Large Commercial Exposures and Tail Risk:

Evidence from the Asia-Pacific P&C Insurance Market

Davide Benedetti
Imperial College Business School,
London, U.K.

Enrico Biffis
Imperial College Business School,
London, U.K.

Andreas Milidonis,
University of Cyprus, Cyprus, &
IRFRC, Nanyang Technological University, Singapore.

September 2015

Acknowledgments: We are grateful to the Insurance Risk and Finance Research Centre (IRFRC; www.irfrc.com), at Nanyang Business School, Nanyang Technological University, for generous financial support. We would like to thank Xiaoxu Liu, Bowen Yang and Wei Yu for excellent research assistance. We are grateful to Luke Armitage, John Buchanan, Rob Caton, Erik Chavez, Lauren Clarke-Wiest, Daniela Collis, Michel Dacorogna, Markus Gesmann, Mike Hood, Chris Kent, Sie Liang Lau, Bronek Masojada, James Slaughter, Paul Wee and Jerry Wu for valuable advice throughout the project. Any errors are our own responsibility.

Large Commercial Exposures and Tail Risk:

Evidence from the Asia-Pacific P&C Insurance Market

Abstract

We study losses in commercial insurance lines in the Asia-Pacific (APAC) region, by using a new dataset resulting from contributions by SCOR Services Asia-Pacific Pte Ltd and two large Lloyd's syndicates, Hiscox and Liberty. We focus on man-made risks for commercial, manufacturing, and on-shore energy exposures, and provide a comparison of tail risk profiles within the APAC region, by allowing for important rating factors such as occupancy type and location. We find evidence of significant tail risk associated with large commercial risks (LCR), in particular for manufacturing exposures and developed APAC countries. The result is confirmed by the analysis of average loss severities for benchmark portfolios of LCR exposures. Developing APAC countries are characterized by lower tail risk than developed APAC countries; this wedge is shown to have increased over time as insurance penetration in the APAC region has increased.

Keywords: Insurance, Reinsurance, Tail Regressions, Large commercial Risks.

1. Introduction

According to Swiss Re (2012), USD 600 billion of direct insurance premiums were written in commercial insurance lines worldwide in 2010. The US is by far the largest market, both in absolute dollar value terms and relative to the size of its economy, but the APAC region is also of strategic importance: Japan, for example, is the second largest market in the world, followed by China. Figures reported in Swiss Re (2012) indicate a premium volume equal to USD 237bn for the US, USD 35.4bn for Japan, and USD 30.7bn for China. In this paper, we focus on commercial property and energy on-shore risk exposures in the APAC region, and on man-made risks such as fire and explosion. We aim to shed light on the heterogeneity of tail risk profiles of different risk exposures within this area, as well as to allow end users to quantify any divergence between the loss experience in this region and that of North America and Europe.

Despite the relevance of commercial lines for the global P&C business and the corporate sector, limited public information is available on their risk characteristics, in particular about extreme loss realizations, which make up a large proportion of the total claims value. The reason is that the heterogeneity of businesses by type and size makes it difficult for individual insurers to develop reliable statistical claims information, and induces those who have larger portfolios and longer claims histories not to disclose information for competitive reasons (see Michaelides *et al.* 1997, Biffis and Chavez, 2014). There is therefore the tendency for underwriters to apply a considerable degree of judgment in pricing decisions, often giving too much weight to the value of reported claims, which may not adequately reflect the risk of the business written, and exacerbating price volatility in response to claims occurrence (e.g., Buchanan and Angelina, 2014, Biffis and Chavez, 2014).

In this paper we address the questions above by using the IRFRC LCR Dataset[1], a combination of Asia Pacific (APAC) LCR data sourced from a global reinsurer (SCOR)

[1] The dataset, referred to as **the IRFRC LCR dataset**, was manually collected by research teams at IRFRC (Nanyang Business School, NTU, Singapore) and at Imperial College Business School (Imperial College London, UK). The dataset includes the APAC component of the Imperial-IICI dataset, a data source developed by Hiscox and Liberty in collaboration with Imperial College Business School, and with the support of the Insurance Intellectual Capital Initiative (IICI) and Lloyd's; see Biffis and Chavez (2014) for an overview of the dataset and some findings on large

and from two large Lloyd's of London syndicates (Hiscox and Liberty). The anonymized database provides information on FGU losses occurred during the period 2000-2013 in the APAC region for commercial, manufacturing, energy on shore, residential and miscellaneous exposures. In addition to FGU loss information, the dataset provides information on risk exposures, including location, occupancy type, and Total Insurable Value (TIV). A more detailed description of the dataset is provided in the next section.

We develop our analysis in four steps. First, we measure tail risk by approximating the tail behaviour of the claims with a power law, and by computing the so called tail index, a measure of the speed of decay of the probability mass attached to larger and larger losses; the smaller the tail index, the larger the probability of occurrence of extreme claims. We also estimate a Generalized Pareto Distribution, which captures the asymptotic behaviour of the probability tail as we consider losses exceeding larger and larger thresholds. We measure tail risk for the entire sample, as well as for subclasses of LCR, focusing in particular on the Commercial (CO), Manufacturing (MA), and Energy On-Shore (EON) classes. These are reviewed more in detail in the next section. We also focus on Developed vs. Developing countries within the APAC region, to understand how economic development might affect the LCR loss profiles. In line with Biffis and Chavez (2014), we find that LCR losses are in general heavy tailed, with manufacturing exposures more so than commercial exposures. In addition to existing studies, we partially extend the analysis to EON exposures, to understand the difference between developed and developing countries, and to the detection of structural breaks in tail risk as insurance penetration increases.

As a second step, instead of considering subsamples of data with homogeneous exposure characteristics, we analyse our data jointly, and use tail index regression as a tool to understand the sensitivity of the tail index to different rating factors, such as occupancy, location, and TIV. We do so by relying on the methodology recently proposed by Wang and Tsai (2009). In a similar vein, following the work of Chavez-Demoulin *et al.* (2015), we extend our estimation of a Generalized Pareto Distribution (GPD) to a setting with covariates. The results provide a natural way to think about how

---

commercial risk exposures. The IRFRC LCR dataset is made freely available through www.irfrc.com.

different rating factors contribute to the overall tail risk of an insurance portfolio. This is formalized in the third step of our analysis. In particular, we construct simple benchmark portfolios of APAC exposures, where we vary the business mix by tilting the portfolio away from MA exposures and towards CO exposures, or away from developed APAC countries while overweighting developing countries. We find that portfolios with higher allocations to MA exposures and developed countries present larger average claim severity and greater tail risk.

As a fourth and final step, we test for the presence of structural breaks in the tail index estimated from our dataset. Whereas the tail index of developed APAC countries is stable over time, the one of developing countries increases. After controlling for gross premium written in the region, we find that the tail risk of developing APAC countries decreases as insurance penetration increases. We conjecture that the wedge between developed and developing countries unveiled by the analysis might be due to a structural difference in LCR exposures. As insurance take up increases, the adoption of international risk mitigation standards and the development of better underwriting strategies reduces risk, making clearer a systematic difference in tail risk, which might be ascribed to structural differences in LCR exposures found in developing and developed countries, the latter presenting typically policy schedules with more complex exposures and larger insured asset values.

In developing our analysis of APAC LCR losses, we face small sample issues originating from the difficulty in collecting good quality data.[2] Small sample size is a notorious problem in tail risk estimation. In particular, there is a delicate trade-off between bias and efficiency of tail index estimators, which has been extensively discussed in the literature since the pioneering works of Hall (1990) and Dacorogna *et al.* (1995) (see also Beirlant et al., 2006, or Embrechts et al., 1997). To provide a better understanding of our results, we use different methodologies to estimate the shape of the tail of APAC LCRs. We rely on both the Pareto and Generalized Pareto Distribution (GPD) model, with and without covariates driving the parameters of interest (see Wang and Tsai, 2009, for the Pareto model, and Chavez-Demoulin et al., 2015, for the GPD

[2] See Biffis and Chavez (2014), and Benedetti *et al.* (2015), for an overview of the data collection challenges encountered in the creation of the Imperial-IICI and IRFRC LCR datasets.

with covariates). We also implement the Weighted Hill estimator, which considers sequences of Hill estimates computed at higher and higher thresholds to provide an estimation methodology which is robust to small sample issues. The wide range of methodologies used allows us to gain a good understanding of LCR losses, despite the challenge of data paucity.

The paper is organized as follows. The next section provides details on the dataset used. Section 3 shows descriptive statistics, and offers a broad discussion of how losses compare across different rating factors. Section 4 reviews the statistical methodologies used in the second half of the paper. In section 5 we develop a systematic analysis of the tail risk profile of APAC LCR losses. Section 6 extends the analysis to portfolio tail risk, by constructing simple benchmark portfolios of LCR exposures that are used to understand the impact of business mix on the severity dimension of loss occurrences.[3] In section 7 we carry out tests for structural breaks in the value of the tail index over time. Section 8 concludes.

2. Dataset

The **IRFRC LCR dataset** is the result of the combination of the APAC subset of the Imperial-IICI dataset, and APAC data provided by SCOR Services Asia-Pacific Pte Ltd**.** The database provides information on FGU losses occurred during the period 2000-2013 in the APAC region for commercial, manufacturing, energy on shore, residential and miscellaneous exposures. In line with the information contained in the Imperial-IICI dataset, the focus is on man-made risks, such as fire and explosion, which are often regarded as un-modeled risks. Natural catastrophes, on the other hand, are excluded, as they are typically covered by catastrophe models. In addition to FGU loss information, the dataset provides information on the risk exposure, including location, occupancy type, and Total Insurable Value (TIV). For anonymization purposes, aggregation of the two data sources was been carried out by bucketing data into three time periods (2000-2003; 2004-2008; 2009-2013), and replacing original currency and country information with the categorical values “developed country” and “developing country”. To define the latter,

[3] The frequency dimension of loss occurrences is beyond the scope of this study, as the dataset used does not provide information on policies for which no claim was recorded.

the World Bank's economic development classification[4] was followed. Details on the data collection, anonymization, and aggregation procedures are discussed in Benedetti *et al.* (2015). In line with the Imperial-IICI dataset, the IRFRC LCR dataset provides aggregate information on indemnities for physical damage and business interruption, as well as claims assessment and settlement fees. Both claims and exposures are expressed in 2013 USD terms; the normalization is obtained by trending claims and exposures at an a notional rate of 2.5% per annum.

In terms of exposure information, in addition to TIV information, Biffis and Chavez (2014) classify occupancy types by developing a classification based on three levels of increasing granularity. The first one broadly classifies exposures into commercial (e.g., offices, banks, stores), manufacturing (e.g., utilities, food processors, mines), residential property (e.g., hotels, hospitals), and energy on shore (e.g., oil refinery). The second level, reported in Table 1 in Section 4, provides some more detail, allowing one to distinguish, for example, a hotel from a hospital, or metals from food producers. The third occupancy level offers a more granular view of the exposures, distinguishing for example between large vs. small hotels, heavy vs. light fabrication infrastructure, and food & drugs vs. chemicals vs. metal & minerals processing plants. Finally, occupancy information is complemented by the claim narrative, which may also provide some more information on the hazard event (e.g., burst of waterpipe, electrical failure, fire from hotel restaurant). Given sample size issues, the IRFRC LCR dataset only uses the first two occupancy classification levels of Biffis and Chavez (2014), and does not allow one to study residential exposures, as they are severely underrepresented in the dataset. Moreover, as TIV information is sometimes not available, the dataset uses the proxy variable TIV*, which relies on Total Sum Insured (TSI), a lower bound for TIV, when the latter is not available.

The anonymized losses of the IRFRC LCR dataset are grouped into different classes, depending on period of loss occurrence, economic development of country of location, and exposure occupancy type. The losses are grouped into 4 periods; period 1 if the loss occurred before or during 2003; period 2 if the loss occurred during and between

[4] See www.worldbank.org.

2004 and 2008; period 3 if the loss occurred during the years 2009 or 2010; finally period 4 if the losses occurred during 2011, 2012 or 2013. The economic development class can take the categorical value Developed (D), if the loss happened in a developed country, or Developing (represented by E, which stands for “emerging”), if the loss occurred in a developing country. The occupancy type of the exposure is classified according to the classification used in Biffis and Chavez (2014): we use categories Commercial (CO), Manufacturing (MA), Residential (RE), Miscellaneous (MI), and Energy on Shore (EON). Moreover, we also use more granular classification based on Table 1 reported below. As the RE class is underrepresented in the APAC dataset, in the following we limit our analysis to the classes MA, CO, and EON. As it turns out, sample size considerations make our analysis more compelling for MA and CO exposures. We still report some results for the EON class when statistical significance is reasonable.

3. Descriptive Statistics

This section presents descriptive statistics and univariate analysis of our dataset in Tables 2 through 4. The focus is FGU loss, our main variable of interest. In the presentation of the results we first show the entire dataset (all collected losses) over the period 2000-2013, and then present results based on three occupancy types: CO, MA, and EON. We then also split the data by economic development.

Panel A of Table 2 shows summary statistics for the whole sample. We can see that the dataset is composed of 526 losses, of which two thirds come from developing economies, after removing losses of classes RE and MI. The average loss for developed countries is almost 4 times larger than the average loss for developing countries, and almost twice as large as the average loss for the total group. The first, second and third quartiles for developed countries are also higher. Volatility, expressed in terms of standard deviation, is also quite large. However, in terms of coefficient of variation (CV, measured as standard deviation divided by the average) this group is less volatile, with a CV of 1.77. The group with the highest variability is the one of developing countries, with a CV of 2.58. Each of the three groups of panel A shows positive skewness and high values of kurtosis, to indicate a high likelihood of positive extremes. The group of developed countries is the one with lower kurtosis. This tell us that although exposures in developed economies tend to have on average higher losses (which might reflect the fact that insurance policy schedules contain high-value insurable assets, such as buildings, infrastructure, and machinery), in terms of kurtosis and likelihood of extremes developing countries have a big cluster of losses in the higher quantiles.

Panels B to D of Table 2 offer a comparison across exposure occupancy types. Panel B considers CO exposures. The considerations are similar to those of panel A, as the developed countries group is the one with higher mean and quartiles, but lower CV (hence variability), lower skewness and lower kurtosis. Panel C focuses on EON exposures. For this type of occupancy not only the average, but also CV, skewness and kurtosis of the developed group are higher. Hence, EON exposures in developed countries not only have higher losses on average, but also a bigger cluster of losses in the high quantiles. Panel D focuses on the MA class, which accounts alone for 75% of the dataset. Results are in line with those obtained for panels A and B. From table 2 it is apparent that the average loss does not change significantly across different occupancy types. What changes is the kurtosis, which tends to be lower for the CO group compared to the others. For developed countries, CO exposures have lower average losses, while EON and MA exposures have average losses slightly above the total average of developed countries in panel A (whole sample). The CO class is also the one with lower kurtosis in general. For the developing group instead, EON exposures are those with higher average losses, while MA exposures have lowest average.

Table 3 shows summary statistics for the total insurable value (TIV), corresponding to the total asset values that could be insured. As this variable was not always available from the data providers, when it is missing the IRFRC LCR dataset uses the total sum insured (TSI) as a proxy for it, and leaves the TIV value blank when no information is available. The resulting variable is indicated as TIV*, to clarify the fact that it differs from the actual TIV. We should also emphasize that even when the TIV is available, it is not consistently defined across the dataset, as it may be derived, for example, from a top location profile or a policy profile, thus overestimating the actual TIV (see Riegel, 2010, for a discussion of different types of exposure information). When the TSI is used instead of the TIV, then the dataset is clearly underestimating the exposure. Note that descriptive statistics are computed based only on the non-zero values of TIV*.

As for table 2, panel A of table 3 shows summary statistics for the whole sample, while panels B to D focus on the classes CO, EON and MA, respectively. Average TIV* of developed countries is always higher than that of developing ones, across all panels of table 3. However, in terms of quartiles the relation is not as clear as it was for FGU losses. In terms of CV, the developing group is the one with highest variability only in the whole sample and the MA subsample (panels A and D). The distribution of this variable across different subsamples is positively skewed with high kurtosis, similar to that of FGU losses. Comparing different occupancy types, we see that MA has higher average TIV* across country types, followed by EON, and lastly CO.

