

PHYSICS FOR ENGINEERING STUDENTS

Propagation of Measurement Uncertainty I

From directly measured quantities to calculated results

Lecture 8 — Book-style student notes

University of Szeged • Institute of Physics

Propagation of Measurement Uncertainty I

Measurement uncertainty must follow the same mathematical path as the measured information. When an output is calculated from uncertain inputs, the measurement model transforms both the best estimates and their probability structure. This lecture develops the practical rules for that transformation, beginning with independence and linear combinations and ending with nonlinear models, uncertainty budgets and complete reporting.

Learning objectives

- Distinguish direct measurements from calculated measurands.
- Recognize independence, correlation and shared uncertainty sources.
- Propagate standard uncertainties through linear and nonlinear models.
- Use sensitivity coefficients, relative forms and uncertainty budgets.
- Identify dominant contributions and report a traceable result.

Central principle

For small uncertainties, the measurement model can be linearized locally. Each input uncertainty is multiplied by a sensitivity coefficient, and statistically independent contributions combine as variances in quadrature.

Further interpretation

The numerical formula is only the last step of a physical argument. The analyst must define the measurand, identify inputs and correlations, justify the uncertainty assigned to each source, and check whether the local approximation represents the actual range of possible values.

1. Direct and indirect measurement

A directly measured quantity is read from an instrument or produced by a measurement channel. An indirectly obtained quantity is calculated from one or more measured inputs through a measurement model. The distinction is practical rather than philosophical: a caliper may provide a length directly, while the acceleration due to gravity, the volume of a sphere, or the density of a body is inferred through a formula. Every input is uncertain, so the calculated result must also carry a quantitatively evaluated uncertainty.

$$Q = f(X_1, X_2, \dots, X_n)$$

The measurement model defines the measurand. For a simple pendulum, the directly observed length l and period T enter the ideal relation $g = 4\pi^2 l / T^2$. The nominal result is obtained by substituting the best estimates of l and T . The remaining question is how small changes in those inputs alter g . This question cannot be answered from the numerical uncertainty values alone; it also requires the mathematical sensitivity of the model to each input.

$$g = 4\pi^2 l / T^2$$

Further interpretation

Uncertainty propagation is therefore part of the measurement model, not an optional calculation added at the end. Before propagating anything, define the output quantity, its unit, the input estimates, the statistical meaning of each input uncertainty, and the conditions under which the model is valid. If an input correction or environmental influence is omitted, the resulting uncertainty can be numerically precise but physically incomplete.

A useful habit is to draw the information path before calculating: indicate which quantities are read, which are corrected, and which are combined to form the result. This exposes hidden inputs such as calibration factors, alignment corrections and environmental coefficients. It also separates uncertainty in the observations from uncertainty in the physical model. In the pendulum example, finite amplitude, an uncertain effective length and timing resolution may all matter. The concise textbook equation is the starting point, while the actual measurement equation must represent the implemented experiment closely enough that residual model error is either negligible or included as an uncertainty component.

In written solutions, introduce symbols before using them and distinguish an estimate from a true or nominal value. A compact result such as $Q = Q_0 \pm u(Q)$ should appear only after the model and uncertainty convention have been stated. This ordering helps a reader see whether uncertainty belongs to a directly observed input, a corrected value, or the calculated measurand.

In practice, identify the directly observed inputs, write their units, and state the calculated output before substituting numbers. If the output definition changes, rebuild the model and budget instead of reusing a propagation formula from a superficially similar experiment.

2. Independence, correlation and shared influences

Statistical independence concerns fluctuations around the expected values. Two quantities can be connected by a physical law and still have independent measurement deviations when they are obtained by different procedures. In a pendulum experiment, the period and length are functionally related through the model, yet reaction-time scatter in T may be independent of scale-reading scatter in l . Conversely, quantities with no direct physical relation can become statistically dependent when they share an instrument, calibration, operator, environment, or data-processing step.

$$X = X_0 + \Delta X, \quad Y = Y_0 + \Delta Y$$

Write each result as a mean plus a deviation. Independence means that knowledge of one deviation does not improve prediction of the other. A common scale-factor error or zero offset produces correlated deviations. In that case covariance terms must be retained. The sign of the covariance matters: positive correlation can increase the uncertainty of a sum and reduce that of a difference, while negative correlation has the opposite effect.

$$\text{Cov}(X, Y) = E[(X - X_0)(Y - Y_0)]$$

Further interpretation

Independence should never be assumed merely because two readings have different names. Trace the measurement chain. If both values use the same ruler, clock, calibration factor, reference pressure, temperature correction, or fitted baseline, they may share an uncertainty component. A useful analysis separates independent repeatability contributions from common systematic components and assigns covariance only where a physical or procedural link justifies it.

