12c	1	*59****	33	27	1	*59****	40*	20
P8u-E8u	17	*43***	*52****	8	16	*44***	*53****	7
P8u-E8c	2	*58****	*42**	18	2	*58****	*43***	17

- The Threshold 8 Algorithms were predicted to beat their Threshold 6 counterparts in all Environments with correlated temporal variation. These predictions were all confirmed at significance level .0001 with the single exception of H6YI *vs* H8YI in P0-E16c, which was confirmed at .01 level. The Threshold 6 Algorithms tend to come out ahead in the Environments with a great deal of uncorrelated temporal variation.

In addition to the Environments listed in table 1, Environments with clumped spatial variation but otherwise identical to those in the second and subsequent rows were used; clumping made little difference to the results. The Threshold 8 Algorithms were also tested against Algorithms with Thresholds of 1, 2, 4 and 12, winning all contests in the listed Environments clearly except for some against the low Threshold variants in P1u-E15u and P2u-E14u. For Environments with BET 8 at least, therefore, HRYI and HYI show similar patterns to HR, with an Aspiration Threshold of around 8 generally giving the best results, but lower Thresholds working well in very unpredictable Environments. Further experiments indicated that the level of Aspiration Threshold generally had more effect than differences *between* the HR, HYI and HRYI families.

3.2 Analysis of Simplified FEARLUS Environments

In analysing simplified versions of FEARLUS Environments and Contests, we use three approaches in various combinations: removing some of the non-CA features of FEARLUS models, considering indefinitely large or very small Environments, and looking at systems involving homogeneous Populations — all following the same Algorithm. Note that if multi-Parcel Estates are not permitted, and all Land Managers in a Population use a non-imitative Selection Algorithm such as HR, each Land Parcel is effectively independent, and we get models without spatial interaction. As this section will show, such models are still worth analysing (and are of course much easier to analyse than models which include spatial interaction). Moreover, results obtained by analysing such models can cast light on models which do include spatial interaction. In this section, we are often concerned with the *proportions* of Land Parcels and/or Managers which will meet a particular fate in the long run and in indefinitely large Environments. The longer a run and the larger an Environment, the more closely our results will tend to be approximated.

Consider an Environment without spatial or temporal heterogeneity, and without the possibility of multi-Parcel Estates — so if a Land Parcel is sold, it always goes to a new Land Manager. Spatial homogeneity requires either a zero-length bitstring, or one which is the same for all Land Parcels. Similarly, temporal homogeneity requires either an External Conditions bitstring of zero-length, or a Flip Probability $f = 0$. If the concatenation of the Land Parcel and External Conditions bitstrings is of zero length, all Land Uses give a zero Yield on all Land Parcels, and all Land Use Selection Algorithms are necessarily equal in performance. Otherwise, each available Land Use will give a fixed Yield across all Years and Parcels, and provided at least one Land Use has a Yield Y_i such that $Y_i < B$ (where B is the BET) and at least one has Yield Y_j such that $Y_j \geq B$, Selection Algorithms may differ in their propensity to enable a Land Manager to survive. We recall that in Year Zero, Land Uses are assigned to Parcels at random, and the Yield does not affect the Manager's Account. This is set to 0 at the start of Year 1, at which point differences between Algorithms begin.

In the type of Environment just described, each member of a homogeneous Population of Land Managers using the Random Selection Algorithm (RS), which uses Random Experimentation every Year, is engaged in a random walk (Grimmett and Stirzaker 1992, sections 3.9, 5.3) independent of its neighbours. There are five qualitatively distinct regimes of behaviour for such a system (provided we set no limit on Account size), depending on the distribution of Land Use Yields:

1. All Land Managers are replaced each Year. This occurs if all Land Uses have Yields $< B$.

2. All Managers will eventually be replaced, there is no limit on lifespan, but the mean lifespan is finite. Sufficient (but not necessary) conditions for this to occur are at least one Land Use with Yield B, at least one with Yield $< B$, and none with Yield $> B$. For example, two Land Uses with Yields B and $B - 1$ give a mean time to replacement of 2.

3. All Managers will eventually be replaced, but the mean lifespan is infinite. This requires at least one Land Use each with Yields $< B$ and $> B$, but these conditions are not sufficient. An example occurs if there are two Land Uses with Yields $B - 1, B + 1$.

4. The proportion of Managers replaced approaches a limit < 1. Again, this requires at least one Land Use each with Yields $< B$ and $> B$. If there are three Land Uses with Yields $B - 1, B + 1, B + 1$, $\frac{1}{2}$ the Managers will eventually be replaced.

5. All Managers live forever. This happens if all Yields are $\geq B$.

Since regimes 1 and 5 occur irrespective of the Algorithms followed by Managers, we ignore them in what follows.

Turning to Aspiration Threshold Selection Algorithms, a homogeneous Population of HR Land Managers will show different dynamics depending on the Aspiration Threshold used, and in particular on its relationship with B, and on the maximum and minimum Yields, Y_M and Y_L. Calling the Aspiration Threshold T, we can distinguish the following cases:

1. If $T \leq Y_L$, HR becomes equivalent to the simple *Habit Selection Algorithm* (H), which always sticks with the current Land Use. Every Land Parcel will retain its original Land Use. Any Parcels where the Yield is $< B$ will see their Land Manager replaced every Year, while the rest will retain their Manager indefinitely.

2. If $T > Y_M$, HR becomes equivalent to RS.

3. Otherwise, all Land Parcels will eventually reach a Land Use with a Yield $\geq T$, and retain it thereafter, but three cases can be distinguished.

(a) If $Y_L < T \leq Y_u$, where Y_u is the greatest Yield $< B$ produced by any Land Use, any Land Parcels on which the Yield Y_i is such that $T \leq Y_i \leq Y_u$ will see their Land Manager replaced every Year (the Land Use remains unchanged as its Yield achieves the Threshold). The remaining Parcels will retain both Land Use and Manager forever. Call the Algorithm in this case HR_a, and its Aspiration Threshold T_a.

(b) If $Y_u < T \leq Y_v$, where Y_v is the least Yield $\geq B$ produced by any Land Use, all Parcels will acquire a fixed Land Use and Manager the first time a Land Use with Yield $\geq B$ is tried. Call the Algorithm in this case HR_b, and its Aspiration Threshold T_b.

(c) If $Y_v < T \leq Y_M$, all Parcels will eventually acquire a fixed Land Use and Manager, but the process may be slowed because the current Manager will abandon a Land Use sufficient to allow them to survive if its Yield is $\geq B$ but $< T$. This risks bankruptcy even if there is a better Land Use than the original to be found, since a run of bad luck in Random Experimentation might lead to the repeated selection of a Land Use with Yield $< B$. Call the Algorithm in this case HR_c, and its Aspiration Threshold T_c.

In a spatio-temporally homogeneous Environment without multi-Parcel Estates, any HR Aspiration Threshold $> Y_u$ maximises the probability that a Land Manager survives their first Year (this probability is 1 for any Land Manager present at the start of the simulation, since none are replaced after Year 0, and after that, a new Land Manager will always inherit a Land Use with a Yield $\leq Y_u$), while an Aspiration Threshold $\leq Y_v$ prevents the Manager losing opportunities to live forever. Note that if $Y_M = B$, (so that no Land Manager's Account can ever exceed 0, and the number of functionally distinct states a Land Parcel can enter will be finite), Algorithm HR_c cannot exist.

