ABSTRACT

CHIVUKULA, SRIVATSA V. Evolution of Trust in Anonymous Interactions. (Under the direction of Munindar P. Singh).

For many applications, the success of multiagent systems depends on cooperation among the participating agents. Autonomous entities often face the economic and social dilemma of whether to act in trustworthy manner. Reputation mechanisms and institutional norms are currently used to make an individual agent responsible for its actions. It is challenging to achieve trust in an anonymous setting. We consider the problem of trust with respect to fairness, especially in an anonymous setting. Decision mechanisms should be built into agents to yield trustworthy behavior. We study conditions under which such cooperation can evolve. Our simulations show that small group interactions make it favorable for fairness to emerge. Techniques such as tagging can generate behavior analogous to small groups.

Evolution of Trust in Anonymous Interactions

by

Srivatsa V. Chivukula

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial satisfaction of the
requirements for the Degree of
Master of Science

Department of Computer Science

Raleigh

2003

Approved By:

Peter Wurman | R. Michael Young

Munindar P. Singh
Chair of Advisory Committee

To my parents

Biography

Srivatsa V. Chivukula was born on July 16th, 1978 in Hyderabad, India. He lived in Bombay until July 2001. He got his Bachelor of Engineering degree in Computer Science and Engineering from Bombay University in 2000. He was a masters student in North Carolina State University in the Department of Computer Science from August 2001 to August 2003.

Acknowledgements

I would like to thank my advisor, Dr. Munindar Singh for his constant support. I would like to thank Dr. Peter Wurman and Dr. Michael Young for serving on my thesis committee.

I thank Ashok and Amit for their encouragement and helpful comments. I thank Pınar for quickly familiarizing me with LaTeX. I thank Guru, Harish, and Lalit for being good and understanding roommates.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Trust is critical for business and social transactions. It is a precondition for the continued existence of any marketplace or society. With users' increasing dependence on the Internet, the importance of computationally modelled trust is increasing.

However, autonomous self-interested agents often either do not trust others or behave in a nontrustworthy manner towards others. As a consequence, economic transactions that could make everybody better off are not carried out. A possible solution, often seen in human societies, is if trust and trustworthiness are reliable (perhaps evolutionarily or culturally acquired) traits [Sen and Dutta, 2002].

Trust can apply to every aspect of an interaction. For example, in business contracts, questions of timeliness, price, and quality of service are crucial. However, for concreteness and simplicity, we will restrict our attention to situations where trust applies only to whether a given party is fair or not, as evidenced by the deal of offers. We also assume that a deal that is offered will be carried out as promised. Thus the sole aspect of trust considered is with respect to the fairness of the deals as it is offered.

In typical business the parties involved benefit only if the deal goes through successfully; otherwise none benefit. For simplicity, consider two-party settings. The party which moves first is often faced with the dilemma of whether it can propose an unfair deal (and if so, how unfair) or whether it must propose a fair deal. Proposing unfair deals is attractive in settings where the responders cannot readily determine if an offer is fair. This problem is particularly aggravated in interactions in which the responder cannot keep track of the proposer, e.g., those where the parties are anonymous.

We are interested in building adaptive mechanisms that promote fairness, i.e., lead to more of the offers made being fair. The mechanisms should apply in anonymous settings and where agents cannot (because of lack of knowledge) or would not (because of lack of computing power or time) determine whether a given offer is fair. In essence, the desired mechanism will involve some agents who do evaluate and prefer fair offers even if it is not immediately beneficial for them to do so. Even for these agents, the mechanism should be simple and impose a small cognitive burden.

1.1 Social Rationality

The problem in designing individual agent decision-making processes lies in determining appropriate actions for a given context or situation [Russell, 1997]. However, individual rational actions can produce inefficient, unscalable, and brittle outcomes [Kalenka and Jennings, 1999]. It is necessary to strike a balance between the preferences of individual agents and those of the society. The guiding principle for socially responsible agents is defined by Jennings and Campos [1997] as follows.

Principle of Social Rationality: If a member of responsible society can perform an action whose collective benefit is positive, then it may select that action.

Since a socially rational agent must consider the collective (or social) costs and benefits of its actions, implementing social rationality requires more computational resources than that for individual rationality. For example, the online auction site eBay, stores some information of the users' old transactional data and employs a feedback system to rate users. Hales and Edmonds [2003] argue that unless agents are endowed with sophisticated induction processes that allowed them to build models of their social context, the implementation requirements of social rationality would have to change every time the composition or specific task domain of the system changes. Also such a system could be unstable. Agents that are malicious, malfunction, or have imperfect information of the states can negate the potential benefits of social rationality. A highly desired feature is for social rationality to emerge from the interactions of autonomous agents, rather than be built in by a designer for every situation.

Our aim is to find minimal decision mechanisms that, when built into agents, produce socially rational solutions to common multiagent trust and fairness problems. One mechanism that can support the emergence of social rationality is tagging, as described in Section 1.3.3.

1.2 Trust

A society's success in producing trust among its members depends on its ability to help agents construct reliable assessments of the fairness of other community members. The following are two well-known approaches to engender trust:

Norms and Institutions Norms are specific allowed behaviorial ranges imposed upon agents. Institutions reduce the uncertainty in the behavior of other agents by prescribing norms. They detect and penalize agents who deviate from the norms. Institutional guarantees reduce the problem of trusting individual agents to that of trusting the institutions. If one trusts that the institutions will do their job, there is less need to assess the trustworthiness of every individual.

The use of third-party norms, however, has some disadvantages. Assessing the effectiveness of such institutions is not always trivial, especially for newcomers to a given agent community. As agent communities span different territorial jurisdictions, the agents might exist in different legal systems. Also, in agent communities, it is easy to change identities and act in a selfish manner.

Reputation This is based on information about or observation of an agent's past behavior in similar situations. Such information is aggregated through direct dealing and distributed by word-of-mouth or by trusted third parties. Reputational information can help agents build estimates of other agents' trustworthiness under the assumption that an agent's past behavior is indicative of its future behavior. However, even under this assumption, one cannot conclude that a deal being offered is fair or if better offers are available.

Reputation systems do not provide reliable information when agents change their identities. Further, it is difficult to translate (or scale) someone's rating of an agent to another agent's reputation-point system. Also, some raters may give unfair ratings intentionally [Dellarocas, 2000].

The above methods for producing trust heavily depend on the identity of the agent to be trusted. Thus, they cannot be applied to anonymous communities.

In open environments there is no central control over agent behaviors. Agents in such systems are primarily driven by self interest. In societies that use norms to filter out unfair deals, the rules are set by third parties and are not flexible.

1.3 Identity

The Internet has given rise to numerous social and business environments that allow frequent and meaningful interactions among strangers. Many games, auction sites, and interactive forums allow users to choose a pseudonym when they register. Even services that identify users through email addresses do not prevent identity changes. Beyond name changes, the Internet enables completely anonymous interactions. For example, anonymizing intermediaries such as remailers and proxy servers can exchange messages between parties without revealing either one's identity to the other.

1.3.1 Anonymity

Anonymous communities are preferred when the interaction is short-lived and it is not required to interact with the same agent again. For our purposes, we assume that the identity of an anonymous agent cannot be guessed by other agents in the system. The infrastructure may know the identity but does not reveal it to anyone.

1.3.2 Pseudonymity

Systems can keep track of individuals by assigning a pseudonym to them. This is commonly seen on the Internet where sites can store certain preference on the client machine as cookies. These are later used to give a customized response to the individual. The cookie acts as a pseudonym. On other sites, registration and log in are required in order to participate. Thus an agent would know if it is dealing with the same agent again, but may never know the true identity of any agent.

1.3.3 Tagging

Human interactions are often limited to observance of similar or almost similar characteristics, as those in themselves, particularly in a new (non-trustworthy) environment. For example, a person on a trip to a foreign country may unhesitatingly trust another person from his native country, even though he may never have seen him before. Such interactions tend to propagate trustworthiness in a society.

