Reinsuring for Catastrophes through Industry Loss Warranties – A Practical Approach

Ali Ishaq, FCAS, MAAA

Abstract:

Within the last couple of decades natural and man-made catastrophes have become a source of increasing concern for the insurance industry. Industry Loss Warranties (ILWs) are reinsurance products whose payout is triggered by catastrophic insured loss. There is a growing market for ILWs because they provide a viable alternative to traditional reinsurance and catastrophe bonds for mitigating losses from such events.

This growing market requires a consistent and sound way of pricing ILWs. The process is made simpler because pricing ILWs does not require knowledge of individual client's exposures but only the expected industry losses. Available catastrophe models provide a ready source of industry loss distributions. Conceptually it is simple to go from a given industry loss distribution to pricing an ILW, but ILWs can vary in their terms and conditions depending on the needs of a particular client. This paper shows how to account for some of these terms and conditions to price ILWs and provides an example of such calculations.

Keywords: Industry Loss Warranty (ILW); OLW; Catastrophe Reinsurance; Pricing Rare Events; Empirical Loss Distributions

---

> "Everything should be made as simple as possible, but not simpler."
>
> --Albert Einstein

1. REINSURANCE AND RARE EVENTS

Extremely rare events, by their very nature, are hard to insure. Insurance companies like to cover events that are rare for the insured, but in aggregate have a distribution that is stable and predictable. In one class of rare events, i.e., natural catastrophes, the loss is either very severe or the damage so wide-spread that there is a real need for primary insurers to spread the risk of loss from these events through some reinsurance mechanism. Until the large losses from natural catastrophes in the last few decades from events like Hurricane Andrew, primary insurers were content to mitigate the risks from natural catastrophes through traditional reinsurance instruments including treaty and facultative reinsurance. Before this time there were a few who had raised the alarm of losses from natural catastrophes large enough to shake the financial foundations of the P&C insurance industry, but not much attention was paid to their fears. Since then, the concern over catastrophic property exposure has been continually growing and newer market mechanisms, such as catastrophe bonds and industry loss warranties, are becoming more established.

2. WHAT IS AN INDUSTRY LOSS WARRANTY?

There are many kinds of instruments gathered under the rubric of industry loss warranties (ILW) also known as original loss warranties (OLW). Essentially they all cover losses from events where the industry-wide insured loss exceeds some pre-agreed threshold[1]. This structure, i.e., where the operative trigger is an industry loss rather than the company's own loss, implies some risk that there could be a loss to the reinsured portfolio without triggering the ILW if the corresponding industry loss is smaller than the industry trigger amount. This is the 'basis risk' for the reinsured. This risk is higher for companies whose exposure concentrations are farther away from the industry averages. Therefore ILW covers are typically bought by companies whose portfolios closely follow the market. This disconnect can be mitigated to some extent by choosing the right kind of trigger. The trigger amount can vary by geography, level, and the kinds of events that contribute to it.

There are many kinds of industry loss warranties available in the market. The variety comes from the kind and level of the industry loss chosen. The industry loss considered as a trigger can vary by amount or geographic scope. For example, an ILW may promise to pay when one of the following happens:

1. A hurricane with industry-wide insured loss in Florida in excess of $15 billion but less than $25 billion.
2. A winter freeze with industry-wide insured loss in North America in excess of $20 billion.
3. An earthquake with industry-wide insured property loss in excess of $35 billion anywhere in the world.
4. Second wind loss with industry-wide insured loss in excess of $10 billion anywhere in the US and territories.

---

[1] In order for these contracts to qualify as insurance there may be a second condition for loss payment that the insured company's loss exceed a pre-determined amount in addition to the industry being in excess of the given threshold. Usually the insured company's loss trigger is set to such a low level that for our analysis we can ignore it.

In the first case if the hurricane does a lot of damage to property that is not insured but the insured amount is small or if the damage is outside the State of Florida, the ILW may not be triggered. In the second example the total of $20 billion may arise from damage in both US and Canada. And in the third instance, the location of the earthquake does not matter but the loss would have to exceed $35 billion after casualty losses have been excluded before the ILW is triggered. There can be industry loss warranties that respond to a second-event of its type. The fourth example above is an instance of such a second event ILW.

