JCGM 101:2008

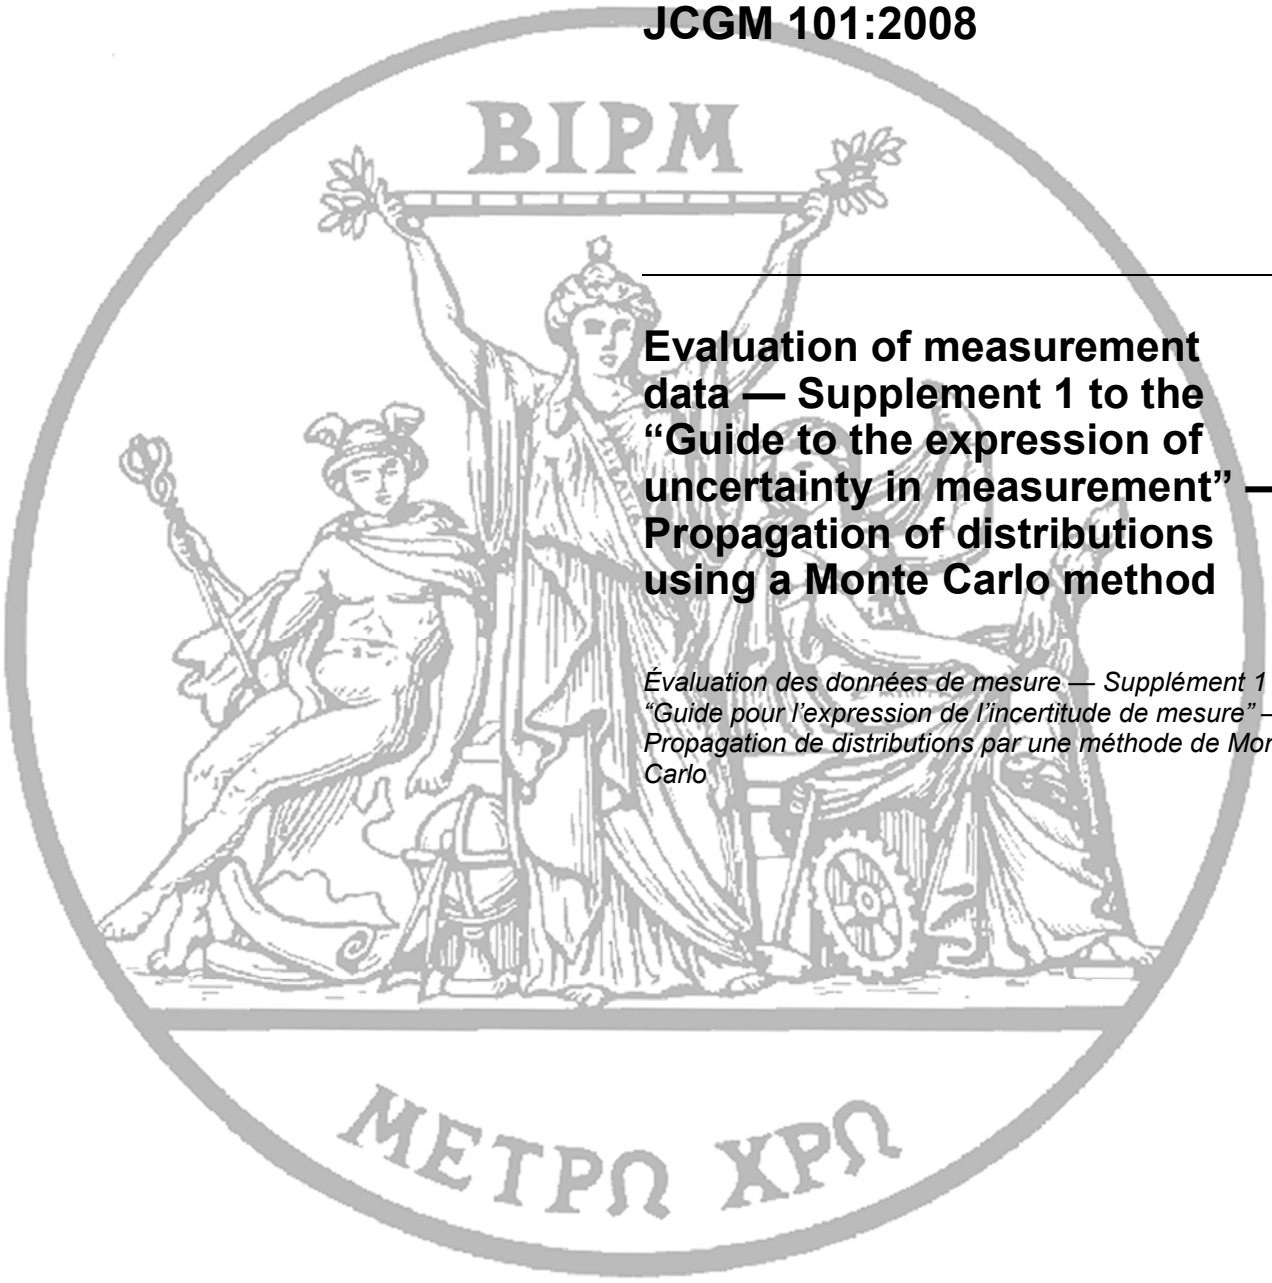

Evaluation of measurement data — Supplement 1 to the "Guide to the expression of uncertainty in measurement" — Propagation of distributions using a Monte Carlo method

Évaluation des données de mesure — Supplément 1 du "Guide pour l'expression de l'incertitude de mesure" — Propagation de distributions par une méthode de Monte Carlo

First edition 2008

Copyright

Even if electronic versions are available free of charge on the website of one or more of the JCGM member organizations, economic and moral copyrights related to all JCGM publications are internationally protected. The JCGM does not, without its written authorisation, permit third parties to rewrite or re-brand issues, to sell copies to the public, or to broadcast or use on-line its publications. Equally, the JCGM also objects to distortion, augmentation or mutilation of its publications, including its titles, slogans and logos, and those of its member organizations.

Official versions and translations

The only official versions of documents are those published by the JCGM, in their original languages.

The JCGM's publications may be translated into languages other than those in which the documents were originally published by the JCGM. Permission must be obtained from the JCGM before a translation can be made. All translations should respect the original and official format of the formulæ and units (without any conversion to other formulæ or units), and contain the following statement (to be translated into the chosen language):

> All JCGM's products are internationally protected by copyright. This translation of the original JCGM document has been produced with the permission of the JCGM. The JCGM retains full internationally protected copyright on the design and content of this document and on the JCGM's titles, slogan and logos. The member organizations of the JCGM also retain full internationally protected right on their titles, slogans and logos included in the JCGM's publications. The only official version is the document published by the JCGM, in the original languages.

The JCGM does not accept any liability for the relevance, accuracy, completeness or quality of the information and materials offered in any translation. A copy of the translation shall be provided to the JCGM at the time of publication.

Reproduction

The JCGM's publications may be reproduced, provided written permission has been granted by the JCGM. A sample of any reproduced document shall be provided to the JCGM at the time of reproduction and contain the following statement:

> This document is reproduced with the permission of the JCGM, which retains full internationally protected copyright on the design and content of this document and on the JCGM's titles, slogans and logos. The member organizations of the JCGM also retain full internationally protected right on their titles, slogans and logos included in the JCGM's publications. The only official versions are the original versions of the documents published by the JCGM.

Disclaimer

The JCGM and its member organizations have published this document to enhance access to information about metrology. They endeavor to update it on a regular basis, but cannot guarantee the accuracy at all times and shall not be responsible for any direct or indirect damage that may result from its use. Any reference to commercial products of any kind (including but not restricted to any software, data or hardware) or links to websites, over which the JCGM and its member organizations have no control and for which they assume no responsibility, does not imply any approval, endorsement or recommendation by the JCGM and its member organizations.

Contents

Annexes

Foreword

In 1997 a Joint Committee for Guides in Metrology (JCGM), chaired by the Director of the Bureau International des Poids et Mesures (BIPM), was created by the seven international organizations that had originally in 1993 prepared the "Guide to the expression of uncertainty in measurement" (GUM) and the "International vocabulary of basic and general terms in metrology" (VIM). The JCGM assumed responsibility for these two documents from the ISO Technical Advisory Group 4 (TAG4).

The Joint Committee is formed by the BIPM with the International Electrotechnical Commission (IEC), the International Federation of Clinical Chemistry and Laboratory Medicine (IFCC), the International Laboratory Accreditation Cooperation (ILAC), the International Organization for Standardization (ISO), the International Union of Pure and Applied Chemistry (IUPAC), the International Union of Pure and Applied Physics (IUPAP), and the International Organization of Legal Metrology (OIML).

JCGM has two Working Groups. Working Group 1, "Expression of uncertainty in measurement", has the task to promote the use of the GUM and to prepare Supplements and other documents for its broad application. Working Group 2, "Working Group on International vocabulary of basic and general terms in metrology (VIM)", has the task to revise and promote the use of the VIM.

Supplements such as this one are intended to give added value to the GUM by providing guidance on aspects of uncertainty evaluation that are not explicitly treated in the GUM. The guidance will, however, be as consistent as possible with the general probabilistic basis of the GUM.

The present Supplement 1 to the GUM has been prepared by Working Group 1 of the JCGM, and has benefited from detailed reviews undertaken by member organizations of the JCGM and National Metrology Institutes.

Introduction

This Supplement to the "Guide to the expression of uncertainty in measurement" (GUM) is concerned with the propagation of probability distributions through a mathematical model of measurement [GUM:1995 3.1.6] as a basis for the evaluation of uncertainty of measurement, and its implementation by a Monte Carlo method. The treatment applies to a model having any number of input quantities, and a single output quantity.

The described Monte Carlo method is a practical alternative to the GUM uncertainty framework [GUM:1995 3.4.8]. It has value when

a) linearization of the model provides an inadequate representation, or

b) the probability density function (PDF) for the output quantity departs appreciably from a Gaussian distribution or a scaled and shifted t-distribution, e.g. due to marked asymmetry.

In case a), the estimate of the output quantity and the associated standard uncertainty provided by the GUM uncertainty framework might be unreliable. In case b), unrealistic coverage intervals (a generalization of "expanded uncertainty" in the GUM uncertainty framework) might be the outcome.

The GUM [GUM:1995 3.4.8] "... provides a framework for assessing uncertainty ...", based on the law of propagation of uncertainty [GUM:1995 5] and the characterization of the output quantity by a Gaussian distribution or a scaled and shifted t-distribution [GUM:1995 G.6.2, G.6.4]. Within that framework, the law of propagation of uncertainty provides a means for propagating uncertainties through the model. Specifically, it evaluates the standard uncertainty associated with an estimate of the output quantity, given

1) best estimates of the input quantities,

2) the standard uncertainties associated with these estimates, and, where appropriate,

3) degrees of freedom associated with these standard uncertainties, and

4) any non-zero covariances associated with pairs of these estimates.

Also within the framework, the PDF taken to characterize the output quantity is used to provide a coverage interval, for a stipulated coverage probability, for that quantity.

The best estimates, standard uncertainties, covariances and degrees of freedom summarize the information available concerning the input quantities. With the approach considered here, the available information is encoded in terms of PDFs for the input quantities. The approach operates with these PDFs in order to determine the PDF for the output quantity.

Whereas there are some limitations to the GUM uncertainty framework, the propagation of distributions will always provide a PDF for the output quantity that is consistent with the model of the measurement and the PDFs for the input quantities. This PDF for the output quantity describes the knowledge of that quantity, based on the knowledge of the input quantities, as described by the PDFs assigned to them. Once the PDF for the output quantity is available, that quantity can be summarized by its expectation, taken as an estimate of the quantity, and its standard deviation, taken as the standard uncertainty associated with the estimate. Further, the PDF can be used to obtain a coverage interval, corresponding to a stipulated coverage probability, for the output quantity.

The use of PDFs as described in this Supplement is generally consistent with the concepts underlying the GUM. The PDF for a quantity expresses the state of knowledge about the quantity, i.e. it quantifies the degree of belief about the values that can be assigned to the quantity based on the available information. The information usually consists of raw statistical data, results of measurement, or other relevant scientific statements, as well as professional judgement.

In order to construct a PDF for a quantity, on the basis of a series of indications, Bayes' theorem can be applied [27, 33]. When appropriate information is available concerning systematic effects, the principle of maximum entropy can be

used to assign a suitable PDF [51, 56].

The propagation of distributions has wider application than the GUM uncertainty framework. It works with richer information than that conveyed by best estimates and the associated standard uncertainties (and degrees of freedom and covariances when appropriate).

An historical perspective is given in annex A.

NOTE 1 Citations of the form [GUM:1995 3.1.6] are to the indicated (sub)clauses of the GUM.

NOTE 2 The GUM provides an approach when linearization is inadequate [GUM:1995 5.1.2 note]. The approach has limitations: only the leading non-linear terms in the Taylor series expansion of the model are used, and the PDFs for the input quantities are regarded as Gaussian.

NOTE 3 Strictly, the GUM characterizes the variable $(Y - y)/u(y)$ by a t-distribution, where Y is the output quantity, y an estimate of Y, and $u(y)$ the standard uncertainty associated with y [GUM:1995 G.3.1]. This characterization is also used in this Supplement. (The GUM in fact refers to the variable $(y - Y)/u(y)$.)

NOTE 4 A PDF for a quantity is not to be understood as a frequency density.

NOTE 5 "The evaluation of uncertainty is neither a routine task nor a purely mathematical one; it depends on detailed knowledge of the nature of the measurand and of the measurement method and procedure used. The quality and utility of the uncertainty quoted for the result of a measurement therefore ultimately depends on the understanding, critical analysis, and integrity of those who contribute to the assignment of its value." [17].

Evaluation of measurement data — Supplement 1 to the "Guide to the expression of uncertainty in measurement" — Propagation of distributions using a Monte Carlo method

1 Scope

This Supplement provides a general numerical approach, consistent with the broad principles of the GUM [GUM:1995 G.1.5], for carrying out the calculations required as part of an evaluation of measurement uncertainty. The approach applies to arbitrary models having a single output quantity where the input quantities are characterized by any specified PDFs [GUM:1995 G.1.4, G.5.3].

As in the GUM, this Supplement is primarily concerned with the expression of uncertainty in the measurement of a well-defined physical quantity—the measurand—that can be characterized by an essentially unique value [GUM:1995 1.2].

This Supplement also provides guidance in situations where the conditions for the GUM uncertainty framework [GUM:1995 G.6.6] are not fulfilled, or it is unclear whether they are fulfilled. It can be used when it is difficult to apply the GUM uncertainty framework, because of the complexity of the model, for example. Guidance is given in a form suitable for computer implementation.

This Supplement can be used to provide (a representation of) the PDF for the output quantity from which

a) an estimate of the output quantity,

b) the standard uncertainty associated with this estimate, and

c) a coverage interval for that quantity, corresponding to a specified coverage probability

can be obtained.

Given (i) the model relating the input quantities and the output quantity and (ii) the PDFs characterizing the input quantities, there is a unique PDF for the output quantity. Generally, the latter PDF cannot be determined analytically. Therefore, the objective of the approach described here is to determine a), b), and c) above to a prescribed numerical tolerance, without making unquantified approximations.

For a prescribed coverage probability, this Supplement can be used to provide any required coverage interval, including the probabilistically symmetric coverage interval and the shortest coverage interval.

This Supplement applies to input quantities that are independent, where each such quantity is assigned an appropriate PDF, or not independent, i.e. when some or all of these quantities are assigned a joint PDF.

Typical of the uncertainty evaluation problems to which this Supplement can be applied include those in which

— the contributory uncertainties are not of approximately the same magnitude [GUM:1995 G.2.2],

— it is difficult or inconvenient to provide the partial derivatives of the model, as needed by the law of propagation of uncertainty [GUM:1995 5],

— the PDF for the output quantity is not a Gaussian distribution or a scaled and shifted t-distribution [GUM:1995 G.6.5],

— an estimate of the output quantity and the associated standard uncertainty are approximately of the same magnitude [GUM:1995 G.2.1],

— the models are arbitrarily complicated [GUM:1995 G.1.5], and

— the PDFs for the input quantities are asymmetric [GUM:1995 G.5.3].

A validation procedure is provided to check whether the GUM uncertainty framework is applicable. The GUM uncertainty framework remains the primary approach to uncertainty evaluation in circumstances where it is demonstrably applicable.

It is usually sufficient to report measurement uncertainty to one or perhaps two significant decimal digits. Guidance is provided on carrying out the calculation to give reasonable assurance that in terms of the information provided the reported decimal digits are correct.

Detailed examples illustrate the guidance provided.

This document is a Supplement to the GUM and is to be used in conjunction with it. Other approaches generally consistent with the GUM may alternatively be used. The audience of this Supplement is that of the GUM.

NOTE 1 This Supplement does not consider models that do not define the output quantity uniquely (for example, involving the solution of a quadratic equation, without specifying which root is to be taken).

NOTE 2 This Supplement does not consider the case where a prior PDF for the output quantity is available, but the treatment here can be adapted to cover this case [16].

2 Normative references

The following referenced documents are indispensable for the application of this document.

JCGM 100 (GUM:1995). Guide to the expression of uncertainty in measurement (GUM), 1995.

JCGM 200 (VIM:2008). International Vocabulary of Metrology—Basic and General Concepts and Associated Terms, VIM, 3rd Edition, 2008.

3 Terms and definitions

For the purposes of this document the terms and definitions of the GUM and the "International vocabulary of basic and general terms in metrology" (VIM) apply unless otherwise indicated. Some of the most relevant definitions, adapted where necessary from these documents (see 4.2), are given below. Further definitions are given, including definitions taken or adapted from other sources, that are important for this Supplement.

A glossary of principal symbols is given in annex G.

3.1
probability distribution
⟨random variable⟩ function giving the probability that a random variable takes any given value or belongs to a given set of values

NOTE The probability on the whole set of values of the random variable equals 1.

[Adapted from ISO 3534-1:1993 1.3; GUM:1995 C.2.3]

NOTE 1 A probability distribution is termed univariate when it relates to a single (scalar) random variable, and multivariate when it relates to a vector of random variables. A multivariate probability distribution is also described as a joint distribution.

NOTE 2 A probability distribution can take the form of a distribution function or a probability density function.

3.2
distribution function
function giving, for every value ξ, the probability that the random variable X be less than or equal to ξ:

$$G_X(\xi) = \Pr(X \leq \xi)$$

[Adapted from ISO 3534-1:1993 1.4; GUM:1995 C.2.4]

3.3
probability density function
derivative, when it exists, of the distribution function

$$g_X(\xi) = \mathrm{d}G_X(\xi)/\mathrm{d}\xi$$

NOTE $g_X(\xi)\,\mathrm{d}\xi$ is the "probability element"

$$g_X(\xi)\,\mathrm{d}\xi = \Pr(\xi < X < \xi + \mathrm{d}\xi).$$

[Adapted from ISO 3534-1:1993 1.5; GUM:1995 C.2.5]

3.4
normal distribution
probability distribution of a continuous random variable X having the probability density function

$$g_X(\xi) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{\xi - \mu}{\sigma}\right)^2\right],$$

for $-\infty < \xi < +\infty$

NOTE μ is the expectation and σ is the standard deviation of X.

[Adapted from ISO 3534-1:1993 1.37; GUM:1995 C.2.14]

NOTE The normal distribution is also known as a Gaussian distribution.

3.5
***t*-distribution**
probability distribution of a continuous random variable X having the probability density function

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

for $-\infty < \xi < +\infty$, with parameter ν, a positive integer, the degrees of freedom of the distribution, where

$$\Gamma(z) = \int_0^\infty t^{z-1}\mathrm{e}^{-t}\,\mathrm{d}t, \qquad z > 0,$$

is the gamma function

3.6
expectation
property of a random variable, which, for a continuous random variable X characterized by a PDF $g_X(\xi)$, is given by

$$E(X) = \int_{-\infty}^{\infty} \xi g_X(\xi)\,\mathrm{d}\xi$$

NOTE 1 Not all random variables have an expectation.

NOTE 2 The expectation of the random variable $Z = F(X)$, for a given function $F(X)$, is

$$E(Z) = E(F(X)) = \int_{-\infty}^{\infty} F(\xi) g_X(\xi)\, \mathrm{d}\xi.$$

3.7
variance
property of a random variable, which, for a continuous random variable X characterized by a PDF $g_X(\xi)$, is given by

$$V(X) = \int_{-\infty}^{\infty} [\xi - E(X)]^2 g_X(\xi)\, \mathrm{d}\xi$$

NOTE Not all random variables have a variance.

3.8
standard deviation
positive square root $[V(X)]^{1/2}$ of the variance

3.9
moment of order r
expectation of the rth power of a random variable, namely

$$E(X^r) = \int_{-\infty}^{\infty} \xi^r g_X(\xi)\, \mathrm{d}\xi$$

NOTE 1 The central moment of order r is the expectation of the random variable $Z = [X - E(X)]^r$.

NOTE 2 The expectation $E(X)$ is the first moment. The variance $V(X)$ is the central moment of order 2.

3.10
covariance
property of a pair of random variables, which, for two continuous random variables X_1 and X_2 characterized by a joint (multivariate) PDF $g_{\boldsymbol{X}}(\boldsymbol{\xi})$, where $\boldsymbol{X} = (X_1, X_2)^\top$ and $\boldsymbol{\xi} = (\xi_1, \xi_2)^\top$, is given by

$$\mathrm{Cov}(X_1, X_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} [\xi_1 - E(X_1)][\xi_2 - E(X_2)] g_{\boldsymbol{X}}(\boldsymbol{\xi})\, \mathrm{d}\xi_1\, \mathrm{d}\xi_2$$

NOTE Not all pairs of random variables have a covariance.

3.11
uncertainty matrix
matrix of dimension $N \times N$, containing on its diagonal the squares of the standard uncertainties associated with estimates of the components of an N-dimensional vector quantity, and in its off-diagonal positions the covariances associated with pairs of estimates

NOTE 1 An uncertainty matrix $\boldsymbol{U}_{\boldsymbol{x}}$ of dimension $N \times N$ associated with the vector estimate $\boldsymbol{x}$ of a vector quantity $\boldsymbol{X}$ has the representation

$$\boldsymbol{U}_{\boldsymbol{x}} = \begin{bmatrix} u(x_1, x_1) & \cdots & u(x_1, x_N) \\ \vdots & \ddots & \vdots \\ u(x_N, x_1) & \cdots & u(x_N, x_N) \end{bmatrix},$$

where $u(x_i, x_i) = u^2(x_i)$ is the variance (squared standard uncertainty) associated with x_i and $u(x_i, x_j)$ is the covariance associated with x_i and x_j. $u(x_i, x_j) = 0$ if elements X_i and X_j of $\boldsymbol{X}$ are uncorrelated.

NOTE 2 Covariances are also known as mutual uncertainties.

NOTE 3 An uncertainty matrix is also known as a covariance matrix or variance-covariance matrix.

3.12
coverage interval
interval containing the value of a quantity with a stated probability, based on the information available

NOTE 1 A coverage interval is sometimes known as a credible interval or a Bayesian interval.

NOTE 2 Generally there is more than one coverage interval for a stated probability.

NOTE 3 A coverage interval should not be termed 'confidence interval' to avoid confusion with the statistical concept [GUM:1995 6.2.2].

NOTE 4 This definition differs from that in the VIM, 3rd Edition (2008), since the term 'true value' has not been used in this Supplement, for reasons given in the GUM [GUM:1995 E.5].

3.13
coverage probability
probability that the value of a quantity is contained within a specified coverage interval

NOTE The coverage probability is sometimes termed "level of confidence" [GUM:1995 6.2.2].

3.14
length of a coverage interval
largest value minus smallest value in a coverage interval

3.15
probabilistically symmetric coverage interval
coverage interval for a quantity such that the probability that the quantity is less than the smallest value in the interval is equal to the probability that the quantity is greater than the largest value in the interval

3.16
shortest coverage interval
coverage interval for a quantity with the shortest length among all coverage intervals for that quantity having the same coverage probability

3.17
propagation of distributions
method used to determine the probability distribution for an output quantity from the probability distributions assigned to the input quantities on which the output quantity depends

NOTE The method may be analytical or numerical, exact or approximate.

3.18
GUM uncertainty framework
application of the law of propagation of uncertainty and the characterization of the output quantity by a Gaussian distribution or a scaled and shifted t-distribution in order to provide a coverage interval

3.19
Monte Carlo method
method for the propagation of distributions by performing random sampling from probability distributions

3.20
numerical tolerance
semi-width of the shortest interval containing all numbers that can correctly be expressed to a specified number of significant decimal digits

EXAMPLE All numbers greater than 1.75 and less than 1.85 can be expressed to two significant decimal digits as 1.8. The numerical tolerance is $(1.85 - 1.75)/2 = 0.05$.

NOTE For the calculation of numerical tolerance associated with a numerical value, see 7.9.2.

4 Conventions and notation

For the purposes of this Supplement the following conventions and notation are adopted.

4.1 A mathematical model of a measurement [GUM:1995 4.1] of a single (scalar) quantity can be expressed as a functional relationship f:

$$Y = f(\boldsymbol{X}), \tag{1}$$

where Y is a scalar output quantity and $\boldsymbol{X}$ represents the N input quantities $(X_1, \ldots, X_N)^\top$. Each X_i is regarded as a random variable with possible values ξ_i and expectation x_i. Y is a random variable with possible values η and expectation y.

NOTE 1 The same symbol is used for a physical quantity and the random variable that represents that quantity (cf. [GUM:1995 4.1.1 note 1]).

NOTE 2 Many models of measurement can be expressed in the form (1). A more general form is

$$h(Y, \boldsymbol{X}) = 0,$$

which implicitly relates $\boldsymbol{X}$ and Y. In any case, to apply the described Monte Carlo method, it is only necessary that Y can be formed corresponding to any meaningful $\boldsymbol{X}$.

4.2 This Supplement departs from the symbols often used for 'PDF' and 'distribution function' [24]. The GUM uses the generic symbol f to refer to a model and a PDF. Little confusion arises in the GUM as a consequence of this usage. The situation in this Supplement is different. The concepts of model, PDF, and distribution function are central to following and implementing the guidance provided. Therefore, in place of the symbols f and F to denote a PDF and a distribution function, respectively, the symbols g and G are used. These symbols are indexed appropriately to denote the quantity concerned. The symbol f is reserved for the model.

NOTE The definitions in clause 3 that relate to PDFs and distributions are adapted accordingly.

4.3 In this Supplement, a PDF is assigned to a quantity, which may be a single, scalar quantity X or a vector quantity $\boldsymbol{X}$. In the scalar case, the PDF for X is denoted by $g_X(\xi)$, where ξ is a variable describing the possible values of X. This X is considered as a random variable with expectation $E(X)$ and variance $V(X)$ (see 3.6 and 3.7).

4.4 In the vector case, the PDF for $\boldsymbol{X}$ is denoted by $g_{\boldsymbol{X}}(\boldsymbol{\xi})$, where $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_N)^\top$ is a vector variable describing the possible values of the vector quantity $\boldsymbol{X}$. This $\boldsymbol{X}$ is considered as a random vector variable with (vector) expectation $E(\boldsymbol{X})$ and covariance matrix $V(\boldsymbol{X})$.

4.5 A PDF for more than one input quantity is often called joint even if all the input quantities are independent.

4.6 When the elements X_i of $\boldsymbol{X}$ are independent, the PDF for X_i is denoted by $g_{X_i}(\xi_i)$.

4.7 The PDF for Y is denoted by $g_Y(\eta)$ and the distribution function for Y by $G_Y(\eta)$.

4.8 In the body of this Supplement, a quantity is generally denoted by an upper case letter and the expectation of the quantity or an estimate of the quantity by the corresponding lower case letter. For example, the expectation or an estimate of a quantity Y would be denoted by y. Such a notation is largely inappropriate for physical quantities, because of the established use of specific symbols, e.g. T for temperature and t for time. Therefore, in some of the examples (clause 9), a different notation is used. There, a quantity is denoted by its conventional symbol and its expectation or an estimate of it by that symbol hatted. For instance, the quantity representing the deviation of the length of a gauge block being calibrated from nominal length (see 9.5) is denoted by δL and an estimate of δL by $\widehat{\delta L}$.

NOTE A hatted symbol is generally used in the statistical literature to denote an estimate.

4.9 In this Supplement, the term "law of propagation of uncertainty" applies to the use of a first-order Taylor series approximation to the model. The term is qualified accordingly when a higher-order approximation is used.

4.10 The subscript "c" [GUM:1995 5.1.1] for the combined standard uncertainty is redundant in this Supplement. The standard uncertainty associated with an estimate y of an output quantity Y can therefore be written as $u(y)$, but the use of $u_c(y)$ remains acceptable if it is helpful to emphasize the fact that it represents a combined standard uncertainty. The qualifier "combined" in this context is also regarded as superfluous and may be omitted: the presence of "y" in "$u(y)$" already indicates the estimate with which the standard uncertainty is associated. Moreover, when the results of one or more uncertainty evaluations become inputs to a subsequent uncertainty evaluation, the use of the subscript "c" and the qualifier "combined" are then inappropriate.

4.11 The terms "coverage interval" and "coverage probability" are used throughout this Supplement. The GUM uses the term "level of confidence" as a synonym for coverage probability, drawing a distinction between "level of confidence" and "confidence level" [GUM:1995 6.2.2], because the latter has a specific definition in statistics. Since, in some languages, the translation from English of these two terms yields the same expression, the use of these terms is avoided here.

4.12 According to Resolution 10 of the 22nd CGPM (2003) " . . . the symbol for the decimal marker shall be either the point on the line or the comma on the line . . . ". The JCGM has decided to adopt, in its documents in English, the point on the line.

4.13 Unless otherwise qualified, numbers are expressed in a manner that indicates the number of meaningful significant decimal digits.

EXAMPLE The numbers 0.060, 0.60, 6.0 and 60 are expressed to two significant decimal digits. The numbers 0.06, 0.6, 6 and 6×10^1 are expressed to one significant decimal digit. It would be incorrect to express 6×10^1 as 60, since two significant decimal digits would be implied.

4.14 Some symbols have more than one meaning in this Supplement. See annex G. The context clarifies the usage.

4.15 The following abbreviations are used in this Supplement:

CGPM	Conférence Générale des Poids et Mesures
IEEE	Institute of Electrical and Electronic Engineers
GUF	GUM uncertainty framework
JCGM	Joint Committee for Guides in Metrology
GUM	Guide to the expression of uncertainty in measurement
MCM	Monte Carlo method
PDF	probability density function
VIM	International vocabulary of basic and general terms in metrology

5 Basic principles

5.1 Main stages of uncertainty evaluation

5.1.1 The main stages of uncertainty evaluation constitute formulation, propagation, and summarizing:

a) *Formulation*:

1) define the output quantity Y, the quantity intended to be measured (the measurand);

2) determine the input quantities $\boldsymbol{X} = (X_1, \ldots, X_N)^\top$ upon which Y depends;

3) develop a model relating Y and $\boldsymbol{X}$;

4) on the basis of available knowledge assign PDFs—Gaussian (normal), rectangular (uniform), etc.—to the X_i. Assign instead a joint PDF to those X_i that are not independent;

b) *Propagation*: propagate the PDFs for the X_i through the model to obtain the PDF for Y;

c) *Summarizing*: use the PDF for Y to obtain

1) the expectation of Y, taken as an estimate y of the quantity,

2) the standard deviation of Y, taken as the standard uncertainty $u(y)$ associated with y [GUM:1995 E.3.2], and

3) a coverage interval containing Y with a specified probability (the coverage probability).

