RECORD, Volume 22, No. 2*

Colorado Springs Meeting
June 26–28, 1996

Session 10TS
Financial Economics: The Option You Can't Refuse

Track: Investment
Key words: Finance, Economics, Pricing

Instructor: DAVID N. BECKER
Recorder: DAVID N. BECKER

Summary: This session overviews modern pricing of assets such as derivatives. It includes theory and practical applications, particularly how they relate to insurance companies. It also covers embedded options and strategies to hedge these options.

Mr. David N. Becker: The Society of Actuaries Foundation has undertaken a project to produce a textbook on financial economics as a resource for the actuarial profession. The research is being underwritten in part by the Foundation and in part by Lincoln National Corporation. The text is expected in 1997. To some extent the purpose of this session is to help explain why this effort was undertaken and provide an introduction to the topics and their utility. The book will be the work of a group of international researchers and actuaries. These include: Phelim Boyle, Daniel Dufresne, Hans U. Gerber, Heinz Mueller, Hal Warren Pedersen, Stanley Pliska, Michael Sherris, and Elias Shiu. The editor is Harry Panjer. The text provides a comprehensive background in the work of theoretical financial economics and also illustrates applications to the insurance industry.

A very extensive background in financial economics would include the areas of utility theory, risk theory, Markowitz's modern portfolio theory, the capital asset pricing model, the arbitrage pricing theory of Ross, and the topic of option-pricing theory. This last topic had its beginning with the work of Black and Scholes in solving the option-pricing problem for equity options. During the 1980s and 1990s this area of research has blossomed and extended to a far greater range of derivative securities, including both equity security-based and fixed income security-based derivatives, i.e., interest rate options.

I won't cover the topics that the financial economics text covers in detail. But I will describe some of the types of derivatives, their current history, risks, and applications. Later, I will focus on one topic—the concept of risk-neutral valuation. This is at the core of option-pricing theory and is one of the most important insights to understand.

A derivative security is a security whose value logically depends on that of another security, but whose cash flows do not arise from that security. A clear example is given by a call option on a common stock. This call option is the right, but not the obligation, to purchase the stock at a pre-arranged price by a certain date. Clearly, if the value of the stock rises, then the value of the call option rises. But there is no relationship between the cash flows of the option and those of the underlying common stock.

There are two other classes of securities that, although not technically derivative securities, do possess many of their traits. These are mortgage-backed securities and asset-backed securities. The key difference between these and true derivatives is that the cash flows of these securities really do depend on the cash flows of the underlying mortgage or other asset pool. Collateralized mortgage obligations (CMOs), which are a subset of mortgage-backed securities, can be purchased with such an array of tranche definitions that many of them are just as volatile as true derivatives. These are included on this list because mortgage-backed securities do have significant embedded options and can behave similarly to derivatives. They have prepayment risk if interest rates fall and extension risk if interest rates rise. Because they include these options, one must use techniques from option-pricing theory to fairly value them.

The securities underlying options include commodities, equities, fixed-income securities, and currencies. By extension, some derivatives are also based on pure interest rates and some on equity indices. It is possible to buy options that are based on the Standard & Poor's (S&P) 500 and on certain foreign equity indices.

Well-known examples of derivatives are forward and futures contracts. These are the earliest. Equity options arose next, basic put and call options on specific common stocks. These basic options come in either the "European" or "American" form. A European option can only be exercised on the day it expires. The American option can be exercised at any time up to and including the date of expiration.

In recent years there's been a whole new class of options that have wonderful names like Asian options, look-back options, barrier options, rainbow options, knock in, knock out.

There are also interest rate caps, floors, and swaps. An interest rate cap is a derivative security that pays off if a specified index, e.g., five-year constant maturity Treasury, exceeds a given threshold or strike level. The payoff is the product of the difference in the index and the strike measured on a specific day and a notional amount of principal. A floor is similar to a cap but pays off if the index falls below the strike level. A swap is usually the exchange of fixed payments for floating payments based on a notional amount of principal.

There are compound structures such as swaptions, which are basically options to enter into a swap; captions, which are options to enter into a cap; and futures options, which are options to enter into a futures agreement.

