Lecture 2 Conditional and Discrete Probability

D. Alex Hughes

Math Camp

September 13, 2015

1. Conditional Probability
   - Introduction and Motivation
   - Conditional Probability Definition
   - Higher-Order Conditional Probability
   - Partitioning Sample Spaces
   - Bayes Theorem
   - Independence

2. Random Variables
   - Expectation of Discrete Random Variables

Conditional Probability Motivation

Often, even without knowing it, we update our assessment of events based on other events having occurred or not.

- How likely were you to have been accepted to *any* PhD program? Did you internal assessment change when you received your first acceptance letter?
- How likely is it that you're eaten by a shark?
- How likely is it that a 10 minute conversation with a gay canvasser *fundamentally* changes your beliefs about marriage equality?

Conditional Probability Motivation

Often, even without knowing it, we update our assessment of events based on other events having occurred or not.

- How likely were you to have been accepted to *any* PhD program? Did you internal assessment change when you received your first acceptance letter?
- How likely is it that you're eaten by a shark?
- How likely is it that a 10 minute conversation with a gay canvasser *fundamentally* changes your beliefs about marriage equality?
  - What if it were published in Nature?

Conditional Probability Motivation

Often, even without knowing it, we update our assessment of events based on other events having occurred or not.

- How likely were you to have been accepted to *any* PhD program? Did you internal assessment change when you received your first acceptance letter?
- How likely is it that you're eaten by a shark?
- How likely is it that a 10 minute conversation with a gay canvasser *fundamentally* changes your beliefs about marriage equality?
  - What if it were published in Nature?
  - What if it were no longer published in Nature?

Preliminary Example

Example

Consider a fair die being tossed, with A defined as the event "6 appears." What is P(A)?

Suppose the die has already been tossed, but we haven't seen it. Now, what is P(A)?

Suppose the die has already been tossed. The shooter who threw it won't tell us what number came up, but does tell us that **B**, "It is an even number," has occurred. What is P(A)?

Preliminary Example

Example

Consider a fair die being tossed, with A defined as the event "6 appears." What is P(A)?

Suppose the die has already been tossed, but we haven't seen it. Now, what is P(A)?

Suppose the die has already been tossed. The shooter who threw it won't tell us what number came up, but does tell us that **B**, "It is an even number," has occurred. What is P(A)?

Answer

If we know no more information, then P(A) is the same $= 1/6$.
However, if we know that B has occurred, we now know that we have a $P(A) = 1/3$. Effectively, by knowing B, the sample space shrunk by $1/2$.
$\Omega = \{2, 4, 6\}$.

Conditional Probability Definition

Definition

Let A and B be any two events defined on Ω such that $P(B) > 0$. The conditional probability of A, given that B has already occurred, is written $P(A|B)$ and is given by

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

Conditional Probability Definition

Definition

Let A and B be any two events defined on Ω such that $P(B) > 0$. The conditional probability of A, given that B has already occurred, is written $P(A|B)$ and is given by

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

Example

A card is drawn from a poker deck. What is the probability that it is a club, given that the card is a king?

Conditional Probability Definition

Definition

Let A and B be any two events defined on Ω such that $P(B) > 0$. The conditional probability of A, given that B has already occurred, is written $P(A|B)$ and is given by

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

Example

A card is drawn from a poker deck. What is the probability that it is a club, given that the card is a king?

Answer

Intuitively, the answer is $1/4$. There are four suits of king, that the card drawn is a club is equally likely. And so,
$P(C|K) = P(C \cap K)/P(K) = (1/52)/(4/52) = \mathbf{1/4}$

Conditional Probability Examples

Example

Suppose we roll two dice. Let A = "the sum is 8" and B = "the first die is 3." What is the $P(B|A)$?

Conditional Probability Examples

Example

Suppose we roll two dice. Let A = "the sum is 8" and B = "the first die is 3." What is the $P(B|A)$?

Answer

A = {(2,6), (3,5), (4,4), (5,3), (6,2)}. So P(A) = 5/36.
B = {(3,1), (3,2), (3,3), (3,4), (3,5), (3,6)}.
But, since A has occured, $P(B \cap A) = \{(3,5)\} = 1/36$. P(A) = 5/36.
So $P(B|A) = 1/5$.

