AIRCURRENTS

MINIMIZING BASIS RISK FOR CAT-IN-A-BOX PARAMETRIC EARTHQUAKE CATASTROPHE BONDS

06.2010

EDITOR'S NOTE: AIR Worldwide has long dominanted the market for catastrophe risk modeling and analytical services related to both the issuance and portfolio management of insurance-linked securities. In this article, AIR principal engineer Dr. Guillermo Franco, provides insight into new techniques for reducing basis risk associated with parametric cat bonds. His research will be published in the November issue of the peer-reviewed journal "Earthquake Spectra."

By Dr. Guillermo Franco

Interest in insurance-linked securities as a means of transferring risk to the capital markets continues to grow. The challenge is to design a bond that will provide coverage for the bond's sponsor while minimizing basis risk and, at the same time, provide the necessary transparency demanded by the investor.

Insurance-linked securities of the "cat-in-a-box" variety typically employ a very simple trigger mechanism. A cat-in-a-box bond is a form of parametric bond; that is, their trigger depends on the physical parameters of the event. In the case of earthquake risk, three fundamental criteria usually determine whether the principal is paid: The epicenter of the earthquake must be located within a specific geographic zone (box), its magnitude must be higher than a certain threshold, and the actual fault rupture must occur within a certain depth. If any of these three criteria are not met, the bond does not trigger and no payment is made.

While parametric catastrophe bonds are appreciated for their transparency (earthquake parameters are publically reported by well-established seismological agencies), they are also known to carry *basis risk*, which can be understood as the difference between the actual losses experienced by the sponsor and the payment the sponsor receives. Basis risk can derive from a variety of sources, chief among them:

1. The design of the catastrophe bond leverages a catastrophe model to calculate the probability of experiencing certain levels of loss. The discrepancy between the losses estimated by the catastrophe model and the actual losses constitute basis risk. This portion of the basis risk is sometimes referred to as *model risk*.
2. Due to the simplicity of the trigger structure, it is also possible that the trigger does not behave as expected; i.e. that it results in payment for events for which the model estimates low losses or in no payment for events for which the model estimates large losses. This error is henceforth referred to as *trigger error*.

AIR Worldwide (AIR) continuously refines and enhances its models in order to reduce model risk. The task, then, is to reduce basis risk deriving from the second source—trigger error—in as much as possible, assuming that basis risk related to the first source cannot be further reduced. Research at AIR (which will be published in the November 2010 issue of Earthquake Spectra) has resulted in a set of novel algorithms that identify triggers that are both simple and transparent and, at the same time, minimize the probability of error of these parametric structures. This article discusses how these algorithms work in the context of a hypothetical catastrophe bond for Costa Rica (which was chosen because the end of the state insurance monopoly, *Instituto Nacional de Seguros*, has led to keen interest in innovation in risk transfer).

AN ILLUSTRATION FOR A HYPOTHETICAL CAT BOND FOR COSTA RICA

For the sake of simplicity, the payment structure of our bond will be represented by a binary outcome: the bond pays the entire principal if an earthquake triggers the bond and it pays zero if it does not. In our example, the bond is designed to trigger for earthquakes that cause a loss equal to or higher than the 100-year return period (1% exceedance probability) loss to a portfolio of insured properties whose geographic distribution is similar to that of the industry's as a whole.

Figure 1 shows the epicenters of simulated earthquakes contained in the 10,000-year catalog of AIR's earthquake model for Costa Rica. The catalog is as large as it is in order to capture the full spectrum of future earthquake experience. Highlighted in red are the simulated events in the catalog that would trigger the bond—i.e., that cause losses to the portfolio equal or higher than the 100-year return period loss. It is no surprise that most of the triggering events are located around the capital San José, since that is where most of the country's exposure is concentrated.

