Advanced Algorithmic Techniques (COMP523)

Randomised Algorithms

Recap and plan

- **Previous lectures:**
  - Approximation Algorithms.
- **Next lectures:**
  - Randomised Algorithms.
- **This lecture:**
  - Probabilities background.

The Poker slide

- Over the weekend, I was playing Texas Hold'em with some friends…
    - (absolutely true story).

https://www.888poker.com/poker/poker-odds-calculator

Heads or Tails

Heads or Tails

You flip a fair coin

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

You flip two fair coins

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

You flip two fair coins

What is the probability of “both heads”?

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

You flip two fair coins

What is the probability of “both heads”?

Possible outcomes: HH, HT, TH, TT

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

You flip two fair coins

What is the probability of “both heads”?

Possible outcomes: HH, HT, TH, TT

1/4

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

You flip two fair coins

What is the probability of “both heads”?

Possible outcomes: HH, HT, TH, TT

1/4

What is the probability of “both the same”?

Heads or Tails

You flip a fair coin

What is the probability of “heads”?

1/2

You flip two fair coins

What is the probability of “both heads”?

Possible outcomes: HH, HT, TH, TT

1/4

What is the probability of “both the same”?

Possible outcomes: HH, HT, TH, TT

Heads or Tails

You flip a fair coin

What is the probability of "heads"?

1/2

You flip two fair coins

What is the probability of "both heads"?

Possible outcomes: HH, HT, TH, TT

1/4

What is the probability of "both the same"?

Possible outcomes: HH, HT, TH, TT

1/2

Heads or Tails

Heads or Tails

You flip two fair coins

Heads or Tails

You flip two fair coins
X = the value of the sum, where
H counts for *1*, T counts for *0*.

Heads or Tails

You flip two fair coins

X = the value of the sum, where H counts for *1*, T counts for *0*.

What is the probability of X=1?

Heads or Tails

You flip two fair coins

X = the value of the sum, where H counts for *1*, T counts for *0*.

What is the probability of X=1?

Possible outcomes: HH, HT, TH, TT

Heads or Tails

You flip two fair coins

X = the value of the sum, where
H counts for *1*, T counts for *0*.

What is the probability of X=1?

Possible outcomes: HH, HT, TH, TT

1/2

Heads or Tails

You flip two fair coins

X = the value of the sum, where H counts for *1*, T counts for *0*.

What is the probability of X=1?

Possible outcomes: HH, HT, TH, TT

1/2

What is the *expected value* of X?

Heads or Tails

You flip two fair coins

X = the value of the sum, where
H counts for *1*, T counts for *0*.

What is the probability of X=1?

Possible outcomes: HH, HT, TH, TT

1/2

What is the *expected value* of X?

Possible outcomes: HH (2), HT (1), TH (1), TT (0)

Heads or Tails

You flip two fair coins

X = the value of the sum, where
H counts for *1*, T counts for *0*.

What is the probability of X=1?

Possible outcomes: HH, HT, TH, TT

1/2

What is the *expected value* of X?

Possible outcomes: HH (2), HT (1), TH (1), TT (0)