In Contests between Subpopulations of HR with different Thresholds, HR_b does best.

1. Against H (equivalent to HR with Threshold $> Y_M$) it will, eventually, capture all Parcels initially assigned to H with a Land Use Yielding $< T_b$ (equivalently, $< B$). While such a Parcel is under any Land Use Yielding $< B$, it will get a new Manager each Year, but if this is another H Manager, the Land Use remains unchanged. If the proportion of Land Uses Yielding $< B$ is r, there is a $(1-r)/2$ chance each Year that the Parcel will acquire an HR_b Manager and a Yield $\geq B$. The chance that it will *not* do so in n Years is thus $((1+r)/2)^n$, which tends to 0 as n grows.

2. A qualitative picture of the possible histories of individual Parcels in contests between HR_b and RS is given in figure 1. The rectangles on the left and right stand for possible states at the end of Year 0, the remainder for possible states immediately after Yields have been calculated and before Land Sales in subsequent Years. The arrows indicate possible transitions. As can be seen, the only state without a transition out of it is that labelled S2 — with an HR_b Manager and a Yield $\geq T_b$. Knowing the proportion of Land Uses with such Yields would not suffice to attach probabilities to all transitions, specifically to those out of state S3, since this may be entered, from state S1, with any amount in the Account. If the Yield distribution is such that RS on its own would be in regime 2 or 3 (see above), indefinitely repeated chances will arise for state S2 to be entered, and all Parcels will end up there; if RS would be in regime 4, some proportion less than 1 will do so, but the possibility of further Parcels moving to S2 will always remain.

3. Figure 2 shows possible parcel histories in contests between HR_a and HR_b. We can note that a larger proportion of Parcels can be expected to enter S2 than S1 in Year 1, that each of the transitions to S1 from non-initial states is matched by a transition from the same state to S2 (which will have the same probability of occurring), and that there are additional transitions to S2 from S3 and S4. We can conclude that a higher proportion of Parcels can be expected to be in S2 than in S1 in all Years. The state transition probabilities have been calculated in terms of the proportions of Land Uses with Yields in each of the three ranges represented by the three rows of state boxes. A closed-form solution for the proportions of Parcels ending up in S1 and S2 has not yet been derived, but numerical work indicates that

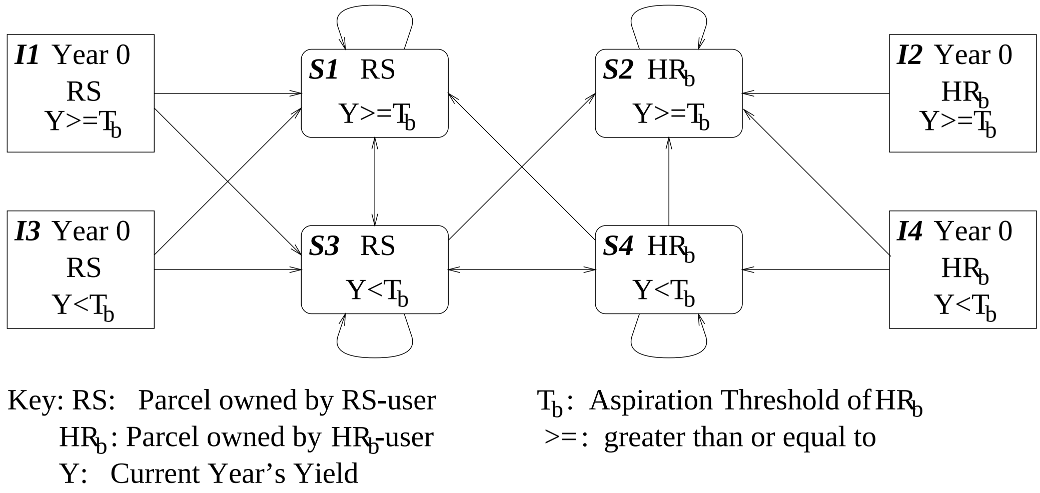

Figure 1: Possible histories of a Land Parcel in a contest between RS and HR$_b$

if y is the proportion of Land Uses with Yields $\geq T_a$ but $< T_b$, then a proportion $\frac{1-y}{2}$ of Parcels will end up in S1 and $\frac{1+y}{2}$ in S2.[3]

4. Figure 3 shows possible Parcel histories in contests between HR$_b$ and HR$_c$. Here, as for state S3 in the RS/HR$_b$ contest, knowing the proportion of Land Uses with Yields in each of the three ranges does not suffice to determine transition probabililites from S5, because S5 may be entered from S3 with any amount in the Manager's Account. However, as for the HR$_a$/HR$_b$ contest, the probability of a Parcel being in one of the final states (S1, S2 and S4) approaches 1 as time approaches infinity, which is not necessarily the case for RS/HR$_b$.

   In this case, the best approach to Parcel histories is to consider the probability that a Land Manager of either type will eventually go bankrupt. Let the probability that a randomly selected Land Use has a Yield $\geq T_b$ be b, and the probability that it has a Yield $\geq T_c$ be c (so $c < b$). The probability that an HR$_b$ Land Manager assigned to a Land Parcel in Year Zero will go bankrupt (this can only happen at the end of Year 1) is then $(1-b)^2$. The corresponding probability for an HR$_c$ Land Manager is higher: it will be $> (1-c)(1-b)$, since the probability the Year Zero Land Use will be rejected by such a Land Manager is $1-c$, the probability that in that case the Year 1 selection will lead to immediate bankruptcy is $1-b$, but if this does not occur, and the Land Use selected does not have a Yield $\geq T_c$, eventual bankruptcy still has a positive probability. Since $c < b$, $(1-c)(1-b) > (1-b)^2$. Hence the chance that the first Manager of a Parcel will hold it for ever is higher if that Manager uses HR$_b$ than if it uses HR$_c$. If the first Manager does go bankrupt, the second is equally likely to be an HR$_b$ or an HR$_c$ (and similarly, if the nth does so, the $n+1$th is equally likely to be of either type). For an HR$_b$ Manager inheriting a Parcel any time after Year 1, bankruptcy can only occur in the first Year of possession, and will do so with probability $1-b$ — if and only if the Yield of the randomly chosen Land Use is $< T_b$. Correspondingly, the chance of holding the Parcel for ever is just b. For an HR$_c$ Manager, however, bankruptcy can occur after any number of Years — since the Manager will keep chosing at random unless and until a Land Use with Yield $\geq T_c$ is found; the probability bankruptcy occurs eventually will be $> 1-b$, since it is $1-b$ in the first Year, but with probability $b-c$ a Yield $\geq T_b$ but $< T_c$ will be gained in the first Year, and in this case the chance of eventual bankruptcy is > 0. Hence for the second and any succeeding Land Manager owning a Parcel, as for the first, the chance of holding the Parcel permanently will be higher for an HR$_b$ than for an HR$_c$ Land Manager. We can conclude that the Parcel is more likely to end in the permanent possession of an HR$_b$ than an HR$_c$ Manager.

[3]This work was performed by our colleague Luis R. Izquierdo.

