Lecture 15 Worksheet

Chrysafis Vogiatzis

Every worksheet will work as follows.

1. You will be asked to form a group with other students in the class: you can make this as big or as small as you'd like, but groups of 4-5 work best.
2. Read through the worksheet, discussing any questions with the other members of your group.
   - You can call me at any time for help!
   - I will also be interrupting you for general guidance and announcements at random points during the class time.
3. Answer each question (preferably in the order provided) to the best of your knowledge.
4. While collaboration between students is highly encouraged and expected, each student has to submit their own version.
5. You will have 24 hours (see gradescope) to submit your work.

Activity 1: Anchoring activity

Let us begin with a small game activity. Consider a (discrete) **uniformly distributed population** X that can take any integer value between 1 and α, where $\alpha > 1$ is sadly unknown. To estimate it, we will try three point estimators. Given a sample of $n = 3$ numbers (call them X_1, X_2, X_3) from the population, we will estimate α as:

a) $\hat{\Theta}_1 = \frac{X_1+X_2+X_3}{3}$.

b) $\hat{\Theta}_2 = \frac{X_1+2X_2+X_3}{3}$.

c) $\hat{\Theta}_3 = \frac{2}{3}\left(X_1 + X_2 + X_3\right) - 1$.

Problem 1: Using estimators

First, navigate yourself to the website

`http://vogiatzis.web.illinois.edu/random.html`

and generate three numbers. Call them X_1, X_2, X_3. What is the point estimate [1] you'd get for α with each of the three estimators?

[1] Recall that we call point estimate the actual value you get by using a point estimator.

Answer to Problem 1.

$X_1 = \quad , \ X_2 = \quad , \ X_3 =$

a) $\hat{\theta}_1 =$

b) $\hat{\theta}_2 =$

c) $\hat{\theta}_3 =$

Problem 2: Estimator bias

If we are trying to estimate α, recall that the bias of some estimator $\hat{\Theta}$ can be calculated as

$$bias\left[\hat{\Theta}\right] = E\left[\hat{\Theta}\right] - \alpha.$$

Moreover, recall that because X_1, X_2, X_3 were obtained from the same population X we know that the following are true:

$$E\left[X_1\right] = E\left[X_2\right] = E\left[X_3\right] = E\left[X\right] = \mu$$
$$Var\left[X_1\right] = Var\left[X_2\right] = Var\left[X\right] = \sigma^2.$$

You do not need the variance in this question, but you may need it later today!

Finally, remember that the mean μ depends on the distribution we have, so in our case it would be $\mu = \frac{1+\alpha}{2}$ (seeing as the numbers obtained come from a discrete uniform distribution between 1 and α).

Putting all of the above together, what is the bias of each of the three estimators? Your bias could possibly be a function of α!

Answer to Problem 2.

a) $bias\left[\hat{\Theta}_1\right] =$

b) $bias\left[\hat{\Theta}_2\right] =$

c) $bias\left[\hat{\Theta}_3\right] =$

You must have gotten that the bias of estimator $\hat{\Theta}_3$ is equal to zero! This is great news. Let us check what its variance is.

Problem 3: Estimator variance

What is the variance of $\hat{\Theta}_3$? [2] Recall that the variance too can be a function of the unknown parameter α. The variance of a uniformly distributed discrete random variable between a and b is $\frac{(b-a+1)^2-1}{12}$. In our case, our population X is between 1 and α so the formula gives us a variance of

$$Var[X] = \frac{\alpha^2 - 1}{12}.$$

[2] You **will** need this variance property. If X and Y are independent random variables, then:

$$Var[aX + bY + c] = a^2 Var[X] + b^2 Var[Y].$$

Answer to Problem 3.

$Var\left[\hat{\Theta}_3\right] =$

Problem 4: Estimator variance for different sample sizes

How would the variance change if we picked a sample of $n = 5$ observations $(X_1, X_2, X_3, X_4, X_5)$ and then calculated the estimator as $\hat{\Theta} = \frac{2}{5}(X_1 + X_2 + X_3 + X_4 + X_5) - 1$? Would it increase, decrease, or stay the same? What happens as n increases more and more?

Answer to Problem 4.

$Var\left[\hat{\Theta}\right] =$

This brings us to our first realization. With an unbiased estimator, larger samples will lead to smaller variances!

Let's make it interesting. Use this estimator to produce some values from the website and note here your best estimate for what α is. What is the biggest number you can possibly obtain from this website? The biggest number that can be produced from the website is...

My best estimate is...

Activity 2: Weird point estimators

Assume that a population is distributed with pdf $f(x) = c\,(1 + \theta x)$, $-1 \leq x \leq 1$, where θ is an unknown parameter, and c a constant. [3]

[3] That means, in English, that c has to be one value and one value alone, whereas θ can be *anything*.

Problem 5: Back to basics

Let's return to the basics for a second! What should c be equal to in order for $f(x)$ to be a valid continuous pdf? [4]

[4] Integrate $f(x)$ over its domain of all values of x allowed and equate to 1.

Answer to Problem 5.

If we are doing things right, θ should disappear after taking the integral. This implies that θ can take any from a series of values, hence being a **parameter** rather than a **constant**. A parameter can be any thing; a constant has to take on a specific value.