NOTE 1 The expectation may not be appropriate for all applications (cf. [GUM:1995 4.1.4]).

NOTE 2 The quantities described by some distributions, such as the Cauchy distribution, have no expectation or standard deviation. A coverage interval for the output quantity can always be obtained, however.

5.1.2 The GUM uncertainty framework does not explicitly refer to the assignment of PDFs to the input quantities. However [GUM:1995 3.3.5], "... a Type A standard uncertainty is obtained from a probability density function ... derived from an observed frequency distribution ..., while a Type B standard uncertainty is obtained from an assumed probability density function based on the degree of belief that an event will occur .... Both approaches employ recognized interpretations of probability."

NOTE The use of probability distributions in a Type B evaluation of uncertainty is a feature of Bayesian inference [21, 27]. Research continues [22] on the boundaries of validity for the assignment of degrees of freedom to a standard uncertainty based on the Welch-Satterthwaite formula.

5.1.3 The steps in the formulation stage are carried out by the metrologist, perhaps with expert support. Guidance on the assignment of PDFs (step 4) of stage a) in 5.1.1) is given in this Supplement for some common cases (see 6.4). The propagation and summarizing stages, b) and c) in 5.1.1, for which detailed guidance is provided here, require no further metrological information, and in principle can be carried out to any required numerical tolerance for the problem specified in the formulation stage.

NOTE Once the formulation stage a) in 5.1.1 has been carried out, the PDF for the output quantity is completely specified mathematically, but generally the calculation of the expectation, standard deviation and coverage intervals requires numerical methods that involve a degree of approximation.

5.2 Propagation of distributions

In this Supplement a generally efficient approach for determining (a numerical approximation to) the distribution function

$$G_Y(\eta) = \int_{-\infty}^{\eta} g_Y(z)\,\mathrm{d}z$$

for Y is considered. It is based on applying a Monte Carlo method (MCM) as an implementation of the propagation of distributions (see 5.9).

NOTE A formal definition [9] for the PDF for Y is

$$g_Y(\eta) = \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} g_{\boldsymbol{X}}(\boldsymbol{\xi})\delta(\eta - f(\boldsymbol{\xi}))\,\mathrm{d}\xi_N \cdots \mathrm{d}\xi_1,$$

where $\delta(\cdot)$ denotes the Dirac delta function. This multiple integral cannot generally be evaluated analytically. A numerical integration rule can be applied to provide an approximation to $g_Y(\eta)$, but this is not an efficient approach.

5.3 Obtaining summary information

5.3.1 An estimate y of Y is the expectation $E(Y)$. The standard uncertainty $u(y)$ associated with y is given by the standard deviation of Y, the positive square root of the variance $V(Y)$ of Y.

5.3.2 A coverage interval for Y can be determined from $G_Y(\eta)$. Let α denote any numerical value between zero and $1-p$, where p is the required coverage probability. The endpoints of a $100p$ % coverage interval for Y are $G_Y^{-1}(\alpha)$ and $G_Y^{-1}(p+\alpha)$, i.e. the α- and $(p+\alpha)$-quantiles of $G_Y(\eta)$.

5.3.3 The choice $\alpha = (1-p)/2$ gives the coverage interval defined by the $(1-p)/2$- and $(1+p)/2$-quantiles, providing a probabilistically symmetric $100p$ % coverage interval.

NOTE When the PDF for Y is symmetric about the estimate y, the coverage interval obtained would be identical to $y \pm U_p$, where the expanded uncertainty [GUM:1995 2.3.5] U_p is given by the product of the standard uncertainty $u(y)$ and the coverage factor that is appropriate for that PDF. This PDF is generally not known analytically.

5.3.4 A numerical value of α different from $(1-p)/2$ may be more appropriate if the PDF is asymmetric. The shortest $100p$ % coverage interval can be used in this case. It has the property that, for a unimodal (single-peaked) PDF, it contains the mode, the most probable value of Y. It is given by the numerical value of α satisfying $g_Y(G_Y^{-1}(\alpha)) = g_Y(G_Y^{-1}(p+\alpha))$, if $g_Y(\eta)$ is unimodal, and in general by the numerical value of α such that $G_Y^{-1}(p+\alpha) - G_Y^{-1}(\alpha)$ is a minimum.

5.3.5 The probabilistically symmetric $100p$ % coverage interval and the shortest $100p$ % coverage interval are identical for a symmetric PDF, such as the Gaussian and scaled and shifted t-distribution used within the GUM uncertainty framework. Therefore, in comparing the GUM uncertainty framework with other approaches, either of these intervals can be used.

5.3.6 Figure 1 shows the distribution function $G_Y(\eta)$ corresponding to an asymmetric PDF. Broken vertical lines mark the endpoints of the probabilistically symmetric 95 % coverage interval and broken horizontal lines the corresponding probability points, viz. 0.025 and 0.975. Continuous lines mark the endpoints of the shortest 95 % coverage interval and the corresponding probability points, which are 0.006 and 0.956 in this case. The lengths of the 95 % coverage intervals in the two cases are 1.76 unit and 1.69 unit, respectively.

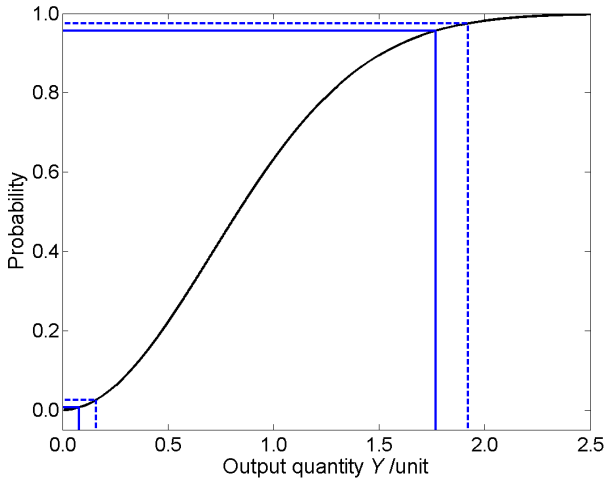

Figure 1 — A distribution function $G_Y(\eta)$ corresponding to an asymmetric PDF and the probabilistically symmetric and shortest 95 % coverage intervals (5.3.6). "Unit" denotes any unit

5.4 Implementations of the propagation of distributions

5.4.1 The propagation of distributions can be implemented in several ways:

a) analytical methods, i.e. methods that provide a mathematical representation of the PDF for Y;

b) uncertainty propagation based on replacing the model by a first-order Taylor series approximation [GUM:1995 5.1.2] — the law of propagation of uncertainty;

c) as b), except that contributions derived from higher-order terms in the Taylor series approximation are included [GUM:1995 5.1.2 note];

d) numerical methods [GUM:1995 G.1.5] that implement the propagation of distributions, specifically using MCM (see 5.9).

NOTE 1 Analytical methods are ideal in that they do not introduce any approximation. They are applicable in simple cases only, however. A treatment and examples are available [8, 13]. These methods are not considered further in this Supplement, apart from in the examples (clause 9) for comparison purposes.

NOTE 2 MCM as considered here is regarded as a means for providing a numerical representation of the distribution for the output quantity, rather than a simulation method *per se*. In the context of the propagation stage of uncertainty evaluation, the problem to be solved is deterministic, there being no random physical process to be simulated.

5.4.2 Approaches to uncertainty evaluation other than the GUM uncertainty framework are permitted by the GUM [GUM:1995 G.1.5]. The approach advocated in this Supplement, based on the propagation of distributions, is general. For linear or linearized models and input quantities for which the PDFs are Gaussian, the approach yields results consistent with the GUM uncertainty framework. However, in cases where the conditions for the GUM uncertainty framework to be applied (see 5.7 and 5.8) do not hold, the approach of this Supplement can generally be expected to lead to a valid uncertainty statement.

5.4.3 An appropriate method has to be chosen for the propagation stage. If it can be demonstrated that the conditions necessary for the GUM uncertainty framework to give valid results hold, then that approach can be used. If there are indications that the GUM uncertainty framework is likely to be invalid, then another approach should be employed. A third situation can arise in which it is difficult to assess whether or not the GUM uncertainty framework will be valid. In all three cases, MCM provides a practical (alternative) method. In the first case, MCM may sometimes be easier to apply due to difficulties in calculating sensitivity coefficients [GUM:1995 5.1.3], for example. In the second, MCM can generally be expected to give valid results, since it does not make approximating assumptions. In the third, MCM can be applied either to determine the results directly or to assess the quality of those provided by the GUM uncertainty framework.

5.4.4 The propagation of the PDFs $g_{X_i}(\xi_i)$, $i = 1, \ldots, N$, for the input quantities X_i through the model to provide the PDF $g_Y(\eta)$ for the output quantity Y is illustrated in figure 2 for $N = 3$ independent X_i. Figure 2 may be compared to figure 3 for the law of propagation of uncertainty. In figure 2, the $g_{X_i}(\xi_i)$, $i = 1, 2, 3$, are Gaussian, triangular, and Gaussian, respectively. $g_Y(\eta)$ is indicated as being asymmetric, as generally arises for non-linear models or asymmetric $g_{X_i}(\xi_i)$.

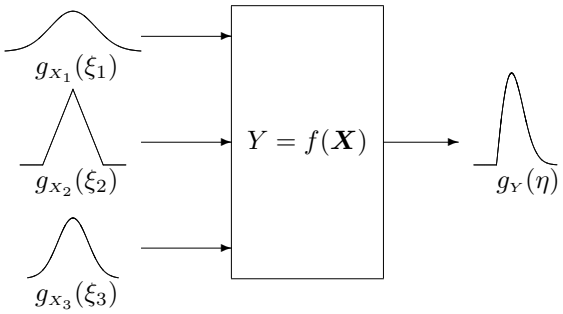

Figure 2 — Illustration of the propagation of distributions for $N = 3$ independent input quantities (5.4.4)

5.4.5 In practice, only for simple cases can the propagation of distributions be implemented without making approximations. The GUM uncertainty framework implements one approximate method, and MCM another. For a small but important subset of problems, the GUM uncertainty framework is exact. MCM is never exact, but is more valid than the GUM uncertainty framework for a large class of problems.

5.5 Reporting the results

5.5.1 The following items would typically be reported following the use of the propagation of distributions:

a) an estimate y of the output quantity Y;

b) the standard uncertainty $u(y)$ associated with y;

c) the stipulated coverage probability $100p$ % (e.g. 95 %);

d) the endpoints of the selected $100p$ % coverage interval (e.g. 95 % coverage interval) for Y;

e) any other relevant information, such as whether the coverage interval is a probabilistically symmetric coverage interval or a shortest coverage interval.

5.5.2 y, $u(y)$ and the endpoints of a $100p$ % coverage interval for Y should be reported to a number of decimal digits such that the least significant decimal digit is in the same position with respect to the decimal point as that for $u(y)$ [GUM:1995 7.2.6]. One or two significant decimal digits would usually be adequate to represent $u(y)$.

NOTE 1 Each reported numerical value would typically be obtained by rounding a numerical value expressed to a greater number of significant decimal digits.

NOTE 2 A factor influencing the choice of one or two significant decimal digits is the leading significant decimal digit of $u(y)$. If this digit is 1 or 2, the deviation of the reported numerical value of $u(y)$ from its numerical value before rounding is large relative to the latter numerical value. If the leading significant decimal digit is 9, the deviation is relatively smaller.

NOTE 3 If the results are to be used within further calculations, consideration should be given to whether additional decimal digits should be retained.

EXAMPLE Reported results corresponding to declaring two significant decimal digits in $u(y)$, for a case in which the coverage interval is asymmetric with respect to y, are

$$y = 1.024 \text{ V}, \qquad u(y) = 0.028 \text{ V},$$
$$\text{shortest 95 \% coverage interval} = [0.983,\ 1.088] \text{ V}.$$

The same results reported to one significant decimal digit in $u(y)$ would be

$$y = 1.02 \text{ V}, \qquad u(y) = 0.03 \text{ V},$$
$$\text{shortest 95 \% coverage interval} = [0.98,\ 1.09] \text{ V}.$$

5.6 GUM uncertainty framework

5.6.1 The GUM provides general guidance on many aspects of the stages of uncertainty evaluation presented in 5.1.1. It also provides the GUM uncertainty framework for the propagation and summarizing stages of uncertainty evaluation. The GUM uncertainty framework has been adopted by many organizations, is widely used, and has been implemented in standards and guides on measurement uncertainty and also in software.

5.6.2 The GUM uncertainty framework comprises the following stages. Each model input quantity X_i is summarized by its expectation and standard deviation, as given by the PDF for that quantity [GUM:1995 4.1.6]. The expectation is taken as the best estimate x_i of X_i and the standard deviation as the standard uncertainty $u(x_i)$ associated with x_i. This information is propagated, using the law of propagation of uncertainty [GUM:1995 5.1.2], through a first- or higher-order Taylor series approximation to the model to provide

a) an estimate y of the output quantity Y, and

b) the standard uncertainty $u(y)$ associated with y.

The estimate y is given by evaluating the model at the x_i. A coverage interval for Y is provided based on taking the PDF for Y as Gaussian or, if the degrees of freedom associated with $u(y)$ is finite [GUM:1995 G], as a scaled and shifted t-distribution.

NOTE The summaries of the X_i also include, where appropriate, the degrees of freedom associated with the $u(x_i)$ [GUM:1995 4.2.6]. They also include, where appropriate, covariances associated with pairs of x_i [GUM:1995 5.2.5].

5.6.3 The propagation and summarizing stages of the GUM uncertainty framework (stages b) and c) in 5.1.1) constitute the following computational steps. Also see figure 3, which illustrates the law of propagation of uncertainty for a model having $N = 3$ independent input quantities $\boldsymbol{X} = (X_1, X_2, X_3)^\top$, which are estimated by x_i with associated standard uncertainties $u(x_i)$, $i = 1, 2, 3$. The output quantity Y is estimated by y, with associated standard uncertainty $u(y)$.

a) Obtain from the PDFs for the input quantities $\boldsymbol{X} = (X_1, \ldots, X_N)^\top$ the expectations $\boldsymbol{x} = (x_1, \ldots, x_N)^\top$ and the standard deviations (standard uncertainties) $\boldsymbol{u}(\boldsymbol{x}) = [u(x_1), \ldots, u(x_N)]^\top$. Use instead the joint PDF for $\boldsymbol{X}$ if pairs of the X_i are not independent (in which case they have non-zero covariance).

b) Set the degrees of freedom (infinite or finite) associated with each $u(x_i)$.

c) For each pair i, j for which X_i and X_j are not independent, obtain from the joint PDF for X_i and X_j the covariance (mutual uncertainty) $u(x_i, x_j)$ associated with x_i and x_j.

d) Form the partial derivatives of first order of $f(\boldsymbol{X})$ with respect to $\boldsymbol{X}$.

e) Calculate y, the model evaluated at $\boldsymbol{X}$ equal to $\boldsymbol{x}$.

f) Calculate the model sensitivity coefficients [GUM:1995 5.1.3] as the above partial derivatives evaluated at $\boldsymbol{x}$.

g) Calculate the standard uncertainty $u(y)$ by combining $\boldsymbol{u}(\boldsymbol{x})$, the $u(x_i, x_j)$, and the model sensitivity coefficients [GUM:1995 formulæ (10), (13)].

h) Calculate ν_{eff}, the effective degrees of freedom associated with $u(y)$, using the Welch-Satterthwaite formula [GUM:1995 formula (G.2b)].

i) Calculate the expanded uncertainty U_p, and hence a coverage interval (for a stipulated coverage probability p) for Y, regarded as a random variable, by forming the appropriate multiple of $u(y)$ through taking the probability distribution of $(Y - y)/u(y)$ as a standard Gaussian distribution ($\nu_{\text{eff}} = \infty$) or t-distribution ($\nu_{\text{eff}} < \infty$).

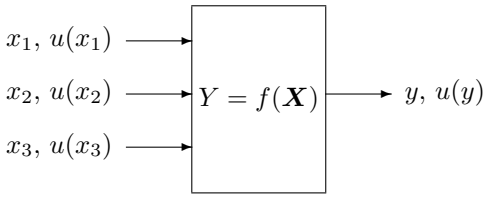

Figure 3 — Illustration of the law of propagation of uncertainty for $N = 3$ independent input quantities (5.4.4, 5.6.3)

5.7 Conditions for valid application of the GUM uncertainty framework for linear models

5.7.1 No condition is necessary for the valid application of the law of propagation of uncertainty to linear models (models that are linear in the X_i).

5.7.2 A coverage interval can be determined, in terms of the information provided in the GUM, under the following conditions:

a) the Welch-Satterthwaite formula is adequate for calculating the effective degrees of freedom associated with $u(y)$ [GUM:1995 G.4.1], when one or more of the $u(x_i)$ has an associated degrees of freedom that is finite;

b) the X_i are independent when the degrees of freedom associated with the $u(x_i)$ are finite;

c) the PDF for Y can adequately be approximated by a Gaussian distribution or a scaled and shifted t-distribution.

NOTE 1 Condition a) is required in order that Y can be characterized by an appropriate scaled and shifted t-distribution.

NOTE 2 Condition b) is required because the GUM does not treat X_i that are not independent in conjunction with finite degrees of freedom.

NOTE 3 Condition c) is satisfied when each X_i is assigned a Gaussian distribution. It is also satisfied when the conditions for the central limit theorem [GUM:1995 G.2] hold.

NOTE 4 The GUM uncertainty framework may not validly be applicable when there is an X_i whose assigned distribution is non-Gaussian and the corresponding contribution to $u(y)$ is dominant.

5.7.3 When the conditions in 5.7.2 hold, the results from the application of the GUM uncertainty framework can be expected to be valid for linear models. These conditions apply in many circumstances.

5.8 Conditions for valid application of the GUM uncertainty framework for non-linear models

5.8.1 The law of propagation of uncertainty can validly be applied for non-linear models under the following conditions:

a) f is continuously differentiable with respect to the elements X_i of $\boldsymbol{X}$ in the neighbourhood of the best estimates x_i of the X_i;

b) condition a) applies for all derivatives up to the appropriate order;

c) the X_i involved in significant higher-order terms of a Taylor series approximation to $f(\boldsymbol{X})$ are independent;

d) the PDFs assigned to X_i involved in higher-order terms of a Taylor series approximation to $f(\boldsymbol{X})$ are Gaussian;

e) higher-order terms that are not included in the Taylor series approximation to $f(\boldsymbol{X})$ are negligible.

NOTE 1 Condition a) is necessary for the applicability of the law of propagation of uncertainty based on a first-order Taylor series approximation to $f(\boldsymbol{X})$ when the non-linearity of f is insignificant [GUM:1995 5.1.2].

NOTE 2 Condition b) is necessary for the application of the law of propagation of uncertainty based on a higher-order Taylor series approximation to $f(\boldsymbol{X})$ [GUM:1995 5.1.2]. An expression for the most important terms of next highest order to be included are given in the GUM [GUM:1995 5.1.2 note].

NOTE 3 Condition c) relates to the statement in the GUM [GUM:1995 5.1.2 note] concerning significant model non-linearity in the case of independent X_i. The GUM does not consider X_i that are not independent in this context.

NOTE 4 Condition d) constitutes a correction to the statement in the GUM [GUM:1995 5.1.2 note] that the version of the law of propagation of uncertainty using higher-order terms is based on the symmetry of the PDFs for the X_i [19, 27].

NOTE 5 If the analytical determination of the higher derivatives, required when the non-linearity of the model is significant, is difficult or error-prone, suitable software for automatic differentiation can be used. Alternatively, these derivatives can be approximated numerically using finite differences [5]. (The GUM provides a finite-difference formula for partial derivatives of first order [GUM:1995 5.1.3 note 2].) Care should be taken, however, because of the effects of subtractive cancellation when forming differences between numerically close model values.

5.8.2 A coverage interval can be determined, in terms of the information provided in the GUM, when conditions a), b) and c) in 5.7.2 apply, with the exception that the content of note 3 in that subclause is replaced by "Condition c) is required in order that coverage intervals can be determined from these distributions."

5.8.3 When the conditions in 5.8.1 and 5.8.2 hold, the results from the application of the GUM uncertainty framework can be expected to be valid for non-linear models. These conditions apply in many circumstances.

5.9 Monte Carlo approach to the propagation and summarizing stages

5.9.1 MCM provides a general approach to obtain an approximate numerical representation $\boldsymbol{G}$, say, of the distribution function $G_Y(\eta)$ for Y [32, page 75]. The heart of the approach is repeated sampling from the PDFs for the X_i

and the evaluation of the model in each case.

5.9.2 Since $G_Y(\eta)$ encodes all the information known about Y, any property of Y such as expectation, variance and coverage intervals can be approximated using $\boldsymbol{G}$. The quality of these calculated results improves as the number of times the PDFs are sampled increases.

5.9.3 Expectations and variances (and higher moments) can be determined directly from the set of model values obtained. The determination of coverage intervals requires these model values to be ordered.

5.9.4 If y_r, for $r = 1, \ldots, M$, represent M model values sampled independently from a probability distribution for Y, then the expectation $E(Y)$ and variance $V(Y)$ can be approximated using the y_r. In general, the moments of Y (including $E(Y)$ and $V(Y)$) are approximated by those of the sampled model values. Let M_{y_0} denote the number of y_r that are no greater than y_0, any prescribed number. The probability $\Pr(Y \leq y_0)$ is approximated by M_{y_0}/M. In this way, the y_r provide a step function (histogram-like) approximation to the distribution function $G_Y(\eta)$.

5.9.5 Each y_r is obtained by sampling at random from each of the PDFs for the X_i and evaluating the model at the sampled values so obtained. $\boldsymbol{G}$, the primary output from MCM, constitutes the y_r arranged in strictly increasing order.

NOTE It is remotely possible that equalities exist amongst the y_r, in which case suitable minute perturbations made to the y_r would enable the y_r to be arranged in strictly increasing order. See 7.5.1.

5.9.6 MCM as an implementation of the propagation of distributions is shown diagrammatically in figure 4 for M provided in advance (see 7.9 otherwise). MCM can be stated as a step-by-step procedure:

a) select the number M of Monte Carlo trials to be made. See 7.2;

b) generate M vectors, by sampling from the assigned PDFs, as realizations of the (set of N) input quantities X_i. See 7.3;

c) for each such vector, form the corresponding model value of Y, yielding M model values. See 7.4;

d) sort these M model values into strictly increasing order, using the sorted model values to provide $\boldsymbol{G}$. See 7.5;

e) use $\boldsymbol{G}$ to form an estimate y of Y and the standard uncertainty $u(y)$ associated with y. See 7.6;

f) use $\boldsymbol{G}$ to form an appropriate coverage interval for Y, for a stipulated coverage probability p. See 7.7.

NOTE 1 Subclause 6.4 and annex C provide information on sampling from probability distributions.

NOTE 2 Mathematically, the average of the M model values is a realization of a random variable with expectation $E(Y)$ and variance $V(Y)/M$. Thus, the closeness of agreement between this average and $E(Y)$ can be expected to be proportional to $M^{-1/2}$.

NOTE 3 Step e) can equally be carried out by using the M model values of Y unsorted. It is necessary to sort these model values to determine the coverage interval in step f).

5.9.7 The effectiveness of MCM to determine y, $u(y)$ and a coverage interval for Y depends on the use of an adequately large value of M (step a) in 5.9.6). Guidance on obtaining such a value and generally on implementing MCM is available [7]. Also see 7.2 and 7.9.

5.10 Conditions for the valid application of the described Monte Carlo method

5.10.1 The propagation of distributions implemented using MCM can validly be applied, and the required summary information subsequently determined, using the approach provided in this Supplement, under the following conditions:

a) f is continuous with respect to the elements X_i of $\boldsymbol{X}$ in the neighbourhood of the best estimates x_i of the X_i;

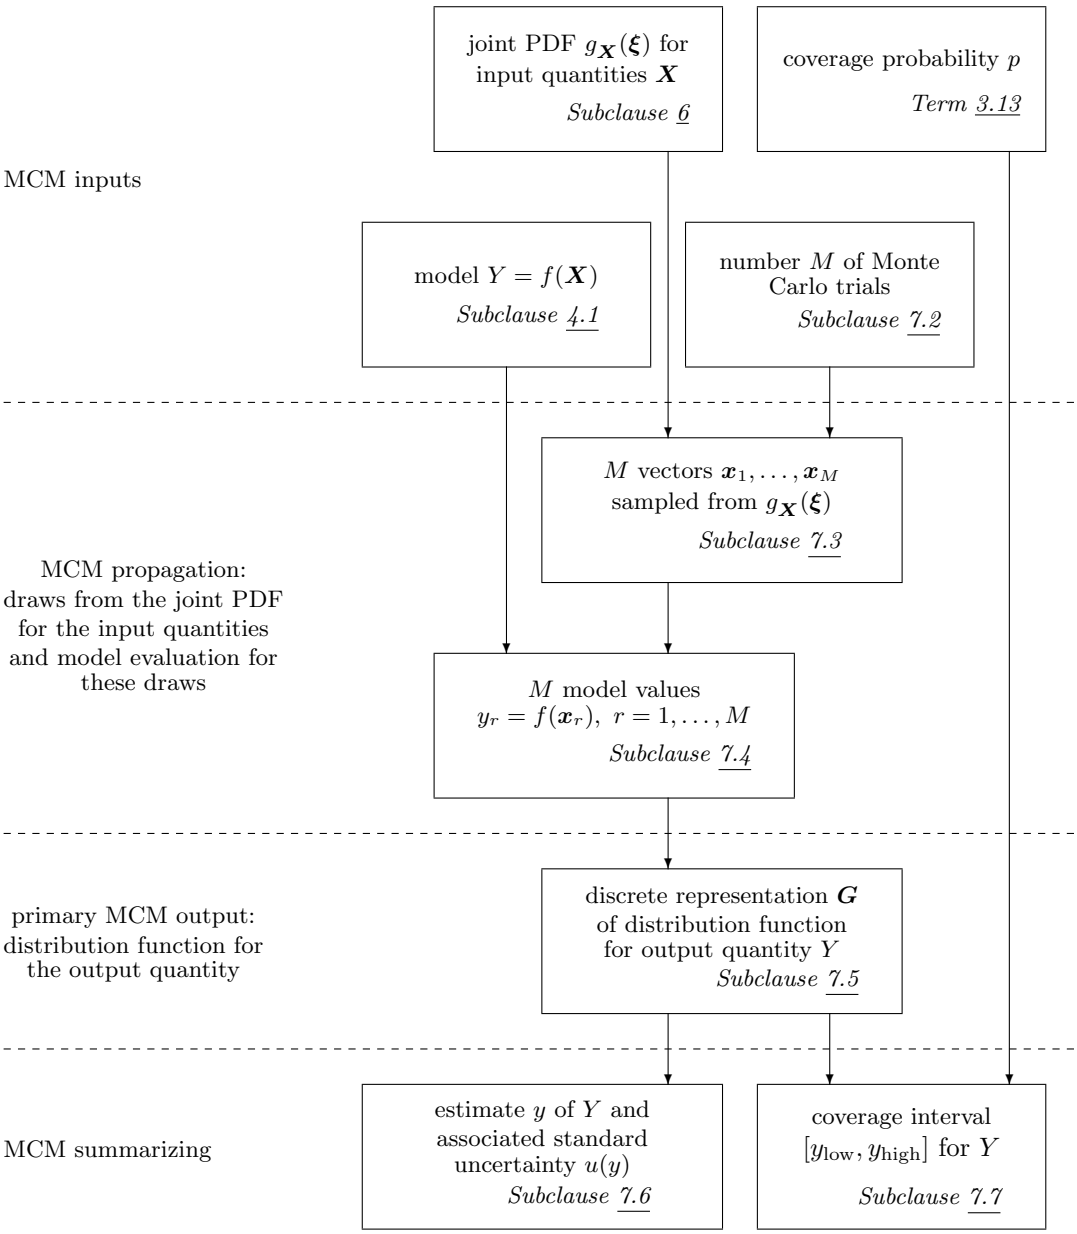

Figure 4 — The propagation and summarizing stages of uncertainty evaluation using MCM to implement the propagation of distributions (5.9.6, 7.1)

b) the distribution function for Y is continuous and strictly increasing;

c) the PDF for Y is

1) continuous over the interval for which this PDF is strictly positive,

2) unimodal (single-peaked), and

3) strictly increasing (or zero) to the left of the mode and strictly decreasing (or zero) to the right of the mode;

d) $E(Y)$ and $V(Y)$ exist;

e) a sufficiently large value of M is used.

NOTE 1 Regarding condition a), no condition on the derivatives of f is required.

NOTE 2 Conditions a) and b) are necessary to ensure that the inverse of the distribution function is unique and hence coverage intervals can be determined. Only condition a) is needed if a coverage interval is not required.

NOTE 3 Condition c) is needed only if the shortest coverage interval is to be determined. In that case, the condition is necessary to ensure that the shortest coverage interval corresponding to a stipulated coverage probability is unique. The mode may occur at an endpoint of the interval over which this PDF is strictly positive, in which case one of the two conditions in 3) is vacuous.

NOTE 4 Condition d) is needed for (stochastic) convergence of MCM as the number M of trials (see 7.2) increases.

NOTE 5 Condition e) is necessary to ensure that the summarizing information is reliable. See 8.2.

5.10.2 When the conditions in 5.10.1 hold, the results from the application of the propagation of distributions implemented in terms of MCM can be expected to be valid. These conditions are less restrictive than those (see 5.7 and 5.8) for the application of the GUM uncertainty framework.

5.11 Comparison of the GUM uncertainty framework and the described Monte Carlo method

5.11.1 The intention of this subclause is to compare the principles on which the GUM uncertainty framework and MCM as an implementation of the propagation of distributions are based. This subclause also provides some motivation for the use of MCM in circumstances where it is questionable whether the application of the GUM uncertainty framework is valid.

5.11.2 For the purposes of comparing the GUM uncertainty framework and MCM, it is helpful to review the considerations in the GUM regarding Type A and Type B evaluations of uncertainty. For Type A evaluation, the GUM provides guidance on obtaining a best estimate of a quantity and the associated standard uncertainty from the average and the associated standard deviation of a set of indications of the quantity, obtained independently. For Type B evaluation, prior knowledge concerning the quantity is used to characterize the quantity by a PDF, from which a best estimate of the quantity and the standard uncertainty associated with that estimate are determined. The GUM states that both types of evaluation are based on probability distributions [GUM:1995 3.3.4], and that both approaches employ recognized interpretations of probability [GUM:1995 3.3.5]. The GUM considers PDFs as underpinning uncertainty evaluation: in the context of the law of propagation of uncertainty, it refers explicitly to input and output quantities as being describable or characterized by probability distributions [GUM:1995 G.6.6]. Also see 5.1.2.

5.11.3 The GUM uncertainty framework does not explicitly determine a PDF for the output quantity. However, the probability distribution used by that framework to characterize the output quantity is sometimes referred to in this Supplement as "provided by" or "resulting from" the GUM uncertainty framework.

