Chapter 8

Co-operation in the StereoLab

The StereoLab (described in chapter 6) consists of boundedly rational satisficing agents playing pairwise games of the single-shot Prisoner's Dilemma (PD) game. Agents can only identify each other by observing surface tags (representing features or social cues) which may not be unique. As a consequence of this it would seem unlikely that reciprocal co-operative strategies would emerge since individual level recognition and punishment is a prerequisite for this [5]. However, since the agents are satisficers, being satisfied with any game outcome other than being "suckered" or "punished", it would appear that co-operation has a better chance of emerging than in scenarios where the agents are strict utility maximisers.

As discussed previously (see chapter 6, table 6.1) the StereoLab has 17 (non-fixed) exogenously defined parameters controlling various aspects of agent behaviour, interaction and environmental structure. In this chapter an empirical investigation of simulation runs is used to answer the question: "Which regions of the parameter space produce highly co-operative and highly non-cooperative societies and why?", with particular emphasis on co-operative societies.

The parameter space was specified over the range of the exogenously defined parameters. Each dimension was quantised into discrete increments. For integer parameters the increment was set to 1. For parameters in the range [0..1] the increment was set to 0.1. This produces a discrete space, or grid, containing 5 x 9 x 10^3 x 11^{12} = 1.41 x 10^{17} possible unique points (one point for each combination of parameter values). In order to locate regions of high and low co-operation between agents in the parameter space, two methods were used. Firstly, *decision tree induction* (using the C4.5 algorithm [139]) was used over a large random sample of points taken from the whole space (see section 8.1). Secondly, *k-means cluster analysis* was combined with points found via hill-climbing in the space (see section 8.2).

Interestingly, it was found that although the agents are satisficers, highly co-operative societies are rare within the space. Analysis suggests that for those rare highly co-operative societies, two distinct forms of game interaction localisation (or "parochialism") promote co-operation (among other mechanisms). These take the form of *spatial parochialism* (agents only playing games with others sharing the same territory) and *cultural parochialism* (agents only playing games with others sharing the same or very similar tags). The C4.5 algorithm combined with random sampling identified the spatial mechanism (see section 8.1.5) as did k-means cluster analysis combined with hill-climbing (see section 8.2.2). The cultural mechanism was found by applying k-means clustering to points from an extended parameter space (see section 8.2.4). The process of cultural parochialism identified is novel and is interpreted as a form of "cultural group selection" (see section 8.2.5). Chapter 9 explores this process further.

8.1 Inducing Regions Using C4.5

The C4.5 algorithm was used to induce regions in the space as follows: Firstly the parameter space was randomly sampled (approximately 10,000 points) in order to gain an empirical measure for the prevalence of co-operation within the space (sample 1). Each point in the space represents a simulation run terminated after 100 cycles. As stated previously (see section 6.5.3 in chapter 6) one cycle is equivalent to 10N game interactions, where N is the number of agents. The number of agents is fixed at 101[1] for all experiments detailed here. Consequently a single simulation run is terminated after 101,000 game interactions[2].

Game interaction between agents involves pairs of agents playing single-shot games of the PD (see section 2.1.1 in chapter 2 and section 6.3 in chapter 6). The amount of co-operation (CC) for a run is calculated as the proportion of mutually co-operative game interactions over the last cycle. If agents selected strategies randomly, 25% of game interactions would be mutually co-operative (i.e. when both agents co-operate), 25% would be mutually defective (i.e. when both agents defect) and 50% would be exploitative (i.e. when one agent co-operates and the other agent defects). Figure 8.1 shows a frequency distribution histogram of the CC measures for the initial sample (sample 1).

Thresholds were then calculated over the CC measures in the sample such that they partitioned the sample into categories of increasing CC value - 10 percentile divisions (see table 8.1). The C4.5 algorithm was then applied using each point as a case and percentile membership as the categorical class for that case. A "weight value" of 600 was used (see section 8.1.1 below for a description of how this value was obtained). An induced decision

[1]The value of 101 agents was selected to be equal to the number of territories in the environment. The number of territories was set to an odd number so that an agent within a given territory had a balanced number of neighbouring territories in either direction around the ring into which game and cultural interaction windows may extend.

[2]This means that this initial sample (sample 1) is a synthesis of 1.01 x 10^9 (one billion, ten million) individual game interactions.

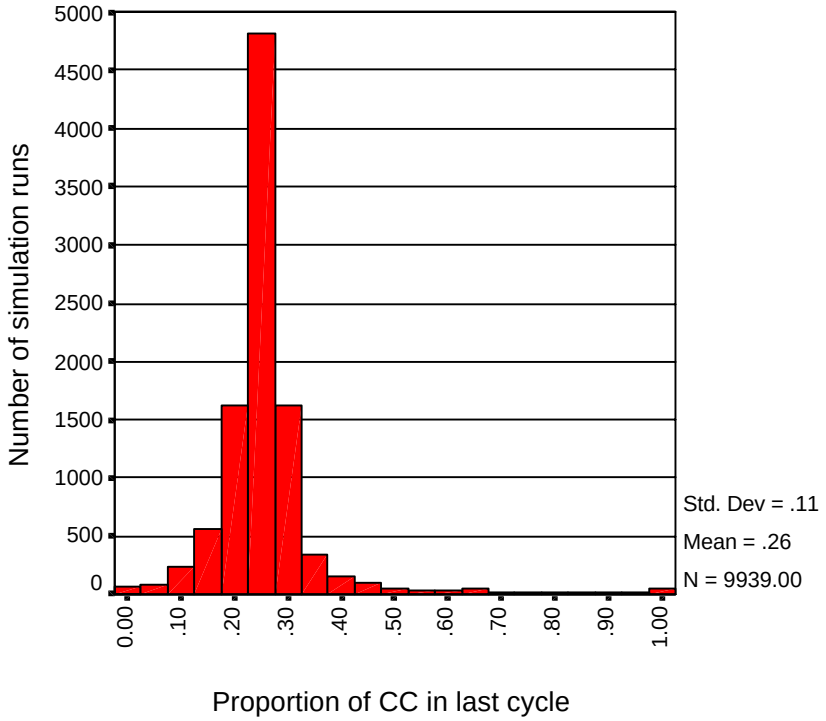

Figure 8.1: Frequency of co-operation over the whole parameter space (sample 1). This is an initial sample comprising approximately 10,000 random points from the parameter space. The chart shows the frequency distribution of co-operation. Co-operation is distributed around the mean of 0.26. Notice the small peak at CC=1.00. This indicates that a small number of points in the sample have maximum co-operation. See section 8.1.2 for details of mechanisms producing high co-operation.

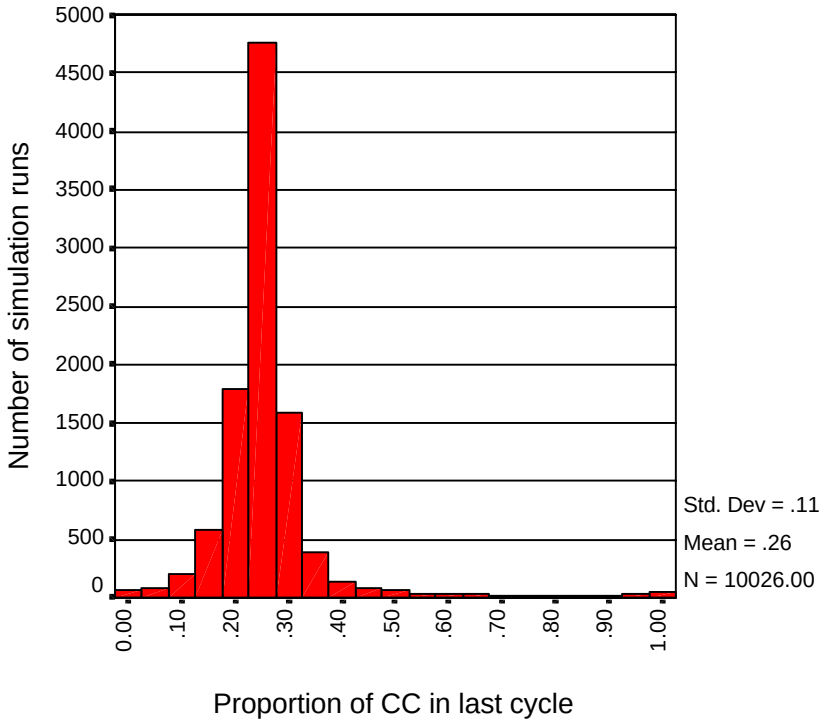

Figure 8.2: Frequency of co-operation over the whole parameter space (sample 2). A second independent sample comprising approximately 10,000 random points from the parameter space. The chart shows the frequency distribution of co-operation. Co-operation is distributed around the mean of 0.26.

tree was produced (see figure 8.3). The tree generalises class membership over parameter intervals. In order to verify the accuracy of the generalisation produced a further sample of approximately 10,000 points was made (sample 2). A tree was independently induced from this sample in order to verify that a stable generalisation was found (see figure 8.4). Figure 8.2 shows a frequency distribution histogram of the CC measures for sample 2.

8.1.1 Coping With Stochasticity

C4.5 allows the setting of a weight parameter to guide the induction process. The weight indicates the minimum number of points that should fall within any region indicated at the leaf node of the tree. Essentially a small value can cause "overfitting"

Percentiles	**Max. Value**
10	.179
20	.215
30	.230
40	.240
50	.249
60	.259
70	.269
80	.284
90	.320

Table 8.1: The percentile ranges over the level of co-operation measure for sample 1.

which is especially troublesome for noisy data. In the case of the simulation data here, "noise" is a result of the non-deterministic nature of the output measures with respect to the input parameters. In this instance a weight of 600 was used. This value was chosen as the best compromise over the three requirements of stability, compactness and accuracy. This value was found empirically by inducing 10 trees with weights from 100 to 1000 in increments of 100 for both samples and then using the independent sample as an unseen test set. From those 20 induced trees the expected, actual and found error along with tree size were calculated. Figures 8.5, 8.6, 8.7 and 8.8 give the results from these weight parameter experiments.

As can be seen, with a weight of 600 a compact tree is produced which is not too far from the error rates attained when much smaller weights are used. Importantly this value also gives a relatively stable induced tree over the two independent samples. Figures 8.3 and 8.4 show the two trees induced from the two independent samples. Figure 8.9 shows the lowest and highest percentile regions found in each tree. The first value in square brackets given after each region indicates the proportion of points within the region which *are not* a member of the class selected for the region. A value of zero would indicate that all points that fall within the region are of the class associated with the region. This value

```
Simplified Decision Tree:

Mt <= 0 : 1 (931.0/516.0) [0.55] [0.35]
Mt > 0 :
|   Cr <= 0 : 1 (795.0/478.1) [0.60] [0.30]
|   Cr > 0 :
|   |   Vg <= 0 : 10 (737.0/499.3) [0.68] [0.22]
|   |   Vg > 0 :
|   |   |   Pm <= 0 : 6 (676.0/567.1) [0.84] [0.06]
|   |   |   Pm > 0 :
|   |   |   |   Ms <= 0.1 : 10 (661.0/558.0) [0.84] [0.06]
|   |   |   |   Ms > 0.1 :
|   |   |   |   |   Mt <= 0.1 : 10 (618.0/512.0) [0.83] [0.07]
|   |   |   |   |   Mt > 0.1 :
|   |   |   |   |   |   Fg <= 0.2 : 10 (1022.0/859.6) [0.84] [0.06]
|   |   |   |   |   |   Fg > 0.2 :
|   |   |   |   |   |   |   Fc <= 0.1 : 6 (832.0/690.0) [0.83] [0.07]
|   |   |   |   |   |   |   Fc > 0.1 :
|   |   |   |   |   |   |   |   M > 7 : 5 (1209.0/1018.3) [0.84] [0.06]
|   |   |   |   |   |   |   |   M <= 7 :
|   |   |   |   |   |   |   |   |   Pm <= 0.3 : 5 (800.0/684.4) [0.86] [0.04]
|   |   |   |   |   |   |   |   |   Pm > 0.3 :
|   |   |   |   |   |   |   |   |   |   Fc <= 0.5 : 3 (748.0/646.0) [0.86] [0.04]
|   |   |   |   |   |   |   |   |   |   Fc > 0.5 : 9 (910.0/781.8) [0.86] [0.04]


Evaluation on training data (9939 items):

         Before Pruning           After Pruning
        ----------------   --------------------------
        Size      Errors   Size      Errors   Estimate

          25  7710(77.6%)     23  7711(77.6%)    (78.6%)   <<

Evaluation on test data (10026 items):

         Before Pruning           After Pruning
        ----------------   --------------------------
        Size      Errors   Size      Errors   Estimate

          25  7835(78.1%)     23  7828(78.1%)    (78.6%)   <<

Confusion Matrix:

          (a)  (b)  (c)  (d)  (e)  (f)  (g)  (h)  (i)  (j)   <-classified as
         ---- ---- ---- ---- ---- ---- ---- ---- ---- ----
          760        13        15    1             31  180   (a): class 1
          342        86       142   80            101  306   (b): class 2
          129       107       260  220            105  286   (c): class 3
           69        81       197  209            102  234   (d): class 4
           65        95       258  240             97  223   (e): class 5
           57        78       331  273            120  274   (f): class 6
           39        87       262  210             94  256   (g): class 7
           44        94       245  175             91  268   (h): class 8
           67        93       195  108            111  394   (i): class 9
          154        34        56    8             85  689   (j): class 10
```

Figure 8.3: The C4.5 induced decision tree generated from the initial sample (sample 1) and evaluated on a second independent sample (sample 1). The Confusion matrix shows how the unseen cases in sample 2 were categorised by the induced tree. This can be compared with figure 8.4 below.

```
Decision Tree:

Mt <= 0 : 1 (877.0/455.0) [0.52] [0.38]
Mt > 0 :
|   Cr <= 0 : 1 (849.0/511.0) [0.60] [0.30]
|   Cr > 0 :
|   |   Vg <= 0 : 10 (759.0/489.0) [0.64] [0.26]
|   |   Vg > 0 :
|   |   |   Pm <= 0 : 5 (724.0/593.0) [0.82] [0.08]
|   |   |   Pm > 0 :
|   |   |   |   Ms <= 0.1 : 10 (717.0/590.0) [0.82] [0.08]
|   |   |   |   Ms > 0.1 :
|   |   |   |   |   Mt <= 0.2 :
|   |   |   |   |   |   Fg <= 0.5 : 10 (610.0/472.0) [0.77] [0.13]
|   |   |   |   |   |   Fg > 0.5 : 2 (620.0/533.0) [0.86] [0.04]
|   |   |   |   |   Mt > 0.2 :
|   |   |   |   |   |   Fg <= 0.2 : 10 (885.0/755.0) [0.85] [0.05]
|   |   |   |   |   |   Fg > 0.2 :
|   |   |   |   |   |   |   Fc <= 0.1 : 6 (723.0/589.0) [0.81] [0.09]
|   |   |   |   |   |   |   Fc > 0.1 :
|   |   |   |   |   |   |   |   B <= 5 : 6 (1270.0/1055.0) [0.83] [0.07]
|   |   |   |   |   |   |   |   B > 5 :
|   |   |   |   |   |   |   |   |   Pm <= 0.5 : 6 (971.0/811.0) [0.84] [0.06]
|   |   |   |   |   |   |   |   |   Pm > 0.5 : 9 (1021.0/880.0) [0.86] [0.04]

Evaluation on training data (10026 items):

         Before Pruning           After Pruning
        ----------------   --------------------------
        Size      Errors   Size      Errors   Estimate

          23  7733(77.1%)     23  7733(77.1%)    (78.1%)   <<

Evaluation on test data (9939 items):

         Before Pruning           After Pruning
        ----------------   --------------------------
        Size      Errors   Size      Errors   Estimate

          23  7821(78.7%)     23  7821(78.7%)    (78.1%)   <<

Confusion Matrix:

          (a)  (b)  (c)  (d)  (e)  (f)  (g)  (h)  (i)  (j)   <-classified as
         ---- ---- ---- ---- ---- ---- ---- ---- ---- ----
          753   23                20             21  191   (a): class 1
          341   75             21  210             92  266   (b): class 2
          125   69             96  375            116  241   (c): class 3
           63   64             95  404            108  218   (d): class 4
           52   58            112  485            105  237   (e): class 5
           47   68            116  436             97  241   (f): class 6
           56   60             97  394             88  229   (g): class 7
           44   48             98  386            130  291   (h): class 8
           71   70             40  300            129  377   (i): class 9
          174   72              1   57             73  613   (j): class 10
```

Figure 8.4: The C4.5 induced tree produced from a second independent sample (sample 2) and evaluated on sample 1. The Confusion matrix shows how the unseen cases in sample 1 were categorised by the induced tree. This can be compared with figure 8.3 above.

