

**CAN GENETIC ALGORITHMS
EXPLAIN EXPERIMENTAL ANOMALIES?
AN APPLICATION TO COMMON PROPERTY RESOURCES**

Marco Casari¹

Universitat Autònoma de Barcelona

First draft: October 2002

This version: April 2003

UFAE and IAE Working Paper number 542.02

Abstract. It is common to find in experimental data persistent oscillations in the aggregate outcomes and high levels of heterogeneity in individual behavior. Furthermore, it is not unusual to find significant deviations from aggregate Nash equilibrium predictions. In this paper, we employ an evolutionary model with boundedly rational agents to explain these findings. We use data from common property resource experiments (Casari and Plott, 2003). Instead of positing individual-specific utility functions, we model decision makers as selfish and identical. Agent interaction is simulated using an individual learning genetic algorithm, where agents have constraints in their working memory, a limited ability to maximize, and experiment with new strategies. We show that the model replicates most of the patterns that can be found in common property resource experiments.

Keywords: Bounded rationality, Experiments, Common-pool resources, Genetic algorithms

JEL Classification numbers: C72, C63, C91, Q2

¹ Correspondence address: Marco Casari, Departament d'Economia i d'Historia Econòmica, CODE, Edifici B, Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain, email: mcasari@pareto.uab.es, tel: ++34.93.581 4068, fax: ++34.93.581 2461.

The paper has benefited from comments of Jasmina Arifovic, Simon Wilkie, Charles Plott, Nelson Mark, Sean Gailmard, Guillaume Frechette, Ben Klemens and of participants at seminars at Ohio State University, the 5th Workshop in Experimental Economics in Siena, Italy, at the ESA meeting in San Diego, CA, USA, University of Guelph, Canada, and the Simposio de Anàlisi Econòmica in Salamanca, Spain. Sharyn Slavin Miller, Maria Satterwhite, and Eloisa Imel from Caltech provided technical support. Financial support from the Division of the Humanities and Social Sciences at Caltech and the EU Marie Curie Fellowship is gratefully acknowledged.

CAN GENETIC ALGORITHMS EXPLAIN EXPERIMENTAL ANOMALIES?

AN APPLICATION TO COMMON PROPERTY RESOURCES

1 Introduction

Even in simple games with a unique equilibrium, experimental results often exhibit patterns inconsistent with the predictions of perfectly rational and selfish agents. It is not unusual to find patterns of heterogeneity in individual behavior when there is a symmetric equilibrium, oscillations in the aggregate outcome, significant differences between inexperienced and experienced players, or systematic deviations from the predicted equilibrium (Kagel and Roth, 1995). In this paper, we employ a model of adaptive learning, based on a genetic algorithm, to explain the results from a common property resource experiment, which, to some degree, exhibits all the mentioned patterns.

Two routes can be followed to explain the above patterns in experimental data. One is to differentiate the goal of the agents from pure personal income maximization to include varying degrees of other-regarding preference. The other route, followed in this paper, is to weaken the perfect rationality assumption. More specifically, we use a model of adaptive learning agents with a limited working memory, inability to maximize, and active experimentation with new strategies. All agents have an identical, although bounded, level of rationality.

Genetic algorithms were first developed by Holland (1975) as stochastic search algorithms by looking at the biological processes of evolution. They have been employed to explain a variety of experimental data, including data from auctions (Andreoni and Miller, 1995, Dawid, 1999), oligopolies (Arifovic, 1994), foreign currency markets (Arifovic, 1996), and Groves mechanisms (Arifovic and Ledyard, 2000). Experimental data offer an attractive test

bed for models of bounded rationality because they present decision-makers with a well-defined environment where decisions are made repeatedly.

In this paper, we focus on common property resource experiments with an emphasis not only on the qualitative findings from human subjects but on the ability of the genetic algorithm to match their quantitative levels as well. There are two main innovative features. One is the study of individual behavior. To the best of our knowledge, no previous study has compared the individual behavior of genetic algorithms with experimental human data. Similar aggregate results can hide a wide diversity in individual actions. The other innovative aspect has to do with analyses of the experimentation process with new strategies. The experimentation process is not simply an additional element of randomness but interacts at a deeper level with the limited cognitive abilities of the agents.

In Section 2, we outline the experimental design and results. In the following Section, we describe the artificial adaptive agents. In Section 4, we present the results of the simulations in reference to the level and variability of aggregate resource use as well as individual heterogeneity. We conclude in Section 5.

2 Experimental design and evidence

This Section first describes the incentive structure of the experiment and then outlines the results. A more detailed description of them can be found in Casari and Plott (2003).

Consider a group of agents $i=1, \dots, 8$. Each agent decides on an effort level $x_i \in [0, 50]$ of a common property resource. An agent i 's payoff function is:

$$\pi_i = \frac{x_i}{X} f(X) - c(x_i) \tag{1}$$

where $c(x_i)=2.5 \cdot x_i$ is the cost of the effort, $X=\sum_1^N x_i$ is the group effort, and $f(X)$ is the group

revenue. Group revenues are shared according to the relative effort $\frac{x_i}{X}$ of each individual.

The function $f(X)$ is continuous in \mathbb{R}^+ , increasing in $X \in [0, 92]$, decreasing for $X > 92$, and with a lower bound at -200 :

$$f(X) = \begin{cases} \frac{23}{2}X - \frac{1}{16}X^2, & \text{if } X \leq 184 \\ 200 \cdot [e^{-0.0575(X-184)} - 1], & \text{if } X > 184 \end{cases} \quad (2)$$

From the first-order conditions to maximize earnings $\frac{\partial \pi_i}{\partial x_i} = 0$, one can derive the best response functions $x_i^* = 72 - \frac{1}{2}X_{-i}$, where $X_{-i} = \sum_{j \neq i}^N x_j$. The Nash equilibrium is unique and

symmetric and leads to an aggregate outcome of $X^*=128$ and an individual outcome of $x_i=16 \forall i$. Group profits at the Nash equilibrium are just 39.5% of the potential profits (128/324). This result is standard in the renewable resource literature (Clark, 1990).

