Distributions for Uncertainty Analysis[1]

Howard Castrup, Ph.D.
President, Integrated Sciences Group
Bakersfield, CA 93306
hcastrup@isgmax.com

Abstract

In performing a measurement, we encounter errors or biases from a number of sources. Such sources include random error, measuring parameter bias, measuring parameter resolution, operator bias, environmental factors, etc. We estimate the uncertainties due to these errors either by computing a standard deviation from a sample of measurements or by forming an estimate based on experience. Estimates obtained by the former method are labeled Type A estimates and those obtained by the latter method are called Type B estimates.

This paper describes statistical distributions that can be applied to both Type A and Type B measurement errors and to equipment parameter biases. Once the statistical distribution for a measurement error or bias is characterized, the uncertainty in this error or bias is computed as the standard deviation of the distribution. For Type A estimates, the distribution or "population" standard deviation is estimated by the sample standard deviation. For Type B estimates, the standard deviation is computed from limits, referred to as *error containment limits* and from probabilities, referred to as *containment probabilities*. The degrees of freedom for each uncertainty estimate can often be determined, regardless of whether the estimate is Type A or Type B.

Background

Until the publication of the Guide to the Expression of Uncertainty in Measurement (GUM) [1], accrediting bodies or auditing agencies for test and calibration organizations did not tend to focus on uncertainty analysis requirements. There were two main reasons for this: (1) a universally accepted methodology was not available, and (2) assessors and auditors did not possess the required expertise. Since the introduction of the GUM, however, accrediting bodies have been increasingly insistent that laboratories implement procedures for uncertainty analysis and be able to demonstrate that these procedures are being competently followed. Since the publication of ISO/IEC 17025 [2], this insistence has intensified. This has placed accrediting bodies and laboratories alike in a "catch-up" mode that has led to some hastily contrived measures, as will be discussed presently.

To induce organizations to estimate uncertainties, it was felt necessary by some to advocate the use of simple algorithms that, while they were not appropriate in most cases, would at least get people on the uncertainty analysis path.

One such algorithm involves the indiscriminate use of the uniform distribution to compute Type B uncertainty estimates. Unfortunately, organizations that not only want to analyze uncertainties but also do the job correctly are sometimes penalized by this ill-advised simplification. On one occasion, a laboratory assessor admitted that the uniform distribution was largely inappropriate but insisted that it still be employed. His reasoning was that it did not matter if uncertainty estimates were invalid as long as everyone produced them in the same way!

This philosophy precludes the development of uncertainty estimates that can be used to perform statistical tests, evaluate measurement decision risks, manage calibration intervals, develop meaningful tolerances and compute viable confidence limits. In other words, apart from providing a number, the uncertainty estimate becomes a useless and potentially expensive commodity.

Obviously, if viable uncertainty estimates are to be produced, the blind acceptance of inappropriate distributions is to be discouraged. Accordingly, we need to elaborate on alternative distributions and discuss the applicability of each

Introduction

Error and Uncertainty

It is axiomatic that the uncertainty in a value obtained by measurement is identical to the uncertainty in the measurement error. Additionally, the uncertainty in the value of a toleranced parameter or a characterized

---

[1] Presented at the 2001 IDW Conference, Knoxville, TN. Revised 27 May 2004, to correct a typographical error in the cubic equation for the quadratic distribution. Revised 11 April 2007 to provide a more tractable form of the lognormal distribution.

reference standard is equal to the uncertainty in the parameter's deviation from its nominal or stated value.

This axiom can be stated mathematically. The notation is the following

- X - the true value of an attribute
- x - a value obtained for the attribute by measurement or the attribute's characterized or nominal value
- ε_x - the error in measurement or deviation from a nominal or characterized value
- U - a mathematical operator that returns the uncertainty in a value
- u_x - the uncertainty in x
- u_{ε_x} - the uncertainty in ε_x.

We begin by saying that

Measured Value = True Value + Measurement Error,

for measured quantities, and

True Value = Nominal Value + Deviation,

for toleranced parameters or characterized reference standards. We now rewrite these expressions using the notation defined above

$$x = X + \varepsilon_x , \qquad (1)$$

for a measured attribute, and

$$X = x + \varepsilon_x , \qquad (2)$$

for a toleranced parameter or characterized standard. Using the uncertainty operator U, we obtain

$$u_x = U(x) = U(X + \varepsilon_x) = U(\varepsilon_x) = u_{\varepsilon_x} , \qquad (3)$$

for a measured attribute, and

$$u_X = U(X) = U(x + \varepsilon_x) = U(\varepsilon_x) = u_{\varepsilon_x} , \qquad (4)$$

for a toleranced parameter or characterized reference. In either case, the uncertainty in the value of interest is equal to the uncertainty in the error or deviation in the value.

Uncertainty Definition

We will now define the operator U. First, however, we need to discuss the nature of measurement errors and deviations. We begin by stating that measurement errors and deviations are random variables that follow statistical distributions.

For certain kinds of error, such as random error, this is easily seen. For other kinds of error, such as parameter bias and operator bias, however, their random nature is not so readily perceived. What we need to bear in mind is that, while a particular error may have a systematic value that persists from measurement to measurement, it nevertheless comes from some distribution of like errors that can be described statistically.

For instance, the diameters of ball bearings emerging from a manufacturing process will vary to some finite amount from bearing to bearing. If one such bearing comes into our possession, it will have a systematic deviation from nominal that is essentially fixed. However, our particular deviation was drawn at random from a population of deviations arising from the manufacturing process. Since this deviation is unknown, we can treat it as a random variable whose uncertainty is a measure of the spread of deviations that characterize the process. The wider this spread, the greater the uncertainty.

