EDUCATION AND EXAMINATION COMMITTEE

OF THE

SOCIETY OF ACTUARIES

INVESTMENT AND FINANCIAL MARKETS STUDY NOTE

ACTUARIAL APPLICATIONS OF OPTIONS AND OTHER FINANCIAL DERIVATIVES

by

Michael A. Bean, FSA, CERA, FCIA, FCAS, PhD

IFM-22-18 **Printed in U.S.A.**

Actuarial Applications of Options and Other Financial Derivatives

Michael A. Bean, FSA, CERA, FCIA, FCAS, PhD

Version: January 21, 2018

Contents

1 Introduction

Options and other financial derivatives are most often associated with investments and financial markets. However, they also arise in many actuarial contexts. This study note provides an overview of the ways in which options and other financial derivatives arise in actuarial applications.

Outside their use as tools of investment management, derivatives arise in two key ways:

- As options or guarantees embedded within retail insurance, savings, lending, or wealth management products; and
- As tools that the providers of these products use to manage the risks associated with such options or guarantees.

Actuaries need to recognize when a retail insurance, savings, lending, or wealth management product contains an embedded option or guarantee so that the product can be appropriately priced and the liabilities associated with it appropriately valued. It is also important for actuaries to have an appreciation for how derivatives can be used to manage the associated risks so that products with embedded options or guarantees can be designed in a way that facilitates risk management and ensures that the costs associated with managing risk are appropriately covered. More than a few insurance companies have experienced unexpected losses on insurance or wealth management products because they failed to recognize the presence of an embedded option or guarantee, appreciate its value, and/or appropriately manage the risk.

Options or guarantees embedded in retail insurance, savings, lending, or wealth management products can be complex and difficult to value, much more so than the plain vanilla options that are typically traded in financial markets. Although the Black-Scholes methodology from financial economics is often a useful starting point for valuing these options or guarantees, it does not always provide an answer that is meaningful, e.g., if one or more of the underlying assumptions, such as the ability to construct a replicating portfolio with a market-observable price, does not hold. Hence, one must be careful when applying the techniques of financial economics to actuarial problems.

In this study note, we describe how options and other derivatives arise in popular insurance and wealth management products such as variable annuities and mortgage guaranty insurance, and consider how insurance companies and other financial providers use derivatives that are traded in financial markets or available over-the-counter to manage the risks in these products. As this study note is only intended to provide an overview of the actuarial applications of derivatives, we do not consider valuation of the options and guarantees embedded in these products in detail. In particular, there is no expectation that readers of this note will have acquired the set of formulas and tools required to do a proper valuation.

2 Options Embedded in Insurance Products

We begin by considering some of the options that are embedded in variable annuities and mortgage guaranty insurance. We also briefly consider some other insurance products with embedded options.

2.1 Guarantees in Variable Annuity Products

A **variable annuity** is a savings or investment product sold by life insurance companies that is similar to a mutual fund but with one or more guarantees on investment performance. Variable annuities are typically purchased as a means of saving for and/or funding retirement.

Like a traditional annuity, a variable annuity has both an accumulation period and a payout period. In the accumulation period, the policyholder makes one or more deposits and the deposited funds grow over time. In the payout period, the policyholder or some other designated person known as the annuitant[1] receives a stream of payments, typically at regular intervals, for a defined period of time, e.g., 20 years or the remaining life of the annuitant. However, unlike a traditional annuity, the size of the payments received is not known in advance and can vary during the payout period depending on the performance of the underlying investments. Variable annuities also generally distribute any funds that remain at the end of the payout period to the policyholder, the annuitant or their respective beneficiaries.

As interest rates declined from historic highs in the early 1980s to historic lows in the 2010s, traditional annuities became more and more expensive, i.e., the principal amount needed to produce payments of a given size became larger and larger. For many people, the amount of money needed to generate their desired level of retirement income was more than they had saved or expected to save. Variable annuities were marketed as an alternative to traditional annuities that enabled people to take advantage of the potentially higher rates of return associated with equity and other financial market investments but still be guaranteed some minimum rate of return.

2.1.1 Types of guarantees

There are four basic types of guarantees associated with variable annuity products:

- Guaranteed minimum death benefit (GMDB);
- Guaranteed minimum accumulation benefit (GMAB), sometimes referred to as a guaranteed minimum maturity benefit (GMMB);
- Guaranteed minimum withdrawal benefit (GMWB); and
- Guaranteed minimum income benefit (GMIB).

To qualify as a life insurance product and be eligible for the tax and statutory benefits typically associated with life insurance products, variable annuity products generally have some sort of a guaranteed minimum death benefit. However, they usually have one or more of the other guarantee benefits as well, either as part of the base product offering or as an add-on that can be purchased at the discretion of the policyholder, albeit at additional cost.

2.1.1.1 Guaranteed minimum death benefit

A **guaranteed minimum death benefit** provides a guarantee on the amount that the beneficiary of the contract receives when the policyholder dies. For example, a variable annuity contract

[1] While the policyholder and the annuitant may be different people, in this note the term policyholder will be used to refer to either.

might guarantee that when the policyholder dies, the beneficiary receives the greater of the account value and the original amount invested. Such a guarantee is known as a **return of premium** guarantee.

2.1.1.2 Guaranteed minimum accumulation benefit

A **guaranteed minimum accumulation benefit** provides a guarantee on the value of the underlying account after a specified period of time has elapsed provided that the policyholder is still alive and the contract still in force at that time. For example, a variable annuity contract might guarantee that 10 years after contract inception, the value of the underlying account will be at least 110% of the original amount invested provided that the contract is still in force and there have been no withdrawals during that time. If the account value at the end of the specified period is less than the guaranteed value then the insurer deposits the difference to the policyholder's account. Depending on the terms of the contract and the age of the policyholder at that time, another guarantee period for the accumulation benefit may begin, funds may continue to accumulate but without a guarantee, or the contract may simply mature.

2.1.1.3 Guaranteed minimum withdrawal benefit

A **guaranteed minimum withdrawal benefit** provides a guarantee on the size of withdrawals that the policyholder can make from the underlying account during the payout period and the length of time these withdrawals can be made provided that withdrawals do not commence until the policyholder reaches a specified age. For example, a variable annuity contract might guarantee that the policyholder can withdraw 5% of the original investment every year for the remainder of the policyholder's life provided that withdrawals do not begin until the policyholder reaches age 65 and do not exceed the specified size. If the account is depleted before the policyholder dies then the insurer continues to make the guaranteed minimum payments to the policyholder until the policyholder death; if it is not depleted before the policyholder dies then the beneficiary of the contract receives the remaining account balance.

