Provenance of the named severity curves

What each published curve in the shipped library is, and where it came from

Stephen J. Mildenhall

2026-10-22 19:57 commit b5ca906+

1 Introduction

`src/aggregate/agg/library.agg` ships eleven reusable severity objects named `<Citation-Key>.<Curve>`, harvested from the candidate assessment in `plan-examples.md` while the reproductions in this folder were being written. Each carries its citation key in `note{}`, and little else. This page is the rest of the record: which paper, which table or exhibit the parameters came out of, what the paper uses the curve for, and every transcription decision that stands between the printed parameters and the DecL.

A twelfth library severity carries the `role:published` tag, `CommAutoMixedExponentialSev`, a five component commercial auto mixed exponential at the 2008 cost level. It came out of my own *Actuarial Review* article rather than out of the assessment, so it has no provenance question to answer and is not covered here.

The curves are working definitions, not reproductions. Where a reproduction exists it is cited, and it is the authority; where none exists the entry says so. Table 1 is what actually ships.

```
%xmode minimal
import warnings

import numpy as np
import pandas as pd
from bookstyle import GT, build

rows = []
with warnings.catch_warnings():
    warnings.simplefilter('ignore')
    # `role:published` is the library's tag for a curve whose parameters were
    # printed somewhere. Twelve entries carry it; the dot in the name is what
    # separates the eleven `<CitationKey>.<Curve>` harvests this page documents
    # from `CommAutoMixedExponentialSev`, which came from elsewhere.
    for name in build.discover(kind='sev', tags='role:published').index:
        if '.' not in name:
            continue
        fz = build(f'agg T 1 claim sev.{name} fixed', update=False).sevs[0].fz
        # `moments=` by keyword: a tabulated curve's `fz` is an unfrozen
        # rv_histogram, which reads a positional 'mv' as a shape parameter.
        m, v = fz.stats(moments='mv')
        dist = getattr(fz, 'dist', None)
        rows.append({'curve': name,
                     'scipy': getattr(dist, 'name', None) or type(fz).__name__,
                     'shape': ', '.join(f'{a:g}' for a in getattr(fz, 'args', ())),
                     'scale': float(getattr(fz, 'kwds', {}).get('scale', np.nan)),
```

```
                    'mean': float(m),
                    'cv': float(np.sqrt(v) / m) if np.isfinite(v) else np.nan})

GT(pd.DataFrame(rows).set_index('curve'), float_format=',.4g')
```

```
Exception reporting mode: Minimal
```

Table 1: Every per-paper severity curve in the shipped library, its `scipy` family, parameters and moments.

curve	scipy	shape	scale	mean	cv
Berens1997.Lognorm	lognorm	0.246221	1.455e+05	1.5e+05	0.25
Berens1997.Pareto	pareto	1.5	5e+04	1.5e+05	—
Bodoff2017.Pareto	lomax	1.5	5e+04	1e+05	—
Jin2016.Gamma	gamma	2.99948	2	6	0.5774
Jin2016.Pareto	lomax	20	100	5.263	1.054
Kang2019.Pareto	lomax	3.5	3.5	1.4	1.528
Mack2003.AmericanPareto	lomax	1.5	1e+04	2e+04	—
Mack2003.Lognorm	lognorm	1.79998	1,979	1e+04	4.953
Mack2003.RiebesellPareto	lomax	0.737	1e+04	—	—
Mata2005.Lognorm	lognorm	2.29	1.105e+04	1.521e+05	13.73
Venter1983.Piecewise	rv_histogram		—	1.88e+04	2.734

A blank coefficient of variation means the second moment does not exist, and an infinite mean means the first does not either. A blank shape and scale means the curve is a tabulation rather than a parametric family, which is why `Venter1983.Piecewise` has neither, and why its mean and CV here are the *unlimited* ones and so do not match the paper: that curve is only itself inside the layer described below. Table 1 is built from the library, so it cannot drift from what actually ships.

2 Summary: paper to curve

Table 2 maps each assessed paper to the curves it contributed.

Table 2: Each assessed paper, the curves it contributed to the library, and why.

