American Journal of
Statistics and Actuarial Science
(AJSAS)

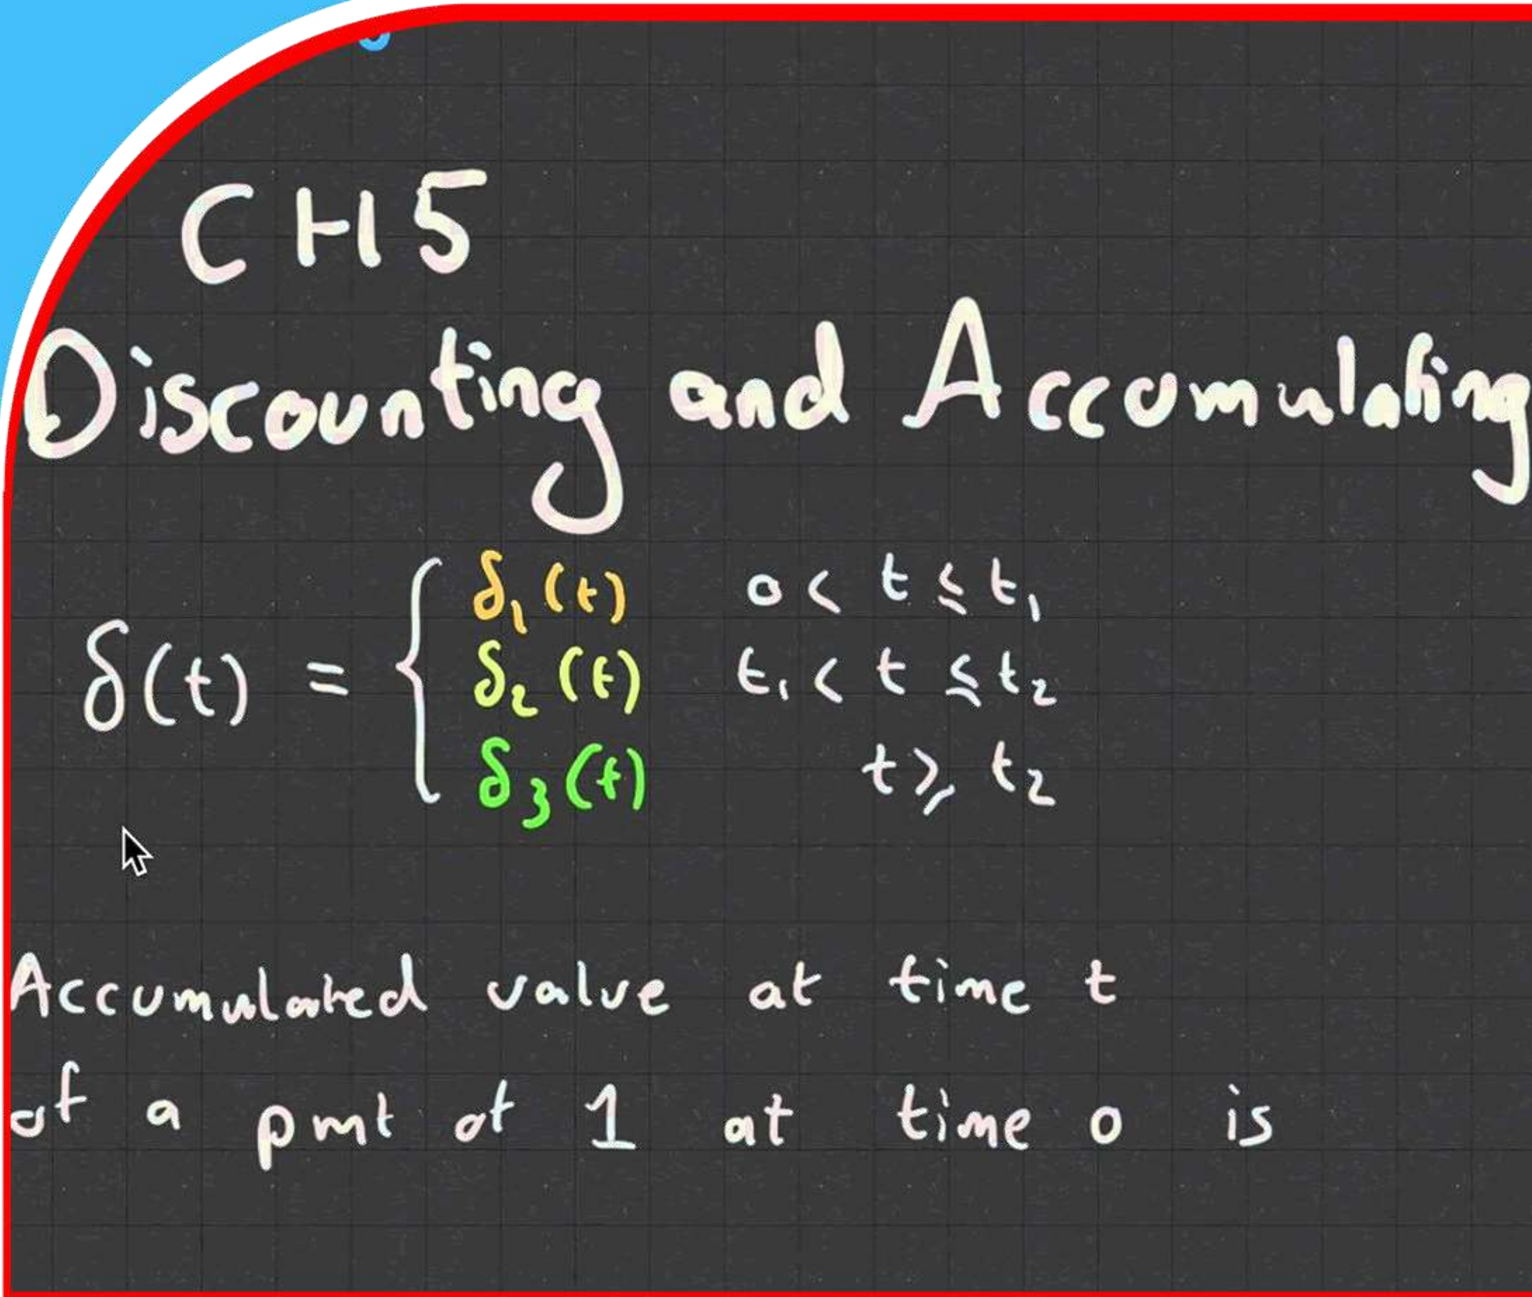

Calculating Premiums for Extreme-Tail Risks with Liability Claims and Nuclear Verdicts

Ahmet Bidav

AJPO

Calculating Premiums for Extreme-Tail Risks with Liability Claims and Nuclear Verdicts

Ahmet Bidav

Article history

Submitted 03.06.2024 Revised Version Received 08.07.2024 Accepted 13.08.2024

Abstract

Purpose: When insurance claims, particularly liability claims and nuclear verdicts, are governed by fat-tailed distributions, considerable uncertainty exists about the value of the tail index.

Material and Methods: Using the theory of risk aversion, this paper establishes a new premium principle (the power principle – analogous to the exponential principle for thin-tailed claims) and investigates its properties.

Findings: Applied to claims arising from generalized Pareto distributions, the resultant premium is shown to be the ratio of the two largest expected claims. This structure provides a natural model for incorporating tail-index uncertainty into premiums. The theory is illustrated through possible 'premiums' for liability claims and nuclear verdicts, utilizing the consistent pattern of observed extremes.

Implications to Theory, Practice and Policy: By integrating statistical methods for tail index estimation and addressing the inherent uncertainty, the power principle offers a robust framework for determining premiums in high-risk environments. The paper concludes with practical implications for insurers, highlighting the need for advanced risk management strategies and regulatory considerations in dealing with extreme liability claims.