Conditional Probability Examples

Example

A person picks 13 cards out of a deck of 52. Let A_1 = "he has at least one ace." Let H = "he has the ace of hearts," and E_1 = "he receives exactly one ace."
What is the probability of $P(E_1|A_1)$?
What is the probability of $P(E_1|H)$? Are they equal? Which is larger?

Conditional Probability Examples

Example

A person picks 13 cards out of a deck of 52. Let $A_1 =$ "he has at least one ace." Let H = "he has the ace of hearts," and $E_1 =$ "he receives exactly one ace."
What is the probability of $P(E_1|A_1)$?
What is the probability of $P(E_1|H)$? Are they equal? Which is larger?

Answer

Let $E_0 =$ "has no ace." Then $p_0 = P(E_0) = \binom{48}{13}/\binom{52}{13}$ and $p_1 = P(E_1) = 4 * \binom{48}{12}/\binom{52}{13}$. Since $E_1 \subset A_1$ and $A_1 = E_0^c$, then, $P(E_1|A_1) = \frac{P(E_1)}{P(A_1)} = \frac{p_1}{1-p_0} \to \ldots$.

Conditional Probability Examples

Answer

And then since $E_1 \subset H$ means you get the ace of heart-s and no other ace: $P(E_1|H) = \frac{P(E_1 \cap H)}{P(H)} = \frac{\binom{48}{12}/\binom{52}{13}}{1/4} = p_1$. Finally, to conclude, we compare the probabilities and see that $P(E_1|A_1) = \frac{p_1}{1-p_0} > p_1 = P(E_1|H)$

Conditional Probability Examples

Answer

And then since $E_1 \subset H$ means you get the ace of heart-s and no other ace: $P(E_1|H) = \frac{P(E_1 \cap H)}{P(H)} = \frac{\binom{48}{12}/\binom{52}{13}}{1/4} = p_1$. Finally, to conclude, we compare the probabilities and see that $P(E_1|A_1) = \frac{p_1}{1-p_0} > p_1 = P(E_1|H)$

Example

Suppose we draw two cards from a poker deck. What are the chances the second card we draw is a Spade, given the first card we drew was a Spade?

Conditional Probability Examples

Answer

And then since $E_1 \subset H$ means you get the ace of heart-s and no other ace: $P(E_1|H) = \frac{P(E_1 \cap H)}{P(H)} = \frac{\binom{48}{12}/\binom{52}{13}}{1/4} = p_1$. Finally, to conclude, we compare the probabilities and see that $P(E_1|A_1) = \frac{p_1}{1-p_0} > p_1 = P(E_1|H)$

Example

Suppose we draw two cards from a poker deck. What are the chances the second card we draw is a Spade, given the first card we drew was a Spade?

Answer

Let A = "the first card is a spade" and B = "the second card is a spade." Then $P(A) = 1/13$, and to compute $P(B|A)$ note that if A has occurred then only $12/51$ remaining cards are spades. Then,
$P(A \cap B) = P(A)P(B|A) = \frac{13}{52}\frac{12}{51}$

More Conditional Probability Examples (!?)

Example

Two events A and B are defined such that:

- The probability A occurs but B does not is 0.2
- The probability B occurs but A does not is 0.1
- The probability that neither occurs is 0.6

What is $P(A|B)$?

More Conditional Probability Examples (!?)

Example

Two events A and B are defined such that:

- The probability A occurs but B does not is 0.2
- The probability B occurs but A does not is 0.1
- The probability that neither occurs is 0.6

What is P(A|B)?

Answer

$P((A \cup B)^C) = 0.6$; therefore $P(A \cup B) = 0.4$.
$0.4 = P(A \cap B^C) + P(B \cap A^C) + P(A \cap B) \rightarrow P(A \cap B) = 0.1$;

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A \cap B)}{P(A \cap B) + P(B \cap A^C)}$$

$$= \frac{0.1}{0.1 + 0.1} = 0.5$$

Can't get enough examples?