Correlation can be investigated experimentally. Measure the two inputs repeatedly in paired order and examine a scatter plot or correlation coefficient, while remembering that a small sample can hide dependence. Controlled reversals are often more informative: use the same reference in both channels, then repeat with separate references, or interchange instruments and observe whether the difference follows the device. Correlation is also created by data reduction. Two reported quantities calculated from the same fitted parameters or baseline are rarely independent even if their final labels differ. The covariance matrix from the fit should accompany those parameters into the next propagation step.

A correlation coefficient describes only linear association and does not by itself prove independence. More importantly, covariance values must match the units and statistical scale of the input uncertainties. If correlation is represented by ρ_{XY} , then $\text{Cov}(X,Y) = \rho_{XY}u(X)u(Y)$. Testing the limiting cases $\rho = +1, 0$, and -1 often reveals whether the implemented formula has the correct sign.

In practice, draw arrows from every shared device or correction to the quantities it affects. Shared arrows indicate possible covariance. Decide the expected sign by asking whether an upward change in the influence makes both reported quantities rise, makes both fall, or moves them oppositely.

3. Linear combinations and variance addition

Many measurement equations are linear combinations. For $Q = a + bX + cY + dZ$, the expected value follows ordinary algebra: substitute the expected values of the inputs. For statistically independent inputs, the variance of the output is the sum of the coefficient-scaled input variances. The standard uncertainty is the square root of that sum. The coefficients b , c , and d convert each input uncertainty into the units of Q and determine its influence on the result.

$$Q_0 = a + bX_0 + cY_0 + dZ_0$$

This rule explains why standard uncertainties are not normally added arithmetically. Independent fluctuations can occur with different signs, so their variances, rather than their signed deviations, accumulate. Geometrically, the individual uncertainty contributions are orthogonal components in an abstract uncertainty space. Their root-sum-of-squares magnitude is smaller than the conservative worst-case sum unless only one contribution is nonzero.

$$u^2(Q) = b^2u^2(X) + c^2u^2(Y) + d^2u^2(Z)$$

Further interpretation

The quadrature rule depends on two conditions: the model is linear in the relevant region, and the retained inputs are uncorrelated. If correlation is present, cross terms must be added. If coefficients themselves are uncertain, they are inputs too. Reporting only the final square root hides the structure, so a transparent calculation lists every sensitivity-weighted contribution before combining them.

The coefficients in a linear combination are sensitivities and should not be treated as decorative algebra. A unit conversion by 1000 multiplies both the value and its absolute uncertainty by 1000, leaving relative uncertainty unchanged. Averaging N equally weighted independent readings uses coefficients $1/N$, which yields the familiar reduction of standard uncertainty by $\sqrt{N}$. Weighted averages use unequal coefficients and require trustworthy input uncertainties. If one coefficient is negative, it changes the nominal combination, but its independent variance contribution remains positive after squaring. These simple cases provide valuable checks on software used for larger budgets.

A useful numerical check compares quadrature with the conservative linear sum. The quadrature result must not exceed the sum of absolute independent contributions. If it does, a square, unit conversion, or coefficient has probably been mishandled. The check is not a substitute for the correct statistical model, but it catches many spreadsheet and calculator errors before they enter a report.

In practice, calculate each absolute contribution $|b_i|u(X_i)$ separately and label its unit. Square only after the conversion into output units. This sequence makes it easier to find a wrong coefficient and to explain the result without hiding all information inside one calculator expression.

