# The Limits of Securitization
*Micro-correlations, Fat Tails and Tail Dependence*
Carolyn Kousky[1] and Roger M. Cooke[2]
January 6, 2010

## Introduction

Aggregation is a key feature of traditional risk management strategies. Financial institutions securitize assets and property insurers hold bundles of policies across many lines of business.  With independent assets and risks with thin tails, the law of large numbers (LLN) and the central limit theorem (CLT) assure risk benefits from aggregation: the standard deviation of the sum increases slower than the mean of the sum and the tails of the aggregation become normal, that is, thin-tailed.  These benefits are the foundation of much of the banking and insurance industry.  Recent events, from Hurricane Katrina in 2005 to the financial crisis of 2008, however, demonstrate that assumptions of independence and thin tails are no longer warranted in a world that is, in the words of Thomas Friedman, "hot, flat, and crowded."  Both globalization and climate change are linking previously independent systems and are altering the frequency and/or the magnitude of extreme events.

The twin influences of globalization and climate change have introduced an "unholy trinity" (Kousky and Cooke 2009) of pitfalls for risk managers relying on aggregation: micro-correlations, fat tails, and tail dependence.  The LLN requires independent aggregands.  Very small global correlations, micro-correlations, are enough to undo the benefits from aggregating independent assets.  The CLT also requires that the aggregands have a finite variance.  Just one aggregand having a very fat tail, such that the variance is infinite, is enough to make the distribution of the sum fat-tailed, as well.  In these cases, aggregation does not ease the risk of catastrophes.  The sample mean, the statistic we often use to predict future impacts, also will have infinite variance, no matter the size of the sample.  Our prognosticator is infinitely wiggly; importing George Burns' analogy, predicting the future is like shooting pool with a rope.  Even though any sample will be finite, suggesting that infinite variance could simply be a theoretical artifact, thick-tailed behavior in finite samples is observable and troubling for risk managers.  Finally, tail dependence is perhaps best described by Shakespeare: "When sorrows come, they come … in battalions."[3]  Failure to consider potential tail dependence by using the normal copulae, has been blamed for Wall Street's woes (Salmon 2009).  The normal copula's popularity derives from the ubiquity of the normal distribution.  This model is "tail independent" (McNiel et al 2005). That is, extremes of variables joined by a normal copula are independent.  This fact has spawned a growth industry of exotic copulae to model tail dependence.  A very recent discovery, however, shows that tail dependence is anything but exotic; indeed it emerges when events are aggregated with the simplest of dependence models using the normal copula (Cooke et al 2010). Tail dependence may be rather the rule than the exception as risky assets are aggregated.

---

[1] Resources for the Future, kousky@rff.org

[2] Resources for the Future and Dept. Mathematics, Delft University of Technology, cooke@rff.org

[3] King Claudius, Hamlet Prince of Denmark, Act 4 Scene 5

This chapter illustrates this unholy trinity with two loss datasets: flood insurance claims data from the United States National Flood Insurance Program (NFIP) and crop insurance indemnities paid data from the United States Department of Agriculture's Risk Management Agency.[4] Both datasets are aggregated by county and year for the years 1980 to 2008. The data are in constant year 2000 dollars. Over this time period there has been substantial growth in exposure to flood risk, particularly in coastal counties. To remove the effect of growing exposure, we divide the claims per county per year by personal income per county per year available from the Bureau of Economic Accounts (BEA)[5]. Thus we study yearly flood claims per dollar income, per year per county. The crop loss claims are not exposure adjusted, as an obvious proxy for exposure is not at hand, and exposure growth was less of a concern. While the unholy trinity may become a greater problem due to globalization and climate change, we do not attempt to extract any temporal trends in these phenomena here, but simply illustrate methods for measuring and detecting our trinity in finite datasets.[6] When data is not available we must turn to experts. A final section offers a perspective on the role of experts in dealing with these new problems.

## Micro-correlations

Micro-correlations are correlations between variables at or beneath the detection limit. These tiny correlations are amplified by aggregation, limiting common diversification strategies. The ballooning under aggregation is illustrated by a very simple formula that should be on the first page of every insurance text book, but isn't. Let $X_1,...X_N$ and $Y_1, ...Y_N$ be two sets of random variables with the same average variance $\sigma^2$ and average covariance C, both among and between these sets. The correlation of the sums of the X's and the sum of the Y's is easily found to be:

$$\rho(\Sigma X_i, \Sigma Y_i) = \frac{N^2 C}{N\sigma^2 + N(N-1)C}. \tag{1}$$

---

[4] We would like to thank Ed Pasterick, Tim Scoville, and Barbara Carter for providing us with this data.

[5] This data goes from 1980 to 2006, so the last two years are dropped from the flood claims data. Income data was not avilable for Guam, Puerto Rico, or St. Croix, so these are dropped from our dataset. Further, the income data for some counties in Virigina was for aggregations of counties. These are also dropped as they cannot match cleanly with our flood claims data.

[6] Previous research has tried to examine changes in loss distributions over time and determine whether such changes can be attributable to climate change. While a time series analysis might seem to be a natural approach to examining our insurance data, the number of years of data in combination with the extreme variability renders this unworkable. For example, if we consider the flood claims data for the entire US, then the correlation with year over the observation period (0.31) is not statistically significant, and is strongly driven by the year 2005 where hundredfold higher claims were registered as a result of Katrina (without exposure adjustment the correlation with year would be significant). More importantly, our goal is not to show that climate is changing and is influencing crop and flood damages. Rather, the goal is to show how this may impact our risk management strategies.

[7] The empirical correlation's limiting distribution is normal with mean $\rho$ and variance $(1-\rho^2)^2/N$, where $\rho$ is the bivariate correlation.

This evidently goes to 1 as N grows, *if* C is non-zero and $\sigma^2$ is finite. If all variables are independent, then C = 0, and the correlation in (1) is zero. The variance of $\Sigma X_i$ is always non-negative; if the $\sigma^2$ and C are constant for sufficiently large N, it is easy to see that $C \geq 0$.

Filling in a few representative numbers, if N = 1000, $\sigma^2 = 1$ and C = .01, then $\rho(\Sigma X_i, \Sigma Y_i)$ = 0.91. If each variance were actually one, then C is the average correlation of any two variables. Based on the sampling distribution of the normal correlation coefficient, a sample size of 27,000 would be required to distinguish a sample correlation of 0.01 statistically from zero at the 5% level[7]. A very small global correlation can make a very big difference when we aggregate, and this small correlation is effectively undetectable at the level of individual variables. This phenomenon was first observed in modeling auto NOx emissions in The Netherlands (van Oorschot et al 2003).

