

A Two-Dimensional Risk Measure

Rick Gorvett, FCAS, MAAA, FRM, ARM, Ph.D.¹

Jeff Kinsey²

Presented at
Enterprise Risk Management Symposium
Society of Actuaries

Chicago, IL

April 23–26, 2006

Copyright 2006 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

¹ Rick Gorvett is State Farm Companies Foundation Scholar in Actuarial Science, and Director of the Actuarial Science Program, at the University of Illinois at Urbana-Champaign.

² Jeff Kinsey is an actuarial analyst at the State Farm Research Center in Champaign, IL, and a master's degree student in actuarial science at the University of Illinois at Urbana-Champaign.

Abstract

The measurement of “risk” is a critical component of enterprise risk management, insurance pricing and reserving, and strategic and operational decision-making. Some commonly used measures attempt to quantify risk by considering volatility or downside “potential.” Other measures examine ratios of upside versus downside potential, or upside potential versus volatility. A general characteristic of such measures is that they describe risk in a single dimension—placing risk along a number line and providing a single quantitative value. In this paper, we suggest that risk is too complex to quantify with a single number. We suggest a two-dimensional risk measure and introduce the concept of iso-risk curves.

“Risk is synonymous with uncertainty—lack of knowledge.”

- Irving Fisher, *The Theory of Interest*, 1930

“But Uncertainty must be taken in a sense radically distinct from the familiar notion of Risk, from which it has never been properly separated.”

- Frank Knight, *Risk, Uncertainty, and Profit*, 1921

1. Introduction

The two quotes above, separated by less than a decade, are exemplary of the difficulty and disagreements which even famous economists have had in defining the concept of “risk.” It is even more sobering to examine the primary works of other authors—e.g., Adam Smith, Alfred Marshall—and observe that there is often no effort at all to define the term³; rather, the word “risk” is frequently assumed as common knowledge, and invested with its popular meaning and connotations. Today, we often use “risk” and “uncertainty,” at least informally, somewhat interchangeably, while recognizing that there is a formal technical difference. In particular, as Knight (1921) acknowledges, in a continuation of the above quote:

The term “risk,” as loosely used in everyday speech and in economic discussion, really covers two things which, functionally at least, in their causal relations to the phenomena of economic organization, are categorically different.... The essential fact is that “risk” means in some cases a quantity susceptible of measurement, while at other times it is something distinctly not of this character.... It will appear that a *measurable* uncertainty, or “risk” proper, as we shall use the term, is so far different from an *unmeasurable* one that it is not in effect an uncertainty at all. We shall accordingly restrict the term “uncertainty” to cases of the non-quantit(at)ive type.

The whole idea of measuring risk is a critical concept for finance, economics and actuarial science⁴—and yet it can prove notoriously difficult. Over the years, a variety of different risk measures has been proposed, several of which have been implemented by various practitioners. Many other risk metrics have been proposed, considered and analyzed in the actuarial, risk management and/or financial literature, in both academic journals and institutional publications. Sometimes, it seems that the “optimal” risk measure may depend upon the characteristics and conditions of the particular situation—e.g., line of business, size, nature of the decision being made and sources of risk.

In this paper, we introduce and consider a new, two-dimensional risk measure. Various other measures attempt to quantify risk in terms of variability or downside “potential.” More elaborate measures created in the last few years have looked at ratios

³ Even in Knight’s book, a chapter entitled “The Meaning of Risk and Uncertainty,” has only three occurrences of the word “risk” in its text.

⁴ While these industries have spent substantial time and effort in trying to quantify risk, it may be that, given the nature and complexity of the world, we will also want to begin achieving breakthroughs in the consideration of *uncertainty* in the future.

of upside versus downside potential, or upside potential versus outcome variability. Generally, however, other measures have attempted to quantify risk in a single dimension: each measure places risk along a number line and represents risk via a single quantitative value. We believe that risk is too complex to be quantified with a single number. In this paper, we present a two-dimensional risk measure and introduce the concept of “iso-risk” curves.

