Insurance and Options

LARRY EISENBERG
University of Southern Mississippi

January, 2005

ABSTRACT

Conditional distribution pricing of derivatives in incomplete markets is introduced by pricing insurance policies with traded and non-traded underlying. A derivative's value is decomposed into hedged and non-hedged values.

Author keywords: insurance, derivatives
G11, G13

College of Business and Economic Development, University of Southern Mississippi, Hattiesburg, MS 39406-5076, email: eisenberg@cba.usm.edu, phone: 601 266-6732, fax: 601 266-4920.

I. Introduction

A. New Class of Derivatives: Insurance Claims as Underlying

Recently reinsurance companies have begun to sell a new class of derivatives. These contracts are hybrid derivatives that include as their underlying not only those found in the financial and commodity markets but also underlying consisting of the claims paid out or filed on insurance policies. It is not clear how many of these contracts have been sold. This information is proprietary and some, if not most transactions, are not reported in the trade journals.

B. No Complete Markets Models for Derivatives with Non-securitized Insurance Underlying

Derivatives traders today for the most part use models that assume markets are complete. The models assume all relevant risks can be perfectly hedged. This assumption is used in the Black-Scholes model.[1] It is an assumption that Harrison and Pliska showed is a necessary condition to obtain a unique martingale measure and hence an unique price for a derivative.

The assumption of complete markets is a useful fiction. This doesn't mean that traders behave as though these models represent reality exactly. If so what option trader would be try to be flat gamma or vega given a neutral view? Nonetheless models based on this assumption are useful to most traders. It is doubtful that investment and money-center banks would invest millions of dollars in these models and the computer systems that do the model calculations if these models did not assist in generating profits. Many desks on the trading floor would probably be out of business without them. In these markets the complete markets assumption is a useful approximation.

How useful is this assumption for hybrids? Suppose that an insurance company wants to sell a derivative where there is no liquid two-way market for some of the underlying such as claims paid when an insured oil rig in the Baltic Sea blows up or claims paid on medical insurance for the employees of a firm? Intuition suggests that the complete markets assumption is a poor approximation of an illiquid market where it may take weeks or months to do a round turn on a trade. Where the bid-ask spread may be as high as 30%. For such a market the assumption of complete markets is not a useful fiction and probably would lead to disastrous results for a trading desk selling such contracts.

C. "Partially" Incomplete Market Models

If a comple markets model won't work for hybrids, will an incomplete markets model work? In such model where only the lack of arbitrage justifies a price, there is no unique martingale measure.

Hence there is no unique price for a hybrid. Given that there is no unique price in these models there appear to be no models that a firm could use to price this contract. What role if any does hedging have in mitigating the risk of these contracts? If one can hedge, does one use the actuarial returns or the risk-neutral ones?

Is the best that a reinsurance company can do is to treat the contract as pure insurance, use only diversification to mitigate the risk, place the contract in the diversified portfolio and forget about it?

No – if we are willing to assume that an insurance or reinsurance company has a pricing function for pure insurance policies. Such a pricing function would most likely price the marginal contribution to the risk of the diversified portfolio by adding a short position in the pure insurance policy to the portfolio. The risk measure as well as the price of this risk is the choice of management. For example the risk might be the probability of ruin, or maybe the probability of losing a AAA credit rating. It might the standard deviation of the returns to diversified portfolio or the loss at the 95th percentile.

For hybrid derivatives, is the assumption that an insurance company has a function to price the insurance it sells an useful fiction, or is it as poor an approximation of reality as the complete markets assumption? The models here assume an insurance company has a pricing function for pure insurance policies. Probably most insurance companies do not have such a firm-wide consistent pricing function. Perhaps they will. It's obvious that they get paid for selling insurance, so they obviously assign prices to the insurance they sell. They price their policies even if the market for these policies is not complete. Perhaps assuming they have a rational firm-wide pricing function is a heroic assumption. On the other hand if an insurance company wants to use models of the kind represented here, it might price these contracts as if they had a firm-wide pricing function. Maybe it will eventually have one.

The purpose of this paper is to show how to price hybrid derivatives given an insurance pricing function. Assuming a pricing function implies assuming a function to price risk. Hence it assumes a measure of risk. We do not propose: a pricing function or a risk measure except to illustrate the pricing method and some of the issues that arise in using it.

