Deborah Duong's 1991 Master's thesis, published "as is" in
Behavioral Science, Volume 40, 1995, pp. 275 - 303.

A SYSTEM OF IAC NEURAL NETWORKS AS THE BASIS FOR SELF-ORGANIZATION IN A SOCIOLOGICAL DYNAMICAL SYSTEM SIMULATION

by Deborah Vakas Duong and Kevin D. Reilly
University of Alabama, Birmingham

This sociological simulation uses the ideas of semiotics and symbolic
interactionism to demonstrate how a simple associative memory in the
minds of individuals on the microlevel can self organize into macrolevel
dissipative structures of societies such as racial cultural/economic
classes, status symbols and fads. The associative memory used is
the IAC neural network. Several IAC networks act together to form
a society by virtue of their human-like properties of intuition and
creativity. These properties give them the ability to create and
understand signs, which lead to the macrolevel structures of society.
This system is implemented in hierarchical object oriented container
classes which facilitate change in deep structure. Graphs of general
trends and a historical account of a simulation run of this dynamical
system are presented.

KEYWORDS: Neural Network, Self-Organization, Object-Oriented System, Container Class, Sociological Model, Dynamical System, Semiotics,
Symbolic Interactionism

INTRODUCTION

Neural Networks, Object Oriented Programming, and Dynamical System Simulation

Neural networks are self organizing dynamical systems reputed to represent meaning better than other Artificial Intelligence methods. An Interactive Activation and Competition (IAC) neural network is an associative memory which has the human capacity to understand new situations and make generalizations. It is by virtue of these traits that IAC networks have another human property when many of them interact together: they create symbols which have shared meaning. These symbols are a source of order for the system as a whole. This simulation is a larger self organizing dynamical system made up of smaller IAC networks interacting. It represents a society, and the orders that the IAC networks develop among themselves represent structures that exist in societies.

To achieve such order, the architecture of individual IAC networks has to be changeable. Traditional IAC models do not have this capability. Object oriented programming with container classes facilitates dynamic network architecture in this system. Thus, this program not only demonstrates how the individual IAC networks can develop order among themselves, but also serves as an example of the value of object oriented programming with container classes in changing deep structure. It demonstrates the value of this implementation method in the modeling of dynamical and self organizing systems in general.

Semiotics

The science of semiotics tells us that human beings continuously search for meaning in signs, and that social life is a system of those signs. Perhaps this is so because of the limitations of biology: we can not see truth directly, but guessing at it through our senses is a matter of survival. We associate people's observable physical features with meaning instinctively, just as we associate words with meaning. Yet, just because we perceive a sign as having a particular meaning and have learned this meaning on our own doesn't mean we are correct, even if we are in agreement with others. Racial prejudice is one such mistake. While race is not a true indicator of ability, we subconsciously classify others on its basis.

This simulation is about the creation of signs, about how symbols come to have meaning to people. It is about people classifying each other and themselves based on outer characteristics. It is a self organizing system: with only a simple neural associative memory in each person and goals "programmed in", properties of society such as racial class and status symbols emerge. The people of the society do not communicate to each other with words, but develop opinions in synchrony from their similar experiences. As a dynamical system, their opinions often develop from "accidents of history": accidents which determine the future of the society. The people in the simulation always look for signs, often choose false signs, and attempt to trick each other by putting false signs upon themselves.

Symbolic Interactionism and Self Organization

As a demonstration of how ideas create societies while societies create ideas, this simulation is an inquiry into the sociology of knowledge. Baert and Schamphliere (1987) pointed out the similarities between the theory of self organization from physics and symbolic interactionism originating in sociology. In symbolic interactionism, man is not only a reactor to society but a creator of it as well. According to G.H. Mead, because man has a self, he is able to put himself into the perspective of other participants in an interaction

and act only after he has predicted the outcome based on past experiences.
People "reconstruct" society by virtue of their capacity to observe and change themselves. That is exactly what happens in this simulation.
Employers try to hire talented workers based on their observable features
while workers modify these features, if they can, based on their perceptions
of whether employers will hire them or not. As in Mead's ideas, this self organizing system goes beyond external/algorithmic determinism and environmentalism: change is not predicted but comes from the individual actions of persons. As the workers change themselves and the employers change their opinions of the workers, properties of societies emerge.

Workers put the history of their traits and whether they were employed or not into an IAC associative memory. This self organizing neural network represents Mead's "me" aspect of the self by incorporating
the employers attitudes into the self. There is stochasm in the employers'
choice of employees and in the workers' choice of traits representing the unpredictability of the "I" response to the community. Both the stochasm and the self organization of the IAC allow for novelty in the system and for status symbols and classes to emerge without being directed to. This simulation differs from symbolic interactionism
in that there is an objective truth: an actual amount of talent each worker has. It is just judged by those whose capacity to judge is limited.

Simulation Purpose

The purpose of this simulation is to demonstrate that a simple associative memory is sufficient to explain the development of social signs, status symbols, racial prejudice and racial class. The individual
actions of workers on the microlevel create the structures of society on the macrolevel: the sum of their actions is greater than the parts.

In this simulation, there are two types of macrolevel order: structures of belief and social structures. We have a unique opportunity
to look into quantified belief because we use neural networks. Possible
belief structures of employers are belief in status symbols and prejudice.
An employer believes in status symbols if he has a relatively higher opinion of people having items purchased with social rewards. In this simulation, money is the social reward for talent. There is some truth to the belief that status symbols are a sign of having desired qualities, more than belief in items which can be acquired freely by anyone or belief in traits such as skin color which no one can change. In this simulation, prejudice occurs when all employers hold a lower opinion of one race than another, even though the members of each race have the same degree of talent. Workers with similar traits may come to have similar beliefs because they are treated the same. This formation of cultural classes by self-fulfilling prophecy is another macrolevel belief structure.

Possible social structures are racial economic classes, the purchasing of status symbols, and fads. Racial economic classes form when one race persistently has more money, material possessions, and higher levels of employment than another. We know status symbols are being purchased when expensive items are purchased more than can be expected by chance by those who can afford them. Fads, in this simulation, are the simultaneous acquiring of traits which cost no money. Social structures result from belief structures while belief structures result from social structures: they co-evolve together.

METHODS

In this system, three employers take their turns laying off employees and hiring replacements from a pool of fifty workers. Each employer has ten employees, making the over-all employment rate 60%. They lay off 70% of untalented workers randomly, and 20% of all workers randomly. With these percentages, it is likely that all the employees that an employer has in any particular cycle will be laid off before five cycles of firing/hiring. Employers hire on the basis of an estimation of the workers' talent, which they can not observe until after employment. They make this guess by an association of workers' traits and talents that they have seen in the past. The first employees are hired randomly. After employees are dismissed, they are not recognized by the employer when they re-apply: he still has to guess their talent.

The three observable worker traits are a fad, a suit and a skin color. The only difference between these three traits is the ease with which they may be acquired by the workers. A fad is free and may be changed to any one of three values. Although one of the suits costs nothing, the other two must be bought with money earned from employment. Their prices are 0 dollars for suit 0, 6 dollars for suit 1 and 9 dollars for suit 2. All workers start with suit 0. One dollar is earned each round that a worker is employed. Skin color is a trait which the worker can not change. It has two possible values: black and white. The proportion of talented persons of each skin color are equal (less than two percentage points difference). The workers' unobservable trait, talent, comes in two values: talented and untalented. Skin color, fad and talent are all assigned randomly at the beginning of the simulation. Once a worker is hired, an employer knows his talent for the purposes of laying off and to update his associative memory.

Workers who are unemployed re-decide what they should be and change their traits, if necessary, in an attempt to impress the employers. Even if they decide to keep the suit they have, they must pay for it again. They change their traits to (or keep them at) the combination which they can afford that they associate most highly with employment. Every time they are judged by an employer, they add knowledge of the result to their associative memories.

The IAC

The IAC neural associative memory holds and processes the knowledge in the simulation. Each person in the simulation has a separate IAC network. Neural networks are good models of human decision making because they process knowledge in the same way that human beings do: they capture meaning with models of human neurons. The IAC network in this model has the same behavior as that described by McClelland and Rummelhart (1989) with the additional property that nodes may be added and deleted as necessary during the simulation, so that memories may be added as new knowledge is learned and deleted as old knowledge is forgotten.

"Interactive activation and competition" is a descriptive name for the IAC network's architecture. This network is divided into several pools of nodes, each node having inhibitory connections with other nodes in its pool. These inhibitions put the nodes in competition with each other: the stronger one node is activated, the more it is able to turn other nodes in the same pool off. Every node in each pool is also connected to a node in a special pool with no intrapool inhibitions, the instance pool. The activation is interactive because the bidirectionally excitatory connections provide the route by which what is going on in one pool influences (while it is influenced by) what is going on in the other pools.

This architecture serves as an inductive associative memory: it can arrive at a general conclusion based on specific examples. The pools represent "traits" : in this simulation they represent the traits of fad, suit, skin color and talent in the employer's memory.
In the worker's memory, they represent fad, suit, and whether the worker was employed or not. The nodes in the pools represent values of traits. For example, in this simulation there are three nodes in the suit pool representing suits 0, 1, and 2. Each instance pool node connects a single set of associations in the pools. In this simulation,
instance nodes represent people, and they have excitatory connections to the values of that person's traits, at most one in each pool (see figure 1). In the employer's memory, each instance node represents a different individual while in the worker's memory, each instance

node represents himself at a different point in time.

