



The Perils of Long-Term Guarantees

By Matthew Easley

The insurance industry is in the business of making promises and keeping them. Historically, it has been one of the most reliable in fulfilling its commitments. But the industry has struggled making money at times when it has issued long-term guarantees in a variety of forms. This article focuses on how to look at long-term guarantees in a way that is less likely to create a future crisis.

Many life insurance contracts have the ability to adjust nonguaranteed elements if necessary, or guaranteed rates may last for a limited period, such as 10–20 years. Others have made permanent promises as to price and coverage that are guaranteed. What this means is that the company is along for the ride with limited actions available to mitigate adverse trends. A lot can change in 40–50 years with regard to economics, regulation, society, and customer behavior.

That said, insurers are good at managing many types of risk. Three in particular are relevant to this discussion:

- **Pooling risks.** Insurers know how to pool risks effectively. They can take a large number of independent events that can be disastrous for the individual and put them into a pool where the collective risk to the insurance company is less. They are able to do this thanks to their underwriting expertise and the law of large numbers. This exists in life insurance, auto insurance and even homeowners insurance apart from catastrophe risk. This is a value-added activity that insurers provide to society at large because they can reduce risk for their customers.
- **Serving as an intermediary.** A second thing that insurers do well is to act as an intermediary to create products that are available in theory, but not in practice, to individuals. One good example is the fixed indexed annuity (FIA). Most



consumers do not have the ability to construct a complex derivatives strategy for themselves. The FIA contract accomplishes this in a consumer-friendly manner and absorbs some of the frictional risk to allow greater simplicity and more flexibility to policy owners. A second example is the early variable annuity (VA) contracts, which bundle mutual fund investing into a single policy, frequently with easy access to multiple mutual fund groups. Insurance companies take some expense and mortality risk but only limited market risk in creating a new service platform for consumers.

- **Leveraging group benefits.** A third area where insurers excel has been the creation of health networks to negotiate better rates for people in their groups. This has many forms, but the power of representing the group in getting better prices and other network benefits has been a major result. The same type of power exists in group underwriting, where normal underwriting standards can be relaxed when providing coverage to most of the members of a group. This is prevalent in payroll deduction programs. Similarly, pension funds can provide book value benefit payments to individual participants. These programs offer participants deals that individuals could not obtain on their own.

There is a recurring theme in these areas of expertise and that is *value creation*. This is a fundamental element of a successful product. The manufacturer must create a product that is valued by the buyer at a price that exceeds its cost of production plus a margin to recover fixed expense and produce a profit margin. Much of what is called pricing within insurance companies is really the determination of the cost of manufacturing a policy, including a cost of capital. A pricing actuary can stay within that paradigm, but the product development actuary must go beyond and seek to understand the value of the product in the hands of the consumer. This requires working with marketing concepts that go beyond traditional actuarial valuation methods.

These activities—pooling, intermediation and leveraging group benefits—attack value creation from the top; that is, insurers are increasing the benefit to the consumer. They create value that is not easily accessed by consumers without the support of an insurance contract. Other efforts can lower the cost of production in various ways. Many automation programs, for example, are efforts to lower the cost of production without increasing fixed cost recovery by a larger amount.

For long-term guarantees, the question becomes, what is the actual cost of manufacturing? Can the company produce the guarantee for less than what it is able to charge the customer? In many cases, the insurance industry has not been able to get this equation right. And the results frequently vary by customer segment and distribution channel due to behavioral differences.

CASE STUDY: VARIABLE ANNUITY

One of the clear examples of selling below the cost of manufacturing happened in the variable annuity industry. During the 1990s, variable annuity products started to incorporate strong guarantees on top of the basic mutual fund format. These guarantees blended products developed to allow flexible investing with promises that behaved like complex derivatives. Although each company had its own experience, many companies underpriced this risk materially. In one particular case, the original price and the updated price differed by a factor of three! It could have been even worse if consumer use of options had been more efficient. The process of getting to adequate reserves and required surplus was expensive in time, money and capital.

The scary part about this story is that it is about smart people who made an expensive series of mistakes. Although it is impossible to know everything that happened, some of the key lessons include the following:

- The individual annuity actuaries were not experts at derivatives pricing. They used the methods that they knew and understood to assess a new type of guarantee structure. Those methods failed to reflect the market value of the guarantees. The lack of relevant experience data made this problem worse.

VALUE-BASED PRICING

The use of value-based pricing is an important departure from traditional pricing methods that are based on discounted cash flows and use of capital. Understanding the cost of producing a product or feature remains important as a comparison to the value. It requires new skills to estimate the value of a product feature from a consumer perspective. The spread between those amounts is the value added.