The amplification of correlation can be seen in the flood insurance claim data. Suppose we randomly draw pairs of counties in the US and compute the correlation in their flood claims. Figure 1 shows the histogram of 500 such correlations. The average correlation is 0.04. A few counties have quite high correlations, but the bulk is around zero. Indeed, based on the sampling distribution for the normal correlation coefficient, correlations less than 0.37 in absolute value would not be statistically distinguishable from zero at the 5% significance level. 91% of the correlations fall into that category.
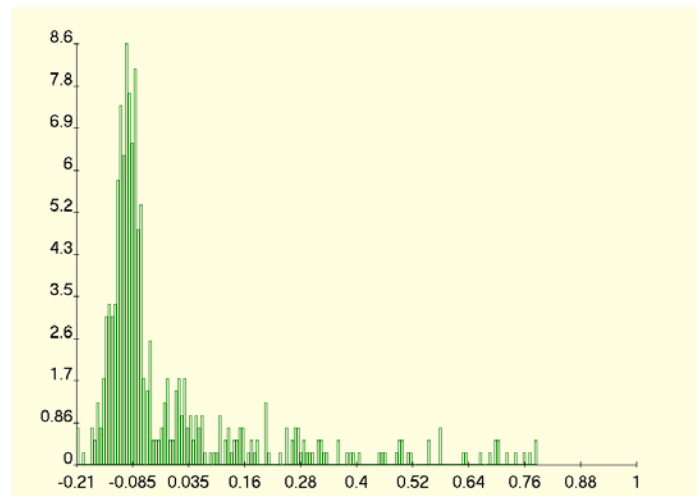


**Figure 1: Histogram of 500 correlations of US flood claims of random pairs of counties**

Instead of looking at the correlations between two randomly chosen counties, consider summing 100 randomly chosen counties, and correlating this with the sum of 100 distinct randomly chosen counties. If we repeat this 500 times, the leftmost histogram in Figure 2 results; the average of 500 such correlations-of-100 is 0.23. The rightmost histogram depicts 500 correlations-of-500, their average value is 0.71. This dramatic increase in correlation is a result of the micro-correlations between the individual variables. Compare Figure 2 with Figure 3, in which each county is assigned an independent (i.e. micro correlations are switched off) uniform

---

[7] The empirical correlation's limiting distribution is normal with mean $\rho$ and variance $(1-\rho^2)^2/N$, where $\rho$ is the bivariate correlation.

variable, for each of 29 years. The correlations-of-1 and correlations-of-500 are effectively the same.
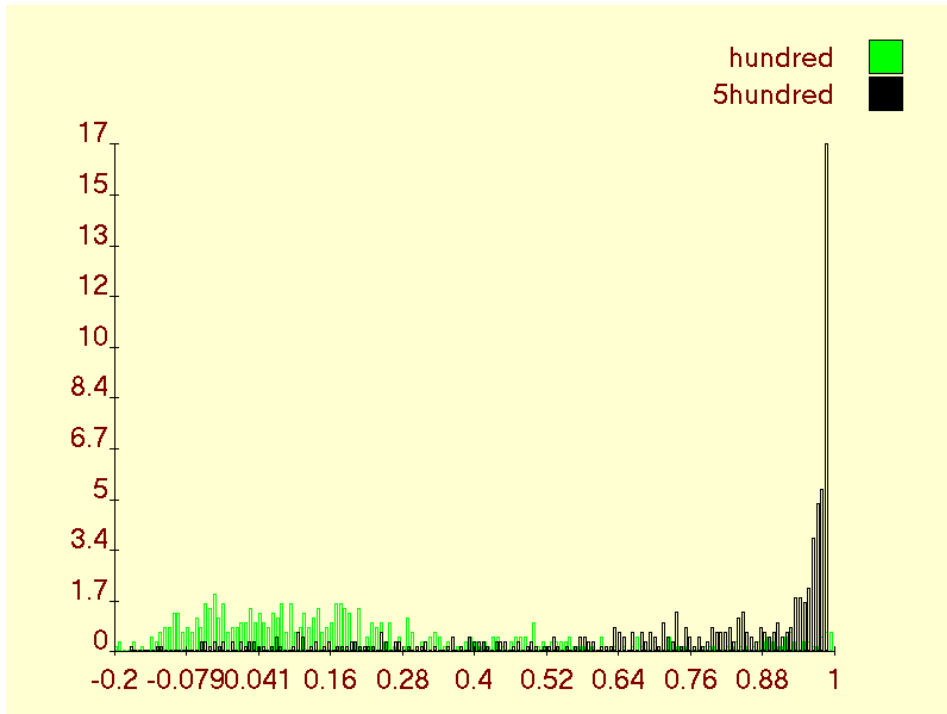


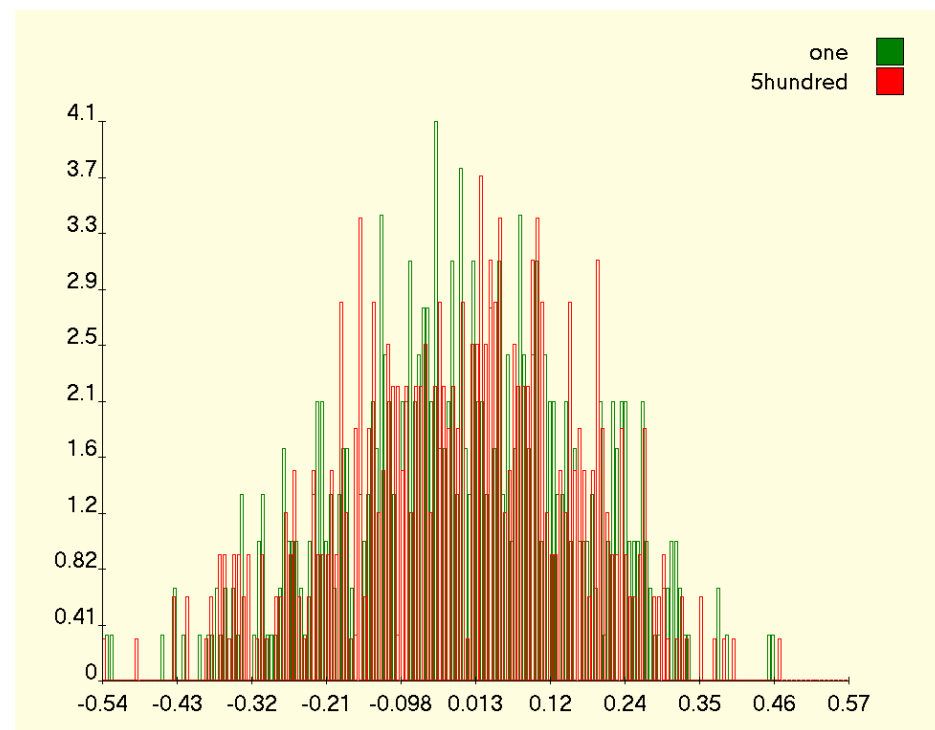**Figure 2: US Flood claims, correlations-of-1 (light) and correlations-of-500 (dark).**



**Figure 3: Correlations-of-1 and correlations-of-500, for 30 realizations of independent uniform variables**

The difficulty with micro-correlations is that they could so easily go undetected. One might not readily assume that fires in Australia and floods in California are correlated, but El Nino events induce exactly this coupling.  Bankers may not have assumed that mortgage default rates around the country were correlated, but the correlation of default rates to general economic conditions creates this micro-correlation.

