Agents, Groups and Institutions:

Experiments with population based simulation models

José Maria Castro Caldas
Hélder Coelho

Janeiro 1999

WP n.º 1999/12

DINÂMIA
CENTRO DE ESTUDOS
SOBRE A MUDANÇA
SOCIOECONÓMICA

FCT
Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA CIÊNCIA, TECNOLOGIA E ENSINO SUPERIOR

Agents, Groups and Institutions:
Experiments with population based simulation models

José Maria Castro Caldas & Helder Coelho

Janeiro 1999

Contents

1. Introduction

During the last decades the role of institutions[1] has been re-emerging as a central topic of economic analysis. Within this movement, that permeates different traditions and economic paradigms, three main trends can be identified: the one with roots in neo-classic economics, labelled as 'New Institutional Economics' (Williamson, 1985; North, 1990; Olson 1965), the one founded in the Austrian tradition and represented by the more recent research of Hayek (for instance, 1973), and the one inspired by the 'old' American institutionalism of Veblen and Commons (Hodgson, 1988, 1993). No unified Institutional Economics therefore exists, and whether or not these different roads tend to converge is controversial. Nevertheless, it is clear that, in spite of the deep methodological and conceptual differences, the problems that are being addressed are the same, and valuable insights can be found in the writings of authors belonging to the different trends.

What is at stake is the institutional deficiency of economics that was diagnosed long ago in the writings of the German Historical School, of Marx and Veblen, but that only now is generally being taken seriously. The sense of urgency in remedying the situation, felt by an increasing number of economists of all schools, is understandable. Three reasons might explain it. First: in spite of the brilliant contributions of many generations of neo-classic economists, the general equilibrium model that explains socio-economic order in the frame of a small set of given institutional arrangements could not get rid of a set of strong assumptions. While many feel that it is still worth trying to develop this model in order to increase its robustness, others consider that this is a lost cause and choose instead to explore alternative paths. Second: to take the institutional framework as given, as comparative statics equilibrium models do, may be legitimate but only in short term analyses. Anyone interested in the processes of change over time (past or future) sooner or later will have to deal with innovation, concerning not only knowledge and technology, but also institutions. Third: recent historical events turned the institutional deficiency of economics into a scandal, eroding the prestige of the profession and the scientific status of economics. The transition to capitalism in the Third Word and in some former socialist countries has an abundance of examples of tragic results of a political action oriented by an institutional deficient theory. Why is it that 'free market' and private property - the recipe for an affluent society in the eyes of many economists - may simply turn into fundamentalism or Mafia dominated economies when imposed on societies and cultures that in the recent past diverged from capitalism?

However, bringing institutions back into economics has not proved to be an easy task. One of the reasons is that including institutions in economic analyses amounts to more than some minor changes in the received models. Even though institutions can be incorporated in a

---

[1] Institutions have been defined by economists in a vast number of (sometimes conflicting) ways. Not wishing to add a new definition to the list we give some of those that convey the meaning in which the term is used in this essay (often interchangeably with *shared rules*): "interconnected and mutually stabilizing configurations of rules" (Viktor Vanberg); "the rules of the game in a society or, more formally, human devised constraints that shape human interaction" (Oliver North); "a way of thought or action of some prevalence and permanence, which is embedded in the habits of a group or the customs of a people" (Walton Hamilton).

rational choice frame as restrictions in the choice set of the agents, this approach leaves out important unresolved problems. The other reason is that the theoretical enterprise of institutional economics demands an articulation of contributions of different disciplines that is made difficult by the prevailing specialisation within the social sciences.[2]

The ultimate goal of an economic theory that focuses on the role of institutions can be stated as "guiding our efforts to improve the social order in which we live by improving the rules of the game [...]" (Vanberg, 1994: 5). However, two question wait to be answered before any progress in this direction is possible: What is the role of institutions and how are they related to the aggregated performance of the socio-economic systems? What is the relationship between institutions and the behaviour of the agents?

The aim of this research is to address those questions by means of a combination of analytical building blocks from behavioural theories of choice and different accounts of institutions and institutional change. In this we are aided by a particular type of formal model, here referred to as 'a population based simulation model'. This essay reports on the results obtained, which must be taken merely as tentative hypotheses and conjectures.

The main effort within the different programs of research on institutional economics has been devoted to theorising at high levels of abstraction by means of verbal models. The role of formal models, compared to other fields of economics, has been modest. One of the reasons for this is that the mathematical modelling toolbox of the economist seems to be insufficient for the present challenge. While expecting that others, sooner or later, will come up with more appropriate mathematical tools, in line with previous work (Caldas and Coelho, 1994), we simulate socio-economic processes incorporating many ideas and techniques from Artificial Intelligence (AI) that have been proving to be of great utility.

AI and Economics have a great experience of co-operation. Herbert Simon (the economist founding father of AI) became interested in Computer Science when he understood that computer simulation models might be useful to analyse and derive the implications of theories of bounded rationality. We now share the same interest but in a context that Simon did not fully explore in his seminal papers, and that AI researchers only recently began to explore - the interaction of bounded rational agents.

The simulation techniques are one of the contributions of AI. The population based simulation models that will be here presented were built upon previous work in Multi-Agent Systems and Simulation, related in particular with Artificial Adaptive Agents, Genetic and other Evolutionary Algorithms (Holland, 1975, 1995; Goldberg 1989) and their use in the simulation of societies (Doran and Gilbert, 1994; Chattoe, 1994, 1998, Arifovic, 1991; Epstein and Axtell, 1996). But the contribution of AI is not restricted to the technical aspects. An important body of theoretical work crossing the borders of AI and social theory (Conte

---

[2] As Hayek sees it, "the rules of just conduct which the lawyer studies serve a kind of order of the character of which the lawyer is largely ignorant; and [...] this order is studied chiefly by the economist who in turn is similarly ignorant of the character of the rules of conduct on which the order that he studies rests" (Hayek, 1973: 4).

and Castelfranchi, 1995; Ferber, 1995) exists that is of great relevance for institutional economics.

The simulation models are built upon a theoretical foundation, the building blocks of which may be summarised as follows:

- human motivations are complex and hardly reducible to a limited set of principles; agents differ as much in character as in capabilities; diversity is an essential feature of human populations;

- the agent has a limited deliberative capability; deliberation and rule following are not necessarily contradictory features of an agent's behaviour;

- for the same agent different levels of choice coexist: the level of situational choice over alternative actions and the level of choice over alternative sets of shared rules;

- institutions are both a result and a pre-condition of the interaction of agents: their emergence and change over time can be a result of a spontaneous process, or of a 'negotiation' among deliberative agents in which power counts;

The model is evolutionary[3]. Not only are we interested in time patterns of stability and change, but also we are looking for explanations based on process (Gilbert, 1995). We want to understand how the unfolding of simultaneous multiple actions that generate aggregate effects, which in turn become causes, in a time sequence that is neither reversible nor predictable, may generate patterns. Stable or unstable, these patterns are the interesting objects for scientific scrutiny; the process through which the interaction of multiple agents leads to them is often called emergence.

Institutions, as Gilbert states, can be viewed as "emergent from individual action", but as he also notes, this may be a misleading view if we fail to acknowledge that "people have the ability to recognise, reason about and react to human institutions", and, we might add, purposefully act in order to strengthen or reshape those institutions. People do not simply adapt to an ever changing environment, they act upon it and they add to the natural environment a social element that accounts for their survival. The type of behaviour that emerges out of emergence is called by Gilbert, second order emergence. The simulation of these second order processes which is "one of the present-day challenges for simulation in the social sciences" (Gilbert, 1995) is precisely the final aim of the modelling exercise that supports the present research.

---

[3] Although not based on a "detailed and coherent analogy between a biological evolutionary process and a social process" (Chattoe, 1998).

The following pages will be devoted to a further clarification of the motivation of the research, to the specification and interpretation of a basic model and to the recombination of the analytical building blocks by means of step by step extensions to the model. The second section will start from a discussion of the reasons why the Smithean and the marginalist approaches to the problem of socio-economic order no longer seem to be a promising road to follow. Briefly, it turns to alternative accounts of society and human action, and concludes with an attempt to sketch a simple model of a bounded rational agent acting in a social context. The third section deals with previous uses and interpretations of Holland's Genetic Algorithm (GA) as a model of a socio-economic process, proposes modifications on the algorithm that lead to an alternative interpretation and to a model which (since it may no longer be considered as a GA) is here referred to as a 'population based simulation model'. In section four, institutions are brought more explicitly into the framework. Viewing institutions as solutions to recurring problems in social interaction the section starts with the implications that can be drawn from three simple game situations related to real world problems and from the simulation of the corresponding processes with the proposed model. These implications lead us to focus, in section five, on the problems of co-ordination and co-operation discussing alternative solutions to the problem of collective action and to an extension of the model that incorporates different levels of choice and a monitoring meta-agent. Conclusions are drawn in section six.

2. Answers to an old enigma

2.1 Smith's no longer satisfactory answer

An old enigma hasn't been answered in a satisfactory way: how is it that more or less stable patterns of social relations are established that ensure the production of the material conditions for the existence of populations of individuals that even though competing for scarce resources are unable to live without one another? Nature seems indeed to have played a trick on human beings - "We cannot survive alone, yet unlike social insects we are not genetically hardwired for co-operation" (Macy, 1998).

The founding fathers of economics, first of all Adam Smith, had an answer to this old paradox. In his time, and for generations of economists, his understanding of the socio-economic order was almost taken as a revelation. In a time when it was generally believed that the welfare of nations rested on the wisdom and benevolence of the rulers, Smith argued instead, that the economic order and welfare existed in spite of the rulers, and should be understood as an unintended result of the actions of a multitude of individuals who pursued their own interests.

In *The Wealth of Nations*, Smith stated clearly that he was not implying that individuals were egoists deprived of any moral sentiments; it simply happened that when it comes to compete over scarce resources, the motive of action that prevails is "self-love". He assumed that these

self-interested individuals were able to recognise that exchange is beneficial for them (given the advantages of the division of labour) and he concluded from a theory of value that universally acceptable terms of exchange might be reached for the satisfaction of all. He concluded also that, in spite of the self-interested orientation of the agents, the economic order resulting from the spontaneous process of exchange was a 'good' one, at least in the sense that it would be preferable to any arrangement devised by a human mind and enforced by any ruler. For Smith, and for most economists after him, the self-interested agents abided by the rules of justice and took them for granted, and this has been interpreted as meaning that, the property rights were perfectly defined and secure and that contracts were enforceable and enforced. Government, markets, property rights and enforceable contracts entered the picture as given - apart from these constraints everyone would be better off if the agents were simply let free to decide.

What Smith couldn't dream of is that in order to prove in a rigorous manner (as general equilibrium theory did) his two main conclusions - (a) the existence of a state of general agreement and satisfaction over a set of relative prices, and (b) the 'good' properties of this state of affairs - so many assumptions (and hard to swallow) had to be imposed on the world. In fact, it is common knowledge among economists that the existence of a market order with efficiency properties can only be proven if perfect rationality, competition, non-increasing returns, absence of public goods, of external effects and of transaction costs... are assumed.

What may then be the problem with Smith's answer to the old paradox? The problem is that it can only be made sense of in a world where no ignorance, novelty and real interdependence among agents would exist. That is, in fact, what is implied by the above listed assumptions. We would not have to worry much about their realism if there were signs encouraging us to believe that the model will in the end turn out to be robust, in the sense that the main conclusions will still prevail when these 'simplifications' are relaxed. While many still believe that this may be the case and keep working on reformulating the model, a growing minority feels that the difficulties of the general equilibrium theory simply show that the explanation for the paradox of society has to be looked for elsewhere. Given the conditions in the real word, the economic order may depend much more than Smith and the general equilibrium theorists assumed on the institutional frame of the interaction; and, this may be taken as an invitation to redirect the research efforts of the economists towards alternative accounts of the social order and human action.

2.2 The unsolved tension between concepts of human action

Explanations of action in terms of social wholes might then seem to the disenchanted economist to be better candidates for alternative accounts. Rather than a result of a spontaneous process of co-ordination of a multitude of actions of self-interested autonomous agents, the economic order might be understood as a structural property reproduced in time by the practices of individuals. The agents would not only be constrained, but would themselves be moulded by institutions, behaving habitually in accordance with social norms

and conventions. However, if this step is given, the disenchanted economist must be prepared to give up at least some of the deliberative and creative capabilities of human beings, which is something that most modern sociologists seem to be even more reluctant to do than the discontented economists.

In fact, over the recent decades sociology seems to have moved in a direction opposed to that of economics. While economics started to shift from rational and self interested choice, as the main explanatory device, to "norm or rule-guided behaviour" (Vanberg, 1994: 13), sociology made the reverse movement, in what Giddens describes as a "rejection of the tendency (...) to see human behaviour as a result of forces that actors neither control nor comprehend" (Giddens, 1984: XVI).

Homo economicus and *homo sociologicus* no longer are the legitimate representatives of the two main concepts of action within the social sciences. *Homo economicus*, the once proud "instrumentally rational and calculating seeker of preference satisfaction" (Heap et al., 1992: 63) is losing faith, as evidence accumulates showing that it lacks empirical content and misses some of the most relevant features of the human being. If most economists still give him shelter, it is only because they refuse to consider other possibilities. *Homo sociologicus*, meanwhile, started to give signs of being tired of living "according to rules, roles and relations" (Heap et al., 1992: 64). Having left home during the sixties, probably for good reasons, he has not returned since. Some sociologists even turned to rational choice to try and understand him.

