DISCUSSION PAPER

Climate Change and Risk Management

Challenges for Insurance, Adaptation, and Loss Estimation

Carolyn Kousky and Roger M. Cooke

RESOURCES
FOR THE FUTURE

Climate Change and Risk Management:
Challenges for Insurance, Adaptation, and Loss Estimation

Carolyn Kousky and Roger M. Cooke

"When sorrows come, they come not single spies
But in battalions."
*King Claudius, Hamlet Prince of Denmark, Act 4 Scene 5*

Executive Summary

Adapting to climate change will not only require responding to the physical effects of global warming, but will also require adapting the way we conceptualize, measure, and manage risks. Climate change is creating new risks, altering the risks we already face, and also, importantly, impacting the interdependencies between these risks. In this paper we focus on three particular phenomena of climate related risks that will require a change in our thinking about risk management: global micro-correlations, fat tails, and tail dependence. Consideration of these phenomena will be particularly important for natural disaster insurance, as they call into question traditional methods of securitization and diversification.

Global micro-correlations between variables—such as investment products, insurance

policies, or catastrophe bonds—obtain when all variables in a relevant class have small positive correlations that would be imperceptible if the variables were considered in isolation. When such variables are aggregated, the correlations of the aggregations will balloon. If micro-correlations are ignored when calculating reserves under a Value-at-Risk management regime, risks will be woefully underestimated, as seen in the figure to the left. A manager may think his residual risk is 0.1%, when in reality it is much higher. Moreover, risk underestimation grows with the number of products securitized.
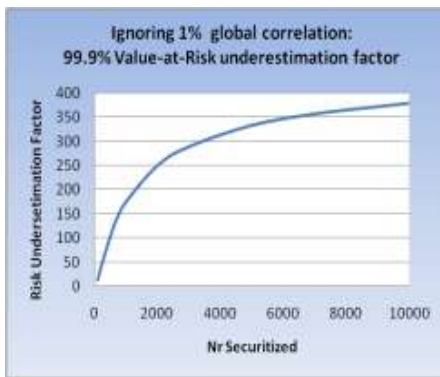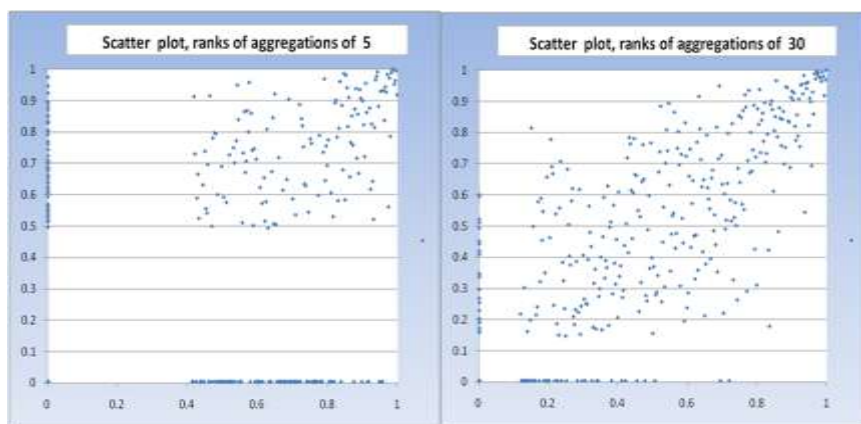
Fat tailed loss distributions resulting from uncertainty in parameters of 'thin tailed' distributions have recently captured the imagination of economists, thanks to the work of Martin Weitzman. Catastrophe modeling firms study the ways in which extreme events compound damages to create fat tails and damages from many extreme events have been shown to fit fat tailed distributions. Financial risk managers identify fat tails by "mean excess" plots showing for each value x, the expected amount by which damage variable X is worse than x, given that it is
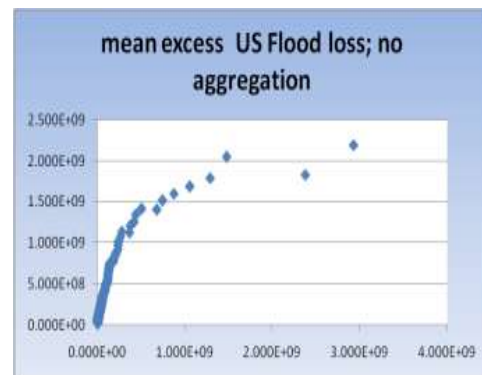
at least *x*. Thin tailed distributions have non-increasing mean excess, but fat tails push this excess ever upward. This increasing mean excess is seen in the figure here, which is based on flood claims in the United States between1980 and 2008. Fat tails increase the solvency risk, thus mooting the 'actuarial fair price' of insurance. Tails may become fatter under aggregation, but they may not; and indeed aggregation is sometimes an effective abatement strategy.

Tail dependence is distinct from simple correlation and tail fatness. It occurs when dependence concentrates in high values. Tail dependence may or may not cohere with tail obesity, and may or may not be aggravated by aggregation. Analyses of flood damages from the state of Florida reveal a tail dependence that is amplified by aggregation. The leftmost picture shows a percentile scatter plot of monthly flood damages for two randomly chosen aggregations of five distinct Florida counties (months without damage give the points on the axes). The right picture shows a similar plot for 30 distinct counties. The amplification of tail dependence has obvious implications for securitization: aggregating policies and marketing tranches of the aggregations may be spreading this dependence over many unsuspecting investors, thereby conscripting a 'battalion of sorrows.' On the other hand, conditioning on the 'southern quadrant' would allow private market securitization to function.

This preliminary project report introduces these three phenomena, provides examples and discusses the implications for estimating losses from extreme events, adaptation policy, and natural disaster insurance markets. We offer an initial foray into the mathematical properties of mirco-correlations, fat tails, and tail dependence; conduct an exploratory analysis of flood and crop insurance claims; and offer some tentative policy recommendations.

Key Words: tail dependence, micro-correlations, fat tails, damage distributions, climate change

JEL Classification Numbers: Q54, G22, C02

Contents

Climate Change and Risk Management:
Challenges for Insurance, Adapatation, and Loss Estimation

Carolyn Kousky[∗] and Roger M. Cooke

1. Introduction

Adapting to climate change will not only require responding to the physical effects of global warming, but will also require adapting the way we conceptualize, measure, and manage risks. Climate change is creating new risks, altering the risks we already face, and also, importantly, impacting the interdependencies between these risks. In this paper we focus on three particular phenomena of climate related risks that will require a change in our thinking about risk management particularly vis-à-vis natural disaster insurance: global micro-correlations, fat tails, and tail dependence.

Global micro-correlations are very small—even undetectable—correlations between variables, such as insurance policies. If unaccounted for, these micro-correlations can undermine common diversification strategies, as shown in sections three and four of this paper. Fat tailed damage distributions are those in which the probabilities of ever more serious damage decrease slowly relative to the extent of the damage. The implications for climate policy are now under active discussion. Less understood is tail dependence, or the possibility that bad events happen together. Research on these topics is needed on three distinct fronts. First, further mathematical research is needed on the properties of and relations between these three issues. Second, data analysis will be required to determine to what extent these three phenomena can be detected and measured in damage data. Finally, policy research on the design and implementation of new risk management approaches will be necessary based on the results of the first two research efforts. This paper offers initial forays into all three of these research areas.

The next section of the paper provides some climate change related examples of mirco-correlations, fat tails, and tail dependence. To motivate our interest in these issues, section three discusses the implications of these phenomena for loss estimation, adaptation policy, and particularly natural disaster insurance markets. Section four formally defines and discusses the mathematical properties of mirco-correlations, fat tails, and tail dependence. Section five analyzes US data on flood and crop insurance claims, demonstrating the presence of positive mirco-correlations in damages across counties, the fat-tailed nature of the distributions, and the

presence of tail dependence that grows as claims are aggregated. This initial empirical work suggests techniques for detecting the presence of these phenomena. Section 6 concludes with policy recommendations. A mathematical appendix shows how aggregation can amplify tail dependence in common deterioration models.

2. Climate Change: Altering Tail Behavior and Dependencies

As stated in the introduction, this paper is concerned with three particular phenomena of climate change related risks: global micro-correlations, fat tails, and tail dependence. Global micro-correlations are tiny, positive correlations between all variables under consideration. These correlations could be so small individually that they would mostly go undetected. Even correlations this tiny, however, can pose problems to risk managers, as discussed in the following two sections. This type of global correlation could arise when every variable is correlated with some latent variable. An example comes from El Niño, which induces correlation between climate events in various parts of the globe. In an El Niño year, precipitation is likely to be more extreme in California, leading to mudslides and floods; nutrient-poor water is likely to cause fish catch declines in Peru; and drier conditions are more likely in Australia, increasing the chance of bushfires. An increase in sea level rise worldwide or a change in the strength of hurricanes could also potentially introduce micro-correlations across losses in geographically diverse areas.

Alternatively, and more realistically, instead of there being a common, small, positive correlation among all the variables, each pair of variables may have a different correlation: small or large, positive or negative. If the average is positive, then similar problems arise. For the US flood and crop insurance data examined in section five, for instance, most correlations are statistically indistinguishable from zero, but their average is positive. In section four, we show that this is the typical situation; if average variances exist, then average covariances must be greater or equal to zero, and "equals zero" is the limiting case. In large finite data sets we will typically see "greater than zero."

The second concept we consider is fat tails. There is growing recognition among economists that, when it comes to climate change, "the tails matter." There is already evidence that damages from many disasters follow distributions with thicker tails (e.g., Malamud and Turcotte 2006). Some research suggests that climate change may be directly fattening the tails of the distributions of many extreme events (e.g., US Climate Change Science Program 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.

Finally, it is not just thick tails that matter for climate change policy, but also tail dependence, or how likely bad outcomes are to occur together. This latter problem has received surprisingly little attention among climate policy scholars and practitioners. A simple example of tail dependence comes from considering the damage distributions associated with computer networks and some highly infectious tropical diseases, which are predicted to spread as the climate warms. Events in the tail of the damage distribution associated with potential computer network problems include network failure and malicious attacks. Events in the extreme tail of the tropical diseases damage distribution not only include rising infection and mortality rates, but

also mass quarantines. These negative outcomes, however, are not independent. If people are quarantined at home, the number of people telecommuting will increase dramatically, stressing computer networks and leading to failures and vulnerabilities that could be exploited. That is, if society is in the tail of disease impacts, it is also more likely to be in the tail of computer security damages.

Another example of tail dependence is seen between property insurance claims and car insurance claims. For low levels of claims there is not much correlation between the two insurance lines. For very large storms, however, there is a dependency between the two variables. When there are many property claims, there tend to be many motor vehicle claims as well. This has been observed for insurance data covering over 700 storm events in France (Lescourret and Robert 2006). What about the flood and crop claims we examine, might they be tail dependent? One can imagine extreme events that would create tail dependence between these losses; a Karkatoa sized volcano would lower global temperatures, acidify the upper atmosphere, and also potentially cause large tsunamis, leading to simultaneous extremes of crop damage and flood damage. The 29 years of US flood damage and crop loss analyzed in section six, however, show no signs of dependence.

