

Debugging Intuition

- How to calculate the probability of at least k successes in n trials?

- X is number of successes in n trials each with probability p
- $P(X \geq k) =$

First clue that something is wrong.
Think about $p = 1$

Not mutually exclusive...

$$\binom{n}{k} p^k$$

Don't care about the rest

Probability that each is success

ways to choose slots for success

Correct:
$$P(X \geq k) = \sum_{i=k}^n \binom{n}{i} p^i (1-p)^{n-i}$$


Variance



Lagunita

CS 109
Lecture 7
April 11th, 2016

Is Peer Grading Accurate Enough?



Peer Grading on Coursera
HCI.

31,067 peer grades for
3,607 students.

Today's Topics

Last time:

Random Variables
Expectation + PDF

Today:

Variance
Bernoulli + Binominal RVs

Next time:

All the other discrete
RVs

St Petersburg

- Game set-up
 - We have a fair coin (come up “heads” with $p = 0.5$)
 - Let n = number of coin flips (“heads”) before first “tails”
 - You win $\$2^n$
- How much would you pay to play?
- Solution
 - Let X = your winnings
 - $$E[X] = \left(\frac{1}{2}\right)^1 2^0 + \left(\frac{1}{2}\right)^2 2^1 + \left(\frac{1}{2}\right)^3 2^2 + \left(\frac{1}{2}\right)^4 2^3 + \dots = \sum_{i=0}^{\infty} \left(\frac{1}{2}\right)^{i+1} 2^i$$
$$= \sum_{i=0}^{\infty} \frac{1}{2} = \infty$$
 - I'll let you play for \$1 thousand... but just once! Takers?

St Petersburg + Reality

- What if Chris has only \$65,536?
 - Same game
 - If you win over \$65,536 I leave the country.
- Solution
 - Let X = your winnings
 - $$E[X] = \left(\frac{1}{2}\right)^1 2^0 + \left(\frac{1}{2}\right)^2 2^1 + \left(\frac{1}{2}\right)^3 2^2 + \left(\frac{1}{2}\right)^4 2^3 + \dots$$
$$= \sum_{i=0}^k \left(\frac{1}{2}\right)^{i+1} 2^i \text{ s.t. } k = \log_2(65,536)$$
$$= \sum_{i=0}^{16} \frac{1}{2} = 8$$

Utility

- Utility is value of some choice
 - 2 options, each with n consequences: c_1, c_2, \dots, c_n
 - One of c_i will occur with probability p_i
 - Each consequence has some value (utility): $U(c_i)$
 - Which choice do you make?
- Example: Buy a \$1 lottery ticket (for \$1M prize)?
 - Probability of winning is $1/10^7$
 - **Buy**: $c_1 = \text{win}$, $c_2 = \text{lose}$, $U(c_1) = 10^6 - 1$, $U(c_2) = -1$
 - **Don't Buy**: $c_1 = \text{lose}$, $U(c_1) = 0$
 - $E(\text{buy}) = 1/10^7 (10^6 - 1) + (1 - 1/10^7) (-1) \approx -0.9$
 - $E(\text{don't buy}) = 1 (0) = 0$
 - *“You can't lose if you don't play!”*

And Then There's This



Lottery: A tax on people who are bad at math.
– Ambrose Bierce

Recall, Geometric Series

$$a^0 + a^1 + a^2 + \dots$$

$$= \sum_{i=0}^{\infty} a^i$$

$$= \frac{1}{1 - a}$$

where $0 < a < 1$

Breaking Vegas

- Consider even money bet (e.g., bet “Red” in roulette)
 - $p = 18/38$ you win $\$Y$, otherwise $(1 - p)$ you lose $\$Y$
 - Consider this algorithm for one series of bets:
 1. $Y = \$1$
 2. Bet Y
 3. If Win then stop
 4. If Loss then $Y = 2 * Y$, goto 2
 - Let $Z =$ winnings upon stopping
 - $$E[Z] = \left(\frac{18}{38}\right)1 + \left(\frac{20}{38}\right)\left(\frac{18}{38}\right)(2 - 1) + \left(\frac{20}{38}\right)^2 \left(\frac{18}{38}\right)(4 - 2 - 1) + \dots$$
$$= \sum_{i=0}^{\infty} \left(\frac{20}{38}\right)^i \left(\frac{18}{38}\right) \left(2^i - \sum_{j=0}^{i-1} 2^j\right) = \left(\frac{18}{38}\right) \sum_{i=0}^{\infty} \left(\frac{20}{38}\right)^i = \left(\frac{18}{38}\right) \frac{1}{1 - \frac{20}{38}} = 1$$
 - Expected winnings ≥ 0 . Use algorithm infinitely often!




Vegas Breaks You

- Why doesn't everyone do this?
 - Real games have maximum bet amounts
 - You have finite money
 - Not able to keep doubling bet beyond certain point
 - Casinos can kick you out
- But, if you had:
 - No betting limits, and
 - Infinite money, and
 - Could play as often as you want...
- Then, go for it!
 - And tell me which planet you are living on

Is $E[X]$ enough?