All the above types together can be thought of as Occurrence ILWs, because they respond to (first or second) occurrence of single large events. ILWs can also be structured so that the coverage applies, not in the case of one large event, but when a series of catastrophes in a year add to exceed a pre-determined amount. For these ILWs, the industry losses contributing to the total are limited so that a single event would not trigger the coverage. Only losses above a certain amount are considered towards the total because there are few industry mechanisms to keep track of smaller losses. For example one could construct an aggregate ILW that pays when all losses in California that cause insured damage of least $100 million but not in excess of $5 billion sum to more than $3 billion in a twelve month period. In either case there may be a provision for reinstatement of the ILW limit upon payment of an agreed premium.

3. INDUSTRY LOSS WARRANTIES COMPARED TO OTHER CATASTROPHE REINSURANCE INSTRUMENTS

Catastrophe bonds, traditional reinsurance, and industry loss warranties each have strengths and weaknesses with respect to their ability to address the risk from catastrophic events. The following table summarizes some of these comparisons:

	Traditional Reinsurance	**Catastrophe Bonds**	**Industry Loss Warranties**
Availability	Wide	Limited	Increasing
Transaction Cost	Medium	High	Low
Risk Charge	High	High	Low
Basis Risk	Small	Small	Variable
Pricing Risk	Medium	Large	Small

Traditional reinsurance is widely available but the risk charge for layers that cover catastrophic risks may be large. Part of this is due to the difficulty in estimating the reinsured company's future losses in these very high layers. For the reinsured the basis risk is small, but the reinsurer has to estimate future losses from a historical portfolio that may be different from the future distribution of insured exposures: there is usually a lag between the information the reinsurer uses to price and the actual exposure that emerges over the insured period. This creates a kind of reverse basis risk and the reinsurer has to charge a larger risk load to cover this risk.

Successful catastrophe bond offers have allowed some large P&C insurers to access the large capital capacity of the world bond market, but these products are still considered non-traditional and complex by many bond traders. These bond offerings are more complex than traditional bond offerings and so this skepticism, on the part of the bond traders, may translate into limited marketability and a higher risk charge. Catastrophe bond offerings are generally not standardized and each offering has to be individually structured and underwritten. This translates into high transaction and fixed costs, which may put these instruments outside the reach off all but the largest insurers.

Industry loss warranties are unfamiliar to many primary companies but, properly structured, may provide an inexpensive solution to many of the catastrophic reinsurance needs that these companies may have. These products have low transaction costs because the pricing risk for the sellers is comparatively low; they do not have to evaluate the expected losses to the reinsured portfolio from a given trigger, but only the loss distribution of the industry portfolio. This lowers the uncertainty and thus the needed risk margin.

The pricing risk for ILWs is lower but not zero. There is still the inherent parameter risk from trying to estimate the loss distribution of events whose frequency is not known and may be changing. Some scientists have postulated that there is a long-term climate cycle which governs the changes in frequencies of large weather events. One study estimates that this climate cycle be as long as 100,000 years[2]. Since the data available to formulate the frequency distributions is rarely longer than a few decades, there can be large error in our estimates of future probabilities of these catastrophic events.

---

[2] See study by Mukul Sharma in the June 10th issue of Earth and Planetary Science Letters (Elsevier, volume 199, issues 3-4)

4. THE MARKET FOR INDUSTRY LOSS WARRANTIES

ILWs can be used to reduce a company's exposure to sharp losses from large events or a collection of events thus controlling the tail of the aggregate loss distribution at a reasonable price. For many companies the tradeoff between cost and stability gained can be very favorable, even when comparing to other reinsurance products. As hinted above, there needs to be care in selecting the appropriate trigger to reduce the basis risk for the ceding company.

ILWs have generally been bought by the larger national insurers but they may be even more useful to the regional insurers. The large insurers have generally been the first ones to utilize this market because they tend to have a more sophisticated view of risk and price within the reinsurance market. This product provides these large buyers another reinsurance option for spreading the risk from large events over a larger market capacity. For the regional companies, the ability to cede the risk from extreme events in a concentrated area may enable them to allow larger geographical accumulations, thus allowing them to concentrate on their areas of expertise while staying within their capital constraints.