D. Applications Not Specific to Insurance Companies

The potential application of the methods described here to price securities with partially hedgeable risk is rather far reaching.[2] Traders know that they cannot apply their complete market models as if they were an exact representation of reality. Even with the seminal Black-Scholes model, traders know that a literal following of the proscriptions of the model could be disastrous. The reason is that the model assumes continuous rebalancing of the duplicating portfolio for an option. Of course in the presence of transactions costs such rebalancing is impossible. Many articles have been written generalizing the Black-Scholes model for the presence of a bid-ask spread. The fact that most trading desks do not use such models is an indication of their limited usefulness. This observation applies to all complete-market models.

Application of the methods described here could be used to model a vanilla option as a hybrid contract with partially hedgeable risk. During the period between adjustments of the duplicating portfolio, the option position would have unhedged risk. Some might argue that another model for contracts with highly liquid markets is not needed. However, credit and emerging markets derivatives generally have markets which are far less liquid than those for vanilla options in mature markets. The literature on credit derivatives has developed to the point that at least two books on credit derivatives may be on the market this year. However, despite the elegance of many credit derivative models, one might ask how realistic it is to apply complete market models to first-to-default swaps on portfolios of corporate bonds.

Beyond complete market models, derivatives models must include diversification as well as hedging techniques. Progress appears to have been more rapid on the trading floor than in the insurance industry, but with the continuing integration of financial markets the learning process might be two-way. If investment and money-center banks learn from the insurance industry to price unhedged risk, then the methods described here might be used in markets other than the hybrid insurance derivative market.[3]

II. Flow of Information and Updates of the Loss Distribution

One of the issues that arises is what is the flow of information regarding updating the loss distribution (a probability density function) for claims paid out on the underlying insurance policy or policies. Unlike information in the financial markets, information in the insurance world can be lagged not seconds or minutes but months.

It is common that for many types of policies that the seller of the policy may not find out what insured events have occurred or what claims have been filed until months later the event or the filing. This is more likely to be the case when a contract is sold my a reinsurance company to an insurance company, and the underlying claims paid out are on policies sold by the insurance company to consumers or non insurance firms. A notable exception would property and casualty insurance when a catastrophic event occurs such as occurred with Union Carbide's Bhopal plant or the toppling of an oildrilling rig in the North Sea. Probably the traders managing the hybrid book at the reinsurance company would find out about the event on CNN that day.

The flow of information that effects updates of the loss distribution is crucial to hybrid models. This information flow determines how the loss distribution changes over time. Four cases of information flow are modeled here: one, the simplest case of no update; two, multiple update and three, continuous update. Here we shall assume the evolution of the sum of claims follows either a Wiener or a compound Poisson process. We will not attempt here to use all the distributions that underwriters and actuaries use in the everyday "nuts-and-bolts" business of estimating loss distributions. Our goal here is to illustrate with relatively simple cases.

A. Simple Case: No Update

The simplest is the case of no update during the covered period. An example of this case would be a hybrid with underlying policies which are a package of polices on many "small" events such as automobile insurance policies. If it were the last quarter of the covered period for a package policy sold to a corporation on its automobile fleet, then the insurance or reinsurance company might well have no information on covered events or claims filed until after the covered period. With a major but currently realistic caveat, if we assume that the only information used to update the loss distribution consists of reports on the covered events or claims filed under the policies, then the loss distribution does not change during the covered period.

The caveat relates to correlation with frequently observed random variables. Frequent updates on correlated variables imply changes in the conditional distribution for the unobserved variate. For the purpose of this case we assume that the insurance company either does not observe changes in the parts of the diversified portfolio which are correlated with the underlying policy, or that it does observe; but it does not use the information.

Correlated observations on other more frequently observed policies would generally not be used unless the required IT infrastructure were in place. Construction and testing of trading system software and a reporting system for keeping track of claims information from thousands of policies, many of them issued by other insurers is not a trivial undertaking. This author would argue that it will be some before such systems are in place. Hence the assumption of correlation with the diversified portfolio (if some of its policies are frequently observed) need not invalidate examining this case.

