

COPYRIGHT NOTICE:

Herbert Gintis: Game Theory Evolving

is published by Princeton University Press and copyrighted, © 2009, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@press.princeton.edu

1

Probability Theory

Doubt is disagreeable, but certainty is ridiculous.

Voltaire

1.1 Basic Set Theory and Mathematical Notation

A *set* is a collection of objects. We can represent a set by enumerating its objects. Thus,

$$A = \{1, 3, 5, 7, 9, 34\}$$

is the set of single digit odd numbers plus the number 34. We can also represent the same set by a formula. For instance,

$$A = \{x | x \in \mathbf{N} \wedge (x < 10 \wedge x \text{ is odd}) \vee (x = 34)\}.$$

In interpreting this formula, \mathbf{N} is the set of natural numbers (positive integers), “|” means “such that,” “ \in ” means “is a element of,” \wedge is the logical symbol for “and,” and \vee is the logical symbol for “or.” See the table of symbols in chapter 14 if you forget the meaning of a mathematical symbol.

The subset of objects in set X that satisfy property p can be written as

$$\{x \in X | p(x)\}.$$

The *union* of two sets $A, B \subset X$ is the subset of X consisting of elements of X that are in *either* A or B :

$$A \cup B = \{x | x \in A \vee x \in B\}.$$

The *intersection* of two sets $A, B \subset X$ is the subset of X consisting of elements of X that are in *both* A or B :

$$A \cap B = \{x | x \in A \wedge x \in B\}.$$

If $a \in A$ and $b \in B$, the *ordered pair* (a, b) is an entity such that if $(a, b) = (c, d)$, then $a = c$ and $b = d$. The set $\{(a, b) | a \in A \wedge b \in B\}$

is called the *product* of A and B and is written $A \times B$. For instance, if $A = B = \mathbf{R}$, where \mathbf{R} is the set of real numbers, then $A \times B$ is the real plane, or the real two-dimensional vector space. We also write

$$\prod_{i=1}^n A_i = A_1 \times \dots \times A_n.$$

A *function* f can be thought of as a set of ordered pairs $(x, f(x))$. For instance, the function $f(x) = x^2$ is the set

$$\{(x, y) | (x, y \in \mathbf{R}) \wedge (y = x^2)\}$$

The set of arguments for which f is defined is called the *domain* of f and is written $\text{dom}(f)$. The set of values that f takes is called the *range* of f and is written $\text{range}(f)$. The function f is thus a subset of $\text{dom}(f) \times \text{range}(f)$. If f is a function defined on set A with values in set B , we write $f: A \rightarrow B$.

1.2 Probability Spaces

We assume a finite *universe* or *sample space* Ω and a set \mathcal{X} of subsets A, B, C, \dots of Ω , called *events*. We assume \mathcal{X} is closed under finite unions (if A_1, A_2, \dots, A_n are events, so is $\cup_{i=1}^n A_i$), finite intersections (if A_1, \dots, A_n are events, so is $\cap_{i=1}^n A_i$), and complementation (if A is an event so is the set of elements of Ω that are not in A , which we write A^c). If A and B are events, we interpret $A \cap B = AB$ as the event “ A and B both occur,” $A \cup B$ as the event “ A or B occurs,” and A^c as the event “ A does not occur.”

For instance, suppose we flip a coin twice, the outcome being HH (heads on both), HT (heads on first and tails on second), TH (tails on first and heads on second), and TT (tails on both). The sample space is then $\Omega = \{HH, TH, HT, TT\}$. Some events are $\{HH, HT\}$ (the coin comes up heads on the first toss), $\{TT\}$ (the coin comes up tails twice), and $\{HH, HT, TH\}$ (the coin comes up heads at least once).

The *probability* of an event $A \in \mathcal{X}$ is a real number $P[A]$ such that $0 \leq P[A] \leq 1$. We assume that $P[\Omega] = 1$, which says that with probability 1 *some* outcome occurs, and we also assume that if $A = \cup_{i=1}^n A_i$, where $A_i \in \mathcal{X}$ and the $\{A_i\}$ are disjoint (that is, $A_i \cap A_j = \emptyset$ for all $i \neq j$), then $P[A] = \sum_{i=1}^n P[A_i]$, which says that probabilities are additive over finite disjoint unions.

