A good training in physics requires a solid understanding of uncertainty[1].

We can imagine a model (which might be wrong) as a function $f(a, b)$ of some input parameters a and b. Those input parameters might be off. The result will be an output with $\pm$ uncertainty. The uncertainty can't account for the model being wrong, but it *can* account for the input parameters having some estimated uncertainty.

Zeroth order analysis You can simply estimate extreme values for the plausible error on the inputs. Then run your calculation with all permutations of $\pm$, and choose the highest and lowest values of your bound.

This approach is *pessimistic*, as it gives you a worst case scenario. Sometimes this can be good, but often it's unrealistic. Error estimation is not about ensuring that the true answer is definitely within your bounds, but rather it is about trying to make a best guess of the likely error.

A better estimation We can refine the zeroth-order estimate of uncertainty propagation by considering that each input parameter comes from a distribution of possible values.

There is a most likely value (the **mode** of the distribution), and a mean value, as well as a **standard deviation** of the distribution. The true value is typically withing one standard deviation of the mean 68% of the time.

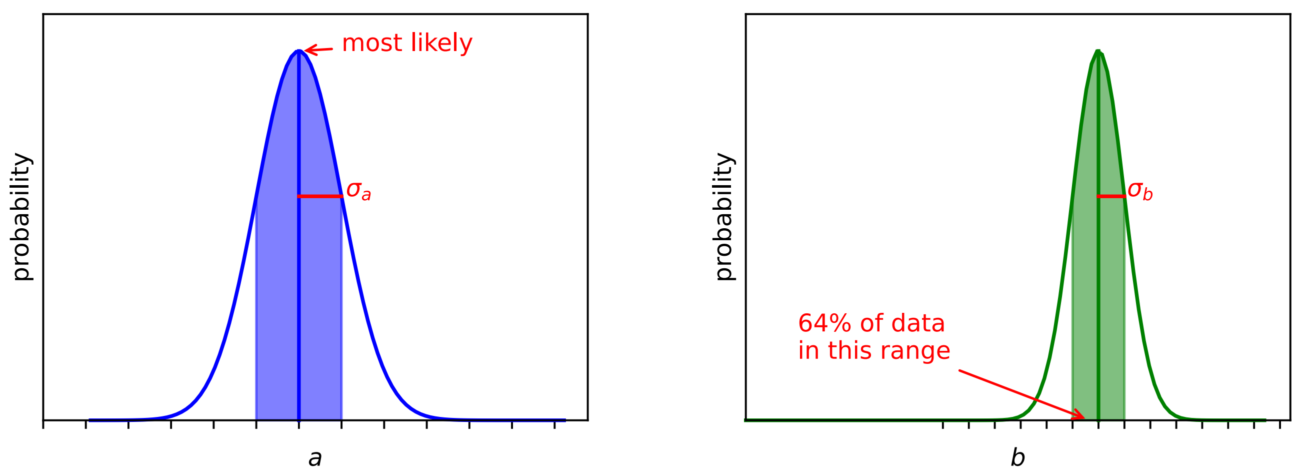

Figure 1: Distributions of values for parameters a and b. The shaded region represents the range $\pm\sigma$ around the mean value, where σ is the standard deviation.

I can imagine choosing a and b randomly from these probability distributions, and then constructing a probability distribution of $f(a, b)$. I could then find the standard deviation of this distribution, and call that my uncertainty in f.

[1]Note: error and uncertainty technically mean different things, but I will use them interchangeably to mean uncertainty, because I'm sloppy. There is a significant uncertainty in the meaning I attach to the word "error." And that is mostly a joke

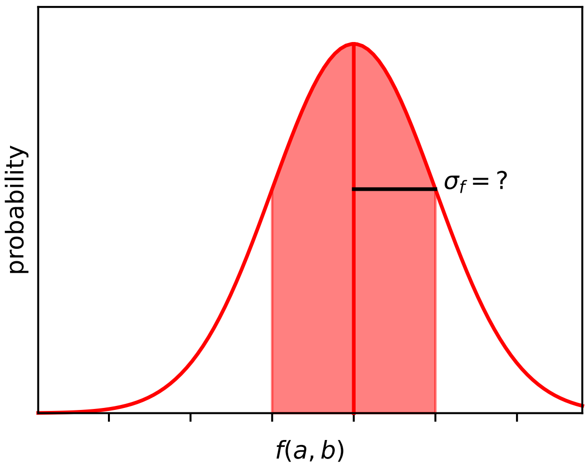

Figure 2: Distributions of the output f obtained by combining the input distributions for parameters a and b.