The relatively low cost of ILWs is due to the lower information asymmetry as compared to most other reinsurance products. Normally as we move from primary insurance to facultative reinsurance to treaty reinsurance to retrocession the information asymmetry increases sharply; the reinsurer has to base its pricing decisions on less and less current information as compared to the information that the reinsured is using to make its buying decisions. This implies that the reinsurer has to build an increasingly larger margin for error (or risk) into their estimates of expected loss under the contract.

Since ILWs are priced based on the industry loss distribution this information asymmetry is suddenly reduced (or even reversed, since the reinsured has to estimate the basis risk it may be taking on[3]), and thus the margin for risk can be correspondingly smaller.

5. PRICING AN INDUSTRY LOSS WARRANTY

It is important that the pricing be consistent as well as accurate. The more consistently a product can be priced, the better the prospects of a successful market. If there is no

---

[3] See the paper 'On the Basis Risk of Industry Loss Warranteis' by Lixin Zeng in the Summer 2000 issue of The Journal of Risk Finance for a discussion of estimating ILW basis risk.

consistent way of pricing a product the market prices may vary widely, which can lead to a fractured market or a race towards the lowest price provided. This can lead to a lack of confidence and consequently a lack of viability of a market for the product. In order to price ILWs consistently we need an acceptable way to calculate the loss distribution for a given industry loss warranty. For the illustration here an industry loss warranty is completely defined by its trigger. We will assume that every time a trigger is met, we need to pay the limit to the insured or that every loss is a full limit loss. This is not too much of a stretch because partial payments are comparatively rare. This assumption helps simplify the analysis and is not essential for the pricing methodologies described. Since we are using the actual industry loss distribution instead of the frequency alone, these methods can be easily adjusted to account for scenarios in which partial payments are allowed.

We start with the industry loss distribution and the definition of the industry loss warranty we want to price. The first step is to extract the conditional distribution of industry losses that meet the requirements of the trigger. Then we can combine the probabilities of triggering the ILW with the payout conditions in the contract to estimate the expected loss distribution for the contract.

6. CONCEPTUAL SIMPLICITY, PRACTICAL DIFFICULTIES

Conceptually the pricing is simple: calculate the distribution of losses under the contract. The expected value from this distribution is the expected loss under the contract and the shape of the distribution can be used to set a risk load. The sum of the expected loss, expenses, and the risk load is the theoretical premium needed.

Even with the conceptual simplicity it is clear that it would require some work to implement these steps in an actual pricing model. The first difficulty is in obtaining the industry loss distribution. This distribution or set of distributions would indicate the probability of industry-wide insured losses from various types of catastrophes that we want to insure. The most obvious possibility may be to use historical industry losses but unadjusted historical losses are not a good predictor of future industry losses. To be used in our calculations, these losses would have to be adjusted for exposure changes, inflation, and any other factors that would make the expected outcome of a historical event different in the prospective year from the historical year. We would also have to adjust for the fact that historical data is limited and incomplete. This adjustment is compounded if we consider that the historical data is a small slice from a possibly changing distribution of extreme events. If we want to price international ILWs we would also have to compensate for the fact that reliable historical loss data for most jurisdictions outside the US and Japan are not generally available.

A more practical source for industry loss distributions is catastrophe modeling data from commercial catastrophe models. This data is already adjusted for changes in exposure and has a good fit to historical record. In addition, the various probabilities coming out of these models are quickly converging to the industry consensus return time estimates for various kinds of events. There is still considerable variation between the various commercial models, but there is great pressure to estimate probabilities that are in line with the consensus estimates. As these models come closer, pricing derived from the inherent distributions would serve to make the pricing of the ILWs more consistent across the market, and this consistency will in turn make the product more marketable.