Example

An election is occuring, and you are consulting for a candidate. In the primary, she faces a chump opponent that she will beat 80% of the time. If she wins the primary, then she faces a stronger candidate and will win only 40% of the time. The election proceeds in the following way: *primary* first, and then *general* second. What is the probability of being elected?

Can't get enough examples?

Example

An election is occuring, and you are consulting for a candidate. In the primary, she faces a chump opponent that she will beat 80% of the time. If she wins the primary, then she faces a stronger candidate and will win only 40% of the time. The election proceeds in the following way: *primary* first, and then *general* second. What is the probability of being elected?

Answer

If A and B are the events of victory in the first and second games then $P(A) = 0.8$ and $P(B) = 0.4$, so the probability of winning the tournament is $P(A \cap B) = P(A)P(B|A) = (0.8)(0.4) = 0.32$.

Manipulating Conditional Probability

As in other identities, the terms in the Conditional Probability are manipulable.

Definition

Most useful is that by multiplying both sides of the Conditional Probability statement by $P(B)$ generates the statement:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

$$\mathbf{P(B)}P(A|B) = \mathbf{P(B)}\frac{P(A \cap B)}{P(B)}$$

$$P(A|B)P(B) = P(A \cap B)$$

Manipulating Conditional Probability Example

Example

Imagine you're out for a surf. You know that the *big kahuna* is out there, and that there is a 10% chance of it rolling through. You know, if you paddle for the wave, there is a 20% chance that you *make it*. What are the chances that you *make it*?

Manipulating Conditional Probability Example

Example

Imagine you're out for a surf. You know that the *big kahuna* is out there, and that there is a 10% chance of it rolling through. You know, if you paddle for the wave, there is a 20% chance that you *make it*. What are the chances that you *make it*?

Kahuna Answer

We're really looking for $P(A \cap B)$, that the wave comes and you catch it. $P(B) = 0.1$, $P(A|B) = 0.2$, and so, $P(A \cap B)$ is just $P(A|B)P(B) \rightarrow 0.2 * 0.1 = 0.02$.

Higher Order Conditional Probability

Consider $P(A \cap B \cap C)$. Substitute D for $P(A \cap B)$ Then,

$$\begin{aligned} P(A \cap B \cap C) &= P(D \cap C) \\ &= P(C|D)P(D) \\ &= P(C|A \cap B)P(A \cap B) \\ &= P(C|A \cap B)P(B|A)P(A) \end{aligned}$$

Definition

More generally,

$$\begin{aligned} P(A_1 \cap A_2 \cap \cdots \cap A_n) &= P(A_n|A_1 \cap A_2 \cap \cdots \cap A_{n-1}) \\ &* \; P(A_{n-1}|A_1 \cap A_2 \cap \cdots \cap A_{n-2}) * \cdots \\ &* \; P(A_2|A_1) * P(A_1) \end{aligned}$$

Example

An urn has five white chips, four black chips, and three red chips. Four chips are drawn sequentially *without replacement*. What is the probability of drawing in the following sequence: White, Red, White, Black?

Example

An urn has five white chips, four black chips, and three red chips. Four chips are drawn sequentially *without replacement*. What is the probability of drawing in the following sequence: White, Red, White, Black?
Hint: Let A be the event that you draw a white chip first. Let B be the event that you draw a red chip second, and so on.

From Larsen and Marx

Answer

$$\begin{aligned} P(A \cap B \cap C \cap D) &= P(D|A \cap B \cap C) \cdot P(C|A \cap B) \cdot P(B|A) \cdot P(A) \\ &= \frac{4}{9} \cdot \frac{4}{10} \cdot \frac{3}{11} \cdot \frac{5}{12} \\ &= \frac{240}{11880} \\ &= 0.02 \end{aligned}$$

Two-Stage Experiments

Example

An urn contains 5 red and 10 black balls. Draw two balls from the urn without replacement. What is the probability the second ball is red?

Two-Stage Experiments

Example

An urn contains 5 red and 10 black balls. Draw two balls from the urn without replacement. What is the probability the second ball is red?

