PRECISION ISO • LAB GUIDE

Calculating Measurement Uncertainty

A step-by-step guide for ISO/IEC 17025 laboratories: key terms, contributors, and a fully worked pH uncertainty budget.

Precision ISO

By Dusty Williams, Founder of Precision ISO | precisioniso.com

Every measurement result your lab reports carries doubt. The value you report is only half of the story. The other half is how much confidence that value carries, or how much quantifiable doubt sits around it. So calculating measurement uncertainty is how you put a number on that doubt.

Many labs pursuing ISO/IEC 17025:2017 feel intimidated by uncertainty calculations. Understandably so, because some cases run complex, with many sources of error to quantify. Even so, many cases stay simple. Once you learn the basics, you can build your own uncertainty budget and hold a ready tool for estimating the uncertainty at the time of measurement.

This guide walks the whole process from the ground up. First, it defines the terms, explains why the value matters, and shows you how to find your contributors. Then it covers Type A and Type B sources, distributions and divisors, sensitivity coefficients, and how to combine everything into one number. Throughout, a single pH probe example runs from start to finish, and a fully worked uncertainty budget ties the pieces together.

Who This Guide Is For

This guide serves owners and staff at labs new to ISO/IEC 17025, in both testing and calibration. You need no statistics background beyond averages and standard deviations. Still, it helps to have a method in mind, a reference standard with a certificate, and an instrument you want to characterize.

For accuracy, we anchor everything to ISO/IEC 17025:2017, Clause 7.6, and to the Guide to the Expression of Uncertainty in Measurement, known as the GUM (JCGM 100:2008). Clause numbers point to the 2017 edition. Because standards get revised, verify each one against your own controlled copy before you rely on it.

What Measurement Uncertainty Really Means

Measurement uncertainty is a parameter that describes the spread of values you could reasonably assign to the measurand. In plain terms, it puts a number on the doubt around a result. That doubt comes from several places at once. For example, the instrument, the environment, the operator, and the method all feed into it.

Picture a digital thermometer that reads 100 °F. You read the display and record 100 °F. However, the true temperature could still sit anywhere in a small band around that reading, perhaps 96 °F to 104 °F. Measurement uncertainty is the size of that band. It does not mean you made a mistake. Instead, it means every measurement has limits, and an honest result states them. In this example, then, the uncertainty of the measured value is ± 4 °F. In short, calculating measurement uncertainty is the disciplined way to size that band.

Why ISO 17025 Requires an Uncertainty Estimate

Clause 7.6 of ISO/IEC 17025:2017 sets the rule. First, your lab must identify the contributions to measurement uncertainty for the method it performs. Calibration labs, for instance, must evaluate the uncertainty for every calibration (Clause 7.6.2). Testing labs must evaluate it too, and where the method makes a rigorous evaluation impractical, you make a reasonable estimate from sound theory and practical experience (Clause 7.6.3).

Remember that every result you send a customer carries doubt. For that reason, ISO/IEC 17025 asks you to quantify that doubt and, where it matters, report it, so the customer understands how much confidence your result deserves. In practice, four reasons make this worth the effort.

- **Decision making.** Sound measurements drive sound choices in production, research, and compliance.
- **Traceability.** A stated uncertainty ties your result back to national and international standards.
- **Comparison.** It lets you compare results between labs, or across time within your own lab.

- **Confidence.** It shows customers and assessors that your numbers are reliable.

Start With the Terms You Need

A handful of terms show up again and again once you start building budgets. Skim these tables now, and return whenever a term needs a refresher. Each term links to the section where the guide puts it to work, so you can jump straight there.

Core Measurement Terms

Term	What it means
Measurand	The quantity you set out to measure. In the thermometer example, it was temperature in °F.
Standard (reference material)	The item with a known value that you compare your measurements against. It can be equipment, a solution, or a material. Testing labs often call it a certified reference material or QC sample.
Device under test (DUT)	The device you calibrate or test to see whether it meets its specification.
Measurement error	The result of a measurement minus the true value of the measurand.
Measurement uncertainty	A value that describes the spread of results you could reasonably attribute to the measurand.
Uncertainty source (contributor)	An individual source of error tied to a measurement input. Each one contributes to the combined standard uncertainty.

