A Critique of the Uniform Distribution

H. Castrup, Ph.D.
President, Integrated Sciences Group
1/28/2000

Before the publication of the GUM,[1] accrediting bodies, such as A2LA, did not focus on the uncertainty analysis requirements of ISO Guide 25. As I learned first hand, this was because (1) a universally accepted methodology was not available, and (2) assessors did not possess the required expertise. Since the introduction of the GUM, however, accrediting bodies have been increasingly insistent on compliance with the Guide's uncertainty analysis requirements. In addition, requirements for performing uncertainty analysis have been accelerating due to international market pressures. As you are well aware, these requirements have been met with strenuous resistance from many test and calibration organizations.

This resistance is largely due to the fact that, in the calibration field, at least, technical tasks that would normally be performed by engineers or analysts are performed mostly by technicians without formal university training in mathematics and statistics. In the early 1990s and even today, this put any hope of obtaining rigorous and valid uncertainty analyses out of reach for many organizations.[2]

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 use of the uniform distribution to compute Category B uncertainty estimates.

As time passed, pressure increased not only to analyze uncertainties but also to employ the uniform distribution for all Category B estimates. Currently, certain individuals are even attempting to institutionalize it.

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 inappropriate but insisted that it be employed. His reasoning was that it did not matter that 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. More will be said on this later.

Obviously, blind acceptance of an inappropriate distribution is to be discouraged. Accordingly, we need to elaborate on various distributions and discuss the applicability of each. We begin by examining the uniform distribution in detail and presenting the criteria for its use.

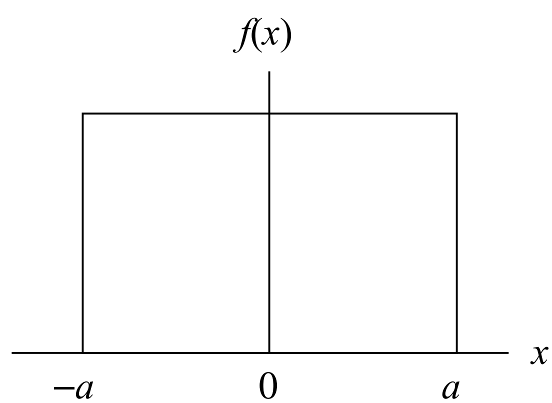

The Uniform Distribution. The probability of lying between *-a* and *a* is constant. The probability of lying outside ±*a* is zero.

Properties of the Uniform Distribution

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

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

[2] Analysis software has since been introduced that removes this obstacle.

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

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

Acceptance of the Distribution

The application of the uniform distribution to obtain Category 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 (standard deviation) estimate. If the limits $\pm a$ of the distribution are known, the uncertainty estimate is just

$$u=\frac{a}{\sqrt{3}}. \tag{1}$$

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 nonspecificity 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?

This argument no longer holds 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.

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. I believe that much of this is due to the fact that rigorous Category 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 me 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 (minimum limits criterion). Second, we must be able to assert that the probability of finding values between these limits is unity (100% containment criterion). Third, we must be able to demonstrate that the probability of obtaining values between the minimum bounding limits is uniform (uniform probability criterion).

Minimum Limits Criterion

It is vital that the limits we establish for the 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 standard deviation (uncertainty) estimate for the uniform distribution is obtained from Eq. (1) above, using a value for a that is not the minimum bounding value will obviously result in an invalid uncertainty estimate.[3]

This alone makes the application of the uniform distribution questionable in estimating bias uncertainty from tolerance limits. 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. I know of a study involving the Fluke 732 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 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 will be elaborated on 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 to a parameter. If we add to this the fact that tolerance limits are often influenced by marketing considerations, equating tolerance limits with minimum bounding limits becomes a very unfruitful and misleading enterprise.

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 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

$$a = \frac{L}{p} . \tag{2}$$

If the uniform probability criterion is not met, however, the uniform distribution would not be applicable, and we should turn to other distributions. Since effort has been expended to obtain values for L and p, this becomes feasible, as will be discussed later.