1

Finite Probability Spaces

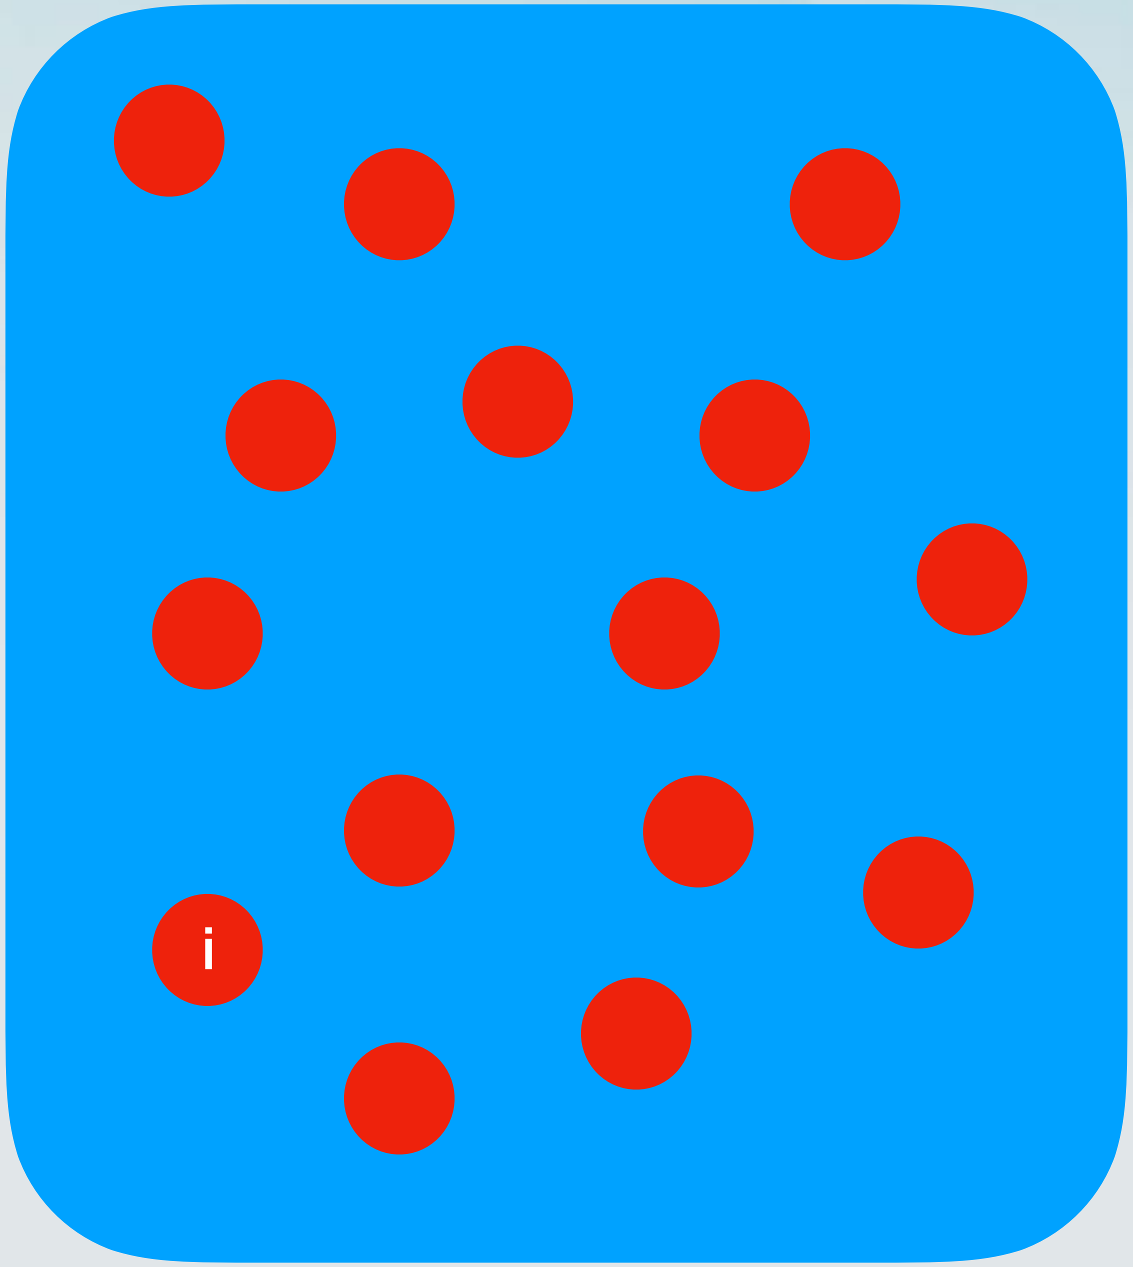

Sample Space Ω

Finite Probability Spaces

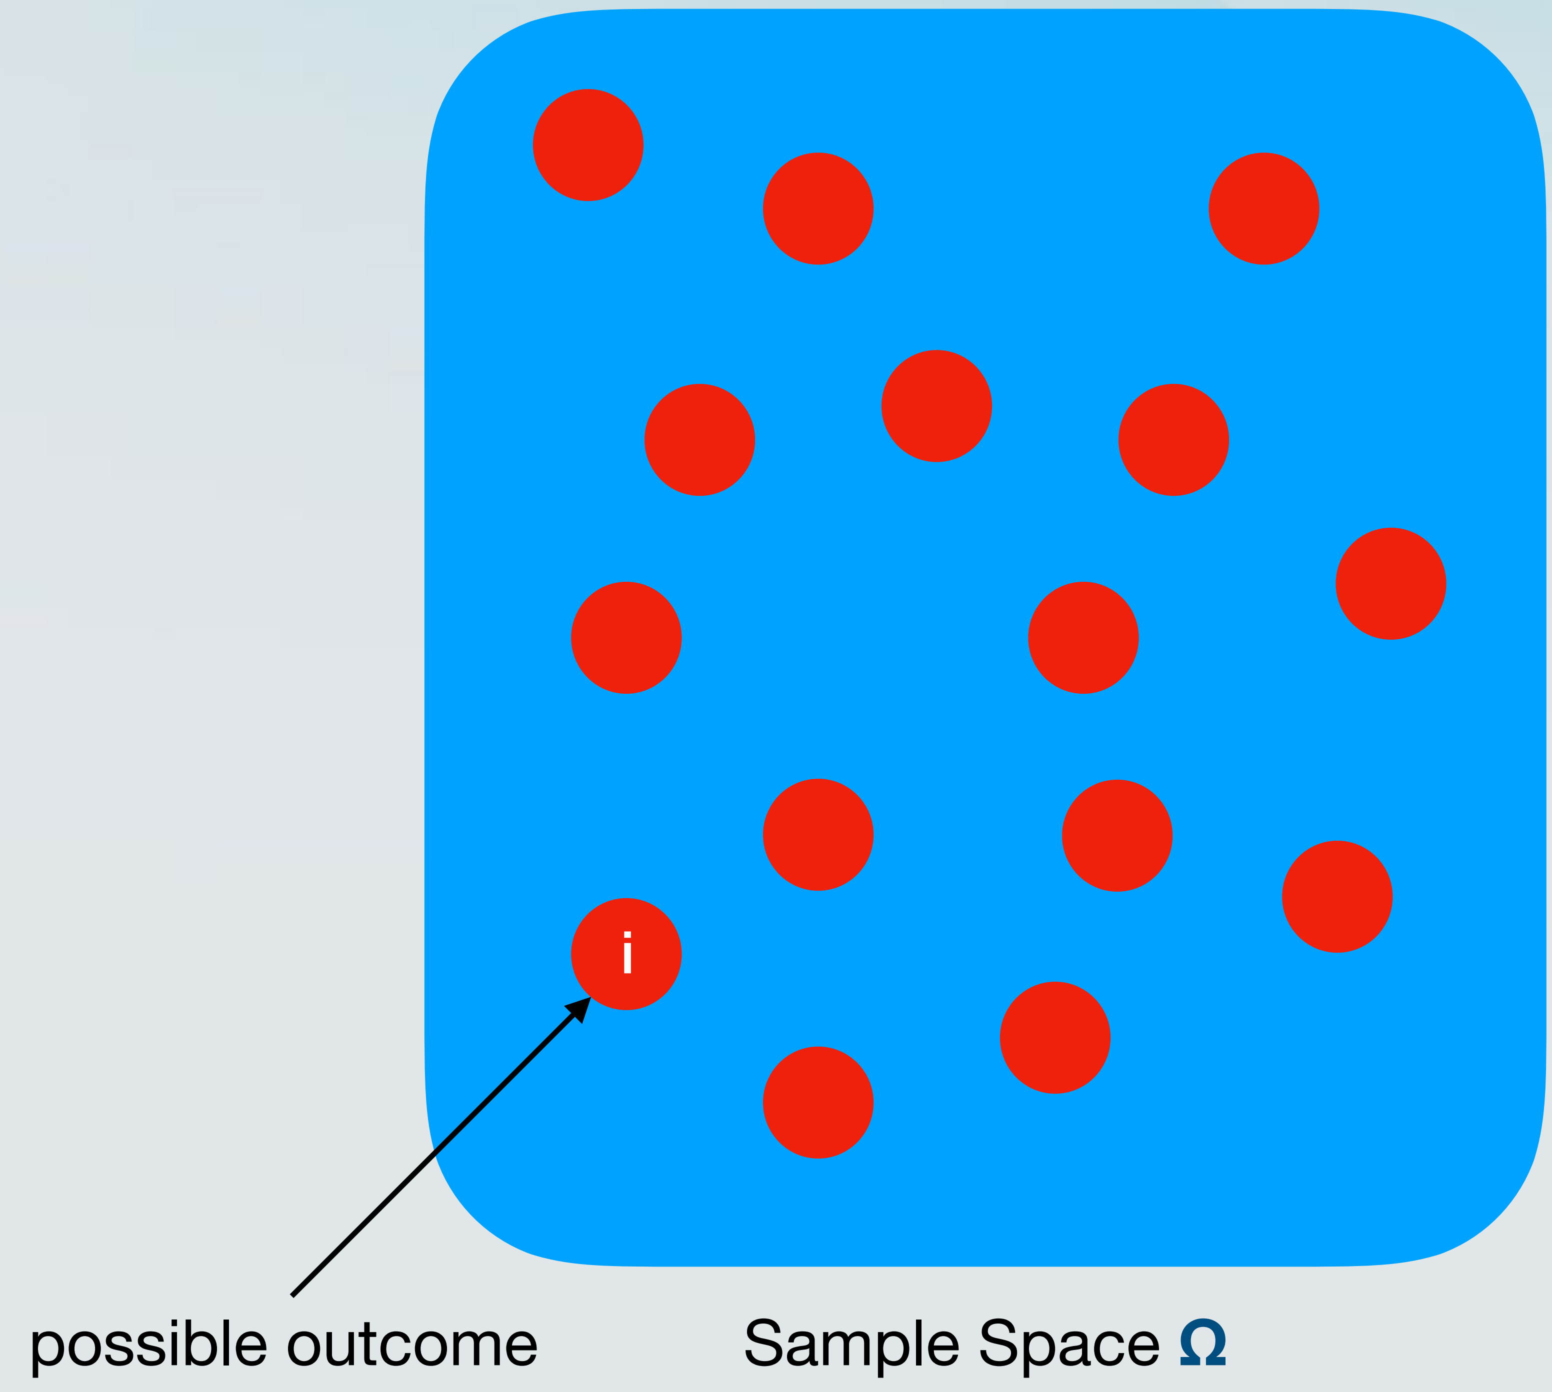

Finite Probability Spaces

