Risk Analysis and Mitigation Tendencies as Interpreted by Cognitive Science

Eric D. Smith
RIMES: Research Institute for Manufacturing and Engineering Systems
IMSE: Industrial, Manufacturing & Systems Engineering Department
University of Texas at El Paso
El Paso, TX 79968
ESmith2@UTEP.edu
(915)747-5205

ABSTRACT
Risk matrices used in industry characterize particular risks in terms of the likelihood of occurrence, and the consequence of the actualized risk. Human cognitive bias research led by Daniel Kahneman and Amos Tversky exposed systematic translations of objective probability and value as judged by human subjects. Applying these translations to the risk matrix allows the formation of statistical hypotheses of risk point placement biases. Industry-generated risk matrix data reveals evidence of biases in the judgment of likelihood and consequence -- principally, likelihood centering, a systematic increase in consequence, and a diagonal bias. Evidence presented could improve risk matrix based risk analysis prevalent in industry.

Key words: risk analysis, risk matrix, cognitive biases, utility function, subjective probability

Sessions:
Systems Engineering: The Systems Engineering of Test and Evaluation

INTRODUCTION

This paper presents a novel cross-disciplinary examination of the risk matrix, combining systems engineering and cognitive psychology, with statistical analysis. As such, there is no existing literature describing a similar statistical study of risk matrix data. As this research began, the authors had disjointed preliminary information pertaining to the conclusions of this paper. Parts of the necessary knowledge, experience and information existed in different forms in the three authors, but was inchoate, and as such, non-conclusive. The first author had knowledge of cognitive biases as summarized in prospect theory (explained later in this paper). The second author had empirical and experiential experience with industrial use of the risk cube. The third author had extensive statistical experience with manufacturing reliability and statistical analysis. All authors had education and experience in engineering as a base.

This paper describes how cognitive biases affect the placement of risk points within a risk matrix when engineers in industry subjectively judge the likelihood of the risk, and, separately, the consequence of a risk. The risk matrix appears in the INCOSE Systems Engineering Handbook [INCOSE, 2006, p. 7.13], and is used widely in industry. Cognitive psychology has produced pertinent summary results for the subjective translation of probability and value that bear directly on the subject of risk assessment and mitigation as practiced with the aid of the risk matrix. Risk managers will benefit from a better understanding of subjective judgments made within the context of the risk matrix.

Terminology: In the industry data analyzed for this paper, the two, mostly human-estimated, parameters of the risk matrix data are likelihood (L) and consequence (C). The cognitive bias literature uses the terms subjective probability ($\pi(p)$) and utility (U), respectively. In this paper, the possible distinctions between these sets of terms will not be developed, but a strong distinction will be drawn against objective probability (p) and objective value (v). Table 1 summarizes the relevant terms.

Subjective Parameters	
Likelihood (L)	Consequence (C)
Subjective Probability, $\pi(p)$	Utility (negative), $U^{-}(v)$
Shown on:	
Ordinate, Y axis	Abscissa, X axis
Objective Parameters	
Objective Probability, p	Objective Value, v

Table 1: Subjective and objective terms.

Risk is usually computed as the product of Likelihood and Consequence:

$$R = L \times C.$$

Usually, likelihood is scaled from 0-1, as is probability, but consequence may be scaled differently.

Risk Matrices in Industry

The most common tool used to track and manage risks is the risk matrix, an example of which is shown in Figure 1, where the high, medium and low risk zones are formed by rectangular hyperbola iso-risk lines, along which the product of likelihood and consequence is constant. In practice, the shape and number of risk zones are arbitrary, and determined by individual companies, according to their particular needs. Allowing the zones to be progressively fuzzy eventually leads to continuous likelihood and consequence numbers. Contrariwise, an increase in discretization results in more risk zones with integer likelihood and consequence designations.