4. Methodology

In this section we review the statistical tools we use to gauge the tail risk profile of APAC LCR losses. The Pareto model and the Generalized Pareto distribution are well known to provide parsimonious proxies for the shape of the tail of the distribution of large losses, or characterize the asymptotic behavior of the losses as exceedances over larger and larger thresholds are considered (see Beirlant *et al.*, 2006, or Embrechts *et al.*, 1997, for example). For the Pareto model we use the standard Hill estimator, while we use the peak-over-threshold (POT) method for the estimation of the Generalized Pareto distribution.

The Hill estimator is a maximum likelihood type of estimator for the single parameter of the Pareto model. However, it strongly depends on the chosen threshold. If the threshold is not high enough, the estimator might be biased if the underlying distribution is a more general Power Law. At the same time, a higher threshold reduces the sample size, leading to a loss in efficiency. This trade-off is well understood since the pioneering works of Hall (1990) and Dacorogna *et al.* (1995). Given the small sample size issues we face, we complement the Hill estimates with those provided by the estimator developed by Huisman *et al.* (2001). As it relies on a weighted average of different Hill estimates computed for increasingly high thresholds, we refer to it as Weighted Hill estimator.

To gain a better understanding of how different exposure characteristics shape the tail risk profile of APAC LCR losses, we then implement tail index regression and Generalized Pareto distributions with parameters driven by covariates, which reflect occupancy and location characteristics of the exposures. We adopt the tail index regression methodology of Wang and Tsai (2009) for the Pareto model, and of Chavez-Demoulin et al. (2015) for the Generalized Pareto distribution.

a. Methodology overview

Our first approach to tail risk quantification relies on the idea of approximating the tail distribution with a Power law, which takes the following form

$$P(X > x) \propto x^{-\alpha} L(x), for\ x > c \geq 0,$$

where $L(x)$ is a slowly varying function satisfying the property that $\lim_{x\to\infty} \frac{L(tx)}{L(x)} = 1$. Both Pareto and Generalized Pareto are special cases of this more general class of distributions, obtained for specific forms of the function $L(x)$.

Given a sample of n positive ordered losses, $x(n) > x(n-1) > \cdots x(i) > \cdots x(1)$, independently generated by an unknown fat-tailed distribution, Hill (1975) assumed a standard Pareto distribution for the losses over a threshold $x(n-k) = c$, i.e.,,

$$P(X > x | X > c) = \left(\frac{x}{c}\right)^{-\alpha}.$$

The estimator proposed by Hill (1975) for the tail index α takes the following form:

$$\hat{\alpha} = \left[\frac{1}{k}\sum_{i=0}^{k-1} \ln \frac{x(n-i)}{x(n-k)} - \ln x(n-k)\right]^{-1},$$

which is nothing else than the maximum likelihood estimator of α. Hence, the Pareto model sets $L(x) = c^{\alpha}$, which is clearly slowly varying as it does not depend on x. Under the above assumptions we have that $\sqrt{k}(\hat{\alpha} - \alpha) \xrightarrow{d} N(0, \alpha^2)$, providing consistency and asymptotic normality of the estimator. The parameter α measures the thickness of the tail of the distribution. A lower α indicates a slower decay of the tail towards zero, and hence a higher probability of extreme events. Moreover, the tail index is in one to one correspondence with the maximal order of finite (centered) moments of the risk considered. For example, for $\alpha < 1$ the first moment (hence the mean) of the distribution does not exist, for $\alpha < 2$ the second moment (hence the variance) is infinite, and so on.

A clear drawback of the above estimator is the choice of k. If the data generating process above the threshold is a more general Power law distribution the estimator will be biased. Dacorogna et al. (1995) show that the bias will generally be increasing in k, meaning that, as we lower the threshold, the Hill estimator $\hat{\alpha}$ will on average be more and more different from the true α. On the other hand, the variance of $\hat{\alpha}$ is decreasing with k, meaning that increasing the threshold will reduce the bias at the expense of efficiency.

Huisman et al. (2001) state that, for a given class of parameterizations of $L(x)$, the bias of the estimator $\hat{\gamma} = \frac{1}{\hat{\alpha}}$ is linear in k, i.e., $E(\hat{\gamma}(k)) = \beta_0 + \beta_1 k$. They therefore propose the following estimator for $\frac{1}{\alpha}$

$$\frac{1}{\hat{\alpha}} = \hat{\gamma} = \sum_{k=m}^{n-1} w(k)\hat{\gamma}(k),$$

where $\hat{\gamma}(k) = \frac{1}{k}\sum_{i=0}^{k-1} \ln \frac{x(n-i)}{x(n-k)} - \ln x(n-k)$ is the estimator of $\frac{1}{\alpha}$ above the threshold $x(n-k)$, and $w(k)$ is a weighting factor. The simple idea behind the Weighted Hill estimator is to pick a threshold $x(m)$ and construct all the possible estimators $\hat{\gamma}(k)$ for $x(k)$ greater than $x(m)$, and then to regress $\hat{\gamma}(k)$ over k and a constant, thus estimating the following linear model $\hat{\gamma}(k) = \beta_0 + \beta_1 k + \varepsilon(k)$.

The coefficient $\hat{\beta}_0$ estimated with this regression is the Weighted Hill estimator for $\frac{1}{\alpha}$. Since the $\hat{\gamma}(k)$ estimators are clearly correlated among themselves, OLS (ordinary least squares) estimation of β_0 is likely to be biased. To deal with this problem Huisman *et al.* (2001) suggest to use WLS (weighted least squares) regression with weighting matrix $W = diag(1, \sqrt{2}, \dots, \sqrt{m})$. So, by defining $\hat{\gamma}$ as the $m \times 1$ ordered vector of $\hat{\gamma}(k)$, and Z as the $m \times 2$ matrix containing a vector of ones in the first column, and the vector $(1,2,\dots,m)'$ in the second one, the Weighted Hill estimator is the coefficient $\hat{\beta}_0$ from $\hat{\beta}_{wls} = (Z'W'WZ)^{-1}(Z'W'W\hat{\gamma})$ – analytical derivation of the standard errors for $\hat{\beta}_0$ can be found in the appendix Huisman *et al.* (2001).

The general theory behind the POT approach is based on a result due to Pickands (1975), which applies to a series of i.i.d. losses, and states that that for a large class of distributions the conditional distribution of the excess losses $Y = X - c$ above the given threshold c, will follow a Generalized Pareto distribution as the threshold grows larger:

$$\frac{P(X-c>y)}{P(X>c)} = P(X-c>y|X>c) = P_c(Y>y) = \begin{cases} \left(1+\frac{\xi}{\beta}y\right)^{-1/\xi}, & \xi \neq 0 \\ e^{-y/\beta}, & \xi = 0 \end{cases}.$$

This distribution reduces to the Exponential distribution as ξ goes to zero, is said to be heavy tailed for $\xi > 0$, and not to have tail for $\xi < 0$. The coefficient ξ can be seen as the inverse of the parameter α of the standard Pareto, which means that for $\xi > 1$ the first moment (hence the mean) of the distribution does not exist, for $\xi > 1/2$ the second moment (hence the variance) is infinite, and so on. The estimation of ξ and β usually relies on maximum likelihood estimation. If we take again the series of ordered realizations $x(n) > \cdots x(i) > \cdots x(1)$, and set $x(n-k) = c$, we can define the excess ordered losses over c as $y(k) = x(n) - c > \cdots x(i) = x(n-i) - c > \cdots y(0) = x(n-k) = 0$, and write the log-likelihood as follows:

$$\boldsymbol{l}(\xi,\beta|y_i) = \begin{cases} \sum_{i=1}^{k}\left\{-\log(\beta) - \left(1+\frac{1}{\xi}\right)\left(1+\frac{\xi}{\beta}y(i)\right)\right\}, & \xi \neq 0 \\ \sum_{i=1}^{k}\left\{-\log(\beta) - {y(i)}/{\beta}\right\}, & \xi = 0. \end{cases}$$

To gain a more structural understanding of tail risk, we implement Pareto and Generalized Pareto models with parameters driven by covariates. For the Paeto model, we follow the tail index regression methodology of Wang and Tsai (2009), who assume that the conditional probability of losses over a threshold, given a vector of covariates Z, can be approximated by

$$P(X > x|Z = z, X > c) = x^{-\alpha(z)}L(x),$$

where $\alpha(z) = \exp(z'\beta)$. In their work, they approximate the above probability as $P(X > x|Z = z, X > c) = \left(\frac{x}{c}\right)^{-\alpha(z)}$, and apply the following approximate negative log-likelihood estimator to derive the coefficients β

$$\hat{\beta} \coloneqq \arg\min_{\beta} \sum_{i=1}^{n}\left\{\exp(z_i'\beta)\ln\frac{x_i}{c} - z_i'\beta\right\} I(\,x_i > c).$$

For the Generalized Pareto distribution, we follow the parameterization of Chavez-Demoulin *et al.* (2015),who assume

$$P(X-c>x|Z=z,X>c)=\begin{cases}\left(1+\frac{\xi(z)(1+\xi(z))}{e^{\nu(z)}}x\right)^{-1/\xi(z)}, & \xi(z)\neq 0\\ e^{-xe^{-\nu(z)}}, & \xi(z)=0\end{cases}$$

where $\nu=\ln((1+\xi)\beta)$ is orthogonal to ξ. They assume a non-parametric relation in Z, and a spline function in time for both ξ and ν. They then apply maximum likelihood estimation to obtain the estimates of the regression coefficients. In our implementation, we assume the linear relations $\xi(z)=z'\theta_\xi$ and $\nu(z)=z'\theta_\nu$, to make it easier to compare the results with those obtained with the approach of Wang and Tsai (2009). The log-likelihood takes therefore the following form:

$$\boldsymbol{l}(\theta_\xi,\theta_\nu|y_i)=$$

$$=\begin{cases}\sum_{i=1}^{k}\left\{\log(1+\xi(z_i))-\nu(z_i)-\left(1+\frac{1}{\xi(z_i)}\right)\left(1+\frac{\xi(z_i)(1+\xi(z_i))}{e^{\nu(z_i)}}y_i\right)\right\}, & \xi(z_i)\neq 0\\ \sum_{i=1}^{k}\left\{-\nu(z_i)-{}^{y_i}/_{\beta}\right\}, & \xi(z_i)=0\end{cases}$$

b. Implementation

We begin our analysis by ordering our dataset, and estimating the model on different subsamples obtained by working on the 5th, 10th,…, 90th and 95th percentiles. We also carry out univariate analyses on different exposure subsamples, by focusing on CO or MA occupancy types (the sample size of EON losses being too small), and on developed vs. developing countries. We then consider different combinations of occupancy types and time periods. We also apply the methodologies to slightly more granular occupancy type classifications (like general industry, processing and commercial exposures vs. metals, mines, chemicals and energy), obtaining similar results.

The analysis with covariates uses the entire sample up to the 85% of the data, and relies on the assumption that losses are conditionally i.i.d., given the regressors. As independent variables we use dummy variables for occupancy type, economic development, time period, and interaction between period and economic development. As a robustness check, we also run the regression including the logarithm of TIV*, in order to control for size effects. This is in line with common pricing approaches relying on rating factors dependent on TIV bands (e.g, Michaelides et al. 1997; Buchanan and Angelina, 2014).

5. Analysis

In part (a) we report results of the curve fitting exercise on the entire sample and over different sub-samples. In part (b) we show how tail parameter estimates vary depending on the different rating factors (economic development, pre- vs. post-crisis, occupancy type). In part (c) we plot the most robust CDFs from the models run in part (b). Part (d) provides a discussion of tail regression results, as well as some robustness checks.

a. Curve Fitting Exercise

Figures 1 and 2 show parameter estimates obtained with Hill, Weighted Hill and POT method (GPD), compared across different subsamples and different threshold levels. The *x*-axis represent different threshold levels measured as percentage of observation above the threshold; for example a value of 20 indicates that the threshold has been set at the 80th percentile. On the *y*-axis we report the estimates of the parameters of interest. To ensure that results are comparable, the parameters reported are $\hat{\alpha}$ for the Hill estimator, $1/\hat{\gamma}$ for the Weighted Hill estimator, and $1/\hat{\xi}$ for the POT method.

Tables 4 and 5 show t-statistics of the estimated coefficients and KS-test for goodness-of-fit estimates obtained with Hill, Weighted Hill and POT method (GPD) compared across different subsamples and different threshold levels. The parameters tested are α for the Hill estimator, γ for the Weighted Hill estimator, and ξ for the Generalized Pareto. The null hypothesis of the KS test is that the distribution being estimated is the correct one.

Figure 1 shows results for the whole period. Panel A reports estimates for the whole sample, while panels B and C focus on developed and developing countries respectively. The GPD parameter tends to be higher than those of the other methods for any given threshold level. Hence GPD tends to assume a lighter tail. In panel A the GPD parameter is above two for higher thresholds, and decreases as we lower the threshold, reaching a value of around two after the 65th percentile (more than 35% of observation). It is worth to remember that an estimate $1/\hat{\xi}$ below two implies infinite variance, and infinite mean when it is below one. The Hill estimator is close to one already at the 80th percentile and falls below one right after, implying extreme heavy tails. Both GPD and Hill estimator are greatly influenced by the choice of the threshold, and both decrease as we lower the threshold. Instead the Weighted Hill estimator proves to be more stable, and does not change significantly in panel A as we consider different thresholds. The Weighted Hill estimator is also the estimator which gives the lowest value of the parameter, and which is around 0.3. Looking at panel A of table 4, we can notice that after the 80th percentile all the estimated parameters are statistically significant. In terms of goodness-of-fit the GPD is the best performing. The Weighted Hill estimator does not seem to fit well the data for the thresholds provided, while the Hill estimator works best for higher percentiles. In Panel B (developed group) the GPD parameter is still the highest; it starts at four and decreases to slightly below three when we use 40% of sample to perform the estimation. The Hill parameter starts around 1.5 for the developed group, and then reaches one at the 60th percentile. The Weighted Hill estimator is still the one which generates the lowest parameter values. Looking at panel B of table 4, we can see that the GPD parameter is not statistically significant. It is important to note that tests for GPD parameters are based on $\hat{\xi}$; when this coefficient is not statistically different from zero it means that $\hat{\alpha}$ (its inverse) approaches infinity and hence the Generalized Pareto tends to the Exponential distribution, which is a light tailed probability model. The KS-test for the Weighted Hill is always close to zero, indicating that this model fits poorly the data at these thresholds. Both Hill estimator and GPD seem to represent the data. For the developing group (panel C) the GPD starts below two this time, and decreases faster than in the other groups as we increase the threshold, and becomes one at the 60th percentile. Both Hill and Weighted Hill estimators are below one. The Hill estimator decreases slightly, while the Weighted Hill is pretty stable and below the Hill one. In panel C of table 4 we can see that all parameters are statistically significant below the 80th percentile. Moreover Hill and GPD fit well the data at any given thresholds, while Weighted Hill works well below the 70th percentile.

Figure 2 shows results for the manufacturing subsample during the whole period, and table 5 relative t-statistics and KS-test. Panel A reports estimates for the whole sample, while panels B and C focus on developed and developing countries, respectively. As before, the GPD parameter is the highest in all panels, and tends to be higher than one or two (meaning finite mean and variance, respectively) for higher thresholds, whereas it decreases as we lower the threshold. GPD seems to fit well the data in terms of KS, and the $\hat{\xi}$ parameter is not statistically different from zero for high thresholds. Hill estimates seem stable as the Weighted Hill estimates in the previous figures. They are lower than the GPD's and higher than the Weighted Hill's. They are generally below one (no mean), and never above two. They also do very well in terms of goodness-of fit as measured with the KS test. Weighted Hill estimates are generally the lowest, are again the most stable and statistically significant, and always lie below 0.5. In these subsamples the estimator seems to perform well in terms of KS-test only in the developing group.

b. Curve Fitting with covariates

Figures 3 to 5 show parameter estimates for the whole period by occupancy type, economic development and time period, obtained with Hill, Weighted Hill and POT estimators. As before, to ensure comparability we focus on estimates $\hat{\alpha}$ for the Hill estimator, $1/_{\hat{\gamma}}$ for the Weighted Hill estimator and $1/_{\hat{\xi}}$ for the POT method (GPD).

Figure 3 refers to the Hill estimator only. Panel A shows results for the entire dataset and different occupancy types. Panel B focuses on the entire sample and economic development class. Looking at occupancy type, we see that MA and total are close to each other; this is due to the fact that our sample contains more MA data than CO. The parameter estimates start slightly above one, and go fast below this value. It is not clear whether CO or MA has heaviest tail, given that for some thresholds the CO parameter is highest and vice versa for other thresholds. Looking at results by economic development, we see that estimates for the developed class appear to be greater than for developing, meaning that it is lighter tailed.