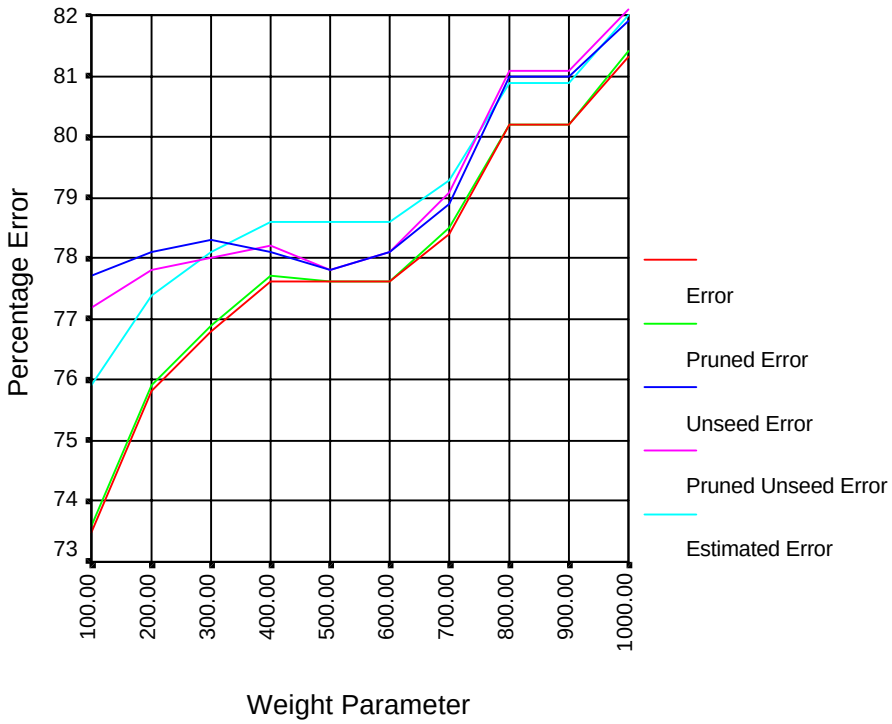

Figure 8.5: The effect of different weight parameter values on the accuracy of the induced decision tree for the initial sample (sample 1). The unseen error values are calculated using the second independent sample (sample 2).

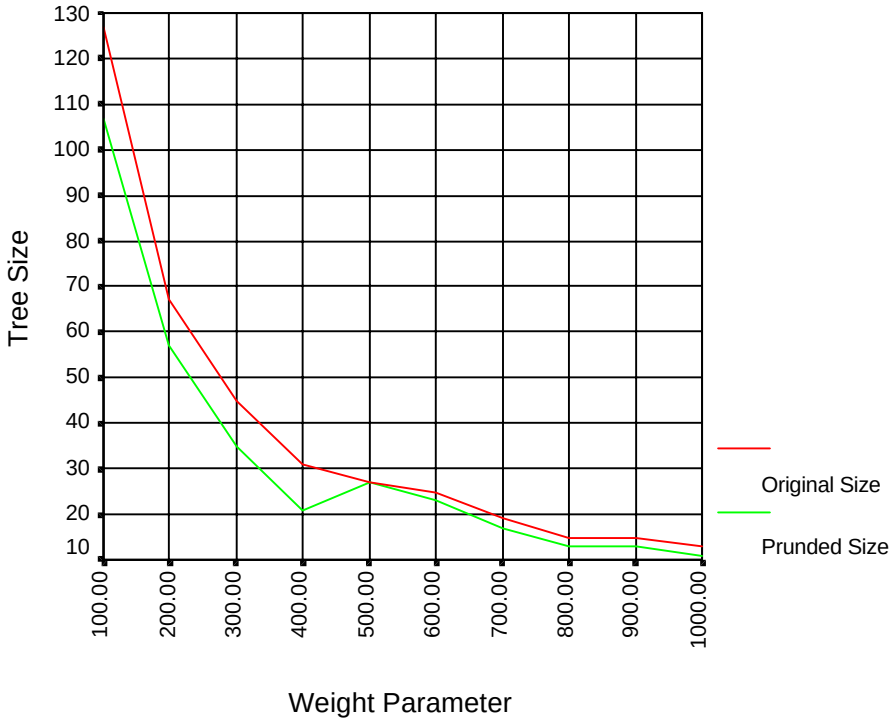

Figure 8.6: The effect of different weight parameters on the size of the induced decision tree for the initial sample (sample 1).

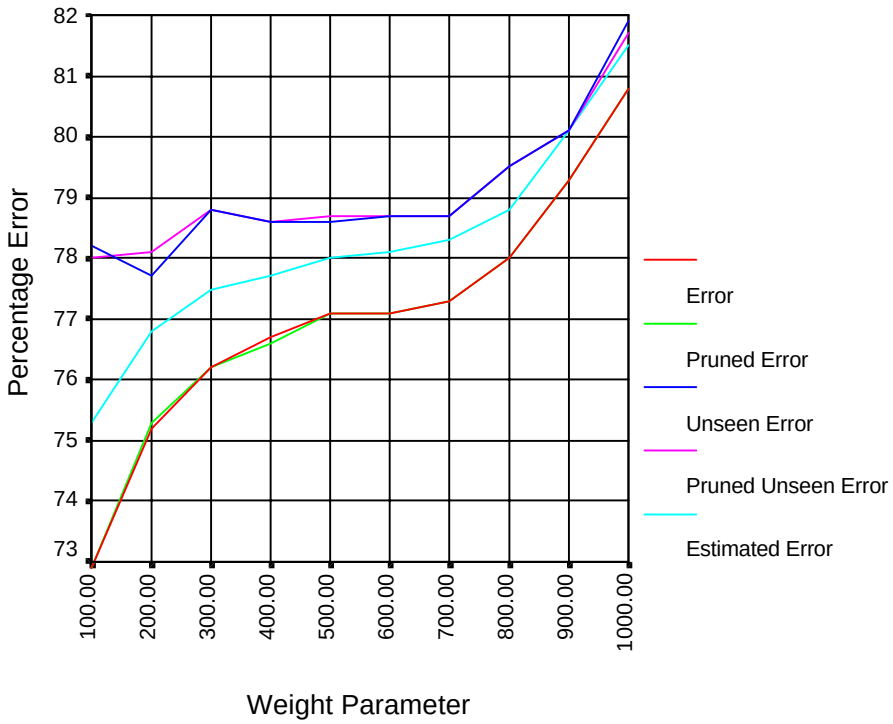

Figure 8.7: The effect of different weight parameter values on the accuracy of the induced decision tree for the second sample (sample 2). The unseen error values are calculated using the previous independent sample (sample 1).

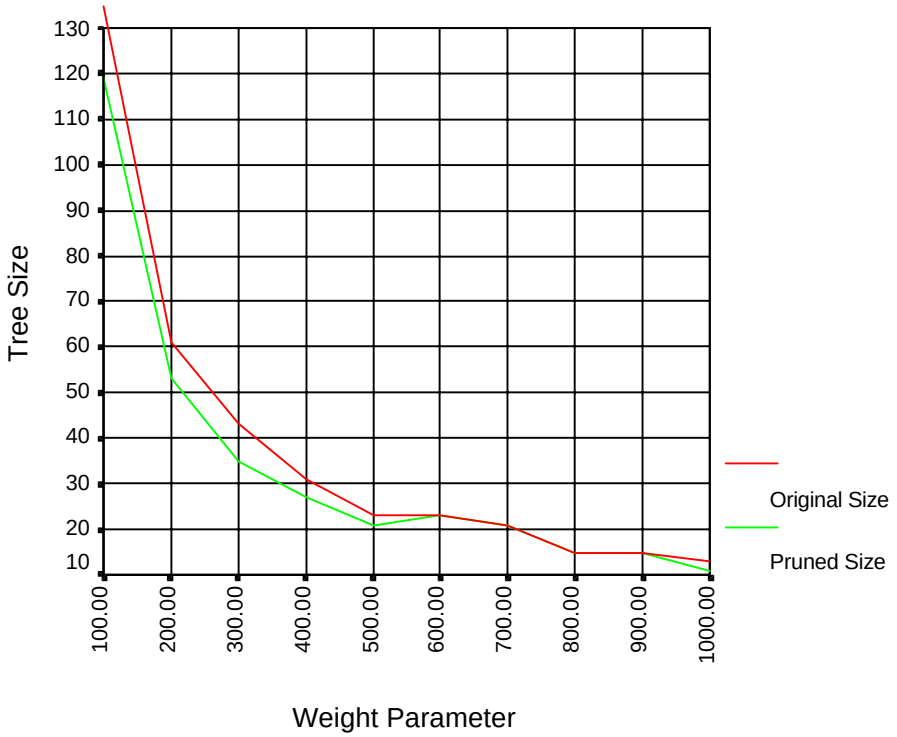

Figure 8.8: The effect of different weight parameters on the size of the induced decision tree for the second sample (sample 2).

Sample 1:

mt=0	class 1, [0.55] [0.35]
mt>0, cr=0	class 1, [0.60] [0.30]
mt>0, cr>0, vg=0	class 10, [0.68] [0.22]
mt>0, cr>0, vg>0, pm>0, ms=0.1	class 10, [0.84] [0.06]
mt>0, mt<=0.1, cr>0, vg>0, pm>0, ms>0.1	class 10, [0.83] [0.07]
mt>0.1, cr>0, vg>0, pm>0, ms>0.1, fg<=0.2	class 10, [0.84] [0.06]

Sample 2:

mt=0	class 1, [0.52] [0.38]
mt>0, cr=0	class 1, [0.60] [0.30]
mt>0, cr>0, vg=0	class 10, [0.64] [0.26]
mt>0, cr>0, vg>0, pm>0, ms=0.1	class 10, [0.82] [0.08]
mt>0, mt<=0.2, cr>0, vg>0, pm>0, ms>0.1, fg<=0.5	class 10, [0.77] [0.13]
mt>0.2, cr>0, vg>0, pm>0, ms>0.1, fg<=0.2	class 10, [0.85] [0.05]

Figure 8.9: The regions of high and low cooperation induced from both samples. Here all high and low regions induced are shown. A weight of 600 was used during the induction process. The numbers in square brackets indicate the quality of the region (see section 8.1.1 for a description of these).

can be viewed as a kind of "error value" indicating the class homogeneity of the region. The higher the value, the lower the quality of the region. The second value in square brackets is 0.9 minus the previous "error value". This value indicates the proportionate increase in predictive power over simply guessing the class of points within the region (since it is assumed that by guessing, a correct prediction would be made 10% of the time). As can be seen, although a weight value of 600 gives a relatively stable and compact tree representation - inducing almost identical regions for the lowest and highest percentile classes from the two independent samples - the quality of the regions found (in terms of "error value") is poor.

In an attempt to find higher quality regions an alternative approach was applied. A weight value of 100 was used (producing very large trees) and the best (in terms of low error value) three lowest and highest class regions were selected from the tree. This process was applied to both independent samples. The results are shown in figure 8.10. As can be seen higher quality regions are found and these are relatively stable across the two samples.

Sample 1

1a) mt>0, cr=0, ms<0.8, ms>0.3	class 1, (402/180) [0.45] [0.45]
1b) mt=0, ms<=0.9	class 1, (829/434) [0.52] [0.38]
1c) mt>0, cr>0, vg=0, fm<=0.1	class 10, (150/37) [0.25] [0.65]
1d) mt>0, cr>0, vg=0, fm>0.1, fc>0.1, mt<0.3	class 10, (148/92) [0.63] [0.27]
1e) mt>0.1, cr>0, vg>0, pm>0.3, fg<=0.1, fc>0.2, ms>0.1	class 10, (256/167) [0.66] [0.24]

Sample 2

2a) mt>0, cr=0, ms<=0.8, ms>0.3	class 1, (427/183) [0.43] [0.47]
2b) mt=0	class 1, (877/465) [0.53] [0.37]
2c) mt>0, cr>0, vg=0, fm<=0.1	class 10, (153/29) [0.20] [0.70]
2d) mt>0, cr>0, vg=0, fm>0.1, fc>0.1, pm>0.2	class 10, (374/259) [0.69] [0.21]
2e) mt>0, cr>0, vg>0, pm>0.4, fg<=0.1, fc>0.1	class 10, (284/187) [0.66] [0.24]

Figure 8.10: The best three regions found for the highest and lowest co-operative classes when the weight value is set to 100 for the two independent samples.

8.1.2 Regions of High and Low Co-operation

The tree induced with weights of 600 from the initial sample (figure 8.3) and the second sample (figure 8.4) located the regions of high and low co-operation shown in figure 8.9. The trees induced with weights of 100 from the two samples located the regions of high and low co-operation shown in figure 8.10. In the following sections explanations, informed by the analysis of typical individual runs falling within these regions, are given of the mechanisms producing high and low co-operation. Time series plots of typical individual runs are shown. Firstly, regions in which co-operation was low are examined. Secondly, regions in which high co-operation was found are considered. From the regions induced by C4.5 (see figures 8.9 and 8.10) it can be seen that when MT=0 or CR=0 co-operation is low. The mechanisms whereby low co-operation is produced in these regions, in addition to some typical individual runs, are given in sections 8.1.3 and 8.1.4. High co-operation appears to result from the relationship between the range of territories over which game interaction takes place and the amount of cultural interaction that takes place relative to game interaction. Sections 8.1.5 and 8.1.6 detail the mechanisms whereby high co-operation

is produced in these regions in addition to some typical individual runs.