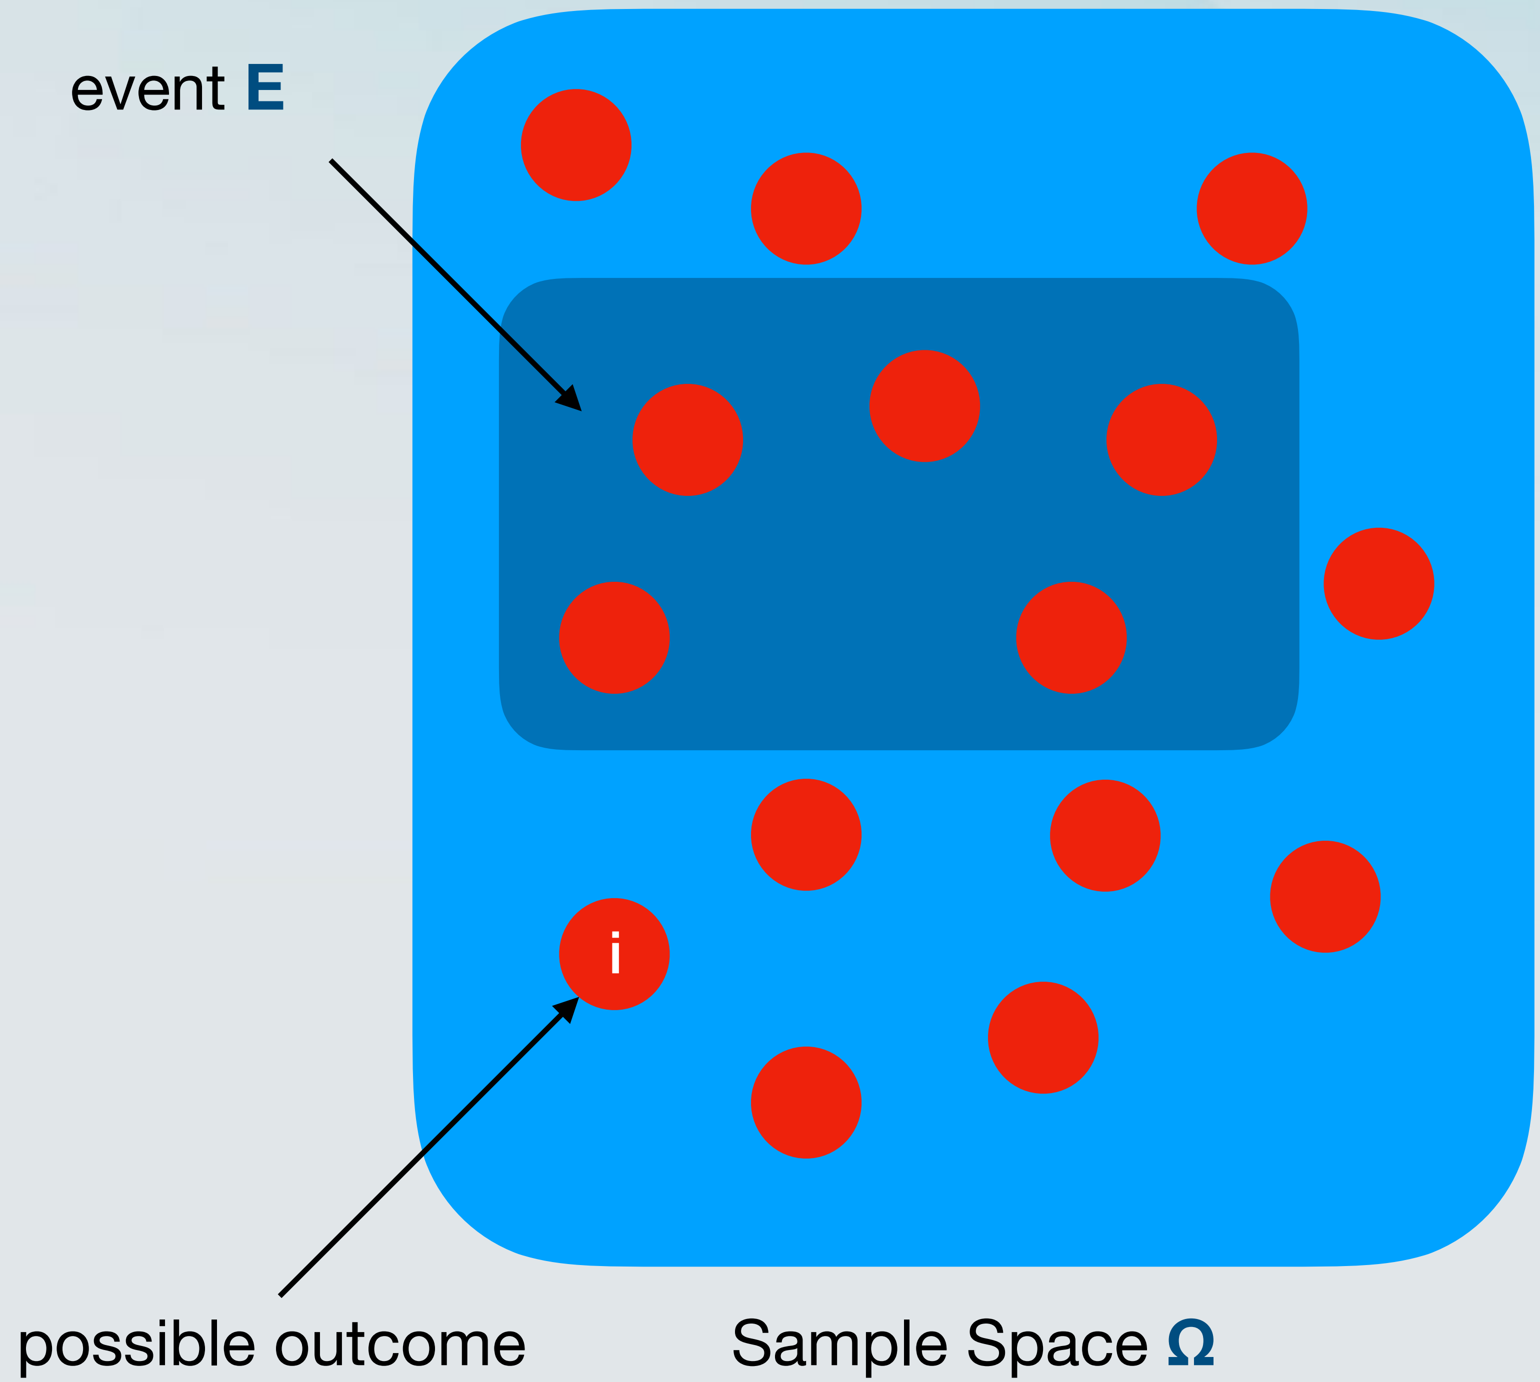

Finite Probability Spaces

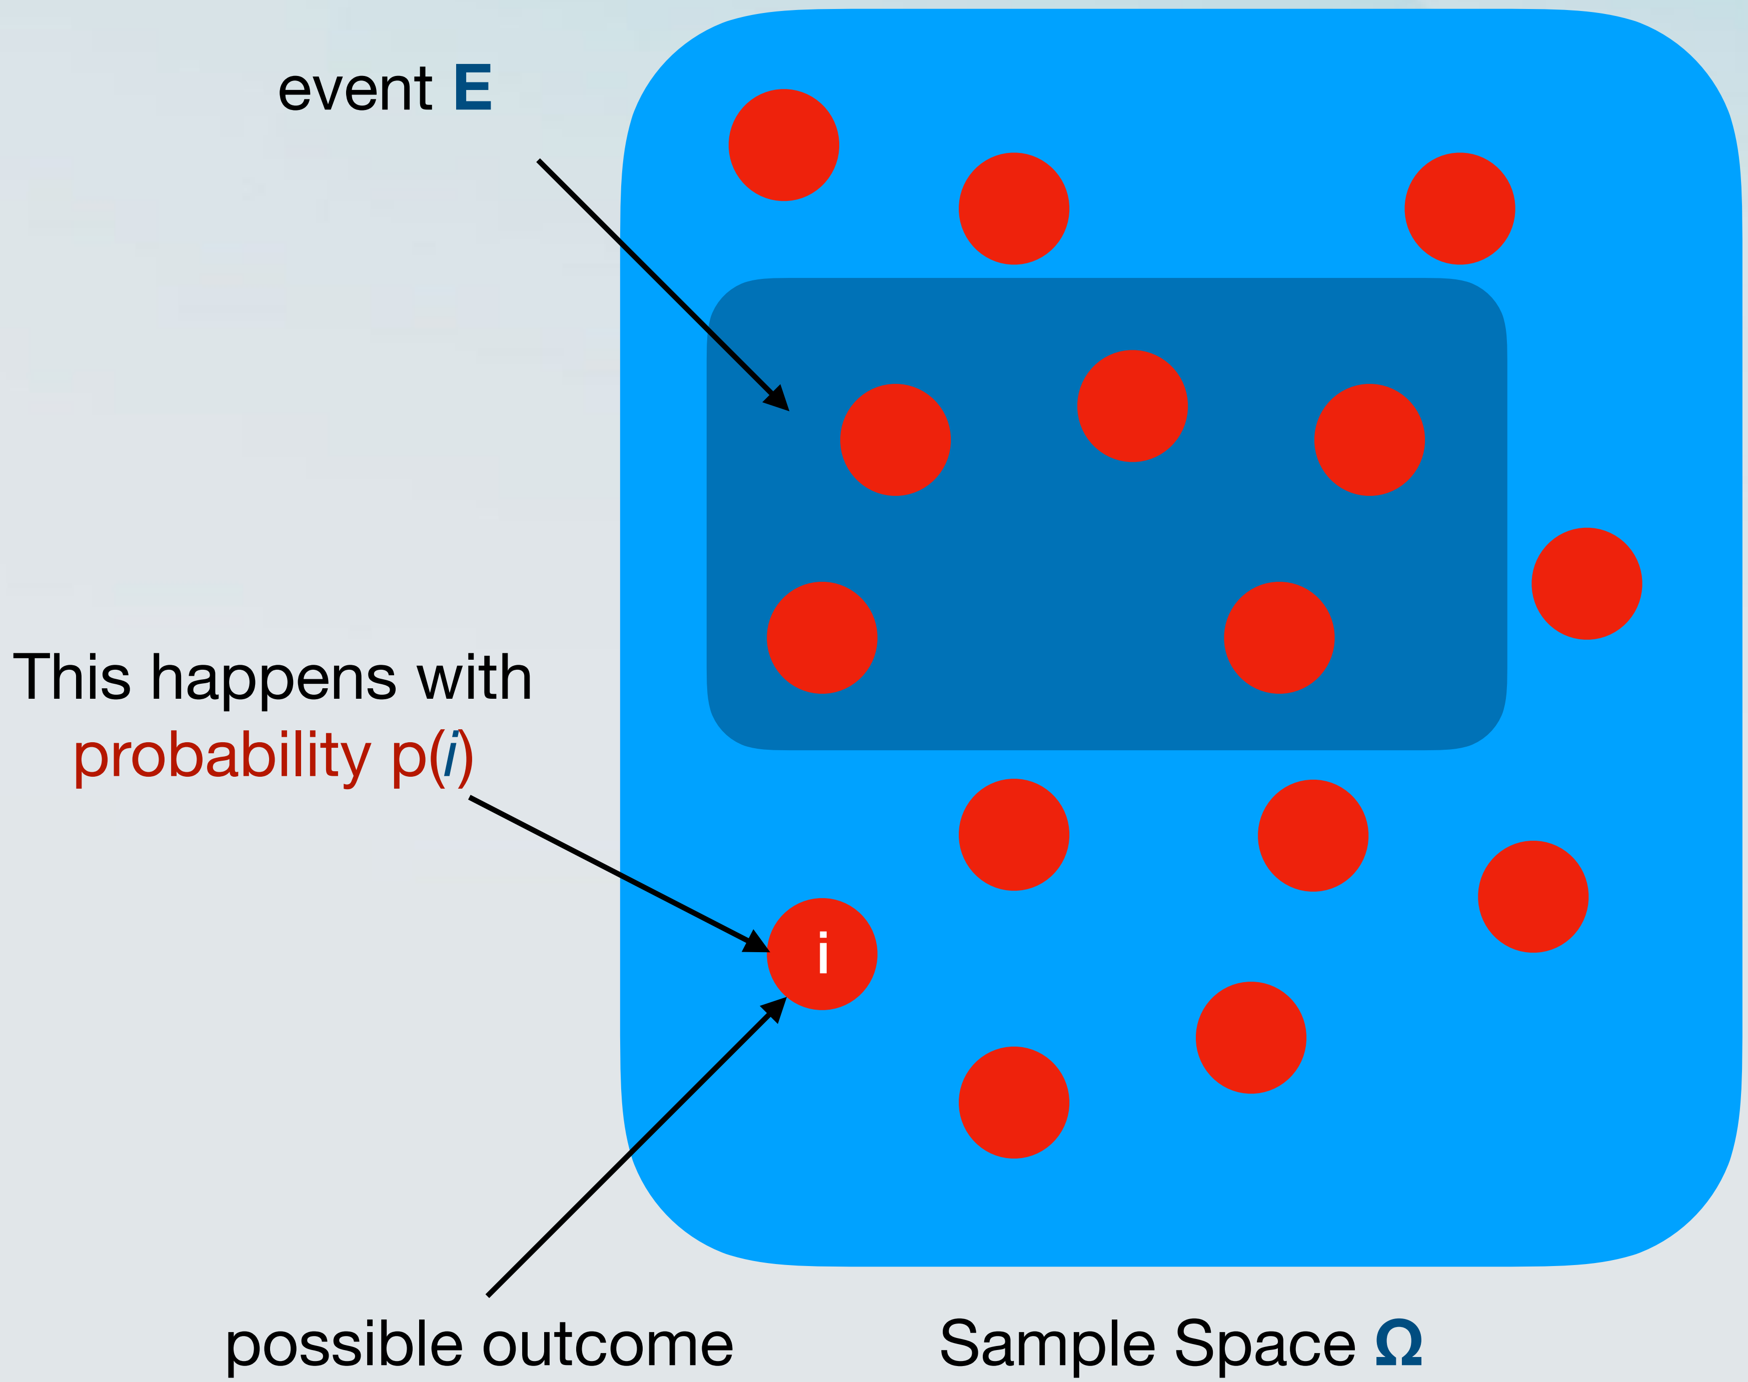

Finite Probability Spaces

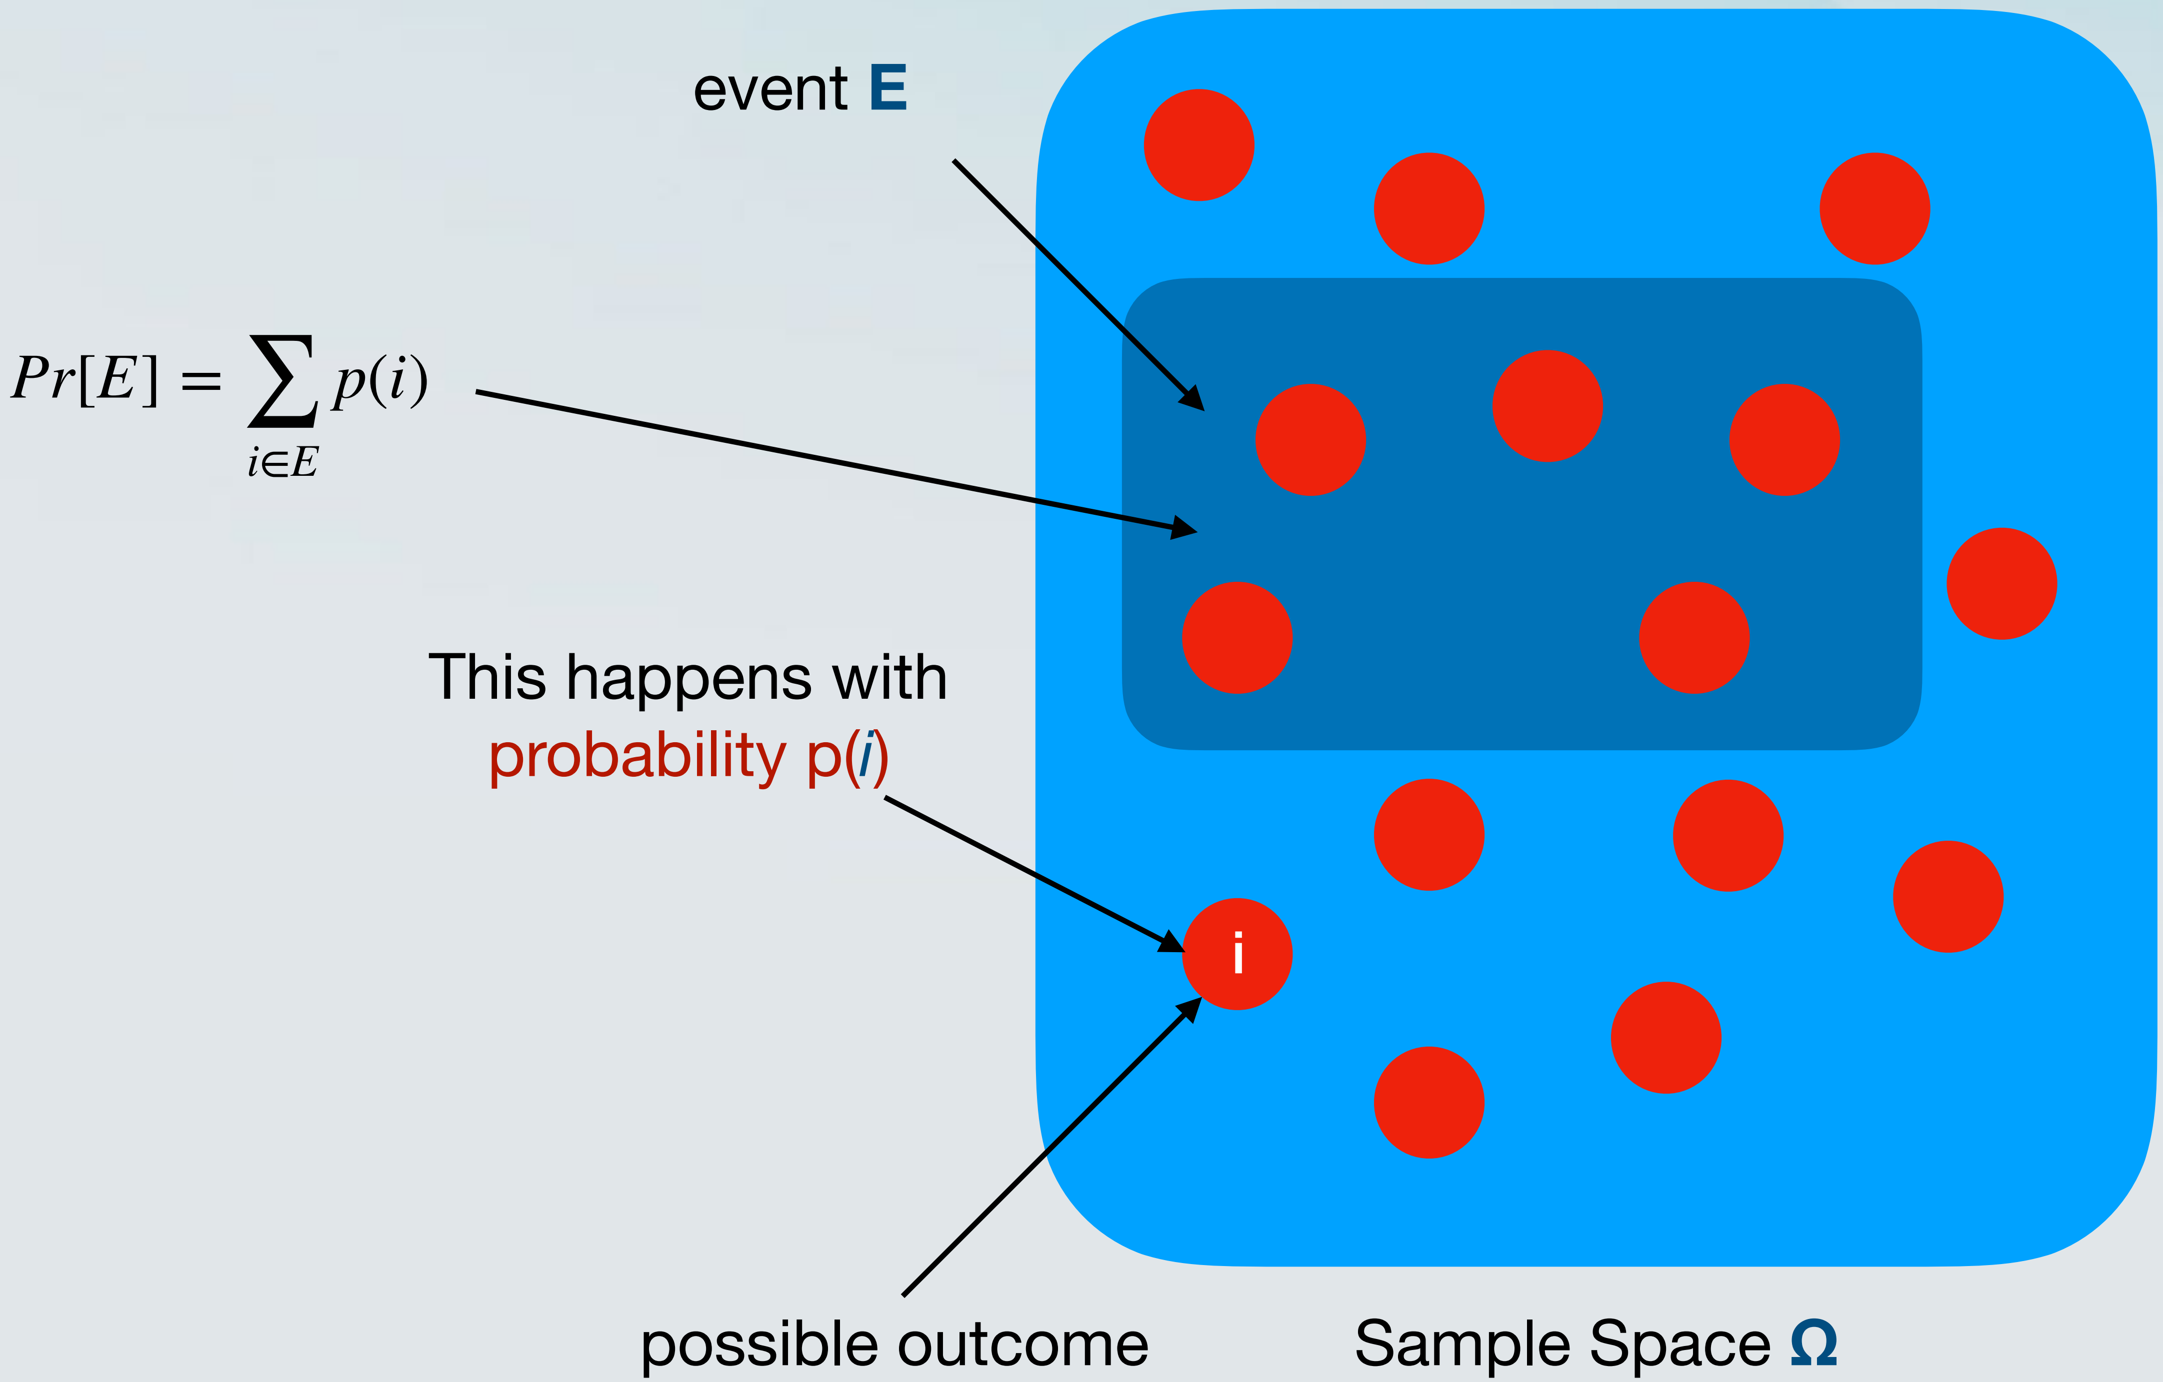

Heads or Tails (fair coin)

Sample Space **Ω**

