Lab 0: Error Analysis Guide

Aaron Dunbrack

Summer 2018

Contents

The most important things to read are sections 1, 2.1, 2.2, 3.1, and 4; these contain the essential details to understanding error. The rest of the main text will help you understand better, but is not necessary. If you really want to understand certain concepts, you are likely to find the appendices helpful.

1 Introduction

Error is one of the most important features of doing quantitative analysis.

Suppose you were doing a lab to check whether $F = ma$. Your measured value of F is 0.99N, and your measured value of ma is 1.02N.

Is this good? Bad? Who knows? I mean, it's close, so perhaps you can say it agrees. If, on the other hand, you did a super-precise measurement, then maybe you can actually say that $F = ma$ is only *approximately* true.

Whether or not you confirm or deny Newton's second law depends on the precision of your experiment! So, in order to do our quantiative analysis properly and come to a conclusion as to whether it is correct or not, we need to know our errors.

1.1 Random vs. Systemic Error

There are two kinds of error that one deals with in the lab: **random error** and **systemic error**.

Random error, per the name, are statistical effects that push your data every which way from the "true" result, with no consistency to its effects. These will result in your data being right on average, but any individual data point will be off. This results in *uncertainty*. Random error won't make you wrong (usually), but will limit how much you can say. Data that is low in random error is called *precise*.

Systemic errors are effects that come from aspects of your system that you have not considered. These result in *bias*, as well as more complicated effects. That is to say, systemic error will "push" your data in some *consistent, reliable* way. It won't (necessarily) increase the scatter of your data, so it won't make your data look "all over the place." Rather, it takes the nice straight line your data should make, and turns it into a line of a different slope (or another curve altogether).

1.2 Quantifying Error: Absolute vs. Relative Error

There are two general ways to talk about error: absolute error and relative error.

If you have a quantity X that you are measuring, the **absolute error** in X, denoted as ΔX, is simply the amount by which you are (or could be) off. For instance, if you measure something to be 5.89m, then you might reasonably

say that your uncertainty is 1cm - that is to say, you could have been off by a centimeter.

You could also report this same uncertainty as a **relative error**, denoted as $\Delta_{\text{rel}}X$. This is when you compare the size of your error to the size of the original quantity.[1] The formula for relative error is:

$$\Delta_{\text{rel}}X = \frac{\Delta X}{|X|} \tag{1}$$

Thus, in the above example, your 1cm uncertainty on your 5.89m measurement would turn into a relative error of 0.0016. In other words, the absolute uncertainty on your measurement is 0.0016 times your original measurement.

You may be more familiar with the idea of relative error from the closely related *percent error*, which is just 100 times larger.

Both relative and absolute errors have advantages and disadvantages as methods of quantifying your error.

Absolute error is straightforward and easier to understand. They are also necessary for direct comparison: if I ask you whether 5 ± 1mm agrees with a theoretical value of 4.8cm, that's easy to tell; if I ask whether 5mm with a relative uncertainty of 0.2 agrees with the same theoretical value, that's not nearly as obvious. As such, when you state a quantity in your lab report, it should typically be written $X \pm \Delta X$ - that is to say, you talk about uncertainties with the absolute error when reporting them. (This will also be the kind of error that these manuals will be referring to if they don't explicitly say "relative error.")

Relative errors give a sense of how "big" an error is in a general sense. That is to say, if I tell you I have a 5% error, you have an immediate sense of how much I'm off by; if I tell you I have a 1cm error, whether that's big or small depends on what I'm measuring. In particular, relative errors can be compared across different sources of errors: if one quantity has a relative uncertainty of 0.2 and another of 0.003, it's easy to see that the contribution of the second to your final amount of error is negligible. If I tell you one has an error of 1cm and the other an error of 0.2s, you have no sense of which is the more relevant contribution. Therefore, relative errors are the better things to look at if you want to compare errors of different varieties.

Absolute errors will have the same units as the original quantity,[2] so a time measured in seconds will have an uncertainty measured in seconds, etc.; there-

[1] If your "original quantity" is zero (or, more generally, you are trying to measure something which should be equal to zero), then you shouldn't use relative error. As a more concise rule of thumb: if your relative error is larger than 1, you should just use absolute error.

[2] Actually, you could change the modifier if you wanted to - for instance, $1.2\text{cm} \pm 1\text{mm}$ - but there's rarely a reason to do so.

fore, they will only be unitless if the original quantity is unitless. Relative errors are *always* unitless.[3]

2 Random Error

2.1 Basic Uncertainties

In order to get a sense of what our uncertainty will be in our final result, we first should know what uncertainty comes from our measurements themselves. We have two ways to do this.