Problem 6: Where did you come up with this?

Assume you obtain a sample of n observations. Consider the sample average $\overline{X} = (X_1 + X_2 + \ldots + X_n)\,/n$. Show that $\hat{\Theta} = 3\overline{X}$ is an **unbiased estimator** for θ. [5]

[5] To do so first you need to calculate $E\left[\hat{\Theta}\right] = E\left[3\overline{X}\right] = 3E\left[\overline{X}\right]$. But, isn't $E\left[\overline{X}\right] = E\left[\frac{\sum X_i}{n}\right] = E\left[X\right]$? We could calculate this as $E\left[X\right] = \int_{-1}^{+1} x f(x) dx$, which may be a function of θ...

Answer to Problem 6.

Problem 7: Variance and standard error

What is the standard error of the point estimator $\hat{\Theta} = 3\overline{X}$? [6] The standard error can be calculated as

$$SE\left[\hat{\Theta}\right] = \sqrt{Var\left[\hat{\Theta}\right]}.$$

[6] To calculate this you will first need to calculate the expectation and the variance of population X. They could very well be a function of θ as you do not know what the parameter is equal to...

Answer to Problem 7.

This is quite common: when dealing with an unknown parameter (in our case, θ), the bias and the variance can depend on the value that the parameter actually has. So, how do we compare estimators for unknown parameters?

Activity 3: Comparing point estimators

Assume we have collected a sample of $n = 3$ observations X_1, X_2, X_3 coming from a population X distributed with *some pdf* with unknown μ and known $\sigma^2 = 16$. We have devised three point estimators for the unknown population mean:

- Get the average from the first two observations omitting the third, i.e.,
$$\hat{\Theta}_1 = \frac{X_1 + X_2}{2}.$$
- Add the "odd" observations once and the "even" observations doubled and divide everything by 4, i.e.,
$$\hat{\Theta}_2 = \frac{X_1 + 2X_2 + X_3}{4}.$$
- Once again omit the third observation and simply add the first two and divide by 4, i.e.,
$$\hat{\Theta}_3 = \frac{X_1 + X_2}{4}.$$

Now, for one last time in this worksheet, we go ahead and calculate the bias and variance of each of the estimators.

- We have for the first estimator, $\hat{\mathbf{\Theta}}_\mathbf{1}$:
$$\begin{aligned} bias\left[\hat{\Theta}_1\right] &= E\left[\hat{\Theta}_1\right] - \mu = E\left[\frac{X_1 + X_2}{2}\right] - \mu = \frac{1}{2}E\left[X_1\right] + \frac{1}{2}E\left[X_2\right] - \mu = \\ &= \frac{1}{2}E\left[X\right] + \frac{1}{2}E\left[X\right] - \mu = E\left[X\right] - \mu = \mu - \mu = 0. \\ Var\left[\hat{\Theta}_1\right] &= Var\left[\frac{X_1 + X_2}{2}\right] = \frac{1}{4}Var\left[X_1 + X_2\right] = Var\left[X_1\right] + \frac{1}{4}Var\left[X_2\right] = \\ &= \frac{1}{4}Var\left[X\right] + \frac{1}{4}Var\left[X\right] = \frac{1}{2}Var\left[X\right] = \frac{1}{2}\sigma^2 = 8. \end{aligned}$$
- Similarly, for the second estimator, $\hat{\mathbf{\Theta}}_\mathbf{2}$, we would end up with: [7]
$$bias\left[\hat{\Theta}_2\right] = 0.$$
$$Var\left[\hat{\Theta}_2\right] = 6.$$
- Finally, for the third estimator, $\hat{\mathbf{\Theta}}_\mathbf{3}$:
$$bias\left[\hat{\Theta}_3\right] = -\frac{\mu}{2}.$$
$$Var\left[\hat{\Theta}_3\right] = 2.$$

[7] These calculations are left as an exercise.

Or, in tabular form:

	$\hat{\Theta}_1$	$\hat{\Theta}_2$	$\hat{\Theta}_3$
bias	0	0	$-\mu/2$
variance	8	6	2

Problem 8: Comparison I

What is the MSE of each of the estimators? Which estimator is the best according to its MSE, if we have been told that $\mu > 4$? [8]

Answer to Problem 8.

Note how the result would have changed if $\mu < 4$. On the other hand, there is no way that the first estimator has the smallest variance among the three.

Problem 9: Observation

Earlier in this activity, we observe that the first two estimators are unbiased (i.e., zero bias). In general, assume you are collecting a sample of n observations $(X_1, X_2, \ldots, X_n)$ and are using $\hat{\Theta} = a_1 X_1 + a_2 X_2 + \ldots + a_n X_n$ to estimate the unknown mean. What condition should $a_1 + a_2 + \ldots + a_n$ satisfy in order for $\hat{\Theta}$ to have bias equal to zero? [9]

Answer to Problem 9.

[8] As a reminder:

$$MSE = bias^2 + variance.$$

[9] Hmmm.. What can you tell about $E\left[\sum_{i=1}^{n} a_i X_i\right]$? Additionally, never forget that $E[X_1] = E[X_2] = \ldots = E[X_n] = E[X]$ because all observations come from the same population X!