Tail dependence among distributions related to climate change could arise in two distinct ways. First, there could be a causal link between two variables, such that when one variable takes on an extreme value, it pushes the other variable to do so, as well. Pandemics and network security provide an example. For policy, it is important to note that this type of dependence will often be unidirectional. That is, a disease pandemic increases the likelihood of network failure, but not vice versa. Tail dependence can also arise because a third variable which is normally dormant or latent pushes both variables into extremes when activated. A very large volcano could correlate flood damage and crop failure, as just mentioned.

In theory, there could be tail dependence across more than two variables, as suggested by the El Niño example. Another telling example of dependencies across multiple loss distributions as driven by climate change is the European heat wave in the summer of 2003. The temperatures for the summer were extremely far out in the tail of the temperature distribution (Schär, Vidale et al. 2004), leading to extreme losses across multiple sectors. The heat wave led to uninsured crop losses of nearly $12.3 billion, extensive fires across Europe burning 647,069 hectares, the shutting down of nuclear power plants in France from lack of river water for cooling, soaring electricity spot market prices, rockfalls from melting permafrost, decreased yields and lower quality harvests, excess deaths of between 22,000 and 35,000 across Europe, and excess mortality from higher ground-level ozone and particulate matter concentrations (De Bono, Giuliani et al. 2004; Schär and Jendritzky 2004; Stedman 2004). With more heat waves, events in the tail of the distributions related to mortality, crop yields, wildfires, and electricity pricing are more likely to occur together.

The European heat wave and the multiple impacts of El Niño events demonstrate that the climate system is linked globally in ways that create multiple dependencies. As yet another example, there is a strong correlation between Atlantic tropical cyclone activity and atmospheric dust from the Saharan Air Layer (Evan, Dunion et al. 2006), as well as strong correlation between the landfall of intense hurricanes in the US and rainfall in the Western Sahel region of West Africa, potentially because both are driven by global oceanic thermohaline processes (Gray

1990). There are other dependencies that are not well understood, such as multiple events tending to occur together, sometimes referred to as "clustering" (RMS 2008). Active hurricane seasons in 1995 and 2004-2005 are good examples, as are the two severe windstorms that caused damage in Europe in 1999 (Riker 2004). These intricate dependencies should make us cautious of simple or ad hoc consideration of climate related risks.

3. Implications for Loss Estimation, Adaptation, and Insurance

There are three areas where micro-correlations, fat tails, and tail dependence as introduced or heightened through climate change could have significant implications: loss estimations, adaptation policies, and of particular emphasis is in this paper, insurance for natural disasters.

3.1. Loss Estimation

Researchers and policymakers have sought estimates on the aggregate costs of climate change to use in considering various mitigation options, as well as estimates of regional or local climate change damages in order to improve planning. The costs of particular extreme events— and how they may be altered under climate change—are also of interest, particularly to insurance companies. All these various exercises in estimating aggregate losses, whether for global changes in mean temperature or for individual extreme events, will be underestimates if fat-tailed distributions are not used when they are appropriate.

If damage variable X is fat tailed, then for k larger than some threshold, mean values of X^k will be infinite. On any finite sample there will always be a finite sample average, but this average increases without bound as sample size increases. The US flood loss data analyzed in section five appear to have a finite mean but infinite variance. The sample mean is an unbiased estimate of the true mean, but with infinite variance, no matter how many samples we draw, the estimate doesn't improve. To appreciate how unreliable the sample mean is in such situations, Table 1 shows five averages based on 100,000 samples of a fat tailed Pareto distribution (see section 4.2 for a discussion) with true mean = 10. The five averages are generated with different random seeds.

100,000 samples of Pareto $1/(1+x)^{1.1}$	
True mean = 10	
random seed 1	5.54
random seed 2	6.26
random seed 3	6.58
random seed 4	6.20
random seed 5	5.42

Table 1: Five Sample Means of 100,000 Pareto Variables with Cumulative Distribution Function 1-1/(1+x)$^{1.1}$

Catastrophe modelers who dutifully report sample means and variances from Monte Carlo simulations without first verifying that the damage tails are sufficiently thin may be

deceiving themselves. The situation is worst for 'super fat' tailed distributions, where the mean is infinite. In these cases, Monte Carlo damage estimates will always be infinitely wrong.

Simulations that neglect dependencies can also be wrong. Too often, impact modeling for changes in climate is performed independently for various sectors of the economy, or specific ecological systems, and interactions between impacts are often ignored (Tol 2002). Similarly, analysis of different events is often done in isolation, without consideration of the correlations between events (Muir-Wood 2004). Many systems are coupled, however, and models of climate change impacts or the costs of disasters that do not address such interconnections, and the possible creation of new interdependencies as the planet warms, could significantly miss-estimate impacts.

For example, insurance companies rely on Monte Carlo catastrophe models, often run by private companies, to assess the expected losses in an area where they have exposure (Grossi, Kunreuther et al. 2005). Many of these models have neglected dependencies across loss distributions leading to an underestimate of losses. Hurricane Katrina estimates, for instance, were too low, in part because there was more damage from the secondary consequences than from the original event (RMS 2006). Catastrophe modelers are attempting to address this and to incorporate another type of tail dependence, event "clustering," whereby severe events occur in close succession (RMS 2006). These improvements in the models will aid companies in developing more accurate pricing and diversification strategies.

3.2. Adaptation

Consideration of tail dependence may suggest policies that could reduce vulnerability to changes in extreme events. Spatio-temporal tail dependence has been calculated for precipitation extremes (Kuhn, Khan et al. 2007). In some situations, though, tail dependencies are not simply a product of the global climate system but have been induced by human actions and it may be possible to mitigate these risks by identifying the underlying cause coupling the tails of the loss distributions.

When we are aware of and understand the nature of the dependence between loss distributions, there will be times when adaptation policy might effectively de-couple risks. For example, take the 1906 earthquake in San Francisco. This event showed the tail dependence between earthquakes and fires; it was both an extreme earthquake and a severe fire that damaged the city. The fire burned for 3 days and devoured 28,000 buildings (Steinberg 2000). The cause of this dependence was gas and water pipes that could not withstand a high magnitude quake. If gas pipes break, fire is more likely, and if water pipes break, flames cannot be controlled, raising the likelihood of a conflagration. De-coupling the tails of these two risks involves the design and installation of pipes that can withstand extreme earthquakes. Surprisingly, there are still many water pipes in the Bay Area that have not been upgraded to withstand earthquakes, although retrofitting is currently ongoing.

Thinking through the response of individuals and all sectors of the economy to an extreme event can also help us identify and understand tail dependencies. An example comes from Hurricane Katrina. Models of potential losses assumed that the pumps in New Orleans would keep flooding in the city to a minimum. However, the extreme nature of Katrina led to an evacuation of people, including pump managers, as well as a power outage, reducing pumping

capacity and leading to much more extensive flood damage in the city than was expected (RMS 2005). This suggests that policies aimed at adapting to the potential increases in extreme weather events from climate change should consider adding redundancies in protective systems and thinking through the impact of not only loss of personnel and power, but damage to all critical infrastructure, including communications.

3.3 Natural Disaster Insurance

Finally, neglect of micro-correlations, fat tails, and tail dependence could have the largest impact on the insurance market for natural disasters (for more on the insurance of natural disasters, seem for example: Woo 1999). The last several decades have seen a rise in the economic costs of disasters, and, of particular concern to the insurance industry, in insured losses (Zimmerli 2003; Wharton Risk Management and Decision Processes Center 2008). While much of the rise in disaster losses has been due to increased exposure in hazardous areas—more buildings and people in harm's way—several weather-related extreme events are predicted to increase in severity and/or frequency with climate change. There is increasing concern about what impact this will have on the insurability of natural disasters (e.g., Mills 2005; Kunreuther and Michel-Kerjan 2007; Charpentier 2008). If risk is increasing over time, such that insurers do not believe they can accurately estimate expected losses, premiums will be higher to compensate for the uncertainty (Kunreuther and Hogarth 1992). If the premiums necessary for insurers to cover a disaster in a climatically changed world are greater than homeowners and businesses are willing or able to pay, the private insurance market will collapse. The impact of climate induced micro-correlations, fat tails, and tail dependence across different lines of insurance or different regions of the world is a further consideration for insurance companies in assessing response strategies to climate change.

Insurance and reinsurance markets rely on the law of large numbers. As independent risks are bundled, fluctuations around the mean die out. That is, for a sum of independent risks, the ratio of the standard deviation to the mean is dramatically smaller than for one risk. With positive dependence, the independence assumption of the law of large numbers is violated and the benefits from holding multiple policies are not necessarily obtained. Laymen don't always appreciate how aggressive such independence assumptions are; a simple numerical example makes this clear. A normal variable (i.e. the familiar bell-shaped type) with mean 10 and standard deviation 5 has a 0.023 chance of taking a negative value. The sum of 10 such variables has mean 100 and standard deviation 15.81, *if the 10 variables are independent.* This sum has a 1.3×10^{-10} chance of taking a negative value. The probability of a negative outcome went down by a factor of 179 **million**, just by adding ten independent copies together. If you add 100 copies, the probability drops to 2.8×10^{-89}. This all happens because the mean goes up with the number summed, whereas the standard deviation goes up with the square root of that number. This is what makes insurance possible.

When risks are correlated, however, the standard deviation relative to the mean does not fall as quickly; when they are perfectly correlated it does not fall at all. Suppose our 10 variables were completely correlated. In that case, the standard deviation of the sum goes up as fast as the mean; instead of 15.81 for the sum of 10 variables, it is 50, and the probability of a negative value remains 0.023. With perfect correlation, summing does not reduce the risk. Perhaps counter intuitively, even a small correlation can prevent bundling from reducing risks. When the

correlation between the variables is equal to the average between flood damages in different Florida counties (0.213; see section five), the probability of a negative outcome with a sum of 10 is not 1.3×10^{-10}, but 1.07×10^{-4}, larger by a factor of nearly one million.

When a disaster hits one home, it will have also hit other homes in the area. Insurers deal with such local correlations by presuming that disasters in disparate regions are independent. If this is true, insurers can again put the law of large numbers to work and swap risk from different areas or purchase reinsurance. With the data analyzed in section five, correlations in US county flood claims are almost always statistically indistinguishable from zero. That doesn't mean they ARE zero; indeed they are not. Bundling actually magnifies these small correlations and eventually defeats the attempt to reduce risk through securitization.

We provide a stylized example to show how this could happen. Suppose an insurance company sells a one-year catastrophic flood insurance policy paying $10,000 in damages should a flood occur. Assume the historical probability of this event is 0.001 per year. The 'actuarial fair price' for this policy would be $10 (probability times loss). The insurance company will, of course, add on administrative costs, a solvency premium, and a profit margin to the premium it charges; we ignore this here. Following the regulatory requirements for insurance companies in the EU, termed Solvency II,[1] the firm's risk management regime stipulates that cash reserves must be capable of covering claims with a 99.9% probability.[2] In this case, the company must keep $10,000 in reserve to guard against insolvency, since this is the loss that would occur with probability 1-.999. This gives a ratio of actuarial fair premiums to required capital reserves of 0.001.