The first method is the statistically rigorous way. We first do the measurement of the same quantity, X, a some number N of (statistically-independent) times. Call these different measurements x_1, x_2, ..., x_N. Then, we take the average of the results, $\overline{X}$:

$$\overline{X} = \frac{x_1 + x_2 + \ldots + x_N}{N} \tag{2}$$

This is then our best estimate of the relevant parameter. The uncertainty of this average is:

$$\Delta\overline{X} = \sqrt{\frac{(x_1 - \overline{X})^2 + (x_2 - \overline{X})^2 + \ldots + (x_N - \overline{X})^2}{N(N-1)}} \tag{3}$$

If you want the uncertainty of an individual measurement (say, of x_1), you instead use the formula:

$$\Delta x_1 = \sqrt{\frac{(x_1 - \overline{X})^2 + (x_2 - \overline{X})^2 + \ldots + (x_N - \overline{X})^2}{N-1}} \tag{4}$$

This is handy if you're trying to use your measurements to estimate the uncertainty of some future measurement.

Note that this means the two quantities are closely related:

$$\Delta x_1 = \sqrt{N}\Delta\overline{X} \tag{5}$$

The second method is to just estimate how much you think you could plausibly be off by. This is the method we will use most often in this lab, simply because it is faster and easier to do. Some good rules of thumb for doing this:

- For a digital device which directly outputs a reading (like a digital scale), you can take the uncertainty to be given by the last digit the device outputs. For instance, if it reads 4.18, then the uncertainty would be 0.01 (in whatever units the scale reads).

[3] This is a part of what makes them directly comparable.

  - Exception: if the reading from the device is significantly fluctuating, you should take the fluctuations to be the typical uncertainty in the reading, instead.
  - Some devices have multiple settings which will change what that last digit is. You should use the last digit from the setting you were using when you made the measurement. (As such, you should also, as a rule, use the setting with the smallest last digit that you can, because this will minimize your uncertainty from that device.)

- For something you are reading by eye, a typical rule is "half the smallest division." This applies *unless* you think there is a reason it should be larger; say, you're measuring a long distance, and are having a hard time getting things to line up right.

- For pre-built devices (like a set of mass weights), there are usually a standard set of tolerances given by the manufacturer, which tells you how much to expect they could be off by (at most). Such tolerances should be specified in your lab manual, where relevant.

At the end of the day, though, these error estimations are just that: estimates. Professional physicists spend good deal of time getting statistically-rigorous uncertainty measurements; for this lab, these rough estimates will suffice for our purposes.

2.2 Uncertainty Propagation Formulas

Once you have these basic measurement uncertainties, how do you determine the uncertainty in your final results? Somehow, you need to turn uncertainties in these quantities into uncertainties in your final measurements.

The first step to doing so is to identify how you calculated your quantity. For instance, if you used $F = ma$ to calculate a mass, you'd take $m = F/a$, and note that the process for calculating m (in the specific case you're considering) is to divide F by a.

The second step is to break this calculation into a sequence of elementary steps. $m = F/a$ is already there, but if you're calculating centripetal force from the formula $F = mv^2/r$, then you might have the steps:

1. Take v and square it.
2. Multiply by m.
3. Divide by r.

Each of those corresponds to a simple step (one that will have a formula we'll give in a minute).

The final step is to propagate error for each step. To do so, we'll need some special formulas that tell us how this is done; for instance, how to determine the error in $A + B$ depending on the errors in A and B individually. The following formulae[4] indicate how to do this, given experimental quantities A and B and constants (i.e., numbers without error) c and n:

$$\Delta(A + c) = \Delta A \tag{6}$$

$$\Delta(cA) = |c|\Delta A \tag{7}$$

$$\Delta(A \pm B) = \sqrt{(\Delta A)^2 + (\Delta B)^2} \tag{8}$$

$$\Delta(A \times B) = |AB|\sqrt{\left(\frac{\Delta A}{A}\right)^2 + \left(\frac{\Delta B}{B}\right)^2} \tag{9}$$

$$\Delta\left(\frac{A}{B}\right) = \left|\frac{A}{B}\right|\sqrt{\left(\frac{\Delta A}{A}\right)^2 + \left(\frac{\Delta B}{B}\right)^2} \tag{10}$$

$$\Delta(A^n) = |n||A|^{n-1}\Delta A \tag{11}$$

$$\Delta\left(\frac{1}{A}\right) = \frac{\Delta A}{A^2} \tag{12}$$

Some of these can be more conveniently expressed in terms of relative errors:

$$\Delta_{\text{rel}}(cA) = \Delta_{\text{rel}}A \tag{13}$$

$$\Delta_{\text{rel}}(A \times B) = \Delta_{\text{rel}}\left(\frac{A}{B}\right) = \sqrt{(\Delta_{\text{rel}}A)^2 + (\Delta_{\text{rel}}B)^2} \tag{14}$$

$$\Delta_{\text{rel}}(A^n) = |n|\Delta_{\text{rel}}A \tag{15}$$

$$\Delta_{\text{rel}}\left(\frac{1}{A}\right) = \Delta_{\text{rel}}A \tag{16}$$

You may recognize that several of them have the same squaring, adding, and square-rooting pattern of different quantities. If we take the quantity $\sqrt{x^2 + y^2}$, that is called "adding x and y *in quadrature*." Using that terminology, we can express the above formulas in plain English:

- When you multiply by a constant, the relative error remains the same (13).
- When you add or subtract two quantities, you add the absolute errors in quadrature (8).
- When you multiply or divide two quantities, you add the relative errors in quadrature (14).

[4]Some are redundant. (6) can be obtained from (8), (7) from (9), and (12) from (10) by setting one error to zero. Also, (12) is a special case of (11) for $n = -1$. However, you will prbably find them easier to use in their simplified form.

- When you take a power of a quantity, you multiply the relative error by the absolute value of the exponent (15).
- When you take one over a quantity, you keep the same relative error (16).

If you're adding and subtracting a bunch of things, or multiplying and dividing a bunch of things, the formulas (8) and (14) extend basically like you would expect:

$$\Delta(A_1 \pm A_2 \pm A_3 \pm \ldots) = \sqrt{(\Delta A_1)^2 + (\Delta A_2)^2 + (\Delta A_3)^2 + \ldots} \tag{17}$$

$$\Delta_{\text{rel}}\left(\frac{A_1 \times A_2 \times \ldots}{B_1 \times B_2 \times \ldots}\right) = \sqrt{(\Delta_{\text{rel}} A_1)^2 + (\Delta_{\text{rel}} A_2)^2 + \ldots + (\Delta_{\text{rel}} B_1)^2 + (\Delta_{\text{rel}} B_2)^2 + \ldots} \tag{18}$$

You may find the above formulas more convenient than breaking such things into a whole bunch of individual multiplications.

2.3 Where do those formulas come from?

Having formulas is nice, but it's even nicer to understand those formulas. While this won't be a complete explanation, nor will it describe how to take errors of more complicated formulas (see Appendix A for that), it should suffice to give you some idea of why the formulas are the way that they are.

Let's begin with adding two things. Let's suppose we make two measurements of the quantities A and B. Let's say A and B are the actual values, but we measure $A + \delta A$ and $B + \delta B$ - that is to say, δA and δB are the amounts we are off by. (We don't actually know these numbers, but they do exist.)

Now, when we take $A + B$, we can ask how far are we off there; i.e., what is $\delta(A + B)$:

$$\delta(A + B) = ((A + \delta A) + (B + \delta B)) - (A + B) = \delta A + \delta B$$

This is true for an individual measurement, but that's not the same thing as (8). Why?

Well, because sometimes, δA and δB will have opposite signs. Sometimes, our errors will cancel. Just adding uncertainties doesn't account for that possibility.

To cut short a longer explanation,[5] if you directly add (or subtract) errors in two different quantities, like the formula above, you add the corresponding uncertainties in quadrature. This accounts for the possibility of cancelling,

[5] We technically quantify uncertainty in the following way: we take the average value of the *square* of our error, and square root that: $\Delta x = \sqrt{\overline{(\delta x)^2}}$. If we square both sides of $\delta(A + B) = \delta A + \delta B$, we get

$$\delta(A + B)^2 = \delta A^2 + \delta B^2 + 2\delta A \delta B$$

mathematically speaking.

Alright. If we have that, then we do get (8). Now, let's look at the multiplication formula, (9).

Let's do the same procedure: take $A+\delta A$ and $B+\delta B$ as our measurements, and calculate $\delta(AB)$:

$$\delta(AB) = (A+\delta A)(B+\delta B) - AB = A\delta B + B\delta A + \delta A\delta B$$

Now, we're going to make a simplification: we're going to assume our errors are relatively "small."[6] If δA and δB are both small, then their product is *really* small, and so we can ignore it. Therefore, approximately, we find:

$$\delta(AB) \approx A\delta B + B\delta A$$

Alright, now we're going apply the same principle as above: we have two errors in different quantities added together, so we're going to add in quadrature to get the uncertainty propagation formula. This gives us:

$$\Delta(AB) = \sqrt{(A\Delta B)^2 + (B\Delta A)^2}$$

Dividing the entire equation by AB and putting that inside the square-root then gives us:

$$\frac{\Delta(AB)}{AB} = \sqrt{\frac{A^2(\Delta B)^2}{(AB)^2} + \frac{B^2(\Delta A)^2}{(AB)^2}}$$

This is the same as equation (14) for multiplication, from which you can get (9) by multiplying by AB.

