Chapter 9

Groups in Tag Space

Experimentation with the StereoLab model in the previous chapter identified several processes producing high co-operation. One such process was characterised by high refusal of game interaction unless tags were very similar. The explanation for why this biasing produced high co-operation was given in the "cultural group selection" theory (see section 8.2.5 of chapter 8).

In this chapter a simplified society (TagWorldII) is constructed which aims to isolate this co-operation forming process. In the TagWorldII, game interactions are modelled as pairwise single-shot Prisoner's Dilemma (PD) games (see section 2.1.1 in chapter 2 and section 6.3 in chapter 6). The cultural spread of strategies and tags is modelled as an asexual population wide evolutionary process where the fitness of each agent is equal to the payoff received from playing rounds of PD. This means that agents are modelled as *optimisers* rather than *satisficers* and *the underlying process of cultural interaction is not modelled.* This method of modelling cultural evolution is a more traditional form used within evolutionary economics [14]. Essentially the method follows the assumptions of the

```
                  Evolutionary Algorithm:

LOOP some number of generations
  LOOP for each agent (a) in the population
    Select a game partner agent (b)
    Agent (a) and (b) invoke their strategies receiving
      the appropriate payoff
  END LOOP
  Reproduce agents in proportion to their average payoff
END LOOP
```

Figure 9.1: The evolutionary algorithm applied to agents within TagWorldII.

replicator dynamics[1] with the addition of mutation, finite population size and non-random mixing. This method was chosen to aid analysis by increasing the simplicity of the model and to test the robustness of the "cultural group selection" theory under the optimising assumption.

As will be shown, the process is robust over a wide range of the parameter space. Contrary to the conclusions drawn from a previous tag model [143] we have found that co-operation arises in the one-shot game.

9.1 The TagWorld II Artificial Society

Agents are represented as fixed length bit strings (of length L+1) comprising a tag of length L bits and a single strategy bit. The strategy bit represents a pure strategy, either unconditional co-operation or unconditional defection. Initially the population of agents are set to random bit strings (with each bit decided by a fair coin toss). Figure 9.1 shows the evolutionary algorithm then applied to agents.

In each generation each agent (a) is selected from the population (of size N) in turn.

[1] See the Glossary section of the thesis for a detailed account.

A game partner is then selected. Partner selection entails the random selection of another agent (b) from the population such that (a) $\neq$ (b) followed by a comparison of tags with agent (a). If the tags are identical a game interaction takes place otherwise (b) is returned to the population without game interaction. If (b) was returned to the population without interaction a second (b) is selected at random from the population and its tag compared with (a). This process is repeated until an agent (b) is found which has an identical tag to (a) or an upper limit F of selections has been reached. If this upper limit is reached then game interaction is forced on the next randomly chosen agent. Consequently (a) will always find a partner even if its tag does not match any other agent because an agent which can not find a matching partner will eventually exhaust its upper limit F of refusals and then be forced to interact with a randomly chosen partner.

During game interaction (a) and (b) invoke their strategies and receive the appropriate payoff. After all agents have been selected in turn and played a game a new population is asexually reproduced. Reproductive success is proportional to average payoff. The entire population of N agents is replaced using a "roulette wheel" selection method [34]. Equation 9.1 and inequality 9.2 outline this method. Equation 9.1 gives the total average payoff for the entire population where a_i is the ith agent from the population, $ap(a)$ is the average payoff obtained by agent a, and N is the size of the population. The inequality 9.2 specifies an agent a_x to select for reproduction from the population where x is the smallest integer that satisfies the inequality and $rnd(0..tap)$ is a uniformally randomly selected value in the range $0..tap$. The inequality is satisfied N times with a different random value. Each time, x gives the index of an agent to reproduce. Using this method the probability that an individual will be reproduced into the next generation is proportional to average payoff.

$$tap = \sum_{i=1}^{N} ap(a_i) \tag{9.1}$$

$$rnd(0..tap) \leq \sum_{i=1}^{x} ap(a_i) \tag{9.2}$$

Mutation is applied to each bit of each reproduced player with probability M = 0.001[2]. Mutation takes the form of flipping a bit with probability M. Consequently *tags and strategies* are mutated in reproduced agents.

