12: Independent RVs

Jerry Cain

April 26th, 2024

Lecture Discussion on Ed

Sums of independent Binomial RVs

Independent discrete RVs

Recall the definition of independent events E and F:

$$P(EF) = P(E)P(F)$$

Two discrete random variables X and Y are **independent** if:

for all x, y:

$$P(X = x, Y = y) = P(X = x)P(Y = y)$$

Different notation, same idea:

$$p_{X,Y}(x, y) = p_X(x)p_Y(y)$$

- Intuitively: knowing value of X tells us nothing about the distribution of Y (and vice versa)
- If two variables are not independent, they are termed **dependent**.

Sum of independent Binomials

$X \sim \text{Bin}(n_1, p)$
$Y \sim \text{Bin}(n_2, p)$
X, Y independent

$$X + Y \sim \text{Bin}(n_1 + n_2, p)$$

Intuition:

- Each trial in X and Y is independent and has same success probability p
- Define $Z =$# successes in $n_1 + n_2$ independent trials, each with success probability p. $Z \sim \text{Bin}(n_1 + n_2, p)$ and $Z = X + Y$ as well.

Holds in general case:

$X_i \sim \text{Bin}(n_i, p)$
X_i independent for $i = 1, \dots, n$

$$\sum_{i=1}^{n} X_i \sim \text{Bin}\left(\sum_{i=1}^{n} n_i, p\right)$$

If only it were always so simple

Coin flips

Flip a coin with probability p of heads a total of $n + m$ times.

Let X = number of heads in first n flips. $X \sim \text{Bin}(n, p)$
Y = number of heads in next m flips. $Y \sim \text{Bin}(m, p)$
Z = total number of heads in $n + m$ flips.

1. Are X and Z independent?
2. Are X and Y independent?

Coin flips

Flip a coin with probability p of heads a total of $n + m$ times.

Let X = number of heads in first n flips. $X \sim \text{Bin}(n, p)$
Y = number of heads in next m flips. $Y \sim \text{Bin}(m, p)$
Z = total number of heads in $n + m$ flips.

1. Are X and Z independent? ❌

Counterexample: What if $Z = 0$?

2. Are X and Y independent? ✅

$$P(X = x, Y = y) = P\left(\begin{array}{c}\text{first } n \text{ flips have } x \text{ heads} \\ \text{and next } m \text{ flips have } y \text{ heads}\end{array}\right)$$

$$= \binom{n}{x} p^x (1-p)^{n-x} \binom{m}{y} p^y (1-p)^{m-y}$$

$$= P(X = x) P(Y = y)$$

# of mutually exclusive outcomes in event: $\binom{n}{x}\binom{m}{y}$
$P(\text{each outcome})$
$= p^x (1-p)^{n-x} p^y (1-p)^{m-y}$

This probability (found through counting) is the product of the marginal PMFs.

Convolution: Sum of independent Poisson RVs

Convolution: Sum of independent random variables

For any discrete random variables X and Y:

$$P(X + Y = n) = \sum_k P(X = k, Y = n - k)$$

In particular, for independent discrete random variables X and Y:

$$\boxed{P(X + Y = n) = \underbrace{\sum_k P(X = k)P(Y = n - k)}_{\text{the convolution of } p_X \text{ and } p_Y}}$$

Insight into convolution

For independent discrete random variables X and Y:

$$P(X+Y=n)=\sum_k P(X=k)P(Y=n-k)$$

the convolution of p_X and p_Y

Suppose X and Y are independent, both with support $\{0, 1, \ldots, n, \ldots\}$:

		X						
		0	1	2	…	n	$n+1$	…
Y	0					✓		
	…				…			
	$n-2$			✓				
	$n-1$		✓					
	n	✓						
	$n+1$							
	…							

- ✓: event where $X+Y=n$
- Each event has probability:
  $P(X=k, Y=n-k)$
  $= P(X=k)P(Y=n-k)$
  (because X, Y are independent)
- $P(X+Y=n) =$ sum of mutually exclusive events