The disenchanted economist will then discover that on both sides of the classical quarrel no one seems to defend the extreme clear-cut positions with passion any more, and that a tendency to agree without hesitation on any terms that are proposed for a peace settlement has replaced the inflamed debates. For instance, Hayek's compromise formulation - "man is as much a rule-following animal as a purpose-seeking one" (Hayek, 1973: 11) - might do. And yet this sentence, although easy to agree with, tells us nothing on how to combine the two terms of the proposition - how exactly can an agent be simultaneously a rule-follower and a purpose-seeker? Our economist may find it desirable to overcome the tension between the two alternative accounts of action, but for this he will have first to try and clarify the terms of the compromise.

2.3 Rule-following and purpose-seeking behaviour

The problem with economic man is not that it is rational. If, as Simon (1987) puts it, rationality simply meant "that people have reasons for what they do" there would be no quarrel about it as a concept of human action since it could mean at least three different things: (a) that people have consistent preferences and that they have the knowledge, the information and the computational powers to select the action that best serves their purposes; (b) that people have objectives but a limited knowledge, information and computational power, so that they use in their deliberations heuristic procedures that lead to 'satisfying'

choices; (c) that following a social rule or convention may be a sufficient motive behind choice. Rationality however has been defined in economics in a much narrower sense and therefore the dispute goes on.

What can make sense of rule-driven behaviour? According to the concept (a) of rationality an agent may follow a rule only to the extent that this rule happens to coincide systematically with a utility maximising choice or if that rule can be thought of as a restriction in the choice set of the agent. Alternatively, in concept (b), the rule may simply amount to an heuristic device to achieve satisfactory levels of a set of objectives. In the case of (c), however, we may find agents that systematically behave in accordance with rules regardless of any deliberation; but rule-following here has a very different meaning: it is a practice driven by a social norm or convention, and an agent may follow a rule even if he believes that there are alternative choices that would lead to higher levels of satisfaction of his personal objectives.

In concepts (a) and (b), the neo-classic and Herbert Simon's concepts of boundless and of bounded rationality, purpose comes first. Rationality is instrumental to purpose. In concept (c), on the contrary, purpose is obscured by rule compliance; but it may be argued that the rule-following agent may be far from the equivalent to a robot, since "all rules provide a space for creative interpretation" (Heap et al., 1992) and therefore for purposeful action.

Concepts (a) and (b) seem to put too much weight on purpose and deliberation, concept (c) seems to put the excessive weight on rule compliance. With exaggeration we could say that instrumental rationality cannot explain why people often behave in accordance with social norms and conventions, and that the alternative concept cannot explain why they often break these rules or try to replace them by new ones. Is it possible to balance purposeful action and rule following in a single model? We believe that Herbert Simon, would be the last to deny it - he (Simon, 1958) spoke of a human being that "is rational within the bounds set by his social role of economic man [whose] calculations rest on his expectations that the others with whom he interacts will accept their corresponding social roles. [And of facts...] that obtain their status as facts by a social process of legitimation...". However, in his seminal papers (Simon, 1955, 1956) where he outlined a model of bounded rationality, the scenario was one of a single agent in interaction with the world or at best two agents over a chessboard. Rule-following as a sufficient motive was not included.

Nevertheless, Simon's critique of neo-classic rationality and a discussion of his 1955 model of bounded rationality may be an appropriate start for the setting up of a model of a bounded rational agent in a social context.

2.4 Bounded rationality in a social context

Let us recall Herbert Simon's basic critique of the rational choice model (Simon, 1955). This model assumes that the agent has: (a) information and knowledge on the relevant aspects of his environment; (b) a well-behaved system of preferences; (c) "skill in computation [and information, and knowledge] that enables him to calculate, for the alternative courses of action that are available to him, which of these will permit him to reach the highest attainable point on his preference scale". As Simon insisted the model lacks empirical content since it demands too much from the agents, ignoring the existing limits to perception, knowledge and computational skills. Looking back on the discussion on bounded rationality and on the experimental evidence on choice behaviour that meanwhile has accumulated, the most relevant points of Simon's critique seem to be the ones related to the limits to perception and to the imperfectness of the mental model which severely restrict the range of opportunities that can be perceived by a human agent.

Recall Simon's tentative model of bounded rational choice. Given, $\mathring{A}$ (a set of alternative choices perceived by the agent), belonging to A (the set of all possible alternative choices), S the set of perceived outcomes of all choices in A, $V(s)$ a real valued scalar or vector function, and k, a real valued scalar or vector *aspiration level*, the choice procedure in this model can be summarised as follows:

- Step 1: search S for a point S_a such that $V(s)$ is satisfactory, that is a point such that $V(s) \geq k$

- Step 2: search for the alternative a in $\mathring{A}$ that maps into S_a

- Step 3: If a can be found in $\mathring{A}$, then stop; Else if there are unexplored alternatives in $\mathring{A}$, then return to Step 1; Else if the whole set $\mathring{A}$ has been explored, then go to Step 4

- Step 4: lower the aspiration level k or search for additional alternatives in A that can be adjoined to $\mathring{A}$ and return to step 1.

In a recursive setting the model might be complemented by an updating procedure of the aspiration level like: if an alternative has been easy to find, then increase the aspiration level, or else decrease it. The weakest point in Simon's simple tentative model is perhaps the aspiration level concept. It clearly serves as a devise aimed at resolving the problems with preference orderings. However its use involves the hard assumption of a process of decision making where points in S are sequentially explored. The obvious problem is that if the initial information set of the agent contains both $\mathring{A}$ and its mapping on S it is unclear why an agent should settle on the first alternative that satisfies an aspiration level. Even though the aspiration level concept could be incorporated in the model that will be next outlined, we will not retain it in this formulation. To drop the aspiration level concept re-opens the problematic

question of the preference orderings. However in the social context of decision that will be set in the next paragraphs the agents may only compare and evaluate the consequences of *observed* actions whereas the problems with preference orderings revealed by experimental evidence are mostly related to expected consequences.

Consider a population of *n* agents living in a world where the systemic results *Y* are determined by the set $\mathring{A} = (a_1, a_2,...,a_n)$ of actions of the individuals in the population. The function that maps *A* into *Y* may be unknown to the agents and may change in time, but given a state *y* of the system and the corresponding set $\mathring{A}$ of actions every agent can assign credit to any action in $\mathring{A}$ using a function f_i that models the current state of his preferences and that, given *y*, maps $\mathring{A}$ into a set of evaluations $S'=(s'_1, s'_2,...,s'_n)$. The agents must recursively pick up an action from the set *A* of all feasible actions in a discrete sequence of time periods t_1, $t_2,...,t_T$.

The situation of an agent in this particular world can be described as one of radical uncertainty: "Future events cannot be associated with probability distributions based on knowledge of the past" (Faber and Proops, 1998). This uncertainty may arise, from the behaviour of the other agents, and from the aggregated behaviour of the system. This means that given any action an agent may be led to consider for time *t*+1, there is no deterministic or probabilistic way by which he can form an expectation of its payoff. However, living in a world of uncertainty we all formulate hypotheses concerning the future based on past experience (Holland, 1986), even if philosophers have taught us that no deterministic or probabilistic truth value can be assigned to these hypotheses. To a certain extent we are inductive agents and induction, in fact, is based on the presupposition that somehow the future will be similar to the past. An inductive agent would therefore form expectations of the future, based on a mental model of the world that includes the belief that the immediate future will be similar (but only similar) to the present, keeping his mind open to constantly revising his beliefs. In what concerns life in society the presupposition of a comparative stability may turn out to be less arbitrary and dangerous for the agents as it may seem. After all we may count on the existence of behavioural patterns in society, a factor that reducing uncertainty, enables choice and action.

This environment differs from the one modelled by Simon in a crucial aspect: while Simon's agent could evaluate any perceived set of actions, our agent cannot. Since *Y* depends on $\mathring{A}$, and the function that maps *A* into *Y* may be unknown to the agent, he may only evaluate actions that belong to an observed set $\mathring{A}$; that is, actions that were simultaneously performed by the agents at a given moment in time. Furthermore, since the function that maps *A* into *Y* may change with time, any evaluation involves an error that (for smooth changes in the world) is larger for evaluation based on remote states as compared to evaluations referred to recent states.

Let us have agent *i* in time period *t* deciding what to do in time *t*+1. What he knows about the world is the system's history, that is, the sequence of systemic results y_j, of actions $\mathring{A}_j$ and of

evaluations S'_j, with $j = 0, 1,...,t$. The simplest possible way to model the decision procedure of one agent taking on board the preceding considerations is possibly the following:

The set $\mathring{A}_t$ is the one evaluated by the agent.

Three reasons for this were given: First, only observed sets of actions may be evaluated. Second, the present is taken as a proxy of the immediate future. Third: the error of the evaluations grows as the agent looks further into the past.

The agent has two alternatives when trying to reach a decision:

(a) to choose an action that looks promising from the set of observed actions $\mathring{A}_t$.

In this case given the set S'_t the agent selects an action using a lottery where the probability of selection is somehow proportional to the credit assigned to each action in $\mathring{A}_t$. (Note that, the agent might instead select the most credited action. However, this would be arbitrary since the credit assigned to each action is taken by the agent as mere indication. The lottery is then a device to model a choice in condition of uncertainty in which 'intuition' guided by credit is in command.)

(b) to choose an action in A not included in $\mathring{A}_t$ in order to test it in $t+1$.

In a world of limited knowledge and information there are reasons to be innovative. Opportunities may be hidden by the fog of uncertainty. This innovative move may be modelled in two ways: the agent may randomly modify the selected action in (a), or he may try the abductive move of recombining the selected action with other actions selected by the same procedure.

The tentative model of the agent's decision procedure is outlined. Seemingly the modelled agent is only a deliberative purpose seeker. In what way may this model help us to reconcile rule-following with purpose-seeking? We will see after the model is implemented in a computable system where multiple agents interact.

3. The Genetic Algorithm as a modelling tool

3.1 Genetic Algorithms

When thinking of an appropriate tool to model social processes, the Genetic Algorithm comes up as a natural candidate, together with other evolutionary algorithms. The Genetic Algorithms, created by John Holland at the University of Michigan during the sixties (Holland, 1975), are defined by Goldberg (1989: 1) as "search algorithms based on the mechanics of natural selection and natural genetics".

In spite of the complexities of the mathematics involved, the procedure of the GA, in itself, is fairly simple, and that is one of its appealing features as a modelling tool. In the context of optimisation the basic idea can be quickly grasped. If we have an arbitrary known function to be maximised, one standard heuristic procedure is the hill-climbing one. We start by a randomly chosen point in the solution space and compute the dependent variable value corresponding to this point. Next, we move to a neighbouring point in the solution space and compute the dependent variable value; if this value is higher then the previously obtained, we move to this new point; if not a new neighbouring point is explored. The procedure is iterated until a point is reached in which no neighbour yields a higher value to the dependent variable. The obvious problem is that the climber that follows this rule on a foggy day may be easily misled to a local peak instead of a mountain top.

The GA is initiated not by a random solution but by a set (*population*) of random solutions (*individuals*). These solutions are encoded by a string of characters of any arbitrary finite alphabet (*chromosomes*)[4]. Once decoded, the value (*fitness*) of each solution is given by the corresponding value of the function to be maximised (*fitness function*). The algorithm operates in a sequence of *generations* in which a new population replaces the old. The transition from one generation to the next includes two phases: (a) in the first phase a *reproduction pool* is formed by adding a given number of individuals with above average fitness to the population, and eliminating an equal number of bellow average fitness individuals; (b) in the second phase the chromosomes of some of these individuals are modified by *crossover* with other chromosomes in the same pool, or by random change (*mutation*). This procedure leads to a new population. The fitness of the individuals in this new population is once again determined and the procedure is repeated until some stopping criteria is met.

Compared with hill climbing the GA procedure searches the solution space not as a lonely individual might do, but as a search party that starts from different points and proceeds with operative communications, would. The empirical results obtained with the use of the GA in function maximisation show that the GA usually converges to 'good' quality solutions. The mathematics of the GA (Holland, 1975, Goldberg 1989) shows that in spite of the stochastic

---

[4] For instance, if a solution is given by an integer, any solution can be coded as a 0/1 string if the base 2 is used to represent this number.

nature of the procedure it has nothing to do with a random search, and it helps to understand why it works effectively in a wide range of situations.

3.2 The GA as a model of socio-economic processes

The other appealing feature of the GA, as compared to standard mathematical models, is that it lends itself to more realistic socio-economic interpretations in the sense that they require less competence, knowledge, and information processing abilities, from the agents. Arifovic (1991) mentions two alternative interpretations referred by Chattoe (1998) as a *mental interpretation* and a *population interpretation*. These alternative interpretations may be presented as follows: (a) Mental interpretation - the population represents a single mind; each chromosome in the population represents a rule; "the frequency with which a given rule is represented in the population indicates the degree of credence attached to it" (Arifovic, 1991:2) (b) Population interpretation - the population represents the active rule of each agent in the population; the frequency of a given rule in the population indicates "the degree to which it is accepted in a population of agents" (Arifovic, 1991:2) .