The PD payoffs are parameterised over T (the temptation payoff for defectors over co-operators) such that T > 1. The reward R for mutual co-operation is 1. The punishment P for mutual defection and the sucker payoff S for co-operation with a defector are both 0.0001[3].

9.2 Results - High Co-operation

A set of runs to 100,000 generations with a population of size N = 100 agents was executed for various values of T and L. The maximum number of refusals was fixed at F = 1000. This high value of F means that it is unlikely that agents will not be paired with other matching agents (if they exist) in the population. For the purposes of analysis co-operation was characterised as the proportion of mutually co-operative interactions occurring over all generations[4].

Figure 9.2 shows results for various values of L and T graphically and figure 9.3

[2]This low value indicates the assumption that agents rarely change their strategy. Since there are 100 agents in the population, we would expect one strategy bit to change over 10 generations (on average).

[3]A value was selected that was small but greater than zero (indicating a very small chance for agents, with Sucker or Punishment payoffs, of reproduction). If a small value is added to P (enforcing T > R > P > S) results are not significantly changed (see section 9.2).

[4]The number of games in which both agents co-operated was counted for the whole run (of 100,000 generations) and then divided by the total number of games played. Thus the level of co-operation for a single run is derived from the results of 10^7 individual games.

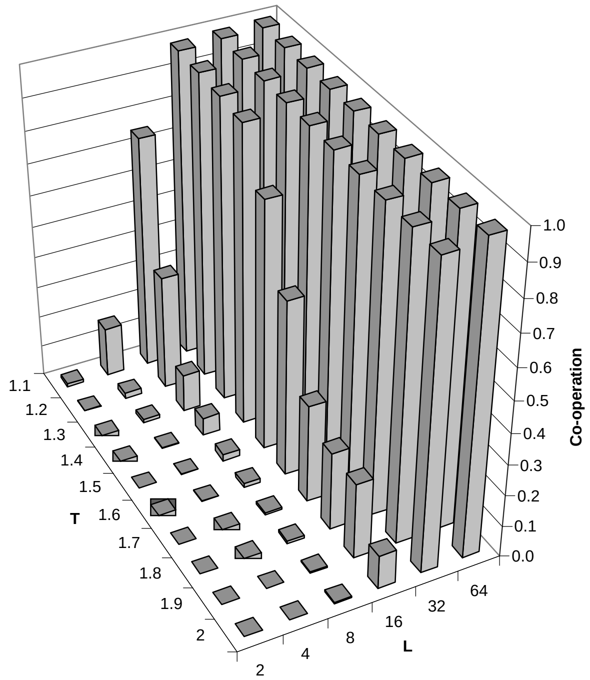

Figure 9.2: Co-operation for various values of L (number of tag bits) and T (the temptation payoff). Co-operation is measured as the proportion of mutually co-operative games over 100,000 generations. Each point is an average of 5 runs. The entire chart represents 3 x 10^9 individual game interactions.

shows the values used to generate the graph. Each bar represents an average of 5 independent runs. As the values of L are increased, and T are decreased, co-operation increases monotonically. As can be seen, where L >= 32 very high levels of co-operation are obtained for all values of T. Figures 9.4 and 9.5 show the results obtained when a small value is added to P (enforcing T > R > P > S). As can be seen the overall result is not significantly changed. Interestingly however, co-operation is increased for most values of T where L = 16 and L = 8[5].

The results obtained indicate that very high levels of co-operation can be sustained between optimising agents in the one-shot PD via simple tag biasing. There is no require-

[5]Why this result is obtained is currently not known. This would be an interesting topic for future investigation.