Answer

$$
\begin{aligned}
P(R_2) &= P(R_2 \cap R_1) + P(R_2 \cap B_1) \\
&= P(R_2|R_1)P(R_1) + P(R_2|B_1)P(B_1) \\
&= (1/3)(4/14) + (2/3)(5/14) = 14/42 = 1/3
\end{aligned}
$$

Partitioned Spaces

Theorem

Let $\{A_i\}_{i=1}^n$ *be a set of events defined over* Ω *such that:*

1. $\Omega = \bigcup_{i=1}^{n} A_i,$
2. $A_i \cap A_j = \emptyset$, *for* $i \neq j$ *and*
3. $P(A_i) > 0 \forall i.$

Then, for any event B *defined on* Ω*,*

$$P(B) = \sum_{i=1}^{n} P(B|A_i)P(A_i)$$

Proof of Partitioning

Proof

The A_i's partition Ω, and so

$$B = (B \cap A_1) \cup (B \cap A_2) \cup \cdots \cup (B \cap A_n), \text{and so,}$$

$$P(B) = P(B \cap A_1) + P(B \cap A_2) + \cdots + P(B \cap A_n).$$

But each $P(B \cap A_i)$ can be written as $P(B|A_i)P(A_i)$, and the result follows.

A Political Example

Example

It is republican primary season again [[yay!]]. Hillary (or Sanders) are the presumptive Democratic candidates, while the Republican party elite are trying to organize their "contenders."

- Ben Carson has a 40% chance of winning, Trump a 35% chance, and, Fiorina a 25% chance.
- The most recent Fox News/Rasmussen poll gives Ben Carson a 45% chance of beating Obama in the general election, Trump a 35% chance, and Fiorina a 20% chance.

What is the probability Hillary will be re-elected?

A Political Example Answer

Answer

- Let B denote the event Hillary wins the election. Let A_{BC}, A_T, A_F denote the events that Ben Carson, Trump, and Fiorina with the primary, respectively.
- $P(A_{BC}) = 0.40$; $P(A_T) = 0.35$; and, $P(A_F) = 0.25$, and
- $P(B|A_{BC}) = 0.55$; $P(B|A_T) = 0.65$; and, $P(B|A_F) = 0.80$. So,

$$\begin{aligned} P(B) &= P(\text{Hillary Wins}) \\ &= P(B|A_{BC})P(A_{BC}) + P(B|A_T)P(A_S) + P(B|A_F)P(A_F) \\ &= (0.55)(0.4) + (0.65)(0.35) + (0.80)(0.25) \\ &= \mathbf{58.25}\%. \end{aligned}$$

Bayes Theorem

Theorem

Let $\{A_i\}_{i=1}^{n}$ *be a set of n events, each with positive probability, that partitions* Ω *in such a way that* $\bigcup_{i=1}^{n} A_i = \Omega$*, and* $A_i \cap A_j = \emptyset$ *for* $i \neq j$*. Then, for any event B also defined on* Ω *where* $P(B) > 0$*,*

$$P(A_j|B) = \frac{P(B|A_j)P(A_j)}{\sum_{i=1}^{n} P(B|A_i)P(A_i)}$$

Bayes Theorem

Theorem

Let $\{A_i\}_{i=1}^n$ be a set of n events, each with positive probability, that partitions Ω in such a way that $\bigcup_{i=1}^n A_i = \Omega$, and $A_i \cap A_j = \emptyset$ for $i \neq j$. Then, for any event B also defined on Ω where $P(B) > 0$,

$$P(A_j|B) = \frac{P(B|A_j)P(A_j)}{\sum_{i=1}^{n} P(B|A_i)P(A_i)}$$

Proof

$$P(A_j|B) = \frac{P(A_j \cap B)}{P(B)} = \frac{P(B|A_j)P(A_j)}{P(B)}$$

$$P(B) = \sum_{i=1}^{n} P(B|A_i)P(A_i)$$

Simple Example

Example

What is $P(A_3|B)$? $P(A_1|B)$? $P(A_1 \cup A_2|B)$?