Generally the industry loss data from the commercial catastrophe models is available in the form of empirical distributions. Our first thought might be to fit a theoretical size of loss distribution to the output from a simulation model. But this may not be the best course because it raises new difficulties. The shape of the distributions for various perils (wind, earthquake, etc) may be very different and may not allow the use of a simple class of loss distributions. Secondly, since we have to censor or otherwise manipulate these distributions to extract the distributions under the various triggers, this may prove to be difficult or lack closed form solutions. So, we have a tradeoff between realism (fit to actual or prospective losses) and ease of computation. Theoretical distributions that are robust may be hard to use and the analysis may be hard to extend as we develop layers of analyses to sum from individual regions to multiple regions and triggers to the distributions of portfolios. Even after the availability of derivative distributions for the various types of triggers we are still left with the task of farther manipulating these distributions to achieve a consistent pricing. The advantage of symbolic distributions would be that a unique expected loss value and corresponding distribution around this mean would presumable exist for each trigger type and size. But the complexity of the process may make this scenario impractical.

7. EMPIRICAL DISTRIBUTIONS

> "God does not care about our mathematical difficulties. He integrates empirically."
> --Albert Einstein

As an alternative to fitting theoretical size of loss distributions to the empirical distributions derived from historical losses or from the output of a catastrophe model, we could use these empirical distributions directly in our pricing. If empirical distributions are used wisely, they can provide sound answers and allow us to get to the answers much more quickly the calculation of the final premium starting with one such empirical distribution is illustrated below.

After we obtain the expected future industry-wide insured loss distribution we can modify it to extract a loss distribution that corresponds to a particular trigger. For example, if we know the prospective industry loss distribution from hurricanes in Florida and we are pricing an ILW with a trigger of $20 billion industry loss, we can censor the distribution at $ 20 billion to determine the industry loss distribution above this point. The expected loss under an ILW can be calculated from the probability alone, but we can use the additional information in the distribution of losses that trigger the contract to estimate the appropriate risk load. To these we add the company's expenses to arrive at an estimate of the premium to be charged.

8. AN EXAMPLE

The first step is the output from either the historical loss analysis or a simulated distribution from a catastrophe model. Exhibit 1 shows what such a distribution might look like. Here the distribution is shown as a list of industry loss amounts along with a description of the geographic scope and peril. This distribution represents the expected losses from a 1000-year simulation period. So, if there are 100 losses above $20 billion in this list, we are assuming that the probability of a $20 billion or larger loss is 10% over the next year. The exact format of the loss distribution is not important and as long as enough information is available, the various formats are convertible from one into another.

The covered event for an ILW may be limited as to geographic area as well as peril. For next step, therefore, we construct a new loss distribution selecting only those events from exhibit 1 that meet the definition of our trigger. In this example we would exclude all events that have a Florida loss of less than $20 Billion. Table 2 illustrates this censored distribution.

The third step is to summarize this distribution by simulation year. In this example this step is used to calculate an annual cost for the ILW. For each year of simulation we calculate the payout as well as the reinstatement premium if any. This calculation can be simplified by our assumption that each time the ILW is triggered there would be a full limit payment and a full reinstatement if there is a reinstatement still available for that year.

Thus for each of the simulated years the model calculates whether a pay-out would happen:

$$PayoutTrigger_{event,year} = \left\{ \begin{array}{ll} 1 & \left\{ \begin{array}{l} \text{if } Loss_{event,industry} > Trigger_{\min} \\ and \\ \text{if } Loss_{event,industry} < Trigger_{\max} \end{array} \right\} \\ 0 & \text{otherwise} \end{array} \right\}$$

Here:

$$Loss_{event,industry} = Industy\ loss\ from\ a\ given\ event$$
$$Trigger_{\min} = Minimum\ Industry\ loss\ that\ will\ trigger\ a\ payout$$
$$Trigger_{\max} = Industry\ loss\ above\ wich\ there\ will\ be\ no\ payout$$

From the above the total number of losses payments for each year of the simulation is calculated as:

$$PayoutTrigger_{year} = \sum_{event} PayoutTrigger_{event,year}$$

The loss payments for each year are limited by the number of reinstatements allowed for the contract and then multiplied by the limit sold by the reinsurer to derive the loss pay-out for each year of the simulation. This 1000-point loss distribution for the expected pay-out can then be directly used to calculate the expected loss and the variance around the expected value for the treaty being priced. These calculations are shown in Exhibit 3.

Other types of ILWs including second-event covers, corridor covers, collar covers, and a consideration for a 'no-claim bonus' can be priced in a similar manner.

