Dynamics of Imitation in a Land Use Simulation

N.M. Gotts
n.gotts@macaulay.ac.uk

J.G. Polhill
g.polhill@macaulay.ac.uk

A.N.R. Law
a.law@macaulay.ac.uk

L.R. Izquierdo
l.izquierdo@macaulay.ac.uk

Macaulay Institute, Craigiebuckler, Aberdeen AB15 8QH, Scotland

Abstract

The paper concerns the socio-spatial dynamics of imitation within a computational model of land use selection and change. Specifically, it reports investigations of the success of imitation in relation to alternative ways of choosing a course of action, in the context of different degrees and kinds of spatio-temporal heterogeneity. Simulation experiments with the model are the main method employed, but analytical work is also reported.

1 Introduction

This paper is about the spatial dynamics of imitation in (relatively) simple models of complex socio-economic systems: specifically, spatially explicit agent-based social simulations of land use change. The work described forms part of the FEARLUS (Framework for Evaluation and Assessment of Regional Land Use Scenarios) project. FEARLUS is aimed at increasing understanding of the processes underlying rural land use change, particularly at the regional scale and in the medium to long term. FEARLUS 'agents' are land managers, who choose land uses, and may copy their neighbours' choices. (Some might prefer the term 'social learning' for the phenomena we discuss — copying others' high-level decisions, rather than specific physical actions — but 'imitation' is used in this way within fields including social psychology, game theory, and agricultural economics.)

Those making decisions on land use may be influenced in various ways by their neighbours (and wider social influences): the most obvious include imitation based on the success of innovative land uses or techniques (and conversely, avoidance of innovations seen to fail). Imitation is one way to economise on computational resources, and/or compensate for an absence of knowledge, and is known to be one method land managers use in choosing what to do (Pomp and Burger, 1995). Our work has focused on the dynamics of different strategies for land use selection, many of them involving imitation, in environments demanding different levels of performance in order to remain solvent, and with different degrees and kinds of spatial and temporal heterogeneity.

There is an extensive literature on formal models of the spatial dynamics of imitation in games such as the 'Prisoner's Dilemma', where neighbours' choices determine an agent's payoffs from their own (Gotts et al., 2003). In land use decision-making, while neighbours may affect each others' payoffs, the local suitability of the land, and the (generally changeable) climatic and economic conditions are of comparable or greater importance.

Most empirical studies of imitation in rural land use change concern the adoption of exogenously produced technical innovations, and, frequently, the 'barriers' to their adoption (Abadi Ghadim and Pannell, 1999). As Cramb et al. (1999) note, this approach tends to assume adoption to be the correct choice; however, resistance may be based on good grounds, including 'objective differences in soil and farming conditions' (p.420). Agent-based studies of the adoption of agricultural innovations include Weisbuch and Boudjema (1999) and Berger (2001). Both these studies focus on the speed and completeness of innovation-adoption processes: how fast and thoroughly techniques new to a region will diffuse through it. By contrast, work with FEARLUS has concentrated on the relative success of land managers with different approaches to choosing among a set of existing land uses, in fluctating economic or climatic conditions (meaning, for example, that shifting from one land use to another and back several times could be the best course of action). This kind of 'fluctuation tracking' among a spatially distributed population capable of imitating neighbours has not previously been modelled, to our knowledge.

2 Methods

The primary approach taken in the FEARLUS project is computer simulation, but some analytical work is also reported here: simulation and analysis can each provide both problems and actual or possible solutions to the other. Our approach to simulation makes considerable use of multi-run simulation *experiments*, with statistical testing of predictions concerning the outcomes. Simulations may quite legitimately be used simply to show that a model system *can* demonstrate a particular form of behaviour. If the model has any stochastic elements, however (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. 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 FEARLUS model consists of a set of *Land Managers*[1], and their *Environment*, which includes a grid of square *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*. These represent a combination of economic and climatic factors. They are encoded as a bitstring, the length of which is a model parameter. The External Conditions can vary from Year to Year but apply across the whole grid. Generally the initial bitstring is determined randomly, and each subsequent bitstring is produced from its predecessor by applying 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 have positive temporal auto-correlation. In the Environments with temporal auto-correlation used here, $f = \frac{1}{8}$.