1.3 De Morgan's Laws

Show that for any two events A and B , we have

$$(A \cup B)^c = A^c \cap B^c$$

and

$$(A \cap B)^c = A^c \cup B^c.$$

These are called *De Morgan's laws*. Express the meaning of these formulas in words.

Show that if we write p for proposition "event A occurs" and q for "event B occurs," then

$$\text{not } (p \text{ or } q) \Leftrightarrow (\text{not } p \text{ and not } q),$$

$$\text{not } (p \text{ and } q) \Leftrightarrow (\text{not } p \text{ or not } q).$$

The formulas are also De Morgan's laws. Give examples of both rules.

1.4 Interocitors

An interocitor consists of two kramels and three trums. Let A_k be the event "the k th kramel is in working condition," and B_j is the event "the j th trum is in working condition." An interocitor is in working condition if at least one of its kramels and two of its trums are in working condition. Let C be the event "the interocitor is in working condition." Write C in terms of the A_k and the B_j .

1.5 The Direct Evaluation of Probabilities

THEOREM 1.1 *Given a_1, \dots, a_n and b_1, \dots, b_m , all distinct, there are $n \times m$ distinct ways of choosing one of the a_i and one of the b_j . If we also have c_1, \dots, c_r , distinct from each other, the a_i and the b_j , then there are $n \times m \times r$ distinct ways of choosing one of the a_i , one of the b_j , and one of the c_k .*

Apply this theorem to determine how many different elements there are in the sample space of

- a. the double coin flip

- b. the triple coin flip
- c. rolling a pair of dice

Generalize the theorem.

1.6 Probability as Frequency

Suppose the sample space Ω consists of a finite number n of equally probable elements. Suppose the event A contains m of these elements. Then the *probability of the event A* is m/n .

A second definition: Suppose an experiment has n distinct outcomes, all of which are equally likely. Let A be a subset of the outcomes, and $n(A)$ the number of elements of A . We define the *probability* of A as $P[A] = n(A)/n$.

For example, in throwing a pair of dice, there are $6 \times 6 = 36$ mutually exclusive, equally likely events, each represented by an ordered pair (a, b) , where a is the number of spots showing on the first die and b the number on the second. Let A be the event that both dice show the same number of spots. Then $n(A) = 6$ and $P[A] = 6/36 = 1/6$.

A third definition: Suppose an experiment can be repeated any number of times, each outcome being independent of the ones before and after it. Let A be an event that either does or does not occur for each outcome. Let $n_t(A)$ be the number of times A occurred on all the tries up to and including the t^{th} try. We define the *relative frequency* of A as $n_t(A)/t$, and we define the *probability of A* as

$$P[A] = \lim_{t \rightarrow \infty} \frac{n_t(A)}{t}.$$

We say two events A and B are *independent* if $P[A]$ does not depend on whether B occurs or not and, conversely, $P[B]$ does not depend on whether A occurs or not. If events A and B are independent, the probability that both occur is the product of the probabilities that either occurs: that is,

$$P[A \text{ and } B] = P[A] \times P[B].$$

For example, in flipping coins, let A be the event “the first ten flips are heads.” Let B be the event “the eleventh flip is heads.” Then the two events are independent.

For another example, suppose there are two urns, one containing 100 white balls and 1 red ball, and the other containing 100 red balls and 1

white ball. You do not know which is which. You choose 2 balls from the first urn. Let A be the event “The first ball is white,” and let B be the event “The second ball is white.” These events are not independent, because if you draw a white ball the first time, you are more likely to be drawing from the urn with 100 white balls than the urn with 1 white ball.

Determine the following probabilities. Assume all coins and dice are “fair” in the sense that H and T are equiprobable for a coin, and $1, \dots, 6$ are equiprobable for a die.

- At least one head occurs in a double coin toss.
- Exactly two tails occur in a triple coin toss.
- The sum of the two dice equals 7 or 11 in rolling a pair of dice.
- All six dice show the same number when six dice are thrown.
- A coin is tossed seven times. The string of outcomes is HHHHHHH.
- A coin is tossed seven times. The string of outcomes is HTHHTTH.

1.7 Craps

A roller plays against the casino. The roller throws the dice and wins if the sum is 7 or 11, but loses if the sum is 2, 3, or 12. If the sum is any other number (4, 5, 6, 8, 9, or 10), the roller throws the dice repeatedly until either winning by matching the first number rolled or losing if the sum is 2, 7, or 12 (“crapping out”). What is the probability of winning?

