

Paradox-Proof Utility Functions for Heavy-Tailed Losses: The Insurance Two-Envelope Problem

Siwei Gao*, Michael R. Powers^{†‡}

July 14, 2015

Abstract

We employ the Insurance Two-Envelope Paradox to explore the disparate effects of bounded versus unbounded buyer utility in the presence of large monetary losses (negative gains). Given that the paradox requires certain losses to have infinite means, we first justify the use of such models in insurance applications. We subsequently argue that the practical impact of using unbounded utility functions in insurance modeling has received inadequate attention in the research literature.

Keywords: von Neumann-Morgenstern utility; two-envelope paradox; St. Petersburg paradox; dominance reasoning; heavy-tailed gains/losses; concave utility; bounded utility.

1 Introduction

The von Neumann-Morgenstern (“VNM”) expected-utility theorem is remarkable both for its theoretical simplicity and broad applicability. In brief, it states that if four intuitive axioms (completeness, transitivity, continuity, and independence) hold with regard to a decision maker’s preferences over the universe of all lotteries, then there must exist a continuous, increasing utility function, $u(\cdot)$, such that, for any two lotteries A and B with respective random monetary gains Y and Z ,

$$A \succeq B \Leftrightarrow E[u(W + Y)] \geq E[u(W + Z)],$$

* Assistant Professor, Department of Accounting, Finance, and Information Systems, School of Business, Eastern Kentucky University.

[†](Corresponding author) Zurich Group Professor of Risk Mathematics, School of Economics and Management, and Professor, Schwarzman Scholars Program, Tsinghua University. Contact information: 386C Weilun Building, School of Economics and Management, Tsinghua University, Beijing 100084, China; phone – (8601) 62795565; e-mail – powers@sem.tsinghua.edu.cn.

[‡]The authors thank the AXA Research Fund for its generous support.

where $W \in (-\infty, \infty)$ denotes the decision maker’s initial wealth. This utility function is unique up to a positive linear transformation. (See von Neumann and Morgenstern, 1947.)

Although VNM’s theorem imposes no further constraints on the shape of $u(\cdot)$, researchers often assume, based upon behavioral observations going back as far as Bernoulli’s solution of the St. Petersburg Paradox (“SPP”; Bernoulli, 1738), that such functions are concave downward (i.e., that decision makers are risk averse). Moreover, by extending the SPP to random gains with arbitrarily heavy tails, it can be argued (as in Menger, 1934) that utility functions must be bounded above. Finally, by constructing the SPP with monetary losses, rather than gains, one can assert (by symmetry) that VNM utility functions must be concave upward and bounded below for sufficiently large losses. Although rarely discussed, these last characteristics of utility functions are of particular importance in the study of buyer behavior in insurance and reinsurance markets.

The classical Two-Envelope Paradox of decision theory and philosophy has been used to show that dominance reasoning (i.e., $A|C \succ B|C \forall C \Rightarrow A \succ B$) does not hold for all VNM utility functions. (See Nalebuff, 1989; Broome, 1995; Chalmers, 2002; and Powers, 2015.) Powers and Zanjani (2013) introduced the conceptually similar Insurance Two-Envelope Paradox (“ITEP”) to demonstrate the difficulty of selecting an appropriate utility function for buyers of insurance and reinsurance in the context of a commonly used actuarial model (i.e., exponential losses with an inverse-exponentially distributed mean loss parameter, yielding unconditionally Pareto losses with an infinite mean). In the present article, we employ the ITEP to explore the disparate effects of bounded versus unbounded buyer utility in the presence of large monetary losses. Since the paradox requires the existence of infinite-mean losses, we first justify the use of such models in insurance applications. We subsequently argue that the practical impact of using unbounded utility functions in insurance modeling has received inadequate attention in the research literature.

