

Balls and Urns

“Balls and urns” problems are paradigmatic. Many problems can be recast as balls and urns problems, once we figure out which are the balls and which are the urns.

How many ways are there of putting b balls into u urns?

- That depends whether the balls are distinguishable and whether the urns are distinguishable

How many ways are there of putting 5 balls into 2 urns?

- If both balls and urns are distinguishable: $2^5 = 32$
 - Choose the subset of balls that goes into the first urn
 - Alternatively, for each ball, decide which urn it goes in
 - This assumes that it's OK to have 0 balls in an urn.
- If urns are distinguishable but balls aren't: 6
 - Decide how many balls go into the first urn: 0, 1, ..., 5
- If balls are distinguishable but urns aren't: $2^5/2 = 16$
- If balls and urns are indistinguishable: $6/2 = 3$

What if we had 6 balls and 2 urns?

- If balls and urns are distinguishable: 2^6
- If urns are distinguishable and balls aren't: 7
- If balls are distinguishable but urns aren't:

$$2^6/2 = 2^5$$

- If balls and urns are indistinguishable: 4
 - It can't be $7/2$, since that's not an integer
 - The problem is that if there are 3 balls in each urn, and you switch urns, then you get the same solution

Reducing Problems to Balls and Urns

Q1: How many different configurations are there in Towers of Hanoi with n rings?

A: The urns are the poles, the balls are the rings. Both are distinguishable.

Q2: How many solutions are there to the equation $x + y + z = 65$, if x, y, z are nonnegative integers?

A: You have 65 indistinguishable balls, and want to put them into 3 distinguishable urns (x, y, z) . Each way of doing so corresponds to one solution.

Q3: How many ways can 8 electrons be assigned to 4 energy states?

A: The electrons are the balls; they're indistinguishable. The energy states are the urns; they're distinguishable.

Distinguishable Urns

How many ways can b distinguishable balls be put into u distinguishable urns?

- By the product rule, this is u^b

How many ways can b indistinguishable balls be put into u distinguishable urns?

$$C(u + b - 1, b)$$

Reducing Problems to Balls and Urns

Q1: How many different configurations are there in Towers of Hanoi with n rings?

A: The urns are the poles, the balls are the rings. Both are distinguishable.

- 3^n

Q2: How many solutions are there to the equation $x + y + z = 65$, if x, y, z are nonnegative integers?

A: You have 65 indistinguishable balls, and want to put them into 3 distinguishable urns (x, y, z) . Each way of doing so corresponds to one solution.

- $C(67, 65) = 67 \times 33 = 2211$

Q3: How many ways can 8 electrons be assigned to 4 energy states?

A: The electrons are the balls; they're indistinguishable. The energy states are the urns; they're distinguishable.

- $C(11, 8) = (11 \times 10 \times 9)/6 = 165$

Indistinguishable Urns

How many ways can b distinguishable balls be put into u indistinguishable urns?

First view the urns as distinguishable: u^b

For every solution, look at all $u!$ permutations of the urns. That should count as one solution.

- By the Division Rule, we get: $u^b/u!$?

This can't be right! It's not an integer (e.g. $7^3/7!$).

What's wrong?

The situation is even worse when we have indistinguishable balls in indistinguishable urns. (See the book.)

Inclusion-Exclusion Rule

Remember the Sum Rule:

The Sum Rule: If there are $n(A)$ ways to do A and, distinct from them, $n(B)$ ways to do B , then the number of ways to do A or B is $n(A) + n(B)$.

What if the ways of doing A and B aren't distinct?

Example: If 112 students take CS280, 85 students take CS220, and 45 students take both, how many take either CS280 or CS220.

A = students taking CS280

B = students taking CS220

$$|A \cup B| = |A| + |B| - |A \cap B| = 112 + 85 - 45 = 152$$

This is best seen using a Venn diagram:

How many numbers ≤ 100 are multiples of either 2 or 5?

Let $A =$ multiples of 2 ≤ 100

Let $B =$ multiples of 5 ≤ 100

Then $A \cap B =$ multiples of 10 ≤ 100

$$|A \cup B| = |A| + |B| - |A \cap B| = 50 + 20 - 10 = 60.$$

What happens with three sets?

$$\begin{aligned} |A \cup B \cup C| = \\ |A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| \\ + |A \cap B \cap C| \end{aligned}$$

Example: If there are 300 engineering majors, 112 take CS280, 85 takes CS 220, 95 take AEP 356, 45 take both CS280 and CS 220, 30 take both CS 280 and AEP 356, 25 take both CS 220 and AEP 356, and 5 take all 3, how many don't take any of these 3 courses?