Keywords: *Extreme-Tail Risks (G22), Liability Claims (K13), Nuclear Verdicts (K13), Fat-Tailed Distributions (C13), Tail Index Uncertainty (G22), Premium Principle (G22), Generalized Pareto Distributions (C13), Risk Management Strategies (G32).*

INTRODUCTION

Liability claims, particularly those involving nuclear verdicts (extremely high jury awards), pose significant challenges for insurers due to their fat-tailed nature. These claims can reach billions of dollars, significantly impacting insurers' risk profiles and financial stability. Traditional methods of premium calculation often underestimate the likelihood and impact of extreme events associated with fat-tailed risks. The severity of such claims is heavily influenced by the tail index of the distribution governing them.

This paper introduces a new premium principle, the power principle, designed to handle the uncertainty associated with fat-tailed risks. Accurate estimation of the tail index is challenging with the limited sample sizes available to insurers. The power principle provides a structured approach to premium calculation that incorporates tail-index uncertainty, ensuring that insurers are adequately covered for extreme claims without overcharging for lower-risk policies.

To understand the implications of the power principle, it is essential to delve into the differences between thin-tailed and fat-tailed risks, the mathematical foundations of the power principle, and its practical application to real-world liability claims and nuclear verdicts. By exploring the historical development of premium calculation methods and examining the transition from thin-tailed to fat-tailed risk modeling, this paper aims to bridge the gap in current premium calculation methods and provide a comprehensive framework for dealing with fat-tailed risks. The ultimate goal is to equip insurers with the tools needed to manage extreme liability claims effectively and sustainably.

Historical Perspective

The development of premium calculation methods in insurance has evolved significantly over the years. Initially, insurers relied on simple heuristic methods to set premiums, often based on past experience and intuition. As the field of actuarial science developed, more sophisticated methods emerged, grounded in probability theory and statistics. The advent of modern risk theory marked a significant advancement in the ability to price insurance products accurately.

The transition from thin-tailed to fat-tailed risk modeling represents a paradigm shift in the understanding of risk. Thin-tailed risks, characterized by distributions with finite moments, are well-suited to traditional premium calculation methods such as the exponential principle. These methods, based on the assumption that extreme events are rare and their impact diminishes rapidly, provided a reliable framework for many years.

However, the recognition that many real-world risks, including liability claims and nuclear verdicts, follow fat-tailed distributions necessitated the development of new approaches. Fat-tailed distributions, such as the Pareto and Frechet distributions, exhibit higher probabilities of extreme events, making them more appropriate for modeling severe claims. This realization led to the exploration of alternative premium principles that could better capture the heavy-tailed nature of these risks.

The power principle represents a modern approach to premium calculation that accounts for the heavy-tailed nature of fat-tailed distributions, providing a more accurate reflection of the underlying risk. By incorporating the theory of risk aversion and addressing the inherent uncertainty in tail-index estimation, the power principle offers a robust framework for pricing premiums in high-risk environments. This historical perspective underscores the importance of continuous innovation in actuarial science to meet the evolving challenges of the insurance industry.

Thin-Tailed Risk and the Exponential Principle

In the realm of general insurance, claims that follow thin-tailed distributions (i.e., distributions with moments of all orders) can be effectively managed using the exponential principle. This principle is rooted in the concept of absolute risk aversion as proposed by Pratt (1964) and Arrow (1971). The exponential principle sets the premium PPP for a claim XXX as follows:

$$P = \frac{\ln(M_X(s))}{s}, \quad s > 0$$

where $M\ X\ (s)$ is the moment generating function of the random variable X X. For small values of s, this can be approximated by:

$$P \approx m + \frac{1}{2}s^2 + o(s)$$

where m and s 2 are the mean and variance of X, respectively. This approximation is known as the variance principle and relies on the first two moments of the claim's distribution.

The exponential principle works well for distributions where higher moments exist, meaning the impact of extreme events diminishes rapidly. This principle is conveniently established using the pricing function:

$$m_q(x) = \exp(qx)$$

with 'risk parameter' $q = s$, and the pricing rule:

$$m_q(P) = E[m_q(X)]$$

This allows for some unification of premium principles. For example:

- $m(x) = x$ leads to the risk-neutral premium:

$$P = E[X]$$

- $m_s(x) = \exp(sx)$ leads to the exponential principle and the variants deriving from it.