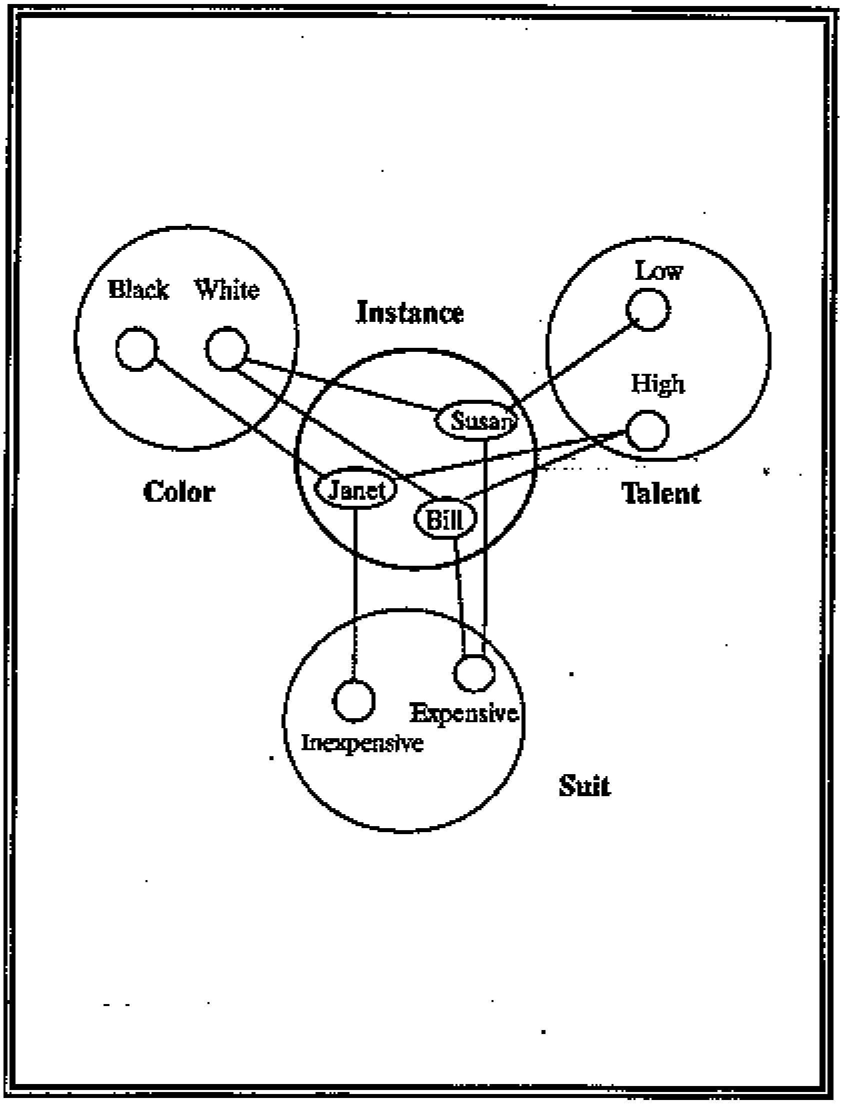

FIGURE 1
An IAC Network. Individuals with a variety of personal attributes.

Intuition and Creativity

To estimate how talented an applicant is, the employer activates the nodes in the pools that represent the applicant's traits and sees how strongly the node representing talent is activated compared to the node representing lack of talent. The answer that he will get will not be a precisely correct logical one, but a more "intuitive", even "creative" answer. Energy flows back and forth through the network fifty times: it goes from the activated traits to the people (represented by instance nodes) who had them, to their traits and to the people who had them while it is being inhibited in the pools, so that each pool has a node activated more than the others. That node, in the talent pool, is the answer.

In this simulation, the amount of talent a person is judged to have is the activation of the node representing talent less the node representing lack of talent in the talent pool. It is a real number. It is an answer which is holistic,

using all of the knowledge. It is intuitive because even if the black people that the employer knows are as talented as the white people he knows, it won't give an answer based on that, but on the other traits that those black people had, and the talent of the people who had those traits, and the talent of the people who were like those people, etc.
Every cycle adds another level of recursion to the answer.
It is intuitive because you can not quite pin down how the answer is arrived at, but it isn't by reasoning. It is almost like a "feeling" about a person. Why don't people use logic instead of this intuition? Perhaps it is just another biological limitation. Our neurons (without the assistance of a paper and pencil and a statistics text book) have difficulty in holding and processing all of the information needed to make a precise answer. If they did, like expert systems, they would not be able to deal with new information because the data to base it on would not be there. If a few examples were there, then the answer would be based on those few examples instead of the whole body of data, leading to an insignificant answer.

In contrast, the IAC network can easily deal with new combinations of trait values it has never seen before and even trait values it has never seen at all using the fact that a new trait value is something other than the trait values now present in the trait pool. If all the trait values in a pool are associated with lack of talent, meaning an employer is dissatisfied with them, then not activating them may lead to a higher estimation of talent in some cases, so that the employer will be willing to try something new.
This creative ability to understand new situations is essential to human beings both in this simulation and in real life. Even though the worker's IAC represents Mead's "me", the values of the culture, it is still creative in imaging what other people (or employers) might think.

A Choice Function

Mead's "I", the decisions made based on information from Mead's "Me", does not follow strictly the output from the association, but has an element of stochasm as well. When an employer chooses an applicant to fill a position, he uses his IAC to measure the talent of each unemployed worker , and then puts all of these values into a choice function which chooses one of them based on a decision parameter.

Possible values of the parameter are from -1.0 to 1.0. If 1.0 is given, the choice function chooses an applicant from the list randomly.
If 0 is given, the values of talent (after normalization) are interpreted as a distribution from which a random variate is picked. The applicants with higher talent have more chance of getting picked. If -1.0 is given, the person with the maximum talent value is chosen.

Other values between -1.0 and 1.0 result in a decision that varies smoothly between these extremes. A new distribution is mapped from the given set of values. The decision parameter actually represents a percentage used in this mapping. For example, decision parameter

-0.5 will let the maximum talent value take up its given percentage of the distribution, as when zero is the decision parameter, plus 50% of the remaining space, the unassigned portion of the new distribution.
The other values are mapped on in descending order: the next value takes on its proportion plus 50% of the remaining space, and so on. If there is not enough space for the given percentage of an applicant, it is then cut off. Of course, 1.0 gives the maximum value 100% of the space.

If the decision parameter is 0.5, the minimum amount of space that any value in the new distribution may take is 50% of 1/n of the total space, where n is the number of applicants. The smallest values are apportioned first, so that values larger than 50% of 1/n are calculated
based on the remaining space. Of course, if the decision parameter is 1.0, each of the n applicants gets 100% of 1/n of the space, making a uniform distribution (see figure 2).

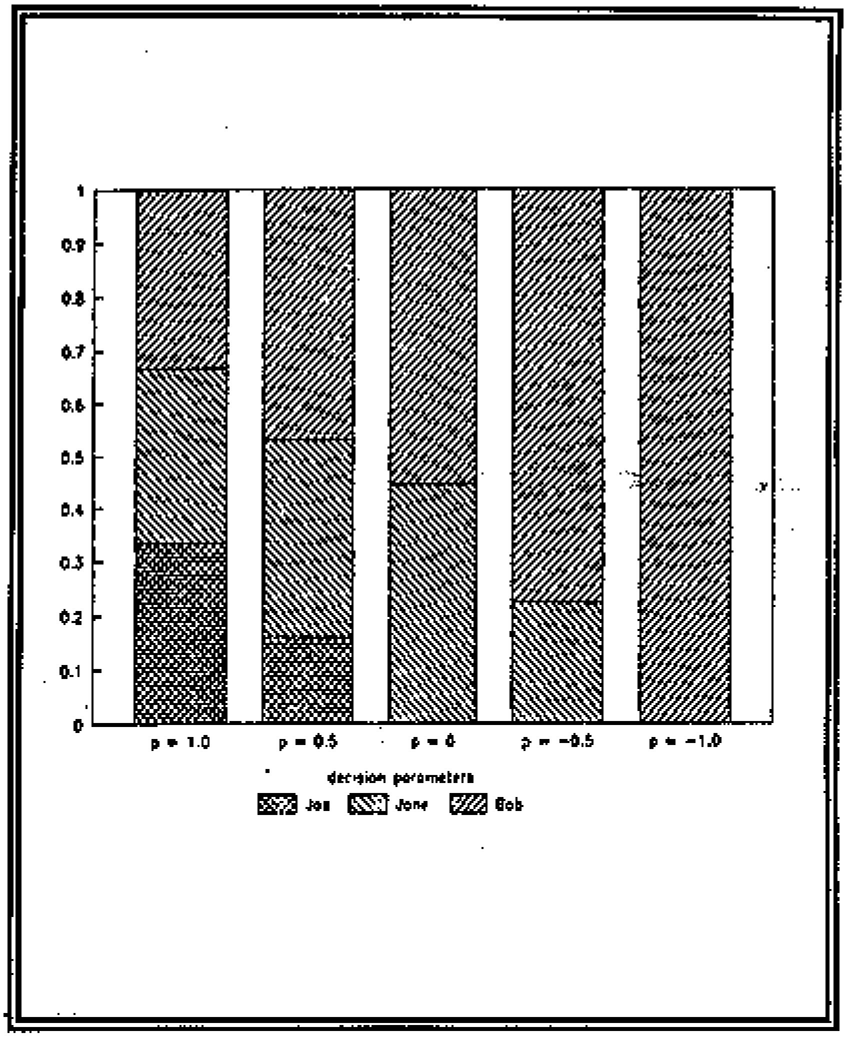

FIGURE 2
Cumulative Distribution of the Choice Function.
Distributions have various decision parameters. The function given Joe =3, Jane = 7, and Bob = 8.

