Graduate students taking STAT 5088 should complete some "grad problems" over the course of the semester. I expect there will be around 5-8 of these problems, and you'll need to do a good job on half of them.

---

The goal of this problem is to create a simulation that uses simulation to confirm the theoretical posterior in the Normal-Normal model.

In group labs for the Beta-Binomial and Gamma-Poisson models, we were able to write a simulation that:

- Sampled a large number of parameter values ($\pi$ or $\lambda$) from the prior distribution.
- For each sampled parameter, generate data according to the data model.
- Restrict to just the parameter values whose data matches hypothetical "observed" data.
- Plot the density of those parameters, add the theoretical posterior probability density to the plot, and observe that they agree.

With the Normal-Normal model, this breaks down because the observed data is of the form $y_1, \ldots, y_n$ and we care about the mean $\bar{y}$. However, the simulated data means will **never** match the observed $\bar{y}$. To get around this problem, instead of matching the observed value exactly we ask that it be close, say within 0.1 of the observed $\bar{y}$.

For this problem, write the simulation so that it will work with arbitrary choices of $\sigma$ (the known sample sd) and $\tau$ (the prior sd). If you can, allow for a variable data size $n$, although I'm ok if you just use $n = 4$.

Test your simulation with various values of $\tau$, $\sigma$ and observed $\bar{y}$.