2.1.1.4 Guaranteed minimum income benefit

A **guaranteed minimum income benefit** provides a guarantee on the future purchase rate for a traditional annuity. For example, a variable annuity contract might guarantee that when the policyholder reaches age 65, the policyholder will be able to purchase a traditional annuity that pays $750 per month per $100,000 of account value at age 65 for the remainder of the policyholder's life. If at the time the policyholder turns age 65, the going rate for a traditional annuity is less than $750 per month then the insurer is obliged to provide the policyholder a traditional annuity for the guaranteed $750 per month rate; however, if the going rate is greater than $750 per month then the policyholder simply purchases a traditional annuity at the going rate.

2.1.2 Guaranteed minimum death benefit with a return of premium guarantee

A guaranteed minimum death benefit with a return of premium guarantee is similar to a European put option. The only difference is that the time the guarantee/option comes due depends on when the policyholder dies, which is uncertain.

To see this, let S_t denote the value of the policyholder's account at time t, let K denote the amount invested, i.e., $K = S_0$ and let T denote the future lifetime of the policyholder. Then, at the time the policyholder or annuitant dies, the beneficiary receives the amount

$$\max(S_T, K) = S_T + \max(K - S_T, 0).$$

Note that the expression $\max(K - S_T, 0)$ is identical to the payoff on a put option with strike price K and time to expiration T. Hence, if the future lifetime of the policyholder were known with certainty then the return of premium guarantee would be a European put option.

However, the future lifetime of the policyholder is not known with certainty; it is a random variable. Hence the same is true of the time to expiration in this case. Consequently, a guaranteed minimum death benefit with return of premium guarantee is not a European put option per se; rather it is a put option with expiration contingent on the death of the policyholder, i.e., it is a *life-contingent* put option.

Using these observations, one can determine a value for the embedded guarantee. To distinguish between the future lifetime of the policyholder and the time to expiration of a generic European put option, and to highlight the dependence of the future lifetime of the policyholder on the current age of the policyholder, henceforth let T_x denote the future lifetime of a policyholder with current age x and let f_{T_x} denote the probability density function for T_x.[2] Further, let $P(T)$ denote the value of a European put option on the given investment with strike price K and time to expiration T. Then the value of the embedded death benefit guarantee is[3]

$$\int_0^\infty P(t) f_{T_x}(t)dt.$$

That is, it is the probability-weighted average of the values of European put options of varying times to expiration with probability weights defined by the random variable T_x. Assuming that all of the conditions of the Black-Scholes model hold, $P(t)$ in this expression is given by the Black-Scholes formula for European put options and, using numerical techniques if necessary to calculate the resulting integral, a value for the embedded death benefit guarantee can be determined.

To determine what this probability-weighted average represents, suppose that the insurer sells a large number of GMDB contracts with return of premium guarantee to policyholders of the same age and that each policyholder invests the same amount of money in exactly the same way. Then the only difference between contracts is the time when the policyholder or annuitant dies, i.e., the time to expiration of the guarantee.

---

[2] For readers familiar with actuarial notation, $f_{T_x}(t) = {}_t p_x \mu_{x+t}$.

[3] This assumes that policies do not lapse. If lapses are possible then the formula for determining the value of the death benefit guarantee is more complex.

If the total number of contracts sold is sufficiently large then, from the law of large numbers of probability theory, the number of contracts that expire at any given time becomes more certain and the portfolio of contracts becomes like a portfolio of European put options. Hence, the probability-weighted average value can be viewed as the average cost of the guarantee per contract for a sufficiently large portfolio of GMDB contracts with identical characteristics. Note, however, that this average value does not represent the cost of the guarantee for any given contract, as that depends on the time when the policyholder associated with the particular contract dies, which is not known in advance.

2.1.3 Earnings-enhanced death benefit

An **earnings-enhanced death benefit** is an optional benefit available with some variable annuity products that pays the beneficiary an additional amount when the policyholder/annuitant dies based on the increase in the account value over the original amount invested. For example, a contract with an earnings-enhanced death benefit might pay the beneficiary an additional benefit equal to 35% of any increase in account value. This additional amount, which is only payable if the account value at the time of the policyholder's death is greater than the original amount invested, can be used to help offset taxes that may become due when the death benefit is paid and the contract is fulfilled.

Assume that an earnings-enhanced death benefit equal to 35% of any increase in account value is added to the return-of-premium death benefit guarantee just discussed. Then, using the earlier notation, the amount paid to the beneficiary at the time of the policyholder's death is

$$S_{T_x} + \max(K - S_{T_x}, 0) + 35\% \times \max(S_{T_x} - K, 0).$$

The latter term in this expression represents the payment from the earnings-enhanced death benefit. Apart from the multiplier, this term is identical to the payoff on a call option with strike price K. Letting $C(T)$ denote the value of a European call option with strike price K and time to expiration T, it follows that the value of the earnings-enhanced death benefit is[4]

$$35\% \times \int_0^\infty C(t) f_{T_x}(t) dt.$$

2.1.4 Guaranteed minimum accumulation benefit with a return of premium guarantee

Like a GMDB with a return of premium guarantee, a guaranteed minimum accumulation benefit with a return of premium guarantee has characteristics similar to a European put option. However, in this case, payment is contingent on the policyholder surviving to the guarantee/option expiration date and the policy still being in force at that time.

Let T_x^* denote the future lifetime of the policy and suppose that there is a single guarantee period, which ends *m* years from now. Note that since lapses are assumed to be possible in this case, the future lifetime of the policy can be less than the future lifetime of the policyholder, so the distribution of T_x^* will generally be different from the distribution of T_x in the previous

[4] As before, this assumes that policies do not lapse. This expression only has meaning for a large portfolio of otherwise identical contracts issued to policyholders of age *x*.

examples. Let $P(T)$ denote the value of a European put option with time to expiration T. Then, the value of the embedded accumulation benefit guarantee is

$$\Pr(T_x^* \geq m) \times P(m)$$

where $\Pr(T_x^* \geq m)$ is the probability that the policyholder is alive and the policy is still in force m years from now. Note that if there is more than one guarantee period then the formula for determining the value of the GMAB is more complex.