The decision parameter may be thought of as representing self-doubt.
If you are absolutely sure of your intuition, then you pick the one you feel the best about every time: this would be represented by a decision parameter of -1.0. If you want to give close seconds a chance sometimes, your decision parameter is a negative number greater than -1.0. If you have little confidence in your ability to decide, you may chose more randomly with a positive decision parameter less than 1.0.

We used a single decision parameter per simulation, usually negative for the employers and slightly negative for the workers, but it is also possible to vary the parameter during simulation, producing a "simulated annealing" effect. This could be useful, for example, to simulate decision making in people through their development. Teenagers are known to try different things before they have much experience, but this gives them experience with which to stabilize their ideas about the world and about themselves. We could represent their decisions with a positive decision parameter to the choice function.
As people gain experience, they gradually become more confident in their decisions and may even become so set in their ways that they no longer give new things a chance. We could represent this by making their decision parameters more negative. According to physics' annealing theory, this is an optimal way to make use of all the facts. Individual points of view are not settled upon before sufficient exposure, and those within a generation come to have similar views in synchrony.

The decision function may be thought of as analogous to the "activation rule" of the nodes of a neural network, persons and employers as analogous to neurons, and the society as analogous to a neural network. As nodes are connected to each other through synapses, so are employers connected to workers by employment. Just as the network dissipates to a solution after many cycles, so does the society as a whole after its cycles. The solutions at the society level are the dissipative structures of racial class, fads, and status symbols.

Container Class Organization

We are able to ask questions about micro-macro interrelations - the relation between the node's state and the network's state, or the relation between the network's state and the society's state - by virtue of the object oriented container class organization of the system. The program is implemented in object-oriented C++. Each entity of this and any other object oriented system is an object of a class. The data and functions which act upon the data are encapsulated together into a class, representing the attributes of an entity and its behavior. Objects are individuals of a class: for example, "dog" would be a class while "Rover" would be an object. A container class, in object oriented programming, is a class which has other classes in its composition (see figure 3).

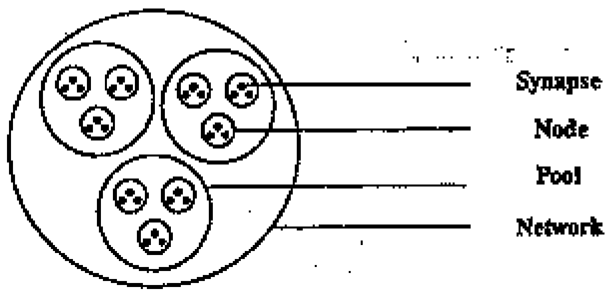

FIGURE 3
Container Relationships. An IAC Network is made of pools. These pools contain nodes and the nodes contain synapses.

Container classes are convenient because they are modular and you know just where to make a modification to ask a question: for example, to ask how a worker's behavior affects the network we know we must change one of the functions in his class. They are important tools in demonstrating how a change in the number, the interrelations among, or even the type of micro-entities a macro-entity is composed of affects the macro-entity. Container classes can do this because the list of objects an object contains may vary in length and in terms of which particular objects they are composed of on-line. Also, inheritance classes may be used to represent variations in the type of object contained object on-line. Possible questions to ask are "How do the number of workers affect the opinions of the workforce as a whole?" or "How do different compositions of the workforce in terms of variation in worker behavior affect the system as a whole?"

In this simulation, container class organization enabled the arrangement of memories in the networks to vary and enabled employers to delete and add employees. Because of this we could ask how changes in the micro-entity brains affect the macro-entity society and how the brain is affected in turn.

Model runs, however, use lots of time and space. By simulating the actual entities themselves instead of using differential equations to represent them in aggregate as in most dynamical system simulations, we have gained the ability to change deep structures on-line, but we have lost the convenience and efficiency of the differential equations. The program has over 2000 lines of code and takes 50 hours of CPU time on a Sun 3/50 workstation. Over 30,000 inquiries to IAC networks are made by the workers and employers. Employers pay, lay off, and hire employees 200 times.

Hierarchy

This simulation is hierarchical, but not in the sense we usually think of when we think of hierarchy in neural networks or in inheritance classes. Its container classes are hierarchical: each class represents a different level in the hierarchy, and may be modified to see the effect on other levels (classes) in the hierarchy.

There are 13 classes in this simulation. From the lowest level to the highest level they are synapse, axon, node, pool, IAC network, trait, memory (characteristics), workforce, employee, employer, companies and society. As depicted in figure 4, the higher levels contain the lower levels.

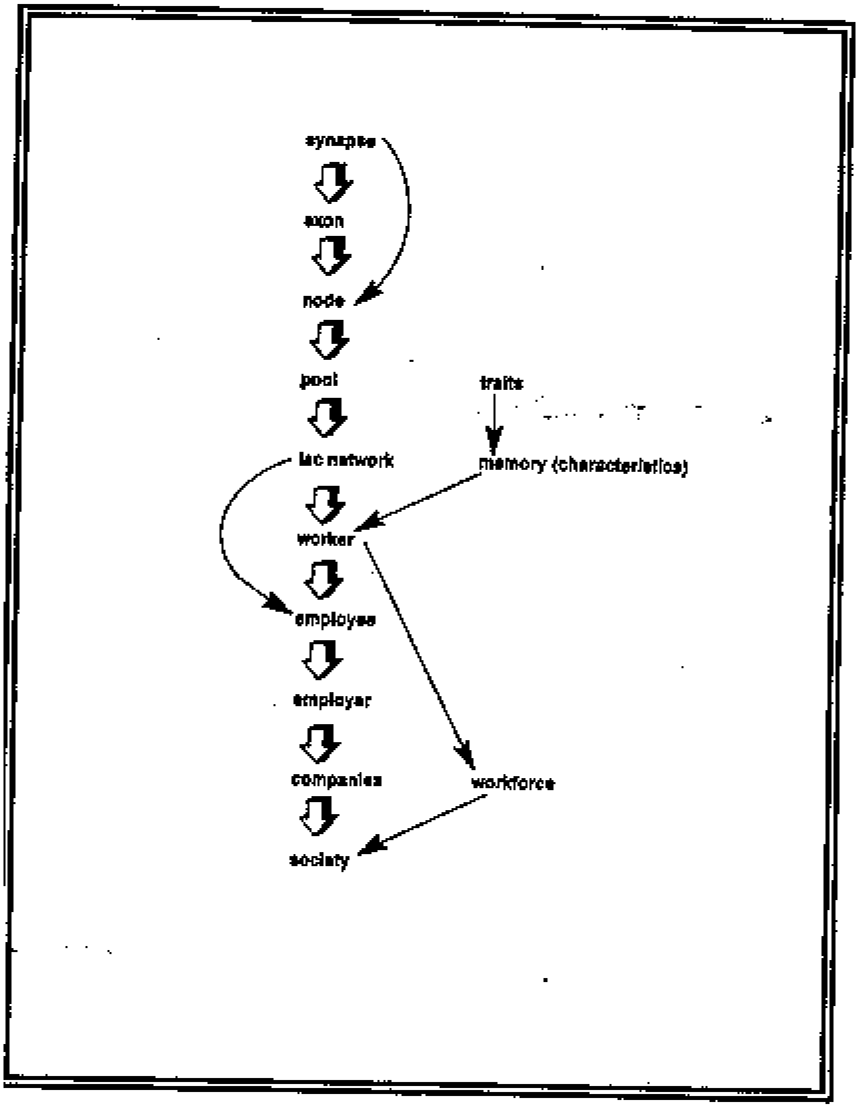

FIGURE 4
Heierarchies of Container Classes. Classes point at the classes which contain them.

A society contains a workforce and companies. A workforce contains a linked list of worker-objects, and the companies-object contains a list of employer-objects. Both the worker and the employer contain an IAC network. Additionally, the employer contains a list of employee-objects and the worker contains a characteristics-object which contains a list of traits. Each employee-object points to a worker in the workforce, so that the companies contain a subset of the workers that the workforce contains. This is one way interrelations may form: when two objects contain the same object they may influence each other through it. The IAC network, which is the brain in both the workers and the employers, contains a list of pools. The pools contain a list of nodes, and the nodes contain a list of synapses and axons. The axons contain synapses of other nodes, creating a connection. From synapse to society there are ten levels of hierarchy.

Reverberation

A change in one of the lowest levels, perhaps in the activation function of the nodes, could reverberate up to the highest hierarchical levels and cause change on the societal level. That is why hierarchical classes are particularly good for asking questions about micro-macro interrelations. The central question of this simulation, "How do people's associations of each other affect society as a whole and how are their associations affected by society?" is really quite complex. To answer it, we might start with changes on the axonal level, the connections in the IAC network, which change when employers and workers learn new sets of associations. We observe those changes reverberate up ten levels of hierarchy to the societal level and back down ten levels to change the associations represented in the axons once again. Or, we might start with changes on the employer level in the employer relationship and watch them reverberate down the hierarchy to the axonal level and back up again. It doesn't matter where we start because perceptual structure co-evolves with social structure: they cause each other.

