Peter J. Deadman
Department of Geography
University of Waterloo
Waterloo, Ontario, N2L 3G1, Canada

fax: (519) 746-0658
e-mail: pjdeadma@fes.uwaterloo.ca

Conference Stream: Theory

Exploring Heterogeneity in Common Pool Resource Experiments with Intelligent Agent Based Simulations

< Working Draft >

Abstract

This work utilizes previously documented common pool resource experiments as a foundation for the construction of a series of computer simulations in which the individuals participating in the experiments are represented as separate intelligent agents. An intelligent agent is an autonomous, self-contained entity that resides within a virtual, computer-based, environment. In this study, agents are created to represent the individual participants in the CPR experiment and the resource that they share in common. By programming the agents with different strategies and endowments, the researcher can allow the agents to interact within a prespecified environment and observe the outcomes. These outcomes may include the performance of individual strategies in a specific environment, or the overall behavior of the group that emerges as a result of the numerous interactions of the individual agents. These models allow the researcher to observe the relative performance, at the individual and group level, of different combinations of individual strategies and to begin to draw connections between individual behaviors and group outcomes.

Group performance in heterogeneous simulations can vary significantly with minor changes in the initial parameters of the environment or the characteristics of the agents. Simulations which allow for simplified communication between agents show that lock-in can occur in which the agents agree on a group wide investment strategy which may or may not be an optimal solution. Some general discussion of the results of these simulations is provided, including a comparison with some observations from experimental economics and game theory. Preliminary observations on the advantages and disadvantages of agent based simulation as a tool for the analysis of the commons dilemma and issues related to heterogeneity are provided, along with some suggestions for future directions in which this work might proceed.

Introduction

Researchers have long been interested influence that heterogeneity in a common pool resource (CPR) management institution can have on the ability of resource users to forge management agreements. In an effort to build better theories, studies of heterogeneity in CPR dilemmas have utilized field studies and laboratory experiments with human subjects. These studies have revealed the complex nature of institutional behavior, and the difficulty of making the connection between individual behavior and group level outcomes.

Common pool resource management institutions may be classified as complex adaptive systems (Holland and Miller 1991). They consist of a network of interacting agents (elements, processes). They exhibit a dynamic aggregate behaviour that emerges as a result of the activities of the agents. This emergent behaviour can be described without detailed knowledge of the behaviour of the individual agents. Furthermore, the agents or individuals operating within this system are described as adaptive if they possess the following criteria: the outcome of the agents actions within its environment can be assigned a value such as utility or fitness; the agent behaves so as to increase this value over time. Agents may possess a number of different mechanisms for adjusting their actions in an effort to improve this fitness value. Complex adaptive systems usually operate far from the global optimum or attractor (Holland and Miller 1991). Indeed, experiments on human subjects in CPR dilemma situations have found that group level performance is frequently far from optimal. Depending upon the design of the system, many different levels of organization and interaction may exist. Agents operating within this system will seek to adapt so as to exploit the local niche to which they have access. This adaptation and evolution in turn creates new niches, or opportunities, to be explored. Such evolution can also result in lock-in, as agents adapt to the actions of other agents pursuing a collective course of action that leads the overall system in a particular direction which may or may not result in the system finding the predetermined global optimum.

The work described in this paper approaches the study of common pool resources from a social simulation perspective. A series of simulations have been developed that attempt to capture the essential elements of individual behavior in CPR dilemmas. The intention of this activity is to shed some light on the connections between individual behavior and group outcomes in CPR dilemmas. Unlike laboratory experiments or field studies, simulations allow the user to specify explicitly the behavior of individuals and the way those individuals interact.

Recent advances in computer hardware and software technologies have led to a growing effort to develop intelligent agent based simulations with resource management applications. Agent based modelling takes a bottom up approach to generating data comparable to that observable in the real system. This bottom up approach focuses on writing instructions to specify the behaviour of the individual component parts of the real world system that is being studied. No instructions specify the overall behaviour of the simulation. Instead, this overall behaviour emerges as a result of the actions and interactions of the individual agents. These intelligent agents contain two basic components, a model of their environment and a model of themselves (Zeigler 1976). The agent's model of itself contains the instructions for generating the behaviour of that agent under different circumstances. The model of the environment is generated by that agent's interactions with the world in which it exists. From this definition, we can observe that agents will seldom possess a complete model of their environment. Similarly, agents possess limited rationality in their ability to function within their environment.

According to Holland and Miller (1991), agent based models have several characteristics that are not available to traditional modelling techniques. Models based on linguistic descriptions, while infinitely flexible, often fail to be logically consistent. Mathematical models, while possessing consistent structure, suffer from reduced flexibility. However, agent based models that are specified by a particular computer language, retain much of the flexibility of linguistic models while having consistency and precision enforced on them by the language. As these models are executed in a simulation, the unfolding behaviour of the individuals and the system can be observed over time.

Through successive runs of these computer experiments, important parameters can be set to different values so as to study variations in outcome. The strategies of individual agents and the resultant outcomes can be carefully analyzed, something not usually possible with human experiments (Holland and Miller 1991). This study starts with previously documented CPR experiments and attempts to reproduce some elements of the behavior of individuals and groups noted in this work. Common pool resource experiments were selected because they are themselves models of the real world, and they have been widely studied.

Common Pool Resource Experiments

In an effort to understand the degree to which predictions about individual and group behaviour, and their resulting outcomes, derived from non-cooperative game theory are supported by empirical evidence, researchers have developed a series of laboratory experiments utilising human subjects. Built upon tight theoretical models of a CPR situation, laboratory experiments serve as a useful mechanism for simulating those models and observing outcomes. Laboratory experiments allow one to control some elements of a commons dilemma, thus facilitating analysis of the relationship between the structure of a resource management institution and the resultant outcomes.