Uniform Probability Criterion

As discussed above, establishing minimum containment limits is nearly always a challenging prospect. It is child's play, however, next to 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 below, assuming uniform probability is just not physically realistic. This is true even in some cases where the distribution would appear to be applicable.

For example, an argument has been recently presented that the distribution of parameters immediately following test or calibration can be said to be uniform. This is not the case. Because of false accept risk (consumer risk), such distributions range from approximately triangular to having a "humped" appearance with rolled-off shoulders.

---

[3] By invalid, I mean physically invalid, not statistically invalid. These two varieties of validity sometimes get confused.

With regard to utilizing parameter tolerances as values for a in Eq. (1), consider the extent to which we have to stretch reality. 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. Moreover, users sometimes discard manufacturer's tolerances in favor of limits modified to meet their needs. One would have to extremely lucky to land on a set of limits that make the uniform distribution appropriate.

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 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 would be

$$\frac{0.005 \text{ V}}{\sqrt{3}} = 0.0029 \text{ V} .$$

RF Phase Angle

RF power incident on a load may be delivered to the load with a phase angle 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

$$\frac{\pi}{\sqrt{3}} \cong 1.814 .$$

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 described below.

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. The position of the bit in the code is determined by which network element originated it.

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

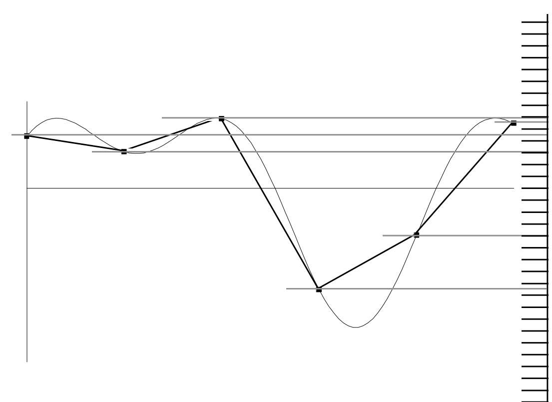

Signal Quantization. The sampled signal points are quantized in multiples of a discrete step size.

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:

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

Other Cases

None are known to me at the present time. I am open to suggestions.

Development of Expanded Uncertainty Limits

NIST Technical Note 1297[4] documents the uncertainty analysis policy to be followed by NIST. In this policy, expanded uncertainty limits for Category 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 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 uncertainty if the underlying distribution is assumed to be uniform.

Since the uncertainty is estimated by dividing the distribution minimum containment limit by the square root of 3, multiplying this estimate by two yields expanded uncertainty limits that are *outside* the distribution containment limits. To be precise, 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?

It can be shown that using the uniform distribution for this purpose 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 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 34% larger than 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, use of the uniform distribution can produce uncertainty estimates that are only about 62% of what they should be. So much for conservatism.

Alternative Distributions

The Normal Distribution

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

[4] 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.

which the sample was drawn. Except in rare cases, we assume that this population follows the normal distribution.

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

Why do we assume a normal distribution? Simply 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 applicable to equipment parameters emerging from manufacturing processes.

An argument has been presented against the use of the normal distribution in cases where the variable of interest is restricted, such as in cases where negative values are not obtainable. In these instances, the value is said to be bound by a physical limit. This condition notwithstanding, the normal distribution is still widely applicable in that, for the most part, the physical limit is located far from the population mean.

In cases where this is not so, other distributions, such as the gamma or the lognormal can be applied.

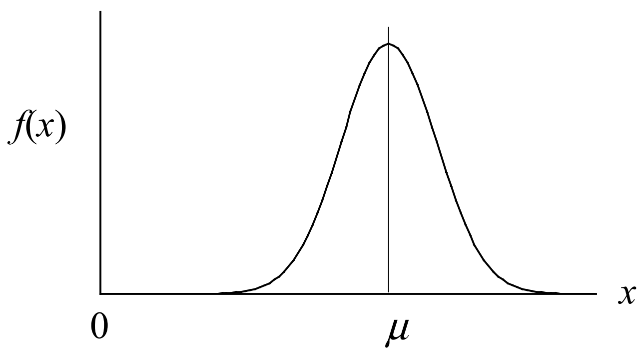

The Normal Distribution. Shown is a case where the population mean μ is located far from a physical limit 0. In such cases, the normal distribution can be used without compromising rigor.