4. Addition, subtraction and pressure examples

Addition and subtraction are special cases of a linear model. For $Q = X + Y$ and $Q = X - Y$, independent uncertainty contributions both enter with positive squares. Values subtract in the second model, but uncertainties do not. This is especially important when two large readings are differenced to obtain a small signal: cancellation in the nominal value can make the relative uncertainty of the difference very large even though the absolute uncertainty remains governed by the two original readings.

$$u(X \pm Y) = \sqrt{u^2(X) + u^2(Y)}$$

Absolute pressure illustrates addition. If gauge pressure is 0.475 ± 0.004 bar and atmospheric pressure is 0.988 ± 0.002 bar, their sum is 1.463 bar, and the standard uncertainty is $\sqrt{(0.004^2 + 0.002^2)} = 0.00447$ bar, assuming independence. Differential pressure illustrates subtraction: the nominal pressures subtract, while their independent variances add. If the same reference pressure occurs in both terms, covariance may lead to partial or complete cancellation and must be modeled explicitly.

$$P_{\text{abs}} = P_{\text{rel}} + P_{\text{atm}} = 1.463 \pm 0.0045$$

Further interpretation

A plus-or-minus symbol has meaning only when its statistical convention is clear. The examples above use standard uncertainties. A worst-case tolerance calculation would add absolute limits linearly and answers a different engineering question. When pressure values share a barometer, reference atmosphere, or calibration curve, the independence assumption must be reconsidered before applying the simple quadrature formula.

Pressure differences demonstrate why measurand definition and correlation belong together. If two absolute pressures are measured by independent instruments, the reference uncertainties add in quadrature. If a differential sensor measures the pressure difference directly, a large common atmospheric component may never enter the measurement model. If two channels share the same atmospheric correction, that shared term can cancel. These arrangements may have identical nominal equations but different covariance structures. The most efficient experiment is often the one that creates common-mode cancellation deliberately, provided the shared influence is stable, and its correlation is documented rather than silently assumed.

Subtraction is particularly sensitive to interpretation when the result is close to zero. A differential pressure consistent with zero can still be an informative outcome, but a percentage uncertainty becomes unhelpful. Report the absolute difference and its standard uncertainty, then state whether zero or a decision threshold lies within an appropriate expanded interval.

In practice, compute the nominal pressure first, then the uncertainty. Keep signs in the nominal equation and covariance terms, but remember that independent variances are nonnegative. Finally, compare the uncertainty with the size of the differential signal rather than with either large absolute pressure alone.

5. Nonlinear functions and local linearization

For a nonlinear function $Q = f(X)$, the same input uncertainty can produce different output uncertainties at different operating points. A quadratic curve is nearly flat near zero and steeper at larger $|X|$, so a fixed horizontal interval corresponds to different vertical changes. The local derivative captures this behavior. Over a sufficiently small interval, the nonlinear curve is approximated by its tangent, and the output deviation is the derivative multiplied by the input deviation.

$$f(X) \approx f(X_0) + (\partial f / \partial X)_0 (X - X_0)$$

The derivative is called the sensitivity coefficient. Its magnitude tells how strongly Q responds to X , its sign tells the direction of change, and its unit is the unit of Q divided by the unit of X . Standard uncertainty uses the magnitude because spread is nonnegative. For asymmetric or large input uncertainty, however, the upper and lower mapped intervals can differ, and a first-order symmetric approximation may be inadequate.

$$u(Q) \approx |(\partial Q / \partial X)_0| u(X)$$

Further interpretation

Local linearization is a controlled approximation. Its adequacy depends on both curvature and the range of plausible input values. A practical check evaluates the full nonlinear model at $X_0 \pm u(X)$ or uses a second-order term. If the mapped changes differ materially, report asymmetric limits or use numerical propagation such as Monte Carlo simulation rather than forcing the result into a symmetric first-order form.

Curvature can introduce both asymmetric uncertainty and a shift in the expected output. For a convex function, the mean of $f(X)$ is generally not exactly $f(E[X])$; this is a consequence of Jensen's inequality and becomes noticeable when relative uncertainty is large. First-order propagation estimates spread around the nominal result but neglects this second-order effect. A second derivative or Monte Carlo calculation can test its importance. Sample plausible input values from their assigned distributions, evaluate the complete model, and summarize the simulated output distribution. This numerical approach naturally handles bounds, asymmetry, and interactions that a single derivative cannot represent.