Budget and Evaluation Terms

Term	What it means
Uncertainty budget	A table that holds every source of uncertainty and its value, then combines them into the total expanded uncertainty. You reuse it to calculate the uncertainty at the time of each measurement.
Type A source	An uncertainty you get from the statistics of repeated measurements, usually the standard deviation of a series.
Type B source	An uncertainty you get from reference material rather than fresh measurements, such as certificates, specifications, or published data.
Divisor	A number that converts a raw error limit into a standard uncertainty, set by the source's distribution.
Standard uncertainty	A single contributor expressed as a standard deviation, the common form every source converts to before you combine them.
Units of measure	The standard amount used to express a quantity, such as °C for temperature or pH for acidity.
Sensitivity coefficient	A multiplier that converts a source in one unit into the units of the measurand.
Relative uncertainty	The standard uncertainty expressed as a ratio of the measured value, often shown as a percentage.

Combining and Reporting Terms

Term	What it means
Combined standard uncertainty	The total standard uncertainty, found by combining every source with the root-sum-of-squares method.
Coverage factor (k)	A multiplier that scales the combined standard uncertainty up to a chosen confidence level. Many labs use k = 2 for roughly 95 percent.
Expanded uncertainty (U)	The final reported uncertainty, the combined standard uncertainty multiplied by the coverage factor.
Repeatability	The random variation when one person measures the same quantity many times with the same standard and method in a short time.
Reproducibility	The variation that appears when conditions change, such as a different operator, instrument, day, or lab.
Guard banding	A quality practice that tightens acceptance limits by a safety margin to account for measurement uncertainty.

For reference, these definitions follow the GUM (JCGM 100:2008). Accreditation bodies expect your method to align with it.

The Four Steps for Calculating Measurement Uncertainty

The cleanest way to approach uncertainty is a step-by-step procedure. In effect, it builds every block you need in the right order. So when you start a budget for a new method or parameter, follow these four steps.

1. **Define the measurement parameter.** Name the measurand and its units. Sometimes you also need the model equation that represents the result.
2. **Build the uncertainty budget.** Identify every contributor in the budget, then decide how to quantify each one.
3. **Collect your measurement data.** Run repeatability and reproducibility studies to see how the method behaves. This is your Type A data.
4. **Calculate the measurement uncertainty.** Populate the budget at the time of calibration or testing, then combine and expand to get the reported value.

Hold this roadmap in mind as you read. Because each step carries real work, we walk through all four using one running example.

The Running Example: A pH Probe

For the rest of the guide, we will calibrate a pH probe against buffer solutions. This case stays simple enough to follow, yet rich enough to show every column of a budget. Once the budget exists, we then use it to calculate the uncertainty at the time of measurement.

The setup is modest. First, we have a pH probe as the DUT and its display. Next, we add a temperature probe to read the buffer temperature, three buffer standards at pH 4, 7, and 10, and beakers to hold the buffers. For this walk-through, we work the numbers at the pH 10.00 point, because contributors like linearity and temperature grow as you move away from pH 7.

Step 1: Define Your Measurement Parameter

Every path to calculating measurement uncertainty starts here, with a clear measurand. First, name the measurand. In our case, it is the pH value the probe indicates when you immerse it in a buffer.

Next, settle the units. Here we measure in pH. The unit pH is defined as the negative log of the hydrogen ion activity, written $pH = -\log(H^+)$. Because our standards are buffer solutions already stated in pH, we keep the whole budget in pH. Converting back to hydrogen ion concentration would add work and change nothing in the result.

Other methods do force a unit change. For example, gravimetric calibration of a pipette weighs dispensed water on a balance, so the reference sits in units of mass. There you need the model equation that converts mass to volume, and then you assign uncertainty sources to each variable in it. Our pH case avoids that, so we can write a simple starting model.

$$pH = pH_B + \delta_{source,1} + \delta_{source,2} + \cdots$$

In words, the display reading equals the nominal buffer value plus a set of small deviations. The probe will not read the buffer value exactly. Rather, sources of error nudge the reading away from the true value, and the model just sums them.

Variable	Meaning
$pH_{(B)}$	The certified value of the buffer solution, for example 10.01 pH
$\delta_{(source,1)}$	The first contributor, or source of uncertainty, in the budget
$\delta_{(source,2)}$	The second contributor in the budget

At this point, we have not named those sources yet. Therefore the next section identifies them, and then we update the model with real labels.

How to Assign Your Uncertainty Sources

Every budget starts from the measurand and its model equation. This is the backbone of calculating measurement uncertainty for any method. Once the model exists, the sources follow from it. Each term is a measured value, a reference value, or a correction, and each one carries uncertainty you must quantify. So you list every physical effect that could push the result away from the true value, then classify each as Type A or Type B.