Figure 4 refers to the Weighted Hill estimator only. Panel A shows results for total (all data) and by occupancy type. Panel B focuses on total and economic development. Looking at results by occupancy type, manufacturing and total are still close to each other, the parameters for both being below one (no mean), and oscillating between 0.3 and 0.36. The CO class is slightly below the MA one; hence the CO class is heavier tailed for these thresholds. However the differences are not great. Looking at results by economic development, we see that the developed group parameter increases slightly as we lower the threshold, going from 0.27 to 0.31, so that the mean is never defined. The developing country parameter is more stable, and around 0.5. Hence, developing countries are less heavy tailed than developed countries according to the Weighted Hill estimator (we had the opposite for Hill).

Figure 5 refers to the GPD estimator only. Panel A shows results for total and occupancy type. Panel B focuses on total and economic development. Looking at the results by occupancy type, we see that MA and total are close to each other; parameters are high for high thresholds, and decrease as we lower the threshold. Parameter for total ranges between 3 and 2; it is between 2.5 and 1.5 for the MA class. The parameter for the CO class is negative for the highest quartile, but then goes above the one obtained for the MA class (meaning thinner tail), and tends to stay above it, meaning that CO is less heavy tailed. So results by occupancy type are neater with GPD than for the other two methods, especially for high thresholds. Looking at results by economic development, we see that estimates for developed countries start at around 4 to 2.8, meaning that we always have finite variance. Developing countries estimate goes from 1.8 to one (variance is never defined), and it's always below the developed one; hence the developing class is more heavy tailed. So GPD is more in line with Hill than with Weighted Hill by economic development.

c. Plots of Cumulative Density Functions (CDFs)

Figure 6 shows empirical and estimated CDFs obtained with the three different methodologies in the whole sample. The *y*-axis reports the forecasted quantiles, while the *x*-axis maps different loss values scaled by the threshold. The threshold chosen is the median. The CDF obtained with GPD is the one with thinnest tail, and is the closest to the empirical CDF. The Weighted Hill CDF is the heaviest tailed, whereas the Hill CDF is in the middle.

d. Regression results

Tables 6 and 7 show results for the Hill and GPD regressions with covariates, respectively. The parameters are estimated for different thresholds ranging from 15% of the observations (85th percentile) to 100% (full sample), and for different set of covariates. All the dummy variables implemented are relative to the constant. The models are estimated on the whole dataset, but with EON and RE excluded. The dummy for MA is relative to CO, the dummy for developing (Emerg.) is relative to developed. The only continuous variable we use is the logarithm of total insured values (logTIV*), to control for size effects. Panel A represents the baseline model, it does not use LogTIV*. Panel B adds LogTIV* to the initial regression. It is worth pointing out that the coefficients estimated for the GPD regression are relative to the ξ parameter here, hence the average marginal effect on $\alpha=1/\xi$ is opposite to their sign.

Table 6 refers to the Hill regression with covariates. For the baseline model (panel A) the dummy for the MA class is generally negative but not statistically significant. The dummy for developing is always positive and significant for thresholds above median, meaning that developing countries tend to have higher α (thinner tail) than developed ones.

When controlling for size (adding log of TIV*, see panel B) the constant loses statistical power. The dummy for manufacturing tends to be more positive, but it is still not statistically significant. The dummy for developing countries is still always positive and statistically significant above 50% of the sample, as in panels A. LogTIV* is generally negative (hence tail becomes heavier as we increase TIV*) but not statistically significant, except for some thresholds.

Table 7 refers to the GPD regression with covariates. For the baseline model (panel A) the dummy for manufacturing is almost always negative, but not statistically significant, except for a high threshold for which it is positive, meaning that MA is slightly heavier tailed. The dummy for developing countries is generally positive and statistically significant for thresholds below the median, meaning that developing countries tend to have lower α (heavier tail) than developed ones, the opposite of the Hill regression. However, for the three highest thresholds the dummy for developing countries is negative, although not statistically significant (except for the highest thresholds), which implies higher α (lighter tail) as in the Hill regression.

When controlling for size (panel B) the constant loses statistical power (as in the Hill case). Same as before, the dummy for manufacturing is negative, but not statistically significant, except for one high threshold for which it is positive. The dummy for developing countries is still generally positive and significant when more than 55% of the sample is used. LogTIV* is negative for higher and lower thresholds (hence the tail becomes thinner as we increase TIV*, the opposite of what we obtained with the Hill regression), positive otherwise (in line with Hill), and it is rarely statistically significant.

6. Price sensitivity analysis

In this section we adopt the perspective of an insurer operating in the APAC region, and explore the impact of business mix on tail risk. In particular, we look at how the risk profile of a portfolio of policies covering risks similar to the ones of our dataset would change as we tilt the underwriting strategy towards specific exposures and away from others. In order to do this, we fix a baseline threshold of USD 1 million and estimate the parameter α of the Pareto model with the Hill estimator. We then perform two types of analyses: one by occupancy type, where we change the percentage of CO and MA risks in the portfolio; the second one by economic development, where we change the percentage of developed and developing countries in the portfolio.

We report the Hill estimates for the Pareto, Pareto with covariates, Generalized Pareto, and Generalized Pareto with covariates models in panel A of tables 8 and 9. The coefficients reported are the estimated α for Pareto and Pareto regression, and the estimated $1/\xi$ for Generalized Pareto, and Generalized Pareto with covariates. Panel A of table 8 reports α parameters for CO and MA, whereas panel A of table 9 refers to developed and developing countries. While the univariate models have been estimated on the subsamples, the models with covariates have been estimated on the whole sample (excluding EON, RE and MI exposures), using both dummies for occupancy and geographic region together. The coefficients reported for the regression models are the average of the two partial effects. E.g., the coefficient relative to the CO class for the Pareto regression (and GPD regression) is the average of the coefficients for the CO class in Developed countries and CO class in Developing countries.

Class CO has a slightly higher α than MA, making it less heavy tailed, except for the Generalized Pareto regression model. Developed countries have an α that is way lower than developing countries for the two Pareto models, meaning that they are heavier tailed. However, developed α is much higher for the two Generalized Pareto models, meaning that they are less heavy tailed for them.

The analysis of portfolio risk is performed via Monte Carlo simulation (MCS). We run 1 million Monte Carlo simulations from a mixture of two parametric models with different weights, representing the exposure of the portfolio to a given type of exposures. We also use Bootstrapping, by simulating 1 million samples of size 50 from the different subsamples of the empirical data. We report the average severity of the losses, where the latter are truncated at USD 10 million, as well as the 80^{th} percentile.

Panel B of table 8 shows results for price sensitivity analysis for different types of portfolios, where we vary the occupancy type. Row 1 simulates a portfolio with only MA risks. Row 2 simulates a portfolio with 25% share in CO and 75% in MA. Row 3 is an equally weighted portfolio between the CO and MA class. Row 4 simulates a portfolio with 75% business written in the CO class, and 25% in the MA class. Finally, row 5 presents a portfolio with only CO risks. Both in the case of MCS and Bootstrap, severity decreases slightly as we rebalance portfolio weights from MA risks to CO risks. Panel C of Table 8 reports the estimated 80%-quantile. The latter is increasing with the fraction of MA exposures in the portfolio.

Average severity for MCS is lower than for Bootstrap, except for the GPD case, for which we find it is highest. However, the MCS quantiles for the two Pareto models are generally higher than for the Bootstrap, and higher than the two Generalized Pareto models, meaning that the Pareto models put more weight on the tail than the empirical and Generalized Pareto distributions. Average severity for MCS is around USD 3.2 million for the Pareto models, and USD 4.2 million for the GPD cases, while it is around USD 4.1 million for the Bootstrap method. The 80% quantile ranges between USD 28 and 40-44 million for Pareto, between USD 22 and 32-34 million for GPD, and between USD 29 and 37.9 million for the Bootstrap method. It is important to mention that these are conditional quantiles, given that the loss exceeds a threshold of USD 1 million. Hence, conditional on a loss overshooting the USD 1 million level, there is a 20% chance that it will be above USD 30 to 40 million, depending on the insurer's portfolio.

Panel B of table 9 shows results for price sensitivity analysis for different types of portfolios where we vary the exposure by different country. Similar to the previous case, different rows represent different business mixes, rebalancing the portfolio away from developing countries and into developed countries. Both in the case of the MCS and Bootstrap method, the average severity increases as we increase the share of developed countries in the portfolio. This is true even for the GPD models, for which the developed country class is less risky. This is due to the scale coefficient β: even if the developing region is more risky (for the GPD models), the size of the losses make both severity and quantiles higher. The effect is more pronounced for the Bootstrap method. Panel C of table 9 reports the 80% quantile, which is also seen to increase as developed country exposures become more important in the insurance portfolio.

In the case of MCS the average severity is generally lower than in the Bootstrap case, except for some simulations of the GPD case without covariates. However, panel C shows that the MCS quantile is in general the highest for the two Pareto models, meaning that Pareto models put more weights on the tail than the empirical distribution. The two GPD models have the lowest quantiles, same as for the occupancy. Average severity in the MCS case, for the two Pareto models, goes from USD 3 to 3.3 million as the ratio of developed to developing country exposures increases, whereas it goes from USD 4.1 to 5.1 million with the Bootstrap method, and it lies between USD 3.6 and 4.9 for the GPD models. Average severity is highest in the portfolio with only developed country exposures. The 80% quantile ranges between USD 18 and 77 million for the MCS method for Pareto models, between USD 22 and 48 million for the Bootstrap method, and goes from USD 16 to 43 million for the GPD models.

7. Structural break analysis

In this section we present structural breaks analysis for the period covered by our dataset, first for the whole sample, and then by conditioning on location. Figures 7 and 8 reports estimated coefficients for Hill estimator, year by year, using respectively USD 140k and USD 10m as thresholds. From both figures we can see that the estimated α's increase on average over time for both the whole sample and the developing country class, while they are almost constant on average for the developed country class.

Tables 10 and 11 report results for structural break tests relative to figure 7 and 8 respectively. The results confirm the visual interpretation that estimated coefficients are significantly increasing over time for the developing country class (and in the overall sample), while they remain stable for the developed country class. Hence, the tail risk of large commercial risk exposures in developing countries seem to decrease over time, presumably due to widespread take up of international risk mitigation standards and use of risk monitoring systems. The emerging wedge in tail risk between developed vs. developing countries may be due to the fact that developing markets' policy schedules have less complex exposures and systematically lower total insurable values (reflecting less sophisticated type and quality of infrastructure and machinery, for example). This interpretation seems to be confirmed by the descriptive statistics presented in table 3 for the CO and EON classes. To gain a better understanding of this result, we present in figure 9 total premiums written in the APAC region by one of the data providers, distinguishing between developed and developing countries. As we can see, insurance penetration is increasing over time for the developing country class, while it has remained stable for the developed country class. This could suggest that, on the one hand, (re)insurers' knowledge of developing countries has improved over time, allowing them to better underwrite risk. On the other hand, however, the emerging wedge in tail risk (associated with lower tail risk for developed countries) suggests that, once better underwriting and risk mitigation strategies are accounted for, a structural difference in LCR exposures for developing APAC countries is suggested by our data.

References

Beirlant, J., Goegebeur, Y., Segers, J., & Teugels, J. (2006). *Statistics of extremes: theory and applications*. John Wiley & Sons.

Benedetti, D., Biffis, E., and A. Milidonis (2015). *Large Commercial Risks (LCR) in Insurance: Focus on Asia-Pacific*, Insurance Risk and Finance Research Centre Technical report, Retrieved from www.irfrc.com.

Bernegger, S. (1997). The Swiss Re exposure curves and the MBBEFD distribution class, *ASTIN Bulletin*, 27(1), pp. 99-111.

Biffis, E. and Chavez, E. (2014). Tail risk in commercial property insurance, *Risks*, *2*(4), 393-410; doi:10.3390/risks2040393.

Buchanan, J. and Angelina, M. (2014). The hybrid reinsurance pricing method: A practitioner's guide. Technical report, to appear in *Variance.*

Chavez-Demoulin, V., Embrechts, P., & Hofert, M. (2015). An extreme value approach for modeling operational risk losses depending on covariates. *The Journal of Risk and Insurance*.

Dacorogna, M.M., Müller, U.A., Pictet, O.V., & De Vries, C.G. (1995). The Distribution of Extremal Foreign Exchange Rate Returns in Extremely Large Data Sets. Technical Report. Tinbergen Institute: Rotterdam, Netherland.

Dacorogna, M. M., Müller, U. A., Pictet, O. V., & De Vries, C. G. (2001). Extremal forex returns in extremely large data sets. *Extremes*, *4*(2), 105-127.

Desmedt, S., Snoussi, M., Chenut, X., & Walhin, J. F. (2012). Experience and exposure rating for property per risk excess of loss reinsurance revisited. *ASTIN bulletin*, *42*(01), 233-270.

Embrechts, P., Klüppelberg, C., & Mikosch, T. (1997). *Modelling extremal events* (Vol. 33). Springer Science & Business Media.

Gabaix, X. and Ibragimov, R. (2011). Rank- 1/2: a simple way to improve the OLS estimation of tail exponents. *Journal of Business & Economic Statistics*.

Gillott, N., Carroll, P., Chamberlin, G., Hudson, B., Malde, S., Masters, G., Taylor, P., and Thomson, A. (1988). *Commercial fire insurance*. Technical report.

Guggisberg, D. (2004). *Exposure Rating*. Swiss Re Technical Publishing Series.
Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, *3*(5), 1163-1174.

Huisman, R., Koedijk, K. G., Kool, C. J. M., & Palm, F. (2001). Tail-index estimates in small samples. *Journal of Business & Economic Statistics*, *19*(2), 208-216.

Ibragimov, R. and Walden, J. (2007). The limits of diversification when losses may be large. *Journal of Banking & Finance*, 31(8):2551–2569.

Klugman, S., H. Panjer, and G E. Willmot. 2012. *Loss models: from data to decisions*. Vol. 715. John Wiley & Sons.

Michaelides, N., Brown, P., Chacko, F., Graham, M., Haynes, J., Hindley, D., Howard, S., Johnson, H., Morgan, K., Pettengell, C., Rodriguez, R., and Simmons, D. (1997). The premium rating of commercial risks. Technical report, Working Party on Premium Rating of Commercial Risks, General Insurance Convention.

Pickands III, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics*, 119-131.

Riegel, U. (2010). On fire exposure rating and the impact of the risk profile type. *ASTIN Bulletin*, *40*(02), 727-777.

Swiss Re (2012). *Insuring ever-evolving commercial risks*. SIGMA, 2012.

Wang, H., & Tsai, C. L. (2009). Tail index regression. *Journal of the American Statistical Association*, *104*(487), 1233-1240.

Figures

Figure 1: The figure shows Parameter Estimates for FGU (losses from the ground up) for the whole dataset from 2000-2013, obtained with the three methodologies. The parameters reported on the y-axes are α for the Hill estimator, 1/γ for the Weighted Hill estimator and 1/ξ for the Generalized Pareto. The x-axes correspond to different threshold levels expressed as percentage of the total sample. Panel A refers to the full group. Panel B refers to developed countries only. Panel C refers to developing countries only.