These derivatives have different tax and accounting treatments depending on the type of derivative and how it is being used by the company. They also differ in their risk-based capital requirements when held in insurance enterprises. A thorough understanding of how derivatives can be used to manage interest rate risk and of their tax, accounting, and risk-based capital requirements is valuable knowledge in managing an insurance enterprise in these times.

But without knowledge, derivatives are dangerous. There are two new risks called B–1 and C–1. The B–1/C–1 risk is the risk that you and your company show up on either page B–1 and/or C–1 of *The Wall Street Journal* in an article describing how a great deal of money was lost using derivatives.

A couple of years ago, you heard about Proctor & Gamble and Gibson Greetings. Both of them got into incredible difficulty because they played with derivatives and they didn't know what they were doing. Now the courts are going to decide whether it was that they didn't know what they were doing or whether it was the person or the company that sold them those derivative securities who didn't disclose the spectrum of risks that they were assuming.

Many mutual funds got burned on derivatives, for example, there was Piper Jaffray Companies. In fact, some of this spilled over to CMOs in mutual funds.

The law requires a mutual fund to make daily valuations of its securities to determine unit values. Some funds couldn't value many of these CMO tranches, so they called the investment banks that sold them the CMOs and asked them for a valuation that could be used for daily reporting. The investment banks operate under a rule that if you're going to sell it to somebody, you need to be willing to buy it back. If the investment banks were asked to put a value on the CMO purchased by a mutual fund, then they might be asked to buy it back for that price. With the volatility and uncertainty in the CMO market, the quotes from investment banks, if they received them, were very low. The problem was the bid/ask spread was astonishingly wide. Thus unit values were dramatically affected. In some cases the funds could not obtain valuations, and situations resulted where certain mutual funds weren't reporting daily values. Or they were using arbitrary values for the CMOs. So when volatility gets high, even securities like CMOs can get in trouble. It isn't limited to derivatives.

The trials and tribulations of Orange County and Baring America Asset Management Co., Inc. are well known. More recently, there has been the "copper meltdown" at Sumitomo, whose loss exposure could range from $1.8 billion to $4–5 billion.

One of the lessons here is that derivatives are very dangerous, although they are a powerful financial tool. They allow for incredible leverage. They can get you in serious difficulty. This doesn't mean, though, that you shouldn't use them. It does mean that you should know what you're doing. You should understand the mechanics of the derivative's cash flows. What are the risks in the derivative? Do you understand the accounting for derivatives? Do you understand the taxation of derivatives? You should understand what exchanges they may be trading on. If over the counter, what about counterparty risk? Are you buying securities for which you can easily liquidate your position?

You need an extensive risk control process inside your company in order to monitor your position and to determine your risk posture on a periodic basis.

There are four basic uses for derivatives. The first is speculation, i.e., you just want to make a bet. In fact, if you listen to some of the infomercials, that's what you think they are really all about.

Second, there is hedging. You've got a risk and you want to lay it off on someone else. There are different ways of hedging. You can, for example, "buy insurance" in case you are worried about interest rates falling. Let's examine how this might be done. Suppose you're going to receive $100 million in six months. You want to protect yourself against the fall in rates between now and the time you actually get the money to invest.

There are two things you could do. You could enter into a futures contract today. This effectively locks the rate in. Alternatively, you can purchase a floor or an option. For example, you could buy a floor, and if rates fall too far, the floor pays off. You could buy a futures option. If rates fell, then you would exercise your option to enter into the futures agreement. The terms of the futures agreement are determined not at the time of exercise of the option, but at the time the option was purchased.

If you hedged with the future and if rates are more attractive in six months, you're stuck with the rate you locked in. If you use the floor, you have protection if rates fall, but you don't give up the gain if rates rise.

The decision between a future or an option depends upon what's more important. Obviously to enter into a futures contract costs no money. You have to pay for the floor or an option.

There can be an advantage in buying the floor or an option, i.e., you allow yourself upside potential. Of course it's reduced by the price of the floor/option you're buying. But you do keep the upside potential.

Third, you can use derivatives to replicate other assets. This will be discussed shortly.

Fourth, there may be certain tax benefits from use of derivatives. For example, at the end of the year if you have securities that have unrealized losses on them, it's possible to overlay derivatives on top of those assets and get a tax deduction as if you sold the underlying security at a loss, even though you don't sell the security. If you do that, however, it changes the cost basis of the asset and may make it difficult to sell later without taking capital gains. Also, it is important to remove the hedge after the necessary holding period so you don't lose money on the hedge itself. One should always be careful of tax gamesmanship! And the tax treatment can change over time.

