DISCUSSION PAPER

September 2009; revised November 2009 ■ RFF DP 09-36-REV

The Unholy Trinity: Fat Tails, Tail Dependence, and Micro-Correlations

Carolyn Kousky and Roger M. Cooke

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

The Unholy Trinity: Fat Tails, Tail Dependence, and Micro-Correlations

Carolyn Kousky and Roger M. Cooke

Abstract

Recent events in the financial and insurance markets, as well as the looming challenges of a globally changing climate point to the need to re-think the ways in which we measure and manage catastrophic and dependent risks. Management can only be as good as our measurement tools. To that end, this paper outlines detection, measurement, and analysis strategies for fat-tailed risks, tail dependent risks, and risks characterized by micro-correlations. A simple model of insurance demand and supply is used to illustrate the difficulties in insuring risks characterized by these phenomena. Policy implications are discussed.

Key Words: risk, fat tails, tail dependence, micro-correlations, insurance, natural disasters

JEL Classification Numbers: Q54, G22, C02

Contents

The Unholy Trinity: Fat Tails, Tail Dependence, and Micro-Correlations

Carolyn Kousky and Roger M. Cooke*

1. Introduction

Managing potential losses is critical to the property insurance company writing homeowners policies, the bank underwriting mortgages, the stock market investor, and the federal government setting rates for the National Flood Insurance Program. If losses are independent and never terribly severe, managers can draw on traditional diversification strategies. It doesn't take the Great Depression or Hurricane Katrina, however, to realize that losses are often catastrophic and dependent; our work focuses on these types of risks, which pose unique challenges to risk managers.

In particular, we examine *fat tails*, *tail dependence*, and *micro-correlations*—an "unholy trinity" of risk management perils. These are distinct aspects of loss distributions, such as damages from a disaster or insurance claims. With fat-tailed losses, the probability declines slowly, relative to the severity of the loss. Tail dependence is the propensity of dependence to concentrate in the tails, such that severe losses are more likely to happen together. Micro-correlations are negligible correlations between risks which may be individually harmless, but very dangerous when aggregated. These three phenomena—types of catastrophic and dependent risks—undermine traditional approaches to risk management.

At the heart of much risk management is aggregation. Firms hold not one insurance policy, or one mortgage, but a portfolio of investments. Holding such bundles offers diversification benefits and stabilizes losses. As we will show here, however, this traditional approach for managing risks can fail when loss distributions are characterized by fat tails, tail dependence, or micro-correlations. If one does not know how to detect these phenomena, it is not possible to manage them, such that firms may unwittingly court insolvency and the government may be exposed to losses of which they are unaware.

* Kousky is a fellow at Resources for the Future and Cooke is Chauncey Starr Senior Fellow at Resources for the Future. Address correspondence to: Carolyn Kousky, Resources for the Future, 1616 P Street, NW, Washington, DC 20036. Email: kousky@rff.org. We would like to thank Mark Heising, Liz Simons, and the Simons Foundation for making this work possible.

While these phenomena could describe a range of risks, in this paper we focus on damages from natural disasters. We use three datasets to show the existence of the unholy trinity in disaster damages and to develop methods for their detection, measurement, and analysis. We then illustrate the challenges for insuring these types of risks with a simple model of the demand and supply of insurance. The datasets include: flood insurance claims data from the National Flood Insurance Program (NFIP), crop insurance indemnities paid data from the United States Department of Agriculture's Risk Management Agency,[1] and the SHELDUS database, maintained by the Hazards and Vulnerability Research Group at the University of South Carolina,[2] which has county-level damages and fatalities from weather events.

The next three sections of the paper describe sequentially the three phenomena of fat tails, tail dependence, and micro-correlations, providing examples of their importance. We use the above mentioned datasets to develop techniques for measuring and analyzing these three types of risks. Section 5 then presents a model of the demand and supply of insurance for catastrophic and dependent risks, demonstrating the pitfalls of insuring distributions characterized by the unholy trinity. Policy implications and potential remedies are discussed.

2. Fat Tails

Fat tails were introduced in mathematical finance in 1963 by Benoit Mandelbrot to describe cotton price changes (Mandelbrot 1963). Since then, evidence has accumulated that many types of damages, from financial losses to natural disasters, are best characterized by distributions with fat tails (e.g., Mandelbrot 2004; Malamud and Turcotte 2006; Latchman, Morgan and Aspinall 2008). The uncertainty surrounding climate change impacts 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 $\alpha, k > 0$. The variable α is referred to as the tail index or tail parameter.

---

[1] We would like to thank Ed Pasterick and Tim Scoville for providing the NFIP data and Barbara Carter for providing the crop data.

[2] Information on SHELDUS is available online: http://webra.cas.sc.edu/hvri/products/SHELDUS.aspx. The damage and fatality estimates in SHELDUS are minimum estimates as the approach to compiling the data always takes the most conservative estimates (for further discussion, see: Cutter, Gall and Emrich 2008).

From this definition, it is clear that there are degrees of tail fatness. The *m*-th moment is infinite if $\alpha \leq m$. If $\alpha \leq 1$, we say 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, with $1 < \alpha \leq 2$, have a defined mean, but an infinite variance. The sample mean for these distributions also has infinite variance no matter how many samples are drawn.

When either the mean or the variance is infinite, two types of problems emerge for risk managers: (1) historical data will be a poor guide for the future, and (2) the tails of aggregations are also fat. To illustrate (1), suppose an insurance company wishes to assess the risk of flood damage to property in Dade County, Florida. They might look to the historical record and tally the total claims divided by the total value insured or the total number of policies. Any such method presumes that the historical average gives a good estimate of the risk. If the loss distribution is really fat, however, they must contend with the fact that the historical average itself has infinite variance, no matter how many years are averaged, and will thus not be stable. For problem (2), consider that a (re-)insurer may traditionally diversify his risk by combining insurance policies from multiple geographic locations. If the losses per area per year are drawn independently from a loss distribution with finite variance, then the sum of many such losses will tend to be normal, and thus thin-tailed, according to the central limit theorem. However, if the loss distributions have infinite variance, then the central limit theorem does not apply; instead the sums converge to a stable law which has the same tail behavior as the summands. Aggregation does not yield thinner tails. We will elaborate further on these points below.

Measuring tail obesity in loss distributions presents challenges. The best-known approach is the Hill estimator, which estimates the tail parameter, α, of a Pareto distributed variable, *X*, or any variable whose tail beyond some threshold x_k follows a Pareto distribution. If $X_i \ldots X_n$ are independent versions of *X*, the Hill estimator based on x_k from a sample $\{x_1, \ldots x_n\}$ is the maximum-likelihood estimator of 1/α given by (see: Resnick 2007):

(1) $(\sum_{i;\, xi > xk} \log(x_i) - \log(x_k)) \,/\, \#[i \mid x_i > x_k]$.

This works well *if* the data are indeed drawn from a Pareto distribution. If the data are not Pareto, but still fat-tailed, the behavior of the Hill estimator is notoriously unstable for large thresholds (Brielant et al. 2005; Resnick 2007), rendering its use in practical risk management problematic. This is seen when using the Hill estimator on property damage data for all natural hazards in the United States over the years 1960 to 2007. Damages are from the SHELDUS database in constant 2007 dollars. Figure 1 shows a Hill plot of this data. The tail index estimated from the top *n* order statistics is plotted against *n* (order 0 is the largest). Is this tail Super Fat or only Really Fat? This Hill index is inconclusive. Bierlant et al. (2005) review attempts to improve the Hill estimator.

Figure 1: Hill plot for US property damages from natural disasters

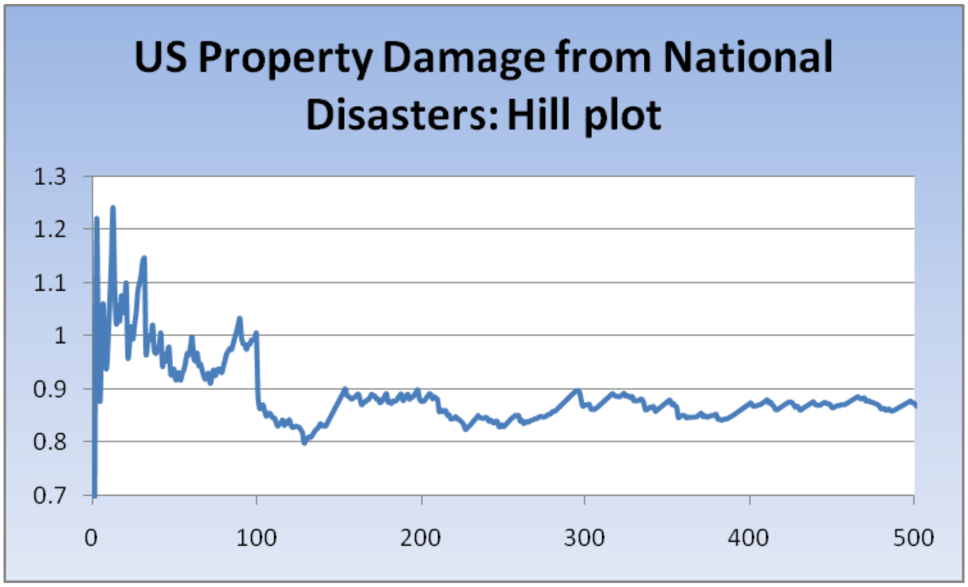

Another diagnostic of tail obesity are mean excess plots. If variable X has cumulative distribution function F then the *mean excess curve* for X is defined as:

(2) $G(x_0) = \mathrm{E}(X - x_0 \mid X > x_0)$.