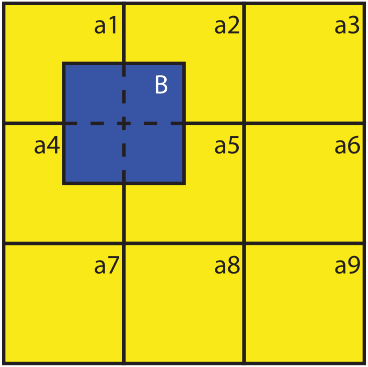

Medical Example

Example

Web M.D. has changed the face of medicine. Now patients are more informed than ever before, which is a major aid to doctors, just ask them. In reality, people now manufacture beliefs that they, with certainty, have exceedingly rare diseases – Ague, fibromyalgia, etc – and demand costly procedures.

Stan goes to his family doc and demands a tests for "Jumping Frenchmen of Maine." Rates of this disease are 1/10,000. The test correctly identifies the presence of JFoM in 90% of cases, and gives false positives in 1/1000 cases. What is the probability he has JFoM, given a positive test result?

Medical Example Answer

Answer

Let B denote the case that a test says he has the disease, A_1 the condition he actually has the disease, and A_2 the case he doesn't have the disease.

$P(B|A_1) = 0.90$ – Correct Positive
$P(B|A_2) = P(B|A_1^C) = 0.001$ – False Positive

We're looking for $P(A_1|B)$ – the probability he has the disease, given the test says he does.

$$
\begin{aligned}
P(A_1|B) &= \frac{P(B|A_1)P(A_1)}{P(B|A_1)P(A_1) + P(B|A_2)P(A_2)} \\
&= \frac{(0.9)(0.0001)}{(0.9)(0.0001) + (0.001)(0.9999)} \\
&= \mathbf{0.08}
\end{aligned}
$$

Feel the _urn

Example

Urn I contains two white chips and one red chip; urn II has one white chip and two red chips. One chip is drawn at random from urn I and placed in urn II. Then one chip is drawn from urn II. Suppose a red chip is drawn from urn II. What is the probability that the chip transferred was white?

Feel the _urn Answer

Answer

Let B = "Draw a red chip from urn II", A_1 = "Drew a red chip from urn I", and A_2 = "Drew a white chip from urn I." We then want to identify $P(A_2|B)$.

$$\begin{aligned} P(A_2|B) &= \frac{P(B|A_2)P(A_2)}{P(B|A_1)P(A_1)+P(B|A_2)P(A_2)} \\ &= \frac{(0.5)(2/3)}{(0.75)(1/3)+(0.5)(2/3)} \\ &= \frac{4/12}{3/12+4/12} \\ &= \frac{\mathbf{4}}{\mathbf{7}}. \end{aligned}$$

Independence

Two events, A and B are independent if the occurence of A has no influence of the probability of the occurence of B.

Definition

A and B are independent if $P(A \cap B) = P(A)P(B)$.

In practice, you will have to be told, or otherwise test that your data are independent.

Example

Draw two cards from a deck. Let A = "first card is a spade" and B = "second card is a spade." What is $P(A)?P(B)?P(A \cap B)?$ with and without replacement.

Introduction

To this point, we have avoided discussion of Random Variables. No longer.

Introduction

To this point, we have avoided discussion of Random Variables. No longer.

Definition

A *random variable* is a function whose domain is a sample space and whose range is some set of real numbers.

If a random variable is denoted X and has as its domain the sample space $\Omega = \{\omega_1, \omega_2, ..., \omega_n\}$, then $X(\omega_j)$ is the value of X at element ω_j.

That is, $X(\omega_j)$ is the number that the function rule (X) assigns to element $\omega_j \in \Omega$.

Random variables are exceedingly useful to store the values of stochastic processes. Most of the time, their use is intuitive, though this is not universally the case.

Vote Buying Example

Example

There are 100 senators. Each of these senators could potentially be involved in a vote buying scandal. There are 2^{100} possible ways the senate could/not be corrupt – {not, not, not, corrupt, corrupt, not, ... }

But we don't really care *who* – we're not Fox News or MSNBC – we just want the data.

