Journal of Risk and Reliability, vol. 224(4), pp.346–355 (2010).
Proceedings of the Institution of Mechanical Engineers, Part O.

Reliability Assessment of Explosive Material Based on Penalty Tests: An Info-Gap Approach

Tanya Mirer[1] and Yakov Ben-Haim[2]

Abstract

A method is developed for experimental assessment of reliability of a system with a stringent safety requirement: explosive material. The focus is on analysis and management of both statistical variability of measurements and non-probabilistic uncertainty in probability distributions (distributional uncertainty). Info-gap theory is used to model the distributional uncertainty in the pdf of the threshold for actuation of the explosive material. The quantitative analysis and the qualitative judgments which accompany the certification of safety are studied. A proposition is proven asserting that the info-gap robustness function, for the class of problems examined, is independent of the experimental design over virtually all of its range.

Keywords: Safety, reliability, penalty tests, info-gaps, distributional uncertainty.

Notation

f_0: normalization coefficient.
$f(x)$: uncertain probability density function of x.
$\widetilde{f}(x)$: estimated probability density function of x.
$\widetilde{F}(x)$: estimated cumulative probability distribution of x.
h: horizon of uncertainty.
$\widehat{h}(S_{\rm c})$: robustness function.
$M(h)$: inverse of robustness function.
n: number of nulls; sample size.
q: actuation probability at load no greater than $x_{\rm a}$.
$\widetilde{q}$: estimated value of q.
$S(f)$: probability of no actuation at or below $x_{\rm s}$.
$S_{\rm a}(f)$: probability of no actuation at or below $x_{\rm a}$.
$S_{\rm c}$: critical value of $S(f)$.
$\mathcal{U}(h)$: info-gap model of distributional uncertainty.
x: random actuation threshold.
$x_{\rm a}$: penalty load.
$x_{\rm s}$: safe load.
α: level of significance.
μ: mean of x.
σ^2: variance of x.
$\rho_{\rm ft}$: fat-tail parameter.
$\rho_{\rm ms}$: mean-shift parameter.
$\rho_{\rm sds}$: standard-deviation-shift parameter.

1 Introduction

An "explosive train" is a chain of explosive components with various sensitivity characteristics. Knowledge of actuation thresholds is critical for the design of explosive trains. The detonation starts

[0]\papers\Explosive\ett06.tex. 8.5.2010
[1]RAFAEL, POB 2250, Haifa, Israel, ztatiana@tx.technion.ac.il
[2]Yitzhak Moda'i Chair in Technology and Economics, Faculty of Mechanical Engineering, Technion—Israel Institute of Technology, Haifa, Israel, yakov@technion.ac.il

from a most sensitive link and transfers to subsequent links until actuation of all links in the chain is obtained.

The explosive tendency of material is random and can be usefully modeled probabilistically. This is done by fitting a probability distribution to data on the explosion threshold. That is, one empirically establishes the probability of actuation as a function of the stress level to which the material is exposed.

The result of a threshold test is either actuation or no actuation, and in either case the test destroys the tested material. Consequently, it is impossible to directly measure the sensitivity threshold. In any individual test one only knows if the stress was higher or lower than the actuation threshold of that sample. That is, the data are censored.

There are two statistical characteristics of the actuation-threshold distribution which are of prime importance in practice and must be confidently estimated: "reliability" and "safety". The *reliability* is the probability that actuation will occur at or above the operational level of stress which is intended to cause actuation. The safety, in contrast, relates to random environmental loads which are far below the operational loads at which actuation of the material is intended. The *safety* is the probability that actuation will occur only at levels of stress which exceed ordinary environmental loads. Safety requirements are typically quite strict, meaning that our estimate of the actuation-threshold distribution must justify high confidence that the safety is nearly one. For instance a typical mil-spec [MIL-HDBK-83578] requires safety of 0.999 which is estimated with 90% confidence. Hence it is necessary to confidently estimate the far tails of the actuation-threshold distribution.

There are several methods for safety and reliability demonstration of explosives.

Binomial tests employ a large number of tests at very low fixed load, designed to avoid actuation. These tests require a large number of test items and, therefore, are very costly. For example, 2300 tests without any failures are required to demonstrate safety of 0.999 with 90% confidence.

Quintal tests employ samples which are tested at various stress levels with both actuation and non-actuation (Bruceton and Probit tests) [Porat, Haim and Markiewicz 1994; Alouaamari, Lefebvre and Perneel 2007; Nance 2008]. Quintal methods are adaptive in the sense that past responses are utilized to more efficiently choose load levels for subsequent tests.

Assuming the type of distribution (usually normal), quintal tests estimate the statistical parameters by fitting the distribution to the obtained censored data. Quintal tests can provide good estimation of the mean, but estimating the standard deviation is more difficult and, consequently, the tails of the distribution may be poorly characterized. It is the tails which are particularly important for safety analysis.

Penalty tests are based on imposing a "penalty" on the actuation process by reducing or amplifying the stress level in relation to nominal values. For instance, the tested material may be exposed to shock waves which are 50% greater than the expected value. Another typical penalty test uses a barrier gap or varies the thickness of the explosion chain interface and the sensitivity of the explosion donor [Ayres *et al.* 1961; Tzidony and Jaeger 1998; De Yong 1986]. The Bruceton or similar techniques are often used to select an appropriate level of penalty. Penalty tests are thus similar to accelerated lifetime tests, though in the present context the purpose is to estimate a threshold rather than a lifetime.

Penalty tests are described further in section 2.1, where they are used in conjunction with binary tests. This enables estimation of the tails of the distribution with a relatively small sample.

Penalty tests are based on a range of assumptions, including:

- The family of probability distributions is known (usually normal).
- The standard deviation is a known fraction of the mean value.
- The penalty amplifies the stress values but does not change the explosion mechanism.