Now, suppose the company pools this policy with 10,000 other independent policies with identical probabilities and losses. The company's ratio of actuarial fair premiums to required capital reserves jumps to 0.506 (see below). If, however, these 10,000 policies had on average a correlation equal to that found in section five for flood claims across all US counties (0.041), this ratio sinks to 0.048. Thus, if the company ignored the tiny average correlation between the policies—which is likely since most of the correlations being averaged are statistically insignificant—it would be undercapitalized by a factor 10.

For another take on this phenomenon, consider a company that "securitizes" their risk; that is, the company trades their flood policy for a 1/10,000-th share in a bundle of 10,000 similar policies. If these other risks are independent, then the bundle is effectively normally distributed with mean 10,000×$10, and standard deviation √(10,000)×stdev(one policy) =100×$10,000×√(0.001×0. 999) = $31,606.96. To cover all claims with 99.9% probability, the bundle holder must have a cash reserve of $197,672.90.[3] The insurance company must

---

[1] Solvency II outlines regulations for insurance companies operating in Europe and is intended to create one European insurance market. The three pillars of solvency II are: (1) quantitative requirements, such as the solvency probability, ultimately put at .005, (the earlier value used here, .001, was based on the Basel II protocol); (2) governance and supervision requirements, and (3) transparency and disclosure requirements.

[2] This is the so-called Value-at-Risk (VaR) risk management scheme.

[3] The formula for this calculation is presented in section 4.1.

contribute its share to this reserve, which is the total reserve divided by 10,000, or $19.77. Instead of $10,000, the insurance company needs to set aside only $19.77 ($10 / $19.77 = 0.506). Such are the apparent benefits of securitization.

What might go wrong with this sanguine business plan? Of course, the historical probability of flooding might not apply in present circumstances subject to global warming. That is already receiving attention in the literature (e.g., Milly, Betancourt et al. 2008). We draw attention instead to the threat of global micro correlations which has so far escaped notice. If each of the 10,000 policies had a small correlation, say 0.01 with every other policy, and if the bundle holder receives $19.77 for each of the 10,000 policies, then the chance that the bundle holder is unable to pay claims is not 0.001, but in the neighborhood 0.38 (the calculation is in the next section). The risk is underestimated by a factor of 380.

All of this is troubling for three reasons: (1) such small correlations would likely be ignored, unless one knew how to look for them (and we have not been looking), (2) ignoring these correlation leads to a large underestimation of the risk, and (3) climate change suggests these micro-correlations are quite plausible. Small global correlations would frustrate attempts to achieve diversification by spreading risks across the globe.

Tail dependence could also be troubling for insurers of natural disaster risk. Risk Management Solutions (RMS), a catastrophe modeling firm, has been exploring notions similar to tail dependence relevant to insurers following Hurricane Katrina: "loss amplification," "Super Cats," and "cat following cat" events (RMS 2005; Muir-Wood and Grossi 2008). Loss amplification refers to large losses that induce even larger losses. Specifically related to insurance company losses, amplification could be caused by increases in the costs of repair when demand for rebuilding overwhelms supply, damage that results when repair is delayed causing further deterioration, relaxation of assessment of claims in the frenzied claims processing after a disaster, or political pressure or litigation requiring expansion of coverage (Muir-Wood and Grossi 2008). Super Cats are those catastrophes that are so large that they trigger other hazards, creating "cascading consequences," and leading to losses across multiple lines of business (RMS 2005)—in the language of this paper, this is the result of tail dependence. So-called 'cat-following-cat' events refer to the unidirectional tail dependence discussed earlier, where one disaster of an extreme magnitude is likely to trigger another, such as the 1906 San Francisco earthquake triggering the massive fire.

Hurricane Katrina created increased interest in dependencies across loss distributions since losses from the storm not only included wind and rain damage, but damage from breached 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 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. While not as related to possible impacts of climate change, 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 exposure of an insurance company can be underestimated.

Amplified losses from Katrina have led modelers to attempt to more fully incorporate the correlations across locations, business lines, coverages, and perils for Super Cats in the catastrophe models insurers use to estimate potential losses (RMS 2005; Muir-Wood and Grossi 2008). These interactions can sometimes be surprising, such as the fact that the shortage of refined petroleum from Katrina led two already troubled airline companies in the US to file for bankruptcy since jet fuel prices rose after Katrina (RMS 2005). Another surprising example comes from the Kobe earthquake, which indirectly led to the collapse of Barings Bank as one of their traders had recklessly bet the Nikkei would recover quickly from the quake, but it did not. RMS is also examining whether there are points—phase transitions—where a system moves from roughly independent damages into the tail region of highly correlated damages. One example they give is a forced evacuation, since once evacuation occurs, properties deteriorate, there is a lack of personnel for response and operation of critical facilities, and rebuilding costs rise (RMS 2005).

Tail dependent risks could lead to catastrophic losses, which could threaten the solvency of insurers (for more on catastrophes and insurance see: Zeckhauser 1996). To cover losses in the tail, insurance companies may turn to reinsurance or to the financial markets through insurance linked securities, such as cat bonds. For holding the risk in the tails of distributions, investors would want a higher return or reinsurers would want higher premiums, potentially leading insurers to reduce coverage or charge higher premiums themselves (Kunreuther and Michel-Kerjan 2007). If particularly extreme, there may not even be a market for these levels of losses. Tail dependence could thus be one of the causes of the low levels of reinsurance for high damage levels and the high prices for reinsurance and cat bonds that have been documented (Froot 2001; Froot 2008; Lane and Mahul 2008). (Froot examines other explanations such as the corporate form of reinsurance, moral hazard and adverse selection in the markets (2001).) It is possible, then, that private (re)insurance may fail to operate in the dependent and thus high-potential-loss region of the distribution, suggesting a possible role for government intervention. Events that stress (re)insurance markets worldwide have been termed cataclysms by Cutler and Zeckhauser (1999). The authors suggest a need for financial markets or government involvement in these cases. This is discussed further in section six.

Further complicating the operation of insurance for tail-dependent loss profiles is the likelihood that the existence or extent of the dependence will never be known with certainty. Insurers have been found to charge an "ambiguity premium" to insure losses where the probabilities are uncertain or to fail to offer coverage at all (Kunreuther and Hogarth 1992). Until the relevant tail dependencies for insurance companies are understood, such aversion to ambiguity could cause low levels of supply in (re)insurance markets for catastrophes.

Disturbingly for the market overall, tail dependence has been found between equity returns across insurance markets in different regions of the world. Slijkerman (2006) finds tail dependence between stock losses of insurers, potentially because their loss holdings are similar or due to similarities in investment risks. Sheremet and Lucas (2008) estimate that for Europe-America and Europe-Australia/Asia, about 60% of the dependence between companies is due to correlated losses, as opposed to correlation in investment portfolios. The percentage is higher for America-Australia/Asia. The authors also find some evidence to suggest that this dependence is increasing over time, potentially from climate change. Interestingly, the authors find far less correlation between American and Australian-Asian markets, suggesting some areas where

diversification can occur (Sheremet and Lucas 2008). Even in these areas, however, diversification could be hindered by the fact that the insured portion of losses from natural hazards are concentrated in only a few countries worldwide (Zimmerli 2003). If large losses are tending to be correlated across the globe, this will fundamentally undermine the ability of (re)insurers to diversify natural disaster risk.

4. Definitions and Properties

This section formally introduces the mathematics of global micro-correlations, fat tails, and tail dependence.

4.1 Global Micro-correlations

As discussed in the previous section, it could be that climate change will begin to introduce small global correlations across some loss variables, such as disaster claims. A curious feature, first noticed in connection with aggregating automobile emissions (van Oorschot et al. 2003) is that aggregations of many weakly correlated random variables become themselves strongly correlated. A small overlooked correlation, whether due to global warming or some other cause, could undermine the common risk diversification practices in the insurance and investment industries of securitizing risks. In fact, this same conclusion holds if the *average* covariance is positive. This may be the result of many statistically insignificant correlations, but if the average is positive, then there is trouble.

To see this, first consider variables $X_1,...X_N$ that are independent with the same standard deviation given by σ (assumed to exist). In this case, the standard deviation of $\Sigma_{i=1...N} X_i = \sigma \sqrt{N}$. If, however, there is actually a small correlation $\rho(X_i, X_j) = \rho$, then the standard deviation of the sum becomes:

$$(1)\ \ StDev(\sum\nolimits_{i=1...N} X_i) = \sigma \sqrt{N} \sqrt{1 + (N-1)\rho}\ .$$

The independence assumption underestimates the standard deviation by a factor $\sqrt{(1 + (N-1)\rho)}$. If the variables are not identical then we can replace (1) with a similar expression using average covariances and variances:

$$(2)\ \ StDev(\sum\nolimits_{i=1...N} X_i) = \sqrt{\sum\nolimits_{i=1...N} \sigma_i^2 + \sum\nolimits_{i \neq j} \mathrm{cov}(X_i, X_j)} = \sqrt{N} \sqrt{\bar{\sigma}^2 + (N-1)\bar{c}}\ .$$

where $\bar{\sigma}^2 = (1/N)\sum\nolimits_{i=1...N} \sigma_i^2$, $\bar{c} = 1/(N(N-1))\sum\nolimits_{i \neq j} \mathrm{cov}(X_i, X_j)$. Since the standard deviation is always non negative, we must have $\bar{\sigma} > -(N-1)\bar{c}$; if $\bar{\sigma}$ is finite and this holds for all N, we

must have $\bar{c} \geq 0$. Hence $\bar{c} = 0$ is a limiting case. $\bar{c}/\overline{\sigma^2}$ is approximately, but not exactly, equal to the average correlation $\Sigma_{i \neq j}\, \rho(X_i, X_j) / (N(N-1))$.

Since catastrophe insurance claims for neighboring properties are correlated, insurance companies seek to hold multiple bundles of correlated policies where they believe the bundles are independent. It is plausible, however, that climate change could induce global micro-correlations. Consider the two sums of policies from different regions: $\Sigma_{i=1...N} X_i$ and $\Sigma_{i=N+1...2N} X_i$. If there is actually a tiny, positive correlation between all of the variables due to climate change, the correlation of the two sums will be:

$$(3) \quad \rho\left(\sum_{i=1...N} X_i, \sum_{i=N+1...2N} X_i\right) = \frac{N^2 \rho \sigma^2}{N\sigma^2 + N(N-1)\rho\sigma^2} = \frac{N\rho}{1+(N-1)\rho}.$$

Evidently, this goes to 1 as $N \to \infty$.