In effect, the model is the contract. If an effect does not appear in the model, it has no path into the budget, which is why getting the model right is the most important move you make. Moreover, if you cannot point to a source document or a data set that justifies a contributor, an assessor will question it.

Common Uncertainty Sources to Consider

Next, a short list of the sources that appear in most pH budgets. For now we only name them, because we quantify them later.

- **Repeatability (δ_r).** The random scatter you see when you repeat the same measurement under the same conditions. This is your one Type A source in most budgets, and it comes from your own data rather than a document.
- **Reproducibility (δ_rp).** The extra variation that shows up when a condition changes, such as a second operator, a different day, or a different instrument.
- **Buffer solution uncertainty (δ_BCoA).** The uncertainty of the reference value, taken from the certificate of analysis for the buffer. In short, this is your traceability link.
- **Resolution (δ_res).** The smallest change the display can show. It sets a floor on how finely you can read the result, no matter how well the probe is calibrated.
- **Temperature (δ_temp).** A correction for the temperature dependence of buffer pH. For instance, each buffer has a published temperature coefficient, such as about −0.03 pH/°C for a pH 10 buffer.
- **Drift (δ_drift).** How much the probe reading moves between calibrations as the electrode ages. Usually you estimate it from historical calibration data.
- **Linearity (δ_lin).** The probe’s departure from an ideal response across the range. A pH electrode follows the Nernst relation, ideally 59.16 mV/pH at 25 °C, but real slope efficiency never hits 100 percent. As a result, the error grows as you move away from pH 7.
- **Contamination (δ_con).** Buffer contamination, mainly CO_2 absorption in alkaline buffers. With single-use sachets, however, you can set this term to zero.

Complete the Model Equation

Linearity deserves a formula, because it depends on where you measure. So if the meter reports a slope efficiency in percent, model the linearity limit as the departure from ideal times the distance from pH 7.

$$\delta_{lin} = \left(1 - \frac{slope}{100}\right) \times |\, pH_{nominal} - 7.00 \,|$$

With the sources named, the model equation for our pH measurement now reads in full.

$$pH = pH_B + \delta_r + \delta_{rp} + \delta_{BCoA} + \delta_{res} + \delta_{temp} + \delta_{drift} + \delta_{lin} + \delta_{con}$$

That completes Step 1. In summary, we know the measurand, the units, and every source that will populate the budget. Now we build it.

Step 2: Build the Uncertainty Budget

An uncertainty budget is a structured table that lists every source of error, quantifies each as a standard uncertainty, and combines them into one expanded uncertainty. In other words, it serves two jobs at once. First, it is the documented record that shows assessors and customers which sources you considered. Second, it is the reusable tool you copy at the time of each calibration to calculate that specific result. So the budget is where calculating measurement uncertainty becomes concrete.

You reuse it because the values that change from one job to the next are few. For example, repeatability changes, because you measure it on each customer device. Likewise, the reference uncertainty can change when you open a new buffer lot, and environmental values can shift. Meanwhile the framework stays fixed. In short, a budget is a live tool, not a one-time calculation you file away.

Start From a Template, Not a Blank Sheet

Rather than starting from a blank sheet, grab a budget template. A good template already holds the framework and the contributors most budgets share, such as repeatability, resolution, and reference uncertainty. Every budget is still unique, however, so you add or remove rows to match your model equation.

The Completed pH Budget

Here is the completed budget for our pH probe at the pH 10.00 point. Each row is one term from the model equation. After the table, the sections that follow explain every column and show where each pH number comes from.

Source (model term)	Type	Distribution	Divisor	Value	Units	Sens. coef.	Sens. units	Std. uncertainty (pH)	% contribution
Repeatability (δ_r)	A	Normal	1	0.012	pH	1	pH/pH	0.0120	12.8%
Reproducibility (δ_rp)	A	Normal	1	0.008	pH	1	pH/pH	0.0080	5.7%
Buffer CoA (δ_BCoA)	B	Normal (k=2)	2	0.010	pH	1	pH/pH	0.0050	2.2%
Resolution (δ_res)	B	Rectangular	1.732	0.005	pH	1	pH/pH	0.0029	0.7%
Temperature (δ_temp)	B	Rectangular	1.732	0.5	°C	0.03	pH/°C	0.0087	6.7%
Drift (δ_drift)	B	Rectangular	1.732	0.020	pH	1	pH/pH	0.0115	11.9%
Linearity (δ_lin)	B	Rectangular	1.732	0.045	pH	1	pH/pH	0.0260	60.0%
Contamination (δ_con)	B	Rectangular	1.732	0.000	pH	1	pH/pH	0.0000	0.0%

Combined standard uncertainty: **0.0335 pH**. Coverage factor: **k = 2**. Expanded uncertainty: **0.067 pH**.