2 Heavy-Tailed Insurance Losses

2.1 A Hierarchy of Definitions

Insurance actuaries and researchers often speak of “heavy-tailed losses” associated with certain types of insurance, especially product liability, environmental/pollution liability, and property catastrophe coverages. Although there is no universally accepted definition of what it means for a continuous loss random variable, $L \sim F_L(\ell)$, $\ell \geq 0$, to be heavy-tailed, several candidate families are commonly suggested:

- $\mathcal{HT}_1^{(k)}$: The loss distribution has an infinite k^{th} moment (i.e., $\int_0^\infty \ell^k dF_L(\ell)$ diverges to infinity for some positive integer k).
- \mathcal{HT}_2 : The loss distribution is subexponential (i.e., $\lim_{\ell \rightarrow \infty} \frac{1-F_{S_n}(\ell)}{1-F_{M_n}(\ell)} = 1$ for

some $n \geq 2$, where $S_n = \sum_{i=1}^n L_i$ and $M_n = \max_{i \in \{1, 2, \dots, n\}} \{L_i\}$ for i.i.d. L_i ; or equivalently, $\lim_{\ell \rightarrow \infty} \frac{1 - F_{S_n}(\ell)}{n[1 - F_L(\ell)]} = 1$ for some $n \geq 2$).

- \mathcal{HT}_3 : The loss distribution has an asymptotic residual waiting-time probability of 1 (i.e., $\lim_{\ell \rightarrow \infty} \frac{1 - F_L(\ell + t)}{1 - F_L(\ell)} = 1$ for all $t > 0$). Although some authors use the term “long-tailed losses” to describe \mathcal{HT}_3 , this is ill-advised in the insurance context, where “long-tailed losses” mean loss payments that require a long time to settle.
- \mathcal{HT}_4 : The loss distribution has an undefined moment-generating function (i.e., $\int_0^\infty e^{t\ell} dF_L(\ell)$ diverges for all $t > 0$; or equivalently, $\lim_{\ell \rightarrow \infty} e^{t\ell} [1 - F_L(\ell)] = \infty$).

These definitions are decreasingly restrictive, so that $\mathcal{HT}_1^{(1)} \subset \mathcal{HT}_1^{(2)} \subset \mathcal{HT}_1^{(3)} \subset \dots$ and $\mathcal{HT}_1^{(k)} \subset \mathcal{HT}_2 \subset \mathcal{HT}_3 \subset \mathcal{HT}_4$ for any k . Another family, \mathcal{HT}' , for which the loss distribution is dominatedly varying (i.e., $\limsup_{\ell \rightarrow \infty} \frac{1 - F_L(\ell/2)}{1 - F_L(\ell)} < \infty$), satisfies the relations $\mathcal{HT}_1^{(k)} \subset \mathcal{HT}' \subset \mathcal{HT}_4$ for any k , but neither contains nor is contained by \mathcal{HT}_2 or \mathcal{HT}_3 . (See, e.g., Embrechts et al., 1997, p. 50.)

In the present article, we are particularly interested in $\mathcal{HT}_1^{(1)}$ – the “heaviest” of the above sets – because infinite-mean losses are necessary for the ITEP to work. However, we recognize that the relevance of such losses to insurance is not universally acknowledged, and therefore offer justification for our focus before proceeding.

2.2 Infinite-Mean Losses in Insurance

Many practitioners and scholars would argue, quite reasonably, that infinite means are impossible because all insurance indemnities are implicitly bounded above by some finite number. Even in the most extreme case of a large, class-action liability suit, a generous judge or jury is certain to award damages less than B , the relevant nation’s GDP – which, for the United States, is currently between 1.5×10^{13} and 2.0×10^{13} USD. (Of course, an insurance company will never pay an indemnity greater than the sum of its total assets and available credit. However, in the present discussion, L denotes the insurance company’s theoretical financial obligation, not its actual loss payment.) Since bounded random variables cannot possess infinite moments, the required conclusion follows immediately.

To counter this argument, we first would acknowledge that it is – in a technical mathematical sense – absolutely correct. However, we also would point out that the same argument inevitably leads to other, equally technical but practically troublesome, conclusions. For example, once we recognize that L is bounded, we also must acknowledge that no continuous distribution defined on $[0, \infty)$ is appropriate for L , simply because the upper bound of the sample space

is $B < \infty$. This excludes the vast majority of commonly used loss models, including the very popular gamma, Weibull, lognormal, and Pareto distributions.