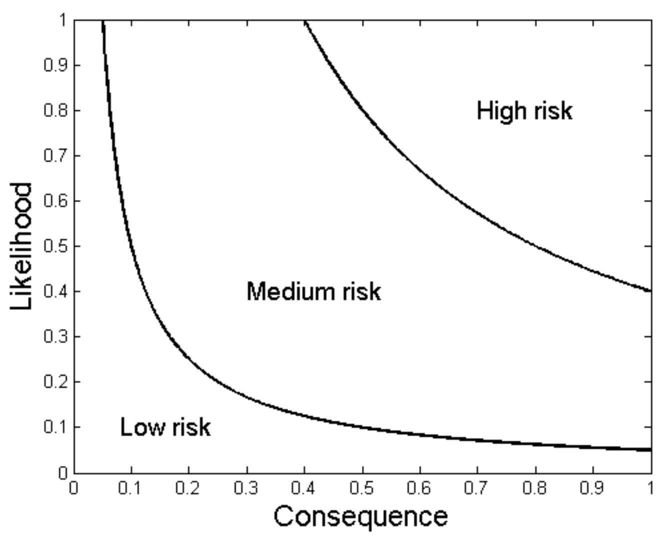

Figure 1: Risk matrix with rectangular hyperbola zones.

Risks are often placed as points inside the risk matrix, and are moved to the lower left as risk likelihood and consequence are reduced. The risk matrix provides a visual communication tool, around which engineers can discuss risk mitigation progress.

The risk matrix in Figure 1 has the weaknesses that a lack of granularity results in the lack of a higher-confidence means of accurately locating and re-locating subjectively assessed risks. The continuous scale in Figure 1 may result in un-resolvable discussion about the location of risk points. Risk managers need risk matrices with sufficient, but not excessive, granularity that aids in the timely placement of risk points, and in risk point re-placement toward the lower left after successful risk mitigation, without encouraging micro-management of risk point placement.

The risk matrix shown in Figure 2 represents an approach that imposes integer determination of likelihood and consequence. Five bins along each axis form a 25-square risk matrix with integer assignment of 1-5 for both likelihood and consequence, providing granularity and reducing discussion in deciding which position is correct. This risk matrix, along with quantitative and qualitative definitions of the ranges, represents the industry standard in major aerospace companies for the categorization of risks. NASA uses this risk matrix, such that only 3 risk levels are available for review for the Ares I project [GAO, 2007].

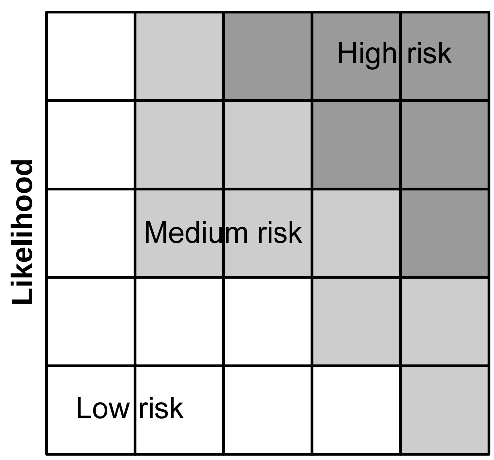

Figure 2: 5x5 risk matrix.

Defined objective criteria, or qualitative descriptions, can be written for each of the 5 bins of consequence and likelihood in the risk matrix. Properly defined criteria add stability to risk point placement, increasing the value of the risk matrix. Table 2 lists a set of qualitative definitions for the five ranges of likelihood and consequence, in this instance, specifically for technology. These are the actually descriptions that guided technology-related risk point placement for the data examined in this paper. Any such definitions should be tailored to company requirements.