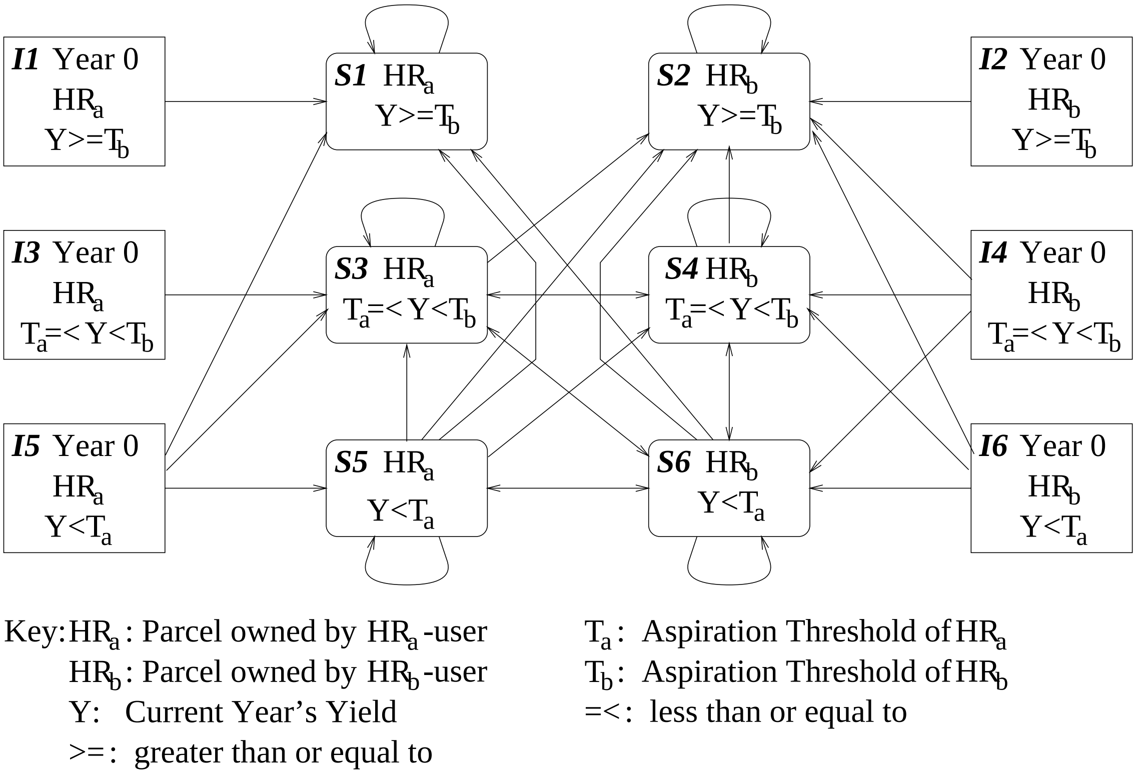

Figure 2: Possible histories of a Land Parcel in a contest between HR_a and HR_b

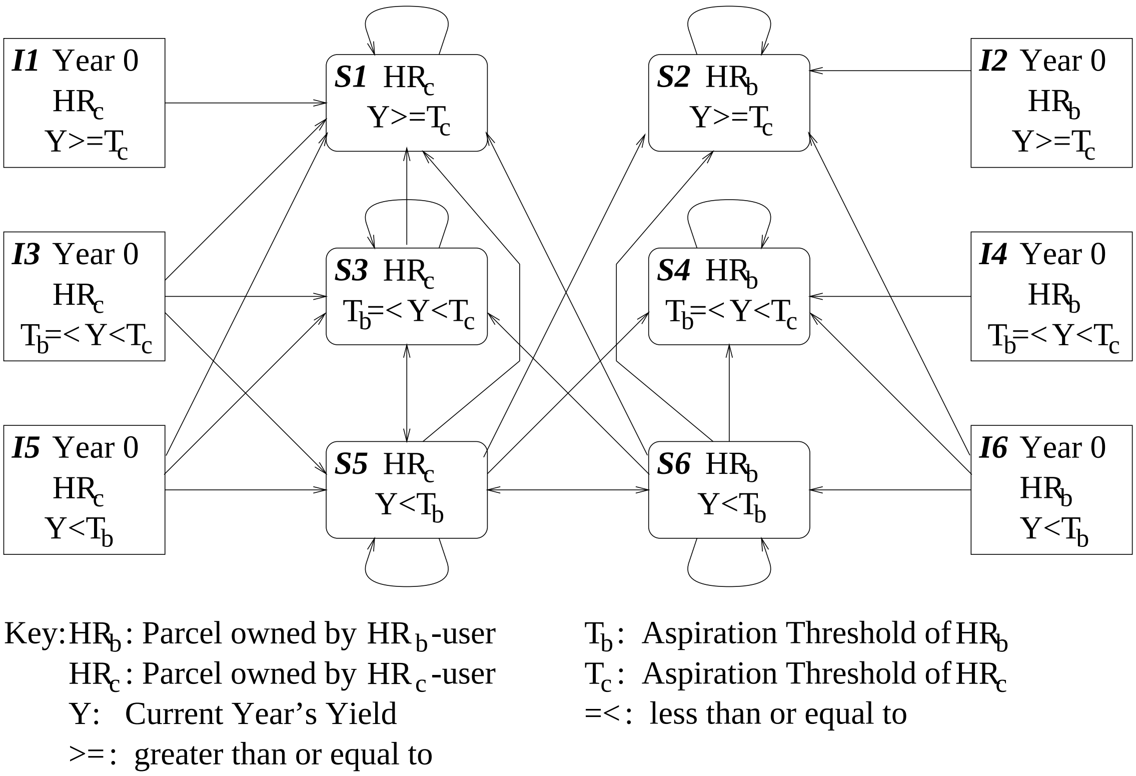

Figure 3: Possible histories of a Land Parcel in a contest between HR_c and HR_b

The classification of variants of HR in terms of the relationship between the Aspiration Threshold, BET and distribution of Land Use Yields, can also be applied to HRYI and HYI, although significant differences arise in subsequent analysis of system behaviour, particularly for HYI.

If the Aspiration Threshold is low enough, HR, HYI and HRYI are all equivalent to H.

For homogeneous HRYI Populations, the possible histories of individual Parcels are qualitatively like those of corresponding HR Populations, even in case (2) where $T > Y_M$, and the Selection Algorithm can be called RYI. For a homogeneous population of RYI Land Managers, every Land Parcel will pass through every Land Use given long enough (as in the case of RS), although there will be a bias toward the higher-Yielding Uses. For any finite Environment, there would eventually come a Year when all Land Managers simultaneously decided on Random Experimentation, returning the system to its Year 0 distribution of possible Land Uses. For individual Land Managers, depending on the distribution of Yields and the probability of using Random Experimentation, it might be a certainty that eventually they would be forced to sell their Land Parcel, or the probability of doing so might approach some limit strictly between 0 and 1. Again, this parallels the RS case. However, there would be spatio-temporal patterns of correlation in Land Use, and in the rate of Land Manager turnover, that would not occur in the RS case, where every Parcel and every Year are independent. A similar pattern of parallels and differences applies to comparisons between homogeneous Populations of HR_a and HRYI_a, HR_b and HRYI_b, HR_c and HRYI_c.

Many qualitative aspects of the descriptions of dynamics of Contests between HR Algorithms with different Thresholds also apply to their HRYI counterparts. Every transition diagrammed in figures 1, 2 and 3 would be possible in corresponding Contests between HRYI variants. The arguments concerning the histories of Land Parcels in such Contests given above can also be extended to their HRYI counterparts, since they depend only on the topology of the transition network, the fact that a new Land Manager is (by the way the Contests are defined) equally likely to belong to either Subpopulation, and the fact that the same procedure is used by both Subpopulations to determine the new Land Use if their respective Aspiration Thresholds are not met. The case is different if we consider Contests of HR against HRYI, on which no similar analysis has yet been attempted.