Each Land Parcel has a set of *Biophysical Characteristics*, encoded as a bitstring and fixed for the duration of a simulation run (again, the length of these bitstrings is a model parameter; it is the same for all Land Parcels). The Biophysical Characteristics of Land Parcels may be either 'clumped' (spatially auto-correlated) 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, 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.

There are also two numerical parameters which do not vary over space *or* time: a *Break Even Threshold* (BET), specifying how much *Yield* (economic return) must be gained from a Land Parcel to break even, and the *Land Parcel Price* (LPP). In all experiments discussed here, the Land Parcels are arranged in a 7×7 chequerboard lattice, with opposite sides joined to produce a toroidal topology; the BET is 8 and, except where otherwise specified, the LPP is 16.

The Environments used in experiments reported here will generally 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' 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' 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 *Spatio-Temporal Heterogeneity Type* (STHT). When an Environment has an LPP other than 16, this will be indicated by adding a suffix: 'P0-E16u-LPP2000' indicates an environment with no Land Parcel Characteristic bits, 16 uncorrelated External Conditions bits, and an LPP of 2000.

A FEARLUS simulation 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 External Conditions, against one representing requirements of the current Land Use, and counting the 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*. A Land Manager whose Account is in deficit sells their worst-performing Parcels to clear the deficit. Land Managers unable to do so while retaining at least one Parcel, leave the simulation.

6. *Sale of Land Parcels*. The selected Land Parcels are sold either to a neighbouring Land Manager, or to a new Land Manager entering the simulation.

Most of the experiments discussed here consisted of a number of simulation runs, pitting two *Subpopulations* against each other. 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 of the 49 Land Parcels is assigned to a different Land Manager. After 200 Years, Subpopulation success is assessed by counting the Land Parcels assigned to members of each. In addition to these 'type 1' experiments, a few 'type 2' experiments are reported. These compare the performance of two Selection Algorithms against a third 'comparison Algorithm' in a given

[1] Some terms used to refer to elements of FEARLUS models could also refer to real-world entities. In the names of FEARLUS model elements, each word begins with an upper-case letter (e.g, 'Land Manager'). Each such term is italicised when first used.

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 significantly better against the comparison algorithm than the other.

In earlier experiments (Polhill et al., 2001), we matched Selection Algorithms involving the imitation of neighbours against Selection Algorithms of other kinds, and each other. We found that competitive advantage between pairs of strategies can depend on the type and extent of spatio-temporal heterogeneity present, and that competitive superiority between strategies is not always transitive even within a single type of Environment. Here, we vary the model Environments more extensively and systematically, and use a non-imitative Algorithm more closely related to its imitative rivals than any we used previously. This allows more accurate assessment of when imitation is useful.

The simplest way to choose a land use for a land parcel from a set of alternatives is to maintain the current one. The main Land Use Selection Algorithms discussed here use an *Aspiration Threshold* (Simon, 1955). The Land Manager looks at whether the Yield a Land Parcel produced in the preceding Year equalled or exceeded the Aspiration Threshold and if so, sticks with the same Land Use for that Land Parcel. Otherwise, some other procedure is used to select the Land Use. In the Selection Algorithms we focus on here, this always involves either *Random Experimentation* (a random choice between the possible Land Uses, all having equal likelihood of being selected), or *Yield-weighted Imitation*.

To apply Yield-weighted Imitation to a Land Parcel, a Land Manager constructs an *Imitation List* of all Parcels it owns, plus those owned by neighbours. It then sums the Yield produced by each Land Use across all the Imitation List Parcels it was used on in the most recent Year, and makes a random choice among those Land Uses, weighted by the Yield totals. This is one of a range of alternative ways to implement imitation of neighbours. In our experiments so far, it is at least as useful to Land Managers, across a wide range of FEARLUS Environments, as any comparably simple approach tried.