The following list displays applications of derivatives that can be beneficial to insurance enterprises. These are:

- Access to other securities and/or security markets
- Decreased transaction costs
- Altering asset mix and tactical asset allocation
- Currency hedging
- Asymmetric returns
- Asset/liability management.

Several of the above applications are intertwined or are different facets of the same transaction using derivatives. By the use of derivatives it is possible to participate in the risk/reward of securities without actually having to purchase those securities. For example, purchasing futures on the S&P 500 allows one to achieve equity returns without having to take a direct equity position. Options on such futures allow one to participate in upside return without the total exposure to a potential decline in market values. There are derivatives that allow one to participate in a similar manner in foreign markets.

A notable benefit of the use of derivatives is the reduction in transaction costs. In the above example there are smaller transaction costs by effecting the equity exposure by derivatives than by directly participating in the equity spot market. These savings can be especially significant in the foreign markets.

Using derivatives allows one to alter asset mix and implement tactical asset allocation strategies on a cost-efficient basis. In addition, by using derivatives one does not actually have to liquidate an existing position in one allocation class and purchase securities in another allocation class. By avoiding large-scale sales and purchases one will not "move the market" in performing such reallocations. If derivatives are used, it is also easier and quicker to change your allocation or mix or to liquidate one's position entirely. Again, transaction costs are minimized.

Foreign investments of suitable quality tend to have returns with low or negative correlation with that of domestic securities. By investing in them one can further reduce the variance of the portfolio's returns via the low or negative correlation of foreign with domestic securities. One does, however, take on currency risk. But there are currency derivatives that allow that risk to be hedged. Clearly, if a company operates in both the domestic and foreign markets, currency risks can be hedged.

Derivative securities often have asymmetric returns, i.e., their risk/return profile is not symmetric. Returns on real securities and portfolios of securities are also not symmetric. By adding suitable derivatives to a portfolio one can change the risk/reward profile of the entire portfolio in a cost-effective manner. Closely related to this is the use of derivatives in asset/liability management. For example, if one has sold an option in a liability, then one might be able to "buy it back" in the financial markets.

The key concept in option-pricing mathematics is known as risk-neutral valuation. This concept is at the heart of the breakthrough in option pricing. It also seems to be one of the most difficult concepts to explain, especially when one is doing interest-rate derivatives or valuation of fixed-income securities with embedded options. This concept is the principal topic of the remainder of this talk. For simplicity and clarity, the concept of risk-neutral valuation will be discussed in the context of a European call option on a common stock.

Before we do that, however, I want to provide a quick summary of important terms and assumptions in option pricing.

In speaking about security pricing models, we talk about the factors that are priced in the risk process. And when we talk about the factors priced, you'll hear the phrase arbitrage-free models. One of the confusions is that much of the option-pricing machinery used for derivatives is referred to as arbitrage-free.

After hearing discussions about option-pricing models, one tends to think that if a model is arbitrage-free, then it is a "good" model, and vice versa if it is not arbitrage-free. Actually, more fundamental is the issue of whether or not the model is risk-neutral. But often the phrase "arbitrage-free" is meant to include both the technical meaning of arbitrage-free and risk neutrality.

An arbitrage-free model is essentially a model that takes a given collection of securities that you essentially decide are fairly priced and uses it to calibrate an option-pricing model. You build the model so that it reprices those securities exactly. Usually this means that one will need a parameter (not a factor) in the model for each of the prices to be replicated. When this model is used to price a new security, the model price is said to be relatively consistent to the set of security prices in the collection. The utility here is that if securities A and B are priced relatively consistent to the underlying collection of securities, then A and B are priced relatively consistent to one another.

If there are errors in the prices of the underlying set, then these errors will be compounded in the pricing of securities outside of the set. When recalibrating an arbitrage-free model every day (or intraday), the parameters are volatile; i.e., their values are not stable. This is due to the fact that one is forcing the model to replicate exactly the prices of some given set of securities whose prices are fluctuating over time.