Tags are such observable social cues attached to agents [Holland, 1993]. An agent can distinguish between agents with different tags, but not necessarily among agents with the same tag. The tag-based selection of partners does not require agents to have extensive cognitive capabilities (as required by reputation mechanisms). The tag-based selection of partners in effect induces an abstract topology in which the agents are located "near" (prefer to interact) or "far" (prefer not to interact) from each other.

1.4 The Ultimatum Game

The ultimatum game [Page and Nowak, 2002] describes a social dilemma inherent in many economic scenarios. Figure 1.1 illustrates the ultimatum game. The top level branching corresponds to the decision by the proposer as to the different ways to share \$20. The second branching corresponds to a decision by the responder as to accept or reject the offer. The payoff for both parties are listed at the bottom.

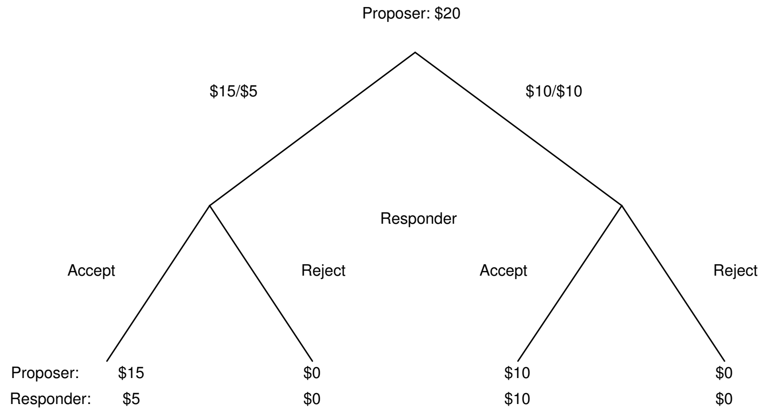

Figure 1.1: The ultimatum game

In general, a proposer can offer a fair split of a resource, by dividing it equally or can offer some amount less than half the resource. The latter deal is unfair to the responder as the proposer tries to get away with a larger share. The responder can now either accept the deal or reject it. If the responder rejects it, then neither benefits since the transaction does not occur.

A rational solution to such a dilemma would be that the responder, bent up on maximizing his payoff, accepts any nonzero offer. A rational proposer should therefore make the smallest nonzero offer, on the assumption that it will be accepted. However, such behavior is not seen among humans. Well-informed players decide to reject the unfair offers, whereas others accept them due to the lack of sufficient knowledge or some other reason such as urgency. But evolutionary dynamics applied to such population shows players quickly start accepting all behaviors irrespective of whether they are fair [Page and Nowak, 2002]. Thus the social rationality of offering fair deals disappears.

Traditional markets are based on the assumption that that there are clear and well-defined measures which allow one to quantify the quality of any item or service with respect to another item or service. This might not be possible in all cases. The relative need for a product or service in a society plays an important role in determining the fairness of a deal, which therefore depends on external factors as supply and demand.

A possible solution to achieve social rationality includes agents that learn about their network neighborhoods, new methods of routing messages efficiently in multi-agent systems, and more accurate methods of making resource selection decisions in environments containing many resources.

However, these goals are difficult to achieve in distributed anonymous networks. Some of the current projects which deal with such networks are:

Gnutella Gnutella is a peer-to-peer (P2P) in which each peer plays the role of a client and a server, and is called a servent [Gnutella]. The strength of Gnutella is its extremely flexible network design. When a search is performed, servers respond with an external IP address or URL where the user can download the document. Since this actual retrieval is done without any privacy protection, using Gnutella is not a good choice if publishers or readers want anonymity.

Freenet The Freenet project is developing a peer-to-peer network that is entirely decentralized and publishers and consumers of information are anonymous [Freenet]. Freenet aims to make files highly accessible and offer some level of anonymity. Since the choice to drop a file is a purely local decision, and since files that are not requested for some time tend to disappear automatically, these systems do not guarantee a specified lifetime for a document.

Free Haven The Free Haven Project is dedicated to designing a system of anonymous storage [Dingledine, 2000]. It assures reader-publisher anonymity, server accountability, and persistent storage for data independent of its popularity. Servers are associated with distinct pseudonyms and trade chunks of files (shares) over the network. Server accountability is achieved by reputation mechanisms. But adversaries can track and target servers by their pseudonyms.

1.5 A P2P Storage Space Bidding Problem

We discuss a peer-to-peer storage space replicating problem as an example scenario for the ultimatum game. Cooper and Garcia-Molina [2002] propose a cooperative replication network where sites trade space, so that each site contributes storage resources to the system and uses storage resources at other sites. A local site wishing to make a copy of some data announces how much remote space it wants and accepts bids for how much of its own space it must give to acquire that remote space. A series of such agreements between pairs of sites builds up a peer-to-peer trading network. Although each site makes a local decision for local benefit, the result is a global network dedicated to replicating data.

For example, consider the following negotiation of agreement between peers. Site A may want to replicate 100 GB of data. Site A may contact site B and request a trade, and site B may respond that it is willing to trade if it receives 150 GB of site A's space in return. If site A contacts multiple sites asking for trade, then site A will receive multiple such bids, and can pick the "lowest" bid. Thus, an agreement may be concluded between site A and some other site C, where site C gives A 100 GB and in return A gives site C 125 GB.

This work differs from other distributed computing systems [Ferguson et al., 1993; Schwartz and Kraus, 1997] using market-oriented principles (like auctions), in the following ways. First, there is no concept of money involved in trading. Sites use the barter system for exchanging resources. Second, there is no distinction between producers and consumers. Producers do not have different incentives and do not follow different policies than consumers. Every peer can be either a producer or a consumer in these transactions.

However, Cooper and Garcia-Molina consider such a system only under a restrictive bidding freedom for the proposers. They experiment with policies in which peers always make fair proposals only considering the space requested and the free space they have. Consider the following free market scenario, in which sites obey the basic rules of trade (i.e., they offer resources which they advertise) but follow their own polices while making proposals. Site A may contact site B asking for 1 GB. In return site B, driven by its own selfishness, may ask for 100 GB. If all other sites whom A contacts also make such offers, site A would have to pay a very high price for replicating its data.

Consider an anonymous system in which sites cannot consult each other before making decisions. Site A has to make a take-it-or-leave-it decision as in the ultimatum game (Section 1.4). Site A might accept the offer and try to recover the loss when it proposes an offer. If all other sites adopt this behavior, the system disintegrates, thus making it almost impossible for fair trade to take place. Resource-rich sites might be able to carry out the trade, but sites with low resources cannot use the system any more.

One way to avoid the disintegration would be for sites to interact with other sites based on their reputation. But maverick sites can change their identities frequently to avoid other sites from building up such information. Also unlike in traditional markets, where repeatedly trading with same agents can be favorable, each site has to look for newer sites to increase reliability and fault tolerance. In other words, sites should store their data on different other sites, so that one site breaking down does not cause much harm. However, reputation approaches are biased towards old sites, discouraging the entry of new sites into the network.

In such a free market scenario the social rational behavior of offering fair deals disappears. This is because the system goes into a state in which every agent makes an individually rational offer, i.e., they ask for the most space that they can get.

Our aim is to evolve socially rational behavior without sacrificing anonymity. We study a theoretical model for decision-making by individual nodes based on the ultimatum game in Section 3.2. We explain the results of our experiments in Chapter 4, and develop a conceptual solution based on our results in Section 4.1.