9. CALCULATING THE LOAD FOR RISK

Once the expected loss has been estimated the next step is to determine how much risk or profit margin to charge for a given exposure. In general, we should charge higher premium for higher risk. This can be done using various measures of risk. One way would be to use the coefficient of variation (CV) of loss derived from the pervious analysis to target a profit margin. Once we have decided the relationship between profit load and the CV based on the market and company appetite for risk and the required return on capital the result could be summarized as a relationship between ROE and the CV. If the result is a linear relationship between risk and ROE and if we want to limit the minimum and maximum return to realistically achievable levels this relationship may look like the following:

$$ROE_{\text{Target}} = \begin{cases} ROE_{\min} & \text{if } CV \le CV_{\min} \\ ROE_{\max} & \text{if } CV \ge CV_{\max} \\ \dfrac{CV - CV_{\min}}{CV_{\max} - CV_{\min}}(ROE_{\max} - ROE_{\min}) & \text{if } CV_{\min} < CV < CV_{\max} \end{cases}$$

$Where:$

$ROE_{\text{Target}} = \text{Target Rate of Return}$

$CV = \textit{Coefficient of Variation}\ \text{of loss under the contract}$

Given the expense and brokerage information, the ROE target can be converted into a loss ratio or a combined ratio target.

Another possible risk load could be based on the risk of loss on the contract once the reinstatement premium is taken into account. One way of calculating this amount is to calculate the profit or loss by year by summarizing the profit or loss from each event in each of years of simulation. If we calculate the expected value and CV of this distribution, this gives us another measure of return that we could target. We could also eliminate all years in which there is a profit which would leave years in which the contract is in a deficit. If our risk appetite is more in line with limiting the loss from a contract in any one year we could use the expected value of this distribution as a measure of expected downside and use this as another constraint on our final price or the desirability of a given contract and the market price.

10. OTHER ISSUES

10.1 International ILWs

Reinsurance markets outside the US are smaller and less developed. Therefore, even though exposure is much higher in US the risk margins in ILWs are generally lower. This is partly due to the higher parameter uncertainty in Asia and Europe. The catastrophe data and modeling are have more inherent uncertainty outside the US and therefore the reinsurer should charge a premium for taking up the risk arising from this uncertainty. As a result, the product may get too expensive compared to the traditional reinsurance and government guarantees. To the extent that ILWs compete on the basis of price with traditional reinsurance, they may be at a disadvantage in the international market, until either the traditional reinsurance prices rise or the catastrophe potential and the corresponding models improve. Some markets are providing worldwide coverage without charging much of a premium for the high parameter risk inherent in the worldwide models. For large buyers this may be an excellent opportunity in the short term, but one wonders if this aggressive stance is sustainable in the long run.

10.2 Aggregate ILWs

An example of pricing Aggregate ILWs has not been provided in this paper, but a methodology very similar to the one used here can be employed to calculate the loss distribution of an Aggregate ILW. Typically these products pay when the sum of industry losses between a minimum and maximum in a year exceed the trigger amount. For example, if an aggregate ILW gets triggered when the sum of industry losses that exceed $100 million but limited to $500 million exceeds $4 billion in a year, one would add losses exceeding $100 million limited to $500 million for each year of simulation and test if this sum exceeds the $4 billion trigger amount. This process would result in a loss and corresponding premium and profit distributions very similar to what result from an occurrence ILW calculation shown herein.

10.3 Lack of Protection from Unforeseen Events

Generally ILW sellers look backwards when deciding what kinds of triggers to offer. Before Hurricane Andrew there was probably not a big market for Florida Hurricane ILWs with triggers in excess of $15 billion. And more recently some buyers look for second and third event coverages, but these generally have triggers above $2 billion. Aggregate ILWs generally exclude losses with industry loss amounts above $500 million. Therefore none of these ILWs would have protected companies in 2004 when a series of hurricanes hit the Southern US Coast, with industry losses in the range of $1 billion each so that none of the ILWs mentioned above would have triggered. This again illustrates how ILWs have less parameter risk than some other reinsurance products where unforeseen losses may account for a large amount of loss in bad years.

10.4 Model Creep

On a related note, a one may notice that the commercial catastrophe models generally get recalibrated after almost every extreme event. This will probably mean that if before 2004 the models indicated four large hurricanes hitting Southern US as extremely improbable, these distributions would be revised.