The rest follow from similar rules - if you really want to get where they come from, you can work through some of them yourself. The multiplication by a constant rule is straightforward. The power rule you can do for positive integer powers using the binomial theorem (and remember that errors squared, cubed, etc. are "very small" and you can ignore them).

For non-integer powers and for the division formula, you'll need some calculus, unfortunately - see Appendix A.3 if you want to see how to get those.

Now, take the average value of both sides, and you find

$$(\Delta(A+B))^2 = (\Delta A)^2 + (\Delta B)^2 + 2\overline{\delta A \delta B}$$

We now make an assumption, known as *statistical independence*: we assume $\overline{\delta A \delta B} = 0$, which gives us (8). Practically speaking, this assumption is true if A and B depend entirely on separate measurements; to make things easy, in this lab, we'll just always pretend that it's true. If it is not true, the quantities A and B are said to be *correlated*, and that makes error analysis significantly harder.

[6] More mathematically precisely: we assume that the relative errors are much less than 1.

2.4 Uncertainty Propagation Examples

2.4.1 Simple Example

Let me begin with a simple example. Let's suppose you have measured a mass m and the acceleration of that mass a, and want to calculate the net force using $F = ma$.

You calculated F by multiplying two things together. Therefore, you want to use the multiplication formula, (9) (or, equivalently, (14)). This gives you:

$$\Delta F = ma\sqrt{\left(\frac{\Delta m}{m}\right)^2 + \left(\frac{\Delta a}{a}\right)^2}$$

This gives you ΔF *if* you calculate F from m and a. If you calculate it some other way (say, $F = mg$), you should propagate through *that* formula instead - even if the two forces are supposed to be equal.

Now, let's use the same formula a little differently. Let's suppose you had measured force and acceleration, and calculated mass.

Then, the first step would be to identify how you calculated mass: you used the formula $m = F/a$. That means you have to use the division rule, now.

Fortunately, it's similarly simple, and therefore we get:

$$\Delta m = \frac{F}{a}\sqrt{\left(\frac{\Delta F}{F}\right)^2 + \left(\frac{\Delta a}{a}\right)^2}$$

This could also have been expressed as:

$$\Delta_{\text{rel}} m = \sqrt{(\Delta_{\text{rel}} F)^2 + (\Delta_{\text{rel}} a)^2}$$

2.4.2 Less Simple Example

Now, let's do a less trivial example, one that requires a little more processing and multiple steps.

Let's suppose you want to calculate the uncertainty in a centripetal force F, where F is calculated according to the formula $F = mv^2/r$. That is to say, you measured m, v and r (and their uncertainties), and want to calculate ΔF.

When we calculated F, we did a series of steps:

- Take v and square it.
- Multiply v^2 by m.
- Divide mv^2 by r.

This means we have three steps that we need to do separately. Each will require its own error propagation step.

For the first step, we are putting v to a power. Thus, we use the power propagation formula (11) to get:

$$\Delta(v^2) = 2v\Delta v$$

Now, we multiply by m, to get:

$$\Delta(mv^2) = mv^2\sqrt{\left(\frac{\Delta m}{m}\right)^2 + \left(\frac{\Delta(v^2)}{v^2}\right)^2} = mv^2\sqrt{\left(\frac{\Delta m}{m}\right)^2 + \left(\frac{2v\Delta v}{v^2}\right)^2}$$

$$\Delta(mv^2) = mv^2\sqrt{\left(\frac{\Delta m}{m}\right)^2 + \left(\frac{2\Delta v}{v}\right)^2}$$

Finally, we divide by r, giving us:

$$\Delta F = \Delta\left(\frac{mv^2}{r}\right) = \frac{mv^2}{r}\sqrt{\left(\frac{\Delta(mv^2)}{mv^2}\right)^2 + \left(\frac{\Delta r}{r}\right)^2}$$

$$\Delta F = \frac{mv^2}{r}\sqrt{\left(\frac{\Delta m}{m}\right)^2 + \left(\frac{2\Delta v}{v}\right)^2 + \left(\frac{\Delta r}{r}\right)^2}$$

This example shows off how equation (18) is useful, becuase it combines the last two steps into one. It also shows off how useful relative errors are if no additions are involved. Let's calculate the same thing again using those methods.