When such systems are in place, this simple case may serve only pedagogical purposes. However all covered periods come to an end. If observations of correlated polices or variables are done weekly, for example; then with six days to go we have a contract with no updates. The reader need only see the flurry of activity on a trading floor when options are coming to expiry to know that the correct modeling of a contract close to expiry can be important. Assuming multiple updates when there are none left could be a problem.

Whether there are updates or not, both Wiener and Poisson processes or mixtures might be used to model policies with no update. However the amplitudes of the jumps would have to small so that we don't read about the jumps in the news that day.

B. Periodic Updates at Preset Dates

This is just the case above except that the time until the end of the covered period is long enough that the hybrid seller gets at least one update and maybe more. Similar to the above case, another example would be that of a number which is known with a lag but whose returns are not independent of those of a periodically observed process. Processes could be jump, diffusion or mixtures of both.

C. Continuous Jump Updates

If we model the blowing up of an oil rig as a Poisson Process, every hour that the hybrid seller gets no news is good news – that CNN has not reported that the oil rig has blown up. Since there is less time until the end of the covered period than one hour ago, the probability of a blow up before the policy expires has decreased.[4]

Another example is that of a hybrid option that includes the price of a commodity with an illiquid market. In this case despite the lack of liquidity for this commodity market a price might be reported hourly or daily. If the commodity returns are more accurately modeled by a jump rather than a diffusion process we have continuous updates with jumps. Admittedly an insurance company might want to avoid selling such a policy as the buyer of the policy may have much more knowledge of this market. Hence the buyer may have a better understanding of how prices change in this market than the insurer does, and thereby may be able to pick off insurers who under price this insurance.

Similar to the above case, another example would be that of a number which is not frequently observed but whose returns are not independent of those of a continuously observed jump process. More examples include any number frequently reported for which there is no traded security indexed to that number. Rainfall, degree days and power usage for regions or time periods that have not been securitized are examples if their time series have jumps.

D. Continuous Diffusion Update

This case is similar to any of the cases immediately above except those where a big jump might occur such as the explosion of an oil rig. These are the four cases of loss distribution updating discussed here. We proceed to the main idea.

III. No Updates on the Loss Distribution

A. A Pedagogical Example

Consider a basket option with two underlying. The first underlying y is the claims paid out for the covered period on some policy which has no update during the covered period. The second x is the US dollar price of Japanese yen (USD/JPY) (times some yen notional) and the claims paid out on some policy with no update. USD/JPY follows a standard geometric diffusion process and has a liquid two-way market.

The contract payoff occurs after the covered period when all the claims information is in and any dubious claims have been investigated and thrown out or accepted. (We assume no indefinite on going litigation.) This payoff occurs at $T+\delta$ where T is the end of the covered period. (In the financial markets one might think of T as the expiry date of the hybrid.)

The payoff is $\mathcal{H}(T+\delta) \equiv max\{y(T)+x(T)-K, 0\}$. We might think of $y(T)$ as the value of y at T known with a lag of δ. Perhaps interest accrued on the payoff over δ is also paid, so that the payoff might be multiplied by $e^{r\delta}$.

It is key to note that at time T the insurance company is now short a standard insurance policy with a fixed retention (strike) of $K' \equiv (K - x(T))$.[5] The key assumption here replacing the assumption that there is a liquid two-way continuously updated market with no lag for y (the complete markets assumption) is that the insurance company not only sells insurance for a price but has a pricing function Q known at time 0. What are the arguments to Q at time T? They are the latest known joint density function on $y(T)$ and any retentions or limits on $y(T)$. In this case there is only a retention K'. K' is a function of the fixed parameter K and $x(T)$. The latter has just been observed and is now fixed. We are also making two more assumptions. First, the insurance pricing function Q prices the incremental risk to the diversified portfolio by including the insurance policy $\mathcal{H}(T+\delta) = max\{y(T)-K', 0\}$ in the diversified portfolio z. Actually with the inclusion of this insurance policy $z(T)$ becomes $z(T)$ plus the price of the insurance policy at T or $z(T)+\mathcal{H}(T)$.[6] Second, this measure of risk is known at T. Note that at T the insurance company is pricing the random payoff at $T+\delta$. This payoff is given by the contract as the basket option payoff. More formally

$$\mathcal{H}_T = Q_T\left[\, g(y(T), z(T)); x(T), \mathcal{H}(x(T), y(T))\right]; \tag{1}$$

