

Full Name: _____

Student No: _____

Grade:

Read before you start:

- Please make sure you write your full name and student number.
- The exam consists of 10 questions, some with multiple parts. All answers require justifications.
- During the exam, you are allowed to use your phone to access the app for probability distributions (and possibly a calculator app), but the internet access of your phone must be switched off.
- The duration of the exam is 2 hours.

You can use the remainder of this page as scratch paper.

1. (5 points)

(a) State the law of large numbers as precisely as you can, and give an example illustrating it.

Solution: The law of large numbers states that the average of a large number of independent, identically distributed RVs is approximately the same as the expected value of each of them.

More specifically, suppose that $X_1, X_2, \dots, X_n$ are i.i.d. with $\mathbb{E}[X_k] = \mu$. If n is large, then

$$\frac{X_1 + X_2 + \dots + X_n}{n} \approx \mu.$$

The approximation can be quantified using Chebyshev's inequality. Namely, let us denote the above average by $\bar{X}_n$. If $\text{Var}[X_k] = \sigma^2$, then for every margin of error $\varepsilon > 0$, the probability that $\bar{X}_n = \mu \pm \varepsilon$ is at least $1 - \frac{1}{n}(\sigma^2/\varepsilon^2)$, hence it is close to 1 when n is large.

As an example, suppose we repeat flipping a fair coin. Let

$$X_k = \begin{cases} 1 & \text{if the } k\text{th flip comes up heads,} \\ 0 & \text{if the } k\text{th flip comes up tails.} \end{cases}$$

Then, $\bar{X}_n$ is the proportion of heads within the first n flips. The law of large number simply formulates the familiar statement that when n is large, the proportion of heads is roughly $1/2$.

(b) Is your answer to part (a) correct?

2. (5 points)

(a) State the central limit theorem as precisely as you can, and give an example illustrating it.

Solution: The central limit theorem is a mathematical formulation of the phenomenon that the average (or sum) of a large number of independent, identically distributed RVs is approximately normally distributed.

Specifically, let $X_1, X_2, \dots, X_n$ be i.i.d. RVs with $\mathbb{E}[X_k] = \mu$ and $\text{SD}[X_k] = \sigma$. The average

$$\bar{X}_n := \frac{X_1 + X_2 + \dots + X_n}{n}$$

has mean μ and standard deviation $\sigma/\sqrt{n}$. According to the central limit theorem, when n is large, $\bar{X}_n$ is approximately normally distributed with the latter parameters.

More precisely, by shifting and scaling $\bar{X}_n$ appropriately, we obtain a RV

$$Z_n := \frac{1}{\sigma/\sqrt{n}}(\bar{X}_n - \mu)$$

which has mean 0 and standard deviation 1. The precise statement of the central limit theorem is that, for every z ,

$$\lim_{n \rightarrow \infty} \mathbb{P}(Z_n \leq z) = \Phi(z)$$

where Φ is the cdf of the standard normal distribution.

As an example, consider again the experiment of repeatedly flipping a fair coin, and let

$$X_k = \begin{cases} 1 & \text{if the } k\text{th flip comes up heads,} \\ 0 & \text{if the } k\text{th flip comes up tails.} \end{cases}$$

Note that $\mathbb{E}[X_k] = 1/2$ and $\text{SD}[X_k] = 1/2$. According to the central limit theorem, the average $\bar{X}_n$ is approximately normal with mean $1/2$ and standard deviation $\frac{1}{2\sqrt{n}}$.

(b) Is your answer to part (a) correct?

3. (18 points) Determine which of the following statements is True and which is False. In each case, give a short justification.

TRUE Saying that the grade of a student is in the upper 25th percentile of her class means that her grade is higher than or equal to at least $\frac{3}{4}$ of her classmates.

Solution: Saying that the grade of a student is in the upper 25th percentile of her class is another way to say that her grade belongs to the top 25% of her class. Put differently, it is the same as saying that the grade of the student is at least as large as the 75th percentile of her class.

TRUE If A and B are two events in a probability model, then necessarily $\mathbb{P}(A \cup B) + \mathbb{P}(A \cap B) = \mathbb{P}(A) + \mathbb{P}(B)$.

Solution: This is simply a rearrangement of the inclusion-exclusion principle, and can be seen true using a Venn diagram.