Chapter 2

Evolutionary Programming

Evolutionary algorithms are stochastic methods that mimic natural evolution [Pohlheim, 1997]. These algorithms operate on a set of potential solutions and apply the principle of the survival of the fittest to produce better and better approximations to a solution. At each generation, a new set of approximations is created by the process of selecting agents according to their level of fitness in the problem domain and letting them interact to produce a new set with different fitness levels. As in natural adaptation, this process leads to the evolution of populations of agents that are better suited to their environment than the agent population they were created from. Evolutionary algorithms model natural processes, such as selection, recombination, mutation, migration, locality and neighborhood. We use evolutionary algorithms to study how agent preferences vary in our experiments. The agents whose preferences fetch high payoffs, survive in the population.

2.1 The Evolutionary Process

Let us first describe the evolutionary process informally. At the beginning of the computation a number of agents (in the population) are randomly initialized by some preset distribution. These agents interact with each other and the fitness levels are evaluated for these agents. The initial generation is produced.

If the termination criteria are not met, the creation of a new generation begins. Individuals are selected according to their fitness for the production of offspring. Parents are recombined to produce offspring. All offspring are mutated with a certain probability. The fitness of the offspring is then computed. The offsprings are inserted into the population replacing the parents and thus producing a new generation. This cycle is repeated until the termination criteria are reached. Algorithm 2.1.1 summarizes the above steps in pseudocode.

Evolutionary algorithms differ substantially from more traditional search and optimization methods. The most significant differences are as follows:

- They search a population of agents in parallel. Only the objective function and corresponding fitness levels influence the directions of search.
- They use probabilistic transition rules, not deterministic ones. Evolutionary algorithms can provide a number of potential solutions to a given problem.

An evolutionary algorithm consists of four stages: selection, recombination, mutation, and reinsertion. Various approaches can be used for each of these stages. The methodology adopted in our simulations is explained below.

Algorithm 2.1.1: EVOLUTIONARYPROGRAMMING(P)

```
/* start with an initial time */
t ← 0;

/* initialize a usually random population of individuals */
INITPOPULATION(P^t)

/* evaluate fitness of all initial individuals of population */
EVALUATE(P^t)

/* test for termination criterion (time, fitness, etc.) */
while notdone
  do
  {
    /* perturb the whole population stochastically */
    P^t ← MUTATE(P^t)

    /* evaluate its new fitness */
    EVALUATE(P^t)

    /* stochastically select the survivors from actual fitness */
    P^{t+1} ← SURVIVE(P^t)

    /* increase the time counter */
    t ← t + 1
  }
```

2.2 Selection

Selection determines how to choose the agents whose strategies have to be adopted. Each agent is associated with a fitness value.

We use the roulette-wheel mechanism for selection [Baker, 1987]. In this mechanism, the agents are mapped to contiguous segments of a line, such that each agent's segment is equal in size to its fitness. A random number is generated and the agent whose segment spans the random number is selected.

Agent	1	2	3	4	5	6	7	8	9	10
Fitness value	2.0	1.8	1.6	1.4	1.2	1.0	0.8	0.6	0.4	0.2
Selection prob	0.18	0.16	0.15	0.13	0.11	0.09	0.07	0.06	0.03	0.02
Segment position	0.18	0.34	0.49	0.62	0.73	0.82	0.89	0.95	0.98	1.0

Table 2.1: Selection probability and fitness value

Table 2.1 shows the selection probability for 10 agents, where the agents are vying for a fitness value of 2. Agent 1 is the most fit and occupies the largest interval (0-0.18), whereas agent 10 as the least fit has the smallest interval (0.98-1.0).

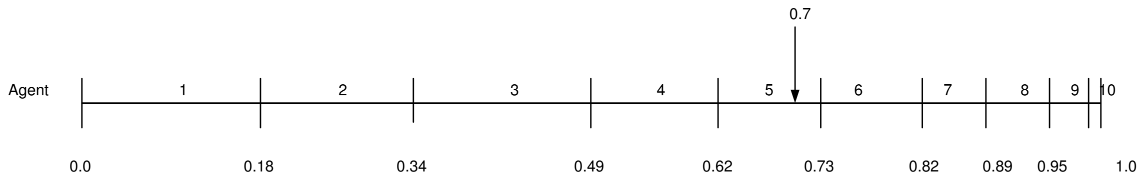

Figure 2.1: Probabilistic selection

For selecting the agent for recombination, a random number (uniformly distributed between 0.0 and 1.0) is independently generated. For example, if the random number generated is 0.7 then agent 5 is selected, as shown in Figure 2.1. The probability of an agent being selected is proportional to its fitness value.

2.3 Recombination

Recombination produces new agents by combining the information contained in the agents from previous generation.

In our experiments, every agent selects two other agents using the roulette-wheel selection mechanism. It compares their fitness values and chooses the strategy of the agent with the highest fitness value.

2.4 Mutation

After recombination every agent undergoes mutation. We used binary mutation where each agent has a binary tag. A tag bit is mutated after every generation with a low mutation rate of 0.01.

2.5 Reinsertion

If fewer agents are produced than the size of the original population, agents have to be reinserted into the old population. Similarly, if more agents are generated than needed, then a reinsertion scheme is used to determine which agents should be inserted into the new population. However, in our approach, the recombination method described above results in keeping the total population of agents constant. Therefore, no special reinsertion method is used.

Chapter 3

Trust Dynamics

We now consider a community of anonymous agents, each of which can play the role of proposer and responder of a transaction, as in the ultimatum game.

We need three basic prerequisites that each autonomous agent should possess.

- Each agent should know its utility function. The utility is usually independent and private to each agent and is some quantity (a real number) related linearly to its benefit from a transaction. We assume linearity for experimental simplicity.
- Each agent should set a minimum threshold of satisfaction relative to a transaction. This usually changes as the agent participates in the transactions but may be fixed in the beginning. However, setting the threshold too high might cause rejection of majority of deals and setting it too low might result in the inability to filter out unfair transactions.
- Each agent should be able to model the behavior of other agents.

3.1 Service Advertisement and User Preferences

Let X be the set of potential issues included in a contract or an advertisement. Let the domain of the values taken by x be D_x. A contract C is a set of issue-value assignments denoted by $\{x_1 = v_1, x_2 = v_2, \ldots, x_n = v_n\}$ where $x_i \in X$ and $v_i \in D_{x_i}$.

We assume that fairness f is a domain-specific monotonically increasing function. The function f maps an advertisement or contract C to a real number in the interval $[0, 1]$. This means that when two deals or services are compared, the better one will get a higher rating (as generated by this function). An example of this function, in the information retrieval domain, is the cosine similarity between the the term vector (e.g., as generated by the TF-IDF approach) of each document and a given query vector [Salton and Buckley, 1988]. Finding such functions for a specific domain is beyond of the scope of this work. In the next section, we deal with two types of agents in the society and, in an idealized case, show that selecting fair deals is a dominant strategy.

The set of agents can be partitioned based on the levels of knowledge or preferences on offer selection, when in the responder role. That is, some agents may be more demanding, some may know the exact requirement, some may be regular users with basic knowledge of the domain, and others may be novices having neither the product know-how nor the domain knowledge to make a selection. We assume that there is a natural ordering among the partitions. Specifically, if there are i partitions of A, then the first partition of agents will accept any offer, the second partition will only accept offers greater than some value f_1, the third will only accept offers greater than some value f_2, and the ith partition will accept only fair offers. That is, $0 = f_1 < f_2 < \ldots < f_i = 1$. In this framework, the first and the ith partitions are important cases representing the extreme behavior in the spectrum of preferences. The agents having these extreme preferences are described as:

Myopic agents They are only interested in their immediate personal payoff. These agents accept the split even if it is unfair. This behavior can be mapped to barter trading systems, in which agents either do not know whether a particular deal is fair or unfair (this information may be expensive to obtain) or may accept the unfair split for some reason such as urgency or necessity.

Prudent agents They are interested in being treated fairly (but may or may not treat others fairly). They would reject an unfair offer. In other words, they assign a *moral premium* to fairness above the direct benefit they derive from a deal.

In the abstract model described in the next section, we define fairness as a deal in which the proposer proposes an equal split of a resource.