In the baseline version of these experiments ,which are simulated here, eight subjects are presented with a situation in which they may choose to invest tokens in two markets. Market 1 provides a constant rate of return on the investment. Market 2 (the CPR) provides a return that varies in relation to the total group investment and the investment of the individual. The total return on the group's investment is determined by a quadratic production function that is concave in form. Through an unknown number of rounds, subjects invest in these markets in an effort to maximise their return. Subjects are aware of their profit in each round and their total profit for the experiment, as well as the total number of tokens invested in Market 2 by the group. Data is collected on the round by round investments of each participant.

Since the primary purpose of these baseline experiments is to evaluate the effectiveness of theoretical predictions, observers are interested in whether or not group investment in Market 2 will approximate a Nash equilibrium, a situation where every player maximises their return given what the other players are doing. The performance of the group is measured as rent as a percentage of optimum, or the return the group receives from Market 2, minus the opportunity costs of investing in market one, compared to the optimal level of investment. In the experiments that are simulated here, the Nash equilibrium level of investment would return rent as a percentage of optimum at 39 percent. Two sets of experiments were run in which participants were given equal endowments of either 10 or 25 tokens each. The optimum and Nash equilibrium levels of investment are not affected by the level of individual token endowments.

Researchers observed that in low-endowment experiments, where participants had 10 tokens each, group rent as a percentage of optimum averaged around 37 percent, whereas in high endowment experiments, group rent as a percentage of optimum averaged -3 percent. The poorer performance of the high endowment groups was attributed to the higher likelihood that individuals would invest too many tokens in Market 2, particularly at the beginning when they were attempting to discover the optimal investment level.

In actual runs of the low and high endowment experiments, researchers found little evidence to support the theoretical prediction that individual investments would settle out around the one-shot Nash equilibrium. Some individuals participating in the experiments did considerably better than others. However, in low endowment experiments, group level behaviour did approximate the Nash equilibrium. In high endowment experiments, group investments were initially far from the Nash equilibrium, but tended to approach it over time. In all experiments, group investment levels tended to follow a pulsing pattern as individuals adjusted their investments up and down in an effort to maximise returns.

The research suggests several factors as contributing to the inability of the subjects to achieve equilibrium levels. The complex nature of the return function and the additional complexity created by the investments of others meant that many individuals were unable to figure out the optimal level of investment, and instead relied on rules of thumb to guide there investments. Such a situation could result in the pulsing patterns observed and in the inability to recognize the equilibrium even if they happened upon it. In all, the issue of why individuals were unable to achieve the Nash equilibrium remains an important unanswered question (Ostrom et al. 1994).

Simulation Structure

The simulation system utilized in this study is Swarm, a multi-agent simulation platform developed at the Santa Fe Institute (Minar et al. 1996). This platform consists of a collection of object oriented libraries of reusable components for building models and analyzing, displaying, and controlling experiments on those models. Swarm adopts a modeling formalism that consists of a collection of autonomous agents, interacting via a time stepped series of discrete events. The basic unit of a Swarm simulation is an agent which is an entity that generates events that can effect itself and other agents (Minar et al. 1996). In a CPR model, individual agents can be created, with their own set of unique characteristics, to represent individual participants in an institution and the natural resource itself. For example, in a simulation of a fishery, agents could be created to represent each individual fisher inside or outside the institution, as well as other actors such as monitors, local authorities, and the stock of fish from which resources are withdrawn.

A typical swarm simulation consists of a hierarchical collection of agents, each capable of performing a given set of tasks that are specified in the methods that are written for its class. The simulations described here contain a number of agent classes which control the simulation and represent the component parts of the model. These classes, and the message passing connections between them, are specified in Figure 1.

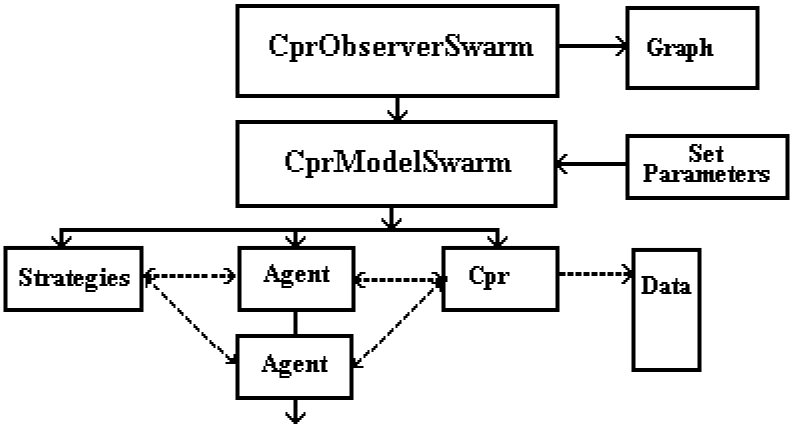

Figure 1: The objects and connections in the CPR model

At the lowest hierarchical level lie the classes which represent the components of the CPR experiment itself, the participants in the experiment and the CPR. In these simulations, two classes of models have been created to handle higher order functions, the CprModelSwarm and the CprObserverSwarm. During initialization, the CprModelSwarm utilizes the parameters that have been set by the user, namely the a, b, and w values, in the attached probe window to create the required number of instances of the necessary classes and to initialize certain parameters or variables within those instances. The CprModelSwarm also creates a total of eight appropriators, from the classes specified for each simulation, as specified by the user in the CprModelSwarm probe window. The CprModelSwarm is also where the schedule of agent actions is specified. This schedule is executed once in every round of the simulation to instruct specific agents in the simulation to perform certain methods in a certain order.

