STA 111: Probability & Statistical Inference

Lecture Six – Review; The Multinomial Distribution
D.S. Section 5.9

Instructor: Olanrewaju Michael Akande

Department of Statistical Science, Duke University

Outline

- Questions from Last Lecture
- Review Questions
- The Multinomial Distribution
- Recap

Introduction

– In the last lecture we extended our discussion of the standard normal distribution to the arbitrary case.

– We also looked at the normal approximation to the binomial distribution.

– Today we will go over more examples to help understand expectations and the normal distribution.

– We will also learn about the generalization of the binomial distribution.

Expectation of a Discrete Random Variable Cont'd

Example 1 (To be done in class): Consider a random variable X which can take on the values: 1, 2, 3 and 4 and define its probability mass function to be $f(x) = c/x^2$.

Also consider another random variable Y which can take on the following values: 5, 10, 15, 20 and 25 and define its probability mass function to be $f(y) = ky$.

Find: (i) c (ii) k (iii) $\mathbb{E}(X)$ (iv) $\mathbb{E}(Y)$ (v) $\mathbb{V}(X)$
(vi) $\mathbb{V}(Y)$ (vii) $\mathbb{E}(X - Y)$ (viii) $\mathbb{E}(2X^2 + 3Y + 5)$

Expectation of a Discrete Random Variable Cont'd

Example 2 (To be done in class): Assume that a policyholder is four times more likely to file exactly two claims as to file exactly three claims. Assume also that the number X of claims of this policyholder is Poisson. What is the expectation $\mathbb{E}(X^2)$?

The Normal Distribution Cont'd

Example 3 (To be done in class – D.S. Chapter 5 Exercises, Question 2): Suppose that X has the normal distribution for which the mean is 1 and variance is 4. Find the value of each of the probabilities:

(a) $\mathbb{P}(X \leq 3)$ (b) $\mathbb{P}(X > 5)$ (c) $\mathbb{P}(X = 1)$ (d) $\mathbb{P}(2 \leq X \leq 5)$

(e) $\mathbb{P}(-1 < X < 0.5)$ (f) $\mathbb{P}(|X| \leq 2)$ (g) $\mathbb{P}(-1 \leq -2X + 3 \leq 8)$

The Multinomial Distribution

Remember that the Binomial distribution describes a random variable that represents success (or failure) based on n independent trials of an experiments with two outcomes, where the probability of success is the same for each trial. Our toy examples have been tossing a coin n times and rolling a die n times.

An extension to the binomial distribution which allows for three or more outcomes (k outcomes) for each trial, where the probability of each outcome is the same for each trial is called the Multinomial Distribution.

The Multinomial Distribution Cont'd

A random variable (vector) $\mathbf{X} = (X_1, \ldots, X_k)$ has the multinomial distribution with parameters n and $\boldsymbol{p} = (p_1, \ldots, p_k)$ if its pmf is given by

$$f(\boldsymbol{x}) = \mathbb{P}(\boldsymbol{X} = \boldsymbol{x}) = \binom{n}{x_1, \ldots, x_k} p_i^{x_1}, \cdots, p_n^{x_k} \quad \text{for} \quad x_1 + \ldots + x_k = n$$

where $\binom{n}{x_1, \ldots, x_k} = \dfrac{n!}{x_1! x_2! \ldots x_k!}$ is called the multinomial coefficient.

The Multinomial Distribution Cont'd

Example 4: Suppose that for a single roll of a loaded die we can observe one with probability 0.25, two or three with probability 0.1 each, four or five with probability 0.2 each and six with probability 0.15.

What is the probability that in 15 independent rolls, we will observe one four times, two thrice, three once, four once, five four times and six twice.

The Multinomial Distribution Cont'd

Let $\boldsymbol{X} = (X_1, X_2, X_3, X_4, X_5, X_6)$ be the observed combination for the die. Then $\boldsymbol{p} = (0.25, 0.1, 0.1, 0.2, 0.2, 0.15)$ and

$$\begin{aligned}\mathbb{P}\left[\boldsymbol{X} = (4,3,1,1,4,2)\right] &= \binom{n}{x_1, \ldots, x_k} p_i^{x_1}, \cdots, p_n^{x_k} \\ &= \binom{15}{4,3,1,1,4,2} 0.25^4 0.1^3 0.1^1 0.2^1 0.2^4 0.15^2 \\ &= \frac{15!}{4!3!1!1!4!2!} 0.25^4 0.1^3 0.1^1 0.2^1 0.2^4 0.15^2 \\ &= 0.0005\end{aligned}$$

The Multinomial Distribution Cont'd

A few important points:

1. The vector of probabilities $\boldsymbol{p}$ should sum to 1.
2. Any $X_i \in (X_1, \ldots, X_k)$ has a binomial distribution with parameters n and p_i. *This is why treating the die roll as a binomial when we only care about one of the sides works!*
3. Collapsing the k different outcomes to two outcomes gets you back to a binomial distribution as well. If $\mathbf{X} = (X_1, \ldots, X_k)$ has a multinomial distribution with parameters n and $\boldsymbol{p} = (p_1, \ldots, p_k)$ and $l < k$ where $i_1, \ldots, i_l$ are distinct elements of the set $\{1, \ldots, k\}$, then $Y = X_{i_1} + \cdots + X_{i_l}$ has a binomial distribution with parameters n and $p_{i_1} + \cdots + p_{i_l}$