Common-pool resource appropriation is very similar to a Cournot oligopoly when x_i is interpreted as the quantity produced and $f(X)$ as the aggregate market profits. As in the adopted design the users of the resource are more than two, a richer set of individual behaviors may be generated. Such individual behavior has been reported in detail in Casari and Plott (2003).

Four sessions of 32 periods were run. Agents face the same incentive structure for the length of a session. No communication was allowed among subjects and at the end of each period they could observe the aggregate outcome but not the individual choices of others. The experimental results are summarized below in three points relating to aggregate resource use, variability in aggregate resource use, and individual heterogeneity, respectively:

- (a) Agents cooperate less than the Nash equilibrium (use the resource more than Nash equilibrium). Average resource use efficiency is 28.4%, which is statistically different than the predicted 39.5% ($p=0.05$).
- (b) Group use fluctuates over time (pulsing patterns). The average standard deviation of group use over time within a session is 12.95 with an average resource use of 131.32. An interval of one standard deviation around the average corresponds to an efficiency range of [0.0%, 58.5%].
- (c) Individual behavior is persistently heterogeneous. For instance, the difference between the average use of the agent who used the resource the most and the average use of the agent who used the resource the least within each session, $[\max_i \{ \bar{x}_i \} - \min_i \{ \bar{x}_i \}] = 28.35$ out of a potential maximum of 50 and a predicted value of 0.

Similar findings in a common property resource environment are documented also by Rocco and Warglien (1996), and Walker, Gardner, and Ostrom (1990). We will compare the simulation results from genetic algorithms with the above results from human subjects.²

3 The artificial adaptive agents

Genetic algorithm (GA) agents interact in the environment that was described in the previous Section. While this Section introduces the GA decision makers along with the parameter values used in the simulations, a full description of the working of a genetic algorithm is given in Holland (1975), Goldberg (1989), Bäck (1996), and Mitchell (1996). For issues specific to Economics see the excellent study of Dawid (1996).

² Other six sessions were run under an experimental design with sanctions, where agents first decided a level of resource use and then had the option to monitor other users and sanction those who exceeded a given threshold of resource use (i.e. free riders). In one sanction treatment the cooperation level is above the Nash equilibrium level (opposite than (a)). In all treatments (b) and (c) are observed. The experimental designs and results are reported in Casari and Plott (2003).

The genetic algorithm decision maker can be described as follow. A strategy is identified by a single real number. It is encoded as a binary string, a so-called chromosome, and has associated with it a score (measure of fitness) that derives from the actual or potential payoff from this strategy. In a social learning (single-population) basic GA, each agent has just one strategy (chromosome) available, which may change from one period to the next. In an individual learning (multi-population) algorithm, which is the version adopted in this study, each agent is endowed with a set of strategies, and each set may change independently from other sets from one period to the next. The changes are governed by three probabilistic operators: a reinforcement rule (selection), which tends to eliminate strategies with lower score and replicate more copies of the better performing ones; crossover, which combines new strategies from the existing ones; and mutation, which may randomly modify strategies. In a basic GA, the strategies (chromosomes) created by crossover and mutation are directly included in the next period's set of strategies (population).

The three operators are stylized devices that are meant to capture elements involved in human learning when agents interact. The reinforcement rule (selection) represents evolutionary pressure that induces agents to discard bad strategies and imitate good strategies; crossover represents the creation of new strategies and the exchange of information; mutation can bring new strategies into a range that has not been considered by the agents.

Most of the parameters of the genetic algorithm were chosen exogenously, based on considerations external to the data here analyzed and not based on fit improvement considerations. On the contrary, the next Section will discuss the two free parameters, mutation and crossover rates.

The description of the exogenous features of the genetic algorithm begins with the reinforcement rule. GA agents are adaptive learners in the sense that successful strategies are

reinforced. Strategies that perform well over time gradually replace poor-performance ones. The most common reinforcement rules in the GA literature are pairwise tournament and biased roulette wheel. We have adopted a pairwise tournament for two reasons. First, it is ordinal, in the sense that the probabilities are based only on “greater than” comparisons among strategy payoffs and the absolute magnitude of payoffs is not important for the reinforcement probability. Being ordinal it does not rely on a “biological” interpretation of the score as a perfect measure of the relative advantage of one strategy over another. As a consequence, the simulation results are robust to any strictly increasing payoff transformation. Second, while in a biased roulette wheel the payoff needs to be positive that is not the case for pairwise tournament. The reinforcement operates by (1) randomly drawn with replacement two strategies, a_{ikt} and a_{iqt} , from a population A_{it} and by (2) keeping for the following interaction only the strategy with the highest payoff in the pair: $a^*_{it} = \operatorname{argmax}\{\pi(a_{ikt}), \pi(a_{iqt})\}$. After each period, these two steps are repeated K times, where K is the population size.

Simulations are run with an individual learning GA, which is discussed in the remainder of this Section. When agents do not consider just one strategy at each period in time, but have a finite collection of strategies from which one is chosen in every period (memory set), the process is called a multi-population GA (Riechman, 1999, Vriend, 2000, Arifovic and Ledyard, 2000). A strategy is a real number $a_{ikt} \in [0, 50]$ that represents the appropriating effort level of agent i in period t . Each agent is endowed with an individual memory set $A_{it} = \{a_{i1t}, \dots, a_{iKt}\}$ composed of a number of strategies K that is constant over time and exogenously given. If a strategy a_{ikt} is in the *memory set*, i.e. it is available, agent i can choose it for play at time t . The individual learning Ga was here adopted because it reproduces the informational conditions of the experiment while the social learning GA does not. Moreover, it is better suited to study individual behavior as in a social learning GA

identifying the evolution of an agent over time is problematic. In the laboratory, an agent could learn from her own experience but not from the experience of others. In fact, an agent could not even observe, let alone copy, the strategy played by others.

The size of the memory set, K , is a measure of the level of sophistication of an agent since it determines how many strategies an agent can simultaneously evaluate and remember. The Psychology literature has pointed out that the working memory has severe limitations in the quantity of information that it can store and process. According to these findings, the memory limitation is not just imperfect recall from one round to the next, but rather an inability to maintain an unlimited amount of information in memory during cognitive processing (Miller, 1956; Daily et al., 2001). The classic article by Miller (1956) stresses the “magic number seven” as the typical number of units in people’s working memory. As the memory set size K needs to be even, both 6 and 8 are viable options. We set $K=6$, which implies that decision-makers have a hardwired limitation in processing information at 6 strategies at a time.