Finite-difference sensitivity checks are valuable when analytic derivatives are cumbersome. Change one input by a small known amount h while holding others fixed and estimate $c_i \approx [f(x_i+h) - f(x_i-h)]/(2h)$. Repeat with a smaller h to check numerical stability. Agreement with the analytic derivative provides an independent verification of the propagation calculation.

In practice, plot or tabulate $f(X_0-u)$, $f(X_0)$, and $f(X_0+u)$. Nearly equal upper and lower changes support a symmetric first-order approximation. A visible asymmetry is evidence to retain nonlinear behavior or report unequal upper and lower uncertainty limits.

6. General propagation law for several inputs

For $Q = f(X_1, \dots, X_n)$, a first-order Taylor expansion around the nominal input vector converts the nonlinear model into a local linear combination. Each partial derivative $c_i = \partial Q / \partial X_i$ is evaluated at the operating point while the other inputs are held fixed. For independent inputs, the combined standard uncertainty is the root-sum-of-squares of $c_i u(X_i)$. This is the general law used throughout experimental science and metrology.

$$u^2(Q) = \sum_i [(\partial Q / \partial X_i) u(X_i)]^2$$

The full correlated form adds covariance terms. Writing the calculation in matrix notation is especially useful for fitted parameters or calibration coefficients, where correlations are common. The gradient of the measurement model multiplies the input covariance matrix on both sides. This formulation makes coordinate changes transparent and prevents the accidental omission of cross terms.

$$u^2(Q) = c^T U c$$

Further interpretation

The formula is applied in four steps: write the model, calculate sensitivities, convert every source to a standard uncertainty, and combine the resulting contributions with the appropriate covariance structure. Derivatives must be evaluated using consistent units. A dimensionless sensitivity is often obtained by multiplying c_i by X_i/Q ; this shows directly how a relative input uncertainty affects the relative output uncertainty.

Matrix propagation is especially valuable for calibration. Suppose a straight-line calibration yields an intercept and slope that are strongly negatively correlated. Using their individual uncertainties while omitting covariance can seriously misstate the uncertainty of a predicted concentration. The covariance matrix preserves this relationship. In computational work, confirm that the matrix is symmetric and positive semidefinite and that variables appear in the same order in the gradient and covariance matrix. A unit test with a simple independent case should reproduce the scalar quadrature result before the method is applied to a complex model.

For vector-valued outputs, the same method produces an output covariance matrix rather than one scalar variance. This is relevant when a calibration returns several concentrations or when position has multiple coordinates. Correlations between outputs should be retained for later calculations; reducing every output to an isolated plus-or-minus value discards information that may be important in subsequent combinations.

In practice, arrange inputs in a fixed order and use that same order in the model, gradient and covariance matrix. Annotating each derivative with units and its evaluation point prevents a symbolic expression from being applied at the wrong operating condition.

7. Uncertainty budgets as analytical tools

An uncertainty budget is a structured summary of all significant uncertainty components affecting a result. For each source it records the input estimate, probability model or evaluation route, standard uncertainty u_i , sensitivity coefficient c_i , and contribution $|c_i|u_i$. Squared contributions are summed to obtain the combined variance when inputs are independent. Type A and Type B labels describe how the uncertainty was evaluated, not whether the physical effect is random or systematic.

$$u_{c(Q)} = \sqrt{\sum_i (c_i u_i)^2}$$

A budget supports both reporting and experimental design. Ranking the variance contributions reveals which improvement can materially reduce $u(Q)$. A component that contributes 70% of the variance deserves more attention than one contributing 1%; reducing a negligible term produces almost no change in the total. The budget also exposes unit inconsistencies, duplicated effects, missing correlations, and corrections that were applied without an accompanying uncertainty.

$$r_i = (c_i u_i)^2 / u_{c(Q)}^2$$

Further interpretation

A good budget is traceable rather than merely tabular. Every line should point to data, calibration information, specifications, or a justified distribution. Components should represent distinct physical effects to avoid double counting. After the calculation, inspect the fractional variance contributions and perform a sensitivity study: the dominant term may come from a large input uncertainty, a large derivative, or both.