Variance

- Consider the following 3 distributions (PMFs)



- All have the same expected value, $E[X] = 3$
- But “spread” in distributions is different
- Variance = a formal quantification of “spread”

Variance

- If X is a random variable with mean μ then the **variance** of X , denoted $\text{Var}(X)$, is:

$$\text{Var}(X) = E[(X - \mu)^2]$$

- Note: $\text{Var}(X) \geq 0$
- Also known as the 2nd **Central** Moment, or square of the Standard Deviation

Computing Variance

$$\begin{aligned}\text{Var}(X) &= E[(X - \mu)^2] \\ &= \sum_x (x - \mu)^2 p(x) \\ &= \sum_x (x^2 - 2\mu x + \mu^2) p(x) \\ &= \sum_x x^2 p(x) - 2\mu \sum_x x p(x) + \mu^2 \sum_x p(x) \\ &= \boxed{E[X^2]} - 2\mu E[X] + \mu^2 \\ &= E[X^2] - 2\mu^2 + \mu^2 \\ &= E[X^2] - \mu^2 \\ &= \boxed{E[X^2] - (E[X])^2}\end{aligned}$$

Ladies and gentlemen, please
welcome the 2nd moment!

Variance of a 6 sided dice

- Let X = value on roll of 6 sided die
- Recall that $E[X] = 7/2$
- Compute $E[X^2]$

$$E[X^2] = (1^2)\frac{1}{6} + (2^2)\frac{1}{6} + (3^2)\frac{1}{6} + (4^2)\frac{1}{6} + (5^2)\frac{1}{6} + (6^2)\frac{1}{6} = \frac{91}{6}$$

$$\begin{aligned}\text{Var}(X) &= E[X^2] - (E[X])^2 \\ &= \frac{91}{6} - \left(\frac{7}{2}\right)^2 = \frac{35}{12}\end{aligned}$$

Properties of Variance

- $\text{Var}(aX + b) = a^2 \text{Var}(X)$

- Proof:


$$\begin{aligned}\text{Var}(aX + b) &= E[(aX + b)^2] - (E[aX + b])^2 \\ &= E[a^2X^2 + 2abX + b^2] - (aE[X] + b)^2 \\ &= a^2E[X^2] + 2abE[X] + b^2 - (a^2(E[X])^2 + 2abE[X] + b^2) \\ &= a^2E[X^2] - a^2(E[X])^2 = a^2(E[X^2] - (E[X])^2) \\ &= a^2 \text{Var}(X)\end{aligned}$$

- Standard Deviation of X , denoted $\text{SD}(X)$, is:

$$\text{SD}(X) = \sqrt{\text{Var}(X)}$$

- $\text{Var}(X)$ is in units of X^2
- $\text{SD}(X)$ is in same units as X

Intuition




Peer Grading on Coursera
HCI.

31,067 peer grades for
3,607 students.

Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade




If we base a final score off 4 peer-grades, 19% of students would get a score that was off by more than 10 percentage points

Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade


$$\text{Var}(X) = E[(X - \mu)^2]$$



Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade

$$\text{Var}(X) = E[(X - \mu)^2]$$




X	$(X - \mu)^2$
25 points	1056 points ²

Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade

$$\text{Var}(X) = E[(X - \mu)^2]$$




True grade = 58
 $E[X] = 57.5$

X	$(X - \mu)^2$
25 points	1056 points ²
80 points	506 points ²

Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade

$$\text{Var}(X) = E[(X - \mu)^2]$$




X	$(X - \mu)^2$
25 points	1056 points ²
80 points	506 points ²
50 points	56 points ²

Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade

$$\text{Var}(X) = E[(X - \mu)^2]$$



X	$(X - \mu)^2$
25 points	1056 points ²

80 points	506 points ²
-----------	-------------------------

50 points	56 points ²
-----------	------------------------


...

$$E[(X - \mu)^2] = 52 \text{ points}^2$$

Peer Grades in Coursera HCI

Let X be a random variable that represents a peer grade

$$\text{Var}(X) = E[(X - \mu)^2]$$



True grade = 58
 $E[X] = 57.5$

X	$(X - \mu)^2$
25 points	1056 points ²
80 points	506 points ²
50 points	56 points ²
...	...


$$E[(X - \mu)^2] = 52 \text{ points}^2$$

$$SD(X) = 7.2 \text{ points}$$

Second Moment

Let X be a random variable that represents a peer grade

$$E[X^2]$$




Recall that Variance is the second **central** moment. What is the second moment?

Second Moment

Let X be a random variable that represents a peer grade

$$E[X^2]$$



Recall that Variance is the second **central** moment. What is the second moment?


Second Moment

Let X be a random variable that represents a peer grade

$$E[X^2]$$

True grade = 58

$$E[X] = 57.5$$



Yea second moment
really isn't that useful a
statistic in this case...

Lots of fun with Random Variables

Moaaahhhh Random Variables

Jacob Bernoulli

- Jacob Bernoulli (1654-1705), also known as “James”, was a Swiss mathematician



- One of many mathematicians in Bernoulli family
- The Bernoulli Random Variable is named for him
- He is my *academic* great¹²-grandfather
- Same eyes as Ice Cube

Bernoulli Random Variable

- Experiment results in “Success” or “Failure”
 - X is random indicator variable (1 = success, 0 = failure)
 - $P(X = 1) = p(1) = p$ $P(X = 0) = p(0) = 1 - p$
 - X is a **Bernoulli** Random Variable: $X \sim \text{Ber}(p)$
 - $E[X] = p$
 - $\text{Var}(X) = p(1 - p)$
- Examples
 - coin flip
 - random binary digit
 - whether a disk drive crashed
 - whether someone likes a netflix movie



Feel the Bern!