So we take the values and store them, *in a random variable*. In this case, we let the RV be 1 if a senator was corrupt, and 0 otherwise. This way, $\Omega = [0, 100]$, not $[0, 2^{100}]$.

Preliminaries

Think back to our initial example of a discrete distribution: flipping a coin three times.

- Recall, what was the:
  - Experiment
  - Sample Outcome
  - Sample Space
  - Event
- If we repeated this experiment many times, we might want a convenient placeholder for our results.
- Create a random variable, denoted X, that takes the value 1 for each experiment in which the event occurs and takes the value 0 for each experiment in which the event doesn't occur.
- Then $\Omega = [0, n]$, not $[0, 2^n]$

Distributional Statistics

There are several measurements we're interested with distributions:

1. Central Tendency:
   - Mean: $\left(\frac{1}{N}\right)\sum_{i=1}^{N} x_i$
   - Median: Select $\left(\frac{N}{2}\right)$ from Ordered Set
   - Mode: Maximally occurrent observation. Useful for nominal-level variables.
2. Dispersion:
   - Ordinal: Comparison of Median and Mode
   - Interval and Ratio: Variance
     - **Variance**: $\sigma^2 = \frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^2$
     - **Standard Deviation**: $\sigma = (\sigma^2)^{\frac{1}{2}}$
3. Skewness: $\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^3$
4. Kurtosis: $\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^4$

Expectation of Discrete Random Variables

Definition

Let X be a numerically-valued (i.e. not True/False, Blue/Green) random variable across Ω with distribution function $f(x)$.

Then, the *expected value*, denoted E[X] or μ, is defined as

$$E[X] = \sum_{x \in \Omega} x \cdot f(x).$$

Here, x refers to the value of the outcome, and $f(x)$ the probability function that maps that outcome. In conversation we refer to the expectation as the *mean*, and can be thought of like the center of gravity of the distribution.

Together with a measure of *dispersion* (and perhaps the other higher-order moments), the mean is a core description we are interested in for any variable.

Expectation of Discrete Random Varialbles

Example

Consider the game of roulette. After bets are placed, the croupier spins, the marble lands, and one of 38 (00, 0, 1, ... 36) to be the winner. Any of the numbers is equally likely. 00 and 0 are neither odd nor even. We bet \$1 on "odd" at even money. What is our expected earnings for each bet?

Expectation of Discrete Random Varialbles

Example

Consider the game of roulette. After bets are placed, the croupier spins, the marble lands, and one of 38 (00, 0, 1, ... 36) to be the winner. Any of the numbers is equally likely. 00 and 0 are neither odd nor even. We bet \$1 on "odd" at even money. What is our expected earnings for each bet?

Answer

$$
\begin{aligned}
p_X(1) &= P(X=1) = \frac{18}{38} = \frac{9}{19} \\
p_X(-1) &= P(X=-1) = \frac{20}{38} = \frac{10}{19} \\
\text{``expected winnings''} &= \$1\left(\frac{9}{19}\right) - \$1\left(\frac{10}{19}\right) \\
&= -\frac{\mathbf{\$1}}{\mathbf{19}}
\end{aligned}
$$

Expectation of Discrete Random Variables

Example

You roll 12 times a fair 3-sided die (Dragonslayer?) that has the values 1, 2, and 4 on the sides. After each roll, you store the number showing in a random variable Q. If you had to place a bet, on what number would you bet is the sum of the rolls?

Expectation of Discrete Random Variables

Example

You roll 12 times a fair 3-sided die (Dragonslayer?) that has the values 1, 2, and 4 on the sides. After each roll, you store the number showing in a random variable Q. If you had to place a bet, on what number would you bet is the sum of the rolls?

Answer

Our expectation for the sum is going to n*E[Q].

$$\begin{aligned} E[Q] &= \sum_{i=1}^{10} x \cdot f(x) \\ &= (1)(1/3) + 2(1/3) + (4)(1/3) \\ &= 7/3 \end{aligned}$$

And so the number most likely to come up is $12 * 7/3 = \mathbf{28}$.