The Sources of Uncertainty Column

This column lists every contributor, one per row, straight from the model equation. In practice, naming each source after its model term keeps the budget and the equation in lockstep, so an assessor can trace one to the other. For our example, the pH budget carries eight rows, from repeatability through contamination. Contamination sits at zero here because we use fresh single-use buffer. Even so, it stays in the table as a documented decision rather than an omission.

The Type Column

The Type column records how you evaluated each source. Type A comes from data you collect. By contrast, Type B comes from information you look up.

To start, Type A uses the statistics of repeated readings. First, take at least ten measurements of a stable quantity. Then compute the standard deviation. That standard deviation is the Type A standard uncertainty, and it carries the units of the measurement.

$$u_A = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2}$$

Type B, on the other hand, uses external information. A certificate lists a value. A manual quotes a specification. When a source arrives as a plus-or-minus limit with no other detail, then you treat it as a rectangular distribution and divide by the square root of three.

$$u_B = \frac{a}{\sqrt{3}}$$

Here *a* is the stated limit or half-width. In our pH budget, therefore, repeatability and reproducibility are the Type A sources. Everything else is Type B, pulled from certificates, specifications, or historical data.

The Distribution Column

The distribution you assign says what you know about where the true value likely sits inside its range. Generally, the less you know, the more conservative the shape. Conversely, when evidence shows values clustering near the center, the shape tightens.

Type A data usually follows a normal distribution, because repeated readings pile up around a center. Type B limits, meanwhile, usually follow a rectangular distribution, because any value inside the stated range looks equally likely. In our budget, then, the two Type A rows are normal, the buffer certificate is normal because it states a coverage factor, and the remaining Type B rows are rectangular.

The Divisor Column

The divisor comes straight from the distribution, and it converts a raw value into a standard uncertainty. Usefully, a good template fills it in automatically once you pick the distribution. For reference, this short table covers the common cases.

Distribution	When it applies	Divisor
Normal	Type A data, or a certificate that states a coverage factor	1 (or k)
Rectangular (uniform)	Tolerances, resolution, manufacturer specifications	1.732 (√3)
Triangular	Values more likely near the center than the edges	2.449 (√6)
U-shaped	Some cyclic effects, such as temperature swings	1.414 (√2)

A calibration certificate is the case to watch. Often, many certificates report an expanded uncertainty at a coverage factor, usually k = 2. To turn that into a standard uncertainty, then, divide by the stated k rather than by the square root of three. That is why the buffer row in our budget uses a divisor of 2.

The Value Column

The Value column holds the raw magnitude of each source, before any divisor or sensitivity coefficient. In other words, it is the number in its original form, exactly as your data or your reference states it.

The values in our pH budget come from four kinds of place. First, repeatability (0.012 pH) and reproducibility (0.008 pH) come from the standard deviations of our own studies. Second, the buffer value (0.010 pH) is the expanded uncertainty printed on the certificate of analysis. Third, resolution (0.005 pH) is half of the display's last digit, since a reading of 0.01 pH resolution can round by half a count either way. Finally, temperature (0.5 °C) is the half-width of the buffer temperature variation, drift (0.020 pH) comes from historical calibration records, and linearity (0.045 pH) comes from the slope formula in Step 1, using a 98.5 percent slope at pH 10.

The Units Column

This column records the unit that goes with the raw value. It matters because sources arrive in different units. For instance, most of our contributors already sit in pH, so their unit is pH. Temperature is the exception. Its raw value is in °C, because it starts as a temperature variation before any conversion.

Tracking units here is what makes the next columns work. Otherwise, you cannot tell which sources need converting into pH, and the final combination would mix units that do not belong together.

The Sensitivity Coefficient Column

Sensitivity coefficients are the part of calculating measurement uncertainty that newcomers find trickiest, so take this column slowly. A sensitivity coefficient is a multiplier that says how much the result changes when a source changes. When a source already sits in the units of the measurand, its coefficient is simply 1, because no conversion is needed. When a source sits in a different unit, however, the coefficient carries the conversion.