Binomial Random Variable

- Consider n independent trials of Ber(p) rand. var.
 - X is number of successes in n trials
 - X is a **Binomial** Random Variable: $X \sim \text{Bin}(n, p)$

$$P(X = i) = p(i) = \binom{n}{i} p^i (1 - p)^{n-i} \quad i = 0, 1, \dots, n$$

- By Binomial Theorem, we know that $\sum_{i=0}^{\infty} P(X = i) = 1$
- Examples
 - # of heads in n coin flips
 - # of 1's in randomly generated length n bit string
 - # of disk drives crashed in 1000 computer cluster
 - Assuming disks crash independently

Bernoulli vs Binomial



Bernoulli is a type of RV



Binomial is the sum of n
Bernoullis

Three Coin Flips

- Three fair (“heads” with $p = 0.5$) coins are flipped
 - X is number of heads
 - $X \sim \text{Bin}(3, 0.5)$

$$P(X = 0) = \binom{3}{0} p^0 (1-p)^3 = \frac{1}{8}$$

$$P(X = 1) = \binom{3}{1} p^1 (1-p)^2 = \frac{3}{8}$$

$$P(X = 2) = \binom{3}{2} p^2 (1-p)^1 = \frac{3}{8}$$

$$P(X = 3) = \binom{3}{3} p^3 (1-p)^0 = \frac{1}{8}$$

Error Correcting Codes

- Error correcting codes
 - Have original 4 bit string to send over network
 - Add 3 “parity” bits, and send 7 bits total
 - Each bit independently corrupted (flipped) in transition with probability 0.1

Error Correcting Codes

Key for 7 bits

Send: 1110?

Receive: 1110000?

Receive: 1010100?

Flip set: $O_i E_j^C$
for all odd sets i
and even sets j

Error Correcting Codes

- Error correcting codes
 - Have original 4 bit string to send over network
 - Add 3 “parity” bits, and send 7 bits total
 - Each bit independently corrupted (flipped) in transition with probability 0.1
 - $X =$ number of bits corrupted: $X \sim \text{Bin}(7, 0.1)$
 - But, parity bits allow us to correct at most 1 bit error
- $P(\text{a correctable message is received})?$
 - $P(X = 0) + P(X = 1)$

Error Correcting Codes

- Using error correcting codes: $X \sim \text{Bin}(7, 0.1)$

$$P(X = 0) = \binom{7}{0} (0.1)^0 (0.9)^7 \approx 0.4783$$

$$P(X = 1) = \binom{7}{1} (0.1)^1 (0.9)^6 \approx 0.3720$$

- $P(X = 0) + P(X = 1) = 0.8503$
- What if we didn't use error correcting codes?
 - $X \sim \text{Bin}(4, 0.1)$
 - $P(\text{correct message received}) = P(X = 0)$

$$P(X = 0) = \binom{4}{0} (0.1)^0 (0.9)^4 = 0.6561$$

- Using error correction improves reliability ~30%!

Genetic Inheritance

- Person has 2 genes for trait (eye color)
 - Child receives 1 gene (equally likely) from each parent
 - Child has brown eyes if either (or both) genes brown
 - Child only has blue eyes if both genes blue
 - Brown is “dominant” (d) , Blue is “recessive” (r)
 - Parents each have 1 brown and 1 blue gene
- 4 children, what is $P(3 \text{ children with brown eyes})$?
 - Child has blue eyes: $p = (1/2) (1/2) = 1/4$ (2 blue genes)
 - $P(\text{child has brown eyes}) = 1 - (1/4) = 0.75$
 - $X = \#$ of children with brown eyes. $X \sim \text{Bin}(4, 0.75)$

$$P(X = 3) = \binom{4}{3} (0.75)^3 (0.25)^1 \approx 0.4219$$

Properties of Bin(n, p)


- Consider: $X \sim \text{Bin}(n, p)$
- $E[X] = np$
- $\text{Var}(X) = np(1 - p)$
- So, to compute $E[X^2]$, we have:

$$\text{Var}(X) = E[X^2] - (E[X])^2$$


$$\begin{aligned} E[X^2] &= \text{Var}(X) + (E[X])^2 \\ &= np(1 - p) + (np)^2 \\ &= n^2p^2 - np^2 + np \end{aligned}$$

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

Galton Board




Galton Board



When a marble hits a pin, it has equal chance of going left or right.


Galton Board



When a marble hits a pin, it has equal chance of going left or right. Each pin represents an independent event.


Galton Board

The bucket index that a marble lands in is equal to the number of times the marble went right




Galton Board

We can define an indicator random variable (R) which represents whether a particular marble goes right as a Bernoulli $R \sim \text{Ber}(0.5)$




Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in.




Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(\text{levels}, 0.5)$




Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$



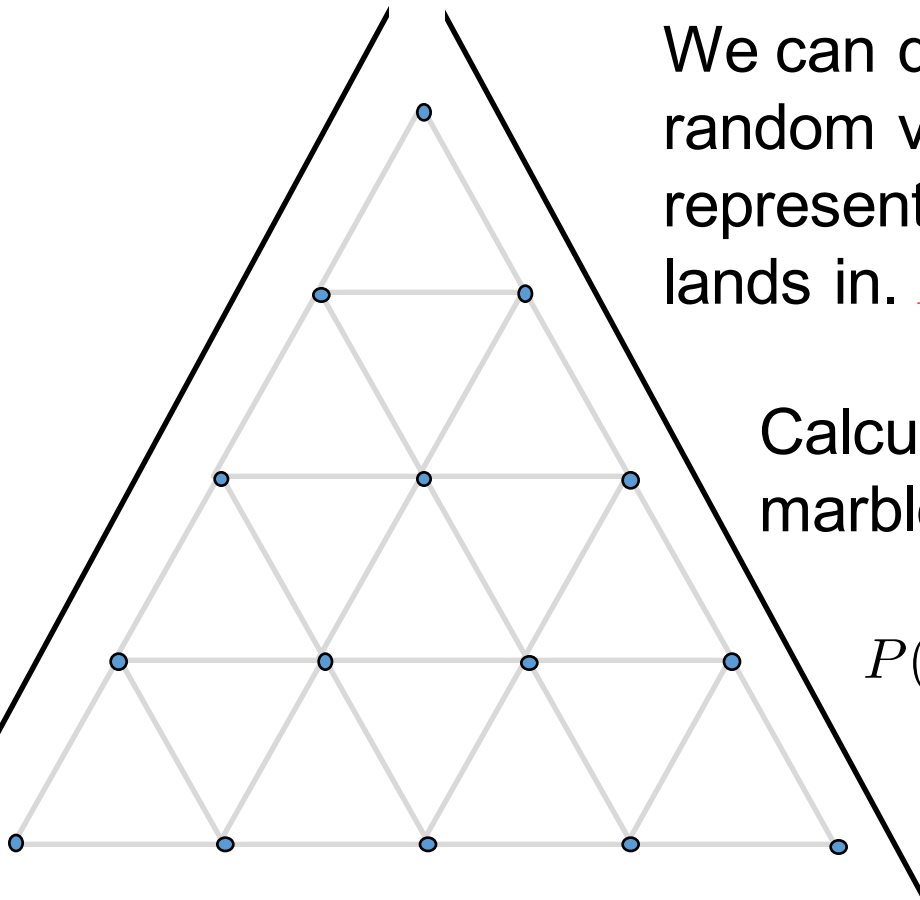
Galton Board



We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$

Calculate the probability of a marble landing in a bucket.

Galton Board



We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$

Calculate the probability of a marble landing in a bucket.

