# The Emergence of Cooperation among Agents using Simple Fixed Bias Tagging

**Enda Howley**
Department Of Information Technology
National University Of Ireland, Galway
Enda.Howley@nuigalway.ie

**Colm O'Riordan**
Department Of Information Technology
National University Of Ireland, Galway
colmor@it.nuigalway.ie

**Abstract-**

**The principle of cooperation influences our everyday lives. This conflict between individual and collective rationality can be modelled through the use of social dilemmas such as the prisoner's dilemma. Reflecting the reality that real world autonomous agents are not chosen at random to interact, we acknowledge the role some structuring mechanisms can play in increasing cooperation. This paper examines one simple structuring technique which has been shown to increase cooperation among agents. Tagging mechanisms structure a population into subgroups and as a result reflect many aspects which are relevant to the domains of kin selection and trust. We will outline some simulations involving a simple tagging system and outline the main factors which are vital to increasing cooperation.**

## 1 Introduction

The people we meet on a daily basis are not randomly chosen from a universal set of individuals throughout the world; instead these interactions are based significantly on our proximity to one another. This is not just limited to physical proximity but also genetic, behavioural, and social proximity. These groups represent the set of individuals whom we are most likely to continuously meet throughout our lives. In reality, an individual may be part of many hundreds of sets or groups during their life, but these group structures will determine their interactions on a daily basis throughout their whole lives. During a lifetime, an individual's membership of certain groups will continuously change and evolve, reflecting their choice to move home, school, job or get married. The main point remains that our interactions are non-random and heavily biased by certain group structures.

Significant research has been conducted involving agents whose interactions are determined by spatial proximity. Nowak and May (1993) [1] describe the significance of group structuring techniques with special attention to spatially determined interactions. This structuring can be represented through the use of tagging which is an abstract method of biasing agent interactions based on membership of certain tag groups. Holland (1993) [2] describes tags as markings or social cues that are attached to individuals (agents) and are observable by others. When determining which agents should interact we calculate the proximity of their tag values, as opposed to their physical proximity. Tagging is an abstraction which allows us reflect the relatedness between agents based on any possible grouping analogy, not simply spatial or genetic proximity. Tagging can be used as a general case to represent all these possible grouping structures without the specific complexities which they entail.

Hales (2004)[3] states that tags can evolve from initially random values into complex ever-changing patterns that serve to structure interactions between individuals. Tagging schemes are highly accurate mechanisms for biasing agent interactions based on their relatedness with each other.

The approach described in this paper, although similar to previous research by Hales and others, differs in that we adopt a simpler approach where agent strategies and tags are not genetically linked. We do not subject the strategies to mutation but focus instead on the effects of varying the numbers of tags and the effects of tag mutation.

This paper will analyse a simple tagging scheme and review some of the factors which contribute to its success. This paper will describe a number of experiments which we have designed involving a simple tagging scheme. These experiments have been specifically designed to display some of the important features of tagging. For example we will review the levels of cooperation achieved using different amounts of tags. We hope to examine the effects of population viscosity[1] on levels of cooperation among a tag using population of agents. These simulations have been conducted on populations of agents competing through the iterated Prisoner's Dilemma.

### 1.1 The Prisoner's Dilemma

The Prisoner's Dilemma (PD) is a simple two player game where each player must make a decision to either cooperate (C) or defect (D). Both players decide simultaneously and therefore have no prior knowledge of what the other has decided. If both players cooperate they receive a specific payoff. If both defect they receive a lower payoff. If one cooperates and the other defects then the defector receives the maximum payoff and the cooperator receives the minimum. The payoff matrix outlined in Table 1 demonstrates the potential payoffs for each player.

Table 1: Payoff Matrix

| Players Choice | Cooperate | Defect |
|---|---|---|
| Cooperate | $(\lambda 1, \lambda 1)$ | $(\lambda 2, \lambda 3)$ |
| Defect | $(\lambda 3, \lambda 2)$ | $(\lambda 4, \lambda 4)$ |

[1]The viscosity of a population relates to the degree of change in that population over time. If individuals do not move far from their place of birth, this may be considered a form of population viscosity.