	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2
2	0.0097	0.0014	0.0011	0.0007	0.0001	0.0002	0.0000	0.0001	0.0001	0.0001
4	0.1591	0.0181	0.0112	0.0020	0.0009	0.0015	0.0006	0.0008	0.0005	0.0006
8	0.7479	0.3694	0.1209	0.0583	0.0220	0.0132	0.0080	0.0085	0.0038	0.0058
16	0.9819	0.9758	0.9690	0.9556	0.7977	0.5659	0.3141	0.2497	0.2429	0.1065
32	0.9882	0.9857	0.9826	0.9796	0.9767	0.9736	0.9705	0.9676	0.9647	0.9618
64	0.9917	0.9903	0.9890	0.9874	0.9858	0.9842	0.9824	0.9811	0.9794	0.9778

Figure 9.3: Co-operation for various values of L (number of tag bits) shown on the vertical and T (the temptation payoff) shown on the horizontal. Co-operation is calculated as the proportion of co-operative interactions over 100,000 generations. Each point is an average of 5 runs.

ment for knowledge of past performance or recognition of individual agents (i.e. the other agents may be viewed as strangers). Unlike some spatial models which have demonstrated co-operation in the one-shot game [127] reproduction is population wide. The next section explains this high co-operation in terms of group formation.

9.3 Group Formation and Dissolution

The tag space can be visualised as an L-dimensional hyper-cube with corners representing unique tag values. Agents sharing a tag, share a corner. Mutation produces movement between corners. Figure 9.6 illustrates a 4-dimensional hyper-cube for L=4. Each corner shows the unique tag value associated with it. Game interaction is therefore taking place in an abstract "tag space". Co-operative groups sharing matching tags will form in corners of the hyper-cube. These groups will out perform non-co-operative groups and hence tend to increase in size over generations. However, if mutation introduces defecting agents into a co-operative group they will tend to out perform the co-operators within the group (by suckering them). From this the seeds of the destruction of the group are planted, since as the number of defectors increases within a group the overall fitness of agents within the group decreases. Other more co-operative groups (if they exist) will tend to expand. While this process is occurring, mutation of tag bits will produce a slow

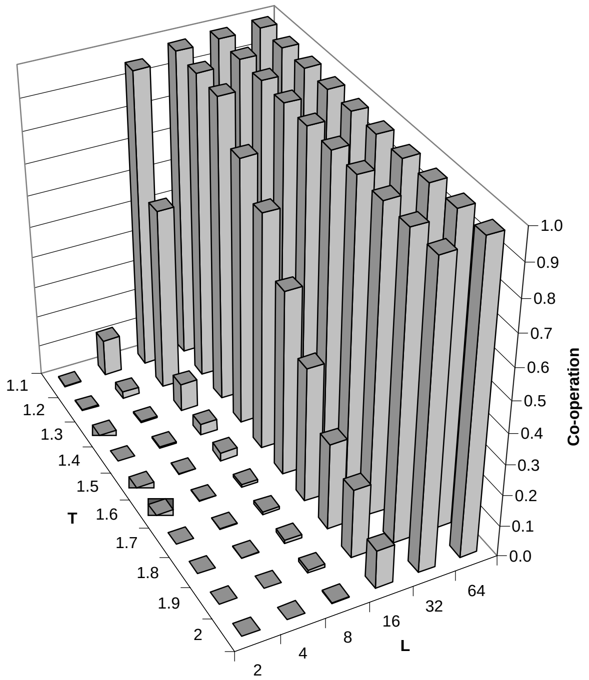

Figure 9.4: Co-operation for various values of L (number of tag bits) and T (the temptation payoff). Co-operation is the proportion of mutually co-operative interactions over 100,000 generations. Each point is an average of 5 runs. Here a small value is added to P (the punishment payoff) enforcing the inequality T>R>P>S. The entire chart represents 3 x 10^9 individual game interactions.

	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2
2	0.0026	0.0021	0.0011	0.0003	0.0007	0.0001	0.0000	0.0000	0.0000	0.0001
4	0.1188	0.0222	0.0047	0.0031	0.0016	0.0017	0.0009	0.0014	0.0002	0.0002
8	0.9539	0.5888	0.0888	0.0362	0.0266	0.0087	0.0103	0.0114	0.0093	0.0038
16	0.9819	0.9757	0.9692	0.8522	0.7582	0.5950	0.4349	0.2802	0.2245	0.1232
32	0.9885	0.9855	0.9827	0.9797	0.9764	0.9731	0.9705	0.9672	0.9648	0.9619
64	0.9917	0.9904	0.9889	0.9874	0.9860	0.9843	0.9826	0.9811	0.9796	0.9778

Figure 9.5: Co-operation for various values of L (number of tag bits) and T (the temptation payoff). Co-operation is the proportion of mutually co-operative interactions over 100,000 generations. Each point is an average of 5 runs. Here a small value is added to P enforcing the inequality T>R>P>S.