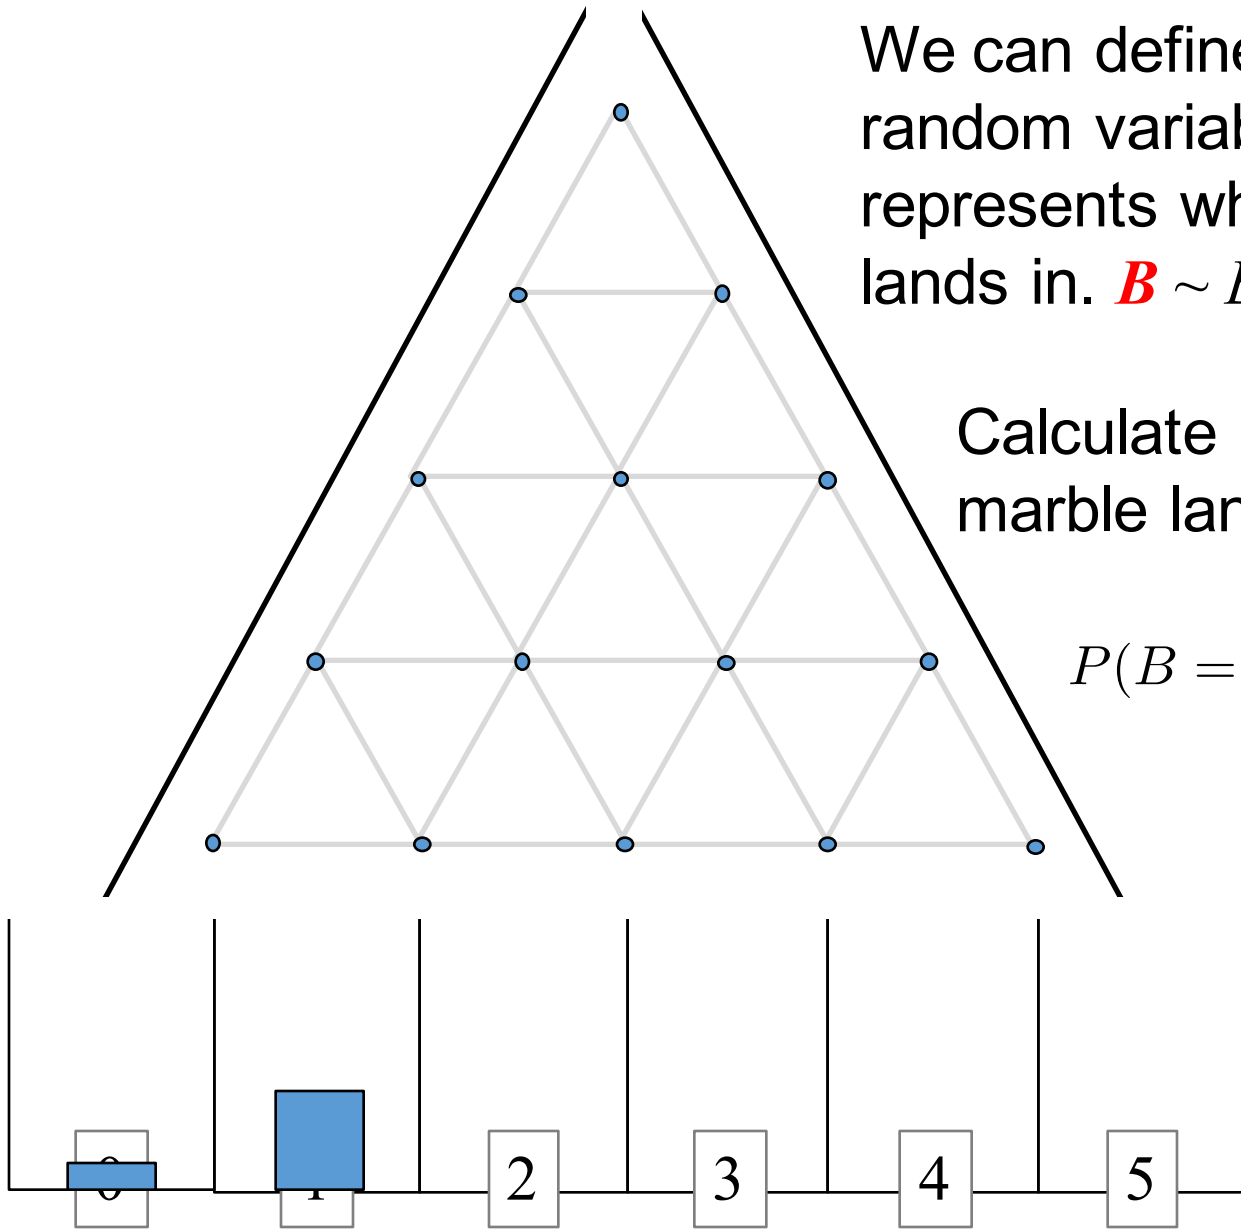
$$P(B = 0) = \binom{5}{0} \frac{1}{2}^5 \approx 0.03$$

Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$

Calculate the probability of a marble landing in a bucket.

