Probability and Randomness

Jay Daigle

jaydaigle@gwu.edu

https://jaydaigle.net/politics

The George Washington University

November 10, 2025

The Dilemma

- If there's a saddle point, we should pick that
- If there's no saddle point, there's no clear best choice
- But if we settle on one strategy, we can be exploited.

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

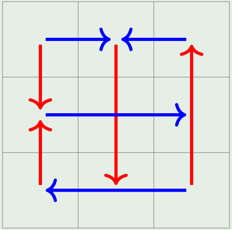

Discussion Question

How do we pick a strategy?

Probability

- Want to talk about randomness.

Definition

- The possible results of a random process are outcomes
- The set of all possible outcomes is the sample space.
- Term comes from statistics—taking a random sample.

Example

- Tossing a coin
  - Possible outcomes: "heads" and "tails".
  - Sample space: $\{h, t\}$.
- Rolling a six-sided die
  - Six possible outcomes.
  - Sample space: $\{1, 2, 3, 4, 5, 6\}$.

Probability

Definition

- Assign each outcome a probability between 0 and 1
- For k outcomes, write $p_1, p_2, \ldots, p_n$
- $0 \leq p_k \leq 1$
  - $p_k = 0$ means outcome k can't happen
  - $p_k = 1$ means outcome k must happen
- $p_1 + p_2 + \cdots + p_n = 1$
  - I.e. exactly one thing will happen
- The list of numbers $P = (p_1, \ldots, p_n)$ is a probability distribution

Probability and Chance

Remark

- In English, often refer to probabilities as chances
- Usually given as a percentage
- A 50% chance is the same as $p = 1/2$.

Example (Coin toss)

- Chance of heads: 50%
- Chance of tails: 50%
- $P = (1/2, 1/2)$.

Example (Die roll)

- Six faces, each is equally likely
- $P = (1/6, 1/6, 1/6, 1/6, 1/6, 1/6)$.

Probability

Remark

- Our definition assumes we have listed every possible outcome.
- In real life that's never quite true.
  - What if a coin lands on edge?
- Useful assumption
- Matches assumption that our matrix contains every possible strategy.

Probabilities are not always equal

- Sometimes one outcome is more likely than another
- Sometimes obvious, sometimes subtle
- In English, "choose at random" usually means every option is equally likely
- In this class, it usually does not.

Definition

- The uniform distribution on n outcomes gives every outcome the same probability
- Get the probability distribution $P = (1/n, 1/n, \dots, 1/n)$.

Non-Uniform Probability Distributions

- Can just give a non-uniform distribution
- $P = (1/3, 2/3)$ means first option is half as likely as the second one.

Example (Weather forecasting)

- Might say 50% chance of sun, 30% chance of clouds, 20% chance of rain.
- Interpret this as a probability distribution: $P = (0.5, 0.3, 0.2)$.

Non-Uniform Probability Distributions

- Many non-uniform distributions come from uniform distributions

Example (Deck of cards)

- 12 face cards
- 40 non-face cards
- Each *card* equally likely
- $P = (12/52, 40/52) = (3/13, 10/13) \approx (0.23, 0.77)$

Non-Uniform Probability Distributions

- We can usually build a non-uniform distribution out of a uniform one.

Example

- Say we want the distribution $P = (1/6, 1/4, 1/3, 1/4)$.
- Check they sum to 1?
  - $\frac{1}{6} + \frac{1}{4} + \frac{1}{3} + \frac{1}{4} = \frac{2}{12} + \frac{3}{12} + \frac{4}{12} + \frac{3}{12} = \frac{2+3+4+3}{12} = 1.$
- Can generate with a uniform distribution over 12 possibilities.
- E.g. take 12 cards, and write A on 2, B on 3, C on 4, D on 3
  - Draw a card uniformly at random: get distribution P
- Or divide a spinner into 12 equal sections
- Or look at seconds on your watch. First 10 seconds gives A, next 15 gives B, next 20 gives C, last 15 gives D.