A = students taking CS 280

B = students taking CS 220

C = students taking AEP 356

$$\begin{aligned} & |A \cup B \cup C| \\ = & |A| + |B| + |C| - |A \cap B| - |B \cap C| - |A \cap C| \\ & + |A \cap B \cap C| \\ = & 112 + 85 + 95 - 45 - 30 - 25 + 5 \\ = & 197 \end{aligned}$$

We are interested in $\overline{A \cup B \cup C} = 300 - 197 = 103$.

The General Rule

More generally,

$$|\cup_{k=1}^n A_k| = \sum_{k=1}^n \sum_{\{I|I\subset\{1,\dots,n\},|I|=k\}} (-1)^{k-1} |\cap_{i\in I} A_i|$$

Why is this true? Suppose $a \in \cup_{k=1}^n A_k$, and is in exactly m sets. a gets counted once on the LHS. How many times does it get counted on the RHS?

- a appears in m sets (1-way intersection)
- a appears in $C(m, 2)$ 2-way intersections
- a appears in $C(m, 3)$ 3-way intersections
- ...

Thus, on the RHS, a gets counted

$$\sum_{k=1}^m (-1)^{k-1} C(m, k) \text{ times.}$$

By the binomial theorem:

$$\begin{aligned} 0 &= (-1 + 1)^m = \sum_{k=0}^m (-1)^k 1^{m-k} C(m, k) \\ &= 1 + \sum_{k=1}^m (-1)^k C(m, k) \end{aligned}$$

Thus,

$$\sum_{k=1}^m (-1)^{k-1} C(m, k) = 1.$$

Each element in $\cup_{i=1}^k A_i$ gets counted once on both sides.

A Hard Example

Suppose $m \geq 10$. How many m -digit numbers have each of the digits 0–9 at least once? (View 00305 as a 5-digit number.)

We need a systematic way of tackling this.

Let A_j be the set of m -digit numbers that have at least one occurrence of j , for $j = 0, \dots, 9$.

We are interested in $|A_0 \cap \dots \cap A_9|$.

The inclusion-exclusion rule applies to unions. Can we use it?

$$\overline{A_0 \cap \dots \cap A_9} = \overline{A_0} \cup \dots \cup \overline{A_9}$$

$$|\overline{A_i}| = 9^m$$

$$|\overline{A_i \cap A_j}| = 8^m$$

...

$$\begin{aligned} |\cup_{i=0}^9 \overline{A_i}| &= 10 \times 9^m - C(10, 2) \times 8^m + \dots \\ &= \sum_{k=1}^9 (-1)^{k-1} C(10, k) \times (10 - k)^m \end{aligned}$$

Thus,

$$\begin{aligned} |\cap_{i=0}^9 A_i| &= 10^m - \sum_{k=1}^9 (-1)^{k-1} C(10, k) \times (10 - k)^m \\ &= \sum_{k=0}^9 (-1)^k C(10, k) (10 - k)^m \end{aligned}$$

The Pigeonhole Principle

The Pigeonhole Principle: If $n + 1$ pigeons are put into n holes, at least two pigeons must be in the same hole.

This seems obvious. How can it be used in combinatorial analysis?

Q1: If you have only blue socks and brown socks in your drawer, how many do you have to pull out before you're sure to have a matching pair.

A: The socks are the pigeons and the holes are the colors. There are two holes. With three pigeons, there have to be at least two in one hole.

- What happens if we also have black socks?

Q2: Bob picks 10 numbers between 1 and 40. Alice wins if she can find two different sets of three of these numbers that have the same sum. Who wins?

A: The holes are the possible sums. The smallest sum is 6 ($1+2+3$), the largest is 117 ($38+39+40$). The pigeons are the possible ways for Alice to choose 3 numbers out of the 10 chosen by Bob.

$$\binom{10}{3} = \frac{10 \times 9 \times 8}{3 \times 2 \times 1} = 120.$$

There's always a way for Alice to win!

Probability

Life is full of uncertainty.

Probability is the best way we currently have to quantify it.

Applications of probability arise everywhere:

- Should you guess in a multiple-choice test with five choices?
 - What if you're not penalized for guessing?
 - What if you're penalized $1/4$ for every wrong answer?
 - What if you can eliminate two of the five possibilities?

- Suppose that an AIDS test guarantees 99% accuracy:
 - of every 100 people who have AIDS, the test returns positive 99 times (very few *false negative*);
 - of every 100 people who don't have AIDS, the test returns negative 99 times (very few *false positives*)

Suppose you test positive. How likely are you to have AIDS?

- Hint: the probability is *not* .99
- How do you compute the average-case running time of an algorithm?
- Is it worth buying a \$1 lottery ticket?
 - Probability isn't enough to answer this question

(I think) everybody ought to know something about probability.

Interpreting Probability

Probability can be a subtle.

The first (philosophical) question is “What does probability mean?”

- What does it mean to say that “The probability that the coin landed (will land) heads is $1/2$ ”?