It is well known that the mean excess curve for the Pareto distribution is linear, $G(x_0) = (x_0+k)/(\alpha-1)$ (McNiel, Frey and Embrechts 2005). Note that $G(x_0)$ is not defined for $\alpha=1$. For a finite ordered sample, $x_1 < x_2, \ldots < x_n$, the sample mean excess *plot* gives the values:

(3) $\{x_i , g(x_i)\};\ g(x_i) = \Sigma_{j = i,\ldots n-1} (x_{j+1}-x_i)/(n-i);\ i = 1,\ldots n-1;\ g(x_n) = 0$.

Evidently for any constant c, $g(cx_i) = cg(x_i)$; it is therefore convenient to standardize such plots by dividing all x_i*'s* by the largest, so that x_n=1 (g is also invariant under shifts of location, but these are not considered here).

Whereas the theoretical mean excess curve for a Pareto variable is a straight line, the empirical mean excess plot is not. Figure 2 shows standardized, sample mean excess plots for Super Fat, Really Fat, and Meso-Fat (infinite third moment) Paretos, where the x_i*'s* are obtained by inverting uniformly spaced percentiles:

(4) $\Pr\{X > x_j\} = (1+x_j)^{-\alpha} = 1 - j / 1001; j = 1,\ldots 1000.$

Evidently, these plots are not linear. For instance, a Pareto with tail index of 2 does not have a sample mean excess plot that is linear with slope 1 as the mean excess curve would be (middle plot in Figure 2). The mean excess *curve* for a Pareto with a tail index of 1 has infinite slope, but of course the sample mean excess plot will not (far right plot of Figure 2). Figure 2 quickly disabuses us of the idea that tail obesity can be measured by eyeballing a sample mean excess plot.

Figure 2: Standardized mean excess plots for inverse percentiles of Pareto variables, with tail index α =3 (left), 2 (center), 1 (right)

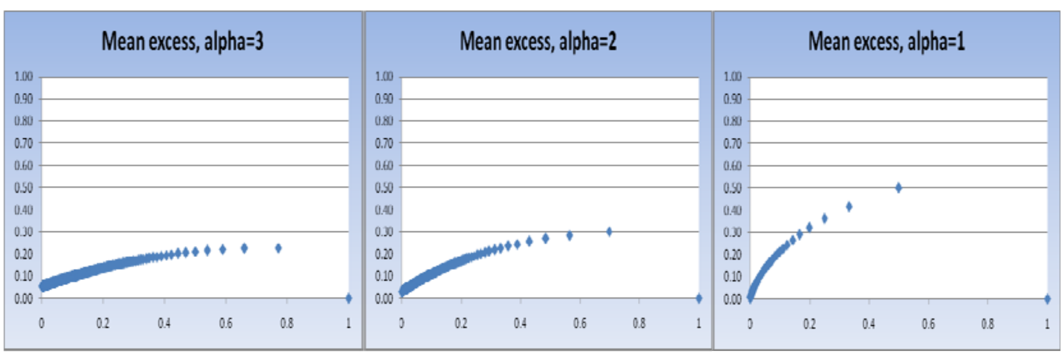

We've noted that whether a tail has infinite variance is of key concern for risk managers as this determines whether the tail will thin under aggregation. As such, a simple diagnostic for whether data comes from a distribution with infinite variance would be useful. Drawing on the behavior of sums of fat-tailed variables, we propose looking at the speed of collapse of the mean excess plot as variables are aggregated. Suppose we draw 5,000 samples from a Pareto distribution with tail index $\alpha = 1$. The slope of the theoretical mean excess plot is infinite, but the sample mean excess plot will look like a noisy version of the right-most plot in Figure 2. Suppose we now form random groups of ten distinct samples and sum the samples in each group. We call this operation "*aggregation by 10*." Since independent sums of these variables converge to a stable law with the same tail index (Resnick 2007), we might expect that the mean excess of the aggregation by 10 should resemble the original mean excess. Indeed it does. Figure 3 compares this standardized mean excess plot for $\alpha=1$ with aggregations by 10 and by 50, for Paretos with tail index $\alpha=1,2,3$, and also exponential variables. Sums of independent exponentials converge quickly to a normal with a descending mean excess plot; indeed for aggregation by 50, the aggregands' values are less disperse. Index $\alpha=2$ is the highest index with infinite variance. Of course the finite sample has finite variance and aggregation by 50 produces a descending mean excess plot. For $\alpha=3$, aggregation by 10 largely eliminates the positive slope.

Figure 3: Standardized mean excess plots, with aggregations by 10 and 50, for Pareto distributions with α= 1,2, 3 and for the exponential distribution

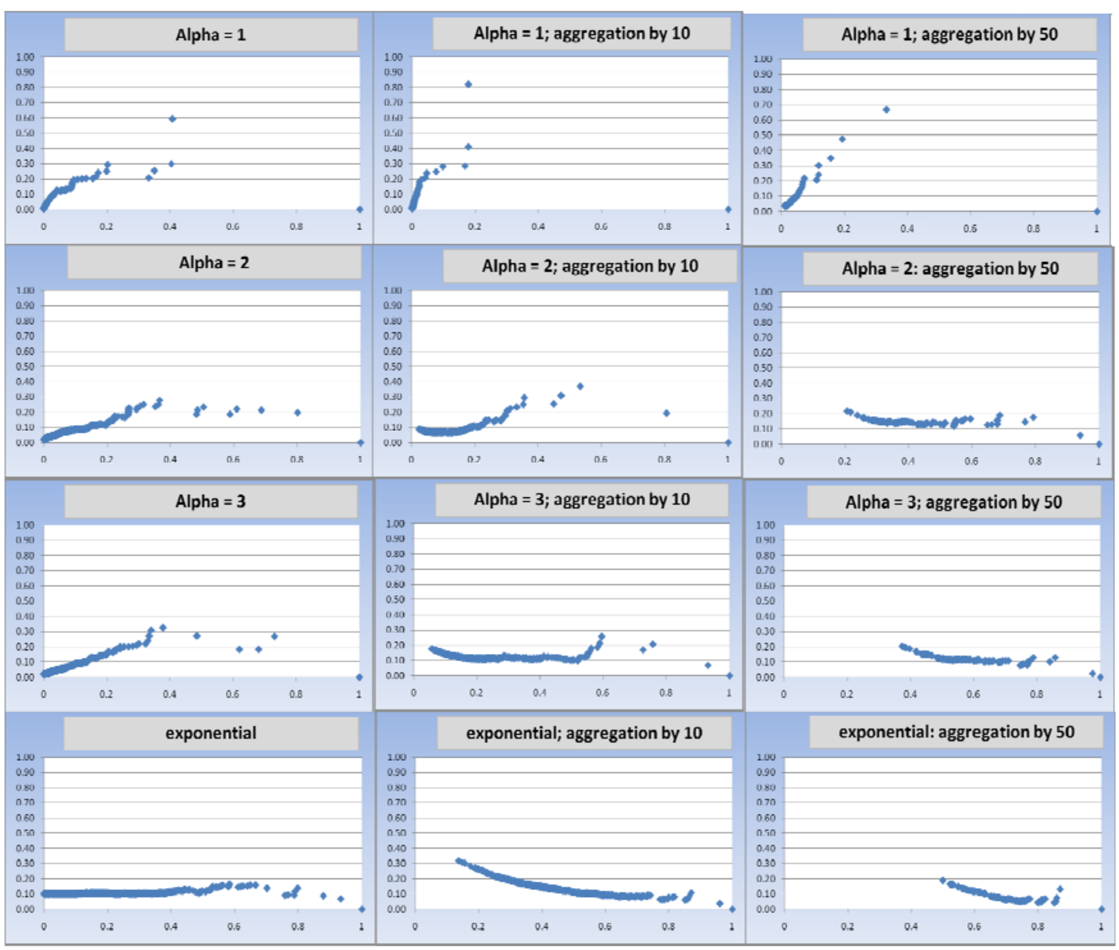

This same approach can be used to examine the tail behavior of actual loss data. We look at crop indemnities paid per county and National Flood Insurance Program claims by county for the years 1980 to 2006. Over this time period, there has been substantial growth in exposure to flood risk, particularly in coastal counties. To remove the effect of growing exposure, we divide the claims by personal income estimates from the Bureau of Economic Accounts (BEA).[3] Thus, we study flood claims per dollar income, by county and year.[4] The crop loss claims are not exposure adjusted, as a proxy for exposure is not obvious, and exposure growth is less of a concern.