The remainder of this paper is organized as follows. Section 2 introduces the two-dimensional risk measure, and describes its essential characteristics. Section 3 provides an illustration of the measure with respect to a hypothetical distribution. Section 4 considers the risk measure with respect to two pairs of illustrative distributions with different characteristics. Section 5 evaluates the relative ranking of three distributions among four risk measures—the measure introduced here and three others commonly used. Section 6 concludes, and briefly mentions directions for future research.

2. A Two-Dimensional Risk Measure

One of the problems associated with a “one-dimensional” risk measure, or any measure that does not have a firm economic and/or mathematical foundation, is that it might not accurately represent the true “riskiness” of outcomes. For an actuary and/or an insurance company, the potential consequences of such a misrepresentation include charging an inappropriate price for an insurance or financial product, not properly estimating required return on equity and even making sub-optimal operational or strategic planning decisions. Often, risk measures do not fulfill their intended purpose because they fail to adequately reflect the characteristics of the potential distribution of outcomes. For example, it is fairly simple to hypothesize a reasonable situation in which traditional measures of risk suggest two distributions are “equivalent” (say, because they have the same standard deviation or value at risk), when in fact one should be considered *more* risky than the other due to its skewness or kurtosis.

The two-dimensional risk measure introduced here is based upon the relationship between both upside and downside conditional expected values. For a particular distribution, and with respect to a selected dollar threshold, the conditional expected downside and upside values are calculated, and the difference between the values is taken. This difference is calculated for all possible thresholds in the range of potential outcomes. The result is an array of upside and downside values for various thresholds, along with their differences. The focus of this risk measure is to find the threshold where the distance between upside and downside is minimized. This

threshold value, along with the corresponding value of the difference between expected upside and downside, represents the two dimensions of the risk measure.⁵

At this point, it is helpful to define a few terms that will be used throughout the remainder of this paper. For a particular outcome distribution, we have the following parameters:

- ***Threshold***: The dollar value at which downside and upside conditional expected values are calculated.
- ***Upside***: The conditional expected value of all outcomes greater than the threshold.
- ***Downside***: The conditional expected value of all outcomes less than the threshold.
- ***Spread***: The difference between the upside and downside values at a particular threshold.
- ***Minimum Spread***: The minimum of all spread values calculated.

These parameters lead to the two dimensions of this risk measure, which can be identified as an ordered pair:

(Minimum Spread, Threshold @ Minimum Spread)

or, notationally,

(s, t) .

This approach has certain similarities to the return-volatility tradeoff in the Capital Asset Pricing Model (CAPM). In that context, an investor desires to maximize expected return for a given level of volatility (or variability or “risk”), or equivalently, minimize volatility for a given level of expected return. An analogous approach applies to this risk measure. In comparing various distributions, it is preferred to have a higher optimal threshold value t for a given value of minimum spread s , and similarly, a smaller s for a given t .

⁵ Among existing risk metrics, this measure is most similar to the Omega Ratio presented by Keating and Shadwick (2002). In the Omega Ratio, the upside and downside expected value relationship involves the *ratio* of upside to downside, as opposed to the *difference* as used in the current measure.

The two dimensions shown above were selected for various reasons. First, the minimum spread value occurs at a threshold where the rate of change of the upside and downside values are equal. This can be seen more clearly in the illustrative example in the following section. The concept of equal rates of change is analogous to many marginal equality concepts found in economics. In addition, it is also very important to include where this minimum value occurs. The threshold value where the minimum occurs will change depending on the mean, variance and skewness of the underlying distribution. That being the case, the threshold value contains various characteristics of the first three central moments of the underlying distribution.