Naturally, one could refine the “no infinite means” argument to overcome this difficulty by asserting that approximating the sample space $[0, B]$ with $[0, \infty)$ is perfectly reasonable as long as it affects only such peripheral items as tail probabilities and high percentiles, but becomes problematic if it substantially alters the mean, which is central to the law of large numbers and (consequently) the calculation of insurance premiums. To address this more refined objection, we will show that the impact of the difference between $[0, B]$ and $[0, \infty)$ on the mean is not of practical significance.

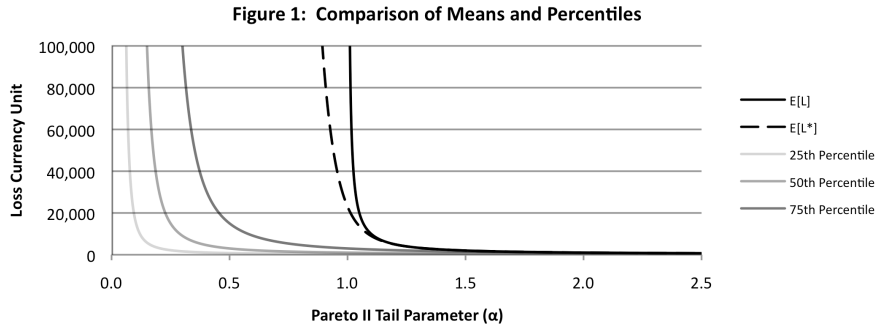
To make matters concrete, let $L \sim \text{Pareto II}(\alpha, \theta)$ (i.e., $F_L(\ell) = 1 - [\theta/(\ell + \theta)]^\alpha$) and $L^* = \min\{L, B\}$, so that

$$E[L] = \begin{cases} \frac{\theta}{\alpha - 1} & \text{for } \alpha > 1 \\ \infty & \text{for } \alpha \leq 1 \end{cases}$$

and

$$E[L^*] = \begin{cases} \frac{\theta}{\alpha - 1} \left[1 - \left(\frac{\theta}{B + \theta} \right)^{\alpha - 1} \right] & \text{for } \alpha \neq 1 \\ -\theta \ln \left(\frac{\theta}{B + \theta} \right) & \text{for } \alpha = 1 \end{cases}.$$

Figure 1 shows how the above expected values, as well as the 25th, 50th, and 75th percentiles of L (and L^*), vary as a function of the distribution parameter α for $\theta = 1000$ and $B = 10^{13}$. One can see that as α approaches 1 from the right, the mean of L diverges to infinity, becoming undefined for all $\alpha \leq 1$. At the same time, because of the relatively large value of B (compared to the scale parameter, θ), the mean of L^* grows more rapidly than does any fixed percentile of the probability distribution. Thus, despite the fact that the latter mean remains finite for all $\alpha > 0$, it is clear that $E[L^*]$ is just as unreasonable – in a practical sense – as $E[L]$ as a basis for calculating the insurance company’s premiums.



Another way to compare the impact of $E[L]$ and $E[L^*]$ is to consider their associated effects on the ability of the insurance company to diversify risk through the law of large numbers. To that end, let

$$\psi = \Pr \left\{ W + nP - \sum_{i=1}^n L_i < 0 \right\}$$

denote the insurance company's one-period probability of ruin for a portfolio of n policyholders, where W is the company's initial net wealth, P is the per-policyholder premium, and the L_i are i.i.d. random variables having the distribution of either L or L^* . Setting $W = 10^{10}$ and $B = 10^{13}$, we use statistical simulation to compute the minimum values of P necessary to maintain $\psi \leq 0.01$ as the number of policyholders, n , increases. To yield informative plots, we choose α to be 1.5, 1.0, and 0.5 (i.e., equally spaced about the transition point $\alpha = 1.0$) and θ to have corresponding values of 10^9 , 10^8 , and 10^3 (to offset, somewhat, the effect of the increasingly heavy tail). Results for $L_i \sim F_L$ are presented in Figures 2(a)-(c), and those for $L_i \sim F_{L^*}$ in Figures 3(a)-(c).