Figure 4: Mean excess plots for US crop loss (above) and exposure corrected flood claims (below)

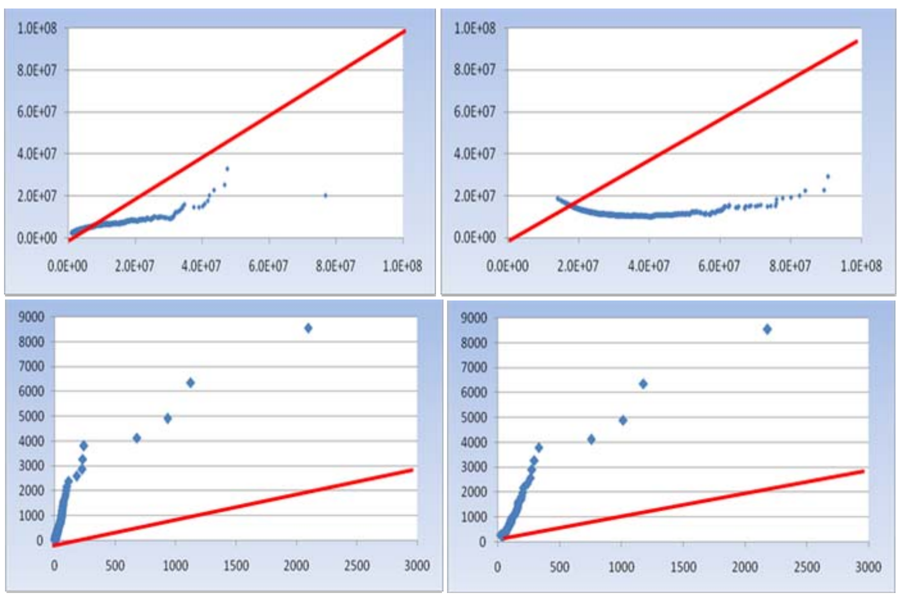

Note: The vertical axis gives mean excess loss, given loss at least as large as the horizontal axis. The upper right picture shows crop payment mean excess for random aggregations of 10 counties. The lower right picture shows flood claims mean excess for random aggregations of 50 counties.

---

[3] 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.

[4] This arguably may not be the most appropriate normalization. The purchase of flood policies has grown over time, for instance, and this is not accounted for here. The mean excess plot for inflation-adjusted claims paid divided by the number of claims in each county, however, shows a similar mean excess plot to the one in Figure 4.

Figure 4 shows each county-year as a realization of a single random variable. The left graphs depict mean excess plots for crop (above) and exposure adjusted flood insurance claims (below). This data is not standardized, and the unit slope lines have been added. Whereas the sample mean excess plot for crop payments suggests a slope less than 1, the sample mean excess plot for flood claims suggests an infinite variance. The mean excess plot of a random aggregation by 10 for crop losses in Figure 4 (upper right) shows a decreasing slope. Compare this with Figure 4 (lower right) showing the mean excess of a random aggregation by 50 of exposure adjusted flood losses. Here, aggregation does not thin the tail significantly; the mean excess plot is self similar. The same picture emerges in when looking again at the SHELDUS data for all natural disasters in the U.S. (Figure 5).

Figure 5: Self similarity under aggregation of US property damage mean excess plots

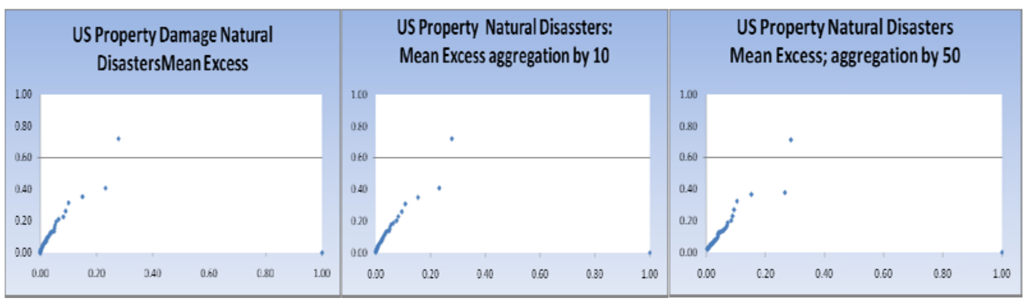

Figures 4 and 5 suggest that the behavior of loss data under random aggregation gives valuable insight into tail obesity. In particular, the speed with which the mean excess curve collapses under random aggregation appears to be a useful indicator of whether the data come from a distribution with infinite variance—the most critical information for risk managers. Unlike the Hill estimator, this indicator is not an estimate of behavior at infinity, but depends only on the finite sample we have before us, making it much more useful in practice. Moreover, it is not an estimate of a parameter in a theoretical model of the data, but a feature of the data itself.

While these tools are useful, applying them to appropriately corrected exposure data is critical. Risk is damage per unit exposure. Naïve examinations of risk will ignore exposure, but changes in exposure can alter risk assessments, sometimes fundamentally. As a simple illustration, Figure 6 shows the mean excess plot of deaths per county-year from natural disasters in the US from 1995-2007.[5] For each point *x* on the horizontal axis, the plot shows the expected excess number of fatalities in county-years having at least *x* fatalities. The left hand curve is very steep, indicating a very fat-tailed distribution. Does this mean that the physical impacts from disasters are fat-tailed, or that population is concentrated in harm's way leading to fat-tailed damages? The graph on the right shows the same data, but now the fatalities per county are divided by the population of the county. It is a different dimensioning of risk and shows a very different result. Whereas the risk per county is very fat tailed, the risk per person is not. The high damages are coming from disasters which hit high population areas, not from the disaster intensity as such. Which figure is most appropriate depends on our risk management problem.

Figure 6: Mean excess plots of fatalities per county-year, without (left) and with (right) correction for population

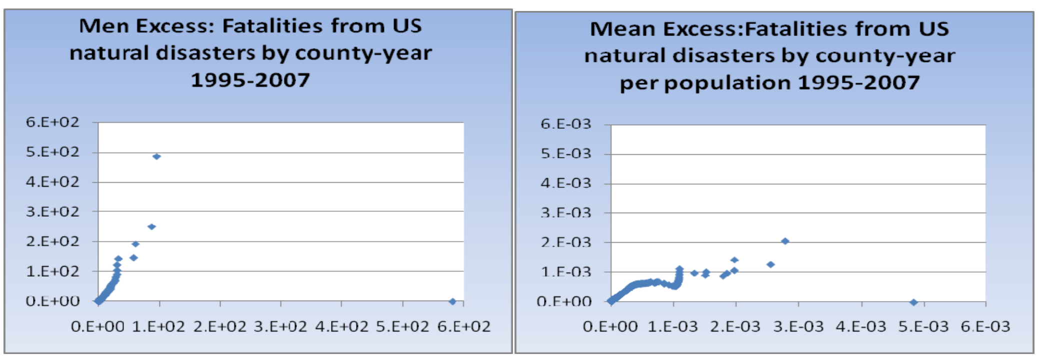

3. Tail Dependence

Tail dependence refers to the tendency of dependence between two random variables to concentrate in the extreme values. For loss distributions, we are interested in dependence of non-negative variables concentrating in the extreme high values, or upper tail dependent (UTD). Upper tail dependence of variables *X* and *Y* is defined as the limit (if it exists) of the probability that *X* exceeds its *r-percentile*, given that *Y* exceeds its *r-percentile,* as *r* goes to 100.

Hurricane Katrina vividly demonstrated tail dependence across damage types and insurance lines. The storm not only caused wind and rain damage, but damage from breached

[5] This data is from SHELDUS and does not include fatality estimates from Katrina. Other estimates in SHELDUS seem to be underestimates of deaths reported in other sources. The data is used here merely for purposes of illustration.

levees and storm surge, power outages, fires that could not be put out, business interruptions, toxic spills, a rise in energy costs from damage to rigs and refineries, and increased costs of reconstruction (this latter effect is referred to as "demand surge," see: Hallegatte, Boissonnade, Schlumberger et al. 2008). After Katrina, some lines of insurance that had not been heavily hit in other catastrophes saw many claims; among them cargo, inland marine and recreational watercraft, floating casinos, onshore energy, automobile, worker's comp, health, and life insurance (RMS 2005), demonstrating the tail dependence across these lines of business. The 9/11 terrorist attacks also demonstrated that for extreme events, multiple lines become affected (Riker 2004). When these tail dependencies are not considered, the tail exposure of an insurance company can be severely underestimated.

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 marginal distributions of X and Y; if we apply any 1-to-1 transformation to X, say $X^* = X^{1/N}$, $N \in \mathbb{N}$, $X > 0$ (which will thin X's tail), then UTD(X^*, Y) = UTD(X, Y). UTD has no simple relation to the standard Pearson correlation coefficient. For example, normal variables with any correlation ρ 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 terrible 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 7 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 7 shows that high flood damages and high wind claims occur together, indicative of tail dependence.

Figure 7: Tail dependence in wind and water claims, Florida 2002–2006

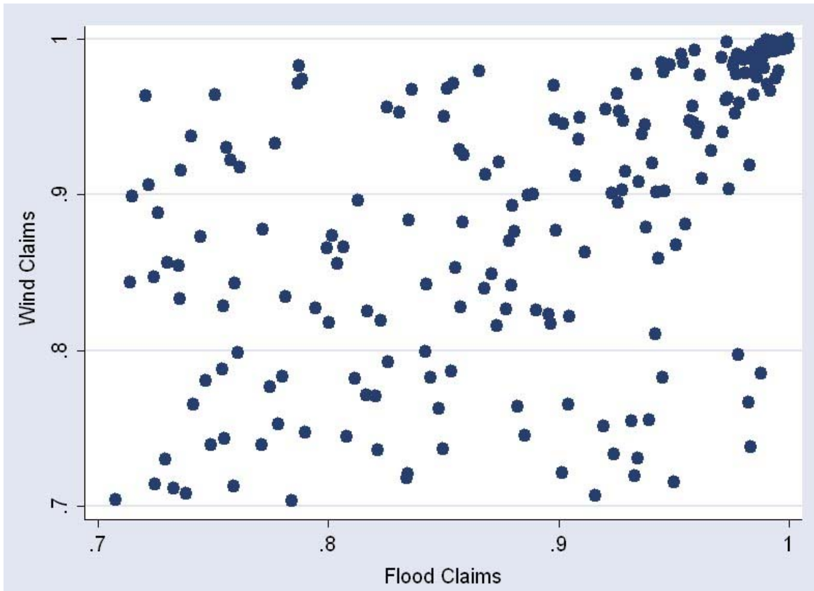

Note: Points of zero damages were removed, so rank axes do not begin at zero.