$\mathcal{H}_T \equiv$ *hybrid price at T*

$Q_T \equiv i$*insurance pricing operator applied at T*

$g(y(T), z(T)) \equiv$ *joint density of* y *and* z

$z(T) \equiv$ *value of the diversified portfolio at T*

$y(T) \equiv$ *total claims to the underlying insurance policy*

$x(T) \equiv$ *USD/JPY at T*

$\mathcal{H}_{T+\delta}(x(T), y(T)) \equiv$ *contractual payoff to hybrid as a function of the underlying.*

Hence at any $t < T$ the insurance company knows that everything that will be needed to price the hybrid at T is already known at t except for $x(T)$. But $x(T)$ is the price of a traded asset. At t the company knows it will be short an insurance policy at T whose price is a function of only one variable: the price of a traded asset. Determining the value today t of a security whose price at T is a known function Q_T of a traded asset $x(T)$ is a complete markets problem which we already know how to solve. The price $\mathcal{H}_t$ is the expectation over the distribution for $x(T)$ conditional on $x(t)$ where $E_t x(T) = e^{(r-r_f)} x(t)$

$$\begin{aligned} \mathcal{H}_t(x(t)) &= e^{-r(T-t)} E_t\left[Q_T(x(T)\,|x(t))\right] \\ &\equiv \int_0^\infty Q_T(x(T)) f_t((x(T)\,|x(t))\, dx(T). \end{aligned} \tag{2}$$

B. The Idea Again

In the previous example the hybrid turned into a contract whose only source of randomness was the insurance payout y. The traded asset x was known and fixed prior to y, and became a parameter of the insurance policy at T through K', the new retention. By assuming that y was revealed after x it was natural to condition the payoff of the hybrid on the value of $x(T)$. Then we calculated statistics on that conditional distribution which is exactly what $Q_T(x(T)\,|x(t))$ does. The next step was to observe that Q_T is a function of $x(T)$. Finally we overved that a claim which delivers Q_T as a function of $x(T)$ is attainable because there is a complete market for x. In other words, if the market for contracts whose payoffs are functions of x is complete, then we can price $\mathcal{H}_t(x(t))$. To use the completeness of markets whose returns are spanned by $x(t)$ we needed to show that the price of $\mathcal{H}_T$ is a function of x.

Apart from a quick review the point is that we do not have to restrict ourselves to cases where y is revealed after x.

Suppose that they are both revealed without a lag. Consider any statistic of the conditional density function at t for $\mathcal{H}_T$ conditional on $x(T)$ $f_t(\mathcal{H}_T|x(T))$. Since the statistic is calculated using $f_t(\mathcal{H}_T|x(T))$, it is not a function of $y(T)$. For example the expectation of $\mathcal{H}_T$ conditioned on $x(T)$

$$\begin{aligned} E_t[\mathcal{H}_T|x(T)] &= \int_{\mathcal{H}(T)=0}^{\infty} \mathcal{H}_T f_t(\mathcal{H}_T|x(T)) d\mathcal{H}_T(T) \\ &= \int_{y(T)=0}^{\infty} \mathcal{H}_T(x(T), y(T)) f_t(y_T|x(T)) d\mathcal{H}_T(T) \end{aligned} \tag{3}$$

is not a function of $y(T)$. Also, statistics on the joint distribution $h(z(T), y(T)|x(T))$ are not functions of $(z(T), y(T))$, but they are functions of $x(T)$, or if they are both independent of x these statistics are at least not functions $(z(T), y(T))$.

If the pricing operator Q_t is calculated integrating over the conditional density for $\mathcal{H}_T$ on $x(T)$, we have a standard complete markets problem. By (0) this is the same as calculating Q_t using $f_t(y_T|x(T))$ or $h(z(T), y(T)|x(T))$, which is also a complete markets problem. As before we can take the expectation of $E_t[\mathcal{H}_T|x(T)]$ over $x(T)$ given $x(t)$ assuming x has a risk-neutral drift and discounting at the risk-free rate. With riskless discounting

$$\begin{aligned} \mathcal{H}_t &= e^{-r(T-t)} E_t[E_t[\mathcal{H}_T|x(T)]|x(t)] \\ &= e^{-r(T-t)} \int_{x(T)=0}^{\infty} f_t(x(T)|x(t)) \left[\int_{\mathcal{H}(T)=0}^{\infty} \mathcal{H}_T f_t(\mathcal{H}_T|x(T)) d\mathcal{H}_T(T) \right] dx(T). \end{aligned} \tag{4}$$