To more fully appreciate the consequences of neglecting micro-correlations, consider the implications for a company's capital holdings, as discussed earlier. If μ is the mean of each X_i, then using the normal approximation and ignoring dependence, we would calculate the probability that $\Sigma_{i=1...N} X_i$ exceeds a threshold T as

$$(4) \quad Pr(\textstyle\sum_{i=1...N} X_i > T; X_i\ independent) \sim 1 - \Phi\left(\frac{T-N\mu}{\sigma\sqrt{N}}\right),$$

where Φ is the standard normal cumulative distribution function (σ still assumed to be finite). With dependence, we cannot compute exceedence probabilities without knowing the whole joint distribution; however a back-of-the-envelope estimation would simply modify the independent estimate by substituting the standard deviation (1) into (4), giving:[4]

$$(5) \quad Pr(\textstyle\sum_{i=1...N} X_i > T;) \sim 1 - \Phi\left(\frac{T-N\mu}{\sigma(\sqrt{N+N(N-1)\rho})}\right).$$

Using the back of the same envelope, we can compute the relative error in the exceedence risk by dividing (5) by (4). Figure 1 shows the results when using the values for T corresponding

---

[4] This estimate is actually quite reasonable when ρ is small and the X_i's are almost independent, for the following reason. For relatively small values of N, the correlation given in (2) will be small and summing 2N variables will be like summing independent variables; the result will look normal. For large N, the correlation is almost one, and adding an additional sum of N is like multiplying the first sum of N by 2, which preserves the shape of the distribution.

to the 99.9 percentile under independence, and $\rho = 0.01$. The estimation error is large, even for modest values of N. Thus, the risk according to (5) for $N = 500$ is 0.1, a factor 100 higher than 0.001.
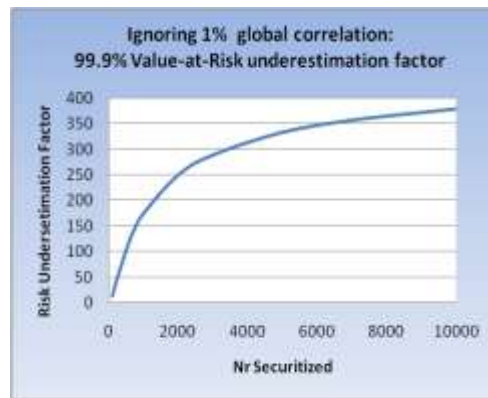

Figure 1: Risk Underestimation Factor

The consequences of the elementary formula (1) are far reaching. Much of the insurance and investment sector runs on "securitization," that is, pooling small investment products, be they insurance policies, mortgages, loans, corporate bonds, etc. so as to reduce the fluctuations of the sum relative to the aggregate mean.[5] This pooled product can then be cut into "tranches" and re-sold. If the products are independent, then the pooling and cutting will yield products with very low risk. However, a small effectively unobservable global correlation means that beyond a certain point, sums of N objects are almost fully correlated, and securitization can no longer reduce risk. Ensuring that private markets will be able to diversify risks evidently requires learning more about micro-correlations. Ironically, upper tail dependence, discussed below, can actually help. If the correlations are in the extreme values of the distributions, then at lower values the variables may be nearly independent.

4.2. Fat Tails

Thanks largely to Weitzman (e.g., 2007), the phenomenon of leptokurtic distributions, commonly known as fat or thick tailed distributions, have captured the imagination of the climate change community. With thick tails, 'extreme outcomes are more likely.' Roughly, a distribution is said to have a fat (right) tail if its density $f(x)$ is asymptotically polynomial as x gets large.[6] The most common example is the Pareto distribution with a density given by:

---

[5] *Recently, various authors express concerns about the limits of diversification, based on the sums of independent, super-fat tailed distributions for which the mean is infinite (Chavez-Demoulin, Embrechts et al. 2005; Ibragimov and Walden 2007). The concern voiced here is based on dependence, not super tail obesity.*

[6] *For a precise definition, see McNiel, Frey et al.(2005).*

(6) $p(x) = \upsilon a^{\upsilon}/(a+x)^{\upsilon+1}$ $\alpha, \upsilon > 0.$

The mean and variance of (1) are (see appendix):

(7) $\mu(X) = a/(\upsilon\text{-}1), \upsilon > 1,$ and

(8) $\sigma^2(X) = a^2\upsilon \, / \, [(\upsilon\text{-}1)^2(\upsilon\text{-}2)], \upsilon > 2.$

The lower the parameter υ, the fatter the tail. For $\upsilon \leq 1$, the expectation is infinite, for $\upsilon \leq 2$, the variance is infinite.

An intuitive diagnostic for assessing tail obesity is the *mean excess*, defined at point x as the mean of $X\text{-}x$, given that $X > x$. If the mean excess is constant, then the distribution has an exponentially decreasing right tail (McNiel, Frey et al. 2005). Upward sloping mean excess plots indicate a fat tail. For the Pareto distribution, the mean excess is linear in x with slope equal to *1/(υ -1)*, indicating a thick tail for *$\upsilon > 1$*.

The relationship between tail obesity and aggregation is not simple. In the appendix, the calculations with L_p symmetric processes show that summing positively correlated leptokurtic variables need not affect tail thickness. Other examples can be found in which independent sums of fat tailed distributions become more obese (Embrechts, Lambrigger et al. 2008). As shown in section five, flood claims and crop loss data have a pronounced fat tail. For crop loss the fatness is attenuated by aggregation, whereas for flood claims this does not appear to be the case. If the tails are not too fat, so that the variance is finite, then aggregation seems to thin the tails. The central limit theorem for sums of independent variables with finite variance says that the sums must approach a normal (thin tailed) distribution. If the variance is infinite, then all bets are off.

4.3 Tail Dependence

Tail dependence 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). In this report, we shall always be concerned with high values or UTD. Technically, upper tail dependence of variables X and Y is defined as the limit (if it exists, which it may not):

(9) $UTD(X,Y) = \lim_{r \to 100} \dfrac{\text{Pr(both X and Y are above their r−percentiles)}}{(100-r)}$

$= Pr(X > r\text{-percentile} \mid Y > r\text{-percentile}).$

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. When thinking about climate change, we may be concerned about tail dependence among distributions of the magnitude of extreme events, such as storm intensity, or among distributions of impacts, such as dollars of damage from a storm event. If *X* and *Y* are leptokurtic, then extreme events could be quite large for either event, but high-loss draws are not necessarily more likely to occur simultaneously. If *X* and *Y* are tail dependent, they are more likely to take on extreme values together, even if their marginal distributions are thin-tailed.

Note that UTD does not depend on the marginal distributions of *X* and *Y*. Thus, if we apply any 1-to-1 transformation to *X*, say $X^* = X^3$, then *UTD(X*,Y) = UTD(X,Y)*. *UTD* has no simple relation to the standard Pearson correlation coefficient. For example, if *(X,Y)* are bivariate normal, with correlation ρ strictly between -1 and 1, the *UTD(X,Y) = 0* (McNiel, Frey et al. 2005). If *(X,Y)ᵢ i = 1,...∞* are bivariate normal with correlation ρ_i, and if $\lim \rho_i \to 1$ as $i \to \infty$, then:

(10) $\text{UTD}(\lim_{i\to\infty}(X_iY_i)) = 1 \quad \neq \quad \lim_{i\to\infty}\text{UTD}(X_i,Y_i) = 0.$

Tail dependence is related to the correlation of extreme event occurrences. For any two events *A* and *B*, the indicator functions *1ₐ, 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:

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

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

(12) $\rho(1_A, 1_B) = \dfrac{P(A \text{ and } B)}{r*(1-r*)} - \dfrac{1-r*}{r*} \to UTD(X,Y) \text{ as } r* \to 1.$

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 multivariate distributions, as they allow us to separate the representation of dependence from the representation of the univariate marginal distributions. 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, as stated earlier, that there is no general relation between fat tails and tail dependence. We could transform the marginal distributions to thick or thin tailed distributions as we like without affecting the tail dependence. These are separate issues. Current research focuses on the relation between tail dependence and

multivariate extreme value copulae (Chavez-Demoulin, Embrechts et al. 2005; Joe, Li et al. 2008).

Just as correlations between sums grow with aggregation, under certain conditions, tail dependence can also grow. 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 $X_1.....X_n$ are symmetrically correlated with a "latent variable." 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 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.

A simple case of this model is treated in the appendix. The X_i's are conditionally independent exponential variables, given the failure rate λ, whose distribution is in the Gamma family with shape factor υ and scale factor a. In this case, the unconditional distribution of each X_i is the Pareto distribution. Any two of the X_i's have Pearson correlation $1/\upsilon$ and have upper tail dependence given by:

$$(13) \quad UTD\big(x_i, x_j\big) = \frac{(\frac{1}{2})^{\upsilon+1}\left[1+\frac{1}{2}+\frac{\upsilon}{2}+ 1/4(\upsilon-1)\right]}{1+\upsilon}; \ \upsilon > 1.$$

If we consider distinct sums of N such variables, the tail dependence expression becomes complicated, but grows with N. Moreover, the lower tail dependence tends toward zero, and the correlation in N-sums conditional on the sums being less than their means, grows very slowly. The thickness of the tails from aggregation does not change appreciably (see appendix).

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; however simulation 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 2 shows four percentile scatter plots using a Gumble copula having weak upper tail dependence. Figure 2a is simply between X_1 and X_2, and the correlation 0.02 is imperceptible. 2b 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 2c and 2d we see the scatter plots of sums of 20 and 40 variables respectively. Now the upper tail dependence becomes quite evident. These plots suggest that the way to detect micro-correlations and tail

dependence is to look at disjunct sums, an idea we will exploit in studying flood and crop insurance data in the next section.
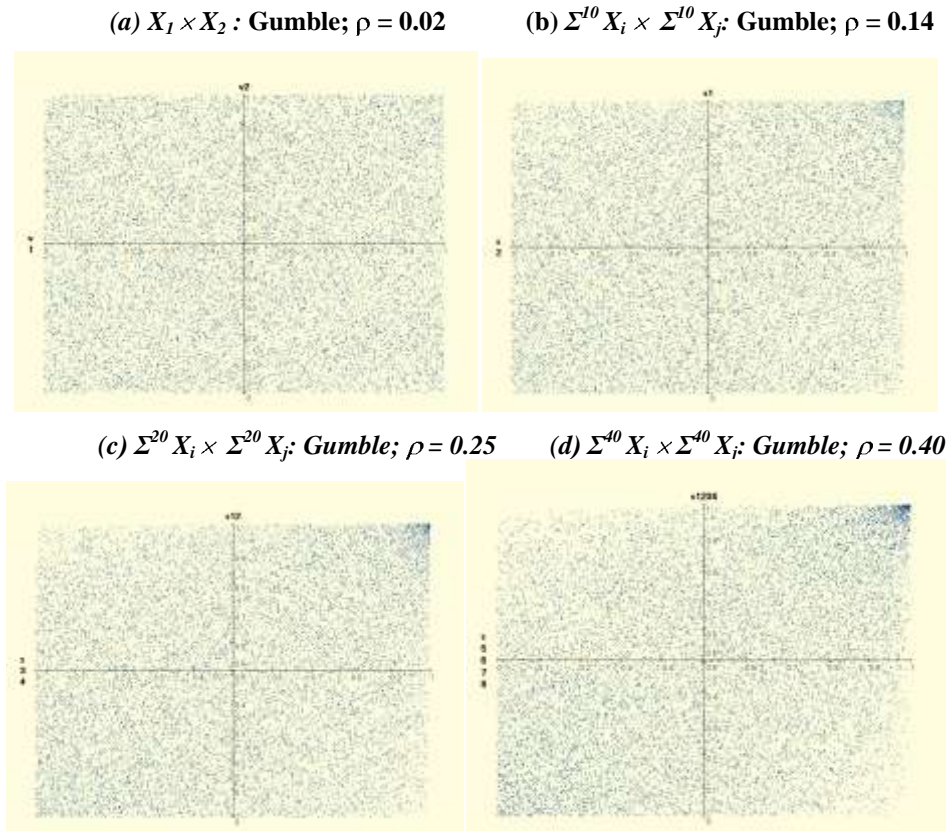
(a) $X_1 \times X_2$: Gumble; $\rho = 0.02$ *(b) $\Sigma^{10} X_i \times \Sigma^{10} X_j$: Gumble; $\rho = 0.14$*

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

Figure 2: Tail Dependence, Gumble Copula

To illustrate how the aggregate tail dependence depends on the copula, Figure 3 shows the same model as in Figure 2, but with the elliptical copula[7] in 3a and the Normal copula in 3b. In both of these, there is no discernable tail dependence.