3.2 Model Definition

A rational agent is one that tries to maximize its utility. Let $A = \{a_1, a_2, \ldots, a_N\}$ be the set of N rational agents belonging to a society over a domain D.

A proposer P is an agent that proposes a deal. A responder Q is an agent that either accepts the offer proposed by the proposer at that instant or rejects it. Formally, at any instant t, $(P^t, Q^t \subseteq A) \wedge (P^t \cap Q^t = \emptyset)$

Let r be the resource, p the proposer, and q the responder. r_{pq} is what the proposer p keeps to himself during a split between p and q. Normalize resource r to 2. Let ϵ

(where $0 \leq \epsilon \leq 2$) be an offer. An offer is fair for the responder when $\epsilon \geq 1$, whereas the proposer who tries to maximize his benefit makes an offer $\epsilon \leq 1$.

A proposer p has two kinds of offers to choose from:

- Unfair (E): $r_{pq} = 2 - \epsilon$ (where $\epsilon < 1$)
- Fair (F): $r_{pq} = 1$

A responder q has two options:

- accept: $\theta = 1$
- reject: $\theta = 0$

Myopic agents can accept any split and prudent agents accept $(2 - \epsilon, \epsilon)$, for all $\epsilon \geq 1$ and reject it for all $\epsilon < 1$.

Agents derive subjective utility from the allocation of resources. We can assume each agent to be maximizing an internal utility function. The utility equals the amount of resources consumed, plus a *moral premium* for rejecting unfair offers (while a responder).

$$\mu_p(r_p) = r_p + (1 - \theta) * \rho$$

In this equation,

$r_p \in \{0, \epsilon, 1, 2 - \epsilon\}$

$r_p = \theta * r_{pq}$ when p is the proposer

$r_p = \theta * (2 - r_{pq})$ when p is the responder

$\rho = 0$ for myopic agents

$\rho =$ moral premium for prudent agents which is strictly larger than ϵ (so that it gets some satisfaction from rejecting the offer)

Although the moral premium helps the responders (by making them happy) while making a decision in the single instance of the game, it has no effect on the evolution of agents. In other words, prudent agents cannot survive by rejecting offers. The agents' fitness values are equal to the amount of resource obtained in the transaction.

The decision of the proposer (on how to split the given resource) will depend on his beliefs about the distribution of the two types among the possible responders.

Let the total number of agents be N. In general, if agents interact in groups, let N_j be the number of agents in the jth group. Each group is a closed agent community in which agents, with certain common characteristics, trade among themselves over items belonging to a particular domain. If an agent belongs to a group of size one, it is isolated from all others.

One round of interaction is a time interval in which an agent proposes different offers to every other agent. One generation is a time interval in which all agents in this domain get to propose exactly once. Thus in every round t every agent in the jth group interacts $2(N_j^t - 1)$ times, where N_j^t is the number of agents at time t.

We assume that the composition (n_M, n_P) is known to all agents, where n_M and n_P are the number of myopic agents and prudent agents respectively. We can weaken this assumption by believing that players can become aware of changes in the population only one period after the changes have happened or that instead of knowing the true distribution of the types, the players have some beliefs about it. We surmise that the latter only has the effect of slowing down the convergence in our experiments.

3.3 Experimental Setup and Results

We now describe the experiments performed to see how the distribution (n_M, n_P) varies in such populations. The main aim of the experiment is to show that fairness can be an evolutionarily stable strategy without any institution enforcing it despite the fact that the agents are anonymous.

3.3.1 Interactions with Complete Anonymity

We have experimented with N = 100 agents. In one round of interaction each agent interacts with every other agent twice, once as a proposer and once as a responder, in a round robin fashion. As a proposer each agent randomly offers a fair deal with probability α or an unfair deal with probability $(1 - \alpha)$ to all other agents where α is the proportion of prudent agents in the system. We measure fairness of the entire population in terms of α. Thus, as α gets closer to 1, fairness evolves. For example, when in the role of a proposer, if an agent generates a number 0.8, it means that it is willing to offer 0.8 units of the resource to the responder and keep 1.2 units to itself. When in the role of a responder myopic agents accept all offers, whereas prudent agents accept only those offers in which the proposers agree to give them a fair deal.

After each generation, we calculate the average performance of the myopic and the prudent agents. Performance of each agent is directly proportional to the amount of resources it obtained in the corresponding round of interaction. New agent behavior assignments are made as described in Sections 2.1 and 2.2: for each agent a_l, two new agents are selected according to their performance. The probability of selecting the jth individual with performance p_j is $\frac{p_j}{\Sigma_{k=1}^{N} p_k}$. Then of these two selected agents, the behavior of the one with higher performance is adopted by agent a_l. This leads to the propagation of successful behaviors or traits in our experiments. As a result, if a behavior produces a better performance in one evaluation period compared to other behaviors, we are likely to see more agents adopting that behavior in the next evaluation period. We run this generational process until the population becomes homogenous (i.e., no new trends emerge) or for a fixed number of evaluation periods (we have used a limit of 10,000) is reached.

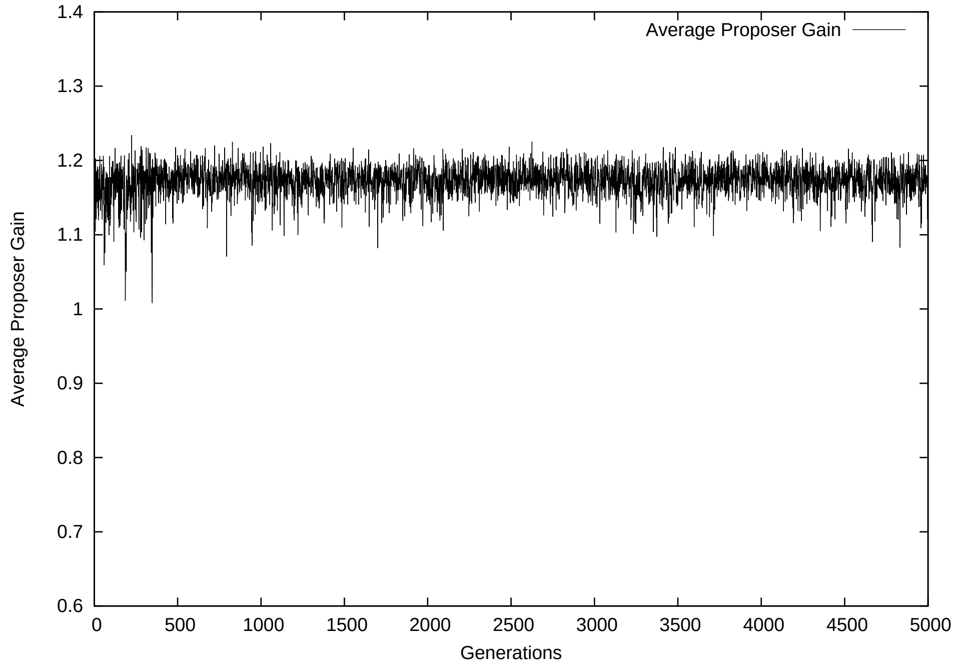

Figure 3.1: Average proposer gain in an anonymous system with 100 agents and the proposers offering either $(1, 1)$ or $(1.3, 0.7)$ split of the normalized resource depending on the proportion of the prudent agents

We ran experiments with no prudent agents initially. This implies that all agents, when in the role of responders, choose any deals that are proposed. Also, since distribution of the types is known to all agents, when in the role of proposers, they offer unfair deals with a high probability. Myopic agents win over the prudent agents by receiving increased payoffs when in the role of responders. As a result, the prudent agents never evolve.