Some or all of these assumptions may be difficult to verify and thus are often untested. The purpose of this paper is to develop a technique for evaluating safety without explicitly validating some of the underlying assumptions.

Section 2 explains in detail the interpretation of penalty tests for actuation-threshold estimation. Section 3 introduces the concept of distributional uncertainty to quantify the analyst's lack of knowledge about the shape of the probability distribution of the actuation threshold. Section 4 introduces and discusses the info-gap robustness function, and section 5 studies an example.

2 Nominal Interpretation of Penalty Tests

Let x denote the threshold for actuation of the explosive material, which is a random variable with probability density function (pdf) $f(x)$.

The safety of the explosive material is the probability that actuation will not occur at or below the operational "safe load" $x_{\rm s}$. The safety depends on the pdf of the actuation threshold, $f(x)$, as:

$$S(f) = \int_{x_{\rm s}}^{\infty} f(x)\,{\rm d}x \qquad (1)$$

Penalty tests are performed at load $x_{\rm a}$ which exceeds the safe load $x_{\rm s}$. The probability of no-actuation at or below the augmented "penalty" load $x_{\rm a}$ is:

$$S_{\rm a}(f) = \int_{x_{\rm a}}^{\infty} f(x)\,{\rm d}x \qquad (2)$$

The purpose of the penalty test is to estimate the safety of the explosive chain, $S(f)$. The nominal interpretation of penalty tests involves two stages. First, $S_{\rm a}(f)$ is estimated from observations with the use of a binary hypothesis test, as described in section 2.1. Then $S(f)$ is estimated by assuming the functional form of the pdf of the actuation threshold, as described in section 2.2.

2.1 Estimation of $S_{\rm a}$ with a Binary Hypothesis Test

Consider a sequence of statistically independent trials of the explosive material, all performed at the augmented load $x_{\rm a}$. A trial is 'null' if no actuation occurred, and 'positive' otherwise. Let q denote the probability of actuation at load no greater than $x_{\rm a}$ which is $1 - S_{\rm a}$ in eq.(2). $x_{\rm a}$ is larger than the safe load $x_{\rm s}$ but still on the left tail of the pdf of actuation thresholds. The only relevant situation in practice is where n null results have been observed in n trials. The safety estimation is irrelevant if actuation occurs during testing. Actuation during testing indicates inadequate design with too small a safety margin, implying that the explosive chain needs to be re-designed.

Let $\widetilde{q}$ be an hypothesized value of q. The analyst would like to test this actuation probability against a smaller value. Having observed n null results in n trials, and the analyst wishes to test between the following two hypotheses:

$$H_0: \qquad q = \widetilde{q} \qquad (3)$$

$$H_1: \qquad q < \widetilde{q} \qquad (4)$$

That is, having observed n nulls in n trials, at what confidence can one reject the hypothesis that $q = \widetilde{q}$ in favor of the hypothesis that actuation probability at the augmented load is smaller than $\widetilde{q}$?

For any binary hypothesis test, the level of significance is the probability, conditioned on H_0, of a result which is at least as extreme as the observation. Stated more explicitly for our case, the level of significance of this test is the probability, conditioned on H_0, of a result which impugns H_0 more than the observation. Denote the number of observed nulls by $m_{\rm o}$. H_0 would be weakened, with respect to H_1, if $m_{\rm o}$ or more nulls were observed. The level of significance for rejecting H_0 is stated formally as the probability of at least $m_{\rm o}$ nulls:

$$\alpha = {\rm Prob}(m \geq m_{\rm o} | H_0) \qquad (5)$$

In our case the number of observed nulls is n, equal to the number of trials, so, using the binomial distribution, eq.(5) becomes:

$$\alpha = (1 - \widetilde{q})^n \qquad (6)$$

Eq.(6) is the level of significance for rejecting H_0 against H_1. It can also be understood as the probability of obtaining n nulls in n trials if H_0 holds. Inverting eq.(6) one finds that the largest value of $\widetilde{q}$ for which H_0 is not rejected at level of significance α is:

$$\widetilde{q} = 1 - \alpha^{1/n} \tag{7}$$

This value of $\widetilde{q}$ is referred to as the estimate of $1-S_{\rm a}$ at level of significance α, given n null observations out of n trials. Speaking less precisely (but more suggestively) one will sometimes say that, given n nulls out of n trials, the probability of actuation is no larger than $\widetilde{q}$ at level of significance α. Any larger value of $\widetilde{q}$ would be rejected in favor of H_1 at level of significance α.

For instance, for $n = 7$, one finds from eq.(7) that H_0 is rejected at level of significance $\alpha = 0.1$ with $\widetilde{q}$ larger than 0.280. If one wants to test H_0 more rigorously to obtain more confidence, for instance $\alpha = 0.05$, one rejects H_0 with $\widetilde{q}$ larger than 0.348 given 7 nulls in 7 trials. Greater confidence is obtained by testing a larger value of the probability of actuation.

2.2 Nominal Estimation of the Reliability

The best estimate of the pdf of the actuation threshold is denoted $\widetilde{f}(x)$, which will be assumed to be unimodal and defined between $-\infty$ and $+\infty$. In preparation for the numerical example in section 5 it will be assumed, in part of the subsequent development, that $\widetilde{f}(x)$ is normal with mean μ and variance σ^2.

The analyst may have further information which constrains the pdf. As a typical example of this additional information it will be assumed that the standard deviation increases linearly with the mean as:

$$\sigma = c\mu \tag{8}$$

where c is known, and c and μ are both positive. Given a sequence of null results it is not possible to independently estimate both mean and variance. Hence an assumption such as eq.(8) is essential to the binary penalty test method which requires mean and variance in order to fully specify the normal pdf (and many others).

