21

ILS Market-Derived Metrics: Implications for Risk Adjustment Transforms and Capital Allocation

Morton Lane; Jerome Kreuser
Lane Financial LLC; RisKontroller Global

The stock of outstanding natural catastrophe insurance-linked securities (ILS) was at an all-time high of nearly US$15 billion as of mid-2012. Nearly half of those ILS were issued in the previous six months and almost all the rest were issued during 2011. All were eligible for trading in the secondary market. A bond issued in 2011 may have a higher or lower coupon than a recently issued new bond. However, the bond's secondary-market yield will be close to that of a new bond on the day it is issued (assuming similar risks). On any particular day, secondary-market ILS yields will vary depending on market perceptions of remaining risk, supply and demand for risk protection in the market as a whole and the set of competitive alternatives.

Taking all the secondary-market yields at a single point in time for each ILS therefore captures the market appetite for risks and compensatory returns that the market collectively demands. It does this for all the risks embedded in the outstanding bonds – whether wind, earthquake or geographic area. Secondary-market prices present a rich trove of embedded information. The purpose of this chapter is to try to extract, ie, make explicit, the risk–return tradeoff that is implicit in the market's pricing. It is no easy task and requires some preparatory reasoning, which we lay out below.

The analysis involves an approach that is suited to insurance or ILS markets that (a) are long-tailed, (b) have expected-loss estimates by scenario and (c) are "incomplete". "Incomplete" is a term used by economists to describe markets where it is not possible to arbitrage away all price discrepancies – perhaps because of an inability to short, or because of a limited discrete set of options or an inability for continuous trading, or simply because of illiquidity. Secondary ILS markets are certainly not liquid. Indeed, quoted prices are often "indications only" and must be observed with caution. Nevertheless we believe useful information can be extracted, provided its use is conditioned by caveats.

In what follows we lay out the rationale for our approach in descriptive terms, lay out the mathematical model in a chapter annex, then apply and illustrate the concept using real data[1] from secondary-market prices on 15/03/2012. Otherwise known as the "Ides of March" in Shakespeare's Julius Caesar, 15/03 may serve as adequate caution for the existence of caveats.

CHOOSING TRANSACTIONS OPTIMALLY

Consider the deals listed in Table 21.1. Six ILS deals are listed together with their issue details and their pricing on 15/03/2012. Assume that an investor is offered these six discrete deals, each in US$25 all-or-none amounts, in the secondary market. Assume the investor has US$100 to invest. What should they do? The answer is fairly obvious. Check that the deals have a positive return and then select those with the biggest expected return. Clearly, each deal has a positive expected return – its secondary yield is greater than expected loss – so the rational investor would rank the deals by expected excess return from best to worst and invest in each until the capital is exhausted. Figure 21.1 shows how the deals would be ranked based on excess expected return.

If the investment question is complicated slightly by saying that capital costs 6%, the same procedure would apply. Rank the deals and invest in every deal with returns over 6%, even if that results in underinvestment of capital.

Presumably, some version of this simple selection procedure – rank-and-choose – is done not just by the single investor, but collectively by the market as a whole, ie, by all its agents and investors. Every deal considered here is quoted in the secondary market, thus the market contains deals with high excess returns down to the deal that is most marginal.

Table 21.1 Deal Statistics for a Sample Set of Six ILS

Sample Set of ILS	Embarcadero Re Ltd. series 2012-I Class A	Kibou Ltd. Series 2012-1 Class A	Multicat Mexico 2009 Ltd Class D	Queen Street V Ltd.	Res Re 2010 Clas 4	Vega Capital Ltd. Series 2010-I Class C
At Issue						
Par Value - Issue Size	150,000,000	300,000,000	50,000,000	75,000,000	117,500,000	150,000,000
Issue Spread	7.25%	5.25%	10.25%	8.50%	13.00%	5.65%
Maturity Date	2/13/2015	2/16/2015	10/19/2012	4/9/2015	6/6/2013	12/30/2013
Original PPM Expected Loss	1.96%	0.95%	2.39%	1.77%	2.42%	1.30%
3/15/2012						
Remaining Term (months)	34	34	7	36	14	21
Secondary Market Price	100.22	99.94	100.21	100.12	101.33	101.18
Outstanding Market Value	$150,322,500	$299,820,000	$50,105,000	$75,090,000	$119,062,750	$151,770,000
Secondary Market Yield	7.17%	5.26%	9.79%	8.45%	11.81%	4.96%
Remodeled* Expected Loss	2.21%	0.85%	2.93%	2.42%	2.00%	2.86%