The dilemma is a non-zero-sum, non-cooperative and simultaneous game. For the dilemma to hold in all cases, certain constraints must be adhered to: $\lambda 2 < \lambda 4 < \lambda 1 < \lambda 3$. These conditions result in $\lambda 2$ being the sucker's payoff, $\lambda 1$ is the reward for mutual cooperation, $\lambda 4$ is the punishment for mutual defection, and $\lambda 3$ provides the incentive or temptation to defect. The dilemma also states $2\lambda 1 > \lambda 2 + \lambda 3$. This constraint prevents players taking alternating turns receiving the sucker's payoff ($\lambda 2$) and the temptation to defect ($\lambda 3$), therefore maximising their score. The following values were used throughout this research: $\lambda 1 = 3, \lambda 2 = 0, \lambda 3 = 5, \lambda 4 = 1$.

In the non-iterated game, the rational choice is to defect, while in the finitely repeated game, it is rational to defect on the last move and by induction to defect all the time. However, if there exists a non-zero probability the two players will play again then cooperation may emerge. Within a society of social groups, repeated meetings are common and as with tag group members who meet repeatedly, significant levels of cooperation can emerge.

### 1.2 Previous Tagging Models

Holland (1993) [2] initially outlined the concept of tags and since then significant numbers of tag models have emerged such as Hales (2004) [3] and Riolo (1997) [4]. Holland (1993) [2] describes tags as markings or social cues that are attached to individuals (agents) and are observable by others. Riolo has described a number of tagging approaches throughout a series of papers, Riolo (1997) [4], Riolo (2000) [5], each focusing on the effects of tagging on levels of cooperation among players in the Iterated Prisoner's Dilemma. These papers outline basic forms of tagging: fixed-bias tagging, variable-bias tagging and evolved-bias tagging.

These models comprise two-player Iterated Prisoner Dilemma (IPD) tournaments. Players are not paired randomly as in traditional tournaments such as those researched by Axelrod (1984) [6]. Instead, each player is given a numeric tag value and this is used to probabilistically bias interactions towards players of similar tag values. In these tagging models significant increases in cooperation can be observed. Simple replicator dynamics are used to reflect the fitness of agents as they move through successive generations. The proportional fitness of a genome is used to determine representation of that genome in subsequent generations.

We hope to gain a greater understanding of how tagging mechanisms successfully increase cooperation among agents playing the IPD. We address the following questions:

- What conditions are required for a tagging scheme to successfully increase cooperation?
- Which individual factors or parameters help determine the scale of cooperation increases?
- How are tagging mechanisms affected by less viscous environments?

In later sections, we outline a number of experiments and the results obtained. In the following section we will discuss the design of our simulator and tagging model.

## 2 Experimental Setup

### 2.1 Strategies

The design of any IPD simulator requires the simple creation of an initial set of player strategies. The research by Nowak (1990) [7] provides a basis for our method of strategy definition. The following is a strategy genome with 3 genes representing 3 possible behaviour values when $P_i$ *( Probability of cooperation in the initial move of a game),* $P_c$ *( Probability of cooperation after opponent has cooperated),* $P_d$ *( Probability of cooperation after opponent has defected ).*

$$Genome = P_i, P_c, P_d, \tag{1}$$

Our initial population of agent genomes is randomly generated with random gene values and therefore the initial population has an initial average fitness[2] of about 2.25[4]. A fourth gene representing a tag value is also given to each genome, this tag value is randomly assigned to all strategies in the population and is a simple integer in the range [1...Maximum number of tags].

### 2.2 Tag Model

Our simulator consists of a population of 100 players; each is randomly generated from a normally distributed set of possible strategies. These strategies are each given random tag values. Games lasting 20 iterations between all members of the same tag value are conducted. Under no circumstances do agents of different tag values play each other; but agents can play themselves. If an agent is the solitary holder of a tag value it may still play itself, even though there are no other members of the population it can interact with.

The representation of a genome in successive generations is based on its fitness in the current generation. Therefore this dictates representation in the next generation and so on, through successive generations. This simple replicator dynamic is similar to that used by Riolo in his simulations. Page (2003) [11] discusses this specific topic in more detail with special emphasis on frequency dependant selection. The essence of this discussion states that the fitness of an individual is determined by the frequency of other agents in the population. This is represented by the following equation:

$$\dot{a_i} = \sum_{j=1}^{n} x_j f_j(x) q_j i - x_i \overline{f} \tag{2}$$

This is a "replicator-mutator equation" which describes both frequency-dependant selection and mutation. This equation has previously been applied to population genetics by Hadeler (1981) [12] and also game theory by Bomze and Buerger (1995) [13]. In our research and the research of

[2] If we use the following IPD payoff matrix values $\lambda 1 = 3, \lambda 2 = 0, \lambda 3 = 5, \lambda 4 = 1$, a initial population of random agents will have an average of these fitness values $(3 + 0 + 5 + 1)/4 = 2.25$

Riolo, all reproduction is asexual. As a result, no crossover between agents occurs.