Data Encapsulation

Change occurs between each level and the next: a change in one level can not change levels further up or down the hierarchy without changing the level next to it (containing or contained by it). This is a trait of natural hierarchies: its simulation is facilitated by the data encapsulation of the container classes. In object oriented programming, data encapsulation means that the data a class is composed of can only be manipulated by the procedures of its class. In our model, each class contains procedures which manipulate only its data, and since the classes are container classes, the data is the next lower level of hierarchy. For example, procedures of class workforce call procedures of class worker to manipulate the worker-objects in its list. Procedures of class worker call procedures of class IAC network to manipulate the IAC network it contains. Every level manipulates the next by calling its procedures, down to the nodes which call procedures of the synapses and axons. Each procedure call is a reverberation to the next level of hierarchy. A class such as synapse or node will contain a pointer to the object which contains it, facilitating change in structure so that reverberations may go back up the hierarchical levels.

Dynamics

We now have enough information to talk about how change is implemented in the simulation. In every cycle, after paying his employees, one of the employers dismisses a random 70% of his untalented employees and 20% of his remaining employees. His linked list of employees decreasesas a result of this action.

Then, he hires the same number back again from the pool of fifty workers, not recognizing the ones that he just dismissed.

His IAC judges each worker's talent and his choice function chooses a worker based on these measures, one at a time. As workers are hired, the memory of each one's traits and talent is added to the employer's network. If an individual was hired by the employer in the past, his representation in the employer's mind is updated.

Memories in the IAC network are updated by the deletion of old nodes and the addition of new ones to the list of nodes representing persons in the instance pool and the creation of axonal connections between the instance nodes and the nodes in the trait pools representing the values of personal traits. As each worker is either accepted or rejected by an employer, he puts this new set of associations between his traits and his employment into his network. If there are ten sets of associations or memories in his network already, he deletes the earliest memory by deleting its instance pool node and its axonal connections to the traits in the pools. He then adds a new instance node and connections to the traits in the pools for the new memory. Keeping only ten memories saves computer memory space and processing time, at the same time it keeps the workers' ideas more or less current. In both the employer's and the worker's IAC networks, new nodes are added to trait pools if none of the memories present have a new memory's trait value and old nodes are deleted in the trait pools if the memory being deleted is the last one to have that trait.

As the workers are hired, the employer's linked list of employees increases. The other two employers repeat this process of laying off and hiring. Then, the unemployed workers all update their traits, one at a time. Each combination of traits which the individual worker can afford is judged by his IAC network in terms of its employment value, then one of these combinations is chosen by his choice function and bought.

These are the changes that occur to the societal connections of the simulation, the employment relations, and the perceptual connections of the simulation, the associations in the IAC. Statistics on each employer's perceptions of the 18 different combinations of worker traits are collected every five frames, along with the numbers and percentages of workers who have those traits. Employment rates for the races and talent rates for each suit are also collected This is done for 200 cycles.in the simulations reported in this paper.

RESULTS

Structure vs. History

Because this is a dynamical system, the results are difficult to report. As Peter Allen, one of the originators of self-organizing social system simulation said, "The richness of human systems, produced by layers of mutual adaptation and initiative and framed by historical circumstances makes each situation `unique'. Case studies may accumulate, but on what principles can general conclusions be drawn?" Graphs of averaged trends, although they do give a good idea of what is happening in a system, do not give a good idea of why it is happening.

Social scientists from Durkheim to Marx, to the dismay of historians, have seen general trends as laws to which specific situations must conform. What happens to an individual is not important: it is determined by laws. What they did not know was that if individuals act on their own, they will drive the general trends while they are influenced by them, causing structure in the general trends to appear even without laws. Under different accidental circumstances, the structure will be different: the particular structure itself is not a law.

This is the idea of self-organization. Nobel Prize winning physicist Ilya Prigogine proposed the idea of dissipative structures which self organize. Peter Allen, Prigogine's co-worker for twenty years, left statistical mechanics to use this idea in social systems. Classical social scientists were influenced by the immutable "laws"of physics of their time, and came to be at odds with historians, who had traditionally believed that men and accident could change the course of history. Prigogine is said to have inserted history into physics: in doing so, he inserted it into sociology as well. So, in reporting results, we will give both graphs of general trends and a historical account of a simulation run.

Chaos

The first thing of note in the results is that, true to dynamical systems, the outcome varies widely from simulation run to simulation run. In some simulations, prejudice and racial class occurred but in others with the same initial conditions and parameters, it did not because of different random numbers to a stochastic choice function. This variation demonstrates how accident can create society.

One example of how it is difficult to make general laws is in trying to determine the effect of stochasm in decision making on the society. In the methods section, we discussed how the decision parameter to the choice function represents the amount of self-doubt. It can also represent open-mindedness or tolerance. The more stochastic the decision parameter is, the more an employer's decisions conform to the population that applied for the job. For example, if the decision parameter is -1.0, then the employer employs all the applicants who have the combination of traits he admires most before he picks any applicants having the next best combination. If there are ten people, two having his favorite combination of traits and eight having his second favorite, and the employer must hire three persons, he will take the two best and one of the others. However, if the decision parameter is zero, although a single one of the best people have greater chance of getting picked than a single one of the second best, there are more of the second best, so that as a class they have more chance of getting picked. The enforcement of more correspondence between who is hired and who applies is an affirmative action mechanism, so we would expect to see less racial class with stochastic decision parameters.

Do stochastic decisions result in less prejudice and racial class?

This does happen most of the time, but not in every case. In a simulation run with only one employer and a maximum-choosing decision parameter of -1.0, with 28 blacks and 22 whites having equal proportions of talent, the employer decided at the beginning that he only liked whites, and kept re-hiring the whites for the entire run of 100 cycles, never giving more than three of the blacks a chance. He never learned the values of suits because color was the only thing important to him. However, in another simulation with decision parameter -1.0, blacks and whites were employed in equal percentages.

Figure 5 shows the trends of prejudice against black people in a society with three employers having decision parameters of -0.5 , and with 17 blacks and 33 whites having talent in equal proportions. This society has racial class: a hypothesis test for the difference between the black employment rate and the white employment rate on typical cycle 180 showed a significant difference at the a = 0.01 level .

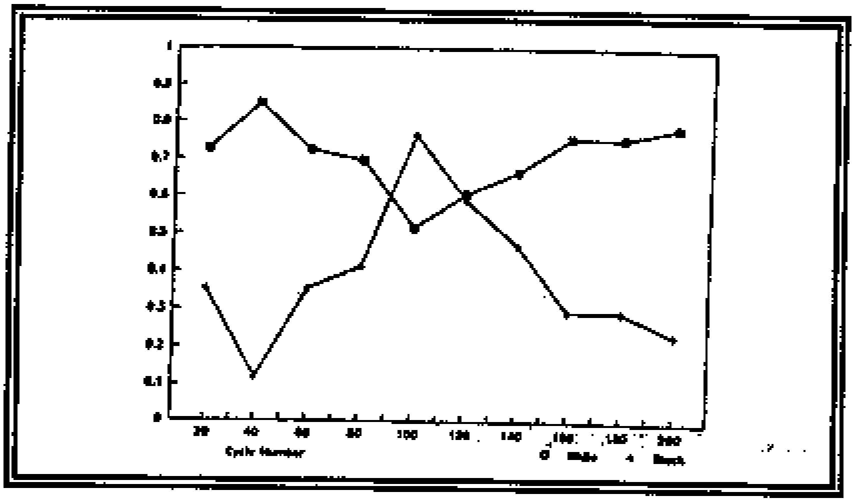

FIGURE 5

Employment Rates for Each Race in Run 1. Decision parameter of the choice function is -0.5.

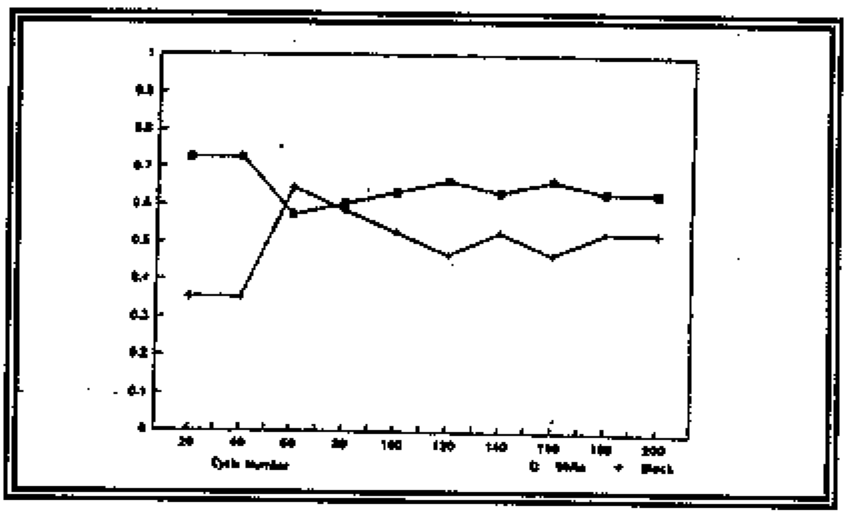

FIGURE 6

Employment Rates for Each Race in Run 2. Decision parameter of the choice function is -0.

Figure 6 shows the same society with a decision parameter of 0 in the employer's choice function. There is no significant difference between the employment rates of white and black (a=0.05) at typical cycle 180: this society has no racial classes. We might be tempted to form a rule, since all else was the same, that the more stochasm there is, the less prejudice and racial class there is. If the world was linear, it would be true. But it is non-linear: When an even more stochastic parameter .15 is used, prejudice and racial class return. It may not even be because of the parameter, but because of the different random accidents that happened in their simulation as a result of a change in the parameter. In another set of simulations using the same parameter but a different set of random numbers and different initial conditions, prejudice occurred with the -0.5 parameter and no prejudice with either the 0 or the .15 (stochastic) parameter.