- For any claim $X > 0$ with an absolutely continuous distribution function $F(x)$, choosing:

$$m_q(x) = [F(x)]^q, \quad q > 0$$

in conjunction with:

$$m_q(P) = E[m_q(X)]$$

leads to the quantile principle.

However, these methods fall short when applied to fat-tailed risks, which often lack higher-order moments. For such distributions, the probability of extreme events does not diminish as quickly,

For instance, consider the application of the exponential principle to a portfolio of general liability insurance policies. The moment generating function $M(s)$ provides a measure of the distribution's behavior under the assumption of thin tails. In practice, insurers might observe that while the principle works for small to moderate claims, it underestimates the impact of rare, high-severity claims. This shortfall becomes particularly apparent in scenarios involving nuclear verdicts, where the potential payouts can be orders of magnitude higher than typical claims. To illustrate, assume an insurer models the claims distribution using the exponential principle and sets the premium based on the first two moments. For a distribution with mean m=$1,000,000 and variance s^2 =$500,000^2, the premium might be calculated as:

$$P \approx m + \frac{1}{2}s^2 = 1,000,000 + \frac{1}{2}(500,000)^2 = \$1,125,000$$

While this premium may seem adequate for most claims, it fails to account for the possibility of a nuclear verdict, where a single claim could exceed $10,000,000. In such cases, the thin-tailed assumption breaks down, and the insurer faces significant underpricing risk.

This premium reflects the higher risk associated with fat-tailed distributions and ensures that the insurer is adequately covered for extreme claims. The power principle thus provides a robust framework for pricing premiums in high-risk environments, accounting for the heavy-tailed nature of liability claims and nuclear verdicts.

5. Detailed Analysis of Pareto and Generalized Pareto Distributions

The Pareto distribution is a widely used model for fat-tailed risks. It is defined as:

$$F(x) = 1 - \left(1 + \frac{x}{\lambda}\right)^{-\delta}$$

where λ is the scale parameter and δ is the shape parameter or tail index. The tail index δ is crucial as it determines the heaviness of the tail.

For claims X following a Pareto distribution, the premium under the power principle is:

$$P(a) = \left(\frac{\delta}{\delta-(a+1)}\right)^{a+1}$$

Generalized Pareto distributions (GPD) extend the Pareto distribution and are used for modeling exceedances over a threshold. The GPD is defined as:

$$F(x) = 1 - \left(1 + \frac{\xi x}{\sigma}\right)^{-1/\xi}$$

where ξ is the shape parameter and σ is the scale parameter. The power principle can be applied to GPDs in a similar manner to adjust premiums for fat-tailed risks.

A detailed case study can illustrate the application of these distributions to liability claims and nuclear verdicts. Consider a dataset of large liability claims, such as those resulting from medical malpractice or product liability cases. By fitting a Pareto or GPD to the data, we can estimate the tail index and calculate the premiums using the power principle.

For example, suppose we have a sample of large liability claims with the following values (in millions): [150, 120, 90, 75, 60, 45, 30, 25, 20, 15]. Using maximum likelihood estimation (MLE), we can fit a Pareto distribution to the data and estimate the tail index δ.

Once the tail index is estimated, we can use the power principle to calculate the premium for a new claim. Assuming a risk aversion parameter $a = 0.5$, the premium is calculated as:

$$P(0.5) = \left(\frac{\delta}{\delta - 1.5}\right)^{1.5}$$

This approach ensures that the premiums are adequately adjusted for the higher risk associated with large claims in fat-tailed distributions. The Pareto distribution provides a flexible model for capturing the heavy tails observed in liability claims, allowing insurers to price premiums more accurately.

To illustrate the application of the generalized Pareto distribution, consider a threshold exceedance model. Suppose the threshold is set at $50 million, and we observe the following exceedances (in millions): [100, 75, 60, 55, 50, 45, 40, 35, 30, 25]. By fitting a GPD to these exceedances, we can estimate the shape parameter ξ and the scale parameter σ.