C. Complete Markets, Attainable Contingent Claims Definition

Before we go to the result which summarizes this idea we need a definition of complete markets. A market is said to be complete if every contingent claim is attainable. A claim is attainable if its payoff can be duplicated using a self-financing trading strategy in traded assets. A contingent claim is a claim whose payoff is a function of the claims upon which it is contingent. These claims are traded assets.

IV. No-Update European Hybrid Option Price

A. Assumumptions

1. Complete Submarket

These are the standard complete market assumptions for the submarket of the partially complete market which is complete. Hence reasonably behaved functions of $\boldsymbol{x}$ can be hedged and priced based on the absence of arbitrage.

1a. There are n traded assets whose prices are $x_1, ..., x_n$. $\boldsymbol{x}(t)$ is the vector of prices of these assets at time t. $\boldsymbol{K}_p$ is a vector of fixed parameters of arbitrary dimension.

1b. The market for contracts with payoff $\mathcal{P}(\boldsymbol{x}(T),\, T; \boldsymbol{K}_p)$ is complete.

2. Incomplete Submarket

These assumptions are twofold: one, they state that there is a submarket of the partially complete market which is not complete; two, they are "lottery complete." Insurance can be purchased on any function of the non-hedgeable risks $\boldsymbol{y}$ as long as the statistics for the distribution of that function used to price lotteries exist and are finite. Each insurance company has a pricing function which it applies to all lotteries. The pricing function's arguments are the observations of the risks $(\boldsymbol{y}, \boldsymbol{t})$ and the joint distribution of observations on the risks at the contracted times and observations on its diversified portfolio of risks on the payout date T.

2a. There are m events whose outcomes can be measured. The measures of these outcomes are y_1, ..., y_m. $\boldsymbol{y}(t)$ is the vector of the measures of these risks at time t.

2b. There is are insurance companies that sells insurance by buying the outcomes of lotteries with payoff $\mathcal{Y}(\boldsymbol{y}(T),\, T; \boldsymbol{K}_y)$ for any function $\mathcal{Y}(\,\cdot\,)$.

2c. There is no liquid two-way market for insurance contracts $\mathcal{Y}(\boldsymbol{y}(T),\, T; \boldsymbol{K}_y)$. That is, the bid-ask spread and the time it takes to do a round-turn in an insurance policy is so large that underlying risks cannot be used to hedge any contracts.

2d. There is no use of updated information between time 0 and time T to change the marginal distribution $\boldsymbol{y}(T)$ or the joint distribution $g(\boldsymbol{y}(T),\, z(T))$ where z is the diversified portfolio of the insurance company.

3. Hybrid Market

3a. Contracts $\mathcal{P}(\boldsymbol{x})$ and $\mathcal{Y}(\boldsymbol{y})$ bought and sold subject to the constraints above on the incomplete insurance market.

3b. A hybrid contract $\mathcal{H}$ has a payoff at some time $T+\delta$ where $\delta \geq 0$ which is a function of $\boldsymbol{x}(T), \boldsymbol{y}(T)$.

4. Pricing Operator

4a. An insurance company determines its price for all lotteries with one pricing operator Q. This operator can be applied to hybrid contracts $\mathcal{H}(\boldsymbol{y}(T), \boldsymbol{x}(T), T; \boldsymbol{K}_h)$ as well as insurance contracts $\mathcal{Y}(\boldsymbol{y}(T), T; \boldsymbol{K}_y)$. Different companies may have different pricing operators.

4b. The pricing operator prices statistics of the distribution of $\mathcal{H}(\boldsymbol{y}(T), \boldsymbol{x}(T), T; \boldsymbol{K}_h)$ conditional on $\boldsymbol{y}(T)$.

B. Result

At T the price of the hybrid is just its payoff $\mathcal{H}(T)$. Consider the density function for $\mathcal{H}(T)$ given $\boldsymbol{x}(T)$ but not given $\boldsymbol{y}(T)$. Rather $\boldsymbol{y}(T)$ is distributed according to the loss distribution function for $\boldsymbol{y}$ which has not been updated since time 0.