Paper	Curves	Why
Berens (1997)	`Berens1997.Lognorm`, `Berens1997.Pareto`	the severities of the paper's two simulation models
Mack and Fackler (2003)	`Mack2003.Lognorm`, `Mack2003.AmericanPareto`, `Mack2003.RiebesellPareto`	two columns of Table 1, plus the Pareto tail index Riebesell's rule implies
Mata and Verheyen (2005)	`Mata2005.Lognorm`	the one severity behind every table in the paper
Bodoff (2017)	`Bodoff2017.Pareto`	Exhibit 2
Jin et al. (2016)	`Jin2016.Gamma`, `Jin2016.Pareto`	the univariate test severities
Kang et al. (2019)	`Kang2019.Pareto`	the running heavy tailed illustration

Paper	Curves	Why
Venter (1983)	`Venter1983.Piecewise`	the claim size law behind Exhibit 3, a tabulated CDF rather than a named family
Bruno et al. (2006)	none	a 14 point discrete severity, likewise not a named family
Clark (2005)	none	three lognormal lines specified by premium and CV, nothing worth naming
Lau (1984)	none	moments only. The severity form is never stated and has to be reverse fitted
Braithwaite and Ware (1997)	none	never assessed. It surfaced only as the conceptual pair to Bodoff (2017)
Goffard and Laub (2020), Nadarajah et al. (2016), Vernic (1999), Bakar et al. (2022), Provost (2005), Hilbe (2014), Lindsay et al. (2000)	none	Tier 3 in the assessment: no reproducible parameters, no compounding, or no insurance content

Bruno et al. (2006) has no curve and is a full reproduction anyway. A severity does not have to be a named family to be worth carrying: `Venter1983.Piecewise` is a tabulated CDF, and Bruno's 14 point law would sit in the library just as well as a `dsev` if anything besides its own chapter needed it.

The rare event and numerical accuracy set (`Asmussen2016a`, `BenAmar2023`, `BenRached2024`, `Liu2026`, `Wilson2016a`, `Wilson2017`, `Homer2003`) contributed no named curves. Those papers are about estimator accuracy in the far tail, and their test severities are unit scale lognormals, Paretos and Lévy laws chosen to stress the estimator rather than published as models of anything.

3 The curves

3.1 Berens1997

Paper. Berens (1997), *Reinsurance contracts with a multi-year aggregate limit*, CAS Forum.

What the paper does with them. It is a simulation study of the surcharge a reinsurer should charge for letting an aggregate limit run over three years instead of one, priced across excess layers of width \$150,000 attaching at \$150,000, \$300,000 and \$450,000. Two frequency and severity pairs are run, deliberately chosen to disagree. The first is Poisson with 30 expected claims and a lognormal severity of mean \$150,000 and standard deviation \$37,500 (Exhibit I, upper). The second is negative binomial with 30 expected claims and $p = 0.10$, and a single parameter Pareto severity with $q = 1.5$ and mean \$150,000 (Exhibit III). Exhibit V plots claim size percentiles from both, and it is the paper's explanation of why the surcharge ranking flips as the attachment rises: under the Pareto far more claims fall below a \$150,000 retention, but at high attachments so many are large enough to exhaust the layer that the layer loss is actually *less* volatile.

`Berens1997.Lognorm` = `lognorm 150000 cv 0.25`. The mean and standard deviation are stated, and \$37,500 / \$150,000 = 0.25 is the CV. The paper then writes "the distribution of the severity was

thus $\Lambda(11.92, .25)$", which is the shortcut $\mu = \ln(\text{mean})$, $\sigma = \text{cv}$, and is not consistent with its own moments: $\Lambda(11.92, 0.25)$ has mean 155,011 and CV 0.2540. The library object follows the stated moments, giving $\sigma = 0.246221$ and $\mu = 11.8881$. Nothing in the paper turns on the difference, but a reproduction that quoted the printed (μ, σ) would run 3.3% high on the severity mean, and further out on excess layers, where a shift in the body of a tight lognormal is amplified.