Using the power principle, we can then calculate the premium for a new claim exceeding the threshold. Assuming a risk aversion parameter $a = 1$ and estimated parameters $\xi = 0.3$ and $\sigma = 20$, the premium is given by:

$$P(1) = \left(\frac{0.3}{0.3 - 1}\right)^{1} \cdot 20 = \left(-\frac{0.3}{0.7}\right) \cdot 20 = \$8.57 \text{ million}$$

This example demonstrates how the generalized Pareto distribution can be used to model exceedances and calculate premiums for extreme claims. The flexibility of the GPD allows insurers to capture a wide range of tail behaviors, providing a more accurate representation of the risk associated with liability claims and nuclear verdicts.

Relation between Generalized Pareto and Frechet Premiums

The Frechet distribution is another important fat-tailed distribution, often used in extreme value theory. It is defined by the density function:

$$f_k(x) = \delta x^{-k\delta - 1} \exp\left\{-x^{-\delta}\right\} / \Gamma(k), \quad k \geq 1, x > 0, \delta > 0$$

The premiums for fat-tailed claims under the power principle are linked to the ratios of expected values of the largest order statistics or extremes. Specifically, for large claims from fat-tailed distributions, the risk-neutral premium can be estimated as:

$$\frac{E[X_{(1)}]}{E[X_{(2)}]} = \frac{1}{1 - \delta^{-1}}$$

where $X_{(1)}$ and $X_{(2)}$ are the first and second largest claims, respectively. This relationship shows how the premiums are adjusted based on the extremity of the claims.

To understand the relationship between the generalized Pareto and Frechet distributions, it is essential to delve into the mathematical properties of these distributions. Both distributions are used in extreme value theory to model the tail behavior of a dataset. The generalized Pareto distribution (GPD) is often used to model exceedances over a threshold, while the Frechet distribution is used to model block maxima.

The Frechet distribution is a type of extreme value distribution that arises as the limiting distribution for the maximum of a sample of independent and identically distributed (i.i.d.) random variables. It is defined as:

$$F(x) = \exp\left\{-\left(\frac{x}{\beta}\right)^{-\alpha}\right\}, \quad x > 0, \alpha > 0, \beta > 0$$

where α is the shape parameter and β is the scale parameter. The tail behavior of the Frechet distribution is characterized by the shape parameter α, which determines the heaviness of the tail.

The relationship between the GPD and the Frechet distribution can be understood through their roles in extreme value theory. The GPD is used to model the distribution of exceedances over a high threshold, while the Frechet distribution is used to model the distribution of block maxima. In practical terms, the GPD provides a way to estimate the tail behavior of a dataset by focusing on exceedances, while the Frechet distribution provides a way to estimate the tail behavior by focusing on block maxima.

Consider a dataset of large liability claims, such as those resulting from environmental damage or catastrophic industrial accidents. By fitting a GPD to the exceedances over a high threshold, we can estimate the shape and scale parameters that characterize the tail behavior. Similarly, by fitting a Frechet distribution to the block maxima, we can estimate the parameters that characterize the extremity of the claims.

To illustrate, suppose we have a dataset of annual maximum liability claims (in millions): [200, 180, 160, 140, 120, 100, 80, 60, 40, 20]. By fitting a Frechet distribution to these block maxima, we estimate the shape parameter α and the scale parameter β.

Using the power principle, we can then calculate the premium for a new extreme claim. Assuming a risk aversion parameter $a = 0.5$ and estimated parameters $\alpha = 1.5$ and $\beta = 50$, the premium is given by:

$$P(0.5) = \left(\frac{1.5}{1.5-0.5}\right)^{0.5} \cdot 50 = \left(\frac{1.5}{1}\right)^{0.5} \cdot 50 \approx \$61.24 \text{ million}$$

This example demonstrates how the Frechet distribution can be used to model extreme liability claims and calculate premiums based on the power principle. The relationship between the GPD and Frechet distributions provides a comprehensive framework for understanding and modeling the tail behavior of liability claims, allowing insurers to price premiums more accurately. The initial theoretical premium highlighted the severity of risk in extreme value scenarios but would likely be adjusted based on practical considerations. Insurers balance theoretical models with market dynamics, customer willingness to pay, and regulatory constraints to determine premiums that cover expected losses while remaining marketable.