In the current context, HRYI is much more like HR than it is like HYI — though simulation experiments reported in Gotts, Polhill, Law and Izquierdo (2003) show the opposite in most 200-Year Contests in spatio-temporally heterogeneous Environments. A classification parallel to that suggested for HR makes sense, but in the case of Populations consisting entirely of HYI Land Managers, must be defined relative to the subset of Land Uses that are present in Year Zero, rather than the total set. (Of course, as the Environment size approaches infinity, the likelihood of all Land Uses being included in the subset approaches 1.) A particular Land Use also vanishes for good if it does not get used in a subsequent Year, even if there are still Managers looking for a Land Use with Yield at or above their Threshold. In HYI Populations where all Managers have the same Threshold, however, all Land Uses present in Year Zero which have a Yield at or above the Threshold will be preserved, as they will be used at least on the Land Parcels they occupied in Year Zero.

The high-Threshold version of HYI (counterpart to RS in the HR classification and RYI in the HRYI case), is YI, which uses Yield-weighted Imitation all the time. Here, the Environment will eventually become, and remain, a monoculture. The higher a Land Use's Yield, the more likely it is to be the single surviving Land Use, but all Land Uses have a finite probability of doing so. If that Land Use has a Yield $\geq B$, all Land Managers survive thereafter; otherwise, all are replaced Yearly. To summarise the differences between RS, RYI and YI: for RS, each Parcel acts independently, changing Land Use at random and Land Managers independently; for YI all Parcels eventually reach the same Land Use, and either all or none of the Land Managers change Yearly; for RYI Land Uses are highly correlated across Land Parcels and Years but given long enough, will change, while the dynamics of Land Manager lifespan are complicated, and depend on the distribution of Yields.

Systems with homogeneous Populations of HYI_a, HYI_b, or HYI_c will all end up with all Parcels having Land Uses giving Yields at or above the Managers' Threshold T, as is the case for corresponding HR or HRYI Populations. In the latter cases, however, a Parcel with a current initial Yield $< T$ (an 'unsettled' Parcel) may end up with any Land Use with Yield $\geq T$. In the HYI case, unsettled Parcels always imitate a neighbour. A currently unsettled Parcel can only end up with its current Land Use, or a Land Use that currently occupies a Parcel which is either a Grid Neighbour, or to which there is a 'Neighbourhood chain' (a sequence in which each Parcel is a Grid Neighbour of the next) via other unsettled Parcels. As further Parcels get settled, the range of possibilities for unsettled Parcels is likely to be progressively reduced. In the HYI_a case, it may be possible to determine, before a Parcel settles, whether or not its final Land Use will have a Yield $\geq B$ and so whether it will get a permanent Manager.

Contests between HYI_a and HYI_b Subpopulations are qualitatively similar to those between HR_a and HR_b, or between HRYI_a and HRYI_b but those between HYI_b and either YI or HYI_c differ from their counterparts.

First, in the RS/HR_b and RYI/HRYI_b cases, the probability of there being Land Managers who are liable to go bankrupt does not necessarily approach 0 with increasing time. The distribution of Yields may be such that there is a non-zero probability of an RS (or RYI) Land Manager remaining in business, but always liable to go bankrupt through a run of bad luck, indefinitely. This cannot be the case in contests between YI and HYI_b, provided *any* HYI_b Land Manager finds a Land Use with Yield $\geq T_b$ (and therefore $\geq B$). Such a Land Use is guaranteed to remain in existence forever, along with the HYI_b Manager using it. Once this is the case, there will always be a non-zero probability that all current Land Uses with Yield $< B$ will disappear within d Years, where d is the 'diameter' of the Environment: the greatest distance between any pair of Parcels, where distance is defined as the number of links in the shortest Neighbourhood chain between a pair. Consequently, the probability that all Land Uses with Yield $< B$ will have vanished approaches 1 as time increases.

Second, in both YI/HYI_b and $\text{HYI}_c/\text{HYI}_b$ contests, it is possible for Land Managers to be guaranteed eternal solvency, yet to keep on changing their Land Use. A YI Manager surrounded by HYI_b Managers who have all found satisfactory Land Uses, will do so if these neighbours do not share a common Land Use. So will an HYI_c Manager surrounded by HYI_b Managers who are all getting Yields $\geq T_b$ but $< T_c$ and do not share a common Land Use. An RS or RYI Manager similarly surrounded by HR_b or HRYI_b Managers remains always liable to bankruptcy, while an HR_c or HRYI_c Manager can only escape this liability by finding a Land Use satisfying their own Threshold.

The analyses reported here show that a small amount of Random Experimentation can make a big difference. Nevertheless, for HYI as for HRYI and HR, the analysis shows that in spatio-temporally homogeneous Environments without multi-Parcel Estates, Aspiration Thresholds at or near the BET outperform either higher or lower ones.

3.3 Further Experiments

The spatio-temporally homogeneous Environments analysed above do not require Land Managers to track changes over time or compensate for differences between Land Parcels — ubiquitous problem for real-life land managers. So while the analysis clarifies why most of our simulation results showed Thresholds close to the BET as optimal, it also suggests we should search for exceptions if we want to understand the interactions between Land Managers' Land Use Selection Algorithms and spatio-temporal heterogeneity. Simulations already described show that there are some kinds of spatio-temporal heterogeneity which give Thresholds lower than the BET an advantage; and some (but only where the BET is lower than the Random Choice Expected Yield) where a Threshold above the BET is better; an obvious target for further simulation experiments was to find circumstances in which the BET is at least equal to this Yield, but the optimum Aspiration Threshold is higher. Given our interest in interactions between neighbours, and specifically in imitation of neighbours' successes, it made sense to seek a simple imitation-based Selection Algorithm for which this was true in at least some Environments.

Since imitation of Neighbours is likely to give the best results when their Land Parcels are the same as a Land Manager's own, and imitation which takes account of recent Yields is likely to

Table 2: Aspiration Thresholds for Random/Stochastic Best-mean Imitation Selection Algorithms.

STHT	H8RSBI/H10RSBI		H8RSBI/H12RSBI		H10RSBI/H12RSBI	
P0-E16c	24	*96****	45	*75**	59	61
P1c-E15u	*81****	39	*80***	40	60	60
P4c-E12c	43	*77**	43	*77**	57	63
P4c-E12u	*98****	22	106^{***}	14	80^{**}	40
P4u-E12c	46	*74**	54	66	62	58
P4u-E4u	*112****	8	*118****	2	*94****	26
P8c-E8c	50	70	57	63	67	53
P12c-E4c	*75**	45	*78***	42	*83****	37

be most useful if these are at least somewhat predictable, the P0-E16c Environment and those similar to it appeared the best place to look. However, we had already found that for HYI and HRYI, an Aspiration Threshold of 8 outperformed higher Thresholds. We therefore decided to investigate the HRSBI family of Land Use Selection Algorithms, described above, in which the imitation target is chosen from among those Land Uses with the highest *mean* Yield across the Social Neighbourhood, with no weight given to the number of Parcels employing each Land Use (Stochastic Best-mean Imitation) rather than from among those with the highest *total* Yield. In a spatially heterogeneous Environment, this could be a high-risk approach, but in a spatially homogeneous one, switching to the Land Use with the highest mean Yield will never give a lower expected Yield than sticking with the current one. However, it was not obvious how much spatial heterogeneity Stochastic Best-mean Imitation could cope with, nor whether, even in the P0-E16c Environment, it would be better than sticking to the current Land Use often enough to make a Threshold above the BET significantly better than one equal to the BET.