The CprObserverSwarm is the highest level of organization utilized in these simulations. The single instance of this class creates the instance of the CprModelSwarm and creates and controls the graphic devices upon which the output of the model is displayed. Included with this is a schedule that specifies how the graphic output devices are to be updated during the simulation. In this case, the CprObserverSwarm creates and updates a line graph showing group rent as a percentage of optimum for each round of the experiment as it is simulated. The CprObserverSwarm also creates the control panel that allows the user to start, stop, step, and quit the simulation.

The CPR Class

This class is designed to represent the actions of the common pool resource program installed on the NovaNet computer system for the laboratory experiments. The methods written for the CPR agent class specify the state of the CPR in relation to actions of the appropriators (see Figure 1). Only one instance of this class is created in these simulations. The quadratic production function for Market 2, as utilized by Ostrom et al. (1994), is embedded in the code of the CPR agent and specified as follows:

$$F\left(\sum x_i\right) = a\left(\sum x_i\right) - b\left(\sum x_i\right)^2 \tag{1}$$

Where $\sum x_i$ is the sum of all the Market 2 bids submitted by the agents. By manipulating the **a** and **b** parameters, the shape and magnitude of the quadratic production function can be controlled.

For all the simulations explored in this work, the **a**, **b**, and **w** parameters of the production function and Market 1 fixed return were set to 23, 0.25, and 0.05 respectively. During the initialization phase of the simulation, the CPR object calculates the optimum group bid as:

$$\sum x_i = (a - w)/2b \tag{2}$$

The Cpr object then calculates the subsequent group return from Market 2 yielded by the optimum group investment. During each round of the simulation, the CPR agent collects the token bids for Markets 1 and 2 from the individual agents. The CPR agent then calculates the total return from Market 2, group rent as a percentage of optimum for Market 2, and the return to each individual appropriator for that round. In addition, during each round of the experiment the Cpr object outputs the bids submitted by each participant, the cumulative rent earned by each participant, and the group rent as a percentage of optimum, to a data file for later analysis.

The Appropriator Classes

The individual appropriators that generate the behavior of these simulations are specified in a class called agents. Writing a set of methods for the individual appropriators is a challenging and potentially controversial endeavor that lies at the heart of any effort to model a common pool resource management institution. In this preliminary investigation, it requires the modeler to specify the mechanism that produces the strategies that individual subjects employ when they are engaged in a CPR experiment. How do subjects determine the bids that they will submit in each round of the experiment? What strategies do they employ? To what extent are they influenced by different factors such as the behavior of others? The greatest challenge, and potentially greatest benefit of this modeling exercise, lies in determining the relationship between individual behavior and the group level performance of the institution being studied.

The appropriator agents who participate in the experiments, do so with limited rationality. The nature of the behavior of complex systems such as economic institutions is such that it is difficult to fully understand the emergent properties that exist in these systems. As a result, we cannot program an intelligent agent with all the information it will require about the actions of the system at the beginning. Even with limited rationality, the agent must be able to learn to some extent which events are of semantic importance, and which strategies will prove most effective, as it goes along.

Induction theory approaches the study of limited rationality by assuming that humans utilize their pattern recognition capabilities to identify simple patterns that exist within the complex world, and construct a collection of models, or schemata, or hypotheses, upon which they act (Arthur 1994). They then utilize feedback from the environment to strengthen or discard them. As such, they learn to identify significant events and respond to them accordingly. In simulations that employ induction, this knowledge base is represented as a collection of rules (Holland et al. 1986). These rules have a condition-action structure of the form, "If such and such, Then so and so" (Holland et al. 1986). Individual entities in a modeling system possess a collection of these rules. Over time in a simulation, the individual proceeds through a cycle in which: 1) it matches incoming information against its set of rules to determine which ones have their conditions satisfied, 2) it selects a subset of these rules to be executed, 3) it enacts these rules, resulting in a specified behavior. In this way, the individual agent has a mechanism by which it can store and update knowledge about its environment and how it should act within it. This mechanism allows the agent to identify and respond to events in its environment which are of importance to the utility of the individual.

These rules can also be clustered together around concepts that are of importance to the individual agent. The agent may execute a combination of rules from different clusters, based on common conditions, in order to respond to environmental stimuli. Therefore, there may exist an implicit or explicit organization that links rules together (Holland et al. 1986). These categories of rules may also be organized in a hierarchical structure in which higher level clusters of rules with more general conditions may be linked to lower level clusters of rules with more specific conditions related to a specific environmental variable. This hierarchical structure of rule clusters nicely parallels the Institutional Analysis and Design framework (Ostrom et al 1994, Kiser and Ostrom 1982) which will be discussed at a later point.

In these simulations, a collection of sixteen possible sets of rules, or strategies, are specified and described in a new class, called the strategies class. Each appropriator agent is assigned a collection of strategies that it may draw upon to determine its bids in each round. At the beginning of each simulation, the agent is randomly assigned a specific number, either four, eight, or all sixteen, of these sixteen strategies which it may access from the strategies class. One of these strategies becomes the current strategy, the rest become alternates. The agents begin by playing the current strategy. But they also keep track of how the alternate strategies would have performed in each round had the agent used them. Every third round, the agents have the opportunity to switch their current strategy with best performing of the alternates. The performance of each alternate strategy is measured as the average return that the agent would have received in each round had it utilized that strategy. In this way, the agent simulates one of the mechanisms of an inductive process whereby it selects from alternative rules, based upon those rules relative strengths, in an effort to find the one with the best performance.