---

[7] *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.*

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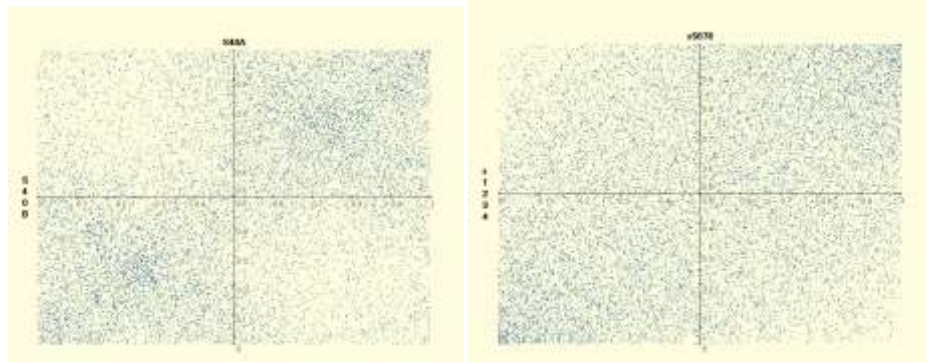

Figure 3: Elliptical and Normal Copulae; No Tail Dependence

The presence of upper tail dependence is very important for two reasons; first, it warns us that our "sorrows come in battalions," and second it tells us that in spite of ballooning correlation, the lower values may remain relatively uncorrelated. Policy implications of this for insurance are discussed in section six.

5. Tails and Dependencies: Disaster Insurance Claims

5.1 Flood Insurance Claims

Flood insurance is available in the US through the National Flood Insurance Program (NFIP), run by the Federal Emergency Management Agency (FEMA). The NFIP was enacted in 1968, partly in response to the concern by private insurers that flood peril was uninsurable.

We start by examining the claims for the all US counties over the years 1980 to 2008. There are 26,604 county-year observations in our data. To first look at whether the data is fat-tailed, we look at the mean excess plot as discussed in Section 3.1. This is shown in Figure 4, where all the data has been pooled. The upward sloping plot suggests that flood insurance claims are indeed fat-tailed.
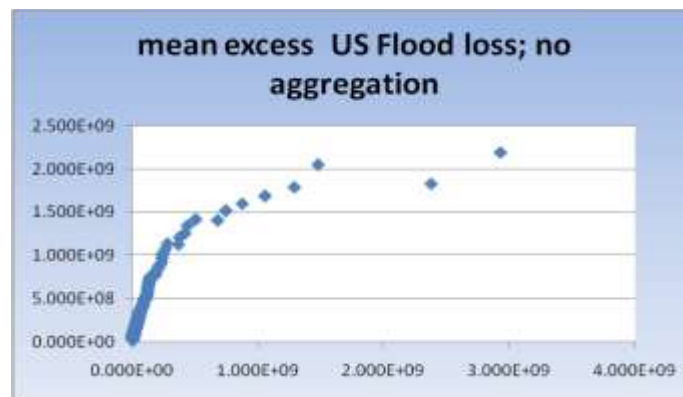

Figure 4: Mean Excess Plot for US Flood Claims

Over the 29 years we examine, exposure in most counties increased considerably. This had the effect of causing an increase in damages over time even after correcting damages for inflation. To remove the effect of growing exposure, we normalize all claims by dividing the claims per county per year by with personal income per county per year available from the Bureau of Economic Accounts (BEA).[8] When examining claims normalized in this way, the distribution is still fat-tailed as show in Figure 5 (left). The slope is not constant but is roughly 4.1 and the Pareto approximation would have a parameter of 1.2, characteristic of infinite variance. Aggregation does not thin the tail as seen in Figure 5 (right), which shows the mean excess for random aggregations of 50 counties. For the rest of our analyses, we continue to work with this data corrected for exposure growth.
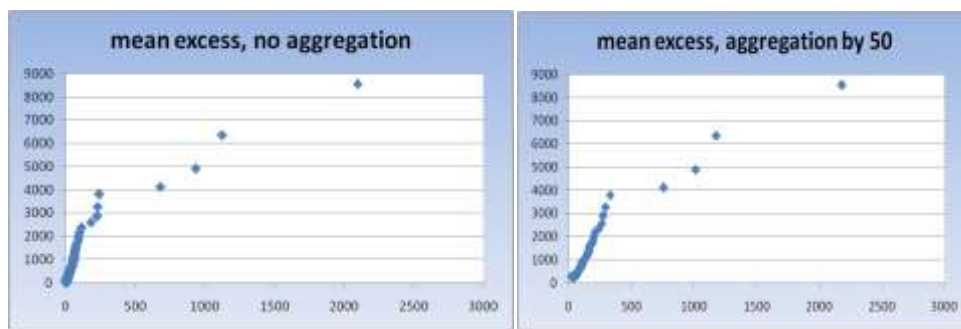

Figure 5: Mean Excess Plot for US Flood Claims with and without Aggregation

We would presume that losses in neighboring counties are correlated. When just randomly drawing any two counties for the entire country, however, we might presume they are more likely to be independent. Figure 6 shows a histogram of the correlation in flood claims in a given year between any two counties randomly drawn 500 times. As predicted, the majority of correlations are around zero, with many correlations actually negative.

---

[8] *This data goes from 1980 to 2006, 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.*
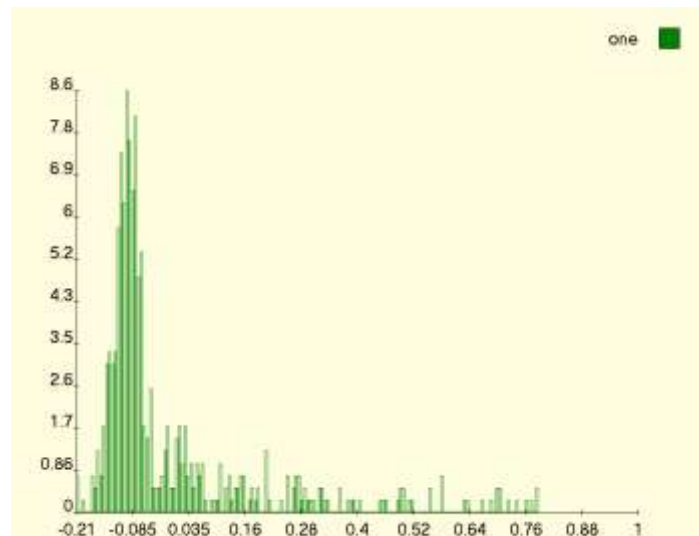
Figure 6: Correlations in Flood Claims between Randomly Selected Counties

Using a simple test based on the sampling distribution of the empirical correlation of the normal distribution with 29 observations, 454 of the 500 correlations of randomly chosen counties would be judged not significantly different from zero. The average correlation in Figure 1 is 0.041, which, by itself, would not be judged significantly different from zero. However, it is not by itself, it is the average over all 2,528 counties with flood losses in this period. That makes the number 0.041 significant. To appreciate this fact, it suffices to aggregate the data into groups and look at the correlations between the aggregations. We randomly drew aggregations of 5, 10, 20, 50, 100, 500, and 1,000 counties. Results are shown in Table 2. The correlations are monotonically increasing in aggregation size. If flood claims were independent, the averages in Table 2 should fluctuate around zero. This is clearly not the case. Figure 7 shows the histograms for 500 random observations of correlations at different levels of aggregation are plotted. By aggregations of one-thousand, the correlation is approaching 1.

Aggregation Size	**1**	**5**	**10**	**20**	**50**	**100**	**500**	**1000**
Average Correlation	0.041	0.055	0.063	0.082	0.17	0.23	0.76	0.92
Standard Deviation	0.23	0.23	0.22	0.22	0.28	0.30	0.30	0.15

Table 2: Correlations of Aggregations for US County Flood Insurance Claims
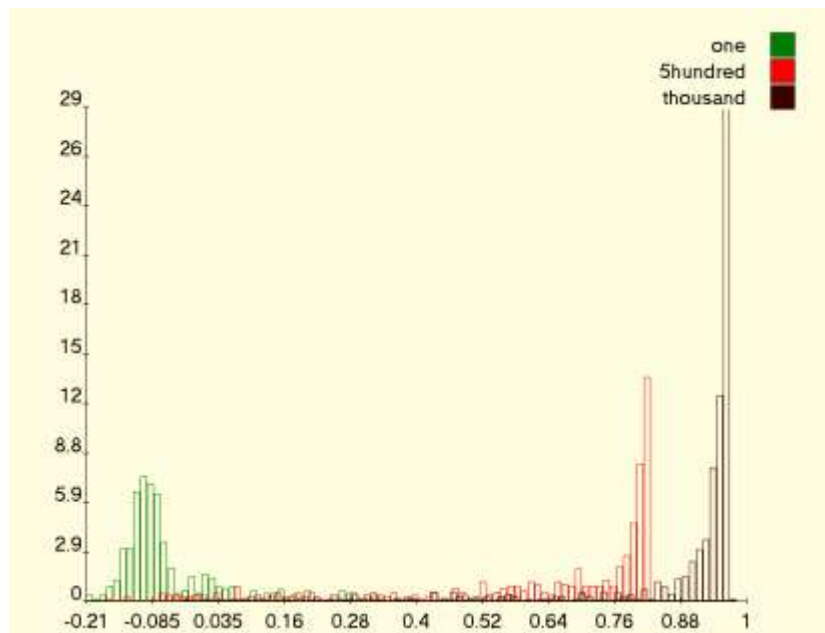
Figure 7: Correlations in Flood Claims at Different Levels of Aggregation

To study tail dependence, 29 observations are not sufficient. Therefore, we turn to the Florida data. Here we examine recorded property damages in each of the 66 counties for each month between 1977 and 2006. This gives 348 observations for each county. As seen in the national data in Table 1, the average correlation rises as counties are aggregated in Florida (shown in Table 3). The correlations are higher, since there is more correlation between the individual counties. The tail dependence also increases, as shown in Table 4 and Figure 8. Table 4 shows that the probability of one correlate exceeding its 90[th] percentile given that the other exceeds its 90[th] percentile, UTD(90), increases with aggregation. In Figure 8, the left graph shows a percentile scatter plot of monthly flood damages for two randomly chosen aggregations of 5 Florida counties (the points on the axes are months with no claims filed.) The right picture shows a similar plot, but for aggregations of 30 counties. The increase in tail dependence is striking.