Apply the lottery pricing operator Q to the conditional density of $\mathcal{H}(T)$ given $\boldsymbol{x}(T)$, $Q_T[\mathcal{H}(T)|x(T)]$. $Q_T[\mathcal{H}(T)|x(T)]$ is clearly in the set of functions $\mathcal{P}(\boldsymbol{x}(T), T; \boldsymbol{K}_p)$. Hence Q is an attainable claim. By standard results for complete markets the price at $t \leq T$ of $Q_T[\mathcal{H}(T)|x(T)]$ is the risk-neutral expectation over $x(T)$ of $Q_T[\mathcal{H}(T)|x(T)]$. Hence

$$\mathcal{H}(t) = e^{-r(T-t)} E_t[Q_T[\mathcal{H}(T)\,|x(T)]\,|\,x(t)]. \tag{5}$$

Note that if the payoff is at $T+\delta$ with $\delta \geq 0$ with $\boldsymbol{y}(T)$ known at $T+\delta$ and $\boldsymbol{x}(T)$ known at T replace the above expression with

$$\mathcal{H}(t) = e^{-r(T+\delta\,-t)} E_t[Q_{T+\delta}[\mathcal{H}(\,T+\delta)\,|x(T)]\,|\,x(t)]. \tag{6}$$

Even if the contract is defined so that the covered insurance period ends at some date T' before the date T that x is fixed one can still use the procedure above for $t \leq T'$. However after T' the hybrid will have become a standard option with completely hedgeable risk.

C. When is It Cheaper to Not Hedge All the Hedgeable Risk?

To answer this question we need to compare prices when we apply the pricing operator to all the risk of a hybrid versus just the risk of the distribution of the hybrid payoff conditional on the traded assets. These prices are of course dependent on the insurance pricing function used by the company. In general for the no-update case the difference in price will be

$$\mathcal{H}_{unhedged}(t) - \mathcal{H}_{hedged}(t) \tag{7}$$

$$= E_t[Q_T[\mathcal{H}(T) \,|\, x(t)] - e^{-r(T-t)} E_t[Q_T[\mathcal{H}(T) \,|x(T)] \,|\, x(t)].$$

To gain further insight into relative costs regarding updates versus no updates and hedging versus not hedging, we put forth a lottery pricing function as part of the model.

V. Conclusion

Conditional distribution pricing of derivatives in incomplete markets is introduced by pricing insurance policies with traded and non-traded underlying. This technique does not require the use of risk premia, and can be used to price credit derivatives with non-traded credit risk and vanilla options when markets are not complete due to stochastic volatility, jumps or trading less frequent than market changes. Also derivatives are priced subject to non-traded risks that have zero correlation with all traded risks without requiring the assumption of perfect diversification. A derivative's value is decomposed into hedged and non-hedged values. When different desks are responsible for hedging and diversification such as underwriters and traders at a reinsurer, this decomposition can be used to track the volatility of these values and the efficacy of each desk in managing risk.

---

[1]Black and Scholes also derived their option pricing equation using CAPM. Rubenstien also derived the Black-Scholes model using an equilibrium model and traders with log-normal utility functions. However most models used on the Street today are not derived with assumptions on utility functions of market participants or that there are equilibria in the financial markets.

[2]The author would like to thank Pat Hagan for his comments on this topic.

[3]The models here assume that some risk is completely hedgeable. They would have to be modified if no risk is completely hedgeable. Instead the models here would allow for residual risk of impferfectly rebalanced duplicating portfolios.

---

[4]If we are using a Poisson or compound Poisson process we have the constant probability of a jump per unit time. If there is less time left until the end of the covered period, then time left times the probability of a jumpe per unit time decreases as time left decreases.

[5]If $K - x(T) \equiv K' < 0$, then we have a negative strike. This means no matter what we are adding a positive number to $y(T)$, so $\mathcal{H}(T) > 0$. The opton on the insurance payout is now a forward on the sum of the payout plus K'.

[6]One can think of z as representing the random value of the diversified portfolio and also as the collection of contracts (legal documents) for all the insurance policies that the company is short. The same for $\mathcal{H}$.