

TJHSST Senior Research Project

Exploring Artificial Societies Through Sugarscape

2007-2008

Jordan Albright

April 3, 2008

Abstract

Agent based modeling is a method used to understand complicated systems through the simple rules of behavior which its agents follow. It can be used to explain simpler systems, such as the pattern in which birds fly, or more complicated systems, such as self-segregating neighborhoods.[4] Though the systems resulting from the interactions of the agents are not perfect replicas of more complicated societies, they lend insight into the way in which they develop. One common application of agent based modeling, Sugarscape, developed by Epstein and Axtell, creates an environment where agents follow simple survival rules within their society. Sugarscape allows for analysis of a variety of trends resulting from the agents interactions, among which is wealth distribution, and is a useful tool for social science.

Keywords: agent based modeling, wealth distribution, Sugarscape, social science

1 Problem Statement and Purpose

Agent based modeling, a bottom-up method of modeling complex situations, has become a useful method for simulating problems in the field of social science. The agents, the main building blocks of the model, are designed to

follow a set of rules or guidelines. Their interactions result in a more sophisticated global result. This approach programming lends itself naturally to social sciences because of simplistic way in which it creates societies through its components which are guided by rules directed at individual interactions rather than the group.

One common simulation using agent based modeling is Sugarscape, designed by Epstein and Axtell which is comprised of a set of agents who make calculated moves through a sugarscape a landscape that varies in the amount of sugar, a renewable source of energy for the agents, available at each square in the grid.[1] The agents, limited by vision, move around the sugarscape grid gathering sugar for energy. As time goes on, the agents continually gather the sugar, gaining energy, may reproduce, and eventually die. Some of the agents are endowed with better vision than others, and tend to be more successful than other nearsighted agents. This, and other factors, such as initial placement, creates an unequal distribution of wealth among the agents.

This behavior, though different for each simulation of Sugarscape, follows certain trends. These trends of wealth distribution naturally lend themselves to analysis using functions common to income distribution and disparity studies. Three functions that lend themselves to this type of problem are the Lorenz curve, the Gini coefficient, and the Robin Hood index (sometimes called the Hoover index).

Although Sugarscape and other forms of agent based modeling lend themselves to social sciences, because the results of such simulations focus on simpler interactions among the agents, with simpler global results, rather than complicating the interactions in favor of more realistic outcomes, sociologists are somewhat hesitant to rely on agent based modeling, favoring differential equations instead.[4] This approach leads to equations that better model the net result, but it is difficult to understand the rules that govern the individuals that make up the simulation. More research, such as the reliability of statistical analysis of the results of the interactions of agents in Sugarscape, needs to be done before agent based modeling will be used more widely in sociology and other social sciences.

2 Background

The application of agent based modeling, specifically Sugarscape, to study wealth distribution and disparity has been undertaken by a number of re-

searchers in economics and social sciences. Sugarscape does not model a typical modern society of today in which production and skill acquisition are factors in the success of agents, but rather more closely models a hunter-gather society in which gathering and trade are the way in which agents accumulate wealth in the form of sugar.

In "An Agent-Based Model of Wealth Distribution", Impullitti and Rebmann used a Netlogo version of Sugarscape to look at wealth distribution from both a classical and a neo-classical approach to economics. Impullitti and Rebmann found that inheritance of non-biological factors increased wealth distribution while inheritance of biologically based factors decreased it. Kunzar did a similar analysis of wealth distribution, though the analysis was heavily concerned with the trend of nepotism. "Simulating the Effect of Nepotism on Political Risk Taking and Social Unrest" showed that descendants of the wealthiest tended to become second class citizens and that the descendants of the lowest class remained so.

Many agent based modeling problems, such as the Impullitti and Rebmann version and this particular problem using sugarscape, are programmed using Netlogo. A version of Sugarscape called MASON programmed in JAVA is also a popular way to implement the Sugarscape simulations.

This program uses three main functions to show wealth distribution: the Lorenz curve, the Gini coefficient, and the Robin Hood index.

The Lorenz curve shows what percent of the population owns what percent of the wealth. It is usually compared to a line of perfect equality, in which 10 percent of the population owns 10 percent of the wealth, 50 percent owns 50 percent of the wealth, and so on.

The Gini coefficient is derived by comparing the area between the Lorenz curve and the line of perfect equality to the integral of the line of perfect equality. It ranges from 1 to 0, with 1 representing perfect inequality, and 0 representing perfect equality.

The Robin Hood index is the greatest vertical distance between the Lorenz curve and the line of perfect equality. Also called the Hoover Index, this is proportional to the amount of wealth that would need to be taken from the rich and given to the poor for perfect equality to be achieved. It often fluctuates in a manner similar to the Gini coefficient.

3 Research Theory and Design Criteria

The Sugarscape agents behaviors are specified by a set of guidelines. One of these guidelines involves searching for food: in each timestep, each agent determines which patch or patches of the Sugarscape would be the best place to move. This is done within each agents scope of vision, a number specified by the user (usually between 1 and 10 patches). The agent looks north, south, east, and west, in the scope of its vision and determines the patches with the most sugar that is not already occupied by another agent. Then the agent randomly selects one of the best patches and moves to that patch. This is done by each agent individually, rather than simultaneously, to prevent two agents from occupying the same patch. The agent then gathers all sugar on the square, which it stores as energy, and subtracts from its energy stores various unit of energyfor metabolism, which varies randomly from turtle to turtleand one unit of energy for each square forward it moved from its previous location.