In Figure 3.2, we plot the number of prudent agents versus the number of generations. After each generation the distribution of the myopic and the prudent agents changes by the probabilistic selection and recombination methods described in Section

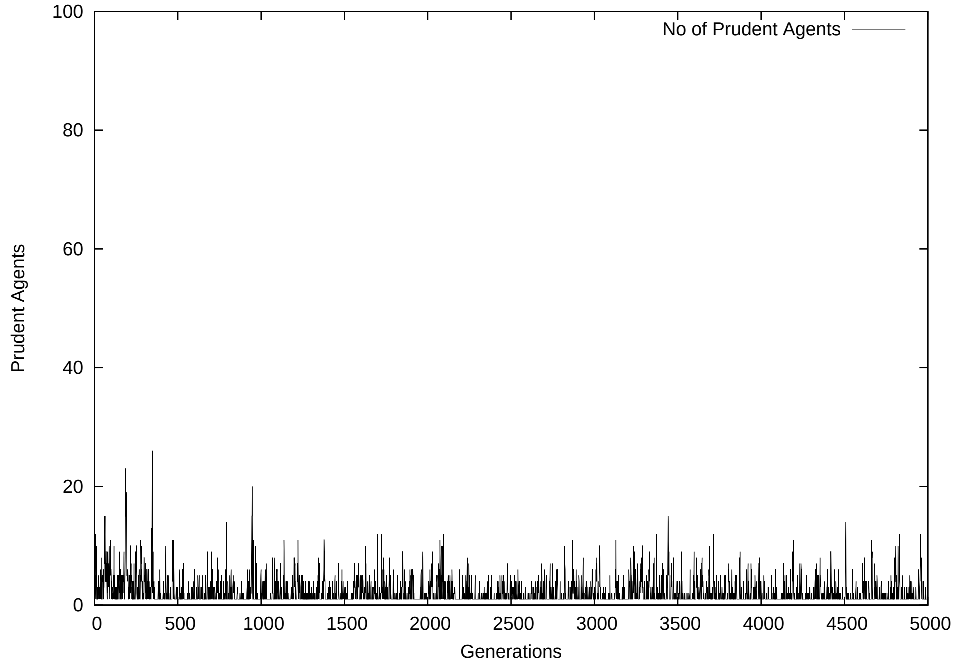

Figure 3.2: Evolution of prudent agents in a totally anonymous system

2.1. The proposers always obtain a resource greater than one, as shown in Figure 3.1. This result is intuitive because the myopic agents always receive a nonzero unit of resource when in the role of a responder, whereas the prudent agents obtain nothing as the number of unfair proposals increase.

Next we ran the simulation with the prudent agents in majority (almost the entire population). Now initially, all the agents propose fair offers with high probability. However, as the myopic agents accept all offers and the prudent agents reject the unfair ones, the fitness of the myopic tends to increase. The population goes into a fixed point where every agent, as proposer, offers an unfair deal and, as responder, accepts every offer. To ensure that fairness is a dominant strategy even in a large agent population, the relative advantage to the myopic agents should be reduced.

However if the agent population is small then the prudent agents, which now

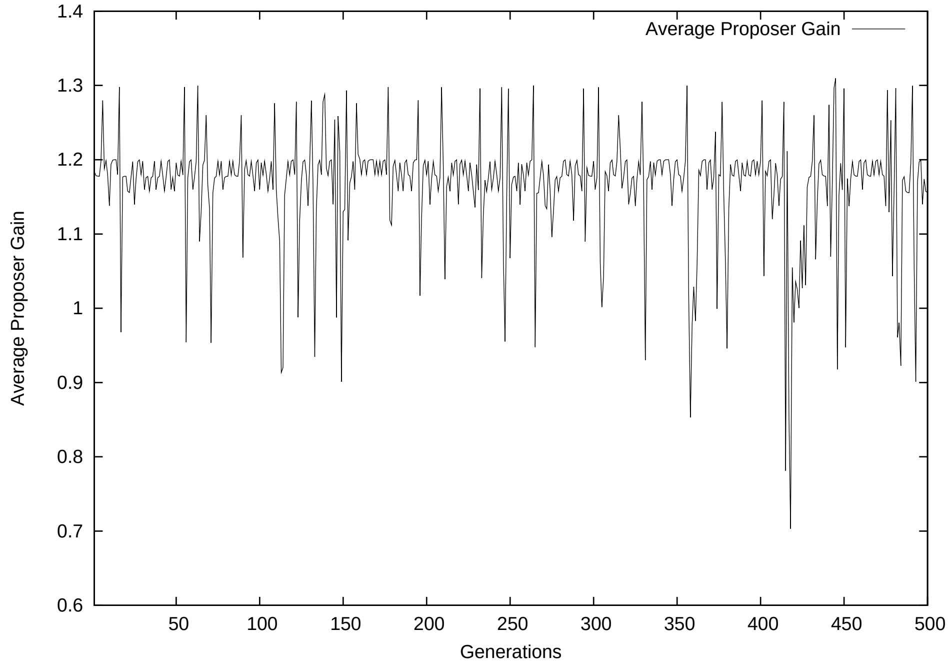

Figure 3.3: Average proposer gain at the end of each round of interaction in an anonymous system with 10 agents

make unfair offers but accept only fair offers, perform better on average. This can force the system out of the previous absorbing state (local maxima). As the prudent agents increase, agents start proposing fair deals to maximize their payoffs, as seen in Figures 3.3 and 3.4.

Figure 3.4 shows the behavior of the system with 10 agents. This graph shows simulations over 500 generations. The system with 10 agents is able to recover from unfair deals but it cannot sustain it.

3.3.2 Interactions with Tags

To produce a socially rational outcome, we now use tags. The basic setup is the same as before, but each agent now possesses a tag, encoded in L bits. The tags have

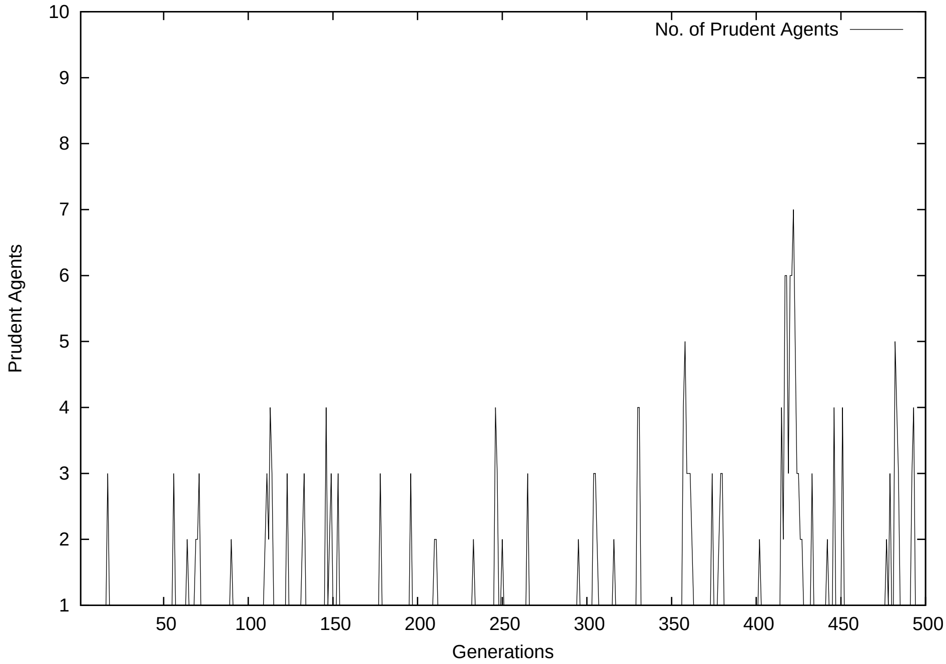

Figure 3.4: Evolution of prudent agents when the agent population is low

no direct effect on the action selected by the agent. However, tags are observable by all agents.