$$P(B = 1) = \binom{5}{1} \frac{1}{2}^5 \approx 0.16$$

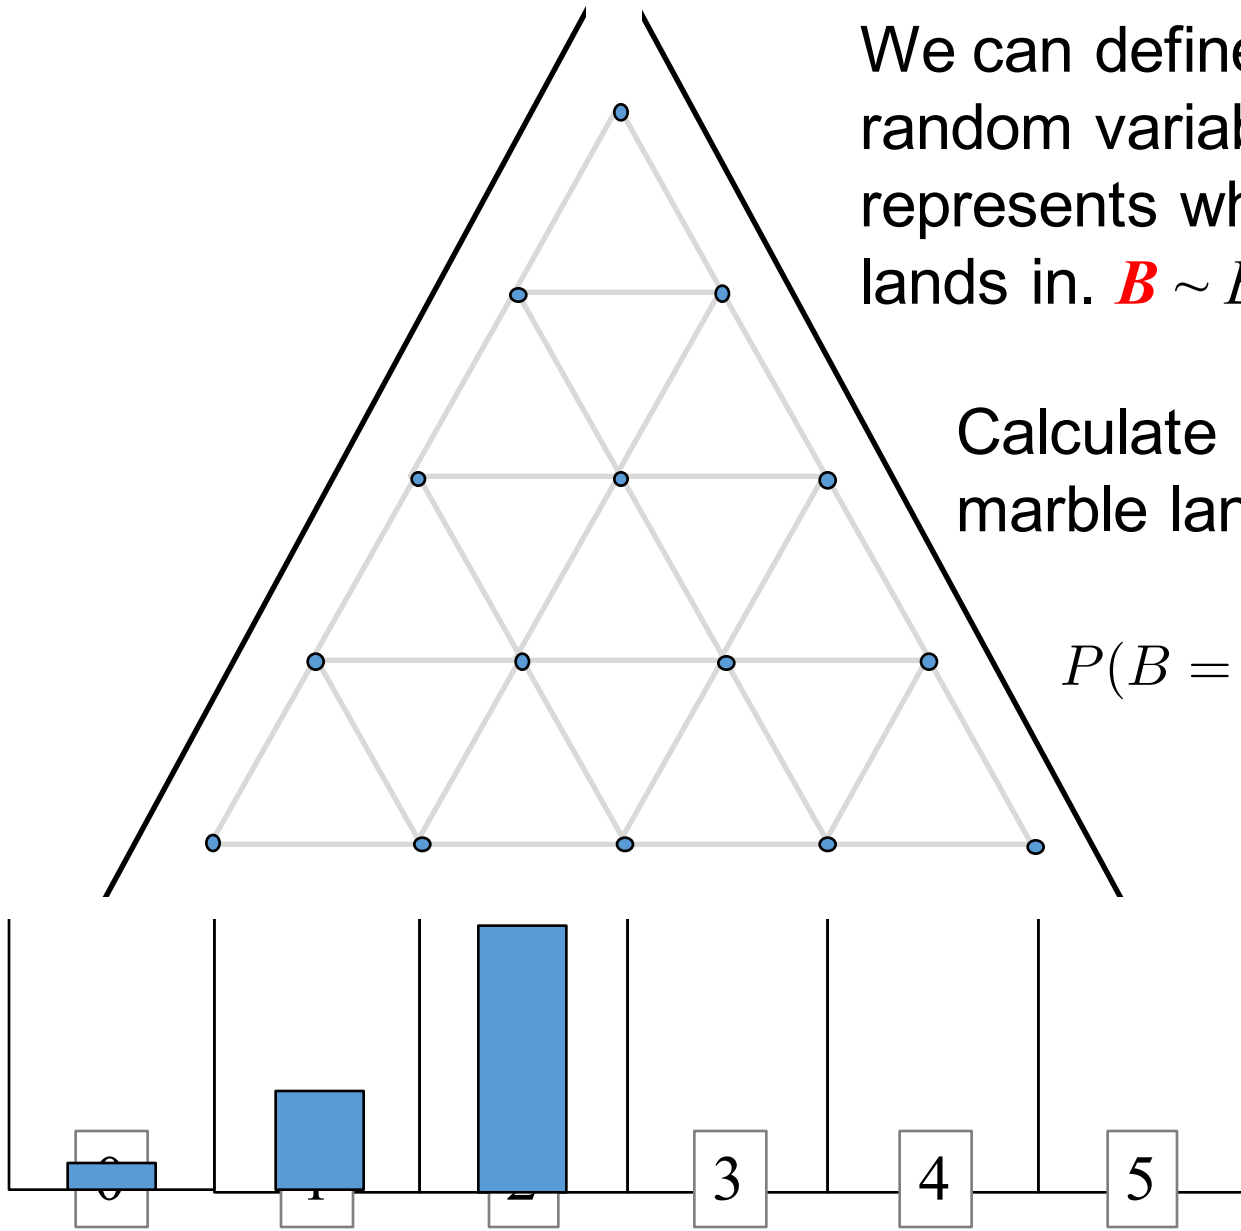


Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$

Calculate the probability of a marble landing in a bucket.