A pool of possible strategies were developed and specified in a separate class called strategies. Only one instance of this class is ever created. Six of the sixteen strategies are based on attempting to maximize the return received in each round. The increment or decrement of bids to Market 2 is varied in each of the strategies. Another six strategies are based on the comparison of average returns from Market 1 and Market 2, a strategy that was reported by some subjects in exit interviews during these experiments (Ostrom et al 1994). The final four strategies are based on a comparison of an individual agent’s bid with group average bid, and then submitting a bid at or above this average level.

Simple Communication Between Agents

A variety of CPR experiments have explored the role of communication in the commons. In experimental settings, communication has been found to increase the frequency with which individuals choose joint income maximizing strategies, even in situations where individual incentives conflict with the cooperative strategy (Ostrom and Walker 1991). In many laboratory experiments, communication opportunities are provided to the group of subjects between predetermined rounds of the experiment. These opportunities may be varied in terms of availability (one-shot versus repeated communication) or in terms of cost (costless versus costly communication). During these communication rounds, subjects are typically given ten minutes in which they may openly discuss the decision problem facing them. No restrictions are placed on these discussions, other than: they are not allowed to discuss side payments, they are not allowed to make physical threats, they are not allowed to see the private information on each others monitor (Ostrom et al. 1994). Communication was also found to improve the performance of groups in experiments where participants had heterogeneous endowments (Hackett et al 1994).

Compared to the relatively wide open communication of the laboratory experiments described by Ostrom et al. (1994), the agents in these simulations employ a very restricted form of information exchange. In the simulations that were developed, the agents were instructed to exchange information as to the bid that each individual felt would yield the highest return. This suggestion is developed by each agent during non-communication rounds of the simulation as the agent keeps track of the Market 2 bid that yields the highest return. Following the submission of these suggestions, they are evaluated in an effort to determine the bid that would yield the highest return. The suggestion yielding the highest apparent return is then selected for incorporation by the agents as an additional strategy. Therefore agents utilizing a pool of four strategies in the non-communication rounds prior to the communication round, adopt the winning suggestion as a fifth strategy for subsequent rounds. Initially this fifth strategy is adopted as the current strategy by each agent. But in later rounds, as in the non-communication induction simulations, the agents evaluate the performance of this new strategy against the alternates, and may switch to one of the alternates if it appears to provide a higher return.

Two alternative approaches to selecting the best of the suggested Market 2 bids are investigated in this set of simulations. Either the agents themselves evaluate the Market 2 bid suggestions to select the one that gives the best return relative to their current knowledge of the simulation, or the CPR itself evaluates all the bids made by the individual agents as if that bid was made uniformly by them all, selecting the one that gives the best group performance and sending that bid to each agent as its new strategy. In the agent based evaluation approach, a mechanism was added to each agent to allow it to evaluate the suggestions made by the other agents and adopt the suggestion that yields the highest return. The individual agents evaluate all the suggestions relative to their own information and current bids, selecting the suggestion that appears to work best for them. In the CPR based evaluation approach, each suggested bid is evaluated by the Cpr agent to determine the return that would be earned if each appropriator in the simulation had submitted that bid. The Cpr agent then selects the bid that results in the best overall group return and instructs the agents in the simulation to adopt that bid as the new group strategy.

Simulation Results

Non-Communication Simulations

Three sets of simulations were run, at 10 and 25 token allotments, in which agents were assigned four, eight, or all sixteen of the possible strategies. All of these simulations were characterized by group behavior that resulted in fluctuating patterns of rent as a percentage of optimum in which the oscillating patterns seen in the earlier fixed strategy simulations do not arise.

When the agents are endowed with 4 strategies and 10 tokens, group rent as a percentage of optimum fluctuates within a range of values between about -11 and 92 percent (see Figure 2). The same general fluctuating pattern is observed when the agents are provided with a 25 token endowment, although the range of values is greater (see Figure 2). Both the 10 and 25 token endowment simulations are characterized by occasional plunges in performance as the agents over invest in the CPR. These dips in performance are more noticeable in the 25 token endowment simulations because of the potential for enormous over investment in Market 2. Over time these large drops in rent as a percentage of optimum tend to disappear as the feedback mechanism in the induction process penalizes strategies which result in poor returns. Strategies which prompt the agent to invest all its tokens in Market 2 are the ones which cause these large drops. As a result, they tend to be selected by individual agents less frequently over time.