An uncertainty budget should distinguish input quantity from uncertainty source. One temperature reading may support several corrections, while one instrument reading may have resolution, repeatability, calibration and drift components. The analyst can either combine independent components at the input level or retain separate source lines, but the choice must avoid double counting. Scenario tests are useful: set one contribution to zero, double it, or change its correlation and observe the impact on u_c . These tests reveal whether the budget behaves physically and whether a claimed dominant source actually controls the final result over realistic conditions.

Budget completeness can be tested by asking what would happen if the experiment were repeated on another day, with another operator or another instrument. Any plausible change not represented in the listed inputs may indicate a missing reproducibility component. The intended scope of the reported result determines how widely the uncertainty model must generalize beyond the immediate run.

In practice, verify that the listed relative variance contributions sum to 100% within rounding. A missing percentage, duplicated source, or contribution expressed in the wrong unit will often appear immediately in this reconciliation.

8. Products, quotients and relative uncertainty

Products and quotients are most naturally expressed in relative form. For $Q = XY$ with independent inputs, differentiation gives contributions $Y u(X)$ and $X u(Y)$. Dividing by $Q = XY$ yields a quadrature sum of the relative input uncertainties. The same result holds for $Q = Y/X$: division changes the nominal value and derivative signs, but squaring removes the signs from independent variance contributions.

$$[u(Q)/Q]^2 = [u(X)/X]^2 + [u(Y)/Y]^2$$

Relative formulas are dimensionless and therefore convenient for scale changes and percentage reporting. They also make dominance easy to see. If $u(X)/X$ is much larger than $u(Y)/Y$, then the relative uncertainty of the result is approximately the larger term. This shortcut is useful for planning, but the exact quadrature calculation should be retained in the reported budget.

$$Q = XY \text{ or } Q = Y/X$$

Further interpretation

Division never implies subtracting uncertainties. That tempting rule can produce an unrealistically small or even zero uncertainty when relative input uncertainties are similar. Correlation changes the expression: a shared scale factor can cancel in a ratio, while independent scale factors do not. The model must distinguish common and separate contributions before choosing the simple relative formula.

Relative uncertainty becomes misleading when a nominal value is close to zero. A differential pressure or background-subtracted signal can have a large or undefined relative uncertainty even though its absolute standard uncertainty is well characterized. In such cases, report the absolute quantity and uncertainty, and use a signal-to-uncertainty ratio only for a clearly defined decision. Products and quotients also require attention to sign and support. If a denominator has appreciable probability near zero, the output distribution can become strongly skewed or heavy-tailed, and first-order relative propagation is no longer a safe description.

When a quotient contains background subtraction in its numerator or denominator, derive the uncertainty from the full unsimplified model. Algebraic cancellation of nominal terms does not always imply cancellation of their uncertainties unless the underlying quantities are identical and correlated appropriately. Maintaining the original input identities prevents accidental loss or duplication of contributions.

In practice, derive the quotient once using partial derivatives, then verify it with logarithmic differentiation. Agreement between the two routes is a strong algebraic check. If they disagree, revisit the definition of numerator, denominator, and any shared variables.

9. Power laws and dimensional sensitivity

For a power law $Q = aX^n$, first-order propagation gives $u(Q)/|Q| = |n|u(X)/|X|$. The exponent is an uncertainty amplification factor. Squaring a measured length doubles its relative uncertainty in area; cubing it triples the relative uncertainty in volume. Negative powers behave similarly in magnitude: inversion changes the direction of the response but not the positive size of the relative uncertainty contribution.

$$u(Q)/|Q| \approx |n|u(X)/|X|$$

For a general product of powers $Q = aX^pY^qZ^r$, independent relative contributions combine in quadrature after multiplication by their exponents. This logarithmic form is compact and dimensionally transparent. It is valid for small relative uncertainties and nonzero nominal inputs. Near zero, relative uncertainty becomes unstable, and the original absolute derivative form should be used.

$$[u(Q)/Q]^2 = p^2[u(X)/X]^2 + q^2[u(Y)/Y]^2 + r^2[u(Z)/Z]^2$$

Further interpretation

Power-law rules are not separate laws; they are consequences of the general derivative method. Their value is speed and insight. They show immediately which exponents amplify uncertainty and help allocate measurement effort. Always return to the original model when inputs are correlated, offsets are present, the output crosses zero, or uncertainties are too large for local linearization.

