Gerber and compound Poisson stop loss premiums

Reproducing Gerber (1982)

Stephen J. Mildenhall

2026-10-22 19:55 commit b5ca906+

1 Context and problem

Gerber (1982) is about evaluating a stop loss premium

$$\pi(t) = \mathsf{E}[(S-t)_+] = \mathsf{E}[S] - \mathsf{E}[S \wedge t]$$

when the aggregate S is compound Poisson with a *continuous* severity. In 1982 the obstacle was that the recursion available at the time, Panjer's, needs an arithmetic severity, so the continuous severity has to be discretized onto a span h first, and the discretization changes the answer.

The paper's contribution is to bracket the error rather than hope about it. It sets out three ways to put a continuous severity onto a lattice of span h:

1. move all the mass in a cell to its **lower** endpoint, which makes the aggregate smaller in the relevant order and produces a **lower bound** on the stop loss premium,
2. move it to the **upper** endpoint, giving an **upper bound**,
3. split the mass between the two endpoints so that the cell's **mean is preserved**, the method of **matching moments**.

Tables 1 to 6 then show all three at $h = 1$ and $h = 0.1$ against an exact column, for expected claim counts 1, 10 and 100. The paper's conclusion is that the bounds do bracket, that the upper bound is the better of the two, and that matching moments beats both by a wide margin: at $\lambda = 100$ it does as well at $h = 1$ as the bounds do at $h = 0.1$.

Forty years later the interesting content is not the ranking, which is settled, but the exact column. It is 28 published four figure numbers for a compound Poisson with a continuous severity, computed by somebody with no FFT, and that makes it a good external check.

2 Paper solution and gold standard

Frequency is Poisson with mean λ and severity is uniform on $(1, 3)$, so $\mathsf{E}[X] = 2$ and $\mathsf{E}[S] = 2\lambda$. Tables 1, 2 and 3, Figure 1, tabulate $\pi(t)$ for $\lambda = 1, 10, 100$ at the retentions Table 1 lists.

The exact column is not a closed form. It is the paper's own method (c) at $h = 0.01$, and the paper says so: "column 2 gives the 'exact' stop-loss premium as calculated by method (c) with $h = 0.01$. (A comparison with the results of methods (a) and (b) with $h = 0.01$ has convinced the author that the displayed digits are accurate.)" So the target is four significant figures, and the reproduction is a comparison of two independent approximations rather than a check against truth.

Table 1
Stop-loss premiums for $\lambda = 1$

t	Exact value	Lower bounds		Upper bounds		Matching moments	
		$h=1$	$h=0.1$	$h=0.1$	$h=1$	$h=1$	$h=0.1$
0	2	100%	100%	100%	100%	100%	100%
2	8.277×10^{-1}	88.9	100.0	100.0	100.0	96.3	100.00
4	2.689×10^{-1}	77.1	99.8	100.0	103.1	96.2	99.96
6	7.184×10^{-2}	65.0	99.9	100.1	105.8	96.1	99.98
8	1.627×10^{-3}	53.5	99.6	100.1	112.2	96.2	99.95
10	3.253×10^{-4}	42.4	99.4	100.2	120.2	95.5	99.94
12	5.815×10^{-5}	32.6	99.3	100.3	128.7	94.8	99.94
14	9.346×10^{-5}	24.6	99.0	100.4	138.9	94.7	99.93
16	1.366×10^{-6}	18.3	98.9	100.5	152.2	94.9	99.92
18	1.840×10^{-6}	13.3	98.5	100.6	167.6	95.2	99.92
20	2.302×10^{-7}	9.5	98.3	100.7	185.2	95.5	99.91

Table 2
Stop-loss premiums for $\lambda = 10$

t	Exact value	Lower bounds		Upper bounds		Matching moments	
		$h=1$	$h=0.1$	$h=0.1$	$h=1$	$h=1$	$h=0.1$
15	5.757	99.0%	99.99%	100.01%	100.5%	99.9%	100.00%
20	2.626	95.3	99.96	100.02	101.7	99.8	100.00
25	9.321×10^{-1}	91.6	99.91	100.04	104.1	99.6	100.00
30	2.563×10^{-1}	80.7	99.82	100.08	108.0	99.4	99.99
35	5.507×10^{-2}	74.7	99.71	100.13	113.5	99.2	99.99
40	9.383×10^{-3}	59.2	99.57	100.20	120.8	99.0	99.99
45	1.289×10^{-3}	53.2	99.40	100.28	130.3	98.8	99.99
50	1.449×10^{-4}	37.9	99.20	100.38	142.1	98.7	99.98
55	1.355×10^{-5}	33.3	98.98	100.48	156.8	98.7	99.98
60	1.067×10^{-6}	21.5	98.74	100.60	174.7	98.7	99.98