Qualitative scales for likelihood and consequence of technology risks				
[L] Likelihood of failure – technology dependence				
Low [L = 1]	Minor [2]	Moderate [3]	Significant [4]	High [5]
No new technology –	Minor modification of	Dependant on innovative use	Dependant on new	Dependant on new

Systems are off the shelf	existing technology	of existing technologies	technologies that are in development	technologies that are not yet funded
[C] Consequence of failure – technical				
Low [C = 1]	Minor [2]	Moderate [3]	Significant [4]	High [5]
Little or no impact on program objectives	Minor reduction in technical performance with little or no impact on program objectives	Reduction in technical performance with limited impact on program objectives	Significant degradation in technical performance with a major impact on program objectives	Major degradation in technical performance that could jeopardize program success

Table 2: Qualitative description for five ranges of likelihood and consequence for technology.

Increased objectivity in the bin definitions provides more confidence in the accuracy of risk point placement. Likelihood determination can be increasingly objective with historical data of frequency of events, while consequence can be increasingly objective with historical data on the magnitude of consequences. Table 3 summarizes how Mil-Std 882d provides more objective descriptions of consequence categories, although only 4 categories are defined [Mil-Std-882d, 2000, p.18]. Note how definitions of consequences necessarily involve ranges, and not point values of damages, because damages are condition and mode dependant. Point values for consequence are a current weakness of risk matrixes.

	$ damage	**Human Impact**	**Environment**	**Law or regulation**
Catastrophic	> $1M	Death, Disability	Irreversible damage	Violate
Critical	$1M - $200K	Hospitalization to >= 3 personnel	Reversible damage	Violate
Marginal	$200K-$10K	Loss of work days by injury	Mitigatible Damage	
Negligible	$10K-$2K	No lost work day by injury	Minimal damage	

Table 3: Mil-Std 882d increasingly objective descriptions for consequence (C).

Perfect objectivity in the determination of likelihood and consequence, though expensive and largely impractical because of the experimentation and testing required, would eliminate the inaccuracies present in qualitative descriptions, and would eliminate the cognitive biases revealed below in the present data. Mil-Std 1629a describes how an increasingly structured and experimentally verified Failure Mode, Effects and Criticality Analysis (FMECA) can provide a virtually objective system risk analysis [Mil-Std-1629a, 1980].

Cognitive Biases in Probability and Value

The most accepted descriptive theory of subjective expected value decision making by humans is Prospect Theory, developed by Amos Tversky and Daniel Kahneman [Kahneman and Tversky, 1979; Tversky and Kahneman, 1992]. Kahneman won the Nobel Prize in Economics in 2002 "for having integrated insights from psychological research into economic science, especially concerning human judgment and decision-making under uncertainty" [Nobel, 2002]. (Since Nobel prizes are not bestowed posthumously, Tversky did not share this prize.) The contention of the Kahneman and Tversky school of "heuristics and biases" is that the presence of cognitive biases – even in extensive and vetted analyses – can never be ruled out. Innate human biases, and exterior circumstances, such as the framing or context of a question, can compromise estimates, judgments and decisions. It is important to note that subjects often maintain a strong sense that they are acting rationally while exhibiting biases [Piattelli-Palmarini, 1994].

Prospect Theory describes the subjective human decision-making process, specifically in the subjective assessment of probabilities and values, and their combination in gambles. Prospect Theory breaks subjective decision making into a preliminary 'screening' stage, and a secondary 'evaluation' stage. In the screening stage, probabilities and values are subjectively assessed, while the evaluation stage combines the subjective probabilities and utilities (subjective values) in the form of gambles. Gambles can be interpreted as multiple risks (L x C) added together. Only the subjective assessment of probabilities and values is of interest in this paper.

Subjective Probability, π(p)

Prospect Theory describes the subjective evaluation of probabilities, $\pi(p)$, according to the experimentally-obtained curve in Figure 3 [Tversky and Kahneman, 1992]. This non-linear transformation, $\pi(p)$, shows how small probabilities are overestimated and large probabilities are underestimated.