A similar chain of reasoning applies to parameters emerging from a test or calibration process and to errors in measurement.

The upshot is that, whether a particular error is random or systematic, it can still be regarded as coming from a distribution of errors that can be described statistically. Moreover, the spread in this distribution is synonymous with the uncertainty in the error. It turns out that there is an ideal statistic for quantifying this spread. This statistic is the standard deviation of the distribution.

Therefore, to define the operator U, we need to define the standard deviation. First, however, we will define the concept of *statistical variance*. Simply put, the variance of a distribution of errors is the distribution's *mean square error*. If $f(x)$ represents the probability density for a population of attribute values or measurement results, and μ_x represents the nominal or mean or value for the population, then the population variance or mean square error var(ε_x) is given by

$$\begin{aligned} u_{\varepsilon_x}^2 &= \mathrm{var}(\varepsilon_x) \\ &= \int_{-\infty}^{\infty} f(\varepsilon_x)\varepsilon_x^2 d\varepsilon_x = \int_{-\infty}^{\infty} f(x)(x-\mu_x)^2 dx \\ &= \mathrm{var}(x) \\ &= u_x^2 . \end{aligned} \qquad (5)$$

Notice that the population variance is a statistic that quantifies the spread of the distribution. That is, the larger the spread, the larger the variance. At first glance, the variance or mean square error would seem to be a good quantity by which to express a population's uncertainty. However, the variance is in the wrong units, namely, the desired units squared. This is rectified by taking the square root of the variance, which yields the standard deviation. Then, by Eq. (3) or (4)

$$
\begin{aligned}
u_x &= U(x) \\
&= U(\varepsilon_x) \\
&= \sqrt{\mathrm{var}(\varepsilon_x)}\,.
\end{aligned}
\tag{6}
$$

So, we see that estimating the uncertainty in measurement is an exercise in which we estimate the standard deviation of the measurement error. If we have a sample of measurements, we can estimate the standard deviation due to random error in the sample using a straightforward expression found in statistics textbooks

$$
u_x = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2}\ , \tag{7}
$$

where n is the sample size and $\bar{x}$ is the sample mean.[2] This is an example of a Type A estimate. For Type B estimates, we work from error containment limits and containment probabilities. The process is described in detail in the literature [4].

Standard and Expanded Uncertainty

To this point, the uncertainty in measurement has been equated with the standard deviation of the population of the measurement error. In the GUM, this uncertainty is called the *standard uncertainty*. If the distribution is known, and the degrees of freedom can be determined [4], the standard uncertainty can be used to develop confidence limits for an uncertainty estimate. The GUM refers to a confidence limit as an *expanded uncertainty*.[3] The factor by which a standard uncertainty is multiplied to yield an expanded uncertainty is called the *coverage factor*.

Unfortunately, in conversation, it is not always clear whether the term "uncertainty" refers to the expanded uncertainty or to the standard uncertainty. In this paper, unless otherwise indicated, it will refer to the standard uncertainty.

[2] Note the formal similarity between Eq. (7) and Eq. (5).

[3] Actually, the terms "standard uncertainty" and "expanded uncertainty" were introduced to supersede the terms "standard deviation" and "confidence limit," respectively, in cases where the degrees of freedom for an uncertainty estimate could not be determined. Before the refinement of methods for estimating degrees of freedom [4], this limitation applied almost universally to Type B estimates, and, by extension to mixed Type A-B estimates.

In this paper,
"uncertainty" = standard uncertainty

Statistical Distributions

In obtaining a Type A uncertainty estimate, we compute a standard deviation using Eq. (7). In obtaining a Type B estimate, we work from a set of bounding limits, referred to as *error containment limits* and a *containment probability*, which is the probability that errors or attribute values lie within these limits. Any one of a variety of distributions may be assumed to represent the underlying distribution of errors or deviations. In this paper, we consider the uniform, normal, lognormal, quadratic, cosine, half-cosine, U-shaped, and the Student's t distribution.

The Uniform Distribution

The uniform distribution is defined by the probability density function (pdf)

$$
f(x) = \begin{cases} \dfrac{1}{2a}, & -a \le x \le a \\ 0, & \text{otherwise}\ , \end{cases}
$$

where $\pm a$ are the limits of the distribution.

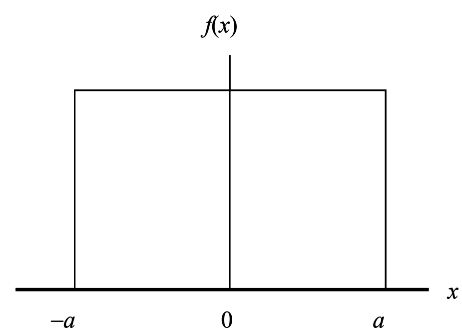

The Uniform Distribution. The probability of lying between -*a* and *a* is constant. The probability of lying outside ±*a* is zero.

Acceptance of the Uniform Distribution

Applying the uniform distribution to obtaining Type B uncertainty estimates is a practice that has been gaining ground over the past few years. There are two main reasons for this:

1. First, applying the uniform distribution makes it easy to obtain an uncertainty estimate. If the limits $\pm a$ of the distribution are known, the uncertainty estimate is just

$$
u = \frac{a}{\sqrt{3}}\,. \tag{8}
$$

It should be said that the "ease of use" advantage has been promoted by individuals who are ignorant of methods of obtaining uncertainty estimates for more appropriate distributions and by others who are simply looking for a quick solution. In fairness to the latter group, they sometimes assert that the lack of specificity of information required to use other distributions makes for crude uncertainty estimates anyway, so why not get your crude estimate by intentionally using an inappropriate distribution?