Table 3
Stop-loss premiums for $\lambda = 100$

t	Exact value	Lower bounds		Upper bounds		Matching moments	
		$h=1$	$h=0.1$	$h=0.1$	$h=1$	$h=1$	$h=0.1$
180	2.177×10^{1}	99.1%	99.99%	100.00%	100.4%	100.0%	100.000%
200	8.304	96.0	99.96	100.02	101.9	100.0	100.000
220	1.959	88.9	99.89	100.05	105.4	100.0	100.000
240	2.647×10^{-1}	77.6	99.77	100.11	111.6	99.9	99.999
260	1.992×10^{-2}	63.5	99.59	100.20	121.6	99.9	99.998
280	8.364×10^{-4}	48.5	99.35	100.32	136.0	99.8	99.998
300	1.994×10^{-5}	34.7	99.06	100.47	156.3	99.8	99.997

Figure 1: Gerber (1982), Tables 1, 2, and 3.

```
%xmode minimal
import numpy as np
import pandas as pd
from bookstyle import GT, build, show

# Transcribed from the 'Exact value' column of Tables 1, 2 and 3.
PAPER = {
    1: {0: 2.0, 2: 8.277e-1, 4: 2.689e-1, 6: 7.184e-2, 8: 1.627e-2,
        10: 3.253e-3, 12: 5.815e-4, 14: 9.346e-5, 16: 1.366e-5,
        18: 1.840e-6, 20: 2.302e-7},
    10: {15: 5.757, 20: 2.626, 25: 9.321e-1, 30: 2.563e-1, 35: 5.507e-2,
         40: 9.383e-3, 45: 1.289e-3, 50: 1.449e-4, 55: 1.355e-5, 60: 1.067e-6},
    100: {180: 2.177e1, 200: 8.304, 220: 1.959, 240: 2.647e-1, 260: 1.992e-2,
          280: 8.364e-4, 300: 1.994e-5},
}
GT(pd.DataFrame([{'table': t, 'lambda': lam, 'E[S]': 2 * lam,
                  'retentions': f'{min(v)} to {max(v)}', 'values': len(v)}
                 for t, (lam, v) in enumerate(PAPER.items(), start=1)]
                ).set_index('table'), float_format=0)
```

```
Exception reporting mode: Minimal
```

Table 1: The three published tables, their claim counts and the retentions each tabulates.

table	lambda	E[S]	retentions	values
1	1	2	0 to 20	11
2	10	20	15 to 60	10
3	100	200	180 to 300	7

The severity supplies an acceptance test before any aggregate is built. A uniform on $(1, 3)$ has mean 2 and variance $4/12 = 1/3$, so $\mathsf{E}[S] = 2\lambda$ and $\mathsf{Var}[S] = \lambda\,\mathsf{E}[X^2] = \lambda(4 + 1/3)$ by the compound Poisson variance formula. Table 2 records the two severity moments the check uses.

```
GT(pd.DataFrame([
    {'quantity': 'severity mean', 'exact': 2.0},
    {'quantity': 'severity second moment', 'exact': 4 + 1 / 3},
]).set_index('quantity'), float_format=8)
```

Table 2: The first two moments of the uniform severity on $(1, 3)$.

quantity	exact
severity mean	2.00000000
severity second moment	4.33333333

3 The `aggregate` reproduction

The severity is `sev 2 * uniform + 1`: a standard uniform scaled to width 2 and shifted to start at 1. Nothing else in the declaration varies across the three tables except λ and the grid, and Table 3 confirms all three builds before any premium is read.

```
def gerber(lam, bs, log2=18):
    return build(f'agg Gerber {lam} claims sev 2 * uniform + 1 poisson '
                 f'hints{{log2={log2}; bs={bs}}}')

a1 = gerber(1, 1 / 1024)
show(a1)
```

```
agg Gerber
  1 claim
  sev 2 * uniform + 1
  poisson
  hints{log2=18; bs=0.0009765625}
```

```
rows = []
for lam in (1, 10, 100):
    a = gerber(lam, 1 / 1024 if lam == 1 else 1 / 128 if lam == 10 else 1 / 64)
    rows.append({'lambda': lam, 'E[S] exact': 2 * lam, 'E[S] agg': a.est_m,
                 'sd exact': np.sqrt(lam * (4 + 1 / 3)), 'sd agg': a.est_sd,
                 'total mass': a.agg_density.sum()})
GT(pd.DataFrame(rows).set_index('lambda'),
   formatters={'E[S] exact': 6, 'E[S] agg': 6, 'sd exact': 6, 'sd agg': 6,
               'total mass': 12})
```