* As remodeled by AIR

Figure 21.1 Six sample deals with expected loss and excess return

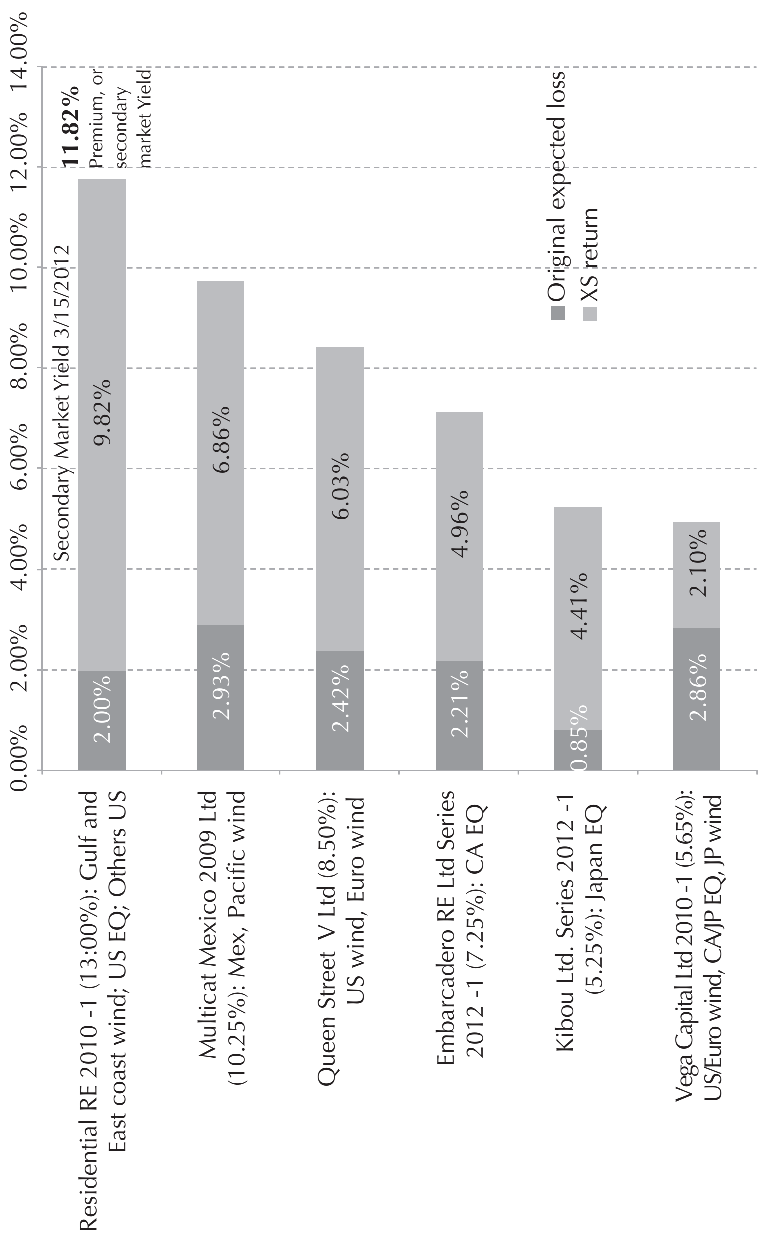

CORRELATIONS AND RISK ADJUSTMENT

There is one other feature that individual investors and the market as a whole will add to the simple selection process: allowing for concentrating or diversifying risks. In other words, the simple selection process will be adjusted if the risk implicit in any particular deal compounds with risks in another deal. When there is a concentration or compounding of risks, it may pay the investor to take a deal with a lesser expected excess return if it can reduce the combined risks of the portfolio. In short, the market will tend to re-rank the investments on a "risk-adjusted" basis and make the selection from the re-ranked list of "risk-adjusted excess returns".

In the equity world, where mean-variance portfolio analysis is pursued, the amount of risk adjustment is determined by quantifying the variance of the deal and its covariance with other deals. There is an agreed-on quantifiable approach. The insurance underwriting world, however, is not blessed with symmetric return distributions, and mean-variance analysis is not justified. There is no generally accepted formula for risk adjustment. Instead, as we have seen in preceding chapters, insurance underwriters (or ILS investors) seek to identify dependencies through "scenario" analysis. Thus, risk modellers will identify a scenario as a set of catastrophic events in any year and will evaluate what the effect of those events is on each of the deals or underwritings under consideration. Cleary, that process can lead to the identification of benign or bad scenarios for a given portfolio; it can conversely allow the investor to choose a portfolio of investments that minimise the consequences.

We will make the assumption that, when we view the outstanding set of ILS priced in the secondary market, the market, collectively through the actions of individual investors, has gone through that selection process. In other words, it has conducted a "risk adjustment" to the known features of a deal that takes account of embedded correlations or diversifications. By this implicit process it prices the relative attractiveness of particular ILS.

If the market for ILS was "complete" – fully arbitrageable and infinitely tradable – we could anticipate that secondary prices would equal risk-adjusted expected losses. Such prices are known as Arrow–Debreu[2] prices (Arrow and Debreu 1954). However, such a perfectly competitive market does not exist in the ILS world. Prices contain excess expected returns prior to any risk adjustment. The prices will probably also contain an excess return, even after risk adjustment. In other words, even after risk adjustment or secondary pricing, some deals will be better than others; some deals will be marginal. Notwithstanding this reality, we assume that the market seeks to set secondary prices as close as possible towards their risk-adjusted expected loss.

The foregoing is a rationalisation for the process we will use here to extract the market's risk adjustment. It is in essence to minimise the (market-weighted-average) distance between the secondary market price of the portfolio and the risk-adjusted expected loss of the portfolio.

There is another way to look at the process we pursue. The expected loss listed in each deal's PPM[3] is usually calculated by a risk-modelling firm generating 10,000 simulated scenarios covering the risk space. In the expected-loss calculation, each of these 10,000 scenarios is considered to be equally likely, ie, with a probability of 0.0001. Clearly, the market views some of these scenarios as more troublesome than others, or else why would they not price all the deals with the same expected loss at the same price? In the mind of "Mr Market" some scenarios are given higher weight than others. What we are seeking to do is find out what those market-weighted probabilities are. Those implied probabilities are what generate implicit risk-adjusted expected losses and differentiate one deal from another. We try to discover what implied probability distribution best explains observed secondary-market prices.