8.1.3 Low Co-operation When Mutation is Zero

When MT=0 simulation runs tend towards lower levels of co-operation compared to that found over the whole parameter space (see figure 8.11). When MT=0 this indicates that no mutation of memes takes place at all. Under such conditions new memes cannot be created since there is no "cultural innovation". There is therefore a battle between only those memes in the initial population, with no new memes being produced. However, although the mean of co-operative interactions decreases from 0.26 to 0.23, the distribution stretches into the higher co-operation levels with a small peak at the highest possible co-operation level (see figure 8.11). This indicates that although there is a lot of low co-operation within this region there are also small amounts of high co-operation. It would seem that given some initial co-operative strategies within the population some runs allow those to spread to the entire population. A process that produces high co-operation within this region is discussed later (see section 8.2.2). Figures 8.12, 8.13 and 8.14 show the dynamics of some typical individual runs from the MT=0 region. In order to illustrate the dynamics of individual runs, over the distribution shown in figure 8.11, runs producing minimum, maximum and modal levels of co-operation were examined. In each case a single run was selected from sample 1 and then a set of replications with different initial pseudo-random number seeds were performed. For all the individual run graphs that follow, CC indicates the proportion of game interactions over a cycle that were mutually co-operative. DD and DC indicate the proportion of games that were mutually defective or mixed respectively. CF indicates the average confidence of all memes in the society. This measure gives an indication of how much satisfaction the society is sustaining. LE indicates the "label entropy" which

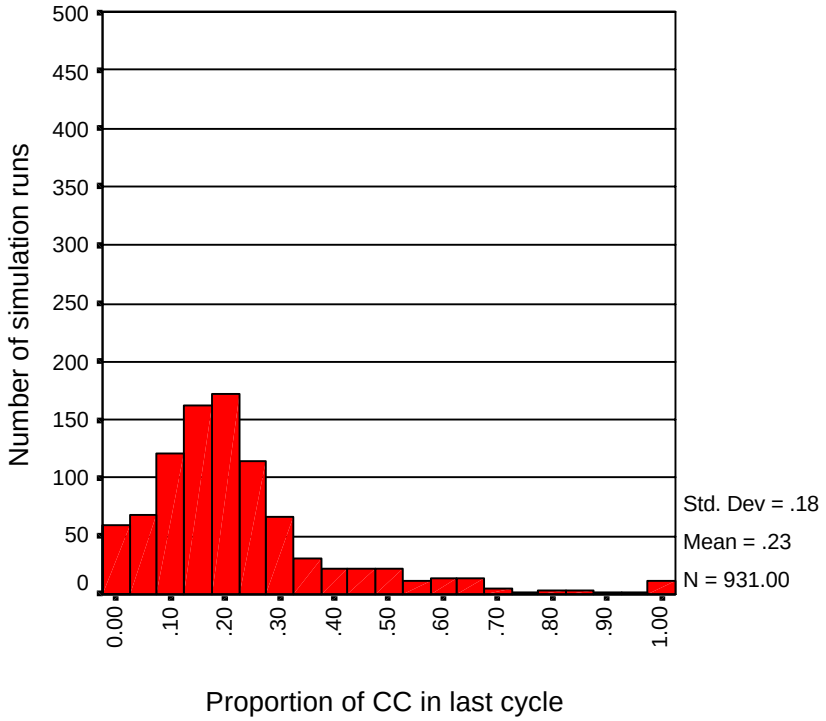

Figure 8.11: Frequency of co-operation in the region where MT = 0. The chart shows the frequency distribution of co-operation. Co-operation is distributed around the mean of 0.23. This can be compared to figure 8.1 showing the frequency of co-operation for the entire parameter space. Notice that there *are* a small number of highly cooperative points within this region. See section 8.2.2 and section 8.2.2 for a discussion these co-operative points.

measures diversity of agent labels (tags) see equations 8.1 and 8.2. Equation 8.1 gives the average number of bits required to identify an agent (as part of a group sharing the same bits) selected at random from the society (normalised into the range 0..1). Equation 8.2 selects the appropriate divisor for the normalisation. An LE value of zero indicates that all agents share the same label bits. An LE value of one indicates that agent labels are maximally distributed over all possible unique label patterns or that all agents possess a different label (if there are more possible unique labels than agents in the society). This entropy measure is basically a normalised form of Shannon's information content equation [139].

$$LE = \left| \frac{\sum_{i=1}^{g} \left(\frac{GS_i}{N} \times \log_2 \left(\frac{GS_i}{N} \right) \right)}{Mg(N,B)} \right| \tag{8.1}$$

$$Mg(N,B) = \left| \begin{cases} \log_2(N) \text{ iff } 2^b > N \\ B \text{ otherwise} \end{cases} \right| \tag{8.2}$$

Where g = number of groups with unique tags

GS_i = number of agents in group i

N = total number of agents in population

B = number of bits in tag

8.1.3.1 Minimum Co-operative Runs for MT=0

Figure 8.12(a) shows a single run from the MT=0 region which resulted in zero co-operative interactions. Notice that although overall agent confidence (CF) is low, no change in behaviour occurs because there is no mutation ("cultural innovation"). CF does not quite reach zero due to repeated reinforcement events as agents mutually replicate the same (non-cooperative) memes. Figures 8.12(b), (c) and (d) show three additional runs with

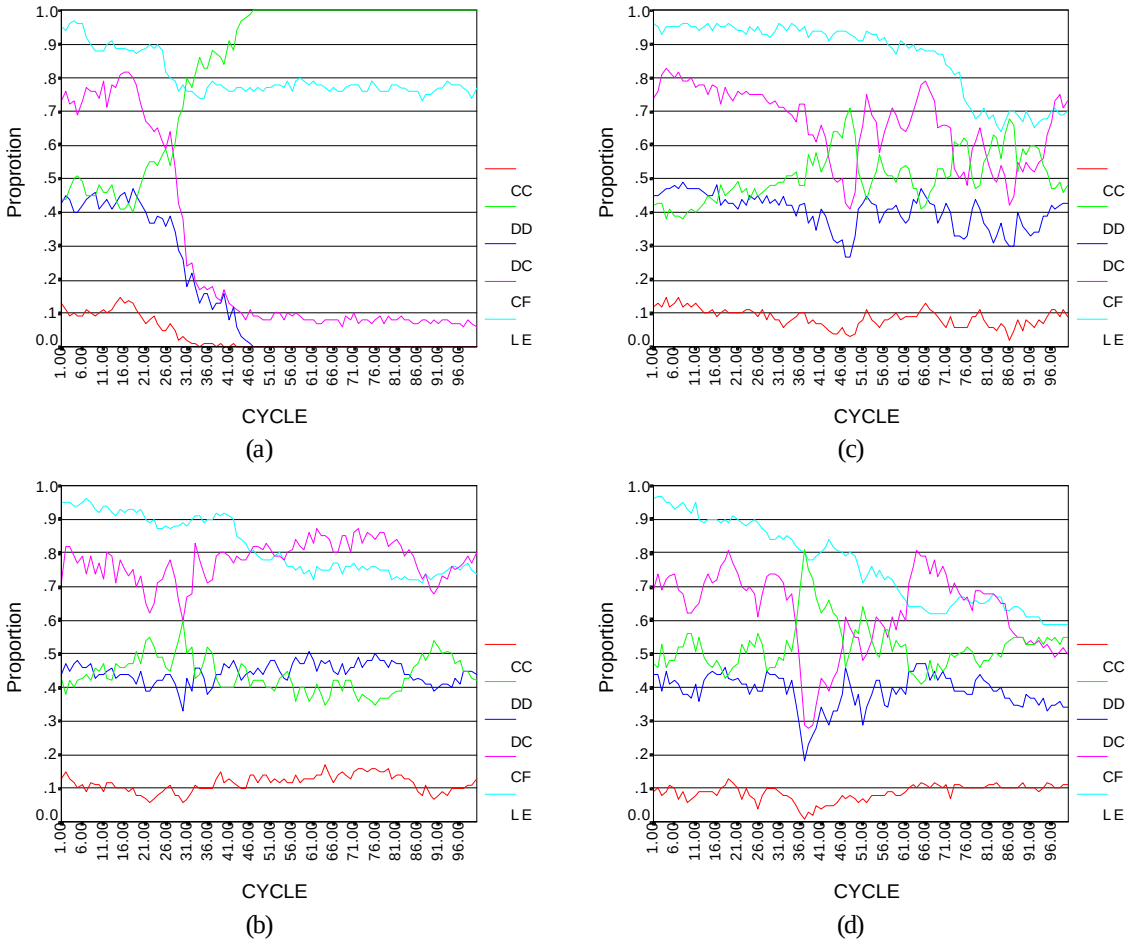

Figure 8.12: A set of runs from a point in region MT=0 producing minimum co-operation. CC, DD and DC show the relative proportions of game interactions. CF shows the average meme confidence level. LE shows the entropy of the label bits. Each run shows a time series for 100 cycles.

identical parameters but with different initial pseudo-random number seeds. Different seeds produce different dynamics (as might be expected). However, each run produces low co-operation levels. Each of the three replications ends with mutual co-operative interactions at approximately 10% of all game interactions.

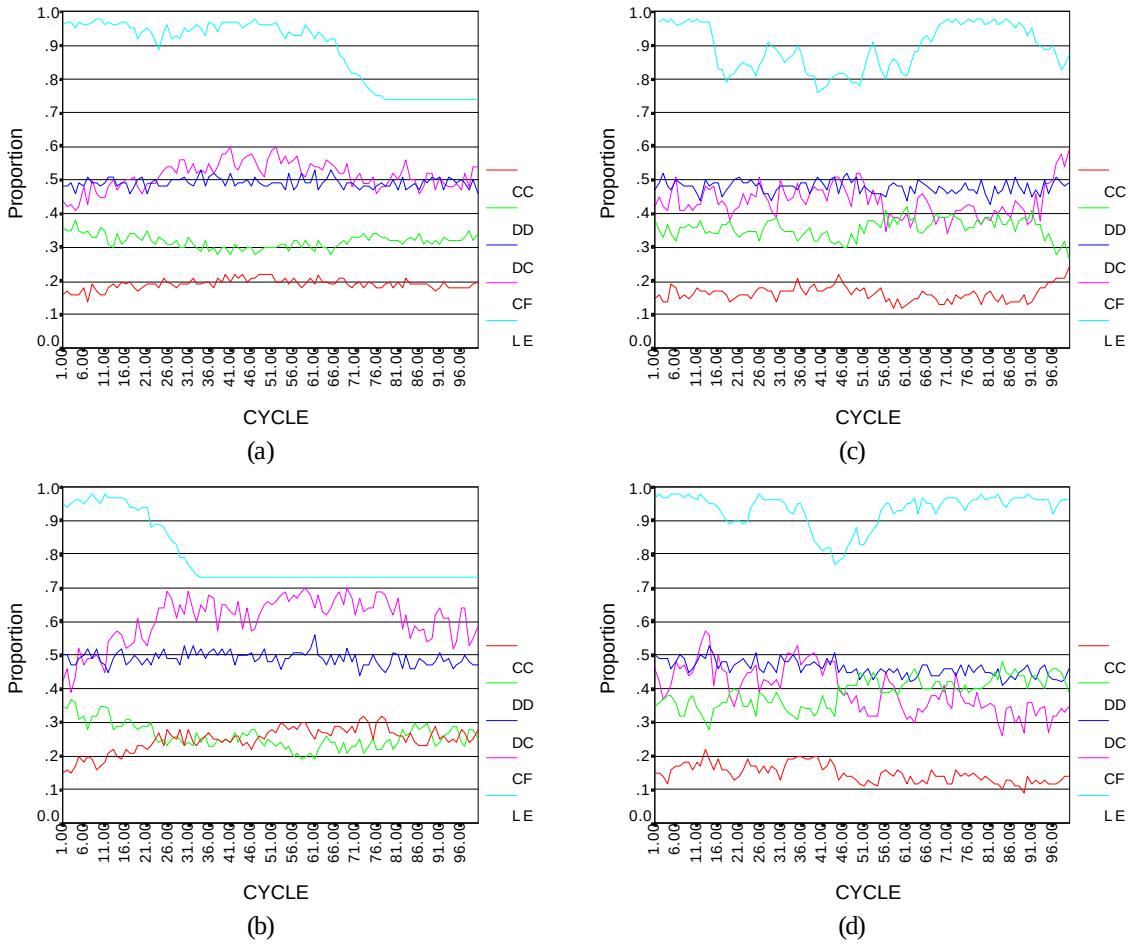

Figure 8.13: A set of runs from a point in region MT=0 producing modal co-operation. Run (a) results in mutual co-oopertive interactions at 20%. Compare this with graphs (b), (c) and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds.

8.1.3.2 Modal Co-operative Runs for MT=0

Figure 8.13(a) shows a single run from the MT=0 region which resulted in 20% mutually co-operative interactions. Figures 8.13(b), (c) and (d) show three additional runs with identical parameters but with different initial pseudo-random number seeds. Each run produces moderate co-operation levels. Each of the three replications ends with mutual co-operative interactions between about 15% and 25%.

8.1.3.3 Maximum Co-operative Runs for MT=0

Figure 8.14(a) shows a single run from the MT=0 region which resulted in 100% mutually co-operative interactions. Figures 8.14(b), (c) and (d) show three additional runs with identical parameters but with different initial pseudo-random number seeds. Each run produces high co-operation levels but there is a larger variation than found with the previous minimum and modal individual runs. Each of the three replications ends with mutual co-operative interactions between about 40% and 70%. This indicates that high co-operation within this region can be sensitive to initial conditions and stochasticities within the model. A mechanism which produces co-operation within this region is discussed later (see section 8.2.2).

8.1.4 Low Co-operation When Confidence Reduction is Zero

When CR=0, simulation runs tend towards lower levels of co-operation. When CR=0 this indicates that agents never reduce the confidence associated with individual memes. In this situation all memes will quickly gain a high confidence (due to infection and reinforcement) even if their hosts are not satisfied.

Figure 8.15 shows a frequency distribution of co-operation for this region. Co-operation is tightly distributed around a low mean and can be compared to the distribution produced over the entire parameter space (figure 8.1).