Assumptions such as these—normality, linearity, etc.—are subject to considerable uncertainty which will be the focus of our attention in section 3.

The analyst wishes to use the experimental trials to determine the parameters of the nominal pdf, $\widetilde{f}(x)$. The analyst assumes that $\widetilde{f}(x)$ is normal with mean μ and variance σ^2. μ is estimated by relating the estimated probability of actuation, eq.(7), to $\widetilde{f}(x)$:

$$1 - \alpha^{1/n} = \int_{-\infty}^{x_{\rm a}} \widetilde{f}(x)\,{\rm d}x \tag{9}$$

μ is evaluated (recall that it is assumed to be positive) by assuming normality of $\widetilde{f}$ and using eq.(8). The result is obtained by noting that the righthand side of eq.(9) can be expressed in terms of the cumulative probability distribution (cpd) of the standard normal variate. This cpd is then inverted to obtain:[3]

$$\mu = x_{\rm a}\left[1 + c\,\Phi^{-1}\Big(1 - \alpha^{1/n}\Big)\right]^{-1} \tag{10}$$

where Φ^{-1} is the inverse of the cpd of the standard normal variate.

With μ from eq.(10) and σ from eq.(8), the nominal estimate of the safety of the system is obtained from eq.(1) with $\widetilde{f}(x)$ as the normal distribution with mean and variance μ and σ^2.

[3] Note that, if $x_{\rm a} > 0$, then $\mu > 0$ if and only if:

$$1 - \alpha^{1/n} > \Phi\left(-\frac{1}{c}\right)$$

There is no positive solution for μ if this condition does not hold. Stated differently, μ jumps discontinuously to large negative values when this relation is violated.

3 Uncertainty of the Nominal PDF

The nominal estimate of the safety, described in section 2, rests on the choice of the nominal pdf of the actuation threshold, $\widetilde{f}(x)$, and on the assumed relation between the mean and the standard deviation, eq.(8), or other similar information. Both assumptions are approximations. Of particular concern is the assumption of knowledge of the precise shape of the distribution. Since very high safety is required, and since new materials are continually introduced but not studied as thoroughly as one might wish, this assumption is subject to considerable uncertainty. Uncertainty in the shape of the pdf of the actuation threshold is referred to as *distributional uncertainty.*

The question which is asked is, How wrong can the nominal pdf be, and the estimated safety is still acceptable? That is, how robust is the safety to distributional uncertainty?

Furthermore, one wishes to know how the robustness to distributional uncertainty varies as one changes the test design. In particular, how does robustness vary with sample size n (number of nulls) and augmented load $x_{\rm a}$? This is the basis for designing the safety assessment.

One first formulates an info-gap model for quantifying the distributional uncertainty in the pdf of the actuation threshold. Then, in section 4, the performance requirement and the robustness function are defined. Finally a procedure for numerical evaluation of the robustness will be formulated. Section 5 illustrates the implications for test design and safety assessment.

The choice of an **info-gap model** to represent uncertainty in the shape of the actuation pdf depends on the analyst's prior knowledge and professional judgment. Many info-gap models are available (Ben-Haim, 2006). The analysis is illustrated with a uniform-bound model, which allows uncertain bumps, wrinkles and multi-modality, without allowing atoms of probability. Let f_0 be a constant, for instance the value of the estimated pdf at its mode. The info-gap model is:

$$\mathcal{U}(h) = \left\{ f(x) : \ f(x) \geq 0, \ \int_{-\infty}^{\infty} f(x)\,{\rm d}x = 1, \ |f(x) - \widetilde{f}(x)| \leq f_0 h \right\}, \quad h \geq 0 \tag{11}$$

$\mathcal{U}(h)$ is the set of pdf's which deviate from the nominal pdf, $\widetilde{f}(x)$, by no more than $f_0 h$. When $h = 0$ the set contains only the nominal pdf, $\widetilde{f}(x)$. The sets become more inclusive as h grows, so h is referred to as the horizon of uncertainty.

4 Robustness to Distributional Uncertainty

4.1 Formulation

The robustness to distributional uncertainty is based on three components: an info-gap model of distributional uncertainty, eq.(11), an expression for system safety, eq.(1), and a performance requirement which is now formulated.

The performance requirement is that the safety, $S(f)$ in eq.(1), is certified—based on tests—to be no less than a critical value $S_{\rm c}$:

$$S(f) \geq S_{\rm c} \tag{12}$$

This is a "satisficing" requirement. Rather than attempting to maximize the safety, one only requires that it satisfy a specified requirement.

Safety assessment is based on the number of trials, n, and the value of the augmented load, $x_{\rm a}$. It is assumed that the safe operational load, $x_{\rm s}$, is fixed. The robustness of a design, $(n, x_{\rm a})$, is the greatest horizon of distributional uncertainty, h, up to which the performance requirement is guaranteed:

$$\widehat{h}(S_{\rm c}) = \max \left\{ h : \ \left(\min_{f \in \mathcal{U}(h)} S(f) \right) \geq S_{\rm c} \right\} \tag{13}$$

4.2 Evaluation

The numerical evaluation of the robustness is implemented by evaluating the inverse of the robustness function, as is now explained.

Let $M(h)$ denote the inner minimum in the definition of the robustness, eq.(13). $\widehat{h}(S_c)$ is the greatest horizon of uncertainty, h, at which $M(h) \geq S_c$. Recall that the uncertainty sets, $\mathcal{U}(h)$, become more inclusive as h increases. This implies that $M(h)$, which is a minimum on the set $\mathcal{U}(h)$, decreases as h increases. Consequently $\widehat{h}(S_c)$ is the greatest horizon of uncertainty at which $M(h) = S_c$. In other words, a plot of h vs. $M(h)$ is the same as a plot of $\widehat{h}(S_c)$ vs. S_c. In short, $M(h)$ is the inverse of $\widehat{h}(S_c)$.