Two standard interpretations:

- Probability is *subjective*: This is a subjective statement describing an individual’s feeling about the coin landing heads
 - This feeling can be quantified in terms of betting behavior
- Probability is an *objective* statement about frequency

Both interpretations lead to the same mathematical notion.

Formalizing Probability

What do we assign probability to?

Intuitively, we assign them to possible *events* (things that might happen, *outcomes* of an experiment)

Formally, we take a *sample space* to be a *set*.

- Intuitively, the sample space is the set of possible outcomes, or possible ways the world could be.

An *event* is a subset of a sample space.

We assign probability to events: that is, to subsets of a sample space.

Sometimes the hardest thing to do in a problem is to decide what the sample space should be.

- There's often more than one choice
- A good thing to do is to try to choose the sample space so that all outcomes (i.e., elements) are equally likely
 - This is not always possible or reasonable

Choosing the Sample Space

Example 1: We toss a coin. What's the sample space?

- Most obvious choice: {heads, tails}
- Should we bother to model the possibility that the coin lands on edge?
- What about the possibility that somebody snatches the coin before it lands?
- What if the coin is biased?

Example 2: We toss a die. What's the sample space?

Example 3: Two distinguishable dice are tossed together. What's the sample space?

- $(1,1), (1,2), (1,3), \dots, (6,1), (6,2), \dots, (6,6)$

What if the dice are indistinguishable?

Example 4: You're a doctor examining a seriously ill patient, trying to determine the probability that he has cancer. What's the sample space?

Example 5: You're an insurance company trying to insure a nuclear power plant. What's the sample space?

Probability Measures

A *probability measure* assigns a real number between 0 and 1 to every subset of (event in) a sample space.

- Intuitively, the number measures how likely that event is.
- Probability 1 says it's certain to happen; probability 0 says it's certain not to happen
- Probability acts like a *weight* or *measure*. The probability of separate things (i.e., disjoint sets) adds up.

Formally, a probability measure \Pr on S is a function mapping subsets of S to real numbers such that:

1. For all $A \subseteq S$, we have $0 \leq \Pr(A) \leq 1$
2. $\Pr(\emptyset) = 0$; $\Pr(S) = 1$
3. If A and B are disjoint subsets of S (i.e., $A \cap B = \emptyset$), then $\Pr(A \cup B) = \Pr(A) + \Pr(B)$.

It follows by induction that if A_1, \dots, A_k are pairwise disjoint, then

$$\Pr(\cup_i^k A_i) = \sum_i^k \Pr(A_i).$$

- This is called *finite additivity*; it's actually more standard to assume a countable version of this, called *countable additivity*

In particular, this means that if $A = \{e_1, \dots, e_k\}$, then

$$\Pr(A) = \sum_{i=1}^k \Pr(e_i).$$

In finite spaces, the probability of a set is determined by the probability of its elements.

Equiprobable Measures

Suppose S has n elements, and we want \Pr to make each element equally likely.

- Then each element gets probability $1/n$
- $\Pr(A) = |A|/n$

In this case, \Pr is called an *equiprobable measure*.

Examples

Example 1: In the coin example, if you think the coin is fair, and the only outcomes are heads and tails, then we can take $S = \{\text{heads}, \text{tails}\}$, and $\Pr(\text{heads}) = \Pr(\text{tails}) = 1/2$.

Example 2: In the two-dice example where the dice are distinguishable, if you think both dice are fair, then we can take $\Pr((i, j)) = 1/36$.

- Should it make a difference if the dice are indistinguishable?

Equiprobable measures on infinite sets

Defining an equiprobable measure on an infinite set can be tricky.

Theorem: There is no equiprobable measure on the positive integers.

Proof: By contradiction. Suppose \Pr is an equiprobable measure on the positive integers, and $\Pr(1) = \epsilon > 0$.

There must be some N such that $\epsilon > 1/N$.

Since $\Pr(1) = \dots = \Pr(N) = \epsilon$, we have

$$\Pr(\{1, \dots, N\}) = N\epsilon > 1 \text{ — a contradiction}$$

But if $\Pr(1) = 0$, then $\Pr(S) = \Pr(1) + \Pr(2) + \dots = 0$.

Some basic results

How are the probability of E and \overline{E} related?

- How does the probability that the dice lands either 2 or 4 (i.e., $E = \{2, 4\}$) compare to the probability that the dice lands 1, 3, 5, or 6 ($\overline{E} = \{1, 3, 5, 6\}$)

Theorem 1: $\Pr(\overline{E}) = 1 - \Pr(E)$.

Proof: E and \overline{E} are disjoint, so that

$$\Pr(E \cup \overline{E}) = \Pr(E) + \Pr(\overline{E}).$$

But $E \cup \overline{E} = S$, so $\Pr(E \cup \overline{E}) = 1$.