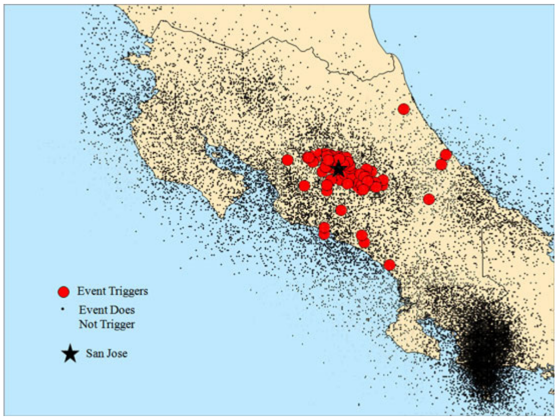

Figure 1. Epicenters of simulated earthquakes in the AIR model's 10-year catalog showing the events that would trigger a cat bond for a threshold loss corresponding to the 100-year return period (in red).

It is, of course, an easy matter to determine which events in the AIR model would trigger the bond based on estimates of their expected loss. The challenge, in terms of minimizing trigger error, lies in using the information in the model to successfully differentiate earthquakes that *should* trigger the bond from those that *shouldn't* based solely on their physical parameters, namely, their magnitude, depth and the location of the epicenter.

To begin, we'll use a set of Z zones to differentiate earthquakes according to their location. For each of these zones, we define distinct magnitude and depth thresholds. If an earthquake in a particular zone fulfills the criteria dictated by the thresholds in that zone, then the earthquake will trigger the bond. Otherwise, it will not.

The problem that needs to be solved, then, is to identify, for each zone, the optimal magnitude and depth thresholds—the thresholds that will minimize trigger error, the second contributor to basis risk.

OBTAINING A "LOWER BOUND" SOLUTION THAT PRODUCES MINIMUM TRIGGER ERROR

Editor's Note: The interested—and technically-inclined—reader is encouraged to read the following paragraphs, which lay out the methodology by which the optimal trigger is achieved. The less technically-inclined reader may wish to skip to the section "The Bottom Line: Design Options for Sponsors".

For purposes of computational simplicity, we divide the geographic domain of interest into K boxes of length d and width d. In Figure 2, for example, K=12. Every box in this grid is associated with one of the aforementioned zones for which appropriate magnitude and depth thresholds are defined to determine if a specific earthquake will trigger the bond. For example, we could define Zone 1 as the region comprised by the four top boxes of Figure 1, Zone 2 as the second row of boxes, and Zone 3 as the bottom four. In each of these zones, different thresholds of magnitude and depth will apply. The more zones we consider, the more flexibility we have to differentiate earthquakes according to their physical characteristics; the more precisely we can differentiate, the lower is the likelihood that the trigger will produce an erroneous outcome. That flexibility is maximized when the number of zones equals the number of boxes (Z=K=12 in Figure 2)—that is, when a pair of thresholds is defined for each box.

An exhaustive search for the best plausible pair of magnitude and depth thresholds for each of the Z=K zones results in a "lower bound" solution, one that will yield the lowest possible number of erroneous outcomes of the trigger for a given box size (length, width = d) and a given loss threshold.

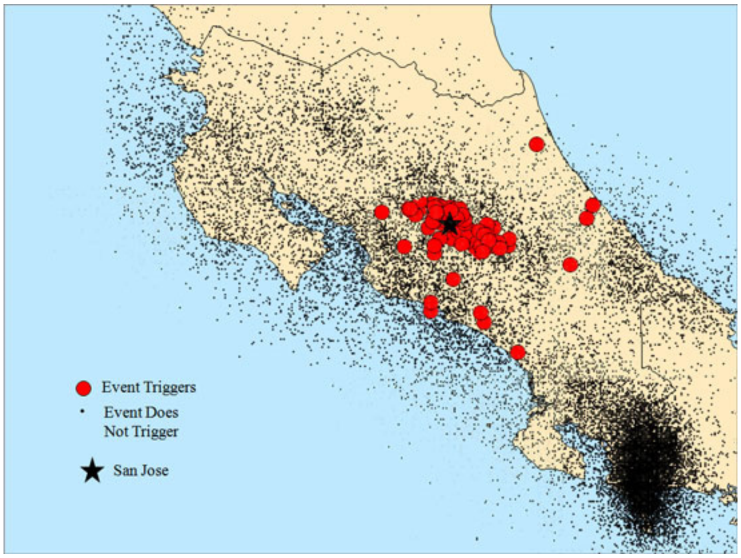

Figure 2. Example of a grid of K=12 boxes covering the domain of interest in Costa Rica showing the events that should trigger a cat bond for a threshold loss corresponding to the 100-year return period. (From Franco, 2010; published with permission from EERI/Earthquake Spectra)