In the event, exploratory experiments suggested that this would be the case, and that an Aspiration Threshold of 10 or even 12 would outperform one of 8 even in the presence of significant spatial heterogeneity. Results of full-scale experiments matching Aspiration Thresholds of 8 and 10, shown in table 2, confirmed these expectations.

For the Environment to which H10RSBI was most obviously suited, we also decided to test it against a range of alternative Aspiration Threshold Selection Algorithms involving Imitation and/or Random Experimentation. Since an exploratory experiment had suggested, somewhat to our surprise, that H10SBI (i.e., H10RSBI without the occasional use of Random Experimentation) at least held its own against H10RSBI in this Environment, we decided to test that Algorithm as well. In this series of experiments H10SBI again did well against H10RSBI, winning 71 out of 120 contests (this would be enough for a two-tailed rejection of the null hypothesis of no difference in performance between the two at the .05 level, but not at the .01 level). The remaining results are shown in table 3, and suggest that while both H10SBI and H10RSBI do well, H10RSBI performs better against those Algorithms which themselves never use Random Experimentation.

To test the impression that H10SBI outperforms H10RSBI in direct contests in P0-E16c, but does worse against some other Algorithms, some 240-run experiments were undertaken. First, a type 1 experiment was run, pitting H10SBI against H10RSBI, with the former predicted to win: H10SBI won 134 of the runs, which would be significant at the .05 level, but does not reach the .01 level. Second, three type 2 experiments were run, comparing the performance of H10SBI and H10RSBI against H8YI, H8SBI, and H10SBI, with H10RSBI expected to do better in all cases. In the first two cases this prediction was confirmed at the .0001 level of significance, and in the third at the .001 level. This is hypothesised to be another instance of 'flattening' due to a phenomenon suggested earlier to account for results concerning HYI and HRYI: progressive loss of Land Uses from the simulation when all Land Managers always select a Land Use currently employed in the

Table 3: H10SBI and H10RSBI tested against other Aspiration Threshold Selection Algorithms, in Environment P0-E16c. In this table, the two rows correspond to H10SBI and H10RSBI, the columns to other Selection Algorithms, and each cell contains the number of runs, out of 120, in which the *row* Selection Algorithm ended up with a majority of Land Parcels. If this number is italicised, the row Selection Algorithm was predicted to win; one, two or three asterisks following the number indicate that the row Algorithm won enough to confirm such a prediction (whether or not one was made) at the .01, .001 or .0001 (one-tailed) significance level.

	H8YI	H8SBI	H8RYI	H8RSBI	H8R	H10YI	H10RYI	H10R
H10SBI	80**	66	*113****	*97****	*118****	77*	*107****	*116****
H10RSBI	*107****	*83****	*114****	*88****	*118****	*110****	*113****	*119****

Social Neighbourhood.

To summarise this section:

1. Experiments showed the importance of Aspiration Threshold level across a range of Environments, and Selection Algorithm families. They indicated that a Threshold equal to the BET was frequently best.

2. Analysis of spatio-temporally homogeneous FEARLUS Environments, without Land Sales, showed that, at least for imitative Selection Algorithms making some use of Random Experimentation, Thresholds near the BET could indeed be expected to outperform both lower and higher Thresholds in the long term.

3. However, the combination of existing simulation results and subsequent analysis suggested a new target for simulation studies: finding a combination of an imitation-based Selection Algorithm, and an Environment, in which Aspiration Thresholds above both the BET and Random Choice Expected Yield would perform well. The combination of Stochastic Best-mean Imitation and a spatially homogeneous but temporally heterogeneous Environment turned out to meet this specification.

Further investigations, both analytical and simulation-based, should reveal how different forms of imitation compare across a wider range of Environments, and perhaps indicate how imitation strategies should themselves be adapted to features of the Environment.

4 The Advantages of Diversity

One Environment used in exploratory experiments, P0-E16u, was included largely as a check that spurious results were not being generated. This Environment is spatially homogeneous, while External Conditions are variable and temporally uncorrelated. Any Land Use (and hence any Selection Algorithm) gives the same expected Yield, and indeed the same expected distribution of Yields over a period of Years, on any Land Parcel. It surprised us when exploratory experiments produced results suggesting that some Selection Algorithms systematically outperformed others in terms of the number of Land Parcels controlled after 200 Years. In particular RS seemed to do well. Good performance by H, the Habit Selection Algorithm, against H8RYI suggested that it was the diversity of choices across Land Parcels within a Year, rather than change in Land Uses from Year to Year, that gave rise to RS's success in this Environment — or at least, that the former was sufficient to give an advantage.

To investigate this phenomenon systematically, we chose four Selection Algorithms: RS, H8R, H8RYI and *Last-year's-optimum-match Deterministic Algorithm* (LD). This Algorithm, unlike those considered so far, relies on 'Innate Knowledge' of Land Uses — i.e., knowledge it does not

Table 4: Effect of LPP on contests in P0-E16u Environments: type 1 experiments.

LPP	LD/H8RYI		LD/H8R		LD/RS		H8RYI/H8R		H8RYI/RS		H8R/RS	
0	195	*285****	147	*333****	137	*343****	199	*281***	162	*318****	210	*270**
16	207	*273**	203	*207***	182	*298****	205	*275***	182	*298****	224	256
2000	241	239	241	239	243	237	232	248	242	238	246	234

Table 5: Effect of LPP on contests in P0-E16u Environments: type 3 experiments. Note that in this table, the two numbers in each cell need not sum to the total number of pairs of runs in the corresponding experiment, since pairs in which the outcome was the same for both members of the pair are not counted.

	LD/H8RYI		H8RYI/H8R		H8R/RS		LD/RS	
LPP0/LPP16: better with LPP0	57	*49*	40	*75***	56	*60*	28	*88****
LPP16/2000: better with LPP16	46	*64*	55	*58*	48	*63*	26	*86****

need to acquire through experience. It uses last Year's bit-values of the External Conditions bitstring, along with those of the Land Parcel's bitstring, to calculate what the Yield from each Land Use would have been in that Year. The Land Uses created at the start of a run are numbered (1 to 8), and LD selects the lowest-numbered among the Land Uses which would have given the best Yield. In any spatially homogeneous Environment, all Land Managers using it will thus select the same Land Use for all their Parcels in any given Year, although in P0-E16u this choice will typically vary from Year to Year.

Since the effect under study appeared quite weak, we matched each of the four Selection Algorithms against the other three in 480-run type 1 experiments in P0-E16u. RS was predicted to beat H8R, both these to beat H8RYI, and all three of these to beat LD. The results are given in table 4 (middle row), and as can be seen, almost all the predictions were confirmed at a significance level of .01 or better.

We hypothesized that the Land Sale process included in the model might somehow be responsible for the phenomenon. If so, setting the LPP so high that no Estate Neighbour can ever afford to buy up a Land Parcel should abolish the effect. If it is set to zero, the phenomenon should if anything be enhanced. We therefore ran the same six Contests between Selection Algorithms in the Environments P0-E16u-LPP2000 and P0-E16u-LPP0 — predicting the same victors in the latter case, but making no predictions in the former. As can be seen from table 4, the predictions were borne out.