Within $s-t$ space,⁶ this single point does not lie in isolation. It lies on a curve of equally risky points. These curves, which we have named “iso-risk” curves, are positive-sloping, non-overlapping curves that fill the entire $s-t$ space.⁷ These iso-risk curves will vary in slope and curvature depending on particular risk appetites. People (or organizations, if one wishes to informally endow them with risk-preference traits) that are more willing to accept risk will have a flatter iso-risk curve at a particular point than those that are more risk-averse. Likewise, the slope of an iso-risk curve will be greater for those that are risk-averse as opposed to those more willing to accept risk.

Analogously to preferring to be on a higher utility curve in economic theory, so too is it preferable to be on a higher (in the northwest direction of the $s-t$ plane) iso-risk curve. In particular, points lying on higher iso-risk curves are perceived to be less risky than those lying on lower iso-risk curves.

3. An Illustrative Example

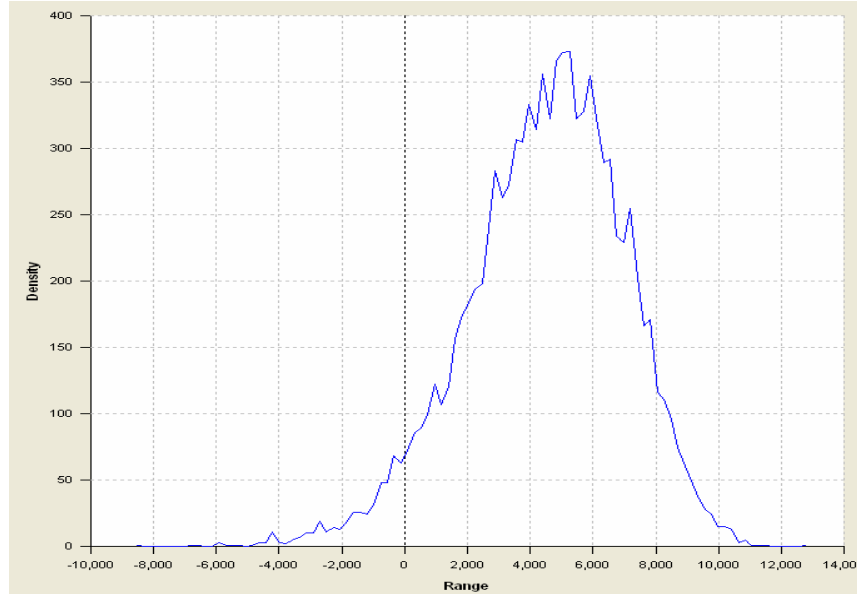
The following example serves to illustrate this risk measure. Figure 1 shows a density plot of a random sample of 10,000 values drawn from a gamma (4.5; 1,000) distribution. Each sample was then transformed in order to achieve the desired skew and scaling according to the formula:

$$[\ln(\text{sample value}) \times 5,000] - 37,000$$

⁶ With the exception of Figures 3, 8 and 11 (which are reversed for clarity and pedagogic purposes), we will adopt the convention that t is plotted along the vertical (or y-) axis, and s is plotted along the horizontal (or x-) axis.

⁷ The obvious analogy is to utility curves in either risk-return or two-consumer-good space.

Figure 1



Upside and downside were both calculated starting at an initial threshold value of -8,100. The threshold was then incremented by +100 up to the value 11,600; at each increment, both the upside and downside expected values were calculated. The formulas for the upside and downside calculations are as follows:

N_u = Number of values greater than the threshold

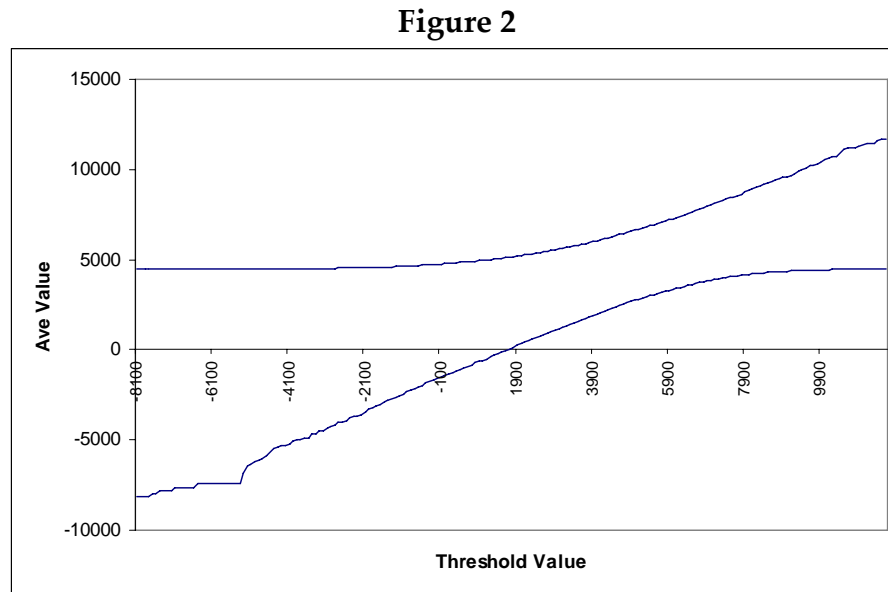
N_d = Number of values less than the threshold

U_i = full amount of a simulated value which is greater than the threshold

D_i = full amount of a simulated value which is less than the threshold

$$Upside = \frac{1}{N_u} \sum_{i=1}^{N_u} U_i \qquad Downside = \frac{1}{N_d} \sum_{i=1}^{N_d} D_i$$

Figure 2 shows a plot of upside and downside values at each threshold point.




This graph in and of itself has some appealing characteristics. First, the two lines are horizontally asymptotic to the mean of the distribution. Second, the vertical distance between the two lines (the spread) gives a relative measure of variability; this will be better illustrated in a later example. Finally, the relative position where the spread is minimized describes the skewness of the distribution. For a normal distribution, spread will be minimized at the 50th percentile, whereas spread for a left-skewed distribution will be minimized at a point greater than the 50th percentile, and for a right-skewed distribution at a point less than the 50th percentile.⁸

⁸ A bit of reflection will suggest that the horizontal point (the threshold value) at which the vertical distance between the upward and downward expected value curves in Figure 2 is minimized, is also the point at which the two curves have the same slope (assuming certain regularity conditions—i.e., assuming that the curves behave “normally”). This is because, if at the minimum-spread threshold the curves do *not* have identical slope, then either immediately to the left or right of that threshold the spread must be smaller (one curve will be “gaining” on the other). This possibility results in a contradiction. So minimum-spread and equality-of-slope occur at the same point. This is analogous to several “marginal equality” concepts in economics—e.g., a consumer will split her/his resources between two consumable goods according to the point at which the marginal benefit provided by each good is the same (otherwise, s/he would favor consuming more of the higher-marginal-benefit good).

Figure 3 shows spread versus the threshold value.


Figure 3



The minimum spread value and the corresponding threshold are the two values of concern for this risk measure. As shown in Figure 3, for this distribution spread is minimized at a threshold value of 5,300 with a minimum spread value of 3,881.


Figure 4 plots this point with the minimum spread value s on the horizontal axis and the threshold at which the minimum spread occurs plotted on the vertical axis (the convention which will be used throughout the remainder of the paper).

Figure 4




The theory underlying this risk measure is that this point is one of a series of points of equal risk that lie on an iso-risk curve. For a particular distribution, the points on a given iso-risk curve are considered to be of equal risk relative to each other (from the standpoint of a person or other entity with a given risk appetite), whereas points above the iso-risk curve are considered less risky and points below the curve are more risky compared to that particular distribution. Figure 5 illustrates this concept.

Figure 5



Conceptually, iso-risk curves can be thought of in two ways. First, the slope of the iso-risk curve can be thought of as how much an entity with a particular risk appetite is willing to sacrifice in the threshold value for a unit increase in certainty (here, a unit increase in certainty would be a unit decrease in spread). Secondly, the vertical axis intercept of the iso-risk curve has an interesting interpretation: the intercept value on the threshold axis indicates the threshold value that would be accepted for absolute certainty (zero spread) of that particular result. Using a simplifying assumption that the iso-risk curves are linear, Figure 6 shows how the slopes of iso-risk curves will change according to varying risk appetites.