At this juncture, it is worth defining trigger error more formally as the number of instances (events) in a 10,000-year stochastic catalog in which the trigger mechanism (the chosen box size and depth and magnitude thresholds) fails to reproduce the desired outcome. This error can be split into two contributions, namely *positive* and *negative error. Positive error* is defined as the error that ensues when an event produces a loss that is lower than the loss threshold but whose parameters nevertheless trigger the catastrophe bond. Positive error means that the sponsor is purchasing unnecessary protection, and therefore may paying an unnecessarily high premium, for events that would not cause catastrohic losses. *Negative error* is defined as the error that ensues when an event produces a loss equal to or higher than the loss threshold but fails to trigger the bond. Negative error means that the sponsor is not getting adequate protection for certain truly catastrophic events.

If we find the lower bound solutions assuming boxes of different sizes (for example, d=1.0, 0.5, 0.25, 0.1, and 0.05 decimal degrees) and for different loss thresholds, T (equivalent, for example, to the 20, 50, 100, 250, and 500-year return period losses), we obtain the graphs shown in row A of Figure 3. Row B of Figure 3 shows the distribution of the events that trigger the structure in a magnitude-depth plot for the different loss threshold assumptions, while row C shows the geographic location of the triggering events.

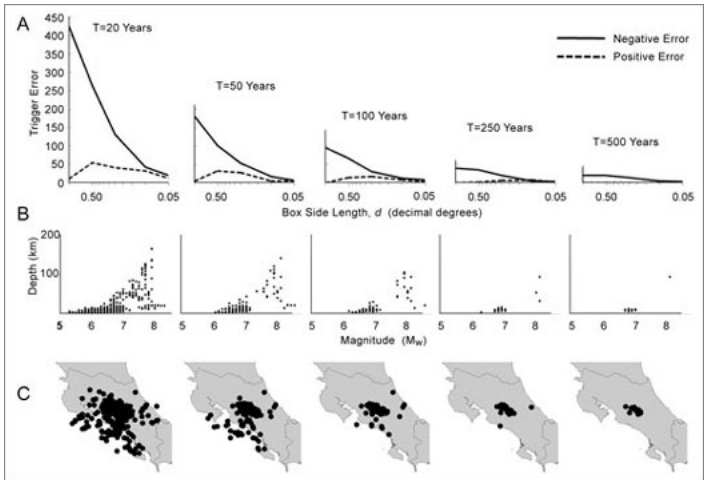

Figure 3. Lower bound error and distribution of triggering events for different loss thresholds and geographic resolutions. (From Franco, 2010; published with permission from EERI/Earthquake Spectra)

What can we learn from this exhibit? First, as was noted before, the total trigger error decreases as the size of the box decreases, regardless of the loss threshold. This is because a decrease in d increases the number of zones. In our example, d=1.0 decimal degrees results in 70 zones (where *K*=*Z*), *d*=0.5 results in 280 zones, *d*=0.25 results in 1,120 zones and *d*=0.1 results in 7,000 zones. The larger number of zones increases the ability to discretize the underlying seismic hazard and better differentiate between earthquakes according to their location. So as the number of zones increases, the precision of the trigger also increases—but of course so does its complexity and computational requirements.

Second, for all loss thresholds we see that the negative error contribution typically dominates the total trigger error for large boxes. This is because bigger boxes will tend to capture many non-triggering events along with the few triggering events, thus increasing the likelihood, by virtue of their sheer numbers, that an event that should not trigger the bond in fact does so. An algorithm programmed to reduce total unbiased error will opt for magnitude and depth thresholds that pay the price of missing some triggering events in exchange for avoiding erroneous outcomes for many non-triggering events. The process of total error minimization therefore introduces negative error, *i.e.*, the possibility of missing events that should trigger the bond.