As each agent is endowed with a memory set, in the individual learning GA (multi-population) there is an additional issue of how to choose a strategy to play out of the K available. This task is performed by a stochastic operator that we will call *choice rule*. The *choice rule* works in a very similar way as the reinforcement rule, i.e. as a one-time pairwise tournament, where (1) two strategies, a_{ikt} and a_{iqt} , are randomly drawn with replacement from the memory set A_{it} and (2) the strategy with the highest score in the pair is chosen to be played: $a^*_{it} = \operatorname{argmax} \{ \pi(a_{ikt}), \pi(a_{iqt}) \}$. A pairwise tournament is different from deterministic maximization, because the best strategy in the memory set is picked with a probability less than one. The choice rule, however, is characterized by a probabilistic response that favors high-score over low-score available strategies. In particular, the probability of choosing a strategy is strictly increasing in its ranking within the memory set. The stochastic element in

the choice captures the imperfect ability to find an optimum, where the probability of a mistake is related to its cost.³

To sum up, this Section has described the genetic algorithm employed in the simulations and motivated the adoption of a pairwise tournament reinforcement rule and of the individual learning design. Within the individual learning design, we discussed the assumed memory size of six strategies for each agent and of a pairwise tournament choice rule.⁴

4 Simulation results with genetic algorithm agents

In this Section, we present the result of the interaction among genetic algorithm agents in a common property resource environment and compare them with the human agent data from the experiment. Extensions to some other experimental designs are also discussed.⁵ Before presenting the analysis of fit, we discuss the choice of some parameter values.

Parameter values. Genetic algorithm agents constantly search for better strategies through active, random experimentation that changes the composition of the memory set. Experimentation is characterized by a level, p , which is the expected share of strategies in the memory set that will randomly change from one period to the next. The value of p is chosen in order to increase the fit between the human data and the simulation results and is set in the following way. First, the strategy space is divided into a grid and coded with binary strings of 0s and 1s of length L . Second, with probability $pm \in (0,1)$ that each digit '0' can flip to '1' or

³ The score of a strategy can be interpreted as the utility of the outcome associated with that strategy. Given the ordinality of pairwise tournaments adopted for reinforcement and choice rule, this GA is based only on the ordinal information of the score, like the utility function of the consumer.

⁴ A score is assigned to every strategy in the memory set, whether the strategy was chosen to be played or not. The score of strategy not chosen to play was assigned under the assumption that all the other agents will not change their actions in the following period (adaptive expectations).

⁵ Simulations with the same GA were run also in common property resource designs with sanctions. The results are reported in Casari (2002).

vice versa. This mutation procedure is quite standard in the GA literature. For a mutation rate pm , the corresponding experimentation level is $p=1-(1-pm)^L$, where L is the number of digits of the binary string. In the simulations we adopt a mutation rate $pm=0.02$ with $L=8$ that corresponds to an expected fraction of new strategies due to experimentation $p=0.1492$ of the total in the memory set. The range of values used in the GA literature is quite wide, and our experimentation level does not appear particularly elevated. Consider for example the following four studies: Arifovic (1996) uses two sets of parameters, $L=30$ with $pm=0.0033$, or $pm=0.033$, which translates into $p=0.0944$ or $p=0.6346$, respectively; Andreoni and Miller (1995), $L=10$, $pm=0.08$ with exponential decay and half-life of 250 generations, which translates into $p=0.5656$ for the first period of the simulation and $p=0.0489$ for period 1000; Bullard and Duffy (1998), $L=21$ with $pm=0.048$: $p=0.6441$; and Nowak and Sigmund (1998) a direct experimentation rate of $p=0.001$.

As noted, the parameter L influences the experimentation rate. Its level was set at $L=8$ before running the simulations in order to establish a reasonably thin grid of the strategy space, and then was maintained constant throughout. The strategy space $[0,50]$ is divided into a grid of 255 points (2^8-1), which corresponds to steps of about 0.2 units. In the experiment with human agents any real number could be chosen. However, in practice, 87% of the actions inputted were integer numbers. The grid chosen can accommodate the level of accuracy in decision making of the laboratory data.

After mutation rate and string length, the third parameter that will be discussed in this Section is the crossover rate. The crossover operator works in two steps: first, it randomly selects two strategies out of a population; second, selects at random an integer number w from $[1, L-1]$. Two new strategies are formed by swapping the portion of the binary string to the right of the position w . In general, not all strategies in the population are recombined using the crossover operator; instead crossover is carried out with some probability, pc ,

which is the crossover rate. Simulations in a common property environment that are not here reported show a rather small influence of crossover on the results. Hence, we decided to set the crossover rate to zero, $pc=0$, and adjust only the mutation rate.

Results. The results of the simulations of resource use with genetic algorithm agents are now presented. Genetic algorithm agents replicate cooperation levels of humans (Result 1), the pulsing patterns (Result 2), and to a large extent individual heterogeneity (Result 3).⁶

The numerical results presented are averages over 100 simulations run with different random seeds 0.005 through 0.995. There are three different lengths T of the simulations in order to mimic the behavior of inexperienced agents ($T=32$, as the actual length of a laboratory session was 32 periods), experienced agents ($T=64$), which have already acquired one session of experience, and of long term behavior ($T=400$).⁷ In all cases, the numerical results presented in Table 1 refers just to the last 32 periods of the simulation and ignore the previous periods. For instance, the aggregate resource use reported when $T=64$, is the average of periods from 33 to 64. The reason of this choice is to be able to perform an homogenous comparison with human agent data, where the length of an experimental session is always of 32 periods.

Result 1 (Aggregate resource use)

The aggregate resource use X of genetic algorithm agents (GAs) is not statistically different from humans agents's levels. In both cases, agents cooperate less than the Nash equilibrium level.