FALSE If A and B are two disjoint (= mutually exclusive) events in a probability model, then necessarily $\mathbb{P}(A \cap B) = \mathbb{P}(A) \mathbb{P}(B)$.

Solution: The identity $\mathbb{P}(A \cap B) = \mathbb{P}(A) \mathbb{P}(B)$ means A and B are independent. Two mutually exclusive events cannot be independent (unless one of them has probability 0).

FALSE The probability density function of every continuous random variable is a continuous function.

Solution: A continuous random variable is a random variable which has a probability density function. The probability density function can be anything as long as it is non-negative and has integral 1. As an example, the pdf of an exponential random variable is not continuous at 0.


FALSE If X and Y are two normal random variables, then the distribution of $X + Y$ is necessarily normal.

Solution: As a trivial counter-example, suppose X is standard normal and $Y = -X$. Then, $X + Y = 0$, which is not normal. There are countless other counter-examples. The statement would have been true if there was an extra assumption of independence between X and Y .


TRUE If X and Y are two independent random variables with $X \sim \text{NBin}(r = 3, p = 0.7)$ and $Y \sim \text{NBin}(r = 5, p = 0.7)$, then $X + Y \sim \text{NBin}(r = 8, p = 0.7)$.

Solution: This is similar to question 2 in Quiz 4. Interpret X and Y in the context of a Bernoulli process with parameter $p = 0.7$. The RV X can be thought of as the number of tails before the 3rd head, and Y as the number of tails between the 3rd head and the 8th head. Then, $X + Y$ is nothing but the number of tails before the 8th head.


4. (5 points) Identify which of the following histograms has the largest and which has the smallest standard deviation.



(i)



(ii)



(iii)

Give a short justification.

Solution: The histogram (ii) has the largest standard deviation, and the histogram (iii) has the smallest standard deviation. The standard deviation of the histogram (iii) is (strictly) in between the other two standard deviations.

All three histograms have a mean of about 5. In the third histogram, all the data values are relatively close to the mean, whereas in the second histogram, all the data values are relatively far from the mean. In the first histogram, the data values are spread over those that are close to the mean and those that are far from it.

5. (5 points) A random sample of 50 female orangutans is used to calculate the 95% confidence interval (33 Kg, 41 Kg) for the mean weight of female orangutans in Brunei. Which of the following statements is a correct interpretation of the latter statement?
- Of the entire population of female orangutans in Brunei, 95% weigh between 33 Kg and 41 Kg.
 - The chance that the mean weight of female orangutans in Brunei is between 33 Kg and 41 Kg is 95%.
 - The chance that a randomly picked female orangutan from Brunei weighs between 33 Kg and 41 Kg is 95%.
 - If we were to repeat taking samples from the female orangutans in Brunei and calculate a 95% confidence interval based on each sample, then about 95% of these intervals would contain the true population mean.

Justify!

Solution: The last option is correct.

The confidence interval is obtained by an interval estimator (L, R) . By definition, the confidence level of an interval estimator is the probability that the random interval (L, R) includes the true value of the parameter it is meant to estimate.

In this case, the confidence level is 95%, which means, in the random experiment of taking a sample from the female orangutans in Brunei, the probability that the resulting confidence interval includes the true mean weight of female orangutans in Brunei is 95%. The last option in the question simply interprets this probability in terms of repeated experiments.

6. (5 points) A coin has probability $p \neq 1/2$ of coming up heads. We flip the coin twice. Which of the following two events is more probable?

$A := \langle \text{the same face comes up in both flips} \rangle$

$B := \langle \text{difference faces come up} \rangle$

Solution: The event A is more probable than the event B .

The sample space is $\Omega := \{HH, HT, TH, TT\}$ and the measure of probabilities is given by

$$\mathbb{P}(HH) = p^2, \quad \mathbb{P}(HT) = p(1-p), \quad \mathbb{P}(TH) = (1-p)p, \quad \mathbb{P}(TT) = (1-p)^2.$$

In this model,

$$A = \{HH, TT\},$$

$$B = \{HT, TH\},$$

hence

$$\mathbb{P}(A) = p^2 + (1-p)^2,$$

$$\mathbb{P}(B) = 2p(1-p).$$

Now, note that

$$\mathbb{P}(A) - \mathbb{P}(B) = p^2 + (1-p)^2 - 2p(1-p) = (p - (1-p))^2 \geq 0$$

with equality if and only if $p = 1-p$. Since $p \neq 1/2$, we have $p \neq 1-p$, hence $\mathbb{P}(A) > \mathbb{P}(B)$.