Note that it is also possible to minimize the difference between positive and negative error; that is, to minimize the bias. This might not lead to a minimum *total* error but it would lead to a trigger that would equally weight the risk of paying for unnecessary protection (positive error) and not getting protection for certain catastrophe events (negative error). The point is that this algorithmic framework allows for flexibility in the design of the cat bond to fulfill other requirements desired by the sponsor.

The graphs from rows B and C provide some clues as to why the negative error bias occurs while searching for the minimum total error. The magnitude-depth plots of row B show that the dispersion of events across the axes is very wide for relatively low loss thresholds. Compare, for example, the plots for T=20 years with that of T=500 years. While in the second case most events are clustered around a small portion of magnitude and depth values, in the first case the events are scattered throughout a variety of magnitude and depth combinations.

Similarly, from row C, the geographic dispersion of triggering events at T=20 is much greater than for T=500. This is due to the fact that smaller losses can ensue from a greater number of combinations of geographic location, magnitude, and depth, while very large losses occur only in locations of very high exposure (as in the San José area) and at relatively large magnitudes and shallow depths. Therefore, the lower the loss threshold and the larger the boxes, the more difficult it is to differentiate the triggering from the non-triggering events based on the physical parameters of location, magnitude and depth.

Since catastrophe bonds typically seek to provide cover for relatively large losses, based on these analyses it is reasonable to conclude that trigger error can be reduced significantly if we use a sufficiently small box. This can be seen, for example, in row A of Figure 3 for the case of T=100, where the error is relatively low for the smaller values of *d*.

However, remember that these solutions have been obtained under the assumption that there are as many zones as boxes—that is, all these solutions represent lower bounds of error at their respective resolution (box size). This is impractical, however, due to the fact that the catastrophe bonds that gain acceptance in the market typically leverage three, four, or five zones—but certainly not thousands, or even dozens.

The question then is: Can we obtain a solution that is as good as the lower bound in terms of minimizing error, but with a much smaller number of zones?

AN OPTIMIZED SOLUTION

Research at AIR has shown that it is indeed possible to attain a level of trigger precision similar to the one shown for the lower bound solutions but using a much smaller number of zones. It is a solution that achieves the best of both worlds: low trigger error and simplicity. It is accomplished through an optimization approach in which the best combinations of magnitude and depth criteria are identified for a specific loss threshold and for a specific resolution while constraining the number of zones.

Figure 4 shows the results obtained by applying this optimization process to our hypothetical cat bond for Costa Rica. The dotted lines show the error of the lower bounds computed for *d*=0.5, 0.25, and 0.1, respectively. As previously noted, the smaller are the boxes, the lower is the error. The (minimum) errors of these lower bound solutions are 80, 47, and 20 (events), respectively. Another consequence of smaller boxes is a higher number of zones (*Z*) for the lower bound solutions (where Z=K), as noted in the graph.

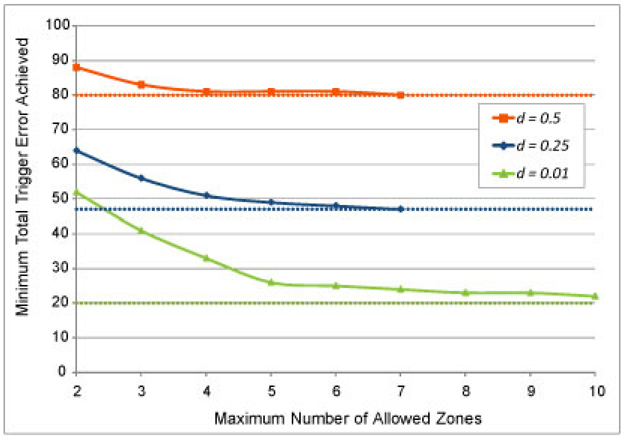

Figure 4. Result of optimization process for different resolutions.

The three solid lines show the errors obtainable using the optimization process, which constrains the number of zones allowed. The red curve shows that the best result we can achieve for *d*=0.5 and Z=2 (we need at least 2 zones to differentiate earthquakes according to location) is an error of 88—just 10% worse (higher) than the lower bound solution of 80. Note that the optimization algorithm was able to find a combination of zones and thresholds at the given resolution (*d*=0.5) that equaled the error of the lower bound solution by using just seven maximum allowable zones. That is a significant improvement over the 280 zones of the lower bound solution.