Uncertainty Estimates

In applying the normal distribution, an uncertainty estimate is obtained from containment limits and an estimated containment probability. The use of the distribution is appropriate in cases where the above considerations apply and the limits and probability are well known. If these quantities are not well known, the Student's t distribution can be applied (see below).

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

$$u = \frac{a}{\Phi^{-1}\left(\frac{1+p}{2}\right)}, \qquad (3)$$

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

$$u = \frac{a}{\Phi^{-1}(p)}. \qquad (4)$$

The Lognormal Distribution

The lognormal distribution can often be used to estimate equipment parameter bias uncertainty 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 invalidate the use of the normal distribution.

The pdf is given by

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

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 $m = 13.1117$. The computed standard deviation for this example is $u = 0.3176$.

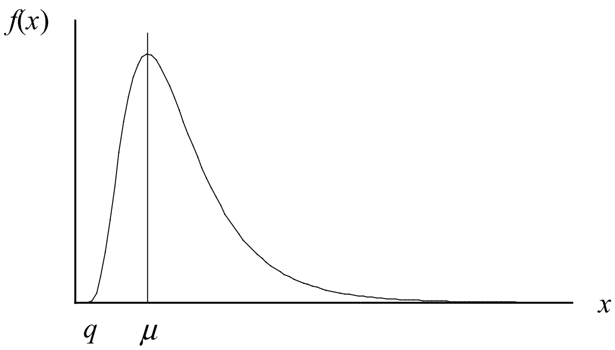

The Lognormal Distribution. Useful for describing distributions for parameters constrained by a physical limit or possessing asymmetric tolerances.

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 and AccuracyRatio, both from Integrated Sciences Group.[5]

The Triangular Distribution

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

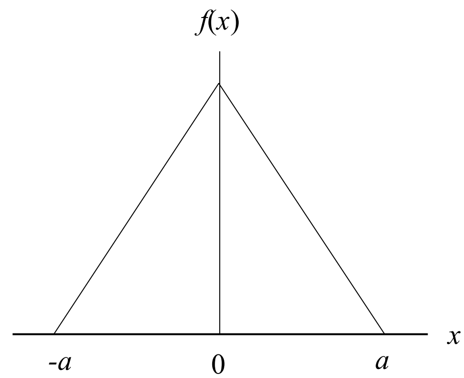

The Triangular Distribution. A distribution that sometimes applies to tested or calibrated parameter immediately following test or calibration.

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

$$u = \frac{a}{\sqrt{6}}. \qquad (5)$$

Like the uniform distribution, using the triangular distribution requires the establishment of minimum containment limits $\pm a$. The same difficulties 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

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

and the uncertainty is estimated using Eq. (5).

Apart from representing post-test distributions under certain conditions, the triangular distribution has limited applicability to physical quantities. 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

[5] See http://www.isgmax.com.

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

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

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

$$u = \frac{a}{\sqrt{5}}\,, \qquad (6)$$

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

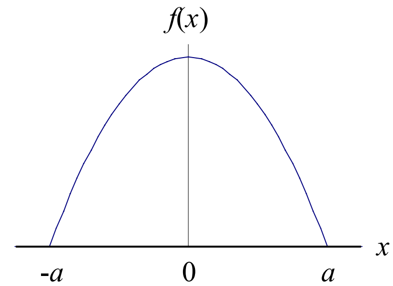

The Quadratic Distribution. Exhibits a central tendency without discontinuities and does not assume linear pdf behavior.

Solving for the limits $\pm a$ when a containment probability p and containment limits $\pm L$ are known involves solving the cubic equation

$$L^3 - 3aL + 2a^3 p = 0, \quad L \le a\ .$$

The solution can be obtained numerically. A routine has been written that solves for a, along with other parameters that represent more robust incarnations of the distribution. This capability will be available in commercial software soon to be released.