For each value of n , the recorded P is a simple average from 1000 independent simulations. Given the intrinsic volatility of Pareto sums, we "borrow strength" across the independent simulations by plotting (in gray) a 21-point moving average of the P s for each n . This approximates the true continuity of the minimum premium as a function of the number of policyholders.

Figure 2(a): $\alpha = 1.5, \theta = 10^9, W = 10^{10}$




Figure 2(b): $\alpha = 1.0, \theta = 10^8, W = 10^{10}$




Figure 2(c): $\alpha = 0.5, \theta = 10^3, W = 10^{10}$

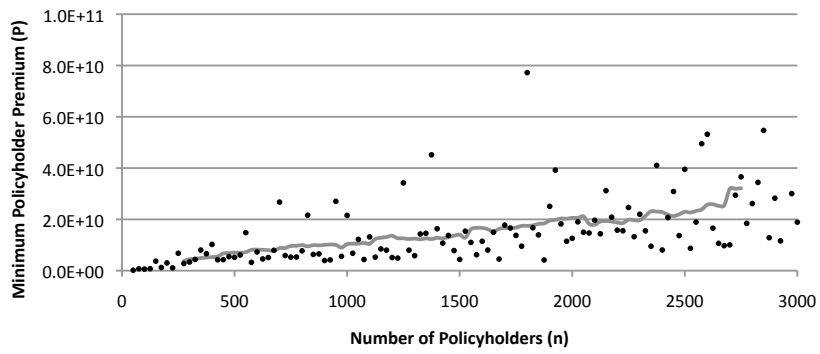

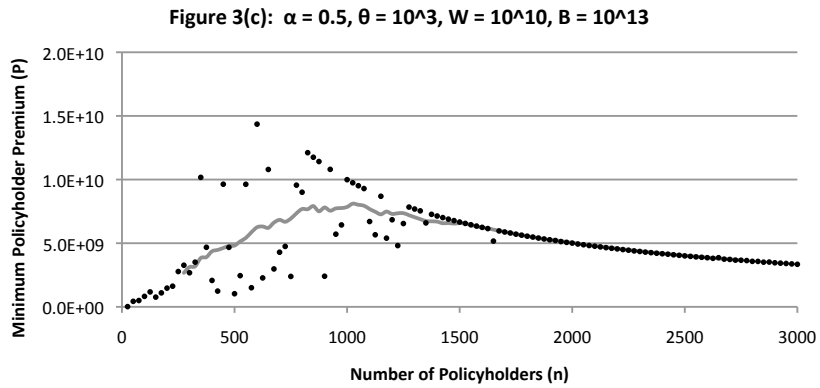
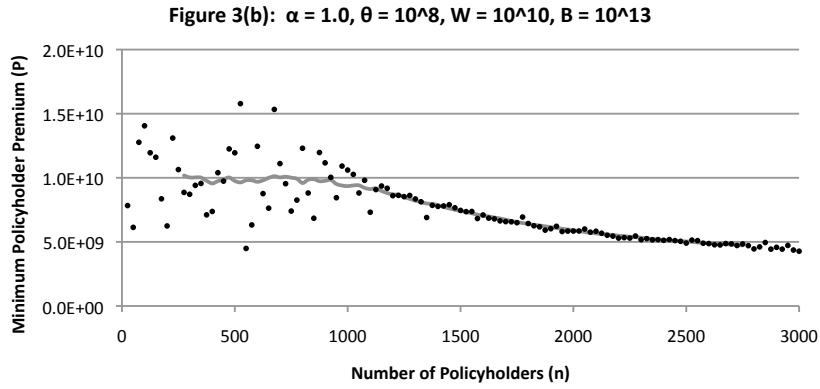


Figure 3(a): $\alpha = 1.5, \theta = 10^9, W = 10^{10}, B = 10^{13}$





Figures 2(a)-(c) show that: (a) when $\alpha = 1.5$ (so that $E[L]$ is finite), the minimum premium decreases toward $E[L]$ as n goes to infinity, indicating positive effects of diversification; (b) when $\alpha = 1$ (at the boundary of the region where $E[L]$ is infinite), the minimum premium remains flat over n , indicating no effects of diversification; and (c) when $\alpha = 0.5$ (so that $E[L]$ is infinite), the minimum premium increases without bound over n , indicating negative effects of diversification. These figures illustrate that an infinite mean violates the law of large numbers, thereby precluding any benefits of diversification.