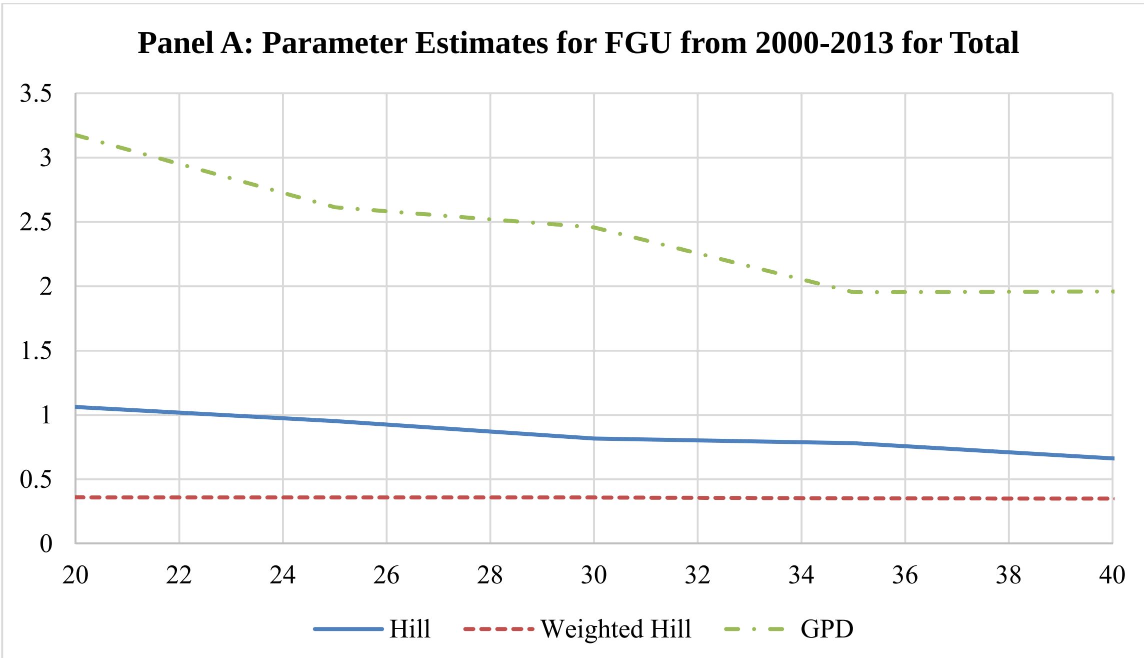

Figure 1: Continued.

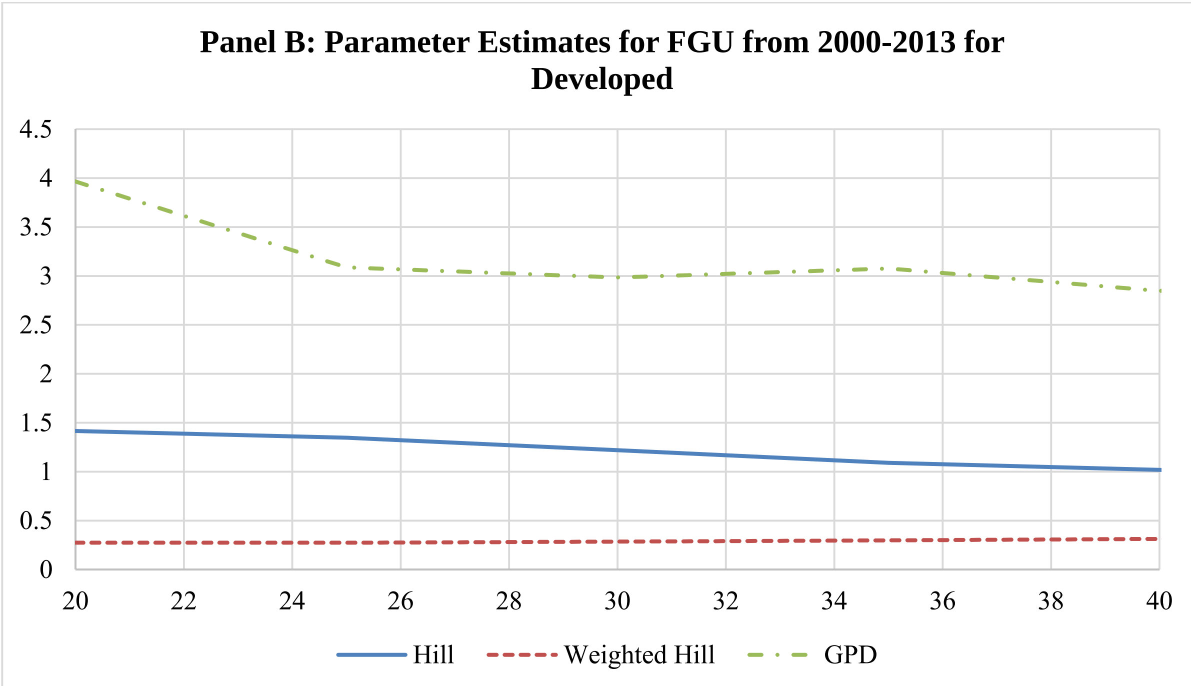

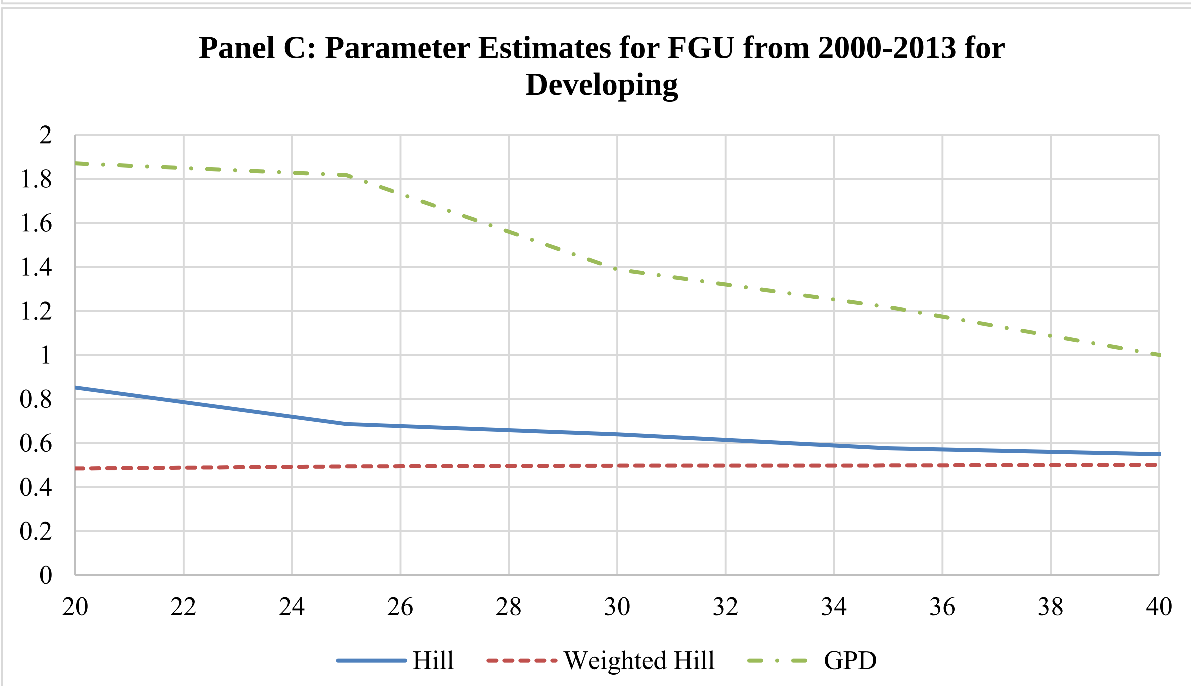

Figure 2: The figure shows Parameter Estimates for FGU (losses from the ground up) for Manufacturing occupancy only from 2000-2013, obtained with the three methodologies. The parameters reported on the y-axes are α for the Hill estimator, 1/γ for the Weighted Hill estimator and 1/ξ for the Generalized Pareto. The x-axes correspond to different threshold levels expressed as percentage of the total sample. Panel A refers to the full group. Panel B refers to developed countries only. Panel C refers to developing countries only.

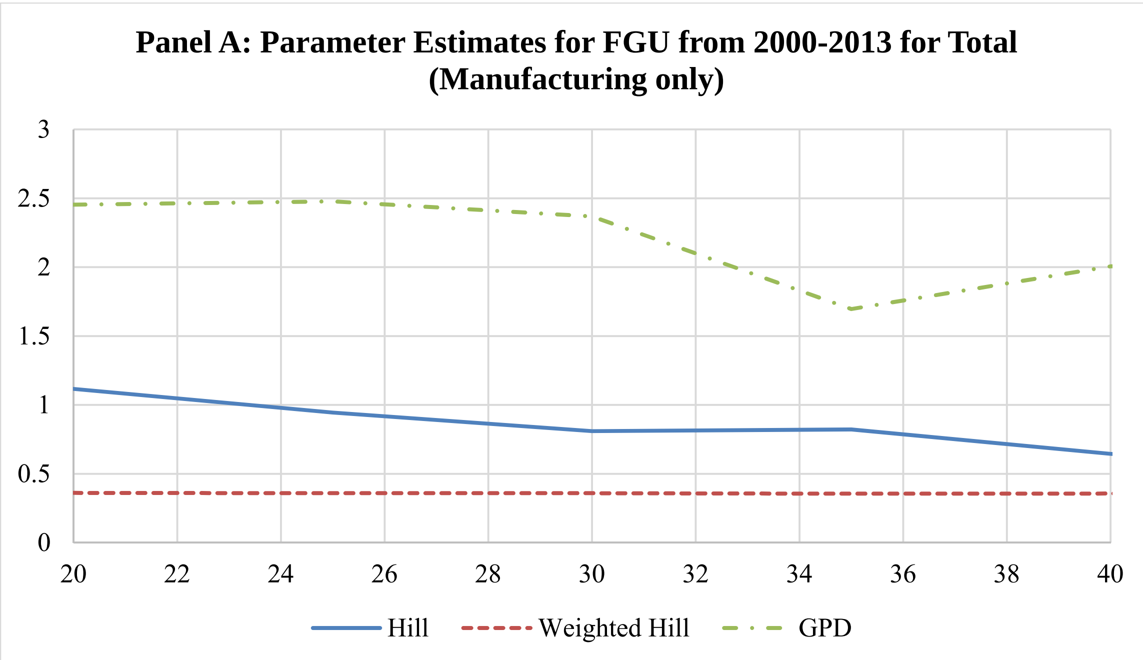

Figure 2: Continued.

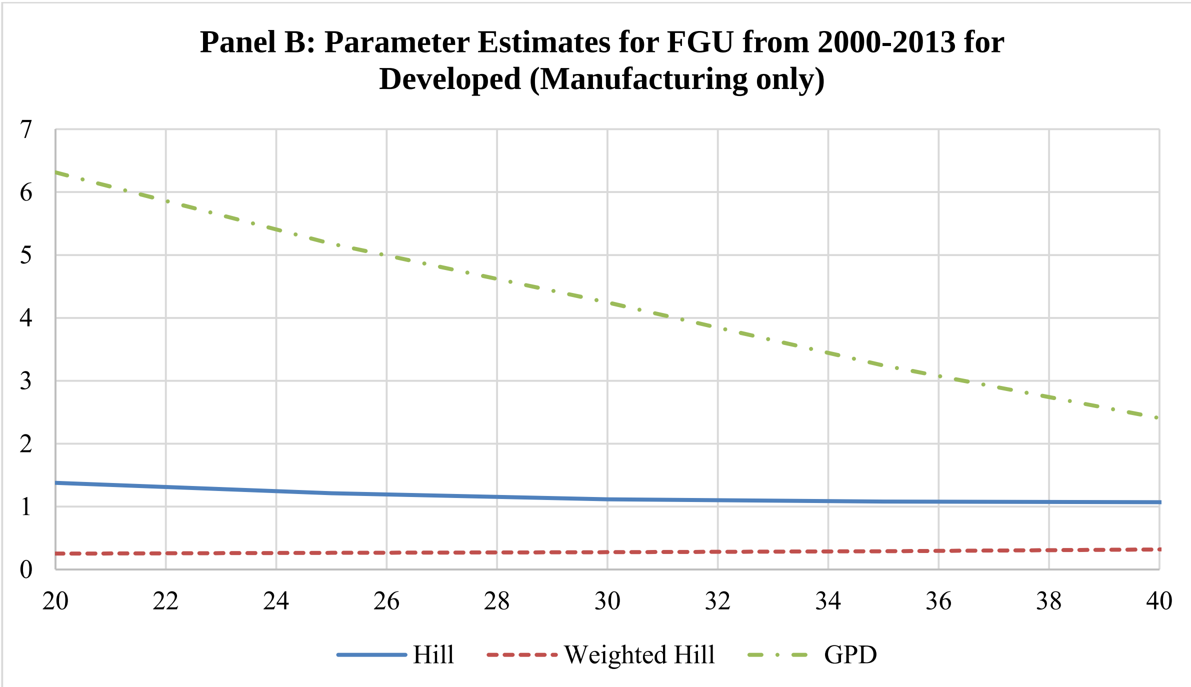

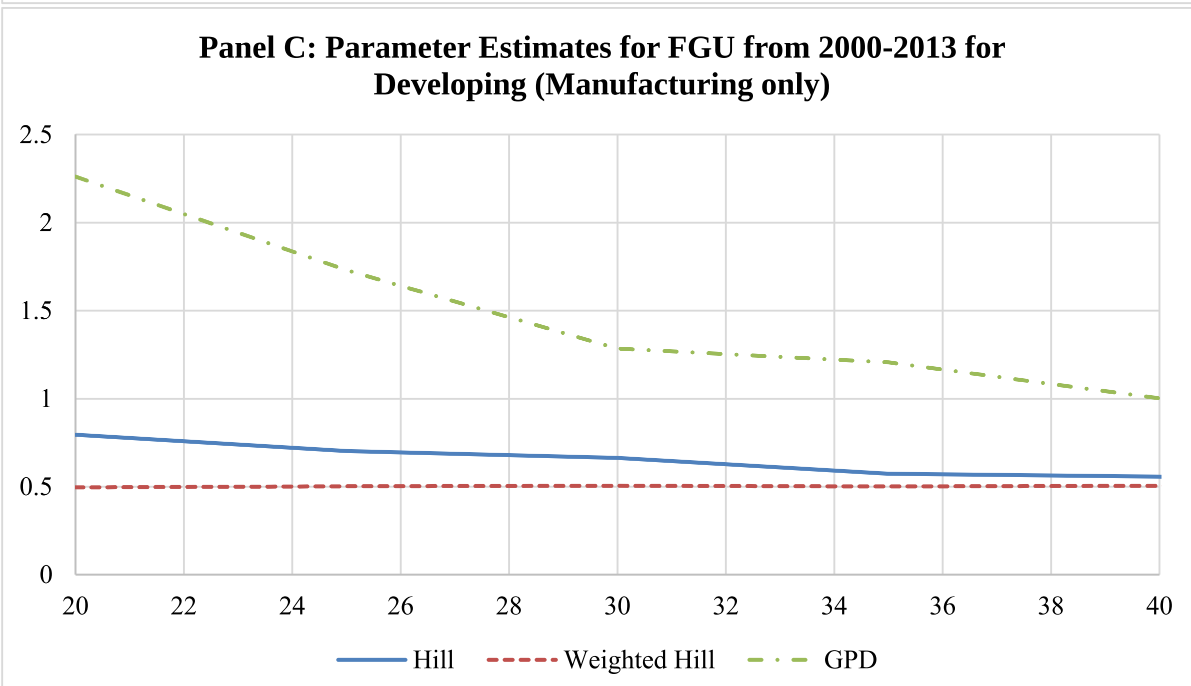

Figure 3: The figure shows Parameter Estimates for FGU (losses from the ground up) from 2000-2013, obtained with the Hill estimator. The parameter reported on the y-axes is α. The x-axes correspond to different threshold levels expressed as percentage of the total sample. Panel A shows results for total sample and split by occupancy type. Panel B shows results for total sample and split by Economic Development.