SETTING UP THE ILLUSTRATIVE EXAMPLE

As already observed, it is typical to evaluate deals over a minimum of 10,000 simulated natural catastrophe scenarios.[4] For the purposes of illustration here we have collapsed that set to seven scenarios. This is done by defining seven tranches or buckets from the ranked outcomes of the market portfolio risk profile. Step 1 is to evaluate the market portfolio for each scenario. Step 2 is to rank those outcomes from best market outcome to worst market outcome. Finally, these ranked scenarios are grouped into the seven tranches, as illustrated in Table 21.2.

Table 21.2 Ranking of scenarios for the market and tranching into buckets

				Market-value-weighted ILS portfolio*		**Embarcadero Re Ltd. Series 2012-I Class A**	**Kibou Ltd. Series 2012-1 Class A**	**Queen Street V Ltd.**
		No. of scenarios	tranche probability	Expected loss by tranche		EL In tranche	EL In tranche	EL In tranche
First tranche	B-T-W 1 – 5,000	5,000	0.500	0.0000	……	0.0000	0.0000	0.0000
Second tranche	B-T-W 5,001 – 8,000	3,000	0.300	0.0073	……	0.0028	0.0001	0.0000
Third tranche	B-T-W 8,001 – 9,000	1,000	0.100	0.0269	……	0.0419	0.0399	0.0060
Fourth tranche	B-T-W 9,001 – 9,500	500	0.050	0.0643	……	0.1297	0.0674	0.1180
Fifth tranche	B-T-W 9,501 – 9,900	400	0.040	0.1944	……	0.2274	0.0225	0.2650
Sixth tranche	B-T-W 9,901 – 9,960	60	0.006	0.3739	……	0.2167	0.0333	0.5333
Seventh tranche	B-T-W 9,961 – 10,000	40	0.004	0.5174	……	0.0500	0.0000	0.9750
		10,000	1					
			A priori expected loss**	2.02%		2.21%	0.85%	2.42%
			Secondary market yield***	8.37%		7.17%	5.26%	8.45%
			Excess return	6.35%		4.96%	4.41%	6.03%

* 92 of the outstanding ILS market portfolio are chosen – certain deals are eliminated to avoid distortions of maturity or impairment.
** A priori expected loss in this table is calculated as the sum product of the expected loss in each tranche multiplied by the probability of each tranche.
*** The portfolio yield is the market-value-weighted yield of all the individual deals.

We note that the aggregation of scenarios into tranches is for ease of exposition. The chapter appendix shows the mathematics for the more general case over all scenarios.

The seven collapsed scenarios are chosen with deliberate purpose. There is more and more concentration on the worst outcomes. The best 5,000 of outcomes form the first tranche. The worst 40 scenarios form the seventh tranche. The division points between tranches are the following value-at-risk (VaR) points: 50%, 80%, 90%, 95%, 99%, 99.6% and 100%. Obviously, on an *a priori* basis, the tranches have different probabilities, as listed in Table 21.2.

For each tranche the multiple scenario losses for the whole portfolio and for each deal are averaged to give expected losses[5] within the tranche. This is an important point to emphasise. The expected loss in any particular tranche is only for the scenarios that are included in that portfolio-defined tranche. These are effectively co-measures and they are instructive in themselves.

Consider Kibou in Table 21.2. Kibou's expected loss in the seventh tranche is 0.0. The seventh tranche is the market portfolio's worst 40 scenarios. Its probability of occurrence is 0.004, or 1 in 250 years. In that tranche the portfolio is expected to lose 51.74% of its value, but Kibou does not contribute to that bad outcome. On the other hand Queen Street V loses 97.5% of its value and is therefore a big contributor to those bad market portfolio outcomes. If the market is focused on losses in the market tail it will risk-adjust Queen Street V Ltd's prices much more than Kibou Ltd's. As already indicated, the mechanism for risk adjustment is changing the probability weights. If the probability weight for the seventh tranche is increased beyond 0.004 then any deals with heavy seventh tranche losses will be have a big risk adjustment. Queen Street will be disadvantaged, Kibou will not. More properly, the risk-adjustment of the expected loss of Queen Street will be greater than that of Kibou.

FINDING THE MARKET RISK ADJUSTMENT

The problem is set up as a linear programming problem as follows. Minimise the difference between the secondary-market price of the market portfolio and the risk-adjusted expected loss of the market portfolio. The variables over which the minimisation must take place are the implied probabilities of each tranche. Collectively the probabilities must add to 1; they must form a distribution.

Table 21.3 Optimised solution probabilities – the market transform

				Market-value-weighted ILS portfolio*		**Embarcadero Re Ltd. Series 2012-I Class A**	**Kibou Ltd. Series 2012-1 Class A**	**Queen Street V Ltd.**
		Individual scenarios probability	tranche probability	Expected loss by tranche		EL In tranche	EL In tranche	EL In tranche
First tranche	B-T-W 1 – 5,000	0.000094	0.4706	0.0000	......	0.0000	0.0000	0.0000
Second tranche	B-T-W 5,001 – 8,000	0.000094	0.2823	0.0073	......	0.0028	0.0001	0.0000
Third tranche	B-T-W 8,001 – 9,000	0.000094	0.0941	0.0269	......	0.0419	0.0399	0.0060
Fourth tranche	B-T-W 9,001 – 9,500	0.000094	0.0471	0.0643	......	0.1297	0.0674	0.1180
Fifth tranche	B-T-W 9,501 – 9,900	0.000094	0.0376	0.1944	......	0.2274	0.0225	0.2650
Sixth tranche	B-T-W 9,901 – 9,960	0.000465	0.0279	0.3739	......	0.2167	0.0333	0.5333
Seventh tranche	B-T-W 9,961 – 10,000	0.001010	0.0404	0.5174	......	0.0500	0.0000	0.9750
			1					
		Risk-adjusted expected loss **		4.62%		2.74%	0.87%	7.03%
		Secondary market yield***		8.37%		7.17%	5.26%	8.45%
		Excess risk-adjusted return		3.74%		4.42%	4.38%	1.42%