Figure 6



One can see that, in exchange for reducing spread (or risk), a risk-averse entity will be willing to accept a smaller (possibly even negative) threshold value compared with a less risk-averse (or more risk-accepting) entity.

4. Comparing Distributions

Often in practice, two or more distributions will be compared, evaluated and even “ranked” according to a particular measure of risk. In this section, we illustrate how the $s-t$ risk measure introduced here behaves when comparing two or more simple distributions.

Example 1: Two Distributions with the Same Mean

The purpose of this first comparison is to show how this risk measure performs with two distributions that have the same mean but different variances. For this comparison, a random sample of 10,000 values was drawn from each of the following distributions:

Normal (0; 1,000)

Normal (0; 10,000)

Few would argue with the suggestion that the $N(0; 10,000)$ is more risky. However, it is interesting to see how the graphical $s-t$ framework illustrates the differences in these two distributions. Figures 7 through 9 show this comparison.

Figure 7




Figure 8


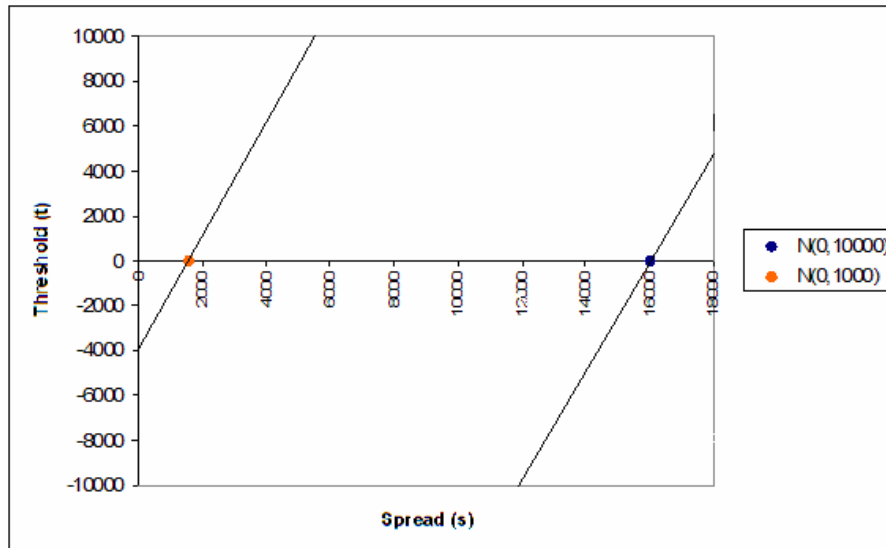


Figure 9



Maintaining the simplifying assumption that iso-risk curves are linear, the $N(0; 1,000)$ distribution is clearly shown to be less risky in Figure 9—it is on a “higher,” or more-left, iso-risk curve. In fact, only a *horizontal* iso-risk curve, which would represent an infinite risk tolerance and fictitious in practice, would place these two distributions as equally preferable in terms of risk. It is also interesting to note the values of minimum spread for each of these distributions. For the $N(0; 1,000)$ distribution, the minimum spread was achieved at a threshold of 0 with a spread value of 1,593. The $N(0; 10,000)$ distribution had a minimum spread at a threshold of 0 with a spread value of 16,042 which is roughly ten times the size of the $N(0; 1,000)$ minimum spread value. It is not coincidental that we began with two distributions in which one had a standard deviation 10 times the other.

Example 2: Two Distributions with the Same Standard Deviation

As a second comparison, a random sample of 10,000 values was drawn from each of the following two distributions:

Normal (-10,000; 1,000)

Normal (10,000; 1,000)

Again, few would argue that the $N(-10000,1000)$ is more risky. For this example, charts analogous to the first example are displayed as Figures 10 through 12.