10.5 Managing a Book of Industry Loss Warranty Business

Once we have appropriately priced a book of ILWs the next question is how to manage this book and the aggregate risk it contributes to the reinsurer's portfolio. If we have used a consistent set of loss distributions to price the ILWs, the aggregate loss distribution for the portfolio can be calculated as the sum of distributions of individual ILWs. This loss distribution can then be combined with loss distributions from other product portfolios to generate a companywide loss distribution. This combined loss distribution could act as a robust input into a DFA or ERM model to gauge the overall risk–return profile of the company.

11. CONCLUSION

Industry Loss Warranties provide an alternative to traditional reinsurance and catastrophe bonds when insurers are trying to smooth their results from the impact of catastrophic events. We have looked at one way to price for these instruments that takes advantage of the empirical industry loss distributions available as an output from many commercial catastrophe reinsurance models to simplify this task.

Exhibit 1: Simulated 1000-year Industry Loss Distribution

Year Number	Industry Loss (Millions)	Catastrophe Description
4	4,679	FL Hurricane
4	2,586	FL Hurricane
5	2,948	Winter Storm
7	19,000	FL Hurricane
8	3,438	FL Hurricane
10	9,242	CA EQ
10	3,304	FL Hurricane
12	3,293	FL Hurricane
14	5,234	Winter Freeze
15	4,636	FL Hurricane
16	2,949	FL Hurricane
17	26,424	CA EQ
17	7,532	FL Hurricane
17	5,419	NY Hurricane
17	4,426	FL Hurricane
19	24,939	CA EQ
20	2,739	FL Hurricane
20	2,603	FL Hurricane
20	2,165	Winter Freeze
23	3,912	FL Hurricane
24	2,441	FL Hurricane
26	20,638	FL Hurricane
26	5,507	FL Hurricane
27	2,573	CA Landslide
28	4,946	FL Hurricane
29	9,626	CA EQ
.	.	.
.	.	.
.	.	.

Exhibit 2: Losses Meeting ILW Trigger

Year Number	Industry Loss (Millions)	Catastrophe Description
26	20638	FL Hurricane
42	24801	FL Hurricane
63	24323	FL Hurricane
153	20977	FL Hurricane
179	30669	FL Hurricane
205	22307	FL Hurricane
232	23976	FL Hurricane
288	27315	FL Hurricane
343	34381	FL Hurricane
431	33108	FL Hurricane
438	20223	FL Hurricane
467	28063	FL Hurricane
467	26904	FL Hurricane
518	70029	FL Hurricane
614	28195	FL Hurricane
640	22597	FL Hurricane
725	29006	FL Hurricane
730	22173	FL Hurricane
779	22259	FL Hurricane
793	20996	FL Hurricane
811	47370	FL Hurricane
866	22261	FL Hurricane
893	56128	FL Hurricane
897	37107	FL Hurricane
908	21207	FL Hurricane
966	20701	FL Hurricane

Terms:

ILW Limit = 100 million
1 reinstatement at 150% initial premium
Epense = 20% of Premium
Initial ROL 5%

Exhibit 3: Simulated Results for the ILW

Year Number	ILW loss	Premium	Profit/ Loss
1	0	5	4
2	0	5	4
3	0	5	4
4	0	5	4
5	0	5	4
6	0	5	4
7	0	5	4
8	0	5	4
9	0	5	4
10	0	5	4
.	.	.	.
.	.	.	.
.	.	.	.
26	100	12.5	-90
.	.	.	.
.	.	.	.
.	.	.	.
467	200	12.5	-190
.	.	.	.
.	.	.	.
.	.	.	.
1000	0	5	4
μ	2.6	5.19	1.55
δ	16.83		
μ/δ	6.47		

Biography of the Author

Ali Ishaq went to school at the University of Texas at Austin where he studied Mathematics. He started his career at the Texas Department of Insurance where he built models to predict insurance company insolvencies. Later, at Aetna Property & Casualty he constructed underwriting models based on Neural Networks. At CNA Reinsurance Company he helped manage the property catastrophe aggregates and built pricing models for various reinsurance products. Currently, Mr. Ishaq is a Regional Actuary at Zurich American Insurance where he works on a diverse book of construction related insurance business.