Math 20, Fall 2017

Edgar Costa

Week 5

Dartmouth College

Last Week

Last week:

- Continuous random variable
- Expected Value

Today:

- Variance

St. Petersburg paradox: Expected value

We can think of the expected value as the "predicted" value of a random variable (if we repeat the experiment infinitely many times).

St. Petersburg paradox

A casino offers a game of chance for a single player in which a fair coin is tossed at each stage. The initial stake starts at 2 dollars and is doubled every time heads appears. The first time tails appears, the game ends and the player wins whatever is in the pot. Thus the player wins 2 dollars if tails appears on the first toss, 4 dollars if heads appears on the first toss and tails on the second, 8 dollars if heads appears on the first two tosses and tails on the third, and so on. Mathematically, the player wins 2^k dollars, where k equals number of tosses until we observe tails.

What would be a fair price to pay the casino for entering the game?

St. Petersburg paradox : Expected value

The payout is $Y = 2^{\mathrm{Geo}(1/2)}$ What is the expected payout?

$$\begin{aligned} E\left[2^{\mathrm{Geo}(1/2)}\right] &= 2\frac{1}{2} + 4\frac{1}{4} + 8\frac{1}{8} + \cdots \\ &= 1 + 1 + 1 + \cdots \\ &= +\infty \end{aligned}$$

What if the casino has finite resources? (or bounds the maximum jackpot)

St. Petersburg paradox : with finite resources

What if the casino has finite resources? (or bounds the maximum jackpot)

If $W =$ total maximum jackpot then

$$E[\text{payout}] = \sum_{k=1}^{L} 2^k \frac{1}{2^k} + W \sum_{k=L+1}^{\infty} \frac{1}{2^k} = L + \frac{W}{2^L} << W,$$

where $L = \lfloor \log_2(W) \rfloor$ = the maximum number of times the casino can play before it no longer fully covers the next bet

St. Petersburg paradox: Some numbers

What if the casino has finite W resources?

W	Expected value
100	7.56
1,000,000	20.91
1,000,000,000	30.86
10^{100}	333.14

You toss a coin, if you get tails you owe me 10^7 dollars, if you get heads I owe you $10^7 + 5$ dollars.

Would you play this game?

Variance

Definition

Let $\mu_X = E[X]$ The variance of X, denoted by $V[X]$, is

$$V[X] = E[(X - \mu_X)^2]$$

.

The standard deviation of X, denoted by $\sigma(X)$ or σ_X is $\sqrt{V[X]}$.

Example

Compute the Expected value and Variance of

$$X = \text{“the outcome of a six faced die roll”}$$

Variance

Theorem

$$V[X] = E[X^2] - \mu_X^2$$

Proof:

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

Example

Compute the Expected value and Variance of

$$X = \text{“the outcome of a six faced die roll”}$$

Variance properties

- $V[cX] = c^2 V[X]$
- $V[X + c] = V[X]$

What about $V[X + Y] = ?$

Variance of sum

Theorem

If X and Y are independent random variables

$$V[X+Y] = V[X] + V[Y]$$

Proof: expand $E[(X+Y)^2]$ and use $E[XY] = E[X]E[Y]$.

Exercise

$V[\mathrm{Bin}(n,p)] = ?$

Variance of sums

Exercise

Let $S_n = X_1 + \cdots + X_n$, with X_i independent, $E[X_i] = \mu$ and $V[X_i] = \sigma^2$

and $A_n = S_n/n$

Compute the expected value, variance and the standard deviation of S_n and A_n

Theorem

$$E[S_n] = n\mu \qquad V[S_n] = n\sigma^2$$

$$E[A_n] = \mu \qquad E[A_n] = \frac{\sigma^2}{n}$$

Important Distributions: Uniform Distribution

- All outcomes of an experiment are equally likely
- we say "X is uniformly distributed"
- If X is discrete and $n = \#\Omega$ then what is the distribution function?

Important Distributions: Binomial Distribution

- Repeat a Bernoulli process n times with probability p of success
- $\operatorname{Bin}(n, p) = X = \sum_{i=0}^{n} X_i$ where X_i are iid to a Bernoulli(p)
- $\Omega = \{0, \ldots, n\}$
- $P(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}$ with $k \in \Omega$
- $E(X) = np$ and $V(X) = np(1 - p)$

- $\text{Geo}(p) = X =$ the number of repeats of a Bernoulli(p) until success
- $P(X = k) = (1-p)^{k-1}p$
- $E(X) = \frac{1}{p}$
- $V(X) = \frac{1-p}{p^2}$ (you prove this in Worksheet #5)

Example

In each time unit a customer arrives with probability p.