One significant feature of our tagging model introduces a mutation operator. This may be applied to random members of the population through successive generations. This is a probability of a genome's tag changing to some random tag value[3].

This simulator extends previous tag models to allow some parameters vary over certain experiments. For example, this simulator can conduct experiments across a number of possible tag parameters.

### 2.3 Tag Generations

To reflect some evolutionary competition within the tag groups there is an intermediate stage replicator dynamic among members of each tag group. This extra evolutionary stage occurs before proceeding onto the next generation. This applies a replicator dynamic to all strategies of the same tag value. Strategies within a tag group are measured based upon fitness and then awarded updated representation within the tag membership. This process may be repeated a number of times within each generation of the overall game. As a result there may exist any number of tag generations within each overall generation of the simulation. The ratio of tag generations to actual generations may prove to be significant. The overall effect of the tag generation replicator dynamics results in all the strategies within the tag group being proportionally represented in the group based on their fitness.

### 2.4 Mutation

A tag may change value from generation to generation. This concept of tag mutation is represented as a probability of an agents tag value changing between generations. Once this occurs, the agent receives a random tag value within the parameter of the valid tag value range. This feature is used in some of our experiments to reflect the behaviour of a less viscous population.

## 3 Results

In this section we outline the results of a number of experiments. Our primary goal is to address the questions which we outlined earlier in this paper.

### 3.1 Performance of Tagging Scheme

The first set of experimental results represents a fitness comparison between two experimental populations. One implements a simple tagging scheme while the other uses no tagging technique. Across 50 simulations, 100 agents are allowed to compete for survival while playing each other. The tagged model randomly distributes up to 50 different tag values to the 100 agents. Throughout each simulation, these players will play IPDs of 20 iterations with peers of the same tag value. 50 generations are simulated and these tag groups are represented through successive generations based on their fitness through the use of a replicator dynamic. Our non-tagged model operates in exactly the same manner but all players play IPD's of 20 iterations with every other member of the population.

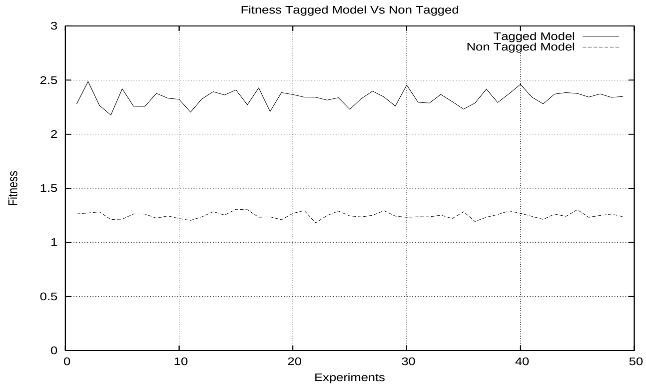


Figure 1: Tagged Model Vs Non-Tagged Model

In Figure 1, we can clearly see the improved fitness achieved through the use of tagging. This data is comparable with previous research by Riolo (1997) [4]. These results are consistent across all experiments while our simulation without tagging displays much lower levels of average fitness. The improved performance of our tagged model is as a result of tag groups insulating themselves against invasion from defecting strategies. In a non-tagged model, the effects of the defecting strategies propagates throughout the population and they dominate very quickly. This is not easily achieved in a tagged model which partitions the population into tag groups, which as a result prevent defection spreading throughout the population. The ability of a tagging scheme to partition populations into groups decreases the chances of groups having a non-cooperative strategy which will undermine the fitness of its peers.

### 3.2 The Significance of Tag Group Size

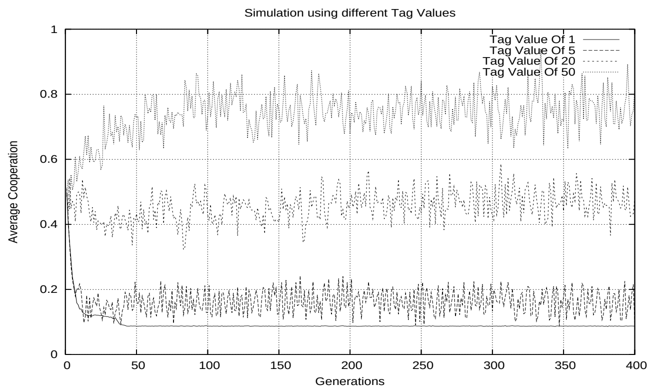


Figure 2: Different Tag Ranges

To fully understand the factors determining the success of a tag model this experiments investigates the effect of