Figures 3(a)-(c), based upon a value of $B = 10^{13}$, are similar to their counterparts in 2(a)-(c) for lower values of n , but all three show the minimum premiums decreasing to $E[L^*] < \infty$ as n goes to infinity. This occurs because the finite means ensure the applicability of the law of large numbers – and thus asymptotically positive effects of diversification – in all cases. The important point here is that, when $\alpha \leq 1$, there are no positive effects of diversification until n becomes rather large. Thus, practically speaking, an insurance company will experience no greater benefits of diversification from L^* than from L , despite the finite mean of the former random variable.

From the above analyses, we conclude that, for sufficiently large B , the apparently meaningful qualitative difference between $E[L] = \infty$ and $E[L^*] < \infty$ for $\alpha \leq 1$ is of no practical significance from the insurance company's perspective. Consequently, one can feel quite comfortable working with the infinite-mean random loss L , as an approximation to L^* .

3 The Insurance Two-Envelope Paradox and Its Resolutions

3.1 The Problem

The ITEP of Powers and Zanjani (2013) may be summarized as follows:

A risk-neutral insurance buyer faces a liability loss amount L such that $L|\mu \sim \text{exponential}(\mu)$, where the mean loss parameter, $\mu = E[L|\mu] > 0$, is initially an unknown random quantity such that $\mu^{-1} \sim \text{exponential}(\theta)$ and $E[\mu^{-1}] = \theta > 0$. The buyer is able to purchase insurance from a risk-neutral insurance company that determines the appropriate insurance premium by assessing the buyer's risk characteristics (i.e., by "underwriting" the buyer) to identify μ , and then computing the premium

$$\begin{aligned}\Pi_\mu &= E[L|\mu] + \pi \\ &= \mu + \pi,\end{aligned}$$

where $\pi > 0$ denotes an expense and profit loading.

It is assumed that, before deciding whether or not to purchase insurance, the buyer may or may not choose to assess his/her own risk, and thus has the option of identifying μ or letting the parameter remain unknown. Without identifying μ , the buyer knows that, with probability 1, Π_μ must be strictly less than the unconditional expected value,

$$E[L] = \infty$$

(because $L \sim \text{Pareto II}(\alpha = 1, \theta)$), and thus will purchase insurance. However, if the buyer identifies μ , then he/she will know that, with probability 1, Π_μ must be strictly greater than the conditional expected value,

$$E[L|\mu] = \mu,$$

and therefore will reject insurance. Paradoxically, the buyer's behavior is the same for all values of μ – implying there is no benefit from identifying the parameter – but learning the value of μ does in fact change the buyer's decision.

3.2 Two Distinct Resolutions

Powers and Zanjani (2013) stated that the “simplest” way to resolve the above paradox is to require that the buyer’s utility function be “sufficiently concave downward” that the conditional certainty equivalent

$$\rho_{L|\mu,W} : E[u(W - L)|\mu] = u(W - \rho_{L|\mu,W})$$

be such that

$$\rho_{L|\mu,W} > \Pi_\mu \tag{1}$$

for some continuous interval of μ . (If (1) does not hold for such an interval, then it constitutes an event “of probability zero” that has no practical impact on the buyer’s behavior.) However, they recognized that this approach works only if $u(\cdot)$ is not bounded below. If the utility function is bounded, then the paradox is resolved in a qualitatively different fashion. This conspicuous difference motivates the present analysis.

Relaxing the assumption that the insurance buyer is risk neutral, let his/her VNM utility function be some continuous, increasing function $u(\cdot)$ defined on the entire real number line, and let

$$\rho_{L|W} : E[u(W - L)] = u(W - \rho_{L|W})$$

denote the unconditional certainty equivalent. It then is easy to see that the ITEP persists for this generalization if and only if both of the following conditions are satisfied:

$$\rho_{L|W} = \infty \tag{2}$$

and

$$\rho_{L|\mu,W} < \Pi_\mu \tag{3}$$

for all μ . In other words, the unconditional certainty equivalent is strictly greater than Π_μ for all μ , but the conditional certainty equivalent is strictly less than Π_μ for all μ , thereby violating dominance reasoning.

