

Review of Elementary Statistics

Economics 143

Jinyong Hahn

Comments by Michael Powell

Random Variables

- Variables that takes on alternative values, each with a probability less than or equal to 1
- Probability distribution associates all possible outcomes with the probability that each will occur
- Continuous random variable: Any value on real line
- Discrete random variable: Finite number of real values

Expected Value

- X : Discrete random variable
- x_1, \dots, x_N : N possible outcomes
- p_1, \dots, p_N : Associated probabilities
- Expected value of the mean is then defined by

$$E[X] = \mu_X = p_1x_1 + \dots + p_Nx_N = \sum_{i=1}^N p_ix_i$$

Variance

- Variance provides a measure of the dispersion of random variable around the mean
- Weighted average of the squares of the deviations of outcomes from its expected value

$$Var[X] = \sigma_X^2 = \sum_{i=1}^N p_ix_i^2 - E[X]^2$$

Comment: There is an alternative definition of variance that is common in most textbooks:

$$Var[X] = \sum_{i=1}^N (x_i - E[X])^2 p_i$$

This is equivalent to the above definition:

$$\begin{aligned}\sum_{i=1}^N (x_i - E[X])^2 p_i &= \sum_{i=1}^N (x_i^2 - 2x_i E[X] + (E[X])^2) p_i \\ &= \sum_{i=1}^N x_i^2 p_i - 2E[X] \sum_{i=1}^N x_i p_i + (E[X])^2 \sum_{i=1}^N p_i \\ &= E[X^2] - 2E[X] E[X] + (E[X])^2 \\ &= E[X^2] - (E[X])^2\end{aligned}$$

- Variance is in itself an expectation
- The (positive) square root of the variance is called the standard deviation, and written as σ_X

Useful Results

Proposition: $E[aX + b] = aE[X] + b$

Proof:

$$\begin{aligned}E[aX + b] &= \sum_{i=1}^N (ax_i + b) p_i \\ &= a \sum_{i=1}^N x_i p_i + b \sum_{i=1}^N p_i \\ &= aE[X] + b \quad \text{since } \sum_{i=1}^N p_i = 1\end{aligned}$$

Proposition: $E[(aX)^2] = a^2 E[X^2]$

Proof:

$$\begin{aligned}E[(aX)^2] &= \sum_{i=1}^N (ax_i)^2 p_i \\ &= \sum_{i=1}^N a^2 x_i^2 p_i \\ &= a^2 \sum_{i=1}^N x_i^2 p_i \\ &= a^2 E[X^2]\end{aligned}$$

Proposition: $Var[aX + b] = a^2 Var[X]$

Proof:

$$\begin{aligned} \text{Var} [aX + b] &= \sum_{i=1}^N ((ax_i + b) - E[aX + b])^2 p_i \\ &= \sum_{i=1}^N (ax_i + b - aE[X] - b)^2 p_i \text{ by above proposition} \\ &= \sum_{i=1}^N (a(x_i - E[X]))^2 p_i \\ &= a^2 \sum_{i=1}^N (x_i - E[X])^2 p_i \\ &= a^2 \text{Var} [X] \end{aligned}$$

Two Random Variables

- (X, Y) : Discrete random variables
- Joint distribution can be described by a list of all possible outcomes (x_i, y_j) $i = 1, \dots, M$; $j = 1, \dots, N$, and associated probabilities p_{ij}
- Example: X is the family income. $Y = 1$ if the head of a household has a college education, and $Y = 0$ otherwise.

	$Y = 1$	$Y = 0$
$X = 5000$	0	1/4
$X = 10,000$	1/8	1/8
$X = 15,000$	1/3	1/6

Covariance

$$\begin{aligned} \text{Cov} (X, Y) &= E [(X - E[X]) (Y - E[Y])] \\ &= E [XY] - E[X] E[Y] \\ &= \sum_{i=1}^M \sum_{j=1}^N p_{ij} x_i y_j - E[X] E[Y] \end{aligned}$$