Aggregation size	1	5	10	15	20	25	30
Average correlation	0.213	0.321	0.401	0.444	0.507	0.551	0.612
Standard deviation of correlations	0.277	0.300	0.286	0.260	0.240	0.216	0.184

Table 3: Florida, correlations of random aggregations, 500 iterations

Aggregation size	1	5	10	15	20	25	30
Average UTD(90)	0.183	0.430	0.498	0.546	0.588	0.626	0.654

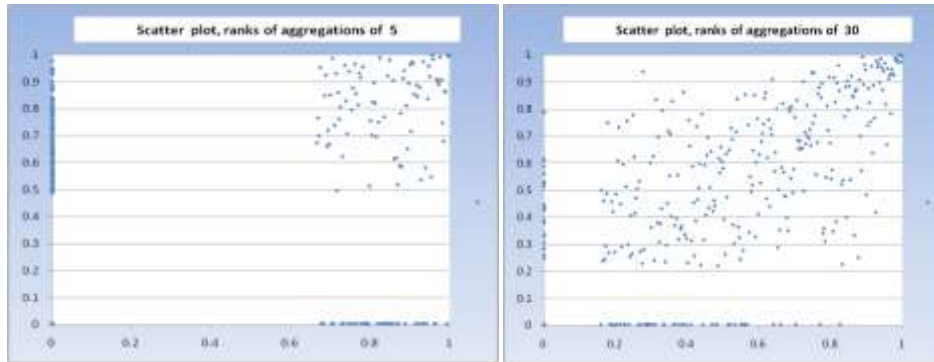**Table 4: Florida, UTD(90), 500 iterations**

Figure 8: Tail Dependence among Florida Counties, with Random Aggregations of 5 (left) and 30 (right)

5.2 National Crop Insurance Data

The US Department of Agriculture maintains data on crop insurance claim payments by county over the years 1979 – 2008. The first year had many rogue entries; useful data exists from 1980-2008 for 2,898 counties across the country. Pooling all the data and considering only those county losses above 10^6 gives 1,172 county loss events. Pooling all these events produces the mean excess plot in Figure 9, (left). The slope is 0.66 and is not constant, but the Pareto approximation would have parameter 2.51, which is fairly fat tailed. Aggregating by 10 (Figure 6, right) shows that the fatness is greatly reduced, the Pareto parameter is now 11.32, effectively wiping out the tail obesity.
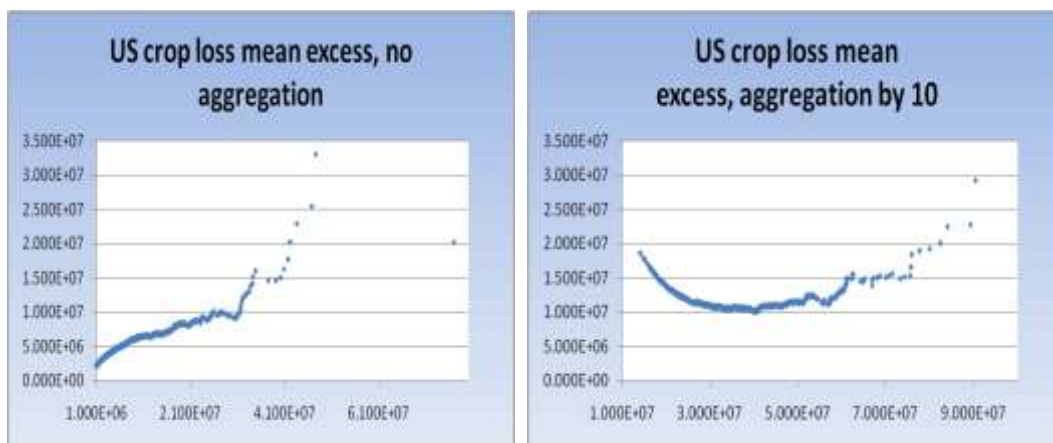

Figure 9: Mean Excess Plots of US Crop Insurance Claims, Without Aggregation (Left) and Counties Aggregated by 10 (Right)

The county-to-county correlations are noisy. Table 5 gives the distribution of 500 correlations between pairs of random aggregations of counties. Aggregation size 1 corresponds to individual county correlations. The mean correlation was 0.13. There are a fair number of negative correlations—this means that there are pairs of counties where 'one does well when the other does badly.'

500 iterations							
Aggregation Size	1	5	10	20	50	100	500
Average Correlation	0.131	0.378	0.499	0.622	0.790	0.882	0.972
Standard Deviation	0.282	0.240	0.202	0.154	0.0956	0.050	0.010

Table 5: Correlations of random aggregations of county crop losses

Using a simple test based on the sampling distribution of the empirical correlation of the normal distribution with 28 observations, 394 of the 500 correlations of randomly chosen counties would be judged not significantly different from zero. Even for the 500 correlations of random aggregations of 5 counties, 244 would be judged not significantly different from zero. At aggregation level 20, only 39 would be judged insignificant. This demonstrates the power of aggregation in discerning micro correlations. If the correlations were indeed zero, then all the average correlations in the table above would be fluctuating about zero, and not increasing monotonically. Figure 10 depicts the distributions of the 500 random aggregations, of size 1, 10 and 50.
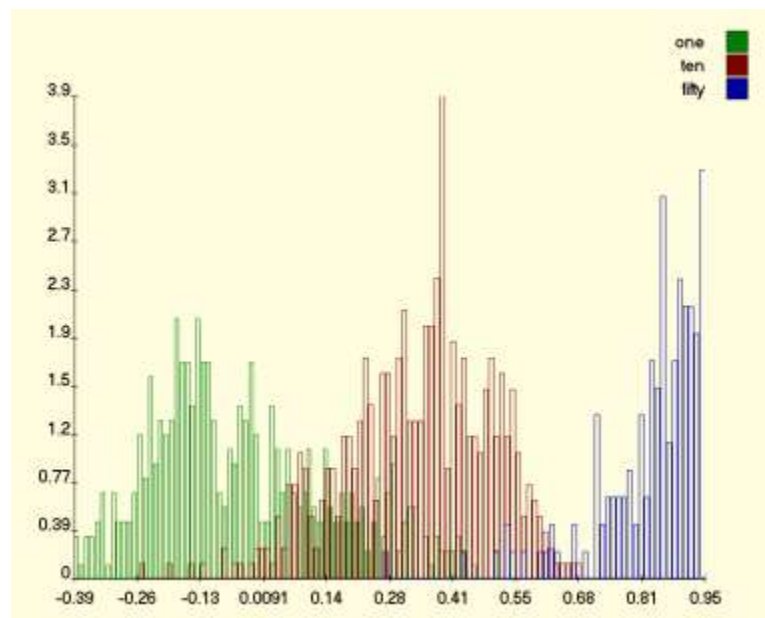

Figure 10: Correlations of Crop Losses of Random Aggregations of 1, 10, and 50 US Counties

It is interesting to examine the correlations of crop losses and flood claims. Here the normalization of exposure is important. If we do not normalize the flood claims, then a correlation of 0.47 between crop loss and flood claims is found. After normalization this becomes 0.087.

It is useful to compare these analyses with a dataset of variables that are sampled independently. Table 6 shows correlations of aggregations from 500 independent uniform variables, each sampled 100 times. The aggregations are sampled 500 times. With these numbers, the individual correlations can vary between -0.3 and 0.3, but the average over all pairs of 500 variables is very small. Even this small sample correlation will balloon when taking aggregations, but the effect is quite small; the correlation of distinct sums of 250 variables produces an average correlation of 0.096. This indicates that average correlations of aggregations is indeed a useful tool for detecting micro correlations.

aggregationn size	1	10	20	50	100	250
av corel	0.000186	0.006413	0.006611	0.018513	0.042434	0.096888
stdev corr	0.099566	0.099455	0.098267	0.102091	0.096133	0.071649

Table 6: Average Correlations Based on 100 Samples from 500 Independent Uniform Variables

6. Policy Implications

The first thing to note is that more research on these topics is clearly needed. Mathematical tools from extreme value theory, financial risk management, and dependence analysis should be developed and brought to bear on these issues. Further analysis of other large loss data sets could teach us more about micro-correlations, fat tails, as tail dependence, beyond our initial findings from the insurance data. Larger, more diverse data sets may reveal properties under random aggregation that are not visible otherwise. Hopefully, features brought to light can be aligned with more physical analyses into phase transitions, Super Cats, and the clustering of catastrophes. Based the mathematical and empirical results, more thought will need to be put into developing robust risk management strategies as the climate changes.

The policy implications from our initial work here are very preliminary and will be elaborated as our research in this area proceeds. Still, while that is ongoing, an initial discussion is warranted. We begin with four issues that stand out for the insurance of natural disaster risk:

(1) Limits to Securitization. In the presence of micro-correlations, a point is reached beyond which further securitization is unhelpful and can actually hurt us. Strategies for identifying potential micro-correlations are needed, as are approaches for identifying and implementing limits to aggregation. What is the role of private markets, what is the role of government oversight, and what are the rights of the public for disclosure? Until such questions are resolved, the results presented here suggest that caution in securitization would be prudent.

(2) Conditional Indemnity. If indeed correlation is found to concentrate at high damage levels, it may well be that conditionalizing on small to modest damage levels could define markets in which diversification is viable. That is, in theory, a cap could be indentified under which tail dependence is at a minimum and private markets could function. This cap may be akin

to the phase transition explored by RMS modelers. A possible role for government could be in covering losses above the cap. Instead of simply assuming the losses, the government could offer potentially subsidized excess-of-loss reinsurance for the higher levels of loss (Cutler and Zeckhauser 1999). This could be a multi-line contract, where losses would have to exceed a trigger in multiple insurance lines (Lescourret and Robert 2006). Such a contract could be similar to the excess-of-loss contracts proposed by Lewis and Murdoch (1996). Catastrophe bonds could also potentially be designed to cover tail dependent risks. 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). In theory, a cat bond that is triggered by losses in multiple lines of coverage or multiple locations could be designed for tail dependent losses. Research on the modeling of extreme dependence that could guide the pricing of such cat bonds or reinsurance is ongoing.

   (3) A Role for Mitigation. Damage distributions, like the insurance claims we examine here, are a function of both a natural hazard and human exposure to that hazard. Addressing exposure, through location decisions, or through mitigation policies, can reduce damages from extreme events. Adherence to strict hurricane building codes, for example, can thin the tail of the damage distribution. With losses lower, insurance supply may increase and prices may fall. Kunreuther has discussed reasons individuals may not invest in mitigation measures, such as myopia, underestimation of losses, and high up-front costs (2006). Mitigation does not just reduce private losses, but also reduces societal losses. As Kunreuther as noted, if I fail to secure my roof, for instance, it could blow off and damage my neighbors' home. But the interconnections are deeper than just spillover property damage. For instance, if someone looses their home and must leave work for an extended period of time, it hurts the local economy. If mitigation has aspects of a public good, it is likely to be underprovided by the private market. Government could ensure provision through strict building codes, tax incentives, and/or public of mitigation projects.