At our present level of analytical development [3, 4], this argument does not hold water. Since the introduction of the GUM, methods have been developed that systematize and rigorize the use of distributions that are physically realistic. These will be discussed presently.

2. Second, it has been asserted by some that the use of the uniform distribution is (uniformly?) recommended in the GUM. This is not true. In fact, most of the methodology of the GUM is based on the assumption that the underlying error distribution is normal. Some of the belief that the uniform distribution is called for in the GUM stems from the fact that several individuals, who have come to be regarded as GUM authorities, have been advocating its use. For clarification on this issue, the reader is referred to Section 4.3 of the GUM.

Another source of confusion is that some of the examples in the GUM apply the uniform distribution in situations that appear to be incompatible with its use. It is reasonable to suppose that much of this is due to the fact that rigorous Type B estimation methods and tools were not available at the time the GUM was published, and the uniform distribution was an "easy out." As stated in item 1 above, the lack of such methods and tools has since been rectified.

The acceptance of the uniform distribution on the basis of its use in GUM examples reminds us of a similar practice that emerged from the application of Handbook 52 to the interpretation of MIL-STD-45662A. In one example in the Handbook, a hypothetical lab was being audited whose nominal operating temperature was 68° F. Some of the 45662A auditors reacted to the example by citing labs that did not maintain this temperature, regardless of whether it was appropriate for the lab's operation. Inevitably, the 68° F requirement actually became institutionalized within certain auditing agencies.

Applicability of the Uniform Distribution

The use of the uniform distribution is appropriate under a limited set of conditions. These conditions are summarized by the following criteria.

The first criterion is that we must know a set of *minimum* bounding limits for the distribution. This is the *minimum limits criterion*. Second, we must be able to assert that the probability of finding values between these limits is unity. This is the *100% containment criterion*. Third, we must be able to demonstrate that the probability of obtaining values between the minimum bounding limits is uniform. This is the *uniform probability criterion*.

Minimum Limits Criterion. It is vital that the limits we establish for the uniform distribution are the minimum bounding limits. For instance, if the limits $\pm L$ bound the variable of interest, then so do the limits $\pm 2L$, $\pm 3L$, and so on. Since the uncertainty estimate for the uniform distribution is obtained by dividing the bounding limit by the square root of three, using a value for the limit that is not the minimum bounding value will obviously result in an invalid uncertainty estimate.

This alone makes the application of the uniform distribution questionable in estimating bias uncertainty from such quantities as tolerance limits, for instance. It may be that out-of-tolerances have never been observed for a particular parameter (100% containment), but it is unknown whether the tolerances are minimum bounding limits. Some years ago, a study was conducted involving a voltage reference that showed that values for one parameter were normally distributed with a standard deviation that was approximately 1/10 of the tolerance limit. With 10-sigma limits, it is unlikely that any out-of-tolerances would be observed. However, if the uniform distribution were used to estimate the bias uncertainty for this item, based on tolerance limits, the uncertainty estimate would be nearly six times larger than would be appropriate. Some might claim that this is acceptable, since the estimate can be considered a conservative one. That may be. However, it is also a useless estimate. This point will be elaborated later.

A second difficulty we face when attempting to apply minimum bounding limits is that such limits can rarely be established on physical grounds. This is especially true when using parameter tolerance limits. It is virtually impossible to imagine a situation where design engineers have somehow been able to precisely identify the minimum limits that bound values that are physically attainable. If we add to this the fact that tolerance limits are often influenced by marketing rather than engineering considerations, equating tolerance limits with minimum bounding limits becomes a very unfruitful and misleading practice.

100% Containment Criterion. By definition, the establishment of minimum bounding limits implies the establishment of 100% containment. It should be said however, that an uncertainty estimate may still be obtained for the uniform distribution if a containment probability less that 100% is applied. For instance, suppose the containment limits are given as $\pm L$ and the containment probability is stated as being equal to some value p between zero and one. Then, if the uniform probability criterion is met, the limits of the distribution are given by

$$a = \frac{L}{p}, \quad L \le a . \qquad (9)$$

If the uniform probability criterion is not met, however, the uniform distribution would not be applicable, and we should turn to other distributions.

Uniform Probability Criterion. As discussed above, establishing minimum containment limits can be a challenging prospect. Harder still is finding real-world measurement error distributions that demonstrate a uniform probability of occurrence between two limits and zero probability of occurrence outside these limits. Except in very limited instances, such as are discussed in the next section, assuming a uniform probability is just not physically realistic. This is true even in some cases where the distribution would appear to be applicable.

For example, a conjecture has recently been advanced that the distribution of parameters immediately following test or calibration can be said to be uniform. While this seems reasonable at face value, it turns out not to be the case. Because of false accept risk (consumer's risk), such distributions range from approximately triangular to having a "humped" appearance with rolled-off shoulders.

As to whether we can treat parameter tolerance limits as bounds that contain values with uniform probability, we must imagine that, not only has the instrument manufacturer managed to miraculously ascertain minimum bounding limits, but has also juggled physics to such an extent as to make the parameter value's probability distribution uniform between these limits and zero outside them. This would be a truly amazing feat of engineering for most toleranced quantities — especially considering the marketing influence mentioned earlier.