Figure 10

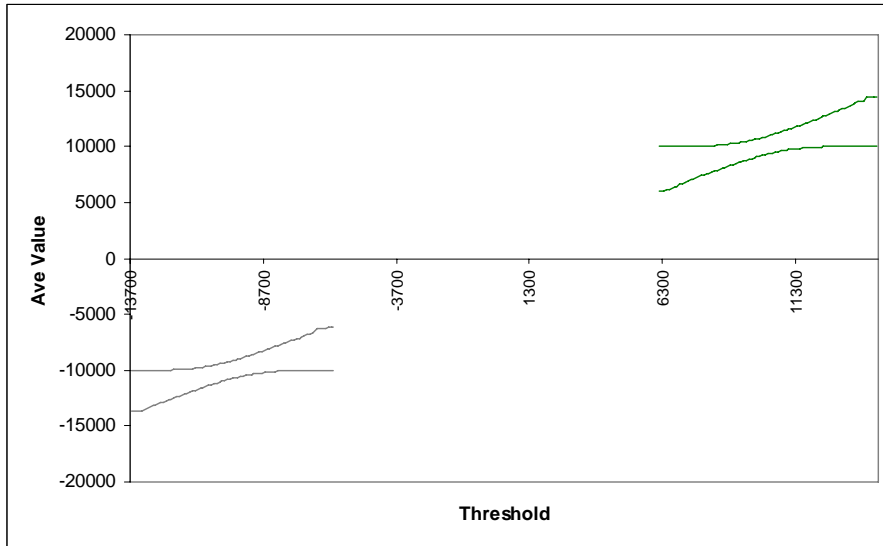


Figure 11

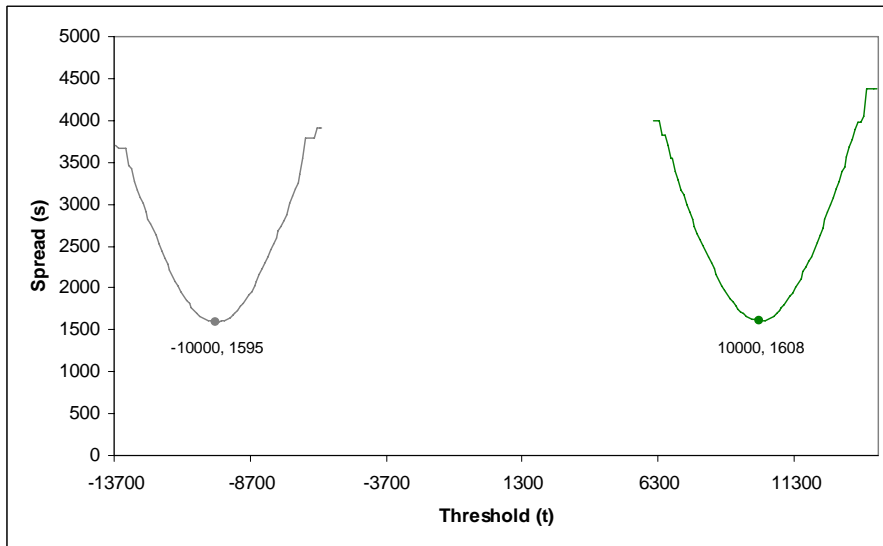
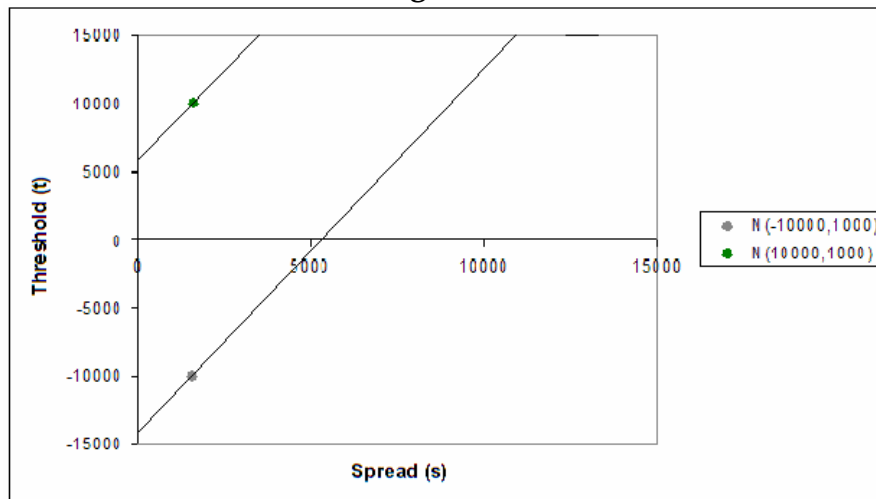