Most rows in our pH budget use a coefficient of 1, since they are already in pH. Temperature, by contrast, is the teaching case. Its value is in °C, and the result is in pH, so we need a coefficient that turns °C into pH. Here the buffer's temperature coefficient does exactly that. For the pH 10 buffer, for example, it runs about 0.03 pH/°C. That coefficient comes from the physics of the buffer, and in harder methods it can come from a partial derivative of the model equation.

The Sensitivity Coefficient Units Column

This column records the unit of the coefficient, and it is the check that your conversion works. Specifically, when you multiply the raw value by the coefficient, the units must cancel to leave the unit of the measurand.

Take temperature again. Its value is in °C and the coefficient is in pH/°C. Multiply them, the °C cancels, and you are left with pH. For every other row, meanwhile, the coefficient is dimensionless, shown as pH/pH, because value and result already share the unit. Therefore, if the units in this column do not resolve to pH, a source has not been converted correctly, and the budget is not ready to combine.

The Standard Uncertainty Column

This step sits at the center of the measurement uncertainty calculation. In effect, this is where every source becomes comparable. The standard uncertainty is the value after you divide by the divisor and multiply by the sensitivity coefficient. Put simply, it states each contributor as a standard deviation in the units of the result.

$$u_i = \frac{\text{value}}{\text{divisor}} \times \text{sensitivity coefficient}$$

To see all three moves at once, work the temperature row. First, start with 0.5 °C. Next, divide by the rectangular divisor of 1.732 to get 0.289 °C. Then multiply by 0.03 pH/°C to land on 0.0087 pH. The resolution row, by comparison, shows the divisor alone: 0.005 pH divided by 1.732 gives 0.0029 pH. The repeatability row needs no conversion at all, so its 0.012 pH passes straight through. Together, these are the values the budget will combine.

The Standard Uncertainty Units Column

Every entry in the standard uncertainty column must share one unit, and this column confirms it. After each source passes through its divisor and coefficient, the result should land in the unit of the measurand. In our budget, accordingly, that unit is pH for every row.

The purpose here is a final check before combining. For instance, if one row still reads °C or millivolts, its conversion failed, and the combination would be invalid. Once every row reads pH, then the sources are ready to merge.

The Relative Uncertainty Column

Relative uncertainty expresses a standard uncertainty as a fraction of the measured value, usually as a percentage. As a result, it lets you compare uncertainty across different measurement values and different methods, because it stays meaningful as the size of the measurement changes.

For our result, the combined standard uncertainty of 0.0335 pH against a value of 10.00 pH gives a relative uncertainty of about 0.34 percent. Admittedly, relative uncertainty carries less weight for pH than for many quantities, because pH is already a logarithmic scale, so most pH reports state the absolute value. Still, we include the column for completeness and for methods where a percentage reads more naturally, such as mass or volume.

The Percent Contribution Column

The percent contribution shows how much each source adds to the combined total. You calculate it by comparing each source's squared standard uncertainty against the squared combined uncertainty.

$$\text{percent contribution}_i = \frac{u_i^2}{u_c^2} \times 100$$

This is the most useful column for improving a method. It shows where calculating measurement uncertainty pays off, by pointing at the source worth fixing. Read down it, and the dominant sources jump out. In our pH budget, for example, linearity accounts for about 60 percent of the total, with repeatability near 13 percent and drift near 12 percent. Resolution and the buffer certificate, by contrast, barely register. So if you wanted a tighter result, a slope recalibration or a slope correction would earn far more than a better buffer. Meanwhile, sources that contribute almost nothing can stay as they are.

Step 3: Collect Your Measurement Data

Estimating measurement uncertainty relies on solid Type A data, so collect it deliberately. The Type A rows in the budget come from studies you run, and Step 3 is where you generate them. Because both repeatability and reproducibility come from repeated measurements, both are Type A.

Repeatability is the scatter between individual readings taken under the same conditions. Same operator, same probe, same buffer, same short time window. First, take at least ten readings of a stable buffer. Then compute the standard deviation of the set. That standard deviation is the repeatability value you enter in the budget. For our example, it came to 0.012 pH.

Reproducibility, by contrast, is the scatter that appears when you change one condition on purpose. A second technician runs the method, or the same technician runs it on a different day. Then you compare the standard deviations of the two sets. For our example, we used 0.008 pH for the additional spread that changed conditions introduced.