Heads or Tails (fair coin)

This happens with
probability 1/2

Sample Space **Ω**

Heads or Tails (fair coin)

This happens with probability 1/2

This happens with probability 1/2

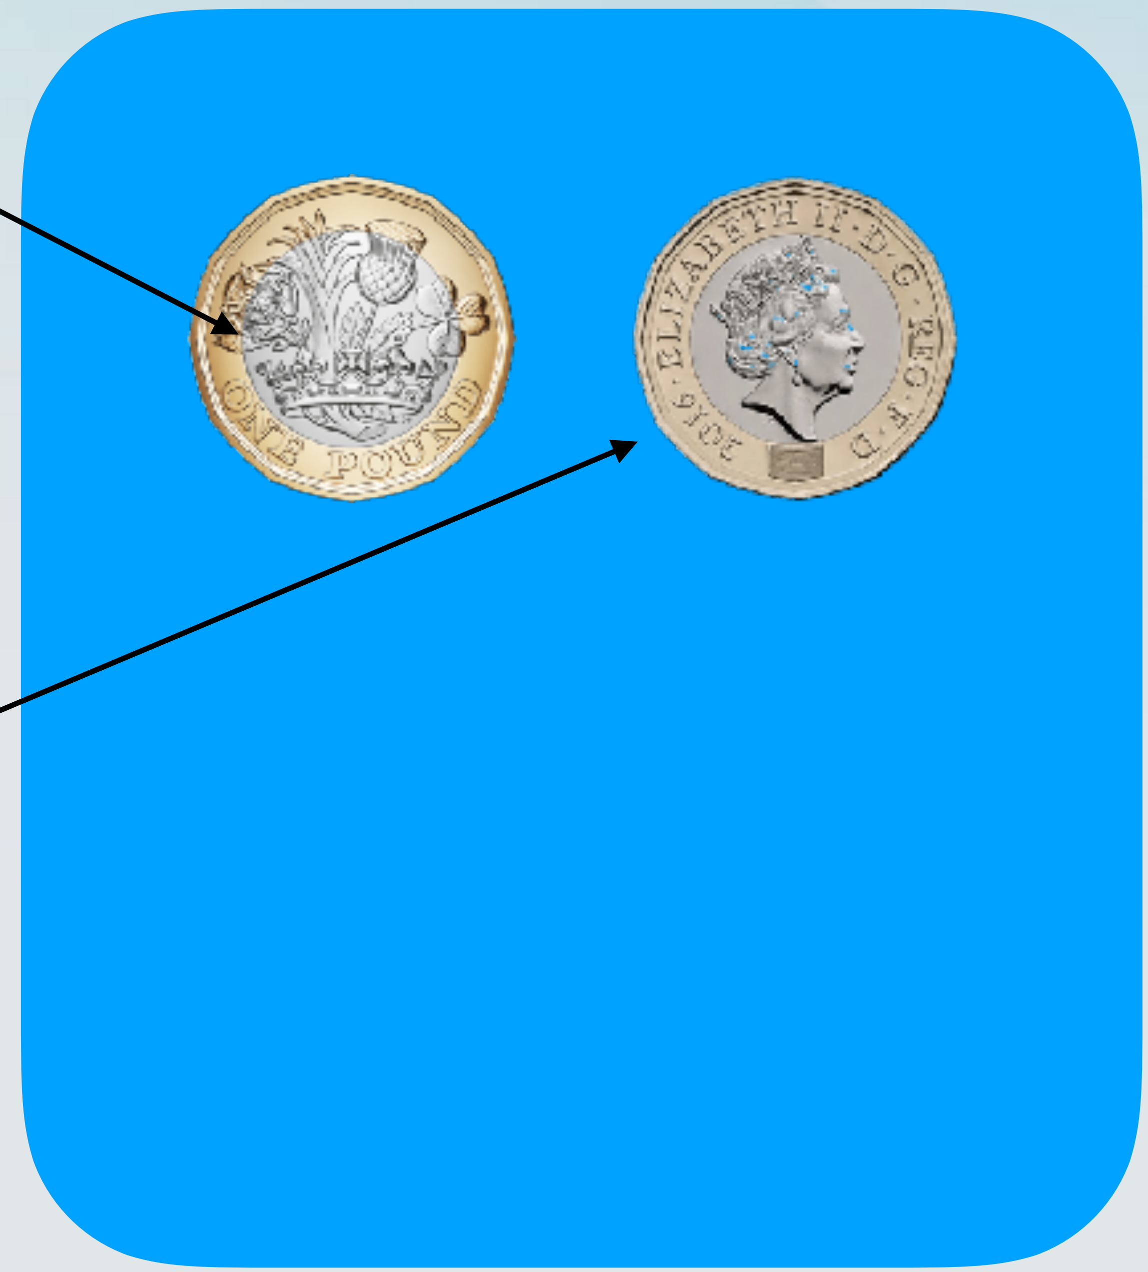

Sample Space **Ω**

Heads or Tails (biased coin)

Sample Space Ω

Heads or Tails (biased coin)

This happens with
probability 2/3

Sample Space **Ω**

Heads or Tails (biased coin)