Banks, insurance companies, mutual funds are all risk dissipaters. Suppose you buy a bond for $1,000 that pays $1,100 in one year, unless it defaults, in which case you get nothing. If the probability of default is 0.01, your expectation is $\mu$= $89, but the standard deviation of your return is $\sigma$ = $109.4 . Your "comfort ratio" $\mu/\sigma$ tells you how by many standard units your mean exceeds zero.  $89/$109.4 = 0.81 is not a good comfort ratio. The CLT shows us that if you buy N identical *independent* bonds, your comfort ratio becomes $\sqrt{N}\mu/\sigma$.  The more bonds you buy, the more comfortable you become. After 100 purchases, your mean is already 8 standard units from zero. Now suppose that all these bonds have a statistically insignificant, but positive correlation $\rho$; then your comfort ratio does *not* keep going up but levels off to $\mu/(\sigma\sqrt{\rho})$[8].  That may not be so bad, *if* you know it…but if you overestimate your comfort, then you might just do something foolish, like leveraging yourself at 30 to 1.

**Fat Tails**
Fat tails were introduced in mathematical finance in 1963 by Benoit Mandelbrot to describe cotton price changes (Mandelbrot 2004).  Outisde of finance, there is evidence that damages from many disasters follow distributions with fat tails (e.g., Malamud and Turcotte 2006). Climate change may be further fattening the tails of the distributions of many extreme events (e.g., U.S. Climate Change Science Program 2008). The uncertainty surrounding climate change may also generate fat tails, as in Weitzman's (2008) analysis, where updating a non-informative prior yields a fat-tailed posterior damage distribution.
The precise mathematical definition of tail obesity is rather subtle (Resnick 2007), but a working notion is that damage variable *X* has a fat tail if, for sufficiently large values *x*, the probability that *X* exceeds *x* is $kx^{-\alpha}$, for some constants $a,k > 0$.   It is easy to see that the *m-th* moment is infinite if $m \geq \alpha$. If $\alpha \leq 1$, the tail is "Super Fat" and the mean or first moment is infinite. Of course, on *N* samples from such a distribution, the average of the *N* sample values will be finite, but it increases with N. "Really fat" tails have infinite variance. The sample mean also has infinite variance, no matter how many samples we draw. This is of particular importance for operational risk, see L. Dalla Valle (2010) in Chapter ???  in this volume.
A good way to gauge tail obesity is by "mean excess" plots (McNiel, et al. 2005). At each damage level x, we consider all damages greater than x, and plot the average by which these exceed x. For fat tailed distributions, the mean excess plot is increasing. If the tail is Really fat, than the slope is greater than one, the tail is Meso Fat if the slope is less than one, but greater than zero. We do not have enough data to study tail behavior within each county. We therefore pool all counties and consider each county-year as a realization of a single loss variable, yielding about 75,000 loss events for crop losses and for flood claims. Figure 4 shows mean excess plots for crop (left) and exposure adjusted flood insurance (right) claims. The unit slope line has been added for reference.

---

[8] For N independent bonds, the comfort ratio is $N\mu/(\sqrt{N}\sigma) = \sqrt{N}\mu/\sigma$.  With constant micro correlation $\rho$ the variance of the sum is $\sqrt{(N\sigma^2+N(N-1)\rho\sigma^2)}$, and the comfort ratio is $\sqrt{N}\mu/(\sigma\sqrt{(1+(N-1)\rho)}) \sim \mu/(\sigma\sqrt{\rho})$.

*Updated version appearing in* Klaus Boecker (ed) Re-Thinking Risk Measurement and Reporting, Uncertainty, Bayesian Analysis and
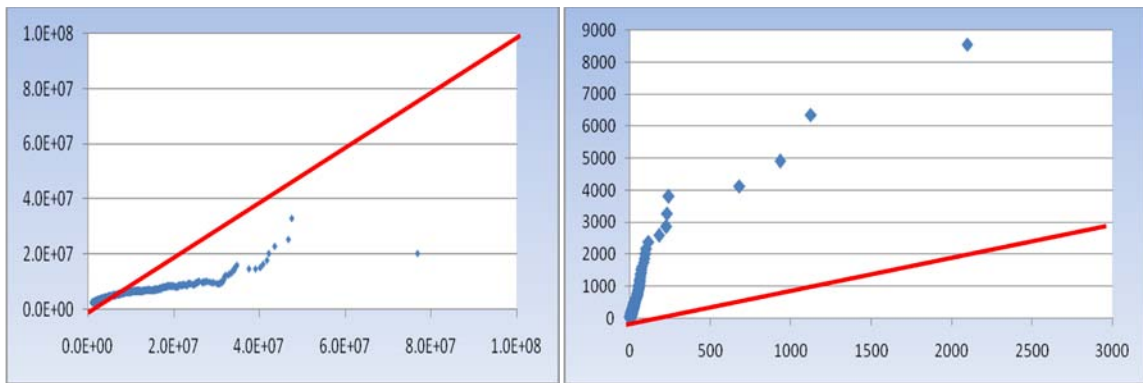Expert Judgement, Risk Books



**Figure 4: Mean excess plots for US crop loss (left) and exposure corrected US flood claims (right).
The vertical axis gives mean excess loss, given loss at least as large as the horizontal axis.**

       The mean excess plot for crop losses shows a slope less than 1, whereas the mean excess plot for exposure-adjusted flood claims has a slope greater than one, indicative of an infinite variance. Whether the data has infinite variance is critical for risk managers considering aggregation. If the variance is finite, the aggregate distribution will have thinner tails; if the variance is infinite, the tail of the aggregation stays fat, which is simply a consequence of the general CLT. This difference can be seen in the crop and flood data. Suppose we form packages of crop losses by randomly grouping the county-years ten at a time. Since the variance exists, the sum will approach a normal (thin-tailed) distribution. This is reflected in the flattened mean excess plot of a random aggregation-by-ten counties of crop losses in Figure 5 (left). Compare this to Figure 5 (right) showing the mean excess of a random aggregation-by-fifty of exposure adjusted flood losses. It has the same shape as the right hand picture in Figure 4. Aggregation does not thin the tail significantly. Fat tailed loss data with finite variance converges to a normal distributions as we form random aggregations, but fat tailed distributions with an infinite variance do not. Instead, they converge to stable laws with the same tail behavior as the original aggregands, see e.g. Embrechts, Klüppelberg & Mikosch (1997), Theorem 2.2.15, for more details. When insurance companies hold really fat tailed policies, the aggregated portfolio is also really fat, and is poorly estimated by historical averages. Clearly, the same holds true for operational risk.
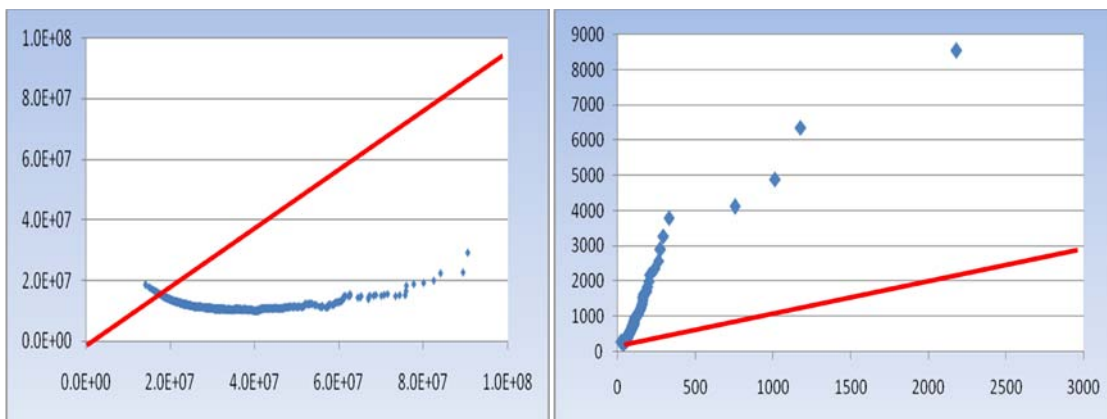