Cases that Satisfy the Criteria

Digital Resolution Uncertainty. We sometimes need to estimate the uncertainty due to the resolution of a digital readout. For instance, a three-digit readout might indicate 12.015 V. If the device employs the standard round-off practice, we know that the displayed number is derived from a sensed value that lies between 12.0145 V and 12.0155 V. We also can assert to a very high degree of validity that the value has an equal probability of lying anywhere between these two numbers. In this case, the use of the uniform distribution is appropriate, and the resolution uncertainty is

$$u_V = \frac{0.0005 \text{ V}}{\sqrt{3}} = 0.00029 \text{ V} .$$

RF Phase Angle. RF power incident on a load may be delivered to the load with a phase angle θ between $-\pi$ and π. In addition, unless there is a compelling reason to believe otherwise, the probability of occurrence between these limits is uniform. Accordingly, the use of the uniform distribution is appropriate. This yields a phase angle uncertainty estimate of

$$u_\theta = \frac{\pi}{\sqrt{3}} \cong 1.814 .$$

It is interesting to note that, given the above, if we assume that the amplitude of the signal is sinusoidal, the distribution for incident voltage is the U-shaped distribution.

Quantization Error. The potential drop (or lack of a potential drop) sensed across each element of an A/D Converter sensing network produces either a "1" or "0" to the converter. This response constitutes a "bit" in the binary code that represents the sampled value. For ladder-type networks, the position of the bit in the code is determined by the location of its originating network element.

Even if no errors were present in sampling and sensing the input signal, errors would still be introduced by the discrete nature of the encoding process. Suppose, for example, that the full scale signal level (dynamic range) of the A/D Converter is a volts. If n bits are used in the encoding process, then a voltage V can be resolved into 2^n discrete steps, each of size $a/2^n$. The error in the voltage V is thus

$$\varepsilon(V) = V - m\frac{a}{2^n} ,$$

where m is some integer determined by the sensing function of the D/A Converter.

The containment limit associated with each step is one-half the value of the magnitude of the step. Consequently, the containment limit inherent in quantizing a voltage V is $(1/2)(a/2^n)$, or $a/2^{n+1}$. This is embodied in the expression

$$V_{quantized} = V_{sensed} \pm \frac{a}{2^{n+1}}.$$

The uncertainty due to quantization error is obtained from the containment limits and from the assumption that the sensed analog value has equal probability of occurrence between these limits:

$$u_V = \frac{a/2^{n+1}}{\sqrt{3}}.$$

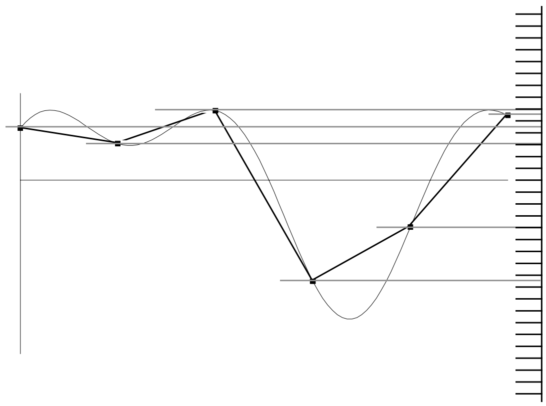

Signal Quantization. The sampled signal points are quantized in multiples of a discrete step size.

Development of Expanded Uncertainty Limits

NIST Technical Note 1297 [6] documents the uncertainty analysis policy to be followed by NIST. In this policy, expanded uncertainty limits for Type B and mixed estimates are obtained by multiplying the uncertainty estimate by a fixed "*k*-factor" equal to two. Assuming an underlying normal distribution, this produces limits that are roughly analogous to 95% confidence limits. The advisability of this practice is debatable, but this is the subject of a separate discussion. For the present, we consider what results from the practice when estimating an uncertainty for a case where the underlying distribution is assumed to be uniform.

Since the uncertainty is estimated by dividing the distribution minimum bounding limit by the square root of three, multiplying this estimate by two yields expanded uncertainty limits that are *outside* the distribution's minimum bounding limits. To be specific, these limits equate to approximately 115% containment probability, which is nonsense.

One way of reconciling the practice is to state that the underlying distribution is actually normal, or approximately normal, and the uniform distribution is used merely as an artifice to obtain an estimate of the distribution's standard deviation. This is a somewhat amazing statement. If the underlying distribution is normal, why not obtain the uncertainty estimate using that distribution in the first place? [4]

It can be shown that using the uniform distribution as a tool for estimating the uncertainty in a normally distributed quantity corresponds to assuming a normal distribution with a 91.67% containment probability. For organizations that maintain a high in-tolerance probability at the unit level, we often see or can surmise 98% or better in-tolerance probabilities at the parameter level. Consequently, for these cases, use of the uniform distribution produces uncertainty estimates that are at least 35% larger than what is appropriate.

As for those who find this acceptable on the basis of conservatism, consider the U.S. Navy's end-of-period reliability target of 72% for general purpose items. For single-parameter items, if the true underlying distribution is normal, use of the uniform distribution can produce uncertainty estimates that are only about 62% of what they should be. So much for conservatism.

The Normal Distribution

When obtaining a Type A estimate, we compute a standard deviation from a sample of values. For example, we estimate random uncertainty by computing

[4] One recommendation that the reader may encounter is that, if all that is available for an error source or parameter deviation is a set of bounding limits, without any knowledge of the nature of the error distribution and with no information regarding a containment probability, then the uniform distribution should be assumed. There are two points that should be made concerning this recommendation.

First, after a little reflection on the difficulty of obtaining minimum containment limits without knowledge of a containment probability, we can see that the recommendation not advisable. The prudent path to follow is to simply put some effort into obtaining a containment probability estimate and ascertaining a most likely underlying distribution. There is really no way around this. Moreover, the author has yet to observe an uncertainty analysis problem where this could not be done.