2.1.5 Guarantee value formulas

The examples just considered suggest that the guarantee value in a variable annuity contract plays the role of an option strike price, at least in the case of death benefit and accumulation guarantees. In basic contracts, the guarantee value is typically defined by a single number such as the original investment amount and remains fixed for the life of the contract provided that there are no withdrawals. However, in more sophisticated contracts the guarantee value is defined by a formula that is periodically recalculated and takes into consideration factors such as the performance of the underlying investments since contract inception, any promised rate of return, and the age of the policyholder.

For example, the guarantee value for a contract could be recalculated every policy anniversary prior to the policyholder/annuitant's 70th birthday according to the following formula: The guarantee value is the maximum of the account value on the anniversary date, the previous guarantee value, and the amount determined by accumulating the initial investment at a rate of 2% per annum, where the guarantee value is initially assumed to be the original amount invested.

Some contracts allow the policyholder to choose when the guarantee value is recalculated; others limit recalculation of the guarantee value to specified dates such as policy anniversaries. Some contracts use the same guarantee value formula for all benefit types; others use different formulas for different benefits. However, in virtually all contracts, withdrawals trigger an adjustment in the guarantee value, either proportionately or dollar-for-dollar, at least in the case of death benefit and accumulation guarantees. A proportionate adjustment preserves the relationship between the account value and the guarantee value whereas a dollar-for-dollar adjustment need not.

Contracts that allow the policyholder to choose when the guarantee value is recalculated generally have some restriction on how frequently this can occur. For example, a contract may limit policyholder-initiated recalculations of the guarantee value to once per calendar year and require that the policyholder be less than 80 years of age at the time the guarantee value is recalculated. A policyholder-initiated recalculation of the guarantee value also generally results in an adjustment to the remaining term of the guarantee. For example, if a contract has a return-of-premium guarantee over 10 years and the policyholder initiates a recalculation of the guarantee value then a new 10-year guarantee term typically begins.

Note that the ability to choose when the guarantee value is recalculated is itself an option. Moreover, this option is an American-style option since the policyholder can choose when it is exercised. As we will soon see, variable annuities can have other types of embedded policyholder options as well.

2.1.6 Guarantee value for a guaranteed minimum withdrawal benefit

The concept of guarantee value is more complicated for a guaranteed minimum withdrawal benefit than it is for a death benefit or an accumulation benefit because with withdrawal benefits, it is the size of the payments, i.e., withdrawals, during the payout period that is guaranteed rather than the size of the account value. Instead of defining a guarantee value directly, these contracts typically introduce a quantity known as the **guaranteed withdrawal base** and define the size of the periodic payments in terms of this base quantity.

For example, the size of a guaranteed annual payment may be defined as 5% of the most recently calculated value of the guaranteed withdrawal base where the guaranteed withdrawal base is calculated as follows: The guaranteed withdrawal base is initially set equal to the amount invested and is recalculated every third policy anniversary as the greater of the account value and the previous value of the guaranteed withdrawal base.

The guaranteed withdrawal base is generally adjusted downward for withdrawals that are larger than the guaranteed size or occur before the policyholder reaches the age required for guaranteed payments to begin. However, it can also be adjusted upward if the policyholder is eligible to make a withdrawal in a given year but elects not to do so. The adjustment particulars depend on the terms of the contract.

Although the guaranteed withdrawal base may appear to be similar to a guarantee value, it is not a guarantee value, nor is it analogous to an option strike price. Indeed, if the value of the guaranteed withdrawal base is greater than the account value, this does not imply that the account value will be increased by the difference at any future time or that the derivative corresponding to the withdrawal benefit guarantee is even in-the-money. The guaranteed withdrawal base is simply a tool for calculating the size of the guaranteed periodic payments.

A more promising candidate for the guarantee value of a GMWB is the actuarial present value of the stream of guaranteed payments, i.e., the present value of the guaranteed payments with discounting for both interest and survivorship. However, other notions of guarantee value may be possible as well.

2.1.7 Underlying on which the guarantee applies

The underlying asset for an option on a stock is the stock itself. The underlying asset for a guarantee associated with a variable annuity product is generally the variable annuity's account value. This raises some interesting challenges when determining the value of a variable annuity guarantee and its cost.

The account associated with a variable annuity typically consists of investments in one or more mutual funds that are selected by the policyholder from a list provided by the variable annuity writer. This list usually includes a large-cap U.S. equity fund that tracks an index such as the S&P 500, a small-cap U.S. equity fund that tracks an index such as the Russell 2000, an international equity fund that tracks an index such as the Europe, Australia, and Far East (EAFE) index, and an investment-grade bond fund. It may also include balanced funds with different mixes of equities and bonds.

The mix of funds in a given policyholder's account depends on the policyholder's investment preferences and tolerance for risk as well as any restrictions that the variable annuity writer may have in place. For example, the variable annuity writer may require each policyholder's account to contain a particular percentage of lower-risk bond funds and may limit the amount that can be allocated to higher-risk funds such as those that track the NASDAQ 100 index. The mix, at a given point in time, also depends on whether the policyholder periodically rebalances the account, i.e., moves money from the better-performing funds to the poorer-performing ones in order to return the mix to what it was at contract inception, or allows it to grow without intervention. In the latter case, allocations to higher-return (likely higher risk) funds will naturally increase and those to lower-return (likely lower risk) funds will decrease.

Consequently, the underlying asset for a variable annuity guarantee is a blend of underlying investments. Moreover, this blend is generally different from one policyholder to another, even for otherwise identical variable annuity products, and can drift over time toward a riskier mix of investments. This makes determining the value of a variable annuity guarantee much more challenging than determining the value of a plain-vanilla option on a stock or stock index.