Random Variables

- Want to talk about probability in context of strategic choices
- Need to think about how *valuable* outcomes are
- The standard terminology is really bad.

Definition

- Suppose we have a sample space with n outcomes
- and a probability distribution $P = (p_1, p_2, \dots, p_n)$
- A random variable X on this sample space is a function that assigns a real number to each of the n possible outcomes.
- Can write as $X = (x_1, x_2, \dots, x_n)$.

Random Variables

Remark

- We think of the random variable as the value of an outcome—how much we like it.
- Similar to difference between "outcome" and "payoff".

Example

- Rolling a six-sided die:
  - $x_1 = 1, x_2 = 2, x_3 = 3, x_4 = 4, x_5 = 5, x_6 = 6$
  - $X = (1, 2, 3, 4, 5, 6)$

Example

- Chance of getting a penny, dime, or quarter
  - $x_1 = 1, x_2 = 10, x_3 = 25$
  - $X = (1, 10, 25)$.

Expected Value

Definition

- Given a sample space with n outcomes
- With probability distribution $P = (p_1, p_2, \ldots, p_n)$.
- Let X be a random variable that assigns the payoff x_k to outcome k.
- We define the expected value of X to be

$$E = E(X) = p_1x_1 + p_2x_2 + \cdots + p_nx_n.$$

- Idea: average payoff per play if you play a bunch of times.
- If you run your process a hundred times, you will get about $100 \cdot E(X)$.

Expected Value

Example

- What is the expected value of rolling a six-sided die?

$$E = \frac{1}{6} \cdot 1 + \frac{1}{6} \cdot 2 + \frac{1}{6} \cdot 3 + \frac{1}{6} \cdot 4 + \frac{1}{6} \cdot 5 + \frac{1}{6} \cdot 6$$
$$= \frac{1}{6}\left(1 + 2 + 3 + 4 + 5 + 6\right) = \frac{1}{6} \cdot 21 = 3.5.$$

Expected Value

Example (Lottery)

- Grand prize: \$100,000,000 with probability 1 in 150,000,000
- Second prize: \$200,000 with probability 1 in 3,000,000
- Third prize: \$10,000 with probability 1 in 150,000
- Fourth prize: \$10 with probability 1 in 300
- What is the expected value of playing?

- How many possible outcomes? Five
- How to compute?

Expected Value

Prize:	Grand	Second	Third	Fourth	Lose
Payoff:	\$100,000,000	\$200,000	\$10,000	\$10	\$0
Probability:	1/150,000,000	1/3,000,000	1/150,000	1/300	0.997

$$E = \frac{100,000,000}{150,000,000} + \frac{200,000}{3,000,000} + \frac{10,000}{150,000} + \frac{10}{300} + 0 \cdot 0.997$$

$$\approx 0.667 + 0.067 + 0.067 + 0.033 + 0$$

$$\approx 0.833.$$

- Average payoff: about 83 cents.

Expected Value

Prize:	Grand	Second	Third	Fourth	Lose
Payoff:	\$100,000,000	\$200,000	\$10,000	\$10	\$0
Probability:	1/150,000,000	1/3,000,000	1/150,000	1/300	0.997

- Average payoff: about 83 cents.

Discussion Question

- Would you play if you got a ticket for free?
- Would you pay a dollar for a ticket?
- What's the most you'd be willing to pay for a ticket?

Expected Value

- Suppose you pay \$1 for a ticket
- New payoffs for each outcome:

Prize:	Grand	Second	Third	Fourth	Lose
Payoff:	\$99,999,999	\$199,999	\$9,999	\$9	−\$1
Probability:	$1/150,000,000$	$1/3,000,000$	$1/150,000$	$1/300$	0.997

$$E = \frac{99,999,999}{150,000,000} + \frac{199,999}{3,000,000} + \frac{9,999}{150,000} + \frac{9}{300} + (-1) \cdot 0.997$$
$$\approx 0.667 + 0.067 + 0.067 + 0.033 + -1 \approx -0.167.$$

- Lose about 17 cents per ticket.