Table 3: The aggregate mean and standard deviation at each claim count, against the compound Poisson formulas.

lambda	E[S] exact	E[S] agg	sd exact	sd agg	total mass
1	2.000000	2.000000	2.081666	2.081666	1.000000000000
10	20.000000	20.000000	6.582806	6.582814	1.000000000000
100	200.000000	200.000000	20.816660	20.816758	1.000000000000

Now the tables, Table 4. The stop loss premium comes straight off `density_df['lev']`, since $\pi(t) = \mathsf{E}[S] - \mathsf{E}[S \wedge t]$.

```
def stop_loss(lam, bs, log2=18):
    a = gerber(lam, bs, log2)
    return {t: a.est_m - float(a.density_df.loc[t, 'lev']) for t in PAPER[lam]}

FINE = {1: 1 / 1024, 10: 1 / 128, 100: 1 / 64}
rows = []
for lam, bs in FINE.items():
    got = stop_loss(lam, bs)
    for t, exact in PAPER[lam].items():
        rows.append({'lambda': lam, 't': t, 'paper exact': exact,
                     'aggregate': got[t], 'rel err': abs(got[t] / exact - 1)})
res = pd.DataFrame(rows)
GT(res.set_index(['lambda', 't']),
   formatters={'paper exact': '{x:.4e}', 'aggregate': '{x:.4e}',
               'rel err': '{x:.2e}'})
```

Table 4: All 28 published stop loss premiums against the rebuilt values, with the relative error on each.

lambda	*t*	paper exact	aggregate	rel err
1	0	2.0000e+00	2.0000e+00	9.84e-14
	2	8.2770e-01	8.2773e-01	3.47e-05
	4	2.6890e-01	2.6891e-01	3.43e-05
	6	7.1840e-02	7.1838e-02	2.83e-05
	8	1.6270e-02	1.6267e-02	2.07e-04
	10	3.2530e-03	3.2535e-03	1.58e-04
	12	5.8150e-04	5.8146e-04	6.50e-05
	14	9.3460e-05	9.3461e-05	8.90e-06
	16	1.3660e-05	1.3664e-05	2.58e-04
	18	1.8400e-06	1.8400e-06	1.75e-05
	20	2.3020e-07	2.3017e-07	1.51e-04
10	15	5.7570e+00	5.7565e+00	8.03e-05
	20	2.6260e+00	2.6255e+00	1.78e-04
	25	9.3210e-01	9.3214e-01	4.51e-05
	30	2.5630e-01	2.5632e-01	8.21e-05
	35	5.5070e-02	5.5073e-02	6.03e-05
	40	9.3830e-03	9.3834e-03	3.90e-05
	45	1.2890e-03	1.2887e-03	2.55e-04
	50	1.4490e-04	1.4495e-04	3.53e-04
	55	1.3550e-05	1.3552e-05	1.45e-04
	60	1.0670e-06	1.0673e-06	2.35e-04
100	180	2.1770e+01	2.1773e+01	1.45e-04
	200	8.3040e+00	8.3045e+00	5.76e-05
	220	1.9590e+00	1.9591e+00	6.48e-05
	240	2.6470e-01	2.6470e-01	4.41e-06
	260	1.9920e-02	1.9919e-02	4.91e-05
	280	8.3640e-04	8.3651e-04	1.28e-04
	300	1.9940e-05	1.9938e-05	1.15e-04

```
GT(pd.DataFrame([
    {'lambda': lam, 'values compared': int((res['lambda'] == lam).sum()),
     'worst relative error': float(res[res['lambda'] == lam]['rel err'].max()),
     'inside 4 significant figures':
         int((res[res['lambda'] == lam]['rel err'] < 5e-4).sum())}
    for lam in (1, 10, 100)]
    + [{'lambda': 'all', 'values compared': len(res),
        'worst relative error': float(res['rel err'].max()),
        'inside 4 significant figures': int((res['rel err'] < 5e-4).sum())}]
).set_index('lambda'), formatters={'worst relative error': '{x:.2e}'})
```

Table 5: How many values reproduce inside the four significant figures the paper prints, table by table.

lambda	values compared	worst relative error	inside 4 significant figures
1	11	2.58e-04	11
10	10	3.53e-04	10
100	7	1.45e-04	7
all	28	3.53e-04	28

By Table 5 all 28 published values reproduce inside the four significant figures the paper prints, over a range of $\pi(t)$ spanning eight orders of magnitude, from 21.77 down to 2×10^{-7}. Two independent approximations to the same quantity, forty years apart, agree everywhere they can be compared.