The evaluation of $M(h)$ is now demonstrated.

The safety depends on the pdf of the actuation threshold, eq.(1). $M(h)$ is the minimal safety from among all the pdf's contained in the info-gap model at horizon of uncertainty h. The pdf which minimizes the safety is the one whose lower tail—below x_s—is as fat as possible subject to the constraints on membership in $\mathcal{U}(h)$. That pdf is now derived and $M(h)$ is evaluated.

First define the points on the upper and lower tails of the nominal pdf, $\widetilde{f}(x)$, at which the lower envelope of the info-gap model vanishes at horizon of uncertainty h, recalling the assumption that $\widetilde{f}(x)$ is unimodal on the infinite interval $(-\infty, +\infty)$:

$$x_u(h) = \max\left\{x : \widetilde{f}(x) - f_0 h = 0\right\} \quad (14)$$

$$x_\ell(h) = \min\left\{x : \widetilde{f}(x) - f_0 h = 0\right\} \quad (15)$$

From the uni-modality of $\widetilde{f}(x)$ one sees that $x_u(h)$ decreases, and $x_\ell(h)$ increases, as h increases. $x_u(h)$ and $x_\ell(h)$ converge on the mode of $\widetilde{f}(x)$ as h increases, and $x_u(h)$ and $x_\ell(h)$ are both defined to equal the mode when h is large enough so that the sets in eqs.(14) and (15) are empty.

Define $x_b(h)$ as:

$$x_b(h) = \max[x_\ell(h),\ x_s] \quad (16)$$

Note that $x_b(h)$ increases (or is constant) as h increases.

Now define the following function which, as will be explained, is a pdf belonging to the info-gap model at horizon of uncertainty h. The basic idea is that the pdf is fattened at low loads as much as possible, subject to the constraint on thinning the distribution at high loads in order to achieve normalization and non-negativity.

$$f(x) = \begin{cases} \widetilde{f}(x) & \text{if } x < x_s - D \\ \widetilde{f}(x) + f_0 h & \text{if } x_s - D \leq x \leq x_s \\ 0 & \text{if } x_s < x \leq x_b(h) \\ \widetilde{f}(x) - f_0 h & \text{if } x_b(h) \leq x \leq x_u(h) \\ 0 & \text{if } x_u(h) < x \end{cases} \quad (17)$$

where D is chosen to achieve normalization as is now explained. The pdf in eq.(17) with a 2% bump ($h = 0.02$), together with the nominal pdf, are illustrated in fig. 1. A 2% bump on the lower tail, below x_s, is much larger than the nominal pdf at those values, as seen in fig. 1. The effect on the bulk of the distribution is hardly noticeable.

****Fig. 1 here.****

1. Lines 3–5: shift $\widetilde{f}(x)$ *down* by $f_0 h$ for all $x > x_s$ and obey non-negativity (which explains the 0's).

2. Line 2: shift $\widetilde{f}(x)$ *up* by $f_0 h$ over an interval of length D sufficient to compensate for the down-shift in lines 3–5.

3. Line 1: retain $\widetilde{f}(x)$ for $x < x_s - D$.

It is evident that $f(x)$ is non-negative, normalized, and within the info-gap envelope at horizon of uncertainty h. No other pdf in $\mathcal{U}(h)$ has lower probability weight at $x \ge x_{\rm s}$. This pdf is the basis for evaluating $M(h)$.

$M(h)$ is evaluated from eq.(1) with $f(x)$ in eq.(17), to obtain:

$$M(h) = \int_{x_{\rm b}(h)}^{x_{\rm u}(h)} \left(\widetilde{f}(x) - f_0 h\right) \,{\rm d}x \tag{18}$$

$$= \widetilde{F}[x_{\rm u}(h)] - \widetilde{F}[x_{\rm b}(h)] - [x_{\rm u}(h) - x_{\rm b}(h)] f_0 h \tag{19}$$

where $\widetilde{F}(x)$ is the cumulative probability distribution of the nominal pdf $\widetilde{f}(x)$. It was noted earlier that $x_{\rm u}(h)$ decreases, and $x_{\rm b}(h)$ increases (or is constant), as h increases. Hence eq.(18) shows that $M(h)$ decreases as h increases. As explained earlier, a plot of h vs. $M(h)$ is identical to a plot of $\widehat{h}(S_{\rm c})$ vs. $S_{\rm c}$. Hence eq.(19) is the basis for numerical evaluation of the robustness function.

4.3 Interpretation of the Robustness

Section 4.2 explained how to numerically evaluate the robustness function, defined in eq.(13). Interpretation of the numerical values which are obtained is now considered. This will be illustrated in the example in section 5.

The questions to be addressed, in interpreting a numerical value of robustness, are: Is this value of robustness large? large enough? too small?, etc. The robustness is interpreted by "calibrating" its numerical value against prior knowledge. This is a qualitative and sometimes subjective process, as is inevitable when dealing with the question "How much robustness is enough?". Three methods will be presented. Further discussion of interpreting robustness based on the concept of analogical reasoning can be found in (Ben-Haim, 2006, chapter 4). Much relevant work has also been done on elicitation of expert opinion (Ayyub 2001, Meyer and Booker 1991) which can be used for calibrating the robustness.