First, we convert all our absolute errors into relative errors using (1).

Then, we calculate $\Delta_{\text{rel}}(v^2) = 2\Delta_{\text{rel}}v$ using (15).

Finally, the rest is multiplying/dividing multiple things together, so we use (18). This gets us the rest of the way:

$$\Delta_{\text{rel}}F = \sqrt{(\Delta_{\text{rel}}m)^2 + (\Delta_{\text{rel}}(v^2))^2 + (\Delta_{\text{rel}}r)^2}$$

$$\Delta F = F\sqrt{(\Delta_{\text{rel}}m)^2 + (2\Delta_{\text{rel}}v)^2 + (\Delta_{\text{rel}}r)^2}$$

That gets you to the same place with much less algebra!

Hopefully, the above examples can serve as guides for your own error propagation procedures.

3 Systemic Error

3.1 Isolating Systemic Errors

Systemic errors are a lot harder to deal with than random errors. This is for two main reasons.

Firstly, systemic errors cannot be reduced by running more trials of the same experiment. They will be present every time, with (roughly) the same effect, and will not average out to zero.

For example, let's suppose the lab was doing an experiment to measure $g = 9.8\text{m/s}^2$. Random errors would make one group measure 10m/s^2, another measure 9.7m/s^2, etc., but on average the whole class would probably get something close to 9.8m/s^2. For a systemic error, this isn't true: it might make everyone in the class measure $g = 9.9\text{m/s}^2$, so that even when you average you don't get the right answer.

Systemic errors *can* be reduced by running multiple entirely *different* experiments to find the same quantity, but that requires coming up with multiple ways to test the same thing. Experimental physicists frequently have multiple independent experiments testing the same thing for this exact reason.

Secondly, systemic errors can't be *measured* by taking the same experiment multiple times and observing how much your results move around, so there's no analogy for formulas (3) and (4). The only way to understand the size of your systemic errors (from your experiment alone) is by understanding what causes them.

Therefore, identifying systemic errors requires knowing the physics behind your system, and understanding your experimental apparatus really well. You have to see where some other physical effect could produce an impact on your results, and what impact that would be.

3.2 Determining the Effects of Systemic Errors

There are two aspects of any systemic error that you would want to know: what kind of effect it has, and how large this effect is.

Both can be determined by re-doing calculations after you include some estimates of any new physical parameters. That's time-consuming, but it can be done.

Often, however, we can make things easier by making quick approximations, so that's what we'll do here. This means we won't actually calculate the quantitative effects of our systemic errors. Instead, we'll get a qualitative sense of

what they're doing and how important they are.

Many systemic errors can be simply stated as making you over- or underestimate some quantity. You can determine, based on that, whether your final quantity would be too small or too large.

For instance, let's suppose you have a mass that tells you it's 1kg, but it's really dirty. Perhaps that dirt is adding extra weight to the mass. This means that the total mass your experiment would be using would be bigger than 1kg, so you underestimated the value of your mass.

Suppose your final goal was to measure the force exerted on this mass, based on the acceleration it undergoes, using $F = ma$. Well, if you underestimate m, then you would also underestimate F.

These kinds of reasoning - "I underestimated this, so I underestimated that" (or, if something is in the denominator, "I underestimated this, so I *over*estimated that") - can tell you what kind of impact your systemic error will have on your final result.[7]

The second thing you can do is estimate size of these errors. This allows you to determine whether or not such an error is important.

To a rough approximation, you can get a sense of their size with the following rule of thumb: the relative error in your final quantity is about the same as the relative error in your original quantity.

So, taking the above example of a dirty mass again, the relative error resulting from the dirt is the mass of the dirt divided by the mass of the original mass (1kg). If you have an estimate of how much the dirt might weigh (say, 1g), you can convert this into a relative error on the mass (in this case, 0.001).

You can then take this relative error as an order-of-magnitude estimate of the relative error in F, so if F was 7.2N, then your absolute error in F would be roughly .0072N. Assuming your random error in F is significantly higher than this, you can neglect this error as insignificant.

[7]Sometimes, your systemic errors don't just increase or decrease everything, or they do so in a complicated way. For instance, let's suppose you fit a line $y = kx$ to a supposedly-linear graph. However, there's a systemic error that makes a corrected formula $y = kx + ax^2$, for some small number a. Will not including this systemic error make your measurement bigger or smaller? There's no easy way to tell, unfortunately - you would have to go through the more detailed analysis. Systemic errors are generally hard!