In the comments that follow we will have in mind both Arifovic's interpretations and use of the GA and Chattoe's critique to her use and interpretations. Three points should be clarified concerning Arifovic's work: (a) Both in the mental interpretation and in the population interpretation the GA populations are collections of rules (not individuals)[5]; (b) The interpretation favoured by Arifovic seems to be the mental interpretation[6]; (c) In Arifovic's thesis there is a (problematic) behavioural interpretation of the GA; (d) Arifovic uses the GA (population interpretation) in contexts where "fitness values depend upon the entire population in existence": "As the environment is not fixed during learning", she writes, "and since the behaviour of the population determines the value of the variables, the fitness of each string depends on the overall behaviour of the population" (Arifovic: 1991:7). This is an important point to keep in mind throughout the discussion that follows.

Arifovic's implementation of the 'GA - population interpretation' for the Cobweb Model may guide us in the discussion. In this model *n* firms compete as price takers in the market for one homogeneous good; production takes time and the production decisions must be made before the market prices are known; in time *t* every agent announces the quantity that he is offering and the market clearing price is determined by a demand function, and announced[7]; trade follows with all items exchanged at this price; after the closing of the market place every agent has to decide what to do in *t*+1; the individual costs of production depend on the quantity individually supplied (the function that determines them is the same for all the

---

[5] It is however true that Arifovic is not absolutely clear about this. On page 24 of her thesis (Arifivic, 1991), we can read in respect to a population interpretation: "a GA population consisting of *n* strings represents *a population of decision rules* of *n* market participants"; however, in page 7 she had written that "a population of strings represents a whole *population of agents* with different opinions" (emphasis added).

[6] "It is richer compared to a single population GA in a sense that we can now think of firms as having a number of different ideas about their possible production quantities..." (Arifovic, 1991:33); "it is not simplistic in a way in which single-population GA is, i.e. where the assumption is that each agent has only a single strategy at his disposal" (Arifovic, 1991:34).

[7] Note that while the prices have to be known to the agents, the demand function may not.

agents and must be taken to be known to all); given this setting when the market place closes for the time period the profits of all firms are common knowledge.

With the population in generation *t* standing for the production decision rules of the *n* individual firms the procedure is as follows:

(1) Decode each string to determine the quantities q_t^i offered by each firm;

(2) Sum up the offers to compute the market clearing price P_t;

(3) Given q_t^i and P_t, compute the profit of each firm at time *t* and take it as a fitness measure of each rule;

Each firm must now decide what to do in time period *t*+1...

(4) Firm *i* selects one of the production decision rules in the population. The probability of one rule being selected is proportional to its relative fitness; (Arifovic interprets this operation as an act of imitation of a successful rival).

(5) The selected rule may be the one to be adopted by firm *i* in *t*+1. But firm *i*, however, may decide to try a "new idea". In this case it "may recombine existing beliefs" selecting a new action for crossover with the one that was previously selected, or it may try a random mutation in the originally selected action.

Having encountered "that the algorithm would not settle to the particular value, no matter how many generations a simulation was run for [...; that it gets] close to the unique [expected] rational expectations equilibrium, but then continues to fluctuate around it without settlement", Arifovic introduced an election operator that works as follows:

(6) The fitness of every 'new idea' generated in (5) is computed assuming price P_t, its fitness is compared to the fitness of the parent string. The 'new idea' will replace the old idea only if it is fitter than the "old" one[8].

What are the questions raised by this interpretation and implementation going from the easiest to the hardest ones:

(a) Why should the agent imitate the choice of a "successful rival" in time *t* and not in time *t-m*?

---

[8] With this Arifovic may have missed one the most interesting features of the GA: the fact that it adapts to changes in the environment. What sense can it make for an agent to try and achieve a better result through mutations after an equilibrium has been found? Is it shear madness, like the one of people who drive on the wrong way of motorways 'for the kiks of it'? Not necessarily. What would happen to a population of sane agents using election if the word would change in *t*+1?

(b) Why shouldn't he simply imitate the rule of the most successful rival?

(c) If individual profit in time t, depends on the set of offers in that period, to select a production decision rule for $t+1$ based on relative profit levels in t, amounts to assuming that the agents believe that the state of the market in both periods will be the same?

(d) If this assumption is taken on board is it compatible with the election procedure? Is the fact that a 'new idea' is being tested consistent with the assumption that price remains constant?

(e) Is the fitness function a system level function that determines the 'global' quality of the decision rules or is it a local device used by the firms to assign credit to the alternative rules?

(f) Even if the agents are aware of each other's actions is it easy to accept that they are aware of the structures that represent the rules 'behind' those actions so that they can copy them, recombine them, or mutate them. What is the interpretation of a string or chromosome?

Although preferred by Arifovic the mental interpretation as implemented by her seems to be even more problematic. In the implemented version of the GA to the Cobweb Model there are as many populations as firms in the market. With population j in generation t standing for alternative production decision rules of the individual firms j, the procedure is the following:

(1) Given population j and the corresponding fitness values, firm j selects one of the production decision rules in the population to be performed in time t (as before the probability of one rule being selected is proportional to its relative fitness). The procedure is repeated for all populations and the offers of the n firms are determined;

(2) Given those offers the market clearing price P_t is computed;

(3) Using this price, the *reproductive fitness* values of all strings of all populations are computed given the profit function. Note that the profit corresponding to all strings is being computed and not only the one corresponding to the strings that were selected generating the equilibrium price P_t. However, Arifovic interprets this *reproductive fitness* as "a profit that firm j would have earned had string i from its population been used for production".

(4) Given this *reproductive fitness* the reproduction pool of each population is selected, the crossover and mutation operator performed in the standard way and thus a new population generated.

(5) The *competition fitness* of the newly generated population is now computed assuming once again P_t step 1 is performed.[9]

Three questions are left unanswered:

(a') why should an agent's decision be built only upon rules that are active in time t and not t-m;

(b') Why shouldn't he simply choose the most successful rule according to its *competition fitness*?

(c') How can P_t be assumed to remain constant throughout the iteration when the rules that are being evaluated differ from the ones that determined the market clearing price?[10]

In spite of the unanswered questions the more problematic aspects of Arifovic's implementations may be related more to the "orthodox economic framework" of her study (Chattoe, 1998) than to the limits of the GA as a modelling tool. For instance, the introduction of the election operator or the use of Gray coding, are only justified by the assumption that the procedure should converge and that it should converge on the rational expectations equilibrium.

The questions raised in respect to Arifovic's implementation, must be answered. Our interpretation of the GA which is based on the model of the decision process outlined in 2.4, tries to do so in respect to the population interpretation[11]:

the reasons why an agent may concentrate on set $Å_t$ given in 2.4 may answer questions (a) and (c);

the reason given for a lottery procedure (see 2.4) for the choice of an action in this set may answer (b);

(d) is simply solved by not using election;
(e) is answered by clearly stipulating that the fitness function is an individual credit assigning device;

In order to approach (f), however, an interpretation must be given to the string (or chromosome). To think of a string as a program (as in Genetic Programming) that generates an action is one of the possible approaches. Since different programs may lead to the same action it is inconceivable to think of an agent that, from the observation of one action, is able

---

[9] Arifovic omits the procedure used to compute the fitness of the initial population in each simulation. We can imagine several alternative procedures.
[10] In other implementations of the mental interpretation (see Price 1997) this problem is not present.
[11] The population interpretation will be dwelt with in future research.

to code it back into a program in a perfect manner, or to build (through crossover) a new feasible program from the combination of parts of parent programs. But if we think of the most simple form of programs which is 'Do This'-'Do That', this might not seem so difficult to accept - in this case one does not have to be able to look inside peoples' heads in order to do the trick. Besides thinking of simple programs we could consider a limited 'coding back' capability: this might give rise to an interpretation of a mutation as a 'recoding' mistake. Crossover (the recombination of partial programs) though, might give rise to unfeasible programs. In this case we could think of the rejection of those programs or alternatively of repair algorithms frequently used in optimisation with GA when such situations occur.

Given the model of the decision process outlined in 2.4 our population interpretation may be summarised as follows[12]:

begin
      $t \leftarrow 0$
      randomly generate a population P(t)
      determine aggregate results for P(t)
      $t \leftarrow t+1$
      while not(stopping criteria) do
            begin
                  for every agent j do
                      begin
                          assign credit to actions in P(t)
                          select an action a_{j1} from P(t)
                          if $r < prob\_cross_j$ then
                              begin
                                  select an action a_{j2} from P(t)
                                  crossover(a_{j1}, a_{j2}) to a'_j
                              end
                          for every bit in a'_j
                              if r< $prob\_mut_j$ then mutate that bit
                  end
                  determine aggregate results for P(t)
                  $t \leftarrow t+1$
              end
      end

With this reformulation the population interpretation might turn out to be less "simplistic" than it seems at first sight. Further, it might as well be taken as a mental interpretation. Note that if all the agents concentrate on the set $Å_t$ this set is a common element that constitutes part of the mental content of all agents, the other part being their possibly differing credit

---

[12] P(t) stands for the population in generation t, r is a uniformly distributed random number between 0 and 1; $prob\_cross_j$ and $prob\_mut_{j}$, the parameters that set the probability of crossover of a selected chromosome and the probability of mutation of each single bit for agent j.

assigning functions and operator parameters, and the corresponding evaluation set *S*'. In this interpretation, in some sense, society would be 'in the mind', and the scope for choice correspondingly socially delimited.

Whether or not this procedure can still be considered a GA is left as an open question. Meanwhile we use the designation 'population based simulation model' and we leave it open for further modifications to be introduced latter in this essay. It is not our intention to argue that this evolutionary procedure and the interpretation we give of it as a social process is the only legitimate or correct one. Other interpretations of the GA and other evolutionary algorithms might even lend themselves to richer and less unrealistic models. In particular other implementations of the mental interpretation of the GA, like Price's (1997) must be considered, and we intend to do so in the future. However, before moving on to more sophisticated models it may be worth exploring further the simpler ones, which, in spite of some strong assumptions, nevertheless seem to have a much more plausible behavioural content than the standard mathematical formalisms of the rational choice models. This exploration will start in the next section.

4. Institutions as solutions to recurring problems in social interaction

Recall one of the questions put forward in the introduction: What is the role of institutions and how are they related to the aggregated performance of the socio-economic systems? It is useful to think of institutions as 'tools' or 'mechanisms' that provide solutions to recurring problems in social interaction (Vanberg, 1994). With the assistance of game theory and experimental economics different types of recurring problem situations occurring in situations that involve no direct communication between agents may be identified. Their discussion may help us to address the question posed.

4.1 Co-ordination: problem 1

Let us take first a very simple co-ordination problem: "choose one of *n* colours; the payoff will grow with the number of players that will choose the colour you picked". If an experiment is carried out recurrently with the same players with choices and payoffs announced at the end of each repetition, we would most certainly observe that the choices would converge to one colour and stabilise there, even if the players were unable to communicate and to agree on a common strategy. Starting from a random choice of colours, every player would understand that he should choose the alternative that, by chance, turned out to be the most often practised. Once choices converge, no single player would have any incentive to move to another alternative. The process of emergence of the equilibrium solution could be described in this case as a spontaneous one - the rule "choose white" or

"choose black" would emerge and become a self-enforcing shared rule (convention) as if an invisible hand had driven the players to a social optimum.

This process may be simulated with the above specified model: each string stands for a choice_of_colour rule and each agent is endowed with the same function that assigns credit to the rules according to their absolute frequency in the population. The results of a typical run (100 players and 16 colours) show that (see figure 1), after an initial period where different colours still compete, the choices converge, as expected, to one of them, which turns out to be one of the most frequent random choices in the initial population. In 9 out of 10 runs of the simulation with different initial populations generated by different random generator seeds the procedure converged to one of the most frequent choices in the initial population and only once did it converge to the second more frequent. Convergence is not perfect due to mutation.[13]

4.2 Co-ordination: problem 2

Consider now another co-ordination problem. The players must now choose among the same *n* colours. The payoffs will depend on the number of players choosing the same colour, but also on the colour chosen. To co-ordinate on black (or on white) is better than any situation where choices are distributed, but to co-ordinate on black is better then to co-ordinate on white. In this case, as before, we could foresee that the players would co-ordinate. However it is much harder to be sure that they would always co-ordinate on the right colour. But once they co-ordinate on one colour (inferior or Pareto optimal), no individual player will have an incentive to move to a different one. If they converge to an inferior colour they might only came to discover the best solution if they all moved simultaneously, and the chance that this might occur spontaneously is small, in particular if the number of players is large.

Simulating this process (100 players and 16 colours) with credit assigned by the function, $payoff(colour_i) = (Abs.\ Freq.\ colour_i) \times (1 + 0.2i)$, we observe, as expected, a process of convergence. In the run reported by Figure 2 (see Appendix), choices that are Pareto inefficient (colours *b* and *e*) compete in the beginning of the simulation, and colour *b* that was not the most frequent in the initial population finally defeats colour *e*. Five out of ten runs of the simulation with different initial populations converged to the Pareto optimal choice, the remaining simulations converged to a Pareto inferior outcome.

Once again an invisible hand guides the agents, except that now it may lead them to a convention that is not the best of all possible. We may conjecture however, that if we let the players discuss and agree on a joint strategy the chances of co-ordination in a Pareto optimal outcome would increase.