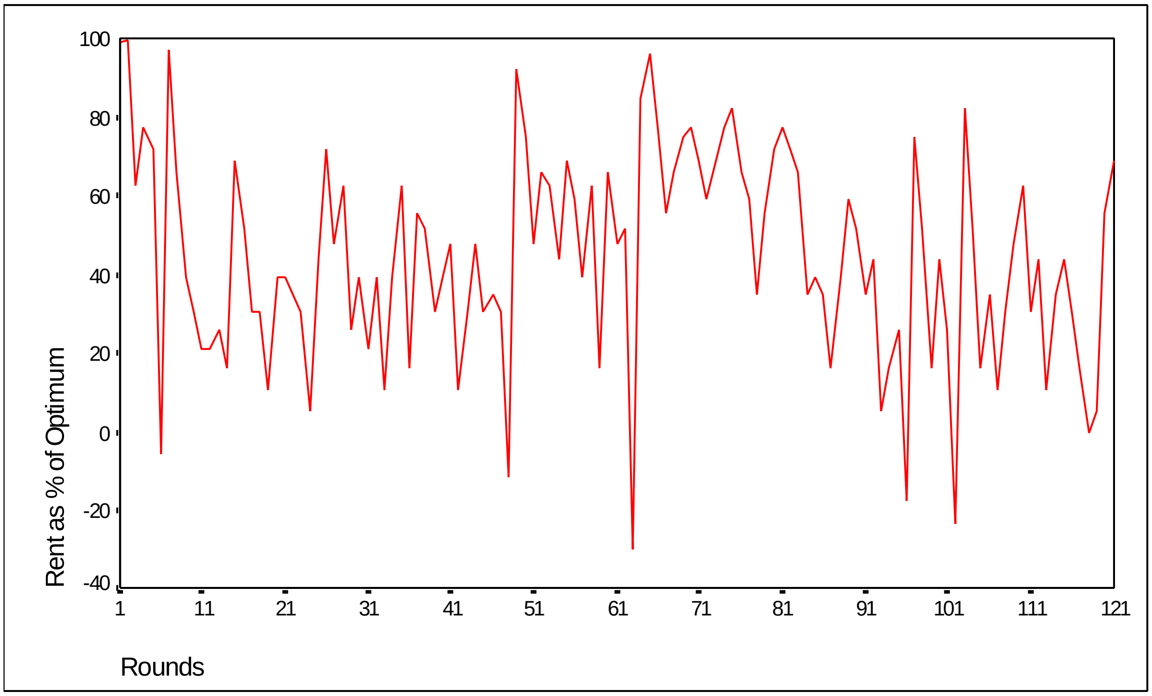

Figure 2: Group performance, non-communication simulation of agents with a 10 token endowment.

The number of strategies available to the agents in these simulations does not significantly alter the group performance. At a 10 token endowment, group performance expressed as rent as a percentage of optimum fluctuates around a mean in the mid 40 percent range. At a 25 token endowment, group performance fluctuates around a mean that is just below zero, about -10 percent. The behavior of these agents is similar to those of groups participating in CPR laboratory experiments as observed by Ostrom et al (1994).

Communication Simulations
Individual Evaluation of Best Group Bid

These simulations explore a communication routine in which the agents individually evaluate the suggestions submitted by the others in the group. The agents select the suggestion that would have given them the highest return, had they submitted that bid in the previous round. These simulations were run to determine if, and at what bid, the members of the group would lock into a uniform, group-wide Market 2 investment level. It should be noted that in these simulations, the common pool resource (Market 2) has a performance function with an optimum group bid of 36 tokens. If evenly distributed, this group bid represents an individual bid of 4.5 tokens from each member. In real world CPR experiments, this could be accomplished by the members of the group dividing themselves in half and agreeing to submit alternating bids of 4 or 5 tokens. However these simulations do not allow for this type of arrangement. Therefore, the closest the simulated agents can come to optimum, is to submit uniform group-wide bids of either 4 or 5 tokens. These group-wide levels of investment return a group rent as a percentage of optimum score of 98.76 percent.

The most important observation in this set of simulations is that the groups do eventually lock in to a uniform Market 2 bid. All agents eventually agree on a bid that results in the best return, given the events that have occurred previously in that particular simulation run. However, this group-wide uniform Market 2 bid frequently results in group performance levels that are sub-optimal. The amount of time required for the agents to lock into this group-wide uniform bids frequently exceeds 100 rounds, and can exceed 200 rounds. Typically, in these simulations, group performance fluctuates as it did in the non-communication simulations. However, unlike the non-communication simulations, eventually a constant group performance level will appear (see Figure 3). Examining the bids data file for the agents reveals that all the members of the group have settled on a uniform Market 2 bid. Furthermore, the size of the token endowment, or the number of strategies possessed by each agent does not appear to influence when the group will lock-in on a uniform Market 2 bid, or what that bid will be.