Thus, we are not talking about what must happen, but what is possible. Even that can be informative, or even surprising. We will concentrate on what happened in the simulation of Figure 5, with a decision parameter of -0.5. We will call this "run 1" and will give both its general trends and its history.

General Trends in Run 1

The first interesting thing about this simulation is that the employers' perceptions of the various traits are synchronous even though they do not tell each other what they believe. They each come to their perceptions through their own experiences alone. The workers, although they do not communicate, also behave in synchrony within their skin-color and even economic classes, because of their similar experiences.

Figures 7 through 13 show the three employers' views of how talented persons having each trait are, comparatively. They all exhibit synchrony: Their opinions go up and down more or less together. The total of the measures of talent for all values of a trait for any employer is 1.0, and the value least associated with talent among the values of one trait is always zero, even if it is highly associated. Values are normalized for comparative purposes.

The synchrony of the perceptions of suits in figures 7 through 9 is not at all surprizing. After all, suits (other than suit 0) are correlated with talent because a worker can only buy them with money obtained from employment, and the untalented are laid off in greater proportions than the talented. Each employer may have observed these facts about suits independently. How might the synchrony of the perceptions of fads in figures 10 through 12 be explained? A fad is free and can be changed as easily by those with talent as by those without it. Perhaps a correlation between the fad and the suit could develop in the case where employers prefer a particular fad in a particular suit and disapprove of the same fad in another suit. Then a correlation between fad and talent could emerge in a self-fulfilling prophecy. Figure 12 is a particularly good example of synchrony. In looking at these, one may be tempted to propose that fads by nature rise and fall quickly, but we have another simulation in which all employers and employees avoided a fad for the entire run.

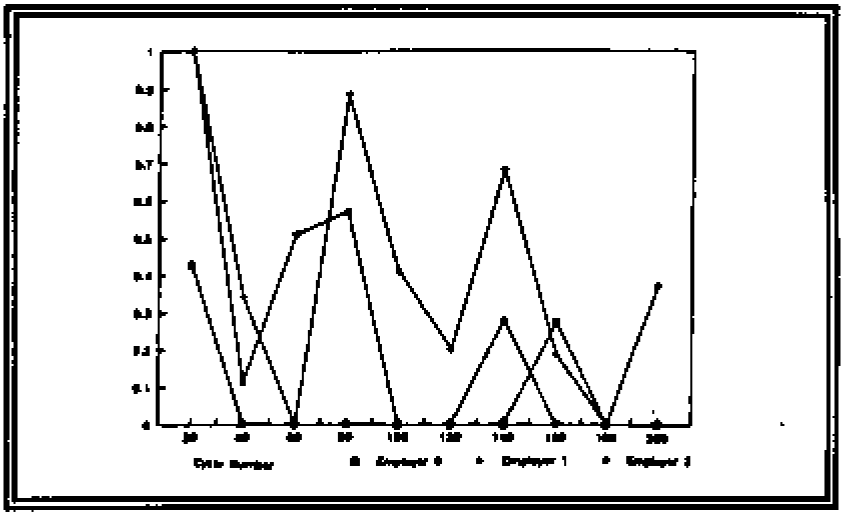

FIGURE 7
Perceptions of Suit 0. The opinions of the three employers are synchronous.

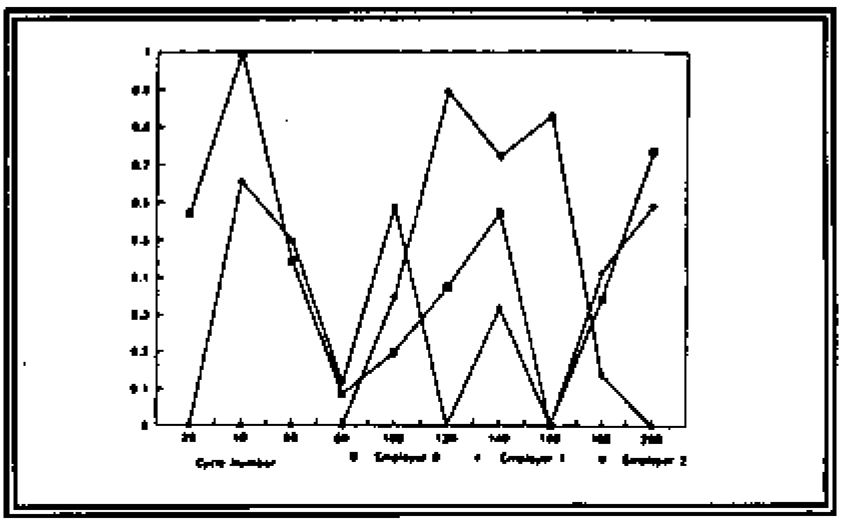

FIGURE 8
Perceptions of Suit 1. The opinions of the three employers are synchronous.

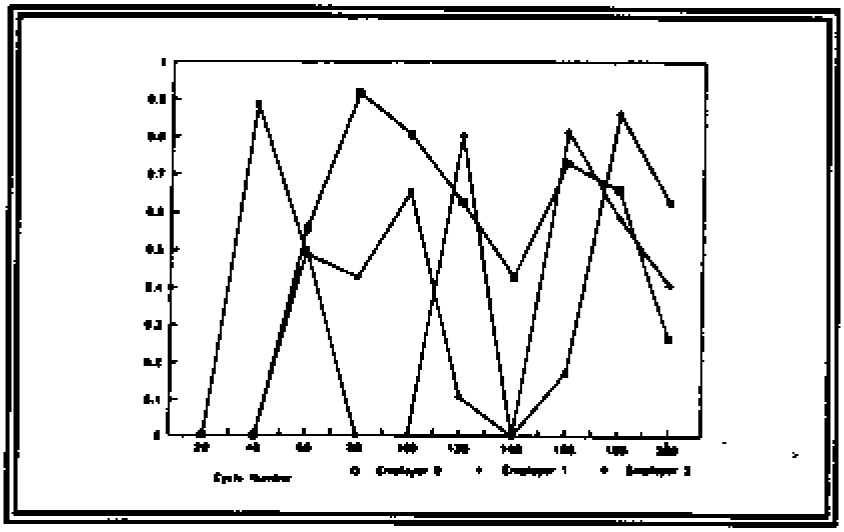

FIGURE 9
Perceptions of Suit 2. The opinions of the three employers are synchronous.

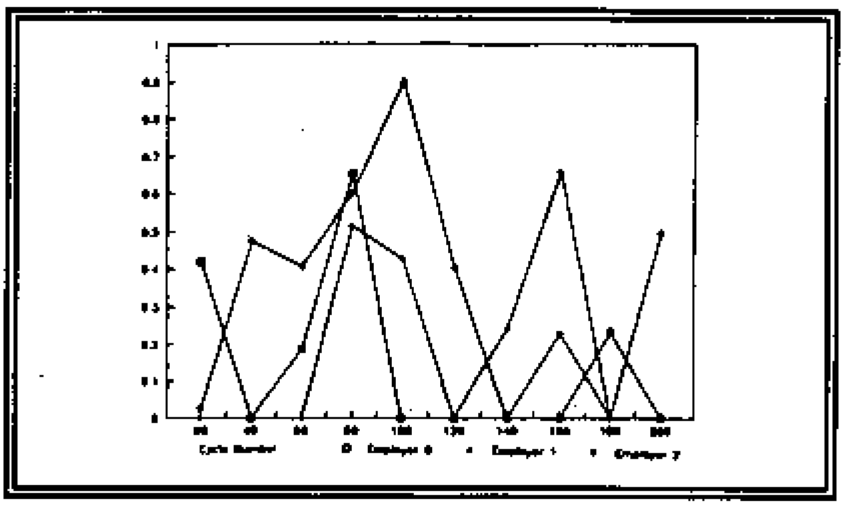

FIGURE 10
Perceptions of Fad 0. The opinions of the three employers are synchronous.

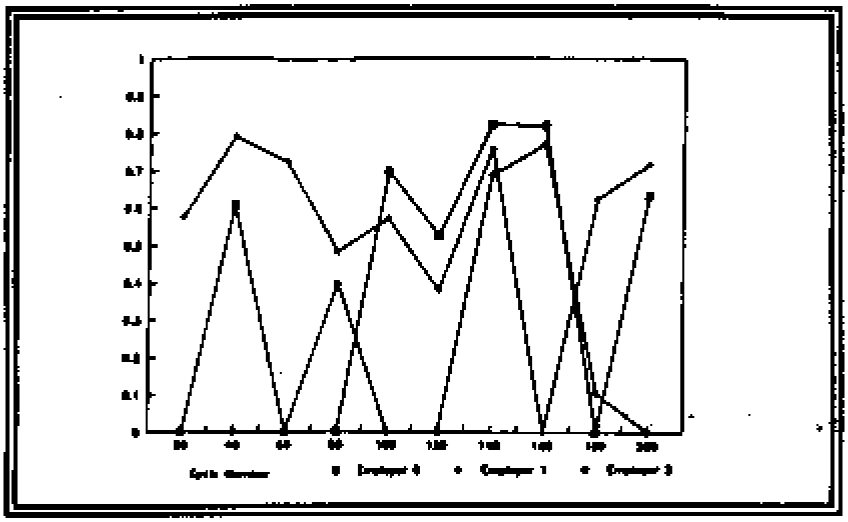

FIGURE 11

Perceptions of Fad 1. The opinions of the three employers are synchronous.