**Figure 5: Mean excess crop loss aggregation by 10 (left) and mean excess flood claims aggregation by 50 (right).**

Since Mandelbrot's 1963 study of cotton prices, we know that fat tails require different methods than thin tails; the sample mean from a Really Fat tailed loss distribution has infinite variance and aggregations of independent fat tailed risks do not become thin tailed. Still, we continue to rely on historical averages, even when we should better. For instance, we know flood claims are really fat tailed but the US National Flood Insurance Program (NFIP) still uses an "historical average loss year" in setting premiums, and, as if that weren't enough, the program only gives a 1% weight to 2005 (with hurricanes Katrina, Rita and Wilma) in "an attempt to reflect the events of 2005 without allowing them to overwhelm the pre-Katrina experience"[9]. NFIP is itself overwhelmed with the claims from 2005, and on its own admission is unlikely to be able to cover the interest on its debt to the treasury.

**Tail Dependence**
Tail dependence is the most worrisome and least well understood member of our unholy trinity. It refers to the tendency of dependence between two random variables to concentrate in the extreme high values (known as upper tail dependence, UTD) or extreme low values (lower tail dependence, LTD). For loss distributions, we are interested in the UTD of non-negative variables. Technically, upper tail dependence of variables $X$ and $Y$ is defined as the limit (if it exists) of the probability that $X$ exceeds its *100 r-percentile*, $x_r$, given that $Y$ exceeds its *100r-percentile,* $y_r$, as $r$ goes to *1*:

$$UTD(X, Y) = \lim_{r \to 1} P(X > x_r \mid Y > y_r). \qquad (2)$$

If $X$ and $Y$ are independent, their tail dependence is zero. If their tail dependence is positive, then when one variable takes on an extreme value, it is more likely the other variable will as well. Note that UTD does not depend on the distributions of $X$ and $Y$; if we apply any monotone transformation to $X$, say $X^* = X^{1/N}$ (which will thin $X$'s tail)*,* then *UTD(X\*,Y) = UTD(X,Y)*. UTD has no simple relation to the standard Pearson correlation coefficient used in (1). For example, normal variables with any correlation $\rho$ strictly between -1 and 1, have zero tail dependence. (McNiel, et al. 2005).

Tail dependence can be seen in loss data. Wind damage and water damage are insured separately in the United States. The former is covered under homeowners policies or state wind pools and the latter is covered by the federal National Flood Insurance Program (NFIP). Flood damage and wind damage are often independent; a rising river does not necessarily mean high winds and a storm with high winds may not have enough rain to cause flood damage. A severe hurricane, however, causes both. This suggests that wind and water insurance payments may be tail dependent in a hurricane-prone state such as Florida. Figure 6 shows this is indeed the case. Wind payments from the state insurer Citizens Property Insurance Corporation were grouped by county and month for the years 2002–2006, as were NFIP flood claims (all are in constant 2007 dollars). Each damage dataset was ranked, with the highest rank standardized to 1, and the ranks plotted against each other. The abundance of points in the upper right quadrant of Figure 6 shows that high flood damages and high wind claims occur together, indicating tail dependence.

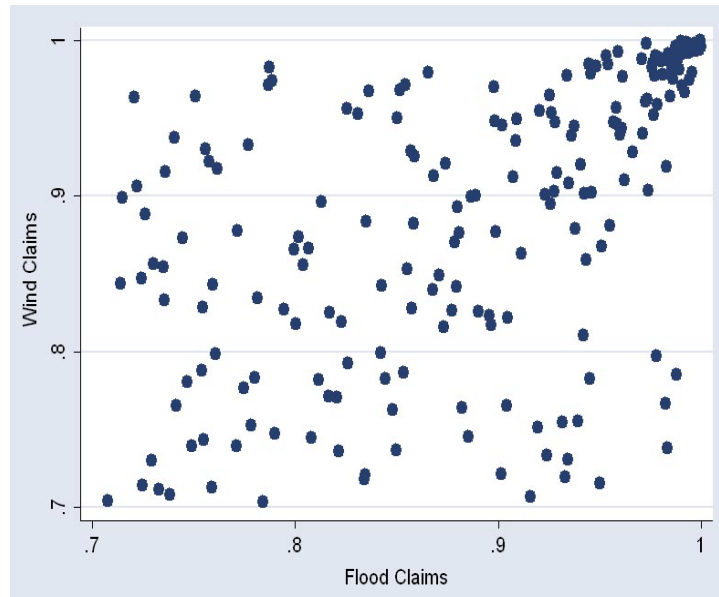---

[9] Hayes and Spafford 2008  p.6

**Figure 6 Percentile scatter plot for Florida wind and flood damages, above the 70 percentile**

Correlations between sums always grow with aggregation, if the average covariance is positive. Under certain conditions, tail dependence can also grow, further foiling diversification. We can see this in data. We examine monthly flood loss data from Florida (the state for which we have monthly observations), per county. If we consider two random groups of five different counties, and make a scatter plot of the percentiles of their monthly losses, the left plot of Figure 7 emerges. The points along the axes correspond to months in which no losses were reported in these counties. We may discern a weak tendency for points to cluster in the upper right corner. This tendency grows appreciably stronger if we take two random groups of 30 different counties, as in the right plot.



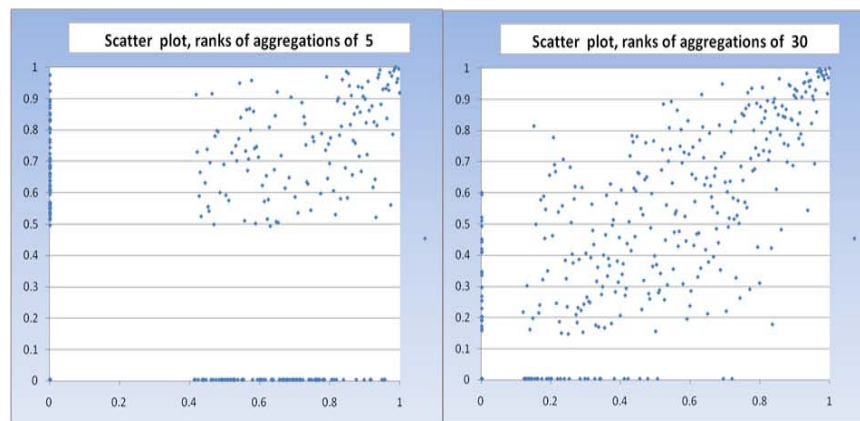**Figure 7: Percentile scatter plots of random aggregations of 5 (left) and 30 (right) Florida counties, monthly flood losses.**