This happens with
probability 2/3

This happens with
probability 1/3

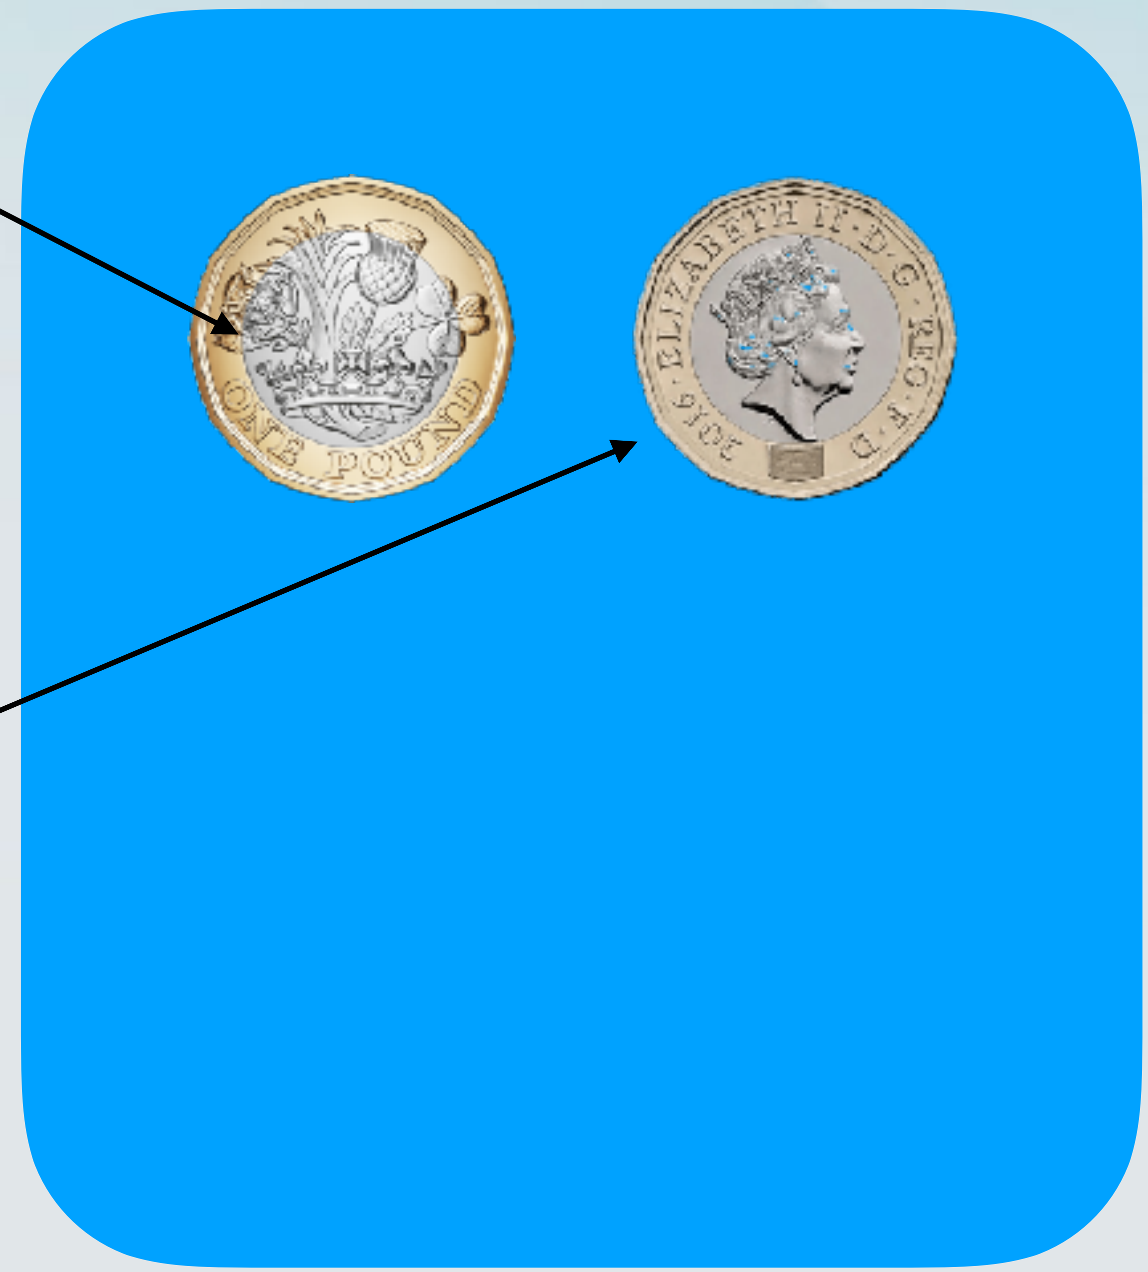

Sample Space **Ω**

Heads or Tails (fair coin)

What is this event?

What is its probability?

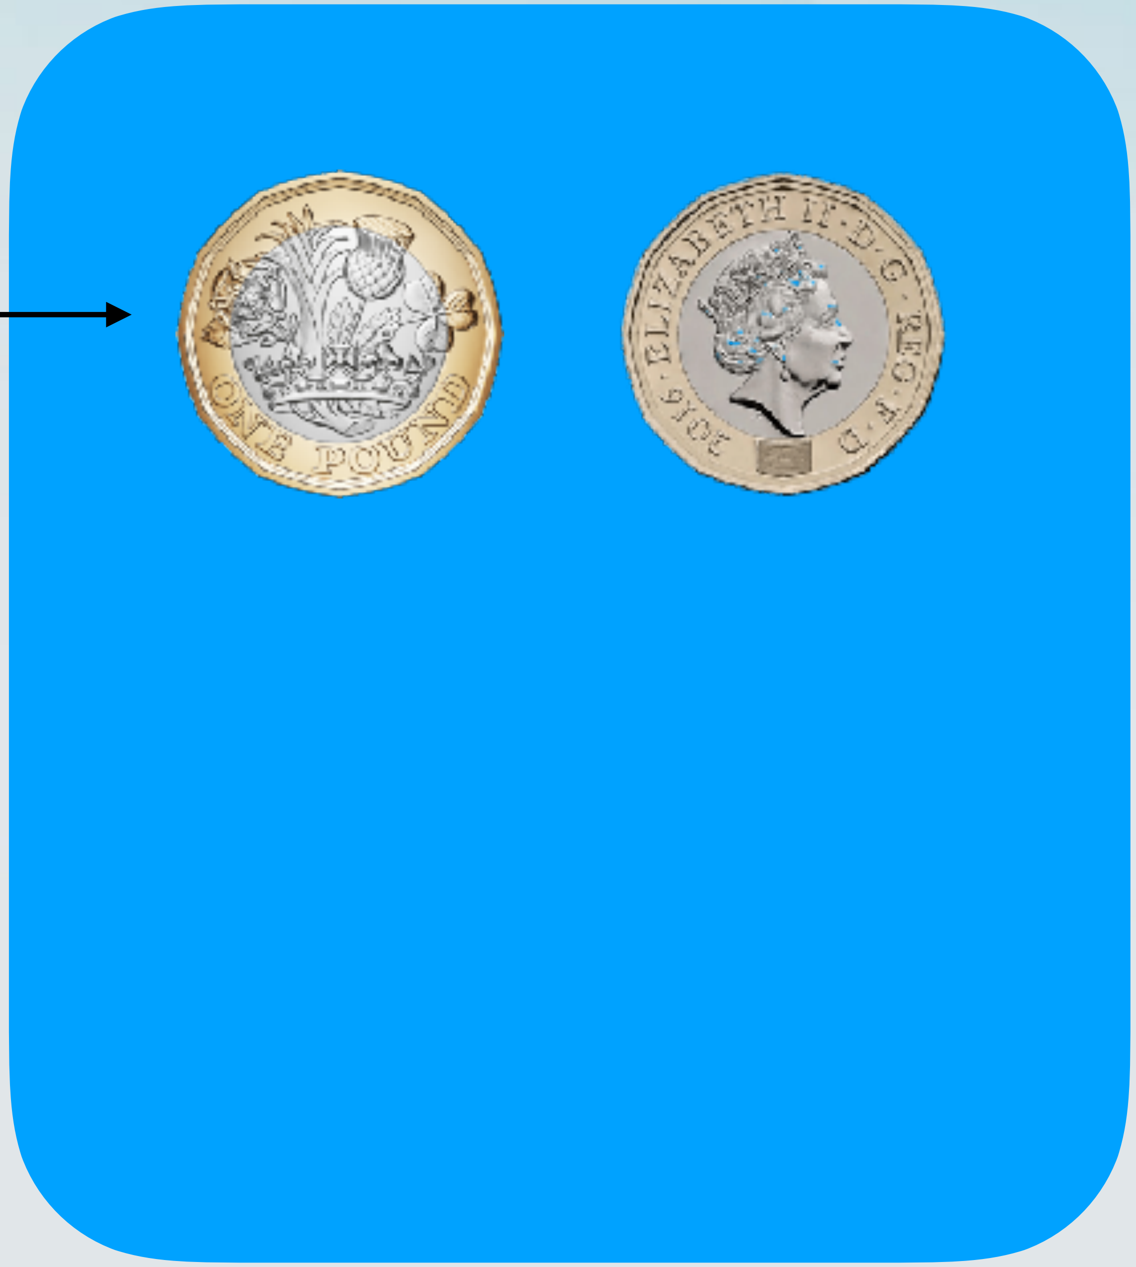

Sample Space Ω

Heads or Tails (fair coin)

What is this event?

What is its probability?