2.1.8 Option to change the investment fund mix

A further complication is that variable annuity contracts typically allow policyholders to change the mix of investments in their account after the contract is written as long as any general restrictions, e.g., a maximum percentage allocated to equities, are respected. Moreover, there are very few restrictions on the number of times a policyholder can change the fund mix. The effect of this is to give the policyholder an option to change the risk characteristics of the underlying asset of the guarantee after the contract is written.[5]

To consider how this might complicate things for a variable annuity writer, consider a GMAB with one year to maturity and an account-to-guarantee-value ratio of 100%, i.e., the embedded option is at-the-money. Suppose, for simplicity, that there are just two fund choices, an equity fund that tracks the S&P 500 index and a short-term bond fund that invests in U.S. Treasury notes, and that the one-year return on the short-term bond fund is 0%. Then since the policyholder is assured that the account value one year from now will be no less than the current account value, the rational thing for the policyholder is to transfer any amounts in the bond fund to the equity fund, i.e., have 100% of the funds in the account allocated to equities, or if there is restriction on the amount that can be allocated to equities, have the maximum allowable percentage allocated to the equity fund. However, this is precisely what the variable annuity writer does not want the policyholder to do, particularly if the premiums charged for the guarantee assumed a lower allocation to the equity fund and there is no mechanism for recouping the lost premiums.

[5] Hence, like policyholder-initiated recalculations of the guarantee value, fund switching is an American-style option embedded within an option.

2.1.9 Contract termination as an option

Another policyholder option that is embedded in variable annuity contracts is the option to withdraw all the funds in the underlying account and terminate the contract prior to maturity. This option has value because the guarantee premiums associated with a variable annuity contract are typically paid over time, e.g., as a percentage of the average account balance, the guarantee value, or something similar, rather than up-front as generally happens with options that are traded in financial markets. Termination of the contract frees the policyholder from the obligation to pay any further guarantee premiums.[6]

Consider, for example, a GMAB with a simple return of premium guarantee. On one hand, if the account value increases to such an extent that there is virtually no chance of the account value being less than the guarantee value at contract maturity, i.e., the guarantee is essentially worthless, then the policyholder can withdraw all the funds in the account and pay no further guarantee premiums. On the other hand, if the account value falls below the initial amount deposited, and there is a good chance that it will be below the guarantee value at maturity then the policyholder can keep the funds in the account, continue to pay the guarantee premiums, and enjoy the full benefit of the guarantee at contract maturity.

Consequently, the option to withdraw all the funds in the underlying account and terminate the contract prior to maturity is essentially an option to stop paying guarantee premiums when it is no longer advantageous to do so. This option is closely related to the option to recalculate the guarantee value, which was discussed earlier.

Indeed, if the account value on an existing contract is greater than the guarantee value then terminating the contract and repurchasing a new one with identical features has the same effect as setting the guarantee value equal to the current account value and beginning a new guarantee term. Hence, a policyholder-initiated recalculation of the guarantee value can be synthetically created, even on a contract that does not allow policyholders to choose when this occurs, by simply terminating and repurchasing contracts of the same type. Of course, this strategy only works if contracts of the same type continue to be offered over time. So, one way companies can limit the synthetic recalculation of guarantee values is to periodically change their product offerings.

2.1.10 Key takeaways

It should be clear from this discussion that the options embedded in even the most basic variable annuity products can be highly complex and difficult to value. Needless to say, it is important for actuaries working with these products to be able to recognize when options are present and have an appreciation for their value. It is also important to have an understanding of what it costs to manage the risk. Later in the study note, we will consider some of the ways in which the risks associated with variable annuity guarantees can be managed using financial market securities.

---

[6] Depending on the terms of the contract and the amount of time that has elapsed since contract inception, the policyholder may be charged a surrender fee. However, this fee is generally intended to cover the costs of any unpaid sales commissions, not future guarantee premiums.

2.2 Mortgage Guaranty Insurance

Another widely available insurance and wealth management product with embedded options is mortgage guaranty insurance. **Mortgage guaranty insurance** is an insurance product that is purchased by mortgage lenders to provide protection against losses that can arise when a borrower defaults on a mortgage loan underwritten by the mortgage lender that is secured by real property such as a house, townhome or condominium.

Mortgage guaranty insurance is different from mortgage life insurance, which is typically purchased by the borrower and pays off the mortgage balance in the event of the borrower's death. Likewise, it is different from mortgage disability insurance, which is typically purchased by the borrower and covers mortgage payments during periods when the borrower is considered disabled, as defined in the insurance contract. Mortgage life and mortgage disability insurance are designed to prevent a mortgage from going into default in the event of the death or disability of the borrower. By contrast, mortgage guaranty insurance is designed to limit losses to the lender, and by extension the depositors and creditors of the lending institution, after a mortgage has gone into default and the foreclosure process has begun. It is not unusual for a mortgage to be covered by all three types of insurance.

Mortgage guaranty insurance is most often purchased for loans that are considered higher risk, either because there is very little equity in the property, i.e., the outstanding loan balance is almost as large as the value of the property or possibly larger, or the borrower is considered less creditworthy, e.g., due to a lack of steady income or a poorer than average credit history. However, it may be purchased for lower-risk loans as well. For example, a lender that considers its expertise to be in mortgage origination and does not wish to have a lot of credit risk on its balance sheet may, as a matter of practice, insure all the mortgages it originates.

Some lenders have a practice of bundling together the mortgages they originate and selling them to investors, either as stand-alone packages or in the form of mortgage-backed securities. A **mortgage-backed security** is a financial security whose cash flows are derived from the cash flows of a defined pool of mortgages. To make such securities attractive to investors, lenders will often insure some or all of the mortgages placed in the underlying pool.

Mortgage guaranty insurance plays an important role in the financial system. It enables borrowers who might not otherwise qualify for credit to obtain mortgages on more affordable terms. It enables lenders, particularly the smaller ones, to better manage their credit risk. Finally, it allows mortgage credit risk to be spread throughout the financial system, which helps to increase liquidity, i.e., the supply of available credit, in the mortgage market. Increased liquidity promotes competition among lenders and helps to lower the cost of mortgage credit for everyone.

2.2.1 Loss in the event of default

The coverage provided to a lender can vary from one mortgage guaranty insurance contract to another but generally includes the following:

- Outstanding loan balance at the time of default;

- Interest on the outstanding loan balance, calculated at the loan interest rate, until the claim is settled;
- Property taxes, property insurance premiums and property maintenance costs, including any that are in arrears at the time of default, until the claim is settled;
- Condominium or related property management fees that the property owner is legally required to pay;
- Legal costs associated with foreclosure; and
- Cost of repairs considered necessary to get the property in saleable condition.