---

[13] Even though in this particular context mutation is hardly justifiable in rational terms, note that instability in convergence is also observable with human beings in similar experimental contexts; boredom or insufficient understanding of the game situation is usually the explanation given by the reseachers.

4.3 The invisible-hand, path-dependency and the information content of $\mathring{A}_t$

In both instances an institution - a convention - emerged as a non-intended result of the interaction of purpose-seeking agents in a typical 'invisible-hand' process. These are interesting processes that could be extended to more complex multiple equilibrium game situations where convergence depends on the existence of 'focal points', but they are not the end of the story.

Both processes are instances of path dependency - a time irreversible outcome arises out of multiple possible causes including chance - and may be related to many processes that are observable in the real world. The first situation is often given as a possible account for the emergence of the 'keep to the right'- 'keep to the left' traffic regulation, the second was dwelt on by Davis (1985) in his famous story of QWERTY and more extensively by Brian Arthur (1994). Both cases have been described as instances of institutional lock-in.

The simulation of these processes may shed some light onto the 'rule-following'-'purpose-seeking' debate. It shows that convergence translates into a loss of diversity in the population of rules. While the number of rules in the population remains the same, the proportion of some rules tends to increase until finally a single rule dominates the whole population (except for occasional mutations). That is, the information content of the population is decreasing as the procedure converges. This means that even though the agents are purpose-seeking 'choosers' their scope for choice tends to get smaller and smaller as the population converges, that is, as a rule emerges and becomes shared.

4.4 Co-operation

A third experimental situation may lead us further: A set of individuals, kept in isolation from each other, must post a contribution (from $0 to a pre-defined maximum) in an envelope, announcing the amount contained in it; the posted contributions are collected, summed up by the experimenter and 'invested', giving rise to a collective payoff that must be apportioned among the individuals; the apportioning rule instituted (known to the agents) stipulates that the share of the collective payoff must be proportional to the announced contributions (not to the posted contributions); the posted contributions and the corresponding announced contributions are subsequently made public (but not attributed to individuals); individual returns on investment are put by the experimenter into the corresponding envelopes and the envelopes are claimed by their owners.

This situation is related to the problem of team production (Alchian and Demsetz, 1972) but it also has common features with instances of public goods provision and collective action (Olson, 1965) and it can, in fact, be generalised to all situations where external effects are

strong. For game theorists it is said to have a N-person repeated Prisoner's Dilemma (PD) structure.

Can we predict what is likely to happen after a number of repetitions of the experiment with the same experimental subjects? Ledyard (1995:112) answers: "There are many theories. One, the economic/game-theoretical prediction, is that no one will ever contribute anything. Each potential contributor will try to "free-ride" on the others. [...] Another theory, which I will call the sociological-psychological prediction, is that each subject will contribute something [...] it is some times claimed that altruism, social norms or group identification will lead each to contribute [...*x*...], the group optimal outcome. [...] Examination of the data reveals that neither theory is right". As a matter of fact, the experimental evidence in similar cases shows that, with large groups, positive posted contributions are observable in the first rounds but free-riding soon emerges leading the group to levels of contribution that all agents consider undesirable. The question therefore is: What might be wrong with the 'economic/game-theoretic' and with the 'sociological-psychological' models?

4.5 The unidimensional mind

"The first principle of economics is that every agent is actuated solely by self-interest", wrote, Edgeworth (1881: 16) long ago. This is still the principle on which the game theoretical/economic standard model is founded. Even though this 'first principle' may seem crystal clear at first sight, the discussion soon becomes confusing as the question of what might be considered to be the interest of the agent is posed: his own well being? the well being of his family? his neighbours? his country? ... In fact, if the scope of self-interest is indefinitely extended, any act, as 'altruistic' as it may seem can be interpreted as self-interested (or even egoistic). A game of words that leads nowhere follows, turning Edgeworth's 'first principle' into a tautology. Jevons (1871: 25), however helps us understand exactly what the marginalists (including Edgworth) had in mind:

> "As it seems to me, the feelings of which a man is capable are of various grades. He is always subject to mere physical pleasure or pain [...]. He is capable also of mental and moral feelings of several degrees of elevation. A higher motive may rightly overbalance all considerations belonging even to the next lower range of feelings; but so long as the higher motive does not intervene, it is surely both desirable and right that the lower motives should be balanced against each other. Starting with the lowest stage - it is a man duty, as it is natural inclination, to earn sufficient food and whatever else may best satisfy his proper and moderate desires. If the claims of a family or of friends fall upon him, it may become desirable that he should deny his own desires and even his physical needs their full customary gratification. But the claims of a family are only a step to a higher grade of duties.
> The safety of a nation, the welfare of great populations, may happen to depend upon his exertions [...]. Motives and feelings are certainly of the same kind to the extent that we are able to weigh them against each other; but they are, nevertheless, almost incomparable in power and authority.
> My present purpose is accomplished in pointing out this hierarchy of feeling, and assigning a proper place to the pleasures and pains with which the Economist deals. It is the lowest rank of feeling which we here treat. [...] Each labourer, in the absence of other motives, is supposed to devote his energy to the accumulation of wealth. A

higher calculus of moral right and wrong would be needed to show how he may best employ that wealth for the good of others as well as himself. But when that higher calculus gives no prohibition, we need the lower calculus to gain us the utmost good in matters of moral indifference."

This long quotation couldn't be avoided because it makes two very clear points: (a) Economics was supposed to deal with the "lowest rank of feeling" under the assumption of the absence of motives arising from any "higher ranks". Since this assumption cannot make sense in a situation of interaction where the consequences of our actions fall not only on our heads but on everybody else's, Jevons could be expected to say that all other situations (including therefore a large section of the subject matter of game theory) are out of the scope of the economic man model. When problems of interaction with a PD structure are present there are "no matters of moral indifference"; (b) The 'feelings' pertaining to different levels are incommensurable, they are "almost incomparable in power and authority": in no way can the hierarchy of feelings be aggregated in a single utility function.

If Jevons was right, if his type of economics was not fit to deal with most situations of interaction, he may have unintendedly discovered another source of trouble of the economic man besides omniscience. To extend Edgeworth's concept of self-interested action beyond "matters of moral indifference" may only lead to paradox. Concerning the experimental evidence, the standard economic/game theoretical approach must be unable to explain why positive contributions are observed in the first rounds of experiments: if I am self-interested, in the sense that I disregard the higher order obligation of contributing to collective goals and the prohibition of not telling the truth, and if I know that all the others disregard it in the same way, why should I be contributive and truthful, bearing the costs alone, and having a benefit that is disproportional to my contribution?

But the unidimentional mind is also present, although differently, in alternative accounts of human action. In what might be called a standard funcionalist sociological explanation for the fact that individuals tend to behave in accordance with social norms, the emphasis would be on *socialisation*, "a process in which, through (positive and negative) sanctions imposed by their social environment, individuals come to abide by norms", and that leads to *internalisation*, "according to which a person's willingness to abide by norms becomes independent of external sanctions and, instead, becomes part of a person's character" (Vanberg, 1994: 14). In reference to the experimental evidence *internalisation* might, in fact, explain the positive contribution in the first rounds. But would it explain the breakdown of the posted contributions observed with repetition?

There is at least one problem with the functionalist explanation: "By invoking at the same time, through the concept of sanctions, that people respond to incentives and, through the notion of internalisation, that their rule compliance is unresponsive to incentives (...) [it] seems to be based on two incompatible conceptions" (Vanberg, 1994: 14). In fact, it is hard to accept that the "willingness to abide by norms" is a second nature, a parameter in the agent model. If place is given in the model for 'deviant behaviour' (as it is in a functionalist explanation), it is hard to believe that those who comply with social norms will continue to

do so when they are confronted with a situation where 'mutant' behaviour is rewarded. Would the socialised agents continue to contribute and tell the truth even when the deviant behaviour is giving rise to higher payoffs?

4.6 The hierarchy of feeling

If economics is to deal with situations of interaction where strong external effects are present, the model of man will have to be reconsidered. When "other motives" matter complying with shared rules may be in itself a motive behind a choice. The agents may be endowed with a *moral disposition* (Vanberg, 1994) that drives them to behave in accordance with the rules that are *believed* to sustain the group existence. Guilt and shame are self-experienced sanctions for rule violation, self-respect is a powerful reward, and the importance of these *emotional* factors is such that self-sacrificing behaviour is not an uncommon feature of human behaviour.

However, even if we accept that a moral disposition has to be taken into account, this propensity cannot be taken as a fixed parameter in the agent's model. The moral disposition is not imprinted once and for all, and in the same degree to all agents, by the process of socialisation. Since a shared rule can only produce the expected benefits if it is generally abided by, it may be pointless not to violate it if most others do[14]. The moral disposition is therefore a variable in two senses: it varies from individual to individual, and it tends to be strengthened with rule compliance and weakened with the spreading of deviant behaviour.

4.7 The model with agents endowed with a moral disposition

Let us recall the experimental situation described in 4.4. In each period of time an agent must decide on his actual contribution and on his announced contribution. We now assume a set shared rules: 'thou shall not lie' and 'thou shall contribute to the collective goal'. To take on board the moral disposition factor, amounts to an important modification in the basic model that has been used. This is implemented by coding one part of the 0/1 string (chromosome) as *announced contribution* and the other part as *moral disposition*; the *announced contribution* part of the string may decode into a real number between 0 and 50, and the *moral disposition* part to a real number between 0 and 1 [15].

The *posted contribution* of agent *i* to the collective investment fund is given by:

(A1) $posted\ contribution_i = announced\ contribution_i \times moral\ disposition_i$

---

[14] In Lisbon, car drivers park their car on the sidewalks. We find ourselves doing the same even though we strongly agree with the rule that forbids it. What would be the point of not parking the car in this manner if the sidewalks would continue to be filled with cars?

[15] In all the simulations that follow the *Population Size* is 20, the *Probability of Crossover* is 0.5 and the *Probability of Mutation* is 0.01, for all agents.

The collective return on investment is given by:

$$(A2) \quad collective\ return = 10 \times \sum_i posted\ contribution_i$$

The apportioning rule is:

$$(A3) \quad return_i = \frac{announced\ contribution_i}{\sum_i announced\ contribution_i} \times collective\ return$$

The credit assignment function (the same across agents) used to evaluate the actions is:

$$(A4) \quad credit_i = return_i - posted\ contribution_i$$

The collective payoff is given by:

$$(A5) \quad collective\ payoff = \sum_i return_i - \sum_i posted\ contribution_i$$

The results of a typical[16] run of the simulation are shown in figure 3 (see Appendix). Until the announced contributions reach their maximum level, the posted contributions increase as well, after this they rapidly tend to zero, while the announced contributions are kept close to the maximum value. Due to the existent incentive to free-riding the initial moral disposition tends to erode with time. As a result the collective payoff deteriorates reaching the zero level around generation 120. After this, only occasional mutations (that might be interpreted as signalling intentions to contribute conditional to the contribution of others) disturb the scenario of collective disaster. The results are therefore consistent with the available experimental evidence mentioned in point 4.4: positive contributions are observed in the first rounds of the experiment but free-riding tends to emerge leading the group to very low levels of contribution.

But perhaps our simulation was biased by the fact that the agents were endowed at the outset with different levels of *moral disposition*. Potential free-riders were therefore included from the beginning in the group. 'Not to lye' would after all be something we all tend to be socialised into. Let us then see what happens if all the agents are created with the highest possible *moral disposition* while mutations that change this original disposition of the agents will continue to be allowed. With this change in the mental setting of the agents (which incidentally is something that an experimental economist is unable to do) the simulation was run. The results reported in figure 4 (see Appendix) reveal that the same overall pattern may be observed. 'Mutant' behaviour quickly invades the population, in spite of the original propensity to tell the truth.

---

[16]Different initial populations were generated using various random generator seeds. The observed overall pattern of outcome is common to all.

It seems therefore that in this context no viable social order will spontaneously emerge and be reproduced in time, the group would simply perish. ... And yet collective action exists in the real world.

5. Alternative solutions to the problem of collective action

5.1 Anarchic Order

Game theorists and other social researchers have invested a huge effort in trying to show that a social order might spontaneously come to existence and be reproduced without the enforcement of social norms by a coercive entity of some kind. In the previous section the first two examples on co-ordination showed how in the context of the present simulations, such a situation might occur, other examples, and possibly more interesting, might extend this result to a wider range of situations. In both cases there were multiple possible equilibria, the final outcome depended both on the initial conditions set for the simulation, the payoff structure and on a multitude of random mutations that might occur in the first generations of a run. Once the process converged it became locked-in. In those situations, an order prevailed that may be considered self-enforcing or anarchic. The interesting open question concerns the second example, in which the simulation converged to a Pareto inferior outcome: If somehow the agents (or some agents) came to suspect that there might be a better outcome, wouldn't they do something about it? Can we imagine in such a situation a *spontaneous* move from one equilibrium to the next, so that the different possibilities are exploited?