Agents are selected to interact based on the tags they possess. If an agent can find others who have the same tag then it interacts with them. However, if an agent cannot find another having the same tag, then it selects agents with tags at a distance of one and which have not interacted. The agent population is varied as before, after each generation but in addition they undergo mutation (see Section 2.3) of their tag bits with some low probability (we used a probability of 0.01). We find that if the number of tag bits is sufficient, the socially rational solution of all agents offering fair deals dominates. This is seen in Figures 3.5 and 3.6.

We found that 5 tag bits were sufficient to evolve fairness in our simulations for 100 agents. In about 2,500 generations, a socially rational system of all agents proposing fair deals evolves. It is interesting to note that the system was started with almost no prudent agents.

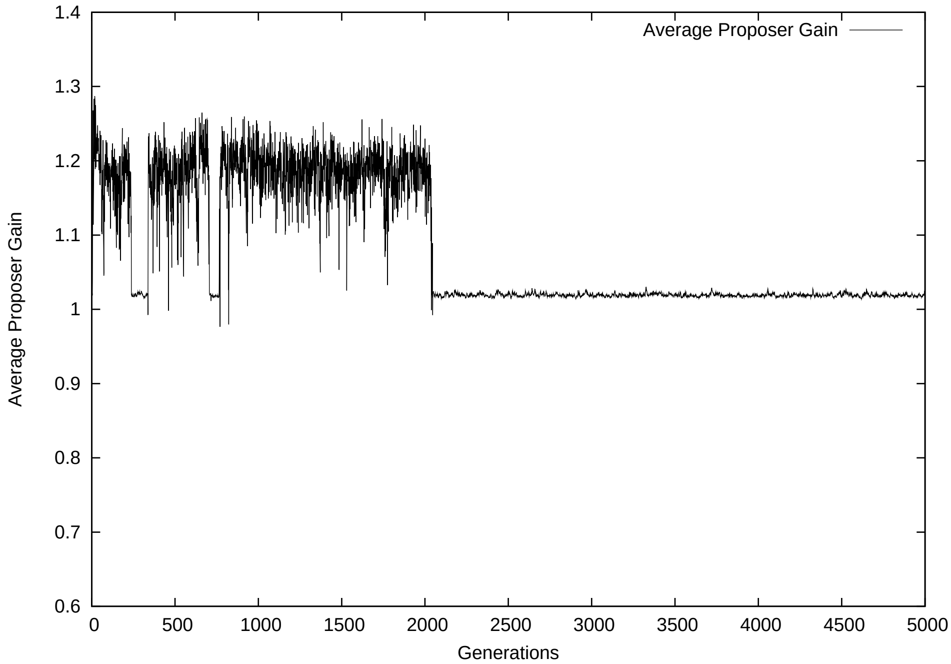

Figure 3.5: Average proposer gain with 100 agents and each agent having a 5 bit tag. Agents interact with others based on their tags. The proposers offer either $(1, 1)$ or $(1.3, 0.7)$ split as before

This evolution can be explained by understanding that in the tagged system, agents with the same tag form a sort of an "individual group." The entire MAS can be considered as a collection of groups. Even though evolution occurs over the entire population, prudent agents benefit when the interactions are in small groups. As a result they evolve slowly, as shown in our experiments.

We studied the average proposer gain with respect to the group size. The results plotted in Figure 3.7 show that as the number of tag bits increases, the proposers tend to offer an equal split of the resources. Note that, in these experiments, the agents do not show any bias to the tags of other agents that they are interacting with.

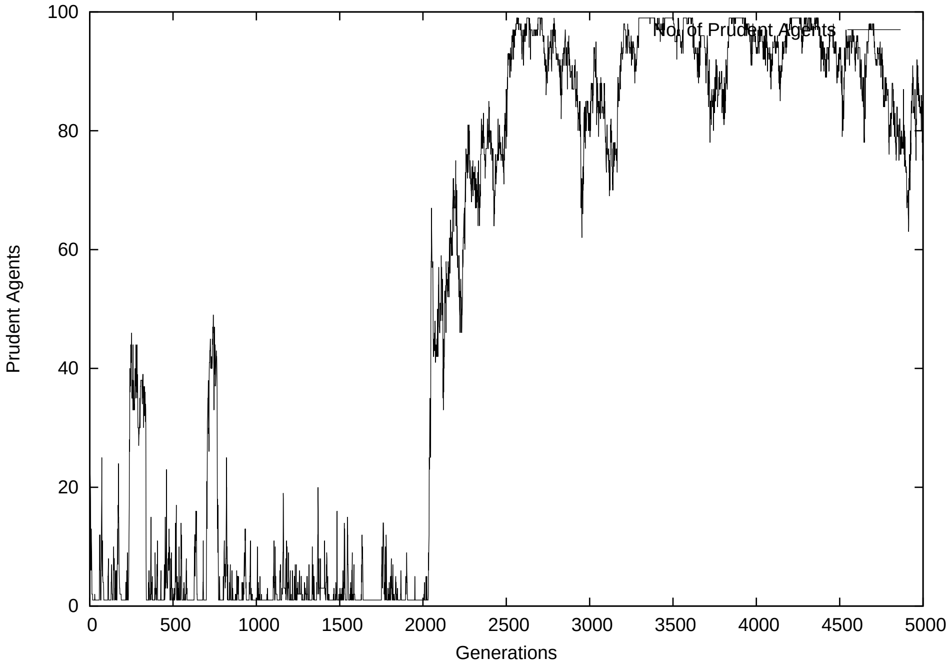

Figure 3.6: Evolution of prudent agents with 100 agents each associated with a 5 bit tag

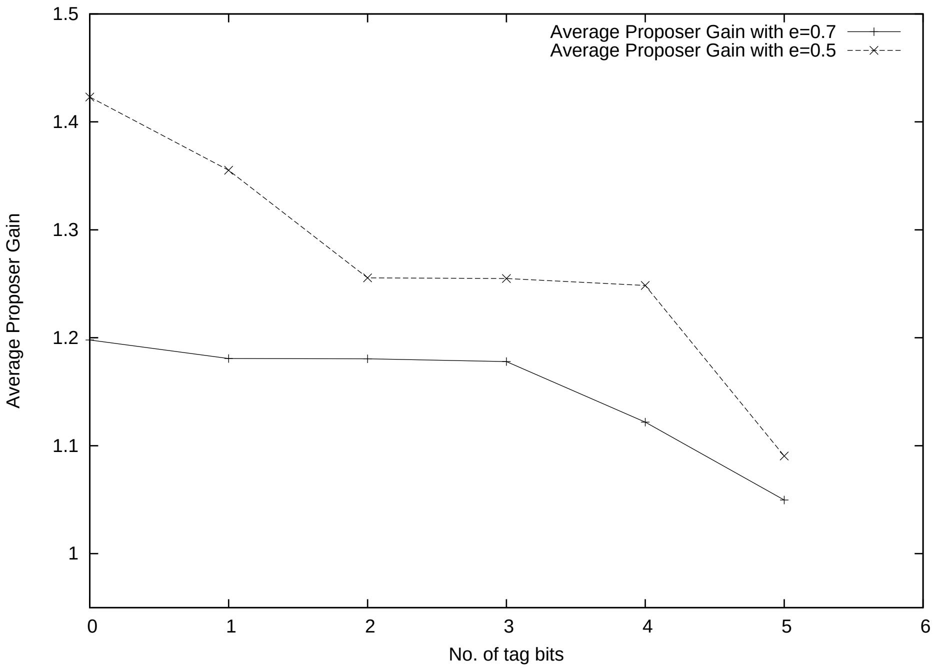

Figure 3.7: Comparison between proposer offer policies (average of 5,000 runs)

3.3.3 Interactions with Tag Bias

Next we experimented with tag-biased agents in addition to a mix of the above agents in the population. The positively tag-biased agents offer fair deals to other agents having the same tag and unfair deals to agents with different agents. These are specialized prudent agents which enforce group bonding. We found that in such an environment agents make fair offers even with fewer the tag bits.