Tail dependence is related to the correlation of extreme event occurrences. For any two events *A* and *B*, the indicator functions $1_A$, $1_B$, are functions which take the value 1 if *A* or *B* occurs respectively, and zero otherwise. The Pearson correlation of the indicator functions is:

*Updated version appearing in* Klaus Boecker (ed) Re-Thinking Risk Measurement and Reporting, Uncertainty, Bayesian Analysis and Expert Judgement, Risk Books

$$\rho(1_A, 1_B) = \frac{P(A \ and \ B) - P(A)P(B)}{\left[P(A)\left(1-P(A)\right)P(B)\left(1-P(B)\right)\right]^{1/2}} \ . \tag{3}$$

Setting *A = {X is above its 100r-percentile}* and *B = {Y is above its 100r-percentile}*, then *P(A) = P(B) = (1-r)* and:

$$\rho(1_A, 1_B) = \frac{P(A \ and \ B)}{r(1-r)} - \frac{1-r}{r} \rightarrow UTD(X,Y) \ as \ r \rightarrow 1. \tag{4}$$

Recently, some progress has been made in understanding conditions under which tail dependence is amplified by aggregation. The gamma density with shape $\upsilon$ and scale $\alpha$ is $f(x) = \alpha^{\upsilon} x^{\upsilon-1} e^{-\alpha x} / \Gamma(\upsilon)$, $x \geq 0$, with mean $\upsilon/\alpha$ and variance $\upsilon/\alpha^2$. $(X_1,..,X_n,…)$ have an $L_1$ symmetric distribution with gamma distributed rate if, for any n, the n-dimensional marginal density is given by:

$$p(x_1,…x_n) = \int (\Pi_{i=1..n} \lambda e^{-\lambda \Sigma x_i}) \lambda^{\upsilon-1} a^{\upsilon} e^{-a\lambda} (1/\Gamma(\upsilon)) \ d\lambda. \tag{5}$$

The tail dependence of two distinct sums of N $L_1$ variables is found to be (Kousky and Cooke, 2009a):

$$\frac{\sum_{k,j=0…N-1} \left(\frac{1}{2}\right)^{\upsilon+j+k} \Gamma(\upsilon+k+j)/[\Gamma(\upsilon)k!j!]}{\sum_{k=0…N-1} \Gamma(\upsilon+k)/[k!\Gamma(\upsilon)]} \ . \tag{6}$$

It has recently been proved that this converges to 1 as $N \rightarrow \infty$ (Cooke et al 2010). This result explains the emergence of tail dependence in terms of familiar random variables. The $L_1$ model has found extensive application in reservoir management and deterioration of civil engineered structures (van Noortwijk 1996, 2009). The dependence model (5) has a simple graphical structure, and is known as a latent variable model (see Figure 8): There is one variable, in this case $\lambda$, which couples all variables $(X_1,…)$ and conditional on the latent variable, all other variables are independent. This is perhaps the simplest of a wide class of complex dependence models.

**Figure 8: Latent variable dependence model**
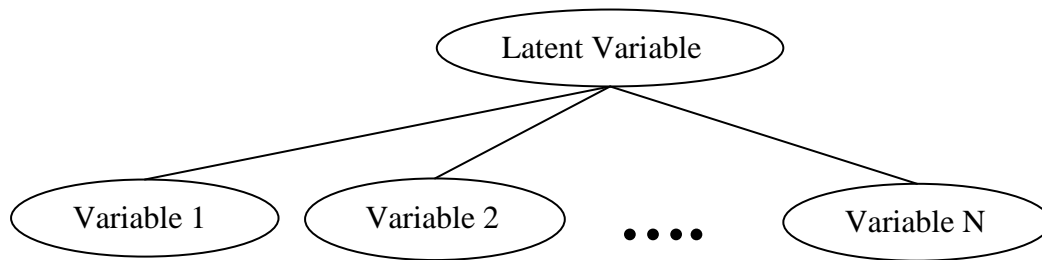


Another recent result explains how sums of events may become tail dependent. Suppose event $Y_i$ occurs when some physical variable $U_i$ exceeds some threshold, r. We may take $U_i$ to be uniformly distributed on the interval [0, 1]. For i = 1,…2N, the variables $U_i$ may be dependent. The simple latent variable model posits a latent (uniform) variable V and assigns a bivariate copula to each $(U_i, V)$. Consider two loss portfolios $S_1(N) = \sum_{i=1..N}Y_i$ and $S_2(N) = \sum_{i=N+1..2N}Y_i$ (the dependence of $S_i(N)$ on r is suppressed in the notation). We are interested in conditions under which $S_1$ and $S_2$ become upper tail dependent as $N \to \infty$. You may think of the $S_i$ as claim amounts of any insurance companies, total loss distributions for operational risk, or the total loss of two bank's credit risky portfolios.

Assume that $P(U_i \leq u \mid V = v)$ is increasing in v, for all u; one says that the (U,V) copula is stochastically increasing. The following has recently been shown (Cooke et al 2010); $P(S_1(N) > \delta N \mid S_2(N) > \delta N) \to 1$ as $N \to \infty$ for any 0 < r < 1, and any 0 < δ < 1, if and only if

$$P(U \leq u \mid V = 1) = 0 \text{ for all } 0 < u < 1. \tag{7}$$

This condition is strictly weaker than UTD(U,V) = 1. The familiar tail dependent copulae, such as the Gumbel and Galambos, satisfy (7). For Archimedean copulae, (7) is equivalent to ordinary upper tail dependence (Theorem 3.12 of Joe (1997)). Significantly, (7) also holds for the normal copula, which is *not* upper tail dependent. If indeed the normal copula is the most frequently occurring "in nature", and if the simple latent variable is indicative of behavior in a wider class of dependence models, then we may expect tail dependence to be the rule. In this case, different banks' credit portfolios, insurance claim portfolios etc may be highly dependent on each other, although, the individual constituents of the portfolios (claims, loans) only moderately depend on the common latent factor. This analysis is a telling argument in favour of a banking supervisions that is not narrowly focused on single banks but rather takes interdependence between different financial players seriously into account. As the title of this chapter makes clear, the benefit of securitization is limited and economic value may turn into economic loss.