The aggregate level or resource use of the GA agents (X_{GA}) closely matches the experimental results ($X_H=131.32$). For inexperienced GA ($T=32$), the cooperation level

⁶ The simulations were run on a PC and the GA agents were programmed in Turbo Pascal. Useful references for the code were Goldberg(1989) and a version given by Jasmina Arifovic.

⁷ Simulation longer than 400 periods were performed (up to 10,000 periods) but do not change the conclusions about long term behavior of GA.

$X_{GA}=131.02$ cannot statistically be distinguished from the human value at a 0.05 level. Similarly for experienced GA, $X_{GA}=130.40$ and long term $X_{GA}=130.02$. (Table 1, columns (3), (4), and (5)).

Result 2 (Variability of aggregate resource use)

Genetic algorithm agents (GAs) exhibit a higher variability over time in aggregate resource use $\sigma(X)$ than human agents; such variability, however, decreases with experience.

When inexperienced GA agents interact ($T=32$), the variability in aggregate group use as measured by the standard deviation of resource appropriation over time is $\sigma(X)_{GA}=17.50$ versus $\sigma(X)_H=12.9$ with humans. With experience the variability decreases to $\sigma(X)_{GA}=15.03$ and $\sigma(X)_{GA}=14.04$. An alternative measure of variability of aggregate resource use is the percentages of periods in which aggregate payoffs are negative. For GA agents this statistics goes from 19.59% ($T=32$), to 16.00% ($T=64$), to 14.97% ($T=400$) while it is 15.5% for human agents (Table 1). A visual comparison between GA agents and human agents is offered by Figure 1. The pattern for GA agents in Figure 1 is an example of four random runs.

The same level of aggregate variability can hide widely different patterns of individual variability. Before proceeding to outline Result 3, an example is presented to introduce the precise definition of individual heterogeneity adopted throughout the paper. Consider scenarios A and B in Table 2 with two players and four periods.

Table 2: Examples of two patterns of individual variability

| Scenari o | Agent | Period | | | | <i>Agent average</i> \bar{x}_i | Indexes of variability of individual actions | | | | |
|--------------|-------|--------|----|----|----|---|---|----------------|------------------------|-------------------------|---------------------|
| | | 1 | 2 | 3 | 4 | | Overall D1 | Overall SD1 | Across agents D2 | Across agents SD2 | Over time SD3 |
| A | x_1 | 12 | 12 | 12 | 12 | 12 | 10 | 5.35 | 10 | 7.07 | 0 |
| | x_2 | 22 | 22 | 22 | 22 | 22 | | | | | |

| | | | | | | | | | | | |
|---|----------------|----|----|----|----|----|----|------|---|---|------|
| B | x ₁ | 12 | 22 | 12 | 22 | 17 | 10 | 5.35 | 0 | 0 | 5.77 |
| | x ₂ | 22 | 12 | 22 | 12 | 17 | | | | | |

Note: D=difference between maximum and minimum, SD=standard deviation

The two scenarios are identical when considering both aggregate production $X_t = \sum_i x_{it}$ and overall indexes of variability of individual actions, such as the mean of the difference, period by period, between the maximum and minimum individual productions,

$$D1 = \frac{1}{T} \sum_{t=1}^T \max_i \{x_{it}\} - \min_i \{x_{it}\},$$

or the standard deviation of individual actions x_{it} (SD1). The differences in the patterns of individual variability between scenario A and B can be captured by splitting the overall individual variability into variability across agents (D2 and SD2) and over time (SD3). In order to calculate agent-specific variability, first we compute the average

$$\text{individual production over time } \bar{x}_i = \frac{1}{T} \sum_{t=1}^T x_{it} \text{ and, using those data, compute the difference}$$

$$D2 = \max_i \{\bar{x}_i\} - \min_i \{\bar{x}_i\} \text{ and the standard deviation for } \bar{x}_i \text{ (SD2) (Table 2). Scenario A rates}$$

highly in terms of variability across agents, and that is referred to here as high individual heterogeneity, while scenario B rates highly in terms of variability over time but exhibits no individual heterogeneity.

When the same statistics developed for the example in Table 2 are applied to the simulation results (Table 1), a remarkable level of individual heterogeneity emerges from the interaction of ex-ante identical genetic algorithm agents (Result 3).

Result 3 (Individual heterogeneity with resource use)

Identical genetic algorithm agents (GAs) use the resource at significantly different rates.

Depending on the level of experience and of the measure adopted, between 45% and 80% of the human individual heterogeneity is reproduced by GA agents. In particular, inexperienced GA agents have an heterogeneity levels in resource use SD2 not statistically different from human agents.

Individual heterogeneity can be measured either with SD2 or D2. Both indexes yields similar conclusions. The standard deviation for human agents $SD2_H=9.05$ is not statistically different than for inexperience and experienced GA ($SD2_{GA}=5.76$ for $T=32$ and $SD2_{GA}=4.79$ for $T=64$) at 0.05 level but is significantly different from the long term value ($SD2_{GA}=4.09$ for $T=400$, Table 1). The same ranking emerges when using the difference between the minimum and the maximum, D2. Individual heterogeneity for Super-experienced GA is smaller than for inexperienced GA ($D2_{GA}=15.66$ with $T=400$ vs. $D2_{GA}=22.78$ with $T=32$); still, human agents are more heterogeneous than inexperienced GA ($D2_H=28.35$).⁸

Had the agents been designed with differentiated goals or variable skills, the heterogeneity of behavior would have not been surprising. Although bounded, the GA agents are endowed with identical levels of rationality. Yet they generate individually distinct behavior. These results are found in several experimental studies, where identical incentives are given and heterogeneous behavior is observed (Laury and Holt, 1998, Cox and Walker, 1998, Palfrey and Prisbrey, 1997, Saijo and Nakamura, 1995).

The only built-in individual diversity among genetic algorithm agents is the random initialization of the strategies. In other words, agents do not have common priors. Moreover, there are four other stochastic operators that might introduce variability in the data: the reinforcement rule, the choice rule, crossover, and mutation. In order to have a benchmark to evaluate the influence of the random element in the results, the GA outcome can be compared with the results of interactions among zero intelligence agents and among noisy Nash agents.