1.8 A Marksman Contest

In a head-to-head contest Alice can beat Bonnie with probability p and can beat Carole with probability q . Carole is a better marksman than Bonnie, so $p > q$. To win the contest Alice must win at least two in a row out of three head-to-heads with Bonnie and Carole and cannot play the same person twice in a row (that is, she can play Bonnie-Carole-Bonnie or Carole-Bonnie-Carole). Show that Alice maximizes her probability of winning the contest playing the better marksman, Carole, twice.

1.9 Sampling

The mutually exclusive outcomes of a random action are called *sample points*. The set of sample points is called the *sample space*. An event A is a subset of a sample space Ω . The event A is *certain* if $A = \Omega$ and

impossible if $A = \emptyset$ (that is, A has no elements). The *probability* of an event A is $P[A] = n(A)/n(\Omega)$, if we assume Ω is finite and all $\omega \in \Omega$ are equally likely.

- Suppose six dice are thrown. What is the probability all six die show the same number?
- Suppose we choose r object in succession from a set of n distinct objects a_1, \dots, a_n , each time recording the choice and returning the object to the set before making the next choice. This gives an ordered sample of the form (b_1, \dots, b_r) , where each b_j is some a_i . We call this *sampling with replacement*. Show that, in sampling r times with replacement from a set of n objects, there are n^r distinct ordered samples.
- Suppose we choose r objects in succession from a set of n distinct objects a_1, \dots, a_n , without returning the object to the set. This gives an ordered sample of the form (b_1, \dots, b_r) , where each b_j is some unique a_i . We call this *sampling without replacement*. Show that in sampling r times without replacement from a set of n objects, there are

$$n(n-1)\dots(n-r+1) = \frac{n!}{(n-r)!}$$

distinct ordered samples, where $n! = n \times (n-1) \times \dots \times 2 \times 1$.

1.10 Aces Up

A deck of 52 cards has 4 aces. A player draws 2 cards randomly from the deck. What is the probability that both are aces?

1.11 Permutations

A linear ordering of a set of n distinct objects is called a *permutation* of the objects. It is easy to see that the number of distinct permutations of $n > 0$ distinct objects is $n! = n \times (n-1) \times \dots \times 2 \times 1$. Suppose we have a deck of cards numbered from 1 to $n > 1$. Shuffle the cards so their new order is a random permutation of the cards. What is the average number of cards that appear in the “correct” order (that is, the k^{th} card is in the k^{th} position) in the shuffled deck?

1.12 Combinations and Sampling

The number of *combinations* of n distinct objects taken r at a time is the number of subsets of size r , taken from the n things without replacement. We write this as $\binom{n}{r}$. (In this case, we do not care about the order of the choices. For instance, consider the set of numbers $\{1,2,3,4\}$. The number of samples of size two without replacement = $4!/2! = 12$. These are precisely $\{12,13,14,21,23,24,31,32,34,41,42,43\}$. The combinations of the four numbers of size two (that is, taken two at a time) are $\{12,13,14,23,24,34\}$, or six in number. Note that $6 = \binom{4}{2} = 4!/2!$. A set of n elements has $n!/r!(n-r)!$ distinct subsets of size r . Thus, we have

$$\binom{n}{r} = \frac{n!}{r!(n-r)!}.$$

1.13 Mechanical Defects

A shipment of seven machines has two defective machines. An inspector checks two machines randomly drawn from the shipment, and accepts the shipment if neither is defective. What is the probability the shipment is accepted?

1.14 Mass Defection

A batch of 100 manufactured items is checked by an inspector, who examines 10 items at random. If none is defective, she accepts the whole batch. What is the probability that a batch containing 10 defective items will be accepted?

1.15 House Rules

Suppose you are playing the following game against the house in Las Vegas. You pick a number between one and six. The house rolls three dice, and pays you \$1,000 if your number comes up on one die, \$2,000 if your number comes up on two dice, and \$3,000 if your number comes up on all three dice. If your number does not show up at all, you pay the house \$1,000. At first glance, this looks like a *fair game* (that is, a game in which the expected payoff is zero), but in fact it is not. How much can you expect to win (or lose)?

