Aspiration Levels in a Land Use Simulation

N.M. Gotts, J.G. Polhill and A.N.R. Law

July 17, 2002

Abstract

The paper describes experiments with FEARLUS, an agent-based model of land use change. The agents are satisficing (rather than optimising) decision-makers. The paper focuses on the effects of varying agents' *aspiration threshold*: the economic return an agent requires from a parcel of land to persist with the current land use. If the threshold is not reached, the agent chooses a land use at random to apply to that parcel in the next year. It is shown that the optimum level for the aspiration threshold is affected by environmental heterogeneity, and the level of return required to break even.

1 Introduction

This paper reports experiments with FEARLUS (Framework for Evaluation and Assessment of Regional Land Use Scenarios), an agent-based social simulation (ABSS) model (Conte, Hegselmann and Terna 1997) of land use change. The paper focuses on the comparative success of variants of an approach to land use selection which differ only in the economic return or *Yield* from a *Land Parcel* which will satisfy the *Land Manager*. (Terms referring to elements of FEARLUS models begin with an upper-case letter, and are italicised when first used.) We refer to the required level of Yield as the Manager's *Aspiration Threshold*. If the Aspiration Threshold is achieved, the Land Manager retains the same *Land Use* for that Land Parcel in the following *Year*.

The concept of aspiration thresholds goes back within economic psychology at least to Simon (1955), who uses the term 'aspiration level' for the minimum price a seller will accept for an item. The concept is linked to that of 'satisficing', introduced in Simon (1957) as an alternative to 'optimising', to indicate that agents frequently continue seeking a solution to a problem only until they find one that is 'good enough', rather than persisting in the hope of finding an optimum solution. This makes sense because the search for a solution is itself, in general, costly, and switching solutions may also involve costs. Moreover, it may not be known whether the new solution will indeed turn out better than the old. All these factors may apply to a real-world land manager's problem of land use selection; current FEARLUS models only deal with the last, which can arise in the real world because of gaps in land managers' knowledge of land uses' yield under various conditions, and the inherent unpredictability of climatic and economic conditions. Simon's concept is little used by those researching agricultural innovation. For example, a recently developed conceptual framework (Abadi Ghadim and Pannell 1999) does not mention it, although noting that farmers are in general somewhat 'risk averse' (given two courses of action with the same *expected* return, they prefer the one which reduces uncertainty).

The use of ABSS models to study land use is showing rapid growth (Berger, Park, Vescovi, Vlek and van de Giesen 2001, Parker, Manson, Janssen, Hoffman and Deadman 2001). Our approach to ABSS is to work initially with quite simple models, exploring their dynamics in some detail before adding additional complexity and realism. Simple models permit multiple simulation runs with each of a range of parameter settings, and hence use of statistical techniques to compare the model's performance under different parameter settings. Thus we can often show that the model *usually* behaves in a specific way, rather than merely that it *can* do so.

Early experiments with FEARLUS are described in Polhill, Gotts and Law (2001). The FEARLUS models discussed both there and here consist of a set of Land Managers (representing households, not individuals), and their *Environment*, which includes a grid of square Land Parcels, and a set of possible Land Uses. Every *Year*, Land Managers use their *Selection Algorithm* to choose a Land Use for each Land Parcel they own. The experiments described concern the competitive performance of Land Managers using different Land Use Selection Algorithms in a range of Environments differing in their spatio-temporal heterogeneity. We showed in Polhill et al. (2001) that the performance of specific Selection Algorithms varies with the environmental context, and with the Algorithms followed by other agents. Here, we vary the model Environments used more extensively and systematically than in our earlier work, and focus on a different and narrower range of Selection Algorithms.

Some Selection Algorithms studied in Polhill et al. (2001) involved imitation of neighbours' choices — a form of adaptation to environmental feedback. While extending our work on imitative Selection Algorithms, we discovered a family of non-imitative Selection Algorithms that use a simpler form of adaptation. This highlighted the strategic significance of the Land Manager's Aspiration Threshold (or 'high yield threshold'), a feature possessed by some of our imitative Selection Algorithms, but not given close attention in Polhill et al. (2001).

A Land Manager using any Selection Algorithm that involves an Aspiration Threshold looks at the Yield that a Land Parcel produced in the preceding Year. If this Yield equalled or exceeded the Aspiration Threshold, the Land Manager sticks with the same Land Use for that Land Parcel. Otherwise, some other procedure is used to select the Land Use. In the *Habit/Random Algorithm* ('HR') family, which is the focus of this paper, this is simply *Random Experimentation* — a random choice between the possible Land Uses, all having equal likelihood of being selected. Members of the HR family differ only in the level of their Aspiration Threshold. They can be differentiated by inserting their Aspiration Threshold into their name: thus H8R is HR with an Aspiration Threshold of 8, H10R is HR with an Aspiration Threshold of 10.

2 Method

The parameters of a FEARLUS model specify the size and shape of the grid of Land Parcels, and the length (e) of the bitstring encoding the *External Conditions* (which vary from Year to Year but apply across the whole grid). Each Land Parcel has a set of *Biophysical Characteristics*, fixed for the duration of a simulation run. These again are encoded as bitstrings; the length of these bitstrings, p, is another model parameter. Randomly generated bitstrings of length $p + e$ represent possible Land Uses (always eight of them in the simulations reported here). Two further Environment parameters are constant over over space *and* time: a *Break Even Threshold* (BET), specifying the Yield required from a Land Parcel to break even, and the *Land Parcel Price* (LPP).

After an initial *Year Zero*, in which Land Parcels are assigned to Land Managers, and there is a random setting of External Conditions and allocation of Land Uses to Land Parcels, the annual cycle is as follows:

1. Land Managers select the Land Use of each Land Parcel they own.
2. The bitstring encoding External Conditions is calculated by flipping each bit in the current bitstring with a *Flip Probability* specified as a model parameter.
3. Yield is calculated for each Land Parcel by matching the bitstring representing its current Land Use against a concatenation of those for its Biophysical Characteristics, and the External Conditions: the Yield is the number of matching bits. The *Account* of each Land Manager, which starts at zero, is updated for each Land Parcel owned by subtracting the BET from its Yield, and adding the result to the Account.
4. Land Managers with Accounts below zero sell their worst-performing Land Parcels one by one (at the LPP) until reaching or exceeding zero. Any obliged to sell all their Land Parcels leave the simulation. Land Parcels for sale are sold in random order. A buyer is chosen stochastically from a list made up of those who owned at least one of the Parcel's eight orthogonal or diagonal *Grid Neighbours* during the preceding Year, and have at least the

LPP in their Account (each Grid Neighbour owned gives its owner one chance to win), plus one potential new Land Manager (given a single chance to win).