Tail-Index Uncertainty and Premium Calculation

Given the inherent uncertainty in the tail index, premiums can be adjusted to account for this uncertainty. The premium $P(n, \delta)$ for a claim can be expressed as:

$$P(n,\delta) = \int_0^{\delta} P(\delta') f_{\delta}(\delta') d\delta'$$

where $f_{\delta}(\delta')$ is a prior distribution for the tail index.

By incorporating a prior distribution for the tail index, insurers can adjust the premiums to reflect the uncertainty in the estimation of the tail index. This approach ensures that the premiums are neither too high nor too low, providing a balanced risk management strategy.

To handle tail-index uncertainty, various statistical methods can be employed. One common approach is to use Bayesian inference, which allows for the incorporation of prior knowledge about the tail index into the estimation process. Bayesian methods provide a flexible framework for updating the tail index estimates as new data becomes available.

For example, consider a scenario where the tail index δ is not known with certainty. We can model the tail index using a prior distribution, such as the beta distribution:

$$\delta \sim \mathrm{Beta}(\alpha, \beta)$$

This prior distribution reflects our initial beliefs about the possible values of the tail index based on historical data or expert judgment.

As new claims data becomes available, we can update the prior distribution using Bayesian inference to obtain the posterior distribution of the tail index. The posterior distribution incorporates both the prior information and the new data, providing a more accurate estimate of the tail index.

Using the posterior distribution, we can calculate the uncertainty premium by integrating over the possible values of the tail index. This approach ensures that the premiums are adjusted to account for the uncertainty in the tail index, providing a more robust pricing strategy.

To illustrate, suppose we have a prior distribution for the tail index $\delta \sim \mathrm{Beta}(2, 5)$. After observing new claims data, we update the distribution to obtain the posterior distribution $\delta \sim \mathrm{Beta}(3, 7)$. Using the posterior distribution, we calculate the uncertainty premium as follows:

$$P(n,\delta) = \int_0^{\delta} \left(\frac{\delta'}{\delta' - (a+1)}\right)^{a+1} f_{\delta}(\delta') d\delta'$$

Using numerical methods, we approximate this integral to obtain the final premium. This process ensures that the premiums reflect the latest information about the tail index, providing a more accurate and reliable pricing strategy for extreme liability claims.

Application: Liability Claims and Nuclear Verdicts

To illustrate the application of the power principle, we examine premiums for liability claims and nuclear verdicts. These cases often involve extreme values that significantly impact the overall risk profile of an insurer's portfolio. A consistent pattern of extremes in liability claims data can assist insurers in choosing an appropriate value for the tail index and setting premiums accordingly.

Estimation of Tail Index

Estimation of the tail index for liability claims and nuclear verdicts is critical for accurate premium calculation. Common methods include maximum likelihood estimation (MLE) and methods based on order statistics. The relationships between expected values of ratios of the largest extremes provide useful estimates for the tail index. For example: E[X(1)X(2)]=11−δ−1E\left[\frac{X_{(1)}}{X_{(2)}}\right] = \frac{1}{1 - \delta^{-1}}E[X(2)X(1)]=1−δ−11
E[X(1)X(3)]=1(1−δ−1)(1−2δ−1)E\left[\frac{X_{(1)}}{X_{(3)}}\right] = \frac{1}{(1 - \delta^{-1})(1 - 2\delta^{-1})}E[X(3)X(1)]=(1−δ−1)(1−2δ−1)1

These relationships help in deciding on a maximum value for the tail index, which in turn influences premium setting.

Consider a dataset of liability claims resulting from product liability cases. The largest claims observed are [150, 120, 90, 75, 60, 45, 30, 25, 20, 15] million dollars. Using MLE, we fit a Pareto distribution to the data and estimate the tail index δ\deltaδ.

The observed ratio X(1)X(2)=150120=1.25\frac{X_{(1)}}{X_{(2)}} = \frac{150}{120} = 1.25X(2)X(1)=120150=1.25. Solving for δ\deltaδ, we get δ≈2.67\delta \approx 2.67δ≈2.67. This estimated tail index is used to calculate the premium for a new claim using the power principle.