Thus $\Pr(E) + \Pr(\overline{E}) = 1$, so

$$\Pr(\overline{E}) = 1 - \Pr(E).$$

Theorem 2: $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$.

$$A = (A - B) \cup (A \cap B)$$

$$B = (B - A) \cup (A \cap B)$$

$$A \cup B = (A - B) \cup (B - A) \cup (A \cap B)$$

So

$$\Pr(A) = \Pr(A - B) + \Pr(A \cap B)$$

$$\Pr(B) = \Pr(B - A) + \Pr(A \cap B)$$

$$\Pr(A \cup B) = \Pr(A - B) + \Pr(B - A) + \Pr(A \cap B)$$

The result now follows.

Remember the Inclusion-Exclusion Rule?

$$|A \cup B| = |A| + |B| - |A \cap B|$$

This follows easily from Theorem 2, if we take \Pr to be an equiprobable measure. We can also generalize to arbitrary unions.

Disclaimer

- Probability is a well defined mathematical theory.
- Applications of probability theory to “real world” problems is not.
- Choosing the sample space, the events and the probability function requires a “leap of faith”.
- We cannot prove that we chose the right model but we can argue for that.
- Some examples are easy some are not:
 - Flipping a coin or rolling a die.
 - Playing a lottery game.
 - Guessing in a multiple choice test.
 - Determining whether or not the patient has AIDS based on a test.
 - Does the patient have cancer?

Conditional Probability

One of the most important features of probability is that there is a natural way to *update* it.

Example: Bob draws a card from a 52-card deck. Initially, Alice considers all cards equally likely, so her probability that the ace of spades was drawn is $1/52$. Her probability that the card drawn was a spade is $1/4$.

Then she sees that the card is black. What should her probability now be that

- the card is the ace of spades?
- the card is a spade?

A reasonable approach:

- Start with the original sample space
- Eliminate all outcomes (elements) that you now consider impossible, based on the observation (i.e., assign them probability 0).
- Keep the relative probability of everything else the same.
 - Renormalize to get the probabilities to sum to 1.

What should the probability of B be, given that you've observed A ? According to this recipe, it's

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

$$\Pr(A\spadesuit | \text{black}) = (1/52)/(1/2) = 1/26$$

$$\Pr(\text{spade} | \text{black}) = (1/4)/(1/2) = 1/2.$$

A subtlety:

- What if Alice doesn't completely trust Bob? How do you take this into account? Two approaches:

(1) Enlarge sample space to allow more observations.

(2) Jeffrey's rule:

$$\Pr(A\spadesuit | \text{black}) \cdot \Pr(\text{Bob telling the truth}) + \Pr(A\spadesuit | \text{red}) \cdot \Pr(\text{Bob lying}).$$

Example

Alice has two coins, a fair one f and a loaded one l .

- l 's probability of landing H is $p > 1/2$.

Alice picks f and flips it twice.

- What is the sample space?

$$\Omega = \{(H, H), (H, T), (T, H), (T, T)\}.$$

- What is \Pr ?
- By symmetry this should be an equiprobable space.

Let $H_1 = \{(H, H), (H, T)\}$ and let $H_2 = \{(H, H), (T, H)\}$.

H_1 and H_2 are independent:

- $H_1 = \{(H, H), (H, T)\} \Rightarrow \Pr(H_1) = 2/4 = 1/2$.
- Similarly, $\Pr(H_2) = 1/2$.
- $H_1 \cap H_2 = \{(H, H)\} \Rightarrow \Pr(H_1 \cap H_2) = 1/4$.
- So, $\Pr(H_1 \cap H_2) = \Pr(H_1) \cdot \Pr(H_2)$.

Alice next picks l and flips it twice.

- The sample space is the same as before:

$$\Omega = \{(H, H), (H, T), (T, H), (T, T)\}.$$

- We now define \Pr by *assuming* the flips are independent:

- $\Pr\{(H, H)\} = \Pr(H_1 \cap H_2) := p^2$
- $\Pr\{(H, T)\} = \Pr(H_1 \cap \bar{H}_2) := p(1 - p)$
- $\Pr\{(T, H)\} = \Pr(\bar{H}_1 \cap H_2) := (1 - p)p$
- $\Pr\{(T, T)\} = \Pr(\bar{H}_1 \cap \bar{H}_2) := (1 - p)^2$.

- H_1 and H_2 are now independent by construction:

$$\begin{aligned} \Pr(H_1) &= \Pr\{(H, H), (H, T)\} = \\ &= p^2 + p(1 - p) = p[p + (1 - p)] = p. \end{aligned}$$

Similarly, $\Pr(H_2) = \Pr\{(H, H), (T, H)\} = p$
 $\Pr(H_1 \cap H_2) = \Pr(H, H) = p^2$.

- For either coin, the two flips are independent.