MODELING ANNUITY POLICYHOLDER BEHAVIOR USING BEHAVIORAL ECONOMICS AND COMPLEXITY SCIENCE

Michael Shumrak, FSA, FCA, MAAA and Marshall Greenbaum, ASA
Ernst and Young, LLP

Vince Darley, Ph.D. and Robert Axtell, Ph.D.
Bios Group, LP

Abstract

This paper describes an innovative approach to stochastic modeling of variable annuity policyholder behavior, combining behavioral economics and agent-based computational modeling. This permits realistic rendering of behavior at a highly *disaggregated* level—the level of individual salespeople and annuity policyholders. The model has been implemented computationally, and is capable of generating lapse rates that are in qualitative agreement with historical data from an annuity lapse study performed by LIMRA. The paper highlights the advantages as well as the challenges inherent in this new approach to modeling customer behavior.

Introduction

Variable annuity product sales continue to be the fastest growing sector of the retail insurance market. For several years now, competitors have tried to distinguish their variable annuity products in the eyes of their distributors through a proliferation of ever enhanced guaranteed benefits. The most recent is the so-called guaranteed income benefit (GMIB). Briefly, a GMIB rider allows a policyholder to annuitize based upon a guaranteed account value in the event that their current account value (a market value) is lower than the guarantee after a specified period of time. Policyholder behavior is a critical element of pricing these benefits. Predicting how many policyholders will elect the benefit dictates prices. After a wave of questionable direct and reinsurance market pricing of another guaranteed benefit, guaranteed death benefits (GMDBs), there is concern about repeating this scenario for GMIBs. Pricing and analyzing the risk (variability of profit) for GMIBs requires estimates of policyholder behavior where there currently is no experience to draw upon. Behavioral assumptions cannot be validated for many years to come. Clearly, statistical modeling techniques relying on historical data cannot be used since there is no such data. This has left the industry in something of a dilemma: while the guaranteed wrap benefits are very popular with the sales force and presumably their customers, have the early entrants to GMIB adequately priced in the cost of this benefit?,

Fortunately there *is* an alternative: an approach that combines behavioral economics and complexity science has the potential to model aggregate customer behavior effectively and accurately even in the absence of historical data. We will present a model we have constructed which illustrates the key concepts for the simple case of annuity policyholder lapse rate prediction. We will also outline the steps which would be necessary to extend this model to predict the behavior of cohorts of customers who have bought any particular product including these new GMIBs or the next wave of “guarantee wrap” product designs.

Behavioral Economics

Much of economic theory is built on the notion that individuals are perfectly rational. This strong assumption makes it relatively straightforward to build precise mathematical models of economic and financial processes. However, there is strong empirical evidence that real humans do not behave perfectly rationally (e.g., Lowenstein and Thaler [1989], Thaler [1994]). The field of behavioral economics is primarily concerned with developing a better understanding of how humans actually behave in economic contexts (see Rabin [1998] for a recent review article). Very recently, behavioral economics has been brought to bear on a variety of questions concerning retirement (Aaron [1999]), including people’s attitudes toward annuities (Lowenstein *et al*. [1999]). Behavioral economics is very much related to and is being supported by marketers’ increased interest in individual specific customer information developed through data mining (also known as database marketing).

Agent-Based Modeling

Agent-based models and related techniques are commonly utilized in the study of complex adaptive systems (CAS), pioneered at the Santa Fe Institute. The Santa Fe Institute (SFI) was founded primarily by scientists from the Los Alamos National Laboratory (LANL), home of the science of atomic and nuclear weaponry. SFI is well known in scientific circles for the development of chaos and complexity theory. In brief, chaos is the study of very complicated dynamics that can result from low dimensional systems of nonlinear equations. It has been applied primarily in the natural sciences to date. By comparison, complexity may be considered the obverse of chaos, for it usually studies the appearance or emergence of relatively low dimensional patterns and statistical regularities in large populations of heterogeneous, interacting individuals

Agent-based models (ABMs) are a tool for studying CAS, and represent a relatively new approach to modeling.[1] Such models are built at the level of the individuals that make up a population and are sometimes called 'individual-based models.' Agent-based models are implemented on a computer and consist of a relatively large number of data structures (called "objects"), each representing a single individual, that are instantiated and permitted to interact according to either fixed or evolving rules of behavior (called "object methods"). While each agent is usually relatively simple in such models, at the aggregate level complex, emergent phenomena may arise. To date, ABMs have been created to study a variety of social processes, from financial markets (Arthur *et al.* [1997], Darley *et al.* [1999]) to firms (Axtell [1999]), retirement patterns (Axtell and Epstein [1999]) and traffic jams (Nagel and Rasmussen [1995]). For general introductions to ABMs see Axelrod [1997] or Epstein and Axtell [1996]