One particular method that helps define the value of particular product features is conjoint analysis. This is particularly useful to look at novel designs compared to traditional designs. It allows companies to look at the value of the traditional feature versus the proposed alternatives from a consumer point of view. Marriott made use of this method in developing the Courtyard brand, but it has been used by many companies that are focused on selling products to consumers.

In developing a valuation tool, consumers are asked to compare different features and indicate which they would value more. Included in these comparisons are different price levels. Statistical methods are then applied to determine the relative value of a feature and even the likely popularity of a given product design.

- Product designs frequently were modified incrementally, not created as new products. It is expensive to allow full investment flexibility with guaranteed benefits. Previously benign features were combined with the new guarantees, which created unanticipated problems. Product development teams often operated under time pressure that made it difficult to consider all of the implications of the new features.
- When products are successful in market, pulling back attractive features is even harder than keeping them out of the initial design. Product updates are expensive, and new products that will sell are likely to get on the calendar ahead of risk management updates except in crisis situations.
- The full impact of regulatory reaction to new features does not happen until later in the product cycle, well after the completion of the design and pricing work. This makes it difficult to anticipate these costs in the product design and pricing.

RISKS

One of the most important jobs of an insurer is to manage risk. Several types of risk are of particular concern when looking at long-term guarantees. These are risks that insurance companies are not well equipped to do in volume.

Systemic Risk

Systemic risk can create difficulties because it does not benefit from diversification (except against other uncorrelated risks). For example, when interest rates go higher or lower, it happens at the same time to all policies. Apart from hedging activity, the company can be impacted by the same risk on all policies at the same time. Floor guarantees in various products have created reduced profits for companies during the current low-interest-rate environment. In the past, companies suffered from the opposite risk. Movement to higher than expected rates created losses and cash flow problems for companies, as parallel issues appeared in multiple products.

One of the most valuable things I did at a prior company was to insist that my employer lower the floor guarantees on new individual annuity sales from 3.5 percent to 3 percent (the minimum permitted at the time). Did I know that rates were going to go this low? Not a chance, but I considered the possibility and determined that the company was not being paid for the extra tail risk. My team studied a number of other adverse scenarios as well. The benefits from this one change more than compensated for all the work we did.

The pain of persistent low interest rates has been felt in many countries. Japan entered the low-rate era sooner than most and its experience has become a new stress scenario for many chief risk officers. Germany has seen margins from prior investments dwindle with minimum crediting rates higher than supportable levels for new investments. Like the U.S., the regulatory scheme helped create the problem through policy requirements that did not allow companies to anticipate a shift to low rates. Today, floor rates in the U.S. can be indexed to rates as low as 1 percent, but that change came after more than 10 years of significant new business exposure to low rates.

A second example of systemic risk is mortality improvement. Although the intensity of a given cause of improvement may impact various ages and socioeconomic groups differently, the direction of movement is likely to be consistent for all insured lives and annuitants. (This does not imply that mortality in total will move the same direction for all segments, as evidenced by the recent statistics on the opioid epidemic; rather, a given cause of death is likely to impact all groups in the same direction.) If guaranteed for the life of the insured, this is a long-term risk with limited diversification and hedging alternatives. For payout annuities, this is a category of risk that is only partially insured today with much of the risk remaining in corporate pension plans.

A third example of systemic risk that has been played out several times is concentration in a particular asset class. If that segment of the market goes bad for a time, it can have serious implications. For instance, one company's investment strategy helped take them down when the junk bond market collapsed. Later

regulatory changes reduced the incentive to take this particular type of risk in the U.S., but not in time to save the company.

Part of the risk for these companies was credit, but liquidity was also a major part of the problem. It doesn't matter if you are right in the long run if you die before you get there. This company's below-investment-grade bonds were ultimately profitable, but with concentrations exceeding 30 percent in their operating companies, they were not able to get past the crisis.

Many of the Guaranteed Investment Contract (GIC) providers ran into a similar problem in the late 1980s when they matched the asset/liability management (ALM) risk for their GIC portfolio with mortgages that had guarantees with similar durations. However, they had to pay off the GICs in cash with no matching liquidity feature in the mortgages. The cost of sourcing cash through additional GIC sales put a large dent in their financials.

Assumption Setting With Limited Experience

Another important risk is setting assumptions with limited experience data. Guessing is part of developing a new product, especially a product that is truly new. However, the need to guess becomes dangerous when that estimate gets embedded into long-term guarantees. Finding ways to inform the new assumption from parallel situations can require some creativity. In addition to looking at similar insurance products, some people have looked outside insurance at other financial products for indications of consumer behavior. One of the more interesting is the use of mortgage refinancing experience to learn about consumer sensitivity to interest rates.