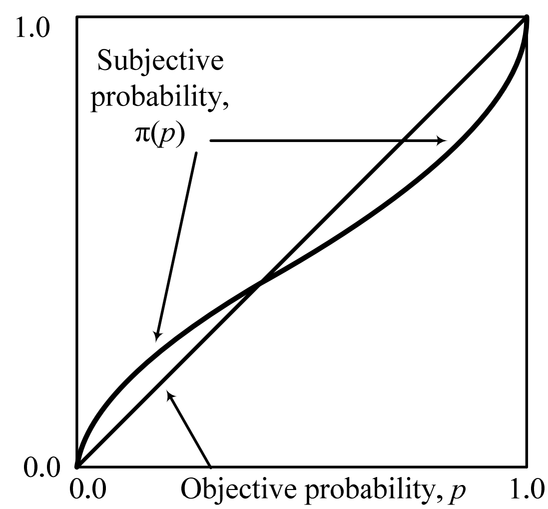

Figure 3: Subjective probability $\pi(p)$.

The modeling equation for subjective probability is [Tversky and Kahneman, 1992].

$$\pi(p) = (p^{\delta}) / [p^{\delta} + (1 - p)^{\delta}]^{(1/\delta)}$$

where p = objective probability, with $0 < \delta \leq 1$. When $\delta = 1$, $\pi(p) = p$ = objective probability. A usual value for δ is $\delta = 0.69$ for losses (and $\delta = 0.61$ for gains). This subjective probability curve can be slightly different not only for losses as compared to gains, but also for different people or industries; however, the shape of the curve is robust.

Subjective Utility

One of the effects of Prospect Theory's screening stage is that values are considered not in an absolute sense (from zero), but subjectively from the reference point established by the subject's monetary state and perspective before a decision. This is an example of the psychological phenomenon called framing. The key graph that shows how objective values translate into subjective utilities is shown in Figure 4. Note the significant disparity in magnitude with which gains and losses are subjectively valued – approximately a 1-to-3 ratio.

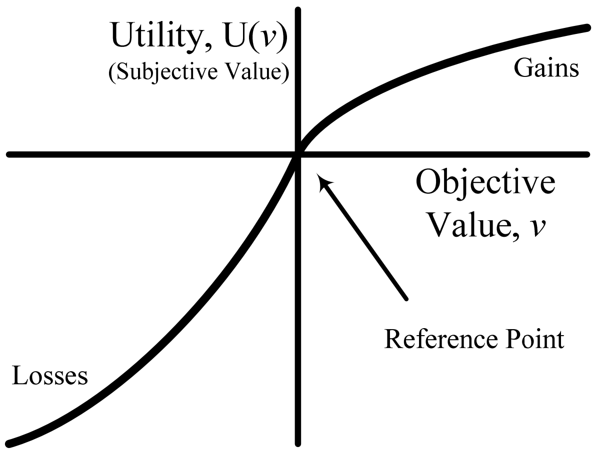

Figure 4: Utility (subjective value) versus objective value according to Prospect Theory.

For gains, $0 \leq v \leq \infty$, the utility function may be: $U^{+}(v) = \mathrm{Ln}(1 + v)$, where v = objective value. While for losses, $-\infty \leq v \leq 0$, the negative utility function may be: $U^{-}(v) = -(\mu)\mathrm{Ln}(1 - cv)$, with $\mu \approx 3$ and c = 1. This negative utility function provides a basis for analysis later in this paper.

Theoretical Implication of Prospect Theory for the Risk Matrix

Figure 5 shows primary hypotheses of probability and value biases for the risk matrix.