Proof:

$$\begin{aligned}
Cov(X, Y) &= E[(X - E[X])(Y - E[Y])] \\
&= \sum_{j=1}^M \sum_{i=1}^N (x_i - E[X])(y_j - E[Y]) p_{ij} \\
&= \sum_{j=1}^M \sum_{i=1}^N (x_i y_j - y_j E[X] - x_i E[Y] + E[X] E[Y]) p_{ij} \\
&= \sum_{j=1}^M \sum_{i=1}^N x_i y_j p_{ij} - E[X] \sum_{j=1}^M \sum_{i=1}^N y_j p_{ij} - E[Y] \sum_{j=1}^M \sum_{i=1}^N x_i p_{ij} + E[X] E[Y] \sum_{j=1}^M \sum_{i=1}^N p_{ij} \\
&= \sum_{j=1}^M \sum_{i=1}^N x_i y_j p_{ij} - E[X] \sum_{j=1}^M y_j \sum_{i=1}^N p_{ij} - E[Y] \sum_{i=1}^N x_i \sum_{j=1}^M p_{ij} + E[X] E[Y] \sum_{j=1}^M p_j^Y \\
&= \sum_{j=1}^M \sum_{i=1}^N x_i y_j p_{ij} - E[X] \sum_{j=1}^M y_j p_j^Y - E[Y] \sum_{i=1}^N x_i p_i^X + E[X] E[Y] \sum_{j=1}^M p_j^Y \\
&= \sum_{j=1}^M \sum_{i=1}^N x_i y_j p_{ij} - E[X] E[Y] - E[Y] E[X] + E[X] E[Y] \\
&= E[XY] - E[X] E[Y]
\end{aligned}$$

Where I defined $\sum_{j=1}^M p_{ij} = p_i^X$ as the marginal distribution of X and $\sum_{i=1}^N p_{ij} = p_j^Y$ as the marginal distribution of Y .

Correlation Coefficient

$$\rho(X, Y) = \frac{Cov(X, Y)}{\sigma_X \sigma_Y}$$

- Unlike covariance, correlation coefficient is scale-free
- Correlation coefficient will always lie between -1 and 1

Useful Results

Proposition: $E[X + Y] = E[X] + E[Y]$

Proof:

$$\begin{aligned}
 E[X + Y] &= \sum_{j=1}^M \sum_{i=1}^N (x_i + y_j) p_{ij} \\
 &= \sum_{j=1}^M \sum_{i=1}^N x_i p_{ij} + \sum_{j=1}^M \sum_{i=1}^N y_j p_{ij} \\
 &= \sum_{i=1}^N x_i \sum_{j=1}^M p_{ij} + \sum_{j=1}^M y_j \sum_{i=1}^N p_{ij} \\
 &= \sum_{i=1}^N x_i p_i^X + \sum_{j=1}^M y_j p_j^Y \\
 &= E[X] + E[Y]
 \end{aligned}$$

Proposition: $Var[X + Y] = Var[X] + Var[Y] + 2Cov[X, Y]$

Proof:

$$\begin{aligned}
 Var[X + Y] &= \sum_{j=1}^M \sum_{i=1}^N ((x_i + y_j) - E[X + Y])^2 p_{ij} \\
 &= \sum_{j=1}^M \sum_{i=1}^N ((x_i + y_j) - E[X] - E[Y])^2 p_{ij} \text{ by previous proposition} \\
 &= \sum_{j=1}^M \sum_{i=1}^N ((x_i - E[X]) + (y_j - E[Y]))^2 p_{ij} \\
 &= \sum_{j=1}^M \sum_{i=1}^N ((x_i - E[X])^2 + 2(x_i - E[X])(y_j - E[Y]) + (y_j - E[Y])^2) p_{ij} \\
 &= \sum_{j=1}^M \sum_{i=1}^N (x_i - E[X])^2 p_{ij} + \sum_{j=1}^M \sum_{i=1}^N (y_j - E[Y])^2 p_{ij} \\
 &\quad + 2 \sum_{j=1}^M \sum_{i=1}^N (x_i - E[X])(y_j - E[Y]) p_{ij} \\
 &= Var[X] + Var[Y] + 2Cov(X, Y)
 \end{aligned}$$

Independence

- Independence roughly means that the probability of an outcome associated with Y will be unrelated with the outcome associated with X , and vice versa
- Knowledge of realization of X does not help updating the probability assessment of Y
- Example: Two consecutive coin tossings

Comment: The formal definition of independence says that if X and Y are independent, their joint distribution can be written as the product of their marginal distributions. Equivalently, $\forall i, j, p_{ij} = p_i^X \cdot p_j^Y$.