The experiments reported here use a 7×7 array of Land Parcels, given a toroidal (wrap-around) topology, and regarded as representing an interior patch of a larger region. The bitstrings defining Land Uses' preferred conditions always contain 16 bits, but Environments differ in the division of these bits between Biophysical Characteristics (variable across space, fixed over time) and External Conditions (uniform across space, variable over time). External Conditions may be either correlated (Flip Probability $\frac{1}{8}$) or uncorrelated (Flip Probability $\frac{1}{2}$) from Year to Year. Each bit of the Biophysical Characteristics bitstring of each Land Parcel is set to 0 or 1 with equal probability and independently. In the real world, a complete lack of temporal or spatial auto-correlation is unlikely to occur; these extreme cases are experimentally useful idealisations.

Experimental Environments are described using the syntax: P$< p >$-E$< e >$[c|u]-BET$< b >$, where p is replaced by the length of the Land Parcel Characteristics bitstrings, e by the length of the External Conditions bitstrings (the 'c' or 'u' indicates whether these are correlated or uncorrelated from Year to Year), and b by the BET. In general, the LPP is set at twice the BET. Thus P12-E4c-BET10 indicates 12 Land Parcel Characteristic bits, 4 temporally correlated External Conditions bits, a BET of 10 and an LPP of 20. The first two parts of such an Environment type description — e.g. 'P12-E4c' — specify its *Heterogeneity Type*.

Each experiment reported consists of a set of simulation runs, pitting two *Sub-Populations* of Land Managers against each other. At the start of a run, each Land Parcel is assigned to a different Land Manager. Whether created at the start of a run, or to take over a Land Parcel being sold, a new Land Manager is equally likely to belong to either Sub-Population. All members of a Sub-Population share the same Selection Algorithm. After 200 Years, the success of the two Sub-Populations is assessed by counting the Land Parcels assigned to members of each. All simulation runs in an experiment use the same two Selection Algorithms, and type of Environment (defined by a fixed set of parameters e.g. P12-E4c-BET10). The runs differ only because a fresh seed is generated for each run, for use in pseudo-random processes within the model. The binomial test is used to determine whether one Sub-Population has finished significantly more runs holding a majority of the Land Parcels than the other. We regard results as significant at the .01 level (1 tailed), although given the number of experiments carried out, caution is advisable in interpreting those results not significant at a higher level.

Exploratory experiments suggested that the use and level of an Aspiration Threshold had an important influence on a wide range of imitative Selection Algorithms. HR performed comparably to imitative Selection Algorithms with the same Aspiration Thresholds in Environments with a lot of spatial variation. For both imitative Algorithms and HR, and across a number of Environments, an Aspiration Threshold near the BET gave the best performance. (The BET was almost always set at the *Random Choice Expected Yield* — the Yield expected from a random choice of Land Use, given the way Land Uses and Environments are constructed when a simulation starts. This is half the maximum possible Yield, i.e. 8 in the models used here.)

The exploratory experiments also led to the development of a 'standard set' of Heterogeneity Types of Environment. This uses seven partitions of a 16-bit Land Use bitstring between bits matched against Land Parcel Biophysical Characteristics, and those matched against External Conditions: $(0 : 16), (1 : 15), (2 : 14), (4 : 12), (8 : 8), (12 : 4), (16 : 0)$. More partitions were chosen with a majority of External Conditions bits than with a minority because experimental results appeared more sensitive to the precise partition used in the former case. For each of the six partitions in which there is any variation in External Conditions, this variation may be either temporally uncorrelated or temporally correlated, giving 13 standard Heterogeneity Types.

Environments with higher inter-Year predictability should make higher Yields easier to obtain, simply by sticking with successful Land Uses. Predictability of Yield from a given Land Use on a given Land Parcel clearly decreases as the number of External Conditions bits increases, and is higher if variation in External Conditions is temporally correlated than if it is uncorrelated. However, it is not a unidimensional property. The variance of Yield for a given combination of Land Parcel and Land Use — one obvious measure of predictability — does not distinguish temporally correlated from temporally uncorrelated Environments. In either, a single External Condition bit produces either 0 or 1 units of Yield each Year, and since the mean Yield produced by that bit will tend toward 1/2 in the long term, the variance will tend toward 1/4. Variance being additive, it will tend toward $e/4$ if there are e External Condition bits. Another measure is the expected variance of Yield over the Years *immediately following* those Years with a given Yield. It can be calculated from the binomial distribution that this will tend toward $ef(1-f)$, where f is the Flip Probability. This is $e/4$ for temporally uncorrelated Environments, $7e/64$ for temporally correlated ones, so it would be the same for a temporally uncorrelated Environment with an External Conditions bitstring of length 7, as for a temporally correlated Environment with one of length 16. However, it does not completely capture the differences between temporally uncorrelated and correlated Environments: in the former, the expected Yield from a given Parcel and Land Use in Year $y+1$ is independent of the Yield in Year y, while in the latter it is not — so in Environments of the two types with equal values for $ef(1-f)$, the current Yield would give a better estimate of next Year's Yield (and, to a decreasing extent, of the Yields of subsequent Years) in the temporally correlated one. This difference can be captured numerically in the expected value of the first autocorrelation (Kendall 1976) of the time-series of Yields produced by a given Land Parcel and Land Use, which will be $(1-2f)/4$, i.e. 0 for temporally uncorrelated Environments, $\frac{3}{16}$ or .1875 for temporally correlated ones.

3 Experiments and Results

Before describing our main experiments and their results, we note four simple analytical points, which have implications for the pattern of results to be expected.

1. In P0-E16u Environments, expected Yield is the same for all Land Uses. One might therefore expect to find no differences in competitive performance among Land Use Selection Algorithms.

2. The Random Choice Expected Yield is the highest Aspiration Threshold level to give the Land Manager a better than even chance of improving Yield when making a random choice. It thus seems a reasonable guess at where optimum Aspiration Thresholds might lie.