3.2.1 Concave-Downward Utility

Given that $u(\cdot)$ is concave-downward everywhere – and thus not bounded below – (2) is always true. Therefore, to preserve dominance reasoning, we must impose (1). This approach is neatly illustrated by the case of constant absolute risk aversion (exponential utility), in which the buyer’s utility is

$$u(W - L) = \frac{1 - e^{-\beta(W-L)}}{\beta}$$

for $\beta > 0$ (with $u(\cdot)$ converging to linear utility as $\beta \rightarrow 0$).

By definition, the conditional certainty equivalent satisfies

$$\begin{aligned}
u(W - \rho_{L|\mu,W}) &= \int_0^\infty \frac{1 - e^{-\beta(W-\ell)}}{\beta} \frac{1}{\mu} e^{-\ell/\mu} d\ell \\
&= \begin{cases} \frac{1}{\beta} - \frac{e^{-\beta W}}{\beta(1-\beta\mu)}, & \beta < 1/\mu \\ -\infty, & \beta \geq 1/\mu \end{cases}.
\end{aligned} \tag{4}$$

Then, since

$$u^{-1}(y) = -\frac{1}{\beta} \ln(1 - \beta y),$$

it follows that

$$\begin{aligned}
\rho_{L|\mu,W} &= \begin{cases} -\frac{1}{\beta} \ln(1 - \beta\mu), & \beta < 1/\mu \\ \infty, & \beta \geq 1/\mu \end{cases} \\
&= \begin{cases} \mu + \frac{\beta\mu^2}{2} + \frac{\beta^2\mu^3}{3} + \frac{\beta^3\mu^4}{4} + \dots, & \beta < 1/\mu \\ \infty, & \beta \geq 1/\mu \end{cases},
\end{aligned} \tag{5}$$

and so (1) always holds if the risk-aversion parameter is sufficiently large; that is, either: (i) $\beta \geq 1/\mu$, or (ii) $\frac{\beta\mu^2}{2} + \frac{\beta^2\mu^3}{3} + \frac{\beta^3\mu^4}{4} + \dots > \pi$ if $\beta < 1/\mu$. These conditions correspond to the “sufficiently concave downward” approach of Powers and Zanjani (2013).

3.2.2 Bounded Utility

Now assume that $u(\cdot)$ is bounded below. In this case, it is easy to see that

$$E[u(W - L)] > -\infty,$$

and therefore that

$$\rho_{L|W} < \infty,$$

immediately contradicting (2).

4 Practical Implications

Considering the ease with which the ITEP is handled for bounded $u(\cdot)$, as well as the argument (based upon the SPP with monetary losses) that VNM utility functions must be bounded below, one might ask: Why worry about the ITEP in the first place?

The answer, of course, is that the boundedness of utility functions for large losses is rarely recognized by insurance researchers. Consequently, the “simplest” resolution of Powers and Zanjani (2013) fits comfortably within a research

literature that frequently employs the unbounded families of linear, exponential, logarithmic, and power utility functions to model real-world problems.

This leads to a second question: What are the costs of using unbounded utility functions in insurance modeling?

Sigmoid (i.e., “S-shaped”) utility functions, with a concave-upward region for lower values of wealth followed by a concave-downward region for higher values of wealth, have been known in the research literature since Friedman and Savage (1948). If bounded sigmoid functions are indeed more reasonable than linear, exponential, and other functional forms, then it is of both theoretical and practical interest to understand to what extent the choice of an unbounded utility function is critical to a buyer’s (1) willingness to purchase full insurance (cf. Mossin, 1968), and (2) preference for deductibles over other insurance contract forms (cf. Arrow, 1971). Certainly, it is easy to construct sigmoid $u(\cdot)$ for which a buyer is effectively risk prone (i.e., such that the inflection point between the concave-upward and concave-downward regions lies above the buyer’s initial wealth), so the results of Mossin (1968) and Arrow (1971) do not apply. However, it is unclear to what extent such utility functions arise in the real world.

