

published in *Intelligent Systems: a Semiotic Perspective. Proceedings of the 1996 International Multidisciplinary Conference Vol II: Applied Semiotics 20-23 October 1996*, eds: J. Albus, A Meystel, R. Quintero pp. 349-354.

Symbolic Interactionist Modeling: The Coevolution of Symbols and Institutions

[Deborah Vakas Duong](#)

Institute for Computational Sciences and Informatics

George Mason University

Fairfax, Virginia 22030

dduong1@gmu.edu

This paper presents a methodology for simulating social systems based on the sociological theory of symbolic interactionism. It is successful in emerging macrolevel institutions such as a role based division of labor and price from microlevel display and reading of signs. This paper presents the philosophy of this type of modeling and an agent based economics simulation which implements it. This simulation uses the artificial intelligence method of coevolving genetic algorithms in its implementation, however, symbolic interactionist simulations may also be implemented with neural networks as was originally done in the author's 1991 master's thesis, a simulation of racial class and status symbols (published in Duong and Reilly, 1995). Symbolic interactionist modeling is an interdisciplinary tool with applications for micro-macro, subjective-objective integration in sociology as well as emergent knowledge representation in artificial intelligence.

Epistemology for Social Science

The most important issue in artificial social simulation is to find an epistemological basis for claiming one simulation is a better model than another. With the computer, we lose the epistemology of holding everything the same and seeing how one thing moves, and gain the epistemology of letting everything move and seeing how they cause each other. One important methodological constraint we should make in our simulations is to not put the answer to the question in the question itself. For example, if the results of the simulation follow directly from the initial conditions, then the simulation is "rigged." To make a simulation without artifact, we should start out with as few assumptions as possible, and make those assumptions as explicit as possible. If we can emerge many known phenomena on a macro level from a few that are known to exist on a micro, we have a good explanation, because it would be unlikely for many known phenomena to be derived from a few incorrect phenomena. Even though there may be several possible models, the ones that follow the "Ockham's razor" standard will be the best explanation. Emergence happens when a small set of assumptions cause phenomena which in turn cause each other, that is, when behavior is explained by the process of the simulation and not directly by the assumptions. The trick to this new methodology is to find the best recipe for primordial soup.

Simulations of institutions which tell the agents to copy each other in their assumptions don't explain well. To simulate institutions with a single genetic algorithm would be an

example of an artifact because convergence is intrinsic to this method. Institutions are not just conformity of behavior: they are interlocked, synchronous behaviors that order individuals to their advantage. How the chaos of self-interested agents may be linearized is not so obvious: Carl Menger called the institution of money a "miracle" and Adam Smith spoke of institutions in terms of an "Invisible Hand" in appreciation of their complexity. To explain institutions we need to explain not only the existence of similar behaviors, but the consonance of different behaviors within the same framework: we need to look at how roles come into being. People see others through "person schema" or roles, more than through the detailed rational strategies that simulation designed around game theory suggest. To simulate society, we need to look more to role, symbol and institution than to game-theoretic strategy.

For the simulation of economic institutions, the known micro phenomena from which we can derive the macro should include the findings of the other social sciences. With the modularity of object-oriented technology we have a new ability to integrate the sciences because different experts can now work on different parts of the same model. It is particularly important to include the findings of interpretive social science, which takes a careful look at the issues of socialness. In their book, *Understanding Computers and Cognition*, Winograd and Flores challenged the foundations of Artificial Intelligence in logic. They used interpretive philosophy to show that artificial intelligence does not properly take into account the social nature of human cognition. This is equally true of economic simulations. What is interpretive social science and how might its findings be used to make a primordial soup for social simulations?

Interpretive social science came from the school of philosophy called Hermeneutics. It influenced sociology in the symbolic interactionist and phenomenological schools, and economics in the Austrian school. Hermeneutics is a school of philosophy which arose out of questions translators of the Bible had about preserving the meaning of texts across different circumstances of life. They saw a paradox about people. People can only understand the world through their own context: information can not be directly copied from one person to another, but is understood differently by those in different contexts. Yet they still come to understand each other through institutions and language. We invent our own meanings, yet we share meaning.