One good example of this risk is setting lapse-rate assumptions on lapse-supported products. With these products, the company expects to realize gains from customers who pay premiums for a time, and then lapse the policy with little or no remaining value. The problem is that these customers have proven to be more savvy shoppers than expected. When they buy a policy with limited surrender value, they appear to make a quality decision about retaining the policy long term. In multiple policy types, including universal life with secondary guarantees, term to 100 (in Canada) and long-term care policies, companies have seen lapse rates of less than 1 percent occur. Original expectations for higher lapse were based on data from prior policies that had higher surrender values. This remained true even when customers were hit with sizable rate increases on long-term care policies.

When dealing with the risk of assumption setting with limited experience data, consider the following actions:

- For new products, the ability to change course during the life of the policy is extremely valuable because of the uncertainty of the initial pricing assumptions. Being wrong is almost certain and should be anticipated.

- Monitoring the early results closely and reacting quickly to deviations can save much of the cost of being wrong. Communicating the issue early and often can help.
- The first reaction to rapid sales growth should be to review the pricing, not have a party.

Policyholder Elections

Policyholder elections represent a different type of risk. For many products, the policyholder has options that can be exercised well into the life of the product. During the design process, it can be challenging to estimate the likelihood of various elections, especially under stress conditions. Utilization rates on VA policies for income guarantees are a good example of this risk. Policy loan utilization within life policies, book value surrenders on fixed annuities and long-term care utilization rates are additional examples of variables that are hard to predict during the pricing process.

Long-term care policies have come under a lot of criticism, but one thing insurers got right was the need to be able to adjust future premiums. Companies have used premiums to adjust for lower interest rates (that they failed to hedge) and low lapse rates as discussed earlier. Although this resulted in large premium increases that have been painful for everyone, the premium option has kept most of these companies solvent in spite of significant losses on key assumptions.

Avoiding Problems

So, how can the product development actuary avoid a problem with mixing long-term guarantees with uncertain assumptions? Here are some questions to ask yourself to test a new product design:

- What happens if I miss the mark on my assumptions? Can I offer a meaningful temporary guarantee instead of a permanent one to allow future adjustments?
- Am I making promises that I cannot hedge? What about the design is making it hard to hedge?
- Am I giving guarantees away where I am not being paid? Or not paid enough?
- What are the scenarios that would cause large losses with this design? (Don't worry about the probability; just think about the level of potential for loss.)
- What would I do if one of these scenarios happened? And what actions would I want to be able to take? How can I retain the right to take those actions?
- And finally, how can I avoid being the next example in a paper 10 years from now?



IDENTIFYING CRITICAL RISKS IN A NEW PRODUCT

What are the major risks to consider in the product under discussion? One way to think about this is to identify what is changing from prior products. What assumptions are you borrowing from experience, and how might they change? What new features are you including? And how sensitive are your projected financial results to variations in different factors?

At the root, the important thing is what can break the product. It may not be clear what the odds are of a given event, but it is important to test the potential of a scenario to create significant losses so you can decide how much effort to focus on that area. If the assumption or feature can create only moderate losses (or reductions in profit), it is unlikely to be one of your key concerns. However, if an unlikely scenario can create large losses, it should stay on your list.

Again, answering several exploratory questions may help you:

- How can the product be modified to reduce the size of the loss in adverse scenarios? Can I design a “circuit breaker” that empowers future management to mitigate losses? How much will that feature impact sales of the product? Will it be more acceptable if I include it as an option for a better price?
- What elements of the guarantee structure really matter to my customers? How can I eliminate risk without impacting those aspects of the product? How can I maintain price stability for them without making an absolute guarantee?
- What distribution channels will be willing to accept a non-standard feature? What features will be regarded as an acceptable variation rather than being considered nonstandard? Will offering the standard version at a higher price help or hurt my efforts in that channel?

LOW-PROBABILITY EVENTS

The Black Swan research indicates that it is difficult to know both the probability and size of tail events. However, people tend to overestimate how far into the tail an event is. Take a minute to look at the past 100 years and think about how many of them you consider unusual. You might think about the following events:

- The Great Depression
- World War II
- The Korean War
- The Vietnam War
- Going off the gold standard
- 1970s oil crisis
- 9/11 and the Iraq War
- The Great Recession

How many years did you consider unusual? You are probably talking about 20–30 years out of 100, which must be considered normal deviations.

Another way to look at this is to think about what constitutes a 1-in-200 event. Thinking longer term, this is the same as considering the 10 most pivotal events in the past 2,000 years. Such a list might include the discovery of the New World, the fall of the Roman Empire, the founding of two major religions, the industrial and digital revolutions, the Black Death, the end of Chinese dynastic rule, the invention of the printing press, and the rise of democracy as a major form of government. How does your model include changes of this magnitude?