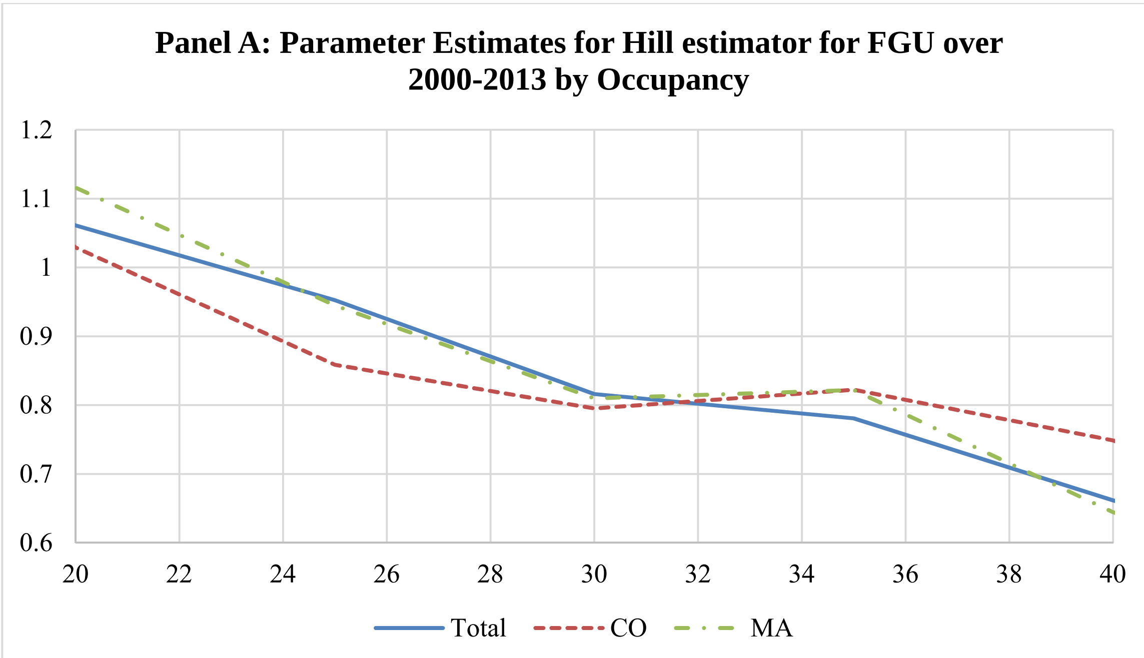

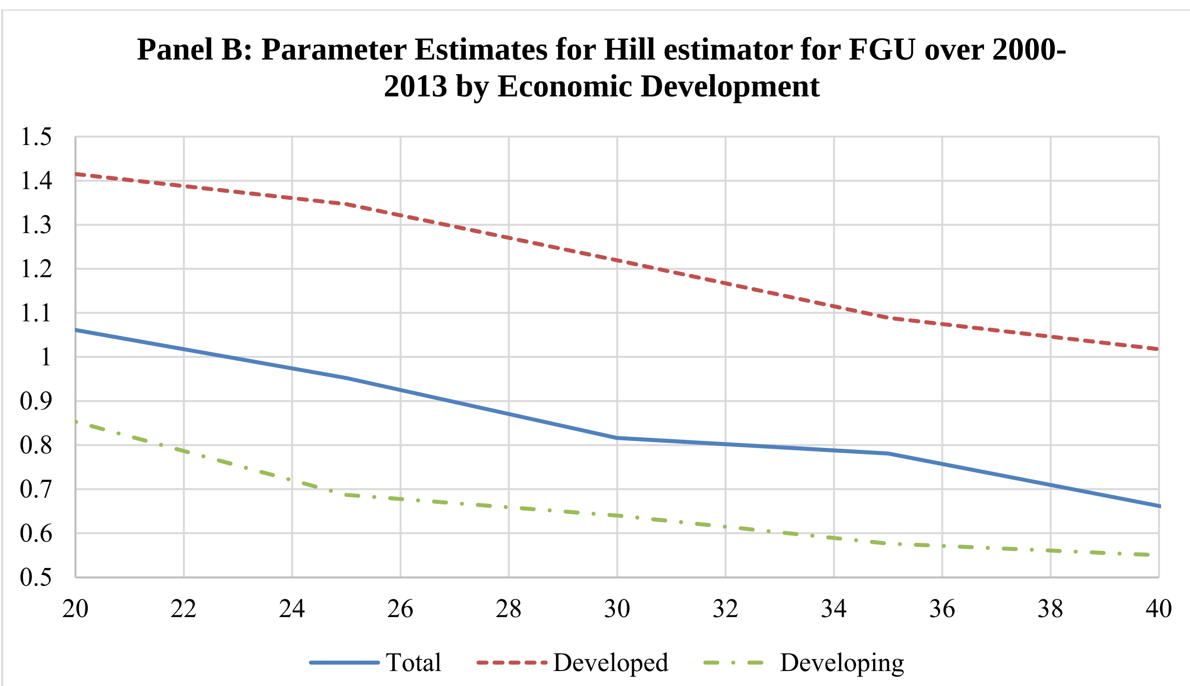

Figure 4: The figure shows Parameter Estimates for FGU (losses from the ground up) from 2000-2013, obtained with the Weighted Hill estimator. The parameter reported on the y-axes is 1/γ. The x-axes correspond to different threshold levels expressed as percentage of the total sample. Panel A shows results for total sample and split by occupancy type. Panel B shows results for total sample and split by Economic Development.

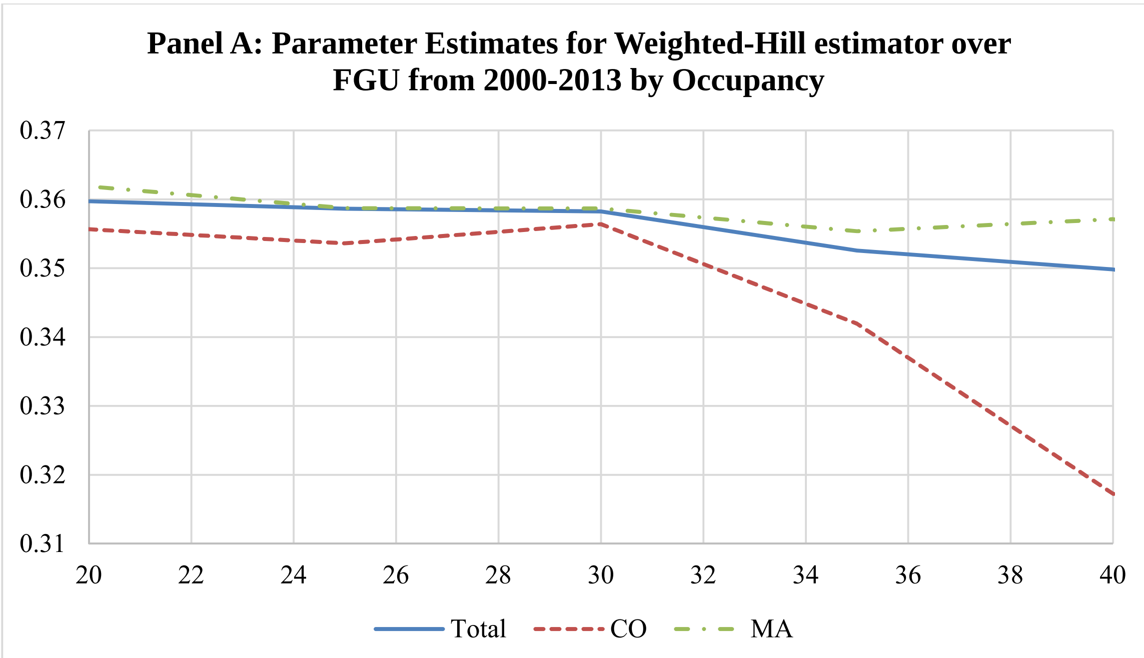

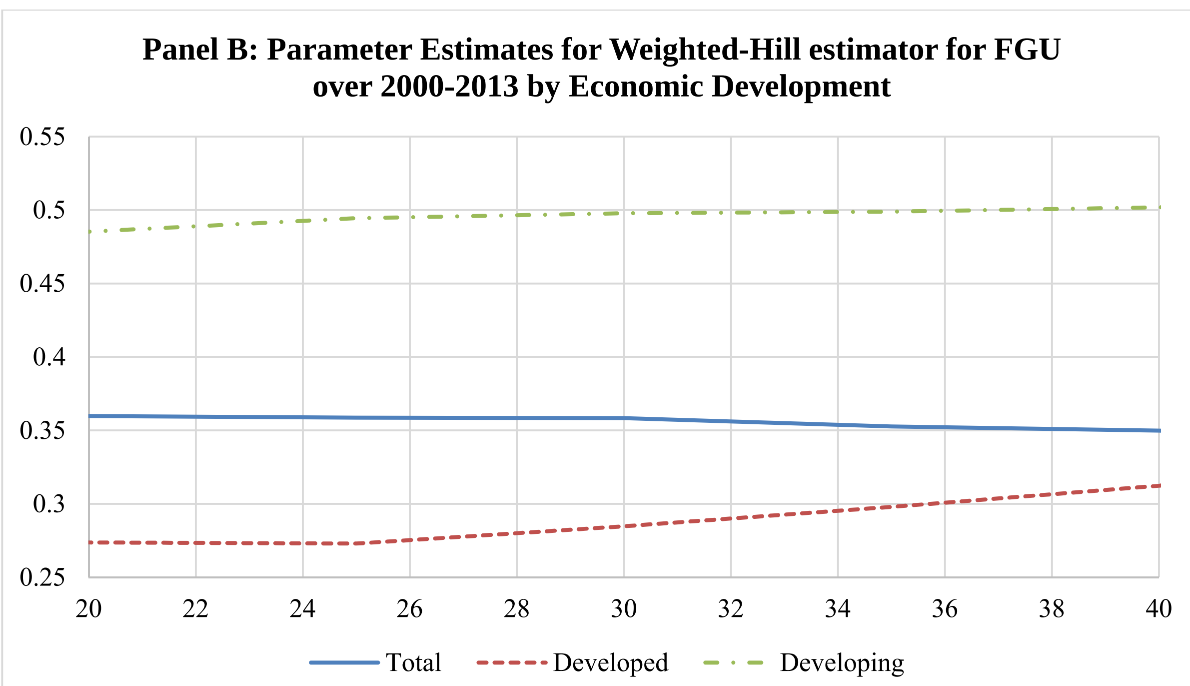

Figure 5: The figure shows Parameter Estimates for FGU (losses from the ground up) from 2000-2013, obtained with the Generalized Pareto distribution. The parameter reported on the y-axes is 1/ξ. The x-axes correspond to different threshold levels expressed as percentage of the total sample. Panel A shows results for total sample and split by occupancy type. Panel B shows results for total sample and split by Economic Development.

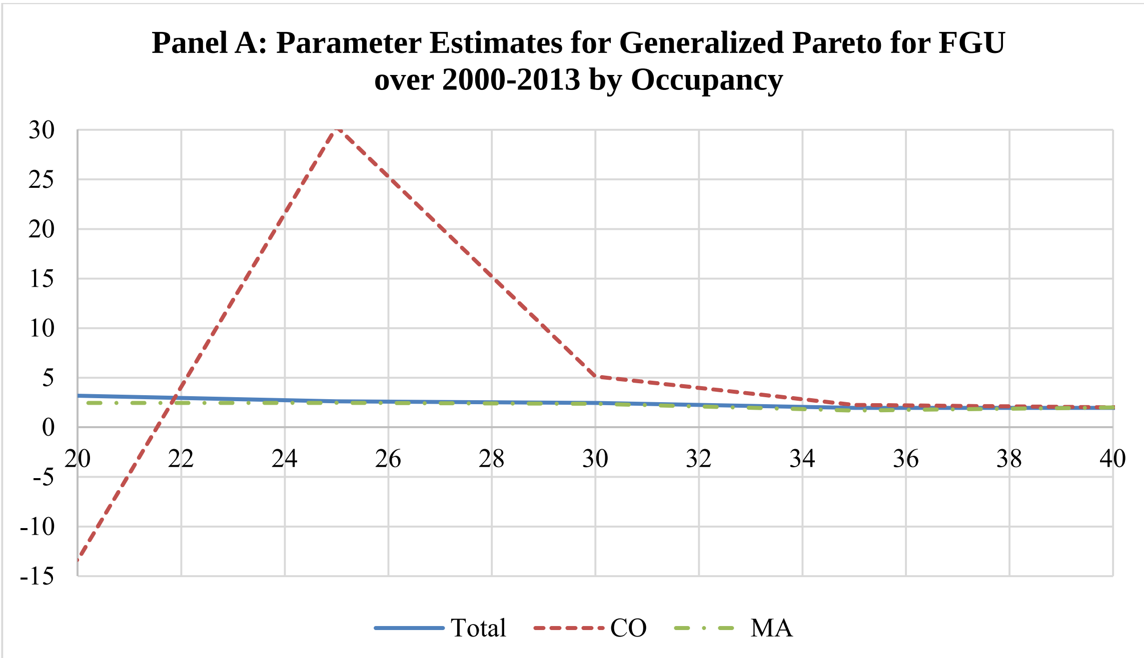

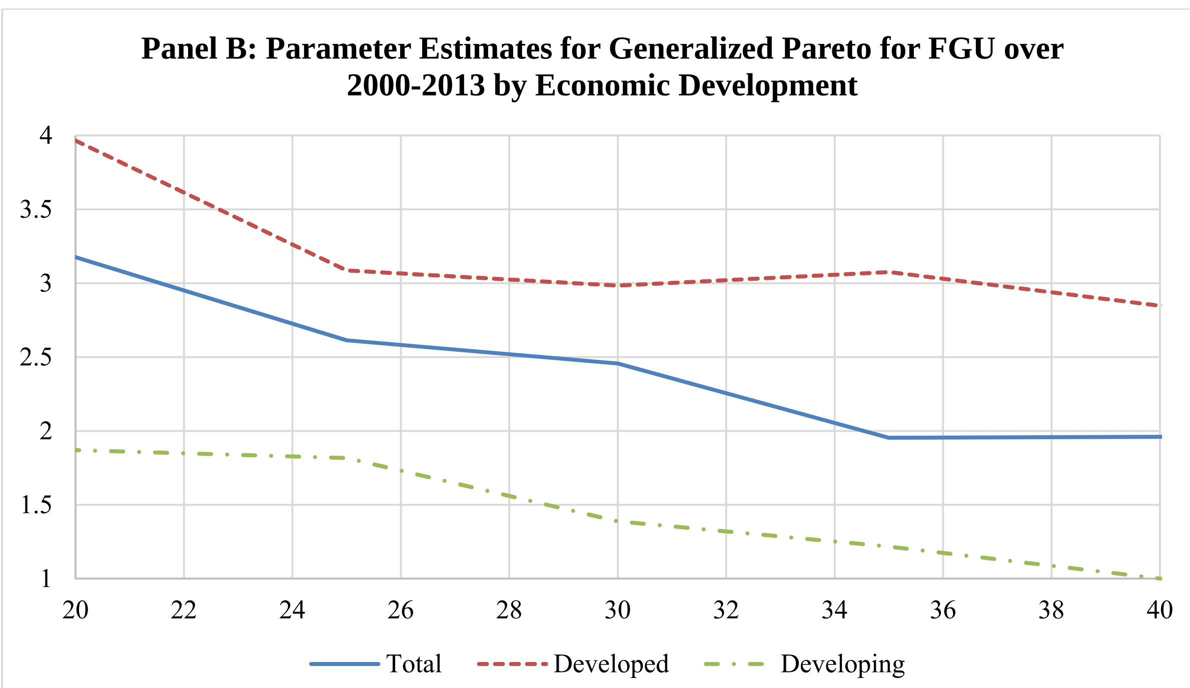

Figure 6: The figure shows empirical and estimated CDFs for FGU (losses from the ground up) from 2000-2013, obtained with the three different methodologies. The y-axes report the forecasted quantiles, while the x-axes correspond to different losses values scaled by the threshold. In all panels the threshold chosen is the median.

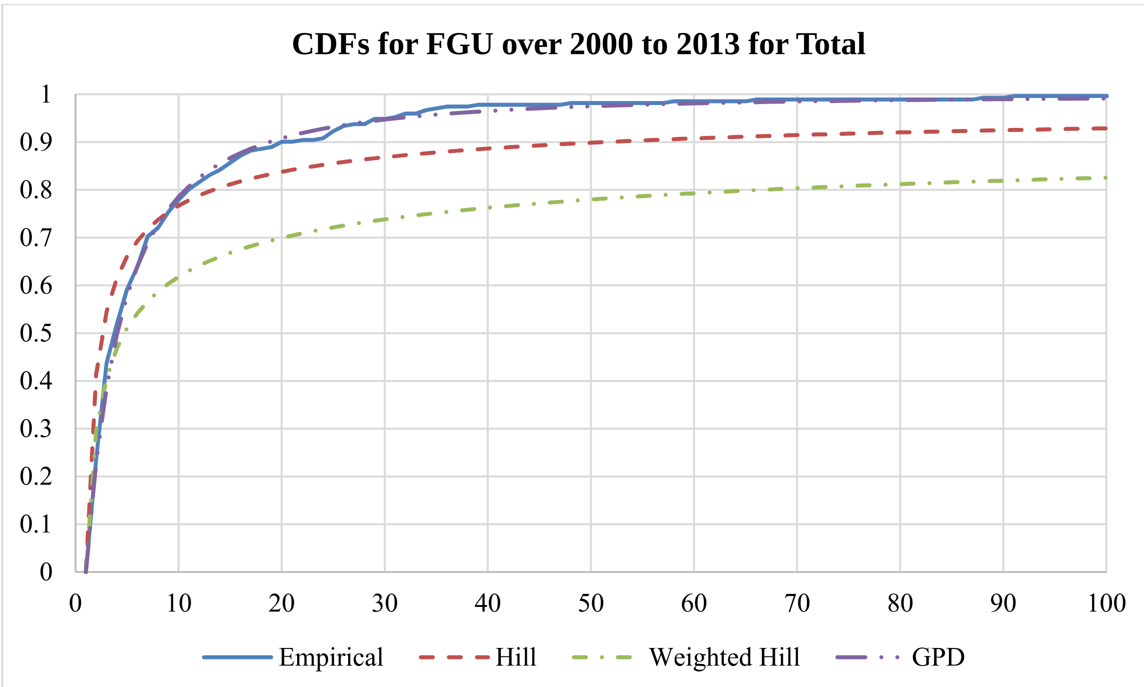

Figure 7: The figure shows Parameter Estimates for FGU (losses from the ground up) year by year, obtained with the Hill estimator, using USD 140 thousand as threshold. The parameter reported on the y-axes is α. The x-axes correspond to different years. Panel A refers to the whole sample. Panel B refers only to the developed region. Panel C refers only to the developing region.