In sociology, the symbolic interactionists describe a process to resolve this paradox. They claim that institutions and language are generated from the interplay of subjective perspectives on the individual level. In other words, shared meaning is an emergent property. The lower order process from which meanings emerge is the display and reading of signs. According to Peter Berger, this occurs in a three phase dialectic of externalization, objectivation and internalization. When we act upon our ideas, we externalize. Once we have done something in the public realm, it becomes separate from ourselves: this is the process of objectivation. Once it is separate from ourselves, it can come back and change ourselves and others: that is internalization. These objectivations are signs that tell us what to expect from each other. They become habitualized and interlocked into a web of meanings that make sense together. This is how we create our ideas. In the process, we create ourselves. An important special case of this is our

categorizations of people, called roles. When our ideas about people are objectified, and these objectifications change others, this is the process of self fulfilling prophecy. The lower level process is that we change ourselves to meet others expectations of us. From this emerges different types of people, or different roles.

The Austrian school of economics uses the ideas of interpretive social science to explain economic behavior. They believe that the best orders are not imposed from the government but arise spontaneously and unintentionally from the acts of people looking out for themselves. These spontaneous orders are the equivalent of institutions from sociology: they are the habitualized interlocked patterns of behavior that have common meaning to people. Some examples of economic institutions are the division of labor, price and money. Money is considered a prime example of an organic institution: unintentionally, through the barter of goods, arises a standard of trade to help everyone trade better.

Social scientists have theorized about the emergence of institutions, but have so far have had no way to test their theories. Computational Social Science is a new way to test these theories. The kind of tests a computer could do would not be tests against an empirical reality, but tests of the consistency of theories. Computers empower researchers to do proofs of sufficiency. Simulations show that a certain set of assumptions are sufficient cause for an emergent phenomenon: that some phenomena follow logically from others. If micro-macro social theories are right, then we should be able to emerge institutions in a way that is not preprogrammed or artifactual, using just the bare principles of interpretation. So, what are the principles to use in modeling agents which can interpret meanings? First of all, the agents should not see directly as though they had direct brain implants from other minds, but should see signs through their senses. They should learn to read each other based on what they find useful to read. They should interpret in the context of what they have seen before. And still, despite this autonomy, they should come to have shared meanings without being preprogrammed to do so.

An Economics Simulation Using the Principles of Interpretive Social Science

This paper presents a simulation which uses these principles of interpretation. It was written in Smalltalk on Kevin Lacobie's Agora platform at the Program on Social and Organizational Learning, George Mason University. Institutions of society which emerge from the reading and display of signs include a division of labor, networks of trade, corresponding trade plans, price, and something close to money.

This simulation is not of a modern economy, but of a simple barter economy, where no wealth may accumulate from day to day. Agents produce in the morning, trade in the afternoon, and consume at night, leaving nothing for the next day. They want to have four goods in equal amounts, and as much as they can get of them. Agents can produce some or all of the goods as they please, but these activities cost effort. An agent only has limited efforts to spend, but if efforts are concentrated on making fewer of the goods, then more will be made. This simulates economies of scale. By this design the agents will be happier if they make lots of one good and trade it away for the others, however, it is

up to the agents to learn how to trade. They develop institutions in the process of learning how to trade. They start out the simulation completely ignorant of what to produce, what to trade, how much to trade, and whom to trade it with, and what sign they should present to others to tell who they are. The knowledge they come to have to get to the answer, the development of interlocked institutions and shared meanings prerequisite to that answer, are the emergent institutions.

Implementation

Every agent has his own genetic algorithm, which is not seeded in any way, neither from other agents nor at the beginning. Each genestring of the genetic algorithm is a plan of production, trade, and a sign to display for a single day. At the beginning of the day, agents produce according to their production plans. They may produce a few of each good, or more of some and less of others as their plans dictates. Next, they trade with agents who display signs that are closest in hamming distance to a sign they seek, if they have a corresponding trade plan. Each trader has several trade plans in his genestrings which tell him an amount of one good to trade for an amount of another. Each of these plans has a sign which the agent uses to know with whom it should attempt the trade. Every agent also has a single sign which he displays to all to attract trade. At the end of the day, each agent consumes all his goods. The fitness of each string is judged by the Cobb-Douglas utility function of these goods. The knowledge representation on the genestrings is detailed in figure 1.