5.11.4 This Supplement attempts to provide an approach that is as consistent with the GUM as possible, especially relating to the use of PDFs for all quantities, but departs from it in a clearly identified way where appropriate. These departures are:

a) PDFs are explicitly assigned to all input quantities X_i (rather than associating standard uncertainties with estimates x_i of X_i) based on information concerning these quantities. The classification into Type A and Type B evaluations of uncertainty is not needed;

b) sensitivity coefficients [GUM:1995 5.1.3] are not an inherent part of the approach, and hence the calculation or numerical approximation of the partial derivatives of the model with respect to the X_i is not required. Approximations to sensitivity coefficients can, however, be provided that correspond to taking all higher-order terms in the Taylor series expansion of the model into account (annex B);

c) a numerical representation of the distribution function for Y is obtained that is defined completely by the model and the PDFs for the X_i, and not restricted to a Gaussian distribution or scaled and shifted t-distribution;

d) since the PDF for Y is not in general symmetric, a coverage interval for Y is not necessarily centred on the estimate of Y. Consideration therefore needs to be given to the choice of coverage interval corresponding to a specified coverage probability.

5.11.5 Since the GUM uncertainty framework explicitly uses only best estimates x_i and the associated uncertainties (and covariances and degrees of freedom where appropriate), it is restricted in the information it can provide about Y. Essentially it is limited to providing an estimate y of Y and the standard uncertainty $u(y)$ associated with y, and perhaps the related (effective) degrees of freedom. y and $u(y)$ will be valid for a model that is linear in X. Any other information about Y, e.g. coverage intervals, is derived using additional assumptions, e.g. that the distribution for Y is Gaussian or a scaled and shifted t-distribution.

5.11.6 Some features of MCM are

a) reduction in the analysis effort required for complicated or non-linear models, especially since the partial derivatives of first- or higher-order used in providing sensitivity coefficients for the law of propagation of uncertainty are not needed,

b) generally improved estimate of Y for non-linear models (cf. [GUM:1995 4.1.4]),

c) improved standard uncertainty associated with the estimate of Y for non-linear models, especially when the X_i are assigned non-Gaussian (e.g. asymmetric) PDFs, without the need to provide derivatives of higher order [GUM:1995 5.1.2 note],

d) provision of a coverage interval corresponding to a stipulated coverage probability when the PDF for Y cannot adequately be approximated by a Gaussian distribution or a scaled and shifted t-distribution, i.e. when the central limit theorem does not apply [GUM:1995 G.2.1, G.6.6]. Such an inadequate approximation can arise when (1) the PDF assigned to a dominant X_i is not a Gaussian distribution or a scaled and shifted t-distribution, (2) the model is non-linear, or (3) the approximation error incurred in using the Welch-Satterthwaite formula for effective degrees of freedom is not negligible, and

e) a coverage factor [GUM:1995 2.3.6] is not required when determining a coverage interval.

6 Probability density functions for the input quantities

6.1 General

6.1.1 This clause gives guidance on the assignment, in some common circumstances, of PDFs to the input quantities X_i in the formulation stage of uncertainty evaluation. Such an assignment can be based on Bayes' theorem [20] or the principle of maximum entropy [8, 26, 51, 56].

NOTE In some circumstances, another approach for assigning a PDF may be useful. In any case, as in any scientific discipline, the reason for the decision should be recorded.

6.1.2 Generally, a joint PDF $g_{\boldsymbol{X}}(\boldsymbol{\xi})$ is assigned to the input quantities $\boldsymbol{X} = (X_1, \ldots, X_N)^\top$. See 6.4.8.4 note 2.

6.1.3 When the X_i are independent, PDFs $g_{X_i}(\xi_i)$ are assigned individually based on an analysis of a series of indications (Type A evaluation of uncertainty) or based on scientific judgement using information [50] such as historical data, calibrations, and expert judgement (Type B evaluation of uncertainty) [GUM:1995 3.3.5].

6.1.4 When some of the X_i are mutually independent, PDFs are assigned individually to them and a joint PDF to the remainder.

NOTE It may be possible to remove some or all dependencies by re-expressing relevant input quantities in terms of more fundamental independent input quantities on which the original input quantities depend [GUM:1995 F1.2.4, H.1.2]. Such changes can simplify both the application of the law of propagation of uncertainty and the propagation of distributions. Details and examples are available [15].

6.1.5 Information relevant to the assignment of PDFs to the X_i is contained in the GUM [GUM:1995 4.3].

6.1.6 Comprehensive guidance on the assignment of PDFs individually or jointly to the X_i is beyond the scope of this Supplement. Such assigned PDFs encode the knowledge and expertise of the metrologist who formulates the model and who is ultimately responsible for the quality of the final results.

6.1.7 A standard text on probability distributions is Evans, Hastings and Peacock [18].

6.2 Bayes' theorem

6.2.1 Suppose that information about an input quantity X consists of a series of indications regarded as realizations of independent, identically distributed random variables characterized by a specified PDF, but with unknown expectation and variance. Bayes' theorem can be used to calculate a PDF for X, where X is taken to be equal to the unknown average of these random variables. Calculation proceeds in two steps. First, a non-informative joint prior (pre-data) PDF is assigned to the unknown expectation and variance. Using Bayes' theorem, this joint prior PDF is then updated, based on the information supplied by the series of indications, to yield a joint posterior (post-data) PDF for the two unknown parameters. The desired posterior PDF for the unknown average is then calculated as a marginal PDF by integrating over the possible values of the unknown variance (see 6.4.9.2).

6.2.2 With the use of Bayes' theorem, the updating is carried out by forming the product of a likelihood function and the prior PDF [20]. The likelihood function, in the case of indications obtained independently, is the product of functions, one function for each indication and identical in form, e.g. to a Gaussian PDF. The posterior PDF is then determined by integrating the product of prior PDF and likelihood over all possible values of the variance and normalizing the resulting expression.

NOTE 1 In some cases (e.g. as in 6.4.11), the random variables, of which the indications are regarded as realizations, are characterized by a PDF with only one parameter. In such cases, a non-informative prior PDF is assigned to the unknown expectation of the random variables, and the posterior distribution for X is given directly by Bayes' theorem, without the need for marginalisation.

NOTE 2 Bayes' theorem can also be applied in other circumstances, e.g. when the expectation and standard deviation are unknown and equal.

6.3 Principle of maximum entropy

6.3.1 When using the principle of maximum entropy, introduced by Jaynes [25], a unique PDF is selected among all possible PDFs having specified properties, e.g. specified central moments of different orders or specified intervals for which the PDF is non-zero. This method is particularly useful for assigning PDFs to quantities for which a series of indications is not available or to quantities that have not explicitly been measured at all.

6.3.2 In applying the principle of maximum entropy, to obtain a PDF $g_X(\xi)$ that adequately characterizes incomplete knowledge about a quantity X according to the information available, the functional

$$S[g] = -\int g_X(\xi) \ln g_X(\xi)\, \mathrm{d}\xi,$$

the "information entropy", introduced by Shannon [48], is maximized under constraints given by the information.

6.4 Probability density function assignment for some common circumstances

6.4.1 General

Subclauses 6.4.2 to 6.4.11 provide assignments of PDFs to quantities based on various types of information regarding those quantities. Given for each PDF $g_X(\xi)$ are

a) formulæ for the expectation and variance of X, and

b) the manner in which sampling from $g_X(\xi)$ can be undertaken.

Table 1 facilitates the use of these subclauses and also illustrates the corresponding PDFs.

NOTE These illustrations of the PDFs are not drawn to scale. The multivariate Gaussian PDF is not illustrated.

6.4.2 Rectangular distributions

6.4.2.1 If the only available information regarding a quantity X is a lower limit a and an upper limit b with $a < b$, then, according to the principle of maximum entropy, a rectangular distribution $\mathrm{R}(a, b)$ over the interval $[a, b]$ would be assigned to X.

6.4.2.2 The PDF for X is

$$g_X(\xi) = \begin{cases} 1/(b-a), & a \le \xi \le b, \\ 0, & \text{otherwise.} \end{cases}$$

6.4.2.3 X has expectation and variance

$$E(X) = \frac{a+b}{2}, \qquad V(X) = \frac{(b-a)^2}{12}. \tag{2}$$

6.4.2.4 To sample from $\mathrm{R}(a, b)$, make a draw r from the standard rectangular distribution $\mathrm{R}(0, 1)$ (see C.3.3), and form

$$\xi = a + (b-a)r.$$

6.4.3 Rectangular distributions with inexactly prescribed limits

6.4.3.1 A quantity X is known to lie between limits A and B with $A < B$, where the midpoint $(A + B)/2$ of the interval defined by these limits is fixed and the length $B - A$ of the interval is not known exactly. A is known to lie in the interval $a \pm d$ and B in $b \pm d$, where a, b and d, with $d > 0$ and $a + d < b - d$, are specified. If no other information is available concerning X, A and B, the principle of maximum entropy can be applied to assign to X a "curvilinear trapezoid" (a rectangular distribution with inexactly prescribed limits).

6.4.3.2 The PDF for X is

$$g_X(\xi) = \frac{1}{4d} \begin{cases} \ln[(w+d)/(x-\xi)], & a-d \le \xi \le a+d, \\ \ln[(w+d)/(w-d)], & a+d < \xi < b-d, \\ \ln[(w+d)/(\xi-x)], & b-d \le \xi \le b+d, \\ 0, & \text{otherwise,} \end{cases} \tag{3}$$

where $x = (a + b)/2$ and $w = (b - a)/2$ are, respectively, the midpoint and semi-width of the interval $[a, b]$ [GUM:1995 4.3.9 note 2]. This PDF is trapezoidal-like, but has flanks that are not straight lines.

Table 1 — Available information and the PDF assigned on the basis of that information (6.4.1, C.1.2)

Available information	Assigned PDF and illustration (not to scale)	Subclause
Lower and upper limits a, b	Rectangular: R(a, b) 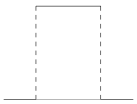	6.4.2
Inexact lower and upper limits $a \pm d$, $b \pm d$	Curvilinear trapezoid: CTrap(a, b, d) 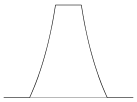	6.4.3
Sum of two quantities assigned rectangular distributions with lower and upper limits a_1, b_1 and a_2, b_2	Trapezoidal: Trap(a, b, β) with $a = a_1 + a_2$, $b = b_1 + b_2$, $\beta = \lvert(b_1 - a_1) - (b_2 - a_2)\rvert/(b - a)$ 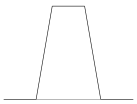	6.4.4
Sum of two quantities assigned rectangular distributions with lower and upper limits a_1, b_1 and a_2, b_2 and the same semi-width $(b_1 - a_1 = b_2 - a_2)$	Triangular: T(a, b) with $a = a_1 + a_2$, $b = b_1 + b_2$ 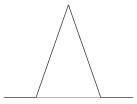	6.4.5
Sinusoidal cycling between lower and upper limits a, b	Arc sine (U-shaped): U(a, b) 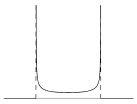	6.4.6
Best estimate x and associated standard uncertainty $u(x)$	Gaussian: N$(x, u^2(x))$ 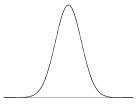	6.4.7
Best estimate $\boldsymbol{x}$ of vector quantity and associated uncertainty matrix $\boldsymbol{U_x}$	Multivariate Gaussian: N$(\boldsymbol{x}, \boldsymbol{U_x})$	6.4.8
Series of indications $x_1, \ldots, x_n$ sampled independently from a quantity having a Gaussian distribution, with unknown expectation and unknown variance	Scaled and shifted t: $t_{n-1}(\bar{x}, s^2/n)$ with $\bar{x} = \sum_{i=1}^{n} x_i/n$, $s^2 = \sum_{i=1}^{n} (x_i - \bar{x})^2/(n-1)$ 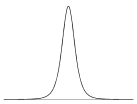	6.4.9.2
Best estimate x, expanded uncertainty U_p, coverage factor k_p and effective degrees of freedom ν_{eff}	Scaled and shifted t: $t_{\nu_{\text{eff}}}(x, (U_p/k_p)^2)$ 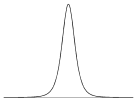	6.4.9.7
Best estimate x of non-negative quantity	Exponential: Ex$(1/x)$ 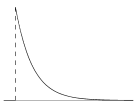	6.4.10
Number q of objects counted	Gamma: G$(q + 1, 1)$ 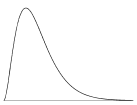	6.4.11

NOTE Formula (3) can be expressed as

$$g_X(\xi) = \frac{1}{4d} \max\left(\ln \frac{w+d}{\max(|\xi - x|, w-d)},\ 0\right)$$

for computer implementation.

6.4.3.3 X has expectation and variance

$$E(X) = \frac{a+b}{2}, \qquad V(X) = \frac{(b-a)^2}{12} + \frac{d^2}{9}. \tag{4}$$

NOTE 1 The variance in expression (4) is always greater than the variance holding for exact limits in expression (2), i.e. when $d = 0$.

NOTE 2 The GUM treats the information about X in 6.4.3.1 by assigning a degrees of freedom to the standard uncertainty associated with the best estimate of X [GUM:1995 G.4.2].

6.4.3.4 To sample from CTrap(a, b, d), make two draws r_1 and r_2 independently from the standard rectangular distribution R$(0, 1)$ (see C.3.3), and form

$$a_\mathrm{s} = (a-d) + 2dr_1, \qquad b_\mathrm{s} = (a+b) - a_\mathrm{s},$$

and

$$\xi = a_\mathrm{s} + (b_\mathrm{s} - a_\mathrm{s})r_2.$$

NOTE a_s is a draw from the rectangular distribution with limits $a \pm d$. b_s is then formed to ensure that the midpoint of a_s and b_s is the prescribed value $x = (a+b)/2$.

EXAMPLE A certificate states that a voltage X lies in the interval 10.0 V ± 0.1 V. No other information is available concerning X, except that it is believed that the magnitude of the interval endpoints is the result of rounding correctly some numerical value (see 3.20). On this basis, that numerical value lies between 0.05 V and 0.15 V, since the numerical value of every point in the interval (0.05, 0.15) rounded to one significant decimal digit is 0.1. The location of the interval can therefore be regarded as fixed, whereas its width is inexact. The best estimate of X is $x = 10.0$ V and, using expression (4) based on $a = 9.9$ V, $b = 10.1$ V and $d = 0.05$ V, the associated standard uncertainty $u(x)$ is given by

$$u^2(x) = \frac{(0.2)^2}{12} + \frac{(0.05)^2}{9} = 0.003\ 6.$$

Hence $u(x) = (0.003\ 6)^{1/2} = 0.060$ V, which can be compared with $0.2/\sqrt{12} = 0.058$ V in the case of exact limits, given by replacing d by zero. The use of exact limits in this case gives a numerical value for $u(x)$ that is 4 % smaller than that for inexact limits. The relevance of such a difference needs to be considered in the context of the application.

6.4.4 Trapezoidal distributions

6.4.4.1 The assignment of a symmetric trapezoidal distribution to a quantity is discussed in the GUM [GUM:1995 4.3.9]. Suppose a quantity X is defined as the sum of two independent quantities X_1 and X_2. Suppose, for $i = 1$ and $i = 2$, X_i is assigned a rectangular distribution R(a_i, b_i) with lower limit a_i and upper limit b_i. Then the distribution for X is a symmetric trapezoidal distribution Trap(a, b, β) with lower limit a, upper limit b, and a parameter β equal to the ratio of the semi-width of the top of the trapezoid to that of the base. The parameters of this trapezoidal distribution are related to those of the rectangular distributions by

$$a = a_1 + a_2, \qquad b = b_1 + b_2, \qquad \beta = \frac{\lambda_1}{\lambda_2}, \tag{5}$$

where

$$\lambda_1 = \frac{|(b_1 - a_1) - (b_2 - a_2)|}{2}, \qquad \lambda_2 = \frac{b-a}{2}, \tag{6}$$

and

$$0 \leq \lambda_1 \leq \lambda_2.$$

6.4.4.2 The PDF for X (figure 5), obtained using convolution [42, p93], is

$$g_X(\xi) = \begin{cases} (\xi - x + \lambda_2)/(\lambda_2^2 - \lambda_1^2), & x - \lambda_2 \leq \xi < x - \lambda_1, \\ 1/(\lambda_1 + \lambda_2), & x - \lambda_1 \leq \xi \leq x + \lambda_1, \\ (x + \lambda_2 - \xi)/(\lambda_2^2 - \lambda_1^2), & x + \lambda_1 < \xi \leq x + \lambda_2, \\ 0, & \text{otherwise}, \end{cases} \tag{7}$$

where $x = (a + b)/2$.

NOTE Formula (7) can be expressed as

$$g_X(\xi) = \frac{1}{\lambda_1 + \lambda_2} \min\left(\frac{1}{\lambda_2 - \lambda_1} \max\left(\lambda_2 - |\xi - x|,\ 0\right),\ 1\right)$$

for computer implementation.

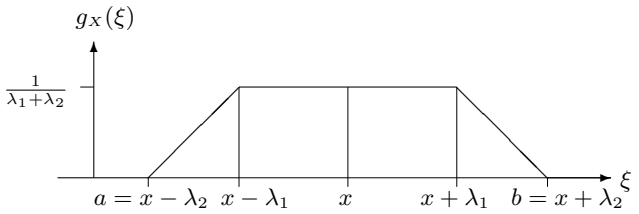

Figure 5 — The trapezoidal PDF for $X = X_1 + X_2$, where the PDFs for X_1 and X_2 are rectangular (6.4.4.2)

6.4.4.3 X has expectation and variance

$$E(X) = \frac{a + b}{2}, \qquad V(X) = \frac{(b - a)^2}{24}(1 + \beta^2).$$

6.4.4.4 To sample from $\mathrm{Trap}(a, b, \beta)$, make two draws r_1 and r_2 independently from the standard rectangular distribution $\mathrm{R}(0, 1)$ (see C.3.3), and form

$$\xi = a + \frac{b - a}{2}[(1 + \beta)r_1 + (1 - \beta)r_2].$$

6.4.5 Triangular distributions

6.4.5.1 Suppose a quantity X is defined as the sum of two independent quantities, each assigned a rectangular distribution (see 6.4.4), but with equal semi-widths, i.e. $b_1 - a_1 = b_2 - a_2$. It follows from expressions (5) and (6) that $\lambda_1 = 0$ and $\beta = 0$. The distribution for X is the trapezoidal distribution $\mathrm{Trap}(a, b, 0)$, which reduces to the (symmetric) triangular distribution $\mathrm{T}(a, b)$ over the interval $[a, b]$.

6.4.5.2 The PDF for X is

$$g_X(\xi) = \begin{cases} (\xi - a)/w^2, & a \leq \xi \leq x, \\ (b - \xi)/w^2, & x < \xi \leq b, \\ 0, & \text{otherwise}, \end{cases} \tag{8}$$

where $x = (a + b)/2$ and $w = \lambda_2 = (b - a)/2$.

NOTE Formula (8) can be expressed as

$$g_X(\xi) = \frac{2}{b-a} \max\left(1 - \frac{2|\xi - x|}{b-a},\ 0\right).$$

for computer implementation.

6.4.5.3 X has expectation and variance

$$E(X) = \frac{a+b}{2}, \qquad V(X) = \frac{(b-a)^2}{24}.$$

6.4.5.4 To sample from $\mathrm{T}(a,b)$, make two draws r_1 and r_2 independently from the standard rectangular distribution $\mathrm{R}(0,1)$ (see C.3.3), and form

$$\xi = a + \frac{b-a}{2}(r_1 + r_2).$$

6.4.6 Arc sine (U-shaped) distributions

6.4.6.1 If a quantity X is known to cycle sinusoidally, with unknown phase Φ, between specified limits a and b, with $a < b$, then, according to the principle of maximum entropy, a rectangular distribution $\mathrm{R}(0, 2\pi)$ would be assigned to Φ. The distribution assigned to X is the arc sine distribution $\mathrm{U}(a,b)$ [18], given by the transformation

$$X = \frac{a+b}{2} + \frac{b-a}{2} \sin \Phi,$$

where Φ has the rectangular distribution $\mathrm{R}(0, 2\pi)$.

6.4.6.2 The PDF for X is

$$g_X(\xi) = \begin{cases} (2/\pi)[(b-a)^2 - (2\xi - a - b)^2]^{-1/2}, & a < \xi < b, \\ 0, & \text{otherwise.} \end{cases}$$

NOTE $\mathrm{U}(a,b)$ is related to the standard arc sine distribution $\mathrm{U}(0,1)$ given by

$$g_Z(z) = \begin{cases} [z(1-z)]^{-1/2}/\pi, & 0 < z < 1, \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

in the variable Z, through the linear transformation

$$X = a + (b-a)Z.$$

Z has expectation $1/2$ and variance $1/8$. The distribution (9) is termed the arc sine distribution, since the corresponding distribution function is

$$G_Z(z) = \frac{1}{\pi} \arcsin\,(2z-1) + \frac{1}{2}.$$

It is a special case of the beta distribution with both parameters equal to one half.

6.4.6.3 X has expectation and variance

$$E(X) = \frac{a+b}{2}, \qquad V(X) = \frac{(b-a)^2}{8}.$$

6.4.6.4 To sample from $\mathrm{U}(a,b)$, make a draw r from the standard rectangular distribution $\mathrm{R}(0,1)$ (see C.3.3), and form

$$\xi = \frac{a+b}{2} + \frac{b-a}{2} \sin 2\pi r.$$

6.4.7 Gaussian distributions

6.4.7.1 If a best estimate x and associated standard uncertainty $u(x)$ are the only information available regarding a quantity X, then, according to the principle of maximum entropy, a Gaussian probability distribution $\mathrm{N}(x, u^2(x))$ would be assigned to X.

6.4.7.2 The PDF for X is

$$g_X(\xi) = \frac{1}{\sqrt{2\pi}u(x)} \exp\left(-\frac{(\xi - x)^2}{2u^2(x)}\right). \tag{10}$$

6.4.7.3 X has expectation and variance

$$E(X) = x, \qquad V(X) = u^2(x).$$

6.4.7.4 To sample from $\mathrm{N}(x, u^2(x))$, make a draw z from the standard Gaussian distribution $\mathrm{N}(0, 1)$ (see C.4), and form

$$\xi = x + u(x)z.$$

6.4.8 Multivariate Gaussian distributions

6.4.8.1 A comparable result to that in 6.4.7.1 holds for an N-dimensional quantity $\boldsymbol{X} = (X_1, \ldots, X_N)^\top$. If the only information available is a best estimate $\boldsymbol{x} = (x_1, \ldots, x_N)^\top$ of $\boldsymbol{X}$ and the associated (strictly) positive definite uncertainty matrix

$$\boldsymbol{U}_{\boldsymbol{x}} = \begin{bmatrix} u^2(x_1) & u(x_1, x_2) & \cdots & u(x_1, x_N) \\ u(x_2, x_1) & u^2(x_2) & \cdots & u(x_2, x_N) \\ \vdots & \vdots & \ddots & \vdots \\ u(x_N, x_1) & u(x_N, x_2) & \cdots & u^2(x_N) \end{bmatrix},$$

a multivariate Gaussian distribution $\mathrm{N}(\boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{x}})$ would be assigned to $\boldsymbol{X}$.

6.4.8.2 The joint PDF for $\boldsymbol{X}$ is

$$g_{\boldsymbol{X}}(\boldsymbol{\xi}) = \frac{1}{[(2\pi)^N \det \boldsymbol{U}_{\boldsymbol{x}}]^{1/2}} \exp\left(-\frac{1}{2}(\boldsymbol{\xi} - \boldsymbol{x})^\top \boldsymbol{U}_{\boldsymbol{x}}{}^{-1}(\boldsymbol{\xi} - \boldsymbol{x})\right). \tag{11}$$

6.4.8.3 $\boldsymbol{X}$ has expectation and covariance matrix

$$E(\boldsymbol{X}) = \boldsymbol{x}, \qquad V(\boldsymbol{X}) = \boldsymbol{U}_{\boldsymbol{x}}.$$

6.4.8.4 To sample from $\mathrm{N}(\boldsymbol{x}, \boldsymbol{U}_{\boldsymbol{x}})$, make N draws z_i, $i = 1, \ldots, N$, independently from the standard Gaussian distribution $\mathrm{N}(0, 1)$ (see C.4), and form

$$\boldsymbol{\xi} = \boldsymbol{x} + \boldsymbol{R}^\top \boldsymbol{z},$$

where $\boldsymbol{z} = (z_1, \ldots, z_N)^\top$ and $\boldsymbol{R}$ is the upper triangular matrix given by the Cholesky decomposition $\boldsymbol{U}_{\boldsymbol{x}} = \boldsymbol{R}^\top \boldsymbol{R}$ (see C.5).

NOTE 1 In place of the Cholesky decomposition $\boldsymbol{U}_{\boldsymbol{x}} = \boldsymbol{R}^\top \boldsymbol{R}$, any matrix factorization of this form can be used.

NOTE 2 The only joint PDFs considered explicitly in this Supplement are multivariate Gaussian, distributions commonly used in practice. A numerical procedure for sampling from a multivariate Gaussian PDF is given above (and in C.5). If another multivariate PDF is to be used, a means for sampling from it would need to be provided.

NOTE 3 The multivariate Gaussian PDF (11) reduces to the product of N univariate Gaussian PDFs when there are no covariance effects. In that case

$$\boldsymbol{U}_{\boldsymbol{x}} = \mathrm{diag}(u^2(x_1), \ldots, u^2(x_N)),$$

whence

$$g_{\boldsymbol{X}}(\boldsymbol{\xi}) = \prod_{i=1}^{N} g_{X_i}(\xi_i),$$

with

$$g_{X_i}(\xi_i) = \frac{1}{\sqrt{2\pi}u(x_i)} \exp\left(-\frac{(\xi_i - x_i)^2}{2u^2(x_i)}\right).$$

6.4.9 *t*-distributions

6.4.9.1 t-distributions typically arise in two circumstances: the evaluation of a series of indications (see 6.4.9.2), and the interpretation of calibration certificates (see 6.4.9.7).

6.4.9.2 Suppose that a series of n indications $x_1, \ldots, x_n$ is available, regarded as being obtained independently from a quantity with unknown expectation μ_0 and unknown variance σ_0^2 having Gaussian distribution $\mathrm{N}(\mu_0, \sigma_0^2)$. The desired input quantity X is taken to be equal to μ_0. Then, assigning a non-informative joint prior distribution to μ_0 and σ_0^2, and using Bayes' theorem, the marginal PDF for X is a scaled and shifted t-distribution $\mathrm{t}_\nu(\bar{x}, s^2/n)$ with $\nu = n - 1$ degrees of freedom, where

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i, \qquad s^2 = \frac{1}{n-1}\sum_{i=1}^{n} (x_i - \bar{x})^2,$$

being, respectively, the average and variance of the indications [20].

6.4.9.3 The PDF for X is

$$g_X(\xi) = \frac{\Gamma(n/2)}{\Gamma((n-1)/2)\sqrt{(n-1)\pi}} \times \frac{1}{s/\sqrt{n}} \left(1 + \frac{1}{n-1}\left(\frac{\xi - \bar{x}}{s/\sqrt{n}}\right)^2\right)^{-n/2}, \tag{12}$$

where

$$\Gamma(z) = \int_0^\infty t^{z-1}\mathrm{e}^{-t}\,\mathrm{d}t, \qquad z > 0,$$

is the gamma function.

6.4.9.4 X has expectation and variance

$$E(X) = \bar{x}, \qquad V(X) = \frac{n-1}{n-3}\frac{s^2}{n},$$

where $E(X)$ is defined only for $n > 2$ and $V(X)$ only for $n > 3$. For $n > 3$, the best estimate of X and its associated standard uncertainty are therefore

$$x = \bar{x}, \qquad u(x) = \sqrt{\frac{n-1}{n-3}}\frac{s}{\sqrt{n}}. \tag{13}$$

NOTE 1 In the GUM [GUM:1995 4.2], the standard uncertainty $u(x)$ associated with the average of a series of n indications obtained independently is evaluated as $u(x) = s/\sqrt{n}$, rather than from formula (13), and the associated degrees of freedom $\nu = n - 1$ is considered as a measure of the reliability of $u(x)$. By extension, a degrees of freedom is associated with an uncertainty

obtained from a Type B evaluation, based on subjective judgement of the reliability of the evaluation [GUM:1995 G.4.2] (cf. 6.4.3.3 note 2). Degrees of freedom associated with the uncertainties $u(x_i)$ are necessary to obtain, by application of the Welch-Satterthwaite formula, the effective degrees of freedom ν_{eff} associated with the uncertainty $u(y)$.

NOTE 2 In the Bayesian context of this Supplement, concepts such as the reliability, or the uncertainty, of an uncertainty are not necessary. Accordingly, the degrees of freedom in a Type A evaluation of uncertainty is no longer viewed as a measure of reliability, and the degrees of freedom in a Type B evaluation does not exist.

6.4.9.5 To sample from $t_\nu(\bar{x}, s^2/n)$, make a draw t from the central t-distribution t_ν with $\nu = n - 1$ degrees of freedom [GUM:1995 G.3] (also see C.6), and form

$$\xi = \bar{x} + \frac{s}{\sqrt{n}} t.$$

6.4.9.6 If instead of a standard deviation s calculated from a single series of indications, a pooled standard deviation s_{p} with ν_{p} degrees of freedom obtained from Q such sets,

$$s_{\text{p}}^2 = \frac{1}{\nu_{\text{p}}} \sum_{j=1}^{Q} \nu_j s_j^2, \qquad \nu_{\text{p}} = \sum_{j=1}^{Q} \nu_j,$$

is used, the degrees of freedom $\nu = n - 1$ of the scaled and shifted t-distribution assigned to X should be replaced by the degrees of freedom ν_{p} associated with the pooled standard deviation s_{p}. As a consequence, formula (12) should be replaced by

$$g_X(\xi) = \frac{\Gamma((\nu_{\text{p}} + 1)/2)}{\Gamma(\nu_{\text{p}}/2)\sqrt{\nu_{\text{p}}\pi}} \times \frac{1}{s_{\text{p}}/\sqrt{n}} \left[1 + \frac{1}{\nu_{\text{p}}} \left(\frac{\xi - \bar{x}}{s_{\text{p}}/\sqrt{n}} \right)^2 \right]^{-(\nu_{\text{p}}+1)/2}$$

and expressions (13) by

$$x = \bar{x} = \frac{1}{n} \sum_{i=1}^{n} x_i, \qquad u(x) = \sqrt{\frac{\nu_{\text{p}}}{\nu_{\text{p}} - 2}} \frac{s_{\text{p}}}{\sqrt{n}} \quad (\nu_{\text{p}} \geq 3).$$

6.4.9.7 If the source of information about a quantity X is a calibration certificate [GUM:1995 4.3.1] in which a best estimate x, the expanded uncertainty U_p, the coverage factor k_p and the effective degrees of freedom ν_{eff} are stated, then a scaled and shifted t-distribution $t_\nu(x, (U_p/k_p)^2)$ with $\nu = \nu_{\text{eff}}$ degrees of freedom should be assigned to X.

6.4.9.8 If ν_{eff} is stated as infinite or not specified, in which case it would be taken as infinite in the absence of other information, a Gaussian distribution $\text{N}(x, (U_p/k_p)^2)$ would be assigned to X (see 6.4.7.1).

NOTE This distribution is the limiting case of the scaled and shifted t-distribution $\text{t}_\nu(x, (U_p/k_p)^2)$ as ν tends to infinity.