The items in the last five bullet points above (that is, everything except the outstanding loan balance at default) represent the costs associated with settling a claim, and are collectively referred to as **settlement costs**.

In the event of borrower default on an insured mortgage, and subject to any deductibles, caps or other forms of risk sharing, the mortgage insurer pays the lender the outstanding loan balance and settlement costs less the amount the lender recovers from the sale of the foreclosed property.

Generally speaking, lenders only file a claim when the sum of the outstanding loan balance and settlement costs exceeds the expected amount recovered on the property. Even in cases where the lender files a claim and the amount recovered on the property turns out to be greater than the outstanding loan balance and settlement costs, the lender is not usually required to pay the mortgage insurer the difference. Hence, ignoring any differences in the timing of settlement costs and property sale, the loss to the mortgage insurer in the event of borrower default is

$$\max(B + C - R, 0),$$

where B denotes the outstanding loan balance at default, C denotes the total settlement cost, and R denotes the amount recovered on the sale of the foreclosed property net of real estate commission.

Letting $K = B + C$ and $S = R$, we recognize this expression as the payoff on a put option with strike price K and underlying asset price S. However, in this case, the strike price is the sum of a quantity that declines over time (the loan balance) and a quantity that varies randomly among claims (the settlement cost), whereas the underlying asset is the recovery value of the property, which depends on market conditions at the time of default. Consequently, a mortgage guaranty insurance contract is a put option with payment contingent on borrower default; that is, it gives the lender the right to put the mortgage loan back to the mortgage insurer in the event of borrower default.

2.2.2 Mortgage loan as a put option

Now consider a mortgage loan that is not insured against borrower default. Then, in the event of borrower default, the loss to the lender is

$$\max(B + C^* - R, 0),$$

where, as before, B is the outstanding loan balance at default and R is the amount recovered on the sale of the property, but now C^* is the lender's total settlement cost, which is generally less than the total settlement cost that a mortgage insurer would pay if the mortgage were insured. The reason the lender's total settlement cost is lower is that, unlike the mortgage insurer, the lender is not responsible for paying arrears interest on the loan. Indeed, from the lender's perspective, arrears interest is simply foregone income rather than an amount the lender must pay to another party.[7]

Consequently, the mortgage loan itself can be considered a put option. However, in this case, it is the borrower who holds the option. Indeed, in the event that the value of the mortgaged property falls below the outstanding mortgage balance, the borrower can put the property back to the lender. More specifically, the borrower can walk away from the mortgage and the property, provided that the lender cannot seek recovery by putting a lien on the borrower's other assets. This is true whether or not the mortgage is insured against default.

Of course, walking away from a mortgage is not without costs. Even if the lender has no recourse to the borrower's other assets, the default will still likely be reported to a credit bureau, which will make it difficult for the borrower to get another loan, at least in the short term. So, although a mortgage loan can be considered a put option, the borrower is unlikely to exercise this option unless it is deep-in-the-money.

2.3 Other Types of Insurance Guarantees

Embedded options arise in many other types of insurance products as well. Two noteworthy examples are guaranteed replacement cost coverage on property and inflation indexing of pension benefits.

Guaranteed replacement cost coverage is a rider, i.e., optional benefit, that can be added to a property insurance policy to provide protection against the risk that the cost of replacing an insured piece of physical property turns out to be greater than the amount of insurance coverage provided in the base policy. This benefit is valuable when the insured property is damaged beyond repair but the cost of replacing it is greater than what was expected at the time the policy was purchased.

For example, a natural disaster such as a tornado or wildfire could hit a densely-populated metropolitan area and destroy or severely damage a large number of homes in the area. Unless there happens to be a surplus of construction workers in the area, the increased demand for recovery and reconstruction services would ordinarily result in higher-than-normal costs, at least at the margin. If some or all of the equipment needed for reconstruction is also damaged in the disaster, the impact on costs could be even more pronounced. Guaranteed replacement cost coverage ensures that a policyholder whose house is damaged beyond repair has the financial resources to rebuild.

[7] From an economic perspective, one could argue that the lender still has to cover the cost of financing the outstanding loan balance during the foreclosure process. However, this cost would generally be less than the arrears interest on the loan. (If it weren't, the lender would not be able to stay in business for long.) So, even if the cost of financing were included, the lender's settlement cost would still be less than the mortgage insurer's.

Guaranteed replacement cost coverage has payout characteristics that are similar to a call option. To see this, let I denote the amount of insurance coverage and C the cost of repairing or replacing the insured property. Then, assuming there are no deductibles or loss-sharing provisions, the amount covered by a policy that does not have guaranteed replacement cost coverage is $\min(C, I)$ while the amount covered by a policy with a guaranteed replacement cost coverage rider is simply C. Since $C = \min(C, I) + \max(C - I, 0)$, the portion of the loss covered by the guaranteed replacement cost rider is $\max(C - I, 0)$. However, this expression is just the payoff on a call option with strike price I and underlying asset price C.

Inflation indexing is a feature in many defined benefit pension plans such as U.S. Social Security that periodically adjusts the amount paid to a person receiving pension payments for losses in purchasing power due to inflation. **Full indexing** ensures that the purchasing power of each payment at the time it is received is no less than the purchasing power of the first payment. **Partial indexing** ensures that any loss in purchasing power over time is no greater than a defined amount. In this context, inflation is typically measured using the consumer price index (CPI).

For example, consider a pension plan with full indexing and let I_t denote the level of the consumer price index t time periods after a particular individual begins receiving pension payments. Further, let P_0 denote the size of the initial pension payment and P_t the size of the pension payment t time periods later. Then

$$P_t = \max(P_0 \times I_t / I_0, P_{t-1})$$

Note that once the size of a pension payment is increased, it cannot be decreased even if the consumer price index subsequently falls. Hence, in the event that there is a decline in the CPI, the purchasing power of the pension payments actually increases; however, this increase is only temporary since the indexing adjustment is calculated using the first payment the pensioner receives rather than the most recent one. From this indexing formula, one sees that there are two embedded guarantees, a guarantee that the purchasing power of each payment will never be less than the purchasing power of the first payment and a guarantee that the absolute size of the payments will never decline.