$$P(B = 2) = \binom{5}{2} \frac{1}{2}^5 \approx 0.31$$

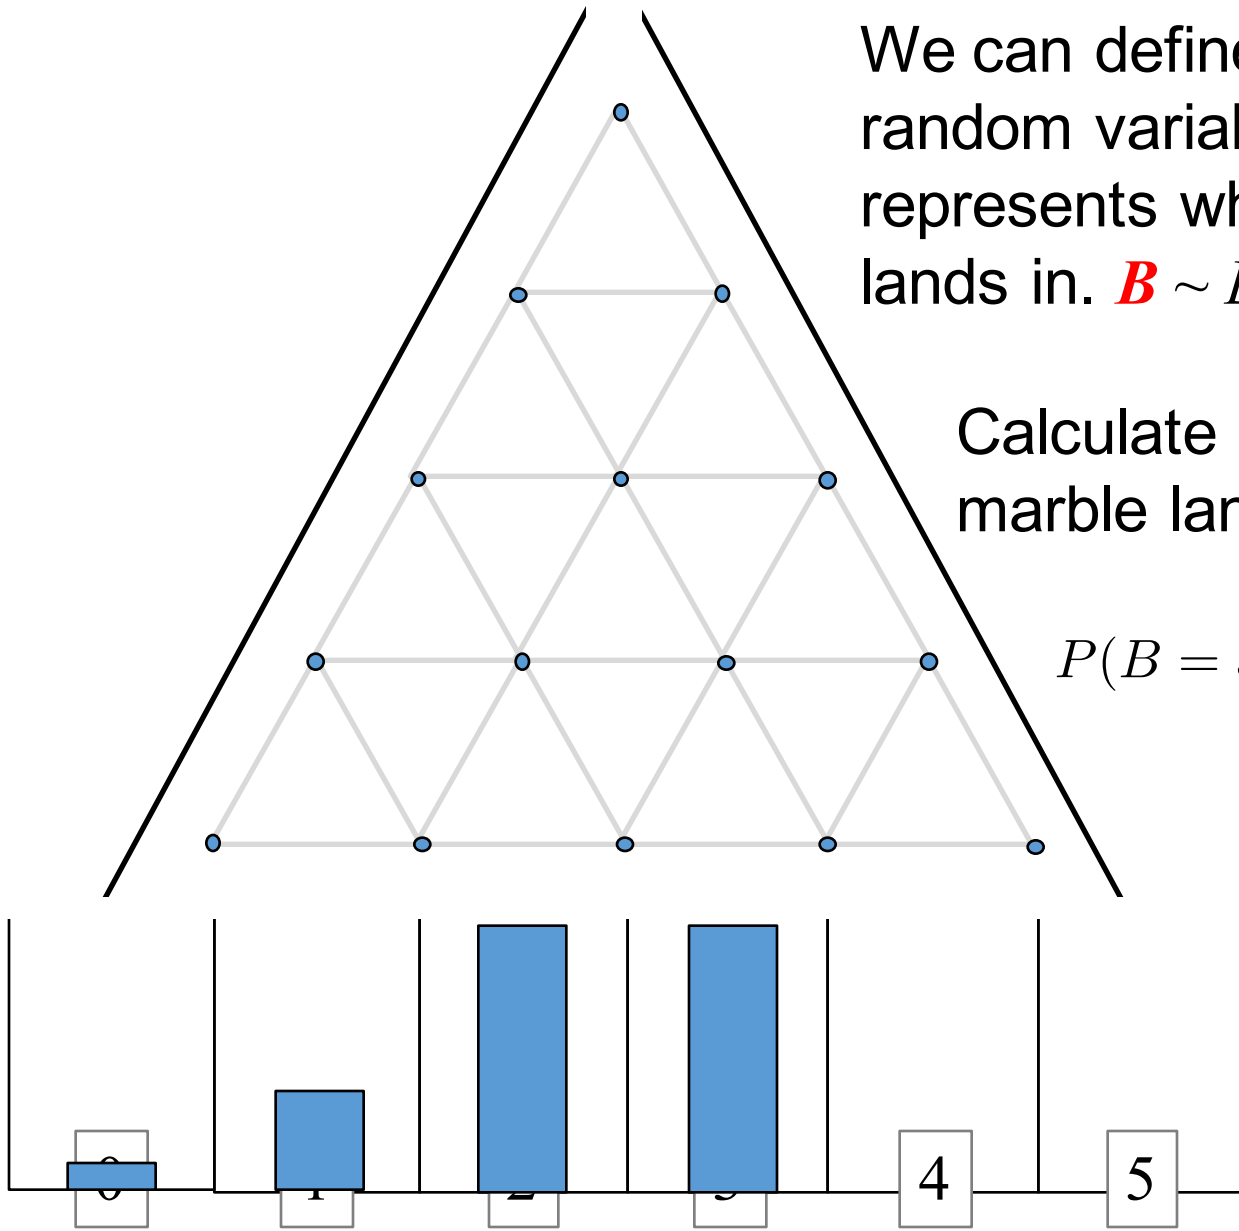


Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$

Calculate the probability of a marble landing in a bucket.

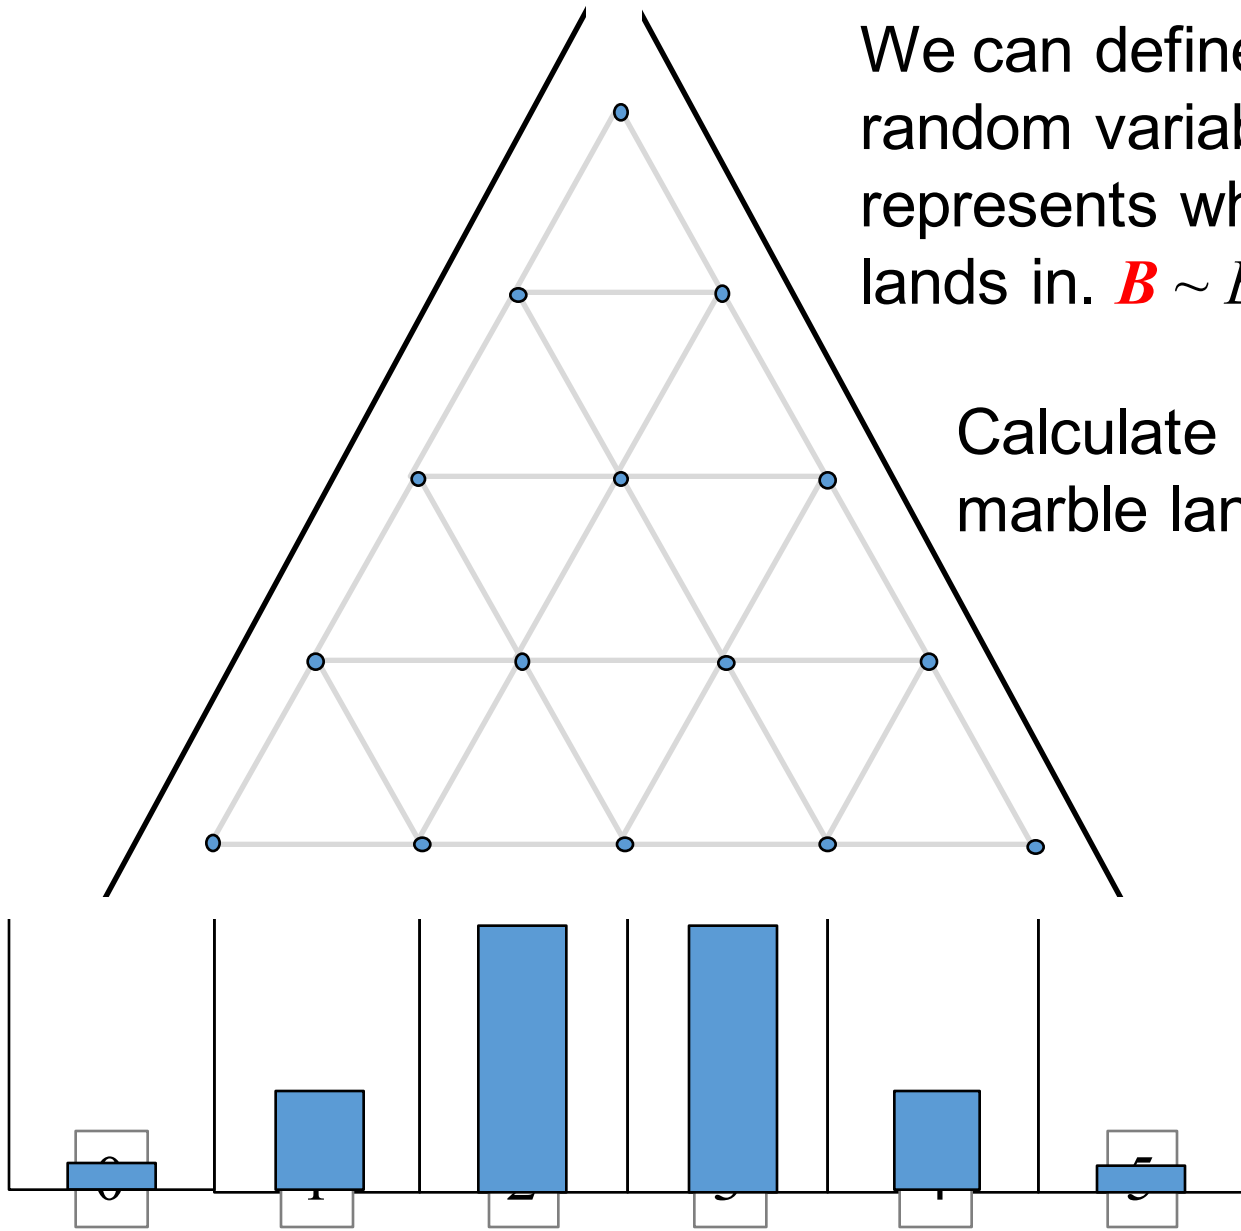
$$P(B = 3) = \binom{5}{2} \frac{1}{2}^5 \approx 0.31$$