1.16 The Addition Rule for Probabilities

Let A and B be two events. Then $0 \leq P[A] \leq 1$ and

$$P[A \cup B] = P[A] + P[B] - P[AB].$$

If A and B are disjoint (that is, the events are mutually exclusive), then

$$P[A \cup B] = P[A] + P[B].$$

Moreover, if A_1, \dots, A_n are mutually disjoint, then

$$P[\cup_i A_i] = \sum_{i=1}^n P[A_i].$$

We call events A_1, \dots, A_n a *partition* of the sample space Ω if they are mutually disjoint and exhaustive (that is, their union is Ω). In this case for any event B , we have

$$P[B] = \sum_i P[BA_i].$$

1.17 A Guessing Game

Each day the call-in program on a local radio station conducts the following game. A number is drawn at random from $\{1, 2, \dots, n\}$. Callers choose a number randomly and win a prize if correct. Otherwise, the station announces whether the guess was high or low and moves on to the next caller, who chooses randomly from the numbers that can logically be correct, given the previous announcements. What is the expected number $f(n)$ of callers before one guesses the number?

1.18 North Island, South Island

Bob is trying to find a secret treasure buried in the ground somewhere in North Island. According to local custom, if Bob digs and finds the treasure, he can keep it. If the treasure is not at the digging point, though, and Bob happens to hit rock, Bob must go to South Island. On the other hand, if Bob hits clay on North Island, he can stay there and try again. Once on South Island, to get back to North Island, Bob must dig and hit clay. If Bob hits rock on South Island, he forfeits the possibility of obtaining the treasure.

On the other hand, if Bob hits earth on South Island, he can stay on South Island and try again. Suppose q_n is the probability of finding the treasure when digging at a random spot on North Island, r_n is the probability of hitting rock on North Island, r_s is the probability of hitting rock on South Island, and e_s is the probability of hitting earth on South Island. What is the probability, P_n , that Bob will eventually find the treasure before he forfeits, if we assume that he starts on North Island?

1.19 Conditional Probability

If A and B are events, and if the probability $P[B]$ that B occurs is strictly positive, we define the *conditional probability* of A given B , denoted $P[A|B]$, by

$$P[A|B] = \frac{P[AB]}{P[B]}.$$

We say B_1, \dots, B_n are a *partition* of event B if $\cup_i B_i = B$ and $B_i B_j = \emptyset$ for $i \neq j$. We have:

- If A and B are events, $P[B] > 0$, and B implies A (that is, $B \subseteq A$), then $P[A|B] = 1$.
- If A and B are contradictory (that is, $AB = \emptyset$), then $P[A|B] = 0$.
- If A_1, \dots, A_n are a partition of event A , then

$$P[A|B] = \sum_{i=1}^n P[A_i|B].$$

- If B_1, \dots, B_n are a partition of the sample space Ω , then

$$P[A] = \sum_{i=1}^n P[A|B_i] P[B_i].$$

1.20 Bayes' Rule

Suppose A and B are events with $P[A], P[B], P[B^c] > 0$. Then we have

$$P[B|A] = \frac{P[A|B] P[B]}{P[A|B] P[B] + P[A|B^c] P[B^c]}.$$

This follows from the fact that the denominator is just $P[A]$, and is called *Bayes' rule*.

More generally, if B_1, \dots, B_n is a partition of the sample space and if $P[A], P[B_k] > 0$, then

$$P[B_k|A] = \frac{P[A|B_k]P[B_k]}{\sum_{i=1}^n P[A|B_i]P[B_i]}.$$

To see this, note that the denominator on the right-hand side is just $P[A]$, and the numerator is just $P[AB_k]$ by definition.

1.21 Extrasensory Perception

Alice claims to have ESP. She says to Bob, “Match me against a series of opponents in picking the high card from a deck with cards numbered 1 to 100. I will do better than chance in either choosing a higher card than my opponent or choosing a higher card on my second try than on my first.” Bob reasons that Alice will win on her first try with probability $1/2$, and beat her own card with probability $1/2$ if she loses on the first round. Thus, Alice should win with probability $(1/2) + (1/2)(1/2) = 3/4$. He finds, to his surprise, that Alice wins about $5/6$ of the time. Does Alice have ESP?

1.22 Les Cinq Tiroirs

You are looking for an object in one of five drawers. There is a 20% chance that it is not in any of the drawers, but if it is in a drawer, it is equally likely to be in each one. Show that as you look in the drawers one by one, the probability of finding the object in the next drawer rises if not found so far, but the probability of not finding it at all also rises.