The same pattern of results is observed for resolutions of d=0.25 and d=0.1. For example, the automated optimization process was able to obtain a solution with just 10 zones that yields very nearly the same error as the lower bound for d=0.01, which uses 7,000 zones.

THE BOTTOM LINE: DESIGN OPTIONS FOR SPONSORS

Designing a successful catastrophe bond entails juggling different requirements: basis risk needs to be minimized, the precision of the trigger needs to be high, and the complexity of the structure needs to be constrained in order to achieve the level of transparency that is demanded by investors. These requirements are often at odds with each other, and may call for solutions that necessitate trade-offs.

Figure 5 shows an example of the array of some of the options available to the sponsors of our hypothetical cat-in-a-box bond for Costa Rica, which is designed to cover earthquake losses equal to or higher than the 100-year return period loss. In Figure 5, resolution (geographic discretization) is increasing from left to right, while the number of zones used increases from top to bottom. Each solution is characterized by a total trigger error (TOT), to which there is a positive error contribution (POS) and a negative error contribution (NEG). and may reintensify.

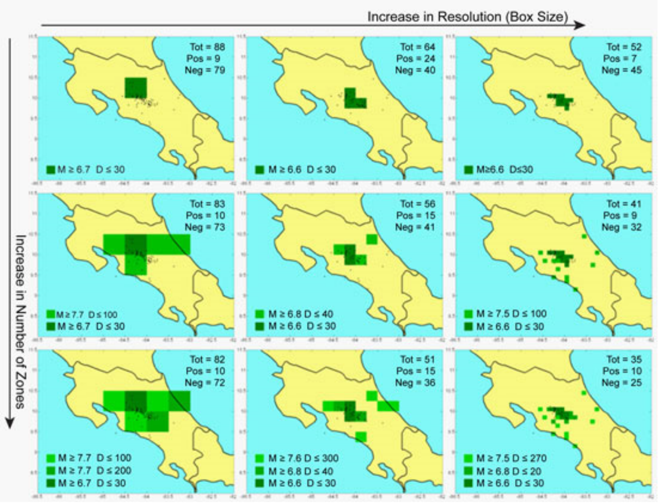

Figure 5. Optimal trigger conditions obtained for increasing number of zones and resolution. The dots depict the events incorrectly triggered (total error); compare to Figure 1 showing all events.

In close collaboration with bond sponsors, AIR is developing sophisticated custom solutions that meet the exacting requirements of stakeholders. That process also entails working with our clients to analyze the sensitivity of bond structures to alternative hypotheses. The optimization algorithms can be made to satisfy different criteria—not only minimum error, as in the example presented here, but also pursuing a balance between risk bias, a specific number of zones, and other criteria that sponsors may demand.

AIR continually strives to offer the most advanced science by way of the most sophisticated catastrophe models to our clients. Now, AIR is also offering the most innovative approaches to the construction of parametric triggers using state-of-the-art optimization techniques, which will ensure that parametric cat bonds of the highest quality will be brought to market.

REFERENCES

FRANCO, G. (2010) “MINIMIZATION OF TRIGGER ERROR IN CAT-IN-A-BOX PARAMETRIC EARTHQUAKE CATASTROPHE BONDS WITH AN APPLICATION TO COSTA RICA” EARTHQUAKE SPECTRA (FORTHCOMING)

ABOUT AIR WORLDWIDE

AIR Worldwide (AIR) is the scientific leader and most respected provider of risk modeling software and consulting services. AIR founded the catastrophe modeling industry in 1987 and today models the risk from natural catastrophes and terrorism in more than 50 countries. More than 400 insurance, reinsurance, financial, corporate and government clients rely on AIR software and services for catastrophe risk management, insurance-linked securities, site-specific seismic engineering analysis, and property replacement cost valuation. AIR is a member of the ISO family of companies and is headquartered in Boston with additional offices in North America, Europe and Asia. For more information, please visit www. air-worldwide.com.