Zero intelligence agents are designed in the spirit of Gode and Sunder (1993) and are essentially pure noise.⁹ The individual strategy for each agent \tilde{x}_i is drawn from a uniform

⁸ Even when the simulation is very long, 10,000 periods, individual heterogeneity does not disappear.

⁹ In Gode and Sunder(1993) they are subject to a budget constraint as well.

distribution on the strategy space [0,50] and then aggregated to compute total resource use, $\tilde{x}_i \sim U[0,50]$ with \tilde{x}_i iid. The outcome from zero intelligence agents is not reported as a viable alternative model to explain the data but to provide a benchmark for the GA results, with special reference to individual heterogeneity. With twice as much aggregate variability ($D2_{ZI}=40.25$ vs. $D2_{GA}=18.10$, Figure 2B), zero intelligence agents are characterized by half as much individual heterogeneity than GA agents ($D2_{ZI}=7.93$ vs. $D2_{GA}=18.97$, Figure 2C).

A fairer evaluation of the impact of randomness in a GA comes from a comparison with Noisy Nash agents. Noisy Nash agents behave in the same fashion as ZI with probability p and are best responders to other Noisy Nash agents with probability $(1-p)$,

$$\bar{x}_i = \begin{cases} \tilde{x}_i, & \text{with prob} = p \\ x_i^*, & \text{with prob} = 1-p \end{cases}. \text{ NN agents - in the same way of the classical model -}$$

understand the concept of Nash equilibrium and are able to compute it, but they occasionally exhibit trembling hand behavior.¹⁰ The level of trembling hand p is set at the same level as the innovation level of GA agents. The comparison between GA and NN is more intriguing. The simulation results for efficiency and aggregate variability are not far from the GA results, but individual heterogeneity is rather small ($D2_{NN}=4.22$), less than one-quarter the GA level and less than one-sixth the human agent level. This latter result suggests that the innovation level is not related to individual heterogeneity in a simple, monotonic fashion. What drives individual heterogeneity in GA agents is not mainly the random element but the individual

¹⁰ For NN, $x_i^* = 72 - \frac{1}{2}(N-1)\left(p\frac{g}{2} + (1-p)x_i^*\right)$, $x_i^* = 14.82$; with $p=0.1492$, $E[\bar{x}_i] = 16.34$ and $E[\bar{X}] = 130.70$.

There are at least two other options to model Noisy Nash. One model involves ZI agents with probability p and $x_i^*=16$, the symmetric individual Nash, with probability $(1-p)$. Unlike the chosen model, the behavioral assumption in this model is that when sane, the agent is not aware that with probability p she is subject to trembling hand and hence $E[x_i] = (1-p)16 + p25 = 17.34$ and $E[X] = 138.74$. Another model involves ZI agents with probability p and $x_i^* = \text{Best response to } E[X_{-i,t-1}]$ with probability $(1-p)$. This latter model is unstable because of the aggregate overreaction to the temporary off-equilibrium situation.

thinking process of each agent along with the inertia built into the decision maker, which leads to path-dependence in choice. In particular, the need to coordinate among many agents might play an important role in generating a diverse behavior across agents. One might also notice that the tendency of GA agents to converge to the Nash equilibrium at the aggregate level seems stronger than at the individual level. In conclusion, Results 1, 2, and 3 are not simply a consequence of the noise built into the GA.

Predictions about other experiments. Besides comparisons with data from baseline common property resource experiments, simulations with genetic algorithm agents allow to make predictions about the effects of different experimental designs. Two changes are here discussed, a modification of the strategy space and the addition of a decentralized monitoring and sanctioning system.

Consider the following three designs: (A) the individual use level strategy space is $[0, 50]$; (B) the individual strategy space is $[0, 20]$; (C) the strategy space is $[0, 16]$. All three designs have the same Nash equilibrium, $x_i = 16$, and differ only in the strategy space. When agents are fully rational, designs A and B simply supply agents with options that are irrelevant to their actions and there is no substantive difference with C. The baseline design considered in this paper is A. In the context of voluntary provision of public good experiments most environment are similar to design C while designs with interior Nash are similar to A and B. Simulations with genetic algorithm show an decrease in aggregate resource use as the individual strategy space reduces from A to B, and then further to C. As Table 3 shows, the efficiency in use achieved by GA increases of about 13 points between A and B.¹¹ The impact of off-equilibrium strategy on the aggregate outcome is driven by the tendency of genetic algorithm agents to experiment with all available strategies. A similar “surprising” efficiency improvement was observed by Walker, Gardner, and Ostrom (1990) in a common

¹¹ Results are less dramatic when GA agents are more experienced (T=400 instead of T=32).

property resource experiment with comparable parameters to the ones set in the simulations run in this study. They report a 40-point increase in efficiency. In designs where rational agents should be unaffected by strategy space choice, the level of resource appropriation is influenced by the strategy space size both for human and genetic algorithm agents. Although to a lower degree than Walker, Gardner, and Ostrom (1990), also public good experiments by Laury and Holt (1998) have revealed a such systematic impact on aggregate cooperation levels of the strategy space. According to them, the most important determinant of the size and direction of these impacts on cooperation appears to be the equilibrium's location relative to the group's potential contributions.¹²

Another design change to the common property resource experiment is the introduction of a decentralized monitoring and sanctioning system. Consider a situation where after having privately decided his own exploitation level of the common property resource, each agent has the option of selecting other individuals for inspection. At a unitary cost, the inspector can view the decision of any individual. If the inspected individual has exploited the resource excessively, relative to a publicly known amount, a fine is imposed and paid to the inspector. In the opposite case, no fine is paid. As the eventual fine is always transferred to the inspector, an agent can make a profit by requesting an inspection on a “heavy” free rider. An experiment in this environment is reported in Casari and Plott (2003) using two sets of parameters values for the sanctioning system. Simulation carried out with genetic algorithm yields aggregate results that are in-between the Nash equilibrium outcome and the human data. Not only GA agents outperform Nash equilibrium predictions at the aggregate level,

¹² “When the Nash equilibrium falls between the lower boundary and the mid-point of the decision space, average contributions typically exceed the equilibrium level. (...) The most important determinant of the size and direction of these deviations appears to be the equilibrium's location relative to the group's aggregate endowment. For example, significant *under-contribution* is observed when the equilibrium is relatively close to the upper boundary.” (Laury and Holt, 1998)

they also reproduce some of the heterogeneity in inspection decisions that one can find in human data. These simulations are not reported in this study.