[3] A tag can only mutate into a tag value within the parameters of the valid tag value range, as predefined in the experimental setup.

varying the number of tags used in the model. The experiment shown illustrates how the number of tags can influence levels of cooperation among the agent population. This experiment represents 400 generations using four distinct simulation parameters. These four simulations represent distinct amounts of tags permissible within the initial population. The number of tags determines how many tag groups exist and also their relative size within a population of 100 players. As in the earlier experiment we identify the ability of tagging to curtail defection propagating throughout a population. As the number of tags increases, the partitioning effect increasingly limits the propagation of defection. The model performs best when the number of tags is high. Therefore in any one tag group there exists a higher probability that all the members will be cooperative and as a result the group will be fitter.

The maximum number of tags in a population of 100 players would be 100 tags (If players are allowed to play themselves) or 50 tags (If players are not allowed play themselves). In our simulations we allow agents play themselves and as a result the maximum number of tags is 100. We discovered that our tagging scheme reached optimum at 50 tags and no further increase in performance was observed for tag amounts up to 100. This can be explained through a closer examination of our replicator dynamic. A single player competing against itself for representation as the solitary player in a tag group, will never increase its representation above the total tag group population of one. (As explained in relation to the replicator dynamic, average agent fitness in $G$ determines agent representation in $G+1$. This is trivial in the case of a single agent as its fitness equals the population's average fitness and therefore must result in 100 percent of representation the population in $G+1$) This significantly curtails the population's ability to increase cooperation any further. This can appear as a ceiling limiting any further increases in cooperation for such a tagging scheme.

### 3.3 Population Viscosity

A population may be considered viscous if over time it does not change significantly. If individuals do not move far from their place of birth, this may be considered as a form of spatial viscosity. Population viscosity was first discussed by Hamilton (1964) [8] and has subsequently been identified as a significant factor determining altruism among individuals by Hamilton (1971) [9] and also Wilson, Pollock and Dugatkin (1992) [10]. Nowak and May (1993) [1] also discussed this concept in relation to a spatially constrained IPD. While population viscosity has been found to increase the relatedness among individuals and therefore increase altruism, would this also be relevant to the tagging model we have discussed.

To reflect some form of population change over time we have introduced a simple form of mutation. Through a simple experimental parameter we can control the probability of an agent tag value changing in a given generation. In abstract terms this could reflect a change in an agents social, behavioural or spatial grouping. As a result this alters which peers an agent it is permitted to interact with.

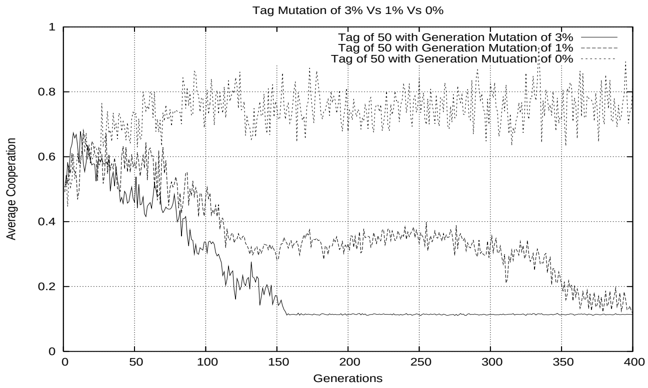


Figure 3: Levels Of Population Viscosity

This experiment depicts three levels of mutation and the resulting average cooperation across 400 generations. As we increase the probability of mutation, the level of population viscosity falls. The population has a greater chance of changing within each agents generational cycle. Reflecting this reduced viscosity the levels of average cooperation fall. This reinforces previous research that viscosity benefits relatedness and altruism between individuals. Simulations using a mutation probability of zero percent represent a totally viscous population. All three simulation used 50 tags and agents competed using a 20 iteration IPD.

### 3.4 The Evolution Of Tags

During our simulations we observed the importance of partitioning an initial population with many tag values and as a result increasing the average cooperation. Throughout this research, we have observed these important behaviours in relation to the size and number of tag groups over time. Throughout successive generations how would the evolutionary process effect tag group dynamics? This experiment examines the number of tags in use over 400 successive generations. The initial population is generated with a possible maximum of 50 random tag values. Throughout the simulation we plot the percentage of these possible 50 tags in use. We also plot the average cooperation throughout the simulation spanning the 400 generations.