What is the probability that no customer arrives in the next n units?

Modified Example

Example

In each time unit a customer arrives with probability p.

What is the probability that it takes n times units for **two** customers to arrive?

Definition

X = number of trials in a sequence of iid Bernoulli trials needed to get r success.

- $\Omega =?$
- $P(X = k) =?$
- $r = 1 \rightsquigarrow X = \text{Geometric}(p)$
- Why $\sum_{k\in\Omega} P(X = k) = 1$?

Poisson distribution

Problem

Assume that you expect λ calls every day out of n possible calls. How would model this?

If we model this as a Binomial distribution we have $np = \lambda \Leftrightarrow p = \frac{\lambda}{n}$

$$
\begin{aligned}
P(X = k) &= \binom{n}{k} \left(\frac{\lambda}{n}\right)^k \left(1 - \frac{\lambda}{n}\right)^{n-k} \\
&= \frac{1}{k!} \frac{n!}{(n-k)! n^k} \lambda^k (1 - \lambda/n)^{n-k}
\end{aligned}
$$

What happens if we take $n \to +\infty$? (Use: $\frac{n!}{(n-k)!n^k} \to 1$ and $(1 - \lambda/n)^{n-k} \to 1$)

$$P(X = k) \approx \frac{1}{k!} \lambda^k e^{-\lambda}$$

See tables in the GS book, page 196!

Comparison

j	Poisson $\lambda = .1$	Binomial $n = 100$ $p = .001$	Poisson $\lambda = 1$	Binomial $n = 100$ $p = .01$	Poisson $\lambda = 10$	Binomial $n = 1000$ $p = .01$
0	.9048	.9048	.3679	.3660	.0000	.0000
1	.0905	.0905	.3679	.3697	.0005	.0004
2	.0045	.0045	.1839	.1849	.0023	.0022
3	.0002	.0002	.0613	.0610	.0076	.0074
4	.0000	.0000	.0153	.0149	.0189	.0186
5			.0031	.0029	.0378	.0374
6			.0005	.0005	.0631	.0627
7			.0001	.0001	.0901	.0900
8			.0000	.0000	.1126	.1128
9					.1251	.1256
					.1251	.1257

Poisson distribution (see GS book for many examples)

- $X = \text{Poisson}(\lambda) =$"the number of events that take place in an interval"
- $P(X = k) = \dfrac{e^{-\lambda}}{k!}\lambda^k$ for $k \in \Omega = \{0, 1, 2, 3, \dots\}$
- $E[X] = \lambda$ (you should be able to do it on your own)
- $V[X] = \lambda$ (a bit more challenging, start by computing $E[X(X-1)]$)

Good assumptions to use the Poisson distribution:

- The rate at which events occur is constant.
- Events occur independently.
- The probability of an event in a interval is proportional to the length of the interval
- Two events cannot occur exactly at the same time

or

The actual probability distribution the binomial distribution and the number of trials is sufficiently larger than the number of successes one is asking about.

Exercise

A typesetter makes, on the average, one mistake per 1000 words. Assume that he is setting a book with 100 words to a page. Let S_{100} be the number of mistakes that he makes on a single page.

- How would you model S_{100}?
- $P(S_{100} = 0) =?$
- $P(S_{100} = 1) =?$
- $P(S_{100} < 10) =?$

The exact probability distribution for S_{100} is $\mathrm{Bin}(100, p = 1/1000)$

However, the Poisson distribution with $\lambda = 100 \cdot \frac{1}{1000} = 0.1$ is also an appropriate model! (see previous tables)

Exercise

In a class of 80 students, the professor calls on 1 student chosen at random for a recitation in each class period. There are 32 class periods in a term.

1. Write a formula for the exact probability that a given student is called upon j times during the term.
2. Write a formula for the Poisson approximation for this probability. Using your formula estimate the probability that a given student is called upon more than twice.

Exercise (Hypergeometric Distribution)

Four balls are drawn at random, without replacement, from an urn containing 4 red balls and 3 blue. Let X be the number of red balls drawn.