Logarithmic differentiation offers a quick derivation for positive product-of-power models: differentiate $\ln Q = \ln a + p \ln X + q \ln Y + r \ln Z$. The differential dQ/Q is then a weighted sum of relative input differentials. This reveals a useful invariance: expressing a length in metres or millimetres changes neither its relative contribution nor the exponent multiplier. The shortcut must still respect correlation. For example, density and volume calculated from the same diameter share inputs; treating them as independent in a later formula can count the same diameter fluctuation more than once.

A log transformation can also stabilize calculations when positive quantities span several orders of magnitude. However, a symmetric uncertainty in log space corresponds to an asymmetric interval in ordinary units. The reporting scale should reflect the distribution and scientific use of the result; do not convert back to a symmetric plus-or-minus form without checking the mapped limits.

In practice, use the exponent rule to predict the approximate answer before detailed calculation. The prediction should match the final relative uncertainty when assumptions hold. A large mismatch signals offsets, correlations, large uncertainty or a model that is not a pure power law.

10. Worked geometry: rectangle, circle and cylinder

Consider a rectangle with $a = 13.00 \pm 0.05$ cm and $b = 25.00 \pm 0.05$ cm. The nominal area is $A = ab = 325.0$ cm². The relative standard uncertainty is $\sqrt{[(0.05/13.00)^2 + (0.05/25.00)^2]} = 0.00434$, giving $u(A) \approx 1.41$ cm². The shorter side contributes more because the same absolute reading uncertainty is a larger fraction of 13 cm than of 25 cm.

$$A = ab, \quad u(A) = A \sqrt{(u(a)/a)^2 + (u(b)/b)^2}$$

For a circle $A = \pi R^2$, the sensitivity is $2\pi R$ and the relative uncertainty is $2u(R)/R$. For a cylinder, $V = \pi R^2 H$, the radius contribution is doubled in relative form. With $R = 0.500 \pm 0.005$ m and $H = 1.200 \pm 0.010$ m, $V = 0.9425$ m³ and the combined standard uncertainty is about 0.0204 m³. The radius term dominates because its 1% relative uncertainty is multiplied by two.

$$[u(V)/V]^2 = [2u(R)/R]^2 + [u(H)/H]^2$$

Further interpretation

A worked calculation should preserve extra digits internally and round only the reported result. The uncertainty normally receives one or two significant digits, and the estimate is rounded to the same decimal place. Numerical dominance should be checked from squared contributions, not visual impressions. If diameter rather than radius is measured, rewrite the model in diameter before differentiating to avoid hidden factors of two.

Geometry examples also show why dimensional checks are indispensable. The derivative $\partial V/\partial R$ for a cylinder has units of area, so multiplying by $u(R)$ produces a volume contribution. If a calculation yields incompatible units, the sensitivity or input uncertainty has been entered incorrectly. Compare absolute contributions in the output unit and relative contributions as dimensionless fractions. For the cylinder, a 1% radius uncertainty contributes 2% to volume, whereas a 0.83% height uncertainty contributes 0.83%. Their quadrature gives about 2.17%, consistent with $0.0204/0.9425$.

If the cylinder radius and height are measured with the same scale, a scale-factor calibration component is correlated. In the relative model, it enters approximately as $2e+e=3e$ for a common fractional error e , rather than $\sqrt{(4e^2 + e^2)} = \sqrt{5}e$ under independence. This example shows that correlation can change both magnitude and physical interpretation of the final uncertainty.

In practice, report enough intermediate values that another student can reproduce the rectangle or cylinder result. Include nominal geometry, relative input uncertainties, exponent multipliers, individual variance contributions, and the rounded final absolute uncertainty.