3. If the LPP exceeds or equals the BET, a Land Manager using any Aspiration Threshold Selection Algorithm with a Threshold also at or above the BET, will never have to sell a Land Parcel on which it has already found a Land Use guaranteed to meet its Threshold: Land Parcels are sold in 'worst Yield first' order, and the LPP gained from any other Land Parcel sold will at least cancel out the loss due to that Land Parcel in the preceding Year. With an LPP twice the BET, as used here, higher BETs should thus favour higher Aspiration Thresholds.

4. Suppose we set the BET at 0, so Land Parcels are never sold, and the Flip Probability f at $\frac{1}{2}$. For a given Land Parcel L_i, call the highest Yield that can be *guaranteed* on that Land Parcel by selecting the right Land Use the *maximin Yield*, m_i (this may differ between Land Parcels). A maximin-guaranteeing Land Use also makes possible the *maximum* Yield for that Land Parcel (gained when the External Conditions are optimal for that Land Use), and produces the best *expected* Yield of any Land Use. Calling the maximum achievable Yield for that Land Parcel M_i, this best expected Yield will be $(m_i + M_i)/2$, and Hm_iR (HR with Aspiration Threshold m_i) will give a mean Yield approaching this (the best achieveable by *any* Selection Algorithm) in the long run. At some point, the Hm_iR Land Manager will switch out of any Land Use that does not guarantee m_i — since, sooner or later, m_i will not be achieved. Once a Land Use guaranteeing m_i has been adopted, it will never be changed.

   Using Hm_iR will give this same expected Yield so long as $f > 0$, but $(m_i + M_i)/2$ is *not* necessarily the best mean Yield achievable on L_i by any Selection Algorithm if $f \neq \frac{1}{2}$: if this Year's External Conditions give some information about next Year's, it may pay to

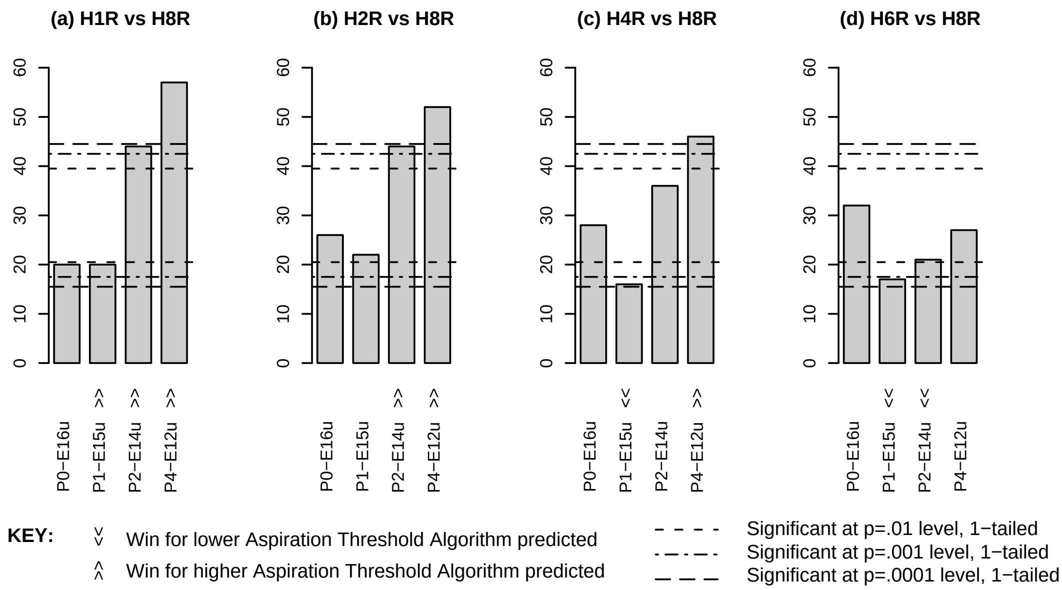

Figure 1: Wins for H8R in contests with H1R, H2R, H4R and H6R in BET=8 Environments with considerable uncorrelated temporal variation. Note the shift in favour of H8R from the second, through the third, to the fourth Environment in each case.

choose a Land Use that risks a Yield below m_i if it also currently offers better chances of getting Yields above m_i. This suggests optimum Aspiration Thresholds might be higher in Environments with $f = \frac{1}{8}$ than in otherwise identical Environments with $f = \frac{1}{2}$. Also, optimum Thresholds should increase with p (the length of the Land Parcel Characteristic bitstrings): the number of matching bits in pairs of random bitstrings of length p will follow a binomial distribution with mean $p/2$, so greater values of p will tend to raise values of m_i.

The experiments reported in sections 3.1 - 3.3 confirm the expectations generated by points 2-4, but reveal additional complexities. They show some evidence *against* those suggested by point 1; evidence supported by experiments and analysis mentioned briefly in 3.4 and to be reported in detail elsewhere.

3.1 Environments with BET Equal to Random Choice Expected Yield

Experiments with a BET of 8 (and LPP of 16) were performed using all 13 members of the standard set of Spatio-Temporal Heterogeneity Types, and 60 runs per experiment. H8R was matched against each of H1R, H2R, H4R, H6R, H10R and H12R. Based on exploratory experiments, it was predicted that an Aspiration Threshold of 8 would always do better than the higher Thresholds tried (except in P0-E16u, where no differences were predicted), and would also do better than the lower Thresholds in Environments of all Heterogeneity Types apart from those with a lot of uncorrelated temporal variation, i.e. P0-E16u, P1-E15u, P2-E14u and P4-E12u. All these predictions were confirmed at a significance level of .0001. Predictions and results for the remaining experiments are shown in figure 1.

In this figure, and in figures 2 and 3, we aim to show both the predictions and results of individual experiments, and the overall patterns which emerge. The latter should be interpreted cautiously, particularly when based to any extent on individual results that are not highly significant. In these figures, a vertical bar is used to indicate the number of contests won by *the HR variant with the higher Aspiration Threshold* in each experiment. Each group of bars covers a particular pairing of HR variants, with the Heterogeneity Type of the Environment for a specific experiment indicated below the corresponding bar. Predictions for individual experiments, and the degree of one-sidedness in either direction necessary to confirm predictions at various significance levels, are indicated as shown in the figure keys.

