1. Let $X_1, X_2, \ldots, X_n$ be IID Uniform$(0, \theta)$ random variables.

   (a) Show that
   $$\frac{X_{(n)}}{\theta}$$
   is a pivotal quantity, where $X_{(n)}$ is the maximum of the sample.

   (b) Find the distribution of that pivotal, and thus construct a 95% confidence interval for θ.

2. Read the *New Yorker* article *The Truth Wears Off* at the url

   `https://www.newyorker.com/magazine/2010/12/13/the-truth-wears-off`

   (a) Let X is a Binomial$(20, p)$ random variable. Suppose that we test the null hypothesis $p = 1/2$ using the rejection region $X \leq 5$ or $X \geq 15$. What is the significance level?

   (b) Let $\hat{p} = X/20$, the usual estimator for p. Supposing that the true value of p is .55, what is $\mathbb{E}(\hat{p} \mid H_0 \text{ } rejected)$? You may compute numerically or estimate by simulation. I recommend at least 1000 trials if you choose the simulation.

   (c) What is the connection, if any, between this result and the *New Yorker* article?

3. Five experiments were conducted independently by five different investigators to test the null hypothesis that a given treatment was no better than a placebo. All the experiments were well designed, randomized double-blind with pairs of subjects matched for relevant factors such as age, sex, and severity of disease. Each pair of subjects was evaluated, and a "success" was recorded if the treatment subject was responding better than the control subject within each pair. The results were:

   ```
   1. 10 successes in 15 trials.
   2. 14 successes in 22 trials.
   3. 21 successes in 38 trials.
   4. 10 successes in 18 trials.
   5. 16 successes in 25 trials.
   ```

   (a) Using the **R** *binom.test()* function, test the null hypothesis $p = 1/2$ for each experiment using the 5% significance level (i.e. reject $H_0: \ p = 1/2$ if the p-value is less than .05).

   (b) Ordinarily it is a bad idea to interpret a failure to reject the null hypothesis as proof that the null hypothesis is true. Here we have repeated failures to reject the null hypothesis. Does the repeated failure to detect a statistically significant difference provide strong evidence that the null hypothesis is in fact true? Explain your reasoning.

4. Read the influential 2005 paper by John Ioannidis
   Why Most Published Research Findings Are False at the URL

   `journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.0020124`

   Suppose that the Journal of Empirical Stuff publishes only those articles submitted which contain a result statistically significant at the 5% level. Authors, knowing the policy, don't bother to submit (or even to waste time writing up) papers on analyses that don't meet this standard. Since prospective authors are desperate to publish, lest their academic career languish, they conduct numerous studies of dubious merit. Thus, the proportion of true null hypotheses in their studies is about 75%. In their rush to publish, they are also typically working with tests of low power (due to small samples, poor design, etc.). To keep things simple, assume that the power for those tests where the null hypothesis is false is only 20%. What fraction of the published papers contain correct results? In other words, what fraction of the published statistically significant results correspond to cases where the null hypothesis was in fact false? Hint: Let the event R be the rejection of a null hypothesis, then the significance level is $\mathbb{P}(R|H_0)$. We want $\mathbb{P}(H_0|R)$ and $\mathbb{P}(H_1|R)$.