11. From diameter to spherical volume

A ping-pong ball exercise connects repeated measurement with propagation. Measure the diameter d at least twenty times under a defined orientation and contact force. Record the raw values, inspect a histogram or dot plot, calculate the sample mean $\bar{d}$ and sample standard deviation s , and use $s/\sqrt{N}$ for the standard uncertainty of the mean when repeatability is the relevant component. Instrument resolution, zero error, deformation, and non-sphericity require separate consideration.

$$V = \pi d^3/6$$

The spherical volume expressed in terms of diameter is $V = \pi d^3/6$. Differentiation gives $dV/dd = \pi d^2/2$ and the relative relation $u(V)/V = 3u(d)/d$. Thus a 0.10% relative uncertainty in mean diameter produces approximately 0.30% relative uncertainty in volume. The factor of three is a direct consequence of the cubic model and does not depend on the unit used for diameter.

$$u(V) \approx (\pi d^2/2)u(d), \quad u(V)/V \approx 3u(d)/d$$

Further interpretation

Repeated readings in one orientation quantify repeatability but not necessarily the uncertainty of the ball's effective diameter. Rotating the ball deliberately samples shape variation; repeating at different contact forces probes deformation. The experimenter must decide whether these effects belong in the measurand definition, require corrections, or contribute additional uncertainty. Increasing N cannot remove an unrecognized geometric or instrumental bias.

A statistically careful ping-pong-ball study separates uncertainty of the mean diameter from variation in the object's shape. If the intended measurand is the volume of that particular nonideal ball, measurements across many orientations characterize geometry rather than mere repeatability. A spherical model may then require an approximation component or a more detailed dimensional model. Caliper contact force can flatten a thin shell, so repeated identical readings can be precise but biased. Photographing orientation, recording force protocol and checking the caliper zero are simple controls that make the propagated volume uncertainty more credible.

The standard deviation of repeated diameter readings should be accompanied by a plot or at least minimum and maximum values. A narrow distribution with two clusters can have the same s as a broad unimodal distribution but indicate a different mechanism. Distribution shape helps decide whether averaging and first-order propagation adequately summarize the observations.

In practice, distinguish the uncertainty of the mean diameter from the scatter of individual diameters. Write both when shape variation is scientifically relevant. The choice determines whether the propagated volume describes an estimated average dimension or variability among orientations.

12. Density and dominant contributions

Density is calculated from mass and volume as $\rho = m/V$. For independent inputs the relative standard uncertainty is the quadrature sum of $u(m)/m$ and $u(V)/V$. If $m = 2.70 \pm 0.01$ g and $V = 1.84 \pm 0.006$ cm³, the nominal density is 1.467 g/cm³. The relative contributions are about 0.370% and 0.326%, leading to $u(\rho)/\rho \approx 0.493\%$ and $u(\rho) \approx 0.0072$ g/cm³.

$$\rho = m/V$$

The largest squared contribution identifies the dominant term, but close contributions may both matter. If the volume is derived from diameter, its uncertainty already contains the factor of three. It would be incorrect to add diameter uncertainty again as a separate line unless it represents another independent effect. Likewise, mass and volume can become correlated if buoyancy correction uses a shared environmental measurement.

$$[u(\rho)/\rho]^2 = [u(m)/m]^2 + [u(V)/V]^2$$

Further interpretation

Dominance is a decision tool. Improving the mass measurement by a factor of ten would reduce the combined uncertainty only to the level set by volume; improving both may expose a previously negligible term such as temperature or model approximation. Recalculate the budget after every proposed improvement rather than assuming that better instrument resolution automatically produces a proportionally better final result.

When contributions are comparable, improvement decisions should consider cost and feasibility as well as ranking. Reducing the mass contribution from 0.370% to 0.100% while leaving volume at 0.326% changes the combined relative uncertainty from 0.493% to about 0.341%; further mass improvement yields little benefit. Reducing volume simultaneously produces a larger gain. This diminishing-return behavior is intrinsic to quadrature. An uncertainty budget therefore supports optimization: choose a target for the combined result, translate it into allowable component levels, and invest effort where the sensitivity-weighted variance can actually be reduced.

An uncertainty-improvement plan should specify a measurable action, such as a finer balance, longer timing interval, better environmental control, or redesigned geometry. Vague intentions to “measure more accurately” cannot be evaluated. Predict the new component value, recompute the budget, and compare the expected reduction with cost and experimental complexity before changing the apparatus.

In practice, compare proposed improvements by recalculating the total, not by comparing instrument specifications alone. The best nominal resolution may be irrelevant if the associated sensitivity is small or if another contribution already sets the uncertainty floor.