Three families of Aspiration Threshold Selection Algorithms were used. Members of a family differ only in the level of their Aspiration Threshold:

HR: Land Managers employing a *Habit/Random Selection Algorithm* always use Random Experimentation if their Aspiration Threshold is not met.

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

HRYI: Those employing a *Habit/Random/Yield-weighted-Imitation Selection Algorithm* choose stochastically whether to use Random Experimentation (with probability $\frac{1}{16}$) or Yield-weighted Imitation (probability $\frac{15}{16}$) if their Aspiration Threshold is not met.[2]

Experiments have been carried out on the optimum value for the Aspiration Threshold (Gotts et al., 2002). Over a wide range of Environments, Thresholds around the level of the BET do best, and a Threshold of 8 is used in all cases here.

3 Results

3.1 Imitation *Vs* Random Choice

We ran type 1, 120-run experiments pitting each of HR, HYI and HRYI against the others across a set of 23 Environments, all with BET 8 and LPP 16. All had 16 bits in the Land Use requirements bitstring, but differing in how these bits were divided between Biophysical Characteristics and External Conditions, whether the Biophysical Characteristics of Parcels were clumped, and whether External Conditions were temporally auto-correlated.

Results are given in table 1. Here, and in table 2, columns are headed by the names of the two competing Selection Algorithms, while rows correspond to Environment types. The number of 'wins' for the two Algorithms in various types of Environment are given in the cells of the column ('wins' for an Algorithm are runs in which the Subpopulation using it ended up with more Land Parcels). If one of the Algorithms was predicted to do better than the other (on the basis of exploratory experiments), 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.

Table 1 results can be summarised as follows. In contests between HYI and HR, HR generally won in environments with both unclumped spatial variation and either uncorrelated, or very little, temporal variation. HYI tended to win in Environments with clumped spatial variation, or none at all, especially if they also had correlated temporal variation. The differences between Environments are quite comprehensible: imitation should be more favoured when nearby Parcels are more similar, and when there is change which can be *tracked* (next Year is likely to be similar to this Year). Results in contests between HRYI and HR were almost the same as for HYI against HR. As between HYI and HRYI, results suggest HRYI has a slight advantage in some Environments: first, those with unclumped spatial variation and a lot of uncorrelated temporal variation (HR outperforms HYI and

[2] The precise value of $\frac{1}{16}$ is arbitrary, but that it is small is not. Earlier work (Polhill et al., 2001) suggested that a small admixture of Random Experimentation could make a big difference to HYI in some circumstances.

Table 1: Type 1 experiments: direct contests between HR, HRYI and HYI Subpopulations

STHT	HR/HRYI		HR/HYI		HRYI/HYI	
P0-E16c	29	*91****	36	*84****	69	51
P1c-E15u	45	*75**	49	71	64	56
P1c-E15c	41	*79***	44	*76**	78**	42
P1u-E15u	*74**	46	*81****	39	70	50
P1u-E15c	43	*77**	50	*70*	*56*	64
P2c-E14u	51	*69*	49	71	55	65
P2c-E14c	35	*85****	28	*92****	63	57
P2u-E14u	*84****	36	*85****	35	72	47
P2u-E14c	35	*85****	50	*70*	*62*	58
P4c-E12u	36	*84****	45	*75**	68	52
P4c-E12c	40	*80***	44	*76**	61	59
P4u-E12u	*86****	34	*79***	41	72	48
P4u-E12c	51	69	56	64	*64*	56
P8c-E8u	49	*71*	41	*79***	58	62
P8c-E8c	34	*86****	46	*74**	60	60
P8u-E8u	*77**	43	*76**	44	58	62
P8u-E8c	*54*	66	*63*	57	58	62
P12c-E4u	50	*70*	41	*79***	51	69
P12c-E4c	54	*66*	60	60	64	56
P12u-E4u	*70*	50	*80***	40	71	49
P12u-E4c	*73*	47	*78***	42	68	52
P16c-E0	66	54	59	61	65	55
P16u-E0	*70*	50	*70*	50	64	56