* 92 of the outstanding ILS market portfolio are chosen – certain deals are eliminated to avoid distortions of maturity or impairment.
** *A priori* expected loss in this table is calculated as the sum product of the expected loss in each tranche multiplied by the probability of each tranche.
*** The portfolio yield is the market-value-weighted yield of all the individual deals.

The problem is further constrained so that all risk adjustments are positive, or at least non-negative. Finally, we have required that the individual scenario probabilities be monotonically rising from the best scenarios to the worst.[6]

The results of this exercise are shown in Table 21.3. The derived probability for the first tranche is 0.4706, which is less than the *a priori* probability of 0.5000. The market is putting less weight on the benign scenarios of the first tranche. In contrast, the market derived probability in the seventh tranche is raised to 0.0404 compared with 0.004 on an *a priori* basis. This is a tenfold increase.

The combined effect of this new probability distribution of the portfolio is a risk-adjusted expected loss of 4.62% versus an *a priori* expected loss of 2.02%. In other words, the process has narrowed the distance to the secondary prices, but not eliminated it completely. There is still some excess risk-adjusted return. Remember also that under ideal, complete market, circumstances the distance or the excess would be zero. The fact that it cannot be reduced beyond the 3.74% shown in Table 21.3 is in some way a measure of the market's incompleteness.

The effect on individual ILS is also shown in the table for the three deals we have chosen to highlight. The *a priori* expected loss of Embarcadero was 2.21% (see Table 21.2). The risk-adjusted expected loss is 2.74%. Similarly for Kibou and Queen Street, the numbers are 0.85% and 2.42%, respectively, before adjustment and 0.87% and 7.03% after adjustment. The effect of risk adjustment on individual deals is dramatic. On an *a priori* basis Queen Street was very attractive in this subset of three. It had the highest excess return (6.03%, shown in Table 21.2) and would by that measure be the best deal of the three. That whole calculation changes after risk adjustment. Now the excess risk-adjusted return for Queen Street drops to 1.42% (see Table 21.3), the worst of the three deals. We had anticipated something like this. Recall that in the worst 40 scenarios for the whole portfolio Queen Street lost 97.5% of its value. That loss is 10 times more problematic with the transformed probabilities, hence the need to adjust expected loss.

TRANSFORMS, DISTORTING DISTRIBUTIONS AND THE MARKET

There is whole literature on distorting original probabilities so that bad outcomes incur higher probabilities and benign ones incur lesser probabilities. Indeed there is healthy debate about which distorting "transform" best explains prices and should be used in pricing exercises. The most popular is the Wang Transform (Wang, 2004), but competing distributions include the T Distribution, the Esscher Transform and many others (Bühlmann *et al* 1998; McNeil, Rüdiger and Embrechts 2005). While these distributions have nice analytical properties and can be summarised by elegant algebra, none dominates the others as the "correct" one. Wang has used his transform to explain prices and Venter has tried to empirically show which is the best fit to a set of prices, but the reason you might choose one form over another remains elusive (Wang 2004).

We are also fitting to prices, but without specifying a closed-form[7] transforming distribution. We require that the implied probabilities provide the best probabilistic explanation of pricing, assuming that the market is optimising as a whole. Table 21.3 gives the transformed distribution for the sample data numerically. It can also be illustrated diagrammatically. Figure 21.2 shows the details as a density function. Actually, it shows the last 10% of the density. The other 90% would show a continuation of the horizontal lines.[8]

Figure 21.2 Probabilities distributions of portfolio outcomes

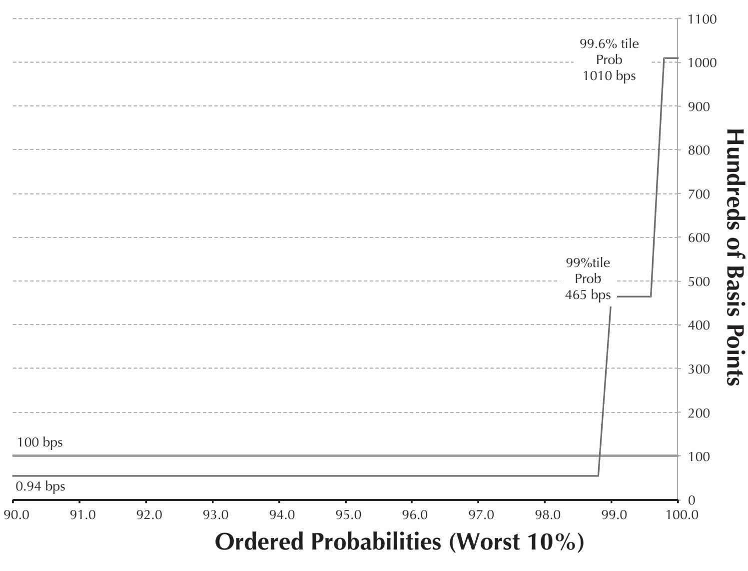

Note: Probabilities shown as 100* basis points - density distributions. The figure is not to scale and is for illustrative purposes only.