4 Things that are NOT errors

There are many things which one would call "errors" in colloquial language, but are not errors in a *scientific* sense. These include:

- Calculation errors: If you think your calculations have a problem, fix the problem.
- Human error: again - if there's a problem, fix it. If it's too late to fix when you realize that you made a significant mistake (say, you realize it after you went home), that is worth mentioning, but it's not a "source of error" in the sense we've been discussing - it's just a mistake.
- Random deviations from procedure that don't affect results: Let's say the lab tells you to take a length of 0.3m for some piece of rope, in order to have the experiment work well. If you have a 31cm rope, and then do your calculations with $L = 31\text{cm}$, then that's not an error, even if it's a deviation from procedure. Usually, if you do all calculations in a way that is accurate to your numbers, then your results should still give you the right result, even if you deviate from procedure, so this won't result in errors!

A Mathematical Details

A.1 What is an uncertainty really?

A proper definition of what an "uncertainty" really is was briefly explained in footnote 5, but here's an elaboration in a bit more detail.

Suppose we are taking a measurement of some quantity X. What, then, really, is ΔX?

Well, this quantity X will have some average value. Let's call this value μ (not $\overline{X}$, for reasons to be explained in the next section).

For any given measurement of X, we can then talk about $X - \mu$ for that measurement. However, on average, this difference will be zero - X will be above μ about as often as it will be below. So that doesn't help.

It might seem reasonable to talk about the average distance of X from μ - that is to say, the average value of $|X - \mu|$. One can do this, but it's very messy, and doesn't follow nice rules.

A better quantity to look at is the *root mean square* distance of X from μ (with "root mean square" often abbreviated as *rms*). This is given by taking the square, then the mean, then the root, yielding the formula:

$$\Delta X = \text{RMS Distance} = \sqrt{\overline{(X-\mu)^2}} \tag{19}$$

This turns out to have much nicer properties than our previous calculation, and so we use this as our definition of ΔX.

Of course, your uncertainty estimates "by eye" won't be this exact quantity, but they'll be vaguely similar, and so we'll use the same formulas that this follows.

A.2 Why $N-1$?

Now that (in the previous section) we defined what the uncertainty actually is, we can see some of the motivation behind the formula for the uncertainty in a single measurement, (4). In fact, this is *almost* the same as the calculation we already have, except for that little -1 in the denominator. What gives?

Well, this is where the distinction between μ and $\overline{X}$ comes into play. The quantity we're measuring has some actual value, μ, that would be the average of infinitely many measurements. However, we only took finitely measurements, and the average of our actual data is $\overline{X}$, which will not necessarily equal μ. So let's go through this carefully (warning: lots of algebra incoming).

Let's suppose we've taken a bunch of measurements x_1, x_2, ..., x_n. This data has an average of $\overline{X}$, and a theoretical value of μ. Now, let's calculate the uncertainty ΔX based on the definition of ΔX given in equation (19). It's easiest to calculate from the square:

$$(\Delta X)^2 = \overline{(X-\mu)^2} \approx \frac{\sum_{i=1}^{N}(x_i-\mu)^2}{N}$$

We'll use a trick of adding and subtracting $\overline{X}$ in each square now:

$$(\Delta X)^2 = \frac{\sum_{i=1}^{N}((x_i-\overline{X})+(\overline{X}-\mu))^2}{N}$$

FOILing out that square gives:

$$\begin{aligned}(\Delta X)^2 &= \frac{\sum_{i=1}^{N}\left[(x_i-\overline{X})^2+(\overline{X}-\mu)^2+2(x_i-\overline{X})(\overline{X}-\mu)\right]}{N} \\ &= \frac{\sum_{i=1}^{N}(x_i-\overline{X})^2+\sum_{i=1}^{N}(\overline{X}-\mu)^2+\sum_{i=1}^{N}2(x_i-\overline{X})(\overline{X}-\mu)}{N} \\ &= \sum_{i=1}^{N}\frac{(x_i-\overline{X})^2}{N}+\frac{N(\overline{X}-\mu)^2}{N}+\frac{\sum_{i=1}^{N}2(x_i-\overline{X})(\overline{X}-\mu)}{N}\end{aligned}$$

$$= \sum_{i=1}^{N} \frac{(x_i - \overline{X})^2}{N} + (\Delta\overline{X})^2 + \cancel{\frac{2(\overline{X} - \mu)\sum_{i=1}^{N}(x_i - \overline{X})}{N}}$$

That last term is zero because the sum vanishes, by definition of $\overline{X}$:

$$\sum_{i=1}^{N}(x_i - \overline{X}) = \left(\sum_{i=1}^{N} x_i\right) - N\overline{X} = N\overline{X} - N\overline{X} = 0$$

Now: we have this nice formula, but don't know either ΔX nor $\Delta\overline{X}$. However, from the general sum formula (17), we *do* know that equation (5) must hold (I leave the algebra to you to do, if you are unconvinced).