Knowledge Representation

- 4 bits per effort:

bit 1: production or trade?

bits 2 - 4: which good to produce

or which trade plan to perform

- 16 bits per trade plan:

bits 1 - 3 good to give

bits 4 - 6 amount to give

bits 7 - 9 good to receive

bits 10 - 12 amount to receive

bits 13 - 16 sign to seek this trade with

- 4 bits per sign to display

tag: what bits mean emerges

Figure 1. Knowledge Representation

These genestrings are long with typically over 150 bits. The sections for efforts come first. The number of efforts each agent is given per day is a parameter of the simulation: a typical setting is 8 efforts. There are four bits per effort: the first bit tells if this effort is to be devoted to trade or production. If it is production, the next three bits tell which good: otherwise they tell which one of a maximum of 8 trade plans. The first bit is a gene switch since it controls the expression of the trade plans in another part of the genestring. It serves to maintain diversity: when the trade is inactive, the trade section gains diversity through mutation, while of course it loses diversity when it is active under the influence of natural selection. There are eight trade plans after the effort sections. Each trade plan encodes a good to give, an amount to give, a good to receive, an amount to receive, a sign to seek in a trade partner. Finally, the sign section encodes a sign to display to attract traders. In a rich environment these bits would come to have meaning just as the inner nodes of neural networks come to have meaning.

This is a coevolution program because the fitness of an agent's plans depends on its correspondence with other agent's plans. Every agent has a single genetic algorithm with 100 genestrings for 100 days of trade plans. All of the first genestrings of agents participate together in a scenario of production and trade and are judged, then all of the second independently of the first and so on until all 100 are judged. On the hundredth day, everyone "thinks": reproduction occurs and the new genestrings are put back into the slots randomly, blending the scenarios together.

The use of a genetic switch to maintain diversity and the method of coevolution are new ways to use genetic algorithms developed for this simulation. A third new development is a new reproduction operator, a type of simulated annealing that augments the implicit annealing that naturally occurs in genetic algorithms. It is used instead of the roulette wheel with slots assigned in proportion to fitness, and instead of De Jong's linearization function. In the beginning of the simulation, when what is a good trade plan is less certain, the roulette wheel of partner selection lies somewhere between the distribution in proportion to fitness and the uniform distribution (though much closer to the fitness distribution). This is so diversity is not lost early on. As the simulation progresses, and more is learned, the distribution comes to be in proportion to fitness. It then starts to bias selection towards the more fit, giving more chance of selection to those with better fitnesses, breaking the symmetry that tends to occur between roughly equal genestrings. This annealed reproduction helps the genestrings to converge.

Results

This program is successful in emerging four institutions: a division of labor, trader networks, price, and something close to money. Each of these institutions are kinds of knowledge which take into account the needs of all: The division of labor gives an agent knowledge of what to produce. The ideal types give the agents knowledge of whom to

trade with so that a trader network may form. Money gives the agents knowledge of what to trade, and price gives the agents knowledge of how much to trade.

The parameters for the simulation are as follows: 10 agents, 100 genestrings each, 4 goods, 4 possible amounts of those goods to trade, 8 efforts given to each agent, 0 efforts required to trade, 1 effort required to make each good, and a utility coefficient of .25 for each good (all goods are the same), 8 trades possible in the genestrings, 4 point crossover, 0.002 mutation rate, and annealing of reproductive choice function from 0.1 to -0.1 in increments of 0.001. The number of each good is cubed in the production process, to simulate economies of scale, and the possible amounts to trade of the goods are 12, 24, 36 and 48 in the story and 6, 12, 18, and 24 in the runs reported on the charts. All of these runs have a different seed for their random number generator.

Figures 2 - 4 illustrate learning to trade. [Figure 2](#) shows that the agents learned to trade in the second half of the simulation. There are four runs shown with all the same parameters except for the seed of the random number generator. We can see that they had varying amounts of success in learning to trade. Early in the simulation, the agents learned to satisfy their needs by producing all of the goods for themselves, but later, they learned to concentrate their efforts on a few goods and trade them. [Figure 3](#) shows the increase in the average utility of agents in a particular run. All of the increase in utility is the result of trade. [Figure 4](#) shows the increase in production resulting from concentration of efforts on a few goods and trading them rather than producing for use. It shows that agents were able to form a division of labor and trader network. Note that run 3 did better than the others in terms of learning production and trade.

[Figure 5](#) shows how the agents knew whom to trade with. The signs they read and displayed came to have meaning. These charts from the last cycle of run 4 show the proportion of trades in a good that were made by people displaying each sign. Sign 12 had oats to offer, sign 8 had peas and barley to offer, and sign 1 offered beans (The goods are not differentiated, but are given names for clarity's sake).