Sum of 2 dice rolls

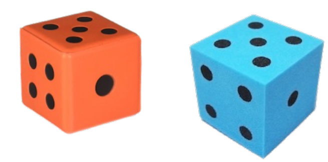

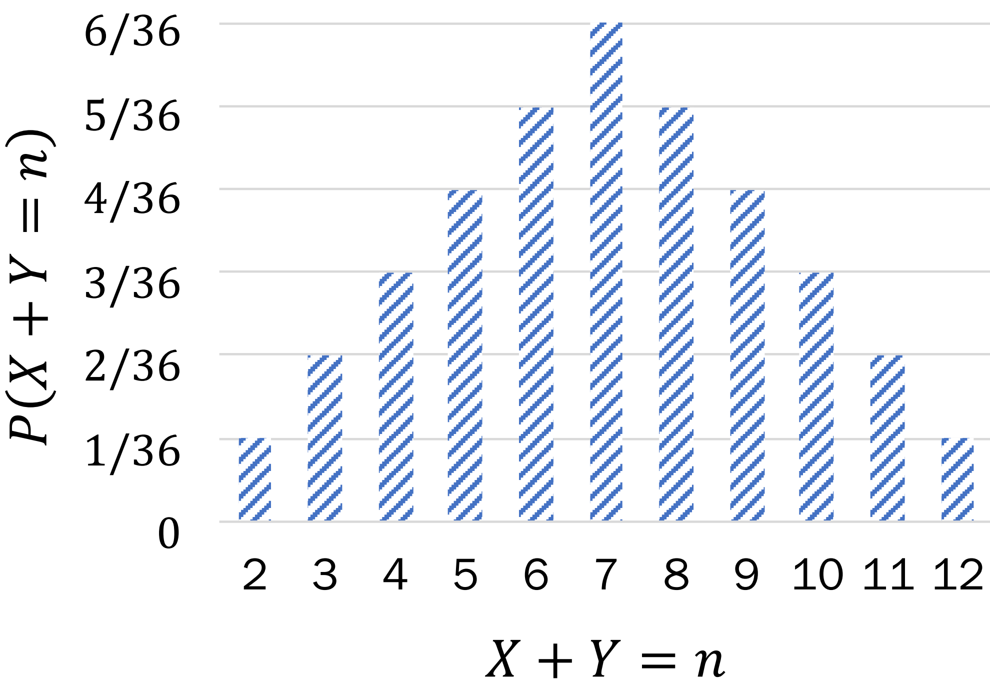

The distribution of a sum of <u>2</u> dice rolls is a convolution of <u>2</u> PMFs.

Example:

$$P(X+Y=4) =$$
$$P(X=1)P(Y=3)$$
$$+\,P(X=2)P(Y=2)$$
$$+\,P(X=3)P(Y=1)$$

Sum of 10 dice rolls (fun preview)