It is our strong conviction that extracting transforms from market actions endorses the idea of distorting probabilities. It does, however, set that search in a different direction. It is a search that is enabled by scenario analysis and may be usefully adopted in any market where such scenario analysis is possible.

IMPLICATIONS FOR THE WHOLE PORTFOLIO OF DEALS

Figure 21.3 Components of premium on 15/03 ILS portfolio

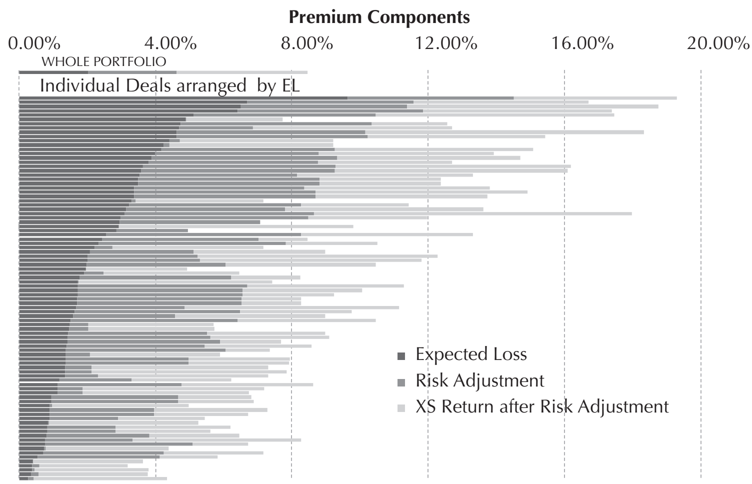

Note: Ordered by expected loss, largest to smallest

Figure 21.3 provides a picture of the effect of the "market transform" that has been extracted from the prices as of 15/03/2012. The top bar in the graph is the portfolio as a whole, its components identical to Table 21.3. Below that, deals are arranged by expected loss, from highest to lowest. Each bar represents a deal, but they are not separately identified. Instead, the six sample deals that we have already examined are isolated in Figure 21.4.

The bar associated with each deal shows the secondary-market yield as of the March date. Each bar is then broken into three components: (i) the *a priori* expected loss, (ii) the incremental expected loss that has been derived, ie, the risk adjustment, and (iii) the remaining part of the yield, which is labelled excess risk-adjusted return. Figure 21.3 provides an overview of the market; Figure 21.4 shows some of its detail.

Figure 21.4 Six sample deals with expected loss and risk-adjusted excess return

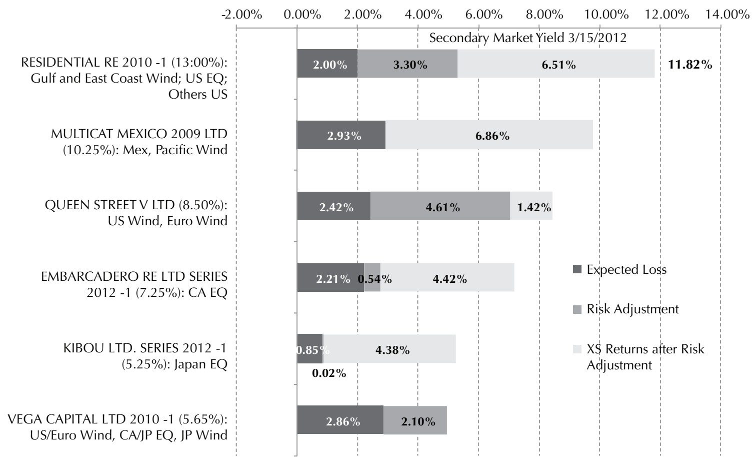

Note: Ordered Highest XS Return to Lowest. Original Spread Noted on Left Axis; Secondary Yield =EL+XS

A number of observations are in order. The first is that the amount of risk adjustment is considerable overall. As already observed, the market as a whole (ie, the market-weighted ILS portfolio) averages a risk-adjusted 4.62%. Effectively, this doubles the prior expected loss of 2.02%. Even with this, however, only about 55% of prices are explained by expected loss or risk-adjusted expected loss.

The second observation is that the risk adjustment varies enormously from deal to deal. Clearly, some deals incur tiny risk adjustments – Kibou from our sample set is one of them. Queen Street V, on the other hand, has a risk adjustment that wipes out most of its seeming excess return.

Notice also that certain deals have no excess return after risk adjustment, Vega Capital Ltd 2010-1 being one in our sample. Another deal not identified here, EOS Wind Ltd, also has no excess return, and these can be viewed as being strictly "marginal" deals. They are fully priced.

Kibou is not the only deal to have a very small risk adjustment: Multicat Mexico (Class A) in our sample also has a tiny adjustment. Although not separately identified, the Pylon series, Calypso and GreenFields would be similar. Anyone familiar with these deals will recognise them as "diversifying" deals: they do not contribute to peak concentrations of US wind and US quake. This feature can be further illustrated by sorting the deals into Perils, Geographies and Structures. This is done in Table 21.4 and Figure 21.5 in the next section.