All the runs shown for this region (see figures 8.16, 8.17, 8.18) follow the same general pattern. The CF value stabilises at the maximum by the end of the first cycle. The EL value (label entropy) stabilises at the initial high value. This is an inevitable consequence because when CR=0 confidence values associated with memes can never be reduced. Within the initial cycle, all memes will quickly reach maximum confidence due to reinforcement

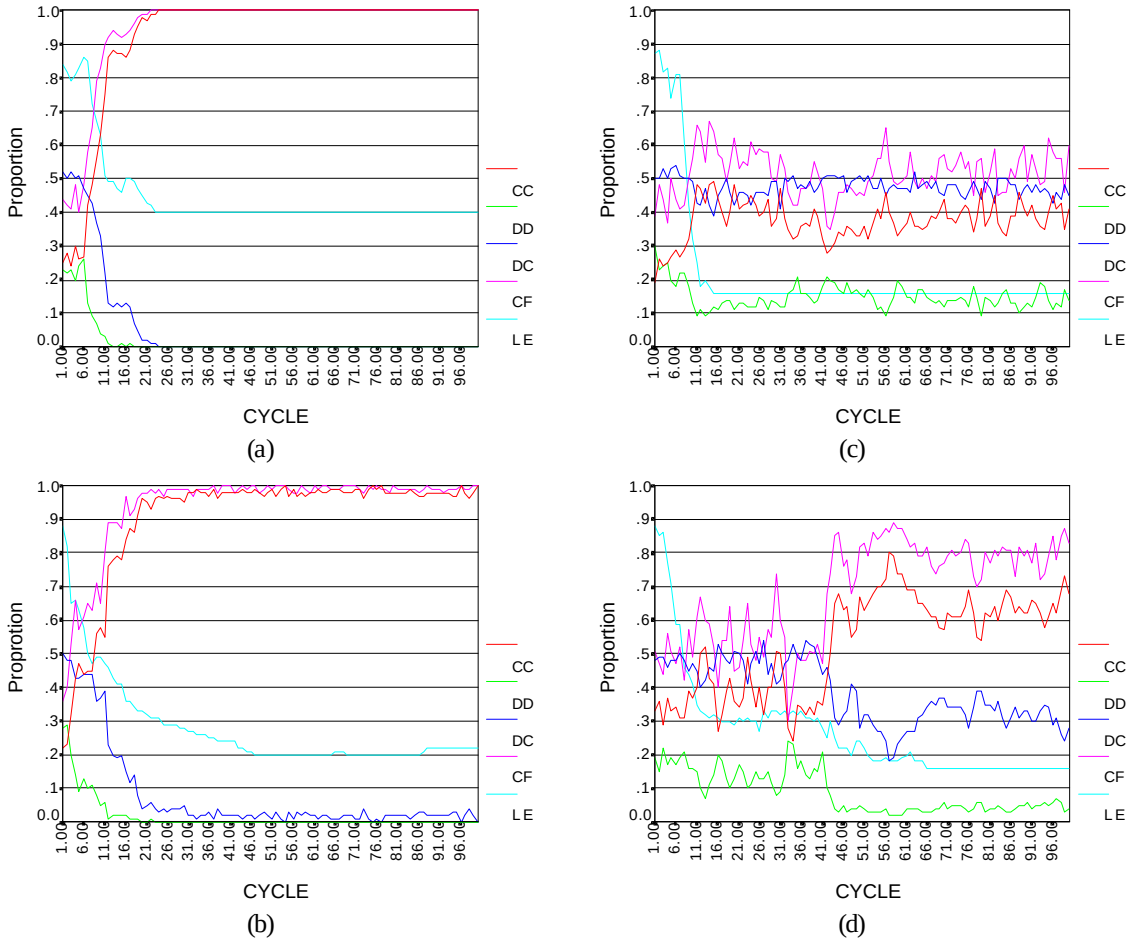

Figure 8.14: A set of runs from a point in region MT=0 producing maximum co-operation. Run (a) resulted in mutual co-operative interactions at 100%. Compare this with graphs (b), (c), and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. Notice that as co-operation increases the LE measure falls.

and satisfaction operations. When all confidence values are at a maximum (CF=1) meme evolution effectively stops for two reasons: Existing memes cannot be changed because *mutation* does not operative on memes with maximum confidence, and memes cannot spread by *replication* because memes with maximum confidence always repel challengers. It is therefore to be expected that runs within this region will result in fixed (though stochastic) behaviours. The random perturbations of game interaction types over time are a result of the stochastic pairing and agent rule selection methods. The runs within this region are of use in order to show what can be achieved (co-operatively) when there is practically no memetic evolution occurring. It is of interest that high and zero levels of co-operation are absent from the cases in this region (figure 8.15). This indicates that high and low co-operation results from processes involving meme evolution (which is not occurring here). Figures 8.16, 8.17, 8.18 show the dynamics of some typical individual runs within the region CR=0. In order to illustrate the dynamics of points within the region, individual runs producing minimum, maximum and modal levels of co-operation were examined. In each case a single run was selected from the region and then a set of three replications with different initial pseudo-random number seeds was performed.

8.1.4.1 Minimum Co-operative Runs for CR=0

Figure 8.16(a) shows a single run from the CR=0 region which resulted in the lowest proportion of co-operative interactions. Figures 8.16(b), (c) and (d) show three additional runs with identical parameters but with different initial pseudo-random number seeds. Each run produces low co-operation levels. Each of the three replications ends with mutual co-operative interactions between approximately 7% and 12% of all game interactions.

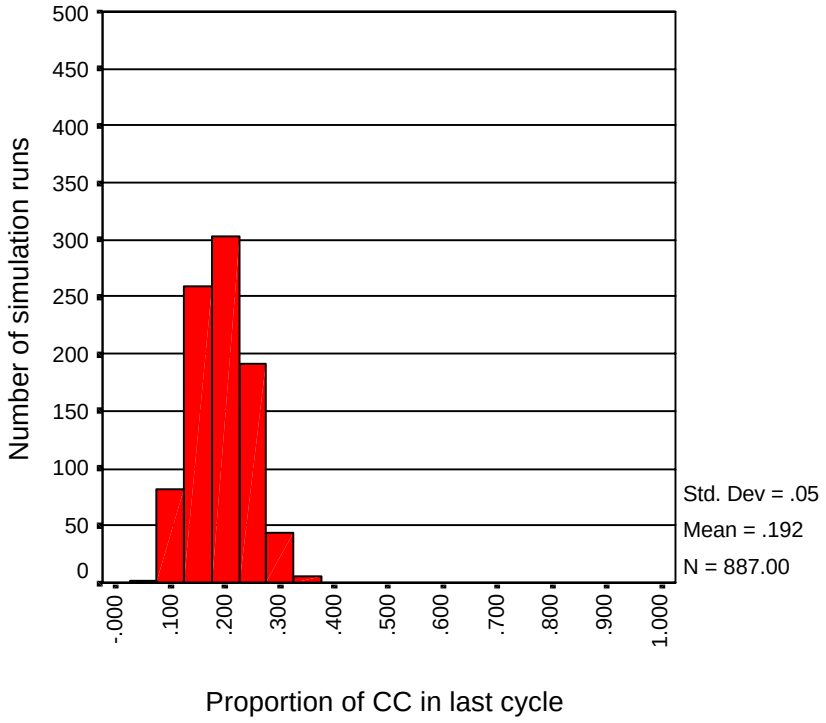

Figure 8.15: Frequency of co-operation within the region where CR=0. The chart shows the frequency distribution of co-operation. Co-operation is tightly distributed around the mean of 0.19. This can be compared to figure 8.1 showing the frequency of co-operation for the entire parameter space and figure 8.11 which shows the frequency of cooperation within the region MT=0.

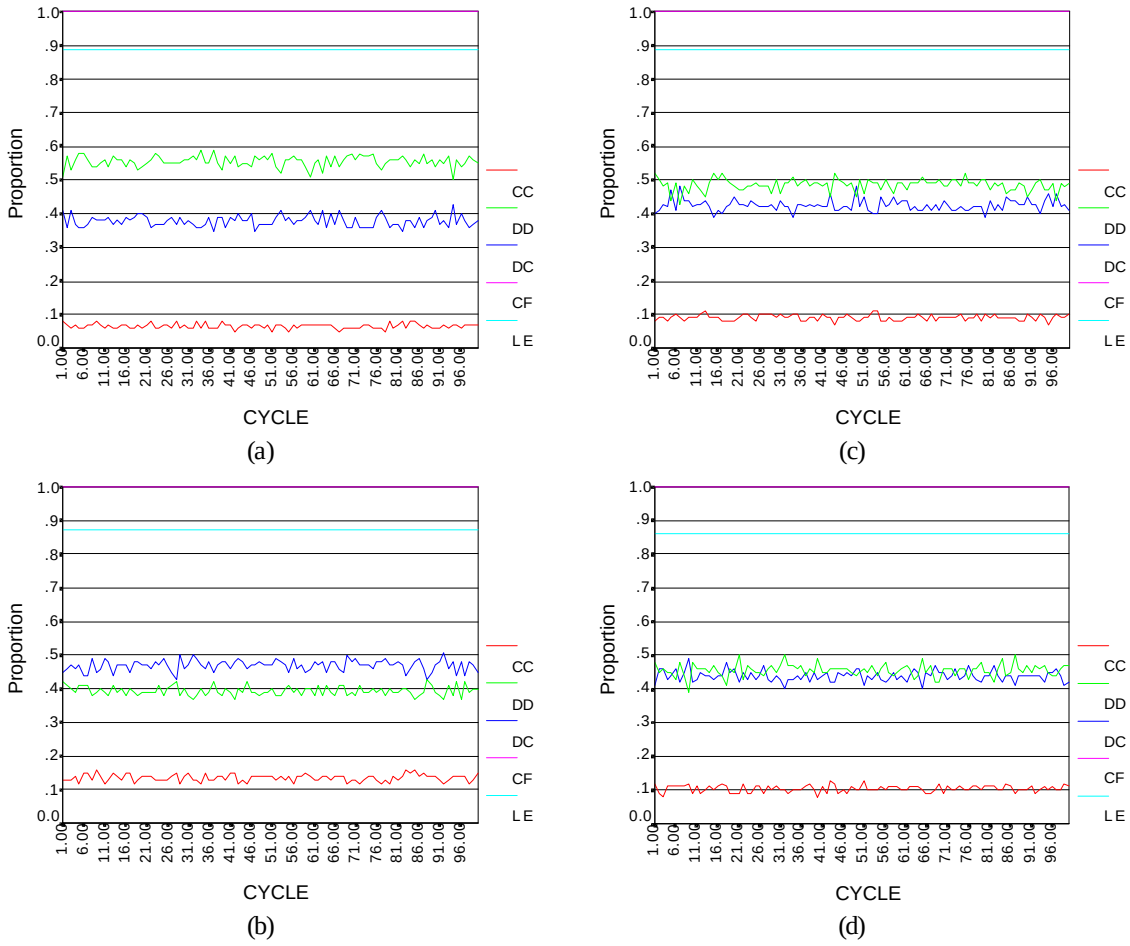

Figure 8.16: A set of runs from a point in region CR=0 producing minimum co-operation. Run (a) resulted in mutual co-operative interactions at 7%. Compare this with graphs (b), (c), and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. As can be seen all runs are flat save for random perturbations.

8.1.4.2 Modal Co-operative Runs for CR=0

Figure 8.17(a) shows a single run from the MT=0 region which resulted in 20% mutually co-operative interactions. Figures 8.17(b), (c) and (d) show three additional runs with identical parameters but with different initial pseudo-random number seeds. Each run produces moderate co-operation levels. Each of the three replications ends with mutual co-operative interactions between approximately 20% and 25% of all game interactions.

8.1.4.3 Maximum Co-operative Runs for CR=0

Figure 8.18(a) shows a single run from the CR=0 region which resulted in the highest levels of mutually co-operative interactions. This is low compared to the levels found within other regions. Figures 8.18(b), (c) and (d) show three additional runs with identical parameters but with different initial pseudo-random number seeds. The runs produce co-operation levels between approximately 35% and 25%.

8.1.5 High Co-operation When Game Interaction Limited to a Single Territory

When VG=0 agent game interaction is strictly limited to a single territory. This precludes agent game interaction outside of the territory. Agents are initially distributed randomly (from an uniform distribution) over all territories. Some territories may contain no agents others may contain several. Under the constraint VG=0 agents which are the sole occupants of a territory never get a game interaction while they remain in the territory. However, if agents subsequently move into an occupied territory they can resume normal game interaction. Agents interacting within their own territory limit their interaction partners substantially. It would appear that such a situation would reduce the space in which

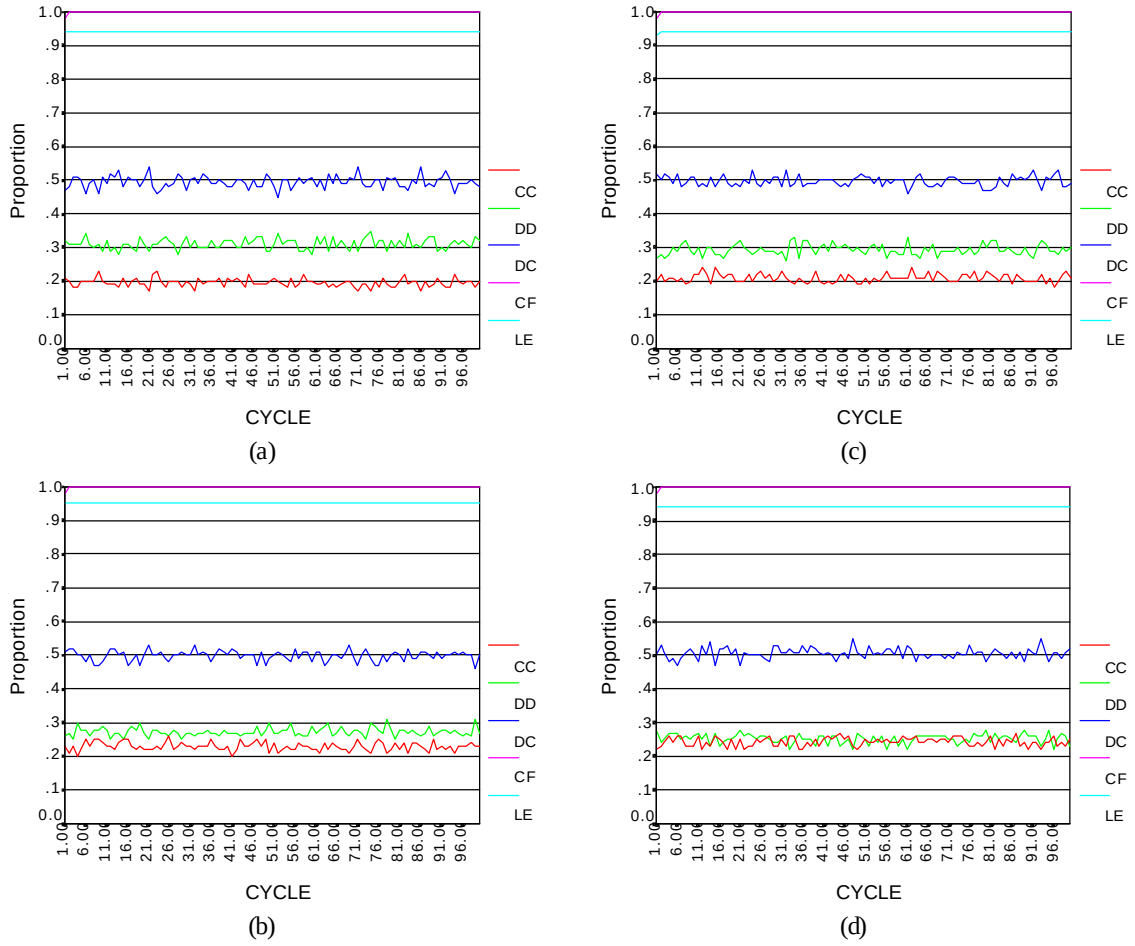

Figure 8.17: A set of runs from a point in region CR=0 producing modal co-operation. Run (a) resulted in mutual co-opertive interactions at 20%. Compare this with graphs (b), (c), and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. All runs are flat save for random perturbations.