The joint distribution of the percentiles of two random variables is called a "*copula,*" and tail dependence is a property of the copula (for more information on copulae, see: Genest and MacKay 1986; Nelson 1999; Embrechts 2007). Copulae are useful tools for studying high dimensional multivariate distributions, as they allow us to separate the representation of dependence from the representation of the univariate marginal distributions (Kurowicka and Cooke 2006). Different marginal distributions can be combined in different dependence structures by choosing different copulae. The fact that tail dependence is a property of the copula immediately shows that there is no general relation between fat tails and tail dependence. These are separate issues. Current research focuses on the relation between tail dependence and multivariate extreme value copulae (Chavez-Demoulin, Embrechts and Nešlehová 2005; Joe, Li and Nikoloulopoulos 2008).

Under certain conditions, tail dependence can grow as variables are aggregated. If the random variables are thought to be insurance policies, this ballooning of tail dependence will again put limits on diversification. As one simple example, consider a basic model of dependence in which a set of random variables $X_1 \ldots X_n$ are symmetrically correlated with a "latent variable." The degree of correlation between the variables will then depend on the chosen copula. If a tail independent copula is chosen, such as the normal copula, aggregation will not increase tail dependence. If a weakly tail dependent copula is chosen, however, then the tail dependence can balloon upon aggregation.

One version of this is called the "L_p symmetric process," which is widely used in reservoir management, maintenance optimization, and deterioration modeling (van Noortwijk 1996). In this case, the latent variable is the scale factor, and the X_i's are conditionally independent gamma transform variables characterized by a fixed shape and a scale factor which is uncertain. Given a scale value, the variables are independent, but lack of knowledge of the scale factor induces a global correlation between the X_i's. In the simple case of conditionally independent exponentials with gamma distributed scale factor with shape υ, the unconditional distribution of each X_i is Pareto.

If we consider distinct sums of N such variables, they have upper tail dependence given by (Kousky and Cooke 2009):

(5) $$\frac{\sum_{k,j=0\ldots N-1}\left(\frac{1}{2}\right)^{\upsilon+j+k}\Gamma(\upsilon+k+j)/([\Gamma(\upsilon)k!\,j!]}{\sum_{k=0\ldots N-1}\Gamma(\upsilon+k)/[k!\Gamma(\upsilon)]};$$

which grows with N. This model is interesting because it is widely applied and it is one of the few in which we can actually compute all relevant quantities. We can see tail dependence emerging from summing familiar random variables.

In general, computations of sums of tail dependent variables are intractable; simulation, however, is rather easy. We choose the correlation $\rho(X_i, Latent) = 0.1$; the correlation between X_i and X_j will depend on the copula chosen to realize this correlation, but will be on the order of 0.01. Figure 8 shows four percentile scatter plots using a Gumbel copula having weak upper tail dependence. Figure 8a is simply between X_1 and X_2, and the correlation 0.02 is imperceptible. Figure 8b shows the scatter plot of distinct sums of 10 variables. The correlation between them of 0.14 is scarcely visible, but we do see some darkening of the upper right corner. In Figure 8c and 8d we see the scatter plots of sums of 20 and 40 variables respectively. Now the upper tail dependence becomes quite evident.

Figure 8: Tail dependence, Gumbel copula

(a) $X_1 \times X_2$ ***:*** **Gumbel;** ρ **= 0.02** **(b)** $\Sigma^{10} X_i \times \Sigma^{10} X_j$**: Gumbel;** ρ **= 0.14**

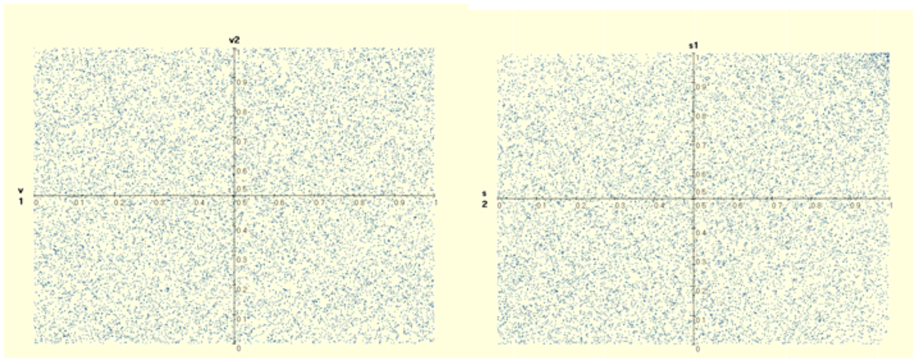

(c) $\Sigma^{20} X_i \times \Sigma^{20} X_j$***: Gumbel;*** $\rho = 0.25$ ***(d)*** $\Sigma^{40} X_i \times \Sigma^{40} X_j$***: Gumbel;*** $\rho = 0.40$

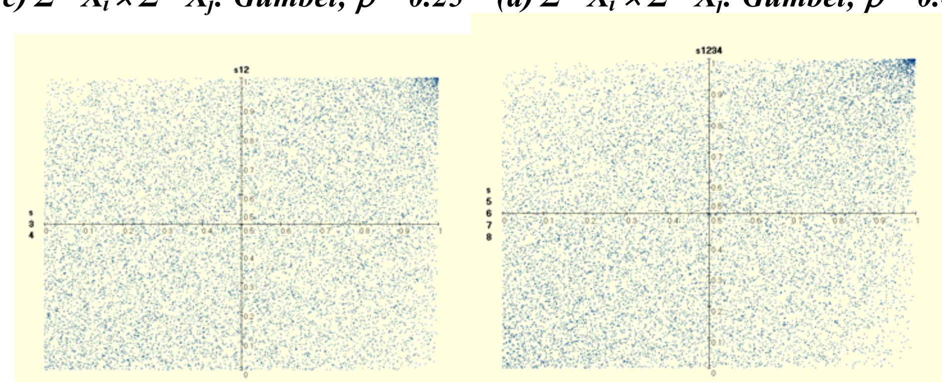

To illustrate how the aggregate tail dependence depends on the copula, Figure 9 shows the same model as in Figure 8, but with the elliptical copula[6] in 9a and the Normal copula in 9b. In both of these, there is no discernable tail dependence.

[6] The elliptical copula concentrates on an elliptical surface to induce the required correlation, see Kurowicka and Cooke (2006). It is of interest mainly because it is analytically tractable, is related to the normal copula, yet has markedly different properties.

Figure 9: Elliptical and Normal copulae – no tail dependence

(a) $\Sigma^{40} X_i \times \Sigma^{40} X_j$: Elliptical; $\rho = 0.40$ (b) $\Sigma^{40} X \times, \Sigma^{40} X_j$: Normal; $\rho = 0.31$

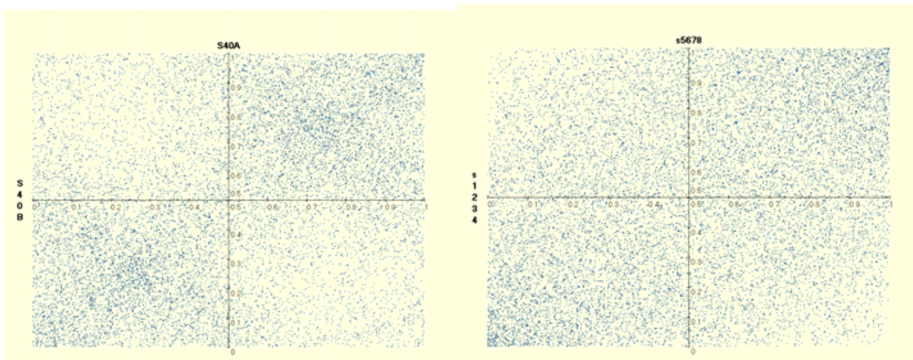

Figures 8 and 9 suggest that the way to detect tail dependence is to look at disjunct sums, as we saw earlier with measures for tail obesity of damage data. Initial work with monthly, county-level NFIP claims data in the state of Florida suggests this approach can be a useful detection method. If we 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 10 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 10: Percentile scatter plots of random aggregations of 5 (left) and 30 (right) Florida counties, monthly flood losses