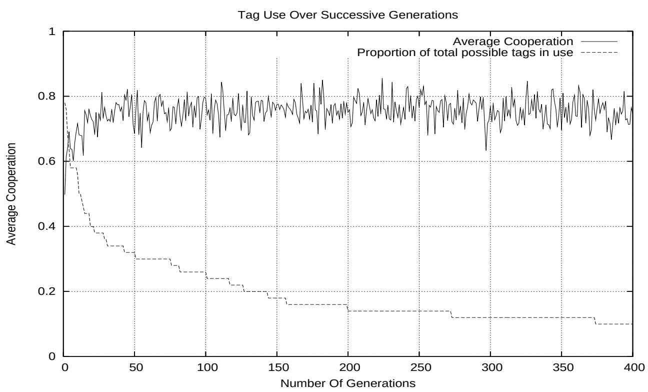


Figure 4: Numbers Of Tags In Use

From this experiment we see that the number of total tag values falls significantly after the initial generations. Significantly the levels of cooperation remain very high. This is explained by the ability of tags to marginalise non-cooperative behaviour over the initial generations. The tagging system results in non-cooperative agents becoming extinct in the earlier generations while the most cooperative tag groups increase rapidly in size. These tag groups with the most altruistic members takeover the total population and lead to the extinction of tag groups which contain any non-cooperative agents. The resulting small number of tag groups all contain highly cooperative agents and are about equal in size. Usually this experiment resulted in about 5 of a possible 50 tag groups surviving, each with an average membership of about 20.

### 3.5 The Effects of Tag Generations

Earlier we described the feature of tag generations which is one of the features of our simulator. Across numerous simulations and experiments we cannot identify any significant differences between experiments using and not using tag generations. Including those which examine the effects of reduced population viscosity. Only one experiment displays any difference in average fitness and is shown in Figure 5. Tag generations occur between actual generation timesteps where we measure fitness, and therefore between every generation timestamp represented in the following graph there are actually a number of tag generations which are not shown. A fitness reading is recorded and displayed at every actual generation. This occurs after the series of tag generations successfully completes.

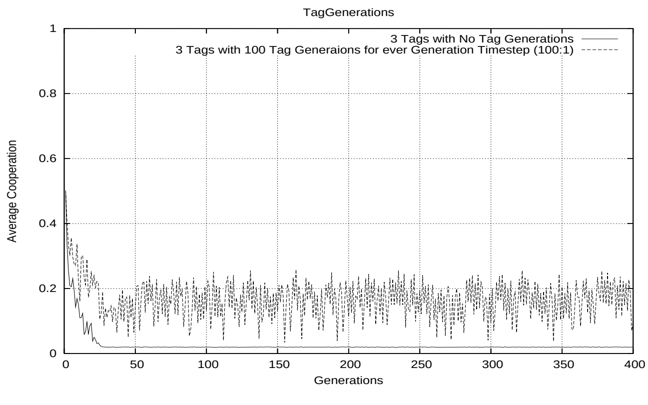


Figure 5: Average Fitness Using Tag Generations

The experiment in Figure 5 represents two simulations with different numbers of tag generations per actual generation. One does not contain any tag generations and moves from generation to generation without any intermediate stages. The other simulation has one hundred tag generations per generation which occur as intermediate stages between actual generations. A tag generation applies a replicator dynamic to the agents within the tag group and awards percentages of the population based on their fitness in the previous tag generation. A universal generation applies a similar replicator to the whole population therefore reflecting the fitness of agents and tag groups alike.

While no significant differences are observed between using tag generation and no tag generations, the above experiment shows one slight difference. When a population falls into total defection, simulations using tag generations maintained a slight level of cooperation which was absent from the other model. This is due to the number of agents receiving the exact same payoffs and fitness values. The tag generations have no other way of distributing available space in the population other then randomly choosing from the panel of equally fit strategies. When filling these last few population positions the algorithm will select random proportions of cooperative and non-cooperative strategies. This element of randomness in the specific case of total population defection has the effect of generating an artificial level of oscillating cooperation in a predominantly non-cooperative environment. This behaviour would not have occurred without our use of probabilistic based strategies.

## 4 Conclusions

Our experiments have highlighted a number of interesting features leading to the success or failure of certain tagging models. For example we have observed the important relationship between the number of tags used and the size of the population. This ratio has a direct effect on the overall levels of cooperation. The success of tagging is based upon preventing invasion from greedy agents. This partitioning effect which is synonymous with tagging and spatial models appears fundamental to their success. A tag group of only one agent is an extreme case but never encounters the "invasion" difficulties which usually jeopardise cooperation among larger groups of agents.