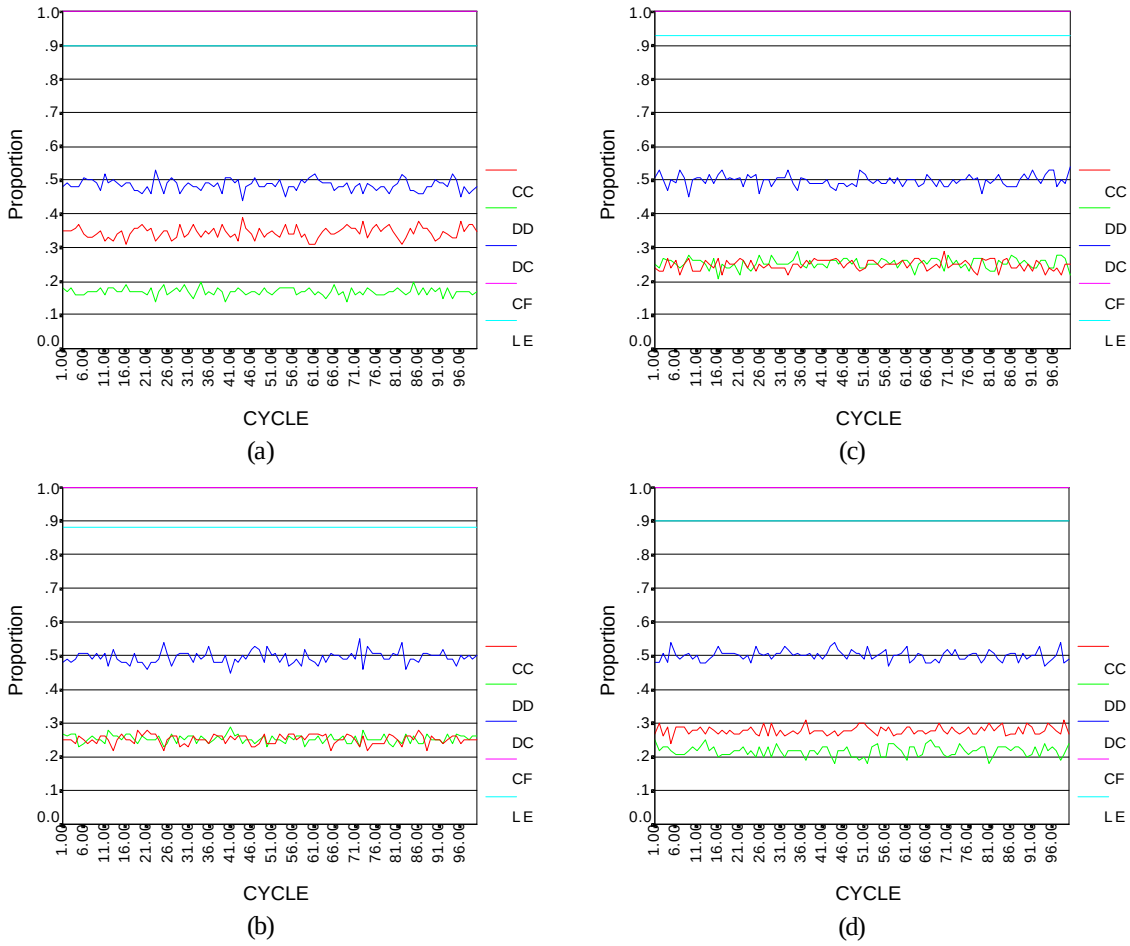

Figure 8.18: A set of runs from a point in region CR=0 producing maximum co-operation. Run (a) resulted in mutual co-opertive interactions at 35%. Compare this with graphs (b), (c), and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. All runs are flat save for random perturbations.

agents mutually search for co-ordinated action. Agent movement is governed by the FM parameter. In region 1c) of figure 8.10 it is notable that FM is low but not zero. This indicates that a small amount of movement does not destroy the formation of co-operative interactions. Region 1d) shows that for higher levels of agent movement FM>0.1 substantially lower levels of high co-operation are found. The region also specifies that cultural interaction should be at least FC>0.1. This appears to indicate that higher movement can be countered by higher levels of cultural interaction. Figure 8.19 shows the frequencies of co-operation within region 1c). Notice that a bimodal distribution is shown, indicating that although high co-operation is prevalent within this region there are still many points with lower levels of co-operation. Figures 8.20, 8.21 and 8.22 show the dynamics of some typical individual runs within the region 1c) from figure 8.10 (VG=0, CR>0, MT>0, FM <= 0.1). In order to illustrate the dynamics of individual runs over the distribution shown in figure 8.19, individual runs producing minimum, maximum and modal levels of co-operation (for this region) were examined. In each case a single run was selected from sample 1 and then a set of three replications with different initial pseudo-random number seeds was performed.

8.1.5.1 Minimum Co-operative Runs for VG=0

Figure 8.20(a) shows a single run from the region VG=0, CR>0, MT>0, FM<= 0.1. This run resulted in the lowest proportion of co-operative interactions found in the region. Figures 8.20(b), (c) and (d) show three additional runs with identical parameters but with different initial pseudo-random number seeds. Different seeds produce different dynamics. Two runs produce low co-operation levels (around 10% or below) the other two produce high co-operation levels (55% and 90%). This indicates a degree of sensitivity to random initial conditions and ongoing stochastic events which may push the evolution of

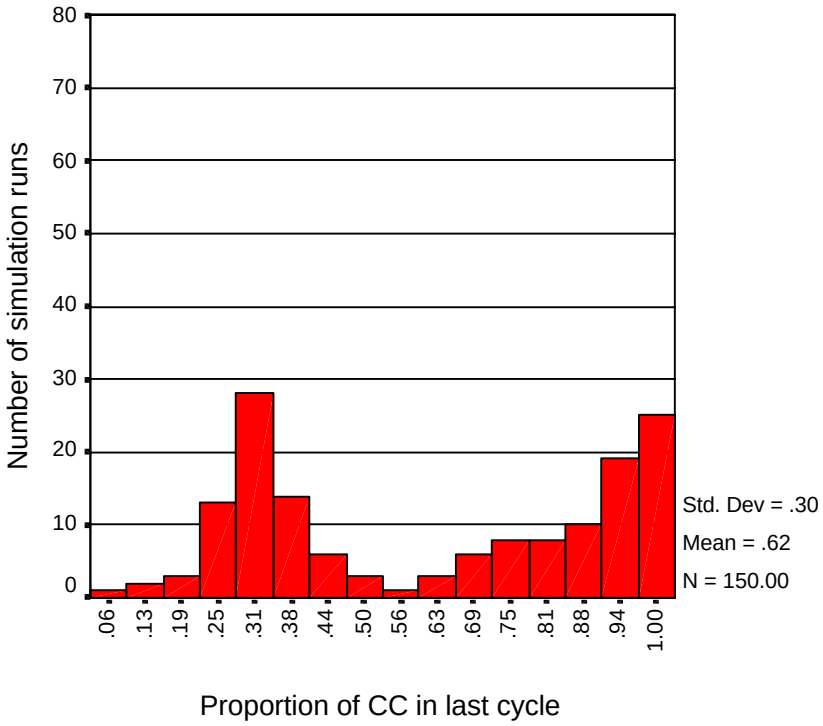

Figure 8.19: Frequency of co-operation within the region 1c (VG=0, CR>0, MT>0, FM<=0.1). The chart shows the frequency distribution of co-operation. Co-operation is distributed around two peaks at approximately 0.30 and 1.00. The distribution appears bimodal. This can be compared to figure 8.1 showing the frequency of co-operation for the entire parameter space.

the system towards low or high levels of co-operation. This indicates that at least some of the low levels of co-operation found within this region may have produced high levels of co-operation with different random number seeds.

8.1.5.2 Modal Co-operative Runs for VG=0

Figure 8.21 shows a set of runs from a single point in the region VG=0, CR>0, MT>0, FM<= 0.1. The original run (a) resulted in co-operative interactions of approximately 30%. This run was selected from the lowest mode of the bimodal distribution shown in figure 8.19. The larger mode value is the same as the maximum value (see below). Figures 8.21(b), (c) and (d) show three additional runs with identical parameters to (a) but with different initial pseduo-random number seeds. Different seeds produce similar dynamics and co-operative interaction levels (between 30% and 35%). The results from this case indicate relative stability against the stochastic nature of the model.

8.1.5.3 Maximum Co-operative Runs for VG=0

Figure 8.22 shows a set of runs from a single point in the region VG=0, CR>0, MT>0, FM<=0.1. The original run (a) resulted in co-operative interactions of 100%. This run was selected from the maximum of those within the region. Figures 8.22(b), (c) and (d) show three additional runs with identical parameters to (a) but with different initial pseudo-random number seeds. Different seeds produce similar dynamics and co-operative interaction levels reach 100% in each case. Results from this case indicate relative stability against the stochastic nature of the model.

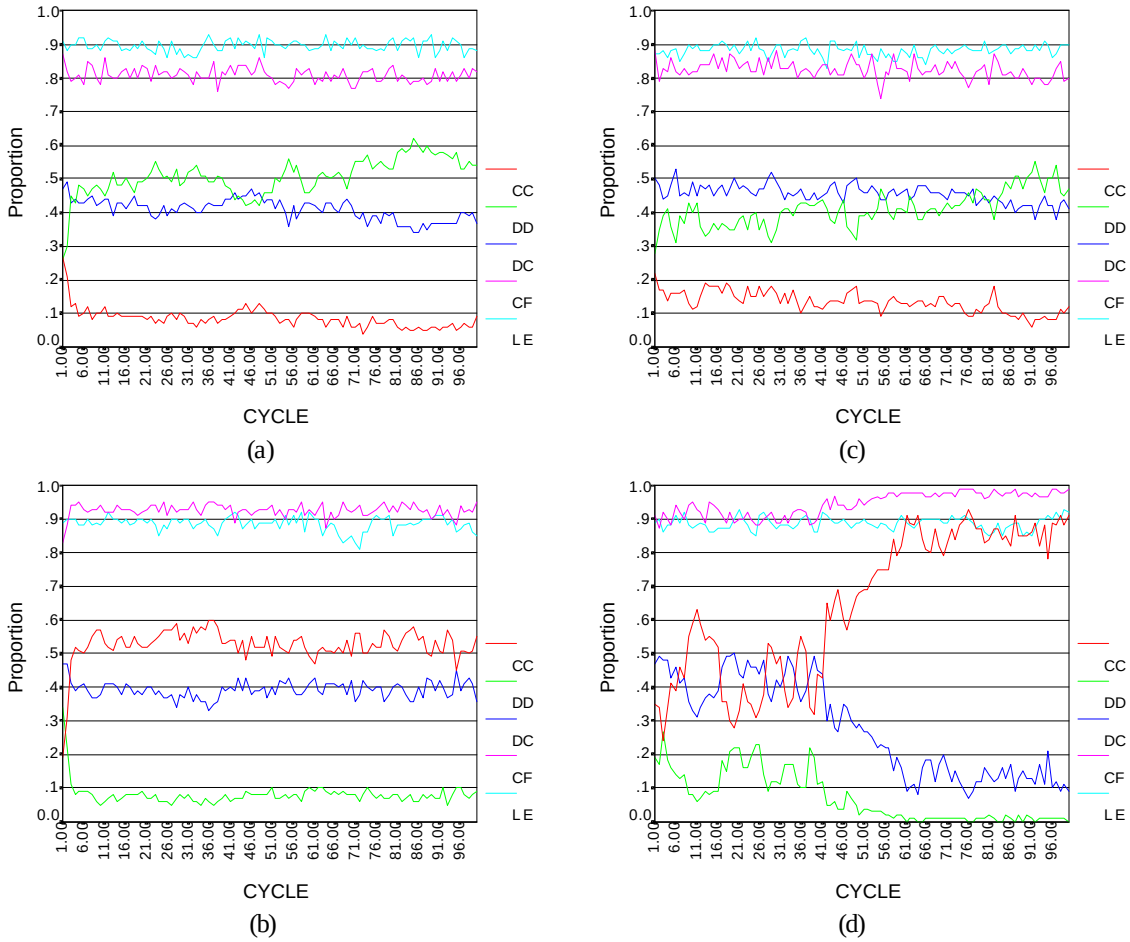

Figure 8.20: A set of runs from a point in region 1c (MT>0, CR>0, VG=0, FM<=0.10) producing minimum co-operation. Run (a) results in mutual co-operative interactions of less than 10%. Compare this with (b), (c) and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. In (b) and (d) a more highly co-operative trajectory is followed. In (d) the run follows a high co-operative trajectory. Notable here is an initial phase (up to about cycle 45) with oscillations of co-operation between about 30% and 60% then the increase in co-operation oscillating between about 80% and 90%.

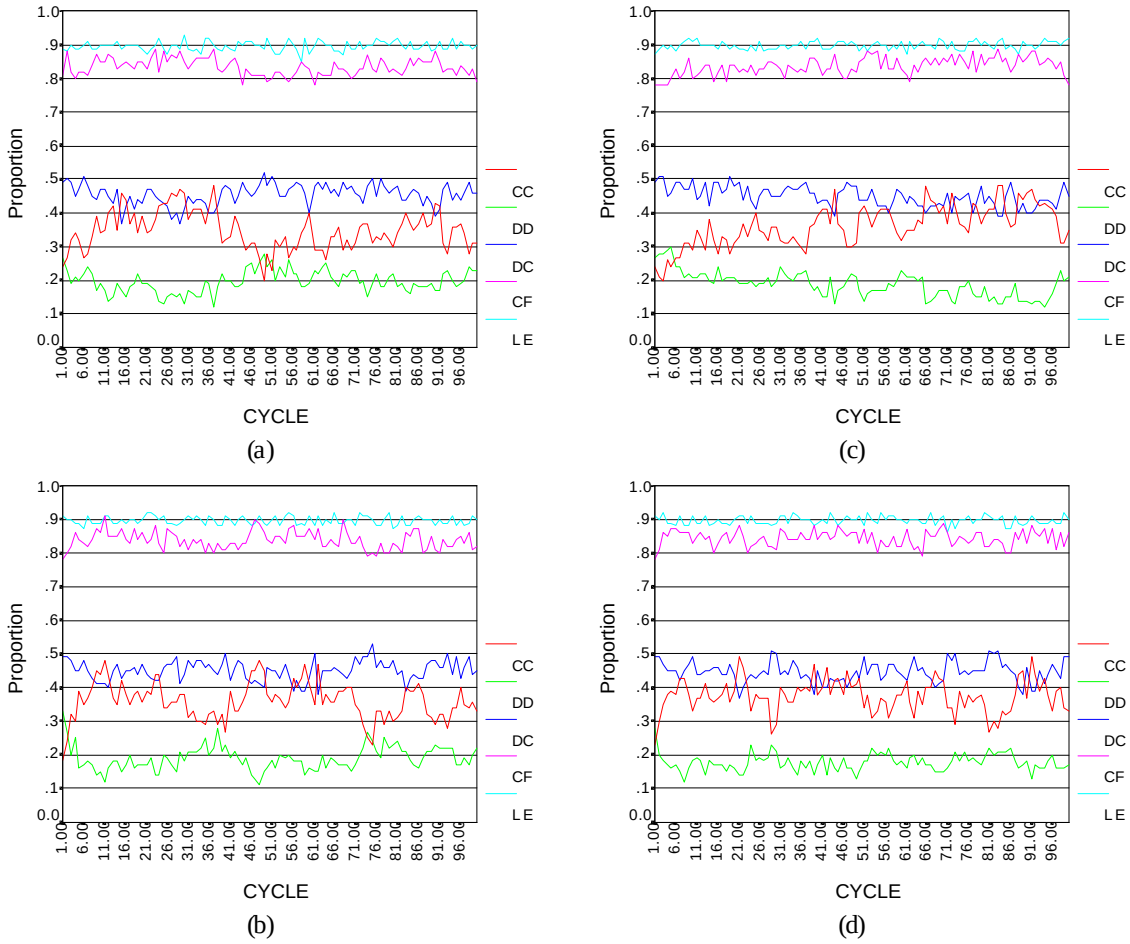

Figure 8.21: A set of runs from a point in region 1c (MT>0, CR>0, VG=0, PM>0.3, FM<=0.1) producing modal co-operation. Run (a) results in mutual co-operative interactions of about 30%. Compare this with (b), (c), (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. As can be seen, the dynamics and co-operative level are relatively stable.