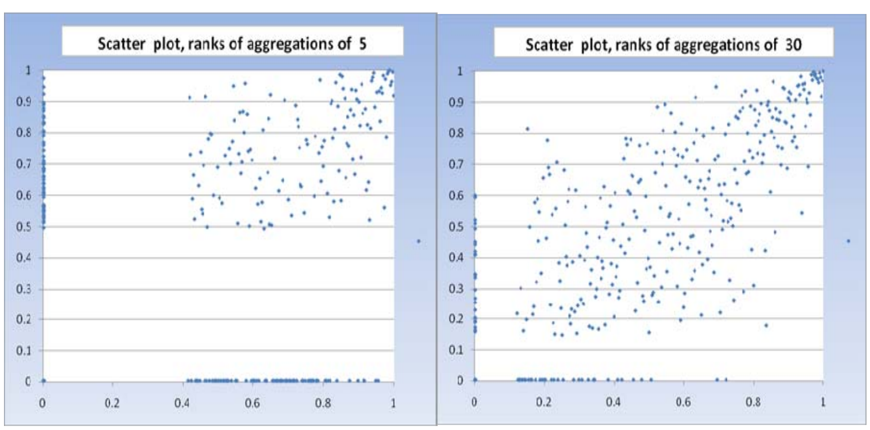

4. Micro-Correlations

Micro-correlations are small, positive correlations between 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, for example, but El Niño events induce exactly this coupling. These tiny correlations are amplified by aggregation, undermining 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 is not. Let $X_1, \ldots, X_N$ and $Y_1, \ldots, Y_N$ be two sets of random variables with the same average variance σ^2 and average covariance C (within and between sets). The correlation of the sums of the X's and the sum of the Y's is easily found to be:

$$(6) \quad \rho\left(\sum X_i, \sum Y_i\right) = \frac{N^2 C}{N\sigma^2 + N(N-1)C} .$$

This evidently goes to 1 as N grows, *if* C is non-zero and σ^2 is finite. If all variables are independent, then C=0, and the correlation in (6) is zero. The variance of ΣX_i is always non-negative; if the σ^2 and C are constant for sufficiently large N, it is easy to see that $C \geq 0$.

The amplification of correlation can be seen most dramatically in the flood insurance claim data. Suppose we randomly draw pairs of US counties and compute their correlation. The green histogram in Figure 11 shows 500 such correlations. The average correlation is 0.04. A few counties have high, positive 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 these correlations fall into that category.

Figure 11: US flood claims, correlations-of-1 (green), correlations-of-100 (blue), and correlations-of-500 (red)

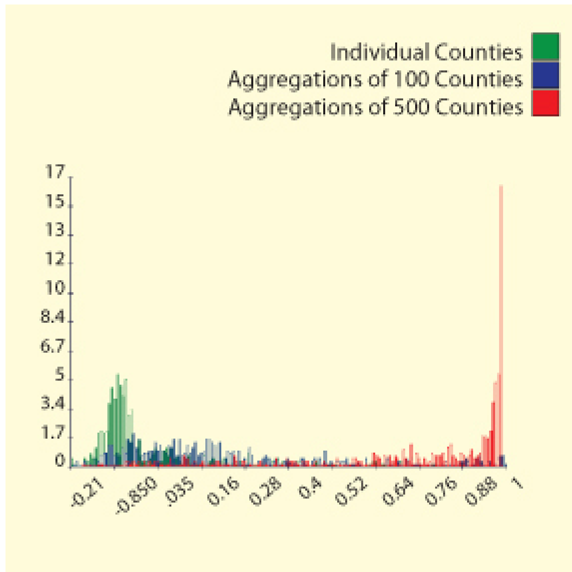

Instead of looking at the correlations between two randomly chosen counties, consider summing 100 randomly chosen counties and correlating this with the sum of another, distinct set of 100 randomly chosen counties. If we repeat this 500 times, the blue histogram in Figure 11 results; the average of 500 such correlations-of-100 is 0.23. The red 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 11 with Figure 12, in which each county is assigned an independent uniform variable, for each of 30 years. The correlations-of-1 and correlations-of-500 are effectively the same. Aggregation amplifies micro-correlations.

Figure 12: Correlations-of-1 and correlations-of-500, for 30 realizations of independent uniform variables

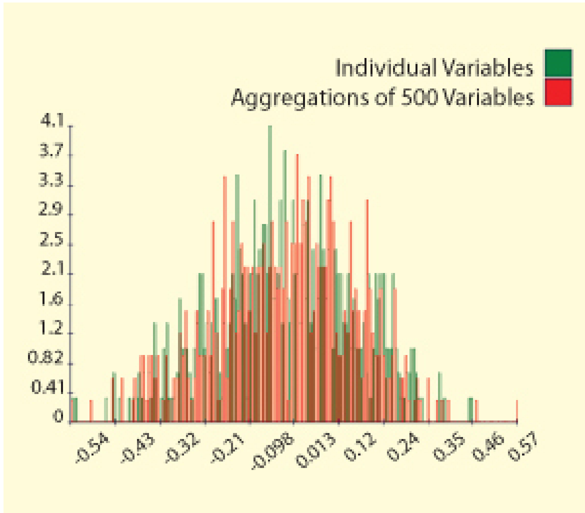

US Crop loss data shows stronger micro-correlations, and comparable levels of amplification are reached at aggregations by 100, as shown in Figure 13.

Figure 13: US Crop losses, correlations-of-1 (green), correlations-of-50 (blue), and correlations-of-100 (red)

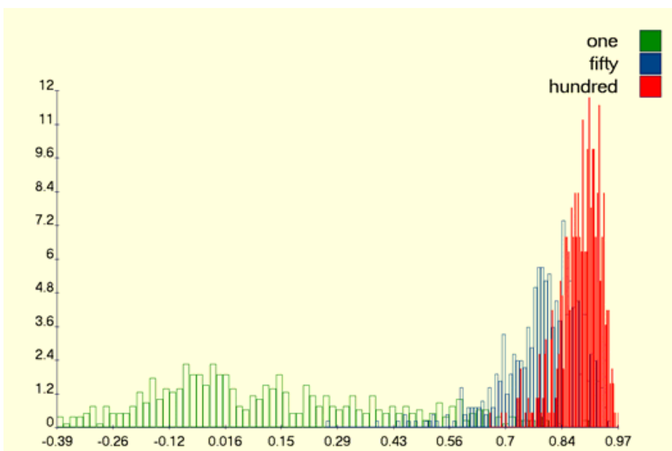

5. Insurance and Policy Implications

Catastrophic and dependent risks are expensive to insure and this can lead to a breakdown in the insurance market. Consider i=1…*N* policyholders in a given region or line of business. While the insurer will be offering other types of coverage besides homeowners policies, such as automobile insurance or business interruption insurance, for example, we focus just on the homeowners policies. Each year a homeowner faces a potential loss of L_i. An insurer offers coverage to homeowners at a price equal to q per dollar of coverage (rates may vary across lines or regions, but individual-specific rates are not possible).

The basic theoretical model of an individual's decision to purchase insurance (e.g., Mossin 1968) is formalized in most microeconomics and decision analysis textbooks. Adapted to our current set-up, let p be the probability of disaster, let w_i be the individual's total wealth, let α_i be the amount of dollars of insurance purchased, and let L_i and q be defined as in the previous paragraph. The expected utility (EU) for a potential consumer of insurance is then given by (deviations from EU theory discussed below):

(7) $EU = (1 - p)u(w_i - \alpha_i q) + pu(w_i - \alpha_i q - L_i + \alpha_i)$.

The risk-averse homeowner chooses α by maximizing expected utility (subject to the constraint that $\alpha_i \geq 0$), giving the first-order condition, assuming an interior solution, (where α^* is the optimal amount of insurance purchased):

(8) $-q(1 - p)u'(w_i - \alpha_i^* q) + p(1-q)u'(w_i - \alpha_i^* q - L_i + \alpha_i^*) = 0$.

Assume insurance is priced actuarially fairly, (ignoring transaction and marketing costs), such that $p = q$. In this case, we get the well-known result that a risk averse consumer, facing actuarial rates, will fully insure: $\alpha_i = L_i$. With this set-up in hand, we turn now to the insurer's problem.

The optimization problem of a solvency-constrained insurer can be modeled as follows. Let $\alpha = \sum \alpha_i^*$ be the portfolio of policies—the total exposure—held by the insurer. Let total claims for the insurance company—which will vary by the portfolio the company holds—be given by $C_\alpha = \sum c_i$. An insurer has access to some level of assets to support losses, given by A, and may also purchase reinsurance, K, at cost r per dollar covered.[7] We ignore marketing and transaction expenses here, although obviously prices will need to be high enough to cover these costs. For a solvency constrained insurer, the firm cannot spend more on claims than the total of revenue ($\Sigma q\alpha_i$), assets, and reinsurance. Following Kleindorfer and Klein (2002), expected profits for the insurer, $E[\Pi(q, A, K)]$, are thus given by:

(9) $E[\Pi(q, A, K) = \Sigma q\alpha_i - rK - E[\text{Min}(C_\alpha, \Sigma q\alpha_i + (1-r)K + A)]$.

The insurer seeks to keep the probability of insolvency below some target level λ.[8] $F(C_\alpha)$ is the cumulative distribution function of claims an insurer faces for a given portfolio. The insurer will then maximize expected profits subject to the following constraint:

(10) $\Pr(C_\alpha > \Sigma q\alpha_i + (1 - r)K + A) \leq \lambda$.

Define $S_{\alpha,\lambda}$ to be the required capital, or surplus, the insurer must have to cover claims that will occur with probability 1- λ when holding a given portfolio α. Then:

(11) $F^{-1}(1-\lambda) = S_{\alpha,\lambda}$.

If the policies within the portfolio are characterized by the three phenomena discussed here, more surplus will be required than if they are independent. This is illustrated in Figure 14.