The third example led to a more pessimistic conclusion: in a situation with a PD structure the invisible hand would lead the group to disaster even if a conditional moral disposition is assumed. To generalise this result, however, might be rushing to conclusions. Game theorists after all showed that even self-interested individuals would have reasons to co-operate if the PD was indefinitely repeated. However, they never generalised this result to N-person games. Axelrod, however did something closer to that. In his simulations where strategies were randomly matched for rounds of PD in each generation of the GA the nice TIT-FOR-TAT strategy spontaneously became dominant in the population. The "emergence of co-operation" was therefore the appropriate title for his research report. But should we take Axelrod's simulations as proof of the possibility of an anarchic order? Commenting Axelrod results Binmore (1997) reports on an interesting study by Probst that ran Axelrod's simulations for not only 50 generations (like the pioneering author) but for much longer periods. Among other findings he concluded that TIT-FOR-TAT in the long run would be defeated by the vicious TAT-FOR-TIT (a strategy that begins by defecting and continues to do so until the opponent also defects, turning then to co-operation). For Binmore this suggests that in the end mean automata might as well triumph...

The possibility of an anarchic order remains therefore open for speculation. It is not completely ruled out, whether we approach it in game theoretical terms, or from the historical

and anthropological record, but it is far from having been proven convincingly. Meanwhile we turn to an alternative solution that is more familiar.

5.2 The Hobbesian solution

Hobbes's argument on the need for a social contract and an enforcing sovereign power has been translated to modern terminology by game theorists: "the Hobbesian argument essentially turns on the claim that the problem of political obligation can only be solved by the creation of a co-operative political game, instead of the non-co-operative game played in the state of nature" (Heap, 1992: 203). Co-operative games are based on pre-play negotiation and binding agreements. In these terms the social contract would be the result of pre-play negotiation and the presence of Hobbes's sovereign the condition to make this contract binding to all.

5.2.1 Extending the model

Avoiding for the moment the complications of the collective negotiation process, we will extend our simulation model introducing a meta-agent with monitoring and sanctioning powers, while we continue to take as given the remaining institutional frame. Our only present aim is to test how the system would behave once the monitoring meta-agent is introduced.

The introduction of the meta-agent amounts to a change in the setting of the experiment described in 4.4: the experimenter (the meta-agent) may now decide to open some (or all) of the envelopes when they are handed to him. If an agent is found to have announced an amount that does not correspond to his contribution he will be sanctioned. His return on investment will now be determined by the rule,

$$\text{(A3')} \quad return_i = \frac{posted\ contribution_i}{\sum\limits_i posted\ contribution_i} \times collective\ return \times moral\ disposition_i\ ,$$

with the implicit penalty reverting to the meta-agent and therefore being included in the collective payoff. The meta-agent chooses the individuals to be inspected by a simple rule: if r (a random real between 0 and 1) is lower than *probability of monitoring* (a parameter of the simulation) then agent i's envelope will be opened.

The results of two runs of the simulation, with probability of monitoring set to 1 (all agents inspected) and the same probability set to 0.5 are shown in figures 5 and 6. For the first run, figure 5 shows that by generation 20 the maximum value for the posted and announced contributions is quickly reached and kept thereafter with some fluctuations. This means that the selective pressures exerted by the monitoring meta-agents successfully counteract free-riding and induce high moral disposition levels. The results with probability of monitoring set

to 0.5 (see figure 6) show that full monitoring (which in reality may be impossible to implement) may even be unnecessary to ensure that co-operation is kept with 'sufficient' levels of the moral disposition.

Bearing in mind the Hobbesian model, Edmund Burke (quoted in Vanberg, 1994:41) moralised: "Men are qualified for civil liberty in exact proportion of their disposition to put moral chains upon their own appetites... Society cannot exist unless a controlling power upon will and appetite be placed somewhere, and the less of it there is within, the more there must be without". With these simulations he might have added ... but the moral chains within only seem to be effective if a power without is placed somewhere. However, to conclude in this moralist vein would be too hasty. The Hobbesian solution may be realistic, but is it the best of all words? There may be other sides to the coin of Hobbes' story.

5.3 Problems for Hobbes (1): mode of monitoring

The existence of a group in a world inhabited by multiple interactive groups can be viewed as a process of constant exploration in search for niches to be exploited. Niches emerge and disappear, in an open-ended process of 'creative destruction'. In this world a tension exists between exploitation and exploration: while the exploitation of a niche may require co-ordination within the group, exploration may require that individuals are granted a freedom of action that includes the possibility of making mistakes and performing choices that may be damaging to the group. Is it possible to model this tension?

Consider the following situation. A group is jointly exploring and exploiting a landscape; exploitation requires a strict co-ordination of the group - if the agents co-ordinate, choosing to act together at a single point in the landscape, the payoffs will be large; if they choose to act separately dispersing themselves on the landscape the payoffs will be much lower. However, to each and everyone it is much better to co-ordinate in some particular unknown points of the landscape than others. An exploration problem is therefore present: the group must discover the best niches in the landscape. Each agent has a private benefit dependent on the spot he chooses. The group has a collective cost that decreases as co-ordination increases. A shared rule may establish that the collective cost is apportioned in equal parts among the group members so that the net payoff of an agent is given by the difference between the individual payoff and the individual share of the collective cost. In this case an individual may be better off if he decides to move away from the group: the variation of his individual payoff may be higher than the variation of his share of the collective cost generated by his choice, even though for the group the situation would worsen. We can admit, therefore, that a new institutional disposition is added: a meta-agent is created to monitor the deviants and impose sanctions, so that no agent would have an incentive to move away from the group.

The model takes the following elements as given:

The action of each agent consists of the choice of a spot represented by an integer between 0 and 100;

The private benefit depends on the spot chosen and is given by,

(B1) $private\ benefit_i = 3000 + 40 \times (spot_i + a) - (spot_i + a)^2$,

where $spot_i$ represents the spot chosen by agent i:

The collective cost depends on the level of co-ordination, represented by a measure of distance among chosen spots and is given by,

(B2) colective $cost = \sum_i |centre - spot_i|$, with $centre = \dfrac{\sum_i spot_i}{population\ size}$

The *apportioning rule* establishes that private benefits remain private and the collective costs are equally shared among agents;

A meta-agent may sanction the agents that deviate from the group imposing on them a penalty that is proportional to the distance of the agent's chosen spot from the centre of the distribution of the group; the monitoring may be more or less frequent as in the previous model;

The credit assignment function for all agents is given by,

(B3) $fitness_i = private\ benefit_i - \dfrac{colective\ cost}{population\ size}$, for non-monitored actions,

and by:

(B4) $fitness_i = private\ benefit_i - 250 \times |center - spot_i|$, for monitored actions.

The purpose of the simulation is to test the effect of monitoring on the capability of the group to explore the landscape, and its adaptability to sudden changes in the world.

The results of three runs of the simulation (where α was set to 0 until the generation 50 and reset to -60 after that) show that the capability of the group to explore the landscape and to exploit it by co-ordination is, in fact, heavily dependent on the *probability of monitoring*. With α=0 the maximising choice for all agents (and for the group) would be to co-ordinate on spot 20 where the private and collective benefits would be maximised; with α=-60, the maximising spot would be 80.

Figure 7, where agents are represented by different markers, reports on the history of the simulation with no monitoring. In the initial generations the agents tend to be dispersed: until

generation 50 they are distributed around the maximising spot 20, after this they move, choosing spots around 80. In this case, exploration dominates exploitation; in spite of heavy collective costs due to insufficient co-ordination, the group is successful in discovering the appropriate niche and in adapting to changes in the environment. When the probabilily of monitoring is set to 1 (see figure 8) the group remains close together since all deviants are heavily sanctioned. The collective costs are low but the group is not able to discover the existing niches and to adapt - exploration dominates exploitation. The results obtained when the probability of monitoring is set to an intermediate level (see figure 9) suggest that there may be a sufficient level of monitoring that allows at the same time co-ordination and exploration.

The implication for the Hobbesian story seems to be that the mode of behaviour of the meta-agent is not neutral in relation to the performance of the group - a social order may not be judged beneficial for the simple fact that co-ordination or co-operation are being enforced. Different modes of enforcement lead to different outcomes and some of them may be judged more desirable then others.

5.4. Problems for Hobbes (2): the origin of shared rules

The easiest explanation for the existence of the shared rules, and the first to be 'discovered' by mankind, is that they were created in the mind and enacted by some kind of real or virtual meta-agent. If this primitive line of explanation is excluded, a second one may be contemplated: they exist because they are *functional* to society or the group. The functionalist approach is problematic because it involves the explanation of a cause (shared rules as the cause of stable behavioural patterns) by its effects (the beneficial effects to the group or to society) (Gilbert, 1995). However it could be reformulated in evolutionary terms, and Hayek's thoughts on the evolution of institutions seem to have developed along this line.

"'Learning from experience'", Hayek writes, is a process of "transmission and development of practices which have prevailed because they were successful - often not because they conferred any recognisable benefit on the acting individual but because they increased the chances of survival of the group to which he belonged". But he also writes: "The rules of conduct have [...] evolved because the groups who practised them were more successful and displaced others" (Hayek, 1973: 18). Hayek's account seems contradictory. On the one hand, it seems to rely on a spontaneous learning process internal to the group that would lead to the emergence and dominance of collectively beneficial shared rules. On the other hand, having himself doubts on the first type of spontaneous process, he assumes a competition between groups that would drive out the inefficient.

Alternative explanations for the origin of shared rules can be built from the development of the following clues: (a) Intelligent individuals are able to recognise and evaluate the effects of shared rules; (b) These intelligent individuals may even formulate theories that enable them to predict the outcome of alternative sets of shared rules and modes of meta-agency and

formulate preferences over these outcomes (c) they may further engage in tacit or formal 'agreements' about institutional arrangements that by influencing individual choices, ensure the group's existence.

In order to elaborate on these points, Buchanan's and Vanberg's (Vanberb,1994:21) distinction between *action interests* and *constitutional interests* is crucial. The action interest concerns personal situational choices within a set of alternatives. The constitutional interest is related to shared rules and modes of meta-agency and may be defined as the individual's interest "in seeing a certain rule implemented in a social community within which he operates" (Vanberg, 1994:21). An intelligent human being articulates action interests and constitutional interests. Different levels in the mind correspond to different levels of rules, and there is no reason to assume a permanent harmony or a permanent conflict between these different levels. They are in harmony when the individual in a given situation chooses to act in accordance with a rule he favours constitutionally; they are in conflict when he chooses to violate that rule. It is not irrational to believe in a constitutional rule and not behave accordingly. Most tax evaders agree with taxes and the provision of public goods; we may consider tax evasion as immoral, but not as irrational. As Vanberg (1974: 62) stresses when a convergence of constitutional interests prevails the key issue of social order consists of the correspondence of constitutional interests and action interests. The key issue is the moral disposition.

Even though the convergence of constitutional interests is a possibility, it seems more interesting to consider the existence of groups where a constitutional framework prevails without what might be defined as a voluntary agreement. In fact it is not difficult to accept that one may *submit* to an order even if his constitutional interests conflict with it. In this case, the negotiation over the constitutional order is an ongoing process, a permanent game with temporary winners and losers in which *power* counts. Morality can still be defined as a disposition to act in accordance with constitutional interest, but (since the constitutional interest may not converge) not as action in accordance with the prevailing shared rules. Institutional change may now be explained not only as a result of constant adaptations in the individual minds, but as a result of changing balances of power. Rule enforcement, now that domination enters the picture, can no longer be thought of as a neutral prerogative of a meta-agent. The Hobbesian story becomes a little more complex. Some of these implications may be explored with further modifications in the simulation model.

5.4.1 Co-operation revisited

In the model presented in 5.2.1 an apportioning rule was assumed: an unspecified deliberation process had led to the enactment of that rule. The model concentrated on the emerging behaviour of the agents under different selective pressures from a meta-agent and the corresponding aggregated outcomes. We are now interested in modelling: (a) the process through which the constitutional preferences of the agents may change as a result of their experience of the aggregated outcomes; (b) how these changes may lead to the replacement of constitutional regimes, and (c) how the distribution of power within the group may be related with to constitutional design and (d) how the constitutional regimes are related to the group welfare.

The following changes were introduced in the experimental setting and in the model:

Instead of announcing a contribution, the agent is now expected to abide by a minimum level of contribution (say 35). The meta-agent as in 5.2.1 may decide to check if the envelope contains the minimum specified. If not the agent will be penalised.

In the model of the agent a variable *size* is introduced representing the agent's *power*, which, in the context of this model, is related to the greater or lesser weight of each agent in the decision process that leads to the adoption of an apportioning rule, and (depending on the apportioning rule) may influence the size of each agent's share of the collective benefits.

The agent's model includes a rule for *contribution* and *values* assigned to two alternative *apportioning rules* that allow the agent to choose among them;

The meta-agent enforces the apportioning rule chosen by a voting procedure in which an agent's *size* determines the weight of his vote, and it monitors the agents' contributions.

The *value* assigned by each agent to the constitutional rules is updated in every generation and is given by the agent's average individual pay-off under each rule's regime. The agent will vote on the rule with greater value to him.

The *size* of the agent is updated assuming that in each generation a part of the individual payoff is 'capitalised'.