Recently, attempts have been made to apply CAS ideas to business problems. In particular, agent-based models have been built in areas involving consumer and capital markets. . As an example, Darley *et al.* [1997] describe an ABM built to simulate the NASDAQ stock market. The artificial stock market is capable of reproducing many of the macroscopic statistical properties observed in the actual market (e.g., price volatility) as well as the emergence of certain characteristic microscopic behaviors (e.g., so-called SOEs bandits). This model seems to be much more capable than even the most sophisticated econometric models at reproducing the qualitative behavior of this market. A very different practical example is agent-based modeling of traffic. The LANL team that pioneered the CAS interpretation of traffic has most recently been involved in the construction of complete models of the road and traffic systems of Albuquerque and Dallas-Fort Worth, for purposes of aiding transportation planners with the next generation of highway investment (see Casti [1996] for an overview). These models have much higher fidelity than conventional statistical or differential equation models of traffic flow. While the conventional approaches focus on average behavior of 'typical' drivers and vehicles over relatively long times, the

[1] Note that in this field, 'agent' is used to refer to any given individual: the units between which interactions take place. So, throughout this paper, we will use the full description 'insurance agents' or the terms 'broker' or 'seller' to refer specifically to the people who sell policies.

ABM approach ultimately resolves the acceleration and deceleration events of individual vehicles over very short time scales (second to minutes) as drivers adapt to ever-changing local conditions.

In the context of the present problem—modeling annuity lapse rates—the notion that customers are distinct yet interdependent is the starting point of the model. Customers interact, directly and indirectly, in the following ways: (i) they interact with each other via social networks; (ii) they interact with insurance agents who may give all of their own clients the same kind of advice; (iii) many read the same news sources which will have information on the economy, markets, retirement etc.

Creating an ABM consists of building behaviorally realistic models of customers and insurance agents by giving each class of individuals explicit rules of behavior.[2] In the model described below we have used rules that are merely plausible, but we foresee using techniques frombehavioral economics, such as surveys, focus groups, conjoint analyses, interviews, and so on, in order to empirically ground such models.[3] Within the model these customer-agents/insurance salespeople-agents act and *interact with each other* under different conditions (such as variations in the capital markets, financial news and advice, etc.). Once the behavioral models have been constructed, the model's group behavior and individual behavior can be calibrated to existing historical data for known scenarios, if desired.

While ABM techniques can be calibrated to perform as well as conventional techniques under previously observed capital markets conditions, ABM provides a more robust approach to the real interaction of factors driving policyholder behavior. Effective use of these techniques can be applied to develop policyholder behavior dynamics for pricing product benefit designs, evaluating policyholder conservation programs and evaluating the impact of economic scenarios not seen before.

---

[2] Note that while ABMs are usually built in terms of individuals having explicit rules, they bear only a superficial resemblance to other rule-based systems such as expert systems. This is because the individuals in ABMs are heterogeneous in important ways (e.g., in terms of income and wealth, preferences) and so in effect have heterogeneous rules. Thus, small changes in rules commonly yield small changes in model output.

[3] It turns out that in a variety of ABMs built to date, the overall results of the model depend quite weakly on details of individual behavior. Stated differently, many different rules of behavior at the individual level yield the same macroscopic outcomes. This phenomenon is known as 'universality' andsuggests, should it be found to manifest itself in the present model, that only a modest effort would be required with behavioral economics techniques in order to build functionally capable ABMs in this domain.

Also, after constructing the behavior model, sensitivity tests can be performed to establish confidence intervals on the output. If we do so, we can immediately see to which of those parameters the new predictions are especially sensitive, and on which they only depend weakly.[4] This is tremendously important in situations for which there is no historical data. In these cases, we get predictions, expected errors in those predictions, and an identification of what we ought to examine if we want to decrease those errors. For example, we might run more focus groups, interviews etc. to gain a better understanding of those behavioral parameters underlying the ABM.

In summary, for known situations, ABM can be a more natural, understandable and convincing model to non-technical senior management than conventional statistical techniques. If large amounts of data are available on policyholders and prospects, the ABM can provide more reliable and intuitively consistent modeling to policyholder behavior than the traditional curve-fitting approach. In those situations where new products are introduced in the absence of experience data, ABM can be a very useful tool. In situations where there is little or no credible policyholder behavioral data, ABMs can provide to be invaluable. Pertinent examples in the variable annuity arena would be the lack of GMIB customer behavior experience and a lack of historical experience for policyholder reactions to new capital market environments that have not occurred since a product's inception.