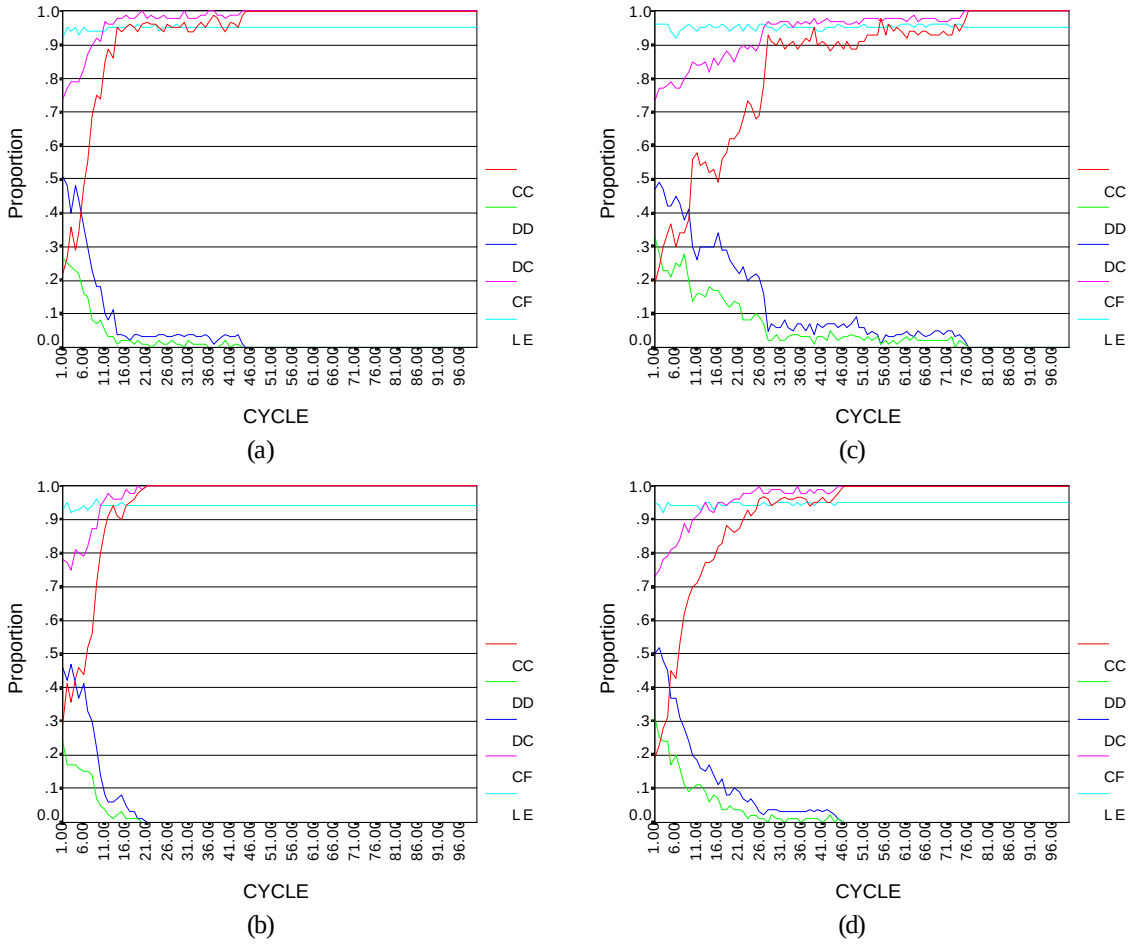

Figure 8.22: A set of runs from a point in region 1c (MT>0, CR>0, VG=0, PM>0.3, FM<=0.1) producing maximum co-operation. Run (a) results in mutual co-operative interactions of 100%. Compare this with (b), (c), (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. As can be seen, the dynamics and co-operative level are relatively stable.

8.1.6 High Co-operation When Cultural Interaction is High, Relative to Game Interaction

When VG>0, game interaction is not limited to a single territory. Region 1e) from figure 8.10 shows that under this condition high co-operation occurs when the proportion of cultural interaction is high, relative to the proportion of game interaction (FC>0.2, FG<=0.1). The amount of meme propagation occurring during cultural interactions is high with PM>0.3. This indicates that a single cultural interaction will involve the attempted propagation of (on average) 30% of the memes held by the proposing agent. The higher the value of PM the more propagation occurs in each cultural interaction. It would appear that *cultural interaction needs to be high relative to game interaction, when game interaction is spatially unconstrained, to produce high co-operation.* Although region 1e) contains the weak constraint MS>0.1 this is not induced in region 2e) from the second independent sample (see figure 8.10). Figure 8.23 shows the frequency of co-operation within this region. Figures 8.24, 8.25 and 8.26 show the dynamics of some typical individual runs within the region 1e) of figure 8.10 (MT>0.1, CR>0, VG>0, PM>0.3, FG<=0.1, FC>0.2, MS>0.1). In order to illustrate the dynamics of individual runs over the distribution shown in figure 8.23, points producing minimum, maximum and modal levels of co-operation for this region were examined. In each case a single run was selected from sample 1 and then a set of replications with different initial pseudo-random number seeds was performed.

8.1.6.1 Minimum Co-operative Runs for Cultural Interaction High, Game Interaction Low

Figure 8.24(a) shows a single run from the region which resulted in the lowest proportion of co-operative interactions. Figures 8.24(b), (c) and (d) show three additional runs

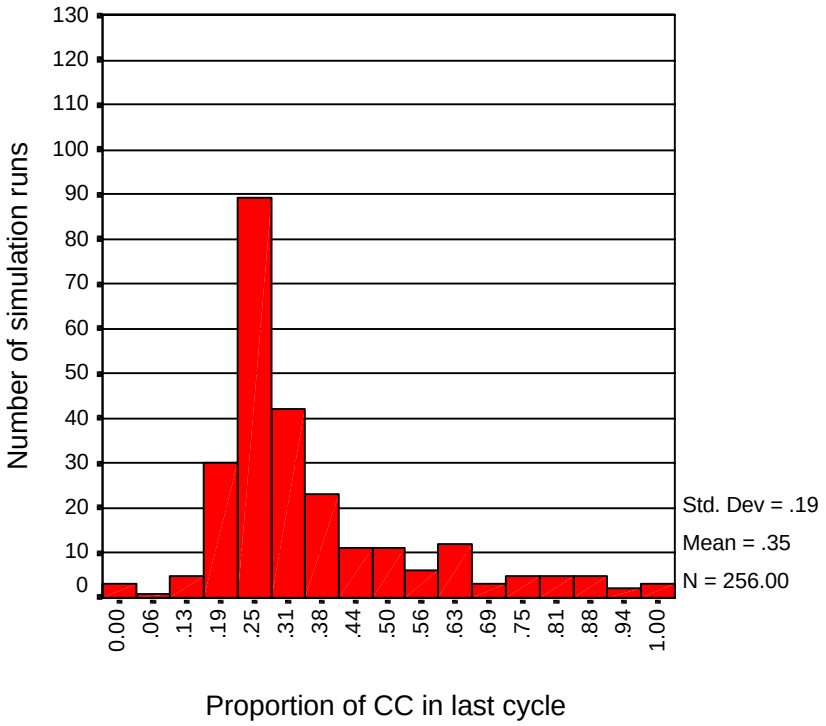

Figure 8.23: Frequency of co-operation within the region MT>0.1, CR>0, VG>0, PM>0.3, FG<=0.1, FC>0.2, MS>0.1. Co-operation is distributed around a mean of 0.35 within this region. This can be compared to figure 8.1 showing the frequency of co-operation for the entire parameter space. Notice that the distribution is skewed toward higher co-operative values.

with identical parameters but with different initial pseudo-random number seeds. Different seeds produce different dynamics. Two runs produce low co-operation levels (around 1% or below) the other two produce high co-operation levels (approximately 35% and 85%). This indicates a degree of sensitivity to random initial conditions and ongoing stochastic events which may push the evolution of the system towards low or high levels of co-operation.

8.1.6.2 Modal Co-operative Runs for Cultural Interaction High, Game Interaction Low

Figure 8.25(a) shows a single run from the region which resulted in a modal proportion of co-operative interactions (calculated in the region) of approximately 25%. Figures 8.25(b), (c) and (d) show three additional runs with identical parameters to (a) but with different initial pseudo-random number seeds. Different seeds produce *similar* dynamics and co-operative interaction levels (between 25% and 45%). Again, the results from this case indicate relative stability against the stochastic nature of the model even though the co-operation level oscillates over a large range.

8.1.6.3 Maximum Co-operative Runs for Cultural Interaction High, Game Interaction Low

Figure 8.26(a) shows a single run from the region which resulted in a high proportion of co-operative interactions (calculated in the region) of 100%. Figures 8.26(b), (c) and (d) show three additional runs with identical parameters to (a) but with different random number seeds. Different seeds produce similar dynamics and all result in 100% co-operation by the end of the run. The results from this case indicate a high degree of stability against the stochasticities of the model.

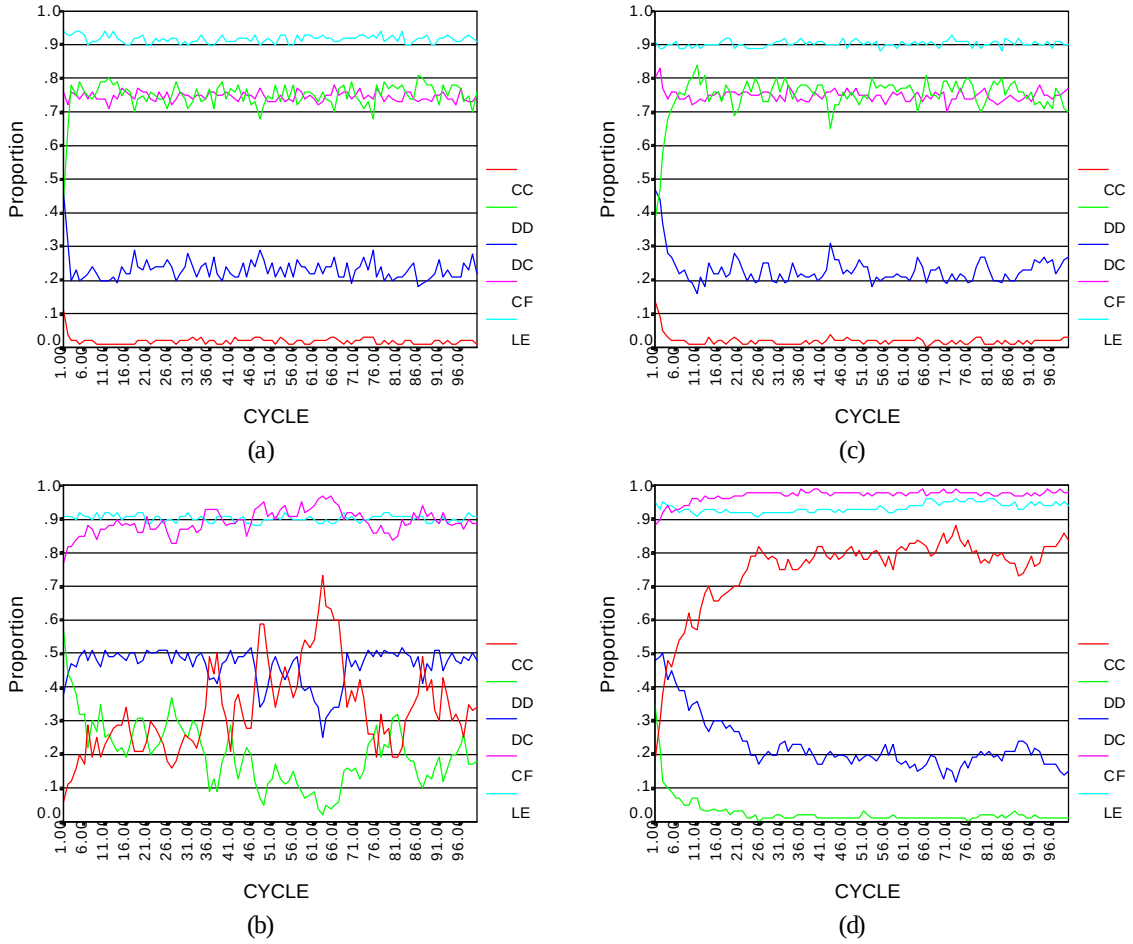

Figure 8.24: A set of runs from a point in region 1e (MT>0.1, CR>0, VG>0, PM>0.3, FG<=0.1, FC>0.2, MS>0.1) producing minimum co-operation. Run (a) results in mutual co-operative interactions of less than 1%. Compare this with (b), (c) and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. In (b) and (d) a more highly co-operative trajectory is followed