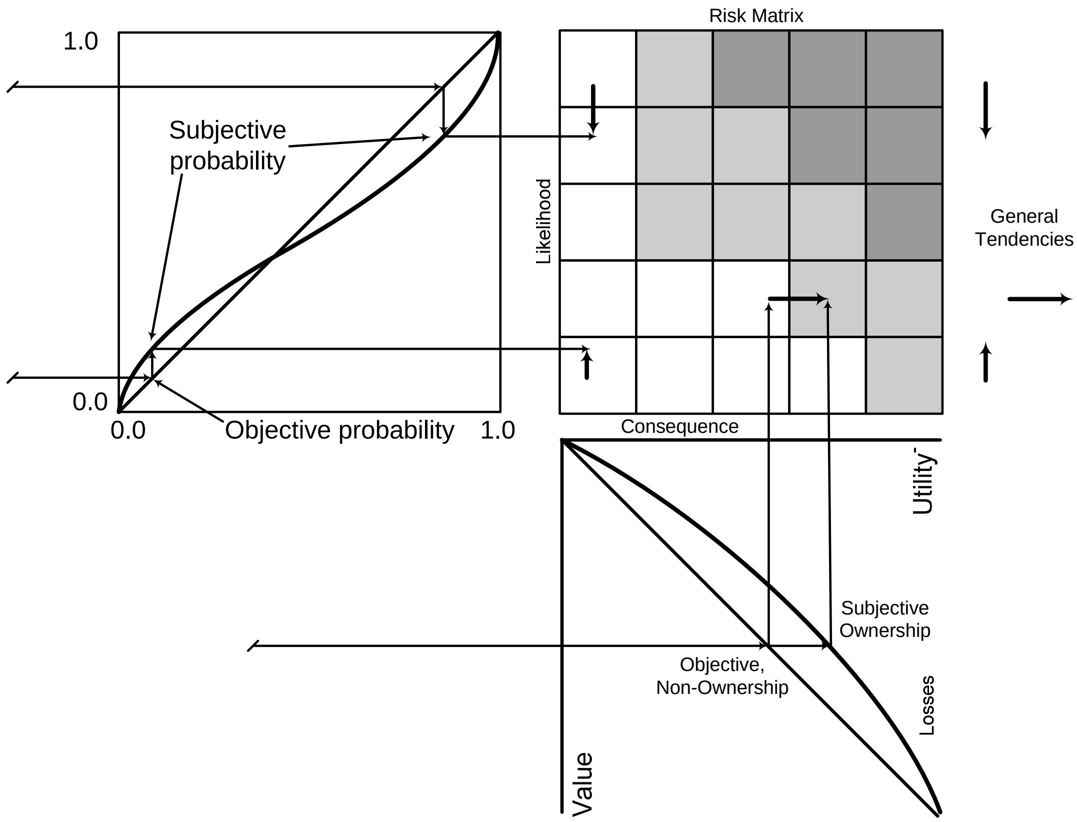

Figure 5: Theoretical implications of prospect theory for the risk matrix.

The most obvious influence of non-symmetric subjective probability assessments will be to compress probability judgments inward toward the risk matrix likelihood of L = 3. Any non-symmetric effect may be difficult to observe.

For the subjective assessment of consequence values, note that only the loss section of the objective value-to-utility (subjective value) graph has been employed, by rotating it counter-clockwise and placing it under the risk matrix. The reasoning for using only half of the utility curve is that risks deal with possible losses, and not with opportunities for gain. The most likely hypothesis arising from the loss section of the utility curve follows.

Human subjects exaggerate the influences of losses when the losses will occur to their personal wealth. In industry, the utility curve predicts that managers with an increased sense of corporate ownership will maintain a heightened awareness as to the monetary loss and corporate destabilization that a risk could cause. In fact, given the limits of any corporation, this managerial view is a correct reflection of corporate vulnerability. Line engineers with less corporate buy-in, and a broader job market, will tend to see, in greater proportion, the objective nature of the consequences of risks, rather than experience the personal feeling of losing a personal enterprise. This is not to say that engineers do not understand the full consequences of failure, only that they are not in managerial positions where failure usually causes personal impacts via the loss of financial bonuses. This hypothesis will be explored given the evidence in the data.