Figure 12



Maintaining the simplifying assumption of linear iso-risk curves, the $N(10,000; 1,000)$ will be on a higher curve in all cases except that of a vertical iso-risk curve. A vertical iso-risk curve would indicate a focus solely on variability reduction.

5. Comparing Risk Measures


The $s-t$ risk measure will now be compared to some of the more popular risk measures currently used in practice. In particular, comparisons will be made to value at risk (VaR), tail value at risk (TVaR) and standard deviation (SD).

For this comparison, a random sample of 10,000 values was drawn from each of the following distributions:

- A. $\text{Lognormal}(0; 0.5) * 3000 - 8000$ (Blue)
- B. $\text{Lognormal}(0; 0.3) * 3000 - 4500$ (Green)
- C. $\text{Normal}(0; 1,500)$ (Red)

Figure 13 is a density plot of the three distributions.

Figure 13

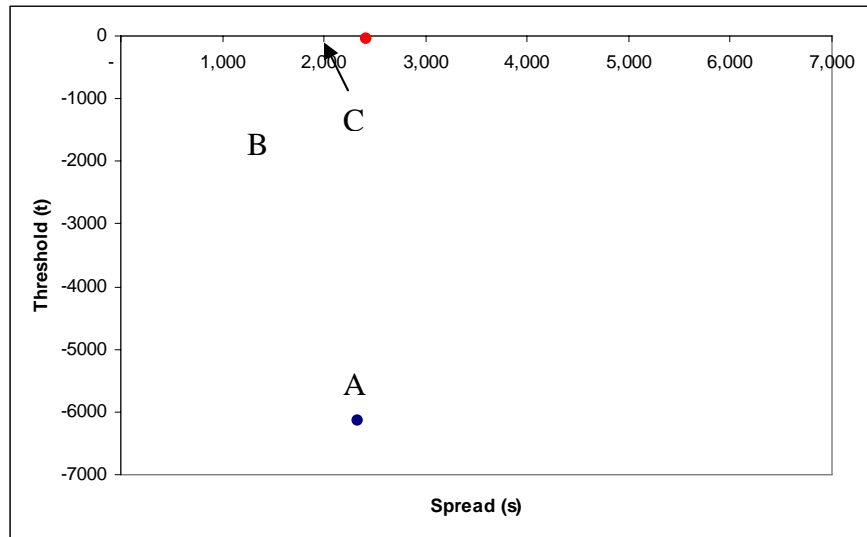


Based on just visual inspection of the density plots, one could possibly make an assessment as to the relative riskiness of these three distributions. Few would argue with the judgment that distribution A (the left-most mode) is the riskiest of the three. However, legitimate arguments could be made as to why either distribution B or C might be considered the least risky distribution. Although distribution C has a higher average value, it is more variable than distribution B and has a higher probability of outcomes less than (approximately) -3,500. However, it also has significantly higher potential for outcomes greater than zero. Distribution B has less variability, but also is likely to be negative a majority of the time. Depending on a particular risk appetite, either distribution could be viewed as the least risky. VaR, TVaR and SD rank these distributions in the following orders (1 = least risky, 3 = most risky):

| Risk Measure | Distribution A | Distribution B | Distribution C |
|---------------------|-----------------------|-----------------------|-----------------------|
| VaR (.01) | 3 | 1 | 2 |
| TVaR (.01) | 3 | 1 | 2 |
| Std. Dev. | 3 | 1 | 2 |