The second point is that, experienced technical personnel nearly always know something about what they are measuring and what they are measuring it with. Except for the cases described above, it is difficult to imagine a scenario where an experienced engineer or technician would know a set of bounding limits and nothing else.

the standard deviation for a sample of repeated measurements of a given value. We also obtain a sample size. The sample standard deviation, equated with the random uncertainty of the sample, is an estimate of the standard deviation for the population from which the sample was drawn. Except in rare cases, we assume that this population follows the normal distribution.

This assumption, allows us to easily obtain the degrees of freedom and the sample standard deviation and to construct confidence limits, perform statistical tests, estimate measurement decision risk and to rigorously combine the random uncertainty estimate with other Type A uncertainty estimates.

Why do we assume a normal distribution? The primary reason is because this is the distribution that either represents or approximates what we frequently see in the physical universe. It can be derived from the laws of physics for such phenomena as the diffusion of gases and is applicable to instrument parameters subject to random stresses of usage and handling. It is also often applicable to equipment parameters emerging from manufacturing processes.

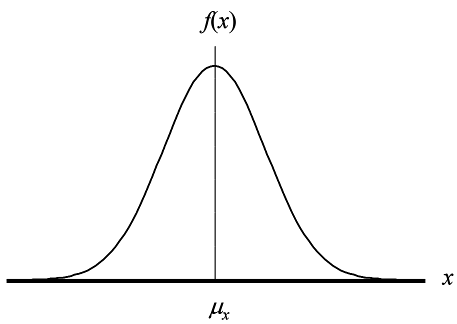

The Normal Distribution. Shown is a case where the population mean μ is located far from a physical limit 0. In such cases, the normal distribution can be used without compromising rigor.

An additional consideration applies to the distribution we should assume for a total error or deviation that is composed of constituent errors or deviations. There is a theorem called the *central limit theorem* that demonstrates that, even though the individual constituent errors or deviations may not be normally distributed, the combined error or deviation is approximately so.

An argument has been presented against the use of the normal distribution in cases where the variable of interest is restricted, i.e., where values of the variable are said to be bound by some physical limit. This condition notwithstanding, the normal distribution is still widely applicable in that, for many such cases, the physical limit is located far from the population mean. In cases where this is not so, other distributions, such as the lognormal distribution can be applied.

Uncertainty Estimates

In applying the normal distribution, an uncertainty estimate is obtained from containment limits and a a containment probability. The use of the distribution is appropriate in cases where the above considerations apply and the limits and probability are at least approximately known.

The extent to which this knowledge is approximate determines the degrees of freedom of the uncertainty estimate [4, 7]. The degrees of freedom and the uncertainty estimate can be used in conjunction with the Student's t distribution (see below) to compute confidence limits.

Let $\pm a$ represent the known containment limits and let p represent the containment probability. Then an estimate of the standard deviation of the population of errors or deviations is obtained from

$$u = \frac{a}{\Phi^{-1}\left(\frac{1+p}{2}\right)}, \tag{10}$$

where $\Phi^{-1}(\cdot)$ is the inverse normal distribution function. This function can be found in statistics texts and in popular spreadsheet programs.

If only a single containment limit is applicable, such as with single-sided tolerances, the appropriate expression is

$$u = \frac{a}{\Phi^{-1}(p)}. \tag{11}$$

The Lognormal Distribution

The lognormal distribution can often be used to estimate the uncertainty in equipment parameter bias in cases where the tolerance limits are asymmetric. It is also used in cases where a physical limit is present that lies close enough to the nominal or mode value to skew the parameter bias pdf in such a way that the normal distribution. is not applicable.

The pdf is given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma\,|\,x-q\,|}\exp\left\{-\left[\ln\left(\frac{x-q}{m-q}\right)\right]^2 \Big/ 2\sigma^2\right\},$$

where q is a physical limit for x, m is the population median and μ is the population mode. The variable σ is not the population standard deviation. It is referred to as the "shape parameter." The accompanying graphic shows a case where $\mu = 10$, $q = 9.6207$, $\sigma = 0.52046$, and $m = 10.8011$. The computed standard deviation for this example is $u = 0.3176$.

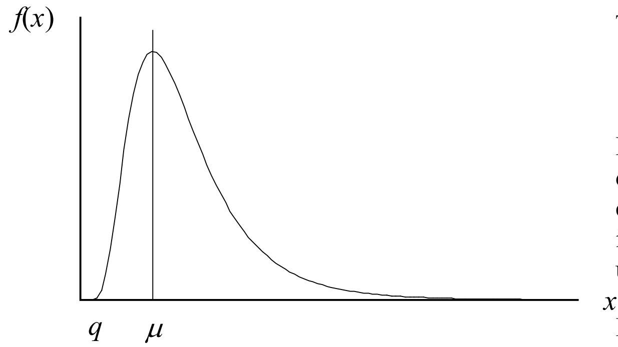

The Lognormal Distribution. Useful for describing distributions for parameters constrained by a physical limit or possessing asymmetric tolerances.

Uncertainty estimates (standard deviations) for the lognormal distribution are obtained by numerical iteration. To date, the only known applications that perform this process are UncertaintyAnalyzer [3] and AccuracyRatio [5].

The Triangular Distribution

The triangular distribution has been proposed for use in cases where the containment probability is 100%, but there is a central tendency for values of the variable of interest [1]. The triangular distribution is the simplest distribution possible with these characteristics.