7. (8 points) Among a sample of 4000 newborns, 2080 turned out to be male. Find a 99% confidence interval for the probability p that a newborn is male.¹

Solution: For a given newborn, we consider a Bernoulli variable

$$X := \begin{cases} 1 & \text{if the newborn is male,} \\ 0 & \text{if the newborn is female.} \end{cases}$$

We would like to estimate $p = \mathbb{P}(X = 1)$.

A sample of 4000 newborns corresponds to a sequence $X_1, X_2, \dots, X_{4000}$ independent copies of X . An unbiased point estimator for p is the sample proportion

$$\hat{p} = \frac{X_1 + X_2 + \dots + X_{4000}}{4000}.$$

By the central limit theorem, $\hat{p}$ is approximately normally distributed. Furthermore,

$$\begin{aligned} \mathbb{E}[\hat{p}] &= p, & \mathbb{V}\text{ar}[\hat{p}] &= \frac{p(1-p)}{4000}, \\ \text{SE}[\hat{p}] &= \sqrt{\frac{p(1-p)}{4000}}. \end{aligned}$$

Thus, for every $z > 0$,

$$\mathbb{P}\left(\underbrace{p - z\sqrt{\frac{p(1-p)}{4000}} < \hat{p} < p + z\sqrt{\frac{p(1-p)}{4000}}}_{\star}\right) \approx \Phi(z) - \Phi(-z) = 2\Phi(z) - 1,$$

¹This problem is from the book *Probability, Random Variables, and Stochastic Processes* (4th ed.) by A. Papoulis and S. U. Pillai.

where Φ denotes the cdf of the standard normal distribution. The inequality $\circledast$ implicitly identifies an interval estimator for p with confidence level $2\Phi(z) - 1$. Setting the confidence level $2\Phi(z) - 1 = 0.99$ and using the app for the (standard) normal distribution, we obtain $z \approx 2.57583$.

In the specific sample at hand, 2080 out of 4000 newborns have been male, hence the observed value of $\hat{p}$ is $2080/4000 = 0.52$. Thus, a 99% confidence interval based on this sample is the interval implicitly determined by the inequality


$$|0.52 - p| < 2.57583 \sqrt{\frac{p(1-p)}{4000}}.$$

Squaring both sides, we find that the equivalent inequality

$$(0.52 - p)^2 < \frac{2.57583^2}{4000} p(1-p),$$

and rearranging and simplifying gives

$$\underbrace{1.0016587p^2 - 1.0416587p + 0.2704}_{f(p)} < 0$$



Solving the quadratic equation $f(p) = 0$, we obtain two solutions $p_1 \approx 0.499636$ and $p_2 \approx 0.540297$. Every p between p_1 and p_2 satisfies $f(p) < 0$. Thus, we obtain the 99% confidence interval $(0.4996, 0.5403)$ for the probability that a newborn is male.

8. (10 points) Let X and Y be two independent, standard normal random variables.

(a) Order the following probabilities from smallest to largest:

$$\mathbb{P}(X + Y > 1)$$

$$\mathbb{P}(X - Y > 1)$$

$$\mathbb{P}(2X > 1)$$


Solution: Recall that every (non-zero) linear combination of independent normal random variables is again normally distributed. Hence, $X + Y$, $X - Y$ and $2X$ are all normal random variables. We have

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y] = 0 + 0 = 0, \quad \text{Var}[X + Y] = \text{Var}[X] + \text{Var}[Y] = 1 + 1 = 2,$$

$$\mathbb{E}[X - Y] = \mathbb{E}[X] - \mathbb{E}[Y] = 0 - 0 = 0, \quad \text{Var}[X - Y] = \text{Var}[X] + \text{Var}[Y] = 1 + 1 = 2,$$

$$\mathbb{E}[2X] = 2\mathbb{E}[X] = 2 \times 0 = 0, \quad \text{Var}[2X] = 4\text{Var}[X] = 4 \times 1 = 4.$$

Thus, $X + Y \sim N(\mu = 0, \sigma^2 = 2)$, $X - Y \sim N(\mu = 0, \sigma^2 = 2)$ and $2X \sim N(\mu = 0, \sigma^2 = 4)$.



