STATS 200: Homework 2

Due Wednesday, October 12, at 5PM

1. **Monte Carlo integration (based on Rice 5.21 and 5.22).** For a given function $f : [a, b] \to \mathbb{R}$, suppose we wish to numerically evaluate

$$I(f) = \int_a^b f(x)dx.$$

One method is the following: Let g be a PDF of a continuous random variable taking values in $[a, b]$, and generate independent random draws $X_1, \ldots, X_n$ from g. Then estimate $I(f)$ by

$$\hat{I}_n(f) = \frac{1}{n} \sum_{i=1}^{n} \frac{f(X_i)}{g(X_i)}.$$

(a) Show that $\mathbb{E}[\hat{I}_n(f)] = I(f)$. Assuming that $\text{Var}[f(X_i)/g(X_i)] < \infty$, show that $\hat{I}_n(f) \to I(f)$ in probability as $n \to \infty$.

(b) Derive a formula for $\text{Var}[\hat{I}_n(f)]$. Show that for some $c_n \in \mathbb{R}$, $c_n(\hat{I}_n(f) - I(f)) \to \mathcal{N}(0, 1)$ in distribution as $n \to \infty$. What is c_n?

(c) Consider concretely the problem of evaluating

$$I(f) = \int_0^1 \cos(2\pi x)dx.$$

Let g be the PDF of the uniform distribution on $[0, 1]$, and consider the above estimate $\hat{I}_n(f)$ using 1000 IID samples from g. Using your result from part (b), compute approximately the probability $\mathbb{P}[|\hat{I}_n(f) - I(f)| > 0.05]$.

(d) Propose a different PDF g on $[0, 1]$ such that the resulting estimate $\hat{I}_n(f)$ using g is more accurate than the estimate in part (c). Compute approximately the probability $\mathbb{P}[|\hat{I}_n(f) - I(f)| > 0.05]$ for your choice of g.

2. **Continuous mapping (Rice 5.7).** Prove the first part of the Continuous Mapping Theorem stated in Lecture 4: If random variables $\{X_n\}_{n=1}^{\infty}$ converge in probability to $c \in \mathbb{R}$ (as $n \to \infty$), and $g : \mathbb{R} \to \mathbb{R}$ is continuous, then $\{g(X_n)\}_{n=1}^{\infty}$ converge in probability to $g(c)$.

3. **Testing gender ratios (based on Rice 9.45).** In a classical genetics study, Geissler (1889) studied hospital records in Saxony and compiled data on the gender ratio. The following table shows the number of male children in 6115 families having 12 children:

Number of male children	Number of families
0	7
1	45
2	181
3	478
4	829
5	1112
6	1343
7	1033
8	670
9	286
10	104
11	24
12	3

Let $X_1, \ldots, X_{6115}$ denote the number of male children in these 6115 families.

(a) Suggest two reasonable test statistics T_1 and T_2 for testing the null hypothesis

$$H_0 : X_1, \ldots, X_{6115} \overset{IID}{\sim} \text{Binomial}(12, 0.5).$$

(This is intentionally open-ended; try to pick T_1 and T_2 to "target" different possible alternatives to the above null.) Compute the values of T_1 and T_2 for the above data.

(b) Perform a simulation to simulate the null distributions of T_1 and T_2. (For example: Simulate 6115 independent samples $X_1, \ldots, X_{6115}$ from Binomial$(12, 0.5)$, and compute T_1 on this sample. Do this 1000 times to obtain 1000 simulated values of T_1. Do the same for T_2.) Plot the histograms of the simulated null distributions of T_1 and T_2. Using your simulated values, compute approximate p-values of the hypothesis tests based on T_1 and T_2, for the above data. For either of your tests, can you reject H_0 at significance level $\alpha = 0.05$? (Include both your code and the histograms with your homework submission.)

In addition to what was reviewed in Question 6 of Homework 1, the following commands may be helpful if you are doing this in R:

To generate a numeric vector of 6115 independent Binomial$(12, 0.5)$ samples:

```
X = rbinom(6115, 12, 0.5)
```

To count the number of elements of a numeric vector `X` that are equal to, say, 8:

```
count = length(which(X==8))
```

To count the number of elements of a numeric vector `Y` that are, say, greater than 0.1:

```
count = length(which(Y>0.1))
```

(c) In this example, why might the null hypothesis H_0 not hold? (Please answer this question regardless of your findings in part (b).)

4. **Most-powerful test for the normal variance.**
(a) For data $X_1, \ldots, X_n \in \mathbb{R}$ and two fixed and known values $\sigma_0^2 < \sigma_1^2$, consider the following testing problem:

$$H_0 : X_1, \ldots, X_n \overset{IID}{\sim} \mathcal{N}(0, \sigma_0^2)$$
$$H_1 : X_1, \ldots, X_n \overset{IID}{\sim} \mathcal{N}(0, \sigma_1^2)$$

What is the most powerful test for testing H_0 versus H_1 at level α? Letting $\chi_n^2(\alpha)$ denote the $1-\alpha$ quantile of the χ_n^2 distribution, describe explicitly both the test statistic T and the rejection region for this test.

(b) What is the distribution of this test statistic T under the alternative hypothesis H_1? Using this result, and letting F denote the CDF of the χ_n^2 distribution, provide a formula for the power of this test against H_1 in terms of $\chi_n^2(\alpha)$, σ_0^2, σ_1^2, and F. Keeping σ_0^2 fixed, what happens to the power of the test as σ_1^2 increases to ∞?

5. **Testing a uniform null (Rice 9.20).** Consider two probability density functions on $[0, 1]$: $f_0(x) = 1$ and $f_1(x) = 2x$. Among all tests of the null hypothesis $H_0 : X \sim f_0(x)$ versus the alternative $H_1 : X \sim f_1(x)$ with significance level $\alpha = 0.10$, how large can the power possibly be?