The Cosine Distribution

While the quadratic distribution eliminates discontinuities within the bounding limits, it rises abruptly *at* the limits. While the distribution is more applicable than the triangular or uniform in most cases, this feature 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

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

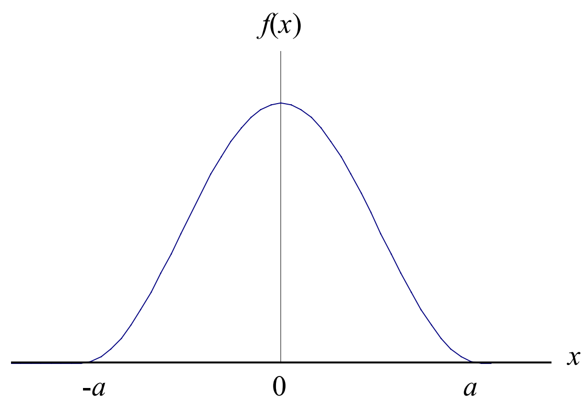

The Cosine Distribution. A 100% containment distribution with a central tendency and lacking discontinuities.

The uncertainty in the variable x is obtained using the expression

$$u = \frac{a\sqrt{1-6/\pi^2}}{\sqrt{3}}\,, \qquad (7)$$

i.e., roughly 63% of the value obtained using the uniform distribution.

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

$$\frac{a}{\pi}\sin(\pi L/a) - ap + L = 0, \quad L \le a \ .$$

The solution algorithm has been implemented in the same software alluded to in the discussion on the quadratic distribution.

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

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

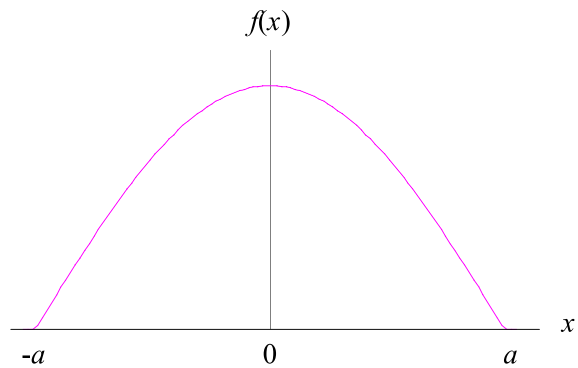

The Half-Cosine Distribution. Possesses a central tendency but exhibits a higher probability of occurrence near the minimum limiting values than either the cosine or the normal distribution.

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

$$u = \sqrt{1 - 8/\pi^2}\, a \ . \tag{8}$$

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

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

The U Distribution

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

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

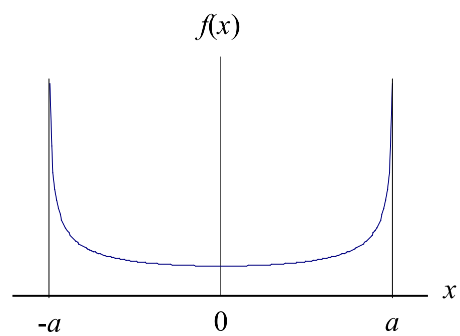

The U Distribution. The distribution is the pdf for sine waves of random phase incident on a plane.

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

$$u = \frac{a}{\sqrt{2}} \ . \tag{9}$$

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

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

The Student's t Distribution

If the underlying distribution is normal, and a Category A estimate and degrees of freedom are available, confidence limits for measured values 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 Γ is the gamma function.

The degrees of freedom quantifies the extent of the knowledge available for estimating uncertainty. This knowledge is incomplete if the limits $\pm a$ are approximate and the containment probability p is estimated from recollected experience. In such cases, the degrees of freedom associated with a Category B estimate is finite. Until recently, this variable has been unobtainable. This has precluded the use of Category B estimates as statistical quantities and has led to such discomforting artifices as fixed coverage factors.