Predictions that H8R would outperform H1R, H2R and H4R in P4-E12u-BET8, and H1R and H2R in P2-E14u-BET8, were confirmed. So were predictions that H4R and H6R would outperform H8R in P1-E15u-BET8. Two further predictions were not confirmed.

The result of the contest between H1R and H8R in P0-E16u-BET8 is interesting: in an Environment of this Heterogeneity Type, as noted above, the expected Yield from any Land Parcel is the same *whatever* Land Use is selected. The result may be a chance occurrence, but as discussed in 3.4, there are reasons to think otherwise.

There appears to be a general shift in favour of H8R from the second through the third to the fourth column of each group — that is, as the amount of uncorrelated temporal variation decreases from P1-E15u-BET8 through P2-E14u-BET8 to P4-E12u-BET8. There also appears to be a shift *against* H8R from figure 1 (a) through to 1 (d), so far as Environments P2-E14u-BET8 and P4-E12u-BET8 are concerned (rightmost columns in each graph), so that H6R does best and H1R worst against H8R in these Environments. The relatively good performance of low-Threshold variants of HR in the Environments of Heterogeneity Types P1-E15u, P2-E14u and P4-E12u prompted a further series of 120-run experiments using all pairings of H1R, H2R, H4R and H6R. The higher-Threshold variant was predicted to win in all contests (based on preliminary experiments), except for those between H1R and H2R, for which no predictions were made. All predictions were confirmed at a significance level of at least .001, with one exception: H4R vs H6R in P1-E15u-BET8, where the result was close to equality.

To summarise, with a BET of 8, equal to the Random Choice Expected Yield, the optimum Aspiration Threshold for HR appears to be equal to the BET, except when there is a lot of uncorrelated temporal variation, when the optimum is somewhat lower.

3.2 Low BET Environments

With a BET of 4, the performance differences between HR Algorithms with different Aspiration Thresholds were more or less completely obliterated. With such a low BET, very few Land Parcels would ever need to be sold. Even with a BET of 6, performance differences were greatly reduced, although H6R and H8R outperformed H1R, H2R and H4R in a wide range of Environments. H6R, H8R and H12R were all pitted against each other in 240-run experiments in Environments of all 13 standard Heterogeneity Types. The clearest results were that H8R outperformed H12R in most Environments with correlated temporal variation, while H6R outperformed H12R in P16-E0-BET6, the Environment with no temporal variation (see 3.4).

3.3 High BET Environments

Environments with BET=10 showed a more complex pattern than those with BET=8. The standard set of 13 Heterogeneity Types was used. Predictions and results of the 60-run experiments are shown in figures 2 and 3. Figure 2 deals with contests between either H8R or H10R on the one hand, and HR variants with lower Aspiration Thresholds on the other, while figure 3 covers contests pitting H4R against H6R, H8R against H10R, and H10R against H12R.

In the Environments with a lot of uncorrelated temporal variability, and hence very low Yield predictability (light grey bars towards the left of each subgroup in figures 2 and 3), there was little difference between Algorithms in *any* of the contests.

The six contests graphed in figure 2 show striking similarities. The patterns seen across the Environments with correlated temporal variation are markedly different from those seen across the Environments with uncorrelated temporal variation. In the former, the higher Threshold variant always has an advantage, and this grows fairly smoothly as temporal variation decreases. In the latter, there is little difference between variants at the highest amounts of temporal variation, the lower Threshold variants do better in intermediate cases, and the higher Thresholds win at the lowest levels of temporal variation. Comparing subfigures (c) and (d) with (e) and (f), the dip in