The alternative to an arbitrage-free model is an equilibrium model. This approach tries to capture the underlying character of the movements in interest rates or security prices. It uses many fewer parameters that are derived over historic time periods using statistical analysis of interest rates. It assumes that the day-to-day prices of securities contain noise. As such, it will not reprice the base securities exactly, but it will be relatively close and provide a good estimate of price over time.

The choice of using an arbitrage-free model or an equilibrium model depends on the purpose. If relative pricing today is the goal, then one wants an arbitrage-free model. This is the way it is used in the derivatives' market or the market for fixed income securities with embedded options. But if one is simulating financial performance over an extended period of time, then an equilibrium model is superior.

Factors refer to the number of random variables in the model. For interest-rate modeling the factors may refer to the short rate, the long rate, volatility of the short rate, or the long-run target mean-reversion rate.

The second item in modeling security prices is the risk process. This comes in two varieties: risk-neutral and realistic. Sometimes the phrase "risk-adjusted" is used for the risk-neutral approach and risk-unadjusted is used for the realistic approach.

A realistic model is one that models the real world, i.e., a world where there is a risk over the short-term default-free rate from investing longer, and investors demand a term premium for accepting that risk. This modeling framework is appropriate for risk analyses. If one is attempting to simulate the risk/return spectrum from holding a security, then a realistic model is appropriate. Recall that borrower behavior, i.e., the borrower calling the bond or the mortgagee prepaying his/her mortgage, is measured in the real world.

A risk-neutral model is one that assumes that there is zero market price of risk, i.e., that all investors are neither risk takers nor risk averse and will accept as a return the risk-free rate. The current term, premia (the additional compensation received for going out further on the yield curve), is embedded in the model via the initial term structure or by incorporating a zero market price of risk. Of course, this is not the real world. But, as noted before, this risk-neutral valuation is the key to option pricing and is examined in detail later. If one is pricing a security on a given day, then risk-neutral valuation is the method of choice.

In general, for pricing purposes, option-pricing models should be risk-neutral and arbitrage-free.

In order to provide a firm, mathematical foundation for option pricing certain assumptions are necessary. These include:

- Information is freely available.
- Unlimited borrowing and lending at the risk-free rate is allowed.
- Securities are infinitely divisible.
- Markets are complete (intuitively, this means there is a sufficient diversity of securities such that their payoffs span all possibilities in all possible future states of the world).
- Markets trade continuously with no transaction costs, taxes, or restrictions on short sales.
- Investors are "price takers," acting rationally on all available information and prefer more wealth to less.
- No riskless arbitrage is possible.

The phrase "no riskless arbitrage" means that it is not possible for an investor to make an investment with a zero net outlay, which has a positive return now or even a positive probability of a positive return in the future. In short, there is no free lunch. The basis for the no-riskless-arbitrage assumption is the high degree of efficiency of the market for securities.

These assumptions lead to the "law of one price" which states that if two securities, A and B, have identical cash flows in all future states of the world, then they must have the same price today. This law is based on the assumption of no riskless arbitrage. If markets are efficient, then if A and B had different prices investors would sell the overpriced security and buy the underpriced security. In this way, they would pocket the difference between the two securities and would use the cash flows on the security they purchased to pay their obligation on the security they sold. But efficient markets do not permit riskless arbitrage. If such an imbalance existed, it would be recognized and the difference in the prices of A and B would vanish rapidly in trading.

Now we are ready to examine the concept of risk-neutral valuation. We'll use a simple binomial example.

Let S be the price of a share of common stock. At the beginning of the month S = \$20. Let r, the riskless rate of return, be 1% per month. Suppose that the stock's price can only assume two values at the end of the month, either \$22 or \$18 per share. Let C be the price of a European call option on the stock, which gives the option holder the right, but not the obligation, to purchase the stock at the end of the month for \$21 per share (the strike price).

Our goal is to arrive at the price of the call option, i.e., to establish a value for C.

Suppose you own H shares of stock and have sold one call option, i.e., your portfolio is long H shares of stock and is short one call option on that stock. At the end of the month the stock will be worth either \$22 or \$18. In the first case the value of the option is \$1 (\$22–1), and in the second case the value of the option is zero, as the price of the stock is below that of the strike price of the option.