To explore what sort of options are associated with these guarantees, first consider the situation at time 1. From the formula for P_t, we have

$$P_1 = P_0 + P_0 \times \max((I_1 / I_0) - 1, 0)$$

and

$$P_1 = P_0 \times I_1 / I_0 + P_0 \times \max(0, 1 - (I_1 / I_0)).$$

Hence, P_1 can be viewed as either the sum of P_0 and P_0 times the payoff on a call option on the scaled inflation index with strike price 1 and scaling such that the value of the index at time 0 is 1, or the sum of the initial payment adjusted for inflation, i.e., $P_0 \times I_1 / I_0$, and P_0 times the payoff on a put option on the scaled inflation index with strike price 1. Note that the equivalence of these interpretations is simply a statement of put-call parity.

Note that, from recursive application of the formula $P_t = \max(P_0 \times I_t / I_0, P_{t-1})$,

$$P_t = \max_{0 \le s \le t}(P_0 \times I_s / I_0).$$

Then, letting $M = \max_{0 \le s \le t}(P_0 \times I_s / I_0)$, we have

$$P_t = P_0 \times I_t / I_0 + (M - P_0 \times I_t / I_0).$$

Now, the parenthetical portion of this expression has a payoff similar to a put option; in fact, as we will see later, this is the payoff on a *lookback put option*. Hence, the payment at time *t* can be viewed as the sum of the initial payment adjusted for inflation, i.e., $P_0 \times I_t / I_0$, and the payoff on a lookback put option. Other interpretations of P_t are also possible. The reader is invited to rewrite the formula for P_t in various ways and discover some of these interpretations.

3 Use of Derivatives to Manage Risk in Insurance and Annuity Products

The remainder of this study note describes some of the ways derivatives can be used to manage the risks associated with operating an insurance enterprise.

Insurance companies, pension plans, and other financial providers use derivatives in three basic ways:

- To transfer some or all of the option risk embedded in an insurance product or pension obligation to the financial markets;
- To modify the characteristics of the asset portfolio, e.g., increase or decrease its duration; and
- To align the asset and liability sides of the balance sheet to protect and/or control the variation in the firm's surplus, i.e., the excess of the firm's assets over liabilities.

As will soon become clear, these uses are closely related with one another.

The usual way of transferring risk associated with an embedded option to the financial markets is to purchase a derivative, either on a derivatives exchange or over-the-counter. For example, if an insurance company has exposure to variable annuity guarantees of the GMAB type and the account values of the underlying variable annuities are all linked to the S&P 500, the company could purchase put options on the S&P 500. If the characteristics of the purchased puts, e.g., strike prices and times to expiration, are similar to the characteristics of the puts embedded in the variable annuities then at least some of the risk associated with the embedded puts will be transferred to financial markets.

A derivative or similar asset that is purchased for the purpose of transferring a risk on a company's balance sheet to the financial markets is known as a **hedge** and the process of buying (or selling) derivatives or other assets for this purpose is referred to as **hedging**.

Hedging does not eliminate risk; it only serves to manage it. For example, an insurance company that has written variable annuity guarantees is still fully responsible for all obligations arising from these guarantees even if the company purchases put options as hedges. Moreover, depending on the characteristics of the company's existing assets, the purchase of put options could actually increase the volatility of the company's asset portfolio. However, if the characteristics of the purchased puts are similar to the characteristics of the options embedded in the variable annuities then the company's surplus will generally be more stable and the risk that the company's assets are insufficient to cover its obligations as they come due should be reduced.

Derivatives traders are well aware that hedging does not eliminate risk. Consequently, the phrase "transferring risk to the financial markets" is a bit misleading. In a strict sense, the only true transfer of risk is an outright sale of a block of business because only an outright sale completely eliminates liabilities from a company's balance sheet. However, phrases such as "transferring risk to the financial markets" are commonly used. As long as one keeps in mind the real meaning of these phrases, there should be little confusion.

The other main uses of derivatives also entail the buying (or selling) of derivatives; however, the reasons for doing so are typically different. For example, a company may have a view on the future direction of interest rates and may wish to modify the portion of the asset portfolio that is not dedicated to backing specific liabilities to reflect this view. By taking long or short positions in interest rate futures or swaps, the company can change the duration of these assets without having to buy or sell the assets themselves. On the other hand, a company with different lines of business, each operating within its defined risk parameters, may discover that when the business lines are combined, the asset and liability sides of the balance sheet for the total enterprise are not well aligned, e.g., the residual exposure to the S&P 500 may be greater than desired. By taking appropriate positions in derivatives, in this case, long positions in S&P 500 puts or short positions in S&P 500 futures, the desired alignment for the total enterprise can be obtained.

In the sections that follow, we consider some of the ways in which derivatives can be used to manage the risks associated with retail insurance. We focus on two particular risks: the equity risk associated with variable annuity guarantees and the catastrophe risk faced by companies selling property insurance.

3.1 Hedging of Variable Annuity Guarantee Risk

In the earlier discussion of variable annuity guarantees, we saw that guarantees such as a GMDB or GMAB with a return-of-premium guarantee have characteristics similar to a European put option but with expiration contingent on the death or survival of the policyholder. Likewise, we saw that an earnings-enhanced death benefit has characteristics similar to a European call option but with expiration contingent on the death of the policyholder.

For simple guarantees such as these, a straightforward way to hedge at least some of the risk is to purchase a collection of European put and/or call options with expiration dates distributed in accordance with the anticipated survival pattern of the policyholders and hold these options until expiration. Provided that the indexes underlying the purchased options closely track the account values on which the guarantees are based, and the strike prices and notional amounts of the options are selected to match the characteristics of the guarantees being hedged, the result should be a lower overall risk position for the company providing the guarantees.

A hedge that is put in place and held until all the contracts associated with it expire is known as a **static hedge**. Hedging strategies that use only static hedges are referred to as **static strategies** or sometimes **hedge-and-forget strategies**.

A slight variation of the strategy just discussed involves periodically reviewing the portfolio exposure that remains after a hedge has been put in place and comparing it to the exposure that was expected based on the policyholder survival pattern originally assumed. If the exposure is less than what was expected, e.g., because more policies have lapsed, then some of the options originally purchased can be sold; if it is greater than what was expected, e.g., because there have been fewer deaths, then some additional options can be purchased. The resulting collection of options is then held to expiration unless a subsequent review reveals that some of the options are no longer needed due to a lower-than-expected exposure. This hedging strategy is still generally considered to be a static strategy because the intent when hedges are put in place is to hold them until expiration.