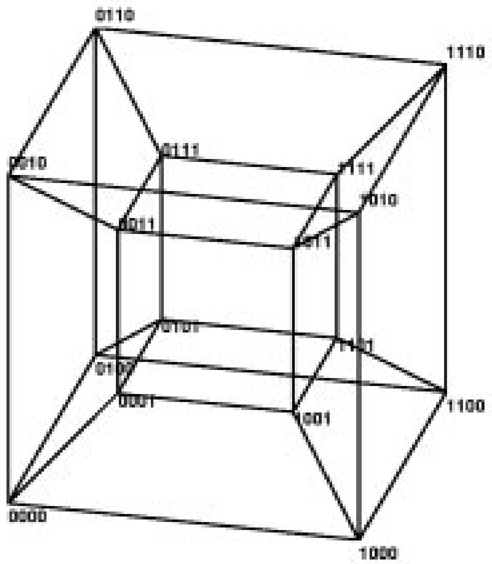

Figure 9.6: For a tag length L = 4, label space can be visualised as a 4D hypercube. The corners present the various possible tag values. Mutation of a single label bit produces movement to a connected corner.

migration of agents between corners of the hyper-cube, possibly founding new groups in previously empty corners.

Figure 9.7 is a visualisation of the process over time taken from a single run. Each line on the vertical axis represents a corner of the hyper-cube (i.e. unique tag value). The horizontal axis represents time in generations. If no agents have a particular tag value in a given generation then the line is left blank (white). Alternatively, if a corner contains all co-operative agents then the line is red. For a mixed group in which there are both co-operators and defectors the line is green. For an entirely defective group the line is coloured blue. Examination of figure 9.7 shows the time evolution of groups in tag space. Initially co-operative groups (red lines) become invaded by defectors producing mixed groups (green) which very swiftly become entirely defective (blue) and then quickly go extinct (white).

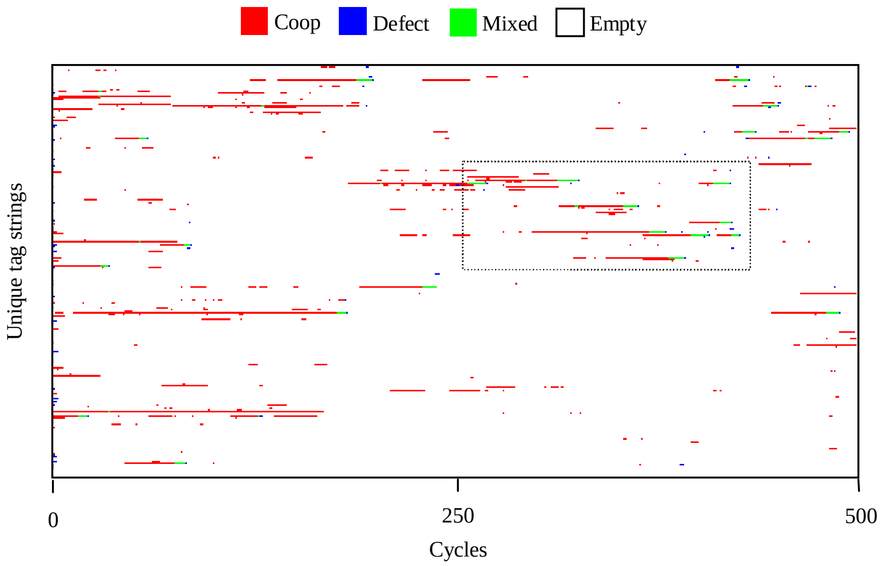

Figure 9.7: Visualisation of 500 cycles from a single simulation run showing co-operative groups coming into and going out of existence. Each line on the vertical axis represents a unique tag value. If all agents sharing a tag value are co-operative then the line is coloured red. If all agents are defectors then the line is coloured blue. A mixed group is shown as green. The horizontal axis represents time in cycles. Here L = 8, N = 100, M = 0.001, T = 1.1 and F = 1000. The highlighted rectangular area is shown expanded in figure 9.8. Although the value L=8 does not produce the highest co-operation the small number of possible tag combinations (2^L=256) allows for easy visualisation of the entire tag space over time.