Real adversity requires more effort to imagine than looking at the historical record. It is necessary to prepare for what might come, not just for what caught the industry by surprise the last time. In setting capital standards, it is common to consider a 1-in-200 standard for capital. But what limitations attach to such a standard?

- First, the output is only as good as the input, and the input is frequently rounded to a few significant digits. The results are only accurate to the same number of digits.
- Given an economic record of just 100 years or so in the modern era, the amount of experience on many economic risks is insufficient to extrapolate the tail well.
- Pricing models contain only those risks that are understood and included. Many tail risks are excluded from the process completely. For example, data from countries where there has been regime change or other major economic adjust-

ments are frequently excluded. However, the experience being excluded may be the tail risk data you really need.

CREATIVE DESIGN

One of the fundamental decisions for a new product is whether to follow a standard product design or to branch out and design something new. The standard product has a better chance of acceptance by producers. Producer resistance is primarily related to retraining efforts. This can be personal or across the system, but the work to bring a new product into the channel is substantial. In contrast, a product that is a virtual copy of other products in market is easier to launch, but will have to compete on price. That can be through a higher interest rate, a higher commission, more generous underwriting or other factors, but finding a place on a crowded shelf may require lower than target profits.

Another way to compete is to add a new guarantee or feature. The cost of a novel feature may be hard to estimate accurately, especially without experience on utilization. But a stronger guarantee or a new feature can attract new business on a spreadsheet comparison basis. Note that this process is frequently producer driven and not focused on customer value.

To escape this trap, it is necessary to get under the surface of the market and learn what customers really value. The complicated products that companies offer frequently charge for guarantees that far exceed what the customer would want to pay given the option. This makes it difficult to sell these products at the target returns if fully priced. Identifying the features that customers really want and focusing the value on those things is an important step toward improving the value equation.

Let's take a simple example. How much would a customer want to pay for a 100 percent guarantee of premium for life versus a 20-year guarantee? This type of trade can be measured quantitatively using marketing analytics, as noted earlier. If you learn that the customer is willing to pay 10 percent and your cost is 20 percent, this is a feature that you should want to eliminate if possible because it is destroying value.

However, there is a threshold issue of a distribution system that is used to selling the lifetime guarantee. Does the feature destroy enough value to be worth the battle? Can you include a substitute guarantee that is acceptable while destroying less value? In this example, the maximum premium after 20 years might be capped at 110 percent of the initial premium. This design introduces risk sharing with the customer that will reduce potential losses in many adverse scenarios and may improve required reserves.

Another example is found in adjustable rate mortgages, through which people accept a period of guarantee (5–10 years commonly) followed by a maximum increase that is indexed and commonly limited to 2 percent per year and 5 percent overall. This

does not appeal to everyone, but it frequently offers lower rates that better fit many budgets. The shorter duration also tends to match the investment needs of most banks.

Both of these examples illustrate the value of controlled risk sharing. A core guarantee is offered and the customer is asked to take some risk. Ideally, the timing of that risk sharing should be when it is expensive for the company and less valuable to the customer. This tends to be true when dealing with benefits in the future. People commonly value current costs that are certain more than distant costs that are uncertain. Conversely, the cost to guarantee frequently rises over time for the company.

A new product does not need to be better for every customer. A majority of people may prefer the standard product and many producers may continue to sell it. To succeed, a new product needs only to carve out a segment of the market. In fact, it may be better if the segment is narrow enough to discourage copycat products that will impact your pricing power. I recall a new product that was predicted to appeal to only 30 percent of consumers based on conjoint modeling. The marketing firm that built the tool was initially disappointed. Then they looked at the results for the market standard product and discovered that it appealed to less than 20 percent of consumers. The team was encouraged to move ahead with the new design.

GETTING STARTED

How can a product development actuary avoid the perils of long-term guarantees?

- Consider the product category where it is hardest to compete, then answer the following questions:

- What can I do to enhance my competitiveness that does not involve lowering price, raising commissions or taking more risk?
- Is there a different way to reach this market that serves customers well without making the risk-return relationship unattractive for my company?
- How can I learn more about what my customers really want?
- How is my current distribution channel impacting my results? And where do I have the best opportunity to introduce something new?
- Will a new product benefit my producer relationships and help with standard product sales? How much of the standard product do I really want?
- Consider where you have special skills and strengths and think how to use them.
- Consider what you could do if you added one or two new skills.
- Consider who might be a good partner to help provide critical skills.

Product development is a most interesting puzzle with its blends of math, intuition and creativity. Best of all, it can be a lot of fun for those involved and create enormous value for the company. ■



Matthew Easley, FSA, MAAA, is a senior vice president at RGA. He can be reached at measley@rgare.com.