1. What is the range of X?
2. What is the probability that $X = 2$? $X = k$?
3. Find $E[X]$.
4. What if each time a ball is drawn, the ball is replaced in the urn.

$\Omega = \{1, 2, 3, 4\}$ and $P(X = k) = \frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}}$ with $N = 7$, $K = 4$, and $n = 4$.
Without replacement is known as Hypergeometric Distribution that depends on N, K, and n.

With replacement X becomes binomial distribution.

Exercise

A bridge deck has 52 cards with 13 cards in each of four suits: spades, hearts, diamonds, and clubs. A hand of 13 cards is dealt from a shuffled deck. Find the probability that the hand has

1. a distribution of suits 4, 4, 3, 2 (for example, four spades, four hearts, three diamonds, two clubs).
2. a distribution of suits 5, 3, 3, 2.

Exponential distribution

- $T = \text{Exp}(\lambda)$ (λ is any positive constant, depending on the experiment.)
- How long until something happens? (that occurs continuously and independently at a constant average rate)
  For example: time between occurrences of a Poisson processes (work it out!)
- $\Omega_T = [0, +\infty]$
- $f_T(t) = \begin{cases} \lambda e^{-\lambda t} & \text{if } t \leq 0 \\ 0 & \text{otherwise} \end{cases}$
- $P(T \leq x) = 1 - e^{-\lambda t}$ if $t \geq 0$.
- $P(T > t + s | T \geq s) = P(T > t)$ (memoryless!)
- $E[T] = \dfrac{1}{\lambda}, \quad V[T] = \dfrac{1}{\lambda^2}$

Normal distribution

- The normal density function of the normal distribution $N(\mu, \sigma)$ with parameters μ and σ is defined as follows:

$$f_X(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-(x-\mu)^2/2\sigma^2}.$$

- The parameter μ represents the “center” of the density.
- The parameter σ is a measure of “spread” of the density, and thus it is assumed to be positive.

Normal distribution: In a picture

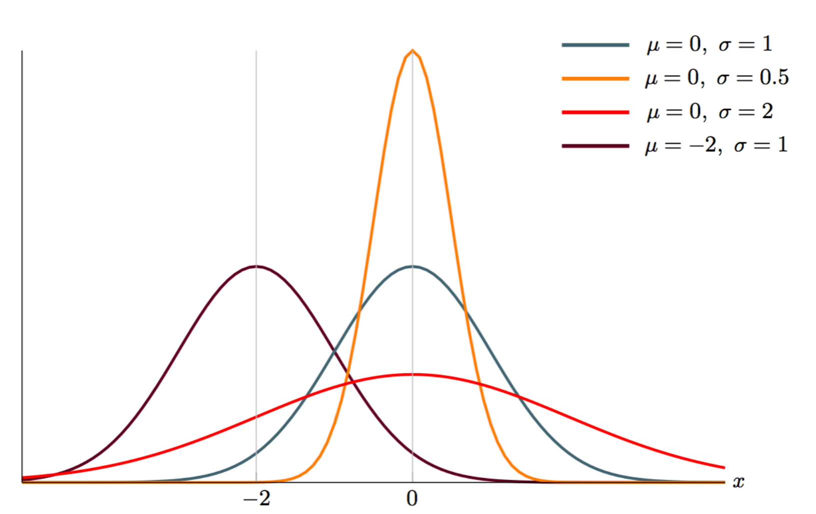

Normal distribution

- We focus mostly on $\mu = 0$ and $\sigma = 1$
- We will call this particular normal density function the *standard normal density*, and we will denote it by $\phi(x)$:

$$\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$$

- There is no nice formula for $\int_a^b \phi(x)\,\mathrm{d}x$
- We instead use numerical tables for $\int_0^d \phi(x)\,\mathrm{d}x$
- Note that

$$\frac{N(\mu,\sigma) - \mu}{\sigma} \sim N(0,1)$$

Exercise

On a test that determines whether an applicant receives a scholarship, the scores are distributed by a normal random variable with $\mu = 500$, $\sigma = 100$. It the top 5% of scores qualify for a scholarship, how high a score do you need to get it?

We seek a such that $P(X \geq a) = 0.05$. Then $P(X < a) = 0.95$.

For $Z = N(0,1)$, we have $P(Z \leq 1.65) \approx 0.95$

So $\frac{a-\mu}{\sigma} = 1.65 \rightsquigarrow a = 665$