`Berens1997.Pareto = 50000 * pareto 1.5`. "Single parameter Pareto" is the European or Type I form, supported on $[\theta, \infty)$, DecL `pareto`. Its mean is $q\theta/(q-1) = 3\theta$, so the stated mean of \$150,000 pins $\theta = 50{,}000$. The paper's second Pareto run has mean \$300,000, hence $\theta = 100{,}000$; that variant is not carried. With $q = 1.5$ the variance is infinite, and because the support starts at 50,000 rather than 0 `aggregate` logs `Pareto not shifted to x>0 range` and falls back to numeric moments. That message is expected here, not a defect.

Status. No reproduction. The single year layer pieces are native, but the rolling multi-year aggregate limit is path dependent and needs a wrapper loop over independent annual builds with cumulative capping.

3.2 Mack2003

Paper. Mack and Fackler (2003), *Exposure-rating in liability reinsurance*, Blätter der DGVFM. Reproduced in `Mack2003.qmd`.

What the paper does with them. It asks whether Riebesell's 1936 increased limits rule, that doubling the sum insured multiplies the net premium by $1 + z$, is consistent with the collective risk model. A power law premium forces a Pareto claim size law with $\alpha = 1 - \text{ld}(1 + z)$, which is below one and so has an infinite mean. Theorem 2 rescues the rule: for any continuous F with a finite mean there is a threshold u such that grafting a Pareto tail onto F above u makes Riebesell exact on $[u, \infty)$. Table 1 tabulates the function that locates that threshold,

$$\alpha_F(u) = 1 - \frac{u\, S(u)}{\mathsf{E}[X \wedge u]},$$

for four severities at seven multiples of the mean. Three of the four columns are here.

`Mack2003.Lognorm = lognorm 10000 cv 4.953`. Table 1's "Lognormal $\sigma = 1.8$" column, and also the claim size distribution of the paper's own worked example in §7. The CV is $\sqrt{e^{\sigma^2} - 1} = 4.953153$ at $\sigma = 1.8$; the library rounds to four figures, which recovers $\sigma = 1.79998$. Table 1 is indexed by $u/\mathsf{E}[X]$ and so is scale free. The 10,000 mean is a placeholder, and the paper remarks only that $\mathsf{E}[X]$ in liability is usually below €50,000.

`Mack2003.AmericanPareto = 10000 * lomax 1.5`. Table 1's "American Pareto $\alpha = 1.5$" column. The paper does not define the term, and the reading that it means Type II, that is Lomax with $S(x) = (\theta/(\theta + x))^\alpha$ and support from zero, is an inference. `Mack2003.qmd` confirms it: all 26 cells of Table 1 come back within 8e-4 under that reading.

`Mack2003.RiebesellPareto = 10000 * lomax 0.737`. Not a curve the paper tabulates. It is the tail index Riebesell implies at the customary $z = 20\%$: $\alpha = 1 - \ln(1.2)/\ln 2 = 0.736966$, which the paper prints as 0.737. **This object carries the exponent, not Mack and Fackler's construction.** Their tail is spliced onto the lognormal above $u \approx 5\,\mathsf{E}[X]$ and is conditional on exceeding u; this is a ground up Lomax that merely has the right index. It has an infinite mean, so use it only inside a finite layer. Unbounded it does not raise, it silently builds on a degenerate two bucket grid and reports a mean of 0.67. For the actual splice, and the refinement that the threshold is $u^* = 5.186\,\mathsf{E}[X]$ where the spliced curve gives $b(2v)/b(v) = 1.200000$ exactly, see `Mack2003.qmd`.

3.3 Mata2005

Paper. Mata and Verheyen (2005), *An Improved Method for Experience Rating Excess of Loss Treaties using Exposure Rating Techniques*, CAS Forum. Reproduced in `Mata2005.qmd`.

What the paper does with it. One severity drives the whole worked example: layer by layer experience and exposure rating of five excess layers over a limits profile running from \$250,000 to \$5,000,000, with an 8% ground up severity trend, a 60% expected gross loss ratio, and the trend then decomposed into frequency and severity components layer by layer. Tables 1 to 5 and 9 to 14 are all limited expected values or ratios of them off this curve.

`Mata2005.Lognorm = 11048 * lognorm 2.29`. The 2005 level lognormal, $\mu = 9.31$ and $\sigma = 2.29$, written in scipy's scale form with scale $= e^{9.31} = 11{,}047.95$ rounded to 11,048. Mean 152,064, CV 13.73. The 2000 level curve, $\mu' = \mu - \ln(1.08^5) = 8.925195$, is not carried; scale it by 1.08^{-5}. The paper prints μ' rounded to 8.93, and using the rounded value moves the limited expected values by about half a percent, more than the effects the paper is measuring, so derive it rather than transcribe it.