1.23 Drug Testing

Bayes’ rule is useful because often we know $P[A|B]$, $P[A|B^c]$ and $P[B]$, and we want to find $P[B|A]$. For example, suppose 5% of the population uses drugs, and there is a drug test that is 95% accurate: it tests positive on a drug user 95% of the time, and it tests negative on a drug nonuser 95% of the time. Show that if an individual tests positive, the probability of his being a drug user is 50%. *Hint:* Let A be the event “is a drug user,” let “Pos” be the event “tests positive,” let “Neg” be the event “tests negative,” and apply Bayes’ rule.

1.24 Color Blindness

Suppose 5% of men are color-blind and 0.25% of women are color-blind. A person is chosen at random and found to be color-blind. What is the probability the person is male (assume the population is 50% female)?

1.25 Urns

A collection of $n + 1$ urns, numbered from 0 to n , each contains n balls. Urn k contains k red and $n - k$ white balls. An urn is chosen at random and n balls are randomly chosen from it, the ball being replaced each time before another is chosen. Suppose all n balls are found to be red. What is the probability the next ball chosen from the urn will be red? Show that when n is large, this probability is approximately $n/(n + 2)$. *Hint:* For the last step, approximate the sum by an integral.

1.26 The Monty Hall Game

You are a contestant in a game show. The host says, “Behind one of those three doors is a new automobile, which is your prize should you choose the right door. Nothing is behind the other two doors. You may choose any door.” You choose door A. The game host then opens door B and shows you that there is nothing behind it. He then asks, “Now would you like to change your guess to door C, at a cost of \$1?” Show that the answer is no if the game show host randomly opened one of the two other doors, but yes if he simply opened a door he knew did not have a car behind it. Generalize to the case where there are n doors with a prize behind one door.

1.27 The Logic of Murder and Abuse

For a given woman, let A be the event “was habitually beaten by her husband” (“abused” for short), let B be the event “was murdered,” and let C be the event “was murdered by her husband.” Suppose we know the following facts: (a) 5% of women are abused by their husbands; (b) 0.5% of women are murdered; (c) 0.025% of women are murdered by their husbands; (d) 90% of women who are murdered by their husbands had been abused by their husbands; (e) a woman who is murdered but not by her husband is neither more nor less likely to have been abused by her husband than a randomly selected woman.

Nicole is found murdered, and it is ascertained that she was abused by her husband. The defense attorneys for her husband show that the probability that a man who abuses his wife actually kills her is only 4.50%, so there is a strong presumption of innocence for him. The attorneys for the prosecution show that there is in fact a 94.74% chance the husband murdered his wife, independent from any evidence other than that he abused her. Please supply the arguments of the two teams of attorneys. You may assume that the jury was well versed in probability theory, so they had no problem understanding the reasoning.

1.28 The Principle of Insufficient Reason

The principle of insufficient reason says that if you are “completely ignorant” as to which among the states A_1, \dots, A_n will occur, then you should assign probability $1/n$ to each of the states. The argument in favor of the principle is strong (see Savage 1954 and Sinn 1980 for discussions), but there are some interesting arguments against it. For instance, suppose A_1 itself consists of m mutually exclusive events A_{11}, \dots, A_{1m} . If you are “completely ignorant” concerning which of these occurs, then if $P[A_1] = 1/n$, we should set $P[A_{1i}] = 1/mn$. But are we not “completely ignorant” concerning which of $A_{11}, \dots, A_{1m}, A_2, \dots, A_n$ occurs? If so, we should set each of these probabilities to $1/(n + m - 1)$. If not, in what sense were we “completely ignorant” concerning the original states A_1, \dots, A_n ?

1.29 The Greens and the Blacks

The game of bridge is played with a normal 52-card deck, each of four players being dealt 13 cards at the start of the game. The Greens and the Blacks are playing bridge. After a deal, Mr. Brown, an onlooker, asks Mrs. Black: “Do you have an ace in your hand?” She nods yes. After the next deal, he asks her: “Do you have the ace of spades?” She nods yes again. In which of the two situations is Mrs. Black more likely to have at least one other ace in her hand? Calculate the exact probabilities in the two cases.

1.30 The Brain and Kidney Problem

A mad scientist is showing you around his foul-smelling laboratory. He motions to an opaque, formalin-filled jar. “This jar contains either a brain

or a kidney, each with probability $1/2$,” he exclaims. Searching around his workbench, he finds a brain and adds it to the jar. He then picks one blob randomly from the jar, and it is a brain. What is the probability the remaining blob is a brain?

