

A Premium Is Not the Expected Loss — And What Sets the Loading Is Not the Accuracy of the Probability but the Number of Policies

— With one policy, keeping the ruin probability at 1 per cent takes a loading of **23.1469** times the expected loss

— With a million policies it falls to **0.0231** — The expected loss has not moved by a single unit — [Paper 383]

Yuuki YAMAGISHI (山岸 勇輝)

Yamagishi Complexity Torus Lab, Japan

September 2026

— *From thinking to syncing.* —

Abstract

It is tempting to think a premium is the probability of loss times its size. This paper asks whether that suffices—— the answer is that it does not, and what sets the shortfall is the number of policies. No new mathematical theorem and no new law is claimed. First, a probability of 0.01 and a loss of 1000000 give an expected loss of **10000** (Section 2). Second, one policy's standard deviation is **99498.74**, which is **9.9499** times the expected loss (Section 2). Third, this is the core of the paper. With one policy the required loading is **23.1469** times the expected loss (Section 3). Fourth, with a million policies the loading falls to **0.0231** times (Section 3). Fifth, multiplying the number of policies by 100 multiplies the loading by exactly **0.1** (Section 3). Sixth, the separator is how many the risk is divided among, not the accuracy of the probability (Section 5).

Scope of this paper (scope note): No new mathematical theorem and no new law is claimed—the law of large numbers, the safety loading, and the ruin probability are all standard. We do not build actuarial science—rigorous ruin bounds from risk theory and reinsurance design are not entered. They are named and no more. Institutions are not discussed—regulatory capital and accounting are not the subject. Independence is assumed—losses are taken as mutually independent. For correlated risks such as earthquakes or epidemics this computation does not hold. A normal approximation is used—the distribution of the total is approximated as normal. No accuracy is claimed for small n . Costs are not treated—administration, profit, adverse selection, and moral hazard are not the subject. Relation to earlier papers: Paper 364 showed selection alone raising a mean and shrinking a spread—there the cropping of the sample made the number; here the size of the sample does. Paper 371 showed a difference appearing with no cause—there too the provenance of an apparent difference was the issue. Paper 300 showed whether two things share a root is decidable—by that test, the expected loss and the safety loading have distinct roots. Paper 367 showed a conserved total with only the distribution free—here the total shrinks as $1/\sqrt{n}$. What is added is giving the loading at four sizes, showing that one policy needs more than 23 times the expected loss, showing that a hundredfold in n gives exactly 0.1 in loading, solving backwards for the 214311 policies that hold the loading to 5 per cent, and putting the separator on how many share the risk.

1. Notation and the Question

Setting: Take a loss probability $p = 0.01$, a loss size $L = 10^6$, and n policies.

$$\text{expected loss} = pL, \quad \sigma_1 = L\sqrt{p(1-p)}, \quad \sigma_n = \frac{\sigma_1}{\sqrt{n}}, \quad \text{loading} = z\sigma_n$$

with $z = 2.3263478740$, the 99th percentile of the normal (a ruin probability of 1 per cent).

Observation 1.1

The expected loss is what comes out on average. $pL = 10000$ —collect that and the books balance in the long run.

Observation 1.2

But “in the long run” is the problem. In the short run it misses—and if one cannot pay when it misses, that is the end.

Observation 1.3

So a loading is needed. Set by the size of the spread——this paper gives that size at each scale.

Question: Is the expected loss enough for a premium?

2. One Policy's Spread Is 9.9499 Times Its Mean

Calculation 2.1: Statistics of one policy (computed in this paper).

Quantity	Value
Loss probability p	0.01
Loss size L	1000000
Expected loss pL	10000.00
Standard deviation σ_1	99498.74
Relative to the expected loss	9.9499

Observation 2.2

The spread is 9.9499 times the mean. This is the numerical face of “rare but large”——because it is 0 in 99 per cent of cases and 1000000 in 1 per cent.