HRYI in such Environments, suggesting that Random Experimentation is superior to Yield-weighted Imitation in such cases), and second, those with very little or no spatial variation and a lot of auto-correlated temporal variation (P0-E16c and P1c-E15c, where HRYI also outperforms HR). A pair of 240-run repeat experiments were performed in these Environments, with HRYI predicted to win: HRYI won in 147 runs in P0-E16c (significant at the .001 level), and 131 runs in P1c-E15c (not significant, but in the expected direction).

HRYI probably outperforms HYI in the second set of Environments because populations composed wholly of HYI Land Managers tend to become locked in to a restricted set of Land Uses. Since an HYI Land Manager *never* adopts a Land Use not in use on any of the Parcels on the Imitation List, an Environment occupied only by such Land Managers tends to lose Land Uses. If HYI Land Managers gain control of all Land Parcels and this occurs, a subsequent change in External Conditions could lead to most or all of them becoming bankrupt within a short time, undermining their dominance; in the analogous situation, a group of HRYI Land Managers would, due to their occasional use of Random Experimentation, be able to track the changing External Conditions and maintain their dominance.

Type 2 experiments were used to compare the performance of HR, HYI, and HRYI against various Selection Algorithms involving neither imitation nor Aspiration Thresholds, but relying on *Innate Knowledge* of which Land Uses do best in various circumstances. Each experiment pitted two of the three against an Innate Knowledge Selection Algorithm in 120 pairs of corresponding Environments. The most interesting results arose in Environments P0-E16c and P1c-E15c, against the *Optimum-match Deterministic Selection Algorithm* (OD). This selects a Land Use from among those with an optimal match to the Land Parcel's Biophysical Characteristics bitstring. However, the Land Uses available in any model run are numbered, and OD always selects the lowest-numbered member of this subset. All Land Managers using OD thus share the same preference order. Results for the OD experiments are given in table 2: this records the number of pairs of runs in which each of the competing Algorithms did better against OD. HR and HRYI both did much better than HYI against OD in P0-E16c and considerably better in P1c-E15c. Neither HR nor HRYI had a clear advantage against the other in these cases. This pattern of differences can presumably be attributed to the permanent loss of Land Uses from the Environment in OD/HYI contests. Imitating OD appears to be positively harmful, at least in P1c-E15c, to judge by the contrast between the direct contest between HR and HYI (which the latter wins), and the performance of each of these two against OD (where HR does considerably better).

Table 2: Type 2 experiments: HR, HRYI, and HYI performance against OD compared

STHT	HR/HRYI		HR/HYI		HRYI/HYI	
P0-E16c	*27*	14	*102****	7	*94****	13
P1c-E15c	51	55	76**	39	*74**	41

In its performance against OD, HRYI thus resembles HR more than it does HYI, in contrast to what was found in direct contests between the three in most Environments examined. Given this complex picture, it is worth asking what analytical approaches can tell us.

Given the way Yields are calculated, every combination of a Land Parcel c and a Land Use u gives rise to a range of possible Yields. The lowest element of this range of integers is $Y_{l(c,u)}$, the number of matches between the Land Parcel's Biophysical Characteristics bitstring and the corresponding bits of the Land Use requirements bitstring; the highest is $Y_{m(c,u)} = Y_{l(c,u)} + e$ (e is the length of the External Conditions bitstring). When all Land Managers have an Aspiration Threshold equal to the BET (as in the table 1 experiments), Land Parcel-Land Use pairs can be classified into three qualitative kinds: those where $Y_{l(c,u)}$ and $Y_{m(c,u)}$ are both below the BET (call these 'BB' pairs), those where $Y_{l(c,u)}$ is below the BET but $Y_{m(c,u)}$ is at or above it ('BA' pairs), and those where both are above ('AA' pairs). This in turn gives rise to a qualitative classification of Land Parcels themselves, according to whether or not there are Land Uses with each of these three relationships with the Parcel. For each of the seven resulting types, we can say a significant amount about what will happen to Parcels of that type, if we adopt the simplification that Land Managers cannot own *Estates* of more than one Parcel, so that a land sale is always to a new Land Manager:

1. BB. The Land Use will be reselected, and ownership will change, every Year.
2. BA. The Land Use will be reselected repeatedly (whenever Yield falls below the BET) but not every Year. (Strictly, the probability that at least n reselections will have occurred, for any n, will approach 1 with time.) Whether this holds for ownership changes depends on details of the Land Uses' Yields, whether External Conditions are temporally auto-correlated, and the probabilities that new Land Managers will belong to each of the three types. There are three qualitatively distinct possibilities: the expected lifespan (time to bankruptcy) of a Land Manager may be finite, this lifespan may be infinite, but with the probability of replacement approaching 1, or the probability of bankruptcy may approach a limit < 1.
3. AA. The Land Use will never be reselected, ownership will never change.
4. BB-BA. As for BA Parcels.
5. BB-AA. The Land Use will be reselected and ownership changed every Year until an AA Land Use is found. Thereafter, Land Use will never be reselected, nor will ownership change.
6. BA-AA. The Land Use will be reselected repeatedly (whenever the Yield falls below the BET) but not every Year, until an AA Land Use is found. Ownership changes may or may not occur before this point. Thereafter, Land Use will never be reselected, nor will ownership change.
7. BB-BA-AA. As for BA-AA Parcels.

Relating this classification to the Spatio-Temporal Heterogeneity Types of the experimental Environments, the P0.., P1.., P2.. and P4.. Environments must all consist entirely of type BA Land Parcels, since the range of Yields of any Land Use will include Yields both above and below the BET. For the P8.. Environments, there can be no BB Land Parcel-Land Use pairs (if all the External Conditions bits are matched, the Yield must at least equal the BET), so all Parcels must belong to types BA, AA, or BA-AA. P12.. Environments can contain Land Parcels of all seven types. Finally, for the P16.. Environments, all Parcels must be of types BB, BB-AA, or AA.

In an Environment consisting wholly of BB, BA and BB-BA Parcels (call these 'Low-Yield'), a Population consisting entirely of HR and/or HRYI Land Managers would result in qualitatively different long-term behaviour from a population consisting entirely of HYI Land Managers. In the former case, the probability that every Land Parcel has been assigned every Land Use at least n times will approach 1 with time. In a Population of HYI Land Managers, however, the probability of all Parcels being assigned the *same* Land Use, and keeping that Land Use thereafter, will approach 1. Related differences will occur in some Environments consisting of a mixture of Low-Yield with 'High-Yield' Parcels (types AA, BA-AA, BB-AA and BB-BA-AA): whatever the Population, each High-Yield Parcel will (with probability approaching 1) settle in one of the Land Uses guaranteeing at least the BET. The Low-Yield Parcels will adopt every Land Use repeatedly, if the Population consists of HR and/or HRYI Land Managers; in a Population of HYI Managers, however, the probability that only those Land Uses permanently assigned to a High-Yield Parcel remain in use will approach 1. The differences described in this paragraph, as they depend only on the relationship between Yield and Land Managers' Aspiration Threshold, hold even if multi-Parcel Estates are permitted, and are also independent of the BET.

Taken together, the simulation and analytical results described above indicate that HR, HRYI and HYI are all qualitatively distinct in their behaviour: there is not a gradation in behaviour from always imitating when reselecting a Land Use, through sometimes doing so, to never doing so. No circumstances have been found in which HYI is clearly superior to HRYI — although one such circumstance did appear in the work reported in Polhill et al. (2001): when matched against a Subpopulation with a more sophisticated (and computationally expensive) approach to the choice of an imitation target ('Intelligent Imitation', or II), which also lacked any random element. In this case, the loss of Land Uses over time 'flattened' the differences between the more and less sophisticated approaches to imitation.