Proposition: If X and Y are independent $E[X \cdot Y] = E[X] \cdot E[Y]$

Proof: Let X and Y be independent. Then $\forall i, j, p_{ij} = p_i^X p_j^Y$.

$$\begin{aligned}
 E[X \cdot Y] &= \sum_{j=1}^M \sum_{i=1}^N x_i y_j p_{ij} \\
 &= \sum_{j=1}^M \sum_{i=1}^N x_i y_j p_i^X p_j^Y \\
 &= \left(\sum_{j=1}^M y_j p_j^Y \right) \left(\sum_{i=1}^N x_i p_i^X \right) \\
 &= E[Y] E[X]
 \end{aligned}$$

Proposition: $Cov(X, Y) = 0$

Proof:

$$\begin{aligned}
 Cov(X, Y) &= E[XY] - E[X] E[Y] \\
 &= E[X] E[Y] - E[X] E[Y] \text{ from previous proposition} \\
 &= 0
 \end{aligned}$$

Estimation

- Means, variances, and covariances can be measured with certainty only if we know the exact probability distribution
- We cannot know the true value of means or variance in general
- We use the sample information to obtain the best possible estimates
- *Estimator:* A rule which will give a sample estimate for each and every possible sample (Because the sample itself is random, an estimator is a random variable)
- *Estimate:* A number

Some decision science

- Let b denote an estimator of an arbitrary parameter β , say expected value. What kind of properties do we usually want from b ?

- *Lack of Bias*: We usually want the expected value of b to be equal to the parameter of interest β . We say that b is unbiased if so. We define

$$\text{Bias} = E[b] - \beta$$

- *Efficiency*: While lack of bias seems to be a desirable property, that alone does not say anything about the dispersion of the estimator about the true parameter. We usually want small dispersion. We say that an unbiased estimator b is efficient if its variance is the smaller than the variance of any other unbiased estimators.
- *Minimum MSE*: Sometimes we have to trade off bias and variance of estimators. It seems useful to minimize the mean squared error

$$MSE = E[(b - \beta)^2]$$

Proposition: $E[(b - \beta)^2] = \text{Bias}^2 + \text{Var}(b)$

Proof:

$$\begin{aligned} E[(b - \beta)^2] &= E[b^2 - 2b\beta + \beta^2] \\ &= E[b^2] - 2\beta E[b] + \beta^2 \\ &= \underbrace{E[b^2] - (E[b])^2}_{\text{Var}(b)} + \underbrace{(E[b])^2 - 2\beta E[b] + \beta^2}_{\text{Bias}^2} \\ &= \text{Bias}^2 + \text{Var}(b) \end{aligned}$$

- *Consistency* (Large sample property): We would like b to be very close to β as sample size increases. We want, as the sample size increases, the probability that b differ from β will be small. We say that $\text{plim } b = \gamma$ if, as sample size approaches infinity, the probability that $|b - \gamma|$ will be less than any arbitrary small number will approach 1. We say that b is consistent if $\text{plim } b = \beta$. Usually, econometricians tend to be more concerned with consistency than with lack of bias

Estimation of μ_X

- Expected value is usually estimated by sample mean

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N X_i$$

- Sample mean is unbiased for μ_X

$$E[\bar{X}] = E\left[\frac{1}{N} \sum_{i=1}^N X_i\right] = \frac{1}{N} E\left[\sum_{i=1}^N X_i\right] = \frac{1}{N} \sum_{i=1}^N E[X_i] = \frac{1}{N} \sum_{i=1}^N \mu_X = \frac{1}{N} N \mu_X = \mu_X$$

Estimation of σ_X^2 and Cov

- Sample variance

$$s_X^2 = \frac{1}{N-1} \sum_{i=1}^N (X_i - \bar{X})^2$$

is unbiased for variance

Proposition: $s_X^2 = \frac{1}{N-1} \sum_{i=1}^N (X_i - \bar{X})^2$ is unbiased.