   (4) Affordability and Equity. It has been shown that the traditional methods of calculating risk exposure may significantly underestimate risk. As with any realization that exposure is greater than previously thought, insurance companies have three response options. They can: reduce exposure, increase the amount of capital they are holding in reserve, or purchase protection from reinsurers or the financial markets (Kunreuther 1998). Any of these options will drive up insurance prices. If prices rise higher than insureds are willing to pay, then the private (re)insurance market may break down, and a broader socio-political discussion of how to manage catastrophe risk will ensue. Florida provides one example of how a collapse in private insurance markets can be poorly handled if the trade-offs are not fully considered. After a series of hurricanes, Florida denied the rate increases private insurers deemed necessary to write policies in an area at such high risk for hurricanes. The state created a residual market mechanism to write property insurance policies to homeowners who could not find a policy in the private market. It has grown to be the largest insurer in the state. The state also has a government reinsurance fund, the Florida Hurricane Catastrophe Fund. To repay debt, it can assess private insurers to recoup costs (certain lines are exempted, such as health and workers' compensation, for example), which would be passed on to all those with an insurance policy from companies operating in Florida (for more on the situation in Florida, see: Grace and Klein 2009). A truly disastrous hit could bankrupt the state, absent a federal bailout. There is concern

that this arrangement is discouraging mitigation; encouraging development in risky areas; subsidizing not just those who need help affording insurance, but also high-income property owners; and distributing the costs of a catastrophe inequitably (Kunreuther and Michel-Kerjan 2009). When catastrophic losses are possible, private insurance may be unable to effectively operate and the question of how society should then manage this risk and how the costs of a disaster should be distributed needs to be carefully discussed.

The line of research we have initiated here also has implications for climate policy beyond insurance markets. Our work has shown that unless risks are properly accounted for, estimates of the costs of single events or of aggregate damages from climate change could significantly underestimate potential losses. Also, as shown in some of the examples, tail dependence between different loss distributions is sometimes created by human activity, suggesting it can also be broken. This was indicated by the example of installing water and gas pipes that can withstand earthquakes. Identifying areas such as this where risks can be reduced will undoubtedly be difficult and potentially incredibly location specific. Still, the work RMS and others are doing in catastrophe modeling will likely continue to shed light on areas for improved disaster response and adaptation policy. This work will also increase our understanding of how high magnitude events trigger further losses, such as the relation between quarantines and network failure, which may otherwise be neglected.

Appendix: L_p Symmetric Measures

An atomless L_p symmetric measure on $\mathbf{R}^n$ is one whose density at $(x_1,\ldots x_n)$ depends only on the L_p norm $[\Sigma|x_i|^p]^{1/p}$. Berman (1980) proved that L_p symmetric measures on $\mathbf{R}^\infty$ can be uniquely represented as condititionally independent gamma transforms. For L_1 measures, we have conditionally independent exponentials given the failure rate.

Recall the gamma integral:

(14) $\int_{x>0} x^{\upsilon-1} e^{-\alpha x}\, dx\ = \Gamma(\upsilon)/\alpha^\upsilon;\ \alpha>0,\ \upsilon > 0;\ \Gamma(\upsilon) = (\upsilon-1)!$.

and the incomplete gamma integral with integer shape υ (Abramowitz and Stegun 1965):

(15) $(1/\Gamma(\upsilon))\int_{x>r}\ x^{\upsilon-1}\alpha^\upsilon\, e^{-\alpha x}\, dx =\ e^{-\alpha r}\, (1 + \alpha r + (\alpha r)^2/2! +\ldots(\alpha r)^{\upsilon-1}/(\upsilon-1)!);\ \upsilon \in \mathbf{N}_+$

The gamma density with shape υ and scale α is $f(x) = \alpha^\upsilon x^{\upsilon-1} e^{-\alpha x} / \Gamma(\upsilon)$, with mean υ/α and variance υ/α^2. $(X_1,..)$ have a L_1 symmetric distribution with gamma prior if, for any n, the n-dimensional marginal density is given by:

(16) $p(x_1,\ldots x_n) = \int (\Pi\lambda e^{-\lambda\Sigma x i})\, \lambda^{\upsilon-1}a^\upsilon\, e^{-a\lambda}\, (1/\Gamma(\upsilon))\, d\lambda$.

Setting n = 1 and integrating over λ, one finds the univariate density and survivor functions

(17) $p(x) = \upsilon a^\upsilon / (a+x)^{\upsilon+1};\ S(x) = (a/(a+x))^\upsilon$

which are Pareto thick tailed (leptokurtic) (Takahasi 1965, Harris 1968). Conditional on λ, the variables in (16) are independent exponentials, and the sum of n such variables is a gamma variable with shape n and scale λ. The density of the sum of the n variables in (16) is therefore

(18) $p(\Sigma_{i=1..n} X_i = r) = \int (1/\Gamma(n))\lambda^n r^{n-1}e^{-\lambda r}\, \lambda^{\upsilon-1}a^\upsilon e^{-a\lambda}(1/\Gamma(\upsilon))d\lambda$

$$= r^{n-1}a^\upsilon\Gamma(\upsilon+n) / (\Gamma(n)\Gamma(\upsilon)(a+r)^{\upsilon+n}).$$

The sums are slightly less leptokurtic, and their tail behavior does not depend on n.

Multiply (15) by x and integrate first over x, then over λ; the result is the mean of X. The variance, covariance and product moment correlation may be obtained in the same way, giving:

(19) $\mu(X) = a/(\upsilon-1); \ \sigma^2(X) = a^2\upsilon \ / \ [(\upsilon-1)^2(\upsilon-2)]$

(20) $COV(X_1X_2) = a^2/ \ [(\upsilon-1)^2(\upsilon-2)]; \ \rho(X_1, X_2) = 1/\upsilon$

(21) $Var(X_1+\ldots X_n) = \sigma^2(X)(n+ n(n-1)/\upsilon)$.

Note that the mean exists only if $\upsilon>1$, and the variance, and covariance require $\upsilon > 2$.

Tail Dependence for sums of L_1 measures

Computations of tail dependence for sums of L_1 measures are tractable. For variables X,Y with the same distribution, the (upper) tail dependence of X and Y is:

(22) $TD(X,Y) = \lim_{r\to\infty} P(X>r \cap Y>r) \ / \ P(X>r)$.

$X\perp Y$ says that X and Y are independent. If $X\perp Y$ then $TD(X,Y) = 0$ but not conversely. A bivariate normal with $\rho^2 < 1$ has $TD = 0$. Tail dependence is invariant under a monotone transformation of X and Y, hence it is a property of the copula joining X and Y. In financial mathematics there is great interest in TD>0, and hence great interest in modeling dependence with copulae which show positive TD. This note shows how TD can arise from L_1 symmetric variables.

Let $\Sigma(\perp)X_i$ denote sums of independent versions of the X_i and $\Sigma(L_1) \ X_i$ denote sums of the L_1 versions of the X_i. Using (15) ad (18) we have:

(23) $P(\Sigma_{i=1\ldots N}(\perp) \ X_i > r) = e^{-\lambda r}(1 + \lambda r + \ldots(\lambda r)^{N-1} \ / \ (N-1)! \ )$

After some math,

(24) $P(\Sigma_{i=1\ldots N}(L_1) \ X_i > r) = K \times [\ \Sigma_{k=0..N-1} \ \Gamma(\upsilon+k)/(k!\Gamma(v)) \ \times (r/(a+r))^k]$

where $K = \alpha^\upsilon/ (\alpha+r)^\upsilon$. As $r \to \infty$, the bracketed term goes to:

(25) $\Sigma_{k=0..N-1} \Gamma(\upsilon+k) / [k!\Gamma(\upsilon)]$.

Note,

(26) $P(X_1 +... X_N > r \cap X_{N+1} + ...X_{2N} > r) =$

$$\int e^{-\lambda 2r} [(1 + \lambda r + (\lambda r)^2/2! +...(\lambda r)^{N-1}/(N-1)!)]^2 \ \lambda^{\upsilon-1}a^{\upsilon} e^{-a\lambda}(1/\Gamma(\upsilon)) \ d\lambda.$$

$$= (a/(a+2r))^{\upsilon} \Sigma_{k,j=0..N-1} \Gamma(\upsilon+k+j)r^{k+j} / [\Gamma(\upsilon)k!j! (a+2r)^{k+j}].$$

Letting $r \rightarrow \infty$, the sum becomes:

(27) $\Sigma_{k,j=0..N-1} (1/2)^{j+k}\Gamma(\upsilon+k+j) / [\Gamma(\upsilon)k!j!].$

The tail dependence of sums of N L_1 variables therefore is

(28) $\dfrac{\Sigma_{k,j=0...N-1}\left(\frac{1}{2}\right)^{\upsilon+j+k}\Gamma(\upsilon+k+j)/[\Gamma(\upsilon)k!j!]}{\Sigma_{k=0...N-1}\Gamma(\upsilon+k)/[k!\Gamma(\upsilon)]}$.

The following table gives some values, comparing the number n of disjunct variables summed. We see that the tail dependence grows in n, and decreases in the shape factor υ.

Upper tail dependence of disjunct sums of n: n =						
Shape	Corr(1,2)	1	3	5	10	50
1		5.00E-01	6.88E-01	7.54E-01	8.24E-01	9.20E-01
2		2.50E-01	4.53E-01	5.49E-01	6.64E-01	8.42E-01
3	3.33E-01	1.25E-01	2.89E-01	3.88E-01	5.23E-01	7.67E-01
4	2.50E-01	6.25E-02	1.80E-01	2.67E-01	4.05E-01	6.94E-01
5	2.00E-01	3.13E-02	1.09E-01	1.80E-01	3.07E-01	6.24E-01
6	1.67E-01	1.56E-02	6.54E-02	1.18E-01	2.30E-01	5.58E-01
7	1.43E-01	7.81E-03	3.86E-02	7.68E-02	1.69E-01	4.97E-01
8	1.25E-01	3.91E-03	2.25E-02	4.90E-02	1.22E-01	4.39E-01
9	1.11E-01	1.95E-03	1.29E-02	3.09E-02	8.72E-02	3.87E-01
10	1.00E-01	9.77E-04	7.39E-03	1.92E-02	6.14E-02	3.38E-01
11	9.09E-02	4.88E-04	4.18E-03	1.18E-02	4.28E-02	2.94E-01
12	8.33E-02	2.44E-04	2.35E-03	7.20E-03	2.94E-02	2.55E-01
13	7.69E-02	1.22E-04	1.31E-03	4.34E-03	2.01E-02	2.19E-01
14	7.14E-02	6.10E-05	7.29E-04	2.60E-03	1.35E-02	1.88E-01
15	6.67E-02	3.05E-05	4.02E-04	1.54E-03	9.04E-03	1.60E-01
16	6.25E-02	1.53E-05	2.21E-04	9.11E-04	5.99E-03	1.35E-01
17	5.88E-02	7.63E-06	1.21E-04	5.34E-04	3.93E-03	1.14E-01
18	5.56E-02	3.81E-06	6.60E-05	3.11E-04	2.56E-03	9.57E-02
19	5.26E-02	1.91E-06	3.59E-05	1.80E-04	1.66E-03	7.98E-02
20	5.00E-02	9.54E-07	1.94E-05	1.04E-04	1.07E-03	6.63E-02

