The Frequency, Severity and Uncertainty of Estimation in a United States Hurricane Loss Model

Noah Cannitelli, Timothy Susca

2026-08-21

Abstract

It is common to model catastrophe losses using a compound Poisson process. In the case of hurricanes, this would involve storms arriving at a constant rate λ, with each carrying an independent random loss. We test both the frequency and severity halves of this model against the NOAA HURDAT2 dataset and Weinkle et al. (2018) normalized damages for the continental USA, from 1900-2017.

Fitting a Poisson/Negative Binomial to the frequency gives $\hat{\lambda} = 1.663$ and $\hat{k} = 7.358$ (per year.) Fitting a lognormal to the n=160 storms using observed damage records gives $\hat{\mu} = 0.682$ and $\hat{\sigma} = 2.204$ (in $Billions USD.) A simulated 95% envelope was constructed from 10,000 samples contains all observations, thus no evidence is found against this model.

Replacing the Poisson model with a negative binomial, and allowing for the arrival rate to vary randomly, gives an expected annual loss that is unchanged and moves the 99.6th percentile by -1.6%. This difference cannot be distinguished from zero by our simulation. Re-sampling and refitting moves the same percentile across a 95% interval of $448 billion to $2051 billion, which is a width that is 156% of the point estimate. This implies that the model is limited mainly by the sample size of the severity dataset.

1 Introduction

An insurer who wishes to write US hurricane coverage would be required to find two main quantities. Firstly, the expected annual payout, which in turn allows the premium to be determined. Secondly, a high percentile of the annual loss distribution, which would allow the insurer to determine how much capital should be reserved in order for the insurer to avoid insolvency in particularly bad years.

Thus we wish to construct a model of total annual loss. This is commonly done by a compound Poisson process with a constant λ, which is estimated by the historical average, and which treats each of the storm's losses as an independent draw from a separate severity distribution.

We choose to define a peril as continental US hurricane landfalls rather than just tropical cyclone counts. This is due to the data we are working with, some of which was collected before satellites existed. Thus, some storms are left unaccounted for in early years. The reason why landfills fix this problem is because these were recorded for the entire time period in question.

2 Data

2.1 Frequency

We chose to use the NOAA HURDAT2 landfall counts. HURDAT2 specifically avoids bias caused by data collected before the existence of satellites. Prior to approximately 1966, storms at sea were often unrecorded as nobody was able to record them, which would cause our analysis to be skewed and ultimately cause bias from undercounting.

2.2 Severity

We chose to use Weinkle et al. (2018) normalized damages. Note that this data needed to be normalized due to hurricane-infrastructure improvements from the beginning of the time period of interest to the end of the time period of interest.

We used two published normalizations, namely PL18 and CL18. PL18 uses county population, and CL18 uses housing units. We used PL18 as the primary series and chose to report CL18 as a sensitivity.

2.3 Data Quality

We noticed some limitations of the severity data:

Imputed Damages

37 of the 197 storms lack recorded base damage. We noticed that in these cases, the normalized losses are represented by 1 of 6 fixed values, which can be reproduced in the R script. Since a continuous distribution cannot recreate this behaviour, we decided to remove these 37 storms and instead fit on the 160 storms with observed damage records.

Left Truncation

The minimal normalized loss was $2.79 million. This is likely due to early-year, low-damage hurricanes not being reported, so this may cause the fitted distribution to be biased by some unknown amount.

Testing for Trends

Muller et al. (2022) warns against testing for a trend in normalized severity. The paper declines to do so on the grounds that the damage estimation methodology changed numerous times across the records. Thus, we will not test for a trend in normalized losses.

3 *Frequency**

** This section was completed by Timothy Susca. Thus, I do not have the corresponding R code for this section on my website.*

3.1 *Models*

We originally tested the Poisson as our model, and additionally tested the fit of a negative binomial. Using R, this yielded:

Metric	*Value*
Chi Squared	1.204
P Value	0.035
Poisson AIC	578.54
NB AIC	577.3

Both provided a good fit to the NOAA data, however Negative Binomial provided a better fit. Here is a summary of the final estimates:

$Parameter$	$Estimate$
$\hat{\lambda}$	1.663
$\hat{k}$	7.358
$Variance-to-mean$	1.204

3.2 *A Rolling Window Analysis*

We calculated 30-year rolling means for the annual hurricane counts. We constructed a 95% null envelope under the null hypothesis of a constant lambda (given above.) The results of which are shown in Figure 4.