We also ran some type 3, 120-run pair experiments, comparing Environment P0-E16u with P0-E16u-LPP2000 and with P0-E16u-LPP0. For each of these two Environment comparisons, we ran four experiments, pitting LD against H8RYI, H8RYI against H8R, and H8R against RS, and also LD against RS, in each case predicting that the second of each pair would do better in P0-E16u than in P0-E16u-LPP2000, and better in P0-E16u-LPP0 than in P0-E16u. The predictions for LD against RS were overwhelmingly confirmed. The remaining results (see table 5) are less clear, but all except one are in the expected direction.

It is clear that in a spatio-temporally homogeneous Environment, RS could not have this kind of advantage over LD: LD would always pick one of the highest-Yielding Land Uses, and if this Yield were at or above the BET, no LD Land Manager would ever have to sell. The experiments described indicate that the phenomenon is also dependent on the Land Sale mechanism, but how does it bring about these effects? There are at least two possibilities:

- Once a Land Manager gains control of more than one Land Parcel, perhaps greater diversity of Land Uses reduces the likelihood of ending up with a negative Account, and thus having to get rid of Parcels again.

- Recall that Land Parcels sold are assigned at random either to a Neighbour with sufficient funds in their Account to pay the LPP, or to a new Land Manager. Each Grid Neighbour owned by an eligible Land Manager gives that Land Manager one 'ticket' in a draw for the right (and obligation) to buy, and one ticket is assigned to a potential new Land Manager, equally likely to belong to either Subpopulation. Hence the probability that a Land Parcel will pass into the control of each Subpopulation depends on the Subpopulation membership of eligible Land Managers, and the numbers of Grid Neighbours of the Parcel for sale that each controls. Perhaps Subpopulations in which different Land Managers favour different Land Uses are better able to maintain a dominant position in a Population once achieved. When a Population is dominated by one of two competing Subpopulations, if a small proportion have to sell Parcels each Year, most of these will be acquired by their neighbours and the dominance of that Subpopulation will persist. If none need to sell in most Years, but a large proportion do so occasionally, the Subpopulation may lose its dominance when that occurs.

Both of these mechanisms can be shown to operate in simple FEARLUS models involving just two Parcels, and a Selection Algorithm devised for the purpose, the *Fickle Specialist Selection Algorithm* (FS). A Fickle Specialist Land Manager chooses a Land Use at random each Year, and applies it on all the Land Parcels they own. Hence on an Estate consisting of a single Land Parcel, FS is equivalent to RS.

Consider a P0-E1u-BET0.5-LPP1 Environment (like P0-E16u Environments, P0-E1u Environments are spatially homogeneous with uncorrelated temporal variation), just two Land Parcels, and two Land Uses. Land Use 1 produces a Yield of 1 if the External Condition bit has value 1 and a Yield of 0 otherwise, Land Use 0 the reverse. (The LPP could in fact take any finite non-negative value, but the argument is particularly simple for the value 1.) Assume two Subpopulations, one using RS, the other FS, and equally likely to provide Land Managers. We can show that in the long run, there will be more Years when both Land Parcels are managed by RS-users than when both are managed by users of FS.

So long as the two Land Parcels have different Managers, as will be the case initially, the dynamics will be exactly the same whichever Subpopulations they belong to. A Manager of either kind coming into ownership of a single Parcel in this model, whether at the start of a run or later, is embarking on a random walk which is certain to end in bankruptcy (when the Account is in deficit by $\frac{1}{2}$, and selling off the Parcel becomes necessary) but with an infinite expected duration (Grimmett and Stirzaker 1992). When this happens, provided the Neighbour has at least 1 in its Account, there is a $\frac{1}{2}$ chance the Parcel will be bought by that Neighbour. Sooner or later, therefore, one Land Manager will own both Parcels, although the *expected* time before this occurs is infinite.

Assume that at this point, the Land Manager owning both Parcels has an amount n in its Account, where n is a non-negative integer, or such an integer plus $\frac{1}{2}$.

Consider first the case where this Land Manager uses FS. Each Year, both Parcels will be assigned the same Land Use, and there will be a $\frac{1}{2}$ probability of gaining a net Yield over the two Parcels of 1, and the same of a net Yield of -1. Sooner or later, it is certain the Account will go into deficit (containing -1 or $-\frac{1}{2}$) — obliging the Manager to sell one Parcel to a new Manager — although the expected time for this to occur is infinite. For any number of Years y, however, if $y > n$ there will be a positive probability p_y that it has occurred, and $p_y \to 1$ as $y \to \infty$. Any possible 'path' to the forced sale of one Parcel can be described simply by listing the amounts in the Manager's Account after each Year — e.g., starting with 3 in the Account, this path could be $[3, 2, 1, 0, -1]$ (taking 4 Years), or $[3, 4, 3, 2, 3, 2, 1, 0, 1, 2, 1, 0, -1]$. All paths of the same length have the same probability of occurrence: 2^{-y} where y is the number of Years required.

Now consider the case in which both Parcels are owned by a single RS Land Manager.

- Since RS has a $\frac{1}{2}$ chance of choosing different Land Uses for the two Parcels, with a consequent net Yield from the two of 0, there is, if the Manager has amount $a \geq 0$ in its Account at Year y, a $\frac{1}{2}$ chance that it will still have a at Year $y + 1$.

- Eventually, however, the Manager will choose the same Land Use for both Parcels, and will either gain or lose 1 — with equal probability.
- For an RS Manager, a path to forced sale of a Parcel can therefore contain steps which leave the amount in the Account unchanged — e.g. $[3, 2, 2, 2, 1, 1, 0, -1]$. We can divide these paths into sets, where each (infinite) set contains all those paths in which the sequence of *changes* of the amount in the Account is the same — so $[3, 2, 2, 2, 1, 1, 0, -1]$ would be in the $[3, 2, 1, 0, -1]$ such, and each set corresponds to a single path for the FS Manager.
- The probability that the path taken by an RS Manager belongs to a particular set is the same as the probability that the FS Manager follows the corresponding path.
- All but one set member, however, will take longer to complete. If we consider all the FS-paths that end in y Years or less, *only* members of the corresponding sets of RS-paths could do so, and in each case, only a finite subset of the infinite set will do so.
- For any $y > n$, therefore, the probability that an RS Manager will lose one of its Parcels will be less than the corresponding probability for an FS Manager.

For a system of this kind running for z Years, for any $z > 2$, the probability of finding both Parcels RS-owned will thus be greater than that of finding both Parcels FS-owned.

It is worth noting that if a *Diversification Selection Algorithm* (DS), guaranteeing that different Land Uses are always employed on the two Parcels when the Land Manager owns both, replaced FS in this system, the certain outcome would be a stable state in which this occurred every Year, and the DS Manager would hold both Parcels for ever once established.

Figure 4 illustrates the dynamics of a contest between FS and LD Subpopulations in a similar FEARLUS Environment, differing in that the BET is 1 (which makes the state diagram finite) and the LPP 0. Land is free (or worthless) once abandoned by its former owner, and any Land Manager can remain in business only by choosing the right Land Use on the Parcel or Parcels owned every Year. Again, any Selection Algorithm will give the same expected Yield on either Land Parcel: $\frac{1}{2}$. In this Environment, FS will hold Land Parcels more often than LD.