Provided the distributions of a and b are not correlated with each other, and the function $f(a, b)$ is does not deviate very much from linear behavior over the distributions of a and b, we can use the following formula to find the uncertainty in f, which we will call σ_f:

$$\sigma_f = \sqrt{\left(\frac{\partial f}{\partial a}\right)^2 \sigma_a^2 + \left(\frac{\partial f}{\partial b}\right)^2 \sigma_b^2} \tag{1}$$

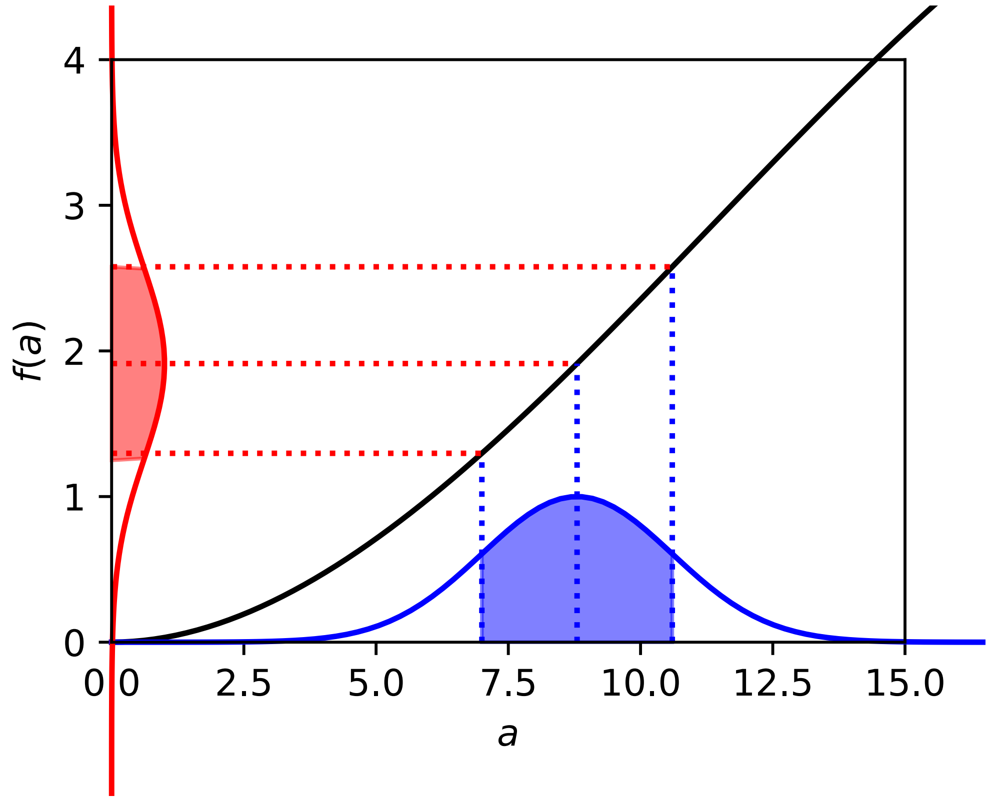

Figure 3: An illustration of why the derivative of f with respect to its input parameters affects its uncertainty.

You won't often need to use this partial derivative relationship, because most often we just add, multiply, subtract, and divide, and you can just use a little table:

function	uncertainty
$f = a + b$	$\sigma_f = \sqrt{\sigma_a^2 + \sigma_b^2}$
$f = a - b$	$\sigma_f = \sqrt{\sigma_a^2 + \sigma_b^2}$
$f = a \cdot b$	$\frac{\sigma_f}{f} = \sqrt{\left(\frac{\sigma_a}{a}\right)^2 + \left(\frac{\sigma_b}{b}\right)^2}$
$f = \frac{a}{b}$	$\frac{\sigma_f}{f} = \sqrt{\left(\frac{\sigma_a}{a}\right)^2 + \left(\frac{\sigma_b}{b}\right)^2}$