Germane finding. `Mata2005.qmd` reproduces Tables 2, 3, 4, 5, 11, 13 and 14 to the last printed digit when driven off the paper's own Table 1, but Table 1's limited expected values are themselves not the exact lognormal ones. They run 0.04% high at \$250,000 drifting to 0.24% low at \$5M, in both years, while `aggregate` matches the closed form to eight figures. The paper's conclusions are unaffected, but a mismatch against Table 1 is the paper's, not the library's.

3.4 Bodoff2017

Paper. Bodoff (2017), *An Actuarial Model of Excess of Policy Limits Losses*, Variance. Reproduced in `Bodoff2017.qmd`. Not the Bodoff already in the docs, which is Bodoff (2007) on percentile layer capital allocation.

What the paper does with it. Exposure rating for excess of policy limits losses, the ones where the insurer pays past the limit. The paper's move is to treat the *effectiveness* of the policy limit as Bernoulli: with probability p the limit caps the loss and with probability $1 - p$ it does not, so the policy limited expected value is a mixture $p\,\mathsf{E}[X \wedge k] + (1-p)\,\mathsf{E}[X]$ of two ordinary limited expected values off the same curve. Exhibit 2 is the severity, Exhibit 1 a nine point limits profile, Exhibit 3 sets $p = 99\%$, and Exhibit 4 is the six layer exposure rating that results.

`Bodoff2017.Pareto = 50000 * lomax 1.5`. Exhibit 2 says only "Pareto, $\theta = 50{,}000$, $\alpha = 1.50$". That it means the two parameter Pareto, $S(x) = (\theta/(\theta + x))^{\alpha}$, is an inference: it is the form in the Clark reinsurance study note the paper cites, and `Bodoff2017.qmd` confirms it by fit. Mean 100,000, infinite variance, so it needs a limit. The XPL model conditions this severity *before* the policy limit and then mixes over the limits profile.

3.5 Jin2016

Paper. Jin et al. (2016), *Moment-based density approximations for aggregate losses*, Scandinavian Actuarial Journal. Assessed in `Jin2016.qmd`, which is **not** a reproduction.

What the paper does with them. They are test severities for a moment based approximation, not models of anything. The method approximates the density of $S = X \mid X > 0$ by a gamma base density times a polynomial adjustment, and needs only moments, so it works from sample moments when frequency and severity are not separately known. The paper's own ground truth is the FFT.

`Jin2016.Gamma = gamma 6 cv 0.5774`. Example 1's $U \sim \gamma(3, 2)$, in the shape and scale reading, so mean 6 and CV $1/\sqrt{3}$. The convention is not stated in the paper but is pinned by arithmetic: Example 1 reports its fitted base gamma as (2.684349, 7.056877), and only the scale reading reproduces those, which `Jin2016.qmd` does to six figures. Transcription wart: 0.5774 is a four figure rounding of $1/\sqrt{3} = 0.5773502691896258$ and yields shape 2.99948 rather than 3. Write `sev 2 * gamma 3` when exactness matters.

`Jin2016.Pareto = 100 * lomax 20`. Examples 5 and 6 use $U \sim \text{Pareto}(\gamma, \delta) = \text{Pareto}(20, 100)$ with density $\gamma\delta^\gamma/(u+\delta)^{\gamma+1}$ on $u > 0$, which is Lomax with shape γ and scale δ. Mean $100/19 = 5.2632$. The paper's figures also show a Pareto(5, 100) variant, which is *heavier* rather than lighter; it is not carried.

Why there is no reproduction. The univariate half of the paper publishes figures and Kolmogorov-Smirnov statistics and nothing else, and those statistics are computed between a simulated empirical distribution and an approximation fitted to that sample's own moments, with no seed given. They are properties of a draw, not of the model. Every genuine reference value in the paper, 96 of them, is in the FFT columns of the bivariate tables, whose dependence is Hesselager common shock *counts* that `aggregate`'s copula coupled `bivariate` cannot express. So these two curves are a parameter transcription that happens to be checkable, not a reproduction target.