4 Comments

4.1 The grid, and a correction to an earlier treatment

The bucket sizes above are not the obvious ones, and the $\lambda = 100$ table is the reason. An earlier treatment of these tables in the technical guides used `bs = 1/16` for $\lambda = 100$, which is ample for the body and not for the tail, as Table 6 shows.

```
sweep = []
for bs in (1 / 16, 1 / 64, 1 / 256):
    got = stop_loss(100, bs)
    for t, exact in PAPER[100].items():
        sweep.append({'t': t, 'bs': f'1/{int(1/bs)}',
                      'rel err': abs(got[t] / exact - 1)})
piv = pd.DataFrame(sweep).pivot(index='t', columns='bs', values='rel err')
piv = piv[[f'1/{d}' for d in (16, 64, 256)]]
GT(piv, formatters={c: '{x:.2e}' for c in piv.columns})
```

Table 6: Relative error on Table 3 at three bucket sizes, showing where the coarse grid fails.

t	1/16	1/64	1/256
180	1.61e-04	1.45e-04	1.44e-04
200	1.27e-04	5.76e-05	5.32e-05
220	2.61e-04	6.48e-05	5.25e-05
240	4.23e-04	4.41e-06	2.17e-05
260	6.95e-04	4.91e-05	9.56e-05
280	1.31e-03	1.28e-04	5.44e-05
300	1.60e-03	1.15e-04	2.23e-04

In Table 6, at `bs = 1/16` the last three rows of Table 3 miss the printed precision, by up to 1.6×10^{-3} at $t = 300$. Refining to `bs = 1/64` brings every row inside 2×10^{-4} and refining again changes little, so the residual was discretization on this side and not an inaccuracy in Gerber's column. That is worth stating the right way round: the first version of this comparison looked like a small failure of the paper's far tail digits, and it was not.

The general point is the one Mack and Fackler (2003) makes from a different direction. A grid sized for the mean is not automatically a grid sized for a tail eight orders of magnitude down, and the quantity that reveals it is the one being read, not the mean or the total mass, both of which are already exact to twelve figures at `bs = 1/16`.

4.2 What became of the three methods

Gerber's three discretizations are still the live choices, and `aggregate` exposes the same distinction as `discretization_calc`, which computes the discrete probabilities from the survival function or the distribution function and so is accurate on the right or the left tail. The library's default is the moment preserving idea rather than either bound, which is Gerber's method (c) and his own recommendation.

What has changed is that the bound is no longer the thing you reach for. Gerber needed bounds because his exact answer cost as much as the bracket did, so bracketing was the cheap way to know he was close. Here the fine grid *is* the cheap option: the whole of Table 3 at `bs = 1/256` costs a fraction of a second, so the honest error control is refinement and comparison, and this page uses it above rather than a bound.

4.3 Why a uniform severity is a good test and a bad model

Uniform on $(1, 3)$ is not a claim size distribution anybody would fit. It is a deliberately awkward test case, and the awkward part is the two hard edges. A severity with compact support and a discontinuous density at both ends is the worst case for any discretization scheme, because the error does not smooth out across a cell boundary the way it does for a density that decays. That the reproduction holds to four figures on a severity chosen to be difficult, rather than on a lognormal, says more than the same result on a friendly curve would.

It also has an exact aggregate at small λ, which is how the $t = 0$ row can be checked to twelve figures: $\pi(0) = \mathsf{E}[S] = 2\lambda$ with no discretization involved at all.

Gerber, H. U. (1982). On the numerical evaluation of the distribution of aggregate claims and its stop-loss premiums. *Insurance Mathematics and Economics*, *1*(1), 13–18. https://doi.org/10.1016/0167-6687(82)90016-6

Mack, T., & Fackler, M. (2003). Exposure-rating in liability reinsurance. *Blätter der DGVFM*, *26*(2), 229–238.