6 Conclusions

In this paper, we study anomalous results from common property resource experiments using a model of artificial adaptive agents. Experimental outcomes show a systematic departure from the Nash equilibrium prediction, do not settle on a steady state, and are characterized by a remarkable individual diversity in behavior. All three results are at odds with the predictions of the unique, symmetric Nash equilibrium (Casari and Plott, 2003, Rocco and Warglien, 1996, and Walker, Gardner, and Ostrom, 1990). Similar features could be found also in public goods (Laury and Holt, 1998) and Cournot oligopoly experiments (Cox and Walker, 1998).

We employ an individual learning genetic algorithm model to simulate behavior in a common property resource game. Their limitations includes inability to maximize, constrained memory, and lack of common knowledge about the rationality level of others. Similar models have been successfully used to replicate experimental behavior in other environments (Arifovic and Ledyard, 2000, Arifovic, 1994).

Simulations are run through individual learning genetic algorithms and evaluated using the experimental results from Casari and Plott (2003) as a benchmark on three dimensions: aggregate cooperation level, aggregate variability, and individual heterogeneity. There are four main conclusions.

First, genetic algorithm agents closely reproduce aggregate level behavior of human agents both in terms of cooperation levels and variability in aggregate cooperation over time.

Second, the interaction of genetic algorithm agents generates about two thirds of the individual heterogeneity in experimental data. This result is remarkable because the artificial

agents are by design identical in their goal of income maximization and in their limited rationality level. Yet, they do not fully account for the individual heterogeneity of human subjects. Hence, the implication to draw is that the experimental data are in fact generated by different types of agents and hence a descriptive model must explicitly include more than one type of agents. Agent diversity can take two non-mutually exclusive dimensions. Agents could intentionally deviate from the maximization of personal income. In particular, they might exhibit varying degrees of other-regarding preferences. On the other hand, agents could differ in their problem-solving skills. For instance, not everybody necessarily has the same memory constraints or computational limitations. The latter path constitutes an interesting extension of this work.

Third, the evolutionary process underlying a genetic algorithm is fundamentally different from noisy best reply. A simple model with trembling hand fares considerably worse than a genetic algorithm in explaining the data. For a start, notwithstanding a comparable level of noise, noisy best reply can explain less than one-sixth of the individual heterogeneity of human data vis-à-vis about two-thirds of the genetic algorithm. Then, it simply makes a static prediction. On the contrary, with genetic algorithm agents, their experimentation through random search interacts with bounded rationality and, with experience, moves the outcome closer to the Nash equilibrium.

Finally, predictions relative to different experimental designs of common property resource appropriation are put forward. When the strategy space is restricted while leaving the Nash equilibrium unchanged, the cooperation level among genetic algorithm agents raises. Experimental results from Walker, Gardner, and Ostrom (1990) support this prediction. Consider also a situation where after having decided his own exploitation level of the common property resource, each agent has the option of selecting other individuals for sanctioning. Simulation of genetic algorithm interactions under two treatments of such a

decentralized sanctioning system were run but not reported in this study. Such simulation results match many of the experimental data pattern reported in Casari and Plott (2003).

To conclude, we find that genetic algorithm agents exhibit many of the same patterns observed in common property resource experiments. Alongside its evolutionary nature, the ability to generate individually distinct patterns of behavior originating from identical goals and identical rationality levels may be the most interesting feature of an individual learning genetic algorithm.

References

- Andreoni, James and John H. Miller (1995): "Auction with Artificial Adaptive Agents," *Games and Economic Behavior*, 58, 211-221.
- Arifovic, Jasmina (1994): "Genetic algorithm learning and the cobweb model," *Journal of Economic Dynamics and Control*, 18, 3-28.
- Arifovic, Jasmina (1996): "The behavior of the exchange rate in the genetic algorithm and experimental economies," *Journal of Political Economy*, 1104, 510-541.
- Arifovic, Jasmina and Curtis Eaton (1998): "The evolution of type communication in a sender/receiver game of common interest with cheap talk," *Journal of Economic Dynamics and Control*, 22, 1187-1207.
- Arifovic, Jasmina and John Ledyard (2000): "Computer Testbeds and Mechanism Design: Application to the Class of Groves-Ledyard Mechanisms for Provision of Public Goods," manuscript, Pasadena, California Institute of Technology.
- Bäck, Thomas (1996): *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*, New York, Oxford University Press.
- Brandts, J. and Schram, A. (2001): "Cooperation and noise in public goods experiments: applying the contribution function approach," *Journal of Public Economics*, 79, 2, 399-427.
- Bullard, James and John Duffy (1998): "A Model of Learning and Emulation with Artificial Adaptive Agents," *Journal of Economic Dynamics and Control*, 22, 179-207.
- Casari, Marco and Charles R. Plott (2003): "Decentralized Management of a Common Property Resource: Experiments with Centuries-Old Institutions," *Journal of Economic Behavior and Organization*, 51, 2, 217-247
- Casari, Marco (2002): "Can bounded rationality explain experimental anomalies? A study with genetic algorithms," Unitat de Fonaments de l'Anàlisi Econòmica, Universitat Autònoma de Barcelona, UFAE and IAE Working Paper, 542.02.
- Clark, Colin W. (1990) *Mathematical Bioeconomics. The Optimal Management of Renewable Resources*, New York, John Wiley & Sons, Second Edition
- Cox, J. and Walker, M. (1998) Learning to play Cournot duopoly strategies, *Journal of Economic Behavior and Organization*, 36, 141-161
- Daily, Larry Z., Lovett, Marsha C., Reder, Lynne M. (2001): "Modeling individual differences in working memory performance: A source activation account," *Cognitive Science*, 25, 3, p.315-353