Two Ways to Handle Reproducibility

In practice, you have two clean ways to handle reproducibility in the budget. The first folds it into the repeatability study. Instead of ten back-to-back readings, you spread the readings across operators and days, so the resulting standard deviation already includes reproducibility. The second, on the other hand, keeps them separate. First you run a short-term repeatability study, then a longer study that varies a condition, and finally you enter each as its own row. Either approach is defensible, so pick the one your data collection can support.

Above all, record every reading on a standardized form tied to the method. A consistent form keeps your Type A data clean and gives an assessor a clear trail from raw readings to the values in the budget.

Step 4: Calculate the Measurement Uncertainty

Step 4 is where calculating measurement uncertainty produces its final number. With every row in standard form, you combine them. Uncertainties add in quadrature. That is, you square each standard uncertainty, sum the squares, and take the square root. Usefully, this root-sum-of-squares keeps independent sources from double-counting.

$$u_c = \sqrt{u_1^2 + u_2^2 + \cdots + u_n^2}$$

Run our eight rows through it, and the combined standard uncertainty lands at 0.0335 pH.

$$u_c = \sqrt{0.0120^2 + 0.0080^2 + 0.0050^2 + 0.0029^2 + 0.0087^2 + 0.0115^2 + 0.0260^2 + 0^2} = 0.0335 \text{ pH}$$

Finally, expand the result. Multiply the combined standard uncertainty by the coverage factor to reach the expanded uncertainty, which is the value you report.

$$U = k \cdot u_c = 2 \times 0.0335 = 0.067 \text{ pH}$$

At k = 2, the expanded uncertainty gives a range with about 95 percent confidence. So you report the calibration point as pH 10.00 ± 0.07 at a coverage factor of k = 2. The percent contribution column then tells you where to aim if that range is too wide. Because linearity dominates, improving the probe's slope would tighten the result far more than any other change.

Using the Budget at the Time of Calibration

Here is the payoff. Once the budget exists for the method, you reuse it for every job rather than starting over. That reuse is what makes calculating measurement uncertainty practical day to day. The framework stays fixed, and only the values that actually change get updated.

Calibration labs frame this through two ideas. The first is the CMC, or calibration and measurement capability. A CMC is the best-case uncertainty a lab can achieve with its own reference equipment under controlled conditions, and it sets the floor a lab advertises. By contrast, the uncertainty you report for a specific customer device usually runs higher, because you measure the customer's item under real conditions with its own repeatability.

The second idea is traceability. ISO/IEC 17025 Clause 6.5 requires your results to trace to the SI through an unbroken chain of calibrations, and the buffer certificate is one link in that chain. So when you open a new buffer lot, its certificate uncertainty replaces the old value in the budget. In turn, that keeps the reported result traceable to the standards actually used on the day. Update the repeatability for the customer's probe, refresh the buffer uncertainty, recombine with the root-sum-of-squares, and the budget hands you the uncertainty for that exact calibration.

Finally, many labs put the finished uncertainty to work through guard banding. In short, they tighten the acceptance limits by part of the expanded uncertainty before they judge a device pass or fail. As a result, guard banding lowers the risk of accepting a device that actually sits outside tolerance.

Putting It Into Practice

Calculating measurement uncertainty gets manageable once you break it into four steps. First, define the measurand and write its model equation. Second, identify every contributor and build the budget around the model. Third, collect Type A data through repeatability and reproducibility, and pull Type B values from your certificates and specifications. Finally, convert each source to a standard uncertainty, combine them with the root-sum-of-squares, apply a coverage factor, and read the expanded uncertainty off the bottom of the table. From there, reuse that budget for each job, and update only the values that change.

Follow that pattern, and the process holds across almost any method your lab runs. So if you want a head start, the Precision ISO uncertainty budget template already carries this framework and the common contributors, so you fill in your model and your data rather than build from scratch. Pair it with the uncertainty of measurement SOP, tailor both to your scope, and you turn the steps above into a repeatable procedure your assessor will recognize as sound.

A note on accuracy: This guide teaches; it does not give legal, regulatory, or accreditation advice, and it guarantees no audit outcome. Clause numbers reference ISO/IEC 17025:2017, and the uncertainty methods follow the GUM (JCGM 100:2008). The pH budget values are illustrative and chosen to demonstrate the method. Before you rely on any clause number, coefficient, or method, verify it against your organization's controlled copy of the standard and your own data.