This depends heavily on a significant presumption that the population is completely viscous in nature. In our simulations examining population viscosity we confirm that less viscous populations promote defection. Slight decreases in viscosity leads to significant falls in average cooperation.

We have identified the emergence of a small set of highly cooperative tag groups which dominate populations after a number of generations. This small number of tag groups are composed of highly cooperative strategies which contribute to rapid growth because of very high group fitness. Tag groups which contain non-cooperative agents will struggle to survive and soon become extinct.

## 5 Summary

In answer to the three questions posed earlier in this paper we recall which conditions are required to successfully increase cooperation. These conditions would include the presence of tags and an almost totally viscous population. The main factor which determines the scale of cooperation increases must be the number of tags used initially in the population. Large numbers will die rapidly as non-cooperative tag groups become extinct but their presence in the initial population is essential for this reason. Finally we can clearly conclude that a less viscous population will suffer significant decreases in cooperation due to the invasion of non-cooperative agents into previously highly cooperative groups.

These tagging systems serve to bias agent interactions based on agents relatedness or proximity to each other. This reflects how real world interactions are biased towards individuals who share common attributes with us. Defining these attributes is extremely difficult due to the complex behaviours of humans. Our tag model implements a very simplistic model which is an abstraction of these social groupings. While this is a very simplistic model it still displays some very complex behaviours which are vital to our understanding of more elaborate tagging models.

Many possible avenues of future work involving this research are possible. One obvious extension would examine more elaborate tagging models. Equally there are numerous aspect of learning and evolution of tags which are possible. This learning could be possible through agents observing their peers behaviour over time. Similarly aspects of inter-agent communication may be explored. Such approaches could investigate the manner agents communicate tag information to each other and the resulting levels of cooperation. Finally, the techniques discussed in this paper can be applied to games other then the Prisoner's Dilemma.

## Acknowledgment

The primary author would like to acknowledge the Irish Research Council for Science, Engineering and Technology(IRCSET) for their assistance through the Embark initiative.

## Bibliography

[1] M. Nowak, R. May, *The Spatial Dilemmas Of Evolution, Int Journal of Bifurcation and Chaos 3, 35-78,* 1993.

[2] J. Holland, *The Effects of Labels (Tags) on Social Interactions, Working Paper, Santa Fe Institute* 93-10-064, 1993.

[3] D. Hales, *Understanding Tag Systems by Comparing Tag Models, Presented at the Second Model-to-Model Workshop (M2M2) co-located with the Second European Social Simulation Association Conference (ESSA'04) at Valladolid, Spain 16-19th of Sept* 2004.

[4] R. Riolo, *The Effects of Tag-Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoner's Dilemma, Working Paper,Santa Fe Institute* 97-02-016, 1997.

[5] R. Riolo, M. Cohen, R. Axelrod, *Evolution of co-operation without reciprocity, Nature 414,* 441-443, 2000.

[6] R. Axelrod, *The Evolution Of Cooperation,* Basic Books, New York, (1984).

[7] M. Nowak, K. Sigmund, *The evolution of stochastic strategies in the prisoner's dilemma, Acta Applicandae Mathematicae,* (20):247-265,1990.

[8] W. D. Hamilton, *The Genetical Evolution of Social Behaviour I and II, Journal of Theoretical Biology,* 7: 1-16, 17-32, 1964.

[9] W.D. Hamilton, *Man and Beast: Comparitive Social Behaviour, Smithsonian Press, Washington DC.* eds., J.F. Eisenberg, W.S. Dillon, 1971.

[10] D.S. Wilson, G.B Pollock, L.A Dugatkin, *Can altruism evolve in purely viscous populations?, Evolutionary Ecology,* , 6, 331-341, 1992.

[11] K. M. Page, *Unifying evolutionary dynamics and a mathematical definition of selection, Proceedings of 5th European Conference on Mathematical Modelling and Computing in Biology and Medicine 2002,* ed. V. Capasso, 303-309, 2003.

[12] K.P. Hadeler, *Stable polymorphisms in a selection model with mutation , SIAM,* J. Appl Math, 41, 1-7, 1981.

[13] I. Bomze, R. Buerger, *Stability by mutation in evolutionary games, Games Econ. Behav,* 11, 146-172, 1995.