The observed rolling mean appears to remain within the simulated envelope throughout the whole period, so this indicates that observed fluctuations around the mean are consistent with random variation.

4 Severity

4.1 Model

We chose to model normalized losses as a lognormal.

Note that if $X \sim Lognormal(\mu, \sigma^2)$, then $lnX \sim N(\mu, \sigma^2)$

We found the following model parameters using R: $\hat{\mu} = 0.682, \hat{\sigma} = 2.204, n = 160$.

Thus, we get:

$Quantity$	$Value$ ($\$Billion$)
$Median\ e^{\hat{\mu}}$	1.977
$Fitted\ Mean\ e^{\hat{\mu}+\hat{\sigma}^2/2}$	22.429
$Standard\ Dev$	253.457
$Sample\ Mean$	12.091

Note that the fitted mean and sample mean differ significantly. Consider the following:

$\frac{\partial lnE[X]}{\partial \sigma} = \frac{\partial(\mu+\sigma^2/2)}{\partial \sigma} = \frac{\partial \mu}{\partial \sigma} + \frac{\partial \sigma^2/2}{\partial \sigma} = \sigma \approx 2.2$

Thus, for each approximate 1% increase in σ causes a 2.2% change in $E[X]$.

$SE(\hat{\sigma}) = \frac{\sigma}{\sqrt{2n}} \approx 0.12$,so a 95% CI is approximately $\hat{\sigma} \pm 1.96 \cdot SE(\hat{\sigma}) \approx \hat{\sigma} \pm 53\%$

Thus, sampling variation in just $\hat{\sigma}$ causes a range in $E[X]$ of $\pm 53\%$. This implies great uncertainty in our estimates.

4.2 The Effect of Censoring Imputed Losses

We compute using R:

$Sample$	n	$\hat{\mu}$	$\hat{\sigma}$
$Observed\ Damages\ Only$	160	0.682	2.204
$All\ Storms\ Inclusive$	197	0.447	2.076

If we chose to include the imputed storms, $\hat{\mu}$ and $\hat{\sigma}$ would both be lowered, which pushes values towards the lower end of the range which aligns with the theory of why we removed them in the first place.

4.3 Testing the Fit of the Lognormal

Q-Q Plot (Figure 1): The observed values closely match the fitted lognormal quantiles with a bit of seperation at the extremes. However this supports the fit of the lognormal model.

Survival Plot (Figure 2): The fit is good for most of the curve, with some of the larger observations falling below.

Simulated Envelope (Figure 3): We drew 10,000 samples of size n=160 from the fitted lognormal and then took the 2.5th and 97.5th percentiles. Since all of the observed losses fall within the 95% bands, we find no evidence against the lognormal.

We also get from this result, that for the largest of 160 losses, the 95% band is $[162.8, 5730.8]$, with an observed value of 235.9. This implies a quite extreme upper tail, which means that although this model fits, there may be a large set of alternatives also.

4.4 Does the Choice of Normalization Method Matter?

Computing using R:

$Method$	$\hat{\mu}$	$\hat{\sigma}$
$PL18$	0.682	2.204
$CL18$	0.770	2.159

It appears that the choice of normalization method does not make a significant difference on the severity distribution.

5 Unifying Our Models

5.1 The Model

The total loss in a year can be modeled by:

$$S = \sum_{i=1}^{N} X_i$$

Where $N \sim Poisson(\lambda)$, each $X_i \sim Lognormal(\mu, \sigma^2)$, and thus S is a compound Poisson random variable. We assume N is independent of each X_i, and that each X_i are i.i.d.

5.2 The Expected Loss is Independent of the Frequency Distribution

By conditioning on N, $E[S|N = n] = n\mu_X$. Thus by the law of total expectation,

$$E[S] = E[E[S|N]] = E[N\mu_X] = E[N]\mu_X$$

Since only E[N] appears, we get that E[S] is independent of the frequency distribution. However, by the law of total variance,

$$Var[S] = E[Var[S|N]] = E[N]Var[X] + Var[N]\mu_X^2$$

So we get that under Poisson, $Var[N] = E[N] = \lambda E[X^2]$.