Tail fattening. Consider a numerical value of the robustness, $\widehat{h}(S_{\rm c})$. Let us denote the pdf in eq.(17) by $f(x|h)$. This pdf is an extreme fat-tail distribution belonging to the info-gap model at horizon of uncertainty h. In particular, $f(x|\widehat{h})$ belongs to the info-gap model at horizon of uncertainty equal to the robustness, $\widehat{h}(S_{\rm c})$, and its value of safety is minimal. This means that the safety will not be less than $S_{\rm c}$ even if the pdf $f(x|\widehat{h})$ occurs. If the analyst can make qualitative judgments about the degree of plausibility of tail-fattening, then this can be used to judge whether robustness $\widehat{h}$ is small or large. If bumps larger than displayed by $f(x|\widehat{h})$ are plausible, then the robustness is low, since $f(x|\widehat{h})$ is an extreme deviation from the nominal pdf which brings the safety to its limiting value. On the other hand, if bumps larger than those of $f(x|\widehat{h})$ are implausible then the robustness is high.

The judgment of the robustness is facilitated by a parameter which expresses the fractional deviation of $f(x|\widehat{h})$ from the nominal pdf, $\widetilde{f}(x)$, at the safe load $x_{\rm s}$:

$$\rho_{\rm ft}(S_{\rm c}) = \frac{f(x_{\rm s}|\widehat{h}) - \widetilde{f}(x_{\rm s})}{\widetilde{f}(x_{\rm s})} \tag{20}$$

$$= \frac{f_0 \widehat{h}(S_{\rm c})}{\widetilde{f}(x_{\rm s})} \tag{21}$$

where eq.(21) is obtained by employing the 2nd line of eq.(17) in eq.(20). If $\rho_{\rm ft}$ is much larger than unity then great tail-fattening can be tolerated without violating the safety requirement, suggesting that the robustness is large. If $\rho_{\rm ft}$ is much smaller than unity then the robustness would seem to be small. The analyst must make the qualitative judgment of what is "much larger" or "much smaller".

Mean shift. Let $\mu_f(h)$ denote the mean of the pdf, $f(x|h)$, in eq.(17). This will be less than μ, the mean of the nominal pdf $\widetilde{f}(x)$, since $f(x|h)$ has a fattened lower tail. Define the fractional shift of the mean:

$$\rho_{\rm ms}(S_{\rm c}) = \frac{\mu - \mu_f(\widehat{h})}{\mu} \tag{22}$$

The analyst may have experience which enables qualitative judgment of the size of a shift in the mean. This refers to shifts in the true population mean which result from substantive change in the underlying processes. If ρ_{ms} is judged to be much greater than plausible mean-shifts then the robustness would seem to be large. If ρ_{ms} is a plausible mean-shift then the robustness would not seem adequate.

Standard-deviation shift. Let $\sigma_f(h)$ denote the standard deviation of $f(x|h)$. This will be greater than σ, the standard deviation of the nominal pdf $\widetilde{f}(x)$, due to the fattened lower tail. Define the fractional shift of the standard deviation:

$$\rho_{\text{sds}}(S_{\text{c}}) = \frac{\sigma_f(\widehat{h}) - \sigma}{\sigma} \tag{23}$$

If ρ_{sds} is judged to be much greater than plausible shifts in the standard deviation then the robustness would seem to be large. If ρ_{sds} is a plausible shift then the robustness would not seem adequate.

4.4 Normal Case

An important special case is now examined: the nominal distribution, $\widetilde{f}(x)$, is normal and f_0 in the info-gap model of eq.(11) is the value of $\widetilde{f}(x)$ at the mode. It will be proven that the robustness curve is independent of the experimental design for most situations of practical interest.

First define a limiting value of the horizon of uncertainty:

$$h_\ell = \exp\left[-\frac{1}{2}\left(\frac{\mu - x_{\text{s}}}{\sigma}\right)^2\right] \tag{24}$$

In all practical cases the operational safe load, x_{s}, is much less than the nominal mean, μ, in units of σ, so h_ℓ will be very close to zero.

Proposition 1 *The robustness curve is independent of the experimental design if h_ℓ is small.*
Given:

- *The nominal distribution, $\widetilde{f}(x)$, is normal with mean μ and variance σ^2.*
- $x_{\text{s}} \le \mu$.
- *The mean and standard deviation are related as in eq.(8).*
- *The info-gap model is eq.(11) with normalization coefficient f_0 at the mode:*

$$f_0 = \frac{1}{\sigma\sqrt{2\pi}} \tag{25}$$

Then:

$$M(h) = \begin{cases} 0 & \text{if } h \ge 1 \\ \Phi(\sqrt{-2\ln h}) - \Phi(-\sqrt{-2\ln h}) - \dfrac{2h\sqrt{-\ln h}}{\sqrt{\pi}} & \text{if } h_\ell \le h \le 1 \end{cases} \tag{26}$$

where $\Phi(\cdot)$ is the cumulative probability distribution of the standard normal variate.

Eq.(26) depends on the moments μ and σ of the nominal pdf $\widetilde{f}(x)$, the level of significance α, the number of null observations n, the operational safe load x_{s} and the augmented penalty load x_{a} *only* through the lower limit h_ℓ. The *shape* of the function is entirely independent of these quantities. In practice h_ℓ is always exceedingly close to zero. This means that, in practical terms, the robustness curve is independent of the experimental design as expressed by α, n, x_{s} and x_{a}.

Proof of proposition 1.

Given the choice of f_0 in eq.(25), note that $x_{\text{u}}(h) = x_\ell(h) = \mu$ for $h \ge 1$, as explained following eq.(15). Thus, since $x_{\text{s}} \le \mu$ by supposition, $M(h) = 0$ for $h \ge 1$ as seen with eqs.(16) and (18). This proves the first line of eq.(26).

The second line of eq.(26) is now proven.