In the instance to which the results that will be presented are referred the collective return on investment is given by equation A2. The apportioning rule 1 is given by,

$$(\text{C1}) \quad return_i = \frac{contribution_i}{\sum_i contribution} \times collective\ return$$

and the apportioning rule 2 by,

$$(\text{C2}) \quad return_i = \frac{size_i}{\sum_i size_i} \times collective\ return$$

Under both regimes the credit of action i to agent j is assigned by the function,

$$(\text{C3}) \quad credit_{i,j} = return_{i,j} - contribution_i, \text{ for non-monitored actions, and by,}$$

$$(\text{C4}) \quad credit_{i,j} = return_{i,j} - contribution_i - 10 \times (35 - contribution_i), \text{ if the action}$$
is monitored and $contribution_i < 35$.

Note that, under rule's 2 regime, agent j evaluates action i as though he himself has performed it. That is, using $size_j$ to compute the return of action i.

The collective payoff is,

$$(\text{C5}) \quad collective\ payoff = \sum_i return_i - \sum_i contribution_i$$

Size is updated according to the following rule:

$$(\text{C6}) \quad size_{i,t+1} = size_{i,t} + \frac{return_{i,t}}{100000}$$

The simulation includes a training period of 100 generations during which no voting takes place and that is used by the agents to experience the regimes of the two rules, assigning values to them. Rule 2 is experienced in the initial fifty generations and rule 1 in the next fifty. In generation 101 and every 20 generations after that a vote takes place that may change the rule regime.

- Simulation 1 - all agents created equal:

The agents are created with size 10 and the probability of monitoring set to 0.9. The results show (see figure 10) that in the first fifty generations (under rule 2) the total contributions and collective payoffs tend to decrease after the initial adjustment: to apportion benefits in a way that is not proportional to contributions leads, as we have seen, to an inefficient outcome if full monitoring is not possible. After generation fifty (under rule 1) the contributions and payoffs start to recover reaching values that are close to maximum. After generation 100, when voting starts, there is unanimity on rule 1. Figure 11 (see Appendix) shows that this

unanimity is kept till the end of the run - the number of votes on rule 1 increases, as the size of the agents increases, but the number of votes on the same rule remains constant. Rule 2 has zero votes and zero voters all through the simulation.

-    Simulation 2 - a more realistic distribution of power

The size of each agent is now randomly generated varying between 0 and 20. The results (see figures 12 on the Appendix) show that the comparatively bad performance of the fifty first generations (under rule 1) tends to improve, between generations 50 and 100. However in generation 101, when it comes to voting, rule 2 wins - rule 2 has a majority of votes even though it does not have a majority of voters (see figure 14). In fact rule 2 performs well for large agents and badly for small ones - the correlation between the *value* of rule 2 and *size* is almost perfect (see figure 15). From this point on the overall pattern of the collective payoffs and contributions is inefficient and rather unstable. Figure 13 shows that the biased apportioning rule 2 generates different growth rates, increasing the size inequalities within the group.

These results are dependent on the level of monitoring. If high levels of monitoring are costly or impossible for some reason, the regime of rule 2 may not be sustained. The results with a lower level of monitoring (see figure 16) show that the contribution levels generated by the biased apportioning rule and insufficient monitoring may lead to the replacement of rule 2's regime. In this simulation (see figure 17), in generation 300 a new majority of small agents successfully imposed rule 1.

The implication suggested by this model is that a rule regime that is not functional to the group may exist and be reproduced in time if the model includes an unbalanced power distribution and high levels of monitoring. It also suggests however that this order which is not based on consent and agreement may be a temporary one. Hobbes order may after all not be the best of all worlds.

6. Conclusion

We started from two initial interrogations: What is the role of institutions and how are they related to the aggregated performance of the socio-economic systems? What is the relation between institutions and the behaviour of the agents? Viewing institutions as 'tools' or 'mechanisms' that provide solutions to recurring problems in social interaction, we explored different environments that have in common the anonymous nature of the interaction and the absence of direct communication between agents.

Concerning the first group of questions, we conjectured that under some circumstances conventions might spontaneously come to exist that provide the basis for social life. These conventions would be self-enforcing; that is, they would be reproduced by the practices of the agents without the interference of any 'external' meta-agency. However, we were led to the conclusion that even in this case no support could be given to the deep rooted belief in a 'naturally' beneficial outcome generated by an invisible hand - the emerged convention could lead to a Pareto inferior result.

We moved then to contexts where the spontaneous processes seem to generate outcomes that are not compatible with the existence of the group. After confronting this possibility with optimistic prospects of spontaneous anarchic order, we moved on to the more familiar reality of a social order based on a contract and an enforcing meta-agency subject to the collective choices of the group. However, the visible hand of meta-agency and the underlying social choice also seem to be no guarantee of efficiency: the mode of meta-agency and the constitutional rules that it enforces affect in a critical way the performance of the group. Worse, the evaluation of the constitutional regimes by the agents might differ, depending not only on the aggregate results, but on the particular situation of an agent within society. The results of simulation with this model suggested that the social choice over alternative constitutional rules, when biased by the distribution of power within the group, may lead to rule regimes that even though generating inferior outcomes for the group, may (or may not) be sustained in time through coercion.

Addressing the second group of questions, we started from a model of a bounded rational agent living in society. This model served us well until we had to consider situations where an agent must balance his situational action interest with his constitutional interest. That led us to consider the possibility of a hierarchy of feeling: agents endowed with a moral disposition may choose to act in accordance with shared rules. But we were forced to conclude that when the incentive structure favours free-riding the moral disposition may not be sustainable. However, intelligent agents may be able to recognise the links between their personal fates and the group welfare, and settle on 'agreements' which, in lacking "moral chains within" may be enforced by "a controlling power, without". For this, a new extension of the model of the bounded rational agent was necessary in order to include preferences and choices over alternative shared rules, and a social choice procedure.

To what extent may these conjectures be useful? Do they enhance in any way the predictive power of economics? Or rather, the relevant question is: should we expect economic theory to have a predictive power? Socio-economic evolution is an open ended and a largely undetermined process; looking at society with the glasses of a naturalist that searches for laws of motion may be hopeless. Society is made of a clay that has a will of its own, it involves an irreducible innovative and partially uncaused element - human (individual and collective) choice. Far from leading to a relativist stance this perspective conveys a strong meaning to the role of social theory -even if it will not allow us to predict, it may helps us to "recognise, reason about and react to human institutions", guiding "our efforts to improve the social order in which we live". No claim is made for a value-free social theory - what is meant by improvement of the social order is in the last resort a matter for individual judgement and choice.

Our wishes concerning the social order must be bounded by the what we perceive as constraining 'facts'. The following may be crucial: (a) The human being is multi-dimensional: a moral disposition exists that may drive the agents to act in accordance with the shared rules of the group; But (b) human beings (as economists have always insisted) respond to incentives; no moral disposition can therefore be taken for granted; no social order is sustainable that is built on shared rules that are not supported by the corresponding rewards and sanctions; (c) In the absence of a spontaneous correspondence between incentives and shared rules, a social order may be 'created' by the agents including an 'external' meta-agency; But, (d) the mode of meta-agency and the enforced shared rules are not neutral in relation to the distribution of power within society.

In this research simulation was instrumental. It was used as a device to check the consistency of the models and to derive their implications. It imposed a discipline on the modelling activity and very often provided insights and gave rise to new ideas to be explored (in particular when the simulator was confronted with unexpected results). The implemented simulation model served its purposes well but it remains open for new extensions to be tested in the future, for instance, situations of interaction of multiple populations, and agents with different preference structures. We are aware of the fact that it relies upon a number of restrictive behavioural assumptions: the agents are supposed to be informed of all the actions performed in a given time period, they are supposed to be able to assign credit to these actions (even though some of these were not directly experienced by them), and they must be able to code an observed action back into a 'program'. To the extent that some of these assumptions may be relaxed by implementations of the mental interpretation of the GA and by other Evolutionary Algorithms we intend to explore them in the future, in particular when face-to-face interaction involving communication will be considered.

References

ALCHIAN, Armen A. and Demsets H. (1972), "Production, Information Costs, and Economic Organization". *American Economic Review*, vol. 62, December 1972, pp 777-795.

ARIFOVIC, Jasmina (1991), *Learning by Genetic Algorithms in Economic Environments*. Doctoral Dissertation, Department of Economics, Chicago: University of Chicago.

ARTHUR, Brian (1994), *Increasing Returns and Path Dependence in the Economy*. Ann Arbor, Michigan: University of Michigan Press.

CALDAS, Jose C. and Coelho H. (1994), "The Simulation of Trade in Oligopolisttic Markets". In Doran J. and Gilbert, N. (Eds.) (1994).

CHATTOE, Edmund (1994), "The Use of Evolutionary Algorithms in Economics: Metaphors or Models for Social Interaction?. In E. Hillebrand and J. Stender (Eds.), *Many-Agent Simulation and Artificial Life*, Amsterdam: IOS Press.

CHATTOE, Edmund (1998), "Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes". *Journal of Artificial Societies and Social Simulation (JASSS)*, vol. 1, no. 3, <http://www.soc.surrey.ac.uk/JASSS/1/3/2.html>.

DAVIS, Paul (1985), "Clio and the Economics of QWERTY". *American Economic Review*, n. 75, pp. 332-337.

DORAN, Jim and Gilbert, N. (Eds.) (1994), *Simulating Societies: The Computer Simulation of Social Phenomena*. London: UCL Press.

EDGEWORTH, F.Y. (1881), *Mathematical Psychics*. London: Kegan Paul.

EPSTEIN, Joshua M. and Axtell, R. (1996), *Growing Artificial Societies: Social Science From the Bottom Up*. Cambridge, Massachusetts: The MIT Press.

FABER, Malte and Proops, John, *Evolution, Time, Production and the Environment*. Third Revised and Enlarged Ed. Berlin: Spriger.

FERBER, Jacques (1995), *Les Systemes Multi-Agents: Vers une Intelligence Collective*. Paris: InterEditions.

GIDDENS, Anthony (1984), *The Constitution of Society: Outline of the Theory of Stucturation*. First paperback edition, Cambridge, UK: Polity Press, 1986.

GILBERT, Nigel (1995), "Simulation: an emergent perspective". 7/28/98,

GOLDBERG, David E. (1989), *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, Massachusets: Addison-Wesley.

HAYECK, Friedrich A. (1973), *Law Legislation and Liberty*, Vol. 1 - *Rules and Order*. Chicago: The University of Chicago Press.

HEAP, S.H., Hollis, M., Lyons, B., Sugden, R., and Weale A. (1992), *The Theory of Choice: A Critical Guide*. Oxford UK: Blackwell.

HODGSON, Geoffrey M. (1988), *Economics and Institutions: A Manifesto for a Modern Institutional Economics*. Cambridge, UK: Polity Press.

HODGSON, Geoffrey M. (1993), *Economics and Evolution: Bringing Life Back Into Economics*. Cambridge, UK: Polity Press.

HOLLAND, John H. (1975), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. Second edition, Cambridge, Massachusetts: The MIT Press, 1992.

HOLLAND, John H. (1995), *Hidden Order: How Adaptation Builds Complexity*. Cambridge, Massachusetts: Addison-Wesley.

HOLLAND, John H., Holyoak K.J., Nisbett, R. E., Thagard, P.R. (1986), Induction: Processes of Inference, Learning and Discovery.

http://www.soc.surrey.ac...rch/simsoc/tuturial.html.

JEVONS, W. S. (1871), *The Theory of Political Economy*. Reprints of Economic Classics, New York: Augustus M. Kelley, 1965.

LEDYARD, John O. (1995), "Public Goods: A Survey of Experimental Research". In John Kagel and Alvin Roth (Eds.), *The Handbook of Experimental Economics*, Priceton, New Jersey: Princeton University Press.

LUCAS, Robert E. (1987), "Adaptive Behaviour and Economic Theory". In Hogarth, Robin M. and Melvin Reder (Eds.), *Rational Choice: The Contrast Between Economics and Psychology*. Chicago: University of Chicago Press.

MACY, Michael W. (1998), "Social Order in Artificial Worlds". *Journal of Artificial Societies and Social Simulation (JASSS)*, vol. 1, no. 1, <http://www.soc.surrey.ac.uk/JASSS/1/3/2.html>.

NORTH, Douglass C. (1990), *Institutions, Institutional Change and Economic Performance*. 5th edition, Cambridge: Cambridge University Press.

OLSON, Mancur (1965), *The Logic of Collective Action*. Cambridge, Massachsets: Harvard University Press.

SIMON, Herbert (1955), "A Behavioral Model of Rational Choice". *Quartely Journal of Economics*, n. 69, pp. 99-118.

SIMON, Herbert (1956), "Rational Choice and the Structure of the Environment". *Psychological Review*, n. 63, pp. 129-138.

SIMON, Herbert (1987), "Rationality in Psychology and Economics". In Hogarth, R.M. and Reder, M.W. (Eds.), Rational Choice: The Contrast Between Economics and Psycology, Chicago: University of Chicago Press.

VANBERG, Viktor J.(1994), *Rules and Choice in Economics*. London: Routledge.

WILLIAMSON, Oliver E.(1985), *The Economic Institutions of Capitalism*. First Paperback Edition, New York: The Free Press, 1987.