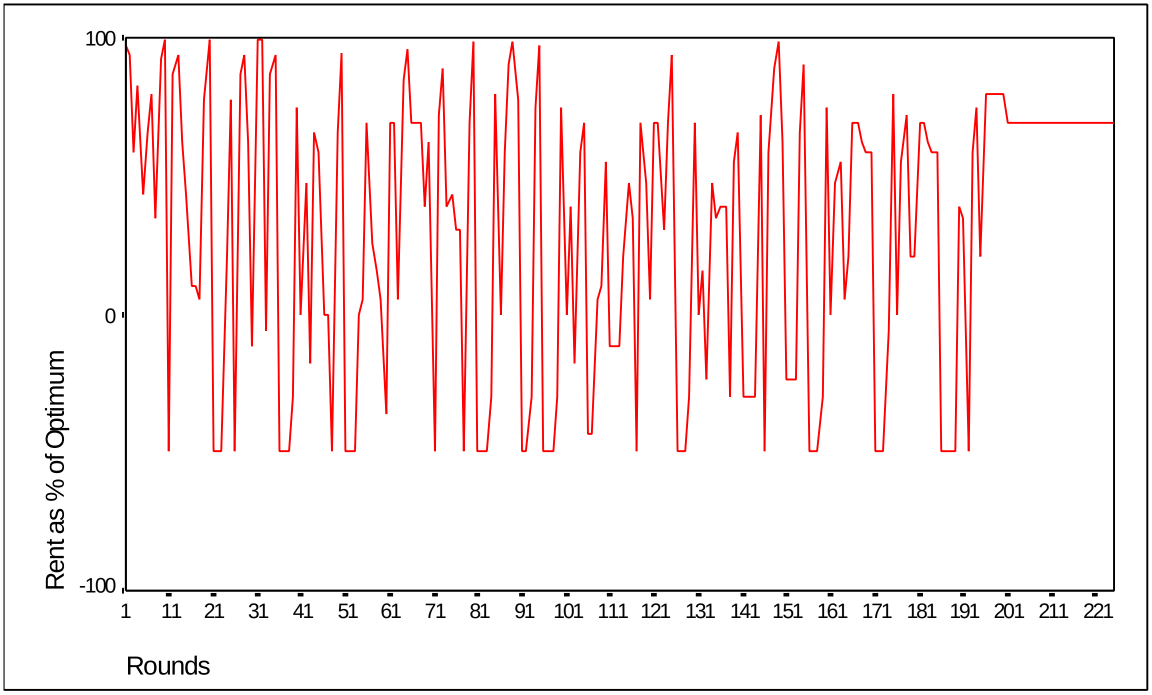

Figure 3: Group performance for 8 strategy agents and a 10 token endowment. The agents lock into a uniform Market 2 bid of 7 tokens each

Because the agents employ a feedback mechanism that evaluates the relative strength of each strategy based on the return that it earns for the agent, the performance of any particular strategy depends upon the actions of the other members of the group in each round. Therefore, a strategy that works well for one agent at a particular point in time, may result in a considerably poorer performance later in the simulation. In these simulations, the agents share their best Market 2 bids every fifth round, and adopt the best one as an additional strategy. During these communication rounds the majority of agents tend to adopt the same Market 2 bid, although this common bid varies from one communication round to the next. However, over time there is a tendency for the agents to converge on a uniform Market 2 bid. Eventually they all adopt a bid for which none of the agents has a better performing alternate, even if this group-wide bid is suboptimal. The actions of the group in previous rounds will determine which bid appears to yield the best performance.

Centralized Evaluation of Best Group Bid

In the previous set of simulations, each agent shared its best Market 2 bid with the group and evaluated the suggestions of others individually. In these simulations we report the behavior of simulations in which each communication round results in a uniform group bid. In these simulations, the Cpr agent evaluates each suggestion as if it had been submitted uniformly by all members, and then instructs the agents to adopt the best performing bid. Although each member adopts this bid as an additional strategy, individuals may choose to switch to an alternate strategy later in the simulation if the alternate appears to provide a higher return.

These simulations differ from the previous communication simulations in two important ways. First, the majority of the members of the group tend to lock into a uniform group-wide bid much earlier than in the previous simulations. This group wide bid is frequently at the near optimal level of 4 or 5 tokens each. Second, one or more members of the group switch away from the group strategy after a few rounds. This behavior results in the establishment of a fluctuating pattern of group performance in which the agents adopt a single group-wide Market 2 investment for the few rounds following communication, followed by a drop in performance as one or two members of the group switch to an alternate strategy (see Figure 4).

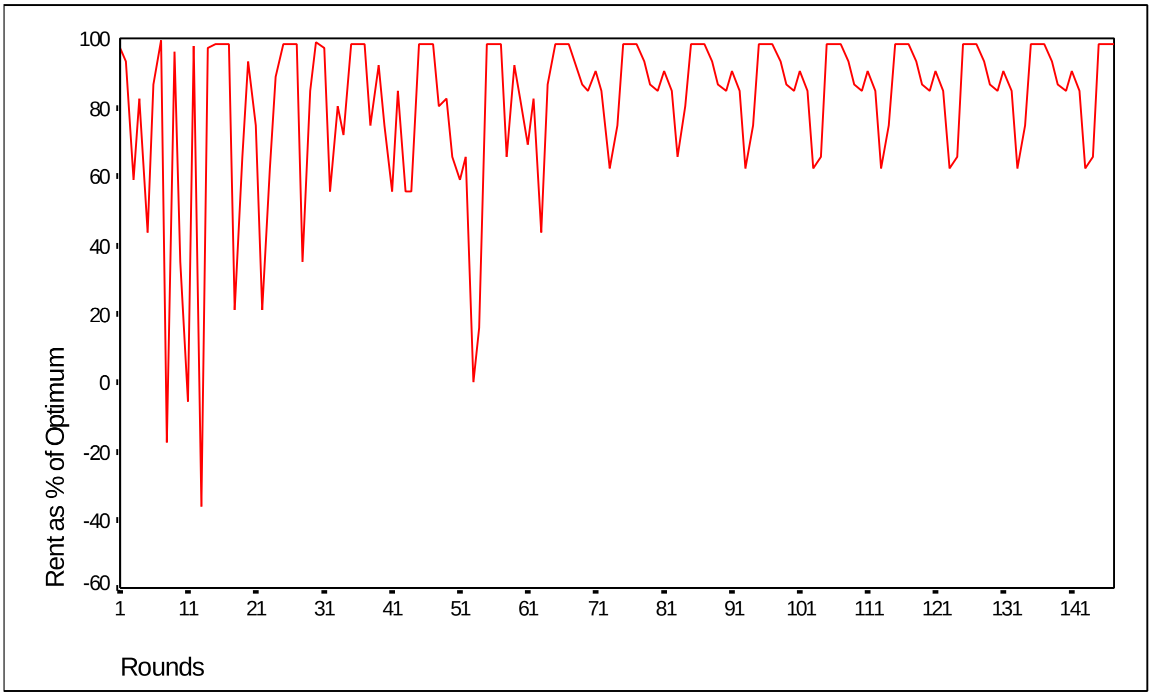

Figure 4: Group performance for communication simulation with centralized evaluation of best bid, 8 strategy agents and a 10 token endowment.