6.4.10 Exponential distributions

6.4.10.1 If the only available information regarding a non-negative quantity X is a best estimate $x > 0$ of X, then, according to the principle of maximum entropy, an exponential distribution $\text{Ex}(1/x)$ would be assigned to X.

6.4.10.2 The PDF for X is

$$g_X(\xi) = \begin{cases} \exp(-\xi/x)/x, & \xi \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

6.4.10.3 X has expectation and variance

$$E(X) = x, \qquad V(X) = x^2.$$

6.4.10.4 To sample from $\mathrm{Ex}(1/x)$, make a draw r from the standard rectangular distribution $\mathrm{R}(0,1)$ (see C.3.3), and form

$$\xi = -x \ln r.$$

NOTE Further information regarding the assignment of PDFs to non-negative quantities is available [14].

6.4.11 Gamma distributions

6.4.11.1 Suppose the quantity X is the average number of objects present in a sample of a fixed size (e.g. the average number of particles in an air sample taken from a clean room, or the average number of photons emitted by a source in a specified time interval). Suppose q is the number of objects counted in a sample of the specified size, and the counted number is assumed to be a quantity with unknown expectation having a Poisson distribution. Then, according to Bayes' theorem, after assigning a constant prior distribution to the expectation, a gamma distribution $\mathrm{G}(q+1,1)$ would be assigned to X.

6.4.11.2 The PDF for X is

$$g_X(\xi) = \begin{cases} \xi^q \exp(-\xi)/q!, & \xi \geq 0, \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

6.4.11.3 X has expectation and variance

$$E(X) = q+1, \qquad V(X) = q+1. \tag{15}$$

6.4.11.4 To sample from $\mathrm{G}(q+1,1)$, make $q+1$ draws r_i, $i = 1, \ldots, q+1$, independently from the standard rectangular distribution $\mathrm{R}(0,1)$ (see C.3.3), and form [18]

$$\xi = -\ln \prod_{i=1}^{q+1} r_i.$$

NOTE 1 If the counting is performed over several samples (according to the same Poisson distribution), and q_i is the number of objects counted in the ith sample, of size S_i, then the distribution for the average number of objects in a sample of size $S = \sum_i S_i$ is $\mathrm{G}(\alpha, \beta)$ with $\alpha = 1 + \sum_i q_i$ and $\beta = 1$. Formulæ (14) and (15) apply with $q = \sum_i q_i$.

NOTE 2 The gamma distribution is a generalization of the chi-squared distribution and is used to characterize information associated with variances.

NOTE 3 The particular gamma distribution in 6.4.11.4 is an Erlang distribution given by the sum of $q+1$ exponential distributions with parameter 1 [18].

6.5 Probability distributions from previous uncertainty calculations

A previous uncertainty calculation may have provided a probability distribution for an output quantity that is to become an input quantity for a further uncertainty calculation. This probability distribution may be available analytically in a recognized form, e.g. as a Gaussian PDF. It may be available as an approximation to the distribution function for a quantity obtained from a previous application of MCM, for example. Means for describing such a distribution function for a quantity is given in 7.5.1 and D.2.

7 Implementation of a Monte Carlo method

7.1 General

This clause gives information about the implementation of a Monte Carlo method for the propagation of distributions: see the procedure given in 5.9.6 and shown diagrammatically in figure 4.

7.2 Number of Monte Carlo trials

7.2.1 A value of M, the number of Monte Carlo trials, i.e. the number of model evaluations to be made, needs to be selected. It can be chosen *a priori*, in which case there will be no direct control over the quality of the numerical results provided by MCM. The reason is that the number of trials needed to provide these results to a prescribed numerical tolerance will depend on the "shape" of the PDF for the output quantity and on the coverage probability required. Also, the calculations are stochastic in nature, being based on random sampling.

NOTE A value of $M = 10^6$ can often be expected to deliver a 95 % coverage interval for the output quantity such that this length is correct to one or two significant decimal digits.

7.2.2 The choice of a value of M that is large compared with $1/(1-p)$, e.g. M at least 10^4 times greater than $1/(1-p)$, should be made. It can then be expected that $\boldsymbol{G}$ will provide a reasonable discrete representation of $G_Y(\eta)$ in the regions near the endpoints of a $100p$ % coverage interval for Y.

7.2.3 Because there is no guarantee that this or any specific pre-assigned number will suffice, a procedure that selects M adaptively, i.e. as the trials progress, can be used. Some guidance in this regard is available [2]. Subclause 7.9 provides such a procedure, a property of which is that the number of trials taken is economically consistent with the expectation of achieving a required numerical tolerance.

NOTE If the model is complicated, e.g. involving the solution of a finite-element model, because of large computing times it may not be possible to use a sufficiently large value of M to obtain adequate distributional knowledge of the output quantity. In such a case an approximate approach would be to regard $g_Y(\eta)$ as Gaussian (as in the GUM) and proceed as follows. A relatively small value of M, 50 or 100, for example, would be used. The average and standard deviation of the resulting M model values of Y would be taken as y and $u(y)$, respectively. Given this information, a Gaussian PDF $g_Y(\eta) = \mathrm{N}(y, u^2(y))$ would be assigned to characterize the knowledge of Y (see 6.4.7) and a desired coverage interval for Y calculated. Although this use of a small value of M is inevitably less reliable than that of a large value in that it does not provide an approximation to the PDF for Y, it does take account of model non-linearity.

7.3 Sampling from probability distributions

In an implementation of MCM, M vectors $\boldsymbol{x}_r$, $r = 1, \ldots, M$ (see 7.2), are drawn from the PDFs $g_{X_i}(\xi_i)$ for the N input quantities X_i. Draws would be made from the joint (multivariate) PDF $g_{\boldsymbol{X}}(\boldsymbol{\xi})$ if appropriate. Recommendations concerning the manner in which this sampling can be carried out are given in annex C for the commonest distributions, viz. the rectangular, Gaussian, t, and multivariate Gaussian. Also see 6.4. It is possible to draw at random from any other distribution. See C.2. Some such distributions could be approximations to distributions based on Monte Carlo results from a previous uncertainty calculation (see 6.5, 7.5 and annex D).

NOTE For the results of MCM to be statistically valid, it is necessary that the pseudo-random number generators used to draw from the distributions required have appropriate properties. Some tests of randomness of the numbers produced by a generator are indicated in C.3.2.

7.4 Evaluation of the model

7.4.1 The model is evaluated for each of the M draws from the PDFs for the N input quantities. Specifically, denote the M draws by $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M$, where the rth draw $\boldsymbol{x}_r$ contains $x_{1,r}, \ldots, x_{N,r}$, with $x_{i,r}$ a draw from the PDF for X_i. Then, the model values are

$$y_r = f(\boldsymbol{x}_r), \; r = 1, \ldots, M.$$

7.4.2 The necessary modifications are made to 7.4.1 if the X_i are not independent and hence a joint PDF is assigned to them.

NOTE Model and derivative evaluations are made when applying the law of propagation of uncertainty, using exact derivatives, at the best estimates of the input quantities. Model evaluations only are made when applying the law of propagation of uncertainty when numerical (finite-difference) approximations to derivatives are used. These evaluations are made, if the GUM recommendation [GUM:1995 5.1.3 note 2] is adopted, at the best estimates of the input quantities and at points perturbed by ± one standard uncertainty from each estimate in turn. With MCM, model evaluations are made in the neighbourhood of

these best estimates, viz. at points that can be expected to be up to several standard uncertainties away from these estimates. The fact that model evaluations are made at different points according to the approach used may raise issues regarding the numerical procedure used to evaluate the model, e.g. ensuring its convergence (where iterative schemes are used) and numerical stability. The user should ensure that, where appropriate, the numerical methods used to evaluate f are valid for a sufficiently large region containing these best estimates. Only occasionally would it be expected that this aspect is critical.

7.5 Discrete representation of the distribution function for the output quantity

7.5.1 A discrete representation $\boldsymbol{G}$ of the distribution function $G_Y(\eta)$ for the output quantity Y can be obtained as follows:

a) sort the model values y_r, $r = 1, \ldots, M$, provided by MCM into non-decreasing order. Denote the sorted model values by $y_{(r)}$, $r = 1, \ldots, M$;

b) if necessary, make minute numerical perturbations to any replicate model values $y_{(r)}$ in such a way that the resulting complete set of $y_{(r)}$, $r = 1, \ldots, M$, form a strictly increasing sequence (cf. condition b) in 5.10.1);

c) take $\boldsymbol{G}$ as the set $y_{(r)}$, $r = 1, \ldots, M$.

NOTE 1 With reference to step a), a sorting algorithm taking a number of operations proportional to $M \ln M$ should be used [47]. A naive algorithm would take a time proportional to M^2, making the computation time unnecessarily long. See 7.8.

NOTE 2 In step a), the term "non-decreasing" rather than "increasing" is used because of possible equalities among the model values y_r.

NOTE 3 With reference to step b), making only minute perturbations will ensure that the statistical properties of the $y_{(r)}$ are retained.

NOTE 4 In step b), it is exceedingly unlikely that perturbations are necessary, because of the very large number of distinct floating-point numbers that can arise from model values generated from input quantities obtained as draws from random number generators. A sound software implementation would make appropriate provision, however.

NOTE 5 With reference to step c), a variety of information can be deduced from $\boldsymbol{G}$. In particular, information supplementary to the expectation and standard deviation can be provided, such as measures of skewness and kurtosis, and other statistics such as the mode and the median.

NOTE 6 If Y is to become an input quantity for a further uncertainty calculation, sampling from its probability distribution is readily carried out by drawing randomly from the $y_{(r)}$, $r = 1, \ldots, M$, with equal probability (see 6.5).

7.5.2 The $y_{(r)}$ (or y_r), when assembled into a histogram (with suitable cell widths) form a frequency distribution that, when normalized to have unit area, provides an approximation to the PDF $g_Y(\eta)$ for Y. Calculations are not generally carried out in terms of this histogram, the resolution of which depends on the choice of cell widths, but in terms of $\boldsymbol{G}$. The histogram can, however, be useful as an aid to understanding the nature of the PDF, e.g. the extent of its asymmetry. See, however, 7.8.3 note 1 regarding the use of a large numerical value of M.

7.5.3 A continuous approximation to $G_Y(\eta)$ is sometimes useful. Annex D contains a means for obtaining such an approximation.

7.6 Estimate of the output quantity and the associated standard uncertainty

The average

$$\widetilde{y} = \frac{1}{M} \sum_{r=1}^{M} y_r \tag{16}$$

and standard deviation $u(\widetilde{y})$ determined from

$$u^2(\widetilde{y}) = \frac{1}{M-1}\sum_{r=1}^{M}(y_r - \widetilde{y})^2 \tag{17}$$

are taken, respectively, as an estimate y of Y and the standard uncertainty $u(y)$ associated with y.

NOTE 1 Formula (17) should be used rather than the mathematically equivalent formula

$$u^2(\widetilde{y}) = \frac{M}{M-1}\left(\frac{1}{M}\sum_{r=1}^{M} y_r^2 - \widetilde{y}^2\right).$$

For the many circumstances in metrology in which $u(y)$ is much smaller than $|y|$ (in which case the y_r have a number of leading decimal digits in common) the latter formula suffers numerically from subtractive cancellation (involving a mean square less a squared mean). This effect can be so severe that the resulting numerical value might have too few correct significant decimal digits for the uncertainty evaluation to be valid [4].

NOTE 2 In some special circumstances, such as when one of the input quantities has been assigned a PDF based on the t-distribution with fewer than three degrees of freedom, the expectation and standard deviation of Y, as described by the PDF $g_Y(\eta)$, might not exist. Formulæ (16) and (17) might not then provide meaningful results. A coverage interval for Y (see 7.7) can, however, be formed, since $\boldsymbol{G}$ is meaningful and can be determined.

NOTE 3 $\widetilde{y}$ will not in general agree with the model evaluated at the best estimates of the input quantities, since, for a non-linear model $f(\boldsymbol{X})$, $E(Y) = E(f(\boldsymbol{X})) \neq f(E(\boldsymbol{X}))$ (cf. [GUM:1995 4.1.4]). Irrespective of whether f is linear or non-linear, in the limit as M tends to infinity, $\widetilde{y}$ approaches $E(f(\boldsymbol{X}))$ when $E(f(\boldsymbol{X}))$ exists.

7.7 Coverage interval for the output quantity

7.7.1 A coverage interval for Y can be determined from the discrete representation $\boldsymbol{G}$ of $G_Y(\eta)$ in an analogous manner to that in 5.3.2 given $G_Y(\eta)$.

7.7.2 Let $q = pM$, if pM is an integer. Otherwise, take q to be the integer part of $pM + 1/2$. Then $[y_{\mathrm{low}},\ y_{\mathrm{high}}]$ is a $100p$ % coverage interval for Y, where, for any $r = 1, \ldots, M-q$, $y_{\mathrm{low}} = y_{(r)}$ and $y_{\mathrm{high}} = y_{(r+q)}$. The probabilistically symmetric $100p$ % coverage interval is given by taking $r = (M-q)/2$, if $(M-q)/2$ is an integer, or the integer part of $(M-q+1)/2$, otherwise. The shortest $100p$ % coverage interval is given by determining r^* such that, for $r = 1, \ldots, M-q$, $y_{(r^*+q)} - y_{(r^*)} \leq y_{(r+q)} - y_{(r)}$.

NOTE Because of the randomness in MCM, some of these $M-q$ interval lengths will be shorter than they would be on average, and some longer. So, by choosing the least such length, (the approximation to) the shortest $100p$ % coverage interval tends to be marginally shorter than that which would have been calculated from $G_Y(\eta)$, with the consequence that the typical coverage probability is less than $100p$ %. For large M, this coverage probability is negligibly less than $100p$ %.

EXAMPLE 10^5 numbers were drawn from a pseudo-random number generator for the rectangular distribution in the interval $[0, 1]$, and the shortest 95 % coverage interval formed as above. This exercise was carried out 1 000 times. The average coverage probability was 94.92 % and the standard deviation of the 1 000 coverage probabilities 0.06 %.

7.8 Computation time

7.8.1 The computation time for MCM is dominated by that required for the following three steps:

a) make M draws from the PDF for each input quantity X_i (or the joint PDF for $\boldsymbol{X}$);

b) make M corresponding evaluations of the model;

c) sort the resulting M model values into non-decreasing order.

7.8.2 The times taken in the three steps are directly proportional to (a) M, (b) M, and (c) $M \ln M$ (if an efficient sort algorithm [47] is used).

7.8.3 If the model is simple and the input quantities are independent, the time in step c) can be expected to dominate, and the overall time taken is typically a few seconds for $M = 10^6$ on a personal computer operating at several GHz. Otherwise, let T_1 be the time taken to make one draw from the PDFs for the input quantities and T_2 that to make one evaluation of the model. Then, the overall time can be taken as essentially $M \times (T_1 + T_2)$, which, if the model is complicated, is dominated by the term MT_2.

NOTE 1 If the model is simple and M very large, e.g. 10^8 or 10^9, the sorting time may be excessive compared with the time taken to make the M model evaluations. In such a case, calculations can instead be based on an approximation to $g_Y(\eta)$ derived from a suitable histogram of the y_r.

NOTE 2 An indication of the computation time required for an application of MCM can be obtained as follows. Consider an artificial problem with a model consisting of the sum of five terms:

$$Y = \cos X_1 + \sin X_2 + \tan^{-1} X_3 + \exp(X_4) + X_5^{1/3}.$$

Assign a Gaussian PDF to each input quantity X_i. Make $M = 10^6$ Monte Carlo trials. The relative computation times for (a) generating $5M$ random Gaussian numbers, (b) forming M model values and (c) sorting the M model values were respectively 20 %, 20 % and 60 %, with a total computation time of a few seconds on a personal computer operating at several GHz.

7.9 Adaptive Monte Carlo procedure

7.9.1 General

A basic implementation of an adaptive Monte Carlo procedure involves carrying out an increasing number of Monte Carlo trials until the various results of interest have stabilized in a statistical sense. A numerical result is deemed to have stabilized if twice the standard deviation associated with it is less than the numerical tolerance (see 7.9.2) associated with the standard uncertainty $u(y)$.

7.9.2 Numerical tolerance associated with a numerical value

Let n_{dig} denote the number of significant decimal digits regarded as meaningful in a numerical value z. The numerical tolerance δ associated with z is given as follows:

a) express z in the form $c \times 10^{\ell}$, where c is an n_{dig} decimal digit integer and ℓ an integer;

b) take

$$\delta = \frac{1}{2} 10^{\ell}. \tag{18}$$

EXAMPLE 1 The estimate of the output quantity for a nominally 100 g measurement standard of mass [GUM:1995 7.2.2] is $y = 100.021\ 47$ g. The standard uncertainty $u(y) = 0.000\ 35$ g, both significant digits being regarded as meaningful. Thus, $n_{\mathrm{dig}} = 2$ and $u(y)$ can be expressed as 35×10^{-5} g, and so $c = 35$ and $\ell = -5$. Take $\delta = \frac{1}{2} \times 10^{-5}$ g $= 0.000\ 005$ g.

EXAMPLE 2 As example 1 except that only one significant decimal digit in $u(y)$ is regarded as meaningful. Thus, $n_{\mathrm{dig}} = 1$ and $u(y) = 0.000\ 4$ g $= 4 \times 10^{-4}$ g, giving $c = 4$ and $\ell = -4$. Hence, $\delta = \frac{1}{2} \times 10^{-4}$ g $= 0.000\ 05$ g.

EXAMPLE 3 In a temperature measurement, $u(y) = 2$ K. Then, $n_{\mathrm{dig}} = 1$ and $u(y) = 2 \times 10^{0}$ K, giving $c = 2$ and $\ell = 0$. Thus, $\delta = \frac{1}{2} \times 10^{0}$ K $= 0.5$ K.

7.9.3 Objective of adaptive procedure

The objective of the adaptive procedure given in 7.9.4 is to provide

a) an estimate y of Y,

b) an associated standard uncertainty $u(y)$, and

c) the endpoints y_{low} and y_{high} of a coverage interval for Y corresponding to a stipulated coverage probability

such that each of these four values can be expected to meet the numerical tolerance required.

NOTE 1 By its stochastic nature, the procedure cannot be guaranteed to provide such an interval.

NOTE 2 y and $u(y)$ generally "converge" considerably faster than y_{low} and y_{high} with respect to the number of Monte Carlo trials.

NOTE 3 Generally, the larger is the coverage probability, the larger is the number of Monte Carlo trials required to determine y_{low} and y_{high} for a given numerical tolerance.

7.9.4 Adaptive procedure

A practical approach, involving carrying out a sequence of applications of MCM, is as follows:

a) set n_{dig} to an appropriate small positive integer (see 7.9.2);

b) set

$$M = \max(J, 10^4),$$

where J is the smallest integer greater than or equal to $100/(1-p)$;

c) set $h = 1$, denoting the first application of MCM in the sequence;

d) carry out M Monte Carlo trials, as in 7.3 and 7.4;

e) use the M model values $y_1, \ldots, y_M$ so obtained to calculate, as in 7.5 to 7.7, $y^{(h)}$, $u(y^{(h)})$, $y_{\text{low}}^{(h)}$ and $y_{\text{high}}^{(h)}$ as an estimate of Y, the associated standard uncertainty, and the left- and right-hand endpoints of a $100p$ % coverage interval, respectively, i.e. for the hth member of the sequence;

f) if $h = 1$, increase h by one and return to step d);

g) calculate the standard deviation s_y associated with the average of the estimates $y^{(1)}, \ldots, y^{(h)}$ of Y, given by

$$s_y^2 = \frac{1}{h(h-1)} \sum_{r=1}^{h} (y^{(r)} - y)^2,$$

where

$$y = \frac{1}{h} \sum_{r=1}^{h} y^{(r)};$$

h) calculate the counterpart of this statistic for $u(y)$, y_{low} and y_{high};

i) use all $h \times M$ model values available so far to form $u(y)$;

j) calculate the numerical tolerance δ associated with $u(y)$ as in 7.9.2;

k) if any of $2s_y$, $2s_{u(y)}$, $2s_{y_{\text{low}}}$ and $2s_{y_{\text{high}}}$ exceeds δ, increase h by one and return to step d);

l) regard the overall computation as having stabilized, and use all $h \times M$ model values obtained to calculate y, $u(y)$ and a $100p$ % coverage interval, as in 7.5 to 7.7.

NOTE 1 Normally n_{dig} in step a) would be chosen to be 1 or 2.

NOTE 2 The choice of M in step b) is arbitrary, but has been found suitable in practice.

NOTE 3 In step g), y can be regarded as a realization of a random variable with standard deviation s_y.

NOTE 4 The standard deviations formed in steps g) and h) tend to reduce in a manner proportional to $h^{-1/2}$ (cf. 5.9.6 note 2).

NOTE 5 In situations where a coverage interval is not required, the test for stabilization of the computation in step k) can be based instead on $2s_y$ and $2s_{u(y)}$ only.

NOTE 6 The factor 2 used in step k) is based on regarding the averages as realizations of Gaussian variables, and corresponds to a coverage probability of approximately 95 %.

NOTE 7 An alternative, non-adaptive approach for a 95 % probabilistically symmetric coverage interval, which can be obtained using the statistics of the binomial distribution [10], is as follows. Select $M = 10^5$ or $M = 10^6$. Form the interval $[y_{(r)}, y_{(s)}]$, where, for $M = 10^5$, $r = 2\,420$ and $s = 97\,581$, or, for $M = 10^6$, $r = 24\,747$ and $s = 975\,254$. This interval is a 95 % statistical coverage interval at the level of confidence 0.99 [GUM:1995 C.2.30] [55], i.e. the coverage probability will be no less than 95 % in at least 99 % of uses of MCM. The average coverage probability of such an interval will be $(s - r)/(M + 1)$, which is greater than 95 % by an amount that becomes smaller as M is increased, viz. 95.16 % for $M = 10^5$ and 95.05 % for $M = 10^6$. (There are other possibilities for r and s; they do not have to sum to $M + 1$. A sufficient condition [10, section 2.6] is that $s - r$ satisfies

$$\sum_{j=s-r}^{M} {}^{M}C_j \, p^j (1-p)^{M-j} < 1 - 0.99,$$

where

$${}^{M}C_j = \frac{M!}{j!(M-j)!},$$

the best result being when this inequality is just satisfied.) These results can be extended to other coverage probabilities (and other choices of M).

8 Validation of results

8.1 Validation of the GUM uncertainty framework using a Monte Carlo method

8.1.1 The GUM uncertainty framework can be expected to work well in many circumstances. However, it is not always straightforward to determine whether all the conditions for its application (see 5.7 and 5.8) hold. Indeed, the degree of difficulty of doing so would typically be considerably greater than that required to apply MCM, assuming suitable software were available [8]. Therefore, since these circumstances cannot readily be tested, any cases of doubt should be validated. Since the domain of validity for MCM is broader than that for the GUM uncertainty framework, it is recommended that both the GUM uncertainty framework and MCM be applied and the results compared. Should the comparison be favourable, the GUM uncertainty framework could be used on this occasion and for sufficiently similar problems in the future. Otherwise, consideration should be given to using MCM or another appropriate approach instead.

8.1.2 Specifically, it is recommended that the two steps below and the following comparison process be carried out:

a) apply the GUM uncertainty framework (possibly with the law of propagation of uncertainty based on a higher-order Taylor series approximation) (see 5.6) to yield a $100p$ % coverage interval $y \pm U_p$ for the output quantity, where p is the stipulated coverage probability;

b) apply the adaptive Monte Carlo procedure (see 7.9.4) to provide (approximations to) the standard uncertainty $u(y)$ and the endpoints y_{low} and y_{high} of the required (probabilistically symmetric or shortest) $100p$ % coverage interval for the output quantity. Also see 8.2.

8.1.3 A comparison procedure has the following objective: determine whether the coverage intervals obtained by the GUM uncertainty framework and MCM agree to within a stipulated numerical tolerance. This numerical tolerance is assessed in terms of the endpoints of the coverage intervals and corresponds to that given by expressing the standard

uncertainty $u(y)$ to what is regarded as a meaningful number of significant decimal digits (cf. 7.9.2). The procedure is as follows:

a) form a numerical tolerance δ associated with $u(y)$ as described in 7.9.2;

b) compare the coverage intervals obtained by the GUM uncertainty framework and MCM to determine whether the required number of correct decimal digits in the coverage interval provided by the GUM uncertainty framework has been obtained. Specifically, determine

$$d_{\text{low}} = |y - U_p - y_{\text{low}}|, \tag{19}$$
$$d_{\text{high}} = |y + U_p - y_{\text{high}}|, \tag{20}$$

viz. the absolute differences of the respective endpoints of the two coverage intervals. Then, if both d_{low} and d_{high} are no larger than δ, the comparison is favourable and the GUM uncertainty framework has been validated in this instance.

NOTE The choice of $100p$ % coverage interval will influence the comparison. Therefore, the validation applies for the specified coverage probability p only.

8.2 Obtaining results from a Monte Carlo method for validation purposes

A sufficient number M of Monte Carlo trials (see 7.2) should be performed in obtaining MCM results for the validation purposes of 8.1. Let n_{dig} denote the number of significant decimal digits required in $u(y)$ (see 7.9.1) when validating the GUM uncertainty framework using MCM. Let δ denote the numerical tolerance associated with $u(y)$ (see 7.9.2). Then it is recommended that the adaptive Monte Carlo procedure (see 7.9.4) be used to provide MCM results to a numerical tolerance of $\delta/5$. Such results can be obtained by replacing δ by $\delta/5$ in step k) of that procedure.

NOTE It can be expected that the use of a numerical tolerance of $\delta/5$ would require a value of M of the order of 25 times that for a numerical tolerance of δ. Such a value of M might present efficiency problems for some computers in operating with vector arrays of dimension M. In such a case, calculations can instead be based on an approximation to $g_Y(\eta)$ derived from a suitable histogram of the y_r, in which the cell frequencies in the histogram are updated as the Monte Carlo calculation proceeds. Cf. 7.8.3 note 1.

9 Examples

9.1 Illustrations of aspects of this Supplement

9.1.1 The examples given illustrate various aspects of this Supplement. They show the application of the GUM uncertainty framework with and without contributions derived from higher-order terms in the Taylor series approximation of the model function. They also show the corresponding results provided by

a) MCM using pre-assigned numbers M of Monte Carlo trials,

b) the adaptive Monte Carlo procedure (see 7.9.4) in which M is determined automatically, or

c) both.

9.1.2 Some of the examples further show whether the MCM results provided in b) in 9.1.1 validate those provided by the GUM uncertainty framework. A numerical tolerance δ (see 7.9.2) associated with $u(y)$, with δ chosen appropriately, is used in comparing MCM and the GUM uncertainty framework. The Monte Carlo results provided in b) were obtained using a numerical tolerance of $\delta/5$ (see 8.2). In some instances, solutions are obtained analytically for further comparison.

9.1.3 Results are generally reported in the manner described in 5.5. However, more than the recommended one or two significant decimal digits are often given to facilitate comparison of the results obtained from the various approaches.

9.1.4 The Mersenne Twister generator [34] was used to generate pseudo-random numbers from a rectangular distribution (see C.3). It passes a comprehensive test for pseudo-random numbers drawn from a rectangular distribution [30] (see C.3.2) and is available within MATLAB [36], the programming environment used to produce the results given here.

9.1.5 The first example (see 9.2) constitutes an additive model. It demonstrates that the results from MCM agree with those from the application of the GUM uncertainty framework when the conditions hold for the latter (as in 5.7). The same model, but with different PDFs assigned to the input quantities, is also considered to demonstrate some departures when not all the conditions hold.

9.1.6 The second example (see 9.3) is a calibration problem from mass metrology. It demonstrates that the GUM uncertainty framework is valid in this instance only if the contributions derived from higher-order terms in the Taylor series approximation of the model function are included.

9.1.7 The third example (see 9.4) is concerned with electrical measurement. It shows that the PDF for the output quantity can be markedly asymmetric, and thus the GUM uncertainty framework can yield invalid results, even if all higher-order terms are taken into account. Instances where the input quantities are independent and not independent are treated.

9.1.8 The fourth example (see 9.5) is that in the GUM concerned with gauge block calibration [GUM:1995 H.1]. The information given there concerning the model input quantities is interpreted, PDFs accordingly assigned to these quantities, and results from the GUM uncertainty framework and MCM obtained and compared. Moreover, this treatment is applied both to the original model and the approximation made to it in the GUM.

9.2 Additive model

9.2.1 Formulation

This example considers the additive model

$$Y = X_1 + X_2 + X_3 + X_4, \tag{21}$$

a special case of the generic linear model considered in the GUM, for three different sets of PDFs $g_{X_i}(\xi_i)$ assigned to the input quantities X_i, regarded as independent. The X_i and hence the output quantity Y have dimension 1. For the first set, each $g_{X_i}(\xi_i)$ is a standard Gaussian PDF (with X_i having expectation zero and standard deviation unity). For the second set, each $g_{X_i}(\xi_i)$ is a rectangular PDF, also with X_i having expectation zero and standard deviation unity. The third set is identical to the second except that the PDF for $g_{X_4}(\xi_4)$ has a standard deviation of ten.

NOTE Further information concerning additive models, such as the model (21), where the PDFs are Gaussian or rectangular or a combination of both, is available [13].

9.2.2 Normally distributed input quantities

9.2.2.1 Assign a standard Gaussian PDF to each X_i. The best estimates of the X_i are $x_i = 0$, $i = 1, 2, 3, 4$, with associated standard uncertainties $u(x_i) = 1$.

9.2.2.2 The results obtained are summarized in the first five columns of table 2, with the results reported to three significant figures in order to facilitate their comparison (see 9.1.3).

NOTE The probabilistically symmetric 95 % coverage interval is determined, because the PDF for Y is known to be symmetric in this case, as it is for the other cases considered in this example.

9.2.2.3 The law of propagation of uncertainty [GUM:1995 5.1.2] gives the estimate $y = 0.0$ of Y and associated standard uncertainty $u(y) = 2.0$, using a numerical tolerance of two significant decimal digits for $u(y)$ ($\delta = 0.05$) (see 5.5). A probabilistically symmetric 95 % coverage interval for Y, based on a coverage factor of 1.96, is $[-3.9,\ 3.9]$.

9.2.2.4 The application of MCM (clause 7) with $M = 10^5$ trials gives $y = 0.0$, $u(y) = 2.0$ and the probabilistically symmetric 95 % coverage interval $[-3.9,\ 3.9]$. Two further applications of the method, with $M = 10^6$ trials, agree with these results to within the numerical tolerance used. These two further applications (different random samplings being made from the PDFs) were made to demonstrate the variation in the results obtained. The fourth and fifth numerical values of M (1.23×10^6 and 1.02×10^6) are the numbers of trials for two applications of the adaptive Monte Carlo procedure (see 7.9) with the use of a numerical tolerance of $\delta/5$ (see 8.2).

9.2.2.5 The PDF for Y obtained analytically is the Gaussian PDF with expectation zero and standard deviation two.

9.2.2.6 Figure 6 shows the (Gaussian) PDF for Y resulting from the GUM uncertainty framework. It also shows one of the approximations (scaled frequency distribution (histogram) of $M = 10^6$ model values of Y) constituting the discrete representation $\boldsymbol{G}$ (see 7.5) to this PDF provided by MCM. The endpoints of the probabilistically symmetric 95 % coverage interval provided by both methods are shown as vertical lines. The PDF and the approximation are visually indistinguishable, as are the respective coverage intervals. For this example, such agreement would be expected (for a sufficiently large value of M), because all the conditions hold for the application of the GUM uncertainty framework (see 5.7).