INDUSTRY CASE STUDY DATA

Airplane industry risk-matrix data were analyzed to test for the presence of the predicted theoretical implications. The data, available as "Original" and "Current" risk points for either one of the three categories of technical, schedule and cost, were collected from two programs within one company: 1,412 database entries from one airplane program, and 665 database entries from a second airplane program. The 1,412 point database forms the basis of analysis throughout the rest of this paper, unless otherwise noted. "Original" refers to the first determination of likelihood and consequence numbers, while "Current" refers to the updated likelihood and consequence numbers. Proprietary information was eliminated from the data. A sample of the scrubbed raw data is given in Appendix A. The complete data set may be small in absolute terms, but it is the first and largest risk matrix data set yet examined for cognitive biases.

The data were collected in the airplane industry over a period of eight years for the purposes of risk assessment, tracking and management. The engineers who entered the data after subjectively estimating likelihood and consequence did not know that the data would ever be used to test for the presence of cognitive biases. A risk manager working with the database estimated that almost none of the data were derived from numerically objective sources. Because individual risk updates occurred at various times -- sometimes months to years after the original datum point determination -- and only the latest update or access time was kept in the database, this study was severely limited in its investigation of time-dependent risk mitigation. Figure 6 shows the original data of 1,412 points. Datum counts at each (C, L) location are proportional to the bubble areas. For example, there are 339 risk points originally placed at (C, L) = (4, 3), while there are only 36 risk points placed at (C, L) = (5, 4). Technical, schedule and cost risk are all combined in Figure 6.

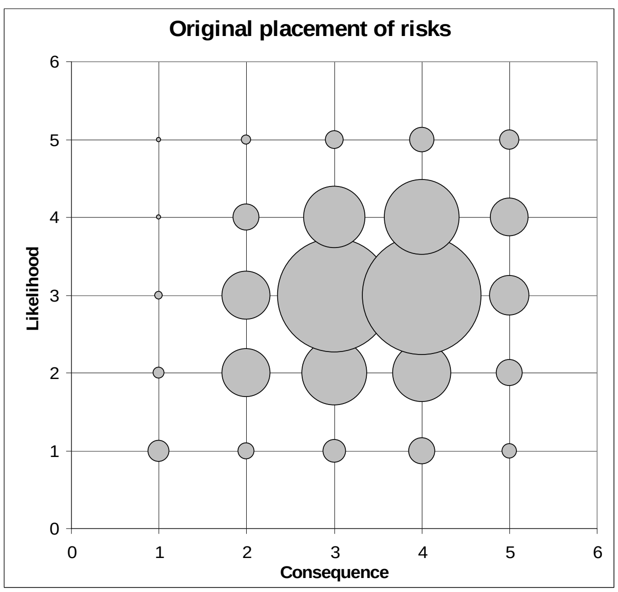

Figure 6: Bubble chart of original data points. Bubble areas proportional to datum counts.

ANALYSES AND OBSERVATIONS OF INITIAL DATA

First, we note that there are two impediments for the appearance of cognitive biases in the industry data: 1) The industry data are granular while the predictions of Prospect Theory are continuous, and 2) the industry-used qualitative descriptions for the five bins of likelihood and consequence may have had some non-linear influence in the placement of risk datum points.

Nevertheless, the evidence of the expected cognitive biases emerges from the data.

Estimation in a Pre-Define Scale Biases

"Given a list of numbers [respondents] are prone to choose those near the middle of the list" [Payne, 1951, p.80]. Schwarz [1990] conducted the following controlled experiment where a response scale effected judgment. The question 'How many hours do you watch television per week?' was asked of Germany survey participants. Half the participants were provided with the Low Category Range and the other half were provided with the High Category Range in Table 4. The estimated TV use by the average German citizen is only slightly more than 2 hours.
With both category ranges, responses were affected by the response scales. In each case, subjects calibrated a number in a way that seemed compatible with the given range of scale. "Even if the behavior under investigation is reasonably well defined … the range of response alternatives may strongly affect respondents' frequency estimates" [Fussell and Kreuz, 1998, p.51].