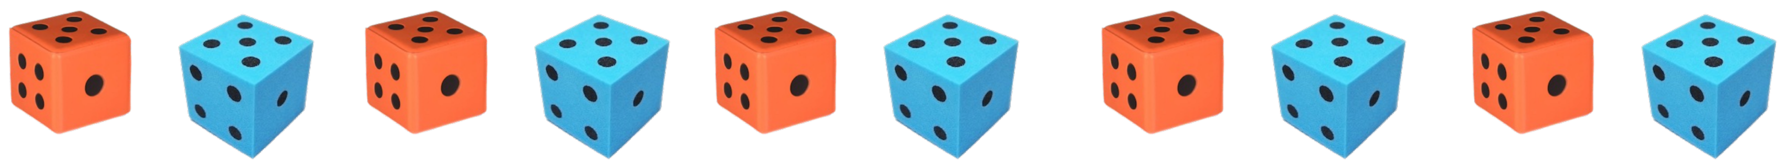

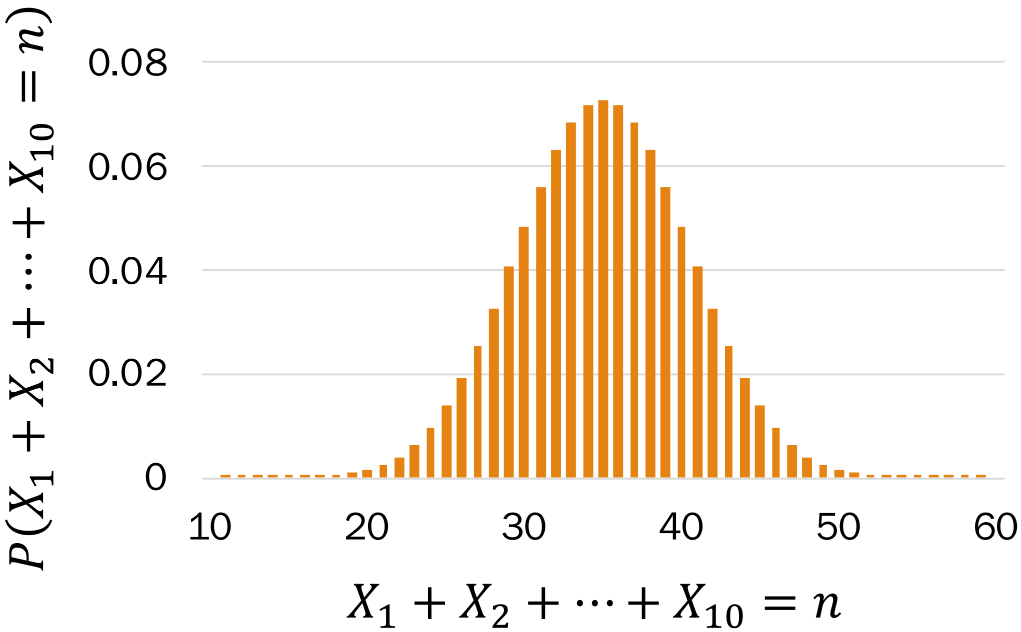

The distribution of a sum of <u>10</u> dice rolls is a convolution <u>10</u> PMFs.

Looks kinda Normal…???
(more on this in a few weeks)

Sum of independent Poissons

$X \sim \text{Poi}(\lambda_1), Y \sim \text{Poi}(\lambda_2)$
X, Y independent

$$X + Y \sim \text{Poi}(\lambda_1 + \lambda_2)$$

Proof (just for reference):

$$P(X+Y=n) = \sum_k P(X=k)P(Y=n-k)$$

X and Y independent, convolution

$$= \sum_{k=0}^{n} e^{-\lambda_1} \frac{\lambda_1^k}{k!} e^{-\lambda_2} \frac{\lambda_2^{n-k}}{(n-k)!} = e^{-(\lambda_1+\lambda_2)} \sum_{k=0}^{n} \frac{\lambda_1^k \lambda_2^{n-k}}{k!\,(n-k)!}$$

PMF of Poisson RVs

$$= \frac{e^{-(\lambda_1+\lambda_2)}}{n!} \sum_{k=0}^{n} \frac{n!}{k!\,(n-k)!} \lambda_1^k \lambda_2^{n-k} = \underbrace{\frac{e^{-(\lambda_1+\lambda_2)}}{n!} (\lambda_1+\lambda_2)^n}_{\text{Poi}(\lambda_1+\lambda_2)}$$

Binomial Theorem:

$$(a+b)^n = \sum_{k=0}^{n} \binom{n}{k} a^k b^{n-k}$$

Sum of independent Poissons

$X \sim \text{Poi}(\lambda_1), Y \sim \text{Poi}(\lambda_2)$
X, Y independent

$$X + Y \sim \text{Poi}(\lambda_1 + \lambda_2)$$

- n servers with independent number of requests/minute
- Server i's requests each minute can be modeled as $X_i \sim \text{Poi}(\lambda_i)$

What is the probability that the total number of web requests received at all servers in the next minute exceeds 10?

Exercises

Independent questions

1. Let $X \sim \text{Bin}(30, 0.01)$ and $Y \sim \text{Bin}(50, 0.02)$ be independent RVs.
   - How do we compute $P(X + Y = 2)$ using a Poisson approximation?
   - How do we compute $P(X + Y = 2)$ exactly?

2. Let $N = \#$ of requests to a web server per day. Suppose $N \sim \text{Poi}(\lambda)$.
   - Each request independently comes from a human (prob. p), or bot $(1 - p)$.
   - Let X be # of human requests/day, and Y be # of bot requests/day.

   Are X and Y independent? What are their marginal PMFs?

1. Approximating the sum of independent Binomial RVs

Let $X \sim \text{Bin}(30, 0.01)$ and $Y \sim \text{Bin}(50, 0.02)$ be independent RVs.

- How do we compute $P(X + Y = 2)$ using a Poisson approximation?