[Figure 6](#) shows the prices on a day in the last cycle of run 3 and run 4. The low standard deviations show that a price has emerged. In run 3, we see that 21 trades between oats and peas all had the same ratio: the odds of this happening by chance are less than one in one trillion. The ratios are constant even though several different amounts of the two goods were traded. Note the few types of trades in run 3: this is because run 3 has begun to emerge oats as money, while the other runs have not emerged a standard of trade.

[Figure 7](#) shows the number and kinds of trades in run 3. The number of trades increases all the time, but the type changes. The red bar indicates trades that directly satisfy need. The other colors are trades involving a particular good as a medium of exchange, where the good obtained is eventually traded for something else. Towards the end of the run, we see that trades that satisfy directly decrease and trades to obtain a good to trade with increase. This is close to the emergence of money: to emerge money there should be even more trades using the standard. Perhaps the coming into existence of money explains the jump in trade activity at the end of this run.

Conclusion

This program illustrates a discourse of trade. The institutional shared meanings which make sense in the context of one run do not make sense in the context of another. The shared meanings of agents in this program are of the sort: "a potato farmer looks like...", "eggs come I boxes of a dozen," "three apples are worth two oranges," "if I offer bread, they will probably take it" and "if I sell corn I should also sell beans." These shared meanings fit together well, creating a consonance of culture. The culture in this program, like in human traditional culture, is annealed into consonance. Actions become interlocked and meaningful in the context of each other so that they come to be for the good of the whole, allowing the utility function of all to increase.

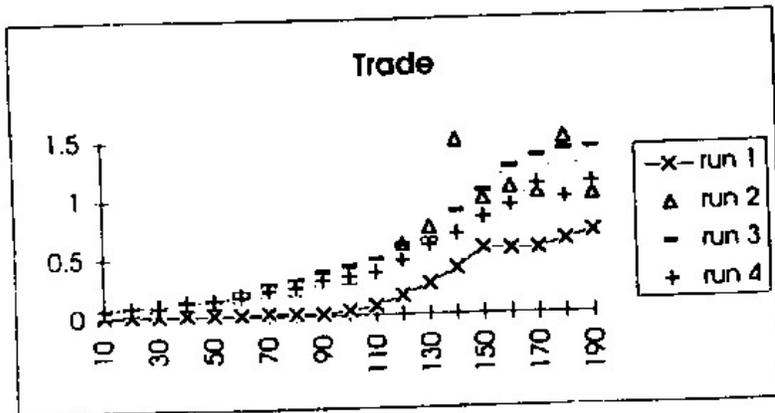


Figure 2. Trades Per Agent over Time. The X axis is cycle number.

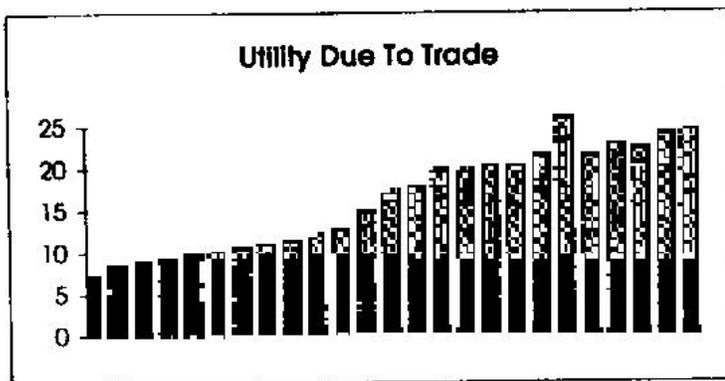


Figure 3. Increase in Utility due to Trade over Time
(The bottom is utility in the absence of trade and the top is utility due to trade.)

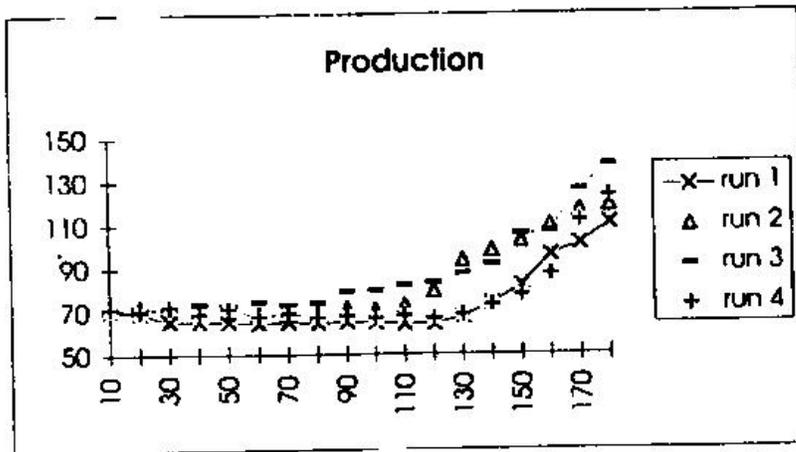


Figure 4. Increase in Production over Time.