Figure 4 represents system states and transitions in this system. The six heavily outlined boxes represent the possible states of the system just before Land Uses are selected — these are distinguished only by whether the two Land Parcels belong to the same Land Manager, and the Subpopulation to which the Manager of each Parcel belongs (there is a line between the Parcels when and only when they are owned by different Managers). The remaining boxes represent the possible transitional states of the system immediately before any Land Parcels without a solvent Manager are assigned one — these Land Parcels are shown as vacant. The system begins in the central box. If one Land Manager (whether FS or LD) owns both Parcels, it will assign the same Land Use to both. Two LD Land Managers each owning one Parcel will always choose the same Land Use in any given Year. In contrast, FS Land Managers will assign Land Use 0 or Land Use 1 with equal probability to the Parcel or Parcels they manage, each Year, and if two each own a Parcel, they are equally likely to choose the same Land Use, or different ones. The labelled arrows show transition probabilities between states. The probability of the system being in each of the six states represented by heavily outlined boxes in Year n can be calculated in terms of the probabilities for Year $n - 1$. Since this produces a system of linear equations, the probability of the system being in each of these states will converge toward a fixed value with increasing time. These values are shown in bold italics, next to each of the six states, and it is straightforward if laborious to check that these values are consistent and unique. It can be seen that the system will in the long run spend a smaller proportion of Years with both Land Parcels owned by LD Land Managers ($\frac{5}{22} + \frac{1}{12} = \frac{41}{132}$), than with both owned by FS Land Managers ($\frac{2}{11} + \frac{23}{132} = \frac{47}{132}$).

Figure 5 shows the dynamics of the same system without multi-Parcel Estates. In this case, although FS and LD are not functionally equivalent — if both cells are owned by LD Land

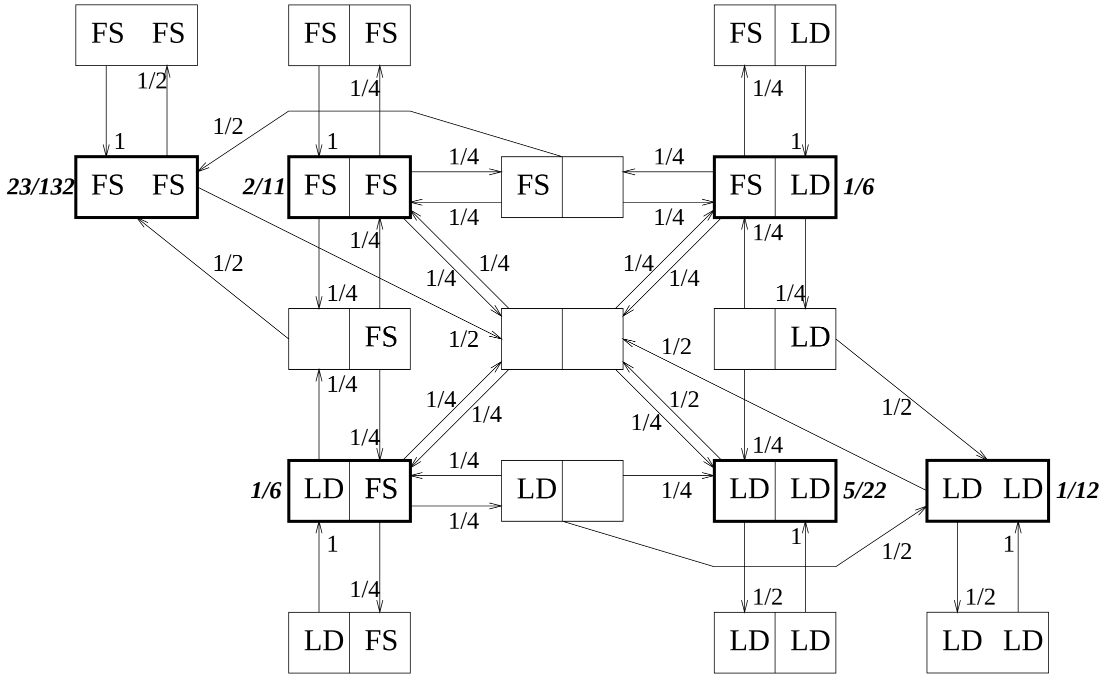

Figure 4: States and state-transitions of an LD/FS contest in a two-cell FEARLUS Environment with multi-Parcel Estates.

Table 6: Effects of Intra-Estate and Inter-Estate Diversity in P0-E16u-LPP0, P0-E16u, and P0-E16u-LPP2000.

LPP	LD/FS		FS/RS	
0	170	*310****	196	*284****
16	183	*297****	224	*256*
2000	232	248	241	239

Managers they always make the same choice, while two FS Land Managers do so only half the time — it turns out that the system spends equal amounts of time in these two states.

The preceding analysis, bringing out the possibility that either Land Use diversity within the Estates of individual Land Managers, or diversity between Estates, could be responsible for the simulation results, prompted a new set of experiments, testing FS against both LD and RS in the same three Environments used in the earlier experiments. Results, shown in table 6, suggest that both factors are operating: all predicted results are in the expected direction and although one (FS against RS in P0-E16u) is short of significance, the corresponding result in P0-E16u-LP0, where the effects are expected to be enhanced, is highly significant.

5 Conclusions

We have aimed in this paper to demonstrate by example that the analysis of simple formal models, and simulation studies on more complex ones, should not be regarded as alternative and even opposed approaches to the formal study of social systems, but as complementary. Each can provide both problems and pointers to solutions for the other. By the same token, work with highly constrained models (such as those obeying all or most of the rules for standard CA in

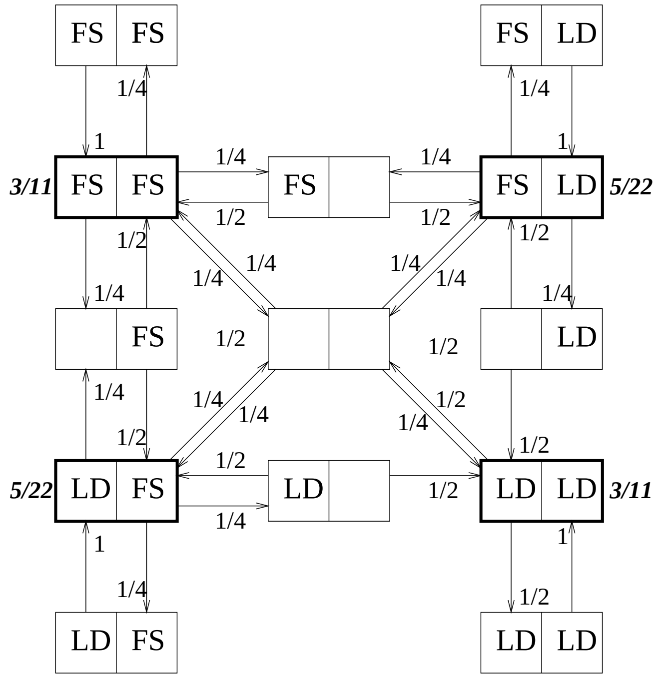

Figure 5: States and state-transitions of an LD/FS contest in a two-cell FEARLUS Environment without multi-Parcel Estates.