1.31 The Value of Eyewitness Testimony

A town has 100 taxis, 85 green taxis owned by the Green Cab Company and 15 blue taxis owned by the Blue Cab Company. On March 1, 1990, Alice was struck by a speeding cab, and the only witness testified that the cab was blue rather than green. Alice sued the Blue Cab Company. The judge instructed the jury and the lawyers at the start of the case that the reliability of a witness must be assumed to be 80% in a case of this sort, and that liability requires that the “preponderance of the evidence,” meaning at least a 50% probability, be on the side of the plaintiff.

The lawyer for Alice argued that the Blue Cab Company should pay, because the witness’s testimonial gives a probability of 80% that she was struck by a blue taxi. The lawyer for the Blue Cab Company argued as follows. A witness who was shown all the cabs in town would incorrectly identify 20% of the 85 green taxis (that is, 17 of them) as blue, and correctly identify 80% of the 15 blue taxis (that is, 12 of them) as blue. Thus, of the 29 identifications of a taxi as blue, only twelve would be correct and seventeen would be incorrect. Thus, the preponderance of the evidence is in favor of the defendant. Most likely, Alice was hit by a green taxi.

Formulate the second lawyer’s argument rigorously in terms of Bayes’ rule. Which argument do you think is correct, and if neither is correct, what is a good argument in this case?

1.32 When Weakness Is Strength

Many people have criticized the Darwinian notion of “survival of the fittest” by declaring that the whole thing is a simple tautology: whatever survives is “fit” by definition! Defenders of the notion reply by noting that we can measure fitness (e.g., speed, strength, resistance to disease, aerodynamic stability) independent of survivability, so it becomes an empirical proposition that the fit survive. Indeed, under some conditions it may be simply false, as game theorist Martin Shubik (1954) showed in the following ingenious problem.

Alice, Bob, and Carole are having a shootout. On each round, until only one player remains standing, the current shooter can choose one of the other players as target and is allowed one shot. At the start of the game, they draw straws to see who goes first, second, and third, and they take turns repeatedly in that order. A player who is hit is eliminated. Alice is a perfect shot, Bob has 80% accuracy, and Carole has 50% accuracy. We assume that players are not required to aim at an opponent and can simply shoot in the air on their turn, if they so desire.

We will show that Carole, the least accurate shooter, is the most likely to survive. As an exercise, you are asked to show that if the player who gets to shoot is picked randomly in each round, then the survivability of the players is perfectly inverse to their accuracy.

There are six possible orders for the three players, each occurring with probability $1/6$. We abbreviate Alice as a , Bob as b , and Carole as c , and we write the order of play as xyz , where $x,y,z \in \{a,b,c\}$. We let $\pi_i(xyz)$ be the survival probability of player $i \in \{a,b,c\}$. For instance, $\pi_a(abc)$ is the probability Alice wins when the shooting order is abc . Similarly, if only two remain, let $\pi_i(xy)$ be the probability of survival for player $i = x,y$ when only x and y remain, and it is x 's turn to shoot.

If Alice goes first, it is clear that her best move is to shoot at Bob, whom she eliminates with probability 1. Then, Carole's best move is to shoot at Alice, whom she eliminates with probability $1/2$. If she misses Alice, Alice eliminates Carole. Therefore, we have $\pi_a(abc) = 1/2$, $\pi_b(abc) = 0$, $\pi_c(abc) = 1/2$, $\pi_a(acb) = 1/2$, $\pi_b(acb) = 0$, and $\pi_c(acb) = 1/2$.

Suppose Bob goes first, and the order is bac . If Bob shoots in the air, Alice will then eliminate Bob. If Bob shoots at Carole and eliminates her, Alice will again eliminate Bob. If Bob shoots at Alice and misses, then the order is effectively acb , and we know Alice will eliminate Bob. However, if Bob shoots at Alice and eliminates her, then the game is cb . We have

$$p_c(cb) = \frac{1}{2} + \frac{1}{2} \times \frac{1}{5} p_c(cb).$$

The first term on the right is the probability Carole hits Bob and wins straight off, and the second term is the probability that she misses Bob ($1/2$) times the probability Bob misses her ($1/5$) times the probability that she eventually wins if it is her turn to shoot. We can solve this equation, getting $p_c(cb) = 5/9$, so $p_b(cb) = 4/9$. It follows that Bob will indeed