Why does the house want 7s in craps?

Example

You throw two dice and add the values. What is the expected value? Why does the house want 7s?

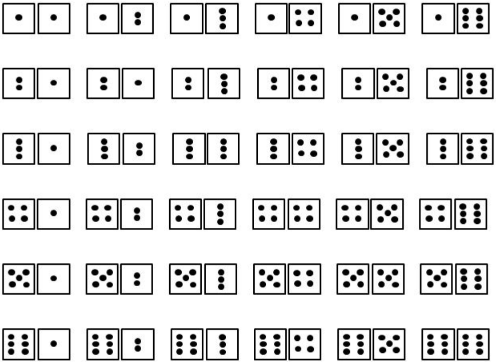

Answer

Answer

The house wants to play the best odds – it wants to make you lose with the highest frequency.

$$
\begin{aligned}
E[X] &= \sum x \cdot f(x) \\
&= (1/36)(2) + (2/36)(3) + (3/36)(4) + (4/36)(5) + (5/36)(6) \\
&+ (6/36)(7) + (5/36)(8) + (4/36)(9) + (3/36)(10) + (2/36)(11) \\
&+ (1/36)(12) = 252/36 = \mathbf{7}
\end{aligned}
$$

Answer

Answer

The house wants to play the best odds – it wants to make you lose with the highest frequency.

$$
\begin{aligned}
E[X] &= \sum x \cdot f(x) \\
&= (1/36)(2) + (2/36)(3) + (3/36)(4) + (4/36)(5) + (5/36)(6) \\
&+ (6/36)(7) + (5/36)(8) + (4/36)(9) + (3/36)(10) + (2/36)(11) \\
&+ (1/36)(12) = 252/36 = \mathbf{7}
\end{aligned}
$$

Example

Suppose X is a binomial random variable with $p = \frac{5}{9}$ and $n = 3$. What is the expected value of X?

Binomial RV Answer

Answer

$$p_X(k) = \binom{3}{k}\left(\frac{5}{9}\right)^k\left(\frac{4}{9}\right)^{3-k}, k = 0, 1, 2, 3$$

Binomial RV Answer

Answer

$$p_X(k) = \binom{3}{k}\left(\frac{5}{9}\right)^{k}\left(\frac{4}{9}\right)^{3-k}, k = 0, 1, 2, 3$$

$$E[X] = \sum_{k=0}^{3} k \cdot \binom{3}{k}\left(\frac{5}{9}\right)^{k}\left(\frac{4}{9}\right)^{3-k}$$

Binomial RV Answer

Answer

$$p_X(k) = \binom{3}{k}\left(\frac{5}{9}\right)^{k}\left(\frac{4}{9}\right)^{3-k}, k = 0, 1, 2, 3$$

$$E[X] = \sum_{k=0}^{3} k \cdot \binom{3}{k}\left(\frac{5}{9}\right)^{k}\left(\frac{4}{9}\right)^{3-k}$$

$$= (0)\left(\frac{64}{729}\right) + (1)\left(\frac{240}{729}\right) + \left(\frac{300}{729}\right) + (3)\left(\frac{125}{729}\right)$$

$$= \frac{1215}{729}$$

$$= 3\left(\frac{5}{9}\right) \text{(remember this result, we will come back to this later)}$$

Make that Chedda'

Example

A fair coin is flipped until the first tail appears. You win $\$2^k$ where k is the toss on which the first tail appears. How much are you willing to pay to play this game?

Make that Chedda'

Example

A fair coin is flipped until the first tail appears. You win \$$2^k$ where k is the toss on which the first tail appears. How much are you willing to pay to play this game?

Answer

This is the St. Petersburg paradox.

$$p_X(2^k) = P(X = 2^k) = \frac{1}{2^k}, k = 1, 2, ...$$

And so,

$$E[X] = \sum_{\forall k} 2^k p_X(2^k) = \sum_{k=1}^{\infty} 2^k \cdot \frac{1}{2^k} = 1 + 1 + \cdots$$