section 3), and with models in which some of these rules are relaxed, can and should be combined. Nor is it necessarily the case that research should begin with analysis, and resort to simulation only when analysis fails: even Binmore (1998), who severely criticises Axelrod's simulation-based work (Axelrod 1984, Axelrod 1997) on the 'Prisoner's Dilemma' for neglecting analytical results already available within game theory, admits that this same work has greatly stimulated work on evolutionary game theory. The direction our own work has taken so far has been to begin primarily with simulation, then to start analysis from those points where phenomena of particular interest became evident. Of course, even our simulation models are a long way from realism. Kliemt (1996) makes a distinction between *thin* and *thick* agent-based simulations. The former are a tool for 'controlled speculation', useful in disciplining theory formation; they often use simplifying and distorting assumptions. The latter are detailed, draw on copious empirical data, and tell the investigator a lot about a specific question (such as the optimum distribution of sales staff in a specific retail outlet), but only about that. Our view is that the thin/thick distinction is a continuum rather than a dichotomy, with mathematically tractable models at one end, and models based on and contributing to detailed case studies at the other. Those using computational tools to study social systems (and indeed, other complex systems) should be ready to slide up and down this continuum: adding complexity to a general model in order to capture key features of a domain of interest, and abstracting away some features (not necessarily the ones recently added), when an interesting or puzzling phenomenon is discovered, in order to find the simplest systems and mechanisms capable of generating it.

Let us briefly review the work reported in this light. The main findings divide into three:

1. The simulation finding that an Aspiration Threshold at or very near the BET works best in most FEARLUS Environments is indicated by analysis to hold in Environments without multi-Parcel Estates or spatial or temporal heterogeneity. Recall, however, that simulation also indicates that lower Aspiration Thresholds work better (for HR, HYI and HRYI) in Environments with much temporally uncorrelated variation, and higher ones (for HSBI and HRSBI) in Environments with considerable *correlated* temporal variation and restricted spatial variation. We plan to extend our analytical work to models including temporal

variability, in order to improve our understanding of this phenomenon.

2. Simulations showed HRYI and HYI as generally more similar to each other in behaviour than either was to HR. The analysis of section 3.2, however, suggests that HRYI is similar to HR and substantially different from HYI in the range of qualitatively different possibilities that can arise in spatio-temporally homogeneous Environments without multi-Parcel Estates — and the same would apply to HRSBI and HSBI. An interesting possibility suggested by this contrast is to design and run simulation experiments to pin down further the respects in which HRYI resembles HR, those in which it resembles HYI, and those in which it is unlike either. The relationships between HR, HRSBI and HSBI could be similarly investigated.

3. The work reported in section 4 perhaps demonstrates most clearly the kind of interplay between simulation and analysis we aim to achieve. Simulation experiments turned up a puzzling phenomenon; analysis suggested two possible mechanisms underlying it; further simulation indicated that both of these were contributing to the effects found. Although the effect itself is relatively weak, and was found in Environments with the unusual property that all Land Uses are equally good, it is important that we be aware of it as a possible factor in other Environments.

The work reported falls near the 'thin' end of the continuum between generality and tractability on the one hand, and realism and applicability on the other. How widely relevant are FEARLUS models and the research we have undertaken with them? Key features of the models implemented so far could be listed as: explicit representation of space, spatial and temporal heterogeneity, ownership or control of territory which may be used or exploited in different ways and potentially lost or gained, and the possibilities of experimenting, and of imitating neighbours. We have concentrated on the ways in which different degrees and kinds of spatio-temporal heterogeneity, and the possibility of acquiring territory, affect the relative success of different strategies for choosing how to use the territory you have. Potentially, we believe this work is relevant to a wide range of areas in urban land use, ethology and animal behaviour, competition between plants, history, anthropology, and political science (see Cioffi-Revilla and Gotts (2003), where such *territorial resource allocation processes* and agent-based models of them are also discussed in general terms).

Acknowledgements

We gratefully acknowledge statistical and mathematical advice from Mark Brewer of Biomathematics and Statistics Scotland. Any errors in this area of course remain our responsibility. This work was funded by the Scottish Executive Environment and Rural Affairs Department, to whom we express our thanks for their support.

References

Axelrod, R.: 1984, *The Evolution of Cooperation*, Basic Books.

Axelrod, R.: 1997, *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton Studies in Complexity, Princeton University Press.

Benson, J. F. (ed.): 1995, *Journal of Environmental Planning and Management. Special Issue: The NERC/ESRC Land Use Modelling Programme (NELUP)*, Vol. 38 (1).

Binmore, K. G.: 1998, Review of "The Complexity of Cooperation" by Robert Axelrod, *Journal of Artificial Societies and Social Simulation.* Online journal, at http://www.soc.surrey.ac.uk/JASSS/JASSS.html.

Chattoe, E.: 1996, Why are we simulating anyway? Some answers from economics, *in* K. G. Troitzsch, U. Mueller, N. Gilbert and J. E. Doran (eds), *Social Science Microsimulation*, Springer, pp. 78–104.

Cioffi-Revilla, C. and Gotts, N. M.: 2003, Comparative analysis of agent-based social simulations: GeoSim and FEARLUS models, *Journal of Artificial Societies and Social Simulation* **6**(4), article 10. Available online at http://jasss.soc.surrey.ac.uk/6/4/10.html.

Gotts, N. M., Polhill, J. G. and Law, A. N. R.: 2002, Aspiration levels in a land use simulation, To appear in Cybernetics and Systems; a draft is available online at http://www.macaulay.ac.uk/fearlus/FEARLUS-publications.html.

Gotts, N. M., Polhill, J. G., Law, A. N. R. and Izquierdo, L. R.: 2003, Dynamics of imitation in a land use simulation, *in* K. Dautenhahn and C. L. Nehaniv (eds), *Proceedings of the AISB '03 Second International Symposium on Imitation in Animals and Artifacts, 7-11 April, The University of Wales, Aberystwyth*, The Society for the Study of Artificial Intelligence and Simulation of Behaviour (SSAISB), pp. 39–46. Also available online at http://www.macaulay.ac.uk/fearlus/FEARLUS-Publications.html.

Grimmett, G. R. and Stirzaker, D. R.: 1992, *Probability and Random Processes*, second edn, Oxford Science Publications.

Kliemt, H.: 1996, Simulation and rational practice, *in* R. Hegselmann, U. Mueller and K. G. Troitzsch (eds), *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View*, Kluwer, chapter 2, pp. 13–28.

Parry, M. L.: 1996, Integrating global and regional analyses of the effects of climate change: A case study of land use in England and Wales, *Climate Change* **32**, 185–198.

Polhill, J. G., Gotts, N. M. and Law, A. N. R.: 2001, Imitative versus non-imitative strategies in a land use simulation, *Cybernetics and Systems* **32**(1-2), 285–307.

Pomp, M. and Burger, K.: 1995, Innovation and imitation: Adoption of cocoa by Indonesian smallholders, *World Development* **23**(3), 423–431.

Simon, H. A.: 1955, A behavioral model of rational choice, *Quarterly Journal of Economics* **69**, 99–118. Reprinted as Ch.14, pp.241-260 in Simon, H.A. (1957) Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting, John Wiley and Sons, New York.

Simon, H. A.: 1957, *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*, John Wiley and Sons, New York.

Veldkamp, A. and Lambin, E. F.: 2001, Predicting land use change, *Agriculture, Ecosystems and Environment* **85**, 1–6.

White, R., Engelen, G. and Uljee, I.: 1997, The use of constrained cellular automata for high-resolution modelling of urban land-use dynamics, *Environment and Planning B: Planning and Design* **24**, 323–343.