SOCIETY OF ACTUARIES

Article from:

Forecasting & Futurism

December 2013 – Issue 8

A Return Visit to the Sugarscape

By Ben Wolzenski

ABSTRACT

This article describes the use of artificial society modeling to gauge the effect of insurance agent population and effectiveness on individual life insurance sales. This is part of an ongoing effort to extend the use of one type of agent-based modeling beyond health care, but is far from a practical at this point.

The December 2012 edition of this newsletter contained the article "Artificial Society Modeling with Sugarscape." In brief, the article described an artificial society as an agent-based model in which the user defines the rules for the agents ("citizens") and the environment. Sugarscape is an artificial society model described by Joshua Epstein and Robert Axtell in their pioneering book, *Growing Artificial Societies.* The article went on to describe how the online applet for Sugarscape could be adapted and interpreted to roughly model the effect of societal changes on future life insurance sales.

For example, a simulated increase in unemployment led to fewer life insurance purchases; delayed household formation produced fewer short-term but greater long-term insurance purchases; and combining increased unemployment, deferred household formation and increased productivity led to greater variability of results over multiple model simulations.

Additional modeling has explored the effect of changes in the relative population of insurance agents and their effectiveness. But first, let's review the Sugarscape model a bit. The Sugarscape model is a large grid, with an initial population of "citizens" who need to move about to gather goods they need to survive. In moving about, citizens can meet each other. When they meet, one of several things can happen. One citizen may cause the other to change to the "cultural group" of the first. Disease may be transmitted from one to the other. The two can mate, creating a child. Most frequently, the two can engage in trade. In trade, the two citizens exchange goods. In our customized version of Sugarscape, one such transaction is a life insurance purchase.

Ben Wolzenski

Ben Wolzenski, FSA, MAAA, is managing member at Actuarial Innovations, LLC in St. Louis, Mo. He can be reached at *bwolzenski@rgare.com.*

One of the features in the customized model is the random designation of a small percentage of citizens to be insurance agents. When one of the citizens is an insurance agent, there is a greater probability that the transaction is a life insurance purchase. What happens when the relative population of agents decreases? To what extent could this be offset by increased sales effectiveness of insurance agents? What outcomes would be produced by greater effectiveness of insurance sales not involving insurance agents?

The baseline Sugarscape model for exploring these questions is one in which the probability that a citizen will buy a life insurance policy in any year is close to recent U.S. experience (about 3.4 percent) and in which the probability that the purchase is from a life insurance agent is 90 percent.[1]

The first simple modification produced close to predicted results. What if there were 50 percent fewer agents, with no other changes? The total number of sales should drop to 56 percent of the previous level.[2] The actual average result was 58 percent of baseline, which was six-tenths of a standard deviation greater than expected.[3] The proportion produced by agents was 81 percent, as expected.

Similarly, would the algebraic prediction of the needed increase in frequency of non-agent sales produce the original number of sales? A 50 percent reduction in the 90 percent of baseline sales from agents equals 45 percent of baseline, so the remaining 55 percent would have to come from non-agents. Non-agents would have produced 10.6 percent of the baseline number without an increase in frequency,[2] so they would need a rate 5.2 times the baseline probability of sale by a non-agent. When that increased probability was tested in the model, the resulting total number of sales was 98 percent of the baseline, with 45 percent coming from agents, as expected.

With these results, I expected that an algebraically predicted increase in agent productivity would make up for the smaller number of agents. However, the model did not produce such a result. When productivity (probability of a sale upon contact with an insurance agent) was nearly doubled[4] to offset the 50 percent decrease in the number of agents, the mean number of sales was only 84 percent of the baseline amount, about three standard deviations lower than expected. The shortfall was entirely due to low agent production; non-agent production was actually slightly higher than predicted. My first reaction was that I must have coded the input incorrectly, so I would have to rerun all the simulations and re-record the results. But when I checked the input, it was correct, so something else was going on in Sugarscape.

In the model, whether any transaction is a life insurance sale is determined by whether a number, obtained by successive multiplication of probability factors, is greater or less than a random number between 0 and 1. The base probability starts out fairly low, but is adjusted based on the citizen's age, cultural group membership, wealth, number of children, and most significantly by whether the transaction is with an insurance agent. It turned out that the effect of successive multiplications with much higher agent productivity (probability of producing a sale) was to produce a significant number of comparison numbers greater than 1. (For example, the baseline probability of an insurance sale if the transaction is with an insurance agent is 0.5, and this was increased to 0.99 with higher productivity. However, if the citizen had three children, both probabilities would be increased 30 percent, to 0.65 and 1.29, respectively.) Any number greater than 1 is effectively wasted productivity, since the random number to which it is compared cannot exceed 1. Thus, doubling agent productivity did not double the number of sales. Of course this was just an idiosyncrasy of the model, but it suggests an analogy to a point of diminishing returns for agent productivity in the real world.

This suggests that it is worthwhile to make other tests about the boundaries on the interpretive use of the model. In one such test, I looked at factors that would influence the size of the population, which for our purposes must be relatively stable after an initial period. Tests showed that if the maximum vision of citizens (how far away they can see to find goods and other citizens) is too small, the population will die out unless the environment is richly endowed with goods, and even then the population may be highly unstable

WHAT HAPPENS WHEN THE RELATIVE POPULATION OF AGENTS DECREASES? TO WHAT EXTENT COULD THIS BE OFFSET BY INCREASED SALES EFFECTIVENESS OF INSURANCE AGENTS?

due to other random variables. On the other hand, the initial population density (the percentage of cells occupied by citizens at time equals zero) appears to have no effect on the ultimate level and stability of the population. Rather, it is determined by how other model variables are set. For example, an initial population density of 5 percent produces about the same ultimate population size and stability as an initial population density of 80 percent if all other model variables are the same. Clearly, I still have much to learn about the artificial world of Sugarscape.

Readers' questions about the Sugarscape model and any suggestions as to future testing are welcome! ▼

ENDNOTES

1. A 2005 article in *The Actuary* reported: "In 2003, independent agents accounted for over half of sales, career agents sold about 40 percent and the remaining 10 percent came from a number of newer channels including brokers, web sales and banks. Newer channels are growing their market share…." While brokers would be included in my definition of agents, Web sales and banks would not. Ten percent is a rough estimate of what non-agent sales may have grown to 10 years later (2013 versus 2003). Better data from any reader would be welcome.
2. Agent sales would be 50 percent of 90 percent, which equals 45 percent of the baseline total. The 10 percent of non-agent sales would increase slightly to 10.6 percent of the baseline total. That is because with the same total population, a decrease in the number of agents would result in an increase in the number of non-agent citizens. The total is rounded to 56 percent in the text above so as not to overstate precision.
3. The results compared were 30 simulations of 1,000 generations each time. The standard deviation of the total number of sales was approximately 4 percent of the baseline number of sales.
4. The increase in the relative population of non-agents at baseline productivity only required 198.7 percent of baseline agent productivity for the algebraic prediction.