SUBDIVIDING BY PERILS, GEOGRAPHIES AND STRUCTURES

Table 21.4 shows how much of the market value of the deals outstanding as of 15/03/2012 covered US wind, how much covered US quake and so forth. It also shows the deals that covered these risks in single-peril format (can be exhausted only by a single peril, eg, US wind events) and in multiple-peril format (eg, can be exhausted by both US wind and US quake events). Multiple-peril structures constitute some 52% of the outstanding ILS. Because they cover either/or events their combined potential limit is greater than the issued limits. In fact, the table shows that potential exhaustible limits by peril constitute 162% of stated actual limits. It is natural to consider that an ILS covering more than one peril will not be seen as diversifying; rather it will be concentrating.

Grouping the deals in the format of Table 21.4 and illustrating them graphically in Figure 21.5 reveals much about the risk adjustment process.

By far and away the biggest risk adjustment is needed for the multi-peril deals, the biggest block of which is deals involving combined US wind and US quake exposures. Clearly they require considerable adjustment. Similarly, the deals involving European wind and US wind have big risk adjustments. The biggest risk adjustments, however, are in transactions that involve worldwide coverages: US, European, Japanese wind and US and Japanese quake. There are not many of these, so the risk adjustment is not so striking graphically. Nevertheless, it is in this group that Vega is found. It is keenly priced to the point of having no excess return after adjustment.

In contrast with these multi-peril deals the single-peril adjustments are more modest. To be sure, some single-peril US wind deals require adjustments; it is after all a peak zone. However the amount of adjustment is quite variable.

Table 21.4 Coverage by type of deal or peril as of 3/15

ILS 3-15-2012										% of
									Total by	Total
Single Peril		**USW**	**USQ**	**EUW**	**EUQ**	**JPW**	**JPQ**	**OTH**	**Type**	**Limit**
	USW	2788							2788	21%
	USQ		1109						1109	8%
	EUW			1763					1763	13%
	EUQ				65				65	0%
	JPW					280			280	2%
	JPQ						560		560	4%
	OTH							290	290	2%
Multi-Peril	52%									
	USW, USQ	5764	5764						11528	88%
	EUW, JPQ			210			210		420	3%
	USW, EUW	435		435					870	7%
	USW, EUW, USQ, JPQ	140	140	140			140		560	4%
	USW, EUW, JPW, USQ, JPQ	124	124	124		124	124		620	5%
	USW, EUW, USQ	150	150	150					450	3%
	Exhaustible Limit	9401	7287	2822	65	404	1034	290	**21303**	
	% Exposure by Peril	71%	55%	21%	0%	3%	8%	2%		162%
Total Limit		**13153**								162%

Figure 21.5 Components of premium, 15/03/2012 ILS portfolio

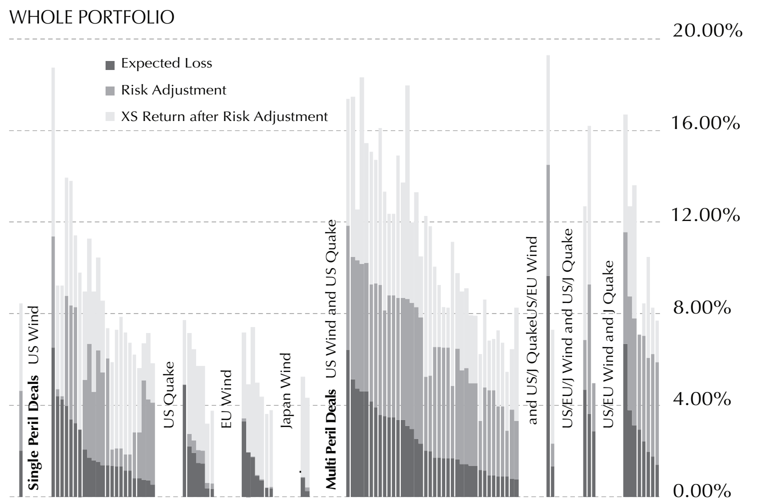

Finally, it is clear that single-peril deals involving European wind, US quake or Japanese wind are rarely heavily risk-adjusted by the market. Kibou, which we previously focused on, is an example of this.

Notwithstanding the relative adjustments and the fact that a few deals are "marginal", the other feature that is quite striking is that most deals involve healthy margins of excess risk-adjusted returns. It averages 3.74%. Perhaps there is a level of return that the market requires whether or not there are adjustments for diversifying and concentrating features. Perhaps there is an implicit cost of capital, a threshold, that the market must exceed before it sets prices. This is not explored here but the example opens up several potentially fruitful avenues for research.

CONCLUDING REMARKS

Warren Buffett famously replied, when asked how the academic community might regard his investment approach, "Well, it may be alright in practice, but it will never work in theory." He probably could have made the same response if he had been asked about how he underwrote the insurance risks that Berkshire Hathaway assumed. We have tried to keep the spirit of those answers in mind when examining the

ILS market. How are ILS prices set in the market? What causes one bond to be priced at twice its expected loss and another to be priced at, say, six times the same expected loss. There have been, and will continue to be, massive explorations of that subject. Some of these, including ours, have involved regression analyses of ILS price histories; others have been empirical analyses of private transactions from the traditional market. Some have involved using "transforms"; other have used fits to risk measures such as standard deviations. Few, however, have sought to unlock pricing via the implied probability distributions embedded in observed secondary-market prices. This is almost certainly true of the ILS market; we cannot vouch for other markets.