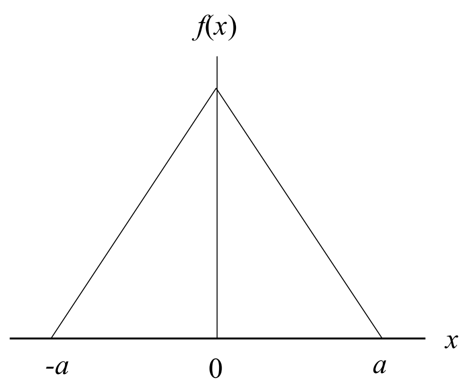

The Triangular Distribution. A distribution that sometimes applies to parameter values immediately following test or calibration.

The pdf for the distribution is

$$f(x) = \begin{cases} (x+a)/a^2, & -a \le x \le 0 \\ (a-x)/a^2, & 0 \le x \le a \\ 0, & \text{otherwise.} \end{cases}$$

The standard deviation for the distribution is obtained from

$$u = \frac{a}{\sqrt{6}} . \qquad (12)$$

Like the uniform distribution, using the triangular distribution requires the establishment of minimum containment limits $\pm a$. The same reservations apply in this regard to the triangular distribution as to the uniform distribution.

In cases where a containment probability $p < 1$ can be determined for limits $\pm L$, where $L < a$, the limits of the distribution are given by

$$a = \frac{L}{1-\sqrt{1-p}}, \quad L \le a .$$

Apart from representing post-test distributions under certain restricted conditions, the triangular distribution has limited applicability to physical errors or deviations. While it does not suffer from the uniform probability criterion, as does the uniform distribution, it nevertheless displays abrupt transitions at the bounding limits and at the zero point, which are physically unrealistic in most instances. In addition, the linear increase and decrease in behavior is somewhat fanciful for a pdf.

The Quadratic Distribution

A distribution that eliminates the abrupt change at the zero point, does not exhibit unrealistic linear behavior and satisfies the need for a central tendency is the quadratic distribution. This distribution is defined by the pdf

$$f(x) = \begin{cases} \frac{3}{4a}\left[1-(x/a)^2\right], & -a \le x \le a \\ 0, & \text{otherwise} \end{cases}$$

where $\pm a$ are minimum bounding limits. The standard deviation for this distribution is determined from

$$u = \frac{a}{\sqrt{5}} , \qquad (13)$$

i.e., about 77% of the standard deviation estimate for the uniform distribution.

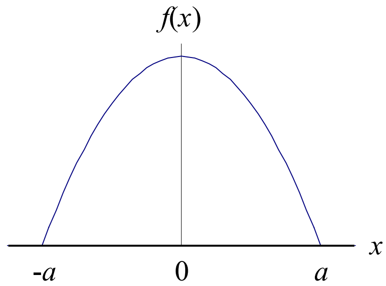

The Quadratic Distribution. Exhibits a central tendency without discontinuities and does not assume linear pdf behavior.

For a containment probability p and containment limits $\pm L$, the minimum bounding limits $\pm a$ are obtained from

$$a = \frac{L}{2p}\left(1 + 2\cos\left[\frac{1}{3}\arccos(1-2p^2)\right]\right) \quad -1 < p < 1 .$$

The Cosine Distribution

While the quadratic distribution eliminates discontinuities within the bounding limits, it rises abruptly *at* the limits. Although the quadratic distribution has wider applicability than either the triangular or uniform distribution, this feature nevertheless diminishes its physical validity. A distribution that overcomes this shortcoming, exhibits a central tendency and can be determined from minimum containment limits is the cosine distribution. The pdf for this distribution is given by

$$f(x) = \begin{cases} \frac{1}{2a}\left[1 + \cos\left(\frac{\pi x}{a}\right)\right], & -a \le x \le a . \\ 0, & \text{otherwise} \end{cases}$$

The uncertainty is obtained from the expression

$$u = \frac{a}{\sqrt{3}}\sqrt{1 - \frac{6}{\pi^2}} , \qquad (14)$$

which translates to roughly 63% of the value obtained using the uniform distribution.

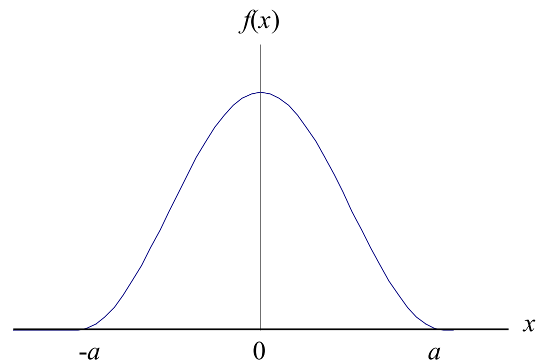

The Cosine Distribution. A 100% containment distribution with a central tendency and lacking discontinuities.

Solving for a when a containment probability and containment limits $\pm L$ are given requires applying numerical iterative method to the expression

$$\frac{1}{\pi}\sin(\pi x) - p + x = 0, \quad x \equiv L/a; \quad L \le a .$$

The solution algorithm has been implemented in the same software alluded to in the discussion on the quadratic distribution. It yields, for the ith iteration,

$$x_i = x_{i-1} - F/F' ,$$

where

$$F = \frac{1}{\pi}\sin(\pi x) - p + x$$

and

$$F' = 1 + \cos(\pi x) .$$

The Half-Cosine Distribution

The half-cosine distribution is used in cases where the central tendency is not as pronounced as when normal or the cosine distribution would be appropriate. In this regard, it resembles the quadratic distribution without the discontinuities at the distribution limits. The pdf is

$$f(x) = \begin{cases} \frac{\pi}{4a}\cos\left(\frac{\pi x}{2a}\right) , & a \le x \le a . \\ 0, & \text{otherwise} \end{cases}$$