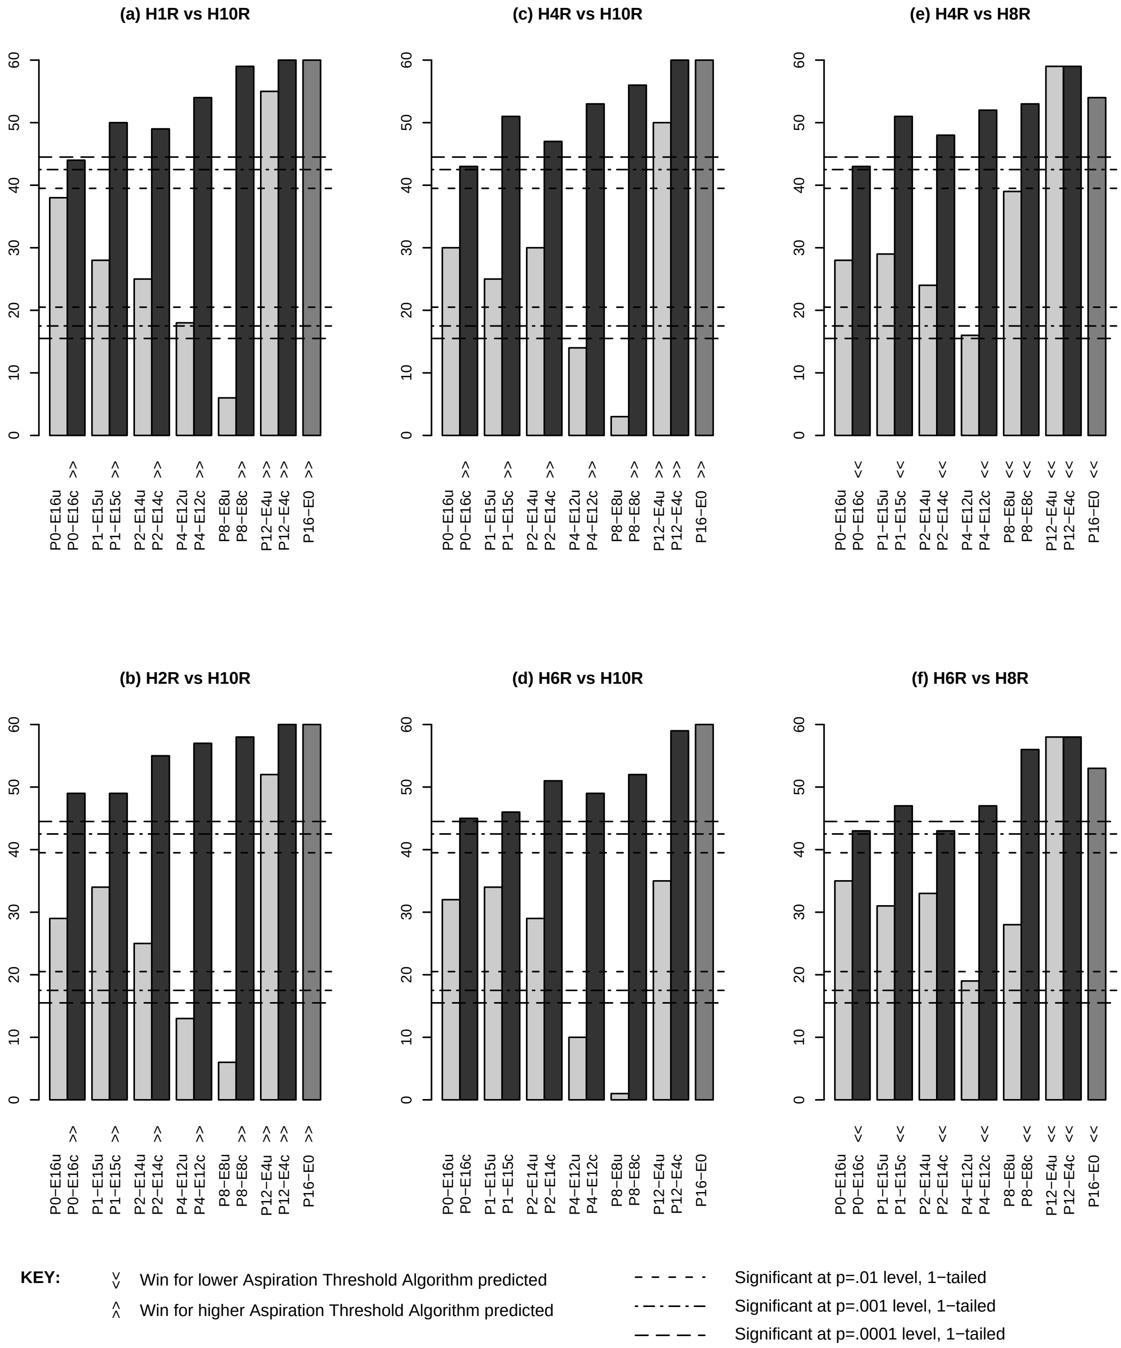

Figure 2: Wins for the **higher** Threshold version of HR, in pairings of H10R against H1R, H2R, H4R and H6R, and of H8R against H4R and H6R, in BET 10 Environments: uncorrelated and correlated temporal variation (left and right members of pairs of bars) show different patterns.

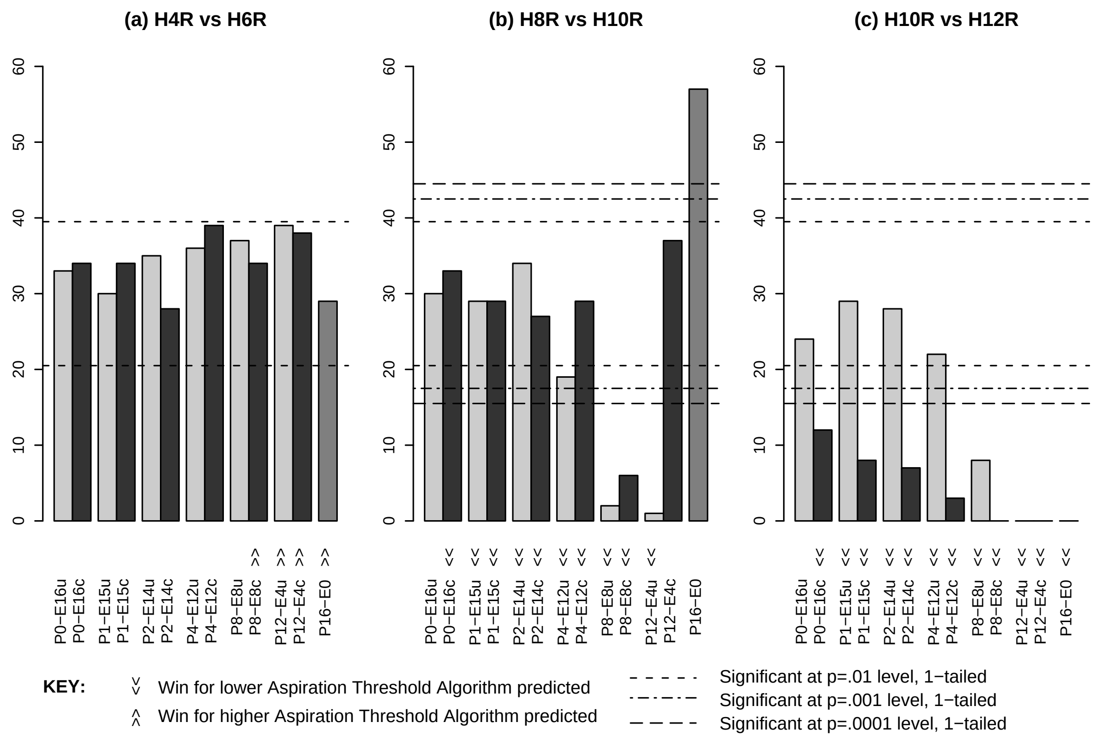

Figure 3: Wins for the **higher** Threshold version of HR, in pairings of H4R with H6R, H8R with H10R, and H10R with H12R, in BET 10 Environments. Note the dip and subsequent rise of wins for H10R against H8R, as the amount of temporal variation decreases.

performance for the lower-Threshold variants appears to be deeper, and to occur at lower levels of temporal variation, when these are pitted against H10R rather than H8R.

Turning to figure 3, H10R outperforms H12R in all Environments other than those with a lot of uncorrelated temporal variation. Subfigure (b) shows that in the contest between H8R and H10R, in contrast to all other contests involving either of these variants, patterns among the Environments with correlated and uncorrelated temporal variation are fairly similar: at high levels of temporal variation there is little to choose between the variants, while at low levels H8R wins easily. In the Environment of perfect Yield predictability (P16-E0), H10R does best, and the major difference between the temporally correlated and uncorrelated cases occurs in Environments close to this one: H8R's advantage disappears in P12-E4c-BET10, but not in P12-E4u-BET10. The results in P4-E12c-BET10, P8-E8c-BET10 and P12-E4c-BET10, where the first and third Environments show near equality while the second shows a clear advantage for H8R, were so unexpected that a second set of experiments was run for these Environments, but with very similar results. The corresponding dip in the performance of H10R in the temporally uncorrelated Environments strengthens the case for this being a real phenomenon.