Table 2 — The application to the model (21), with each X_i assigned a standard Gaussian PDF, of (a) the GUM uncertainty framework (GUF), (b) MCM, and (c) an analytical approach (9.2.2.2, 9.2.2.7, 9.2.3.4)

Method	M	y	$u(y)$	Probabilistically symmetric 95 % coverage interval	d_{low}	d_{high}	GUF validated ($\delta = 0.05$)?
GUF		0.00	2.00	[–3.92, 3.92]			
MCM	10^5	0.00	2.00	[–3.94, 3.92]			
MCM	10^6	0.00	2.00	[–3.92, 3.92]			
MCM	10^6	0.00	2.00	[–3.92, 3.92]			
Adaptive MCM	1.23×10^6	0.00	2.00	[–3.92, 3.93]	0.00	0.01	Yes
Adaptive MCM	1.02×10^6	0.00	2.00	[–3.92, 3.92]	0.00	0.00	Yes
Analytical		0.00	2.00	[–3.92, 3.92]			

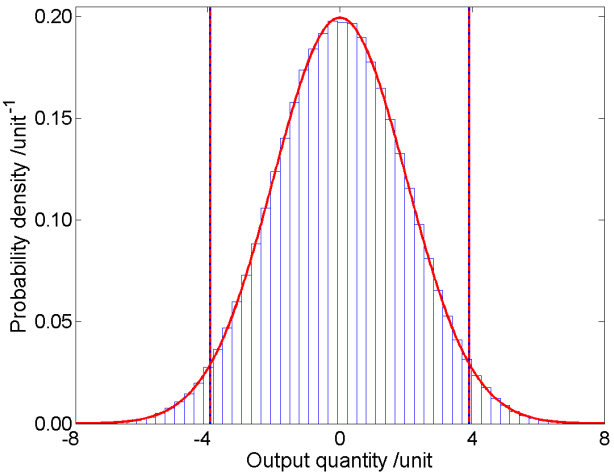

Figure 6 — Approximations for the model (21), with each X_i assigned a standard Gaussian PDF, to the PDF for Y provided by (a) the GUM uncertainty framework and (b) MCM (9.2.2.6, 9.2.3.3). "Unit" denotes any unit

9.2.2.7 Columns 6 to 8 of table 2 also shows the results of applying the validation procedures of 8.1 and 8.2. Using the terminology of 7.9.2, $n_{\text{dig}} = 2$, since two significant decimal digits in $u(y)$ are sought. Hence, $u(y) = 2.0 = 20 \times 10^{-1}$, and so $c = 20$ and $\ell = -1$. Thus, according to 7.9.2, the numerical tolerance is

$$\delta = \frac{1}{2} \times 10^{-1} = 0.05.$$

The magnitudes d_{low} and d_{high} of the endpoint differences (expressions (19) and (20)) are shown in table 2 for the two applications of the adaptive Monte Carlo procedure. Also shown is whether the GUM uncertainty framework has been validated for $\delta = 0.05$.

9.2.2.8 Figure 7 shows the length $y_{\text{high}} - y_{\text{low}}$ of the 95 % coverage interval for Y (see 7.7), as a function of the probability at its left-hand endpoint, determined from $\boldsymbol{G}$. As expected for a symmetric PDF, the interval takes its shortest length when symmetrically located with respect to the expectation.

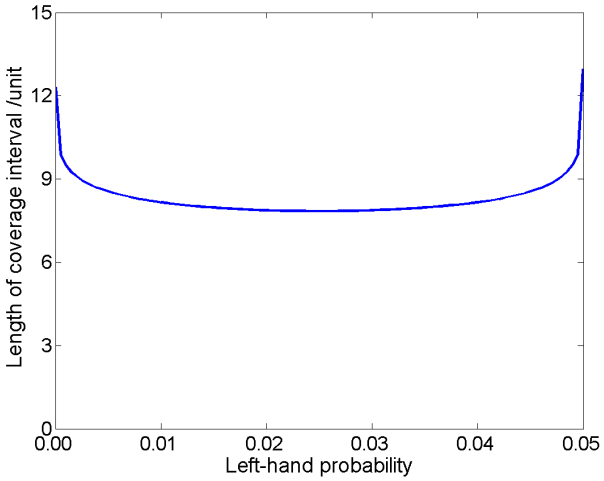

Figure 7 — The length of the 95 % coverage interval, as a function of the probability at its left-hand endpoint, for the discrete representation G of the distribution function obtained by applying MCM to the model (21) (9.2.2.8, 9.4.2.2.11)

9.2.2.9 Subclause 9.4 provides an example of an asymmetric PDF for which the shortest coverage interval differs appreciably from the probabilistically symmetric coverage interval.

9.2.3 Rectangularly distributed input quantities with the same width

9.2.3.1 Assign a rectangular PDF to each X_i, so that X_i has an expectation of zero and a standard deviation of unity (in contrast to 9.2.2.1 where a Gaussian PDF is assigned). Again, the best estimates of the X_i are $x_i = 0$, $i = 1, 2, 3, 4$, with associated standard uncertainties $u(x_i) = 1$.

9.2.3.2 By following the analogous steps of 9.2.2.3 to 9.2.2.5, the results in table 3 were obtained. The analytic solution for the endpoints of the probabilistically symmetric 95 % coverage interval, viz. $\pm 2\sqrt{3}[2 - (3/5)^{1/4}] \approx \pm 3.88$, was obtained as described in annex E.

Table 3 — As table 2, but for rectangular PDFs, with the X_i having the same expectations and standard deviations (9.2.3.2, 9.2.3.3, 9.2.3.4)

Method	M	y	$u(y)$	Probabilistically symmetric 95 % coverage interval	d_{low}	d_{high}	GUF validated ($\delta = 0.05$)?
GUF		0.00	2.00	[–3.92, 3.92]			
MCM	10^5	0.00	2.01	[–3.90, 3.89]			
MCM	10^6	0.00	2.00	[–3.89, 3.88]			
MCM	10^6	0.00	2.00	[–3.88, 3.88]			
Adaptive MCM	1.02×10^6	0.00	2.00	[–3.88, 3.89]	0.04	0.03	Yes
Adaptive MCM	0.86×10^6	0.00	2.00	[–3.87, 3.87]	0.05	0.05	No
Analytical		0.00	2.00	[–3.88, 3.88]			

9.2.3.3 Figure 8 shows the counterpart of figure 6 in this case. By comparison with figure 6, some modest differences between the approximations to the PDFs can be seen. The GUM uncertainty framework provides exactly the same PDF for Y when the PDFs for the X_i are Gaussian or rectangular, because the expectations of these quantities are identical,

as are the standard deviations, in the two cases. The PDF provided by MCM takes smaller values than those provided by the GUM uncertainty framework in the neighbourhood of the expectation and to a smaller extent towards the tails. It takes slightly greater values in the flanks. The endpoints of the coverage intervals provided are again almost visually indistinguishable, but table 3 shows small differences.

9.2.3.4 The probabilistically symmetric 95 % coverage interval determined on the basis of the GUM uncertainty framework is in this case slightly more conservative than that obtained analytically. As for normally distributed quantities, the validation procedure was applied (columns 6 to 8 of table 3). As before, $n_{\mathrm{dig}} = 2$, $u(y) = 20 \times 10^{-1}$, $c = 20$, $\ell = -1$ and $\delta = 0.05$. The endpoint differences d_{low} and d_{high} are larger than for the case of normally distributed quantities (table 2). For the first of the two applications of the adaptive Monte Carlo procedure, the GUM uncertainty framework is validated. For the second application, it is not validated, although d_{low} and d_{high} for this application are close to the numerical tolerance $\delta = 0.05$ (seen if more decimal digits than in table 3 are considered). Different validation results such as these are an occasional consequence of the stochastic nature of the Monte Carlo method, especially in a case such as that here.

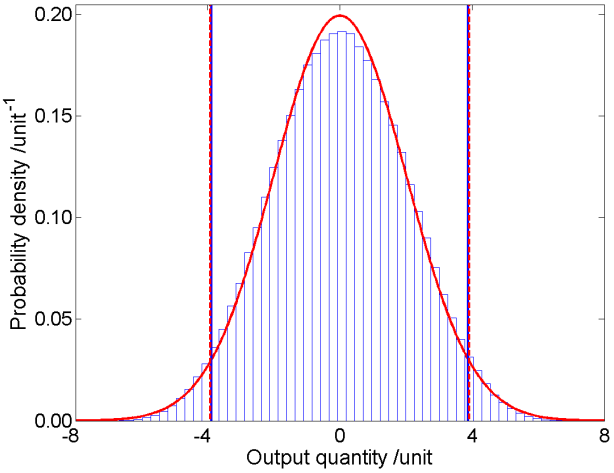

Figure 8 — The counterpart of figure 6 for quantities having the same expectations and standard deviations, but rectangular PDFs (9.2.3.3)

9.2.4 Rectangularly distributed input quantities with different widths

9.2.4.1 Consider the example of 9.2.3, except that X_4 has a standard deviation of ten rather than unity. Table 4 contains the results obtained.

9.2.4.2 The numbers M of Monte Carlo trials taken by the adaptive procedure (0.03×10^6 and 0.08×10^6) are much smaller than they were for the two previous cases in this example. The main reason is that, in this case, $\delta = 0.5$, the numerical tolerance resulting from requesting, as before, two significant decimal digits in $u(y)$, is ten times the previous value. Were the previous value to be used, M would be of the order of 100 times greater.

Table 4 — As table 3, except that the fourth input quantity has a standard deviation of ten rather than unity, and no analytic solution is provided (9.2.4.1, 9.2.4.5)

Method	M	y	$u(y)$	Probabilistically symmetric 95 % coverage interval	d_{low}	d_{high}	GUF validated ($\delta = 0.5$)?
GUF		0.0	10.1	[–19.9, 19.9]			
MCM	10^5	0.0	10.2	[–17.0, 17.0]			
MCM	10^6	0.0	10.2	[–17.0, 17.0]			
MCM	10^6	0.0	10.1	[–17.0, 17.0]			
Adaptive MCM	0.03×10^6	0.1	10.2	[–17.1, 17.1]	2.8	2.8	No
Adaptive MCM	0.08×10^6	0.0	10.1	[–17.0, 17.0]	2.9	2.9	No

9.2.4.3 Figure 9 shows the two approximations obtained to the PDF for Y. They differ appreciably. The dominance of the PDF for X_4 is evident. The PDF for Y resembles that for X_4, but there is an effect in the flanks resulting from the PDFs for the other X_i.

9.2.4.4 Figure 9 also shows the endpoints of the probabilistically symmetric 95 % coverage interval for Y obtained from these approximations. The inner pair of vertical lines indicates the endpoints of the probabilistically symmetric 95 % coverage interval determined by MCM. The outer pair results from the GUM uncertainty framework, with a coverage factor of $k = 1.96$.

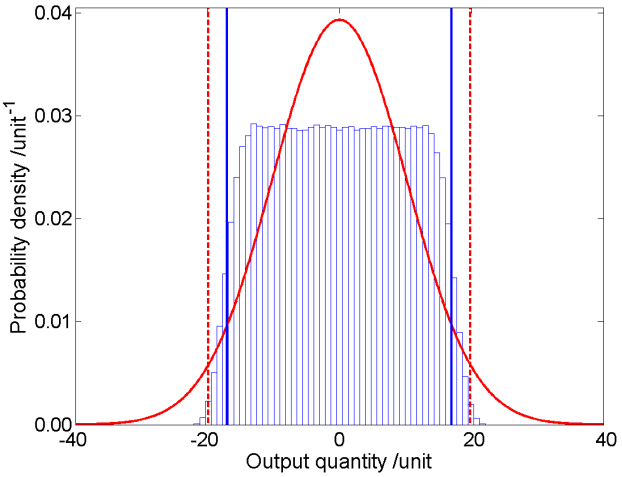

Figure 9 — As figure 8, except that the fourth input quantity has a standard deviation of ten rather than unity (9.2.4.3, 9.2.4.4)

9.2.4.5 The probabilistically symmetric 95 % coverage interval determined on the basis of the GUM uncertainty framework in this case is more conservative than that obtained using MCM. Again, the validation procedure was applied (columns 6 to 8 of table 4). Now, $n_{\mathrm{dig}} = 2$, $u(y) = 1.0 \times 10^1 = 10 \times 10^0$, $c = 10$, $\ell = 0$ and $\delta = 1/2 \times 10^0 = 0.5$. For the two applications of the adaptive Monte Carlo procedure, the GUM uncertainty framework is not validated. For a numerical tolerance of one significant decimal digit in $u(y)$, i.e. $n_{\mathrm{dig}} = 1$, for which $\delta = 5$, the validation status would be positive in both cases, the 95 % coverage intervals all being $[-2 \times 10^1,\ 2 \times 10^1]$. See 4.13.

NOTE The conditions for the central limit theorem to apply are not well met in this circumstance [GUM:1995 G.6.5], because of the dominating effect of the rectangular PDF for X_4 (see 5.7.2). However, because these conditions are often in practice assumed to hold, especially when using proprietary software for uncertainty evaluation (cf. 9.4.2.5 note 3), the characterization of Y by a Gaussian PDF on the assumption of the applicability of this theorem is made in this subclause for comparison purposes.

9.3 Mass calibration

9.3.1 Formulation

9.3.1.1 Consider the calibration of a weight W of mass density ρ_{W} against a reference weight R of mass density ρ_{R} having nominally the same mass, using a balance operating in air of mass density ρ_{a} [39]. Since ρ_{W} and ρ_{R} are generally different, it is necessary to account for buoyancy effects. Applying Archimedes' principle, the model takes the form

$$m_{\mathrm{W}}(1 - \rho_{\mathrm{a}}/\rho_{\mathrm{W}}) = (m_{\mathrm{R}} + \delta m_{\mathrm{R}})(1 - \rho_{\mathrm{a}}/\rho_{\mathrm{R}}), \tag{22}$$

where δm_{R} is the mass of a small weight of density ρ_{R} added to R to balance it with W.

9.3.1.2 It is usual to work in terms of conventional masses. The conventional mass $m_{\mathrm{W,c}}$ of W is the mass of a (hypothetical) weight of density $\rho_0 = 8\ 000\ \mathrm{kg/m^3}$ that balances W in air at density $\rho_{\mathrm{a}_0} = 1.2\ \mathrm{kg/m^3}$. Thus,

$$m_{\mathrm{W}}(1 - \rho_{\mathrm{a}_0}/\rho_{\mathrm{W}}) = m_{\mathrm{W,c}}(1 - \rho_{\mathrm{a}_0}/\rho_0).$$

9.3.1.3 In terms of conventional masses $m_{\mathrm{W,c}}$, $m_{\mathrm{R,c}}$ and $\delta m_{\mathrm{R,c}}$, the model (22) becomes

$$m_{\mathrm{W,c}}(1 - \rho_{\mathrm{a}}/\rho_{\mathrm{W}})(1 - \rho_{\mathrm{a}_0}/\rho_{\mathrm{W}})^{-1} = (m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}})(1 - \rho_{\mathrm{a}}/\rho_{\mathrm{R}})(1 - \rho_{\mathrm{a}_0}/\rho_{\mathrm{R}})^{-1}, \tag{23}$$

from which, to an approximation adequate for most practical purposes,

$$m_{\mathrm{W,c}} = (m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}}) \left[1 + (\rho_{\mathrm{a}} - \rho_{\mathrm{a}_0}) \left(\frac{1}{\rho_{\mathrm{W}}} - \frac{1}{\rho_{\mathrm{R}}}\right)\right].$$

Let

$$\delta m = m_{\mathrm{W,c}} - m_{\mathrm{nom}}$$

be the deviation of $m_{\mathrm{W,c}}$ from the nominal mass

$$m_{\mathrm{nom}} = 100\ \mathrm{g}.$$

The model used in this example is given by

$$\delta m = (m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}}) \left[1 + (\rho_{\mathrm{a}} - \rho_{\mathrm{a}_0}) \left(\frac{1}{\rho_{\mathrm{W}}} - \frac{1}{\rho_{\mathrm{R}}}\right)\right] - m_{\mathrm{nom}}. \tag{24}$$

NOTE Applying the law of propagation of uncertainty to the "exact" model (23) is made difficult by the algebraic complexity of the partial derivatives. It is easier to apply MCM, because only model values need be formed.

9.3.1.4 The only information available concerning $m_{\mathrm{R,c}}$ and $\delta m_{\mathrm{R,c}}$ is a best estimate and an associated standard uncertainty for each of these quantities. Accordingly, following 6.4.7.1, a Gaussian distribution is assigned to each of these quantities, with these best estimates used as the expectations of the corresponding quantities and the associated standard uncertainties as the standard deviations. The only information available concerning ρ_{a}, ρ_{W} and ρ_{R} is lower and upper limits for each of these quantities. Accordingly, following 6.4.2.1, a rectangular distribution is assigned to each of these quantities, with limits equal to the endpoints of the distribution. Table 5 summarizes the input quantities and the PDFs assigned. In the table, a Gaussian distribution $\mathrm{N}(\mu, \sigma^2)$ is described in terms of expectation μ and standard deviation σ, and a rectangular distribution $\mathrm{R}(a, b)$ with endpoints a and b ($a < b$) in terms of expectation $(a + b)/2$ and semi-width $(b - a)/2$.

NOTE The quantity ρ_{a_0} in the mass calibration model (24) is assigned the value $1.2\ \mathrm{kg/m^3}$ with no associated uncertainty.

Table 5 — The input quantities X_i and the PDFs assigned to them for the mass calibration model (24) (9.3.1.4)

X_i	Distribution	Parameters			
		Expectation μ	Standard deviation σ	Expectation $x = (a + b)/2$	Semi-width $(b - a)/2$
$m_{\mathrm{R,c}}$	$\mathrm{N}(\mu, \sigma^2)$	100 000.000 mg	0.050 mg		
$\delta m_{\mathrm{R,c}}$	$\mathrm{N}(\mu, \sigma^2)$	1.234 mg	0.020 mg		
ρ_{a}	$\mathrm{R}(a, b)$			$1.20\ \mathrm{kg/m^3}$	$0.10\ \mathrm{kg/m^3}$
ρ_{W}	$\mathrm{R}(a, b)$			$8 \times 10^3\ \mathrm{kg/m^3}$	$1 \times 10^3\ \mathrm{kg/m^3}$
ρ_{R}	$\mathrm{R}(a, b)$			$8.00 \times 10^3\ \mathrm{kg/m^3}$	$0.05 \times 10^3\ \mathrm{kg/m^3}$

9.3.2 Propagation and summarizing

9.3.2.1 The GUM uncertainty framework and the adaptive Monte Carlo procedure (see 7.9) were each used to obtain an estimate $\widehat{\delta m}$ of δm, the associated standard uncertainty $u(\widehat{\delta m})$, and the shortest 95 % coverage interval

for δm. The results obtained are shown in table 6, in which GUF_1 denotes the GUM uncertainty framework with first-order terms, MCM the adaptive Monte Carlo procedure, and GUF_2 the GUM uncertainty framework with higher-order terms.

9.3.2.2 0.72×10^6 trials were taken by the adaptive Monte Carlo procedure with the use of a numerical tolerance of $\delta/5$ (see 8.2) with δ set for the case where one significant decimal digit in $u(\widehat{\delta m})$ is regarded as meaningful (see 9.3.2.6).

9.3.2.3 Figure 10 shows the approximations to the PDF for δm obtained from the GUM uncertainty framework with first-order terms and MCM. The continuous curve represents a Gaussian PDF with parameters given by the GUM uncertainty framework. The inner pair of (broken) vertical lines indicates the shortest 95 % coverage interval for δm based on this PDF. The histogram is the scaled frequency distribution obtained using MCM as an approximation to the PDF. The outer pair of (continuous) vertical lines indicates the shortest 95 % coverage interval for δm based on the discrete representation of the distribution function determined as in 7.5.

Table 6 — Results of the calculation stage for the mass calibration model (24) (9.3.2.1, 9.3.2.6)

Method	$\widehat{\delta m}$ /mg	$u(\widehat{\delta m})$ /mg	**Shortest 95 % coverage interval** /mg	d_{low} /mg	d_{high} /mg	**GUF validated** ($\delta = 0.005$)**?**
GUF_1	1.234 0	0.053 9	[1.128 5, 1.339 5]	0.045 1	0.043 0	No
MCM	1.234 1	0.075 4	[1.083 4, 1.382 5]			
GUF_2	1.234 0	0.075 0	[1.087 0, 1.381 0]	0.003 6	0.001 5	Yes

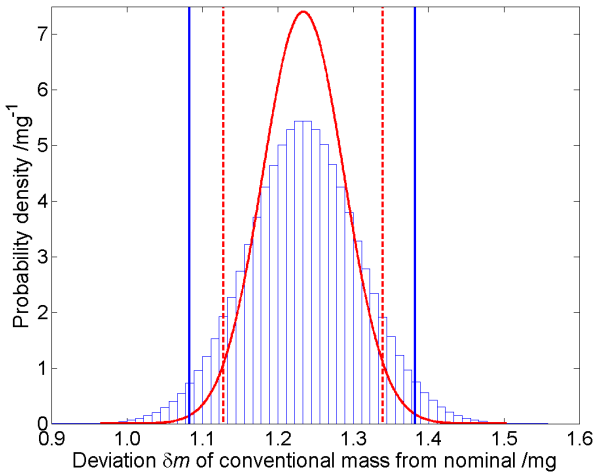

Figure 10 — Approximations to the PDF for the output quantity δm obtained using the GUM uncertainty framework with first-order terms and MCM (9.3.2.3)

9.3.2.4 The results show that, although the GUM uncertainty framework (first order) and MCM give estimates of δm in good agreement, the numerical values for the associated standard uncertainty are noticeably different. The value (0.075 4 mg) of $u(\widehat{\delta m})$ returned by MCM is 40 % larger than that (0.053 9 mg) returned by the GUM uncertainty framework (first order). The latter is thus optimistic in this respect. There is good agreement between $u(\widehat{\delta m})$ determined by MCM and that (0.075 0 mg) provided by the GUM uncertainty framework with higher-order terms.

9.3.2.5 Table 7 contains the partial derivatives of first order for the model (24) with respect to the input quantities together with the sensitivity coefficients, viz. these derivatives evaluated at the best estimates of the input quantities. These derivatives indicate that, for the purposes of the GUM uncertainty framework with first-order terms, the model for this example can be considered as being replaced by the additive model

$$\delta m = m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}} - m_{\mathrm{nom}}.$$

MCM makes no such (implied) approximation to the model.

Table 7 — Sensitivity coefficients for the mass calibration model (24) (9.3.2.5)

X_i	Partial derivative	Sensitivity coefficient
$m_{\mathrm{R,c}}$	$1 + (\rho_{\mathrm{a}} - \rho_{\mathrm{a_0}})(1/\rho_{\mathrm{W}} - 1/\rho_{\mathrm{R}})$	1
$\delta m_{\mathrm{R,c}}$	$1 + (\rho_{\mathrm{a}} - \rho_{\mathrm{a_0}})(1/\rho_{\mathrm{W}} - 1/\rho_{\mathrm{R}})$	1
ρ_{a}	$(m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}})(1/\rho_{\mathrm{W}} - 1/\rho_{\mathrm{R}})$	0
ρ_{W}	$-(m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}})(\rho_{\mathrm{a}} - \rho_{\mathrm{a_0}})/\rho_{\mathrm{W}}^2$	0
ρ_{R}	$(m_{\mathrm{R,c}} + \delta m_{\mathrm{R,c}})(\rho_{\mathrm{a}} - \rho_{\mathrm{a_0}})/\rho_{\mathrm{R}}^2$	0

9.3.2.6 Table 6 also shows in the right-most three columns the results of applying the validation procedure of 8.1 and 8.2 in the case where one significant decimal digit in $u(\widehat{\delta m})$ is regarded as meaningful. Using the terminology of that subclause, $n_{\mathrm{dig}} = 1$, since a numerical tolerance of one significant decimal digit in $u(\widehat{\delta m})$ is required. Hence, $u(\widehat{\delta m}) = 0.08 = 8 \times 10^{-2}$, and so the c in 7.9.2 equals 8 and $\ell = -2$. Thus $\delta = 1/2 \times 10^{-2} = 0.005$. d_{low} and d_{high} denote the magnitudes of the endpoint differences (19) and (20), where y there corresponds to $\widehat{\delta m}$. Whether the results were validated to one significant decimal digit in $u(\widehat{\delta m})$ is indicated in the final column of the table. If only first-order terms are accounted for, the application of the GUM uncertainty framework is not validated. If higher-order terms are accounted for [GUM:1995 5.1.2 note], the GUM uncertainty framework is validated. Thus, the non-linearity of the model is such that accounting for first-order terms only is inadequate.

9.4 Comparison loss in microwave power meter calibration

9.4.1 Formulation

9.4.1.1 During the calibration of a microwave power meter, the power meter and a standard power meter are connected in turn to a stable signal generator. The power absorbed by each meter will in general be different because their complex input voltage reflection coefficients are not identical. The ratio Y of the power P_{M} absorbed by the meter being calibrated and that, P_{S}, by the standard meter is [43]

$$Y = \frac{P_{\mathrm{M}}}{P_{\mathrm{S}}} = \frac{1 - |\Gamma_{\mathrm{M}}|^2}{1 - |\Gamma_{\mathrm{S}}|^2} \times \frac{|1 - \Gamma_{\mathrm{S}}\Gamma_{\mathrm{G}}|^2}{|1 - \Gamma_{\mathrm{M}}\Gamma_{\mathrm{G}}|^2}, \tag{25}$$

where Γ_{G} is the voltage reflection coefficient of the signal generator, Γ_{M} that of the meter being calibrated and Γ_{S} that of the standard meter. This power ratio is an instance of "comparison loss" [1, 28].

9.4.1.2 Consider the case where the standard and the signal generator are reflectionless, i.e. $\Gamma_{\mathrm{S}} = \Gamma_{\mathrm{G}} = 0$, and measured values are obtained of the real and imaginary parts X_1 and X_2 of $\Gamma_{\mathrm{M}} = X_1 + \mathrm{j}X_2$, where $\mathrm{j}^2 = -1$. Since $|\Gamma_{\mathrm{M}}|^2 = X_1^2 + X_2^2$, formula (25) becomes

$$Y = 1 - X_1^2 - X_2^2. \tag{26}$$

9.4.1.3 Given respectively are best estimates x_1 and x_2 of the quantities X_1 and X_2 from measurement and the associated standard uncertainties $u(x_1)$ and $u(x_2)$. X_1 and X_2 are often not independent. Denote by $u(x_1, x_2)$ the covariance associated with x_1 and x_2. Equivalently [GUM:1995 5.2.2], $u(x_1, x_2) = r(x_1, x_2)u(x_1)u(x_2)$, where $r = r(x_1, x_2)$ denotes the associated correlation coefficient [GUM:1995 5.2.2].

NOTE In practice the electrical engineer may sometimes have difficulty in quantifying the covariance. In such cases, the uncertainty evaluation can be repeated with different trial numerical values for the correlation coefficient in order to study its effect. This example carries out calculations using a correlation coefficient of zero and of 0.9 (cf. 9.4.1.7).

9.4.1.4 On the basis of 6.4.8.1, $\boldsymbol{X} = (X_1, X_2)^{\top}$ is assigned a bivariate Gaussian PDF in X_1 and X_2, with expectation and covariance matrix

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \qquad \begin{bmatrix} u^2(x_1) & ru(x_1)u(x_2) \\ ru(x_1)u(x_2) & u^2(x_2) \end{bmatrix}. \tag{27}$$

9.4.1.5 Because the magnitudes of X_1 and X_2 in expression (26) are in practice small compared with unity, the resulting Y is close to unity. Results are accordingly expressed in terms of the quantity

$$\delta Y = 1 - Y = X_1^2 + X_2^2, \tag{28}$$

taken as the model of measurement. For physical reasons, $0 \leq Y \leq 1$, and hence $0 \leq \delta Y \leq 1$.

9.4.1.6 The determination of an estimate δy of δY, the associated standard uncertainty $u(\delta y)$, and a coverage interval for δY will be considered for choices of x_1, x_2, $u(x_1)$, $u(x_2)$ and $r(x_1, x_2)$. All quantities have dimension 1.

9.4.1.7 Six cases are considered, in all of which x_2 is taken as zero and $u(x_1) = u(x_2) = 0.005$. The first three of these cases correspond to taking $x_1 = 0$, 0.010, and 0.050, each with $r(x_1, x_2) = 0$. The other three cases correspond to taking the same x_1, but with $r(x_1, x_2) = 0.9$. The various numerical values of x_1 (comparable to those occurring in practice) are used to investigate the extent to which the results obtained using the considered approaches differ.

9.4.1.8 For the cases in which $r = r(x_1, x_2) = 0$, the covariance matrix given in formulæ (27) reduces to $\mathrm{diag}(u^2(x_1), u^2(x_2))$ and the corresponding joint distribution for X_1 and X_2 to the product of two univariate Gaussian distributions for X_i, for $i = 1, 2$, with expectation x_i and standard deviation $u(x_i)$.

9.4.2 Propagation and summarizing: zero covariance

9.4.2.1 General

9.4.2.1.1 The evaluation of uncertainty is treated by applying the propagation of distributions

a) analytically (for purposes of comparison),

b) using the GUM uncertainty framework, and

c) using MCM.

NOTE These approaches do not constrain the PDF for δY to be no greater than unity. However, for sufficiently small uncertainties $u(x_1)$ and $u(x_2)$, as here, the PDF for δY may adequately be approximated by a simpler PDF defined over all non-negative values of δY. A rigorous treatment, using Bayesian inference [51], which applies regardless of the magnitudes of $u(x_1)$ and $u(x_2)$, is possible, but beyond the scope of this Supplement. Also see clause 1 note 2.

9.4.2.1.2 δy and $u(\delta y)$ can generally be formed analytically as the expectation and standard deviation of δY, as characterized by the PDF for δY. See F.1. The PDF for δY can be formed analytically when $x_1 = 0$ and, in particular, used to determine the endpoints of the shortest 95 % coverage interval in that case. See F.2.

9.4.2.1.3 The GUM uncertainty framework with first-order terms and with higher-order terms is applied for each of the three estimates x_1 in the uncorrelated case. See F.3. An estimate δy of δY is formed in each case [GUM:1995 4.1.4] from

$$\delta y = x_1^2 + x_2^2.$$

9.4.2.1.4 MCM is applied in each case with $M = 10^6$ trials.

9.4.2.2 Input estimate $x_1 = 0$

9.4.2.2.1 For the input estimate $x_1 = 0$, higher-order terms must be used when applying the law of propagation of uncertainty, because the partial derivatives of δY with respect to X_1 and X_2, evaluated at $X_1 = x_1$ and $X_2 = x_2$, are identically zero when $x_1 = x_2 = 0$. Thus, if the law of propagation of uncertainty with first-order terms only were applied, the resulting standard uncertainty would incorrectly be computed as zero.

NOTE A similar difficulty would arise for x_1 close to zero.