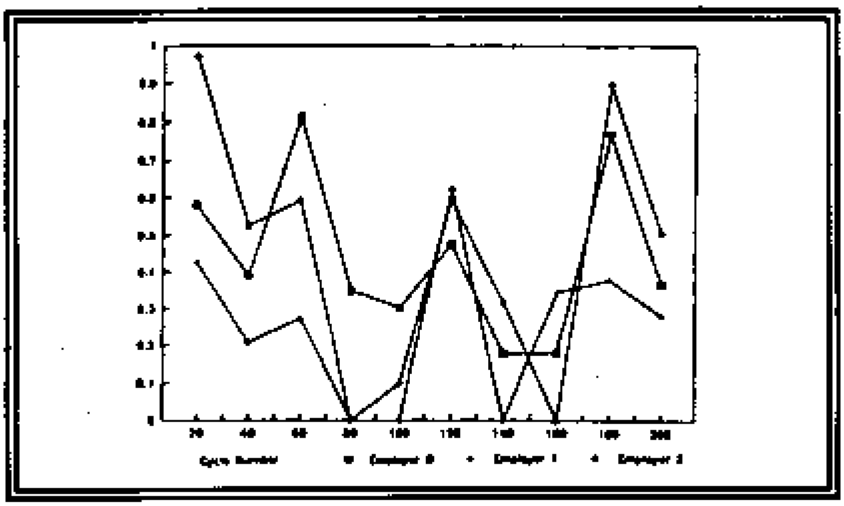

FIGURE 12

Perceptions of Fad 2. The opinions of the three employers are synchronous.

What is surprising is the synchrony of the employers' opinions about skin color in figure 13. Prejudice against black is more persistent than fad preferences. It was lifted and reinstated in the employers in synchrony. These perceptions are clearly wrong: we know for a fact that black and white have the same proportion of talented persons. We can't claim that it is simply because they are a minority: we have other simulations in which prejudice is against the majority. Figure 14 gives the average perception of white people, which, because of normalization, is just the average perception of black people inverted. Notice that the perceptions of the three employers did not cancel out, but averaged into a wave of high amplitude because they were synchronous.

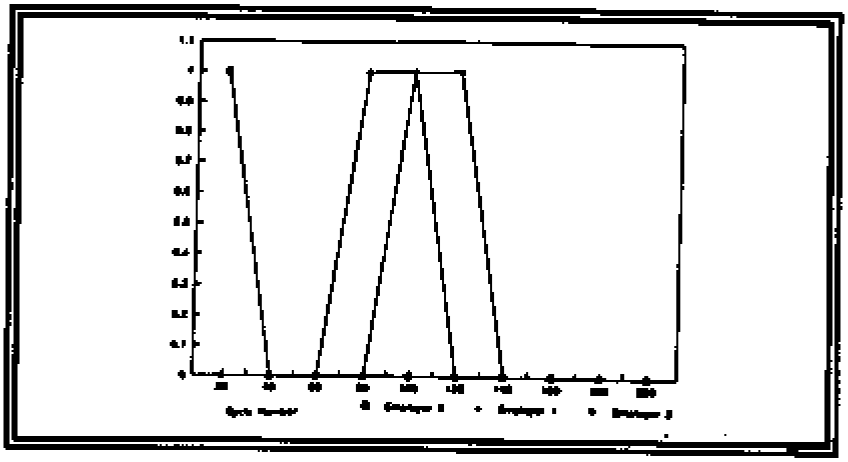

FIGURE 13

Perceptions of Black. The opinions of the three employers are synchronous. Prejudice against black is persistant.

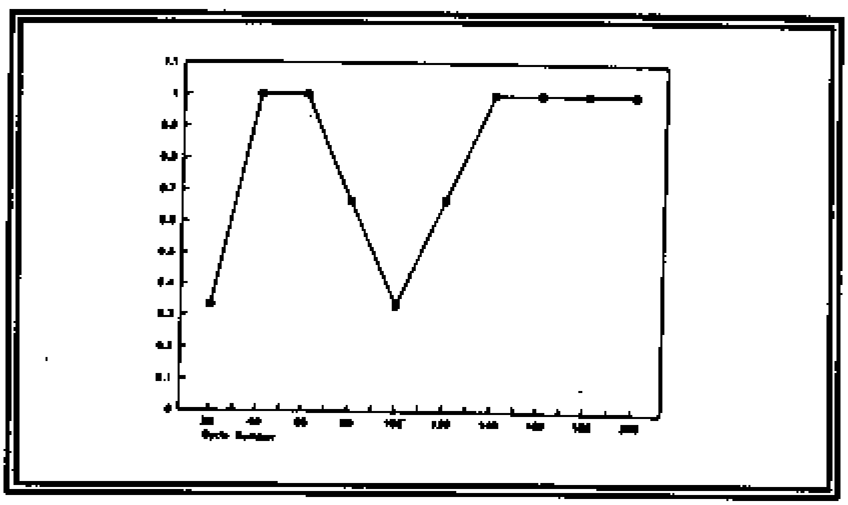

FIGURE 14

Average Perception of White. The opinions of the three employers are averaged together.

Synchrony occurs because the employers all change and are changed by the same pool of workers. The traits of the workers are the only communication between the employers, and are an effective one. To know why the synchrony occurred, refer to the details about how color is perceived as being related to other traits, and how it comes to be related to the other traits in a self-fulfilling prophecy, that is given in the historical account.

Figures 15 through 20 illustrate the workers' response to the employers. They contain the three employers' responses to a trait averaged together and the percentage of all workers having that trait.

In figures 15 through 17, the workers respond to the employers preference for fads but it is delayed, perhaps because employed workers are not required to change their traits or because the workers who do change base their choices on their past ten experiences with having their traits judged by an employer.

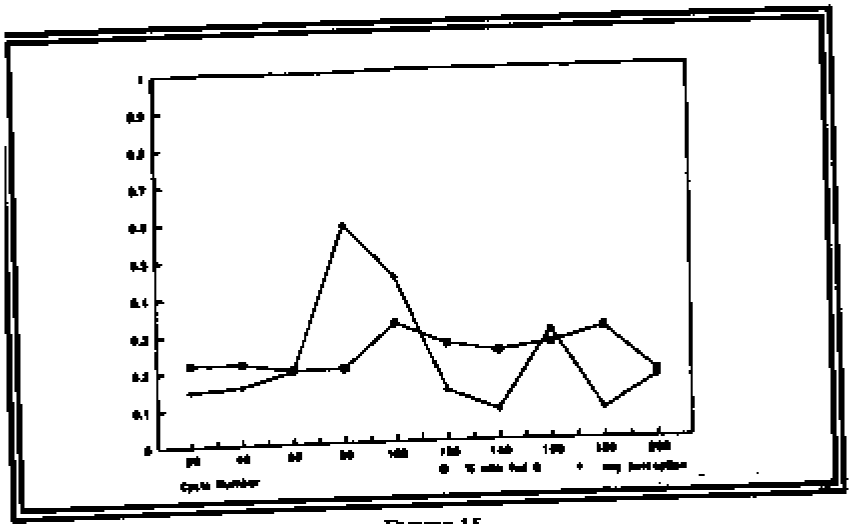

FIGURE 15

Perceptions of Fad 0 and People who Have it. The three employers' perceptions are averaged. The percentage of workers having the fad follows the employers' opinions of the fad.

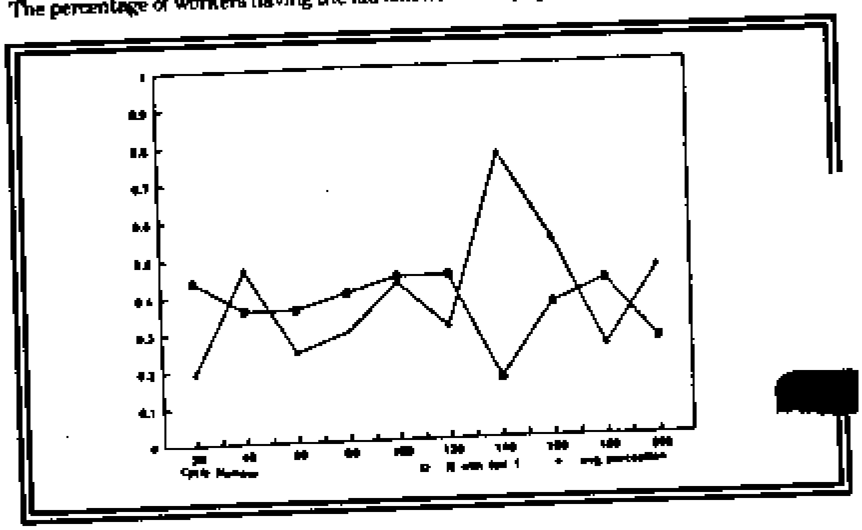

FIGURE 16

Perceptions of Fad 1 and People who Have it. The three employers' perceptions are averaged. The percentage of workers having the fad follows the employers' opinions of the fad.

Figures 18 through 20 show the average employers perception of each suit as well as the actual percentage of each suit that are talented.
They are correct in their perceptions: The perceptions of suits rise and fall with the percentage of talented persons wearing those suits. The rising curve for expensive items shows belief in status symbols. The percentage of persons wearing the suits follows the employers preference as well. This shows the purchasing of status symbols. We know that fads and suits have become symbols with shared meaning because the workers respond to the employers' perceptions of them.