Please note that in ABMs practitioners use the term 'agent' to refer to any given individual - the units between which interactions take place. This is not to be confused with insurance agents who sell the products. So, throughout the balance of this paper, we will use the terms 'insurance agents' or 'broker' to refer specifically to the people who sell policies and the term 'agent' in its ABM context.

A Model of Customer Lapse Behavior

We have built a simple agent-based model of annuity customer behavior. There are two distinct populations of agents in this model of annuity lapse: policyholders ("customers") and policy-brokers ("sellers"). The model can simulate the actions of millions of individual policyholders, and some thousands or tens of thousands of brokers. These two types of agents have distinct rules of behavior, which we have modeled very simply at present. In practice, any particular insurer who wanted to calculate the benefit costs based upon policyholder behavior and/or or risk attached to these benefits within their own particular customer base, should use information derived from market research

[4] Such tests can be used to expose 'universality' in ABMs (see previous footnote).

of their own particular distributors and customers. They would use techniques such as surveys, focus groups and conjoint analysis, in order to calibrate the behavior in the model with their own customers and brokers. This would establish the behavioral economics for each particular product-channel-customer situation.[5]

Customers

The number of customers is user-specifiable in the model (as many as 1,000,000 on a Pentium portable appears feasible). Each customer holds one policy, which can have any duration (policy year) in the model. The customer object keeps track of the age and sex of each customer. Each customer has a particular asset allocation between variable and fixed products.

In each period of the model, each customer decides whether or not to lapse its policy according to certain rules. In the particular model described here each agent makes an annuity lapse decision once annually. Before a customer is activated for decision-making purposes, there is first some probability that it will die. If it dies it is removed from further lapse decision-making. If it does not die there is then some small probability that is lapses its policy simply for liquidity reasons. If a customer has not died and not lapsed it then checks for messages from the broker who sold it its policy, supplying advice on whether or not to lapse the policy. The customer acts upon the advice with certain probabilities and decision criteria. If the customer has not received advice, they either decide on their own what to do with some probability and decision criteria, or perhaps look around at others in the customer population—their 'friends' or reference group—to see whether or not they are keeping their policies – and they imitate them. Finally, each period each active customer has its age incremented by one year. The precise parameterization of the model we have developed is described in the appendix.

Sellers

The number of sellers studied in our model can be specified by the user. We have looked at a range of 100 to 10,000. We configured the seller population such to affect a skewed distribution of the number of customers per seller. Each period each seller contacts a portion of their customers giving more attention to policyholders with higher account values (i.e., the seller's more valuable customers). For each customer contacted, the seller provides an optimistic

[5] It should be noted that there exists no systematic procedure today for turning the results of such efforts at eliciting actual customer and distributor behavior into the kinds of rule specifications needed by the model. Presumably, such procedures will be developed over time as experience with this modeling methodology grows.

assessment of returns the customer might get by lapsing its current annuity policy and buying a new policy from their broker. The sellers' optimistic forecasts vary by broker depending on their assessment of expected market returns and current competitor fixed annuity rates levels. A typical surrender penalty schedule is used. In the years of high penalty, the seller's advice is unlikely to be accepted. As with customers, more details regarding seller behavior is described in the appendix. And as with customers, effective data mining is key to developing robust behavioral economic behavior patterns for the sellers.

In essence, the interaction between customers and sellers is strategic self-interest on the part of both agents in the ABM. The sellers have incentives to move their customers into new products. The customers will go along with seller recommendations in most cases if they believe there is net gain from doing so. As such, this is a game theoretic model, a kind of principal-agent set-up with asymmetric information among the players; for a review of game theoretic models from a behavioral perspective, see Camerer [1997].

Scenarios That Drive the Model

External market scenarios are incorporated into the model, through interest rates and rates of return (equity and credited rates as well as competitor rates). We performed tests using three scenarios. Simplistically, these may be described as (1) a 'flat market' scenario; (2) an inflationary scenario in which real returns to equity are low, a 'bad market'; and (3) a low inflation scenario in which equity returns are good, the 'good market.' These scenarios generate distinct output.

Typical Model Output

The primary output of the model is the annual lapse rate by policy duration shown both weighted by policyholder and by account value. Multiple realizations of the model (10 is the default) for identical parameters yield distinct lapse rate time series, results that can be averaged or overlaid on one another with the run-to-run variation characterized. This variation can of course be used to estimate risk, and to estimate the model's sensitivity to particular behavioral parameters / characteristics.