Therefore, we can replace $(\Delta\overline{X})^2$ with $\frac{(\Delta X)^2}{N}$ and finish off our calculation:

$$(\Delta X)^2 = \frac{1}{N}\sum_{i=1}^{N}(x_i - \overline{X})^2 + \frac{(\Delta X)^2}{N}$$

$$N(\Delta X)^2 = \sum_{i=1}^{N}(x_i - \overline{X})^2 + (\Delta X)^2$$

$$N(\Delta X)^2 - (\Delta X)^2 = \sum_{i=1}^{N}(x_i - \overline{X})^2$$

$$(N-1)(\Delta X)^2 = \sum_{i=1}^{N}(x_i - \overline{X})^2$$

$$(\Delta X)^2 = \frac{\sum_{i=1}^{N}(x_i - \overline{X})^2}{N-1}$$

$$\Delta X = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \overline{X})^2}{N-1}}$$

And there you have it - equation (4), with the $N-1$ factor and all. From this, equation (3) is a result of (5), so now we have all of our basic formulas for uncertainty.

TL;DR: The -1 is a result of the fact that the average of our data, $\overline{X}$, is not necessarily the same as the actual true value, μ, when we said the uncertainty was our average distance from the true value.

A.3 A Fuller Explanation of Error Propagation Formulas

(Warning: much of this section uses calculus. If you don't know calculus, you are unlikely to find this helpful.)

Using the above definition and some basic notions from calculus, we can get a better idea of where the error propagation formulas come from, and figure out how to make our own.

Following the notation of section 2.3, we'll denote by δX the difference in our measurement from the true value of X, and similarly for δY when we need it.

A.3.1 1 Variable Error Formulas

Let's begin with a function of one variable, $f(X)$. The actual error of our measurement of $f(X)$ is given by:

$$\delta[f(X)] = f(X + \delta X) - f(X)$$

Well, that's not very helpful in general. Fortunately, we have some calculus tools on our side. Let's make the assumption that δX is small, and Taylor-expand it:

$$f(X + \delta X) \approx f(X) + f'(X)\delta X + \dots$$

Now, here's the key approximation we'll make: we'll drop the ... from the above expression. We assume that the second-order errors are negligible, and so use the first order errors. This is good enough for most purposes, although if you were really careful you might want to do something more specific.

Therefore, we can approximate the earlier formula as:

$$\delta[f(X)] \approx [f(X) + f'(X)\delta X] - f(X) = f'(X)\delta X$$

Now, we just plug into the definition for ΔX, (19), and find the result:

$$\Delta[f(X)] = \sqrt{\overline{(f'(X)\delta X)^2}} = \sqrt{(f'(X))^2\overline{(\delta X)^2}} = |f'(X)|\Delta X$$

Thus, our fully general one-variable error formula is simply:

$$\Delta[f(X)] = |f'(X)|\Delta X \tag{20}$$

From there, using the fact that if $f(X) = X^n$, the derivative is $f'(X) = nX^{n-1}$, you can directly get the power error propagation formula (11).

A list of error propagation formulas for a variety of functions can be found in Table 1.

Calculation	Error Formula
X^n	$\lvert n\rvert\lvert X\rvert^{n-1}\Delta X$
e^X	$e^X\Delta X$
$\ln(X)$	$\frac{\Delta X}{\lvert X\rvert}$
$\sin(X)$	$\lvert\cos(X)\rvert\Delta X^*$
$\cos(X)$	$\lvert\sin(X)\rvert\Delta X^*$

Table 1: A list of more exotic error functions.
*NOTE: For trig functions, ΔX must be measured in radians

A.3.2 2+ Variable Error Formulas

Now, let's consider functions of two variables, like addition and multiplication: $f(X,Y) = X+Y$ and so on. We're going to need a bit of multivariable calculus here, but conceptually, it's the same. We begin with a Taylor-expansion:

$$f(X+\delta X, Y+\delta Y) \approx f(X,Y) + \frac{\partial f}{\partial X}\delta X + \frac{\partial f}{\partial Y}\delta Y + \dots$$

Therefore, again dropping the higher-order terms, we similarly find the expression:

$$\delta[f(X,Y)] \approx \frac{\partial f}{\partial X}\delta X + \frac{\partial f}{\partial Y}\delta Y$$

Again, we plug into our formula (19) and (after a bit of algebra) see what we get:

$$\Delta[f(X,Y)] = \sqrt{\left(\frac{\partial f}{\partial X}\right)^2 \Delta X^2 + \left(\frac{\partial f}{\partial Y}\right)^2 \Delta Y^2 + 2\left(\frac{\partial f}{\partial X}\right)\left(\frac{\partial f}{\partial Y}\right)\overline{\delta X \delta Y}}$$

Now, we assume that X and Y are uncorrelated, so $\overline{\delta X \delta Y} = 0$. Otherwise, you need that number, too, to do your calculation properly; that number is called the *covariance* of X and Y.