Observation 2.3

Nothing clusters near the mean. A loss of exactly 10000 never occurs——it is either 0 or 1000000.

Observation 2.4

So collecting only the expected loss ruins you in the 1 per cent case. A claim for 1000000 against 10000 in hand——it is not enough.

3. The Loading Runs from 23.1469 to 0.0231

Calculation 3.1: Loading against the number of policies (computed in this paper, ruin probability 1 per cent).

Policies n	Std. dev. per policy	Relative to mean	Loading (times the mean)
1	99498.74	9.9499	231468.69 (23.1469)
100	9949.87	0.9950	23146.87 (2.3147)
10000	994.99	0.0995	2314.69 (0.2315)
1000000	99.50	0.0099	231.47 (0.0231)

Observation 3.2

This is the core of the paper. With one policy the loading is 23.1469 times the expected loss—the premium becomes 241469 rather than 10000. Nobody buys it.

Observation 3.3

With a million it is 0.0231. A premium of 10231—2.3 per cent added to the expected loss.

Observation 3.4

The expected loss has not moved. 10000 in every row—only the number of policies changed.

Observation 3.5

A hundredfold in n gives exactly 0.1 in loading. $1/\sqrt{100} = 0.1$ —not approximately: the law of large numbers appearing directly as $1/\sqrt{n}$.

Observation 3.6

It inverts too. Holding the loading to 5 per cent of the expected loss takes **214311** policies——substituting back returns exactly 5 per cent. The formula closes.

4. Insurance Is Not the Business of Guessing Probabilities

Observation 4.1

Knowing p perfectly does not make one policy work. **23.1469** times——this is not a problem of estimation but of the spread itself.

Observation 4.2

Conversely, an imperfect p still works when n is large. Because the loading has room—— n is what makes the number.

Observation 4.3

So the core of the business is gathering. Increasing the number of parties to share with——“collect independent counterparties” acts more strongly than “estimate the probability”.

Observation 4.4

Without independence the computation fails. In an earthquake or an epidemic everyone claims at once——the benefit of $\sqrt{n}$ vanishes. This paper has not computed that case.

Observation 4.5

Distinct roots by the criterion of Paper 300. The expected loss is fixed by probability and size, the loading by variance and count——both enter “the premium” from different roots.

5. The Separator: How Many Share the Risk

Observation 5.1

The separator——how many the risk is divided among; not the accuracy of the probability.

	Expected loss	Loading	Viable as insurance
$n = 1$	10000	231469	no
$n = 10000$	10000	2315	yes
$n = 1000000$	10000	231	yes

Observation 5.2

The left column is identical in all three rows. An expected loss of 10000——the same risk, probability and size, and yet viability differs.

Observation 5.3

The separator can be moved. Cede to a reinsurer and the effective n grows——if one cannot gather enough oneself, one joins somebody else’s pool.

Observation 5.4

Sometimes it cannot. Correlated risks do not shrink their variance with n ——which is why catastrophes are hard to insure.

Observation 5.5

To be stated honestly——23.1469 and 214311 are values under stated assumptions: $p = 0.01$, $L = 10^6$, independence, and a normal approximation. Other conditions give other numbers——what does not move is the structure: the loading shrinks as $1/\sqrt{n}$.

Observation 5.6

One last thing—look only at the expectation and an unviable business looks viable. The expected loss is 10000 in all three rows—a mean says nothing about how far that mean can be trusted. Without looking at the variance, one cannot tell whether the business stands.

6. The Census and the Close

1. Expected loss 10000 (Calculation 2.1): $p = 0.01$, $L = 10^6$.
2. Standard deviation 99498.74 (Calculation 2.1): 9.9499 times the mean.
3. A loss of 10000 never occurs (Observation 2.3): it is 0 or 10^6 .
4. $n = 1$ needs a loading of 23.1469 (Calculation 3.1): a premium of 241469.
5. $n = 10^6$ needs 0.0231 (Calculation 3.1): a premium of 10231.
6. The expected loss never moves (Observation 3.4): 10000 in every row.
7. $\times 100$ in n gives $\times 0.1$ (Observation 3.5): $1/\sqrt{n}$.
8. 5 per cent takes 214311 policies (Observation 3.6): the inversion closes.
9. A perfect p does not save $n = 1$ (Observation 4.1): not an estimation problem.
10. Correlation removes the $\sqrt{n}$ (Observation 4.4): catastrophes.