An interesting question is what value of H makes the portfolio riskless? In other words, how many shares of stock should be owned such that the long position of H shares coupled with the short position of one call option is risk-free? To be risk-free, the portfolio must have the same value no matter what the price of the stock is at the end of the month. This means the value of the portfolio in the up state and the down state must be the same. The value of the portfolio in the up state is

\$22 * H – 1; the value in the down state is \$18 * H. If these two values are equal one can solve for H. H = 0.25, or the portfolio must consist of one-quarter of a share of stock and short one call.

The value of this riskless portfolio at the end of the month is (\$22 * 0.25) – \$1 = \$18 * 0.25 = \$4.50. At the beginning of the month the value of the portfolio is (\$20 * 0.25) – C or \$5 – C. In the absence of riskless arbitrage a riskless portfolio earns the risk-free rate. Therefore the following holds: (1.01) * (\$5 – C) = \$4.50. Solving for C, C = \$0.5445. This is the value of the call option assuming no riskless arbitrage.

Note that the probabilities of the stock price moving up and down were not used. Therefore, the value of the call option is independent of the actual or real probabilities of upward and downward movement in the price of the underlying stock.

Now let's shift to the Black-Scholes analysis of the call option on a common stock. Here, S, the price of the common stock, is assumed to follow a lognormal distribution, which is an example of an Ito process or generalized Weiner process. This is a random walk with mean or drift μ and standard deviation σ. Let r be the risk-free rate and assume that this rate is constant over all maturities and constant over time. Black and Scholes found that the price of the call option is related to the price of the stock in the following way:

$$\frac{\partial C}{\partial t} + rS\frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} = rC$$

The Black-Scholes differential equation doesn't involve the mean of the distribution of stock prices; therefore, the solution does not involve the mean. Because the mean of the distribution of security prices reflects the necessary reward for risk, this indicates that neither the differential equation nor the solution depends on the particular risk preferences of individual investors. The solution is independent of the risk preferences of investors; i.e., you obtain the same solution no matter what risk preferences are held. In that case, any set of risk preferences may be used to evaluate C and you will get the same price. The simplest choice of risk preference is the risk-neutral preference. For a risk-neutral investor the expected return on a security is the risk-free rate, r. The present value of any cash flow in a risk-neutral world is obtained by discounting the expected value at the risk-free rate.

It is important to realize that the risk-neutral assumption is merely an artificial device to simplify the computations of obtaining the solution to the differential equation. The solution, the price C of the call option, is the same in all worlds, not just the risk-neutral world. When one moves from a risk-neutral world to a risk-averse world, three things happen:

- The expected growth rate in the stock price changes
- The discount rate used for any payoffs from the call must change
- The two effects always offset exactly.

Now let's revisit the simple binomial example. In a risk-neutral world the expected return on any security must be the riskless rate, 1% per month. Choosing the security to be our share of common stock, the probability of an upward movement in this risk-neutral world must satisfy ($22 * p) + [$18 * (1 – p)] = $20 * (1.01). Thus the "risk-neutral probability" is 0.55. Applying this to the call option, the value of the call option at the end of the month equals ($1 * 0.55) + ($0 * 0.45) = $0.55. Therefore, the price of the call at the beginning of the month is C = $0.55/(1.01) = $0.5445.

Thus the no-riskless-arbitrage assumption and the risk-neutral valuation assumption yield the same answer. This will always be true if the security price follows an Ito process and the underlying security is a traded security, i.e., not a commodity.

The theory is not claiming that the real world is risk neutral. Risk-neutral valuation simply provides a special frame of reference wherein the computation of the price is particularly simple. The argument above demonstrates that the price determined in the risk-neutral world will equal the price in a risk-averse world. The price determined by risk-neutral valuation also will equal the price in a no-riskless-arbitrage world.

The above also applies to interest-rate derivatives and the pricing of fixed-income securities with embedded options. Option-pricing models for these securities involve one of two approaches. One can generate risk-neutral interest rate paths and assign them the same probability, or one can generate realistic interest rate paths and assign them risk-neutral probabilities. (This is theoretically supported by Girsanov's theorem.) The difficulty that arises is that if equiprobable risk-neutral interest rate paths are used, many people cannot accept the results as they represent the mean across interest rate paths that "just cannot occur." There are often many paths of the yield curves that do not behave at all as yield curves have behaved historically. This behavior is stated, for example, in terms of the frequency of rates being high or low, the rapid rise or fall of rates, and the frequency and duration of inversions. As a result, there is a strong tendency for people not to accept the results of option-pricing theory when applied to nontraded insurance liabilities and the value of the insurance enterprise itself.