Such asymptotic results do not tell the whole story. For finite sums, we can observe tail dependent-like behavior even if the copula in a latent variable model does not satisfy (7). To illustrate, we use the latent variable with $U_i$ and V joined by Frank's copula $(U_i,V)$, event $Y_i = $

**Resources for the Future**                                                          **Kousky and Cooke**
*Updated version appearing in* Klaus Boecker (ed) Re-Thinking Risk Measurement and Reporting, Uncertainty, Bayesian Analysis and
Expert Judgement, Risk Books

$U_i > 0.9$, $P(Y_i) = 0.1$, $\rho(U_i,V) = 0.9$, and $\rho(Y_1,Y_2) = 0.36$. Frank's copula is an Archimedean copula, which does not satisfy (7), and thus has zero tail dependence. Letting $S_1(N) = \Sigma_{i=1...N} Y_i$ and $S_2(N) = \Sigma_{i=N+1...2N} Y_i$, we find that $\rho(S_1(100), S_2(100)) = 0.98$. (All calculations are in done in Maple). Since $S_i(N)$ has a discrete distribution, the standard definition of tail dependence is not useful. The following graphs show two types of tail dependent behavior. Figure 9 (left) looks at $P\{S_1(N) > 0.7N \mid S_2(N) > 0.7N\}$ as N = 5….500, and compares this with $P\{S_1(N) > 0.7N\}$. We see that while the probability of exceeding 70% of N decreases in N, in this range, the conditional probability of such exceedence stays constant. This means that although the diversification within each portfolio increases with growing number of atoms, the conditional loss probability remains constant and diversification goes awry. Even though this is an artificial example, the conclusions drawn have great similarities with what happens in the financial industry during 2007/2008. Put another way, similar mechanisms as described above may create the notorious systemic risk that may lead to a collapse of the entire financial system . The saw-tooth aspect of these curves is caused by the discreteness of $S_i(N)$. Figure 9 (right) compares $P\{S_1(N) > i\}$ and $P\{S_1(N) > i \mid S_2(N) > i \}$, as function of i, for N = 100. The curious non-monotonic feature of the conditional probability is observed in many simulations. There are two forces at work; as i increases $P\{S_1(N) > i\}$ decreases, and as i increases, conditioning on $S_2(N) > i$ pushes the conditional distribution of V given $S_2(N) > i$ toward one, which raises the probability of $S_1(N) > i$. The conclusion is that tail dependent-like behavior may arise for finite sums even in situations not covered by (7).

**Figure 9 Tail dependent –like behavior for latent variable with Frank's copula, conditional and unconditional fractional occurrence rates (left); conditional and unconditional exceedence rates (right).**
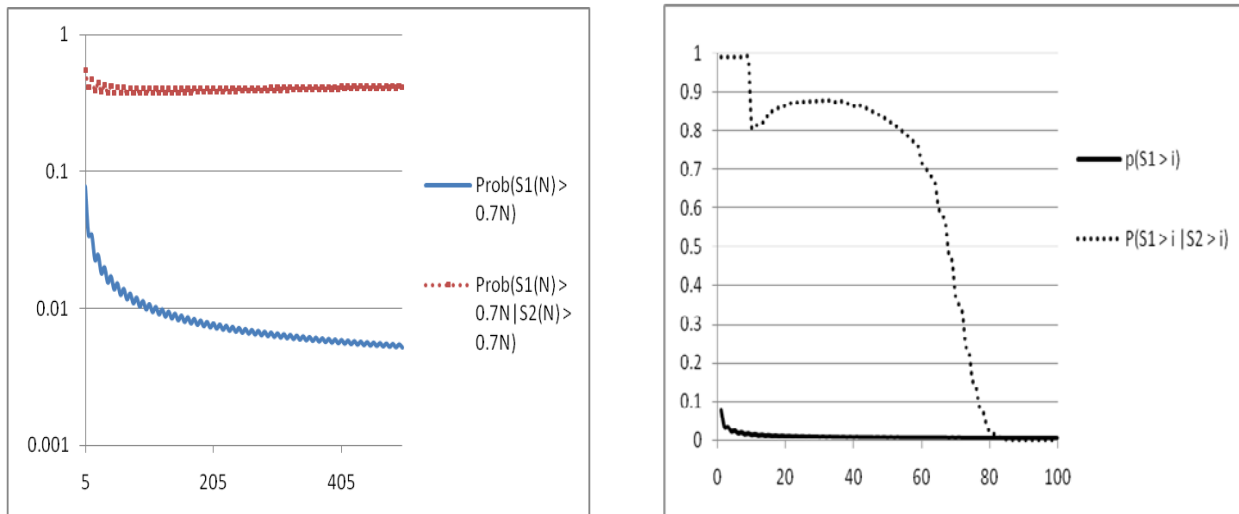


Tail dependence is not just theoretical. After Hurricane Katrina, heretofore independent lines of insurance saw many large claims; among them property, cargo, inland marine and recreational watercraft, floating casinos, onshore energy, automobile, worker's compensation, health, and life insurance (RMS 2005), pointing to tail dependence across these lines of business. Similar tail dependence, specifically between property and motor vehicle claims, was also observed in France (Lescourret and Robert 2006). Such tail dependence has the potential to wipe out insurance companies if not properly considered and addressed.

**The Role of Experts**

It is evident that dependence is important in modeling risks. Whenever possible, models should be based on dependence observed in the actual data. If we know how to look, we can uncover these relationships. The preceding sections suggest that random aggregations provide a useful tool in detecting and measuring our unholy trinity. There are many cases, however, when data can be difficult to come by and yet we still must make risk management decisions. In these cases, we may have no alternative but to plumb the opinions of experts.

Structured expert judgment has been used for uncertainty analysis for more than 35 years (US Nuclear Regulatory Commission,Wash-1400, 1975) and quantitative dependence elicitation for at least 20 (Cooke and Goossens, 2000). These efforts usually assume a given copula type, usually normal or maximum entropy, and use experts to obtain bivariate rank correlations, something much simpler than trying to obtain estimates on our unholy trinity.

Extensive experience with bivariate dependence elicitation was obtained in the Joint USNRC-EU uncertainty analysis of accidents with nuclear power plants[10]. Prior to this study, the question of probabilistic dependence was largely ignored, as if all important dependences were captured in functional relationships. Since this study, however, the technique for bivariate dependence elicitation has been deployed many times and returns usable results with expert buy-in.

The joint study had to break new ground in dependence modeling and dependence elicitation, and these subjects are treated extensively in the reports. The format for eliciting dependence was to ask about joint exceedence probabilities: "Suppose the effectiveness of supportive treatment for radiation exposure in people over 40 was observed to be below the median value, what is your probability that also the effectiveness of supportive treatment in people under 40 would be below its median value?" Experts became quickly familiar with this format. Dependent bivariate distributions were found by taking the minimally informative copula which reproduced these exceedence probabilities, and linking these together in a Markov tree structure. Further developments, generalizing both the choice of copula and the tree dependence structure are found in Kurowicka and Cooke (2006). Recently, Morales et al. (2007) extended these techniques to more complex dependence structures.

The use of structured expert judgment to quantify the persons of the unholy trinity poses special problems. The processes we are examining, such as globalization and climate change, are influenced by both natural and socio-economic processes, raising a key question for expert judgment: who are the experts? An early expert elicitation study by Nordhaus (1994) undertaken to understand the probability of various levels of damage from climate change, found that natural scientists were distinctly more pessimistic about the tail of the climate change damage distribution than economists. For a doubling of $CO_2$ by the mid twenty-first century, natural scientists estimated the probability ranging from 0.2 to 0.95 that global incomes would fall by more than 25%. Economists estimated the same probability to be between 0.003 to 0.09. With

---

[10] Reports may be obtained from
http://dutiosc.twi.tudelft.nl/~risk/index.php?option=com_docman&task=cat_view&gid=89&&Itemid=13

**Resources for the Future**                                                                                **Kousky and Cooke**
*Updated version appearing in* Klaus Boecker (ed) Re-Thinking Risk Measurement and Reporting, Uncertainty, Bayesian Analysis and
Expert Judgement, Risk Books

such a striking difference, the question of who the experts are in considering complex processes becomes central. Averaging across experts is unsatisfying when the estimates are so bi-modal. Use of performance based expert combination algorithms can help in such situations (Aspinall 2010).