We have used a practical example (the data of 15/03/2012) to illustrate our approach. Thus we know it works "in practice". Lest we be subject of the Buffet jest, we have also laid out a theory in the chapter appendix listed below. The fact that we have model and data actually gives some comfort in the approach. However, we do not think the analysis is complete by any means. There are implications in the model and the data for capital allocations and for performance measurement within a portfolio. Those we leave for future research.

Nor do we suppose this approach is without qualification. As already noted, the prices we have used may be suspect; so may the valuations of deals. What is presented here is also just a single data point at the time of a relatively hard market. But market conditions change. The risk adjustment made in a hard market may be quite different from the one made in a soft market. Over time, however, the shifts in probability distributions could be extraordinarily important. This is the point to conclude on. Every insurance company has its own underwriting methodologies for evaluating the transactions it is presented with. They form part of the culture and *raison d'être* of a company. The question those underwriters are continually confronted with is, Is this consistent with the way the rest of the market will evaluate the deal? We believe that the approaches herein will facilitate an easier answer to that question.

APPENDIX

A probability transform from the ILS market – the discrete case

We give the following definitions:

M is the total market of Insurance Linked Securities;
p_i is the secondary-market yield associated with the ith deal;
$l_{i,j}$ is the fractional loss for deal i in scenario j provided by the modelling firm;
r_j is the probability of scenario j as defined by the risk modelling firm usually 1 in 10,000;
$D_{i,j}$ is cashflow of deal i in scenario j;
w_i is the market share for ideal I;
L_j is the market portfolio loss for scenario j.

In this section we provide the mathematical framework for the results obtained in the paper. It follows on the work of Kreuser and Lane (2006 and 2007).

Under these definitions, the cashflow fraction in scenario j is given by $D_{i,j} = p_i - l_{i,j}$ and the total expected return of the portfolio defined by the w_i is

$$\sum_i Mw_i \left(p_i - \sum_j r_j l_{i,j} \right) = \sum_i Mw_i \sum_j r_j D_{i,j}$$

We assume that the market, through the individual agents acting together, adjusts the prices and market values, p_i, along with a maximum allowable loss limit (or risk appetite of the market) of L_j for each scenario such that the chosen share w_i is optimal. Specifically, the market obtains an allocation solving the following problem:

$$\underset{w_i}{MAX} \quad \sum_i Mw_i \left(p_i - \sum_j r_j l_{i,j} \right)$$

subject to

$$\begin{aligned} &\sum_i Mw_i \left(l_{i,j} - p_i \right) \le L_j \quad : \pi_j \\ &\sum_i Mw_i = M \quad : \omega \\ &Mw_i \ge 0 \quad : \mu_i \end{aligned} \tag{21.1}$$

The variables following the colon (:) are the dual variables corresponding to each constraint.

The value Mw_i is the amount of the market value of deal i. We can drop the M from the objective function and divide through on the constraints so that the problem becomes:

$$\underset{w_i}{MAX} \quad \sum_i w_i \left(p_i - \sum_j r_j l_{i,j} \right)$$

subject to

$$\begin{aligned} &\sum_i w_i \left(l_{i,j} - p_i \right) \leq \frac{L_j}{M} \quad : \pi_j \\ &\sum_i w_i = 1 \quad : \omega \\ &w_i \geq 0 \quad : \mu_i \end{aligned} \tag{21.2}$$

We note that $\frac{L_j}{M}$ is the share of the market portfolio loss in scenario j.

This problem has a dual because the solution to equation 21.2 is feasible and bounded. It is:

$$\underset{\pi_j, \omega}{MIN} \quad \sum_j \pi_j \frac{L_j}{M} + \omega$$

subject to

$$\begin{aligned} &p_i - \sum_j r_j l_{i,j} = \sum_j \pi_j \left(l_{i,j} - p_i \right) - \mu_i + \omega \quad : w_i \\ &\pi_j \geq 0 \quad \omega \geq 0 \quad \mu_i \geq 0 \end{aligned} \tag{21.3}$$

The goal here is to solve for π_j. Since we know the "optimal" solution from equation 21.2, we can use $\frac{L_j}{M} = \sum_i w_i \left(l_{i,j} - p_i \right)$ and obtain the optimal by solving equation 21.3.
Then we have:

$$\begin{aligned} &\sum_j \left(p_i - l_{i,j} \right) \left(r_j + \pi_j \right) = \sum_j D_{i,j} \left(r_j + \pi_j \right) = \omega - \mu_i \\ &\pi_j \geq 0 \quad \omega \geq 0 \quad \mu_i \geq 0 \end{aligned} \tag{21.4}$$

Let $R = \sum_j \pi_j$ and $\lambda_j = \frac{r_j + \pi_j}{1+R}$. The λ_j are then probabilities and sum to one. We note that when solving equation 21.2 we may just as well have $\sum_i w_i \leq 1$. The reason is that we assume that $p_i - \sum_j r_j l_{i,j} > 0$. Therefore all $w_i > 0$ and the inequality will be satisfied as an equality and we will have $\mu_i = 0$ and we have $\omega \geq 0$. So we have:

$$\frac{\omega}{1+R} = \sum_j D_{i,j} \lambda_j \quad \forall i \tag{21.5}$$

This is essentially the problem we solved in this paper but we used tranches of scenarios instead of every possible scenario to make the presentation simpler.