Proof:

$$\begin{aligned} s_X^2 &= \frac{1}{N-1} \sum_{i=1}^N (X_i - \bar{X})^2 \\ &= \left(\frac{1}{N-1} \right) \sum_{i=1}^N (X_i^2 - 2\bar{X}X_i + \bar{X}^2) \\ &= \left(\frac{1}{N-1} \right) \left[\sum_{i=1}^N X_i^2 - 2\bar{X} \sum_{i=1}^N X_i + \sum_{i=1}^N \bar{X}^2 \right] \\ &= \left(\frac{1}{N-1} \right) \left[\sum_{i=1}^N (X_i^2) - 2N\bar{X}^2 + N\bar{X}^2 \right] \\ &= \left(\frac{1}{N-1} \right) \left[\sum_{i=1}^N X_i^2 - N\bar{X}^2 \right] \end{aligned}$$

This gives us:

$$\begin{aligned} E(s_X^2) &= E \left[\left(\frac{1}{N-1} \right) \left[\sum_{i=1}^N X_i^2 - n\bar{X}^2 \right] \right] \\ &= \left(\frac{1}{N-1} \right) E \left[\sum_{i=1}^N X_i^2 - n\bar{X}^2 \right] \\ &= \left(\frac{1}{N-1} \right) \left(\sum_{i=1}^N E[X_i^2] - n\bar{X}^2 \right) \\ &= \left(\frac{1}{N-1} \right) (NE[X^2] - NE[\bar{X}^2]) \\ &= \left(\frac{N}{N-1} \right) (E[X^2] - E[\bar{X}^2]) \end{aligned}$$

Recall that:

$$\begin{aligned} Var(X) &= E[X^2] - (E[X])^2 \\ E[X^2] &= Var(X) + (E[X])^2 \end{aligned}$$

And (shown below):

$$\begin{aligned}\frac{\text{Var}(X)}{N} &= E[\bar{X}^2] - (E[X])^2 \text{ since } \bar{X} \text{ is unbiased} \\ E[\bar{X}^2] &= \frac{\text{Var}(X)}{N} + (E[X])^2\end{aligned}$$

This gives us:

$$\begin{aligned}E(s_X^2) &= \left(\frac{N}{N-1}\right) \left(\text{Var}(X) + (E[X])^2 - \frac{\text{Var}(X)}{N} - (E[X])^2\right) \\ &= \left(\frac{N}{N-1}\right) \left(\frac{N\text{Var}(X) - \text{Var}(X)}{N}\right) \\ &= \text{Var}(X)\end{aligned}$$

- More intuitive estimator

$$\frac{1}{N} \sum_{i=1}^N (X_i - \bar{X})^2$$

is unbiased

Comment: Actually, this estimator IS biased.

- The right denominator $N - 1$ is sometimes called the degree of freedom
- Likewise, sample covariance is an unbiased estimator of covariance

$$\frac{1}{N-1} \sum_{i=1}^N (X_i - \bar{X})(Y_i - \bar{Y})$$

and sample correlation is defined by

$$\frac{\frac{1}{N-1} \sum_{i=1}^N (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\frac{1}{N-1} \sum_{i=1}^N (X_i - \bar{X})^2} \sqrt{\frac{1}{N-1} \sum_{i=1}^N (Y_i - \bar{Y})^2}}$$

Law of Large Numbers

- The variance of the sample mean

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N X_i$$

is equal to

$$\frac{1}{N^2} \sum_{i=1}^N \text{Var}(X_i) = \frac{1}{N^2} \sum_{i=1}^N \sigma^2 = \frac{\sigma^2}{N}.$$

- Therefore, the variance of the sample mean is negligible if N is sufficiently large. Note that $\text{Var}(\bar{X}) \rightarrow 0$ as $N \rightarrow \infty$
- Roughly speaking, the sample mean almost becomes a nonstochastic number with very small dispersion around its mean μ
- When the sample size is large enough, the sample mean is virtually equal to μ

Normal Distribution

- Symmetric bell-shaped density
- Fully characterized by mean and variance
- If two or more random variables are normally distributed, any weighted sum of these variables will be normally distributed.
- In particular, if X_1, \dots, X_N constitute a random sample from $N(\mu, \sigma^2)$, i.e., they are independent and identically distributed, then

$$\bar{X} \sim N\left(\mu, \frac{\sigma^2}{N}\right)$$

and therefore

$$\frac{\bar{X} - \mu}{\sigma/\sqrt{N}} \sim N(0, 1)$$

- If $X \sim N(\mu, \sigma^2)$, then

$$\Pr[\mu - 1.96\sigma < X < \mu + 1.96\sigma] = .95$$

$$\Pr[\mu - 2.57\sigma < X < \mu + 2.57\sigma] = .99$$

- If $X \sim N(0, 1)$, then we say that X has a standard normal distribution

Chi-Square distribution

- The sum of the squares of N independently distributed standard normal random variables is distributed as chi square distribution with N degrees of freedom
- We sometimes write it as $X \sim \chi^2(N)$