Appendix
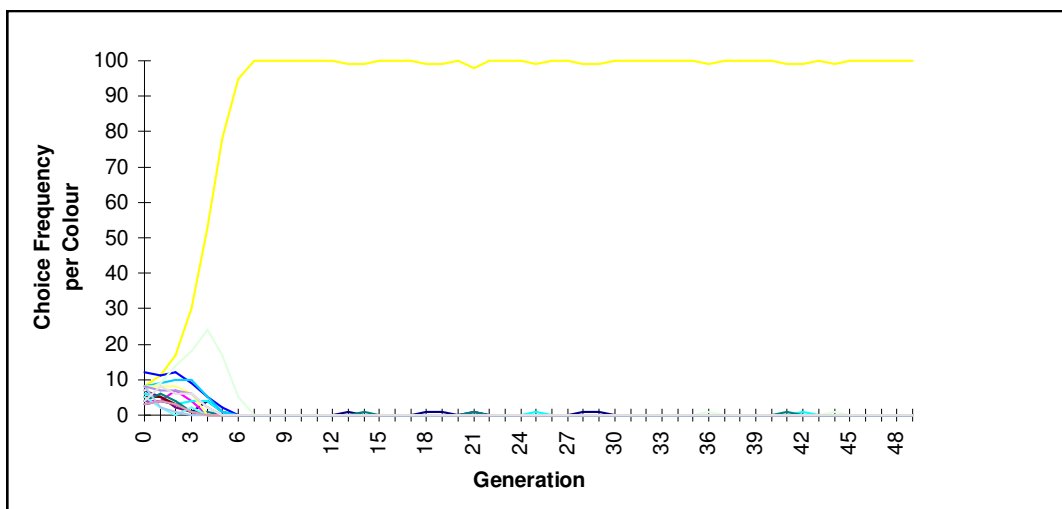
Fig 1: Co-ordination (problem 1): Choice frequency per colour through the simulation
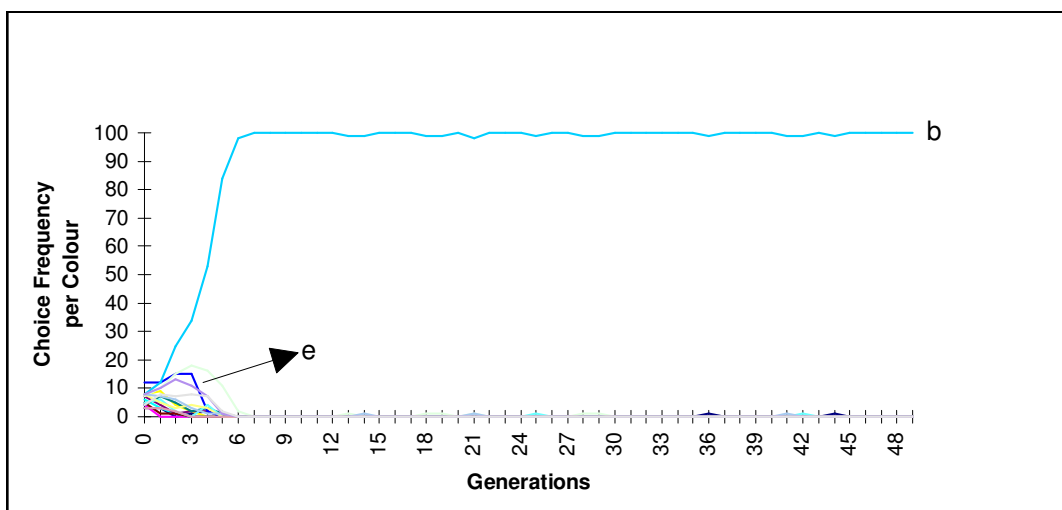
Fig 2: Co-ordination (problem 2): Choice frequency per colour through the simulation
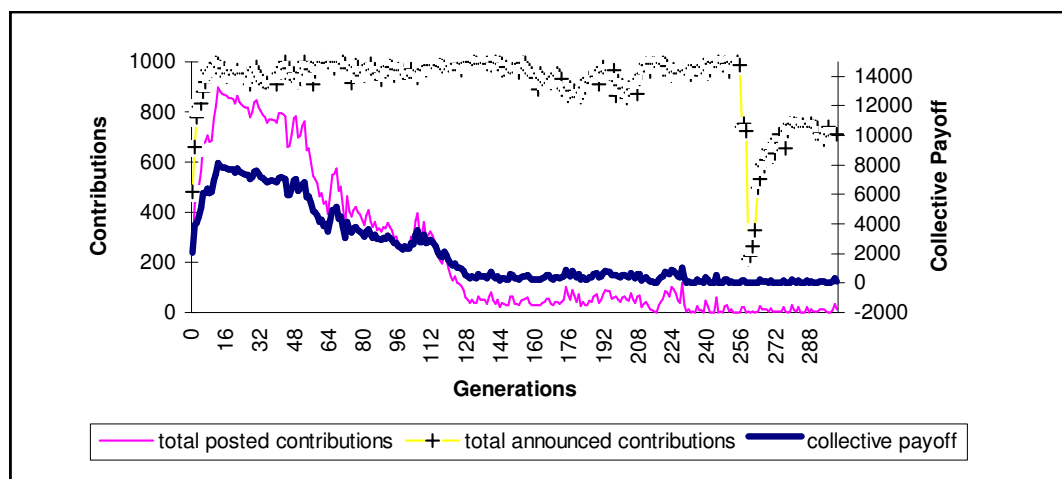
Fig 3: Co-operation: Posted and announced contributions, and collective payoffs through the simulation
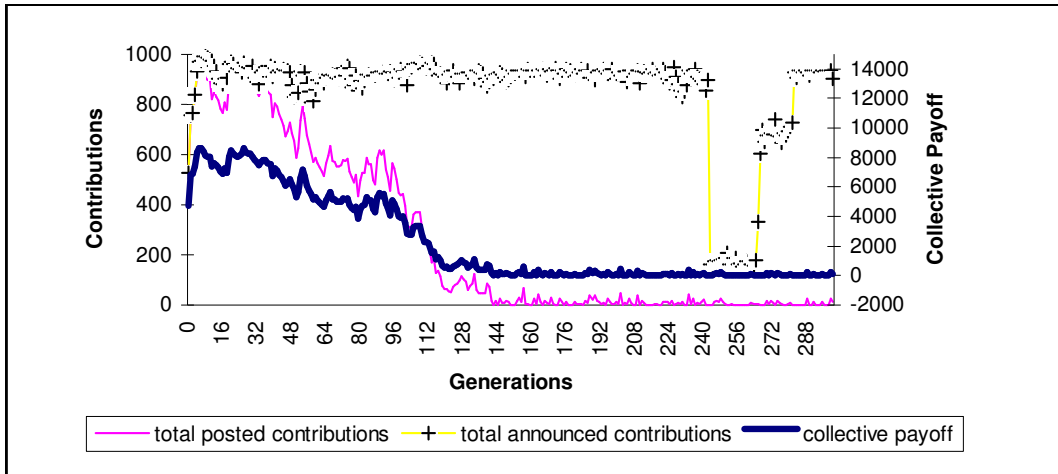
Fig 4: Co-operation - agents created with maximum degree of truth: Posted and announced contributions, and collective payoffs through the simulation
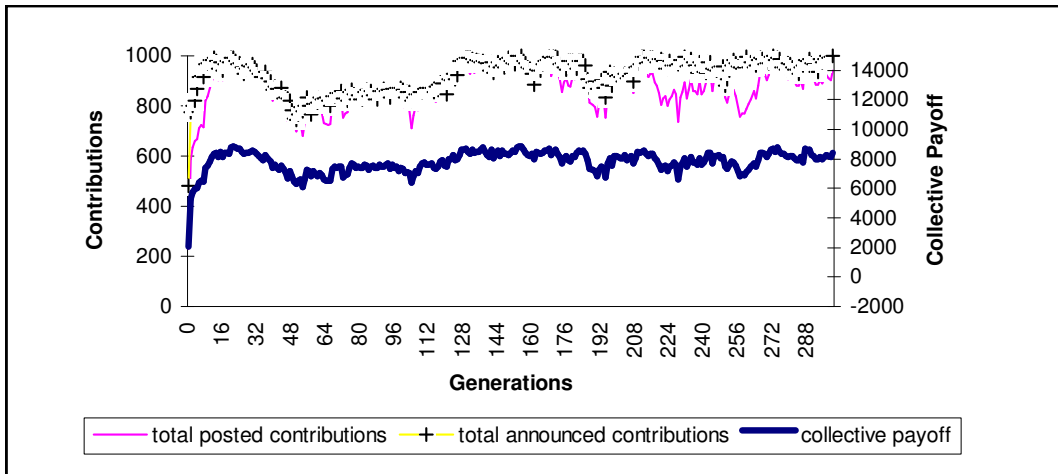
Fig. 5: Co-operation with monitoring meta-agent; *probability of monitoring* = 1: Posted and announced contributions, and collective payoffs through the simulation
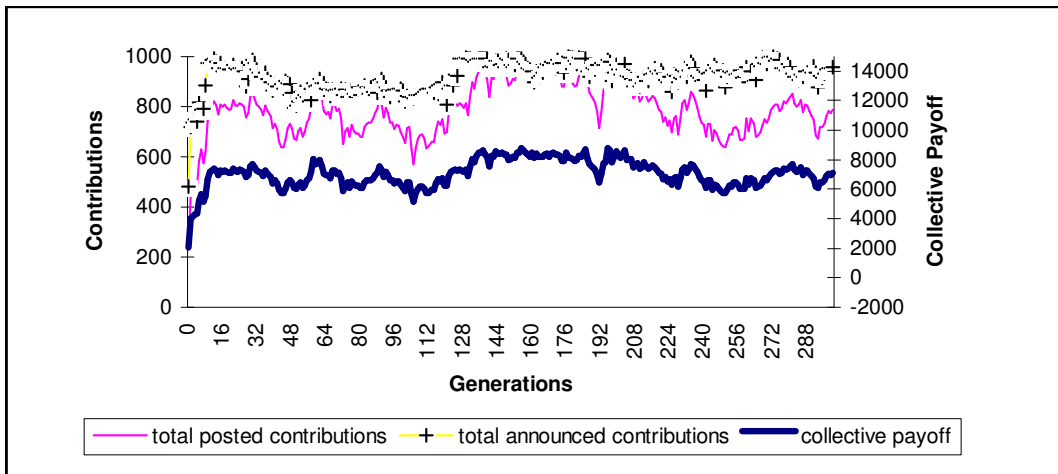
Fig. 6: Co-operation with monitoring meta-agent; *probability of monitoring* = 0.5: Posted and announced contributions, and collective payoffs through the simulation
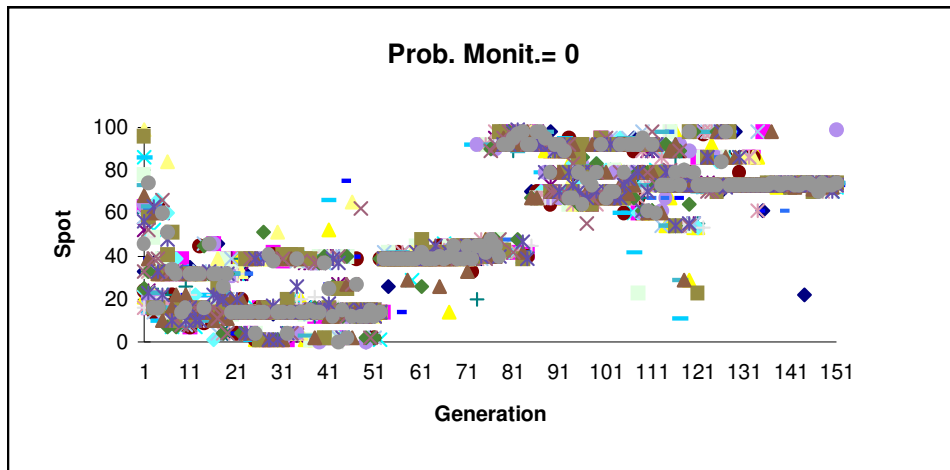
Fig. 7: Exploring and exploiting a landscape (no monitoring)
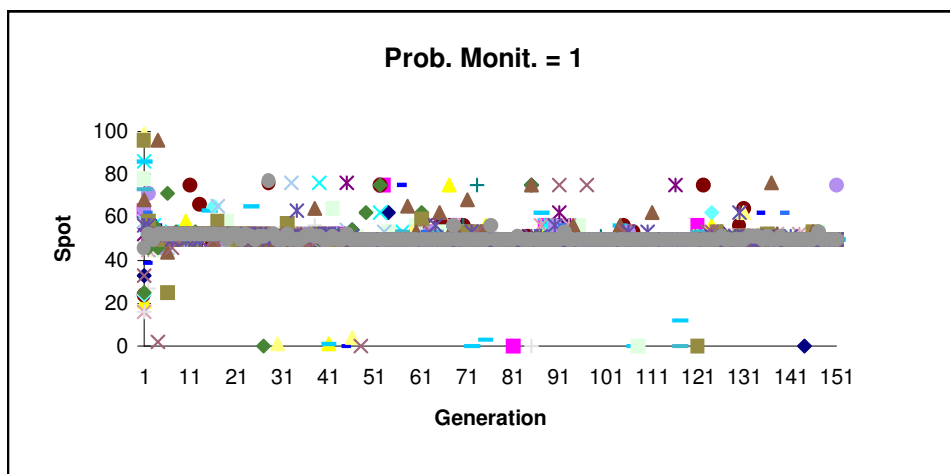
Fig. 8: Exploring and exploiting a landscape (maximum monitoring)
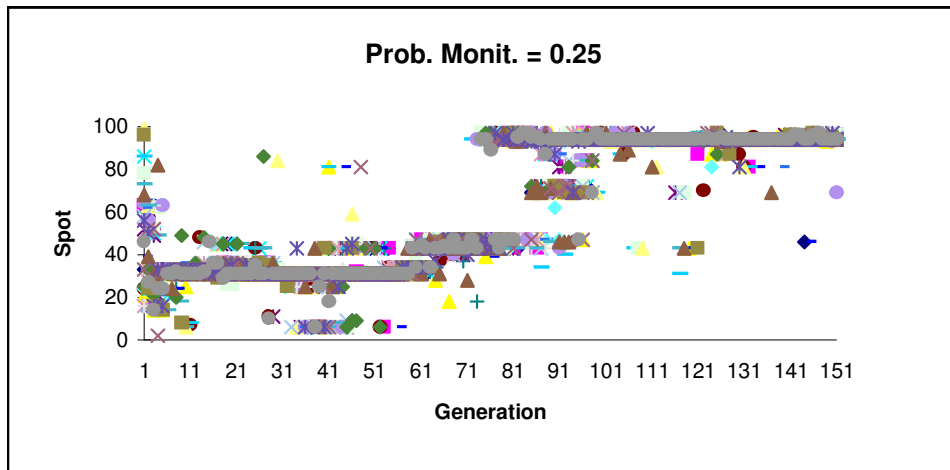