Expected Value

Discussion Question

- When rolling a six-sided die, do you expect to get a 3.5?
- When entering that lottery, do you expect to lose seventeen cents?

Remark

- "Expected value" is a dangerous phrasing
- You do not (usually) expect to get that exact value
- Computes the *average* result over many attempts
- Useful for analyzing repeated decisions
- Sort of useful for analyzing one-off decisions

Expected Value

Definition

The expected value principle says a rational player will choose the option with the largest expected value.

- Can be seen as a definition of "rational"
- Often taken as prescriptive: a player *should* choose . . .
- Has many caveats! Just one model of decision-making.

Limitations to Expected Value Principle

- Useful framework for analyzing decisions
- Most useful when making repeated decisions
- Most useful when all payoffs are on the same scale

Discussion Question

Should you pay $1 for a 1 in 100,000 chance of getting $100,001?

- Expected value: $\frac{100{,}001}{100{,}000} = 1.00001$
- More than a dollar, so expected value principle says yes
- But should you?

Limitations to Expected Value Principle

Discussion Question

Should you pay \$1,000,000 for a 1 in 1,000 chance of getting \$2,000,000,000?

- Expected value: $\frac{2{,}000{,}000{,}000}{1{,}000} = 2,000,000$
- Expected value principle says yes
- Is this a good idea?
- Do you have a million dollars?
- What would you do with 2 billion dollars?
- What will you do if you lose?

Limitations to Expected Value Principle

Diminishing marginal utility of money

- First dollar is more valuable than the millionth dollar
- People worry about "risk of ruin"—chance of going broke
- This is the point of insurance—negative expected value, but reduces risk of ruin.

- Can avoid some problems by giving "utility" instead of dollars
  - Kind of fake, not clear what we're measuring
  - But useful mathematical abstraction
  - (In the real-life applications we care about, the numbers are kind of made up anyway.)
- This is just one way to analyze decisions—but it's often useful and it's the framework we'll use.

Mixed Strategies in Games

- Now have the tools to analyze games like Rock Paper Scissors.

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- No saddle points
- Every strategy has a counter-strategy
- Any predictable strategy will lose.
- Need unpredictability—which means randomnness.

Mixed Strategies

Definition

- Consider an $m \times n$ zero-sum two-player game
- The strategies corresponding to rows are columns are now called pure strategies
- A mixed strategy for Row is a probability distribution $P = (p_1, \ldots, p_m)$ on their set of pure strategies
  - Row chooses a pure strategy at random, choosing row k with probability p_k
- A mixed strategy for Column is a probability distribution $Q = (q_1, \ldots, q_n)$ on their set of pure strategies.

Mixed Strategies

Example

- Column plays Rock Paper Scissors with $Q = (1/4, 1/2, 1/4)$
  - Rock 1/4 the time, Paper 1/2 the time, Scissors 1/4 the time
- Now Row can view each strategy as a lottery

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0
	1/4	1/2	1/4

- If Row plays Rock: $E = \frac{1}{4} \cdot (0) + \frac{1}{2} \cdot (-1) + \frac{1}{4} \cdot (1) = -\frac{1}{4}$
  - On average, Row will lose one point per four games.

Mixed Strategies

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0
	1/4	1/2	1/4

- If Row plays Rock: $E = \frac{1}{4} \cdot (0) + \frac{1}{2} \cdot (-1) + \frac{1}{4} \cdot (1) = -\frac{1}{4}$
  - On average, Row will lose one point per four games.
- If Row plays Paper: $E = \frac{1}{4} \cdot (1) + \frac{1}{2} \cdot (0) + \frac{1}{4} \cdot (-1) = 0$
- If Row plays Scissors: $E = \frac{1}{4} \cdot (-1) + \frac{1}{2} \cdot (1) + \frac{1}{4} \cdot (0) = \frac{1}{4}$
- What should Row do? Play Scissors.