9.4.2.2.2 Figure 11 shows the PDFs for δY determined by applying the propagation of distributions

a) analytically (the exponentially decreasing curve for $\delta Y \geq 0$ and zero elsewhere),

b) using the GUM uncertainty framework with higher-order terms in order to characterize the output quantity by a Gaussian PDF (bell-shaped curve), and

c) using MCM (scaled frequency distribution).

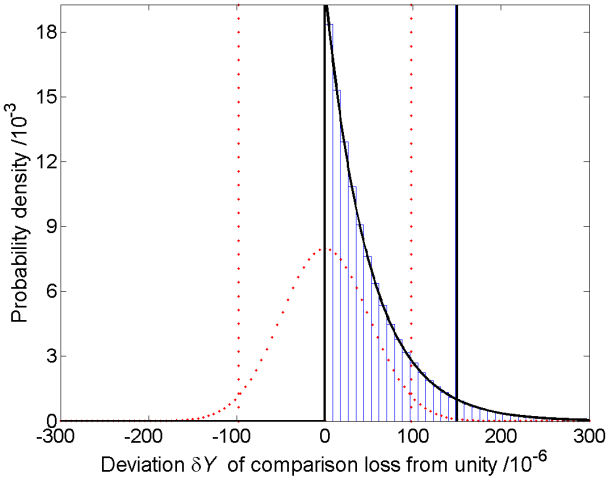

Figure 11 — Results for the model of comparison loss in power meter calibration in the case $x_1 = x_2 = 0$, with $u(x_1) = u(x_2) = 0.005$ and $r(x_1, x_2) = 0$ (9.4.2.2, 9.4.2.6, 9.4.2.9 and 9.4.2.11)

9.4.2.2.3 It is seen in the figure that the use of the GUM uncertainty framework with higher-order terms in order to characterize the output quantity by a Gaussian distribution yields a PDF that is very different from the analytic solution. The latter takes the form of a particular chi-squared distribution—the sum of squares of two standard Gaussian variables (see F.2).

9.4.2.2.4 Since the partial derivatives of the model function (28) of order higher than two are all identically zero, the solution obtained essentially corresponds to taking all Taylor-series terms, i.e. the full non-linearity of the problem, into account. Thus, the particular Gaussian distribution so determined is the best that is possible using the GUM uncertainty framework to characterize the output quantity by such a distribution.

9.4.2.2.5 It can therefore be concluded that the reason for the departure from the analytic solution of the results from the approach based on the GUM uncertainty framework is that the output quantity is characterized by a Gaussian PDF. No Gaussian PDF, however it is obtained, could adequately represent the analytic solution in this case.

9.4.2.2.6 It is also seen in figure 11 that the PDF provided by MCM is consistent with the analytic solution.

9.4.2.2.7 The estimates δy determined as the expectation of δY described by the PDFs obtained

a) analytically,

b) using the GUM uncertainty framework, and

c) applying MCM

are given in columns 2 to 4 of the row corresponding to $x_1 = 0.000$ in table 8. Columns 5 to 8 contain the corresponding $u(\delta y)$, with those obtained using the GUM uncertainty framework with first-order terms (G_1) and higher-order terms (G_2).

Table 8 — Comparison loss results, for input estimates with associated zero covariance, obtained analytically (A), and using the GUM uncertainty framework with first-order terms (G_1) and higher-order terms (G_2) and MCM (M) (9.4.2.2.7, 9.4.2.2.10, 9.4.2.3.4, 9.4.2.4.2)

x_1	Estimate δy $/10^{-6}$			Standard uncertainty $u(\delta y)$ $/10^{-6}$				Shortest 95 % coverage interval for δY $/10^{-6}$			
	A	**G**	**M**	**A**	**G_1**	**G_2**	**M**	**A**	**G_1**	**G_2**	**M**
0.000	50	0	50	50	0	50	50	[0, 150]	[0, 0]	[–98, 98]	[0, 150]
0.010	150	100	150	112	100	112	112	—	[–96, 296]	[–119, 319]	[0, 367]
0.050	2 550	2 500	2 551	502	500	502	502	—	[1 520, 3 480]	[1 515, 3 485]	[1 590, 3 543]

9.4.2.2.8 The estimate $\delta y = 0$ obtained by evaluating the model at the input estimates is invalid: the correct (analytic) PDF for δY is identically zero for $\delta Y < 0$; this estimate lies on the boundary of the non-zero part of that function. The estimate provided by MCM agrees with that obtained analytically. The law of propagation of uncertainty based on first-order terms gives the wrong, zero, value for $u(\delta y)$ already noted. The value (50×10^{-6}) from the law of propagation of uncertainty based on higher-order terms agrees with that obtained analytically and by MCM.

NOTE When MCM was repeated several times the results obtained were scattered about 50×10^{-6}. When it was repeated a number of times with a larger numerical value of M the results were again scattered about 50×10^{-6}, but with a reduced dispersion. Such dispersion effects are expected, and were observed for the other Monte Carlo calculations made. Reporting the results to greater numbers of significant decimal digits would be necessary to see the actual numerical differences.

9.4.2.2.9 Figure 11 also shows the shortest 95 % coverage intervals for the corresponding approximations to the distribution function for δY. The 95 % coverage interval, indicated by dotted vertical lines, as provided by the GUM uncertainty framework is infeasible: it is symmetric about $\delta Y = 0$ and therefore erroneously implies there is a 50 % probability that δY is negative. The continuous vertical lines are the endpoints of the shortest 95 % coverage interval derived from the analytic solution, as described in F.2. The endpoints of the shortest 95 % coverage interval determined using MCM are indistinguishable to graphical accuracy from those for the analytic solution.

9.4.2.2.10 The endpoints of the shortest coverage intervals relating to the standard uncertainties in columns 5 to 8 of the row corresponding to $x_1 = 0.000$ in table 8 are given in columns 9 to 12 of that table.

9.4.2.2.11 Figure 12 shows the length of the 95 % coverage interval (see 7.7), as a function of the probability value at its left-hand endpoint, for the approximation to the PDF provided by MCM shown in figure 11. The 95 % coverage interval does not take its shortest length when symmetrically located with respect to the expectation in this case. Indeed, the shortest 95 % coverage interval is as far-removed as possible from a probabilistically symmetric coverage interval, the left and right tail probabilities being 0 % and 5 %, respectively, as opposed to 2.5 % and 2.5 %. This figure can be compared with that (figure 7) for the additive model of 9.2, for which the PDF for Y is symmetric about its expectation.

9.4.2.3 Input estimate $x_1 = 0.010$

9.4.2.3.1 For the input estimate $x_1 = 0.010$, with correlation coefficient $r(x_1, x_2) = 0$, figure 13 shows the PDFs obtained using the GUM uncertainty framework with first-order terms only and with higher-order terms, and using MCM.

9.4.2.3.2 The PDF provided by MCM exhibits a modest left-hand flank, although it is truncated at zero, the smallest possible numerical value of δY. Further, compared with the results for $x_1 = 0$, it is closer in form to the Gaussian PDFs provided by the GUM uncertainty framework. These Gaussian PDFs are in turn reasonably close to each other, δY having expectation 1.0×10^{-4} and standard deviations 1.0×10^{-4} and 1.1×10^{-4}, respectively.

9.4.2.3.3 Figure 13 also shows the endpoints of the shortest 95 % coverage intervals obtained by the three approaches. The continuous vertical lines denote the endpoints of the interval provided by MCM, the broken vertical lines those resulting from the GUM uncertainty framework with first-order terms, and the dotted vertical lines from the GUM uncertainty framework with higher-order terms. The intervals provided by the GUM uncertainty framework are shifted to the left compared with the shortest 95 % coverage interval for MCM. As a consequence, they again

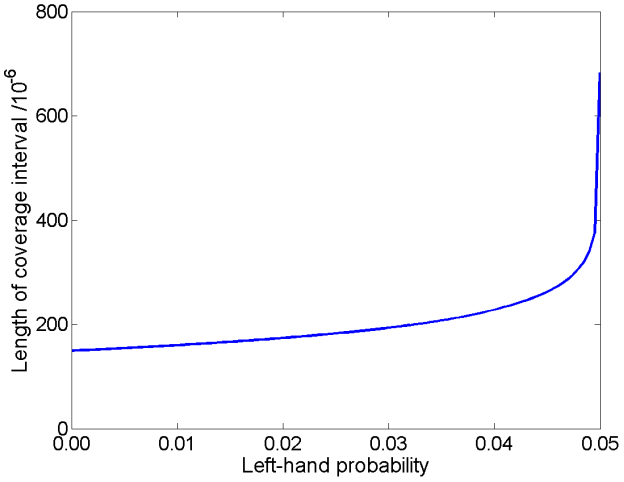

Figure 12 — The length of the 95 % coverage interval, as a function of the probability value at its left-hand endpoint, for the approximation to the distribution function obtained by applying MCM to the model (28) (9.4.2.2.11)

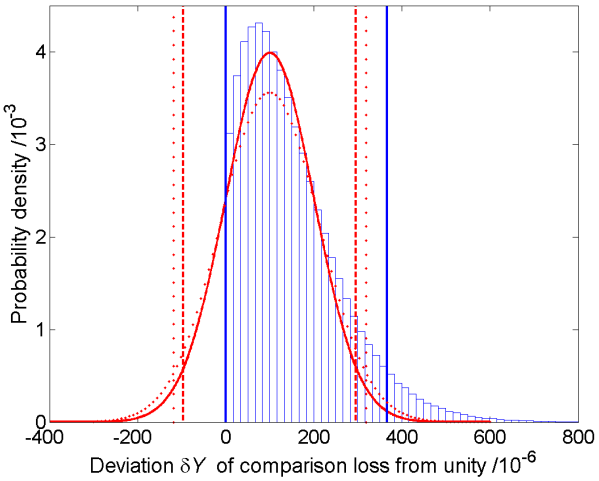

Figure 13 — As figure 11 except that $x_1 = 0.010$, and the PDFs resulting from the GUM uncertainty framework with first-order (higher-peaked curve) and with higher-order terms (lower-peaked curve) (9.4.2.3.1, 9.4.2.3.3, 9.4.2.4.1, 9.4.3.3)

include infeasible values of δY. The shift is about 70 % of the standard uncertainty. The interval provided by MCM has its left-hand endpoint at zero, the smallest feasible value.

9.4.2.3.4 The corresponding results are given in the penultimate row of table 8.

9.4.2.4 Input estimate $x_1 = 0.050$

9.4.2.4.1 Figure 14 is similar to figure 13, but for $x_1 = 0.050$. Now, the PDFs provided by both variants of the GUM uncertainty framework are virtually indistinguishable from each other. Further, they are now much closer to the approximation to the PDF provided by MCM. That PDF exhibits a slight skewness, as evidenced in the tail regions. The coverage intervals provided by the two variants of the GUM uncertainty framework are visually almost identical, but still shifted from those for MCM. The shift is now about 10 % of the standard uncertainty. The intervals provided by the GUM uncertainty framework are now feasible.

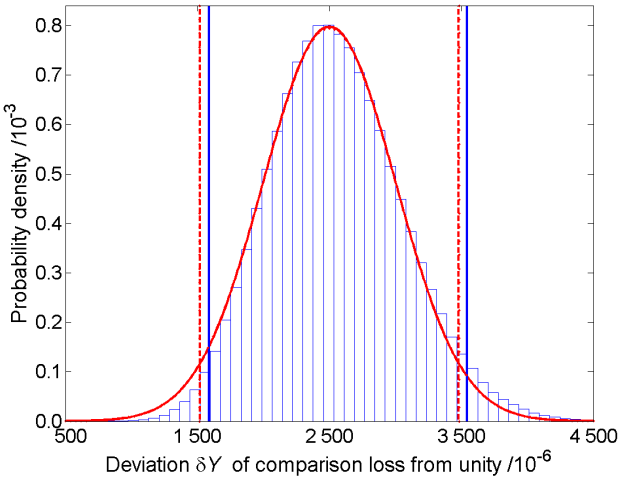

Figure 14 — As figure 13 except that $x_1 = 0.050$ (9.4.2.4.1, 9.4.3.3)

9.4.2.4.2 The corresponding results are given in the final row of table 8.

9.4.2.5 Discussion

As x_1 becomes increasingly removed from zero, the results given by the GUM uncertainty framework, with first-order and with higher-order terms, and those for MCM become closer to each other.

NOTE 1 The numerical values $x_1 = x_2 = 0$ lie in the centre of the region of interest to the electrical engineer, corresponding to the so-called "matched" condition for the power meter being calibrated, and thus in no sense constitute an extreme case.

NOTE 2 Because of the symmetry of the model in X_1 and X_2, exactly the same effect would occur were x_2 used in place of x_1.

NOTE 3 One reason why the GUM uncertainty framework with first-order terms (only) might be used in practice is that software for its implementation is readily available: results obtained from it might sometimes be accepted without question. For the case where $x_1 = x_2 = 0$ (figure 11), the danger would be apparent because the standard uncertainty $u(\delta y)$ was computed as zero, and consequently any coverage interval for δY would be of zero length for any coverage probability. For $x_1 \neq 0$ (or $x_2 \neq 0$), $u(\delta y)$ and the length of the coverage interval for δY are both non-zero, so no such warning would be available without prior knowledge of likely values for $u(\delta y)$ and this length. Thus, a danger in implementing software based on the GUM uncertainty framework for these calculations is that checks of the software for x_1 or x_2 sufficiently far from zero would not indicate such problems, although, when used subsequently in practice for small values of x_1 or x_2, the results would be invalid, but conceivably unwittingly accepted.

9.4.3 Propagation and summarizing: non-zero covariance

9.4.3.1 General

9.4.3.1.1 The three approaches used in the cases where the X_i are uncorrelated (see 9.4.2) are now applied for the three cases in which they are correlated, with $r(x_1, x_2) = 0.9$. However, the GUM uncertainty framework with first-order terms only is used. Unlike the cases where the X_i are uncorrelated, the GUM uncertainty framework with higher-order terms is not applied, no counterpart being provided in the GUM for the formula containing higher-order terms when the x_i have associated non-zero covariances (see 5.8). Other aspects match those in 9.4.2.

9.4.3.1.2 For the GUM uncertainty framework with first-order terms, $u(\delta y)$ is evaluated as described in F.3.2. Expression (F.7) in that subclause gives, for $x_2 = 0$,

$$u^2(\delta y) = 4x_1^2 u^2(x_1).$$

Consequently, $u(\delta y)$ does not depend on $r(x_1, x_2)$ and the GUM uncertainty framework with first-order terms gives identical results to those presented in 9.4.2. In particular, for the case $x_1 = 0$, $u(\delta y)$ is (incorrectly) computed as zero, as in 9.4.2.2.1.

9.4.3.1.3 MCM was implemented by sampling randomly from $\boldsymbol{X}$ characterized by a bivariate Gaussian PDF with the given expectation and covariance matrix (expressions (27)). The procedure in C.5 was used.

NOTE Apart from the requirement to draw from a multivariate distribution, the implementation of MCM for input quantities that are correlated is no more complicated than when the input quantities are uncorrelated.

9.4.3.2 Input estimates $x_1 = 0$, 0.010, and 0.050

9.4.3.2.1 Table 9 contains the results obtained. Those from MCM indicate that although δy is unaffected by the correlation between the X_i, $u(\delta y)$ is so influenced, more so for small x_1. The 95 % coverage intervals are influenced accordingly.

Table 9 — Comparison loss results, for input estimates with associated non-zero covariance ($r(x_1, x_2) = 0.9$), obtained analytically, and using the GUM uncertainty framework (GUF) and MCM (9.4.3.2.1)

x_1	Estimate δy $/10^{-6}$			Standard uncertainty $u(\delta y)$ $/10^{-6}$			Shortest 95 % coverage interval for δY $/10^{-6}$		
	Analytical	**GUF**	**MCM**	**Analytical**	**GUF**	**MCM**	**Analytical**	**GUF**	**MCM**
0.000	50	0	50	67	0	67	—	[0, 0]	[0, 185]
0.010	150	100	150	121	100	121	—	[−96, 296]	[13, 398]
0.050	2 550	2 500	2 551	505	500	504	—	[1 520, 3 480]	[1 628, 3 555]

9.4.3.2.2 Figures 15 and 16 show the PDFs provided by the GUM uncertainty framework with first-order terms (bell-shaped curves) and MCM (scaled frequency distributions) in the cases $x_1 = 0.010$ and $x_1 = 0.050$, respectively. The endpoints of the shortest 95 % coverage interval provided by the two approaches are also shown, as broken vertical lines for the GUM uncertainty framework and continuous vertical lines for MCM.

NOTE Strictly, the conditions under which δY can be characterized by a Gaussian PDF do not hold following an application of the GUM uncertainty framework in this circumstance (see 5.8) [GUM:1995 G.6.6]. However, this PDF and the endpoints of the corresponding 95 % coverage interval are shown because such a characterization is commonly used.

9.4.3.3 Discussion

In the case $x_1 = 0.010$ (figure 15), the effect of the correlation has been to change noticeably the results returned by MCM (compare with figure 13). Not only has the shape of (the approximation to) the PDF changed, but the corresponding coverage interval no longer has its left-hand endpoint at zero. In the case $x_1 = 0.050$ (figure 16), the differences between the results for the cases where the input quantities are uncorrelated and correlated (compare with figure 14) are less obvious.

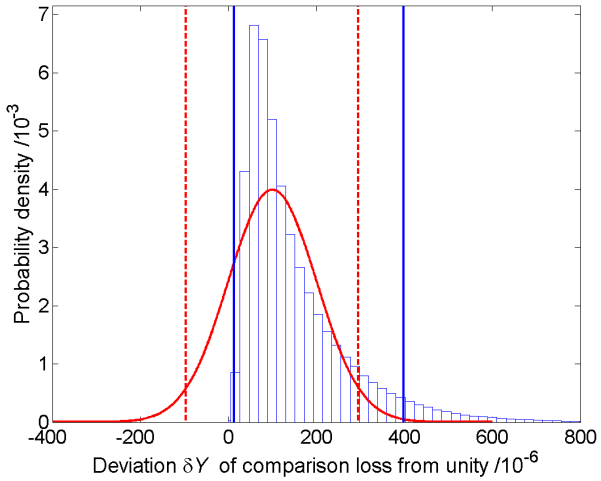

Figure 15 — Results for the model of comparison loss in power meter calibration in the case $x_1 = 0.010$, $x_2 = 0$, with $u(x_1) = u(x_2) = 0.005$ and $r(x_1, x_2) = 0.9$ (9.4.3.2.2, 9.4.3.3)

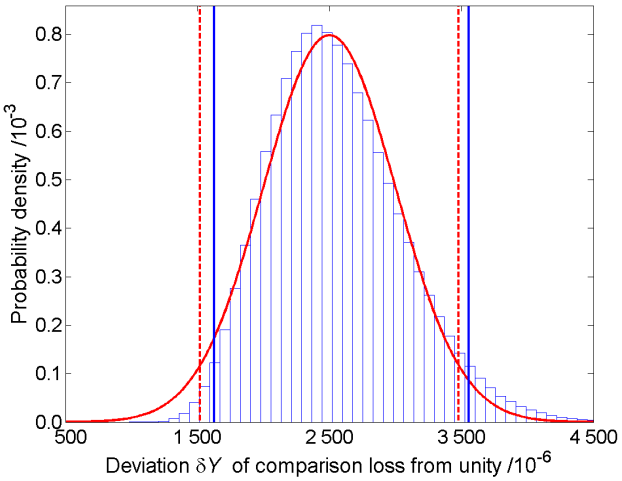

Figure 16 — As figure 15 except that $x_1 = 0.050$ (9.4.3.2.2, 9.4.3.3)

9.5 Gauge block calibration

9.5.1 Formulation: model

9.5.1.1 The length of a nominally 50 mm gauge block is determined by comparing it with a known reference standard of the same nominal length. The direct output of the comparison of the two gauge blocks is the difference d in their lengths given by

$$d = L(1 + \alpha\theta) - L_{\mathrm{s}}(1 + \alpha_{\mathrm{s}}\theta_{\mathrm{s}}), \tag{29}$$

where L is the length at 20 °C of the gauge block being calibrated, L_{s} is the length of the reference standard at 20 °C as given in its calibration certificate, α and α_{s} are the coefficients of thermal expansion, respectively, of the gauge being calibrated and the reference standard, and θ and θ_{s} are the deviations in temperature from the 20 °C reference temperature, respectively, of the gauge block being calibrated and the reference standard.

NOTE 1 The GUM refers to a gauge block as an end gauge.

NOTE 2 The symbol L for the length of a gauge block is used in this Supplement in place of the symbol ℓ used in the GUM for that quantity.

9.5.1.2 From expression (29), the output quantity L is given by

$$L = \frac{L_{\mathrm{s}}(1 + \alpha_{\mathrm{s}}\theta_{\mathrm{s}}) + d}{1 + \alpha\theta}, \tag{30}$$

from which, to an approximation adequate for most practical purposes,

$$L = L_{\mathrm{s}} + d + L_{\mathrm{s}}(\alpha_{\mathrm{s}}\theta_{\mathrm{s}} - \alpha\theta). \tag{31}$$

If the difference in temperature between the gauge block being calibrated and the reference standard is written as $\delta\theta = \theta - \theta_{\mathrm{s}}$, and the difference in their thermal expansion coefficients as $\delta\alpha = \alpha - \alpha_{\mathrm{s}}$, models (30) and (31) become, respectively,

$$L = \frac{L_{\mathrm{s}}[1 + \alpha_{\mathrm{s}}(\theta - \delta\theta)] + d}{1 + (\alpha_{\mathrm{s}} + \delta\alpha)\theta} \tag{32}$$

and

$$L = L_{\mathrm{s}} + d - L_{\mathrm{s}}(\theta\delta\alpha + \alpha_{\mathrm{s}}\delta\theta). \tag{33}$$

9.5.1.3 The difference d in the lengths of the gauge block being calibrated and the reference standard is determined as the average of a series of five indications, obtained independently, of the difference using a calibrated comparator. d can be expressed as

$$d = D + d_1 + d_2, \tag{34}$$

where D is a quantity of which the average of the five indications is a realization, and d_1 and d_2 are quantities describing, respectively, the random and systematic effects associated with using the comparator.

9.5.1.4 The quantity θ, representing deviation of the temperature from 20 °C of the gauge block being calibrated, can be expressed as

$$\theta = \theta_0 + \Delta, \tag{35}$$

where θ_0 is a quantity representing the average temperature deviation of the gauge block from 20 °C and Δ a quantity describing a cyclic variation of the temperature deviation from θ_0.

9.5.1.5 Substituting expressions (34) and (35) into expressions (32) and (33), and working with the quantity δL representing the deviation of L from the nominal length

$$L_{\text{nom}} = 50 \text{ mm}$$

of the gauge block, gives

$$\delta L = \frac{L_{\text{s}}[1 + \alpha_{\text{s}}(\theta_0 + \Delta - \delta\theta)] + D + d_1 + d_2}{1 + (\alpha_{\text{s}} + \delta\alpha)(\theta_0 + \Delta)} - L_{\text{nom}} \tag{36}$$

and

$$\delta L = L_{\text{s}} + D + d_1 + d_2 - L_{\text{s}}[\delta\alpha(\theta_0 + \Delta) + \alpha_{\text{s}}\delta\theta] - L_{\text{nom}} \tag{37}$$

as models for the measurement problem.

9.5.1.6 The treatment here of the measurement problem is in terms of the models (36) and (37) with output quantity δL and input quantities L_{s}, D, d_1, d_2, α_{s}, θ_0, Δ, $\delta\alpha$ and $\delta\theta$. It differs from that given in GUM example H.1 in that in the GUM the models (34) and (35) above are treated as sub-models to models (32) and (33), i.e. the GUM uncertainty framework is applied to each model (34) and (35), with the results obtained used to provide information about the input quantities d and θ in models (32) and (33). The treatment here avoids having to use the results obtained from MCM applied to the sub-models (34) and (35) to provide information about the distributions for the input quantities d and θ in expressions (32) and (33).

9.5.2 Formulation: assignment of PDFs

9.5.2.1 General

In the following subclauses the available information about each input quantity in the models (36) and (37) is provided. This information is extracted from the description given in the GUM, and for each item of information the GUM subclause from which the item is extracted is identified. Also provided is an interpretation of the information in terms of an assignment of a distribution to the quantity. Table 10 summarizes the assignments made.

Table 10 — PDFs assigned to input quantities for the gauge block models (36) and (37) on the basis of available information (9.5.2.1). Table 1 provides general information concerning these PDFs

Quantity	PDF	Parameters					
		μ	σ	ν	a	b	d
L_{s}	$t_\nu(\mu, \sigma^2)$	50 000 623 nm	25 nm	18			
D	$t_\nu(\mu, \sigma^2)$	215 nm	6 nm	24			
d_1	$t_\nu(\mu, \sigma^2)$	0 nm	4 nm	5			
d_2	$t_\nu(\mu, \sigma^2)$	0 nm	7 nm	8			
α_{s}	$\text{R}(a, b)$				9.5×10^{-6} °C^{-1}	13.5×10^{-6} °C^{-1}	
θ_0	$\text{N}(\mu, \sigma^2)$	−0.1 °C	0.2 °C				
Δ	$\text{U}(a, b)$				−0.5 °C	0.5 °C	
$\delta\alpha$	$\text{CTrap}(a, b, d)$				-1.0×10^{-6} °C^{-1}	1.0×10^{-6} °C^{-1}	0.1×10^{-6} °C^{-1}
$\delta\theta$	$\text{CTrap}(a, b, d)$				−0.050 °C	0.050 °C	0.025 °C

9.5.2.2 Length L_{s} of the reference standard

9.5.2.2.1 Information

The calibration certificate for the reference standard gives $\widehat{L_{\text{s}}} = 50.000\,623$ mm as its length at 20 °C [GUM:1995 H.1.5]. It gives $U_p = 0.075$ µm as the expanded uncertainty of the reference standard and states that it was obtained using a coverage factor of $k_p = 3$ [GUM:1995 H.1.3.1]. The certificate states that the effective degrees of freedom associated with the combined standard uncertainty, from which the quoted expanded uncertainty was obtained, is $\nu_{\text{eff}}(u(\widehat{L_{\text{s}}})) = 18$ [GUM:1995 H.1.6].

9.5.2.2.2 Interpretation

Assign a scaled and shifted t-distribution $t_\nu(\mu, \sigma^2)$ (see 6.4.9.7) to L_s, with

$$\mu = 50\ 000\ 623 \text{ nm}, \qquad \sigma = \frac{U_p}{k_p} = \frac{75}{3} \text{ nm} = 25 \text{ nm}, \qquad \nu = 18.$$

9.5.2.3 Average length difference D

9.5.2.3.1 Information

The average $\widehat{D}$ of the five indications of the difference in lengths between the gauge block being calibrated and the reference standard is 215 nm [GUM:1995 H.1.5]. The pooled experimental standard deviation characterizing the comparison of L and L_s was determined from 25 indications, obtained independently, of the difference in lengths of two standard gauge blocks, and equalled 13 nm [GUM:1995 H.1.3.2].

9.5.2.3.2 Interpretation

Assign a scaled and shifted t-distribution $t_\nu(\mu, \sigma^2)$ (see 6.4.9.2 and 6.4.9.6) to D, with

$$\mu = 215 \text{ nm}, \qquad \sigma = \frac{13}{\sqrt{5}} \text{ nm} = 6 \text{ nm}, \qquad \nu = 24.$$

9.5.2.4 Random effect d_1 of comparator

9.5.2.4.1 Information

According to the calibration certificate of the comparator used to compare L with L_s, the associated uncertainty due to random effects is 0.01 μm for a coverage probability of 95 % and is obtained from six indications, obtained independently [GUM:1995 H.1.3.2].

9.5.2.4.2 Interpretation

Assign a scaled and shifted t-distribution $t_\nu(\mu, \sigma^2)$ (see 6.4.9.7) to d_1, with

$$\mu = 0 \text{ nm}, \qquad \sigma = \frac{U_{0.95}}{k_{0.95}} = \frac{10}{2.57} \text{ nm} = 4 \text{ nm}, \qquad \nu = 5.$$

Here, $k_{0.95}$ is obtained from table G.2 of the GUM with $\nu = 5$ degrees of freedom and $p = 0.95$.

9.5.2.5 Systematic effect d_2 of comparator

9.5.2.5.1 Information

The uncertainty of the comparator due to systematic effects is given in the certificate as 0.02 μm at the "three sigma level" [GUM:1995 H.1.3.2]. This uncertainty may be assumed to be reliable to 25 %, and thus the degrees of freedom is $\nu_{\mathrm{eff}}(u(\widehat{d_2})) = 8$ [GUM:1995 H.1.6].

9.5.2.5.2 Interpretation

Assign a scaled and shifted t-distribution $t_\nu(\mu, \sigma^2)$ (see 6.4.9.7) to d_2, with

$$\mu = 0 \text{ nm}, \qquad \sigma = \frac{U_p}{k_p} = \frac{20}{3} \text{ nm} = 7 \text{ nm}, \qquad \nu = 8.$$

9.5.2.6 Thermal expansion coefficient α_s

9.5.2.6.1 Information

The coefficient of thermal expansion of the reference standard is given as $\widehat{\alpha_\mathrm{s}} = 11.5 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}$ with possible values of this quantity represented by a rectangular distribution with limits $\pm 2 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}$ [GUM:1995 H.1.3.3].

9.5.2.6.2 Interpretation

Assign a rectangular distribution $\mathrm{R}(a, b)$ (see 6.4.2) to α_s, with limits

$$a = 9.5 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}, \qquad b = 13.5 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}.$$

NOTE There is no information about the reliability of the limits and so a rectangular distribution with exactly known limits is assigned. Such information may have been omitted from the description in the GUM because the corresponding sensitivity coefficient is zero, and so the quantity makes no contribution in an application of the GUM uncertainty framework based on first-order terms only.

9.5.2.7 Average temperature deviation θ_0

9.5.2.7.1 Information

The temperature of the test bed is reported as $(19.9 \pm 0.5)\ ^\circ\mathrm{C}$. The average temperature deviation $\widehat{\theta_0} = -0.1\ ^\circ\mathrm{C}$ is reported as having an associated standard uncertainty due to the uncertainty associated with the average temperature of the test bed of $u(\widehat{\theta_0}) = 0.2\ ^\circ\mathrm{C}$ [GUM:1995 H.1.3.4].

9.5.2.7.2 Interpretation

Assign a Gaussian distribution $\mathrm{N}(\mu, \sigma^2)$ (see 6.4.7) to θ_0, with

$$\mu = -0.1\ ^\circ\mathrm{C}, \qquad \sigma = 0.2\ ^\circ\mathrm{C}.$$

NOTE There is no information about the source of the evaluation of the uncertainty and so a Gaussian distribution is assigned. Also see 9.5.2.6.2 note, regarding such information.

9.5.2.8 Effect Δ of cyclic temperature variation

9.5.2.8.1 Information

The temperature of the test bed is reported as $(19.9 \pm 0.5)\ ^\circ\mathrm{C}$. The stated maximum offset of 0.5 °C for Δ is said to represent the amplitude of an approximately cyclical variation of temperature under a thermostatic system. The cyclic variation of temperature results in a U-shaped (arc sine) distribution [GUM:1995 H.1.3.4].