[7] This is, of course, a very simplified construction of reinsurance. For alternate forms of reinsurance, see Ladoucette and Teugels (2006).

[8] Here, this probability is taken to be exogenous, perhaps set by capital regulations. For instance, in the EU, beginning in 2012, insurance companies will be regulated through the Solvency II regulations. The Solvency Capital Requirement component of the regulations takes a Value-at-Risk (VAR) approach, whereby an insurer must keep the probability of insolvency below a certain level, currently set at 0.5 percent.

Figure 14a compares the cumulative distribution functions of the sum of 100 lognormal variables (as a proxy for an insurance portfolio) when they are mutually independent (green) and mutually correlated at 0.01 (red), without tail dependence (using the normal copula). This shows that for any conceivable target solvency probability, the required surplus is greater for the micro-correlated variables. Figure 14b shows similar graphs, except that the correlation is realized by a tail dependent copula (Gumbel). The amplification of tail dependence noted in section 3 causes the high quantiles in figure 14b to increase dramatically. This means that much more capital will be required, even for very low values of ρ, when the variables are tail dependent compared to when they are independent. For fat-tailed lines, more capital is also needed as shown in figures 14c and 14d. These figures show sums of fat-tailed variables (Pareto variables with tail index of 2), with a micro correlation (red, c) and tail dependence, (red, d).

Figure 14: Surplus for independent and dependent sums of lognormal variables (a, b) and Pareto variables (c, d); mutual correlation 0.01 realized with normal copula (a, c) and gumbel copula (b, d)

a

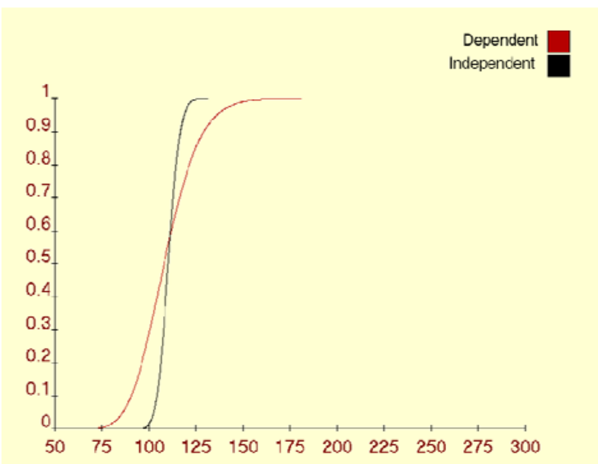

b

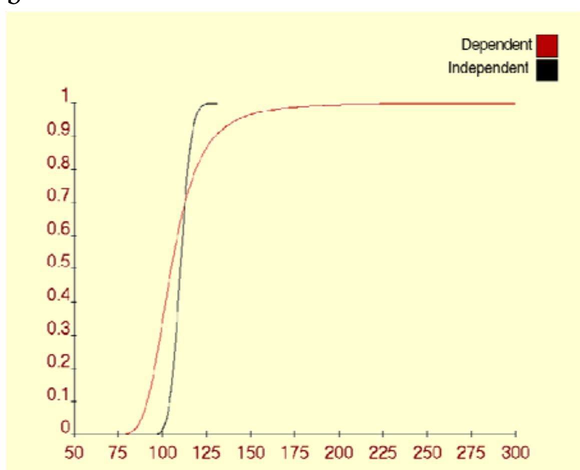

c d

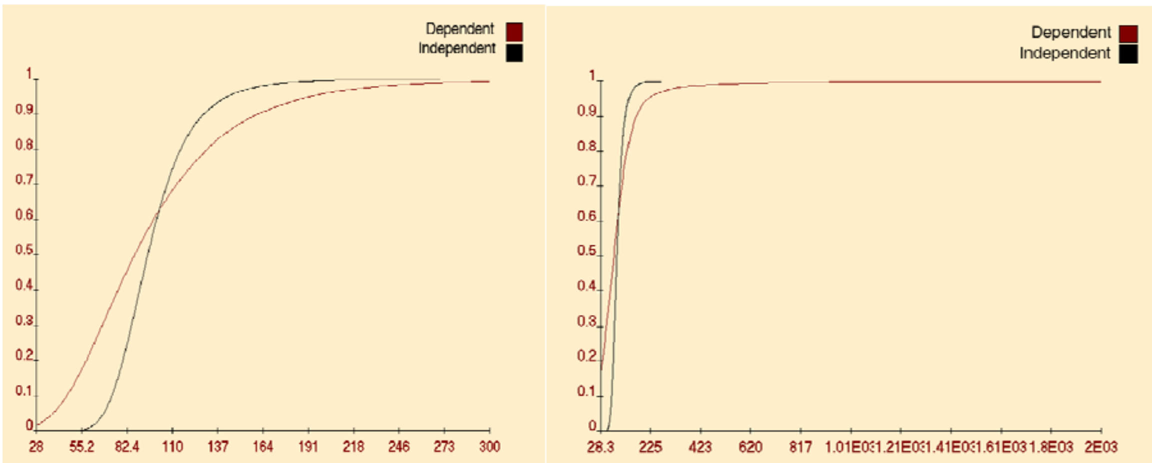

In these cases of catastrophic and dependent insurance lines, maintaining the same probability of solvency means increasing premiums, either to raise capital directly to meet the surplus requirement or to finance the purchase of reinsurance. Two potential difficulties emerge. First, state regulators may limit q. It has been noted by insurance scholars that state insurance commissioners in the US—who have the power to regulate premiums—tend to weight low prices and availability of policies more heavily than solvency considerations or management of catastrophe risk (Klein and Wang 2007). If insurers cannot charge prices that they feel are sustainable, they will leave the market (Klein 2005). This creates the need for so-called residual market mechanisms. These are programs set up by states to provide insurance policies to those people who cannot find a policy in the voluntary market. Many residual market mechanisms have underpriced insurance, perhaps out of ignorance of the fat tailed or dependent nature of a risk, due to a lack of foresight, or from political pressure, creating higher exposures for states. Florida has been particularly troubled in this regard.

Second, and more problematically for insurance markets, homeowners might not be willing to pay the higher premiums required for catastrophic and dependent lines. To get a first-order estimate, assume the homeowner's utility function is given by ln(wealth). Assume also that the homeowner fully insurers or does not insure at all. Of course, in reality, insurance is not an all or nothing decision, and the homeowner can choose to partially insure. Indeed, when $q > p$, from equation (8), a risk averse homeowner will choose $\alpha_i^* < L_i$.[9] Still, to get a better understanding of how catastrophic and dependent risks could lead to a break down in the insurance market, assume, for sake of illustration, that the homeowner chooses $\alpha_i^* = L_i$, and that the loss is equal to the value of their home. The homeowner faces two constraints. The first is that the utility from insurance must be greater than the utility without insurance or:

(12) $\ln(w_i - L_i q) \geq (1 - p)\ln(w_i) + p\ln(w_i - L_i)$.

Note that when the homeowner fully insurers, their total wealth is the same with or without a disaster, since loss is fully compensated by the insurance company. Let $q = xp$. This allows us to focus on x, or the multiple of expected loss rate that a homeowner would be willing to pay for insurance. Also, set w_i=1 and let L_i thus be the fraction of wealth that the homeowner's home represents, which we assume is the same as the fraction of wealth at risk. Making these substitutions and solving for x gives what we refer to as the *utility constraint*:[10]

(13) $$x \leq \frac{1-(1-L_i)^p}{L_i p}.$$

Second, the owner cannot spend more on insurance than their disposable wealth, which is simply their total wealth, w_i, minus the value of their home, which is equivalent to the potential loss they face. Again, set q=xp and w_i=1. Thus, their budget constraint is given by:

(14) $L_i x p \leq 1 - L_i$.

---

[9] This can be seen by rearranging the first order condition in (8) to give: $u'(w - \alpha_{i,s}^* q_s) = [p_s(1-q_s)/ q_s(1-p_s)]u'(w- \alpha_{i,s}q_s - L_i + \alpha_{i,s})$. For $q_s>p_s$, the term in the brackets on the right-hand side of the equation is less than 1. This implies that $u'(w- \alpha_{i,s}q_s - L_i + \alpha_{i,s}) > u'(w - \alpha_{i,s}^* q_s)$. Since for a risk averse consumer, u' is decreasing in wealth, we have: $w- \alpha_{i,s}q_s - L_i + \alpha_{i,s} < w - \alpha_{i,s}^* q_s \rightarrow \alpha_{i,s}^* < L$.

[10] Note that if L_i=0, then this expression is undefined. Using L'Hôpital's Rule, the limit of this expression as the loss approaches zero is 1. This makes intuitive sense. As the loss gets smaller and smaller, homeowners will be willing to pay less and less above the expected loss. At the limit, x goes to 1, indicating an unwillingess to pay more than the expected loss (it just so happens that at a loss of zero, the homeowner would not actually insure anyway).