3.6 Kang2019

Paper. Kang et al. (2019), *Moment-Based Density Approximation Techniques as Applied to Heavy-tailed Distributions*, International Journal of Statistics and Probability.

What the paper does with it. Pareto II(3.5, 3.5) is the running illustration for approximating a heavy tailed density by exponential tilting: tilt the density so it becomes exponentially bounded, approximate the tilted density by a base density times a polynomial adjustment, then untilt. Figure 1 plots the tilted density at $\theta = 1.5$, Table 1 gives integrated squared errors over a grid of tilting parameters and polynomial degrees, and Tables 2 and 3 give VaR and TVaR for the exact distribution beside three approximants. The exact column is a clean external check on a heavy tail: $\text{VaR}_{0.99} = 9.54658$, $\text{TVaR}_{0.99} = 14.7652$, $\text{TVaR}_{0.999} = 31.7646$.

`Kang2019.Pareto = 3.5 * lomax 3.5`. The paper defines Pareto II(α, β) explicitly, density $\alpha\beta^\alpha/(x+\beta)^{\alpha+1}$ on $x \geq 0$, and names it Lomax, so the mapping is unambiguous: shape $\alpha = 3.5$, scale $\beta = 3.5$. Mean 1.4, CV 1.5275, skewness 11.78. Moments exist up to order 3 only, so there is no kurtosis. Watch the parameters: §6.3.1 of the paper switches to a simulated sample from Pareto II(3.5, **35**), rescaled, and the VaR and TVaR figures there are not comparable. The library object is the analytic Pareto II(3.5, 3.5) of §6.1.

Status. No reproduction. The paper's "aggregate" is severity only, with no compounding, so it exercises the quantile machinery and nothing else. It is a good heavy tail VaR and TVaR accuracy benchmark, which is why the curve was kept.

3.7 Venter1983

Paper. Venter (1983), *Transformed beta and gamma distributions and aggregate losses*, PCAS. Reproduced in `Venter1983.qmd` and in the highlights chapter.

What the paper does with it. It is the claim size law of Exhibit 3, the paper's test of the transformed gamma as a three moment approximation to the aggregate. Exhibit 3 Part 2 states it as a piecewise linear CDF from \$0 to a \$250,000 policy limit, tabulated at 23 points, with the residual 2.41% of

probability sitting as a point mass at the limit. Frequency is Poisson with $\lambda = 13.7376$, and the compound distribution has mean \$250,000, CV 0.7667 and skewness 1.0744.

`Venter1983.Piecewise` = `chistogram xps [...] [...]`. The 23 knots are Venter's own, and the 23 probabilities are the successive CDF increments with the limit mass appended. **The top bin is not part of the paper.** The atom at the limit is easier to obtain than to declare, so the limit mass is carried in one wide bin from \$250,000 to \$300,000 and a `250000 xs 0` layer then folds it back onto the limit. That makes the curve correct only inside that layer; used unlimited it spreads 2.41% of probability uniformly over a \$50,000 interval that Venter never wrote. Like `Mack2003.RiebesellPareto`, this one carries a usage constraint rather than a free standing distribution.

The transcription has an acceptance test built in, which is why the curve can be trusted: read the knots wrongly and the severity mean, CV and skewness (18,198.195, 2.6600, 3.6746) stop matching Exhibit 3 Part 2. They match, as the moments table in the Venter1983 chapter records.

What is not carried. Exhibit 3 Part 2 also prints Venter's own \$500 discrete PMF, whose first five rows could not be decoded. Nothing depends on it, since `aggregate` discretizes the piecewise linear CDF itself.

4 Reference curves with no paper attached

Alongside the eleven published curves the library carries 18 single claim aggregates named `Curve*`, listed by `build.discover(kind='agg', tags='topic:severity role:reference')`. These are vocabulary, not reproductions. Each is the simplest possible carrier, one fixed claim so the aggregate *is* the severity, and the shapes are illustrative with a common scale near 10,000 so they can be read against each other. No paper prescribes any of these parameters.