A direct consequence of this methodology is that we can isolate those aspects of the behaviors on which the model (a) doesn't depend very much where we can use aggregate statistics, simpler surveys, etc. to estimate these aspects of the given customer cohort and (b) depends very sensitively – in this case we may wish to gather more data and conduct more extensive surveys to ensure the model's accuracy.

Typical output from 10 realizations of the model for the 'flat market' case, parameterized according to the values given in the appendix, is shown in figure 1 below, a time series plot of lapse rate over time. Note the modest run-to-run variation.

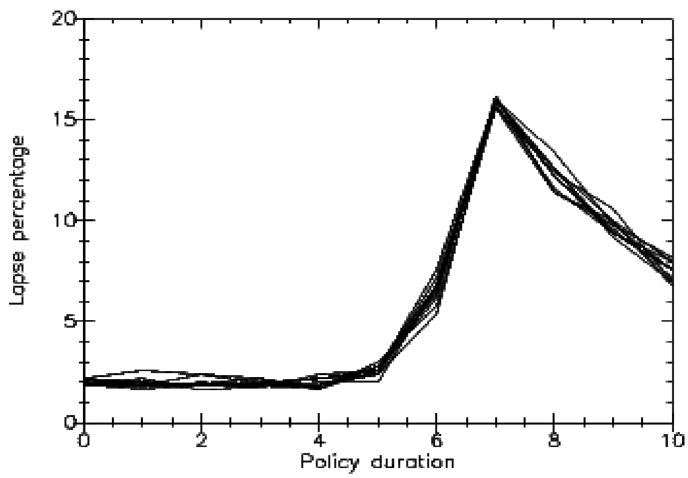

Figure 1: Lapse rates for the 'flat market' case

These results are in qualitative agreement with LIMRA data. The model also reproduces another characteristic feature described in the LIMRA annuity persistence study data, namely, large accounts have higher lapse rates.

The calculations have been repeated for a statistically identical population of agents but in the inflationary scenario. The lapse rate time series results are shown in figure 2.

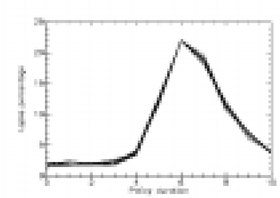

Figure 2: Lapse rates for the inflationary case

Here we see somewhat higher lapse rates in years 4 - 7, and somewhat lower rates thereafter. This is in accord with intuition since competitor rates are relatively high in this scenario and market performance has been poor, dissatisfying the customer.

Lastly, the 'low inflation' / good market scenario was run, yielding the lapse rate time series shown in figure 3.

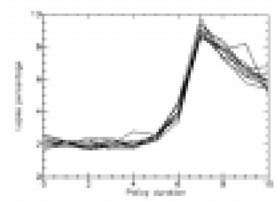

Figure 3: Lapse rates for the low inflation case

Here we see relatively low lapse rates across the board, due to the high returns to equity and relatively low competitor rates.

The results from an agent-based model policyholder behavior model such as this can be incorporated directly into any existing actuarial pricing and risk analysis modeling approach to determine ranges of profitability and risk associated with a given product. In the case of the GMIB benefit, the behavioral economics and agent-based modeling approach would be extended to consider the required policyholder behavior elements such as the policyholder's propensity to convert their annuity from pay-in to pay-out (annuitize) when the GMIB is "in the money" to trigger the economic benefit of the GMIB.

Summary

A model of policyholder lapse for variable annuities has been synthesized using ideas from (1) *behavioral finance and economics* (e.g., agents are boundedly rational) and (2) *complexity* (e.g., interactions between agents are important). The model has been implemented computationally as a so-called agent-based model. Initial exploration of the model has demonstrated that it is capable of generating lapse rates that are in qualitative agreement with historical annuity experience data.

We anticipate that these models (and extensions thereof to other elements of customer behavior such as policy option election) will be valuable for three reasons:

1. Insofar as the behavior of policy-holders and brokers is calibrated properly for recent market conditions, these models should provide guidance for appropriate product pricing in capital market and / or product design regimes that are unlike those experienced recently, such as a protracted bear market and / or the new GMIBs.

2. ABMs are more robust than existing models for any insured product for which there is experience data on sellers and customers. This is critical when re-pricing and/or assessing the effectiveness of post surrender charge conservation programs geared to influence future customer behavior.

3. They can be deployed stochastically to provide confidence intervals to drive risk analysis and to evaluate the effectiveness of risk management techniques such as reinsurance and capital markets hedge product approaches to risk reduction.