9.5.2.8.2 Interpretation

Assign an arc sine distribution $\mathrm{U}(a, b)$ (see 6.4.6) to Δ, with limits

$$a = -0.5\ ^\circ\mathrm{C}, \qquad b = 0.5\ ^\circ\mathrm{C}.$$

NOTE There is no information about the reliability of the limits and so a U-shaped distribution with exactly known limits is assigned. As in 9.5.2.6.2 note, such information may have been omitted from the description in the GUM.

9.5.2.9 Difference $\delta\alpha$ in expansion coefficients

9.5.2.9.1 Information

The estimated bounds on the variability of $\delta\alpha$ are $\pm 1 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}$, with an equal probability of $\delta\alpha$ having any value within those bounds [GUM:1995 H.1.3.5]. These bounds are deemed to be reliable to 10 %, giving $\nu(u(\widehat{\delta\alpha})) = 50$ [GUM:1995 H.1.6].

9.5.2.9.2 Interpretation

Assign a rectangular distribution with inexactly prescribed limits (see 6.4.3) to $\delta\alpha$, with

$$a = -1.0 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}, \qquad b = 1.0 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}, \qquad d = 0.1 \times 10^{-6}\ ^\circ\mathrm{C}^{-1}.$$

The stated reliability of 10 % on the estimated bounds provides the basis for this value of d.

9.5.2.10 Difference $\delta\theta$ in temperatures

9.5.2.10.1 Information

The reference standard and the gauge block being calibrated are expected to be at the same temperature, but the temperature difference $\delta\theta$ could lie with equal probability anywhere in the estimated interval -0.05 °C to 0.05 °C [GUM:1995 H.1.3.6]. This difference is believed to be reliable only to 50 %, giving $\nu(u(\widehat{\delta\theta})) = 2$ [GUM:1995 H.1.6].

9.5.2.10.2 Interpretation

Assign a rectangular distribution with inexactly prescribed limits (see 6.4.3) to $\delta\theta$, with

$$a = -0.050\ ^\circ\mathrm{C}, \qquad b = 0.050\ ^\circ\mathrm{C}, \qquad d = 0.025\ ^\circ\mathrm{C}.$$

The stated reliability of 50 % provides the basis for this value of d.

9.5.3 Propagation and summarizing

9.5.3.1 The GUM uncertainty framework

The application of the GUM uncertainty framework is based on

— a first-order Taylor series approximation to the model (36) or (37),

— use of the Welch-Satterthwaite formula to evaluate an effective degrees of freedom (rounded towards zero) associated with the uncertainty obtained from the law of propagation of uncertainty, and

— assigning a scaled and shifted t-distribution with the above degrees of freedom to the output quantity.

9.5.3.2 Monte Carlo method

The application of MCM

— requires sampling from a rectangular distribution (see 6.4.2.4 and C.3.3), Gaussian distribution (see 6.4.7.4 and C.4), t-distribution (see 6.4.9.5 and C.6), U-shaped distribution (see 6.4.6.4), and rectangular distribution with inexactly prescribed limits (see 6.4.3.4), and

— implements adaptive MCM (see 7.9) with a numerical tolerance ($\delta = 0.5$) set to deliver $n_{\text{dig}} = 2$ significant decimal digits in the standard uncertainty.

9.5.4 Results

9.5.4.1 Table 11 gives the results obtained for the approximate model (37) using the information summarized in table 10. Figure 17 shows the PDFs for δL obtained from the application of the GUM uncertainty framework (solid curve) and MCM (scaled frequency distribution). The distribution obtained from the GUM uncertainty framework is a t-distribution with $\nu = 16$ degrees of freedom. The endpoints of the shortest 99 % coverage intervals for δL obtained from the PDFs, which are indicated as vertical lines, are visually indistinguishable.

9.5.4.2 1.26×10^6 trials were taken by the adaptive Monte Carlo procedure. The calculations were also carried out for a coverage probability of 95 %, for which 0.53×10^6 trials were taken.

Table 11 — Results obtained for the approximate model (37) using the information summarized in table 10 (9.5.4.1, 9.5.4.3)

Method	$\widehat{\delta L}$ /nm	$u(\widehat{\delta L})$ /nm	Shortest 99 % coverage interval for δL /nm
GUF	838	32	[745, 931]
MCM	838	36	[745, 932]

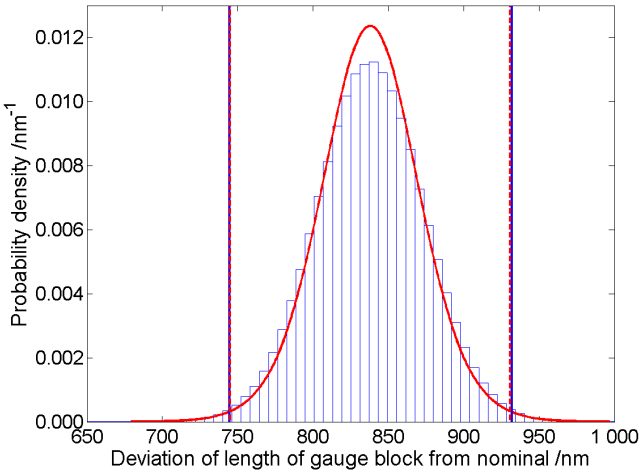

Figure 17 — PDFs for δL obtained using the GUM uncertainty framework (continuous bell-shaped curve) and MCM (scaled histogram) for the approximate model (37) using the information summarized in table 10 (9.5.4.1)

9.5.4.3 Results obtained for the non-linear model (36) are identical to the results in table 11 to the number of decimal digits given there.

9.5.4.4 There are some modest differences in the results obtained. $u(\widehat{\delta L})$ was 4 nm greater for the application of MCM than for the GUM uncertainty framework. The length of the 99 % coverage interval for δL was 1 nm greater.

These results apply equally to the non-linear and the approximate models. Whether such differences are important has to be judged in terms of the manner in which the results are to be used.

Annex A
Historical perspective

A.1 The GUM is a rich document, covering many aspects of uncertainty evaluation. Although it does not refer explicitly to the use of a Monte Carlo method, such use was recognized during the drafting of the GUM. The ISO/IEC/OIML/BIPM draft (First Edition) of June 1992, produced by ISO/TAG 4/WG 3, states [G.1.5]:

> If the relationship between Y and its input quantities is nonlinear, or if the values available for the parameters characterizing the probabilities of the X_i (expectation, variance, higher moments) are only estimates and are themselves characterized by probability distributions, and a first-order Taylor expansion of the relationship is not an acceptable approximation, the distribution of Y cannot be expressed as a convolution. In this case, a numerical approach (such as Monte Carlo calculations) will generally be required and the evaluation is computationally more difficult.

A.2 In the published version of the GUM, this subclause had been modified to read:

> If the functional relationship between Y and its input quantities is nonlinear and a first-order Taylor expansion of the relationship is not an acceptable approximation (see 5.1.2 and 5.1.5), then the probability distribution of Y cannot be obtained by convolving the distributions of the input quantities. In such cases, other analytical or numerical methods are required.

A.3 The interpretation made here of this re-wording is that "other analytical or numerical methods" cover any other appropriate approach. This interpretation is consistent with that of the National Institute of Standards and Technology of the United States [50]:

> [6.6] The NIST policy provides for exceptions as follows (see Annex C):
>
> It is understood that any valid statistical method that is technically justified under the existing circumstances may be used to determine the equivalent of u_i, u_c, or U. Further, it is recognized that international, national, or contractual agreements to which NIST is a party may occasionally require deviation from NIST policy. In both cases, the report of uncertainty must document what was done and why.

Annex B
Sensitivity coefficients and uncertainty budgets

B.1 Neither the propagation of distributions nor its implementation using MCM provides sensitivity coefficients [GUM:1995 5.1.3]. However, by holding all input quantities but one fixed at their best estimates, MCM can be used to provide the PDF for the output quantity for the model having just that input quantity as a variable [8]. The ratio of the standard deviation of the resulting model values (cf. 7.6) and the standard uncertainty associated with the best estimate of the relevant input quantity can be taken as a sensitivity coefficient. This ratio corresponds to that which would be obtained by taking all higher-order terms in the Taylor series expansion of the model into account. This approach may be viewed as a generalization of the approximate partial-derivative formula in the GUM [GUM:1995 5.1.3 note 2]. Both the sensitivity coefficients and the contributions for each input quantity to the uncertainty associated with the estimate of the output quantity will in general differ from those obtained with the GUM.

B.2 In many measurement contexts it is common practice to list the uncertainty components $u_i(y) = |c_i|u(x_i)$, $i = 1, \ldots, N$, where c_i is the ith sensitivity coefficient and $u(x_i)$ the standard uncertainty associated with the ith input estimate x_i, contributing to the standard uncertainty $u(y)$. Usually these are presented in a table, the "uncertainty budget". This practice may be useful to identify the dominant terms contributing to $u(y)$ associated with the estimate of the output quantity. However, in cases for which (a valid implementation of) the propagation of distributions is more appropriate, an uncertainty budget should be regarded as a qualitative tool.

Annex C
Sampling from probability distributions

C.1 General

C.1.1 This annex provides technical information concerning sampling from probability distributions. Such sampling forms a central part of the use of MCM as an implementation of the propagation of distributions. A digital library of mathematical functions [38] and a repository of relevant software [37] may also be consulted.

C.1.2 A generator for any distribution, such as the distributions considered in 6.4 (also see table 1), can in principle be obtained from its distribution function, together with the use of a generator for the rectangular distribution, as indicated in C.2. A generator for a rectangular distribution is provided in C.3.3. For some distributions, such as the Gaussian distribution and the t-distribution, it is more efficient to use specifically developed generators, such as those provided in this annex. Subclause 6.4 also gives advice on sampling from probability distributions.

NOTE Generators other than those given in this annex can be used. Their statistical quality should be tested before use. A testing facility is available for pseudo-random number generators for the rectangular distribution. See C.3.2.

C.2 General distributions

A draw from any strictly increasing, univariate continuous distribution function $G_X(\xi)$ can be made by transforming a draw from a rectangular distribution:

a) draw a random number ρ from the rectangular distribution R(0, 1);

b) determine ξ satisfying $G_X(\xi) = \rho$.

NOTE 1 The inversion required in step b), that is forming $\xi = G_X^{-1}(\rho)$, may be possible analytically. Otherwise it can be carried out numerically.

EXAMPLE As an instance of analytical inversion, consider the exponential PDF for X, with X having expectation x (> 0), viz. $g_X(\xi) = \exp(-\xi/x)/x$, for $\xi \geq 0$, and zero otherwise (see 6.4.10). Then, by integration, $G_X(\xi) = 1 - \exp(-\xi/x)$, for $\xi \geq 0$, and zero otherwise. Hence $\xi = -x \ln(1 - \rho)$. This result can be simplified slightly by using the fact that if a variable Q has the rectangular distribution R(0, 1), then so has $1 - Q$. Hence, $\xi = -x \ln \rho$.

NOTE 2 Numerically, ξ can generally be determined by solving the "zero-of-a-function" problem $G_X(\xi) - \rho = 0$. Upper and lower bounds for ξ are typically easily found, in which case a recognized "bracketing" algorithm such as bisection or, more efficiently, a combination of linear interpolation and bisection [11], for example, can be used to determine ξ.

NOTE 3 If the pseudo-random number generator for the rectangular distribution is to be used as a basis for generating numbers from another distribution, a draw of ρ equal to zero or one can cause failure of that generator. An example is the exponential distribution (see 6.4.10). Its PDF (expression (9)) is not defined for ρ equal to zero or one. The use of the generator given in C.3.3 would not give rise to a failure of that type.

C.3 Rectangular distribution

C.3.1 General

C.3.1.1 The ability to generate pseudo-random numbers from a rectangular distribution is fundamental in its own right, and also as the basis for generating pseudo-random numbers from any distribution (see C.2, C.4 and C.6) using an appropriate algorithm or formula. In the latter regard, the quality of the numbers generated from a non-rectangular distribution depends on that of the generator of numbers from a rectangular distribution and on the properties of the algorithm employed. The quality of the numbers generated from a non-rectangular distribution can therefore be expected to be related to those generated from the rectangular distribution. Only a generator that can faithfully provide rectangularly distributed numbers used in conjunction with a good algorithm can be expected to constitute a generator that can faithfully provide non-rectangularly distributed numbers.

C.3.1.2 It is thus important that the underlying facility for generating rectangularly distributed numbers is sound [31]. Unless the user is sure of its pedigree, a generator should not be used until adequate testing has been carried out. Misleading results can otherwise be obtained. The use of a testing facility [30] is recommended. A procedure for generating rectangularly distributed numbers, which has been shown to perform well in these tests and is straightforward to implement, is given in C.3.3.

C.3.1.3 Table C.1 defines relevant aspects of the functioning of a procedure for generating pseudo-random numbers from the rectangular distribution R(0, 1), specifying the input, input-output and output parameters associated with their determination.

NOTE 1 By setting the seeds in table C.1 to seeds previously used, the same sequence of random numbers can be produced. Doing so is important as part of software regression testing, used to verify the consistency of results produced using the software with those from previous versions.

NOTE 2 Some pseudo-random number generators provide a single draw at each call and others several draws.

Table C.1 — Generation of pseudo-random numbers from a rectangular distribution (C.3.1.3, C.3.2.2)

Input parameter	
q	Number of pseudo-random numbers to be generated
Input-output parameter	
$\boldsymbol{t}$	Column vector of parameters, some of which may be required as input quantities, that may be changed as part of the computation. Subsequent values of these parameters are not usually of immediate concern to the user. The parameters are needed to help control the process by which the pseudo-random numbers are generated. The parameters may be realized as global variables and thus not explicitly appear as parameters of the procedure. One or more of these parameters may be a seed, used to initiate the sequence of random numbers produced by successive calls of the procedure
Output parameter	
$\boldsymbol{z}$	Column vector of q draws from the rectangular distribution R(0, 1)

C.3.1.4 A pseudo-random number x drawn from R(a, b) is given by $a + (b - a)z$, where z is a pseudo-random number drawn from R(0, 1).

C.3.2 Randomness tests

C.3.2.1 Any pseudo-random number generator used should

a) have good statistical properties,

b) readily be implemented in any programming language, and

c) give the same results for the same seed on any computer.

It is also desirable that it is compact, thus rendering its implementation straightforward. One such generator that comes close to satisfying these requirements is that due to Wichmann and Hill [52, 53]. It has been used in many areas including uncertainty computations. However, its cycle length (the number of random numbers generated before the sequence is repeated) is 2^{31}, today considered inadequate for some problems. Moreover, not all tests of its statistical properties were passed [35]. Further, the generator was designed for 16-bit computers, whereas today 32-bit and 64-bit computers are almost universally used.

NOTE The period of the sequence of numbers produced by a pseudo-random number generator is the number of consecutive numbers in the sequence before they are repeated.

C.3.2.2 An extensive test of the statistical properties of any generator submitted to it is carried out by the test suite TestU01 [30]. This suite is very detailed, with many individual tests, including the so-called Big Crush. Several generators passing the Big Crush test are listed by Wichmann and Hill [54]. An enhanced Wichmann-Hill generator (see C.3.3) also passes the test, and has the properties [54] that

a) it is straightforward to code in any programming language. It does not depend upon bit manipulation used by some generators,

b) the state (the amount of information preserved by the generator between calls to it) is small and easy to handle (cf. the parameter $\boldsymbol{t}$ in table C.1),

c) it can readily be used to provide multiple sequences needed for highly parallel applications, likely to be a feature of future uncertainty calculations, and

d) there are variants of the generator for 32- and 64-bit computers.

C.3.3 Procedure for generating pseudo-random numbers from a rectangular distribution

C.3.3.1 Like the previous generator, the enhanced Wichmann-Hill generator is a combination of congruential generators. The new generator combines four such generators, whereas the previous version combined three. The new generator has a period of 2^{121}, acceptable for any conceivable application.

C.3.3.2 Table C.2 defines the enhanced Wichmann-Hill generator for producing pseudo-random numbers from R(0, 1) for a 32-bit computer.

Table C.2 — The enhanced Wichmann-Hill generator for pseudo-random numbers (C.3.3.2, C.3.3.3) from a rectangular distribution on the interval $(0,\ 1)$ **for 32-bit computers.** $\lfloor w \rfloor$ **denotes the largest integer no greater than** w**.** $i_j \bmod b_j$ **denotes the remainder on division of** i_j **by** b_j

	Input parameter
	None
	Input-output parameter
i_1, i_2, i_3, i_4	Integer parameters required as input quantities and that are changed by the procedure. Set to integers between 1 and 2 147 483 647 before the first call. Do not disturb between calls. Subsequent values of these parameters are not usually of concern to the user. The parameters provide the basis by which the pseudo-random numbers are generated. They may be realized as global variables and thus not appear explicitly as parameters of the procedure
	Constant
$\boldsymbol{a}$, $\boldsymbol{b}$, $\boldsymbol{c}$ $\boldsymbol{d}$,	Vectors of integer constants of dimension 1×4, where $\boldsymbol{a} = (a_1, \ldots, a_4)$, etc., given by: $\boldsymbol{a} = (11\ 600, 47\ 003, 23\ 000, 33\ 000)$, $\boldsymbol{b} = (185\ 127, 45\ 688, 93\ 368, 65\ 075)$, $\boldsymbol{c} = (10\ 379, 10\ 479\ , 19\ 423, 8\ 123)$, $\boldsymbol{d} = 2\ 147\ 483\ 123 \times (1,1,1,1) + (456, 420, 300, 0)$. Do not disturb between calls
	Output parameter
r	Pseudo-random number drawn from R(0, 1)
	Computation
	a) For $j = 1, \ldots, 4$: i) Form $i_j = a_j \times (i_j \bmod b_j) - c_j \times \lfloor i_j / b_j \rfloor$ ii) If $i_j < 0$, replace i_j by $i_j + d_j$ b) Form $w = \sum_{j=1}^{4} i_j / d_j$ c) Form $r = w - \lfloor w \rfloor$

C.3.3.3 For 64-bit computers, step a) of Computation, including (i) and (ii), in the generator of table C.2 is to be replaced by the simpler step:

a) For $j = 1, \ldots, 4$, form $i_j = (a_j \times i_j) \bmod d_j$

C.4 Gaussian distribution

The procedure in table C.3 provides draws from the standard Gaussian distribution N(0, 1) using the Box-Muller transform [3]. A draw from the Gaussian distribution $N(\mu, \sigma^2)$ is given by $\mu + \sigma z$, where z is a draw from N(0, 1).

Table C.3 — The Box-Muller Gaussian pseudo-random number generator (C.4)

Input parameter	
	None
Output parameter	
z_1, z_2	Two draws, obtained independently, from a standard Gaussian distribution
Computation	
	a) Generate random draws r_1 and r_2 independently from the rectangular distribution R(0, 1) b) Form $z_1 = \sqrt{-2 \ln r_1} \cos 2\pi r_2$ and $z_2 = \sqrt{-2 \ln r_1} \sin 2\pi r_2$

C.5 Multivariate Gaussian distribution

C.5.1 The most important multivariate distribution is the multivariate (or joint) Gaussian distribution $N(\boldsymbol{\mu}, \boldsymbol{V})$, where $\boldsymbol{\mu}$ is a vector of expectations of dimension $n \times 1$ and $\boldsymbol{V}$ a covariance matrix of dimension $n \times n$.

C.5.2 Draws from $N(\boldsymbol{\mu}, \boldsymbol{V})$ [45, 49] can be obtained using the procedure in table C.4.

NOTE 1 If $\boldsymbol{V}$ is positive definite (i.e. all its eigenvalues are strictly positive), the Cholesky factor $\boldsymbol{R}$ is unique [23, page 204].

NOTE 2 If $\boldsymbol{V}$ is not positive definite, perhaps because of numerical rounding errors or other sources, $\boldsymbol{R}$ may not exist. Moreover, in cases where one or more of the eigenvalues of $\boldsymbol{V}$ is very small (but positive), the software implementation of the Cholesky factorization algorithm used may be unable to form $\boldsymbol{R}$ because of the effects of floating-point errors. In either of these situations it is recommended that $\boldsymbol{V}$ is "repaired", i.e. as small a change as possible is made to $\boldsymbol{V}$ such that the Cholesky factor $\boldsymbol{R}$ for the modified matrix is well defined. The resulting factor is exact for a covariance matrix that numerically is close to the original $\boldsymbol{V}$. A simple repair procedure is available [49, page 322] for this purpose, and is embodied in the MULTNORM generator [45].

NOTE 3 If $\boldsymbol{V}$ is semi-positive definite, the eigendecomposition $\boldsymbol{V} = \boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}^\top$, where $\boldsymbol{Q}$ is an orthogonal matrix and $\boldsymbol{\Lambda}$ a diagonal matrix, can be formed. Then $\boldsymbol{\Lambda}^{1/2}\boldsymbol{Q}^\top$ can be used to obtain draws from $N(\boldsymbol{0}, \boldsymbol{V})$, even if $\boldsymbol{V}$ is rank deficient.

Table C.4 — A multivariate Gaussian random number generator (C.5.2)

Input parameter	
n	Dimension of the multivariate Gaussian distribution
$\boldsymbol{\mu}$	Vector of expectations of dimension $n \times 1$
$\boldsymbol{V}$	Covariance matrix of dimension $n \times n$
q	Number of multivariate Gaussian pseudo-random numbers to be generated
Output parameter	
$\boldsymbol{X}$	Matrix of dimension $n \times q$, the jth column of which is a draw from the multivariate Gaussian distribution
Computation	
	a) Form the Cholesky factor $\boldsymbol{R}$ of $\boldsymbol{V}$, i.e. the upper triangular matrix satisfying $\boldsymbol{V} = \boldsymbol{R}^\top \boldsymbol{R}$. (To generate q pseudo-random numbers, it is necessary to perform this matrix factorization only once.) b) Generate an array $\boldsymbol{Z}$ of standard Gaussian variates of dimension $n \times q$ c) Form $\boldsymbol{X} = \boldsymbol{\mu}\boldsymbol{1}^\top + \boldsymbol{R}^\top \boldsymbol{Z}$, where $\boldsymbol{1}$ denotes a vector of ones of dimension $q \times 1$

C.5.3 Figure C.1 shows 200 points generated using the MULTNORM generator [45] from $\mathrm{N}(\boldsymbol{\mu},\ \boldsymbol{V})$, where

$$\boldsymbol{\mu} = \begin{bmatrix} 2.0 \\ 3.0 \end{bmatrix}, \qquad \boldsymbol{V} = \begin{bmatrix} 2.0 & 1.9 \\ 1.9 & 2.0 \end{bmatrix},$$

i.e. in which the two quantities concerned are positively correlated. Similar generators are available elsewhere [12].

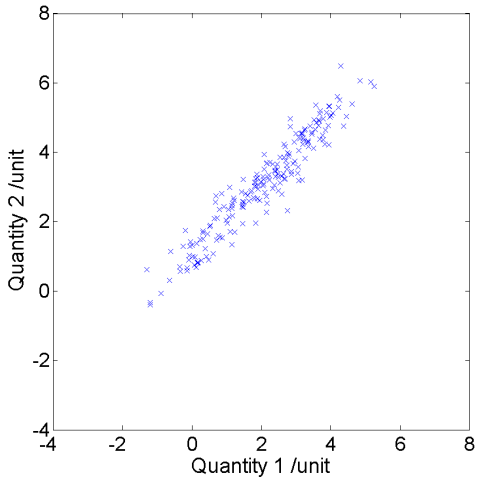

Figure C.1 — Points sampled from a bivariate Gaussian distribution with positive correlation (C.5.3, C.5.4)

C.5.4 In figure C.1, the points span an elongated angled ellipse. Were the off-diagonal elements of $\boldsymbol{V}$ to be replaced by zero, the points would span a circle. Were the diagonal elements made unequal, and the off-diagonal elements kept at zero, the points would span an ellipse whose axes were parallel to the axes of the graph. If the diagonal elements were negative, and hence the quantities concerned negatively correlated, the major axis of the ellipse would have a negative rather than a positive gradient.

C.6 *t*-distribution

The procedure in table C.5 provides an approach [29], [44, page 63] to obtain draws from the t-distribution with ν degrees of freedom.

Table C.5 — A t-distribution pseudo-random number generator (C.6)

Input parameter	
ν	Degrees of freedom
Output parameter	
t	Draw from a t-distribution with ν degrees of freedom
Computation	
	a) Generate two draws r_1 and r_2 independently from the rectangular distribution $\mathrm{R}(0,1)$ b) If $r_1 < 1/2$, form $t = 1/(4r_1 - 1)$ and $v = r_2/t^2$; otherwise form $t = 4r_1 - 3$ and $v = r_2$ c) If $v < 1 - \lvert t\rvert/2$ or $v < (1 + t^2/\nu)^{-(\nu+1)/2}$, accept t as a draw from the t-distribution; otherwise repeat from step a)

NOTE ν must be greater than two for the standard deviation of the t-distribution with ν degrees of freedom to be finite.

Annex D
Continuous approximation to the distribution function for the output quantity

D.1 It is sometimes useful to work with a continuous approximation $\widetilde{G}_Y(\eta)$, say, to the distribution function for the output quantity Y, rather than the discrete representation $\boldsymbol{G}$ of 7.5.

NOTE Working with a continuous approximation means, for instance, that

a) sampling from the distribution function can be carried out without the need for rounding, as in the discrete case, and

b) numerical methods that require continuity for their operation can be used to determine the shortest coverage interval.

D.2 In order to form $\widetilde{G}_Y(\eta)$, consider the discrete representation $\boldsymbol{G} = \{y_{(r)}, r = 1, \ldots, M\}$ of $G_Y(\eta)$ in 7.5.1, after replacing replicate model values of $y_{(r)}$ as necessary (step b) in that subclause). Then, carry out the following steps:

a) assign uniformly spaced cumulative probabilities $p_r = (r - 1/2)/M$, $r = 1, \ldots, M$, to the $y_{(r)}$ [8]. The numerical values p_r, $r = 1, \ldots, M$, are the midpoints of M contiguous probability intervals of width $1/M$ between zero and one;

b) form $\widetilde{G}_Y(\eta)$ as the (continuous) strictly increasing piecewise-linear function joining the M points $(y_{(r)}, p_r)$, $r = 1, \ldots, M$:

$$\widetilde{G}_Y(\eta) = \frac{r - 1/2}{M} + \frac{\eta - y_{(r)}}{M(y_{(r+1)} - y_{(r)})}, \qquad y_{(r)} \le \eta \le y_{(r+1)}, \qquad r = 1, \ldots, M - 1. \tag{D.1}$$

NOTE The form (D.1) provides a convenient basis for sampling from $\widetilde{G}_Y(\eta)$ for purposes of a further stage of uncertainty evaluation. See C.2 for sampling inversely from a distribution function. Some software libraries and packages provide facilities for piecewise-linear interpolation. Since $\widetilde{G}_Y(\eta)$ is piecewise linear, so is its inverse, and such facilities can readily be applied.

D.3 Figure D.1 illustrates $\widetilde{G}_Y(\eta)$ obtained using MCM based on $M = 50$ sampled values from a Gaussian PDF $g_Y(\eta)$ with Y having expectation 3 and standard deviation 1.

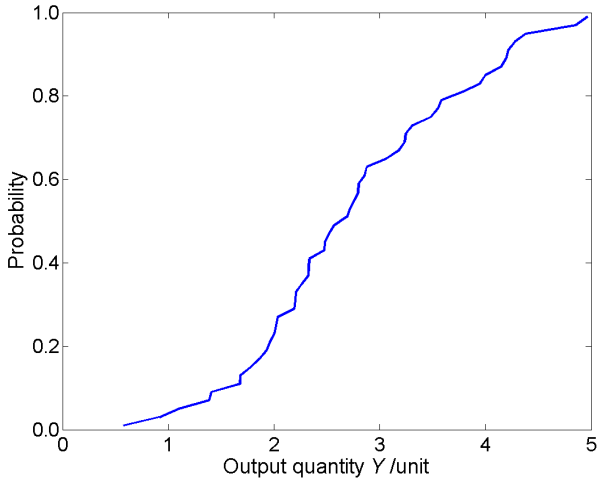

Figure D.1 — An approximation $\widetilde{G}_Y(\eta)$ to the distribution function $G_Y(\eta)$ (D.3). "Unit" denotes any unit

D.4 Consider $\widetilde{g}_Y(\eta) = \widetilde{G}'_Y(\eta)$, with $\widetilde{G}_Y(\eta)$ given in expression (D.1). The function $\widetilde{g}_Y(\eta)$ is piecewise constant with breakpoints at $\eta = y_{(1)}, \ldots, y_{(M)}$. The expectation $\widetilde{y}$ and standard deviation $u(\widetilde{y})$ of Y, described by $\widetilde{g}_Y(\eta)$, are taken, respectively, as an estimate of Y and the standard uncertainty associated with that estimate. $\widetilde{y}$ and $u(\widetilde{y})$ are given by

$$\widetilde{y} = \frac{1}{M} \sum_{r=1}^{M}{}'' \; y_{(r)} \tag{D.2}$$

and

$$u^2(\widetilde{y}) = \frac{1}{M}\left(\sum_{r=1}^{M}{}'' \, (y_{(r)} - \widetilde{y})^2 - \frac{1}{6}\sum_{r=1}^{M-1}(y_{(r+1)} - y_{(r)})^2\right), \tag{D.3}$$

where the double prime on a summation symbol indicates that the first and the last terms in the sum are to be taken with weight one half.

NOTE For a sufficiently large numerical value of M (10^5, say, or greater), $\widetilde{y}$ and $u(\widetilde{y})$ obtained using formulæ (D.2) and (D.3) would generally be indistinguishable for practical purposes from those given by formulæ (16) and (17), respectively.

D.5 Let α denote any value between zero and $1 - p$, where p is the required coverage probability (e.g. 0.95). The endpoints of a $100p$ % coverage interval can be obtained from $\widetilde{G}(\eta)$ by inverse linear interpolation. To determine the lower endpoint y_{low} such that $\alpha = \widetilde{G}_Y(y_{\text{low}})$, identify the index r for which the points $(y_{(r)}, p_r)$ and $(y_{(r+1)}, p_{r+1})$ satisfy

$$p_r \leq \alpha < p_{r+1}.$$

Then, by inverse linear interpolation,

$$y_{\text{low}} = y_{(r)} + \left(y_{(r+1)} - y_{(r)}\right)\frac{\alpha - p_r}{p_{r+1} - p_r}.$$

Similarly, the upper endpoint y_{high}, determined such that $p + \alpha = \widetilde{G}_Y(y_{\text{high}})$, is calculated from

$$y_{\text{high}} = y_{(s)} + \left(y_{(s+1)} - y_{(s)}\right)\frac{p + \alpha - p_s}{p_{s+1} - p_s},$$

where the index s is such that the points $(y_{(s)}, p_s)$ and $(y_{(s+1)}, p_{s+1})$ satisfy

$$p_s \leq p + \alpha < p_{s+1}.$$

D.6 The choice $\alpha = 0.025$ gives the coverage interval defined by the 0.025- and 0.975-quantiles. This choice provides the probabilistically symmetric 95 % coverage interval for Y.

D.7 The shortest coverage interval can generally be obtained from $\widetilde{G}_Y(\eta)$ by determining α such that $\widetilde{G}_Y^{-1}(p + \alpha) - \widetilde{G}_Y^{-1}(\alpha)$, $= H(\alpha)$, say, is a minimum. A straightforward numerical approach to determining the minimum is to evaluate $H(\alpha)$ for a large number of uniformly spaced choices $\{\alpha_k\}$ of α between zero and $1 - p$, and choose α_ℓ from the set $\{\alpha_k\}$ that yields the minimum from the set $\{H(\alpha_k)\}$.