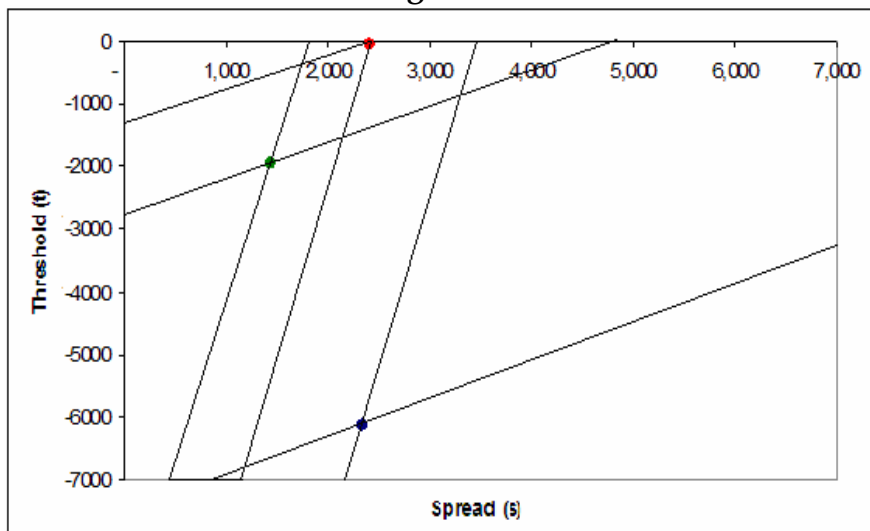
Each risk measure has ranked distribution A as the riskiest and distribution B as the least risky. Now let's examine the *s-t* risk measure's indications. Figure 14 shows how each of these distributions appears when the minimum spread and the corresponding threshold are plotted.

Figure 14



The iso-risk curve through each point will vary depending on particular risk appetites. For example, a risk-averse appetite would shy away from volatile distributions and therefore would have a steeper iso-risk slope, and thus distribution B would be considered less risky than distribution C. In contrast, a risk appetite that has a higher tolerance for risk would be more accepting of a volatile distribution as long as higher outcomes are more probable and would have flatter iso-risk curves, thus assessing distribution C as less risky than B. There also exists a unique iso-risk curve where these two points are considered equally risky. Figure 15 shows these various iso-risk curves plotted on a single graph. Again, the simplifying assumption of linear iso-risk curves is used.

Figure 15



6. Conclusion and Future Research

In this paper, we have presented a two-dimensional risk measure. It is suggested that this measure reflects several moments of the distribution of possible outcomes, and thus has the potential to be a useful and effective tool in considering the relative preferences of different distributions.

This is a young idea, and there are numerous paths in which further research can and will be conducted. Much of that research will involve potential applications of this risk measure concept, including:

- Determination of appropriate reinsurance prices (via the spread-threshold “tradeoff” inherent in iso-risk curves).
- Capital allocation decisions (via relevant distances between iso-risk curves).
- Strategic and operational decision-making (based on hypothesized distributions of outcomes before and after various alternatives).

Another avenue for future exploration is to consider various formulas for the iso-risk curves, which would permit parametric expression and analysis of different risk appetites.

It is hoped that this paper will inspire additional research and comments in these and related areas.

References

Fisher, I. 1930. *The Theory of Interest*, New York: The Macmillan Company.

Keating, C., and Shadwick, W. 2002. An introduction to Omega. The Financial Development Centre Limited

Knight, F. 1921. *Risk Uncertainty, and Profit*. Boston: Houghton Mifflin Co.