Appendix

Customer Behavior:

Age is uniformly distributed from 55 to 65 in the customer population, and sex is assigned randomly—at this point mortality is not tied to gender so this customer data is not used in any of what follows. It was desired to have the distribution of variable annuity asset allocation in the population near a 70% variable/30% fixed mix initially. This was accomplished by giving each agent an allocation sampled from a uniform distribution (over [50%, 100%] variable, thus 75% is the mean value). The probability that an agent will die in any particular period is 0.01 in the current model, although any hazard function can be employed. The probability that an agent lapses for reasons of liquidity is 0.02 at present. If an agent has not received advice from the seller, it either (1) with probability 0.40 decides on its own if it makes economic sense to lapse, or (2) queries a small sample of other customers and imitates their lapse actions. Overall, this behavior can be summarized in the following 'pseudo-code.'

Pseudo-code of customer behavior

```
Each period each active (non-dead, non-lapsed) customer is activated:
   With some fixed probability it dies;
   With another probability it lapses due to liquidity needs
   If it is neither dead nor lapsed:
      If it has received advice it lapses with probability p, an
         increasing function of the excess returns net of penalty
         suggested by the seller that it can receive from another
         annuity product;
     If it has not received advice:
         It decides on its own whether or not to lapse with a fixed
            probability (0.40);
         It imitates the lapse behavior of the overall population with a
            fixed probability (0.50);
```

Broker Properties

Given the "80/20 rule", i.e., 80% of annuities are sold by 20% of agents, the seller population is configured such that a skewed distribution of the number of customers/agent obtains. To accomplish this, a triangular distribution of customers/agent was used. As it turned out, this distribution is insufficiently skewed (the top 50% of sellers control 75% of the business), but is more realistic than a uniform distribution. (Companies interested in instantiating the model for their own purposes could use their own data to specify this distribution.) The fraction of its customers each seller contacts is uniformly distributed over [0,1]. The seller's recommendations to its customers are based on market scenarios, with a mean value equal to the current overall market rate of return, and a 2 percentage point variance. This means that about 95% of the broker recommendations (±2 standard deviations) are within 4 points of the base return. Overall, the pseudo-code of seller behavior can be succinctly stated as follows:

Pseudo-code of seller behavior

```
Each period each seller selects some fraction of its client base to
      contact:
   It ranks its clients by account value and contacts high value clients
      first;
   It makes a noisy recommendation of returns available in the market to
      lapsing for each contacted client based on current returns
```

References

Aaron, H. 1999. *Behavioral Dimensions of Retirement*. Russell Sage Foundation/Brookings Institution Press: New York, N.Y./Washington, D.C.

Arthur, W.B., S. N. Durlauf, and D. Lane. 1997. *The Economy as an Evolving Complex System II*. Addison-Wesley: Reading, Mass.

Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer, and P. Tayler. 1997. "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market." In Arthur, Durlauf and Lane, eds.

Axelrod, R. 1997. *The Complexity of Cooperation*. Princeton University: Princeton, New Jersey.

Axtell, R.L. 1999. "The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions." Sante Fe Institute working paper #99-03-019.

Axtell, R.L. and J.M. Epstein. 1999. "Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement." In Aaron, ed.

Brooks, R., and P. Maes, eds. 1995. *Artificial Life IV*. MIT Press: Cambridge, Mass.

Camerer, C. 1997. "Progress in Behavioral Game Theory." *Journal of Economic Perspectives, 11*: 167-188.

Casti, J. 1996. *Would-Be Worlds : How Simulation Is Changing the Frontiers of Science.* John Wiley & Sons - New York..

Darley, V.

Epstein, J.M. and R. Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up.* Brookings Institution Press / MIT Press: Washington, D.C. / Cambridge, Mass.

LIMRA International, Inc. 1997. "Annuity Persistence Study." *In TSA 1995-96 Reports.* LIMRA International and the Society of Actuaries.

Lowenstein, G. and R. Thaler. 1989. "Anomalies: Intertemporal Choice." *Journal of Economic Perspectives, 3*: 181-193.

Loewenstein, G., D. Prelec, and R. Weber. 1999. "A Psychological Perspective on Economic Aspects of Retirement." In Aaron, ed.

Nagel, K. and S. Rasmussen. 1995. "Traffic at the Edge of Chaos." In Brooks and Maes, eds.

Rabin, M. 1998. "Psychology and Economics." *Journal of Economic Literature, vol. XXXVI* (March): 11-46.

Thaler, R. 1994. *Quasi-Rational Economics.* Russell Sage Foundation: New York, New York.