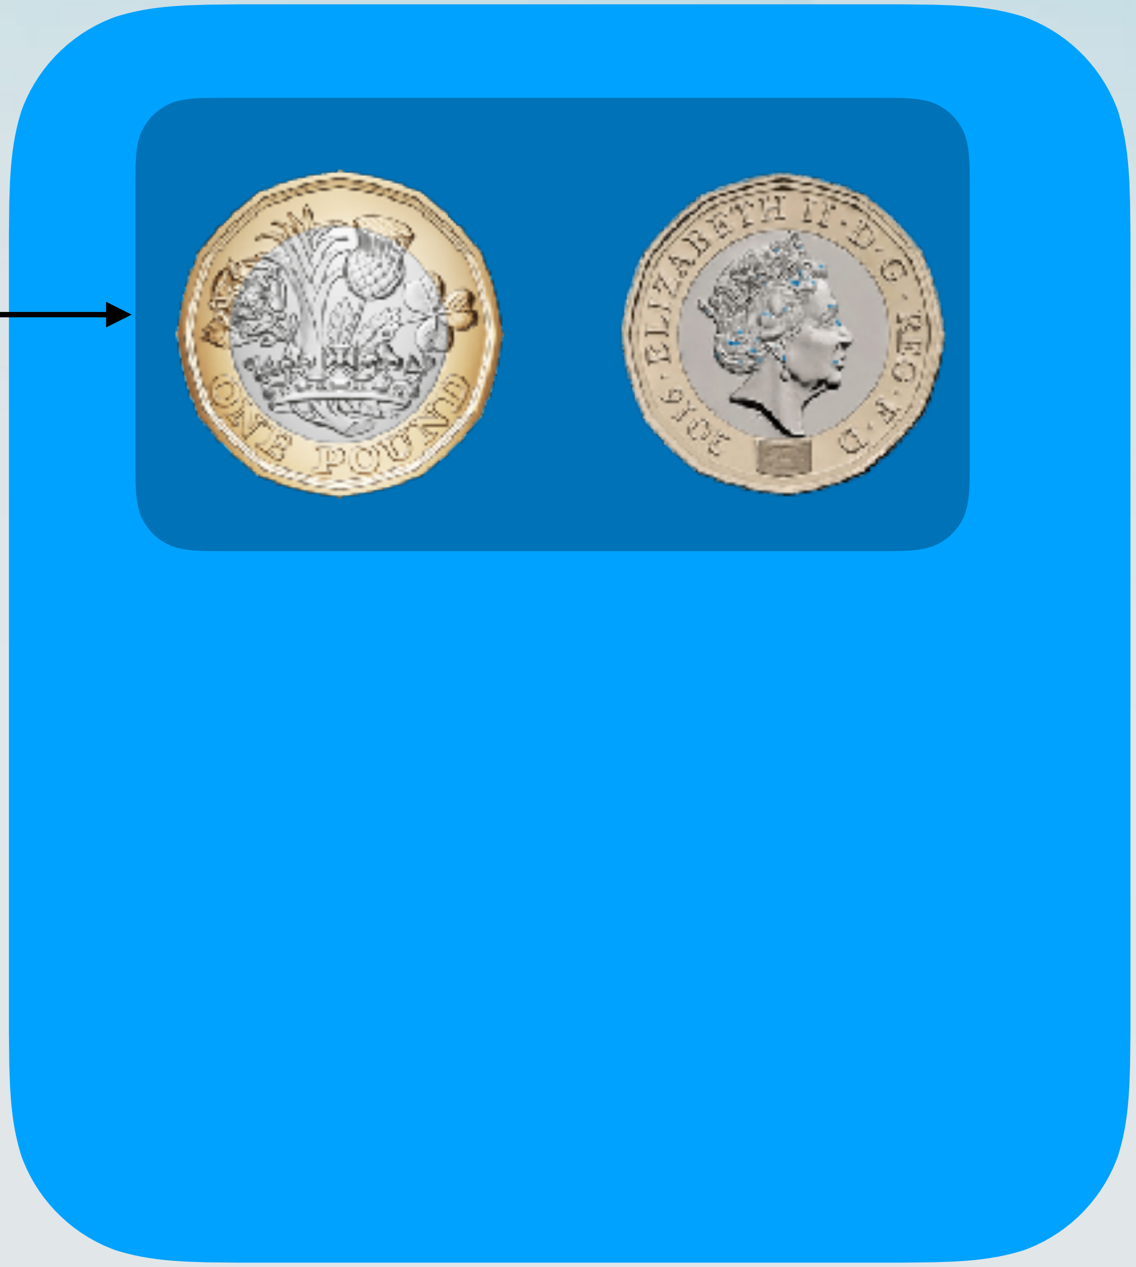

Sample Space **Ω**

Heads or Tails (two fair coins)

Sample Space **Ω**

Heads or Tails (two fair coins)

This happens with
probability 1/4

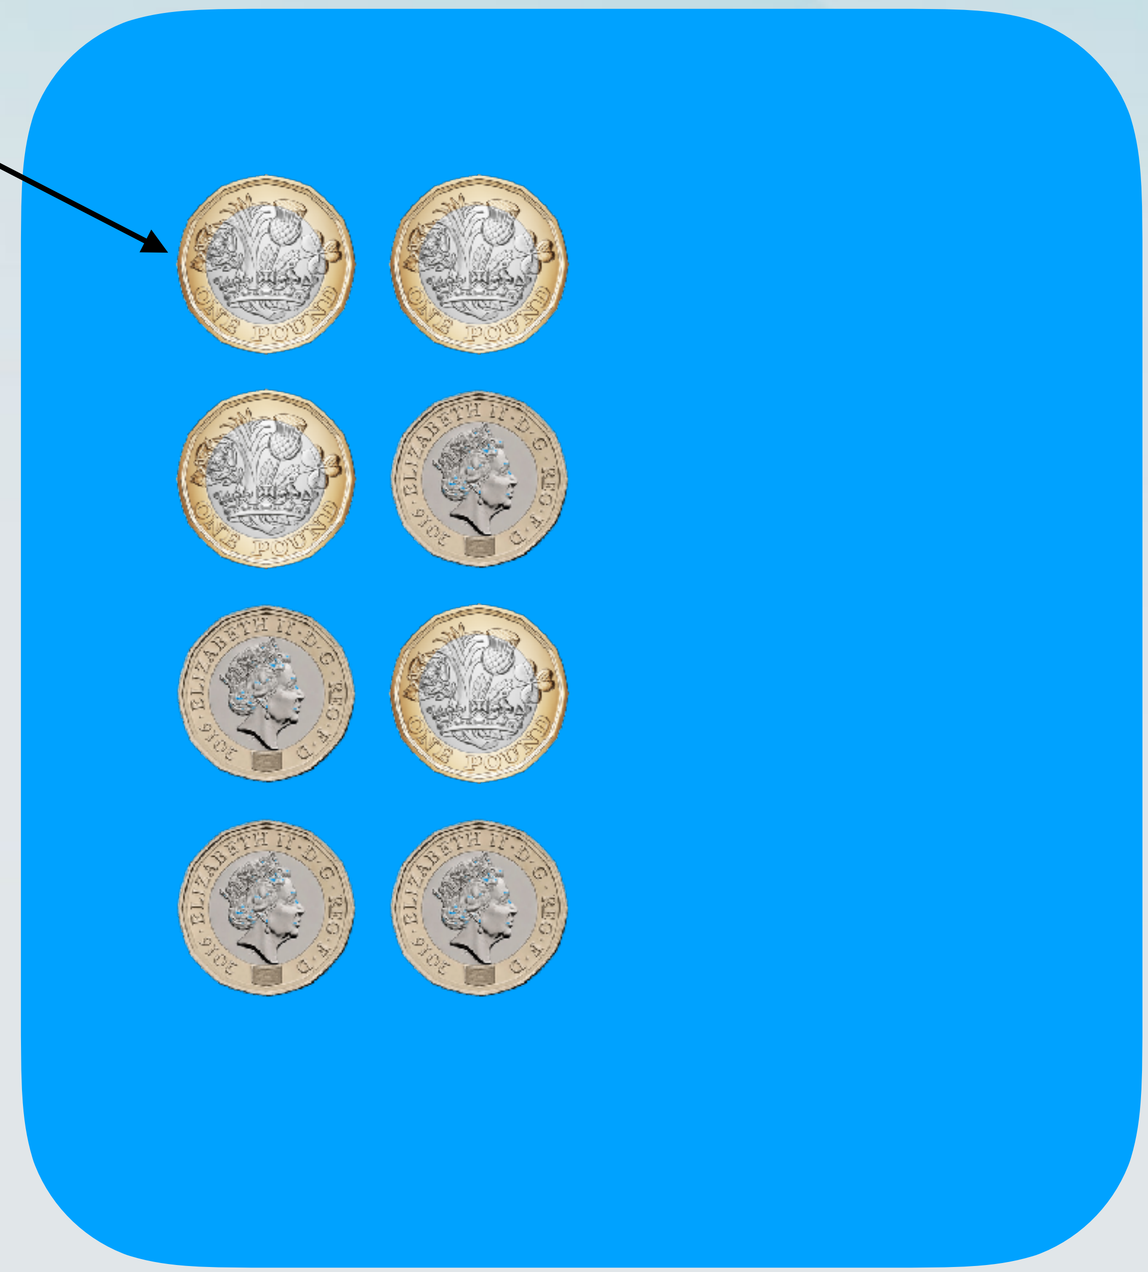

Sample Space **Ω**

Heads or Tails (two fair coins)