For the normal distribution with mean and variance μ and σ^2 one can readily show that, for $h \le 1$, $x_{\rm u}(h)$ and $x_\ell(h)$ defined in eqs.(14) and (15) can be expressed:

$$x_{\rm u}(h) = \mu + \sigma\sqrt{-2\ln h} \quad (27)$$
$$x_\ell(h) = \mu - \sigma\sqrt{-2\ln h} \quad (28)$$

Using the explicit form of the normal pdf for $\widetilde{f}(x)$, it is readily shown that $h \ge h_\ell$ implies that $x_\ell(h) \ge x_{\rm s}$. Thus, from eq.(16) one finds that $x_{\rm b}(h) = x_\ell(h)$. In other words, $M(h)$ in eq.(19) is evaluated at $x_{\rm b}(h) = x_\ell(h)$.

Let $\widetilde{F}(x)$ denote the cpd of the nominal pdf. Thus, employing eq.(27):

$$\widetilde{F}(x_{\rm u}) = \text{Prob}(x \le x_{\rm u}) = \text{Prob}\left(\frac{x-\mu}{\sigma} \le \frac{x_{\rm u}-\mu}{\sigma}\right) = \Phi\left(\frac{x_{\rm u}-\mu}{\sigma}\right) \quad (29)$$
$$= \Phi\left(\sqrt{-2\ln h}\right) \quad (30)$$

Similarly, using eq.(28):

$$\widetilde{F}(x_\ell) = \Phi\left(-\sqrt{-2\ln h}\right) \quad (31)$$

The second line of eq.(26) results from combining eqs.(19), (25), (27), (28), (30) and (31). ■

5 Example and Discussion

****Fig. 2 here.****

Fig. 2 shows a curve of robustness, $\widehat{h}(S_{\rm c})$, vs. critical safety, $S_{\rm c}$, for a specific experimental configuration. The operational safe load (in shifted and normalized units) is $x_{\rm s} = -0.4$ and the augmented penalty load is $x_{\rm a} = 1$. The nominal pdf of the actuation threshold, $\widetilde{f}(x)$, is normal, the ratio between the standard deviation and the mean in the nominal pdf is $c = 0.3$, the level of significance is $\alpha = 0.1$ and the number of observations, all nulls, is $n = 7$.

In this configuration the estimated mean at level of significance $\alpha = 0.1$, eq.(10), is $\mu = 1.212$ so the estimated standard deviation is $\sigma = 0.363$ according to eq.(8). Thus the augmented penalty load, $x_{\rm a}$, is slightly more than half a standard deviation below the estimated mean, while the safe load, $x_{\rm s}$, is more than 4 standard deviations below the estimated mean. The limiting horizon of uncertainty, h_ℓ in eq.(24), is 5.38×10^{-5}.

Two basic properties of all robustness curves—trade off and zeroing—are demonstrated in fig. 2, as is now explained.

The negative slope of the curve expresses the **trade off** between robustness against distributional uncertainty on the one hand, and the statistical safety of the system on the other hand. Large robustness entails low safety. The evaluation of the statistical safety depends on knowledge of the pdf of the actuation threshold. However, this pdf is uncertain. The robustness, $\widehat{h}(S_{\rm c})$ on the vertical axis, is the horizon of uncertainty in the pdf up to which the safety is guaranteed to be no less than the corresponding critical value, $S_{\rm c}$ on the horizontal axis. In the insert in fig. 2 one sees that statistical safety of 0.995 (or more) is guaranteed at robustness of 0.0016. Greater robustness, say 0.0034, is obtained only by accepting lower statistical safety of 0.99. In short, robustness against distributional uncertainty trades off against statistical safety. (The intuitive interpretation of these numerical values of robustness will be discussed shortly.)

This trade off can be understood as a conflict between the two foci of uncertainty which are present in this problem. On the one hand the threshold of actuation is a random variable which is represented by a pdf. On the other hand, this pdf is uncertain. The statistical safety assesses performance vis á vis the random threshold, while the info-gap robustness assesses the immunity against distributional uncertainty in the pdf. The trade off results from the conflict between these two foci of uncertainty.

S_c	$\widehat{h}(S_c)$	$\rho_{ft}(S_c)$	$\rho_{ms}(S_c)$	$\rho_{sds}(S_c)$
0.9997	8.1×10^{-5}	1.50	0.0023	0.011
0.999	2.9×10^{-4}	5.50	0.0030	0.013
0.998	6.1×10^{-4}	11.3	0.0040	0.015
0.99	3.4×10^{-3}	63.9	0.012	0.041
0.95	2.0×10^{-2}	372	0.052	0.17
0.90	4.3×10^{-2}	812	0.10	0.31

Table 1: Parameters for interpreting the robustness values.

Zeroing is the property that the robustness curve hits the horizontal axis at the best estimate of the safety. That is:

$$\widehat{h}(S_c) = 0 \quad \text{if} \quad S_c = S(\widetilde{f}) \tag{32}$$

$\widetilde{f}(x)$ is the estimated pdf of the actuation threshold at level of significance α, which in fig. 2 is normal with mean 1.212 and standard deviation 0.363. With these values the estimated safety, eq.(1), is $S(\widetilde{f}) = 0.999\,995$ which is exceedingly close to one. However, the robustness against distributional uncertainty is zero at this estimate (that's the zeroing property), which means that one cannot have much confidence in this extraordinarily high statistical safety.

Combining the properties of zeroing and trade off, one concludes that only *lower* statistical safety has *positive* robustness against distributional uncertainty.