Most guarantees embedded in variable annuity products are not simple return-of-premium guarantees and cannot be hedged with simple index options. Indeed, as noted earlier, the guarantee value in a variable annuity contract is typically defined by a formula that is periodically recalculated, either on a regular schedule or at the discretion of the policyholder, and takes into consideration factors such as the performance of the underlying investments since contract inception, any promised rate of return, and the age of the policyholder. Moreover, the underlying asset is typically a blend of investments rather than a single index and can vary from one policyholder to another as well as over time. A company wishing to hedge such a guarantee can still use a static strategy but will need to purchase options that are considerably more complex and sophisticated than the standardized ones traded on an exchange.

3.1.1 Static strategies using exotic options

Options that have non-standard characteristics and are only available over-the-counter are known as **exotic options**. Exotic options are typically sold by investment banks and customized to meet the needs of the particular client. Given their degree of customization, exotic options are not actively traded and unlike exchange-traded options, their prices are not published on a regular basis. However, for companies purchasing these options as static hedges and planning to hold them until expiration, this is not usually an issue.

A number of exotic option designs have become quite popular and, over time, have acquired names of their own. We briefly consider five types of exotic options that are particularly useful for hedging variable annuity guarantees: lookback options, shout options, chooser options, rainbow options and forward start options.

Note that the exotic options purchased for hedging variable annuity guarantees are generally customized to the needs of the variable annuity writer and have a combination of features. For example, the purchased option could be a combination of a chooser option and a rainbow option with some lookback features. Hence, one should not interpret the assignment of names to these particular exotic options as constituting a strict classification scheme; rather, these names should be viewed as shorthand for identifying popular option features in an intuitive way.

3.1.1.1 Lookback options

A **lookback option** is an option whose payoff at expiration depends on the maximum or minimum value reached by the option's underlying asset during the life of the option. The form of the payoff at expiration depends on whether the option is a **standard lookback call**, a **standard lookback put**, an **extrema lookback call**, or an **extrema lookback put**.

Let S_t denote the value at time *t* of the index on which the option is written, let K denote the option strike price and let T denote the time that corresponds to option expiration. Further, let M denote the maximum value of S_t over the life of the option and let m denote the corresponding minimum. Then the payoffs at expiration for the various types of lookback option are as follows:

- Standard lookback call: $S_T - m$;
- Standard lookback put: $M - S_T$;
- Extrema lookback call: $\max(M - K, 0)$; and
- Extrema lookback put: $\max(K - m, 0)$.

Note that from the definitions of M and m, $m \le S_T \le M$, so the payoff on a standard lookback call or put is automatically non-negative. Moreover, the payoff in these two cases does not depend on the strike price.

Standard lookback options are sometimes referred to as **lookback options with a floating strike price**. When this terminology is used, the strike price is defined to be the value of the underlying asset reached during the life of the option that results in the largest payoff to the holder of the option; since this value is not known until the maturity of the option, the strike price is considered to be "floating." In the same vein, extrema lookback options are sometimes referred to as **lookback options with a fixed strike price**.

Lookback options are useful for hedging variable annuity guarantees where the guarantee value is periodically recalculated as the greater of the account value and the existing guarantee value. Consider, for example, a variable annuity with account values that track the S&P 500 index and guarantee values that are recalculated daily as the greater of the account value and the existing guarantee value. Such a guarantee can be hedged by purchasing a standard lookback put option on the S&P 500 index with expiration equal to the maturity of the variable annuity guarantee.

Most variable annuities do not recalculate guarantee values daily; for variable annuities with periodic recalculation of guarantee values, annual or triennial frequencies are more common. Hedging such guarantees with a standard lookback put option would provide more protection than necessary and be quite costly. An alternative would be to have an investment bank design a customized lookback put option that only considers the value of the underlying index on specific dates.

3.1.1.2 Shout options

A **shout option** is an option that gives the purchaser the right to lock in a minimum payoff amount exactly once during the life of the option at a time of the purchaser's choosing. When the purchaser exercises this right and a minimum payoff amount for the option is locked in, the purchaser is said to "shout" these instructions to the option writer, hence the reason for the name.

As before, let S_t denote the value at time t of the index on which the option is written, let K denote the option strike price and let T denote the time that corresponds to option expiration. Let S^* denote the value of the index at the time that the purchaser shouts to the writer. Then the payoffs at expiration for a **shout call option** and a **shout put option** are as follows:

- Shout call option:
  - $\max(S_T - K, S^* - K, 0)$ if the shout is exercised; and
  - $\max(S_T - K, 0)$ if the shout is not exercised.
- Shout put option:
  - $\max(K - S_T, K - S^*, 0)$ if the shout is exercised; and
  - $\max(K - S_T, 0)$ if the shout is not exercised.

Note that if the shout is not exercised, the payoffs are the same as for standard call and put options.

A shout call option allows the purchaser to lock in a one-time gain in the underlying index whereas a shout put option allows the purchaser to lock in a one-time loss in the underlying index. Note that if $S_t \le K$ for all t, the shout on a shout call option will never be exercised and if $S_t \ge K$ for all t, the shout on a shout put option will never be exercised.

Shout options are useful for hedging variable annuity guarantees in situations where the guarantee value is recalculated at the discretion of the policyholder. Consider, for example, a variable annuity with account values that track the S&P 500 index and a guarantee value that can be recalculated only once during the life of the contract at the discretion of the policyholder. Suppose that the guarantee value is initially equal to K and that in the event the policyholder requests a recalculation, the new guarantee value is equal to the greater of K and the account value at that time. Then if S_t denotes the account value at time t and S^* denotes the account value at the time the policyholder requests recalculation of the guarantee value, the value of the policyholder's account at the maturity of the guarantee including any payment needed to honor the guarantee is $\max(S_T, S^*, K) = K + \max(S_T - K, S^* - K, 0)$. However, we recognize the latter term in this expression as the payoff on a shout call option. Hence, the combination of the account value and guarantee can be hedged by implementing all of the following actions:

- Holding (the discounted present value of) the initial guarantee value K in cash (or a risk-free asset that matures when the guarantee does);
- Purchasing a shout call option with strike price equal to the initial guarantee value K; and
- Exercising the shout when the policyholder elects to have the guarantee recalculated.