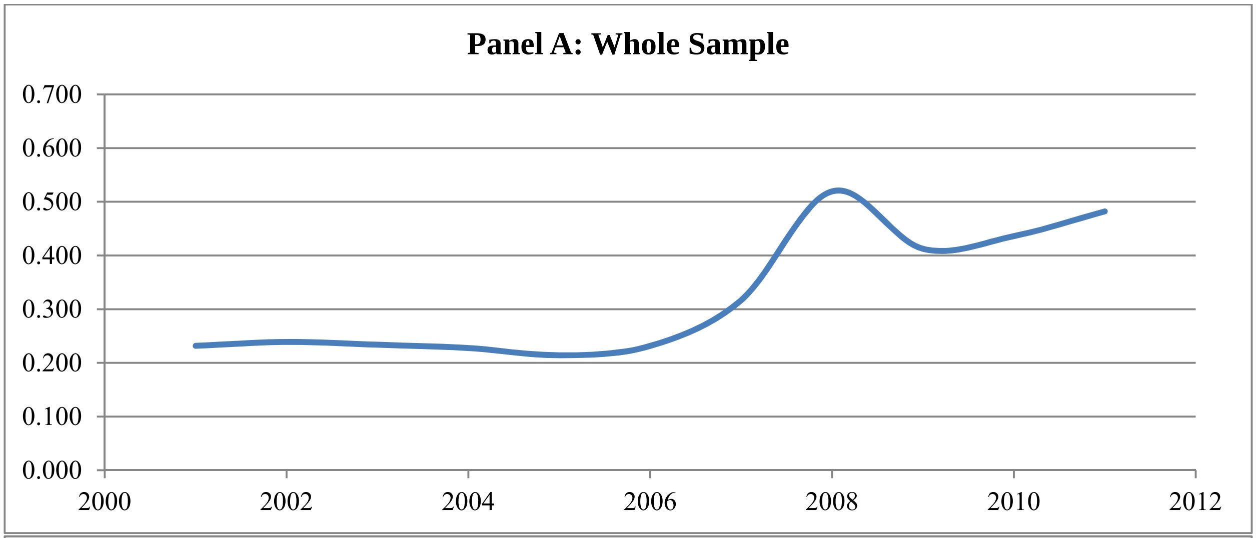

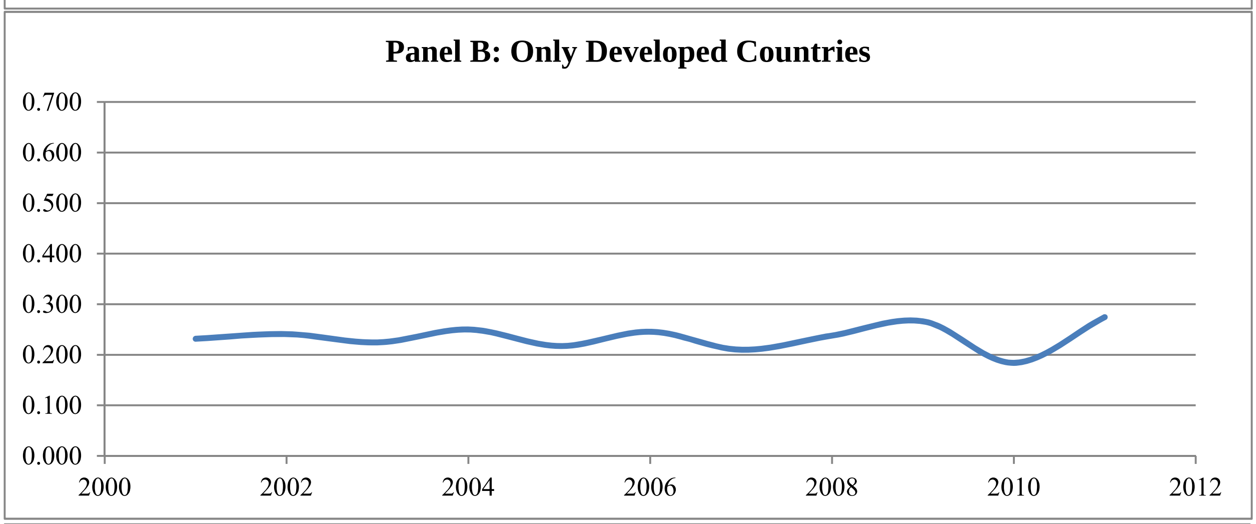

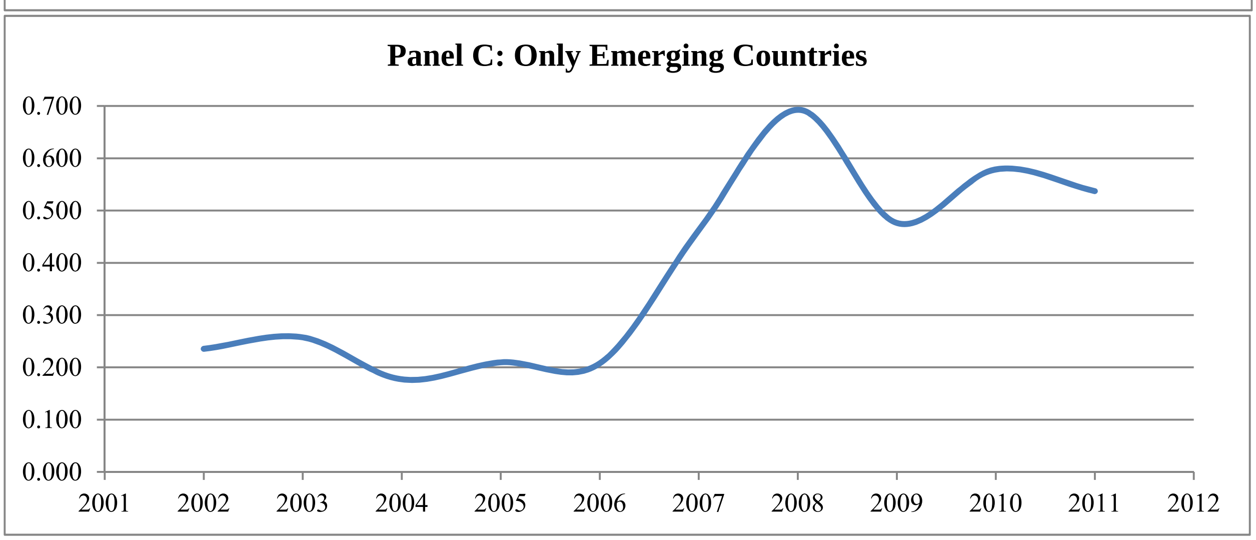

Figure 8: The figure shows Parameter Estimates for FGU (losses from the ground up) year by year, obtained with the Hill estimator, using USD 10 million as threshold. The parameter reported on the y-axes is α. The x-axes correspond to different years. Panel A refers to the whole sample. Panel B refers only to the developed region. Panel C refers only to the developing region.

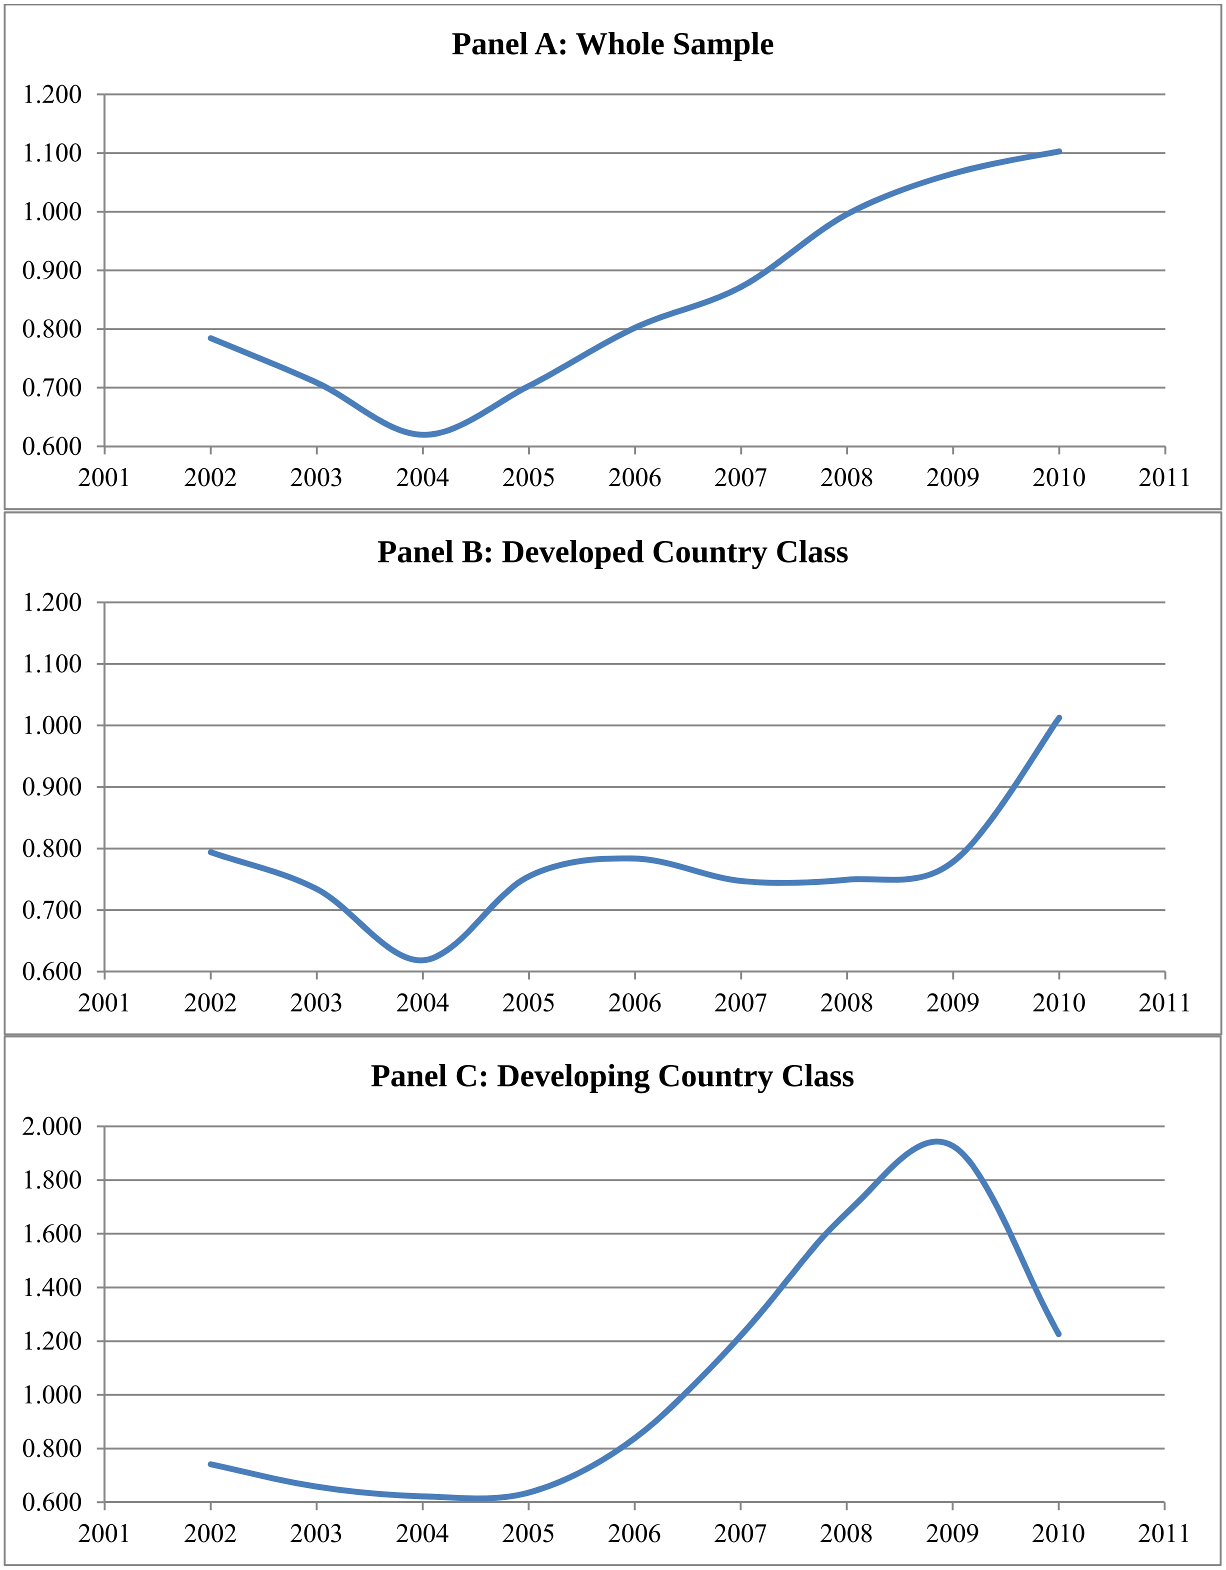

Figure 9: The figure shows Total Premium Written in the APAC region, and split by geographic region, for the period from 2000 to 2014.

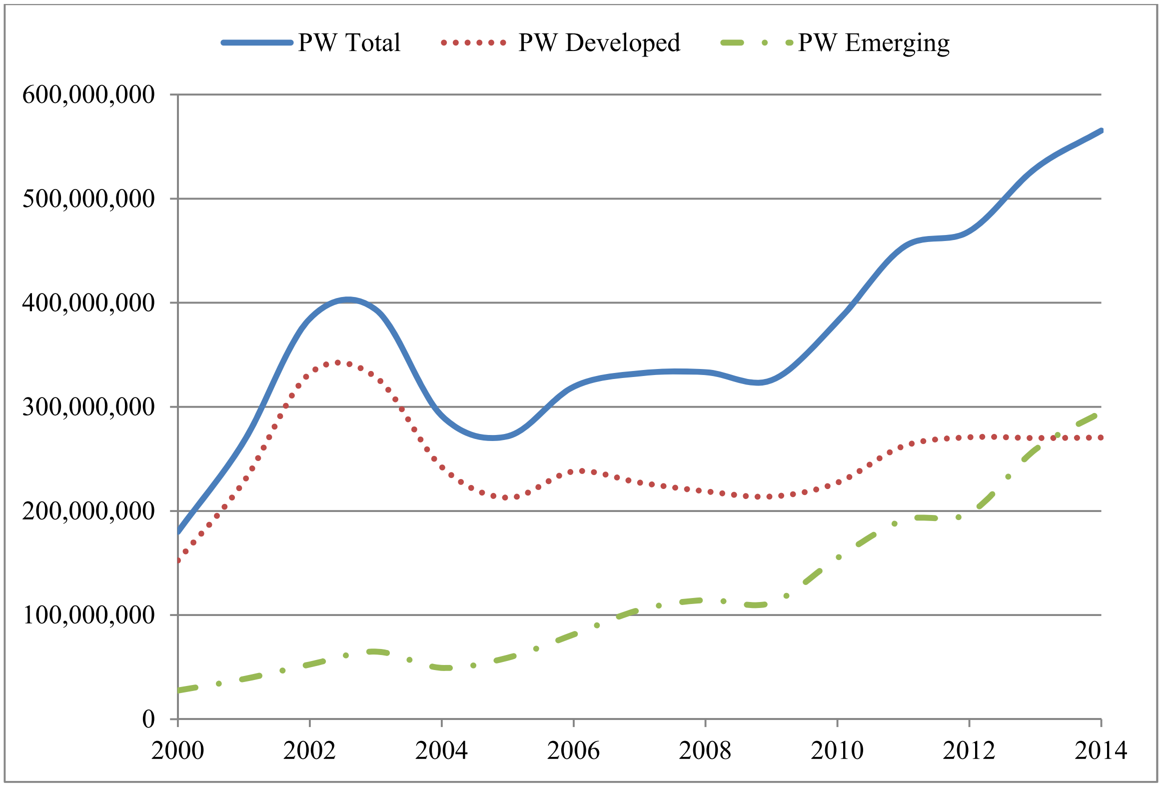

Tables

Table 1: The table shows the categories for the second level of exposure information provided in Biffis and Chavez (2014).