- Dawid, Herbert (1996): *Adaptive Learning by Genetic Algorithms, Analytical Results and Applications to Economic Models*, Springer, Berlin.
- Dawid, Herbert (1999): "On the Convergence of Genetic Learning in a Double Auction Market," *Journal of Economic Dynamics and Control*, 23, 1545-1569.
- Franke, Reiner (1997): "Behavioural Heterogeneity and Genetic Algorithm Learning in the Cobweb Model," IKSF - Institut für Konjunktur- und Strukturforchung, University of Bremen, no.9.
- Gode, Dhananjay K., and Shyam Sunder (1993): "Allocative Efficiency of Markets with zero-Intelligence Traders: Market as Partial Substitute for Individual Rationality," *Journal of Political Economy*, 1, 101, 119-137
- Goldberg, D. (1989): *Genetic Algorithms in Search Optimization, and Machine Learning*, New York, Addison-Wesley.
- Gordon, H. Scott (1954): "The economic theory of a common property resource: the fishery," *Journal of Political Economy*, 62, 124-42.
- Holland, John H. (1975): *Adaptation in Natural and Artificial Systems*, Ann Arbor, University of Michigan Press.
- Holland, John H. and John H. Miller (1991): "Artificial Adaptive Agents in Economic Theory," *American Economic Review, Papers and Proceedings*, 81, 365-70.
- Kagel, J.H. and A.E. Roth (1995), *The Handbook of experimental economics*, Princeton University Press, Princeton, N.J.
- Kollman, Ken, John H. Miller, and Scott E. Page (1997): "Political institutions and sorting in a Tiebout model," *American Economic Review*, 87, 5, 977-992.
- Laury, Susan K., and Charles A. Holt (1998) "Voluntary Provision of Public Goods: Experimental Results with Interior Nash Equilibria," in *Handbook of Experimental Economics Results*, edited by C. R. Plott and V. L. Smith, New York: Elsevier Press, forthcoming
- LeBaron, Blake (2000): "Agent-based Computational Finance: Suggested Readings and Early Research," *Journal of Economic Dynamics and Control*, 24, 679-702.
- Marimon, Ramon, Ellen McGrattan, and Thomas J. Sargent (1990): "Money as a medium of Exchange in an Economy with Artificially Intelligent Agents," *Journal of Economic Dynamics and Control*, 14, 329-373.
- Miller, John H. and James Andreoni (1991): "Can Evolutionary Dynamics Explain Free Riding Experiments?" *Economics Letters*, 36, 9-15.

- Moir, Robert (1999): "Spies and Swords: behavior in environments with costly monitoring and sanctioning," *manuscript*, Dept. of Economics, University of New Brunswick, Canada.
- Nowak, Martin A. and Sigmund Karl (1998): "Evolution of Indirect Reciprocity by Image Scoring," *Nature*, 393, 573-576.
- Ostrom, Elinor, Roy Gardner, and James Walker (1994): *Rules, Games, and Common-Pool Resources*, Ann Arbor, University of Michigan.
- Palfrey, T.R., and Prisbey, J.E., 1997. Anomalous Behavior in Public Goods Experiments: How Much and Why?. *American Economic Review* 87, 5, 829-846.
- Riechmann, Thomas (1999): "Learning and Behavioral Stability: An economic interpretation of genetic algorithms," *Journal of Evolutionary Economics*, 9, 225-242.
- Rocco, Elena and Massimo Warglien (1996): "Computer mediated communication and the emergence of 'electronic opportunism,'" *Working paper* 1996-01, University of Trento.
- Rubinstein, Ariel (1998): *Modeling bounded rationality*, Cambridge, MIT Press.
- Saijo, T., and Nakamura, H., 1995. The 'Spite' Dilemma in Voluntary Contribution Mechanism Experiments. *Journal of Conflict Resolution* 39, 3, 535-560.
- Vriend, Nicolaas J. (2000): *An Illustration of the Essential Difference between Individual and Social Learning, and its Consequences for Computational Analysis*, *Journal of Economic Dynamics and Control*, 24, 1,1-19.
- Walker, James M., Roy Gardner, and Elinor Ostrom (1990): "Rent Dissipation in a Limited-Access Common-Pool Resource: Experimental Evidence," *Journal of Environmental Economics and Management*, 19, 203-211.

Table 1: Simulations of Resource Use Without Sanctions –Dynamic Over Time

| | Human agents H (1) | Nash Equilibrium m (2) | Artificial agents, pm=0.02, pc=0.00 | | | Artificial agents, T=64 | |
|---|--------------------------|------------------------------|--|------------------------|------------------------|----------------------------|---------------------------|
| | | | T=32 (3) | T=64 (4) | T=400 (5) | pm=0.02 pc=0.40 (6) | pm=0.01 pc=0.00 (7) |
| GROUP RESULTS | | | | | | | |
| X - Resource use | 131.32 | 128.00 | 131.02 | 130.40 | 130.02 | 130.52 | 129.82 |
| 0.95 confidence interval on X | [127.25, 135.39] | - | [130.96, 131.08] | [130.35, 130.45] | [129.98, 130.06] | | |
| $\sigma(\mathbf{X})$ -Standard deviation of use over time | 12.95 | 0.00 | 17.50 | 15.03 | 14.04 | 14.53 | 9.51 |
| Efficiency | 28.4% | 39.5% | 26.85% | 29.75% | 31.16% | 29.75% | 33.55% |
| Periods with negative earnings | 15.5% | 0.00% | 19.59% | 16.00% | 14.97% | 16.12% | 8.22% |
| INDIVIDUAL RESULTS | | | | | | | |
| MAX2 -Agent with maximum use (2) | 37.92 | 16.00 | 28.71 | 27.72 | 27.01 | 25.04 | 25.89 |
| MIN2 - Agent with minimum use (3) | 9.57 | 16.00 | 5.93 | 8.75 | 11.35 | 9.31 | 9.81 |
| D2 - Individual difference, (2)-(3) | 28.35 | 0.00 | 22.78 | 18.97 | 15.66 | 15.73 | 16.08 |
| SD2 -Individual use standard deviation | 9.05 | 0.00 | 5.76 | 4.79 | 4.09 | | |
| 0.95 confidence interval of SD2 | [5.13, 33.42] | | - [5.04, 6.65] | [4.19, 5.53] | [3.58, 4.72] | | |
| Test of $H_0 \{X_H=X_{GA}\}$ | | | Cannot reject H_0 | Cannot reject H_0 | Cannot reject H_0 | | |
| Test of $H_0 \{SD2_H=SD2_{GA}\}$ | | | Cannot reject H_0 | p-value= 0.05 | H_0 rejected | | |