3.2 Diversity When All Yields Are Equal

One Environment type used in exploratory experiments, P0-E16u, was included as a check that spurious results were not being generated. It is spatially homogeneous, since Parcels have no Biophysical Properties; External Conditions are variable and temporally uncorrelated. Any Land Use (hence any Selection Algorithm) gives the same expected Yield, and the same expected distribution of Yields over a period of Years, on any Parcel. However, exploratory experiments suggested that some Selection Algorithms systematically outperformed others. RS, a 'Selection Algorithm' that always uses Random Experimentation, did particularly well. Further experiments 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 this success.

To investigate this phenomenon systematically, we chose four Selection Algorithms: RS, HR, HRYI and *Last-year's-optimum-match Deterministic Algorithm* (LD). This Algorithm 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 available in a run are numbered, and LD selects the lowest-numbered among those which would have maximised 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. As the effect under study appeared weak, we matched each of the four Selection Algorithms against the other three in 480-run type 1 experiments. As predicted, RS beat HR, both these beat HRYI, and all three of these beat LD.

We hypothesized that the land sale process included in the model might underlie the phenomenon. Setting the LPP so high (at 2,000) that no Land Manager can ever afford to buy up a neighbour's Land Parcel (thus any Parcel sold goes to a new Manager) abolished the effect; setting the LPP to zero enhanced it. This confirms that it does depend on the land sale mechanism, but how is it caused? There are at least two possibilities. First, once Land Managers gain control of multi-Parcel Estates, perhaps greater diversity of Land Uses reduces the likelihood of an *individual* ending up with a negative Account, and thus losing Parcels. Second, *Subpopulations* in which different Land Managers favour different Land Uses may be 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 no Managers 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 analytically 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 Environment (which is spatially homogeneous with uncorrelated temporal variation, like P0-E16u), with a BET of $\frac{1}{2}$, 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, while these values are reversed for Land Use 0. The LPP need not be specified — any finite and non-negative value will do. As in the experimental simulations, assume there are two Subpopulations, equally likely to provide Land Managers both initially and as replacements. If one of these Subpopulations uses RS and the other FS, there will in the long run 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 owners, the dynamics will be the same whether both use RS, both use FS, or one uses each. If both Parcels have the same owner, however, matters are different. This can only occur as a result of the owner of one Parcel buying up the other when its Land Manager goes bankrupt. Assume that at that point it has an amount n in its Account. We can assume without loss of generality that n is an integer, since if it has a fractional part, this will make no difference to subsequent events.

If the Land Manager uses FS, choosing a Land Use at random each Year to apply to both Land Parcels, it will either gain or lose 1 every Year. Random walk theory (Grimmett and Stirzaker, 1992) shows that the Land Manager's Account is certain to go into deficit eventually (forcing the sale of both Parcels and returning the system to its starting state), but that the mean time for it to do so is infinite.

Figure 1 shows part of the state diagram of the system described, including (at the top) the first few of the infinite set of states in which both Parcels are owned by a single FS Land Manager, and (in the middle row) the four states outside this set but adjacent to FS-0, the state in which the FS Land Manager has an empty Account. (Transition into one of these states corresponds to bankruptcy for the FS Land Manager and the consequent recruitment of two new Managers; they are not endpoints, and the existence of transitions *from* them, and other transitions into them, is indicated by the dashed arrows.) Taking, say, the state labelled FS-3 as a starting point, consider how the probability of the state having entered one of the four states in the middle changes over time. For the first three Years it is 0, after four Years it is $\frac{1}{16}$, and thereafter it increases every two Years, approaching 1 over time. In general, from state FS-n, the probability will rise above 0 in Year $n+1$, then increase in Years $n+3, n+5, n+7 \ldots$.