5.3 An Alternative Frequency Model

The Poisson process used so far has assumed a fixed λ, which can cause trends in the climate to be unaccounted for. We consider a model that lets the rate vary from year to year in an attempt to alleviate these problems.

Suppose $N|\Lambda \sim Poisson(\Lambda)$ with $\Lambda \sim Gamma$ giving $E[\Lambda] = \lambda$ and $Var[\Lambda] = \frac{\lambda^2}{k}$. Then N is negative binomial, giving

$$E[N] = E[E[N|\Lambda]] = E[\Lambda] = \lambda$$

$$Var[N] = E[Var[N|\Lambda]] = E[\Lambda] + Var[\Lambda] = \lambda + \frac{\lambda^2}{2}$$

Thus the negative binomial gives exactly what a Poisson process gives when the rate is allowed to be random. Computing using R gives:

n	$Poisson$	NB
0	0.1896	0.2233
1	0.3153	0.3029
2	0.2621	0.2333
3	0.1453	0.1342
4	0.0604	0.0640
5	0.0201	0.0268
6	0.0056	0.0102

Note how the negative binomial model yields larger values for extremes and less on normal years.

5.4 Computing Moments

Computing using R:

$Quantity$	$Poisson$	NB
$E[S]$	37.299	37.299
$Var[S]$	$107,668.7$	$107,857.7$
$Ratio_{Poi}$	1	1.00176

This seems to suggest that the two models should gives similar results.

5.5 Simulating Moments

For 1000000 years, we draw N from the Frequency model and draw N losses from the Severity model. Using R:

Quantity	*Simulated*	*Algebraic*
$E[S]$	37.577	37.299
$Var[S]$	123,349.8	107,668.7
$P(S=0)$	0.1894	$e^{-\lambda}=0.1896$

The means and probability of a no-loss year seem to agree, however the simulated variance does not. This can be explained by the sampling variance result found in 4.1., and thus it seems the variance of S is essentially unestimatable via simulation in general.

5.6 Summary of Results

All figures in $US Billions

Model	*Mean*	*99th Percentile*	*99.6th Percentile*	*Change in Mean*	*Change in 99.6th Percentile*
Poisson	37.577	526.680	1,027.546	–	–
NB	37.392	525.939	1,010.798	−0.49%	−1.63%

Note that the changes in mean and percentiles can be explained entirely by simulation variability, as 5.2 proves that they are identical.

Another interesting finding is the uncertainty in the estimates which are a consequence of the sample of severity. Since the estimates come from only 160 storms and are not known, we decided to test how this actually matters. We drew 1000 samples of 160 storms with replacement, fitted a lognormal to each, and then recombined the models. We get computing with R:

Quantity	*Value*
99.6th percentile point estimate	1027.55
95% interval across samples	[447.70, 2050.75]
Interval Width as % of PE	156%

So we get that uncertainty from the size of the sample for severity gives a roughly two orders of magnitude difference.

6 Some Limitations

As discussed in 4.1., although a lognormal fits the data, we cannot estimate the tails without uncertainty. If an insurer were to use this information to reserve capital for devastating, rare catastrophes, it would be done with significant uncertainty as to whether that amount reserved is sufficient.

Theoretically our model does not include an upper bound on loss, which is clearly not a reflection of reality as a hurricane is bounded by the total property value of its range.

We assumed that N is independent of each X_i and that each X_i are i.i.d. This is a little bit questionable; for example two hurricanes occurring in the same year may be related to higher temperatures in a given year.

7 References

NOAA National Hurricane Center. HURDAT2 Atlantic hurricane database.

Weinkle, J., C. Landsea, D. Collins, R. Musulin, R. P. Crompton, P. J. Klotzbach, and R. Pielke Jr. (2018). Normalized hurricane damage in the continental United States 1900–2017. Nature Sustainability, 1, 808–813.