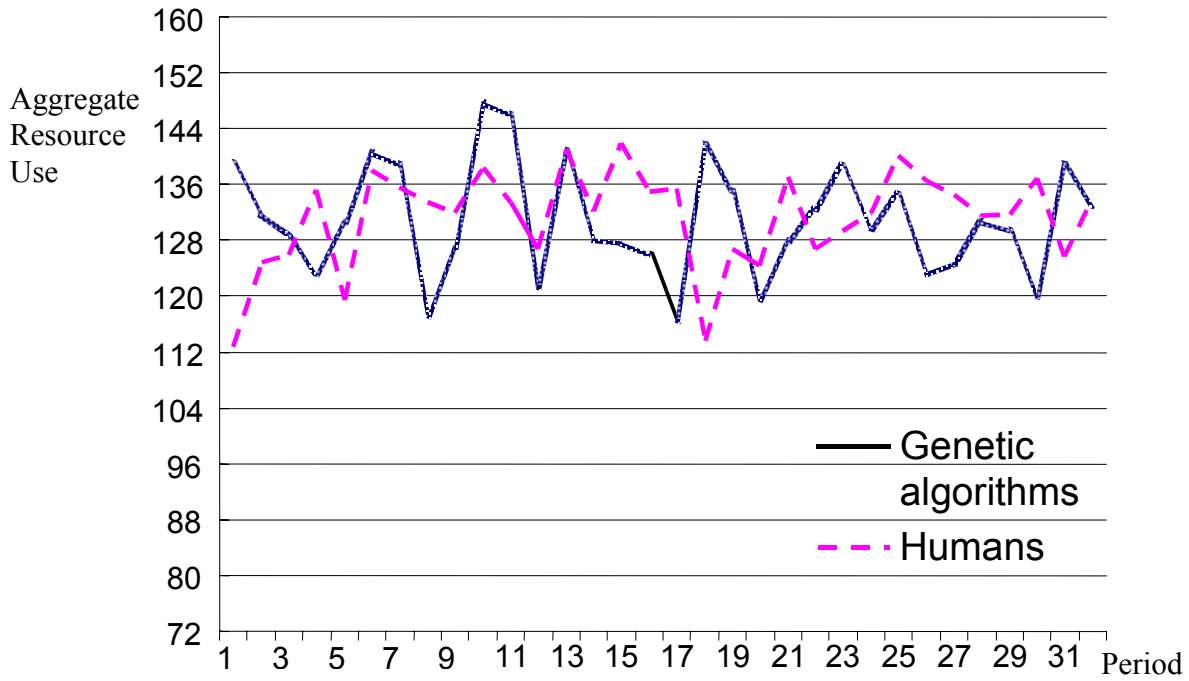
Notes to Table 1: T=total periods of simulation; 32 periods considered in the statistics, T-32,...,T. If T=32 all the periods of the simulation are included in the statistics. The basic action is the use level of the resource, x_{itk} , by agent $i=1,...,8$ at period $t=1,...,T$ for run $k=1,...,100$ (random seeds 0.005 through 0.995). The aggregate resource use $X_{tk}=\sum_{i=1}^8 x_{itk}$ and its average for each run is $\bar{X}_k = \frac{1}{\tau} \sum_{t=T-\tau+1}^T X_{tk}$. The significance tests are carried out, using \bar{X}_k as a single observation, under the null hypothesis H_0 that the random variables Z_i are iid and normally distributed. Parameters of the GA: K=6, N=8, L=8, pm=0.02; GA v.5.0.

Table 3: Simulations of Resource Use – Impact of Strategy Space

| | Human agents | Nash Equilibrium | Artificial agents, pm=0.02, pc=0.00, T=64 | | |
|--|--------------|------------------|--|---------------|---------------|
| | H (1) | m (2) | A [0, 50] (3) | B [0, 20] (4) | C [0, 16] (5) |
| GROUP RESULTS | | | | | |
| X - Resource use | 131.32 | 128.00 | 130.40 | 126.21 | 123.90 |
| $\sigma(\mathbf{X})$ - Standard deviation of use over time | 12.95 | 0.00 | 15.03 | 5.59 | 4.24 |
| Efficiency | 28.4% | 39.5% | 29.75% | 42.69% | 47.68% |
| Periods with negative earnings | 15.5% | 0.00% | 16.00% | 0.03% | 0.00% |
| INDIVIDUAL RESULTS | | | | | |
| MAX2 - Agent with maximum use (2) | 37.92 | 16.00 | 27.72 | 17.65 | 15.94 |
| MIN2 - Agent with minimum use (3) | 9.57 | 16.00 | 8.75 | 10.65 | 14.81 |
| D2 - Individual difference, (2)-(3) | 28.35 | 0.00 | 18.97 | 7.00 | 1.03 |


Notes: K=6, N=8, L=8, pm=0.02; GA v.5.0. See notes to Table 1.

Figure 1: Aggregate resource use: human versus genetic algorithm agents




Note: Humans, average of four sessions; GAs, average of simulations with four different random seeds.


Figure 2: Genetic algorithms and randomness



Notes: **Nash equilibrium**: prediction with selfish, perfectly rational agents; **Human subjects**: average of 4 experimental sessions; **Genetic algorithm agents**: selfish, boundedly rational agents ($T=64, \tau=32$, average over 100 simulated runs, v.5.0); **Zero-intelligence agents**: random draws from a uniform distribution (average over 100 simulated runs v.5.6); $\tilde{x}_i \sim U[0,9]$ with \tilde{x}_i iid, $\vartheta=50$; **Noisy Nash agents**: are ZI with probability p and are best responders to other Noisy Nash agents with probability $(1-p)$.



2B - GROUP VARIABILITY



2C - INDIVIDUAL HETEROGENEITY