Table 1: Upper tail dependence

The three figures following show rank scatter plots for sums of L_1 measures with shape (υ) = 3. The first shows two variables, the second two sums of 10 variables, and the third shows two sums of 50 variables.
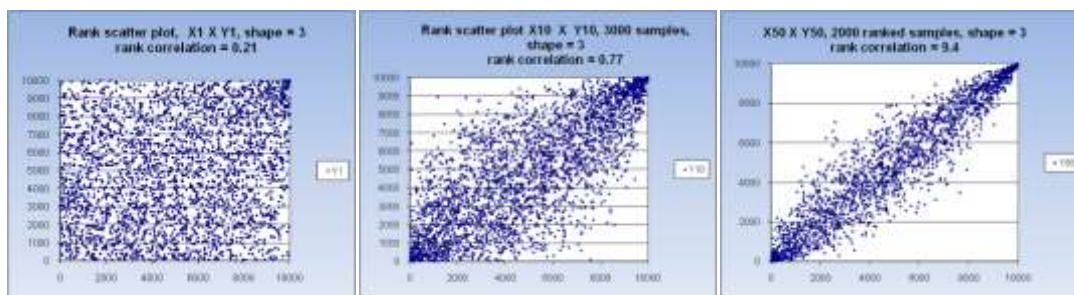

Figure 7: Scatter plots of sums with indicated correlation

Lower Tail Dependence

We may suspect that as the upper tail dependence gets stronger, the lower tail dependence gets weaker. The formulas for calculating this are complex and the limiting formula cannot be derived as easily as for upper tail dependence. Further, the calculations for small values of r with large values of N, encounter numerical problems. The trend, however, can be seen in the table below.

$PROB(Z_1,Z_2 <r) / Prob(Z_1<r)$			
Z_i = sum of n; a=3, v=3			
r= mean * fraction:			
n / **0.5**	**0.25**	**0.1**	**0.01**
1 / 0.56	0.36	0.17	0.020
4 / 0.57	0.27	0.048	2.66E-05
8 / 0.63	0.31	0.0371	3.31E-06

Table 6: Lower tail dependence

The fact that correlation is concentrated in the upper tail means that by conditioning on lower values of each variable, the dependence can be restricted. Such lower quadrant conditioning is comparable to liability indemnity of insurance policies.

References

Abramowitz, M. and I. A. Stegun, Eds. 1965. *Handbook of Mathematical Functions*. Mineola, NY, Dover Publications.

Berman, S. M. 1980. Stationarity, Isotropy and Sphericity in lp. *Probability Theory and Related Fields* 54(1): 21-23.

Charpentier, A. 2008. Isurability of Climate Risks. *The Geneva Papers* 33: 91-109.

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.

Cutler, D. M. and R. J. Zeckhauser. 1999. Reinsurance for Catastrophes and Cataclysms. *The Financing of Catastrophe Risk*. K. A. Froot. Chicago, University of Chicago Press**:** 233-269.

De Bono, A., G. Giuliani, S. Kluser and P. Peduzzi. 2004. Environmental Alert Bulletin: Impacts of Summer 2003 Heat Wave in Europe. Châtelaine, Switzerland, United Nations Environment Programme, Global Resource Information Database (GRID) - Europe

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

Embrechts, P., D. D. Lambrigger and M. Wüthrich. 2008. Multivariate Extremes and the Aggregation of Dependent Risks: Examples and Counter-Examples. Zurich, Switzerland, ETH Zurich, Department of Mathematics.

Evan, A. T., J. Dunion, J. A. Foley, A. K. Heidinger and C. S. Velden. 2006. New Evidence for a Relationship Between Atlantic Tropical Cyclone Activity and African Dust Outbreaks. *Geophysical Research Letters.* 33(L19813): doi:10.1029/2006GL026408.

Froot, K. A. 2001. The Market for Catastrophe Risk: A Clinical Examination. *Journal of Financial Economics* 60: 529-571.

Froot, K. A. 2008. The Intermediation of Financial Risks: Evolution in the Catastrophe Reinsurance Market. *Risk Management and Insurance Review* 11(2): 281-294.

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

Grace, M. F. and R. W. Klein. 2009. The Perfect Storm: Hurricanes, Insurance, and Regulation. *Risk Management and Insurance Review* 12(1): 18-124.

Gray, W., M. 1990. Strong Association Between West African Rainfall and US Landfall of Intense Hurricanes. *Science* 249: 1251-1256.

Grossi, P., H. Kunreuther and D. Windeler. 2005. An Introduction to Catastrophe Models and Insurance. *Catastrophe Modeling: A New Approach to Managing Risk*. P. Grossi and H. Kunreuther. New York, NY, Springer**:** 23-42.

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."

Harris, C. M. 1968. The Pareto Distribution as a Queue Service Discipline. *Operations Research* 16(307-316).

Ibragimov, R. and J. Walden. 2007. The Limits of Diversification when Losses May be Large. *Journal of Banking & Finance* 31: 2551-2569.

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

Kuhn, G., S. Khan, A. R. Ganguly and M. L. Branstetter. 2007. Geospatial-Temporal Dependence Among Weekly Precipitation Extremes with Applications to Observations and Climate Model Simulations in South America. *Advances in Water Resources* 30: 2401-2423.

Kunreuther, H. 1998. Insurability Conditions and the Supply of Coverage. *Paying the Price: The Status and Role of Insurance Against Natural Disasters in the United States*. H. Kunreuther and R. J. Roth. Washington, D.C., Joseph Henry Press.

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. and R. Hogarth. 1992. How Does Ambiguity Affect Insurance Decisions? *Contributions to Insurance Economics*. G. Dionne. Dordrecht, The Netherlands, Kluwer Academic Publishers.

Kunreuther, H. C. and E. Michel-Kerjan. 2007. Climate Change, Insurability of Large-Scale Risks and the Emerging Liability Challenge. *Penn Law Review* 155(6): 1795-1842.

Kurowicka, D. and R. M. Cooke. 2006. *Uncertainty Analysis with High Dimensional Dependence Modeling.* New York, New York, Wilely.

Lane, M. and O. Mahul. 2008. Catastrophe Risk Pricing: An Empirical Analysis. *Policy Research Working Paper 4765*. Washington, D.C., The World Bank, Global Capital Market Non Banking Financial Institutions Division, Finance and Private Sector Development Vice Presidence.

Lescourret, L. and C. Y. Robert. 2006. Extreme Dependence of Multivariate Catastrophic Losses. *Scandinavian Actuarial Journal* 4: 203-225.

Lewis, C. M. and K. C. Murdock. 1996. The Role of Government Contracts in Discretionary Reinsurance Markets for Natural Disasters. *The Journal of Risk and Insurance* 63(4): 567-597.

Malamud, B. D. and D. L. Turcotte. 2006. The Applicability of Power-Law Frequency Statistics to Floods. *Journal of Hydrology* 322: 168-180.

McNiel, A. J., R. Frey and P. Embrechts. 2005. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton, NJ, Princeton University Press.

Mills, E. 2005. Insurance in a Climate of Change. *Science* 309: 1040-1044.

Milly, P. C. D., J. Betancourt, M. Falkenmark, R. M. Hirsch, Z. W. Kundzewicz, D. P. Lettenmaier and R. J. Stouffer. 2008. Stationarity is Dead: Whither Water Management? *Science* 319(February 1): 573-574.

Muir-Wood, R. 2004. Top 10 Risks: Paradise Painted Black. *Risk & Insurance* April 15.

Muir-Wood, R. and P. Grossi. 2008. The Catastrophe Modeling Response to Hurricane Katrina. *Climate Extremes and Society*. H. F. Diaz and R. J. Murnane. Cambridge, UK, Cambridge University Press.

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

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

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

———. 2006. Press Release: RMS Study Estimates $50-$80 Billion of Insuerd Losses From a Repeat of the 1906 San Francisco Earthquake and Fire; Insured Losses Expected to Amplify Due to 'Super Cat' Effects. Newark, California, Risk Management Solutions.

———. 2008. Press Release: Advanced RMS Simulation Platform Allows Re/Insurers to Model "Clustering" of Weather Events. Newark, California, Risk Management Solutions.

Schär, C. and G. Jendritzky. 2004. Hot News From Summer 2003. *Nature* 432: 559-560.

Schär, C., P. L. Vidale, D. Lüthi, C. Frei, C. Häberli, M. A. Liniger and C. Appenzellar (2004). The Role of Increasasing Temperature Variability in European Summer Heatwaves. *Nature* 427: 332-336.

Sheremet, O. and A. Lucas. 2008. Global Loss Diversification in the Insurance Sector. Tinbergen Institute Discussion Paper. Amsterdam, The Netherlands, Tinbergen Institute.

Slijkerman, J. F. 2006. Insurance Sector Risk. Tinbergen Institute Discussion Paper. Amsterdam, The Netherlands, Tinbergen Institute.

Stedman, J. R. 2004. The Predicted Number of Air Pollution Related Deaths in the UK During the August 2003 Heatwave. *Atmospheric Environment* 3(8): 1087-1090

Steinberg, T. 2000. Acts *of God : the Unnatural History of Natural Disaster in America*. New York, Oxford University Press.

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

Takahasi, Koiti. 1965 A note on the multivariate Burr's distribution. *Annals of the Institute of Statistical Mathematics* (*1965*) 17:257-260,

Tol, R. S. J. 2002. Estimates of the Damage Costs of Climate Change. *Environmental Resource Economics* 21: 47-73.

U.S. Climate Change Science Program. 2008. Weather and Climate Extremes in a Changing Climate. Washington, D.C., US Climate Change Science Program.

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

van Oorschot, M. M. P., B. C. P. Kraan, R. M. M. van den Brink, P. H. M. Janssen and R. M. Cooke. 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

Weitzman, M. 2007. A Review of The Stern Review on the Economics of Climate Change. *Journal of Economic Literature* 45(3): 703-724.

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

Wharton Risk Management and Decision Processes Center. 2008. Managing Large-Scale Risks in a New Era of Catastrophes: Insuring, Mitigating and Financing Recovery from Natural Disasters in the United States. Philadelphia, PA, Wharton Risk Management and Decision Processes Center, Georgia State University, and Insurance Information Institute.

Woo, G. 1999. *The Mathematics of Natural Catastrophes*. London, Imperial College Press.

Zeckhauser, R. 1996. Insurance and Catastrophes. *The Geneva Papers on Risk and Insurance Theory* 20: 157-175.

Zimmerli, P. 2003. Natural Catastrophes and Reinsurance. Zurich, Switzerland, Swiss Re.