How good is your answer?

Suppose you compute the answer to some problem. You know that your answer is only an approximation, but **is it a good approximation?** Please give this some thought: You don't know the true answer, so how can you know your error?

Let's model the problem so that the answer is the value of a function: Let x represent all of the information we need to describe our problem, let y be the answer, and let $y = f(x)$ describe the relationship between the problem and the solution (see, e.g. the *implicit function theorem*). If we call our approximate solution $\bar{y}$, then our question becomes:

$$\bar{y} \approx y \ ? \tag{Q}$$

One of the best ways of answering (Q) is to first ask a weaker question:

$$\text{Is there an } \bar{x} \approx x \text{ such that } \bar{y} = f(\bar{x})? \tag{q}$$

This question, (q), can often be answered by analyzing the method independently of the specific data (backward error analysis), or by computing some extra quantity (a residual) that gives a bound on the distance between x and $\bar{x}$. This question is fundamentally a question about the computation itself, in fact, if $\bar{y}$ is the solution to a nearby problem, then we say that it is a *backward stable result* or that the computation was *backward stable.*

Unfortunately, a backward stable result is not necessarily a good result. We want to know if our answer is good, and backward stability is not enough. Recall that the problem is *well conditioned* if

$$\bar{x} \approx x \Longrightarrow f(\bar{x}) \approx f(x)? \tag{p}$$

Fortunately, a backward stable computation applied to a well conditioned problem always gives a good answer ($p \wedge q \Rightarrow Q$). Here is how it comes together: We have that $\bar{y} = f(\bar{x})$ and $y = f(x)$, so a well conditioned f means $\bar{x} \approx x \Longrightarrow \bar{y} \approx y$, and backward stable $\bar{y}$ means $\bar{x} \approx x$, therefore $\bar{y} \approx y$.

Let's look at this in the context of finding a zero of a function: "find p^* such that $g(p^*) = 0$". If you like, this can be described as a function evaluation by "evaluate $f(0)$, where $f(y) = g^{-1}(y)$. Suppose we have a computed approximation to p^*, say $\bar{p}$. If $g(\bar{p}) = \epsilon \approx 0$, then $\bar{p}$ is the exact solution to the nearby problem "find p such that $g(p) - \epsilon = 0$", and so is a backward stable approximation. This does not guarantee that $\bar{p} \approx p^*$, for $\bar{p} - p^* = g^{-1}(\epsilon) - g^{-1}(0)$, which may be large. When can this be large? Only when our problem is illconditioned.

When a problem is very illconditioned, there is no method that can guarantee a good solution. Even the magic method that takes the finite precision approximation to the problem, finds its exact solution, and then rounds it to finite precision, is only giving the approximate solution to a nearby problem.