Closing: It is tempting to think a premium is the probability of loss times its size—but on that alone the business does not stand. The expected loss is 10000 and one policy's standard deviation is 99498.74—9.9499 times the mean. Since it is 0 in 99 per cent of cases and 1000000 in 1 per cent, a loss of exactly 10000 never occurs. Collect only the expected loss and the 1 per cent case ruins you. So a loading is needed — and its size is set by the number of policies—holding the ruin probability at 1 per cent, $n = 1$ needs a loading of 23.1469 times the expected loss, making the premium 241469 rather than 10000. Nobody buys it. $n = 10^6$ needs 0.0231, a premium of 10231. The expected loss has not moved — 10000 in every row — and only the number of policies changed. A hundredfold in n gives exactly 0.1 in loading, the law of large numbers appearing directly as $1/\sqrt{n}$. Inverted, holding the loading to 5 per cent takes 214311 policies. So insurance is not the business of guessing probabilities—knowing p perfectly still needs 23.1469 times with one policy; this is not estimation but spread. Conversely an imperfect p works when n is large. The core is “collect independent counterparties” rather than “estimate the probability”. One thing separates them—how many the risk is divided among; not the accuracy of the probability. All three rows share one risk, one probability, one loss size, with an expected loss of 10000 — and only viability differs. Ceding to a reinsurer raises the effective n , but correlated risks do not shrink their variance with n , which is why catastrophes are hard to insure. To be said honestly—23.1469 and 214311 rest on stated assumptions: $p = 0.01$, $L = 10^6$, independence, and a

normal approximation. What does not move is the structure: the loading shrinks as $1/\sqrt{n}$. One last thing—look only at the expectation and an unviable business looks viable. The expected loss is 10000 in all three rows. A mean says nothing about how far that mean can be trusted, and without the variance one cannot tell whether the business stands.

References

1. Lundberg, F. (1903). *Approximerad framställning av sannolikhetsfunktionen*. Uppsala: Almqvist & Wiksell.
2. Cramér, H. (1930). *On the Mathematical Theory of Risk*. Stockholm: Skandia.
3. Bühlmann, H. (1970). *Mathematical Methods in Risk Theory*. Berlin: Springer.
4. Klugman, S. A., Panjer, H. H., Willmot, G. E. (2012). *Loss Models: From Data to Decisions* (4th ed.). Hoboken: Wiley.
5. Borch, K. (1962). Equilibrium in a Reinsurance Market. *Econometrica*.
6. Yamagishi, Y. (2026). *The Three Hundredth Census: Whether Two Things Share a Root Is Decidable* [Paper 300]. Zenodo. 10.5281/zenodo.22171914
7. Yamagishi, Y. (2026). *Selection Alone Raises the Mean by 1.52 Standard Deviations and Shrinks the Spread to 0.45* [Paper 364]. Zenodo. 10.5281/zenodo.22697530
8. Yamagishi, Y. (2026). *Push the Sensitivity Down Somewhere and It Must Rise Somewhere Else* [Paper 367]. Zenodo. 10.5281/zenodo.22697709
9. Yamagishi, Y. (2026). *With No Cause at All, the Bottom Rises 1.76 Standard Deviations and the Top Falls 1.76* [Paper 371]. Zenodo. 10.5281/zenodo.22698226

— *From thinking to syncing.* —

Yuuki YAMAGISHI (山岸 勇輝)

© 2026 Yamagishi Complexity Torus Lab, Japan · Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0)