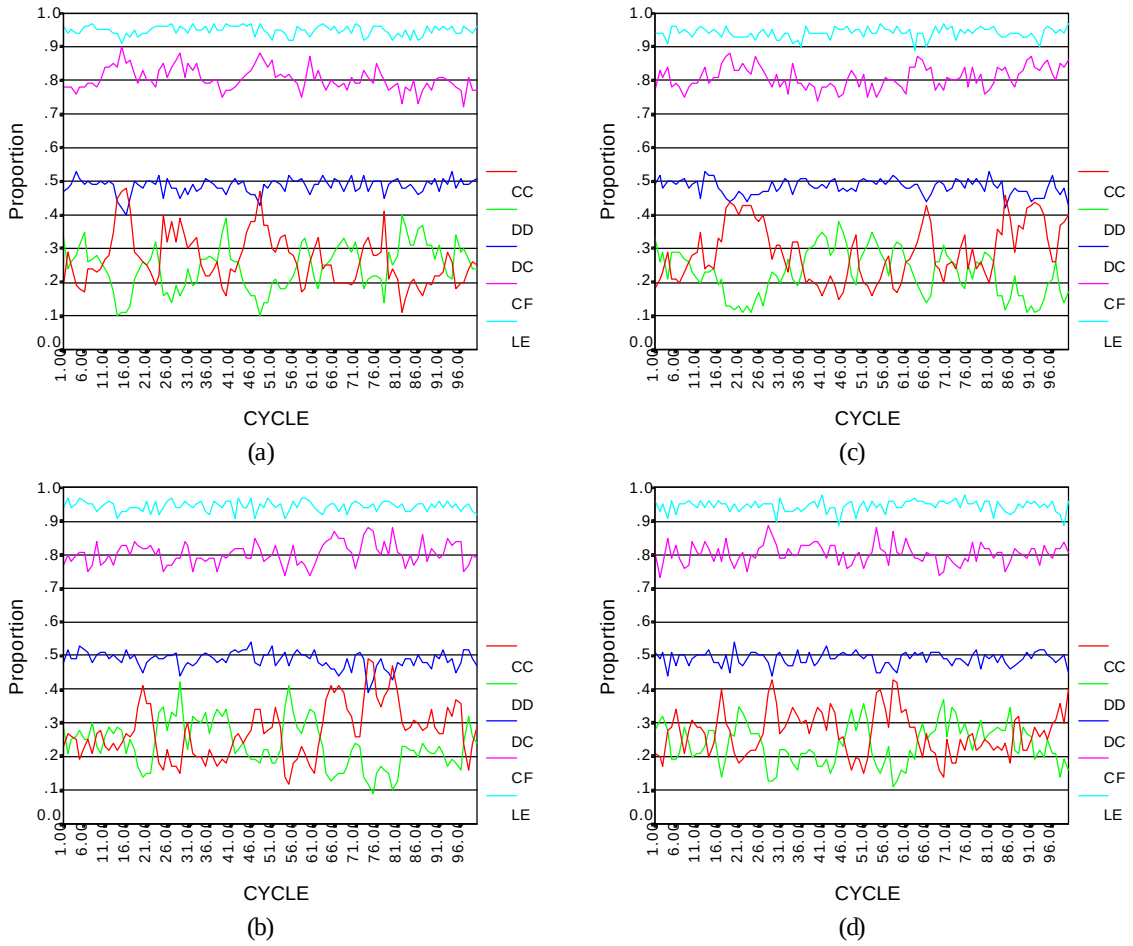

Figure 8.25: A set of runs from a point in region 1e (MT>0.1, CR>0, VG>0, PM>0.3, FG<=0.1, FC>0.2, MS>0.1) producing modal co-operation. Run (a) results in mutual co-operative interactions of less than 25%. Compare this with (b), (c) and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. The dynamics are similar. Co-operation level oscillates between about 15% and 45% in all of the runs.

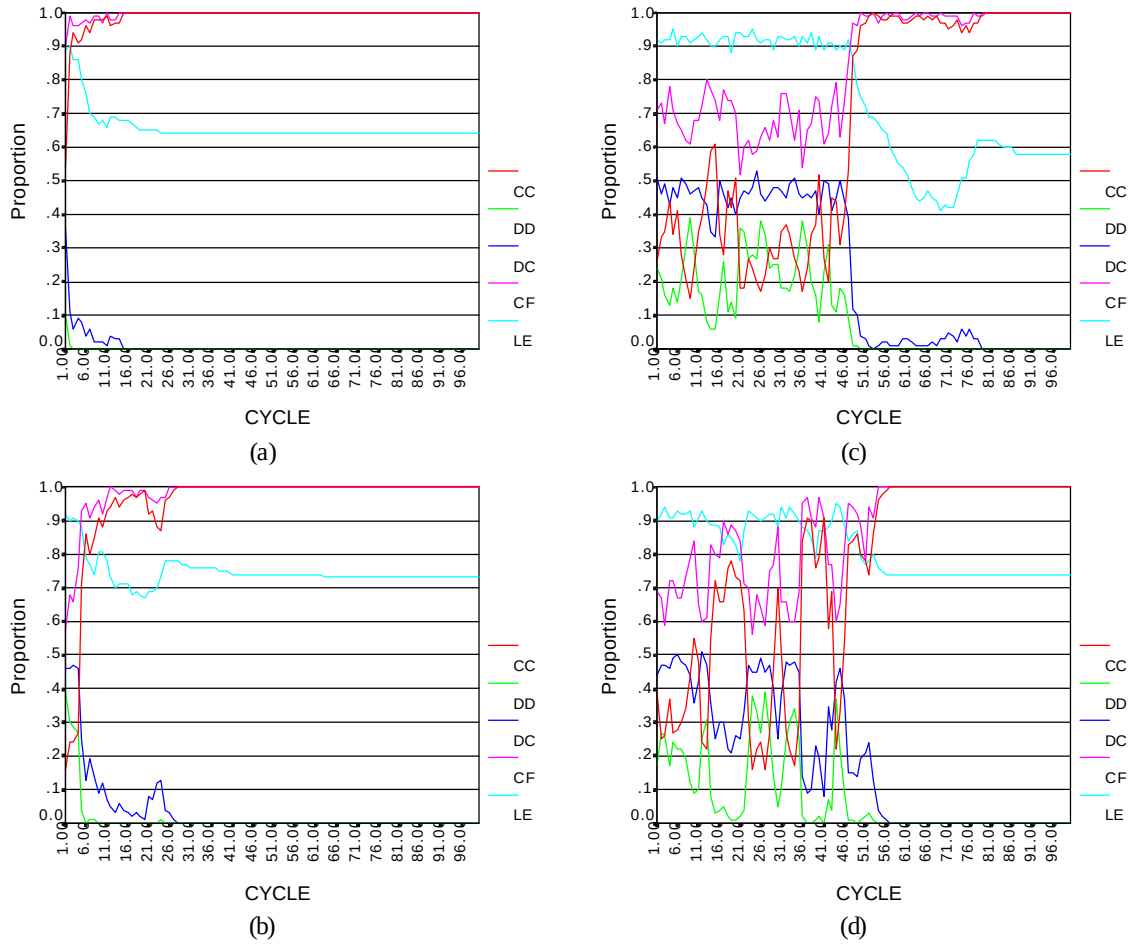

Figure 8.26: A set of runs from a point in region 1e (MT>0.1, CR>0, VG>0, PM>0.3, FG<=0.1, FC>0.2, MS>0.1) producing maximum co-operation. Run (a) quickly results in mutual co-operative interactions of 100%. Compare this with (b), (c) and (d) which show alternative runs with the same parameters but different initial pseudo-random number seeds. The dynamics are similar. There is a period of oscillation before lock-in of co-operation. Also the label entropy (LE) measure falls when co-operation reaches a high level. For each run, 100% co-operation is produced by the end of the run.

8.1.7 Summary of Results of C45 Induced Regions

The following is a summary of the results from the individual runs taken from within the induced regions given in figure 8.10:

- Co-operation is low when there is no "cultural innovation" (MT=0). Without cultural innovation in the form of meme mutation agents cannot adapt and co-ordinate their game interactions.

- Co-operation is low when there is no "cultural interaction" (CR=0). Again, without cultural interaction agents cannot adapt. Also when CR=0, mutation will also be low (whatever the value of MT).

- Co-operation is high when agents game interact with others in the same territory (VG=0). Agents have only a small set of game interaction partners. This makes the search space for co-ordinated game interactions smaller and so it is more likely that agents will find a co-operative convention with their small number of partners.

- Co-operation is high when the proportion of cultural interaction to game interaction is high (PM>0.3, FG<=0.1, FC>0.2). Under these constraints the amount of cultural interaction which takes place (measured in attempted individual meme propagations by agents) is at least one order of magnitude higher than the number of games played. A high amount of cultural interaction between games gives the agents more scope to adapt and hence co-ordinate their game interactions.

The results here may be reasonably extended to runs with longer numbers of cycles. When MT and CR are at zero longer runs would not produce more co-operation since in both these cases agents do not adapt their memes and therefore will not change

their behaviour. Also, in the high co-operation cases, since all agents are satisfied when all interactions are co-operative, maximum co-operation will tend to "lock-in" and stay at that maximum permanently[3]. This "lock-in" phenomena can be seen in figures 8.22 and 8.26. However, for many high co-operation points within these regions, maximum co-operation is not reached and so longer runs could produce different results.

8.2 Locating Co-operative Regions Using Hill-Climbing and Cluster Analysis

In order to locate regions of highest co-operation, in the StereoLab parameter space, a simple form of hill-climbing was used (see section 7.2 in chapter 7) to find points in the parameter space producing full (100%) co-operation over the last cycle of a simulation run. This means that *all* game interactions were mutually co-operative over the last cycle.

The hill-climbing involved a random start point followed by successive (single unit) random movements to adjacent points (including diagonally adjacent points) within the parameter space. Consequently, movement involved increasing or decreasing each parameter value by a single unit (or leaving the value the same). The specific action was selected randomly. If a movement increased the level of co-operation then the new position was accepted otherwise it was retracted. Climbing continued until either a global maximum was found (every game played over the last cycle was mutually co-operative) or until a computational limit had been reached (specified as a number of simulation runs).

In each of the following cases, 200 hill-climbers were executed for 100 steps, hence 20,000 simulation runs were executed in all (an identical amount of computational effort

[3]Lock-in will only occur if all game interactions continue long enough to push all meme confidence values to the maximum. Given this, agents will not mutate or take on new memes. Agent movement could change such a co-operative equilibrium since movement in this model is not linked to agent satisfaction. It is not possible therefore to be certain that lock-in has occurred if agent movement is occurring.

to that used in the previous random samples). Duplicate and sub-optimal points (i.e. those which did not produce full co-operation over the last cycle) were removed from the samples. The points were normalised into the range [0..1] on each dimension by dividing by the full range of each parameter (shown in table 6.1 in chapter 6). The points were then analysed using *k-means cluster analysis*. In order to use k-means cluster analysis the number of required clusters must be specified *a priori*. The number of clusters used in the final analysis was determined by executing k-means with cluster numbers from 2 to 10 and inspecting the reduction produced in the objective function (sum of the squared distances of each point from its cluster centroid). Each k-mean run was reproduced 3 times with different initial random starting centroids. In each case the best (i.e. lowest objective function value) result was chosen.

The k-means clustering method was chosen over a hiearachical method of clustering since it allows for analysis of the objective function relative to the number of clusters, thus enabling an informed choice to be made for the number of clusters finally selected for analysis.

The hill-climbing and clustering technique was applied to the same parameter space used in section 8.1 with the C4.5 algorithm and to an extended parameter space (see section 8.2.3 below) which identified another co-operative region with novel group formation processes leading to co-operation.

8.2.1 Clustering the Points

The hill-climbing produced 115 points producing maximum co-operation. In order to determine the effectiveness of different numbers of clusters for the k-means analysis several values were compared. Figure 8.27 shows the final value of the objective function

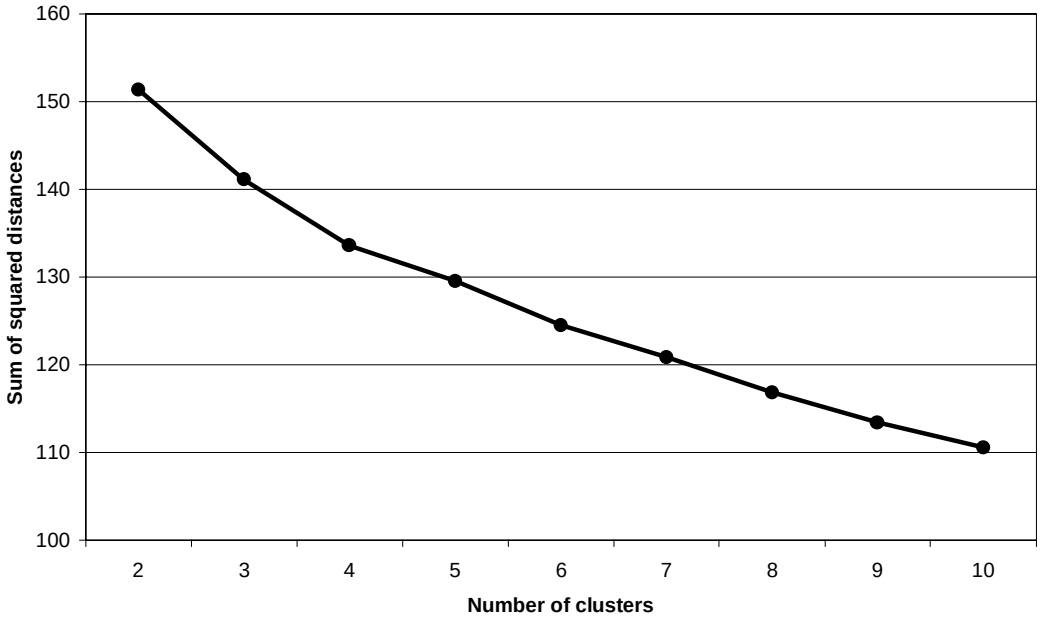

Figure 8.27: Final sum of squared distances of each point from its cluster centroid after k-means clustering for various numbers of clusters. Each point represents the best of three runs with different initial random centroids. The sample comprises 115 points.

for 2 to 10 clusters. In each case the best of three runs with different random starting centroids was selected. From examination of these results it was decided to select 4 clusters for the analysis. As can be seen in figure 8.27 the gradient of the line flattens slightly, indicating reduced returns for cluster numbers above 4. Also it was considered beneficial for reasons of comparison with the previous C4.5 analysis to keep the number of clusters low.

Table 8.2 shows the details of each of the 4 clusters. Interestingly, all clusters have high MS values indicating that overall high MS promotes co-operation. We can see this weakly captured by C4.5 over the random runs in regions 1b) and 1e) (see figure 8.10).

Clusters 2, 3 and 4 have very low VG and FM centroids. For cluster 3 both VG and FM centroids are at zero meaning that all members of the cluster have a value of zero for these parameters. This is similar to region 1c) identified by C4.5 (see figure 8.10). All of

Cluster Centres and Sizes				
	cluster 1	cluster 2	cluster 3	cluster 4
	size 18	size 26	size 30	size 41
param	ave (stdv)	ave (stdv)	ave (stdv)	ave (stdv)
B	.69 (.26)	.25 (.25)	.79 (.23)	.74 (.22)
M	.43 (.30)	.35 (.35)	.61 (.26)	.24 (.27)
PM	.69 (.27)	.76 (.23)	.47 (.32)	.62 (.24)
MT	.07 (.19)	.24 (.26)	.66 (.30)	.61 (.26)
CI	.39 (.30)	.64 (.25)	.41 (.28)	.51 (.28)
CR	.54 (.23)	.62 (.24)	.64 (.21)	.58 (.30)
MS	.79 (.26)	.74 (.25)	.77 (.27)	.63 (.26)
FG	.33 (.25)	.48 (.28)	.74 (.25)	.50 (.29)
FC	.71 (.21)	.51 (.26)	.47 (.30)	.68 (.29)
FM	.59 (.33)	.15 (.18)	.00 (.00)	.02 (.08)
BF	.24 (.22)	.25 (.22)	.44 (.26)	.56 (.28)
BG	.47 (.32)	.66 (.26)	.42 (.29)	.50 (.31)
BC	.56 (.33)	.47 (.32)	.17 (.15)	.63 (.27)
TG	.53 (.33)	.42 (.25)	.51 (.30)	.50 (.28)
TC	.64 (.33)	.43 (.29)	.66 (.29)	.48 (.27)
VC	.47 (.33)	.60 (.32)	.41 (.28)	.46 (.25)
VG	.47 (.35)	.09 (.20)	.00 (.00)	.01 (.04)

Table 8.2: Cluster centres and sizes for four clusters identified by the k-means clustering method.

the 30 points that comprise the cluster fall within region 1c). The cluster analysis confirms this region as one producing high co-operation. Example runs and analysis for this region are presented in section 8.1.5 above. Cluster 2 and 4 have low values for VG and FM with the majority of points that comprise these clusters having zero values for both these parameters.