Reported Behavior			
Low Category Range		High Category Range	
Up to ½ hour	7.4% (5) [7%]		
½ - 1 hour	17.7% (12) [25%]		
1 – 1 ½ hours	26.5% (18) [52%]		
1 ½ - 2 hours	14.7% (10) [66%]		
2 – 2 ½ hours	17.7% (12) [84%]	Up to 2 ½ hours	62.5% (40) [63%]
More than 2 ½ hours	16.2% (11) [100%]	2 ½ - 3 hours	23.4% (15) [86%]
		3 – 3 ½ hours	7.8% (5) [94%]
		3 ½ - 4 hours	4.7% (3) [98%]
		4 -4 ½ hours	1.6% (1) [100%]
		More than 4 ½ hours	0.0% (0) [100%]

Table 4: Reported television watching per day for two response scales [Schwarz, 1990].

Statistical Evidence for Effect of Estimation in a Pre-Defined Scale

Figure 7 shows the original likelihood data centered in their 1-5 scale around L =3. "Given a list of numbers [respondents] are prone to choose those near the middle of the list" [Payne, 1951, p.80]. In the case of the subjective assessment of likelihood (L), this centering behavior may be compounded because: 'People estimate probabilities poorly' [Cosmides and Tooby, 1996].

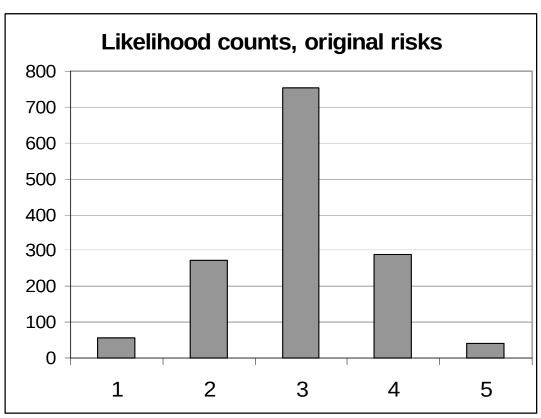

Figure 7: Likelihood data centered at L = 3.

Figure 8 shows the original consequence data centered around C = 3.5, roughly balanced on C =3 and C = 4. An explanation for the shift toward higher consequence is provided in Section 3.5.

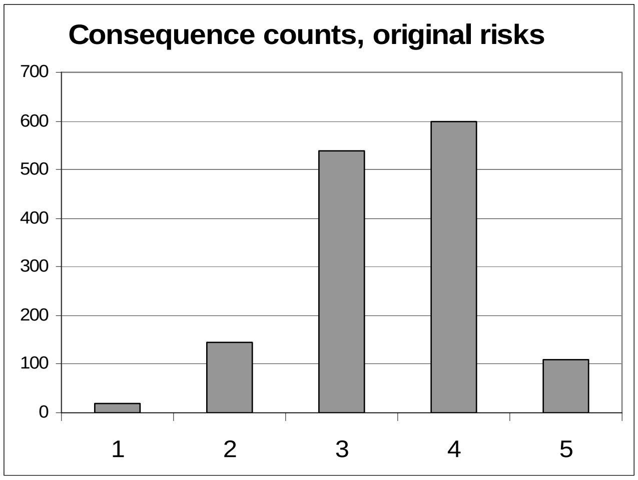

Figure 8: Consequence counts predominately at C = 4 and C = 3.

Infrastructure characteristics help explain some centering of data: A production airplane with a good safety record will have few extreme risks, and risk programs will not focus on low magnitude risks. Thus, some centering effect is expected. However, the original likelihood counts show significant heightened centering, while the original consequence counts show significant increases in magnitude.

Diagonal Bias

In anticipation of later moving the risk point toward the origin, engineers seem to withdraw risk points from the origin, upward and rightward along the diagonal.

Statistical Evidence for Diagonal Bias

The linear regression of all original risk points in Figure 9 shows a diagonal bias.