Galton Board

We can define an indicator random variable (B) which represents what bucket a marble lands in. $B \sim \text{Bin}(5, 0.5)$

Calculate the probability of a marble landing in a bucket.




PDF




<https://www.youtube.com/watch?v=p65aYYuAz-s>



PMF for $X \sim \text{Bin}(10, 0.5)$



PMF for $X \sim \text{Bin}(10, 0.3)$



Power of Your Vote

- Is it better to vote in small or large state?
 - Small: more likely your vote changes outcome
 - Large: larger outcome (electoral votes) if state swings
 - a ($= 2n$) voters equally likely to vote for either candidate
 - You are deciding $(a + 1)^{\text{st}}$ vote

$$P(2n \text{ voters tie}) = \binom{2n}{n} \left(\frac{1}{2}\right)^n \left(\frac{1}{2}\right)^n = \frac{(2n)!}{n!n!2^{2n}}$$


- Use Stirling's Approximation: $n! \approx n^{n+1/2} e^{-n} \sqrt{2\pi}$

$$P(2n \text{ voters tie}) \approx \frac{(2n)^{2n+1/2} e^{-2n} \sqrt{2\pi}}{n^{2n+1} e^{-2n} 2\pi 2^{2n}} = \frac{1}{\sqrt{n\pi}}$$

- Power = P(tie) * Elec. Votes = $\frac{1}{\sqrt{(a/2)\pi}} (ac) = \frac{c\sqrt{2a}}{\sqrt{\pi}}$
- Larger state = more power

Is Peer Grading Accurate Enough?

Looking ahead




Peer Grading on Coursera
HCI.

31,067 peer grades for
3,607 students.

Is Peer Grading Accurate Enough?

Looking ahead



1. Defined random variables for:
 - True grade (s_i) for assignment i
 - Observed (z_i^j) score for assign i
 - Bias (b_j) for each grader j
 - Variance (r_j) for each grader j
2. Designed a probabilistic model that defined the distributions for all random variables


$$s_i \sim \text{Bin}(\text{points}, \theta)$$

$$z_i^j \sim \mathcal{N}(\mu = s_i + b_j, \sigma = \sqrt{r_j})$$

Problem param
↙

Is Peer Grading Accurate Enough?

Looking ahead




1. Defined random variables for:
 - True grade (s_i) for assignment i
 - Observed (z_i^j) score for assign i
 - Bias (b_j) for each grader j
 - Variance (r_j) for each grader j
2. Designed a probabilistic model that defined the distributions for all random variables
3. Found the variable assignments that maximized the probability of our observed data



Inference or Machine Learning

Yes, With Probabilistic Modelling

Before:



After:



Voilà, c'est tout