At each timestep, the agent may also reproduce. This occurs if the agent has enough energy to do so; this amount of energy (between 1 and 100 units of energy) is determined by the user. If the agent reproduces, it subtracts the birth energy from its energy store, and another agent is hatched on the same square as the agent. The user may choose what attributes of the parent agent will be inherited by its offspring. There is a switch that allows for the inheritance of vision and metabolism. There is not a switch that allows for the inheritance of wealth, though this may be a switch that allows for this in the next version.

At each timestep, the agents may also die. This happens either after 80 timesteps to simulate death due to ageor if an agent cannot maintain an energy surplus.

Each timestep, the amount of sugar in the patches adjusts to reflect the consumption by the turtles. If a turtle moves to a specific patch, that turtle removes all sugarenergyfrom that patch. Every other timestep, patches regrow their sugar by one increment.

While the turtles are moving throughout the sugarscape, a number of different mathematical analyses run in the background and graphical representations of these analyses are shown as well.

3.1 Algorithms

This version of Sugarscape utilizes three different algorithms to analyze wealth distribution: the Lorenz curve, the Gini coefficient, and the Robin Hood index. Both the Gini coefficient and the Robin Hood index are derived in relation to the Lorenz curve, but they offer different information regarding wealth distribution.

The Lorenz curve is usually plotted in relation to the line of perfect equality. The line of perfect equality describes a population whose wealth is distributed evenly among individuals. For instance, ten percent of the population would own ten percent of the wealth, fifty percent would own fifty, and so on. The Lorenz curve plots the actual distribution of the wealth. For instance, sixty percent of the population may own forty percent of the wealth, and seventy may own forty-five percent. The Lorenz curve is usually calculated using the cumulative distribution and the average size, μ :

$$L(y) = \frac{\int_1^y x dF(x)}{\mu} \quad (1)$$

The Gini coefficient represents the ratio of the area of the Lorenz curve to the area of the triangle of perfect equality (the integral of the line of perfect equality). It is usually calculated using the mean difference between every possible pair of data points:

$$G = \frac{\sum_{i=1}^n \sum_{j=1}^n |x_i - x_j|}{2n^2\mu} \quad (2)$$

The Robin Hood index represents the amount of wealth that would need to be redistributed (taken from the wealthy individuals and given to the poorer ones) in order for there to be perfect equality. It is calculated by finding the greatest vertical distance between the Lorenz curve and the line of perfect equality.

$$H = \frac{1}{2} \sum_{i=1}^n \left| \frac{E_i}{E_{total}} - \frac{A_i}{A_{total}} \right| \quad (3)$$

Robin Hood index is also a good indicator of public health, though that is not the purpose for which it is used here.

4 Results

The goal of this project is to provide insight into how wealth is distributed in a free trade society. The society is limited in its production and resembles more of a hunter-gatherer society in which each agent gathers as much food as it can. This model is developed using a Sugarscape society written in Netlogo, whose agents are limited by age, metabolism, and vision.

The wealth distribution in this version of sugarscape, as in the simulation created by Impullitti and Rebmann, varies greatly depending on inheritance. If metabolism and vision are inherited, the Gini coefficient varies by an average of 0.8, with the average Gini coefficient over 800 timesteps of the non-inheritance simulation at 0.37 and the average Gini coefficient over 800 timesteps of the inheritance simulation at 0.44. This reflects a much greater inequality when the agents are able to inherit the "genes"—good or bad—of their parent agents. It is important to note, however, that the wealth distribution during inheritance simulations is much more stable than the wealth distribution of the non-inheritance simulations. However, the average metabolism of the group over time behaves differently than the average vision. The average metabolism and vision of the group are very sporadic in the non-inheritance simulation; it simply falls somewhere in the range of possible metabolism or vision. In the inheritance simulations, the average metabolism of the group tends to approach a lower limit. However, although high vision often contributes to the acquisition of energy, it is not directly related as metabolism is, thus it does not – and cannot – approach an upper limit; the success of individuals with high vision is dependent largely on location.

5 Conclusion

This project, as well others like it, is attempting to make simulation models more useful to social sciences. Small disturbances and changes in initial conditions can be quickly quantified here, and though the resulting interactions are much more simplistic than real interactions in societies and organizations, the insight taken from simulation models can be used to make improvements in real societies and organizations.

6 Recommendations

Most of the results in this project are derived from simulations with the same initial conditions. It would be useful to analyze the effects of initial conditions on wealth distribution. Also, it would be useful to measure how inheritance of wealth – not just genes – affects wealth distribution. This is something I intend to explore fourth quarter.

References

- [1] J. M. Epstein and R. Axtell, Growing Artificial Societies: Social Science from the Bottom Up, MIT Press and Brooking Institution Press, Washington, D.C., 1996.
- [2] G. Impullitti and C. Rebmann, "An agent based model of wealth distribution", New School University, September 2006.
- [3] L. Kunzar and W. Fredrick, "Simulating the effect of nepotism on political risk taking and social unrest, June 2005.
- [4] M. Macy and R. Willer, "From factors to actors: Computational sociology and agent based modeling", Annual Review of Sociology, pp. 143-66, August 2001.
- [5] L. Smith, "Complexity Meets Development - A Felicitous Encounter on the Road of Life", pp. 151-160, October 2007.