Solving for x gives:

$$(15)\quad x \le \frac{1-L_i}{L_i p}.$$

For a given probability of disaster, we can plot these two constraints to determine the range of values for x and L_i for which a homeowner would insure. This is shown in Figure 15 for $p = 0.1$ (left) and $p = 0.01$ (right). A homeowner will only insure when values of x (y-axis) and L_i (x-axis) are below both curves. As the probability of a disaster decreases, the income constraint relaxes, as seen in 15 (right), since insurance is cheaper for lower probability risks for a fixed loss level.[11] Figure 16 shows our constraints also as a function of p. For a homeowner with more wealth apart from their home, the fraction at risk will be lower and thus the budget constraint will be less binding. On the other hand, the utility constraint becomes more binding since wealthier homeowners will prefer to self-insure. Note that the low slope of the utility constraint for multipliers in the region 1 to 3 means that a small change in this multiplier corresponds with a large shift in fractional wealth at risk for which insurance is rational. Similar results emerge for other risk-averse utility functions.

[11] As the probability decreases, the utility constraint shifts out very slightly. This is because even though as the probabiltiy decreases a homeowner is willing to pay less for insurance, the graphs in Figure 15 are plotting x, not q. For a given q, as p decreases, c must increase.

Figure 15: Homeowner budget and utility constraints; p=0.1 (left) and p=0.01 (right)

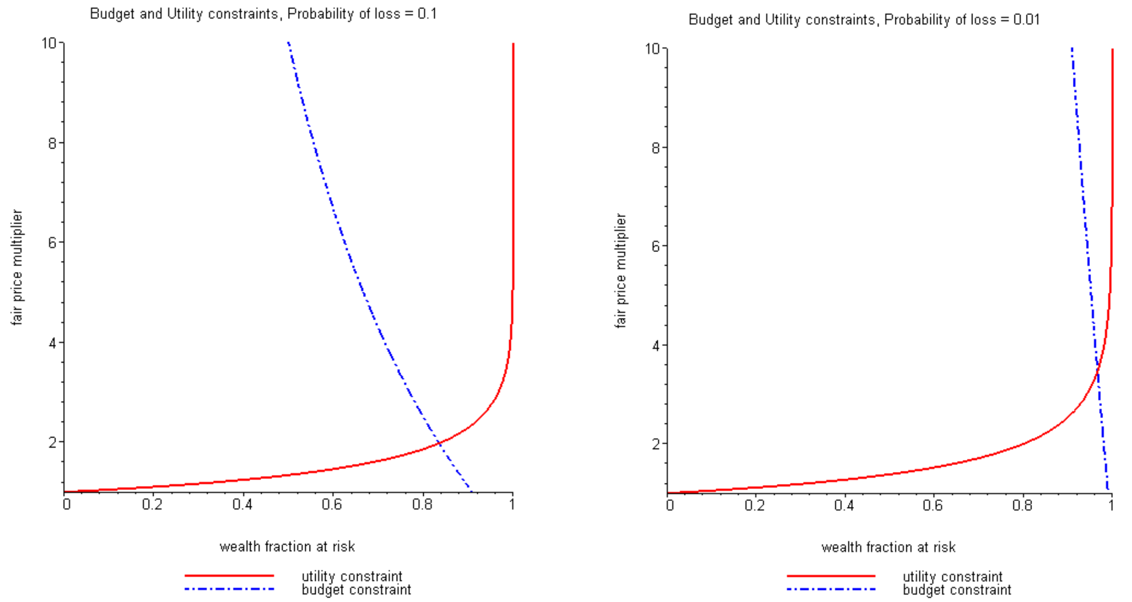

Figure 16: Homeowner budget and utility constraints as a function of wealth fraction at risk and probability of loss

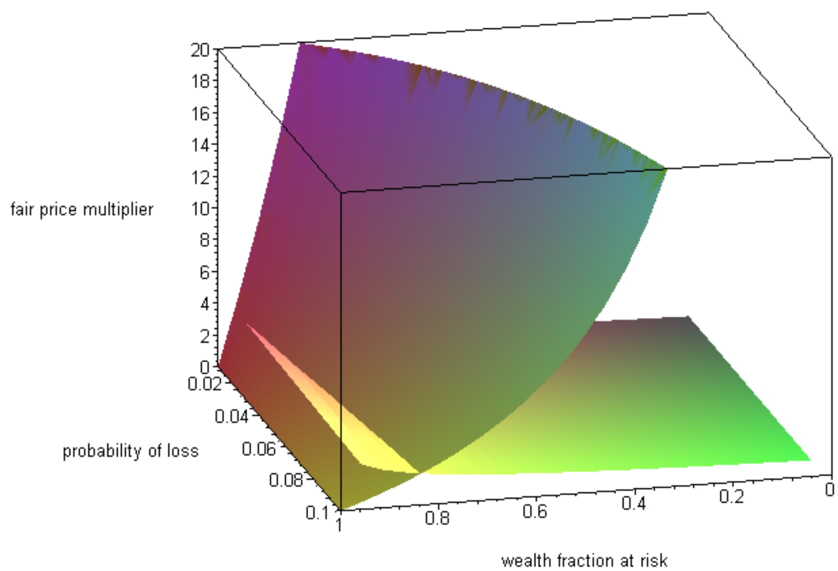

Figures 15 and 16 show that our risk averse homeowner will not pay much more than a couple times expected loss for insurance—and this is when a very large portion of wealth is at risk. How does this compare to the rates needed to be charged by insurers facing risks characterized by our three phenomena? Let us fix $p = 0.01$, and consider an insurance company that writes 10,000 policies. Assume that the insurance company has set $\lambda=0.005$ (this is the draft Solvency II requirement). The portfolio of 10,000 policies will follow a distribution that is approximately normal. If the policies are independent, then it is easy to calculate that the company must charge each policy a fair price multiplier of 1.26 times the expected loss in order to meet their solvency constraint. Referring to Figure 15(b), the utility constraint is satisfied if the fractional wealth at risk is greater than 0.375, and the budget constraint is satisfied if it is less than 0.988. Individuals whose fractional wealth at risk is outside this interval are either too rich or too poor to render insurance a rational option. Suppose that each pair of insurance policies had a small correlation of 0.001. Such a small correlation would be virtually imperceptible; however, the multiplier needed to meet the surplus requirement increases to 1.85 and the interval of fractional wealth at risk for which insurance is rational shrinks to [0.755, 0.982]. If the pair wise correlation rose to 0.01—still negligible by any reasonable standard—the multiplier is 3.58 and interval for which insurance is rational is empty.

While a great deal of literature lays blame for failing insurance markets on the irrationality of those at risk, this simple analysis suggests that under even micro-correlated risks, insurance may not be rational for many people at risk, if solvency constraints are imposed. Since tail dependent and fat-tailed risks also require more surplus, similar results can be expected for these types of risks, as well. The extra loading required to cover fat tailed, micro-correlated, or tail dependent lines can be substantial since for these types of risks, bundling policies does not offer the benefits it does for thin tailed and independent risks. In some cases, therefore, insurance does not provide enough of a benefit to homeowners to pay the required loadings or becomes too expensive as homeowners bump against their budget constraint. In these situations, no amount of homeowner education or outreach activities will increase demand.

Of course, the situation will be exacerbated by the well-documented biases individuals exhibit when evaluating low-probability risks (e.g., Kahneman, Slovic and Tversky 1982). In the simple model presented here, homeowners are assumed to know the probability and loss they face with certainty. In reality, these are often unknown and in forming subjective assessments about low-probability risks, homeowners have been found to exhibit a set of biases, such as treating low-probability risks as zero-probability and being overly optimistic about losses (e.g., Camerer and Kunreuther 1989). For fat-tailed risks, homeowners could also be estimating expected loss based on only a few years of their past experience. If these were years with low losses, it will lead them to incorrectly believe premiums should be low. There is already a documented bias toward individuals assuming that small samples are representative (Tversky and Kahneman 1982); the impact of this will be even more pernicious with fat tails. These types of biases can lead to sub-optimally low levels of insurance, but the simple model here suggest a bigger problem may simply be the inherent difficulty in insuring catastrophic and dependent risks due to the required loadings.

Thus, either because the homeowner does not value insurance at the price offered, or due to income constraints or behavioral biases, the insurance market for catastrophic and dependent risks may break down. Insurance, however, generates positive externalities. Insured home- and business-owners are more likely to have the funds to rebuild and to do so quickly, generating economic spillover effects in the community. As homeowners do not consider the benefits to neighbors of insuring, too few people may insure, suggesting a social interest in helping bring down the cost of insurance for these risks in order to increase take-up rates. If insurance could be provided more cheaply for these risks without threatening the solvency of insurers, it could provide both private and public benefits.

One way the cost of insurance for dependent risks could be brought down is by shifting some of the risk to the financial markets, for example, through the use of a catastrophe (cat) bond. Cat bonds are issued by (re)insurance companies that set up a separate legal structure called a special purpose vehicle (SPV) to issue the bond and invest the proceeds in low-risk securities. Investors in the bond receive the interest on the investment as well as some fraction of premiums paid by the (re)insurer. If the particular catastrophe the bond is designed for, called the trigger, does not occur, investors get their principal back at the end of the time period of the bond. If the trigger occurs, the investors lose their money as it is given to the (re)insurer to cover claims.

There are a few reasons, however, to be skeptical about demand for cat bonds and other insurance linked securities (ILS). The possibility of total loss means cat bonds are usually given a non-investment grade rating, discouraging some investors. This can at least be partially overcome by issuing the bond in tranches (Michel-Kerjan and Morlaye 2008), where, for example, one tranche has the principle at risk and the other does not, so it receives a higher rating (in this layer, repayment of the principle may be delayed if a disaster occurs). Second, the modeling used for the pricing of ILS is often difficult for lay people to follow, which again might discourage some investors. Finally, it had been argued cat bonds would be attractive to investors since they were likely uncorrelated with the market. This cannot be assumed to be true; a cat bond failed to meet an interest payment when Lehman Brothers failed (Hartwig 2009). It likely appears that cat bonds and the stock market generally may be tail dependent themselves.