shoot at Alice, so

$$p_{\mathbf{b}}(\text{bac}) = \frac{4}{5} \times \frac{4}{9} = \frac{16}{45}.$$

Similarly, we have $p_{\mathbf{b}}(\text{bca}) = 16/45$. Also,

$$p_{\mathbf{a}}(\text{bac}) = \frac{1}{5} p_{\mathbf{a}}(\text{ca}) = \frac{1}{5} \times \frac{1}{2} = \frac{1}{10},$$

because we clearly have $p_{\mathbf{a}}(\text{ca}) = 1/2$. Similarly, $p_{\mathbf{a}}(\text{bca}) = 1/10$. Finally,

$$p_{\mathbf{c}}(\text{bac}) = \frac{1}{5} p_{\mathbf{c}}(\text{ca}) + \frac{4}{5} \times p_{\mathbf{c}}(\text{cb}) = \frac{1}{5} \times \frac{1}{2} + \frac{4}{5} \times \frac{5}{9} = \frac{49}{90},$$

because $p_{\mathbf{c}}(\text{ca}) = 1/2$. Similarly, $p_{\mathbf{c}}(\text{bca}) = 49/90$. As a check on our work, note that $p_{\mathbf{a}}(\text{bac}) + p_{\mathbf{b}}(\text{bac}) + p_{\mathbf{c}}(\text{bac}) = 1$.

Suppose Carole gets to shoot first. If Carole shoots in the air, her payoff from cab is $p_{\mathbf{c}}(\text{abc}) = 1/2$, and from cba is $p_{\mathbf{c}}(\text{bac}) = 49/90$. These are also her payoffs if she misses her target. However, if she shoots Alice, her payoff is $p_{\mathbf{c}}(\text{bc})$, and if she shoots Bob, her payoff is $p_{\mathbf{c}}(\text{ac}) = 0$. We calculate $p_{\mathbf{c}}(\text{bc})$ as follows.

$$p_{\mathbf{b}}(\text{bc}) = \frac{4}{5} + \frac{1}{5} \times \frac{1}{2} p_{\mathbf{b}}(\text{bc}),$$

where the first term is the probability he shoots Carole (4/5) plus the probability he misses Carole (1/5) times the probability he gets to shoot again (1/2, because Carole misses) times $p_{\mathbf{b}}(\text{bc})$. We solve, getting $p_{\mathbf{b}}(\text{bc}) = 8/9$. Thus, $p_{\mathbf{c}}(\text{bc}) = 1/9$. Clearly, Carole's best payoff is to shoot in the air. Then $p_{\mathbf{c}}(\text{cab}) = 1/2$, $p_{\mathbf{b}}(\text{cab}) = p_{\mathbf{b}}(\text{abc}) = 0$, and $p_{\mathbf{a}}(\text{cab}) = p_{\mathbf{a}}(\text{abc}) = 1/2$. Also, $p_{\mathbf{c}}(\text{cba}) = 49/90$, $p_{\mathbf{b}}(\text{cba}) = p_{\mathbf{b}}(\text{bac}) = 16/45$, and $p_{\mathbf{a}}(\text{cba}) = p_{\mathbf{a}}(\text{bac}) = 1/10$.

The probability that Alice survives is given by

$$\begin{aligned} p_{\mathbf{a}} &= \frac{1}{6}(p_{\mathbf{a}}(\text{abc}) + p_{\mathbf{a}}(\text{acb}) + p_{\mathbf{a}}(\text{bac}) + p_{\mathbf{a}}(\text{bca}) + p_{\mathbf{a}}(\text{cab}) + p_{\mathbf{a}}(\text{cba})) \\ &= \frac{1}{6} \left(\frac{1}{2} + \frac{1}{2} + \frac{1}{10} + \frac{1}{10} + \frac{1}{2} + \frac{1}{10} \right) \left(= \frac{3}{10}. \right. \end{aligned}$$

The probability that Bob survives is given by

$$\begin{aligned} p_{\mathbf{b}} &= \frac{1}{6}(p_{\mathbf{b}}(\text{abc}) + p_{\mathbf{b}}(\text{acb}) + p_{\mathbf{b}}(\text{bac}) + p_{\mathbf{b}}(\text{bca}) + p_{\mathbf{b}}(\text{cab}) + p_{\mathbf{b}}(\text{cba})) \\ &= \frac{1}{6} \left(0 + 0 + \frac{16}{45} + \frac{16}{45} + 0 + \frac{16}{45} \right) \left(= \frac{8}{45}. \right. \end{aligned}$$