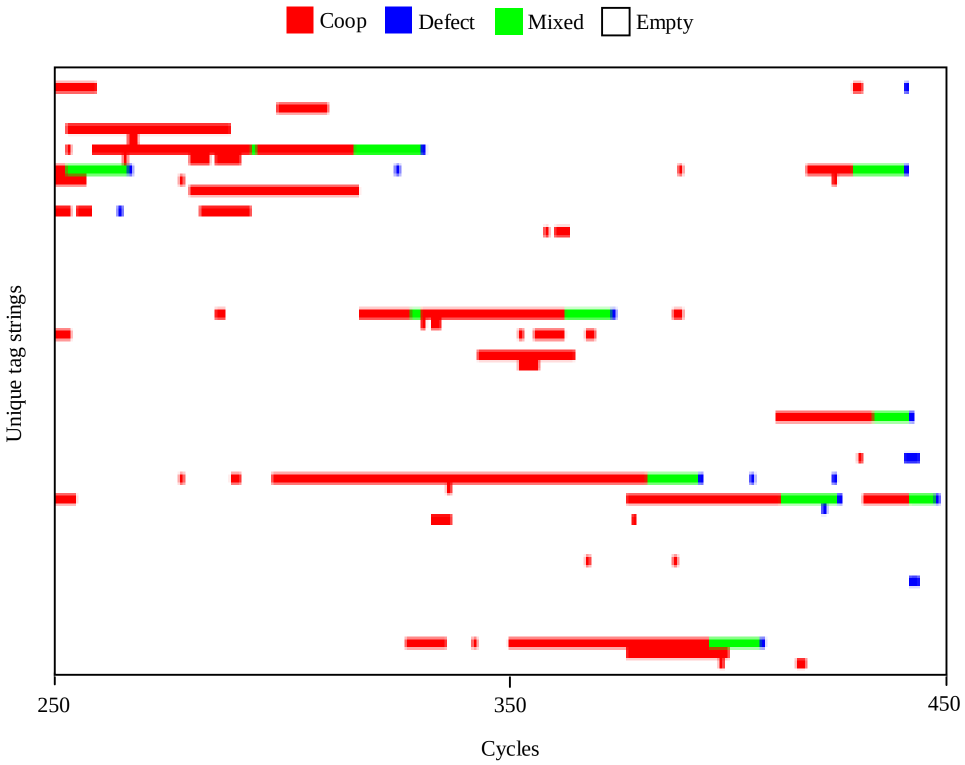

Figure 9.8: An expanded view of an individual simulation run. The highlighted part of figure 9.7 is here expanded to show detail. Notice that co-operative groups form (red lines) then become invaded with defectors (green lines) and ultimately become completely composed of defectors (blue lines), then go out of existence.

9.4 Co-operation From Complete Defection

Although high levels of co-operation are demonstrated over many contiguous generations, starting from the random initialisation of agents, these results do not indicate if a society can recover co-operation from a state of complete defection. In order to test this, experiments were conducted in which the initialisation of agents was modified so that all agent strategy bits were set to defection. Figure 9.9 shows a set of runs for various values of N (population size) against the number of generations before mutual co-operation emerges. These empirical results are compared to a simplified analytical (probability based) treatment given in equations 9.3 to 9.8.

Equation 9.8 gives the expected average number of generations required (from an initial society of all defectors) before two co-operative agents perform a game interaction. This is dependent on the population size N and the mutation rate M. In all of the 50 runs used to form the empirical results in figure 9.9, high co-operation immediately (in the next generation) followed the first occurrence of a mutually co-operative encounter. It was also observed that drift over the tag bits tended to lead a society of all defectors toward sharing the same tag bits. These empirically observed (rather than derived) phenomena were used to simplify the analytical treatment.