Setting the BET at 12 was found, in exploratory experiments, to reduce differences between Selection Algorithms, but mainly in Environments with fewer Land Parcel bits than External Conditions bits. In the remaining Environments, the pattern of results was somewhat similar to that found in corresponding Environments with BET=10. In P16-E0-BET12, H12R clearly defeated all other variants of HR tried except H10R (it was predicted to defeat them all), and this Environment differed very markedly from all others. In the P12- Environments, and most markedly in P12-E4u-BET12, variants with lower Aspiration Thresholds did better, inflicting clear and predicted defeats on H12 in two cases. In P8- Environments (as in all those with a preponderance of External Condition bits), differences between Algorithms were not significant.

3.4 Overall Pattern of Results

Figure 4 summarises the effects of varying the Aspiration Threshold on the performance of HR, as revealed by simulation experiments. It illustrates how Thresholds equal either to the BET, or to the Random Choice Expected Yield of 8, perform against higher and lower Thresholds, showing whether Environments of a particular BET and Heterogeneity Type lead to differences between Thresholds, and if so, whether either or both reference values (the BET and 8) lie within the *range* of Thresholds giving optimal or near-optimal performance. The shading given to each cell is based on the overall tendency of results; some are more firmly based than others.

Heterogeneity Types with temporally uncorrelated variation in External Conditions occupy the upper rows, those with temporally correlated variation the lower (with P16-E0, which includes no External Conditions, in between). Empty cells indicate possible kinds of Environment that have not been used in experiments.

Note first that if the BET is 8, equal to the Random Choice Expected Yield, Aspiration Thresholds approximately equal to the BET appear to work best, except in Environments with very low Yield predicability. In Environments with very low but non-zero predictability, a rather lower Aspiration Threshold does better; reasons to expect the best Threshold to be lower in these Environments were explained at the start of section 3. P0-E16u is a special case, discussed below.

With a BET of 4, the Aspiration Threshold has little apparent effect on Land Managers' ability to gain or retain control of Land Parcels (it would still be expected to affect the accumulation of wealth in Land Managers' Accounts, but this has not been tested). The same flattening of differences between Selection Algorithms occurs to a lesser degree with BETs 6, 10, and 12, being most marked when Yields are most unpredictable from Year to Year. All this makes good sense: differences between Aspiration Thresholds are reduced when the choice of Land Use makes little difference to the Land Manager's probability of achieving the BET.

With a BET of 6, the optimum Aspiration Threshold is 8 in most Environments; if temporal variability is removed (P16-E0), an Aspiration Threshold of 6 seems superior, although evidence for this is not conclusive (the similar result for P2-E14c is based on slight and conflicting evidence). It is noteable that a threshold equal to the BET is in the optimal range for all BETs, in Environments