The minimum $\omega = 0$. The λ_j are the Arrow-Debreu risk-neutral probabilities or the risk-adjusted probabilities and $\frac{\omega}{1+R}$ are the risk-adjusted prices (Arrow and Debreu 1954). In this formulation we should have the prices equal to zero for an arbitrage free argument.

We can think of equation 21.5 as the result of minimising the risk-adjusted probabilities to closely match the risk-adjusted prices while also minimising a weighted sum of the π_j.

The analogy to the continuous case and Girsanov's Theorem (Avellaneda and Laurence 2002) is that the probability transform here converts the process to a martingale when $\omega = 0$. This process therefore gets to as "close" to a martingale as possible and $\frac{\omega}{1+R}$ can be considered a measure of the "incompleteness" of the sub-market of Insurance Linked Securities.

Further, since:

$$\frac{\omega}{1+R} = \sum_j D_{i,j} \frac{\lambda_j}{r_j} r_j \quad \forall i \tag{21.6}$$

We have that $\frac{\lambda_j}{r_j}$ are what are called the state-price deflators and these weights reflect investor's preference towards the different states.

If we expand (21.5), we get,

$$p_i - \frac{\omega}{1+R} = \sum_j l_{i,j} \frac{r_j + \pi_j}{1+R} = \sum_j l_{i,j} \lambda_j \tag{21.7}$$

So that the risk-adjusted ILS prices are $p_i - \frac{\omega}{1+R}$ for equation 21.1.

This may seem odd at first. However, the solution to the problem made assumptions about the market shares w_i all being non-zero. This was the case for the market solution but it may not be the case in general when we use a risk-preference constraint.

If we replace the loss-limit constraint in equation 21.2 by a risk-preference constraint, as we did in Kreuser and Lane, 2007, we obtain the formulation used in the body of that paper. These risk-preference constraints were limits on the expected losses in the tail of the distribution.

In that case we then get a solution where some $w_i = 0$ and subsequently the $\mu_i > 0$. The same analysis holds for more than one tail constraint.

We omitted the cost of capital, but that is also easy to include in the objective function and obtain a similar formulation.

1 There were 110 ILS outstanding as of 15/03/2012. Several of these were eliminated if they were either impaired or under three months to maturity. This avoids including ILS where the original EL may have shifted after issue estimate. The scenarios used are courtesy of AIR Worldwide. Similarly, we have used mid-market prices from the AON price sheet. We are grateful to each company, but in both cases any errors that may have been generated are our sole responsibility.

2 The Arrow–Debreu model and state-prices are discussed in the chapter appendix.

3 PPM is the private placement memorandum, also known as the offering circular, accompanying each deal.

4 Both AIR and RMS use this number, although each will generate different numbers on demand and there are many more thousands of event behind the scenario generation. Individual investors may evaluate deals over even greater number of scenarios – 40,000 or 100,000.

5 The expected loss within the portfolio tranches can also be thought of as the differences between the tail value at risk (TVaR) values at the incepting and exiting part of each tranche.

6 Actually, this is imposed within the tranches rather than on tranche probabilities. Thus, the individual probabilities of the 5,000 scenarios in the first tranche must be equal and they must be less than the individual 3,000 equal probabilities in the second tranche, etc. While this is intuitive, we believe it is related to the overlapping and therefore monotonic nature of the TVaR expression of tranche expected losses. The chapter appendix describes a model that is more general and does not require these monotonicity constraints.

7 The "form" is derived from the linear programme.

8 The horizontal axis is the scenario number, ranked B-T-W, it is not numerical loss. Since each scenario is equally likely, the distribution of scenarios is a uniform distribution, ie, the horizontal line.

REFERENCES

Arrow, Kenneth, and Gerard Debreu, 1954, “Existence of an Equilibrium for a Competitive Economy”, *Econometrica* 22(3), July, pp. 265–90.

Bühlmann, Hans, Freddy Delbaen, Paul Embrechts and Albert N. Shiryaev, 1998, “On Esscher Transforms in Discrete Finance Models”, ASTIN Bulletin, Vol. 28, No. 2, pp. 171-186.

Avellaneda, Marco, and Peter Laurence, 2000, *Quantitative Modelling of Derivative Securities* (Boca Raton, FL: Chapman & Hall/CRC).

Kreuser, Jerome, and Morton Lane, 2006, “An Introduction to the Benefits of Optimization Models for Underwriting Portfolio Selection”, available at http://www.lanefinancialllc.com, *Proceedings of the 28th International Congress of Actuaries*, Paris, June.

Kreuser, Jerome, and Morton Lane, 2007, “Optimal Insurance and Reinsurance Portfolios, Implied Pricing, Allocating Retrocessional Costs and Capital Allocation”, available at http://www.lanefinaciallc.com, presented at the “Adventures in Risk” conference in Christchurch, New Zealand, September.

McNeil, Alexander J., Rüdiger Frey and Paul Embrechts, 2005, *Quantitative Risk Management: Concepts, Techniques, and Tools* (Princeton, NJ, and Oxford, UK, Princeton University Press, Princeton, New Jersey).

Wang, Shaun S., 2002, “A Universal Framework for Pricing Financial and Insurance Risks”, *ASTIN Bulletin* 32(2), pp. 213–34.

Wang, Shaun S., 2004, “Cat Bond Pricing Using Probability Transforms”, Geneva Papers: Etudes et Dossiers on “Insurance and the State of the Art in Cat Bond Pricing” (278), pp. 19–29, January