Code	Definition	Code	Definition
A	Miscellaneous	Q	Offices/Banks
B	Manufacturers/Processors	R	Residential
C	Chemicals/Pharmaceuticals	T	Transport
D	Bridges/Dams/Tunnels/Piers	U	Utilities
E	Conglomerates	V	Telecoms and Data Processing
F	Food	W	Woodworkers (Sawmills, Papermills)
G	Grain	X	Onshore Crude
H	General Mercantile/Shops	Y	Onshore GasPlants
J	Mines	Z	Onshore Construction
K	Crops	2	Hospital/Health care centres
L	Auto	4	Semiconductor/Fabs
M	Metals	5	Motor Manufaturers
O	Municipal Property	6	Warehouses
P	Energy (Oil Refineries/Petrochemicals)		

Table 2 (part 1): The table shows descriptive Statistics for FGU (from the ground up) losses for the entire sample period 2000-2013. The summary statistics reported are: Number of observations (N), Minimum and Maximum values (Min and Max), Average (Avg), Standard deviation (Std), Skewness (Skw), Kurtosis (Kur), First, Second and Third Quartile (Q1, Q2, Q3). Panel A refers to the whole dataset: total and split by Economic Development. Panel B refers to CO occupancy only: total and split by Economic Development. Panel C refers to EON occupancy only: total and split by Economic Development. Panel D refers to MA occupancy only: total and split by Economic Development.

	Panel A: Total		
	Total	Developed	Developing
N	508	180	328
Avg	15,635,610	30,323,265	7,575,311
Std	37,118,930	53,541,941	19,572,282
Skw	5	4	5
Kur	41	22	29
Min	140,503	157,431	140,503
Q1	438,854	4,370,947	306,261
Q2	2,543,246	10,647,249	766,226
Q3	13,485,010	34,863,759	5,920,833
Max	365,698,968	365,698,968	171,771,593

	Panel B: Commercial		
	Total	Developed	Developing
N	81	33	48
Avg	11,422,102	15,210,288	8,817,725
Std	22,420,610	20,923,211	23,250,065
Skw	3	2	4
Kur	13	6	19
Min	142,156	157,431	142,156
Q1	341,340	2,095,801	258,593
Q2	2,127,620	5,952,546	653,844
Q3	9,563,530	16,230,101	5,840,551
Max	112,806,811	82,730,286	112,806,811

Table 2 (part 2): The table shows descriptive Statistics for FGU (from the ground up) for the entire sample period 2000-2013. The summary statistics reported are: Number of observations (N), Minimum and Maximum values (Min and Max), Average (Avg), Standard deviation (Std), Skewness (Skw), Kurtosis (Kur), First, Second and Third Quartile (Q1, Q2, Q3). Panel A refers to the whole dataset: total and split by Economic Development. Panel B refers to CO occupancy only: total and split by Economic Development. Panel C refers to EON occupancy only: total and split by Economic Development. Panel D refers to MA occupancy only: total and split by Economic Development.

	Panel C: Energy on Shore		
	Total	Developed	Developing
N	44	20	24
Avg	23,770,287	32,801,006	16,244,688
Std	57,012,275	80,107,204	25,410,331
Skw	5	4	3
Kur	35	21	17
Min	533,472	533,472	535,567
Q1	5,971,383	4,632,006	6,386,209
Q2	8,240,584	9,182,496	8,179,591
Q3	15,413,632	28,222,177	14,266,922
Max	365,698,968	365,698,968	121,515,749
	Panel D: Manufacturing		
	Total	Developed	Developing
N	383	127	256
Avg	15,592,185	33,860,064	6,529,604
Std	36,689,582	54,040,150	18,025,050
Skw	5	3	5
Kur	34	17	37
Min	140,503	218,268	140,503
Q1	409,063	5,634,221	295,398
Q2	1,882,356	14,057,501	683,688
Q3	13,501,743	35,865,893	3,379,646
Max	327,487,391	327,487,391	171,771,593

Table 3 (part 1): The table shows descriptive Statistics for TIV* (total insurable value) for the entire sample period 2000-2013. The summary statistics reported are: Number of observations (N), Minimum and Maximum values (Min and Max), Average (Avg), Standard deviation (Std), Skewness (Skw), Kurtosis (Kur), First, Second and Third Quartile (Q1, Q2, Q3). Panel A refers to the whole dataset: total and split by Economic Development. Panel B refers to CO occupancy only: total and split by Economic Development. Panel C refers to EON occupancy only: total and split by Economic Development. Panel D refers to MA occupancy only: total and split by Economic Development.

	Panel A: Total		
	Total	Developed	Developing
N	452	133	319
Avg	2,166,960,565	2,425,663,982	2,059,099,893
Std	15,289,334,724	9,998,003,110	17,029,396,779
Skw	5	4	5
Kur	41	22	29
Min	787,252	1,069,725	787,252
Q1	268,716,208	42,571,677	441,748,792
Q2	780,976,256	334,003,896	900,762,469
Q3	1,454,611,208	1,218,786,562	1,494,971,080
Max	304,794,042,837	100,546,318,805	304,794,042,837
	Panel B: Commercial		
	Total	Developed	Developing
N	72	24	48
Avg	865,121,287	935,732,856	829,815,502
Std	917,555,393	1,130,830,023	801,383,175
Skw	3	2	4
Kur	13	6	19
Min	787,252	1,851,096	787,252
Q1	174,860,492	38,082,152	292,019,131
Q2	689,759,816	288,215,111	705,774,330
Q3	1,127,453,593	1,885,655,377	1,014,072,721
Max	4,001,683,538	3,664,345,555	4,001,683,538

*TIV was replaced by TSI (total sum insured) when unavailable.

Table 3 (part 2): The table shows descriptive Statistics for TIV* (total insurable value) for the entire sample period 2000-2013. The summary statistics reported are: Number of observations (N), Minimum and Maximum values (Min and Max), Average (Avg), Standard deviation (Std), Skewness (Skw), Kurtosis (Kur), First, Second and Third Quartile (Q1, Q2, Q3). Panel A refers to the whole dataset: total and split by Economic Development. Panel B refers to CO occupancy only: total and split by Economic Development. Panel C refers to EON occupancy only: total and split by Economic Development. Panel D refers to MA occupancy only: total and split by Economic Development.

	Panel C: Energy on Shore		
	Total	Developed	Developing
N	39	15	24
Avg	1,688,299,126	2,438,517,314	1,219,412,758
Std	3,388,105,740	5,336,299,469	1,017,834,334
Skw	5	4	3
Kur	35	21	17
Min	1,069,725	1,069,725	97,688,983
Q1	244,006,261	39,015,793	526,578,931
Q2	912,685,266	256,525,203	1,006,514,269
Q3	1,452,157,006	2,576,308,696	1,388,220,841
Max	20,706,158,419	20,706,158,419	3,883,766,841
	Panel D: Manufacturing		
	Total	Developed	Developing
N	341	94	247
Avg	2,496,579,991	2,804,020,865	2,379,578,201
Std	17,553,569,883	11,685,813,090	19,344,047,633
Skw	5	3	5
Kur	34	17	37
Min	3,044,195	4,988,207	3,044,195
Q1	312,265,655	43,111,663	463,227,247
Q2	804,911,129	346,898,230	957,255,016
Q3	1,493,503,149	1,087,026,440	1,611,459,713
Max	304,794,042,837	100,546,318,805	304,794,042,837

*TIV was replaced by TSI (total sum insured) when unavailable.

Table 4: The table refers to figure 1 and shows t-Stats of the estimated coefficients and KS-test for goodness-of-fit on FGU (losses from the ground up) for the whole dataset from 2000-2013, obtained with the three methodologies on different thresholds. The parameters tested are α for the Hill estimator, γ for the Weighted Hill estimator and ξ for the Generalized Pareto. The null hypothesis of the KS test is that the distribution being estimated is the correct one. Statistically significant values are shown in bold and italics. Panel A refers to the full group. Panel B refers to developed countries only. Panel C refers to developing countries only.

			Panel A					
			Hill		Weighted Hill		GPD	
%Obs	#Obs	Threshold	t-Stat	KS pvalue	t-Stat	KS pvalue	t-Stat	KS pvalue
20	93	18,205,761	***9.6***	***0.930***	***9.1***	0.000	***2.1***	***0.978***
25	116	13,097,962	***10.8***	***0.243***	***9.4***	0.000	***2.7***	***0.985***
30	139	8,895,495	***11.8***	0.049	***9.7***	0.000	***3.1***	***0.977***
35	162	6,980,850	***12.7***	0.050	***10.2***	0.000	***3.8***	***0.974***
40	186	4,581,297	***13.6***	0.006	***10.8***	0.000	***4.2***	***0.767***

			Panel B					
			Hill		Weighted Hill		GPD	
%Obs	#Obs	Threshold	t-Stat	KS pvalue	t-Stat	KS pvalue	t-Stat	KS pvalue
20	32	42,504,805	***5.7***	***0.699***	***6.9***	0.000	1.0	***0.999***
25	40	34,863,759	***6.3***	***1.000***	***8.0***	0.000	1.4	***1.000***
30	48	27,776,326	***6.9***	***0.968***	***8.8***	0.000	1.6	***1.000***
35	56	21,965,627	***7.5***	***0.969***	***9.9***	0.000	1.7	***1.000***
40	64	18,175,338	***8.0***	***0.924***	***11.0***	0.000	1.9	***1.000***

			Panel C					
			Hill		Weighted Hill		GPD	
%Obs	#Obs	Threshold	t-Stat	KS pvalue	t-Stat	KS pvalue	t-Stat	KS pvalue
20	61	6,206,870	***7.8***	***0.548***	***9.1***	0.005	***2.3***	***0.996***
25	76	3,484,732	***8.7***	***0.161***	***9.3***	0.033	***2.8***	***0.998***
30	91	2,372,768	***9.5***	***0.226***	***9.5***	***0.074***	***3.5***	***1.000***
35	106	1,554,141	***10.3***	***0.070***	***9.8***	***0.168***	***4.1***	***0.939***
40	122	1,110,381	***11.0***	***0.157***	***10.1***	***0.380***	***4.8***	***0.808***

Table 5: The table refers to figure 2 and shows t-Stats of the estimated coefficients and KS-test for goodness-of-fit on FGU (losses from the ground up) for the whole period, 2000-2013, obtained with the three methodologies on different thresholds in the subsample of manufacturing occupancy only. The parameters tested are α for the Hill estimator, γ for the Weighted Hill estimator and ξ for the Generalized Pareto. The null hypothesis of the KS test is that the distribution being estimated is the correct one. Statistically significant values are shown in bold and italics. Panel A refers to the full group. Panel B refers to developed countries only. Panel C refers to developing countries only.

			Panel A					
			Hill		Weighted Hill		GPD	
%Obs	#Obs	Threshold	t-Stat	KS pvalue	t-Stat	KS pvalue	t-Stat	KS pvalue
20	77	19,814,119	***8.8***	***0.832***	***8.3***	0.000	***2.2***	***0.968***
25	96	13,501,743	***9.8***	***0.497***	***8.6***	0.000	***2.5***	***0.968***
30	115	9,171,755	***10.7***	***0.138***	***8.8***	0.000	***2.9***	***0.963***
35	134	7,761,958	***11.6***	***0.307***	***9.3***	0.000	***3.7***	***0.913***
40	153	4,616,407	***12.4***	0.022	***9.6***	0.000	***3.9***	***0.882***

			Panel B					
			Hill		Weighted Hill		GPD	
%Obs	#Obs	Threshold	t-Stat	KS pvalue	t-Stat	KS pvalue	t-Stat	KS pvalue
20	25	47,705,187	***5.0***	***0.977***	***6.6***	0.000	0.6	***1.000***
25	32	35,865,893	***5.7***	***0.909***	***7.5***	0.000	0.8	***1.000***
30	38	28,819,445	***6.2***	***0.571***	***8.5***	0.000	1.1	***0.927***
35	44	24,560,687	***6.6***	***0.823***	***9.3***	0.000	1.4	***1.000***
40	51	21,180,193	***7.1***	***0.842***	***9.8***	0.000	1.8	***0.921***

			Panel C					
			Hill		Weighted Hill		GPD	
%Obs	#Obs	Threshold	t-Stat	KS pvalue	t-Stat	KS pvalue	t-Stat	KS pvalue
20	51	5,425,747	***7.1***	***0.803***	***8.5***	0.015	1.9	***0.998***
25	64	3,379,646	***8.0***	***0.322***	***8.7***	***0.051***	***2.6***	***0.999***
30	77	2,344,706	***8.8***	***0.405***	***9.0***	***0.096***	***3.3***	***0.995***
35	90	1,428,214	***9.5***	***0.291***	***9.4***	***0.221***	***3.8***	***0.953***
40	102	1,085,713	***10.1***	***0.245***	***9.6***	***0.317***	***4.4***	***0.926***

Table 6: The table shows results for the Hill regression with covariates. The parameters are estimated for different levels of thresholds ranging from 15% of observation (85th percentile) to 100% (full sample). Dummy variables implemented are relative to the constant. The dummy for manufacturing (Manufact.) is relative to commercial, the dummy for developing (Emerg.) is relative to developed. The only continuous variable is the logarithm of total insured values (LogTIV*), to control for size effects. Significant t-Statistics (above 1.96) are in bold and italics. The models are estimated with the exclusion of energy-on-shore. Panel A represents the baseline model (does not use LogTIV*). Panel B adds LogTIV*.