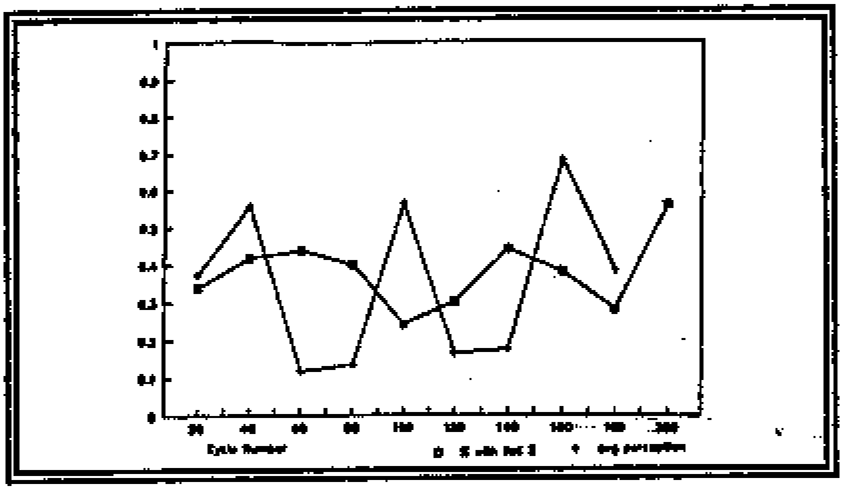

FIGURE 17

Perceptions of Fad 2 and People who Have it. The three employers' perceptions are averaged. The percentage of workers having the fad follows the employers' opinions of the fad.

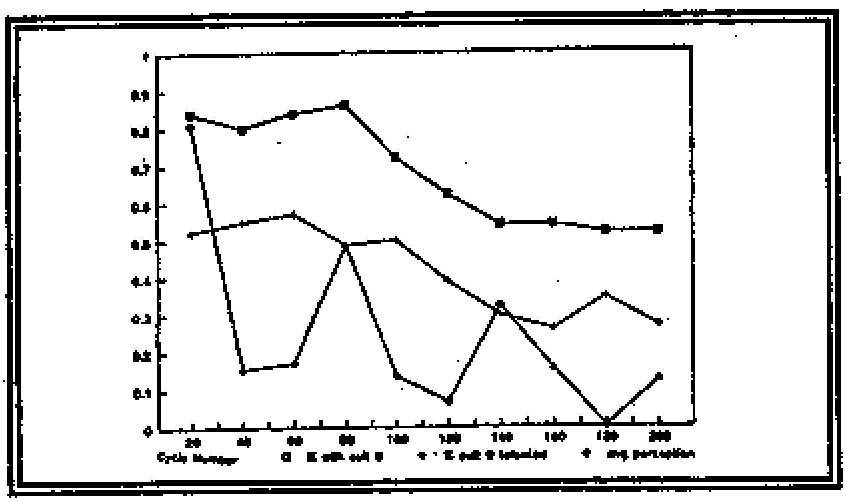

FIGURE 18

Perceptions of Suit 0 and Reality. The three employers' perceptions are averaged. Their opinions, the actual talent of the people wearing the suit, and the percent of the people having the suit are all synchronous.

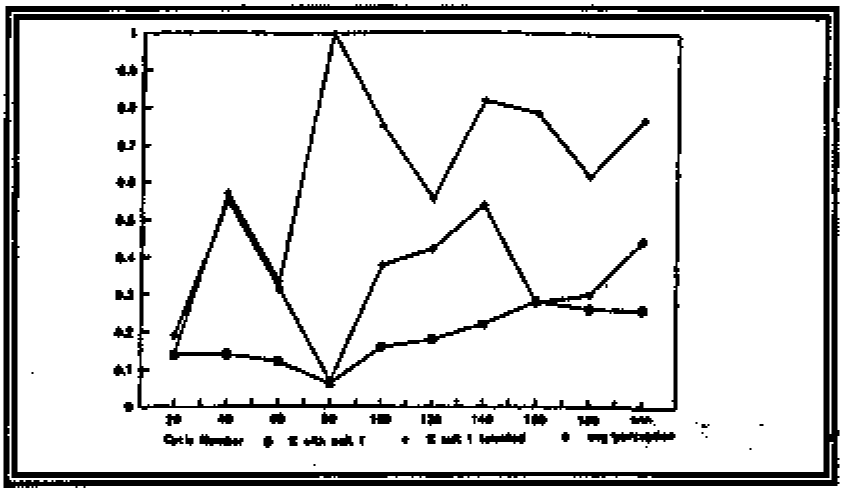

FIGURE 19

Perceptions of Suit 1 and Reality. The employers' opinions, the actual talent of the people wearing the suit, and the percent of people having the suit are all synchronous.

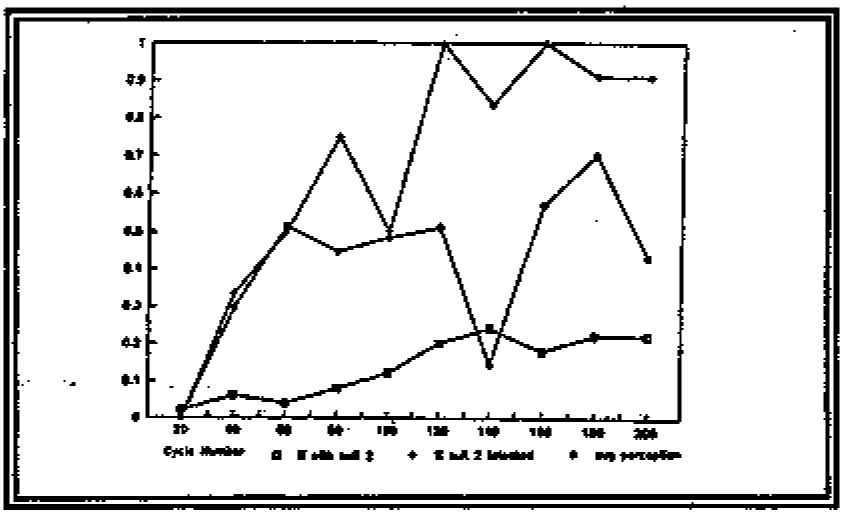

FIGURE 20

Perceptions of Suit 2 and Reality. The employers' opinions, the actual talent of the people wearing the suit, and the percent of the people having the suit are all synchronous.

The employers adapt to the workers, and the workers adapt to the employers: as they do, society unfolds. But something is missing from these averaged descriptions of society's unfolding. It is the complexity: the relations between the traits are left out. This is important: for example, a particular fad could be favored in a black person wearing a particular suit but disliked in a white person wearing the same suit. Black people could use this fact to open the door to employment and change the employers' mind about them. Trying to understand a complex system by looking at an averaged graph of each of its components alone is almost like trying to understand a paragraph by looking at graphs of the average number of each word. Important information about the interrelations of the components is missing. Without the details, the causes are missing.

The following historical description of the simulation gives some causal information, but is by no means complete. We have only

snapshots of what is happening every five cycles, and do not know what is happening between the cycles. Nor can we pinpoint which accident,
had it happened slightly differently, could blow the entire society to a different state. We can never understand "why" something happens in any complex system perfectly, but must be content with general impressions and guesses.

The History of Run 1

At cycle 5, we already see prejudice. 35% of blacks are employed while 72% of whites are, although at this early stage employers
think fads are the most important factor. By cycle 10, all employers prefer white for everything, and fads are no longer important to them. Employer 0 has decided that he likes suits, and despite his prejudice hired a talented black wearing suit 1. It was a lucky accident that there happened to be a black in a suit when there was no other whites in the same suit to get the position, and fortunately he was talented. Perhaps this lucky accident started a trend so that by cycle 15 employer
0 prefers black for most suits (by suit, we refer to suits which cost money, suits 1 and 2).

Employers 1 and 2 still preferred white for suits, however. They did not hire blacks at all. Employer 0's preference for blacks in suits over whites was a window of opportunity for black, but because he was not hired by the other 2 employers, he did not have enough money to buy any suits. By cycle 20, the window had closed: employer 0 preferred white again for most combinations of suits and fads. This is an example of how synchrony was enforced: employer 0 could not change his prejudice because employers 1 and 2 did not. Black still made do with what he had: two out of three employers preferred him for the combination of suit 0 and fad 1. Seven out of thirteen applicants responded by presenting this combination and five were employed. This is more than expected by chance because there are always between 3 and 9 possible combinations of traits any individual can choose. This is an example of class cultural structure because white did not present this trait disproportionately. Luckily, they were talented: the few blacks that were hired were representative of their race in the employers' minds, and because they were talented, they opened the door for the rest.

By cycle 25 blacks had 70% employment and whites had only 54%. Employers 0 and 1 preferred blacks for most combinations of suits and fads. Six out of nine blacks present the suit 0 fad 1 combination which is still preferred by employers 0 and 1. Again five were hired, but this time the lot was rather untalented. Perhaps this is why by cycle 30, only 11% of blacks are employed again. Employer 0 no longer prefers him for suit 0 fad 1, and employer 1 does, but now fad 1 is less popular with him. Black does what he was previously awarded for: he presents more of the suit 0 fad 1 combination than any thing else, and contributes to his downfall in doing so. Employer 1 hires whites with more popular fads instead.

In cycle 35, blacks are preferred for suits in 2 out of 3 employers,
but they again can not afford the suits because of the 17% employment rate. By cycle 40, all employers have learned that suits are a good indicator of talent. This is so when black has an 11% employment

rate and very few suits. He is only preferred in an unpopular fad in suit 0, which he presents in greater proportion, but is rarely chosen.