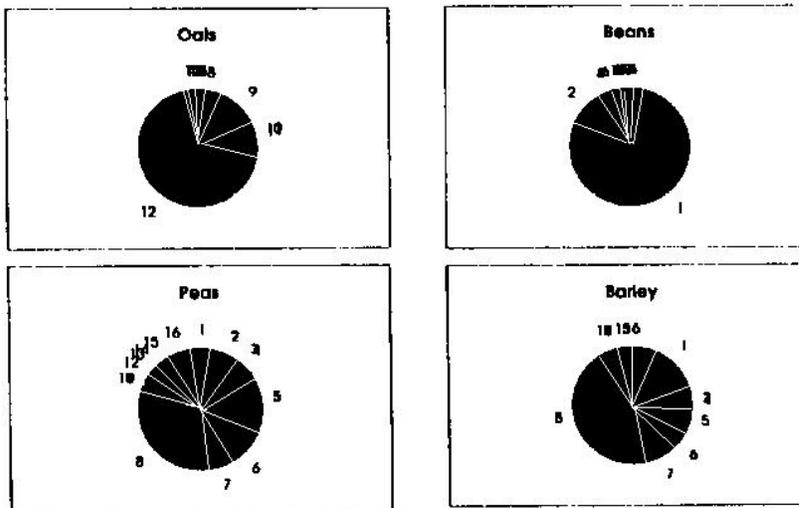


Figure 5. Who Sells a Good, According to Sign Displayed, in One Cycle

run 3			
trade type	number of trades	mean ratio	standard deviation
oats / peas	21	1	0
oats / beans	2	1	0
oats / barley	4	2.5	0.58
beans / barley	3	1.33	0.58

run 4			
trade type	number of trades	mean ratio	standard deviation
oats/ peas	1	1	0
oats/ beans	3	1.5	0
oats/ barley	12	2	0
peas/ barley	5	1.6	1.34
beans/ barley	2	0.875	1.77

Figure 6. Prices on a day in the last cycle of run 3 and run 4.

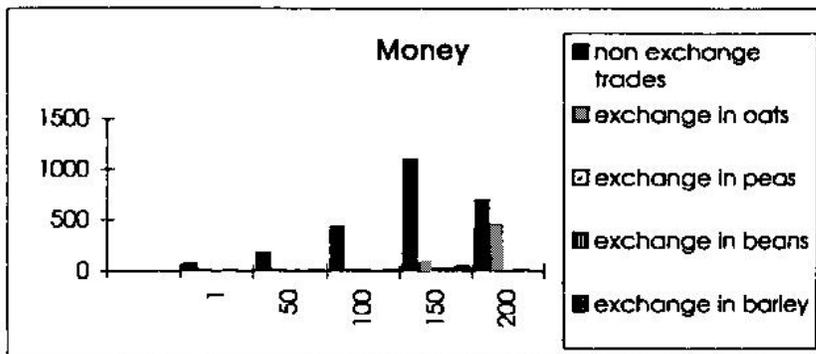


Figure 7. One Good Becomes a Standard of Trade

References

Berger, Peter and Thomas Luckmann. *The Social Construction of Reality*. New York: Anchor Books, 1966.

Duong, Deborah Vakas and Kevin D. Reilly. "[A System of IAC Neural Networks as the Basis for Self Organization in a Sociological Dynamical System Simulation.](#)" *Behavioral Science*, 1995,40,4,275-303.

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/Industrial/NASA/Defense San Diego, CA: Society for Computer Simulation, 1992, pp. 425-429.*

Goldberg, David. *Genetic Algorithms in Search, Optimization, and Machine Learning*. New York: Addison-Wesley, 1989.

Lacobie, Kevin J. Documentation for the Agora. Unpublished document, 1994.

Tesfatsion, Leigh. "A Trader Network Game with Endogenous Partner Selection" _Iowa State University Economic Report Series No. 36, 1995.

Winograd, Terry and Fernando Flores. _Understanding Computers and Cognition_ New York:Addison-Wesley, 1987.