It has been argued that another mechanism for bringing down the cost of catastrophe insurance is allowing for tax-deferred catastrophe reserves (Harrington and Niehaus 2001; Milidonis and Grace 2007). Currently, insurance companies must keep catastrophe funds in general surplus accounts where they may be depleted, regulators may treat the extra funds as reasons for more stringent price regulations, and additional surplus and the investment earnings on it are taxed as income (Klein and Wang 2007). Davidson (1998) uses a simple example to explain the problem with taxing capital used to pay catastrophe losses as income. He assumes, for simplicity, an insurer charges $1,000 (the expected loss) of a 1/10 possibility of a $10,000 loss. If the $1,000 is taxed in year 1—assuming no losses—then, at a 35% rate, only $650 remains. This continues through year 9. Then in year 10, assume the loss occurs. The insurer does not have enough to cover the loss. Some can be recouped through carry-back provisions, and in reality, premiums will increase to account for the taxation, but it means costs are higher.

To help overcome this problem, insurers could choose to allocate cat funds to a trust or separate account where they could accumulate tax-free, and only be withdrawn for payment of claims following pre-defined triggers (Davidson 1998). The trigger could be based on specific events or firm-specific catastrophic loss levels. Such a reserve could have firm-specific caps based on tail Value-at-Risk (VaR) measures or some other tool. There could be provision for use of the funds for other purposes conditional on paying the appropriate taxes.

These mechanisms can potentially lower the cost of policies to a point where more homeowners will choose to insure. For fat-tailed and tail dependent risks, however, the expected losses that occur with probability less than λ could be quite substantial and these risks are not managed by an insurer using VaR approaches. That is, the insurance company will keep the solvency probability at λ, but does not care, given that their loss exceeds VaR_λ, by how much it does so. There is no pressure from consumers for insurers to manage risks beyond VaR_λ either, since they are protected by consumer guarantee funds—state programs that pay the claims of insolvent insurers. As an example of the perverse incentive the funds create, it has been found that property-casualty insurers have excessive premium growth the year before insolvency (Bohn and Hall 1999). This is because premiums can be a cheap form of borrowing, since claims will be paid later, that can be used to make risky investments. This creates, as Bohm and Hall note, a “heads-I-win-tails-somebody-else-loses” situation. The excess tail risk is essentially transferred to the public.

If regulators chose to do so, the approaches just discussed for lowering the costs of insurance could also be used to cover events beyond the solvency probability. Tail dependent losses, for example, may create damages in excess of the VaR_λ. Cat bonds can be used to cover these losses. For instance, in April 2007, Swiss Re structured a cat bond covering flood risk in the UK that is triggered if there is flooding in at least four of fifty reference locations (Swiss Re 2008). Such specially designed ILS may prove useful in managing very low-probability, high-magnitude losses.

Another approach to address high loss layers or to overcome the breakdown in the insurance market when it is too severe to be handled by the mechanisms just discussed, is for the federal government to act as a reinsurer of state residual market mechanisms or private insurers. Lewis and Murdoch (1996) have argued for federal excess-of-loss contracts for very high layers of coverage because (re)insurers have not been able to effectively diversify catastrophe risk and the government can offer greater inter-temporal diversification. The authors suggest industry-wide triggers to minimize moral hazard. Litan (2006) proposes a pre-funded program administered by a quasi-independent group in the Treasury Department to provide reinsurance to insurance companies or state plans. Premiums would be risk-based and could be reduced if the entity adopted risk reduction measures. Government would only cover losses above a relatively high threshold, with lower level losses being covered by private insurance and state and local governments. This layered approach, it is argued, will provide incentives for all parties to adopt mitigation measures. Cutler and Zeckhauser (1999) note that federal reinsurance could allow private market functioning for smaller losses, internalize some loss costs that the government has control over, and raise revenues for a service the government may provide anyway.

Pricing for a federal program designed for losses characterized by our trinity would have to be set carefully and vary by state and/or exposure of the insurer to avoid moral hazard. Consider Figure 17, which uses the SHELDUS data to examine property damages for the states of Alaska (a), Florida (b), and Montana (c). In Alaska, the tail of damages is extremely fat, with infinite variance. Florida too, has fat-tailed losses, and they can get to be an order of magnitude higher than Alaska. The magnitude of losses in Montana, on the other hand, is quite smaller, and thin tailed, as indicated by the decreasing mean excess plot after about $4 million. If each state did not face prices appropriate for the risk, there could be inequitable cross-subsidization across states and insufficient motivation to adopt mitigation.

Figure 17: Mean Excess of Property Damage for (a) Alaska, (b) Florida, (c) Montana

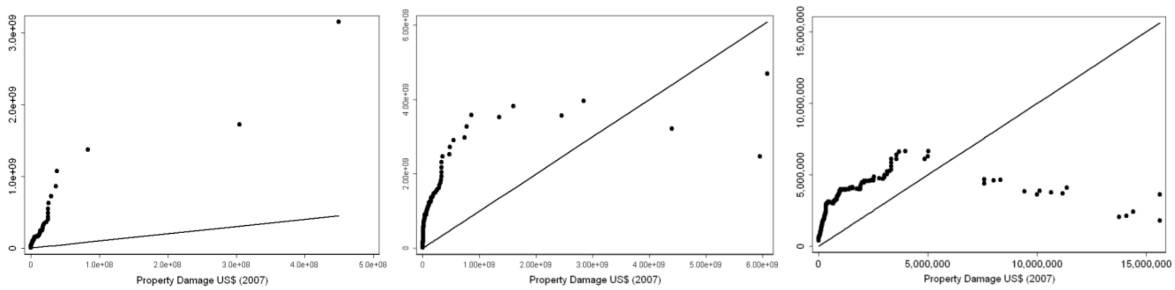

Mitigation should be a priority since it can bring down damage levels. The government should prefer to lower damage costs as it reduces ex-post aid, minimizes economic impacts, and reduces pain and suffering. Homeowners, too, given the uninsurable costs of a disaster, would prefer, all else equal, to protect their home against damage rather than face a risk of disaster and insure. Many mitigation measures have also proven to be cost effective, paying for themselves in reasonable timeframes. Far fewer homes are fortified against disasters, however, than this would suggest. Homeowners may fail to mitigate for a variety of reasons: they underestimate or dismiss the probability of a disaster, are myopic, do not see or understand the mitigation in place when purchasing a home, do not have the necessary upfront costs, do not consider the benefits to their neighbors, and/or are not as concerned about disaster losses due to federal aid (Lewis and Murdoch 1999; Kunreuther 2006).

Creative solutions for overcoming these barriers would provide public benefits. One option is to require buildings to conform to stringent building codes when a property is purchased. The cost of the mitigation can be incorporated into the mortgage. Another option is tax breaks for homeowners who mitigate, as is currently being considered by Congress, or increasing funding to mitigation grant programs run by the Federal Emergency Management Agency or states. A more novel approach could be to have government finance fortification, with homeowners repaying only when a certain magnitude event occurs or slowly over time, like an interest-free loan. Mitigation to thin tails and de-couple risks, however promoted, is likely the preferred risk management approach to catastrophic and dependent risks.

6. Conclusion

Examination of several datasets of damages from natural disasters in the United States shows that they are often characterized by fat tails, tail dependence, and/or micro-correlations.

We cannot effectively manage what we do not measure, however, and neglect of these phenomena can lead insurers—whether the insurer is a private company, a state, or the federal government—to unwittingly expose themselves to much greater risk levels. This paper is a first step toward the development of tools to detect, measure, and analyze these three phenomena. If the analysis in this paper is correct, the insurance sector may already be bumping up against the limit of securitizing natural disaster risks. The current structure of private insurance markets in the US may have outlived itself, challenging economists and risk analysts to come up with new innovative ways of harnessing market forces to combat risk.

References

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.

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.

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., 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. SCOR-JRI Conference on Insurance, Reinsurance and Capital Market Transformations. Paris, France. September 20-21.

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

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." Journal of Risk and Uncertainty **33**(1): 73-100.

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

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

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

Ladoucette, S. A. and J. L. Teugels (2006). "Analysis of Risk Measures for Reinsurance Layers." Insurance: Mathematics and Economics **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." Earth-Science Reviews **87**: 94-112.

Lewis, C. M. and K. C. Murdock (1999). Alternate Means of Redistributing Catastrophic Risk in a National Risk-Management System. The Financing of Catastrophe Risk. 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." Journal of Hydrology **322**: 168-180.

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

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

McNiel, A. J., R. Frey and P. Embrechts (2005). Quantitative Risk Management: Concepts, Techniques, and Tools. 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"." The Geneva Papers **33**: 153-176.

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

Mossin, J. (1968). "Aspects of Rational Insurance Purchasing." The Journal of Political Economy **76**(4): 553-568.

Nelson, R. (1999). An Introduction to Copulas. New York, Springer.

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

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

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

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. Judgment under Uncertainty: Heuristics and Biases. D. Kahneman, P. Slovic and A. Tversky. Cambridge, UK, Cambridge University Press**:** 23-31.

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

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