It follows that

$$\mathbb{P}(X + Y > 1) = \mathbb{P}(X - Y > 1) < \mathbb{P}(2X > 1)$$

(why?).

- (b) Find $\text{Cov}[2X + Y + 1, X - 3Y - 2]$.

Solution: We have

$$\text{Cov}[2X + Y + 1, X - 3Y - 2] = 2\text{Cov}[X, X] - 6\text{Cov}[X, Y] + \text{Cov}[Y, X] - 3\text{Cov}[Y, Y].$$

Since X and Y are independent, $\text{Cov}[X, Y] = \text{Cov}[Y, X] = 0$. Moreover, $\text{Cov}[X, X] = \text{Var}[X] = 1$ and $\text{Cov}[Y, Y] = \text{Var}[Y] = 1$. Therefore,

$$\text{Cov}[2X + Y + 1, X - 3Y - 2] = 2 - 0 + 0 - 3 = -1.$$

9. (20 points) An elevator can safely carry a load of up to 800 Kg. The weight of an individual can be modeled as a normal random variable with mean 75 Kg and standard deviation 10 Kg. We are wondering about ways to ensure that the safety cap on the load is complied with. As first glance, we suspect that loading the elevator with 10 passengers should be relatively safe, because that would allow roughly 80 Kg per person.

- (a) What is the probability that a single passenger weighs more than 80 Kg?

Solution: Let W denote the weight of the passenger in kilograms. By assumption, $W \sim N(75, 10^2)$. Using the app for the normal distribution,

$$\mathbb{P}(W > 80) = 1 - F_{N(75, 10^2)}(80) \approx 0.30854.$$

- (b) What is the probability that of 10 passengers, at least one weighs more than 80 Kg?

Solution: Let $W_1, W_2, \dots, W_{10}$ denote the weights of the 10 passengers in kilograms. These

are independent RVs with $N(75, 10^2)$ distribution. We have

$$\begin{aligned}
 & \mathbb{P}(\text{at least one passenger weighs more than 80 Kg}) \\
 &= \mathbb{P}(W_1 > 80 \text{ or } W_2 > 80 \text{ or } \dots \text{ or } W_{10} > 80) \\
 &= 1 - \mathbb{P}(W_1 \leq 80 \text{ and } W_2 \leq 80 \text{ and } \dots \text{ and } W_{10} \leq 80) \\
 &= 1 - \mathbb{P}(W_1 \leq 80) \mathbb{P}(W_2 \leq 80) \cdots \mathbb{P}(W_{10} \leq 80) && \text{(by independence)} \\
 &= 1 - (F_{N(75,10^2)}(80))^{10} \\
 &\approx 1 - 0.69146^{10} && \text{(using the app)} \\
 &\approx 0.97502 .
 \end{aligned}$$

- (c) What is the probability that 10 passengers together weigh more than the safe limit?

Solution: Let $S_{10} := W_1 + W_2 + \cdots + W_{10}$. Recall that any (non-zero) linear combination of independent normal RVs is again a normal RV, hence S_{10} is normal with

$$\begin{aligned}
 \mathbb{E}[S_{10}] &= \mathbb{E}[W_1] + \mathbb{E}[W_2] + \cdots + \mathbb{E}[W_{10}] \\
 &= 10 \times 75 = 750 , \\
 \mathbb{V}\text{ar}[S_{10}] &= \mathbb{V}\text{ar}[W_1] + \mathbb{V}\text{ar}[W_2] + \cdots + \mathbb{V}\text{ar}[W_{10}] && \text{(by independence)} \\
 &= 10 \times 10^2 = 1000 .
 \end{aligned}$$

Using the app for the normal distribution,

$$\begin{aligned}
 & \mathbb{P}(\text{10 passengers together weigh more than 800 Kg}) \\
 &= \mathbb{P}(S_{10} > 800) \\
 &= 1 - F_{N(750,1000)}(800) \\
 &\approx 0.05692 .
 \end{aligned}$$

- (d) Explain why the probability calculated in part (c) is smaller than the probability calculated in part (b)

Solution: If the sum of the weights of 10 passengers is larger than 800 Kg, then at least one of them must weigh more than 80 Kg (why?). However, if one or more passengers weigh more than 80 Kg, it is still possible that other passengers weigh less than 80 Kg, in such a way that the total weight remains less than 800 Kg.