What they *do* record is a selection: scipy ships more than a hundred continuous distributions and actuaries use a small, well trodden subset. The picks and the literature behind them are:

- **Light and moderate tails**, `CurveExponential`, `CurveGamma`, `CurveWeibull`, `CurveInverseGaussian`. The standard working severities. Appendix A of Klugman et al. (2012) is the canonical actuarial table and the parameterizations follow its families; Parodi (2015) and Parodi et al. (2024) use lognormal, gamma and Weibull as the everyday attritional curves, and Parodi and Watson (2019) derives them from a property graph model of fire and explosion losses.
- **Lognormal and the GB2 family**, `CurveLognormal`, `CurveGenGamma`, `CurveBetaPrime`, `CurveBurr12`, `CurveBurrIII`, `CurveLogLogistic`. The flexible three and four parameter curves that nest most of the others. Venter (1983) introduced the transformed beta and gamma to actuaries, McDonald and Richards (1987) studied the hazard rates of the generalized beta family, and Cummins et al. (1990) is the insurance loss application. The GB2 umbrella covers the log-t, generalized gamma, B2, Burr III and XII, Lomax, Fisk, gamma, Weibull, and the lognormal as a limit.
- **Heavy and power law tails**, `CurvePareto`, `CurveLomax`, `CurveGPD`, `CurveInverseGamma`. Large loss and reinsurance work, plus the peaks over threshold tail of extreme value theory. The two power laws both use shape 2.5, chosen so the variance stays finite: the auto sizer refuses an infinite variance without an explicit grid.
- **Extreme value laws**, `CurveGumbel`, `CurveFrechet`, `CurveGEV`. Block maxima, for a largest event rather than a ground up attritional severity.
- **Bounded severity**, `CurveBeta`. A loss as a fraction of a known sum insured, and the gateway to exposure curve models.

`CurveBeta` also carries a deliberate absence. The Swiss Re standard for property per risk exposure rating is the two parameter MBBEFD class of Bernegger (1997) on the destruction ratio $[0, 1]$. It is an exposure *curve* family rather than a scipy continuous random variable, so there is no `sev mbbefd` form and the bounded beta is the closest stand in currently available.

5 Using the curves

`library.agg` is the default database, so the curves resolve from a bare `build` with no explicit load, as Table 3 shows:

```
a = build('agg XPL 1 claim 1000000 xs 0 sev.Bodoff2017.Pareto fixed')
b = build('agg J 3 claims sev.Jin2016.Gamma poisson')

GT(pd.DataFrame([{'program': 'Bodoff2017.Pareto limited to 1M', 'mean': a.est_m},
                 {'program': 'Poisson(3) x Jin2016.Gamma', 'mean': b.est_m}]
                ).set_index('program'), float_format=',.2f')
```

Table 3: Two programs resolving a named curve straight out of the shipped library.

program	mean
Bodoff2017.Pareto limited to 1M	78,178.21
Poisson(3) x Jin2016.Gamma	18.00

Four of the eleven have no finite variance and need a limit or an explicit grid. `Bodoff2017.Pareto` and `Mack2003.AmericanPareto` have $\alpha = 1.5$ and raise `InfiniteVarianceError` unbounded, which is the behavior you want. The other two fail less cleanly, and both are worth knowing about:

- `Berens1997.Pareto`, also $\alpha = 1.5$, is a Type I Pareto supported from 50,000 up, so it takes the numeric moment path and never reaches the infinite variance check. It raises `ValueError: Inadmissible value passed to round_bucket, inf` instead.
- `Mack2003.RiebesellPareto` at $\alpha = 0.737$ has no mean either, and it fails silently. The auto sizer cannot size a grid from an infinite mean, falls back to a two bucket grid, and returns 0.67 with no error at all. Always give this one a layer.

Bakar, S. A. A., Nadarajah, S., & Ngataman, N. (2022). A family of density-hazard distributions for insurance losses. *Communications in Statistics - Simulation and Computation*. https://doi.org/10.1080/03610918.2020.1784430

Berens, R. M. (1997). Reinsurance contracts with a multi-year aggregate limit. *Casualty Actuarial Society Forum*.