Interpreting the numerical robustness values is an important matter involving professional judgment. Table 1 shows robustness values and other parameters for a range of critical safeties. For instance, the robustness for critical safety $S_c = 0.999$ is $\widehat{h}(0.999) = 2.9 \times 10^{-4}$. Is this robustness large? large enough? or too small? To understand what '$\widehat{h} = 2.9 \times 10^{-4}$' means one must return to the info-gap model of distributional uncertainty in eq.(11). When the horizon of uncertainty is $h = 2.9 \times 10^{-4}$ the uncertainty set $\mathcal{U}(h)$ contains, among its pdf's, some pdf's with bumps—excess probability density—on the lower tail as large as 2.9×10^{-4} times the nominal pdf at its maximum. A bump of this size on the lower tail at x_s is much larger than the nominal pdf at x_s as seen by $\rho_{ft}(0.999) = 5.5$ in table 1. If the analyst makes the judgment that fattening as much as $\rho_{ft}(0.999) = 5.5$ is 'very large' or 'extraordinary', then robustness of 2.9×10^{-4} is 'large'. On the other hand if bumps only 5.5 times the nominal pdf at x_s are felt to be small or plausible corruptions of the pdf, then robustness of 2.9×10^{-4} is small. These are value judgments like all judgments of questions such as 'How safe is safe enough?'.

Table 1 also shows values of the mean shift and standard-deviation shift, ρ_{ms} and ρ_{sds}. Looking again at the row for $S_c = 0.999$ one sees that the mean of the extreme pdf is reduced fractionally with respect to the nominal mean by $\rho_{ms}(0.999) = 0.0030$, and the standard deviation is enlarged fractionally by $\rho_{sds}(0.999) = 0.013$. Not surprisingly, the shift parameters ρ_{ms} and ρ_{sds}—which are averages over the entire distribution—are much less sensitive than the fat-tail parameter ρ_{ft} which is localized at the value of x_s. Robustness of $\widehat{h}(0.999) = 2.9 \times 10^{-4}$ is large if shifts in the true population mean as large of 0.3% are extraordinary and would result only from process changes or material contamination and so on. A similar interpretation applies to ρ_{sds}.

Suppose the parameters ρ_{ft}, ρ_{ms} and ρ_{sds} do not indicate sufficient robustness at critical safety $S_c = 0.999$. Table 1 shows parameters for a range of safety values. The robustness increases as S_c decreases, which is the trade off property. The analyst can use results such as Table 1 to identify the value of critical safety at which the robustness is judged to be large. This critical safety can be confidently ascribed to the system in question.

Discussion of proposition 1. The safe operating load, x_s, will always be much less than the mean actuation threshold, μ, in units of the standard deviation σ. This means that the limiting horizon of uncertainty, h_ℓ in eq.(24), will always be very close to zero. This in turn implies that the robustness function, $\widehat{h}(S_c)$, is independent of the experimental configuration—number of observations n, level of significance α, and augmented penalty load x_a—over virtually all of its range. Proposition 1 is an explicit expression for the (inverse) of this universal robustness function, and fig. 2 shows its shape.

Why can't the robustness against distributional uncertainty be improved by, say, increasing the sample size where all measurements are nulls (no actuation)?

Part of the answer lies is noting that the pdf of the actuation threshold could be estimated, thus reducing or eliminating the distributional uncertainty. However, this would require non-null observations (actuations) at various loads. Large numbers of measurements would result in better estimation than low sample sizes. Small samples are considered, in which all measurements are nulls: no actuation occurs.

In addition, the robustness curve is independent of the sample size and level of significance because these variables relate to the *statistical uncertainty* of the probability of no-actuation, as expressed in the binary hypothesis test of eqs.(3) and (4). The robustness curve deals with the *distributional uncertainty* surrounding the pdf of the actuation threshold. Speaking metaphorically one might say that the statistical uncertainty (represented by a pdf) is orthogonal to the distributional uncertainty (represented by an info-gap model). This orthogonality obtains some support (graphical, at least) by noting that the axes of fig. 2 (which are orthogonal to one another) are statistical safety and info-gap robustness. These are two distinct types of uncertainty, touching two different aspects of the overall problem. Nonetheless, this metaphor should not be carried too far since proposition 1 depends on specific conditions such as the linear relation between mean and standard deviation in eq.(8) and the uniform-bound info-gap model in eq.(11).

****Fig. 3 here.****

Very high safety. Proposition 1 holds for robustness greater than h_ℓ, defined in eq.(24), which is necessarily close to zero. The proposition asserts that, for robustness greater than h_ℓ, the robustness curve is universal: independent of sample size, level of significance, etc. For robustness less than h_ℓ—at which the critical safety is very large—the robustness curves become differentiated for different parameter values. Fig. 3 shows robustness curves for two sample sizes, $n = 3$ and $n = 7$, whose h_ℓ values are 1.6×10^{-5} and 5.4×10^{-5} respectively. These robustness curves indeed diverge at very low robustness and very high critical safety. This divergence is limited by the fact that the nominal safeties of these two configurations—the points of intersection with the S_c axis—are both close to unity, and by the fact that the h_ℓ values themselves are small. The nominal safeties, $S(\widetilde{f})$, for $n = 3$ and $n = 7$ are 0.999 999 and 0.999 995 respectively. If one requires very large safety, say in excess of 0.9999, then one faces the fact that the robustnesses for these values are exceedingly small precisely because h_ℓ is small. The robustnesses of different sample sizes also cannot differ very much, as seen in fig. 3, again because they are bounded above by h_ℓ which is so small.

Similar tightly-constrained divergence of robustness curves is observed when comparing different values of other design parameters such as level of significance, α, or augmented penalty load x_a.

In conclusion, a method for experimental assessment of reliability of a system with a stringent safety requirement has been demonstrated. The focus has been on analysis and management of both statistical variability and distributional uncertainty. Info-gap theory has been used to model the distributional uncertainty in the pdf of actuation. The quantitative analysis and the qualitative judgments which accompany the certification of safety have been illustrated through an example. The example studied here—determination of safety of explosive material—involves augmented loads. The use of augmented loads is closely related to accelerated lifetime testing, and the methodology developed here can be adapted and applied to these and other applications. It has been proven, for the class of problems formulated here, that the info-gap robustness function is independent of the experimental design over virtually all of its range.

Acknowledgement

The authors are indebted to Miriam Zacksenhouse for valuable comments.