- (e) What is the maximum number of passengers in order for the probability of overloading to be at most 0.1%?

Solution: Let $W_1, W_2, \dots, W_n$ be the weights of n passengers in kilograms, and $S_n := W_1 + W_2 + \cdots + W_n$. As in part (c), S_n is normally distributed with

$$\mathbb{E}[S_n] = 75n , \quad \mathbb{V}\text{ar}[S_n] = 10^2 n .$$

In order for the probability of overloading to be at most 0.1%, we must have

$$0.001 \geq \mathbb{P}(S_n > 800) = \mathbb{P}\left(\frac{S_n - 75n}{10\sqrt{n}} > \frac{800 - 75n}{10\sqrt{n}}\right)$$


Note that $\frac{S_n - 75n}{10\sqrt{n}}$ has the standard normal distribution. Thus,

$$\mathbb{P}\left(\frac{S_n - 75n}{10\sqrt{n}} > \frac{800 - 75n}{10\sqrt{n}}\right) = 1 - \Phi\left(\frac{800 - 75n}{10\sqrt{n}}\right),$$

where Φ denotes the cdf of the standard normal distribution. According to the app for the normal distribution, $1 - \Phi(z) \leq 0.001$ if and only if $z \geq 3.09023$. Hence, to ensure $\mathbb{P}(S_n > 800) \leq 0.001$ we must have

$$\frac{800 - 75n}{10\sqrt{n}} \geq 3.09023.$$

In particular, n must satisfy $75n + 30.9023\sqrt{n} - 800 \leq 0$.



Solving the quadratic equation $75x^2 + 30.9023x - 800 = 0$, we obtain two solutions $x_1 \approx -3.54282$ and $x_2 \approx 3.01079$. Hence, n must satisfy $-3.54282 \leq \sqrt{n} \leq 3.01079$, or equivalently, $n \leq 3.01079^2 \approx 9.06486$. It follows that the largest n for which $\mathbb{P}(S_n > 800) \leq 0.001$ is $n = 9$.

10. (20+ points) A *Geiger-Müller counter* (GM counter for short) is an instrument used for detecting ionized radiations. In normal circumstances, a GM counter records background radiation at an average rate of 1 tick per minute. While testing a GM counter in the backyard of your house, you notice occasional sudden surges of ticks followed by periods of silence. To understand this phenomenon, you decide to model the timing of the ticks on the GM counter with a Poisson process.


(a) What is the probability that within a 1-minute interval, the GM counter ticks more than 5 times?

Solution: Let $\lambda = 1$ tick/min denote the rate of the Poisson process. The number of ticks in a 1-minute interval is a Poisson random variable N with mean $\mu = \lambda \times 1 = 1$. Hence,

$$\begin{aligned} \mathbb{P}(N > 5) &= 1 - \mathbb{P}(N \leq 5) \\ &= 1 - \sum_{n=0}^5 \mathbb{P}(N = n) \\ &= 1 - e^{-\mu} \sum_{n=0}^5 \frac{\mu^n}{n!} \\ &= 1 - e^{-1} \left(\frac{1}{0!} + \frac{1}{1!} + \frac{1}{2!} + \frac{1}{3!} + \frac{1}{4!} + \frac{1}{5!} \right) \\ &\approx 0.00059418 \end{aligned}$$

Alternatively, this probability can be found using the app for the Poisson distribution.

- (b) Divide a 30-minute interval into 30 1-minute sub-intervals. What is the probability that during at least one of these sub-intervals, the GM counter ticks more than 5 times?


Solution:

Let N_i (for $i = 1, 2, \dots, 30$) denote the number of ticks in the i th sub-interval. Each N_i is a Poisson random variable with mean $\mu = 1$. According to part (a), for each i we have $\mathbb{P}(N_i > 5) \approx 0.00059418$. Since the sub-intervals are disjoint, the random variables $N_1, N_2, \dots, N_{30}$ are independent. It follows that

$$\begin{aligned}
 & \mathbb{P}(N_i > 5 \text{ for at least one sub-interval}) \\
 &= 1 - \mathbb{P}(N_1 \leq 5 \text{ and } N_2 \leq 5 \text{ and } \dots \text{ and } N_{30} \leq 5) \\
 &= 1 - \mathbb{P}(N_1 \leq 5) \mathbb{P}(N_2 \leq 5) \cdots \mathbb{P}(N_{30} \leq 5) \\
 &\approx 1 - (1 - 0.00059418)^{30} \\
 &\approx 0.017673.
 \end{aligned}$$

- (bonus) (c) Use the result of the previous part to argue that the chance that within 30 minutes, there is a 1-minute interval during which more than 10 ticks are recorded is less than 2%.