Under the assumption that that's zero, though, we get our formula in two variables:

$$\Delta[f(X,Y)] = \sqrt{\left(\frac{\partial f}{\partial X}\right)^2 \Delta X^2 + \left(\frac{\partial f}{\partial Y}\right)^2 \Delta Y^2 + 2\left(\frac{\partial f}{\partial X}\right)\left(\frac{\partial f}{\partial Y}\right)\overline{\delta X \delta Y}} \quad (21)$$

If you know enough multivariable calculus to understand the partial derivatives, you should now be able to derive the addition and multiplication rules with that formula.

The derivation and result also extend straightforwardly to three or more variables: you add more partial derivatives times matching errors in quadrature, under the assumption that all errors are pairwise-uncorrelated.

A.3.3 Statistical Independence

It's worth noting what happens when the statistical independence breaks down.

For instance, consider if we were find the error in $\frac{X}{X+Y}$. Clearly, the numerator and denominator are not independent - when the numerator increases, the denominator increases too, so an error in X produces less error in this ratio than it otherwise would, because the error in the numerator and denominator will cancel.

There are two solutions to this. The first is to try to get your equation into a form where the error can be evaluated properly, and each variable only shows up once. For instance, the above fraction is equal to $\frac{1}{1+Y/X}$. There, if X and Y are uncorrelated, we don't have a computational issue. This is neat when you can do it; unfortunately, it's more of a trick than an easy-to-use procedure.

The other is to use the full calculus formulas. This is how to do such analyses properly, if you ever find yourself in such a situation, and is - if you know calculus - often easier than trying to find a trick procedure.

A.3.4 Higher Order Effects

What if our Taylor series approximation breaks down? For instance, what if $f'(X) = 0$ (or is small enough that the second-order correction matters)?

A number of complicated things can happen, led by our error propagation formulas being no longer approximately correct. That aside, though, more happens than just having different formulas: random errors have biasing effects.

Let's take a simple example: suppose we want to measure X^2 when the true value, μ, is equal to 0. Now: *any* sort of X error is going to result in a positive estimate for X^2. This means that a really small random error in X doesn't make X^2 go up and down by the same amount - it always makes it go up.

This makes our random error in X into a systemic one in X^2. That's quite the messy issue!

Fortunately, it rarely comes up, but if you're at a place where $f'(X) = 0$ (in any variable X), you'll need to look very closely at your experiment to make sure these systemic errors don't matter. (Fortunately, the "errors-are-small" assumption usually means these are still very small, and you can neglect them

altogether... but error aren't always small enough that you can do that!)

Let's take a more concrete example: a warped meter stick. The meter stick can be warped up or down; let's say the amount it is warped up or down is an amount X.

It doesn't matter if X is positive or negative; either way, the meter stick is going to take a curved path between your two points rather than a straight one, and so it will always be longer than the actual distance. You are going to consistently overestimate the actual distance, regardless of which direction it was warped.

B Significant Figures

In previous science classes, a great deal of attention may have been paid attention to significant figures (how many decimal places you should work to) and rules for manipulating them, trailing zeros, etc.

In general, significant figures serve two roles:

1. A stand-in for uncertainties. Significant figure rules are easy to follow - easier than uncertainty propagation, certainly - and serve as a functional replacement for them if actual uncertainties are not being calculated.

2. Aesthetic reasons. There's no practical reason to go to the twentieth decimal point if your error picks up at the second; those last decimal places aren't important, because your error is larger than them.

Since we'll be using actual uncertainties, the first reason isn't relevant. Therefore, significant figures are, in this class, largely a matter of taste. In the same sense as English grammar: you are allowed to break the rules if you have a good reason for doing so. The rules serve as a guide to good communication, but should be ignored where they inhibit it.

As a general rule of thumb, work to three significant figures in your quantities in this class, with a number on your uncertainty to match. You can go further if you have a very-precisely measured quantity, or two if your error is very large (say, bigger than 10%). Generally, though, three sig figs will be both enough for rounding errors to not matter and few enough that it won't be aesthetically abhorrent.