6 References

1. MIL-HDBK-83578, Department of Defense Handbook Criteria for Explosive Systems and Devices Used on Space Vehicles, Jan. 1999, Superseding DOD-E-83578.

2. Porat Z., Haim, M. and Markiewicz, Y., 1994, Estimation of $P(x > y)$ or $P(x < y)$, *IEEE Tran. on Reliability,* vol.43, No.3, pp.466–469.

3. Alouaamari, M., Lefebvre, M.H. and Perneel, C., 2007, Statistical Assessment Methods for the Sensitivity of Energetic Materials, 10th Seminar on New Trends in Research of Energetic Materials, University of Pardubice, Pardubice, Czech Republic, April 25–27, 2007, p.60.

4. Nance, Douglas V., 2008, Analysis of Sensitivity Experiments—A Primer, Technical Report, Air Force Research Laboratory, Munitions Directorate, AFRL/RWAC, November 2008.

5. Ayres, J.N., Hampton, L.D., Kabik, I. and Solem, A.D., 1961, *Varicomp, A Method for Determining Detonation Transfer Probabilities,* NavWeps Report 7411, U.S. Naval Ordnance Laboratory, White Oak MD.

6. Tzidony, D. and Jaeger, M., 1998, Detonation pressure of donors as a parameter for assessing detonation transfer probabilities, propellants, *Explosives, Pyrotechnics,* Vo.23, Issue 1, pp.14–16.

7. De Yong, Leo V., 1986, *Prediction of Ignition Transfer Reliability in Pyrotechnic Systems Using the Varicomp Technique,* Australian Dept. of Defense, Materials Research Laboratories, Ascot Vale, Victoria.

8. Ben-Haim, Yakov, 2006, *Info-gap Decision Theory: Decisions Under Severe Uncertainty,* 2nd edition, Academic Press, London.

9. Ayyub, Bilal M., 2001, *Elicitation of Expert Opinions for Uncertainty and Risks,* CRC Press.

10. Meyer, Mary A. and Jane M. Booker, 2001, *Eliciting and Analyzing Expert Judgment: A Practical Guide,* ASA-SIAM, reprint of 1991 edition by Academic Press.

Figure and Table Captions

Figure 1. Nominal pdf $\widetilde{f}(x)$ (solid) and fat-tail pdf $f(x)$ with 2% bump ($h = 0.02$) (dash).

Figure 2. Robustness curve, $\widehat{h}(S_{\rm c})$ vs. $S_{\rm c}$. $x_{\rm s} = -0.4$, $x_{\rm a} = 1$, $c = 0.3$, $\alpha = 0.1$, $n = 7$.

Figure 3. Robustness curves, $\widehat{h}(S_{\rm c})$ vs. $S_{\rm c}$, for two sample sizes. $x_{\rm s} = -0.4$, $x_{\rm a} = 1$, $c = 0.3$, $\alpha = 0.1$.

Table 1. Parameters for interpreting the robustness values.

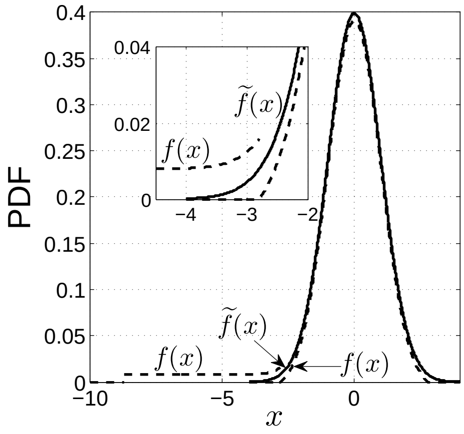

Figure 1: Nominal pdf $\widetilde{f}(x)$ (solid) and fat-tail pdf $f(x)$ with 2% bump ($h = 0.02$) (dash).

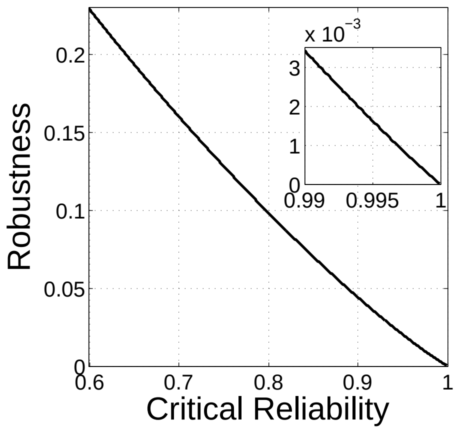

Figure 2: Robustness curve, $\widehat{h}(S_{\rm c})$ vs. $S_{\rm c}$. $x_{\rm s} = -0.4$, $x_{\rm a} = 1$, $c = 0.3$, $\alpha = 0.1$, $n = 7$.

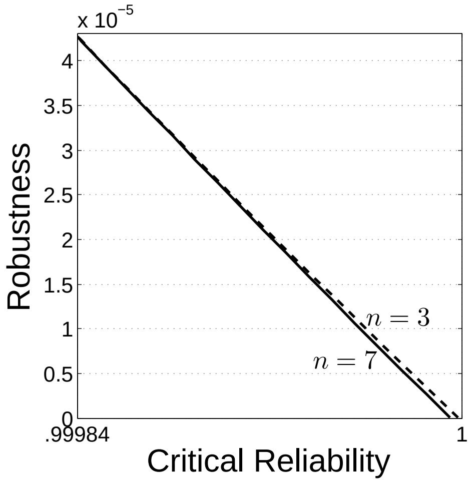

Figure 3: Robustness curves, $\widehat{h}(S_{\rm c})$ vs. $S_{\rm c}$, for two sample sizes. $x_{\rm s} = -0.4$, $x_{\rm a} = 1$, $c = 0.3$, $\alpha = 0.1$.