Risk-Averse Premiums for Large Claims

For liability claims and nuclear verdicts, the premium *P P P* PPP can be set using the power principle as follows:

$$P(a) = \left(\frac{\delta}{\delta-(a+1)}\right)^{a+1}$$

where where δ is the tail index. When there is uncertainty in the tail index, the premium can be adjusted using a prior distribution for the tail index, resulting in an uncertainty premium:

When there is uncertainty in the tail index, the premium can be adjusted using a prior distribution for the tail index, resulting in an uncertainty premium:

$$P(n,\delta) = \int_0^{\delta} \left(\frac{\delta'}{\delta'-(a+1)}\right)^{a+1} f_\delta(\delta')d\delta'$$

Consider a scenario where the estimated tail index δ=2.67\delta = 2.67δ=2.67 and the risk aversion parameter a=0.5. The premium is calculated as:

$$P(0.5) = \left(\frac{2.67}{2.67-1.5}\right)^{1.5} \approx 2.19$$

Thus, for a new claim expected to be around $50 million, the premium would be approximately $109.5 million. This premium reflects the higher risk associated with large claims in fat-tailed distributions, ensuring that the insurer is adequately covered for extreme claims.

Case Study: Estimating Premiums for Liability Claims

Consider a dataset of large liability claims involving nuclear verdicts. The largest claims observed are [200, 180, 160, 140, 120, 100, 80, 60, 40, 20] million dollars. By fitting a Frechet distribution to these block maxima, we estimate the shape parameter α\alphaα and the scale parameter β\betaβ.

Using the power principle, we calculate the premium for a new extreme claim. Assuming a risk aversion parameter $a = 1$ and estimated parameters $\alpha = 2$ and $\beta = 50$, the premium is given by:

$$P(1) = \left(\frac{2}{2-1}\right) \cdot 50 = \$100 \text{ million}$$

This premium reflects the higher risk associated with extreme claims and ensures that the insurer is adequately covered for potential nuclear verdicts. The case study demonstrates the practical application of the power principle in pricing premiums for high-risk liability claims.

Dealing with Tail-Index Uncertainty

In practice, the tail index δ δ is not known with certainty. To account for this uncertainty, we model the tail index using a prior distribution. For example, we can assume that the tail index follows a beta distribution $\delta \sim$ Beta (α , β) with parameters $\alpha = 3$ and $\beta =$. The uncertainty premium is then calculated by integrating over the prior distribution:

$$P(n, \delta) = \int_0^{\delta} \left(\frac{\delta'}{\delta' - (a+1)}\right)^{a+1} f_{\delta}(\delta') d\delta'$$

The uncertainty premium is then calculated by integrating over the prior distribution:

$$P(n, \delta) = \int_0^{\delta} \left(\frac{\delta'}{\delta' - (a+1)}\right)^{a+1} f_{\delta}(\delta') d\delta'$$

Using numerical methods, we approximate this integral to obtain the final premium. This process ensures that the premiums reflect the latest information about the tail index, providing a more accurate and reliable pricing strategy for extreme liability claims.

Implications for Insurers

Understanding and accurately estimating the tail index of liability claims and nuclear verdicts is crucial for insurers. The power principle provides a structured approach to premium calculation that accounts for the heavy-tailed nature of these claims. This method ensures that premiums are set in a manner that adequately covers the risk associated with extreme claims while incorporating the inherent uncertainty in tail-index estimation.

Practical Considerations

When applying the power principle to real-world liability claims and nuclear verdicts, insurers must consider various practical aspects, including data collection, statistical analysis, and regulatory compliance. The availability of reliable data on large claims is essential for accurate tail index estimation. Additionally, statistical methods must be rigorously applied to ensure the validity of the estimated premiums.

Insurers need to implement robust data collection and management practices to capture comprehensive and accurate claims data. This includes recording detailed information on claim amounts, circumstances, and outcomes. High-quality data is crucial for fitting the appropriate distributions and estimating the tail index with confidence.