Muller. et al. (2022). Hurricane Normalization 2022. GitHub repository. https://github.com/DrJoMuller/Hurricane-Normalization-2022

8 Figures

Figure 1:

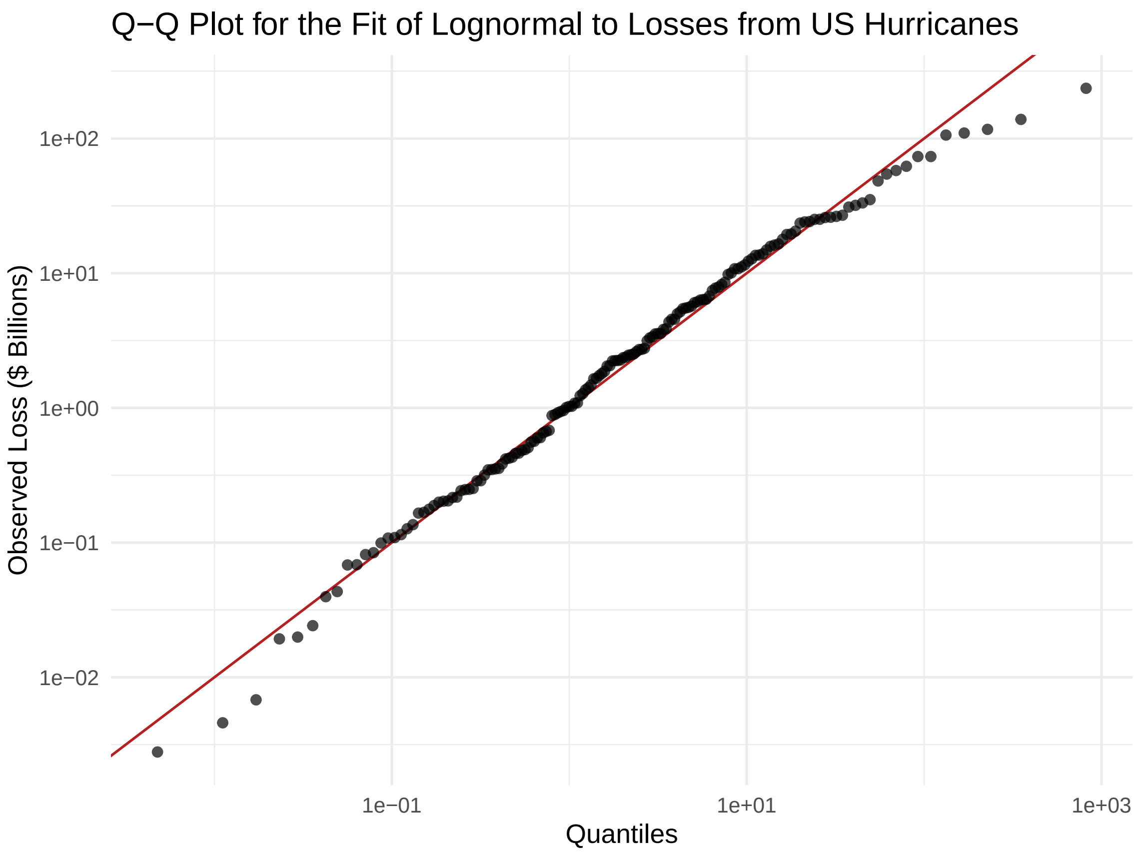

Figure 2:

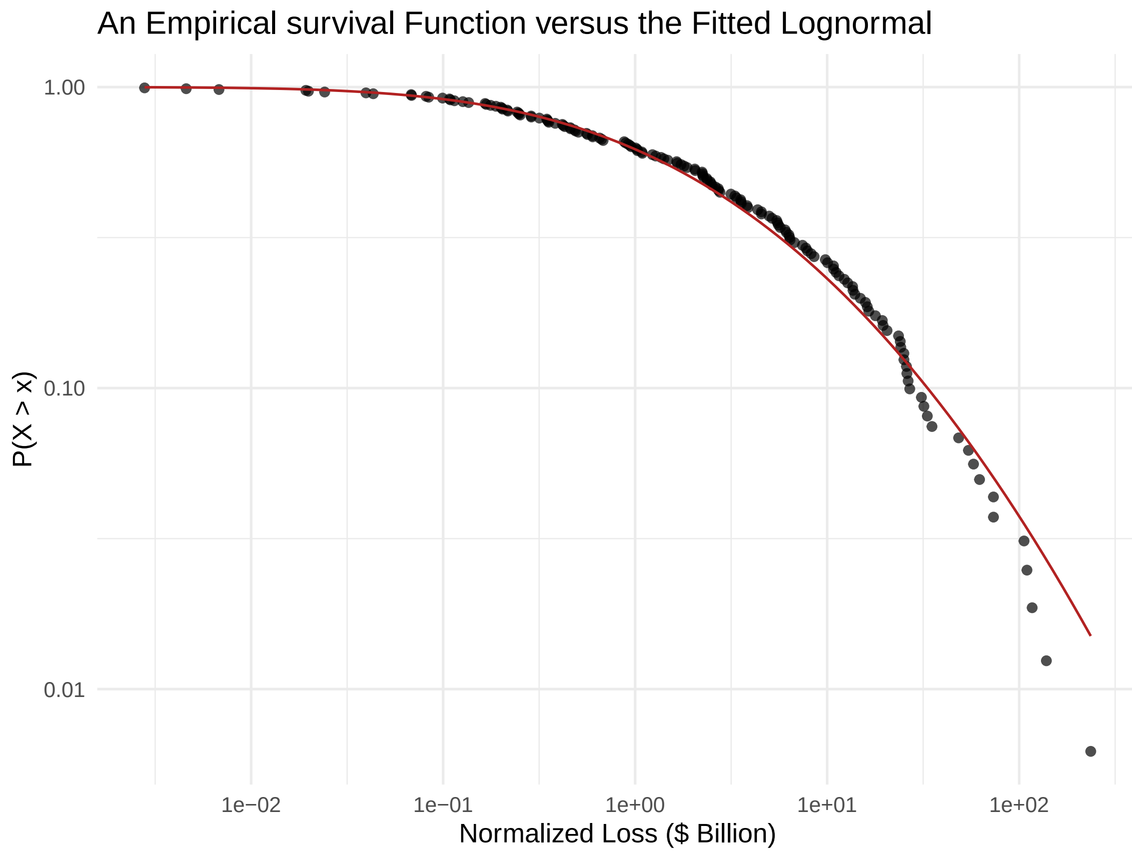
An Empirical survival Function versus the Fitted Lognormal
P(X > x)
1.00
0.10
0.01
1e−02
1e−01
1e+00
1e+01
1e+02
Normalized Loss ($ Billion)

Figure 3:

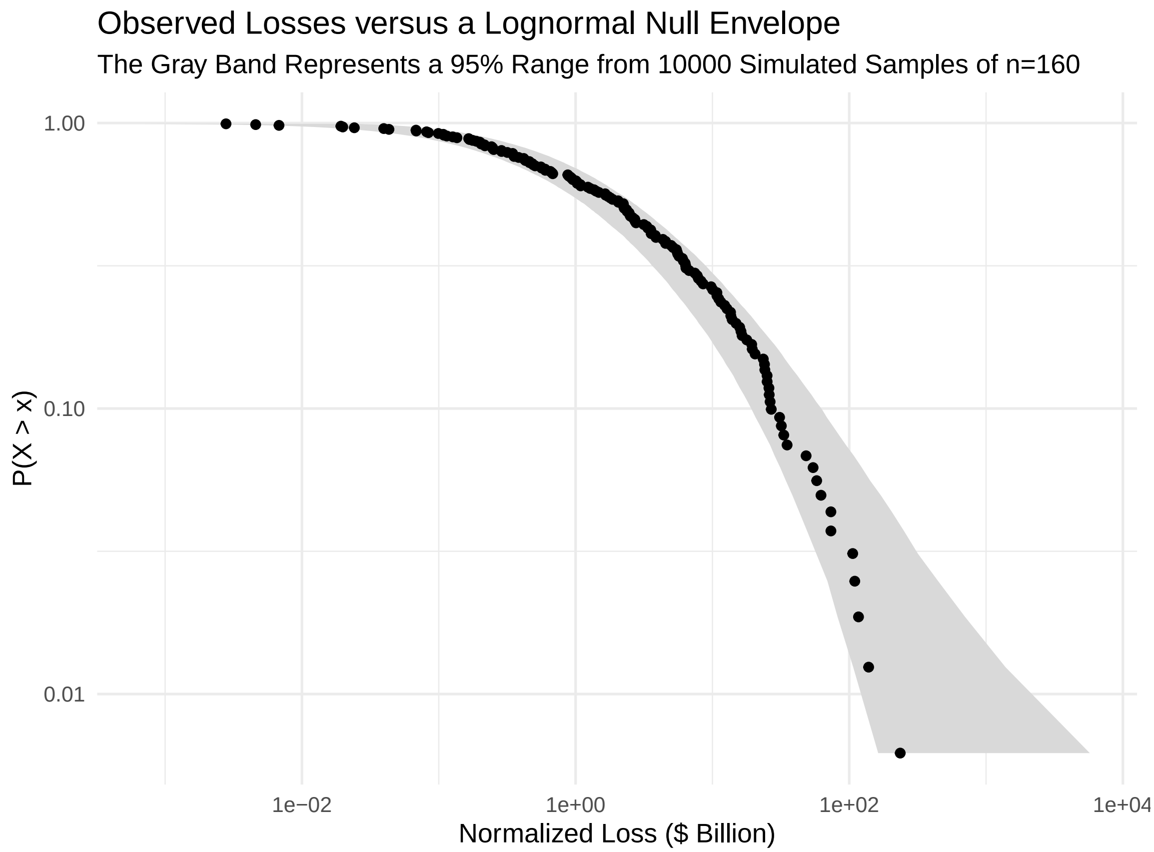
Observed Losses versus a Lognormal Null Envelope
The Gray Band Represents a 95% Range from 10000 Simulated Samples of n=160
1.00
0.10
0.01
P(X > x)
1e−02
1e+00
1e+02
1e+04
Normalized Loss ($ Billion)