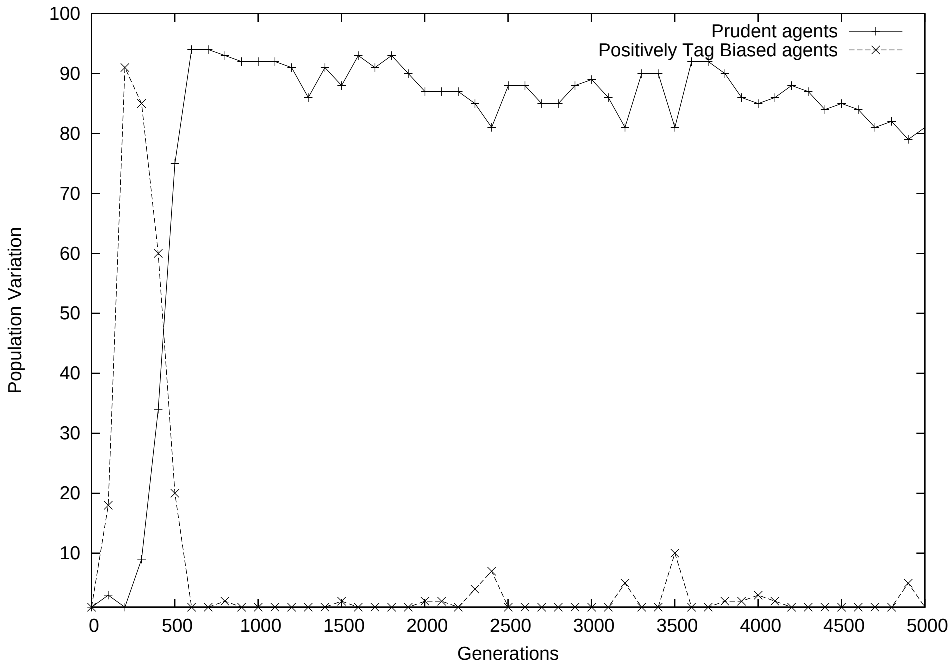

Figure 3.8: Comparison in the evolution when myopic, prudent, and positively-tag based agents are present in the population

Figure 3.8 shows the evolution of the population in the presence of myopic, prudent, and positively tag-biased agents with 3 tag bits. It is interesting to note that the tag-biased agents help in the quick evolution of the prudent agents. The tag-biased agents themselves do not survive when the prudent agents increase. This is because though they are able to make cooperative relationships with agents of the same tag, their performance suffers badly when they interact with the evolved prudent agents of different tags.

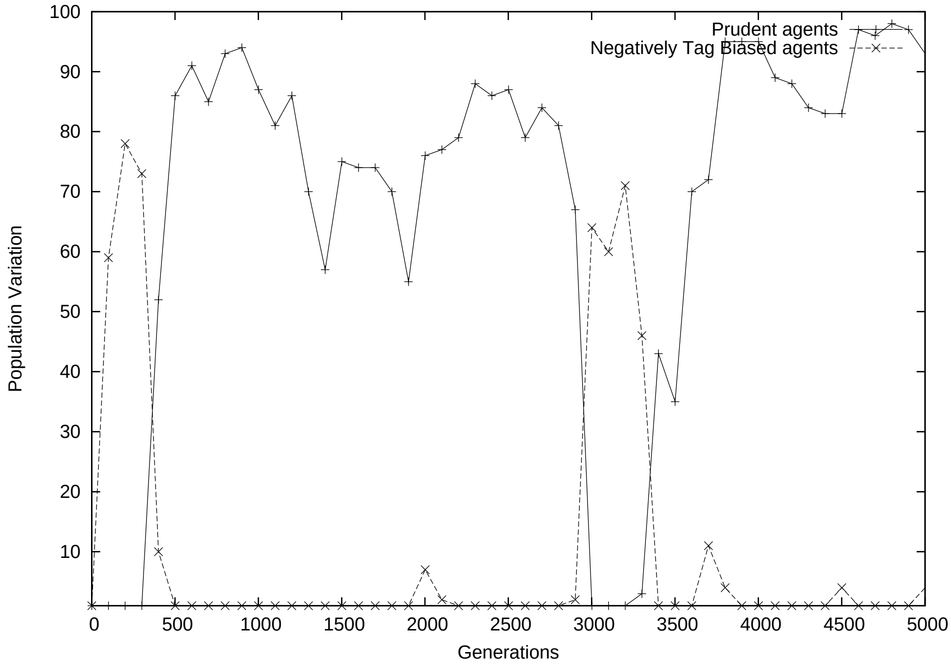

Figure 3.9: Comparison in the evolution when myopic, prudent, and negatively-tag based agents are present in the population

We observe a similar behavior with negatively tag-biased agents. Negatively tag-biased agents make unfair offers to agents of the same tag and fair offers to others. Figure 3.9 shows the population variance with myopic, prudent, and negatively tag-biased agents. Although the negatively tag-biased agents help in the quick evolution of the prudent agents, they do not allow the latter to exist.

Chapter 4

Discussion

Our approach is different from those discussed by Sen and colleagues [Sen, 1996; Sen and Dutta, 2002], where the agents cooperate with others at a cost. Sen and colleagues experiment with situations where agents can help other agents by sharing work such that the cost to the helper is less than the cost saving of the helped agent. In our experiments agents react rationally to the environment. That is, they can condition their behavior to the state of the system. We assume that the evolution of certain preferences in the population depends on the relative success of agents endowed with those preferences and no exogenous incentives are involved. We see four preferences developing in the system.

- Offer a fair deal when in the role of the proposer and accept only fair deals when in the role of responder (F-F).
- Offer a fair deal but accept all deals (F-A).
- Offer an unfair deal and accept all deals (U-A).
- Offer an unfair deal but accept only fair deals (U-F).

Consider an initial state where all players have the preference F-F; i.e., they offer fair deals and accept only fair deals. Agents offering unfair deals are punished by other agents, who reject their offers, thereby yielding a lower payoff to the unfair proposer. Therefore, fair deals are promoted by the effective threat that unfair deals will be rejected. However, agents with the preference F-A can easily enter the system. When there is a sufficient number of agents having the F-A preference, agents with the U-A preference enter the environment, and receive higher payoffs than agents with the previous two preferences. The system drifts away to a state where all agents offer unfair deals.

Agents having the U-F preference can pull the system out of the previous state. If agents interact in small groups, agents with the U-F preference get a higher payoff than those with the U-A preference. Once there is a large number of agents with the U-F preference, agents learn (prefer) to offer fair deals. The system gets back to the state where everyone offers a fair deals and accepts fair deals. Thus agents in a small group tend to give fair offers because the loss for the proposer, who gives out bad deals to a responder is higher than the loss to the responder if it gets rejected.

4.1 Tag Based Agent Interaction Protocol

We now describe a skeleton for a communication protocol based on tagging in which agents do not reveal their identity to other agents. More than one agent can be mapped to a single tag, thus supporting a many-to-one mapping. A public-private key is associated with each agent. Each agent is assigned an identity (agentId) by a broker agent who stands for the infrastructure. The agentId is based on a cryptographic hash of the agent's public key. Communication between the participating agents and the broker agent can be encrypted by public-private key encryption. And, authentication can be provided by exchanging digital certificates. Tags are assigned by the broker agent to every agent when it joins the anonymous multiagent system. The routing protocol between participating agents makes use of the tags to route messages only between specific agents. The system should be able to organize itself, even when agents are not able to maintain reputational information of other agents.