3.1.1.3 Chooser options

A **chooser option** is an option that gives the purchaser the right to choose, after a specified period of time has passed, whether the underlying option is a call or a put. Chooser options are useful hedging tools for variable annuities with two-sided guarantees, e.g., a GMDB with a return-of-premium guarantee and an earnings-enhanced death benefit equal to 35% of any account value gains. The value of a chooser option at the time the choice is made is the greater of the value of a call and an otherwise equivalent put.

3.1.1.4 Rainbow options

A **rainbow option** is an option whose payoff depends on two or more risky assets. For example, the option's underlying asset could be a blend of the S&P 500 index, the NASDAQ 100 index, the Russell 2000 index and the EAFE index. Rainbow options are useful hedging tools when policyholders can hold multiple assets in their accounts and the guarantee applies to the account as a whole rather than individual assets in the account.

3.1.1.5 Forward start options

A **forward start option** is an option that will come into effect at some future time and be at-the-money at that time. An example would be a put option that will come into effect two years from now with strike price equal to the index value at that time and mature one year after that. Forward start options are useful for hedging guarantees that will come into effect during the payout period of a GMWB while the variable annuity is still in the accumulation period.

3.1.2 Dynamic hedging strategies

The hedging strategies discussed thus far have all been static strategies, i.e., the hedges are put in place and generally held until expiration. Static strategies generally focus on matching the cash flows and maturity dates of the hedge contracts with those of the variable annuity guarantees. Static strategies can be very effective and have the advantage of being relatively simple to understand. However, they cannot always be put in place, e.g., if the exotic options required to put the hedge in place are not available.

An alternative is to use a dynamic strategy. A **dynamic strategy** is a hedging strategy that involves frequent buying and selling of assets in the portfolio that backs a particular liability with the objective of matching any changes in the value of the liability with corresponding changes in the value of the assets, so that when the liability comes due, there are assets of sufficient value that can be liquidated to cover any payments associated with the liability. Note that the focus is on the value of the assets over time rather than the cash flows they generate.

Dynamic strategies can be very difficult to implement. At a minimum, one must be able to value the liability being hedged at every point in time and calculate the sensitivities of this liability to the key variables on which it depends. For variable annuity guarantees, this usually means calculating the first and second derivatives of the guarantee liability with respect to the underlying index (the delta and gamma hedge parameters), the first derivative of the guarantee liability with respect to interest rates (the rho hedge parameter), and the first derivative of the guarantee liability with respect to volatility (the vega hedge parameter). The challenge is that the values of some variables, such as volatility, cannot be directly observed and must be estimated. If the realized values of these variables turn out to be different than what was estimated then the sensitivities of the guarantee liability to these variables, i.e., the hedge parameters, will be different than expected. In that case, the hedge portfolio, which is constructed based upon expected sensitivities, will not be an appropriate match for the guarantee liability, and changes in the value of the guarantee liability will not be matched with corresponding changes in the value of the hedge portfolio.

Dynamic strategies are often deceptively simple to describe on paper and can give the appearance of providing almost complete protection for little or no cost. However, practitioners should beware. More than a few companies have discovered that the protection they thought they had in place disappeared when market conditions suddenly changed or the hedge wound up costing much more than expected. Companies considering using a dynamic strategy should keep in mind that if a dynamic strategy appears to be a lot cheaper than an equivalent static one, it probably means that an important risk has been overlooked.

3.2 Hedging of Catastrophe Risk

Property insurance is based on the premise that losses on individual properties are mostly independent. For example, a short in the furnace of one house that results in an electrical fire has no effect on the likelihood of a fire in a house on the other side of town. However, losses on individual properties are never completely independent. For example, a tornado that hits one house in a neighborhood is likely to do damage to other houses in the neighborhood and a hurricane that hits one city or state is likely to impact other cities and states in the vicinity as well.

Property insurers can manage most of the risks associated with property insurance by holding well-diversified, homogeneous portfolios of policies, e.g., insuring a large number of homes of similar type in many different geographic areas of the country. However, the risks associated with natural disasters such as hurricanes, floods, tornadoes, earthquakes and wildfires cannot be managed in this way because these risks are inherently non-diversifiable. Indeed, when a natural disaster strikes, it impacts all of the properties in a geographic area, not just one or two. Absent an effective way to manage catastrophe risks such as these, an otherwise healthy, well-diversified insurer could find itself insolvent.

The traditional way of managing catastrophe risk has been for an insurer to purchase **reinsurance**, which is simply insurance sold by a usually much larger insurance company to a smaller insurance company. It is assumed that the company selling the reinsurance, known as the **reinsurer**, has a strong enough balance sheet to be able to withstand the large losses typically caused by a natural disaster. However, in recent years, other approaches for managing this risk, which involve transferring some of the risk to the financial markets, have emerged. Two notable examples are weather derivatives and catastrophe bonds.

A **weather derivative** is a financial derivative whose payoff depends on the occurrence and severity of a defined weather event. For example, a derivative could pay the amount by which a particular loss index exceeds a defined base amount if a category 2 or greater hurricane makes landfall on the continental U.S. by a specified date. A few weather derivatives are available on derivatives exchanges; however, compared to most other derivative offerings, they tend not be very actively traded.

A **catastrophe bond** (or **cat bond** for short) is a bond issued by an insurance company where the repayment of principal and interest is contingent on there not being a catastrophe that results in large losses for the insurer. More specifically, it is a bond issued by an arms-length entity known as a **special purpose vehicle** that is established by an insurance company and with which the insurance company has a reinsurance treaty to cover a defined catastrophe. In the event that the catastrophe occurs, the special purpose vehicle makes a payment to the insurance company based upon the terms of the reinsurance treaty; otherwise, it returns the proceeds of the bond with interest to investors at the bond's maturity, which is usually the same time as the reinsurance treaty expires. This structure effectively gives the insurance company the option of defaulting on the bond without triggering a general default on the company's other indebtedness; moreover, if the company exercises this option, the holder of the bond has no recourse to the company's other assets. In exchange for providing the company this option, holders of cat bonds receive a higher interest rate.

Weather derivatives and cat bonds have been around in some form since the 1990s and continue to grow in popularity. However, they still represent a relatively small portion of the total capacity for covering catastrophe risk, and this is likely to remain the case for some time.