Apparently employers are becoming dissatisfied with untalented whites' ability to buy suits, and they become disillusioned with suits. In cycle 55, employer 2 preferred black for a popular fad and hired three of them, so that they were more popular with everybody by cycle 60. By cycle 65, blacks employment rate had grown to 29%, while he is preferred in most categories in employers 1 and 2. Correlation between suit and talent hits an all time low because of prejudice: 60% of suit 0 are talented while only 20% of suit 1 and 40% of suit 2 are. As black gains to 41% employment by cycle 70, the correlation between suit and talent returns. But that doesn't matter to black: the employers haven't relearned the suit correlation yet, they just want to employ black. Black gains to a high of 88% employment by cycle 95, compared to only 15% white employment. White is preferred for some suits, and because of his past experiences knows that he should buy them, but he can not afford to. Black is doing well, but never in his history has he been rewarded for buying suits, so he does not buy them.

First employer 2 relearns the value of suits. Then by cycle 100, all employers know the value of suits. Maybe they have learned that just because someone is black doesn't mean they are talented, either. By cycle 105, black employment is reduced to 64%, which is even with white employment. But these figures can not remain constant: even though black has more money than white, only 17% of blacks buy suits while 45% of whites do. As a class, white has learned to buy suits while black has not. White is preferred slightly across the board.

At cycle 110, black has 41% employment and begins to learn about suits. Unfortunately, he will lose in the end because the suit he prefers is too expensive for him. He has 29% suit 2's, which is quite a lot considering how expensive they are. Employers are now preferring him for suit 2. In cycle 120, 58% of blacks were employed, but white out did him in buying 33% suit 2's compared to black's 11%, so that white regained preference in most categories by cycle 125.

By cycle 135, black is 41% employed and is buying suits disproportionately higher than whites with his fewer funds, but white is still preferred for suit 1. Despite low employment, black continues to buy suit 2 which he is preferred for. In cycle 150, 4 black people are employed, three of which have expensive suit 2. Black is unable to gain in this and eventually loses to white who can better afford suit 2.

By cycle 180, 29% of blacks are employed. White is preferred in suit 0 across the board except that black is preferred by two employers for suits he can no longer afford. In cycle 190, black employment is at 35% and he is liked a little more. He uses all his funds to buy all suit 2's and no suit 1's, because he used to be preferred for suit 2. However, he is no longer preferred for suit 2. Perhaps a few did get in, because in cycle 200, two of the employers who liked him before in suits show some preference for him in some fads of suit 0 as well.

In summary, black was caught in a vicious circle: he was discriminated against because he did not have a suit, and could not afford a suit because he was discriminated against. Black was called upon to prove himself with suits, but often did not know the correct strategies to take because of past prejudice. Sometimes he was lucky enough to gain, but lost in the end to whites money and experience in buying suits. Suits reinforced prejudice, but this was self limiting because prejudice lowered the correlation between suit and talent.

DISCUSSION

Emergent Order

The results from run 1 show every type of macrolevel order mentioned in the introduction: belief in status symbols and prejudice in employers, cultural class beliefs in workers, racial economic classes, purchasing of status symbols and the existence of fads. We have achieved the objective of macrolevel social structures from microlevel associative memory, even though the simulation does not include many important factors in real societies.

A likely objection to this simulation is that it is unrealistic and too simple. Employers do not really throw out a quarter of their workforce in turn. Talent doesn't come in only two values. The real world is much more complex. With so many things not taken into account, just what does this simulation prove? Even if it were possible to take into account all the complexity of the real world, it could not predict what would happen next because it would be sensitive to initial conditions.

The point is not to predict, but to understand how societal structures may result from individual decisions, even if they are unrealistic. This simulation does not tell how prejudice comes into being as much as it tells how, once it is there, it might be "recreated" with the help of status symbols. It shows how status symbols can become a sign of what a society values (in this simulation, talent) and develop a rough correlation with their signified value even though the symbol itself has nothing to do with the value. It shows how what is perceived as a different class of persons actually becomes a different class of persons with members that develop different understandings of the world and how to deal with it because they are treated differently. It shows how shared meanings develop. It shows how all of these things can come from simple associative memory.

Semiotics

This simulation can do all of this despite its differences with real societies because it has one important thing in common with them: it treats people's traits like a set of symbols in a language. Like a language, the meanings of the symbols are common to all persons of the same society, but each member learns it on their own, by experience.
Which language, or societal structure, comes to be meaningful can not be predicted, but we know that people's traits will come to have shared meaning. Unfortunately, when traits that can not be changed become a sign of the negative, the injustice of prejudice results.

This model also instructs us about the nature of prejudice. Perhaps prejudice is so hard to change because it is a "gut level" feeling. Prejudice originates in the primitive levels of subconscious association: even the most liberal white woman in American society will fear a black man in inexpensive clothes that she passes in a dark alley at night. Prejudice is part of the systems of meaning that everyone learns in a society. A prejudiced person feels (correctly) that he has come to his opinion on his own and feels (incorrectly) that since others agree with him he must be right. Like language, prejudice in this simulation and in real society self organizes and recreates itself.

The Observable and the True

The use of employers in this simulation is not essential: we could have developed symbol systems with only people interacting, as long as they have goals and rewards. Still, one could object to the use of intuitive associative mechanisms in explaining employment practices because there are "objective" assessments of ability on which basis employers may hire. In making such an objection, we forget that subjective appearance is the only information we usually have access to. We assume the signs we use are true only out of convenience and necessity, not because they actually point to the truth. For example, when a woman applies for a job, all the information an employer usually has to judge her with is her resume and her appearance and mannerisms during an interview. Suppose the resume says that she received good grades in school. This is not objective information, because the employer does not know if she received them through conscientious effort or if her teachers, knowingly or unknowingly accepted a false representation of her ability as true (with the teacher shortage, the latter is likely). Nor does the employer know the politics involved in her recommendations. Getting good grades can be like buying an expensive suit: they prove you had enough ability to earn them, but not getting them doesn't prove that you don't have ability or that you haven't learned. Buying an expensive suit proves that you have enough money to buy it and perhaps that you were useful enough to society to afford it, while not buying it doesn't prove you don't have enough money or are not useful. Both good grades and a new suit are rewards for having done what society values. The suit in this simulation which may be bought with social rewards may represent good grades as well.

Even though many things that we take for granted as being objective are really subjective, we are not lost: correlations do develop between the observable and the true. People in this simulation and in real life learn to "show off": to get those good grades (even at the sacrifice of their learning) so as to display their abilities. If they possibly can, they buy suits and other status symbols to communicate their talent and social power, even though suits inherently have nothing to do with talent. Yet, these correlations will never be perfect: most likely they will be far from perfect. There will always be many who gain by making false representations of themselves and many who lose because they are unable to play the game, namely, the lower class.

Another objection that may be raised is that if we are so limited in our ability to perceive truth, are not even the categories we frame our social science inquiries into, including in this simulation, themselves prejudiced? Is not all social science doomed to failure, reflecting opinion more than truth? The answer is that it is not a law that there can be no good signs of objective truth and that we must be trapped by our senses or by the web of meanings of our culture. It is a fact of life we can seek methods to overcome. The natural sciences and mathematics have succeeded in making useful indicators of what is objectively true, even if truth is only revealed in the jumps of Kuhnian scientific revolutions. One thing which has the ability to reveal truth is a model that works. A model which can produce one known phenomenon from another has explanatory power which transcends our human limitations. Computer simulation is an excellent tool for developing such "models that work".

Computer Simulation, Determinism, and Dynamical Systems

Some sociologists believe that the use of computer simulation in sociology implies a strict determinism and lack of free will in human beings (Muir 1986). This may be true of single - outcome equilibrium simulations, but it is not true of self organizing simulations which are sensitive to initial conditions: their entire future may be changed at the bifurcation point by the smallest individual decision. Dynamical systems are deterministic, but because of their fractal basin boundaries, what develops during the simulation depends upon the exact arrangement of every atom of the system. Thus, if our "free will" can do so much as blow an atom from one unpredictable quantum state to another, it can change the course of history. Dynamical systems are deterministic yet unpredictable because we can not know the state of every atom. If chaos and self organization are good models of society, then the individual is not just the pawn of the environment.

REFERENCES

Allen, Peter M. "Towards a New Science of Human Systems." *Inte rnational Social Science Journal*, February 1989, 41, 81 - 91.

Baert, Patrick and Jan De Schampheleire. "Autopoiesis, Self Organization, and Symbolic Interactionism: Some Convergences." *Kybernetes*,1988 17, 1, 60 - 69.

Duong, Deborah Vakas and Kevin D. Reilly. "Neural Network and Self-Organizing

System Simulation through Container Classes in Object Oriented Programming-
Studies on Hierarchical Neural Networks-Part I."
Proceedings:Third Workshop on Neural Networks: Academic/Insustrial/NASA/Defense
(WNN 92-selected papers), pp 425-429. SanDiego, CA:Society for Computer Simulation;
Also:SPIE Vol 1721.

McClelland, James L. and David Rummelhart. Explorations in Distributed Processing. Cambridge: MIT Press, 1987.

Muir, Donal E. "A Mathematical Model/Computer Simulation of Adaptive System Interaction." Behavioral Science, 1986, 31.

Nicolis, Gregoire and Ilya Prigogine. Exploring Complexity: An Introduction. New York:W.H.Freeman and Company, 1989.

Pagels, Henry P. The Dreams of Reason: The Computer and the Rise of the Sciences of the Complexity. New York:Bantam Books, 1988.

Reilly, Kevin; Hayashi,Yochi;Duong,Deborah; and Krishnamraju,Penmatcha. "Cooperative Schemes for Conventional and Neural Expert Systems." *Proc. IJCNN (International Joint Conference on Neural Networks).* VolII. A-888, 1991.

White, Roger. "Structural Evolution in Urban Systems." Systems Research,1989, 6,3, 245 - 253.
(Manuscript received September, 1993)