t distribution

- Assume that $X \sim N(0, 1)$ and $Z \sim \chi^2(N)$
- Further assume that they are independent
- Then,

$$\frac{X}{\sqrt{Z/N}} \sim t(N)$$

- It has a t -distribution with N degrees of freedom

F distribution

- Assume that $X \sim \chi^2(N_1)$ and $Z \sim \chi^2(N_2)$
- Further assume that they are independent
- Then,

$$\frac{X/N_1}{Z/N_2} \sim F(N_1, N_2)$$

Confidence Interval and Hypothesis Test for μ

- Suppose that $X_1, \dots, X_N \sim N(\mu, \sigma^2)$
- Suppose that σ^2 is known
- 95% confidence interval for μ is given by

$$\bar{X} - 1.96 \frac{\sigma}{\sqrt{n}} < \mu < \bar{X} + 1.96 \frac{\sigma}{\sqrt{n}}$$

This is because

$$\frac{\bar{X} - \mu}{\sigma/\sqrt{N}} \sim N(0, 1)$$

and hence

$$\Pr \left[\left| \frac{\bar{X} - \mu}{\sigma/\sqrt{N}} \right| < 1.96 \right] = .95$$

or

$$\Pr \left[-1.96 \frac{\sigma}{\sqrt{n}} < \mu - \bar{X} < 1.96 \frac{\sigma}{\sqrt{n}} \right] = .95$$

or

$$\Pr \left[\bar{X} - 1.96 \frac{\sigma}{\sqrt{n}} < \mu < \bar{X} + 1.96 \frac{\sigma}{\sqrt{n}} \right] = .95$$

- For example, if $N = 100$, $\bar{X} = 3$, and $\sigma = 10$, then it would be 3 ± 1.96
- Confidence intervals can be used to test hypotheses. Suppose that the null hypothesis is given by $H_0 : \mu = 0$. Because the 95% confidence interval is given by $(1.04, 4.96)$, it is unlikely that the null hypothesis is true. We may therefore reject the null hypothesis.
- Throughout the course, we implicitly assume that the alternative hypothesis is two sided
- Formally, we reject the null (under 5% significance level) $H_0 : \mu = \mu_0$ if

$$\left| \frac{\bar{X} - \mu_0}{\sigma/\sqrt{N}} \right| \geq 1.96$$

It can be verified that this procedure is equivalent to the intuitive procedure based on 95% confidence intervals

- Roughly speaking, the hypothesis testing strategy is based on following intuition: We believe that, if $\mu = \mu_0$, an estimator of μ , i.e., \bar{X} should be close to μ_0 . If not, there is a significant evidence against the null. Therefore, it is sensible to reject the null when $|\bar{X} - \mu_0|$ is large. Now, the problem is that, there is a noise in the estimation of μ . Therefore, our intuition has to be modified a little bit. If we believe that there is a large noise, then it may be sensible to accept the null even when $|\bar{X} - \mu_0|$ is reasonably large. The ratio $\left| \frac{\bar{X} - \mu_0}{\sigma/\sqrt{N}} \right|$ provides a sensible compromise between these two conflicts. It is statistically convenient because we know $\frac{\bar{X} - \mu_0}{\sigma/\sqrt{N}}$ has a standard normal distribution.
- When σ is unknown, which is a more realistic situation, we reject instead when

$$\left| \frac{\bar{X} - \mu_0}{s_X/\sqrt{N}} \right| \geq t_c(N - 1)$$

where $t_c(N - 1)$ is the critical value from $(N - 1)$. The ratio $\left| \frac{\bar{X} - \mu_0}{s_X/\sqrt{N}} \right|$ is called the t-statistic, and has a $t(N - 1)$ distribution under the null

- As in normal distribution based confidence interval, t-distribution based inference is based on the fact that

$$\frac{\bar{X} - \mu}{s_X/\sqrt{N}} \sim t(N - 1)$$

- Finding the critical value: See Table 3, designed for two-tail tests. It reports the probability that the t-value will exceed the number in the table in absolute value. Usually, we only care about the column under .05 (= usual significance level). The title “Percentiles...” is actually a misnomer (In some other books, we would want to look at the column under .025).