If the minimum limiting values $\pm a$ are known, the uncertainty is obtained from the expression

$$u = \sqrt{1 - 8/\pi^2}\, a . \qquad (15)$$

If containment limits $\pm L$ and a containment probability p are known, the limiting values may be obtained from the relation

$$a = \frac{\pi L}{2\sin^{-1}(p)}, \quad L \le a .$$

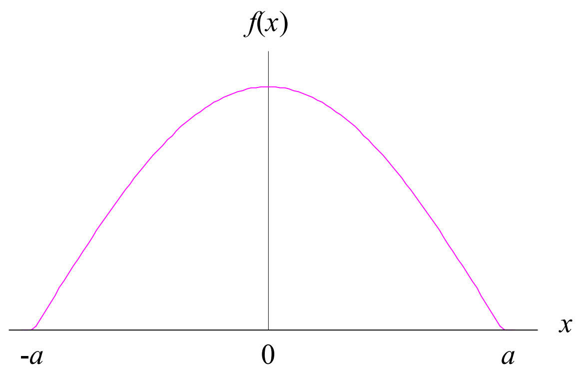

The Half-Cosine Distribution. Possesses a central tendency but exhibits a higher probability of occurrence near the minimum limiting values than either the cosine or the normal distribution.

The U Distribution

The U distribution applies to sinusoidal RF signals incident on a load. It has the pdf

$$f(x) = \begin{cases} \dfrac{1}{\pi\sqrt{a^2 - x^2}}, & -a < x < a \\ 0, & \text{otherwise}, \end{cases}$$

where a represents the maximum signal amplitude. The uncertainty in the incident signal amplitude is estimated according to

$$u = \frac{a}{\sqrt{2}}\,. \qquad (16)$$

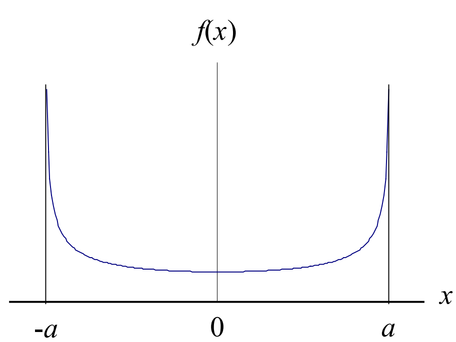

The U Distribution. The distribution is the pdf for sine waves of random phase incident on a plane.

If containment limits $\pm L$ and a containment probability p are known, the parameter a can be computed according to

$$a = \frac{L}{\sin(\pi p/2)}, \quad L \le a\,.$$

The Student's t Distribution

If the underlying distribution is normal, and a Type A estimate and degrees of freedom are available, confidence limits for measurement errors or parameter deviations may be obtained using the Student's t distribution. This distribution is available in statistics textbooks and popular spreadsheet applications. Its pdf is

$$f(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\pi\nu}\,\Gamma\left(\frac{\nu}{2}\right)}(1 + x^2/\nu)^{-(\nu+1)/2},$$

where ν is the degrees of freedom and $\Gamma(\cdot)$ is the gamma function.

The degrees of freedom quantifies the amount of knowledge used in estimating uncertainty. This knowledge is incomplete if the limits $\pm a$ are approximate and the containment probability p is estimated from recollected experience. Since the knowledge is incomplete, the degrees of freedom associated with a Type B estimate is not infinite. If the degrees of freedom variable is finite but unknown, the uncertainty estimate cannot be rigorously used to develop confidence limits, perform statistical tests or make decisions. This limitation has often precluded the use of Type B estimates as statistical quantities and has led to such discomforting artifices as fixed coverage factors.

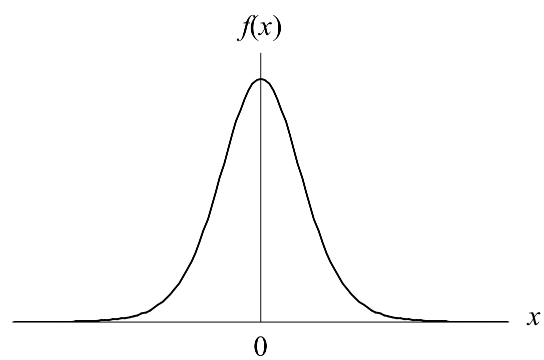

Student's t Distribution. Shown is the pdf for 10 degrees of freedom.

Fortunately, the GUM provides an expression for obtaining the approximate degrees of freedom for Type B estimates. However, the expression involves the use of the variance in the uncertainty estimate, and a method for obtaining this variance has been lacking until recently [4]. A rigorous method for obtaining this quantity has been implemented in commercially available software [3] and in a freeware application [7].

Once the degrees of freedom has been obtained, the Type B estimate may then be combined with other estimates and the degrees of freedom for the combined uncertainty can be determined using the Welch-Satterthwaite relation [1]. If the underlying distribution for the combined estimate is normal, the t distribution can be used to develop confidence limits and perform statistical tests.

The procedure is to first estimate the uncertainty using Eq. (10) and then estimate the degrees of freedom from the expression

$$\begin{aligned} \nu_B &\cong \frac{1}{2}\left(\frac{\sigma^2(u)}{u^2}\right)^{-1} \\ &\cong \frac{3\varphi^2 a^2}{2\varphi^2(\Delta a)^2 + \pi a^2 e^{\varphi^2}(\Delta p)^2}, \end{aligned} \tag{17}$$

where

$$\varphi = \Phi^{-1}\left(\frac{1+p}{2}\right).$$

The variables Δa and Δp represent "give or take" values for the containment limits and containment probability, respectively.