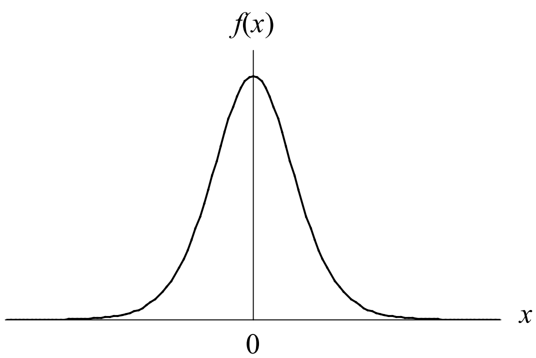

Student's t Distribution. Shown is the pdf for 10 degrees of freedom.

Fortunately, the GUM provides an expression for obtaining the approximate degrees of freedom for Category B estimates. However, the expression involves the use of the variance in the estimate, and a method for obtaining this variance has been lacking until now. A rigorous method for obtaining this quantity was presented at the last Measurement Science Conference[6] in Anaheim. The method has been incorporated into a freeware application available from http://www.isgmax.com.

Once the degrees of freedom has been obtained, the Category B estimate may be combined with other estimates and the resulting degrees of freedom can be determined using the Welch-Satterthwaite relation. 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 approach for employing the t distribution is to first estimate the uncertainty using Eq. (3) and then estimate the degrees of freedom from the expression

$$\begin{aligned} \nu_B &\cong \frac{1}{2}\left(\frac{\sigma^2(u)}{u^2}\right)^{-1} \\ &\cong \frac{3\varphi^2 a^2}{2\varphi^2(\Delta a)^2 + \pi a^2 e^{\varphi^2}(\Delta p)^2}, \end{aligned} \tag{10}$$

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. (10) may seem to be anything but rigorous. However, several data input formats have been developed that rigorize the process of estimating Δa and Δp. They are discussed in the cited paper and are available in the freeware application.

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

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 is acceptable. Such distributions can yield statistically valid estimates, regardless of whether or not these estimates are physically valid.

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 estimates is as irresponsible as employing intentionally biased instruments to obtain measurements that are favorably skewed in one direction or another.

Moreover, 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.

If conservatism is desired, it can be implemented by insisting on high degrees of confidence *after* a valid uncertainty estimate is obtained. While this was not strictly possible in the past, since the degrees of freedom for Category B estimates was unobtainable, it is now within reach.

Another consideration that argues against computing intentionally incorrect (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, if an uncertainty estimate is obtained from containment limits alone, without any consideration of containment probability it is, in effect, a zero-information estimate -- an intentionally obscured quantity that cannot be employed to evaluate measurements, compute risks, tolerance parameters, or make decisions.

Recommendations for Selecting Distributions

Unless information to the contrary is available, the normal distribution should be applied as the default distribution. If possible, 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 distributions, 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.

A general Procedure for Obtaining Category B Estimates

The following was presented by James Jenkins of Quametec, Inc. It is reproduced verbatim below.

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

1. Take a sample of data representative of the population of interest.
2. Compute a sample standard deviation, u.
3. Assume an underlying distribution, e.g., normal.
4. Develop a coverage factor (e.g., a t-statistic) based on the degrees of freedom associated with the sample standard deviation and a desired level of confidence.
5. Multiply the sample standard deviation by the coverage factor to obtain $L = tu$ and use $\pm L$ as confidence limits.

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

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

Dr. Castrup of Integrated Sciences Group has refined the above procedure so that standard deviations can be estimated for non-normal populations and in cases where the confidence limits are asymmetric or even single-sided.

He has also developed a method for estimating the "delta u" variable in Equation G.3 of the GUM so that the degrees of freedom can be computed for Category B uncertainties. Both the refinement of the procedure and the delta u method are too complicated for elaboration here. Both methodologies have been built into uncertainty analysis software. For a description of the delta u methodology, as well as a freeware application that applies it, go to http://www.isgmax.com, click on 'Category B Degrees of Freedom Estimator,' and elect to either view the methodology or download the software.