Using expert judgment to quantify dependence requires that analysts and experts know where to look for dependence relations. The difficulty with micro-correlations is that they are often overlooked and tail dependence may not be observed until a very extreme event occurs, by which point it is a bit late for risk managers who would have liked to been able to reduce the threat. Many did not think to explore how tail dependent mortgages were until it was too late. Similarly, risk modelers over looked tail dependence related to hurricane damages until Katrina had hit.

Still, there will be many cases where we will have some idea about the catastrophic or dependent nature of a risk and will need to rely on expert judgment to quantify the dependence. This can be done by sticking to the cardinal rule of expert elicitation: *Elicit only about observable phenomena with which the experts are familiar.* This has two implications: First experts should only be asked questions which we could in principle ask of nature: Describe the controlled experiment or measurement which you would do to determine the quantity of interest and ask the expert to quantify his/her uncertainty on the outcome of this experiment or measurement. Second, the experiments or measurements must be similar to those which have been performed in the past, to ensure the experts are familiar with the question.

This suggests that experts will be able to play a role in countering the unholy trinity as more experience is gained in recognizing and measuring these phenomena. The first order of business is to gather more relevant data and find more ways detecting micro correlations, fat tails and tail dependence. Mathematical modeling will be needed to explore features of complex dependence models, define plausible model types and generate elicitation questions to quantify these models which experts can answer. The latent variable model presented above is simple, but suitable in many cases. Assuming the copula coupling the latent variable to each physical variable is the same for all physical variables, we could quantify such models by asking experts to quantify one typical copula. If tail dependence for the copula is excluded, it may be sufficient to quantify the correlation in one copula. This can be leveraged by quantifying the rank correlation between any two physical variables, and using a minimal information copula between the latent and the physical variables which induces the elicited correlation between observable variables. This technique was used extensively in the USNRC-EU joint uncertainty analysis.

The state of play is currently moving beyond such simple models. A recent example uses complex relationships between bivariate and conditional bivariate distributions to model dependence in four financial time series (Aas and Berg 2010), namely Norwegian stock index (T), the MSCI world stock index (M), the Norwegian bond index (B) and the SSBWG hedged bond (S). The structure of bivariate and conditional bivariate distributions is called a D-Vine (Figure 10), and is the simplest of a wide class of dependence structures which represent high dimensional distributions in terms of conditional bivariate distributions (Kurowicka and Cooke 2006).

*Updated version appearing in* Klaus Boecker (ed) Re-Thinking Risk Measurement and Reporting, Uncertainty, Bayesian Analysis and Expert Judgement, Risk Books

**Figure 10 D-Vine dependence structure for Norwegian stock index (T), the MSCI world stock index (M), the Norwegian bond index (B) and the SSBWG hedged bond (S) (Aas and Berg, 2010).**
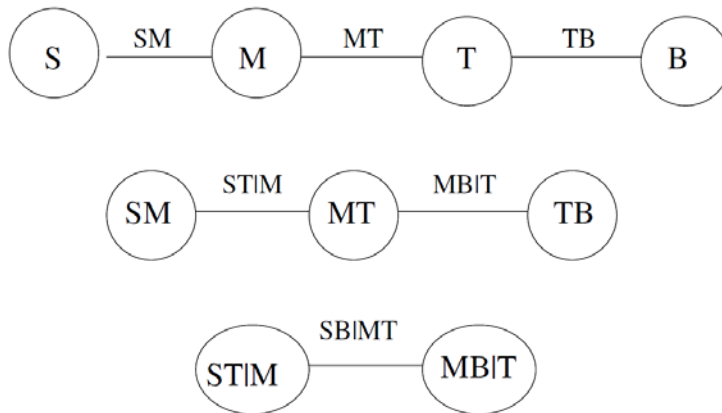


In Figure 10, dependence in the joint distribution is represented as bivariate distributions (S,M),  (M,T), and (T,B), and conditional bivariate distributions (ST | M), (M,B | T), and (S,B | M,T). A nested maximal likelihood procedure was used to quantify the (conditional) copulae, based on historical data. Tail dependent bivariate copulae in combination with this dependence structure gave demonstrably better results than more simple models. In particular, it outperformed the multivariate Student T copula, which is the current favorite in some circles.

Such examples are illuminating and challenging; illuminating in that they show that capturing tail dependence in combination with complex dependence structures can pay off, and challenging in that they make it more difficult to ignore these features just because multivariate historical data is not available. Developing protocols for using structured expert judgment to shore up data shortfalls will place greater demands on the mathematical sophistication of analysts, experts, and problem owners.

## References

Aas, K.  and Berg D. (2010) "Modelling dependence between financial returns using pair-copula constructions"  Handbook on Vine Copulae  Kurowicka and Joe (eds) World Scientific, 2010.

Aspinall, W. (2010) "Improving expert advice" appearing in *Nature.*

Bierlant, J., G. Dierckx and A. Guillou (2005). "Estimation of the Extreme-Value Index and Generalized Quantile Plots." Bernoulli **11**(6): 949-970.

Camerer, C. F. and H. Kunreuther (1989). "Decision Processes for Low Probability Events: Policy Implications." Journal of Policy Analysis and Management **8**(4): 565-592.

Chavez-Demoulin, V., P. Embrechts and J. Nešlehová (2005). Quantitative Models for Oeprational Risk: Extremes, Dependence and Aggregation. Implementing an AMA for Operation Risk. Federal Reserve Bank of Boston. May 18-20.

Cooke, R.M., and Goossens, L.J.H. Nuclear science and  technology. Procedures guide for structured expert judgment,  EUR 18820EN, European Commission. Luxembourg  2000.

Cooke, R.M., Kousky, C. and Joe, H. (2010) Micro-correlations and tail dependence, appearing in Dependence Modeling, Handbook on Vine Copulae  Kurowicka and Joe (eds) World Scientific, 2010.

David, H. A. and H. N. Nagarja (2003). Order Statistics, Third Edition. New York, New York, John Wiley & Sons, Inc.

Davidson, R. J. J. (1998). "Working Toward a Comprehensive National Strategy for Funding Catastrophe Exposure." Journal of Insurance Regulation **17**(2): 134-170.

Doherty, N. A. and H. Schlesinger (2002). "Insurance Contracts and Securitization." The Journal of Risk and Insurance **69**(1): 45-62.

Embrechts, P. (2007). Copulas: A Personal View. Zurich, Switzerland, ETH Zurich.

Genest, C. and J. MacKay (1986). "The Joy of Copulas: Bivariate Distributions with Uniform Marginals." The American Statistician **40**(4): 280-283.

Embrechts P., C. Klüppelberg and T. Mikosch (1997). Modelling extremal events for insurance and finance, Springer, Berlin

Hallegatte, S., A. Boissonnade, M.-E. Schlumberger and R. Muir-Wood (2008). "Demand Surge and Worker Migrations in Disaster Aftermaths: Application to Florida in 2004 and 2005."

Harrington, S. E. and G. Niehaus (2001). "Government Insurance, Tax Policy, and the Affordability of Catastrophe Insurance." Journal of Insurance Regulation **19**(4): 591-612.