In the simulation depicted in Figure 4, the members of the group adopt the near-optimum investment level of 5 tokens each after each communication round. However, shortly thereafter four members of the group switch to an alternate strategy which appears to provide a higher return, thereby lowering group performance. The fluctuating pattern we see is the result of this cycle of group induced compliance at a communication round, followed by subsequent strategy changes by one or more members of the group. The number of strategies provided to each agent or the size of token endowment does not appear to be correlated with the group wide investment level, or the number of agents that will subsequently change strategies.

Discussion

All the agents employ a return maximising strategy here, with a variety of possible techniques for achieving that goal. The agents select from amongst different strategies based upon a score that is assigned to each strategy and calculated based upon the average return that the strategy achieved whenever it was used. Because each strategy is scored on the basis of its average return over time, the learning magnitude of this mechanism starts out very high for each strategy and decreases over time as each strategy is used. Returns earned by these strategies have a large impact on the average return early in the simulation, but decrease in relative importance over time.

The most interesting observation of the non-communication simulations is the fact that they perform similarly to groups of human subjects in CPR baseline laboratory experiments. As in CPR experiments, the group performance for the simulations follows an oscillating pattern in which high performance leads to over investment in the CPR and the resultant drop in performance causes a reduction in group wide investment in the CPR. In addition, the mechanism that allows agents to switch strategies is based on a goal of utility maximization. Agents will switch to another strategy if it appears to return a higher rent. Such a mechanism is likely to cause over investment, as agents seek higher returns from Market 2, followed by reduced investment, as the agents react to the reduced returns that the caused by over investment.

Still more interesting is the observation that the simulations perform similarly to subjects in lab experiments in terms of average performance over time. At the ten token endowment, the simulations perform near the Nash equilibrium over time. At the 25 token endowment, the simulations perform near zero percent of optimum over time. Has enough human rationality been captured in these agents to represent the actions of humans in this highly simplified environment? We know that some students in the lab experiments reported following a strategy similar to the unit return strategy described earlier. We know that students would attempt to maximise returns from one round to the next, and submit a variety of bids in an attempt to maximise utility. Perhaps in capturing these behavioural patterns, we have reproduced the essence of human behaviour in this simplified game. Although clearly no claim can be made that the agents are reproducing the thought processes of human beings, it appears that in such a simplified environment the simulations do a achieve a reasonable degree of replicative validity at the group level.

Heterogeneity in Communication Simulations

If non-communication lab experiments were reproduced with some validity in the previous simulations, no attempt is made in these simulations to reproduce communication lab experiments. Replicating the wide open nature of the communication process in CPR lab experiments, and the degree of complexity in potential human actions that is added to a CPR experiment by communication rounds is far beyond the capabilities of this current modeling effort. Instead, these simulations attempt to address the potential effects of increased agent interaction, in this case by sharing information, on overall group behavior and the tendency for common stable group-wide behaviors to emerge.

Simulations in which the agents evaluate the suggested bids of the other agents independently are characterized by eventual convergence of the group to a stable group wide uniform investment in Market 2. The emergence of this stable condition usually occurs somewhere between 100 and 250 rounds of the simulation. The length of time required to achieve tacit collusion is a product of the limited rationality of the agents and the mechanism they use to select from amongst the different strategies. Near the beginning of these simulations the agents suggest a wide variety of bids during the communication round. Frequently these bids are higher than the bids that would result in the group optimum level of investment because they were recorded when the agent submitted a bid higher than the group average. However, when all members of the group implement this bid, group performance drops and the bid is discarded. Over time, the agents continue to implement a variety of individual and group strategies, evaluating them as to their performance as they go along. Eventually, a bid is suggested and adopted by some members of the group that provides a return that is higher than the score of any alternate strategy. At this point the agent will continue to submit this communicated bid indefinitely.

Over a few communication rounds, this bid is adopted by more and more agents until all agents find that it performs better than any alternate strategy.

This simple communication mechanism changes the simulation significantly from the previous non-communication simulations, for it creates a simple self-reinforcing mechanism as discussed by Arthur (1988). According to Arthur, researchers have discovered that systems in many different fields of study, from theoretical biology to physics, tend to possess a multiplicity of asymptotic states, or "emergent structures". The initial configuration of the system, and some early, often random, events tend to push these dynamic systems into the domain of one of these asymptotic states, or attractors, and thus select a state that the system eventually "locks into" (Arthur 1988).

Arthur points out that such states exist in economic systems as well and cites examples from international trade theory, spatial economics, and industrial organization. The evolution of silicon valley in California is one such example from spatial economics. According to Arthur(1988), the self-reinforcing mechanisms that drive these systems to lock-in are derived from four generic sources: large set-up or fixed costs (which give the advantage of falling unit costs to increased output), learning effects (which act to improve products or lower their cost as their prevalence increases), coordination effects (which confer advantages to "going along" with other economic agents taking similar action), and adaptive expectations (where increased prevalence on the market enhances beliefs of further prevalence). Further, Arthur notes that these systems display four properties; multiple equilibria, possible inefficiencies, lock-in, and path dependence. The communication simulations in which the evaluation of best bid is conducted heterogeneously by each agent display these properties.