Ultimately, the determination of a buyer’s certainty equivalent is an empirical issue that depends heavily on where the inflection point between concave-upward and concave-downward utility occurs. Nevertheless, it is possible to explore the sensitivity of the change in certainty equivalent to the introduction of unbounded utility through two variations of the loss model presented above.

For the original loss model of Section 3.1, we saw that the certainty equivalent changes dramatically when moving from a bounded utility function (for which $\rho_{L|W}$ is finite) to an unbounded exponential utility function (for which $\rho_{L|W}$ is infinite). However, this model involves an unconditional Pareto II loss random variable whose expected value is infinite, and therefore may be suspected of yielding somewhat extreme results. As an alternative, one can consider an example with a much lighter-tailed unconditional exponential loss.

To this end, let $L \sim \text{exponential}(\mu)$, where $\mu = E[L] > 0$ is known, and note that the unconditional certainty equivalent for exponential utility is given implicitly by

$$u(W - \rho_{L|W}) = \int_0^\infty \frac{1 - e^{-\beta(W-\ell)}}{\beta} \frac{1}{\mu} e^{-\ell/\mu} d\ell.$$

Since this is exactly the same expression as (4), except that $\rho_{L|\mu,W}$ has been replaced by $\rho_{L|W}$, it follows that the unconditional certainty equivalent in the present case must be identical to the right-hand side of (5), except that μ is now fixed. Consequently, and somewhat surprisingly, this means that, even for the light-tailed exponential loss, a change from bounded utility to unbounded exponential utility yields vastly different values of the certainty equivalent for sufficiently large β .

5 Conclusions

In the preceding analysis, we employed the ITEP to explore the disparate effects of bounded versus unbounded buyer utility in the presence of large monetary losses. We subsequently argued that the practical impact of using unbounded utility functions in insurance modeling has received inadequate attention in the research literature.

Although a simple theoretical example (using exponential losses) shows that differences implied by bounded and unbounded functions are potentially extreme, we recognize that a proper analysis of these differences requires empirical study. In further work, we will investigate experimentally how the location of the inflection point between concave-upward and concave-downward utility varies among decision makers in the real world.

References

- [1] Arrow, K.J. *Essays in the Theory of Risk Bearing*; North-Holland: Amsterdam, the Netherlands, 1971.
- [2] Bernoulli, D. *Specimen Theoriae Novae de Mensura Sortis*; 1738.
- [3] Broome, J. The Two-Envelope Paradox. *Analysis* **1995**, *55*, 1, 6-11.
- [4] Chalmers, D.J. The St. Petersburg Two-Envelope Paradox. *Analysis* **2002**, *62*, 2, 155-57.
- [5] Embrechts, P., Klüppelberg, C., and Mikosch, T. *Modelling Extremal Events for Insurance and Finance*; Springer-Verlag: Berlin, 1997.
- [6] Friedman, M. and Savage, L.J. The Utility Analysis of Choices Involving Risk. *Journal of Political Economy* **1948**, *56*, 4, 279-304.
- [7] Menger, K. Das Unsicherheitsmoment in der Wertlehre. *Journal of Economics* **1934**, *5*, 4, 459-485.
- [8] Mossin, J. Aspects of Rational Insurance Purchasing. *Journal of Political Economy* **1968**, *76*, 4, 553-568.
- [9] Nalebuff, B. The Other Person's Envelope Is Always Greener. *Journal of Economic Perspectives* **1989**, *3*, 1, 171-181.
- [10] Powers, M.R. Paradox-Proof Utility Functions for Heavy-Tailed Payoffs: Two Instructive Two-Envelope Problems. *Risks* **2015**, *3*, 1, 26-34.
- [11] Powers, M.R. and Zanjani, G. Insurance Risk, Risk Measures, and Capital Allocation: Navigating a Copernican Shift. *Annual Review of Financial Economics* **2013**, *5*, 201-223.

- [12] von Neumann, J. and Morgenstern, O. *Theory of Games and Economic Behavior*; 2nd Edition; Princeton University Press: Princeton, NJ, USA, 1947.