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

Bodoff, N. M. (2017). An Actuarial Model of Excess of Policy Limits Losses. *Variance*, *11*. https://doi.org/10.66573/001c.116059

Braithwaite, P., & Ware, B. C. (1997). Pricing extra-contractual obligations and excess of policy limits exposures in Clash Reinsurance Treaties. *Casualty Actuarial Society Forum*. https://doi.org/10.1016/s0167-6687(97)89196-2

Bruno, M. G., Camerini, E., Manna, A., & Tomassetti, A. (2006). A new method for evaluating the distribution of aggregate claims. *Applied Mathematics and Computation*. https://doi.org/10.1016/j.amc.2005.09.038

Clark, D. R. (2005). Reinsurance Applications for the RMK Framework, 353–366.
Cummins, J. D., Dionne, G., McDonald, J. B., & Pritchett, B. M. (1990). Applications of the GB2 family of distributions in modeling insurance loss processes. *Insurance: Mathematics and Economics*, *9*(4), 257–272. https://doi.org/10.1016/0167-6687(90)90003-V
Goffard, P.-O., & Laub, P. J. (2020). Orthogonal polynomial expansions to evaluate stop-loss premiums. *Journal of Computational and Applied Mathematics*, *370*. https://doi.org/10.1016/j.cam.2019.112648
Hilbe, J. M. (2014). *Modeling count data*. Cambridge University Press.
Jin, T., Provost, S. B., & Ren, J. (2016). Moment-based density approximations for aggregate losses. *Scandinavian Actuarial Journal*, *2016*(3), 216–245. https://doi.org/10.1080/03461238.2014.921640
Kang, J. S. J., Provost, S. B., & Ren, J. (2019). Moment-Based Density Approximation Techniques as Applied to Heavy-tailed Distributions. *International Journal of Statistics and Probability*. https://doi.org/10.5539/ijsp.v8n3p1
Klugman, S. A., Panjer, H. H., & Willmot, G. E. (2012). *Loss Models: from data to decisions* (Fourth.). John Wiley & Sons, Inc.
Lau, H.-S. (1984). An Effective Approach for Estimating the Aggregate Loss of an Insurance Portfolio. *The Journal of Risk and Insurance*. https://doi.org/10.2307/252798
Lindsay, B. G., Pilla, R. S., & Basak, P. (2000). Moment-based approximations of distributions using mixtures: Theory and applications. *Annals of the Institute of Statistical Mathematics*, *52*(2), 215–230.
Mack, T., & Fackler, M. (2003). Exposure-rating in liability reinsurance. *Blätter der DGVFM*, *26*(2), 229–238.
Mata, A. J., & Verheyen, M. A. (2005). An Improved Method for Experience Rating Excess of Loss Treaties using Exposure Rating Techniques. *CAS Forum*.
McDonald, J. B., & Richards, D. O. (1987). Hazard Rates and Generalized Beta Distributions. *IEEE Transactions on Reliability*, *R36*(4), 463—–466.
Nadarajah, S., Chu, J., & Jiang, X. (2016). On moment based density approximations for aggregate losses. *Journal of Computational and Applied Mathematics*. https://doi.org/10.1016/j.cam.2015.11.048
Parodi, P. (2015). *Pricing in General Insurance*. CRC Press. file:///C:/Users/youhe/Downloads/kdoc/_o/_00042/_01.pdf
Parodi, P., Thrumble, D., Watson, P., Ji, Z., Wang, A., Bhatia, I., et al. (2024). Loss modelling from first principles. *British Actuarial Journal*, *6*, 1–39. https://doi.org/10.1017/s1357321724000163
Parodi, P., & Watson, P. (2019). Property Graphs - a Statistical Model for Fire and Explosion Losses Based on Graph Theory. *ASTIN Bulletin*, *49*(2), 263–297. https://doi.org/10.1017/asb.2019.4
Provost, S. B. (2005). Moment-based density approximants. *Mathematica Journal*, *9*(4), 727–756.
Venter, G. G. (1983). Transformed beta and gamma distributions and aggregate losses. *Proceedings of the Casualty Actuarial Society*, *70*(133), 289–308.
Vernic, R. (1999). Recursive Evaluation of Some Bivariate Compound Distributions. *ASTIN Bulletin*, *29*(2), 315–325. https://doi.org/10.2143/ast.29.2.504618