Solution: Consider again the partitioning of a 30-minute interval into disjoint 1-minute sub-intervals, indicated in red in the following diagram.



Now, consider an arbitrary 1-minute sub-interval. An example is drawn in green. Every such sub-interval intersects one or two of the red sub-intervals. If the green sub-interval contains more than 10 ticks, then at least one of the red sub-intervals it intersects must have more than 5 ticks (why?). Therefore,

$$\begin{aligned}
 & \mathbb{P}(\text{there is a 1-minute sub-interval with more than 10 ticks}) \\
 & \leq \mathbb{P}(\text{at least one of the red sub-intervals has more than 5 ticks}) \\
 & \approx 0.017673 \quad (\text{by the result of part (b)}) \\
 & < 2\%.
 \end{aligned}$$

Still wondering about the observed occasional surges of ticks, you decide to test the hypothesis that the amount of radiations in your backyard is higher than ordinary. You start counting the ticks and notice that the 50th tick arrives in less than 30 minutes.

- (d) Write down the precise statements of the two competing hypotheses.

Solution: Let λ denote the average rate of the GM counter ticks in your backyard.

H_0 $\lambda = 1$ tick/min; The amount of radiation in your backyard is ordinary.

H_1 $\lambda > 1$ tick/min; The amount of radiation in your backyard is higher than ordinary.

(e) Compute the p-value.

Solution: Let M denote the random number of ticks from when you start counting until 30 minutes later. According to the null hypothesis, M is a Poisson random variable with mean $\lambda \times 30 = 30$. We have observed that $M \geq 50$. Therefore,

$$\begin{aligned} \text{p-value} &= \mathbb{P}(\text{a result which is at least as extreme as what is observed} \mid H_0) \\ &= \mathbb{P}(M \geq 50 \mid H_0) \\ &\approx 0.00052. \quad (\text{using the app for the Poisson distribution}) \end{aligned}$$

Two slightly different approaches to find the p-value:

1. Recall the normal approximation to the Poisson distribution discussed in class. Since the parameter of M is relatively large, we may want to approximate its distribution under the null hypothesis with a normal distribution with the same mean and variance. Let M' be a normal random variable with mean 30 and variance 30. Then,

$$\begin{aligned} \text{p-value} &= \mathbb{P}(M \geq 50 \mid H_0) \\ &\approx \mathbb{P}(M' \geq 49.5) \quad (\text{using the histogram correction}) \\ &\approx 0.00019. \quad (\text{using the app for the normal distribution}) \end{aligned}$$

Although this approximation is quite off the mark, it has still captured the order of magnitude of the p-value.

2. The observation can equivalently be expressed in terms of a gamma random variable. Namely, let S denote the random time at which the 50th tick arrives. According to the null hypothesis, S has the $\text{Gamma}(r = 50, \lambda = 1 \text{ tick/min})$ distribution. We have observed that $S < 30 \text{ min}$. Therefore,

$$\begin{aligned} \text{p-value} &= \mathbb{P}(\text{a result which is at least as extreme as what is observed} \mid H_0) \\ &= \mathbb{P}(S < 30 \mid H_0) \\ &\approx 0.00052. \quad (\text{using the app for the gamma distribution}) \end{aligned}$$

(f) At a significance level of 1%, how should you conclude?

Solution: Since the calculated p-value, 0.00052, is (much) smaller than the significance level, we must reject the null hypothesis. In other words, we have a statistically significant evidence that the amount of radiation in your backyard is higher than ordinary!

Remark. “Statistical significance” should not be mistaken for “medical/physical significance”. In particular, saying that the evidence is statistically significant does not entail that the amount of radiation is significantly high for it to raise health concerns. In general, the term “statistical significance” only refers to the strength of the *evidence* in our judgment between the two competing hypotheses.