		Panel A: Hill Regression		
%Obs		Const.	Manufac.	Emerg.
15	Coeff	0.457	-0.323	0.117
	t-stat	1.46	1.00	0.44
20	Coeff	0.091	-0.081	0.127
	t-stat	0.32	0.28	0.56
25	Coeff	-0.039	-0.079	0.184
	t-stat	0.15	0.31	0.93
30	Coeff	-0.265	0.028	0.117
	t-stat	1.11	0.11	0.65
35	Coeff	-0.295	-0.004	0.155
	t-stat	1.35	0.02	0.93
40	Coeff	-0.371	-0.132	0.199
	t-stat	1.95	0.68	1.30
45	Coeff	-0.580	-0.122	0.230
	t-stat	***3.27***	0.67	1.62
50	Coeff	-0.654	-0.181	0.314
	t-stat	***3.99***	1.08	***2.36***
55	Coeff	-0.870	-0.088	0.347
	t-stat	***5.43***	0.54	***2.75***
60	Coeff	-0.936	-0.115	0.384
	t-stat	***6.18***	0.75	***3.19***
65	Coeff	-1.060	-0.048	0.468
	t-stat	***7.11***	0.32	***4.06***
70	Coeff	-1.219	0.036	0.519
	t-stat	***8.17***	0.24	***4.67***
75	Coeff	-1.293	0.067	0.577
	t-stat	***8.80***	0.46	***5.34***
80	Coeff	-1.291	-0.007	0.622
	t-stat	***9.22***	0.05	***5.91***
85	Coeff	-1.370	0.040	0.654
	t-stat	***9.84***	0.30	***6.35***
90	Coeff	-1.372	0.000	0.700
	t-stat	***10.20***	0.00	***6.93***
95	Coeff	-1.398	-0.007	0.719
	t-stat	***10.62***	0.05	***7.25***
100	Coeff	-1.416	-0.059	0.716
	t-stat	***11.13***	0.48	***7.32***

		Panel B: Hill Regression			
%Obs		Const.	Manufac.	Emerg.	LogTIV*
15	Coeff	1.399	0.119	0.139	-0.068
	t-stat	1.08	0.33	0.53	1.03
20	Coeff	1.901	0.103	0.170	-0.103
	t-stat	1.73	0.34	0.75	1.82
25	Coeff	1.987	0.200	0.083	-0.118
	t-stat	***2.03***	0.72	0.41	***2.35***
30	Coeff	1.529	0.051	0.140	-0.093
	t-stat	1.74	0.21	0.76	***2.08***
35	Coeff	1.263	-0.078	0.093	-0.082
	t-stat	1.59	0.36	0.55	***2.02***
40	Coeff	0.402	-0.067	0.190	-0.049
	t-stat	0.54	0.33	1.21	1.30
45	Coeff	-0.186	-0.132	0.277	-0.022
	t-stat	0.26	0.72	1.88	0.63
50	Coeff	-0.485	-0.016	0.321	-0.019
	t-stat	0.72	0.09	***2.30***	0.57
55	Coeff	-0.580	-0.010	0.376	-0.020
	t-stat	0.90	0.06	***2.82***	0.60
60	Coeff	-0.695	0.020	0.423	-0.018
	t-stat	1.12	0.12	***3.30***	0.58
65	Coeff	-0.955	0.103	0.481	-0.011
	t-stat	1.59	0.63	***3.87***	0.38
70	Coeff	-1.239	0.119	0.521	-0.001
	t-stat	***2.12***	0.75	***4.31***	0.03
75	Coeff	-1.558	0.106	0.569	0.013
	t-stat	***2.73***	0.68	***4.81***	0.46
80	Coeff	-1.681	0.054	0.606	0.019
	t-stat	***3.01***	0.37	***5.23***	0.69
85	Coeff	-1.733	0.036	0.650	0.021
	t-stat	***3.17***	0.25	***5.70***	0.78
90	Coeff	-1.705	0.060	0.681	0.017
	t-stat	***3.18***	0.43	***6.06***	0.64
95	Coeff	-1.749	0.040	0.678	0.018
	t-stat	***3.31***	0.29	***6.13***	0.69
100	Coeff	-1.752	-0.011	0.690	0.018
	t-stat	***3.36***	0.08	***6.30***	0.67

Table 7: The table shows results for the GPD regression with covariates. The parameters are estimated for different levels of thresholds ranging from 15% of observation (85th percentile) to 100% (full sample). Dummy variables used are relative to the constant. The dummy for manufacturing (Manufact.) is relative to commercial, the dummy for developing (Emerg.) is relative to developed. The only continuous variable is the logarithm of total insured values (LogTIV*). Significant t-Statistics (above 1.96) are in bold and italics. Models are estimated with the exclusion of energy-on-shore. Panel A shows the baseline model (does not use LogTIV*). Panel B adds LogTIV*. Coefficients are estimated relative to the ξ parameter here, hence the average marginal effect to $\alpha=1/\xi$ is opposite to their sign.

		Panel A: GPD Regression		
%Obs		Const.	Manufac.	Emerg.
15	Coeff	0.548	-0.306	-0.133
	t-stat	0.74	0.42	0.29
20	Coeff	-0.936	1.271	-0.033
	t-stat	***31.68***	***7.31***	1.53
25	Coeff	0.154	0.221	0.075
	t-stat	0.34	0.50	0.19
30	Coeff	0.117	0.349	-0.075
	t-stat	0.31	0.94	0.23
35	Coeff	0.247	0.281	0.030
	t-stat	0.63	0.74	0.09
40	Coeff	0.500	-0.042	0.080
	t-stat	1.34	0.11	0.30
45	Coeff	0.478	-0.042	0.149
	t-stat	1.45	0.13	0.61
50	Coeff	0.709	-0.298	0.370
	t-stat	1.92	0.80	1.46
55	Coeff	0.536	-0.080	0.450
	t-stat	1.70	0.25	1.84
60	Coeff	0.687	-0.218	0.584
	t-stat	***2.10***	0.65	***2.37***
65	Coeff	0.607	-0.107	0.871
	t-stat	1.90	0.32	***3.31***
70	Coeff	0.539	-0.014	0.956
	t-stat	1.73	0.04	***3.81***
75	Coeff	0.610	-0.078	1.119
	t-stat	1.87	0.23	***4.42***
80	Coeff	0.792	-0.299	1.201
	t-stat	***2.28***	0.84	***5.00***
85	Coeff	0.717	-0.187	1.207
	t-stat	***2.18***	0.56	***5.17***
90	Coeff	0.772	-0.247	1.274
	t-stat	***2.36***	0.75	***5.57***
95	Coeff	0.805	-0.264	1.248
	t-stat	***2.49***	0.81	***5.63***
100	Coeff	0.876	-0.344	1.147
	t-stat	***2.76***	1.09	***5.50***

		Panel B: GPD Regression			
%Obs		Const.	Manufac.	Emerg.	LogTIV*
15	Coeff	2.224	1.156	-0.288	-0.135
	t-stat	***2.28***	***7.11***	1.50	***5.66***
20	Coeff	1.294	0.201	-0.059	-0.055
	t-stat	0.86	0.41	0.18	0.77
25	Coeff	1.718	0.865	-0.449	-0.097
	t-stat	0.36	1.12	0.96	0.48
30	Coeff	0.775	-0.131	-0.085	-0.009
	t-stat	0.59	0.27	0.28	0.14
35	Coeff	0.624	-0.317	-0.103	0.007
	t-stat	0.52	0.78	0.38	0.12
40	Coeff	-0.495	-0.306	0.124	0.062
	t-stat	0.48	0.81	0.49	1.13
45	Coeff	-1.101	-0.514	0.364	0.104
	t-stat	1.03	1.22	1.35	1.86
50	Coeff	-1.598	-0.185	0.382	0.116
	t-stat	1.41	0.48	1.48	***2.00***
55	Coeff	-1.316	-0.285	0.546	0.107
	t-stat	1.16	0.75	***2.01***	1.75
60	Coeff	-0.842	-0.238	0.696	0.083
	t-stat	0.69	0.64	***2.41***	1.25
65	Coeff	-0.799	-0.073	0.813	0.075
	t-stat	0.59	0.20	***2.77***	1.06
70	Coeff	-0.714	-0.109	0.849	0.074
	t-stat	0.50	0.29	***2.93***	1.01
75	Coeff	-0.670	-0.177	0.902	0.075
	t-stat	0.45	0.46	***3.09***	1.00
80	Coeff	0.177	-0.306	1.005	0.037
	t-stat	0.12	0.78	***3.55***	0.51
85	Coeff	1.108	-0.381	1.157	-0.007
	t-stat	0.81	0.96	***4.30***	0.10
90	Coeff	1.727	-0.288	1.200	-0.042
	t-stat	1.32	0.77	***4.71***	0.66
95	Coeff	1.818	-0.311	1.081	-0.044
	t-stat	1.42	0.86	***4.44***	0.71
100	Coeff	2.103	-0.365	1.055	-0.057
	t-stat	1.75	1.04	***4.56***	0.98

Table 8: The table shows results for price sensitivity analysis to insurer's portfolios with different shares in occupancy risk. All analyses are carried out above a common threshold of USD 1 million. Estimated parameters of the four models used are reported in panel A. Pareto and Generalized Pareto with Covariates have been estimated on both Occupancy and Geographic dummies. The reported coefficients for these models are averages of the partial effects. Panel B reports estimated severity - truncated average loss between 1 and 10 million. Panel C reports the simulated 80% quantile. The 'MCS' columns report results from 1 million Monte Carlo simulations from a mixture of two parametric models (one for commercial and one for manufacturing risk) with different weights, representing the exposure of the portfolio to a given combination of risks. The models used are, Pareto, Pareto with covariates, Generalized Pareto (GPD) and Generalized Pareto with covariates (GPD with Cov.) .The 'Bootstrap' columns instead, show results for 1 million bootstrap samples of size 50 from different subsamples of the empirical data. The severity column reports the simulated truncated average loss between 1 and 10 million.

Panel A

	Alpha	
Model	CO	MA
Pareto	0.478	0.424
Pareto with Covariates	0.495	0.445
Generalized Pareto	1.295	1.287
Generalized Pareto with Covariates	1.200	1.540

Panel B

		Diversification by Occupancy (Average severity)				
%CO	%MA	Bootstrap	MCS Pareto	MCS Pareto with Cov.	MCS GPD	MCS GPD with Cov.
0	100	4,287,400	3,266,903	3,227,243	4,582,149	4,279,877
25	75	4,224,224	3,246,818	3,210,691	4,500,714	4,173,538
50	50	4,168,877	3,228,545	3,193,642	4,441,841	4,068,739
75	25	4,119,124	3,207,118	3,177,158	4,372,426	3,978,081
100	0	4,072,025	3,191,670	3,159,435	4,309,621	3,895,888

Panel C

		Diversification by Occupancy (80%-Quantile)				
%CO	%MA	Bootstrap	MCS Pareto	MCS Pareto with Cov.	MCS GPD	MCS GPD with Cov.
0	100	37,966,350	44,641,057	40,955,890	32,059,110	34,086,130
25	75	35,291,569	39,524,307	37,467,241	29,599,960	31,281,521
50	50	33,050,342	35,721,306	34,112,022	27,273,712	28,537,134
75	25	30,995,377	32,207,574	30,891,311	24,906,015	25,698,466
100	0	29,054,178	28,747,761	28,469,690	22,685,403	22,903,126

Table 9: The table shows results for price sensitivity analysis to insurer's portfolios with different shares in occupancy risk. All analyses are carried out above a common threshold of USD 1 million. Estimated parameters of the four models used are reported in panel A. Pareto and Generalized Pareto with Covariates have been estimated on both Occupancy and Geographic dummies. The reported coefficients for these models are averages of the partial effects. Panel B reports estimated severity - truncated average loss between 1 and 10 million. Panel C reports the simulated 80% quantile. The 'MCS' columns report results from 1 million Monte Carlo simulations from a mixture of two parametric models (one for developed and one for emerging risk) with different weights, representing the exposure of the portfolio to a given combination of risks. The models used are, Pareto, Pareto with covariates, Generalized Pareto (GPD) and Generalized Pareto with covariates (GPD with Cov.) .The 'Bootstrap' columns instead, show results for 1 million bootstrap samples of size 50 from different subsamples of the empirical data. The severity column reports the simulated truncated average loss between 1 and 10 million.

Panel A

Model	Alpha	
	Developed	Emerging
Pareto	0.371	0.537
Pareto with Covariates	0.385	0.555
Generalized Pareto	1.912	0.982
Generalized Pareto with Covariates	1.822	0.918

Panel B

		Diversification by Economic Development (Average severity)				
%Devel.	%Emerg.	Bootstrap	MCS Pareto	MCS Pareto with Cov.	MCS GPD	MCS GPD with Cov.
0	100	3,684,803	3,114,574	3,087,442	3,827,749	3,627,838
25	75	3,916,553	3,162,850	3,136,416	4,002,781	3,819,117
50	50	4,194,612	3,220,951	3,194,982	4,234,933	4,071,976
75	25	4,551,945	3,279,662	3,254,007	4,535,838	4,393,894
100	0	5,048,906	3,346,243	3,322,266	4,960,328	4,825,248

Panel C

		Diversification by Economic Development (80%-Quantile)				
%Devel.	%Emerg.	Bootstrap	MCS Pareto	MCS Pareto with Cov.	MCS GPD	MCS GPD with Cov.
0	100	22,589,717	19,874,070	18,316,413	17,691,883	16,060,119
25	75	28,713,565	27,571,373	24,697,726	24,600,626	22,363,006
50	50	35,132,779	38,076,341	34,168,122	31,659,466	28,584,786
75	25	41,533,709	54,158,731	47,408,106	37,998,539	34,299,695
100	0	47,921,745	76,779,846	66,397,462	43,747,248	39,233,637

Table 10: The table shows structural breaks tests for the Pareto model and Hill estimator. All models are estimated with a threshold of USD 140k. The three α's reported are the coefficient before, after and during the tested year. The test is performed using the Likelihood Ratio test statistic. Significant t-Statistics (above 1.96) are in bold and italics. Panel A refers to the whole sample, panel B to developed countries and panel C to developing countries.

Panel A: Whole Sample					
Year	Alpha before	Alpha after	Alpha during	Chi2-Test	P-value
2001	0.242	0.346	0.232	5.5	***0.019***
2002	0.241	0.355	0.239	9.5	***0.002***
2003	0.239	0.369	0.234	16.6	***0.000***
2004	0.237	0.382	0.228	22.8	***0.000***
2005	0.234	0.398	0.214	30.6	***0.000***
2006	0.233	0.418	0.232	39.9	***0.000***
2007	0.248	0.441	0.317	41.5	***0.000***
2008	0.282	0.423	0.520	20.5	***0.000***
2009	0.298	0.427	0.412	14.4	***0.000***
2010	0.313	0.423	0.436	7.8	***0.005***
2011	0.330	0.366	0.482	0.5	0.477
Panel B: Developed region					
Year	Alpha before	Alpha after	Alpha during	Chi2-Test	P-value
2001	0.234	0.231	0.232	0.0	0.950
2002	0.236	0.230	0.241	0.0	0.875
2003	0.232	0.231	0.224	0.0	0.971
2004	0.235	0.228	0.250	0.0	0.831
2005	0.233	0.229	0.217	0.0	0.915
2006	0.235	0.226	0.246	0.1	0.801
2007	0.231	0.232	0.210	0.0	0.977
2008	0.232	0.231	0.238	0.0	0.981
2009	0.234	0.220	0.266	0.1	0.756
2010	0.230	0.240	0.184	0.0	0.862
2011	0.232	0.223	0.274	0.0	0.883
Panel C: Developing region					
Year	Alpha before	Alpha after	Alpha during	Chi2-Test	P-value
2001	0.265	0.453	-	4.3	***0.039***
2002	0.252	0.466	0.235	9.3	***0.002***
2003	0.254	0.479	0.257	13.9	***0.000***
2004	0.240	0.494	0.177	20.6	***0.000***
2005	0.234	0.511	0.210	27.8	***0.000***
2006	0.230	0.533	0.208	37.2	***0.000***
2007	0.280	0.543	0.462	29.5	***0.000***
2008	0.365	0.514	0.693	9.3	***0.002***
2009	0.388	0.526	0.476	7.4	***0.006***
2010	0.418	0.503	0.579	2.3	0.129
2011	0.438	0.461	0.537	0.1	0.764

Table 11: The table shows structural breaks test for Pareto model, estimated via Hill, for several years. All the models have been estimated over a threshold of USD 10 million. The three α's reported are the coefficient before, after and during the tested year. The test is performed using the Likelihood Ratio test statistic. Significant t-Statistics (above 1.96) are in bold and italics. Panel A refers to the whole sample. Panel B refers only to the developed region. Panel C refers only to the developing region.

Panel A: Whole Sample					
Year	Alpha before	Alpha after	Alpha during	Chi2-Test	P-value
2002	0.918	0.849	0.784	1.6	0.213
2003	0.820	0.885	0.709	-1.0	1.000
2004	0.757	0.944	0.620	0.1	0.747
2005	0.742	0.997	0.703	2.1	0.143
2006	0.767	1.011	0.802	1.0	0.312
2007	0.776	1.105	0.872	6.1	***0.013***
2008	0.784	1.139	0.996	9.8	***0.002***
2009	0.828	1.030	1.065	1.1	0.301
2010	0.830	1.172	1.103	5.8	***0.016***
Panel B: Developed Country Class					
Year	Alpha before	Alpha after	Alpha during	Chi2-Test	P-value
2002	0.964	0.775	0.794	2.0	0.155
2003	0.837	0.809	0.734	0.0	0.834
2004	0.780	0.851	0.618	0.6	0.432
2005	0.774	0.869	0.755	0.8	0.358
2006	0.806	0.839	0.784	0.2	0.671
2007	0.781	0.926	0.747	1.2	0.279
2008	0.765	1.030	0.749	3.3	0.069
2009	0.789	0.960	0.778	1.5	0.220
2010	0.781	1.265	1.013	4.0	***0.044***
Panel C: Developing Country Class					
Year	Alpha before	Alpha after	Alpha during	Chi2-Test	P-value
2002	0.775	0.997	0.742	1.0	0.315
2003	0.768	1.029	0.658	2.1	0.145
2004	0.700	1.130	0.622	9.0	***0.003***
2005	0.671	1.277	0.636	23.2	***0.000***
2006	0.690	1.444	0.838	40.6	***0.000***
2007	0.763	1.483	1.220	39.7	***0.000***
2008	0.829	1.323	1.677	18.4	***0.000***
2009	0.913	1.188	1.927	5.2	***0.023***
2010	0.945	1.057	1.226	0.5	0.477