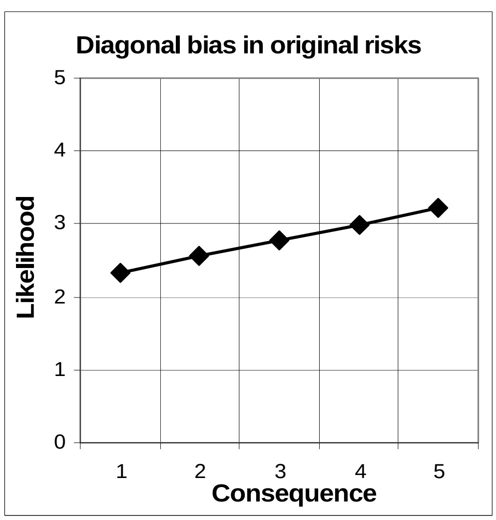

Figure 9: Linear regression of original datum points showing a diagonal bias.

Table 5 gives the linear regression coefficients associated with Figure 9.

Linear regression on 1412 Original Points		
Likelihood Intercept	Slope Coefficient	R^2
2.2	0.22	0.051

Table 5: Linear regression parameters showing diagonal bias.

Estimation of the linear regression coefficients as computed here might be biased because both the predictor and the response have been categorized into integer bins; for example, a risk category of 2 could correspond to an actual risk assessment anywhere from 1.5 to 2.5. This data censoring could be severe at the borders of the risk matrix where an assignment of 5 could have been derived from an elevated risk assessment of '5.7,' or even '6,' resulting from the consequence of a risk point exceeding the average consequence of risk points previously placed in the C = 5 bin.

CONCLUSION

To the authors' knowledge, this is the first time that the effects of cognitive biases have been documented within the risk matrix. The data show clear evidence that probability and value translations, as likelihood and consequence judgments, are present in industry risk matrix data. The straightforward steps in the development of this conclusion were as follows: 1) the translations were predicted by prospect theory, and, 2) historical data from 2 programs within 1 company confirmed the main predictions. Risk matrices have thus been shown to be, not objective number grids, but subjective, albeit useful, means to verify that risk items have received risk-mitigating attention. Confirmation of the presence of probability and value biases in risk data from other industries or companies can be the subject of future papers. The real world effects of using biased risk mitigation data should also be explored.

REFERENCES

L. Cosmides, and J. Tooby, Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty, Cognition 58 (1996), 1-73.

S. R. Fussell, and R. J. Kreuz, Social and cognitive approaches to interpersonal communications, Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 1998.

GAO, "NASA agency has taken steps toward making sound investment decisions for Ares I but still faces challenging knowledge gaps," United States Government Accountability Office, Washington, DC, 2007.

INCOSE, "International Council on Systems Engineering," 2006. Retrieved April 2006 http://www.incose.org.

D. Kahneman, and A. Tversky, Prospect theory: An analysis of decision under risk, Econometrica 46(2) (1979), 171-185.

Mil-Std-882d, Standard practice for system safety, Department of Defense, 2000.

Mil-Std-1629a, Procedures for performing a failure mode, effects and criticality analysis, Department of Defense, 1980.

Nobel, "The Bank of Sweden Prize in Economic Sciences in memory of Alfred Nobel 2002," 2002. Retrieved March, 2006 from Nobel Foundation: http://nobelprize.org/economics/laureates/2002/index.html.

S. L. Payne, The art of asking questions, Princeton University Press, Princeton, 1951.

M. Piattelli-Palmarini, Inevitable illusions: How mistakes of reason rule our minds, John Wiley & Sons, New York, 1994.

N. Schwarz, Assessing frequency reports of mundane behaviors: Contributions of cognitive psychology to questionaire construction, Review of Personality and Social Psychology 11 (1990), 98-119.

A. Tversky, and D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty, Journal of Risk and Uncertainty 5 (1992), 297-323.