Hartwig, R. P. (2009). Financing Catastrophe Losses Amid a Financial Catastrophe. Presentation at the Bahamas Weather Conference. Nassau, Bahamas, Insurance Information Institute. Online at: http://www.iii.org/media/presentations/bhccats.

Joe H. (1997), Multivariate Models and Dependence Concepts, Chapman &Hall, London.

Joe, H., H. Li and A. K. Nikoloulopoulos (2008). Tail Dependence Functions and Vine Copulas. Pullman, WA, Washington State University, Department of Mathematics.

Kahneman, D., P. Slovic and A. Tversky, Eds. (1982). Judgment under Uncertainty: Heuristics and Biases. Cambridge, UK, Cambridge University Press.

Klein, R. W. (2005). A Regulator's Introduction to the Insurance Industry. Kansas City, Missouri, National Association of Insurance Commissioners.

Klein, R. W. and S. Wang (2007). Catastrophe Risk Financing in the US and EU: A Comparative Analysis of Alternative Regulatory Approaches. <u>SCOR-JRI Conference on Insurance, Reinsurance and Capital Market Transformations</u>. Paris, France. September 20-21.

Kleindorfer, P. R. and R. W. Klein (2002). Regulation and Markets for Catastrophe Insurance. <u>Advances in Economic Design</u>. M. R. Sertel and S. Koray. Berlin, Germany, Springer-Verlag**:** 263-280.

Kousky, C. and R. Cooke (2009). The Unholy Trinity, Fat Tails, Tail Dependence and Micro-Correlations, RFF Discussion Paper 09-36-REV. Washington, D.C., Resources for the Future.

Kousky, C. and R. Cooke (2009). Climate Change and Risk Management: Challenges for Insurance, Adaptation, and Loss Estimation. RFF Discussion Paper 09-03-REV. Washington, D.C., Resources for the Future.

Kousky, C., E. Luttmer and R. Zeckhauser (2006). "Private Investment and Government Protection." <u>Journal of Risk and Uncertainty</u> **33**(1): 73-100.

Kunreuther, H. (2006). "Disaster Mitigation and Insurance: Learning from Katrina." <u>The ANNALS of the American Academy of Political and Social Science</u> **604**: 208-227.

Kunreuther, H. (2008). "Reducing Losses from Catastrophic Risks through Long-Term Insurance and Mitigation." <u>Social Research</u> **75**(3): 905-930.

Kurowicka, D. and R. Cooke (2006). <u>Uncertainty Analysis with High Dimensional Dependence Modelling</u>. West Sussex, England, John Wiley & Sons, Ltd.

Ladoucette, S. A. and J. L. Teugels (2006). "Analysis of Risk Measures for Reinsurance Layers." <u>Insurance: Mathematics and Economics</u> **38**(3): 630-639.

Latchman, J. L., F. D. O. Morgan and W. P. Aspinall (2008). "Temporal Changes in the Cumulative Piecewise Gradient of a Variant of the Gutenberg-Richter Relationship, and the Imminence of Extreme Events." <u>Earth-Science Reviews</u> **87**: 94-112.

Lewis, C. M. and K. C. Murdock (1999). Alternate Means of Redistributing Catastrophic Risk in a National Risk-Management System. <u>The Financing of Catastrophe Risk</u>. K. A. Froot. Chicago, University of Chicago Press**:** 51-85.

Malamud, B. D. and D. L. Turcotte (2006). "The Applicability of Power-Law Frequency Statistics to Floods." <u>Journal of Hydrology</u> **322**: 168-180.

Mandelbrot, B. (1963). "The Variation of Certain Speculative Prices." <u>The Journal of Business</u> **36**(4): 394-419.

Mandelbrot, B. (2004). <u>The (Mis)Behavoir of Markets</u>. New York Basic Books.

Mas-Colell, A., M. D. Whinston and J. R. Green (1995). <u>Microeconomic Theory</u>. New York, Oxford University Press.

Morales, O. Kurowicka, D. and Roelen, A. (2007) Eliciting Conditional and Unconditional Rank Correlations from Conditional Probabilities, Reliability Engineering & System Safety, 93, 600-710.  Available online 12 March 2007,  Volume 93, Issue 5, May 2008.

McNiel, A. J., R. Frey and P. Embrechts (2005). <u>Quantitative Risk Management: Concepts, Techniques, and Tools</u>. Princeton, NJ, Princeton University Press.

Michel-Kerjan, E. and F. Morlaye (2008). "Extreme Events, Global Warming, and Insurance-Linked Securities: How to Trigger the "Tipping Point"." <u>The Geneva Papers</u> **33**: 153-176.

Milidonis, A. and M. F. Grace (2007). Tax-Deductible Pre-Event Catastrophe Loss Reserves: The Case of Florida. <u>Working Paper</u>. Manchester, UK, The University of Manchester.

Nelson, R. (1999). <u>An Introduction to Copulas</u>. New York, Springer.

Nordhaus, W. D. (1994). "Expert Opinion on Climatic Change." <u>American Scientist</u> **82**(1): 45-52.

Resnick, S. (2007). <u>Heavy Tailed Phenomena: Probabilistic and Statistical Modeling</u>. New York, Springer.

Riker, W. (2004). "Using Models for Effective Risk Management." <u>Exposure</u> **Summer**.

RMS (2005). Hurricane Katrina: Profile of a Super Cat: Lessons and Implications for Catastrophe Risk Management. Newark, California, Risk Management Solutions.

Salmon, F. (2009) "Recipe for Disaster: The Formula That Killed Wall Street" Wired, Feb 23, 2009.  http://www.wired.com/techbiz/it/magazine/17-03/wp_quant?currentPage=all,

Swiss Re (2008). Disaster Risk Financing: Reducing the Burden on Public Budgets. Swiss Re Focus Report. Zurich, Switzerland, Swiss Re.

Tversky, A. and D. Kahneman (1982). Belief in the Law of Small Numbers. <u>Judgment under Uncertainty: Heuristics and Biases</u>. D. Kahneman, P. Slovic and A. Tversky. Cambridge, UK, Cambridge University Press**:** 23-31.

Van Oorschot, M.M.P., Kraan, B.C.P., van den Brink, R.M.M. , Janssen, P.H.M.,  and Cooke, R.M., (2003) Uncertainty Analysis for NOx Emissions from Dutch passenger cars in 1998, Applying structured expert elicitation and distinguishing different types of uncertainty, RIVM report 550002004/2003, Bilthoven, The Netherlands

US Nuclear Regulatory Commission Reactor Safety Study: An Assessment of Accident Risks in the US Commercial Nuclear Power Plants, Wasah-1400, NUREG-75/014, 1975.

van Noortwijk, J. M. (1996). Optimal Maintenance Decisions for Hydraulic Structures. Dept. of Mathematics. Delft, The Netherlands, Delft University of Technology. PhD dissertation.

van Noortwijk, J. M. (2009).A survey of the application of gamma processes in maintenance Reliability Engineering & System Safety, 94(1):2-21.

Weitzman, M. (2008). On Modeling and Interpreting the Economics of Catastrophic Climate Change, forthcoming in the Review of Economics and Statistics.