The probability that Carole survives is given by

$$\begin{aligned} p_c &= \frac{1}{6}(p_c(abc) + p_c(acb) + p_c(bac) + p_c(bca) + p_c(cab) + p_c(cba)) \\ &= \frac{1}{6} \left(\frac{1}{2} + \frac{1}{2} + \frac{49}{90} + \frac{49}{90} + \frac{1}{2} + \frac{49}{90} \right) = \frac{47}{90}. \end{aligned}$$

You can check that these three probabilities add up to unity, as they should. Note that Carole has a 52.2% chance of surviving, whereas Alice has only a 30% chance, and Bob has a 17.8% chance.

1.33 The Uniform Distribution

The *uniform distribution* on $[0, 1]$ is a random variable that is uniformly distributed over the unit interval. Therefore if \tilde{x} is uniformly distributed over $[0, 1]$ then

$$P[\tilde{x} < x] = \begin{cases} 0 & x \leq 0 \\ x & 0 \leq x \leq 1 \\ 1 & 1 \leq x \end{cases}.$$

If \tilde{x} is uniformly distributed on the interval $[a, b]$, then $(\tilde{x} - a)/(b - a)$ is uniformly distributed on $[0, 1]$, and a little algebra shows that

$$P[\tilde{x} < x] = \begin{cases} 0 & x \leq a \\ \frac{x-a}{b-a} & a \leq x \leq b \\ 1 & b \leq x \end{cases}.$$

Figure 1.1 depicts this problem.

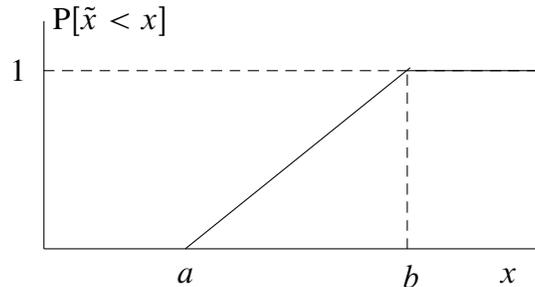


Figure 1.1. Uniform distribution

Suppose \tilde{x} is uniformly distributed on $[a, b]$ and we learn that in fact $\tilde{x} \leq c$, where $a < c < b$. Then \tilde{x} is in fact uniformly distributed on $[a, c]$.

To see this, we write

$$\begin{aligned} P[\tilde{x} < x | \tilde{x} \leq c] &= \frac{P[\tilde{x} < x \text{ and } \tilde{x} \leq c]}{P[\tilde{x} \leq c]} \\ &= \frac{P[\tilde{x} < x \text{ and } \tilde{x} \leq c]}{(c - a)/(b - a)}. \end{aligned}$$

We evaluate the numerator as follows:

$$\begin{aligned} P[\tilde{x} < x \text{ and } \tilde{x} \leq c] &= \begin{cases} \emptyset & x \leq a \\ P[\tilde{x} < x] & a \leq x \leq c \\ P[\tilde{x} \leq c] & c \leq x \end{cases} \\ &= \begin{cases} \emptyset & x \leq a \\ \frac{x-a}{b-a} & a \leq x \leq c \\ \frac{c-a}{b-a} & c \leq x \end{cases}. \end{aligned}$$

Therefore,

$$P[\tilde{x} < x | \tilde{x} \leq c] = \begin{cases} \emptyset & x \leq a \\ \frac{x-a}{c-a} & a \leq x \leq c \\ 1 & c \leq x \end{cases}.$$

This is just the uniform distribution on $[a, c]$.

1.34 Laplace's Law of Succession

An urn contains a large number n of white and black balls, where the number of white balls is uniformly distributed between 0 and n . Suppose you pick out m balls with replacement, and r are white. Show that the probability of picking a white ball on the next draw is approximately $(r + 1)/(m + 2)$.

1.35 From Uniform to Exponential

Bob tells Alice to draw repeatedly from the uniform distribution on $[0, 1]$ until her current draw is less than some previous draw, and he will pay her $\$n$, where n is the number of draws. What is the average value of this game for Alice?