Fig. 9: Exploring and exploiting a landscape (intermediate monitoring)
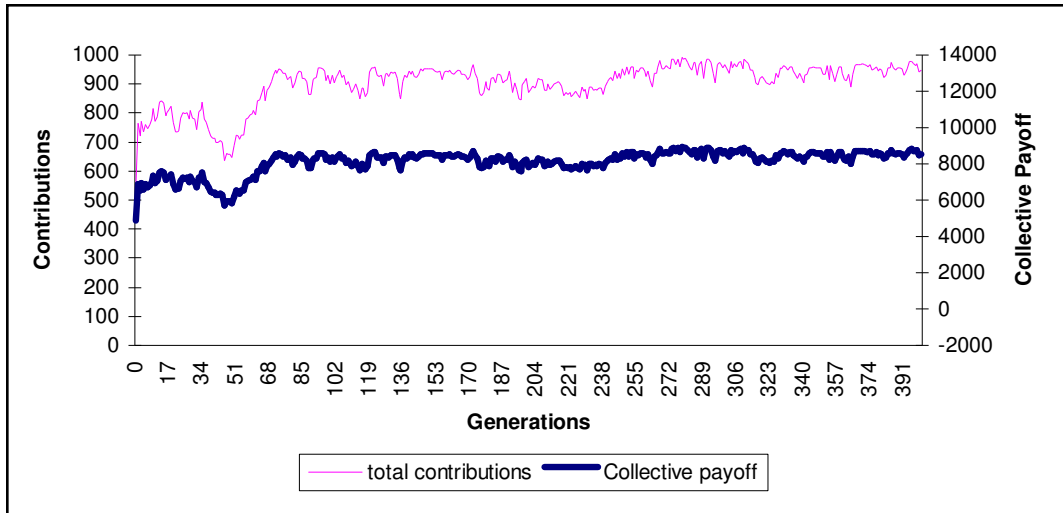
Fig. 10: Co-operation revisited (all agents created equal): contributions and collective payoffs through the simulation
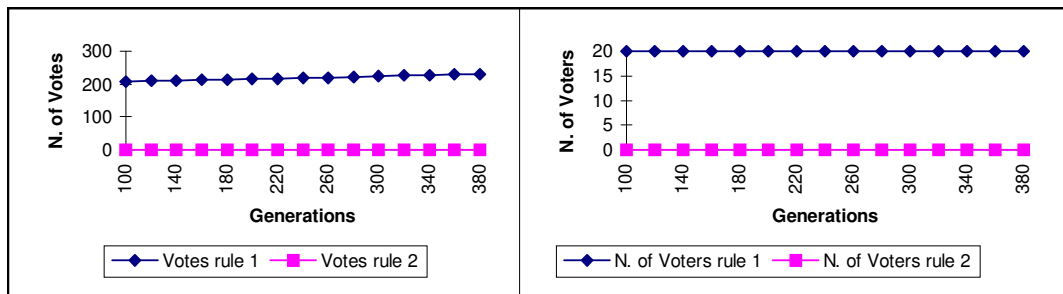
Fig. 11: Co-operation revisited (all agents created equal) - n. of votes and voters
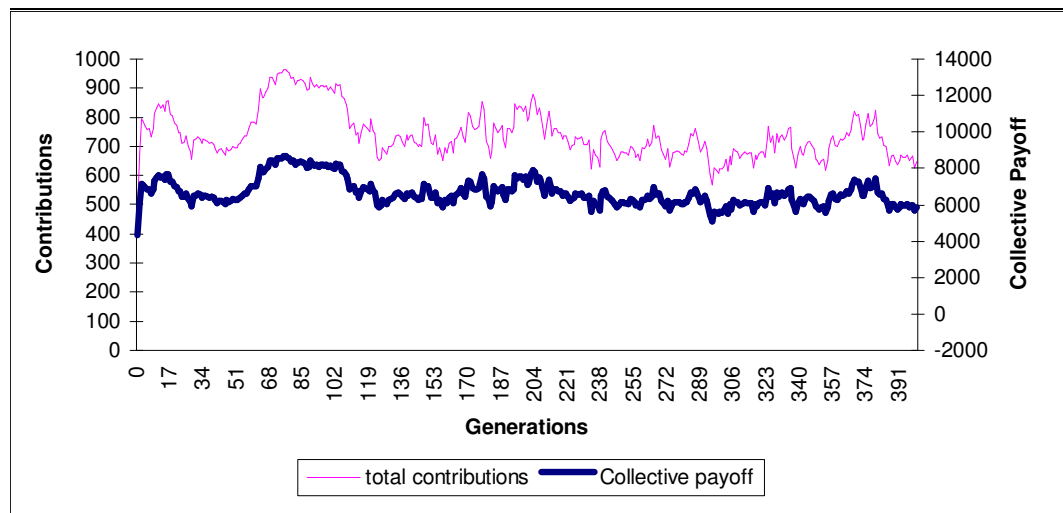
Fig. 12: Co-operation revisited (distributed power): contributions and collective payoffs through the simulation
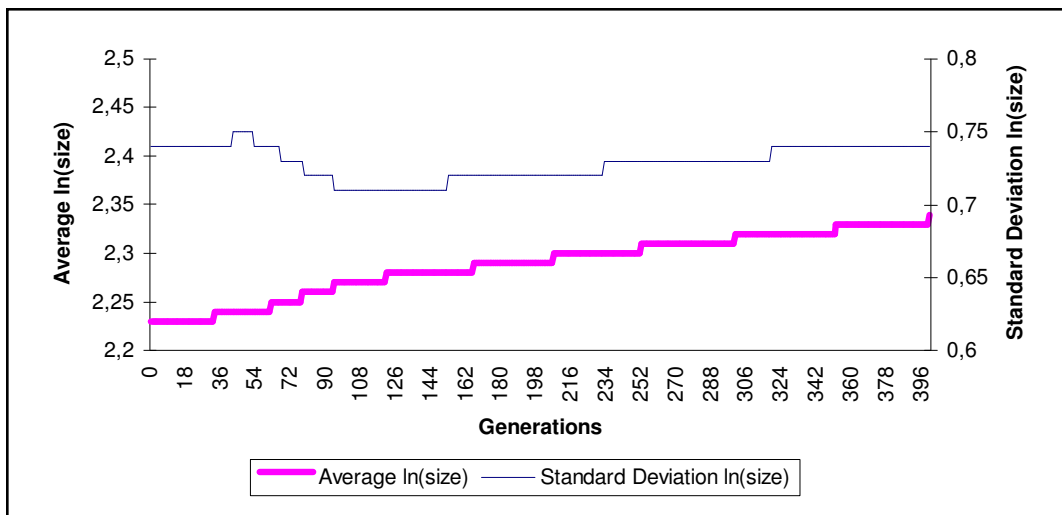
Fig. 13: Co-operation revisited (distributed power) - size distribution
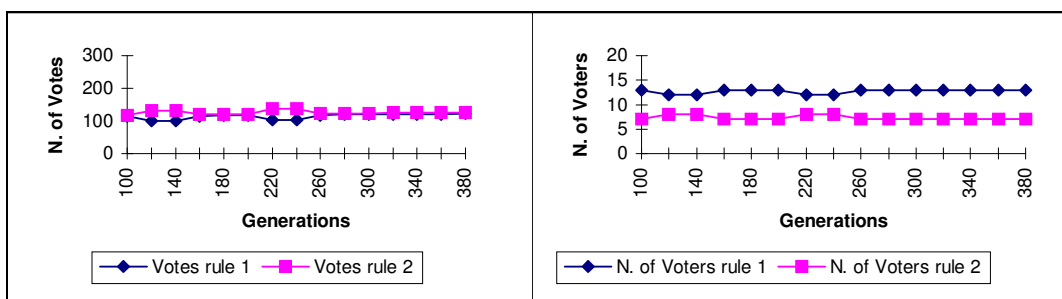
Fig. 14: Co-operation revisited (distributed power) - n. of votes and voters
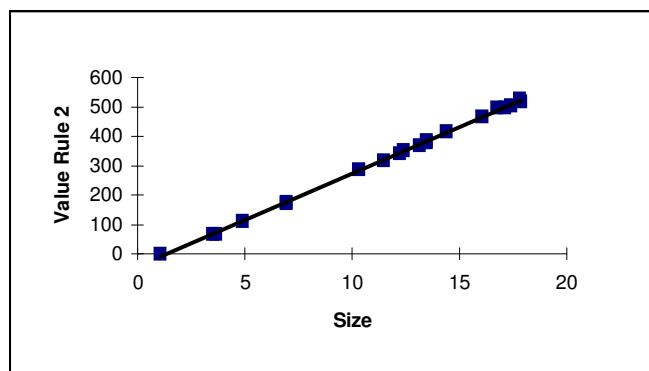
Fig. 15: Co-operation revisited (distributed power),
correlation size - value of rule 2, in generation 101
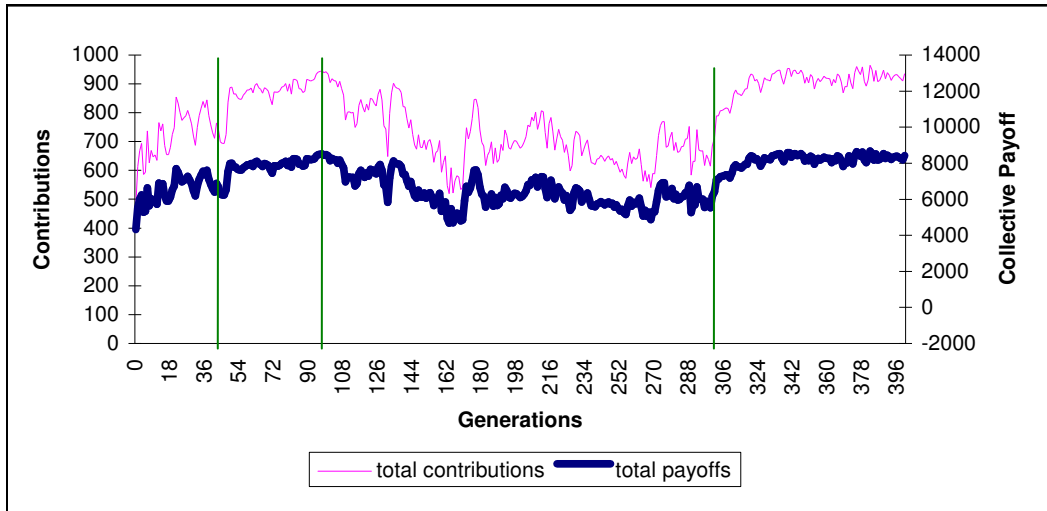
Fig. 16: Co-operation revisited (distributed power), probability of monitoring=0.59): contributions and collective payoffs through the simulation
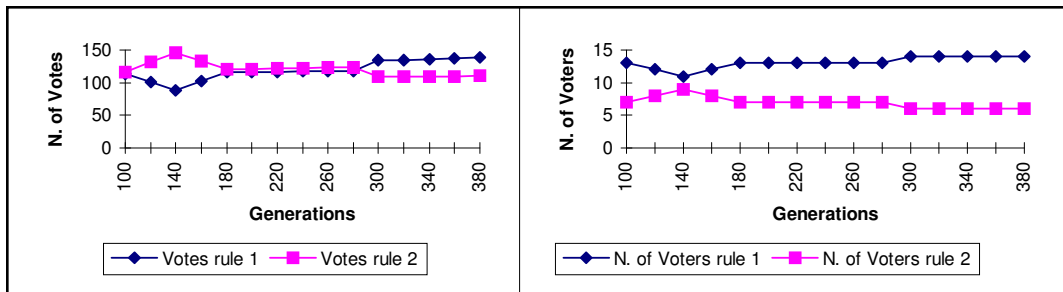
Fig. 17: Co-operation revisited (distributed power), probability of monitoring=0.59, n. of votes and voters

Contents 0