At first glance, Eq. (17) may seem to be anything but rigorous. However, several data input formats have been developed that rigorize the process of estimating Δa and Δp [4]. They are available in the referenced software applications cited above [3, 7].

Striving for Conservative Estimates

If an uncertainty estimate is viewed as an end product that will be filed away without application of any kind, then employing unrealistic distributions and fixed coverage factors may be considered acceptable by some. Such distributions can yield statistically valid estimates, regardless of whether or not these estimates are physically valid.

However, if an uncertainty estimate is to be employed in making decisions, such as may result from hypothesis testing or decision risk analyses, employing a physically unrealistic distribution is to be discouraged. In these cases, advocating the use of such a distribution on the grounds that it yields conservative uncertainty estimates is as irresponsible as employing intentionally biased instruments to obtain measurements that are favorably skewed in one direction or another.

In addition, the use of unrealistic distributions may yield estimates that are considerably *smaller* than what is appropriate under certain conditions. The example of estimating bias uncertainty for single-parameter Navy general purpose items, mentioned earlier, is a case in point.

Another consideration that argues against employing conservative uncertainty estimates is that this practice sometimes leads to "reckless" conclusions. This is the case when a measurement from one laboratory is tested against a measurement from another to assess equivalence between laboratories. If conservative estimates are used, the test actually becomes *less* stringent than otherwise.

The bottom line is that conservative uncertainty estimates are essentially zero-information quantities that have no legitimate use. If conservatism is desired, it can be implemented by insisting on high confidence levels in estimating confidence limits *after* a valid uncertainty estimate is obtained. The higher the confidence level, the wider (more conservative) the confidence limits.

Recommendations for Selecting Distributions

Unless information to the contrary is available, the normal distribution should be applied as the default distribution. For Type B estimates, the data input formats alluded under the discussion of the Student's t distribution should also be employed to estimate the degrees of freedom. If it is suspected that the distribution of the value of interest is skewed, apply the lognormal distribution.

In using the normal or lognormal distribution, some effort must be made to estimate a containment probability. If a set of containment limits is available, but 100% containment has been observed, then the following is recommended:

1. If the value of interest has been subjected to random usage or handling stress, and is assumed to possess a central tendency, apply the cosine distribution. If it is suspected that values are more evenly distributed, apply either the quadratic or half-cosine distribution, as appropriate. The triangular distribution may be applicable, under certain circumstances, when dealing with parameters following testing or calibration.
2. If the value of interest is the amplitude of a sine wave incident on a plane with random phase, apply the U distribution.
3. If the value of interest is the resolution uncertainty of a digital readout, apply the uniform distribution. This distribution is also applicable to estimating the uncertainty due to quantization error and the uncertainty in RF phase angle.

General Procedure for Obtaining Uncertainty Estimates

Type A Estimates

In making a Type A estimate and using it to construct confidence limits, we apply the following procedure taken from the GUM and elsewhere:

1. Take a random sample of size n representative of the population of interest. The larger the sample size, the better. In many cases, a sample size less than six is not sufficient.
2. Compute a sample standard deviation, u using Eq. (7).
3. Assume an underlying distribution, e.g., normal.
4. Develop a coverage factor based on the degrees of freedom ($n - 1$) associated with the sample standard deviation and a desired level of confidence. If the underlying distribution is assumed to be normal, use either t-tables or Student's t spreadsheet functions. In Microsoft Excel, for example, a two-sided coverage factor can be determined using the TINV function: $t = \text{TINV}((1-p),\nu)$, where p is the confidence level and ν is the degrees of freedom.
5. Multiply the sample standard deviation by the coverage factor to obtain $L = tu$ and use $\pm L$ as $p \times 100\%$ confidence limits.

Type B Estimates

In making a Type B estimate, we reverse the process. The procedure is

1. Take a set of confidence limits, e.g., parameter tolerance limits $\pm L$ (containment limits).
2. Estimate the confidence level, e.g., the in-tolerance probability (containment probability).
3. Estimate the degrees of freedom using Eq. (17).
4. Assume an underlying distribution, e.g., normal.[5]
5. Compute a coverage factor, t, based on the containment probability and degrees of freedom.
6. Compute the standard uncertainty for the quantity of interest (e.g., parameter bias) by dividing the confidence limit by the coverage factor: $u = L / t$.

[5] The Type B estimation procedure has been refined so that standard deviations can be estimated for non-normal populations and in cases where the confidence limits are asymmetric or even single-sided [3, 5].

References

[1] ISO/TAG4/WG3, *Guide to the Expression of Uncertainty in Measurement*, International Organization for Standardization (ISO), Geneva, 1993.

[2] ISO/IEC 17025 1999(E), *General Requirements for the Competence of Testing and Calibration Laboratories*, ISO/IEC, December 15, 1999.

[3] UncertaintyAnalyzer, ©1994-1997, Integrated Sciences Group, All Rights Reserved.

[4] Castrup, H., "Estimating Category B Degrees of Freedom," *Proc. Measurement Science Conference*, January 2000, Anaheim.

[5] AccuracyRatio, © 1992-2001, Integrated Sciences Group, All Rights Reserved.

[6] Taylor, B. and Kuyatt, C., *NIST Technical Note 1297*, "Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results," U.S. Dept. of Commerce, 1994.

[7] ISG Category B Uncertainty Calculator, © 2000, Integrated Sciences Group, All Rights Reserved. Available from http://www.isgmax.com.