Now consider the lower part of figure 1, which shows the corresponding part of the system 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 a $\frac{1}{2}$ chance that if the system is in state RS-n, it will follow the 'loop' transition and remain there.

- Consider all the paths which start at FS-n and end with the first visit to the middle row. There are an infinite number, but each has a non-zero probability of occurring. We can calculate the probability that the system will have completed such a path after no more than y Years, for any y. (This will exceed 0 if

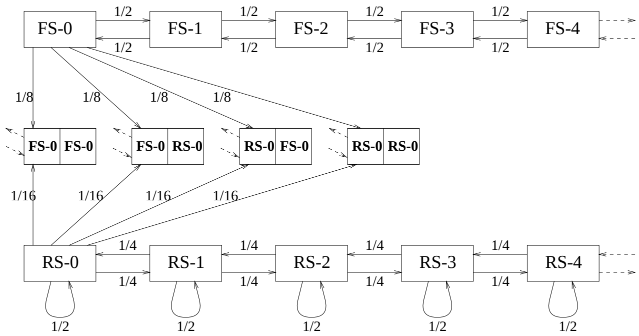

Figure 1: States and state-transitions of an FS/RS contest in a two-cell FEARLUS Environment with multi-Parcel Estates.

and only if $y > n$. Random walk theory shows that the probabilities tend to 1 as y tends to infinity.)

- To each of these paths, there corresponds an infinite set of paths starting at RS-n and ending with the first visit to the middle row. Each member of this set includes the same sequence of non-loop steps as the corresponding path from FS-n; they differ in the number of loop transitions which occur at the start, and/or between the other steps of the sequence.
- The probability that the system's path to the middle row from RS-n will lie within any one of these infinite sets is equal to the probability that it would follow the corresponding path from FS-n to the middle row.
- However, only one member of such an infinite set takes the same time to reach the middle row as the corresponding path from FS-n; all the rest take longer, and for any specified amount of extra time, all but a finite number take more than that amount extra.
- If we consider *all* the paths from FS-n that reach the middle row in y or fewer Years, *only* paths among the corresponding infinite sets from RS-n could possibly do the same, and for each of the paths from FS-n, all but a finite number fail to do so.
- Thus for any $y > n$, the probability of the system reaching the middle row in y or fewer Years is less from RS-n than from FS-n.

If land sales to existing Managers were disallowed, RS and FS would be functionally equivalent, so this would not be the case.

Figure 2 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.

The six heavily outlined boxes represent the possible states of the system just before Land Uses are selected — distinguished only by whether the two Land Parcels belong to the same Land Manager (a line between the Parcels shows they are owned by different Managers), and the Subpopulation to which each Manager belongs. The remaining boxes represent the possible transitional states of the system immediately before any Land Parcels without a solvent Manager (shown as empty) are assigned one. The system begins in the central box. LD Managers will always choose the same Land Use — say Land Use 0. FS Land Managers will assign Land Use 0 or 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 state will converge toward a fixed value with increasing time. These values are shown in bold italics, beside the six states. 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 it will spend with both owned by FS Land Managers ($\frac{2}{11} + \frac{23}{132} = \frac{47}{132}$).

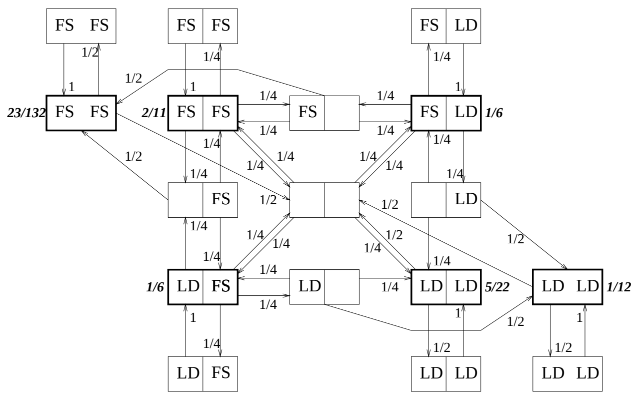

Figure 2: States and state-transitions of an LD/FS contest in a two-cell FEARLUS Environment with multi-Parcel Estates.