Figure 4:

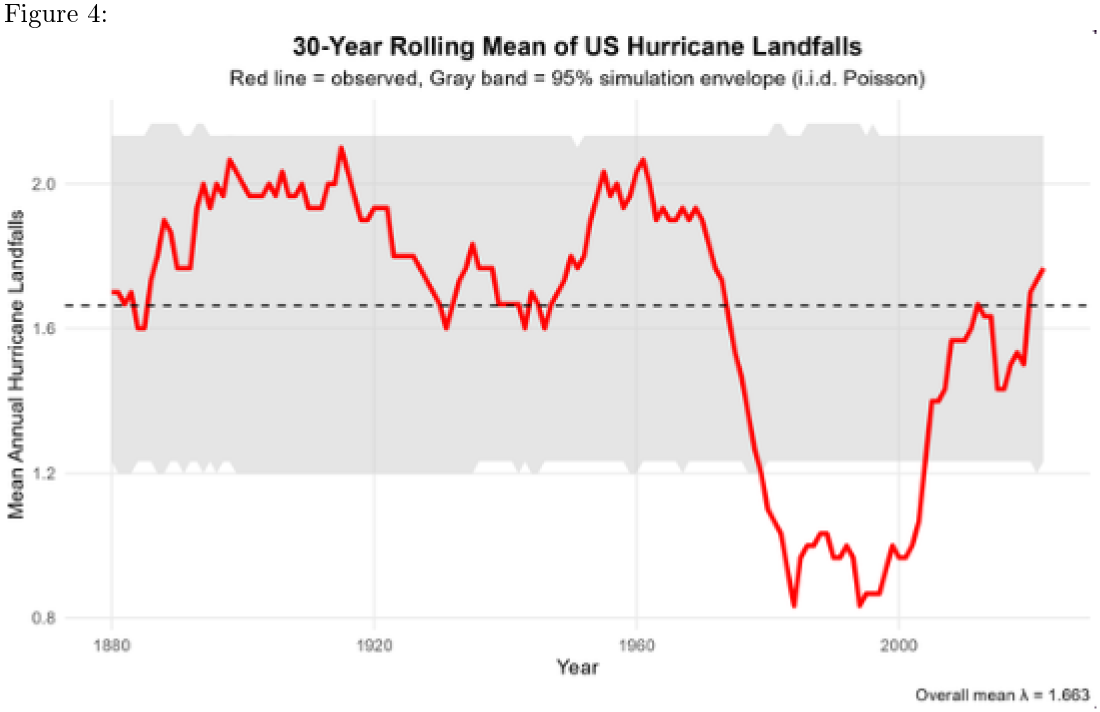
30-Year Rolling Mean of US Hurricane Landfalls
Red line = observed, Gray band = 95% simulation envelope (i.i.d. Poisson)
Mean Annual Hurricane Landfalls
2.0
1.6
1.2
0.8
1880
1920
1960
2000
Year
Overall mean λ = 1.663