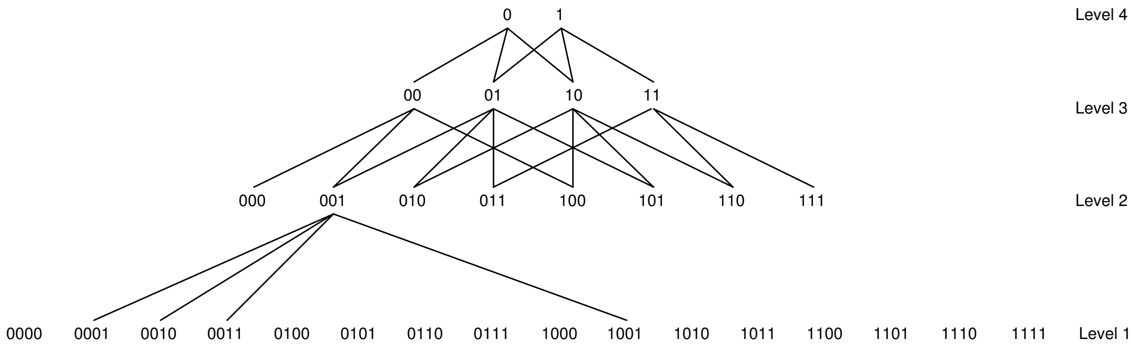

Figure 4.1: An example for a hierarchy of tags issued by the broker agent according to the subsequence relationship

Initially all agents are assigned tags at level 1, and they interact with other agents in the same tag group. There are no restrictions on the offers made or accepted. Agents can either make fair offers or unfair offers maximizing their net gain in their transactions.

An agent's request messages are multicast to agents with tags with which the agent's tag forms a common subsequence. For example, as shown in Figure 4.1, the message of an agent with tag 001 is multicast to other agent with the same tag by the underlying routing protocol, and to agents with tags 0001, 0010, 0011, and 1001. Agents fulfilling the threshold of the number of accepted offers for a particular level keep progressing to the next higher level by decrementing their tag length.

The broker agent maintains two counters for each agent. These are *totalOffers* and *numAcceptedOffers*, which keep track of the number of offers they made and those that got accepted, respectively. Before requesting a service, agents have to contact the broker agent to obtain a tag. The broker agent checks these counters of the agent requesting for a new tag.

If the *numAcceptedOffers* count exceeds some predetermined constant N for that level, an agent obtains a new tag with reduced length, i.e., an agent with tag 0101 and *numAcceptedOffers* = N acquires a new tag 010 and *numAcceptedOffers = 0*. Note that this counter is not observable by other agents. The main motivation for an agent to reduce its tag length is that doing so enables it to multicast its messages to a larger set of agents. It is also interesting to note that an agent being promoted to a higher level need not have offered fair deals only. Promotion simply means that more of its deals were accepted. The actual identity of the agent is not included in the message being multicast. Also the agent cannot distinguish between requests directed at it from a higher level agent or an agent from the same group.

As the number of agents in the lowermost level decreases, it becomes costly to offer unfair deals. Limited agent interactions between anonymous entities automatically gives rise to fairness as shown by our earlier experiments.

An agent whose ratio of *numAcceptedOffers* to *totalOffers* is low is demoted when it asks for a new tag. The demotion is done by randomly appending or prepending the agent's tag with a 0 or 1 bit. Agents with zero tag length broadcast their requests to all other agents in the community offering the service.

The broker agents could append or prepend or remove a bit from beginning or end depending on the number of agents in a particular group at that time. This helps in load balancing of the system.

4.2 Concluding Remarks

We have not experimented with the protocol described in Section 4.1. Interacting in small groups affects the supply-demand economics in the population. An agent cannot obtain good treatment (greater reliability in the storage space problem) if it has not done something good (offered good deals) in the community. There are many issues involved in the practical implementation of the protocol. We rely on a centralized tag-issuing broker agent that, depending on how many good offers an agent made, varies the group size to which the agent belongs.

There are other subtle issues related to fairness. Kahneman et al. [1986] present an interesting discussion on fairness. Game theorists and mathematical biologists experimenting with the ultimatum game model consider an equal split of a resource between the proposer and responder as indicating fairness [Sigmund et al., 2001, 2002; Page and Nowak, 2002; Guth et al., 1982]. Fairness depends on the preferences and values assigned to the goods by the user. In our experiments, we have considered users with two kinds of preferences as responders and making two kinds of offers as proposers. This has simplified the experiments, and has enabled us to isolate the factors responsible for fairness to evolve. However, it would be interesting to study the evolutionary dynamics with different combinations of user preferences and offers.

In a barter trade, fairness can vary with the changes in supply and demand for a particular good. We have not factored in social competition that may be present in a community. One could argue that in competition, better offers might be prevalent if the population is large. This would be true if proposers and responders were distinct in an anonymous setting. Because agents can play both roles, an agent as responder could combine the offers it receives to determine an estimate of the current fair market price. It could then propose that fair price to others. Further, competition makes sense if the goods can be substituted for others, which is possible only if they are somehow alike.

Bibliography

James E. Baker. Reducing bias and inefficiency in the selection algorithm. In *Proceedings of the 2nd International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 14–21. Lawrence Erlbaum Associates, 1987.

Brian F. Cooper and Hector Garcia-Molina. Bidding for storage space in a peer-to-peer data preservation system. In *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS)*, pages 372–384, 2002.

Chrysanthos Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 150–157. ACM Press, 2000.

Roger Dingledine. The free haven project: Design and deployment of an anonymous secure data haven. Master's thesis, MIT, June 2000. URL `http://www.freehaven.com/`.

Donald Ferguson, Christos Nikolau, and Yechiam Yemini. An economy for managing replicated data in autonomous decentralised systems. In *Proceedings of the International Symposium on Autonomous and Decentralised Systems*, 1993. URL `citeseer.nj.nec.com/ferguson93economy.html`.

Freenet. URL http://freenet.sourceforge.com/.

Gnutella. URL http://www.gnutella.com/.

Werner Guth, Rolf Schmittberger, and Bernd Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization*, pages 367–388, 1982.

David Hales and Bruce Edmonds. Evolving social rationality for MAS using "Tags". In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (to appear)*. ACM Press, 2003.

John Holland. The effect of labels (tags) on social interactions. In *Santa Fe Institute Working Paper 93-10-064*, 1993.

Nicholas R. Jennings and J. R. Campos. Towards a social level characterization of socially responsible agents. In *IEE Proceedings on Software Engineering*, pages 11–25, 1997.

Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Fairness and assumptions of economics. *Journal of Business*, 59(4):S285–S300, 1986.

Susanne Kalenka and Nicholas R. Jennings. Socially responsible decision making by autonomous agents. In *Cognition, Agency and Rationality*, pages 135–149, 1999.

Karen M. Page and Martin A. Nowak. Empathy leads to fairness. In *Bulletin of Mathematical Biology*, pages 1101–1116, 2002.

Hartmut Pohlheim. Genetic and evolutionary algorithm toolbox, July 1997. Available from URL http://www.geatbx.com/docu/index.html.

Stuart Russell. Rationality and intelligence. In *Artificial Intelligence*, number 94, pages 57–77, 1997.

Gerard Salton and Chris Buckley. Term-weighting approaches in automatic text retrieval. *Information Processing and Management*, 24(5):513–523, 1988.

Rina Schwartz and Sarit Kraus. Bidding mechanisms for data allocation in multi-agent environments. In *Agent Theories, Architectures, and Languages*, pages 61–75, 1997. URL `citeseer.nj.nec.com/schwartz97bidding.html`.

Sandip Sen. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *International Conference on Multi-Agent Systems*, pages 322–329, Menlo Park, CA, 1996. AAAI Press.

Sandip Sen and Partha Sarathi Dutta. The evolution and stability of cooperative traits. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1114–1120. ACM Press, 2002.

Karl Sigmund, Ernst Fehr, and Martin A. Nowak. The economics of fair play. *Scientific American*, pages 80–85, January 2002.

Karl Sigmund, Christoph Hauert, and Martin A. Nowak. Reward and punishment. In *Proceedings of the National Academy of Sciences*, pages 10757–10762, 2001.