Figure 3 shows the dynamics of the same system without multi-Parcel Estates. FS and LD are not functionally equivalent — if both cells are owned by LD Land Managers they always make the same choice, while two FS Managers do so only half the time — but it turns out that the system spends equal amounts of time in these two states.

The preceding analysis, bringing out the possibility that

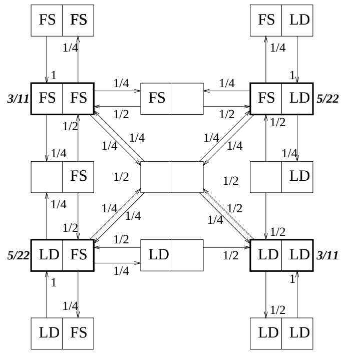

Figure 3: States and state-transitions of an LD/FS contest in a two-cell FEARLUS Environment without multi-Parcel Estates.

either Land Use diversity within the Estates of individual Land Managers, or diversity between Estates, could be responsible for the simulation results, prompted more experiments, testing FS against both LD and RS in P0-E16u Environments, with each of the three LPP settings used earlier. Results suggested that both factors are operating.

4 Conclusions

The work described here suggests that imitation of neighbours, even if weighted toward land uses that have been successful in the neighbourhood, will not invariably prove superior to random choice among all possible alternatives as an approach to land use selection. Rather, its appropriateness depends on the nature of any spatio-temporal heterogeneity in the environment: spatial variation between nearby parcels will reduce the usefulness of imitation, while trackable temporal variation in conditions will favour it. Moreover, in some environments, mixing imitation and random experimentation may be superior to using either on its own. These findings suggest that empirical studies of imitation of neighbours should find corresponding differences in the prevalence of imitation, depending on the environment's patterns of heterogeneity. In our own work, studies of agents able to adapt their probability of imitation (and other factors such as their Aspiration Threshold) are planned.

More generally, within the wider context of research on imitation in animals and artifacts, the work reported draws attention to the need to study the processes involved in imitation at the level of social dynamics, as well as that of individual cognition; and to study its interactions with other ways of acquiring and selecting among possible behaviours or courses of action. Imitation is likely to be most important when it is most beneficial to the imitator, and this means when good targets for imitation are available: agents facing similar problems, but knowing more about the possible solutions.

The paper also illustrates the interplay between simulation and analysis we aim to achieve, particularly in the work on P0-E16u Environments. Here, simulation experiments turned up a puzzling phenomenon; analysis showed that two possible mechanisms underlying it could both occur in some cases; further experiments indicated that both were contributing to the effects found.

Acknowledgements

This work was funded by the Scottish Executive Environment and Rural Affairs Department, to whom we express our thanks for their support.

References

Amir K Abadi Ghadim and David J Pannell. A conceptual framework of adoption of an agricultural innovation. *Agricultural Economics*, 21:145–154, 1999.

Thomas Berger. Agent-based spatial models applied to agriculture: a simulation tool for technology diffusion, resource use changes and policy analysis. *Agricultural Economics*, 25:245–260, 2001.

R A Cramb, J N M Garcia, R V Gerrits, and G C Saguiguit. Smallholder adoption of soil conservation technologies: evidence from upland projects in the Philippines. *Land Degradation and Development*, 10:405–423, 1999.

Nicholas M Gotts, J Gary Polhill, and Alistair N R Law. 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, 2002.

Nicholas M Gotts, J Gary Polhill, and Alistair N R Law. Agent-based simulation in the study of social dilemmas. *Artificial Intelligence Review*, 19:3–92, 2003.

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

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

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

Herbert A Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, February 1955.

G Weisbuch and G Boudjema. Dynamical aspects in the adoption of agri-environmental measures. *Advances in Complex Systems*, 2:11–36, 1999.