Equation 9.3 gives the probability that no strategy bits are mutated in the whole population over a generation (where n is the population size and m is the mutation rate).

$$pnc(n,m) = (1-m)^n \tag{9.3}$$

Equation 9.4 gives the probability that one and only one strategy bit in the population is mutated.

$$poc(n,m) = nm(1-m)^{n-1} \tag{9.4}$$

So, given a population of all defectors, equation 9.5 gives the probability that more than one agent will be mutated into a co-operator over a single generation.

$$pmo(n,m) = 1 - pnc(n,m) - poc(n,m) \tag{9.5}$$

Assuming that two (and only two) newly mutated co-operative agents share the same tags[6], equation 9.6 gives the probability that those two agents will interact (during the game playing phase) producing a mutually co-operative outcome.

$$pti(n) = \frac{2}{n-1} \tag{9.6}$$

Equation 9.7 combines the probabilities that more than one agent will become co-operative with that of two agents interacting. This gives the probability that a mutually co-operative game will be played from a starting condition of all agents being defectors over one generation.

$$pmoti(n,m) = pti(n) * pmo(n,m) \tag{9.7}$$

Finally, equation 9.8 gives the average number of generations that can be expected before a mutually co-operative game interaction occurs (incorporating all the previous equations).

$$\begin{aligned} ang(n,m) &= \frac{1}{pmoti(n,m)} \\ &= \frac{1}{\frac{2}{n-1}(1-(1-m)^n - nm(1-m)^{n-1})} \end{aligned} \tag{9.8}$$

In figure 9.9, the predictions of equation 9.8 are compared empirically with actual runs. As can be seen the analytical treatment tends to under-estimate the number of generations before co-operation emerges for low population sizes and over-estimate for large population sizes. It can be hypothesised that this is due to the simplification in equation 9.6. Drift will

[6]This assumption is a simplification but appears to produce a good fit with the empirical data (see figure 9.9). The assumption is made since estimating the number of distinct tag groups as they evolve over time appears very complex.

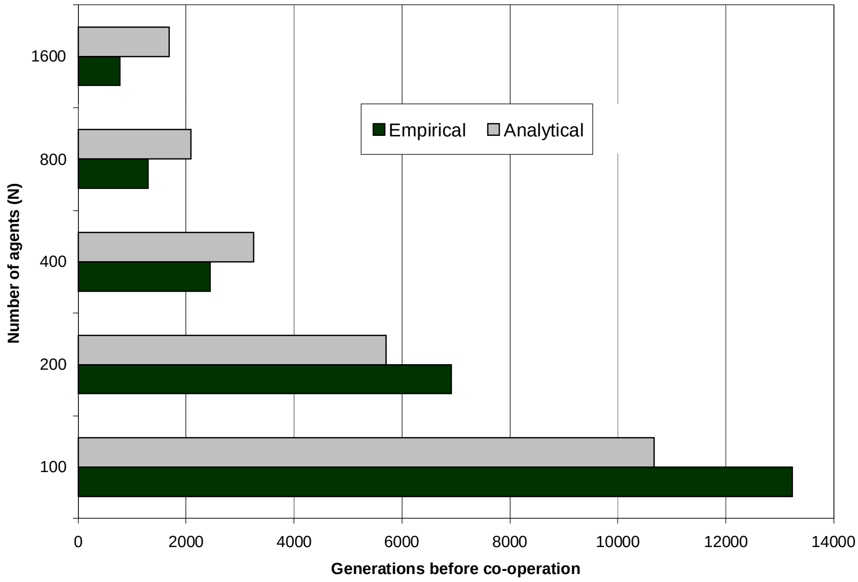

Figure 9.9: Number of generations before mutual co-operation emerges. Here a comparison is made between the analytical model (A) (given in equation 9.8) and the empirical results (E). Each bar for the emprical test comprises an average of 10 runs. Mutation is set to M = 0.001, tag size L = 32 and refusals F = 10N for all runs shown.

rarely produce populations in which all agents have identical tags. When the population is split between different tag values the simplification will under-estimate the number of generations required. This is because two co-operators may be produced but they may not share the same tag and hence will not game interact. Conversely when the population size is large it will often be the case that more than two agents will be mutated to co-operators. A more detailed analysis would need to capture the dynamics of the tag bits over the population and the effect of those dynamics on producing co-operative interactions.