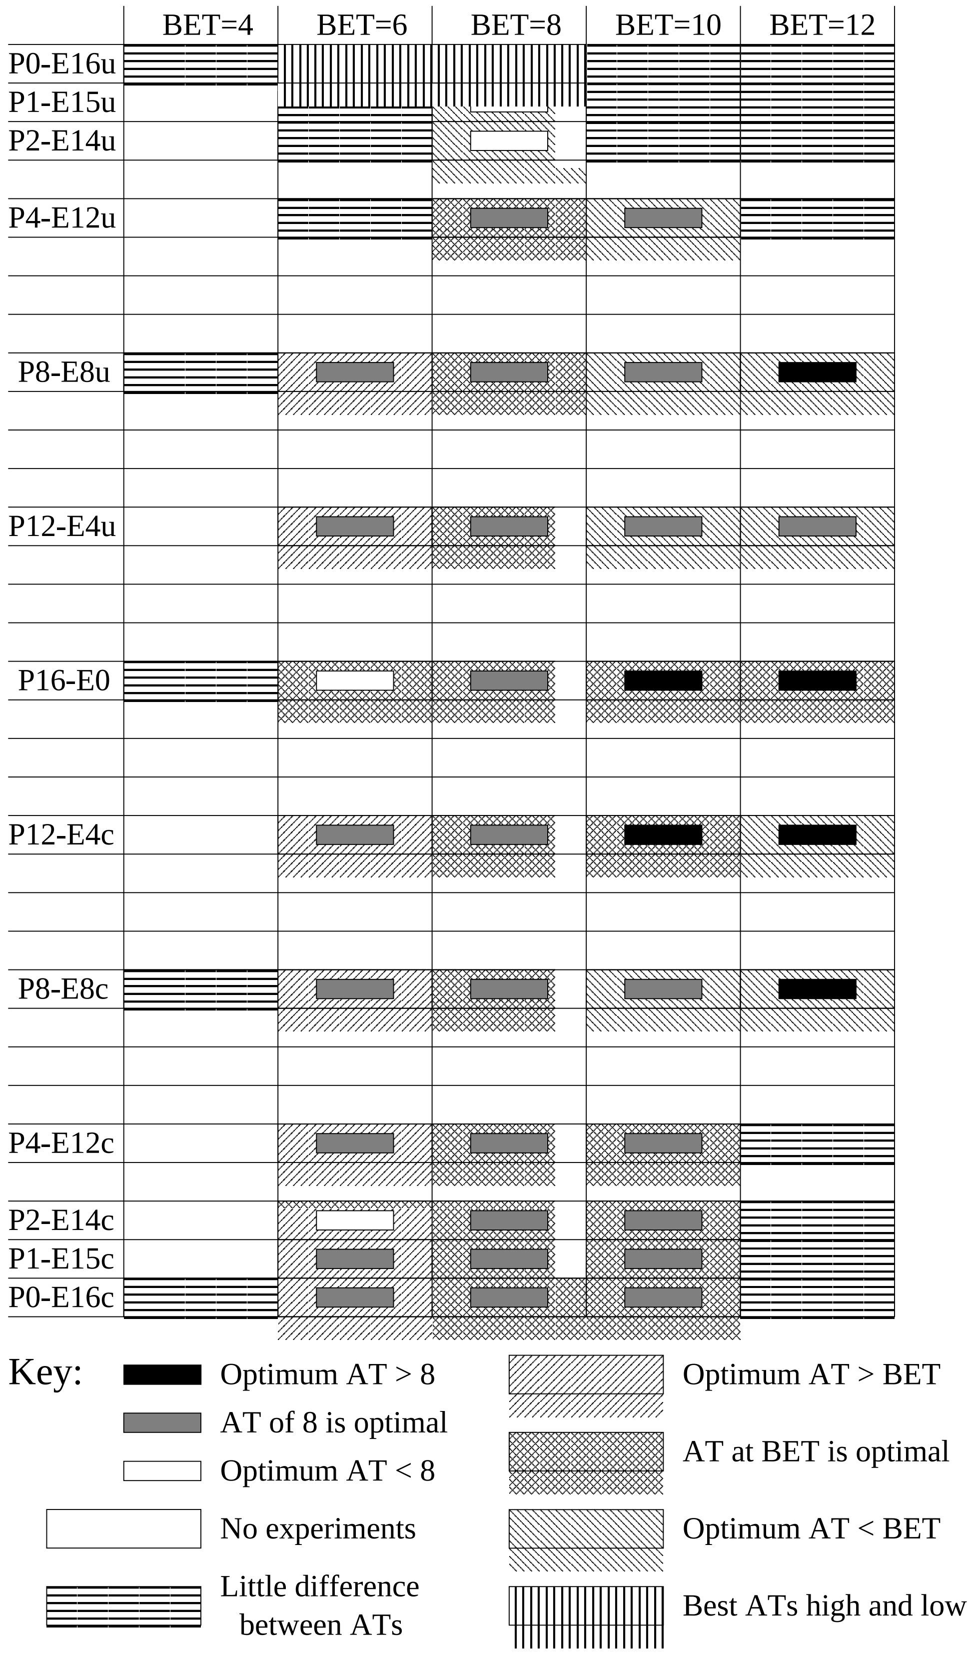

Figure 4: Summary of results across all Environments

of this Heterogeneity Type. In a temporally unvarying Environment, a Land Use that produces a Yield at or above the BET on a given Land Parcel when first used will always do so, and the Land Manager can hold that Parcel permanently by retaining that Land Use. Setting the Threshold below the BET would risk sticking with an economically unsustainable Land Use, while setting it higher would risk abandoning a sustainable one.

For a BET of 10, optimal Aspiration Thresholds are very similar to those for a BET of 8 at intermediate levels of Year-to-Year unpredictability — specifically, for all the temporally correlated Heterogeneity Types except P12-E4c, and for two temporally uncorrelated Heterogeneity Types — P8-E8u and P12-E4u. At both extremes of predictability, differences appear. With Yield predictability sufficiently low, Aspiration Threshold makes little difference, presumably because no Aspiration Threshold would give Yields at or above the BET, so most Land Managers are replaced in most Years. At just above this level (P4-E12u), Thresholds rather below 8 do best; this phenomenon occurred only at somewhat lower levels of predictability with a BET of 8, but the difference between the two BET levels in P4-E12u Environments is slight, and may be due to chance. At the opposite extreme, H10R does best when there is no temporal unpredictability (P16-E0-BET10). In this type of Environment the probability of at least one of eight Land Uses having a Yield of at least 10 on any given Land Parcel (so that H10R will eventually find a Land Use guaranteeing the BET) is above 0.87. To a lesser degree, a Threshold of 10 also does best in P12-E4c Environments, which are the next most predictable.

A simpler pattern is found with a BET of 12: Thresholds at least equal to the Random Choice Expected Yield but below the BET do best, except when temporal variation is completely absent (even here, H10R appears as good as H12R — the chance of at least one of the eight Land Uses having a Yield of at least 12 on any given Land Parcel is just below 0.27), or highly unpredictable (when performance differences vanish).

A special shading has been used for two top-row cells (P0-E16u-BET6 and P0-E16u-BET8). This is not based wholly on results reported here. The Heterogeneity Type P0-E16u makes all Land Uses equal in terms of expected Yield, and expected variance of Yield over time. However, results to be reported in full elsewhere show that Sub-Populations using Selection Algorithms that maximise the *diversity* of Land Uses employed do best in P0-E16u-BET8. This is due to the land sales mechanism described in section 2 (if land sales are blocked by using a very high LPP, the effect disappears, while a zero LPP accentuates it). If the diversity of Land Uses on a Sub-Population's Land Parcels is low in all or most Years, many of them will tend to go out of business at the same time, allowing those members of a competing Sub-Population with greater Land Use diversity who happen to have done well, to buy up much of their land. Among Sub-Populations using HR Algorithms, the greatest diversity will be found if either an unattainable or a zero Threshold is used: in the former case random choices are made each Year; in the latter, differences in Land Use present when the Managers gain the Parcels concerned persist.

Setting this special case aside, the optimal Threshold appears never to exceed whichever is the greater of the Yield needed to break even (the BET), and the Random Choice Expected Yield, but varies systematically with the characteristics of the Environment. Important determinants are the BET, and the degree and pattern of temporal heterogeneity in the Environment — which affects the degree to which knowledge of one Year's Yield from a given Land Parcel and Land Use allows prediction of the Yield in the following Year. These determinants interact in a moderately complex pattern, but greater BET and greater predictability both tend to raise the optimal Threshold, as might be expected: the former requires higher Yields, the latter makes them easier to attain.

4 Discussion

The primary finding of the simulation experiments reported is the influence of Aspiration Thresholds on the dynamics of strategic interactions within our ABSS model of Land Use selection. The Aspiration Threshold has important effects across a wide range of FEARLUS Environments. Results to be reported elsewhere show that these persist if the alternative to retaining the current Land Use involves imitation of neighbouring Land Managers rather than random choice.

If we are right in thinking that real-world land managers are often satisficing rather than optimising in their land use decisions, the findings reported here have considerable relevance for empirical studies, and for theoretical work such as Abadi Ghadim and Pannell (1999) and Pender (1998), which takes a neoclassical approach, as well as for the agent-based studies described below. In particular, they suggest that comparative studies of innovation in environments differing in temporal heterogeneity or in the difficulty of breaking even would be useful. It should be possible to investigate whether real-world land use systems show effects similar to our model by studying the relationships between rates of land use change, land sales and farmer bankruptcies in environments differing in respect of the unpredictability of economic returns and the difficulty of breaking even. Farmers with high rates of land use change should tend to do relatively better as returns become more predictable, and as yields required to break even get higher.