There is a problem in using risk-neutral interest rate paths coupled with behavioral assumptions, e.g., borrower call or prepayment, interest sensitive lapse, company interest crediting rate strategy, reinvestment and disinvestment strategies, that are based on a statistical analysis of behavior in the real world. If the risk-neutral interest rate paths are used, will the behavioral assumptions be valid? Probably not. And it is unknown how to adjust them with any degree of accuracy. By Girsanov's theorem one could generate risk-neutral paths having equal probability and, theoretically, convert them to realistic paths with risk-neutral probabilities. But work by Steve Craighead and Mark Tenney indicates that the number of paths needed would be enormous. There is more work that needs to be done in applying finance theory to insurance enterprises.

Insurance liabilities contain many options and guarantees that must be considered in light of option pricing. Examples include guaranteed surrender values, floor crediting rates, temporary guarantees in the crediting rate, flexible premium provisions, partial surrender and loan provisions, fixed/variable fund transfers in variable products, and minimum death benefit guarantees.

An insurance enterprise has embedded options in the assets/liabilities. There is potentially severe interest rate risk because of reinvestment and disintermediation. This risk may be quantified by applying the tools of financial economics on the insurance enterprise in either of two ways: apply to the value of the insurance firm or block of business in an integrated way, or apply to the asset/liabilities separately.

When doing so it is important to keep the following in mind. What assumptions must be made to apply the methods of financial economics to insurance enterprises? How credible are these assumptions? How do you deal with the uncertainty? Do all assumptions of financial economics apply in the real world? If not, how might the approach have to be changed?

Remember: liabilities in an insurance enterprise are not equivalent to bonds!

The question essentially is, isn't the merger/acquisition market a secondary market for insurance liabilities? Often people think of mergers and acquisitions as a secondary market for insurance liabilities. But that's not a real secondary market in the sense of a secondary market for assets.

For example, suppose you have an account at Charles Schwab and you own a 2002 GM bond, semi-annual coupon of 8% interest. If you go to Charles Schwab, it can sell that bond for you. Does GM participate in any way in that transaction? From GM's point of view, it doesn't really care whether you sold that bond to me or someone else. GM is totally uninvolved in it. It's a security. You've got a security you want to sell. You find a buyer and sell it. That's the secondary market for assets.

To have a similar secondary market for liabilities, you'd have to be able to come up to me directly or via some broker and say, "Would you like to buy my single-premium deferred annuity? Can we strike a deal?" And the answer is, you can't even ask the question because there is no such market for insurance liabilities. (There is a small market for life insurance policies on terminally ill individuals. It is not an efficient market.) In fact, it is not likely that there will be a true secondary market. First, policyholder's purchase insurance products to meet their particular needs. It is unlikely that another individual will have the same need. Second, to a great degree insurance products are very similar. Therefore, one can simply buy a new product rather than buy someone else's existing product. Third, there are adverse tax consequences to such sales. Recall the transfer-for-value rules on life insurance and annuities. Fourth, there are public policy issues, for example, insurable interest. Also note that there is no secondary market for savings accounts, bank certificates of deposit, etc.

In merger and acquisition situations, it isn't just the liability that is being sold. It's the liability and the supporting assets. And they are sold in exchange for the earnings (free cash flows) that result out of the interaction between the assets/liabilities that were transferred and however that combination is to be managed into the future by the buyer. And the buyer must be another insurance company.

This highlights one of the difficulties with treating assets/liabilities separately. With regard to liabilities, if there is no market, then there are no market values. Without market values one cannot calibrate option-pricing models to solve for an appropriate option-adjusted spread. Therefore, in applying option pricing directly to liabilities, one is forced to make an arbitrary choice of the spread to Treasuries to use in discounting liability cash flows. As a consequence, market values are arbitrary and so are option-adjusted duration and convexity numbers derived from those market values. If those values are blindly used in managing the business, there is a significant chance for suboptimization. Another problem is that the proposed spreads to be used in discounting liabilities lead to interpretative difficulties.