- How do we compute $P(X + Y = 2)$ exactly?

$$P(X + Y = 2) = \sum_{k=0}^{2} P(X = k)P(Y = 2 - k)$$

$$= \sum_{k=0}^{2} \binom{30}{k} 0.01^k (0.99)^{30-k} \binom{50}{2-k} 0.02^{2-k} 0.98^{50-(2-k)} \approx 0.2327$$

2. Web server requests

Let N = # of requests to a web server per day. Suppose $N \sim \text{Poi}(\lambda)$.

- Each request independently comes from a human (prob. p), or bot ($1 - p$).
- Let X be # of human requests/day, and Y be # of bot requests/day.

Are X and Y independent? What are their marginal PMFs?

$$P(X = x, Y = y) = P(X = x, Y = y \mid N = x + y)P(N = x + y) + P(X = x, Y = y \mid N \neq x + y)P(N \neq x + y)$$

Law of Total Probability

$$= P(X = x \mid N = x + y)P(Y = y \mid X = x, N = x + y)P(N = x + y)$$

Chain Rule

$$= \binom{x+y}{x} p^x (1-p)^y \cdot 1 \cdot e^{-\lambda} \frac{\lambda^{x+y}}{(x+y)!}$$

Given $N = x + y$ indep. trials, $X \mid N = x + y \sim \text{Bin}(x + y, p)$

$$= \frac{(x+y)!}{x!\,y!} e^{-\lambda} \frac{(\lambda p)^x \big(\lambda(1-p)\big)^y}{(x+y)!} = e^{-\lambda p} \frac{(\lambda p)^x}{x!} \cdot e^{-\lambda(1-p)} \frac{\big(\lambda(1-p)\big)^y}{y!}$$

$$= P(X = x)P(Y = y)$$

where $X \sim \text{Poi}(\lambda p), Y \sim \text{Poi}\big(\lambda(1-p)\big)$

Yes, X and Y are independent!

Expectation of Common RVs

Linearity of Expectation: Important

Expectation is a linear mathematical operation. If $X = \sum_{i=1}^{n} X_i$:

$$E[X] = E\left[\sum_{i=1}^{n} X_i\right] = \sum_{i=1}^{n} E[X_i]$$

- Even if you don't know the **distribution** of X (e.g., because the joint distribution of $(X_1, \dots, X_n)$ is unknown), you can still compute **expectation** of X.

- Problem-solving key: Define X_i such that $X = \sum_{i=1}^{n} X_i$

Most common use cases:
- $E[X_i]$ easy to calculate
- Sum of dependent RVs

Expectations of common RVs: Binomial

Review

$X \sim \text{Bin}(n, p) \quad E[X] = np$

# of successes in n independent trials with probability of success p

Recall: $\text{Bin}(1, p) = \text{Ber}(p)$

$$X = \sum_{i=1}^{n} X_i$$

Let $X_i = i$th trial is heads
$X_i \sim \text{Ber}(p), E[X_i] = p$

$$E[X] = E\left[\sum_{i=1}^{n} X_i\right] = \sum_{i=1}^{n} E[X_i] = \sum_{i=1}^{n} p = np$$

Expectations of common RVs: Negative Binomial

$Y \sim \text{NegBin}(r, p) \quad E[Y] = \frac{r}{p}$

# of independent trials with probability of success p until r successes

Recall: $\text{NegBin}(1, p) = \text{Geo}(p)$

$$Y = \sum_{i=1}^{?} Y_i$$

1. How should we define Y_i?

2. How many terms are in our summation?

Expectations of common RVs: Negative Binomial

$Y \sim \text{NegBin}(r, p) \quad E[Y] = \frac{r}{p}$

# of independent trials with probability of success p until r successes

Recall: $\text{NegBin}(1, p) = \text{Geo}(p)$

$$Y = \sum_{i=1}^{r} Y_i$$

Let Y_i = # trials to get ith success (after $(i-1)$th success)

$Y_i \sim \text{Geo}(p), E[Y_i] = \frac{1}{p}$

$$E[Y] = E\left[\sum_{i=1}^{r} Y_i\right] = \sum_{i=1}^{r} E[Y_i] = \sum_{i=1}^{r} \frac{1}{p} = \frac{r}{p}$$