9.5 Discussion

In a previous study of tag based partner selection in the PD, Riolo [143] concluded that tags produced little increase in co-operation in the one-shot game. In his work a tag was represented as a single real number attached to each agent. The abstract topology of the tag space was therefore one dimensional. The matching of tags was based on a probabilistic function applied to the distance between two tags, meaning that agents with similar but not identical tags could engage in game interactions. The number of refusals allowed before forced interaction was low (50 refusals in a population of 200 agents). Additionally a fitness cost was attached to each refusal made by an agent (although this was reduced to zero in an attempt to get high co-operation in the one-shot game). Under these conditions it was demonstrated that high co-operation emerged when agents engaged in the Iterated PD (IPD) but not in the single-shot PD game. The interpretation placed on this previous study was of agents representing animals, searching for game partners and evolving genetically[7]. In the work presented here allowable refusals is high (10 times the population size) and there is no associated cost. Also the tag is represented as a bit string which must match *exactly* with a partner for game interaction to be selected by an agent. Under these assumptions a different kind of tag space topology is possible and high co-operation is produced in the single-shot game.

Intuitively it would seem that the *exact tag matching* constraint is not necessary to produce high co-operation in all cases. For example, consider a situation in which mutation was zero and two groups of agents existed in the tag space such that there was no inter-group game interaction. If one group consists of all co-operators and the other contains some defectors then the co-operative group would expand at the expense of the non-co-operative group. This would even hold if agents were applying some partial matching scheme - *so long as there is an interaction boundary between the two groups.* By "interaction boundary" is meant that some mechanism partitions the agents into strict game interaction groups so that games can only take place between individuals sharing a group.

[7] Riolo R. (personal communication).

9.5.1 Interpretations

The interpretation placed on this work is of agents representing a "bundle" or "complex" of culturally learned and transmissible behaviours, a so-called "meme-complex" [40], [15]. The assumption is that in a population of hosts meme-complexes which produced high utility (for their hosts) would be more likely to be copied (in proportion to their relative utilities). One key to such a process producing high altruism and co-operation is the packaging of the tag with the beneficent strategy or behaviour as a cultural unit. Allison in his theory of altruism [2] echoes this assumption with reference to the importance of "cultural packaging techniques". In the work presented here "packages" of tags and defective strategies do not dominate the population because such packages destroy the very groups that they are a part of.

Recent work by Bowles & Gintis [18] gives a detailed analytical treatment of the value of groups in the promotion of co-operation in the PD when binary social cues convey useful predictive information concerning a game partner's strategy. However, they do not address the issue of *how* social cues come to be have such predictive utility. The results produced within the TagWorldII society show that even simple mechanisms can produce this kind of correlation because groups which contain defectors quickly die out. In contrast to the examples used by Bowles & Gintis (who focus on racial groups) the mechanism which produces this quick extinction of non-cooperative groups within the TagWorldII society requires that cultural interaction, in the form of individuals moving between groups easily, is high. It is this ability of individuals to quickly swap cultural groups, by taking on new tags from others, which drives the co-operation producing process. Strong group boundaries which prevent easy entry and exit from a group would hinder or even destroy co-operation forming by the process illustrated in TagWorldII.

9.5.2 Future Directions

The TagWorldII society was parameterised over a number of dimensions. A scan was made (see section 7.2 in chapter 7) over a restricted part of the parameter space demonstrating (through an existance proof, see figure 3.1 in chapter 3) that high co-operation is present over a large area of that space. However, other dimensions of the space have not been explored. The role of mutation and refusals would be an interesting area of investigation. Also, measures other than just co-operation, such as group sizes over time and migration rates between groups, would be of interest and could be used to elaborate an analytical model. This could link this work with patch based models of altruism developed within evolutionary biology [163]. Another interesting area of further investigation would be to make refusals an endogenous parameter encoded into each agent and able to evolve. In such a scenario would high refusal rates evolve? If so then the assumption of high refusals could be justified via endogenous evolutionary processes rather than as an exogenous assumption of the model.