It does seem likely that real-world land managers are often satisficing rather than optimising when making decisions. For a subsistence farmer, it is advisable to maximise the chance of growing at least enough to eat (i.e., of reaching a certain threshold harvest), rather than the *expected* harvest. Even for a small farmer in an industrialised country, who is unlikely to starve, maximising the chance of producing enough to stay in business for another year may often be the priority. Some recent studies of agricultural innovation and the obstacles that prevent it, such as Fujisaka (1994) discuss these obstacles in ways which can readily be interpreted in terms of satisficing and aspiration thresholds.

The existing ABSS work most closely related to that described here falls into two categories: spatially explicit models of rural land use, and work (with or without a spatial aspect) concerning the effects of aspiration levels on the dynamics of strategies in the 'Prisoner's Dilemma' and other 'social dilemmas' — situations in which two or more agents can each do best individually by acting selfishly, but where all will end up better off if all act cooperatively. We mention a few examples of each, choosing from among those closest to our own work.

Work in the first of these categories which also deals with aspiration levels is almost nonexistent, but in a notable exception Rouchier, Bousquet, Barreteau, Page and Bonnefoy (2000) describe the SHADOC model of irrigation systems in the middle Senegal valley. Water and common infrastructure constitute shared resources, and farmers belong to organisations set up to run these. After every growing season, each agent assesses its success, and if its *criterion of satisfaction* is not met, may adopt the rules of another agent. The 'criterion of satisfaction' is clearly similar to our Aspiration Threshold, but no systematic study of the effects of varying this criterion is reported, and the focus is on overall outcomes rather than competition between strategies.

There are a number of other ABSS studies of land use which concern imitation (Weisbuch and Boudjema 1999, Berger et al. 2001). For all such studies, it might be useful to compare the dynamics of imitation with those resulting from an HR-like strategy: use of an aspiration threshold, together with random experimentation when the threshold is not met. This would help to distinguish the effects of agents changing their course of action when (and only when) the current one gives poor results, from those specific to imitation.

In the Prisoner's Dilemma (for full descriptions of this and similar abstract games see Hargreaves Heap and Varoufakis (1995)), each of the two players has two possible strategies, called 'C' and 'D', and in a single game, will always do better by playing 'D' whatever the other player does — but if both play 'D', both do worse than if both had played 'C'. In a single game, both players, if self-interested and rational, will play 'D', but if the two are to repeat the game an unknown number of times, various complicated strategies are feasible. Posch (1999) gave players of this 'repeated Prisoner's Dilemma' different aspiration levels, in a study without any spatial aspect. Players scoring below the aspiration level switched from C to D or vice versa, with a probability which varied between players. With aspiration levels fixed for each player, the most successful had very low aspirations — although mean scores were well above the aspiration level, unless a good deal of 'noise' was added to the system. Allowing aspiration levels to adapt (moving toward actual scores) led to higher aspiration levels and, in a noisy environment, to higher overall scores. In contrast, Kirchkamp (Kirchkamp 1999, Kirchkamp 2000), pitting agents against their neighbours on a two-dimensional lattice, found that allowing agents to update learning rules which included a kind of aspiration threshold *reduced* the range of two-player games in which repetition produced stable cooperation. Macy and Flache (2002) found a similar effect. These findings do not map in any obvious way onto our own; our agents are competing primarily against an exogenous standard, the BET, rather than each other. Taken with our own results, however, these studies indicate that aspiration levels may be important across a broad range of agent-based simulation work and related empirical and theoretical studies.

5 Acknowledgement

The authors would like to acknowledge the financial support of the Scottish Executive Environment and Rural Affairs Department.

References

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

Berger, T., Park, S., Vescovi, F., Vlek, P. L. G. and van de Giesen, N.: 2001, Sustainable water use under changing land use, rainfall reliability, and water demands in the Volta basin, *LUCC Newsletter* (6), 14–15.

Conte, R., Hegselmann, R. and Terna, P. (eds): 1997, *Simulating Social Phenomena*, number 456 in *Lecture Notes in Economics and Mathematical Systems*, Springer-Verlag.

Fujisaka, S.: 1994, Learning from six reasons why farmers do not adopt innovations intended to improve sustainability of upland agriculture, *Agricultural Systems* **46**, 409–425.

Hargreaves Heap, S. P. and Varoufakis, Y.: 1995, *Game Theory: A Critical Introduction*, Routledge, London.

Kendall, M.: 1976, *Time-Series*, Charles Griffin, London.

Kirchkamp, O.: 1999, Simultaneous evolution of learning rules and strategies, *Journal of Economic Behavior and Organization* **40**, 295–312.

Kirchkamp, O.: 2000, Evolution of learning rules in space, *in* R. Suleiman, K. G. Troitzsch and G. N. Gilbert (eds), *Tools and Techniques for Social Science Simulation*, Physica-Verlag, Berlin, chapter 10, pp. 179–195.

Macy, M. W. and Flache, A.: 2002, Learning dynamics in social dilemmas, *Proceedings of the National Academy of Sciences* **99** (Supplement 3), 7229–7236.

Parker, D., Manson, S., Janssen, M., Hoffman, M. and Deadman, P.: 2001, Multi-agent systems for the simulation of land use and land cover change: A review, manuscript available via http://www.csiss.org/events/other/agent-based/participants.htm.

Pender, J. L.: 1998, Population growth, agricultural intensification, induced innovation and natural resource sustainability: an application of neoclassical growth theory, *Agricultural Economics* **19**, 99–112.

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.

Posch, M.: 1999, Win stay-lose shift strategies for repeated games — memory length, aspiration levels and noise, *Journal of Theoretical Biology* **198**, 183–195.

Rouchier, J., Bousquet, F., Barreteau, O., Page, C. L. and Bonnefoy, J.-L.: 2000, Multi-agent modelling and renewable resource issues: The relevance of shared representations for interacting agents, *in* S. Moss and P. Davidsson (eds), *Multi-Agent-Based Simulation: Second International Workshop MABS 2000*, number 1979 in *Lecture Notes in Artificial Intelligence*, Springer, Berlin, pp. 181–197.

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.

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