When suggested bids are evaluated by a central authority, the efficiency of the communication process is improved because the Cpr object has the ability to consider the group-wide implications of each bid. This reduces the tendency for the agents to adopt a bid following communication that results in over investment. As a result the group locks into a uniform bid, at or near the optimal levels of 4 or 5 tokens each, sooner. However, despite the near optimal level of the group bids following the communication rounds, the group is unable to maintain this arrangement and one or more members diverge from the prearranged bid in subsequent rounds. This results in the fluctuating pattern seen in figure 4 in which near optimal group level performance is repeatedly imposed during every communication round and subsequently declines as members of the group diverge from the imposed bid.

This form of imposed group level bid differs from the previous set of simulations, in which suggested bids were evaluated individually, in that the individuals do not come to a form of tacit agreement, but instead have bids imposed upon them. The agents have not ruled out their alternative strategies in these simulations. Those that defect do so because they find an alternative strategy that appears to yield a higher return. Also, because these bids are imposed early in the simulation, the learning magnitude of all alternate strategies is still relatively high. Strategies which appear to return a higher yield than the imposed bid have their relative strength reinforced in comparison to the near optimal imposed bid. As a result the agent will continue to turn to that strategy after each subsequent communication round.

The agents in these simulations possess no mechanism to represent a social norm favoring cooperation. Therefore they are not encouraged to cooperate with the imposed group bid by any mechanism other than an objective evaluation of the potential earnings that may be earned by their alternate internal strategies. Therefor if it appears that an alternate strategy will yield a higher return in the next round, they switch away from the group strategy without any consideration of the potential actions of the other agents. Future simulations may allow for the threat of sanctions, if the individual diverges from a group strategy.

Agent based modelling provides an interesting new opportunity for the study of common pool resource problems because it permits theories of individual action and group behaviour to be tested through computer simulations. A considerable body of work has been published in the learning and game theory literature that examines individual action in two player and N-player prisoner's dilemma games. This theory has some general applicability to CPR dilemmas, however the different structure of CPR games will require some modifications to existing theory. Mathematical solutions to these more complicated situations may be non-trivial.

Clearly the simple nature of these early simulations leaves a great deal of avenues open to further investigation. Some future directions in which this work might proceed include the further development of the limited rationality learning model discussed here or the exploration of alternative models. Alternative communication routines could also be explored in an effort to more closely replicate the communication procedures followed in the lab experiments.

This initial effort has addressed the potential of intelligent agent based modelling and simulation systems for understanding more about common pool resource management institutions. This work has focused on modelling CPR laboratory experiments as a prelude to the development of other resource management or institutional models. But eventually we will wish to extend these models to link human systems and natural systems models in resource management applications. Some examples of this already exist. Simulations such as Phoenix (Cohen et al. 1989) combine dynamic models of natural processes (in this case forest fire spread) with dynamic models of human action (the movement of the fire fighters and equipment).

Other models could be developed to explore the interrelationship between human resource management activities and dynamic changes in natural systems. Such resource management models may find applications in water resources management, fisheries management, and forest management. A water resources management model could tie a sub-model of a natural system such as a watershed and its associated aquifers to a model of human action under different institutional configurations. A similar model of a fishery could also be constructed in which human agents develop and modify the rules governing the management of the natural resource, under differing natural and institutional conditions. This relationship between the natural system and the human institution that interacts with it has been termed system-dependant selection (Lansing et al 1998). System-dependent selection considers what can happen when feedback from the environment affects the behaviour of an individual or human institution, and visa versa. This feedback loop results in some form of functional organisation emerging at the level of both systems as a whole.

In these models, agents representing human individuals or organisations will have to deal with constantly changing conditions in the natural system. The equivalent effect in the simulations developed for this paper would have been if the quadratic production function built into the Cpr had changed over time. In addition these models will have to capture the essential components and actions of resource management institutions. Although some theoretical tools exist, such as the IAD framework and A grammar of institutions, to assist in this effort, such models will be considerably more complex than the ones discussed here. However, if these challenges can be met in a series of incremental efforts, there is considerable potential for modelling and simulation as a tool for assisting us in our understanding of these systems.

References

Arthur, B. W. (1994). Inductive Reasoning and Bounded Rationality. *American Economic Review*. 84(2), 406-411.

Hackett, S., Schlager, E., Walker, J. (1994). The Role of Communication in Resolving Commons Dilemmas: Experimental Evidence with Heterogeneous Appropriators. *Journal of Environmental Economics and Management*. 27, 99-126.

Holland, J.H., Holyoak, K.J., Nisbett, R.E., Thagard, P.R. (1986). *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, Ma.

Holland, J.H. and Miller, J. H. (1991). Artificial Adaptive Agent in Economic Theory. *American Economic Review*, 81(2):365-370.

Kiser, L.L., and Ostrom, E. (1982). The Three Worlds of Action: A Metatheoretical Synthesis of Institutional Approaches. *in*: *Strategies of Political Inquiry*. ed. by E. Ostrom. Sage Press. pp. 179-222.

Lansing, J.S., Kremer J. N., Smuts B. B. (1998). *System-Dependant Selection, Ecological Feedback, and the Emergence of Functional Structure in Ecosystems*. Santa Fe Institute Working Paper 98-01-014.

Minar, N., Burkhard R., Langton C., Askenazi M. (1996). *The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations*. Overview paper. Santa Fe Institute, Santa Fe, NM.

Ostrom E., Gardner R., Walker J. (1994). *Rules, Games, & Common Pool Resources.* The University of Michigan Press.

Zeigler, B.P. (1976). *Theory of Modelling and Simulation*. John Wiley, New York.