D.8 The computation of a coverage interval is facilitated if pM is an integer. Then, the numerical value of α, such that $H(\alpha)$ is a minimum, is equal to r^*/M, where r^* is the index r such that the interval length $y_{(r+pM)} - y_{(r)}$, over $r = 1, \ldots, (1 - p)M$, is least.

Annex E
Coverage interval for the four-fold convolution of a rectangular distribution

E.1 In 9.2.3.2, the analytic solution

$$\pm 2\sqrt{3}[2 - (3/5)^{1/4}] \approx \pm 3.88 \tag{E.1}$$

is stated. It constitutes the endpoints of the probabilistically symmetric 95 % coverage interval for the output quantity Y in an additive model having four input quantities with expectations of zero and standard deviations of unity, the PDFs for which are identical rectangular distributions. This result is established in this annex.

E.2 The rectangular distribution $\mathrm{R}(a, b)$ (see 6.4.2) takes the constant value $(b - a)^{-1}$ for $a \leq \rho \leq b$ and is zero otherwise. The n-fold convolution of $\mathrm{R}(0, 1)$ is the B-spline $B_n(\rho)$ of order n (degree $n - 1$) with knots $0, \ldots, n$ [46]. An explicit expression is [6]

$$B_n(\rho) = \frac{1}{(n-1)!} \sum_{r=0}^{n} {}^nC_r (-1)^r (\rho - r)_+^{n-1},$$

where

$${}^nC_r = \frac{n!}{r!(n-r)!}, \qquad z_+ = \max(z, 0).$$

In particular,

$$B_4(\rho) = \frac{1}{6}\rho^3, \qquad 0 \leq \rho \leq 1$$

(with different cubic polynomial expressions for $B_4(\rho)$ in other intervals between adjacent knots), and hence

$$\int_0^1 B_4(\rho)\,\mathrm{d}\rho = \left[\frac{1}{24}\rho^4\right]_0^1 = \frac{1}{24} \approx 0.0417.$$

E.3 The left-hand endpoint y_{low} of the probabilistically symmetric 95 % coverage interval lies between zero and one, since

$$0.025 = \frac{1}{40} < \frac{1}{24}$$

of the area under the PDF lies to the left of y_{low}, which is therefore given by

$$\int_0^{y_{\mathrm{low}}} B_4(\rho)\,\mathrm{d}\rho = \frac{1}{24} y_{\mathrm{low}}^4 = \frac{1}{40},$$

i.e.

$$y_{\mathrm{low}} = (3/5)^{1/4}.$$

By symmetry, the right-hand endpoint is

$$y_{\mathrm{high}} = 4 - (3/5)^{1/4}.$$

Thus, the probabilistically symmetric 95 % coverage interval is

$$\left[(3/5)^{1/4}, 4 - (3/5)^{1/4}\right] \equiv 2 \pm \left(2 - (3/5)^{1/4}\right).$$

The corresponding coverage interval for the four-fold convolution of the rectangular PDF $R(-\sqrt{3}, \sqrt{3})$ (which has zero expectation and unit standard deviation) is given by shifting this result by two units and scaling it by $2\sqrt{3}$ units, yielding expression (E.1).

Annex F
Comparison loss problem

This annex is concerned with some details of the comparison loss problem (see 9.4). Subclause F.1 provides the expectation and standard deviation of δY (see 9.4.2.1.2). Subclause F.2 provides the PDF for δY analytically when $x_1 = x_2 = r(x_1,\ x_2) = 0$ (see 9.4.2.1.2). Subclause F.3 applies the GUM uncertainty framework for uncorrelated and correlated input quantities (see 9.4.2.1.3 and 9.4.3.1.1).

F.1 Expectation and standard deviation obtained analytically

F.1.1 The variance of a quantity X can be expressed in terms of expectations as [42, page 124]

$$V(X) = E(X^2) - [E(X)]^2.$$

Thus,

$$E(X^2) = [E(X)]^2 + V(X) = x^2 + u^2(x),$$

where x is the best estimate of X and $u(x)$ the standard uncertainty associated with x. Thus, for the model (28), viz. $\delta Y = 1 - Y = X_1^2 + X_2^2$,

$$\delta y = E(\delta Y) = x_1^2 + x_2^2 + u^2(x_1) + u^2(x_2).$$

This result applies

a) regardless of which PDFs are assigned to X_1 and X_2, and

b) whether X_1 and X_2 are independent or not.

F.1.2 The standard uncertainty associated with δy can be obtained from

$$u^2(\delta y) = u^2(x_1^2) + u^2(x_2^2) + 2u(x_1^2, x_2^2),$$

where, for $i = 1$ and $i = 2$, $u^2(x_i^2) = V(X_i^2)$, and $u(x_1^2, x_2^2) = \mathrm{Cov}(X_1^2, X_2^2)$. Then, applying Price's Theorem for Gaussian distributions [40, 41],

$$u^2(\delta y) = 4u^2(x_1)x_1^2 + 4u^2(x_2)x_2^2 + 2u^4(x_1) + 2u^4(x_2) + 4u^2(x_1, x_2) + 8u(x_1, x_2)x_1x_2. \quad \text{(F.1)}$$

When $x_2 = 0$ and $u(x_2) = u(x_1)$, and replacing $u(x_1, x_2)$ by $r(x_1, x_2)u^2(x_1)$,

$$u(\delta y) = 2\left\{x_1^2 + [1 + r^2(x_1, x_2)]u^2(x_1)\right\}^{1/2} u(x_1).$$

F.1.3 When X_1 and X_2 are uncorrelated, i.e. $u(x_1, x_2) = 0$, expression (F.1) becomes

$$u^2(\delta y) = 4u^2(x_1)x_1^2 + 4u^2(x_2)x_2^2 + 2u^4(x_1) + 2u^4(x_2). \quad \text{(F.2)}$$

Expression (F.2) can be verified by applying formula (10) of the GUM [GUM:1995 5.1.2] and the immediately following GUM formula [GUM:1995 5.1.2 note].

F.2 Analytic solution for zero estimate of the voltage reflection coefficient having associated zero covariance

F.2.1 For the case $x_1 = x_2 = r(x_1, x_2) = 0$ and $u(x_1) = u(x_2)$, the PDF $g_Y(\eta)$ for Y can be obtained analytically. It is valuable to have such a solution for further validation purposes. In the above circumstances,

$$\delta Y = u^2(x_1)\left[\frac{X_1^2}{u^2(x_1)} + \frac{X_2^2}{u^2(x_2)}\right].$$

F.2.2 The term in square brackets is the sum, Z, say, of the squares of two independent quantities, each of which is distributed as a standard Gaussian PDF. Thus the sum is distributed as chi-squared with two degrees of freedom [42, page 177], so that

$$\delta Y = u^2(x_1)Z,$$

where Z has PDF

$$g_Z(z) = \chi_2^2(z) = \mathrm{e}^{-z/2}/2.$$

F.2.3 The application of a general formula [42, pages 57–61] for the PDF $g_Y(\eta)$ of a differentiable and strictly decreasing function of a variable (here Z) with a specified PDF yields

$$g_Y(\eta) = \frac{1}{u^2(x_1)}\chi_2^2\left(\frac{\eta}{u^2(x_1)}\right) = \frac{1}{2u^2(x_1)}\exp\left(-\frac{\eta}{2u^2(x_1)}\right), \qquad \eta \geq 0.$$

F.2.4 The expectation of δY is given by

$$\delta y = E(\delta Y) = \int_0^\infty \eta g_Y(\eta)\,\mathrm{d}\eta = 2u^2(x_1)$$

and the variance

$$u^2(\delta y) = V(\delta Y) = \int_0^\infty (\eta - y)^2 g_Y(\eta)\,\mathrm{d}\eta = 4u^4(x_1),$$

i.e. the standard deviation is $2u^2(x_1)$, results that are consistent with those in F.1.

F.2.5 By integration, the corresponding distribution function is

$$G_Y(\eta) = 1 - \exp\left(-\frac{\eta}{2u^2(x_1)}\right), \qquad \eta \geq 0. \tag{F.3}$$

F.2.6 Let δy_α be that η in expression (F.3) corresponding to $G_Y(\eta) = \alpha$ for any α satisfying $0 \leq \alpha \leq 1 - p$. Then

$$\delta y_\alpha = -2u^2(x_1)\ln(1-\alpha)$$

and a $100p$ % coverage interval for δY (see 7.7) is

$$[\delta y_\alpha,\ \delta y_{p+\alpha}] \equiv [-2u^2(x_1)\ln(1-\alpha),\ -2u^2(x_1)\ln(1-p-\alpha)] \tag{F.4}$$

with length

$$H(\alpha) = -2u^2(x_1)\ln\left(1 - \frac{p}{1-\alpha}\right).$$

F.2.7 The shortest $100p$ % coverage interval is given by determining α to minimize $H(\alpha)$ (see 5.3.4). Since $H(\alpha)$ is a strictly increasing function of α for $0 \leq \alpha \leq 1 - p$, $H(\alpha)$ is minimized when $\alpha = 0$. Thus, the shortest $100p$ %

coverage interval for δY is

$$[0, \ -2u^2(x_1)\ln(1-p)].$$

For $u(x_1) = 0.005$, the shortest 95 % coverage interval is

$$[0, \ 0.000\ 149\ 8].$$

F.2.8 The 95 % probabilistically symmetric coverage interval for δY is given by setting $\alpha = (1-p)/2$ (see 5.3.3):

$$[-2u^2(x_1)\ln 0.975, \ -2u^2(x_1)\ln 0.025] = [0.000\ 001\ 3, \ 0.000\ 184\ 4],$$

which is 20 % longer than the shortest 95 % coverage interval.

NOTE The above analysis is indicative of an analytical approach that can be applied to some problems of this type. In this particular case, the results could in fact have been obtained more directly, since $g_Y(\eta)$ is strictly increasing and the shortest coverage interval is always in the region of highest density.

F.3 GUM uncertainty framework applied to the comparison loss problem

F.3.1 Uncorrelated input quantities

F.3.1.1 The comparison loss problem considered in 9.4 has as the model of measurement

$$\delta Y = f(\boldsymbol{X}) = f(X_1, X_2) = X_1^2 + X_2^2,$$

where X_1 and X_2 are assigned Gaussian PDFs having expectations x_1 and x_2 and variances $u^2(x_1)$ and $u^2(x_2)$, respectively.

F.3.1.2 The application of GUM subclause 5.1.1 gives

$$\delta y = x_1^2 + x_2^2$$

as the estimate of δY. The only non-trivially non-zero partial derivatives of the model are, for $i = 1, 2$,

$$\frac{\partial f}{\partial X_i} = 2X_i, \qquad \frac{\partial^2 f}{\partial X_i^2} = 2.$$

F.3.1.3 Hence the application of GUM subclause 5.1.2 gives, for the standard uncertainty $u(\delta y)$,

$$u^2(\delta y) = \left[\left(\frac{\partial f}{\partial X_1}\right)^2 u^2(x_1) + \left(\frac{\partial f}{\partial X_2}\right)^2 u^2(x_2)\right]\Bigg|_{\boldsymbol{X}=\boldsymbol{x}} = 4x_1^2u^2(x_1) + 4x_2^2u^2(x_2), \tag{F.5}$$

based on a first-order Taylor series approximation of $f(\boldsymbol{X})$. If the non-linearity of f is significant [GUM:1995 5.1.2 note], the term

$$\frac{1}{2}\left[\frac{\partial^2 f}{\partial X_1^2} + \frac{\partial^2 f}{\partial X_2^2}\right]\Bigg|_{\boldsymbol{X}=\boldsymbol{x}} u^2(x_1)u^2(x_2)$$

needs to be appended to formula (F.5), in which case formula (F.5) becomes

$$u^2(\delta y) = 4x_1^2u^2(x_1) + 4x_2^2u^2(x_2) + 4u^2(x_1)u^2(x_2). \tag{F.6}$$

F.3.1.4 A 95 % coverage interval for δY is given by

$$\delta y \pm 2u(\delta y),$$

as a consequence of δY having a Gaussian PDF.

F.3.2 Correlated input quantities

F.3.2.1 When the input quantities are correlated, the uncertainty matrix associated with the best estimates of the input quantities is given in formulæ (27).

F.3.2.2 The application of GUM subclause 5.2.2 gives

$$\begin{aligned} u^2(\delta y) &= \left[\left(\frac{\partial f}{\partial X_1} \right)^2 u^2(x_1) + \left(\frac{\partial f}{\partial X_2} \right)^2 u^2(x_2) + 2 \frac{\partial f}{\partial X_1} \frac{\partial f}{\partial X_2} r(x_1, x_2) u(x_1) u(x_2) \right] \Bigg|_{\boldsymbol{X}=\boldsymbol{x}} \\ &= 4x_1^2 u^2(x_1) + 4x_2^2 u^2(x_2) + 8r(x_1, x_2) x_1 x_2 u(x_1) u(x_2). \end{aligned} \tag{F.7}$$

Annex G
Glossary of principal symbols

A	random variable representing the lower limit of a rectangular distribution with inexactly prescribed limits
a	lower limit of the interval in which a random variable is known to lie
a	midpoint of the interval in which the lower limit A of a rectangular distribution with inexactly prescribed limits is known to lie
B	random variable representing the upper limit of a rectangular distribution with inexactly prescribed limits
b	upper limit of the interval in which a random variable is known to lie
b	midpoint of the interval in which the upper limit B of a rectangular distribution with inexactly prescribed limits is known to lie
$\mathrm{CTrap}(a, b, d)$	rectangular distribution with inexactly prescribed limits (curvilinear trapezoid distribution) with parameters a, b, and d
$\mathrm{Cov}(X_i, X_j)$	covariance for two random variables X_i and X_j
c	n_{dig}-decimal digit integer
c_i	ith sensitivity coefficient, obtained as the partial derivative of the model f for the measurement with respect to the ith input quantity X_i evaluated at the vector estimate $\boldsymbol{x}$ of the vector input quantity $\boldsymbol{X}$
d	semi-width of the intervals in which the lower and upper limits A and B of a rectangular distribution with inexactly prescribed limits are known to lie
d_{high}	absolute value of the difference between the right-hand endpoints of the coverage intervals provided by the GUM uncertainty framework and a Monte Carlo method
d_{low}	absolute value of the difference between the left-hand endpoints of the coverage intervals provided by the GUM uncertainty framework and a Monte Carlo method
e	base of the natural logarithm
$E(X)$	expectation of a random variable X
$E(\boldsymbol{X})$	vector expectation of a vector random variable $\boldsymbol{X}$
$E(X^r)$	rth moment of a random variable X
$\mathrm{Ex}(\lambda)$	exponential distribution with parameter λ
f	mathematical model of measurement, expressed as a functional relationship between an output quantity Y and the input quantities $X_1, \ldots, X_N$ on which Y depends
$\boldsymbol{G}$	discrete representation of the distribution function $G_Y(\eta)$ for the output quantity Y from a Monte Carlo procedure
$\mathrm{G}(\alpha, \beta)$	gamma distribution with parameters α and β

$g_X(\xi)$	probability density function with variable ξ for the input quantity X
$g_{\boldsymbol{X}}(\boldsymbol{\xi})$	joint (multivariate) probability density function with vector variable $\boldsymbol{\xi}$ for the vector input quantity $\boldsymbol{X}$
$g_{X_i}(\xi_i)$	probability density function with variable ξ_i for the input quantity X_i
$G_Y(\eta)$	distribution function with variable η for the output quantity Y
$\widetilde{G}_Y(\eta)$	continuous approximation to the distribution function $G_Y(\eta)$ for the output quantity Y
$g_Y(\eta)$	probability density function with variable η for the output quantity Y
$\widetilde{g}_Y(\eta)$	derivative of $\widetilde{G}_Y(\eta)$ with respect to η, providing a numerical approximation to the probability density function $g_Y(\eta)$ for the output quantity Y
J	smallest integer greater than or equal to $100/(1-p)$
k_p	coverage factor corresponding to the coverage probability p
ℓ	integer in the representation $c \times 10^{\ell}$ of a numerical value, where c is an n_{dig}-decimal digit integer
M	number of Monte Carlo trials
N	number of input quantities $X_1, \ldots, X_N$
$\mathrm{N}(0,1)$	standard Gaussian distribution
$\mathrm{N}(\mu, \sigma^2)$	Gaussian distribution with parameters μ and σ^2
$\mathrm{N}(\boldsymbol{\mu}, \boldsymbol{V})$	multivariate Gaussian distribution with parameters $\boldsymbol{\mu}$ and $\boldsymbol{V}$
n	number of indications in a series
n_{dig}	number of significant decimal digits regarded as meaningful in a numerical value
$\Pr(z)$	probability of the event z
p	coverage probability
q	integer part of $pM + 1/2$
q	number of objects counted in a sample of specified size
$\boldsymbol{R}$	upper triangular matrix
$\mathrm{R}(0,1)$	standard rectangular distribution over the interval $[0,\ 1]$
$\mathrm{R}(a,b)$	rectangular distribution over the interval $[a,\ b]$
$r(x_i, x_j)$	correlation coefficient associated with the estimates x_i and x_j of the input quantities X_i and X_j
s	standard deviation of a series of n indications $x_1, \ldots, x_n$
s_{p}	pooled standard deviation obtained from several series of indications
$\top$	superscript denoting matrix transpose

s_z	standard deviation associated with the average z of the values $z^{(1)}, \ldots, z^{(h)}$ in an adaptive Monte Carlo procedure, where z may denote an estimate y of the output quantity Y, the standard uncertainty $u(y)$ associated with y, or the left-hand endpoint y_{low} or right-hand endpoint y_{high} of a coverage interval for Y
$\mathrm{T}(a, b)$	triangular distribution over the interval $[a,\ b]$
$\mathrm{Trap}(a, b, \beta)$	trapezoidal distribution over the interval $[a,\ b]$ with parameter β
t_ν	central t-distribution with ν degrees of freedom
$t_\nu(\mu, \sigma^2)$	scaled and shifted t-distribution with parameters μ and σ^2, and ν degrees of freedom
$\mathrm{U}(0, 1)$	standard arc sine (U-shaped) distribution over the interval $[0,\ 1]$
$\mathrm{U}(a, b)$	arc sine (U-shaped) distribution over the interval $[a,\ b]$
U_p	expanded uncertainty corresponding to a coverage probability p
$\boldsymbol{U}_{\boldsymbol{x}}$	uncertainty matrix associated with the vector estimate $\boldsymbol{x}$ of the vector input quantity $\boldsymbol{X}$
$\boldsymbol{u}(\boldsymbol{x})$	vector $(u(x_1), \ldots, u(x_N))^\top$ of standard uncertainties associated with the vector estimate $\boldsymbol{x}$ of the vector input quantity $\boldsymbol{X}$
$u(x_i)$	standard uncertainty associated with the estimate x_i of the input quantity X_i
$u(x_i, x_j)$	covariance associated with the estimates x_i and x_j of the input quantities X_i and X_j
$u(y)$	standard uncertainty associated with the estimate y of the output quantity Y
$u(\widetilde{y})$	standard uncertainty associated with $\widetilde{y}$
$u_{\mathrm{c}}(y)$	combined standard uncertainty associated with the estimate y of the output quantity Y
$u_i(y)$	ith uncertainty component of the standard uncertainty $u(y)$ associated with the estimate y of the output quantity Y
$\boldsymbol{V}$	covariance (variance-covariance) matrix
$V(X)$	variance of a random variable X
$V(\boldsymbol{X})$	covariance matrix for the vector random variable $\boldsymbol{X}$
w	semi-width $(b - a)/2$ of an interval $[a,\ b]$
X	input quantity, regarded as a random variable
$\boldsymbol{X}$	vector $(X_1, \ldots, X_N)^\top$ of input quantities, regarded as random variables, on which the output quantity Y depends
X_i	ith input quantity, regarded as a random variable, on which the output quantity Y depends
x	estimate (expectation) of X
$\boldsymbol{x}$	vector estimate (vector expectation) $(x_1, \ldots, x_N)^\top$ of $\boldsymbol{X}$
$\bar{x}$	average of a series of n indications $x_1, \ldots, x_n$

x_i	estimate (expectation) of X_i
x_i	ith indication in a series
$x_{i,r}$	rth Monte Carlo draw from the probability density function for X_i
$\boldsymbol{x}_r$	rth Monte Carlo draw, containing values $x_{1,r}, \ldots, x_{N,r}$, drawn from the probability density functions for the N input quantities $X_1, \ldots, X_N$ or from the joint probability density function for $\boldsymbol{X}$
Y	(scalar) output quantity, regarded as a random variable
y	estimate (expectation) of Y
$\widetilde{y}$	estimate of Y, obtained as the average of the M model values y_r from a Monte Carlo run or as the expectation of Y characterized by the probability density function $\widetilde{g}_Y(\eta)$
y_{high}	right-hand endpoint of a coverage interval for Y
y_{low}	left-hand endpoint of a coverage interval for Y
y_r	rth model value $f(\boldsymbol{x}_r)$
$y_{(r)}$	rth model value after sorting the y_r into increasing order
$z^{(h)}$	hth value in an adaptive Monte Carlo procedure, where z may denote an estimate y of the output quantity Y, the standard uncertainty $u(y)$ associated with y, or the left-hand endpoint y_{low} or right-hand endpoint y_{high} of a coverage interval for Y
α	probability value
α	parameter of a gamma distribution
β	parameter of a trapezoidal distribution equal to the ratio of the semi-width of the top of the trapezoid to that of the base
β	parameter of a gamma distribution
$\Gamma(z)$	gamma function with variable z
δ	numerical tolerance associated with a numerical value
$\delta(z)$	Dirac delta function with variable z
η	variable describing the possible values of the output quantity Y
λ_1	top semi-width of the trapezoid for a trapezoidal distribution
λ_2	base semi-width of the trapezoid for a trapezoidal distribution
μ	expectation of a quantity characterized by a probability distribution
ν	degrees of freedom of a t-distribution or a chi-squared distribution
ν_{eff}	effective degrees of freedom associated with the standard uncertainty $u(y)$
ν_{p}	degrees of freedom associated with a pooled standard deviation s_{p} obtained from several series of indications

ξ	variable describing the possible values of the random variable X
$\boldsymbol{\xi}$	vector variable $(\xi_1, \ldots, \xi_N)^\top$ describing the possible values of the vector input quantity $\boldsymbol{X}$
ξ_i	variable describing the possible values of the input quantity X_i
σ	standard deviation of a quantity characterized by a probability distribution
σ^2	variance (squared standard deviation) of a quantity characterized by a probability distribution
Φ	phase of a quantity that cycles sinusoidally
χ^2_ν	chi-squared distribution with ν degrees of freedom

Bibliography

[1] Beatty, R. W. Insertion loss concepts. *Proc. IEEE 52* (1964), 663–671.

[2] Berthouex, P. M., and Brown, L. C. *Statistics for Environmental Engineers*. CRC Press, USA, 1994.

[3] Box, G. E. P., and Muller, M. A note on the generation of random normal variates. *Ann. Math. Statist. 29* (1958), 610–611.

[4] Chan, A., Golub, G., and LeVeque, R. Algorithms for computing the sample variance: analysis and recommendations. *Amer. Stat. 37* (1983), 242–247.

[5] Conte, S. D., and de Boor, C. *Elementary Numerical Analysis: An Algorithmic Approach*. McGraw-Hill, 1972.

[6] Cox, M. G. The numerical evaluation of B-splines. *J. Inst. Math. Appl. 10* (1972), 134–149.

[7] Cox, M. G., and Harris, P. M. Software specifications for uncertainty evaluation. Tech. Rep. DEM-ES-010, National Physical Laboratory, Teddington, UK, 2006.

[8] Cox, M. G., and Harris, P. M. SSfM Best Practice Guide No. 6, Uncertainty evaluation. Tech. Rep. DEM-ES-011, National Physical Laboratory, Teddington, UK, 2006.

[9] Cox, M. G., and Siebert, B. R. L. The use of a Monte Carlo method for evaluating uncertainty and expanded uncertainty. *Metrologia 43* (2006), S178–S188.

[10] David, H. A. *Order Statistics*. Wiley, New York, 1981.

[11] Dekker, T. J. Finding a zero by means of successive linear interpolation. In *Constructive Aspects of the Fundamental Theorem of Algebra* (London, 1969), B. Dejon and P. Henrici, Eds., Wiley Interscience.

[12] Devroye, L. *Non-Uniform Random Number Generation*. Springer, New York, 1986.

[13] Dietrich, C. F. *Uncertainty, Calibration and Probability*. Adam Hilger, Bristol, UK, 1991.

[14] Dowson, D. C., and Wragg, A. Maximum entropy distributions having prescribed first and second order moments. *IEEE Trans. IT 19* (1973), 689–693.

[15] EA. Expression of the uncertainty of measurement in calibration. Tech. Rep. EA-4/02, European Co-operation for Accreditation, 1999.

[16] Elster, C. Calculation of uncertainty in the presence of prior knowledge. *Metrologia 44* (2007), 111–116.

[17] EURACHEM/CITAC. Quantifying uncertainty in analytical measurement. Tech. Rep. Guide CG4, EURACHEM/CITEC, 2000. Second edition.

[18] Evans, M., Hastings, N., and Peacock, B. *Statistical distributions*. Wiley, 2000.

[19] Frenkel, R. B. Statistical background to the ISO 'Guide to the Expression of Uncertainty in Measurement'. Tech. Rep. Monograph 2, NML Technology Transfer Series, Publication number TIP P1242, National Measurement Laboratory, CSIRO, Australia, 2002.

[20] Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. *Bayesian Data Analysis*. Chapman and Hall, London, 2004.

[21] Gleser, L. J. Assessing uncertainty in measurement. *Stat. Sci. 13* (1998), 277–290.

[22] Hall, B. D., and Willink, R. Does "Welch-Satterthwaite" make a good uncertainty estimate? *Metrologia 38* (2001), 9–15.

[23] Higham, N. J. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, 1996.

[24] ISO. ISO 3534-1. Statistics — Vocabulary and symbols — Part 1: Probability and general statistical terms. International Standards Organization, Geneva.

[25] Jaynes, E. T. Information theory and statistical mechanics. *Phys. Rev 106* (1957), 620–630.

[26] Jaynes, E. T. Where do we stand on maximum entropy? In *Papers on Probability, Statistics, and Statistical Physics* (Dordrecht, The Netherlands, 1989), R. D. Rosenkrantz, Ed., Kluwer Academic, pp. 210–314. `http://bayes.wustl.edu/etj/articles/stand.on.entropy.pdf`.

[27] Kacker, R., and Jones, A. On use of Bayesian statistics to make the Guide to the Expression of Uncertainty in Measurement consistent. *Metrologia 40* (2003), 235–248.

[28] Kerns, D. M., and Beatty, R. W. *Basic Theory of Waveguide Junctions and Introductory Microwave Network Analysis.* Pergamon Press, London, 1967.

[29] Kinderman, A., Monahan, J., and Ramage, J. Computer methods for sampling from Student's t-distribution. *Math. Comput 31* (1977), 1009–1018.

[30] L'Ecuyer, P., and Simard, R. TestU01: A software library in ANSI C for empirical testing of random number generators. `http://www.iro.umontreal.ca/~simardr/testu01/tu01.html`.

[31] Leydold, J. Automatic sampling with the ratio-of-uniforms method. *ACM Trans. Math. Software 26* (2000), 78–98.

[32] Lira, I. *Evaluating the Uncertainty of Measurement. Fundamentals and Practical Guidance.* Institute of Physics, Bristol, UK, 2002.

[33] Lira, I. H., and Wöger, W. Bayesian evaluation of the standard uncertainty and coverage probability in a simple measurement model. *Meas. Sci. Technol. 12* (2001), 1172–1179.

[34] Matsumoto, M., and Nishimura, T. Mersenne Twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Modeling and Computer Simulation 8* (1998), 3–30.

[35] McCullough, B. D., and Wilson, B. On the accuracy of statistical procedures in Microsoft Excel 2003. *Computational Statistics and Data Analysis* (2004).

[36] Moler, C. B. *Numerical computing with MATLAB.* SIAM, Philadelphia, 2004.

[37] Netlib. The Netlib repository of freely available software, documents, and databases of interest to the numerical, scientific computing, and other communities contains facilities for sampling from probability distributions, `http://www.netlib.org`.

[38] NIST. The NIST Digital Library of Mathematical Functions contains facilities for sampling from probability distributions, `http://dlmf.nist.gov`.

[39] OIML. Conventional value of the result of weighing in air. Tech. Rep. OIML D 28, Organisation Internationale de Métrologie Légale, Paris, 2004.

[40] Papoulis, A. On an extension of Price's theorem. *IEEE Trans. Inform. Theory IT-11* (1965).

[41] Price, R. A useful theorem for nonlinear devices having Gaussian inputs. *IEEE Trans. Inform. Theory IT-4* (1958), 69–72.

[42] Rice, J. R. *Mathematical Statistics and Data Analysis*, second ed. Duxbury Press, Belmont, Ca., USA, 1995.

[43] Ridler, N. M., and Salter, M. J. Propagating S-parameter uncertainties to other measurement quantities. In *58th ARFTG (Automatic RF Techniques Group) Conference Digest* (2001).

[44] Robert, C. P., and Casella, G. *Monte Carlo Statistical Methods.* Springer-Verlag, New York, 1999.

[45] Salter, M. J., Ridler, N. M., and Cox, M. G. Distribution of correlation coefficient for samples taken from a bivariate normal distribution. Tech. Rep. CETM 22, National Physical Laboratory, Teddington, UK, 2000.

[46] Schoenberg, I. J. Cardinal interpolation and spline functions. *J. Approx. Theory 2* (1969), 167–206.

[47] Scowen, R. S. QUICKERSORT, Algorithm 271. *Comm. ACM 8* (1965), 669.

[48] Shannon, C. E. A mathematical theory of information. *Bell Systems Tech. J. 27* (1948), 623–656.

[49] Strang, G., and Borre, K. *Linear Algebra, Geodesy and GPS*. Wiley, Wellesley-Cambridge Press, 1997.

[50] Taylor, B. N., and Kuyatt, C. E. Guidelines for evaluating and expressing the uncertainty of NIST measurement results. Tech. Rep. TN1297, National Institute of Standards and Technology, USA, 1994.

[51] Weise, K., and Wöger, W. A Bayesian theory of measurement uncertainty. *Meas. Sci. Technol. 3* (1992), 1–11.

[52] Wichmann, B. A., and Hill, I. D. Algorithm AS183. An efficient and portable pseudo-random number generator. *Appl. Statist. 31* (1982), 188–190.

[53] Wichmann, B. A., and Hill, I. D. Correction. algorithm AS183. An efficient and portable pseudo-random number generator. *Appl. Statist. 33* (1984), 123.

[54] Wichmann, B. A., and Hill, I. D. Generating good pseudo-random numbers. *Computational Statistics and Data Analysis 51* (2006), 1614–1622.

[55] Willink, R. Coverage intervals and statistical coverage intervals. *Metrologia 41* (2004), L5–L6.

[56] Wöger, W. Probability assignment to systematic deviations by the Principle of Maximum Entropy. *IEEE Trans. Instr. Measurement IM-36* (1987), 655–658.

Alphabetical index