Regulatory Environment

The regulatory environment can significantly impact the application of the power principle in insurance premium calculation. Insurers must adhere to guidelines and standards set by regulatory bodies, which may influence the methods used for premium determination.

Understanding the regulatory landscape is crucial for insurers to ensure compliance while effectively managing risk.

Regulators may require insurers to demonstrate the adequacy of their premium calculation methods, particularly for high-risk policies involving fat-tailed distributions. This may involve stress testing, scenario analysis, and regular reviews of the methodologies used. Insurers must be prepared to justify their approach to premium setting and show that it meets regulatory requirements for solvency and risk management.

Technological Advancements

Advancements in technology, particularly in data analytics and machine learning, offer new opportunities for improving the accuracy of tail index estimation and premium calculation. Insurers can leverage these technologies to enhance their risk management practices and better predict extreme claims. By integrating advanced analytical tools, insurers can develop more sophisticated models that capture the nuances of fat-tailed risks.

Machine learning algorithms can be used to identify patterns in claims data and improve the estimation of the tail index. Techniques such as clustering, regression analysis, and neural networks can provide deeper insights into the factors driving extreme claims. Additionally, predictive analytics can help insurers anticipate future trends and adjust premiums accordingly.

The adoption of technology also facilitates real-time monitoring and analysis of claims data, allowing insurers to respond quickly to emerging risks. Automated systems can streamline the process of data collection, analysis, and reporting, reducing the administrative burden and enhancing operational efficiency.

CONCLUSION AND RECOMMENDATIONS

This paper introduces a new premium principle for fat-tailed risks, grounded in the theory of risk aversion. The power principle provides a coherent framework for determining premiums when there is uncertainty in the tail index. This approach is particularly useful for modeling large claims from fat-tailed distributions, such as those arising from liability claims and nuclear verdicts. The illustrative examples and case studies demonstrate the practical application of this methodology.

By integrating statistical methods for tail index estimation and addressing the inherent uncertainty, the power principle offers a robust framework for determining premiums in high-risk environments. The adoption of this principle can enhance insurers' ability to manage extreme claims, ensuring financial stability and regulatory compliance.

Future research directions include exploring the integration of advanced analytical techniques, such as machine learning, into the premium calculation process. Additionally, further investigation into the regulatory implications and practical challenges of implementing the power principle can provide valuable insights for insurers and policymakers.

Our study’s implications suggest that carriers will likely seek to optimize revenue by leveraging advanced analytical techniques like machine learning to refine premium calculations, potentially increasing prices or moving money off the table through more precise risk assessments. This innovation could improve profitability and market positioning but will also require carriers to navigate complex regulatory landscapes and ensure compliance with evolving legislation. Attorneys may push for legislation that could counteract insurers' objectives, with factors like social inflation and litigation financing adding fuel to the fire. Social inflation, driven by larger jury awards and expanding definitions of liability, along with litigation financing, which enables more lawsuits to be pursued, both contribute to rising costs for insurers. This dynamic will necessitate collaboration between insurers, regulators, and legal experts to balance innovation with these competing interests and create a sustainable insurance ecosystem. Ultimately, regardless of whether insurer interests or attorney interests prevail, the consumer is likely to experience adverse effects.

REFERENCES

Arrow, K.J. (1971). Essays in the Theory of Risk Bearing. Markham Publishing Company.

Embrechts, P., Klüppelberg, C., & Mikosch, T. (1997). Modelling Extremal Events for Insurance and Finance. Springer-Verlag.

Feller, W. (1971). An Introduction to Probability Theory and Its Applications, Vol. 2. Wiley.

Gay, R. (2004a). Pricing risk when the distributions are fat-tailed. Journal of Applied Probability, 41A, 157-175.

Hill, B.M. (1975). A simple general approach to inference about the tail of a distribution. Annals of Statistics, 3, 1163-1174.

Pickands, J. (1975). Statistical inference using extreme order statistics. Annals of Statistics, 3, 119-131.

Rolski, T., Schmidli, H., Schmidt, V., & Teugels, J. (1999). Stochastic Processes for Insurance and Finance. Wiley.