The centroid for cluster 1 has VG and FM values that are near the median of the range. FC and PM are high relative to FG, indicating high levels of cultural interaction relative to game interaction. However, here we see something different from the regions induced by C4.5. This cluster differs from regions 1d) and 1e) shown in figure 8.10 because the centroid for MT is low. The majority of points within the cluster have an MT of zero and hence fall within region 2b) (see figure 8.10). But this previously induced region was

identified as a low co-operative region. Interestingly the points do not fall into region 1b) which excludes points with MS=1, but this constraint was only induced by C4.5 for one of the random samples (sample 1). However, some evidence of high co-operation occurring within the region MT=0 can be seen in figure 8.11 and figure 8.14 (see section 8.1.3). The runs shown in figure 8.14 represent a point that falls within cluster 1. Inspection of the runs reveals something distinctive over all previous individual runs. The label entropy (LE) measure (see section 8.1.3 and equations 8.1 and 8.2 above) falls dramatically over the course of each run. As the LE measure falls the amount of co-operation given by the CC measure tends to rise. It appears therefore that *some process involving the spread of tags between agents is related to the spread of co-operation.* The cluster 1 centroid shows a low BF value indicating that many of the bits in the label strings of the agents are changeable via cultural interaction (few are fixed). Also a marginally higher value for TG over the other cluster centres is seen. This suggests some kind of tag based game biasing process might be producing co-operation. However, this explanation is speculative. In order to explore the proposed explanation further the parameter space was extended and a new search made (see section 8.2.3 below) which found that a tag based process was indeed associated with high co-operation.

8.2.2 Summary of Results from Hill-Climbing and Clustering

The following is a summary of the results from the hill-climbing and cluster analysis of the points found (table 8.2):

- Co-operation is high when agents only game interact within a single territory (VG=0). This confirms what was found by application of the C4.5 algorithm (see section 8.1.7).

- Co-operation is high when cultural interaction is high, game interaction is low and

mutation is zero (MT=0). A cluster is found which produces high co-operation even though mutation is mainly zero within that cluster. However, the number of fixed bits in the tag is low and biasing of game interaction by tag is marginally higher than other clusters, suggesting some kind of tag based process promoting co-operation. Further support for a tag based process is indicated by an individual run that falls within the cluster (see figure 8.14) which shows a large decrease in label entropy (LE) as co-operation increases.

8.2.3 Clustering the Points in an Extended Parameter Space

In a previous section (section 8.2.1 above) hill-climbing and cluster analysis produced a cluster which identified high co-operation within a region where MT was low (mainly zero). This is at odds with the regions induced by C4.5. Inspection of an individual run and the cluster centroid suggested that tag biasing and harmonisation processes may have played some process in producing co-operation. The biasing of cultural and game interactions over tags are mediated by BG, BC, TG and TC parameters. These parameters indicate the number of refusals or "trys" an agent gets before forced game (TG) or cultural (TC) interaction with an agent which does not have a label string with at least a proportion of BG or BC shared label bits. If say, BG=1 and TG=10 this indicates that an agent will refuse game interaction if the partner selected does not share all the same label bits. Since TG=10 the agent will be allowed ten consecutive refusals before a forced interaction occurs. The contribution of label biasing for increasing co-operation has been discussed and demonstrated in work by other authors [143] and [29]. Also the formation of groups (based on tag similarity) appears intuitively to offer the possibility of a form of "group selection" (i.e. groups which co-operate do well and spread that co-operative behaviour to others).

To explore the validity of this intuitive explanation the parameter space was extended.

The range of TG and TC was extended from [1..10] to [1..200]. Since the number of agents within the society was left set at 101, this indicates that with high values of TG and TC, there is a high probability that agents can locate a compatible agent for interaction if one exists and is reachable spatially. The biasing process is therefore much stronger. The space was also quantised more finely. In the extended space the unit increment of real valued parameters was reduced from 0.1 to 0.01. This vastly increases the size of the space but allows for parameters to take on a greater range of values. Mutation (MT) in the previous space could not fall between 0 and 0.1. This does not allow for societies where agents have low but non-zero mutation rates. In this more finely quantised extended space the unit movement which hill-climbers make is smaller by a factor of ten. This means that more hill-climbing steps are required (compared to the previous space) to cover the same distance. As in the previous hill-climbing runs, 100 hill-climbers were executed for 200 steps (making a total of 20,000 individual simulation runs).

As expected the space was more difficult to search. After duplicate and sub-maximum points were removed, 39 points were found that produced maximum co-operation. As before, various different numbers of clusters were evaluated with the k-means method (see figure 8.28). From examination of these results it was decided to select 5 clusters for the analysis. As can be seen in figure 8.28, the gradient of the line reverses and then flattens, indicating reduced returns for cluster numbers above five[4]. It is disconcerting to note that even with 3 independent k-mean runs for 6 clusters the result is actually marginally worse than the best found for 5 clusters. This is a strong reminder that the k-means algorithm is far from exhaustive or optimal[5].

[4]For reasons of comparison with the previous results from C4.5 and cluster analysis it was considered beneficial to keep the number of clusters low.

[5]Although anomalous such a result is possible. Clustering was performed using the SPSS software.

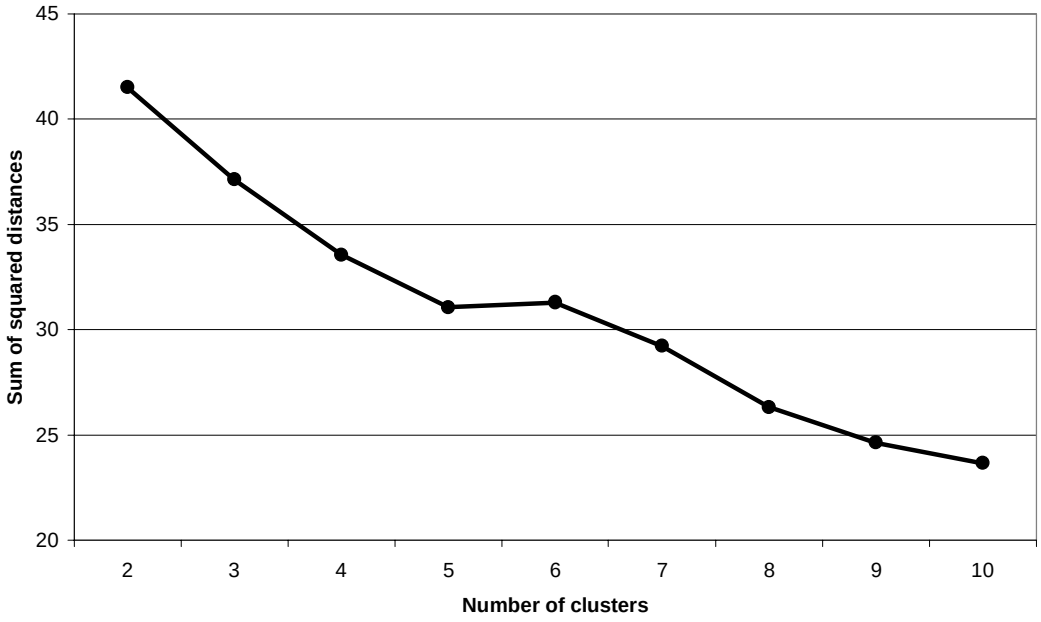

Figure 8.28: Final sum of squared distances of each point from its cluster centroid after k-means clustering for various numbers of clusters. Each point represents the best of three runs with different initial random centriods. The sample comprises 39 points.

Table 8.3 shows the cluster details obtained from clustering the 39 points with maximum co-operation (found via hill-climbing) from the extended parameter space into 5 clusters using k-means. The clusters can be divided into two broad categories: cluster 3 in which FM and VG are zero indicating game interaction limited to single territory (as discussed previously in section 8.1.7) and clusters 1, 2, 4 and 5 which have high PM and FC values relative to FG indicating high amounts of cultural interaction relative to game interaction. These latter clusters also have low BF and high BG and TG values. The high BG and TG values indicate high amounts of game biasing based on a high degree of tag similarity. Essentially game interaction is refused if agents do not share almost identical tags. The low BF means that the majority of the tags are learned and changed culturally since they are not fixed. MT is also low (but not zero) and interestingly M is low indicating that agents have a small memory and store few stereotypes.

Cluster Centres and Sizes					
	cluster 1	cluster 2	cluster 3	cluster 4	cluster 5
param	size 12	size 4	size 4	size 12	size 7
	ave (stdv)	ave (stdv)	ave (stdv)	ave (stdv)	ave (stdv)
B	.35 (.27)	.06 (.13)	.88 (.14)	.71 (.28)	.86 (.28)
M	.28 (.26)	.09 (.19)	.50 (.43)	.09 (.19)	.09 (.09)
PM	.73 (.22)	.78 (.15)	.65 (.29)	.85 (.17)	.71 (.23)
MT	.05 (.03)	.20 (.14)	.80 (.01)	.12 (.09)	.05 (.04)
CI	.59 (.27)	.32 (.23)	.24 (.12)	.33 (.21)	.54 (.31)
CR	.74 (.16)	.42 (.21)	.12 (.06)	.66 (.27)	.60 (.24)
MS	.60 (.24)	.50 (.25)	.55 (.13)	.76 (.21)	.59 (.26)
FG	.39 (.23)	.13 (.01)	.88 (.09)	.51 (.24)	.32 (.30)
FC	.73 (.21)	.70 (.29)	.71 (.27)	.66 (.16)	.69 (.13)
FM	.34 (.31)	.70 (.23)	.00 (.00)	.45 (.28)	.71 (.23)
BF	.36 (.18)	.19 (.10)	.51 (.36)	.27 (.16)	.39 (.29)
BG	.86 (.11)	.83 (.18)	.29 (.19)	.88 (.10)	.42 (.28)
BC	.40 (.25)	.53 (.31)	.60 (.24)	.64 (.18)	.37 (.31)
TG	.72 (.21)	.64 (.29)	.83 (.00)	.62 (.23)	.27 (.26)
TC	.45 (.35)	.34 (.34)	.50 (.44)	.78 (.15)	.34 (.29)
VC	.68 (.18)	.29 (.14)	.59 (.37)	.43 (.24)	.31 (.16)
VG	.65 (.30)	.38 (.24)	.00 (.00)	.45 (.39)	.62 (.34)

Table 8.3: Cluster centres and sizes for five clusters identified by the k-means clustering method in the extended parameter space.

8.2.4 Summary of Results from Hill-Climbing and Clustering in the Extended Space

The following is a summary of the results from the hill-climbing and cluster analysis of the points found (table 8.3):

- As found previously, co-operation is high when agents only game interact within a single territory (VG=0).

- Co-operation is high when BG and TG are high and BF is low. These values indicate that game interaction is highly biased towards agents sharing the same tags which are culturally learned.

These latter observations suggest more strongly that some form of tag biasing and learning is promoting high co-operation. However, it is still unclear from the results how this mechanism might operate. In the next section an explanation (or theory) is given.

8.2.5 Cultural Group Selection?

It is argued that a form of "cultural group selection" is operating in the runs identified in clusters 1, 2, 4 and 5 (table 8.3). Such a process might operate in the following way:

- Agents bias game interaction strongly towards those sharing the same label bits (tags).
- Game interaction therefore takes place within "game interaction groups" based on tag similarity.
- Groups which satisfy agents (i.e. produce co-operative outcomes in game interactions) cause agents within them to increase confidence in their memes. The memes include the tag bits which form the interaction group and the co-operative behavioural rules.
- Groups which do not satisfy agents (by definition) contain dissatisfied agents who decrease confidence in their memes (rules and tag bits).
- Cultural interaction between agents will tend to spread those memes with high confidence.
- Consequently agents within low co-operation "game interaction groups" will tend to be "recruited" into high co-operation "game interaction groups" due to take-up of those memes (tag bits) which identify those groups.
- An agent can therefore be seen as moving between interaction groups until sustained satisfaction is attained.

- Low co-operation groups will tend to lose agents to high co-operation groups.

The substantive claim in the above is that *tags combined with biasing create interaction boundaries around groups sharing the same tags.* If those tags can be changed by cultural learning from other agents then agents can "move" between such groups. Essentially then, tags combined with biasing create an "abstract space" in which co-operative "ghettos" can form. The process is therefore analogous to co-operation produced from the limitation of game interaction to a single spatial territory. In both cases *a method of distinguishing in-group members from out-group members allows for co-operative niches to emerge by the biasing of game interactions towards in-group members.*

Both the spatial and tag based limitation of game interaction may be viewed as mechanisms facilitating "parochialism", allowing agents to develop co-operation within exclusive in-groups. Smaller game interaction groups reduce the mutual space of strategies the agents need to search to find maximum co-operation and hence "lock-in" of that co-operation.

There are two major differences between spatial and tag based "parochialism". Spatial localisation cannot operate unless the spatial movement of agents is very low. In the StereoLab, movement is not related to satisfaction level but is random. Tag localisation allows agents to "move" in *tag space* by culturally interacting and receiving memes (tags) from other agents or by mutation. The topology of the tag space however is not one dimensional but rather a B-dimensional hypercube (where B is the number of binary tag bits). This latter point is discussed in chapter 9.

The outlined process above appears plausible but is difficult to separate from the other dimensions of the StereoLab society. Consequently the theory, as presented in this chapter, remains speculative. In order to capture the essence of the process, and express the

theory in a more concrete form, a simplified artificial society was constructed (TagWorld II). The next chapter details this simplified society and demonstrates empirically the power of the above process to select for group-functional behaviours in an evolutionary process, even when the agents are boundedly rational *optimisers* rather than satisficers. The optimising assumption simplifies the model since the underlying process of cultural transmission is not modelled. Also it can be determined if the group selection process identified in this chapter is dependent on the satisficing assumption.

8.3 Conclusion

In this chapter the StereoLab society outlined in chapter 6 was explored for regions of high co-operation within the parameter space. Two techniques were used: random sampling combined with decision tree induction (using C4.5) and hill-climbing combined with cluster analysis. C4.5 and random sampling is particularly suited to characterising the entire space when points of interest are not rare in the space. Hill-climbing and cluster analysis is more applicable when points of interest are rare but can be found by hill-climbing. Both techniques were usefull in analysing the data sampled from the parameter space. Finally, the parameter space was extended to allow for a higher level of interaction biasing based on label bits. Hill-climbing and cluster analysis suggested that a *group formation process based on tags* was producing co-operation (see section 8.2.3). A theory was advanced to elaborate the possible operation of this process (see section 8.2.5). But the complexity of the model makes analysis and testing of the theory difficult. In order to separate this process from the other dimensions of the StrereoLab and capture it in a tangible form, a simplified society (TagWorldII) was constructed to examine tag based limitation of game interaction. This is described in the next chapter.