13. Integrated measurement workflow and worked applications

A complete propagation analysis starts with the measurand and a defensible measurement equation. List every input, correction and influence quantity, then attach an estimate and standard uncertainty to each. Identify shared references and correlations before calculation. Evaluate partial derivatives at the nominal operating point, form sensitivity-weighted contributions in the units of the output, and combine them with the covariance structure. Check whether local linearization is adequate by testing the full model near the uncertainty limits.

$$X_i \pm u(X_i) \rightarrow Q = f(X_1, \dots, X_n) \rightarrow c_i = \partial Q / \partial X_i \rightarrow u_c(Q)$$

Next build an uncertainty budget and inspect dominance. Verify units, rounding, statistical scale, and the source of every component. Where a result depends on a sum or difference, look for common-mode cancellation or reinforcement. For products, quotients and power laws, relative forms provide rapid checks. For strongly nonlinear models, bounded inputs, asymmetric distributions or large uncertainties, numerical propagation is preferable to a first-order formula.

$$U = k u_c(Q)$$

The final statement should give the estimate, unit, combined standard uncertainty or expanded uncertainty, coverage factor where applicable, and the measurement conditions. It should also identify the model and important correlations. This reporting makes the result reproducible and comparable. The central discipline is to connect each numerical uncertainty to a physical source and each physical source to the sensitivity of the measurand.

Further interpretation

The complete picture is a chain: measurement inputs, mathematical model, sensitivity coefficients, individual uncertainty contributions, covariance-aware combination, and transparent reporting. A weak link anywhere in that chain limits the credibility of the final interval. Propagation does not repair poor data or an incomplete model; it translates the quality and structure of the inputs into the uncertainty of the result.

A final independent check should reproduce the result by a second route whenever consequences justify it. Dimensional analysis checks units; limiting cases check whether the formula behaves sensibly as one uncertainty tends to zero; a numerical perturbation checks derivative signs and magnitudes; and Monte Carlo simulation checks linearization. The reported number should survive these tests within expected approximation error. Documentation should make the chain auditable by another student: raw estimates, uncertainty conversions, correlations, sensitivities, intermediate contributions, rounding rule, and final coverage statement should all be recoverable without guessing.

The final report should distinguish combined standard uncertainty u_c from expanded uncertainty U . If U is quoted, give k and the basis for coverage. If the output distribution is asymmetric, report interval endpoints rather than implying symmetry. These conventions allow another laboratory to compare like with like and reuse the result in a later propagation calculation.

In practice, read the final statement without the preceding calculation. It should still identify the measured quantity, unit, uncertainty convention, coverage information and important conditions. If interpretation depends on hidden assumptions, the report is incomplete.

One additional example joins several rules. Suppose electrical power is calculated as $P=VI$ from independently measured voltage and current. Relative uncertainties combine in quadrature because the model is a product. If the same data-acquisition system introduces a shared gain error, that component is correlated and acts twice in the product: both V and I move upward together. The independent resolution and noise terms still combine in quadrature. Separating these components avoids the common mistake of assigning one overall uncertainty to each input and then treating the totals as independent.

A second application concerns a calibrated sensor result $Q=(S-S_0)/a$, where S is the signal, S_0 is a zero reading, and a is sensitivity. The numerator is a difference, so its absolute uncertainty depends on the uncertainties and covariance of S and S_0 . The division by a then introduces the relative uncertainty of the calibration sensitivity. If S and S_0 are recorded close together on the same drifting channel, positive covariance can reduce the uncertainty of their difference. If they are obtained on different days, that cancellation may not apply. The measurement schedule therefore changes the uncertainty structure even though the algebraic equation remains the same.

These examples demonstrate a general design principle: seek measurement arrangements that reduce important sensitivities or create defensible common-mode cancellation. A ratio can remove a shared scale factor; a differential measurement can suppress a common offset; measuring over a longer interval can reduce relative timing resolution. Such choices improve the model before numerical propagation. They should be documented so that the final uncertainty is understood as a consequence of both data quality and experimental architecture.

For revision, reconstruct any worked result from a blank page. Define the measurand, write the full equation, list input estimates and standard uncertainties, mark correlations, calculate sensitivities, form contributions, combine them, check nonlinear behavior, and report the rounded result. If this sequence can be followed without relying on memorized special cases, the central method of uncertainty propagation has been mastered.