This happens with probability 1/4

This happens with probability 1/4

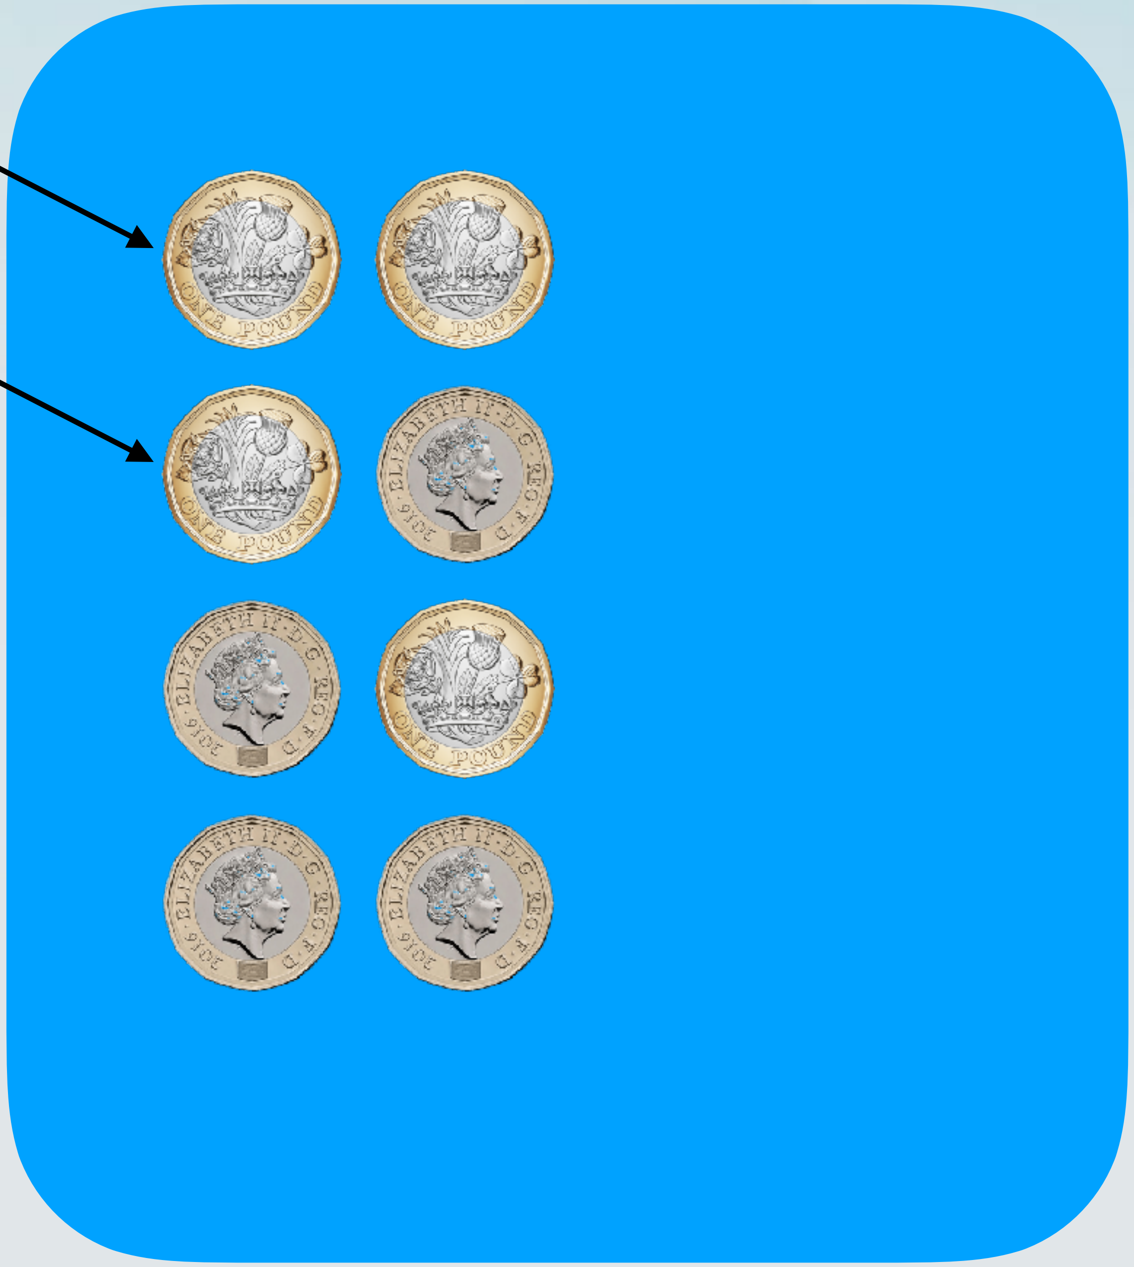

Sample Space **Ω**

Heads or Tails (two fair coins)

This happens with probability 1/4

This happens with probability 1/4

This happens with probability 1/4

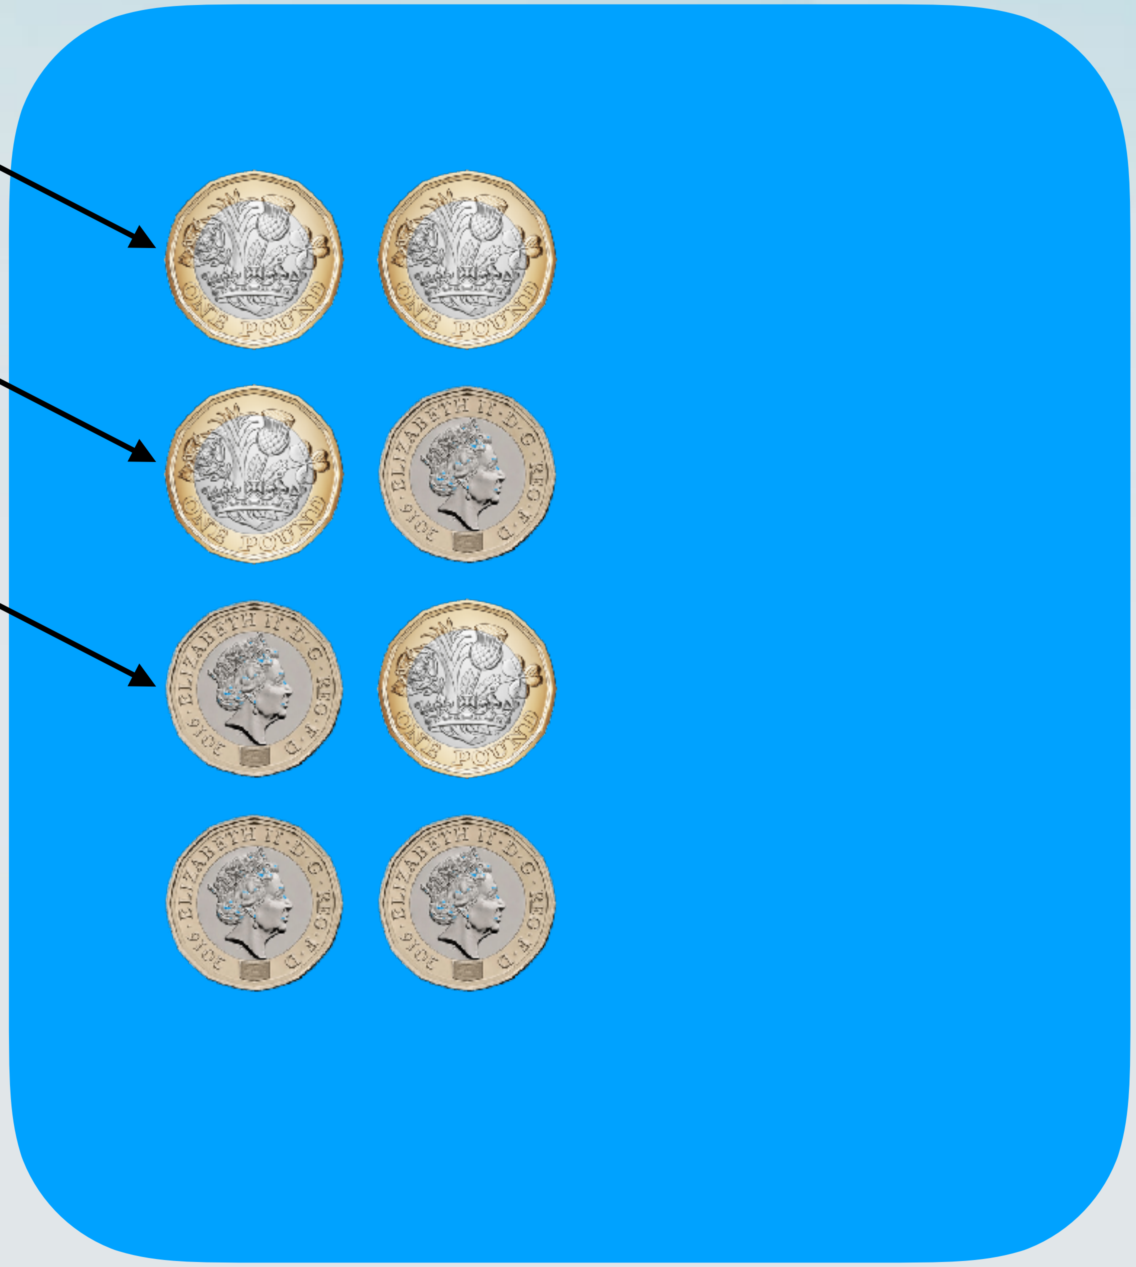

Sample Space Ω

Heads or Tails (two fair coins)

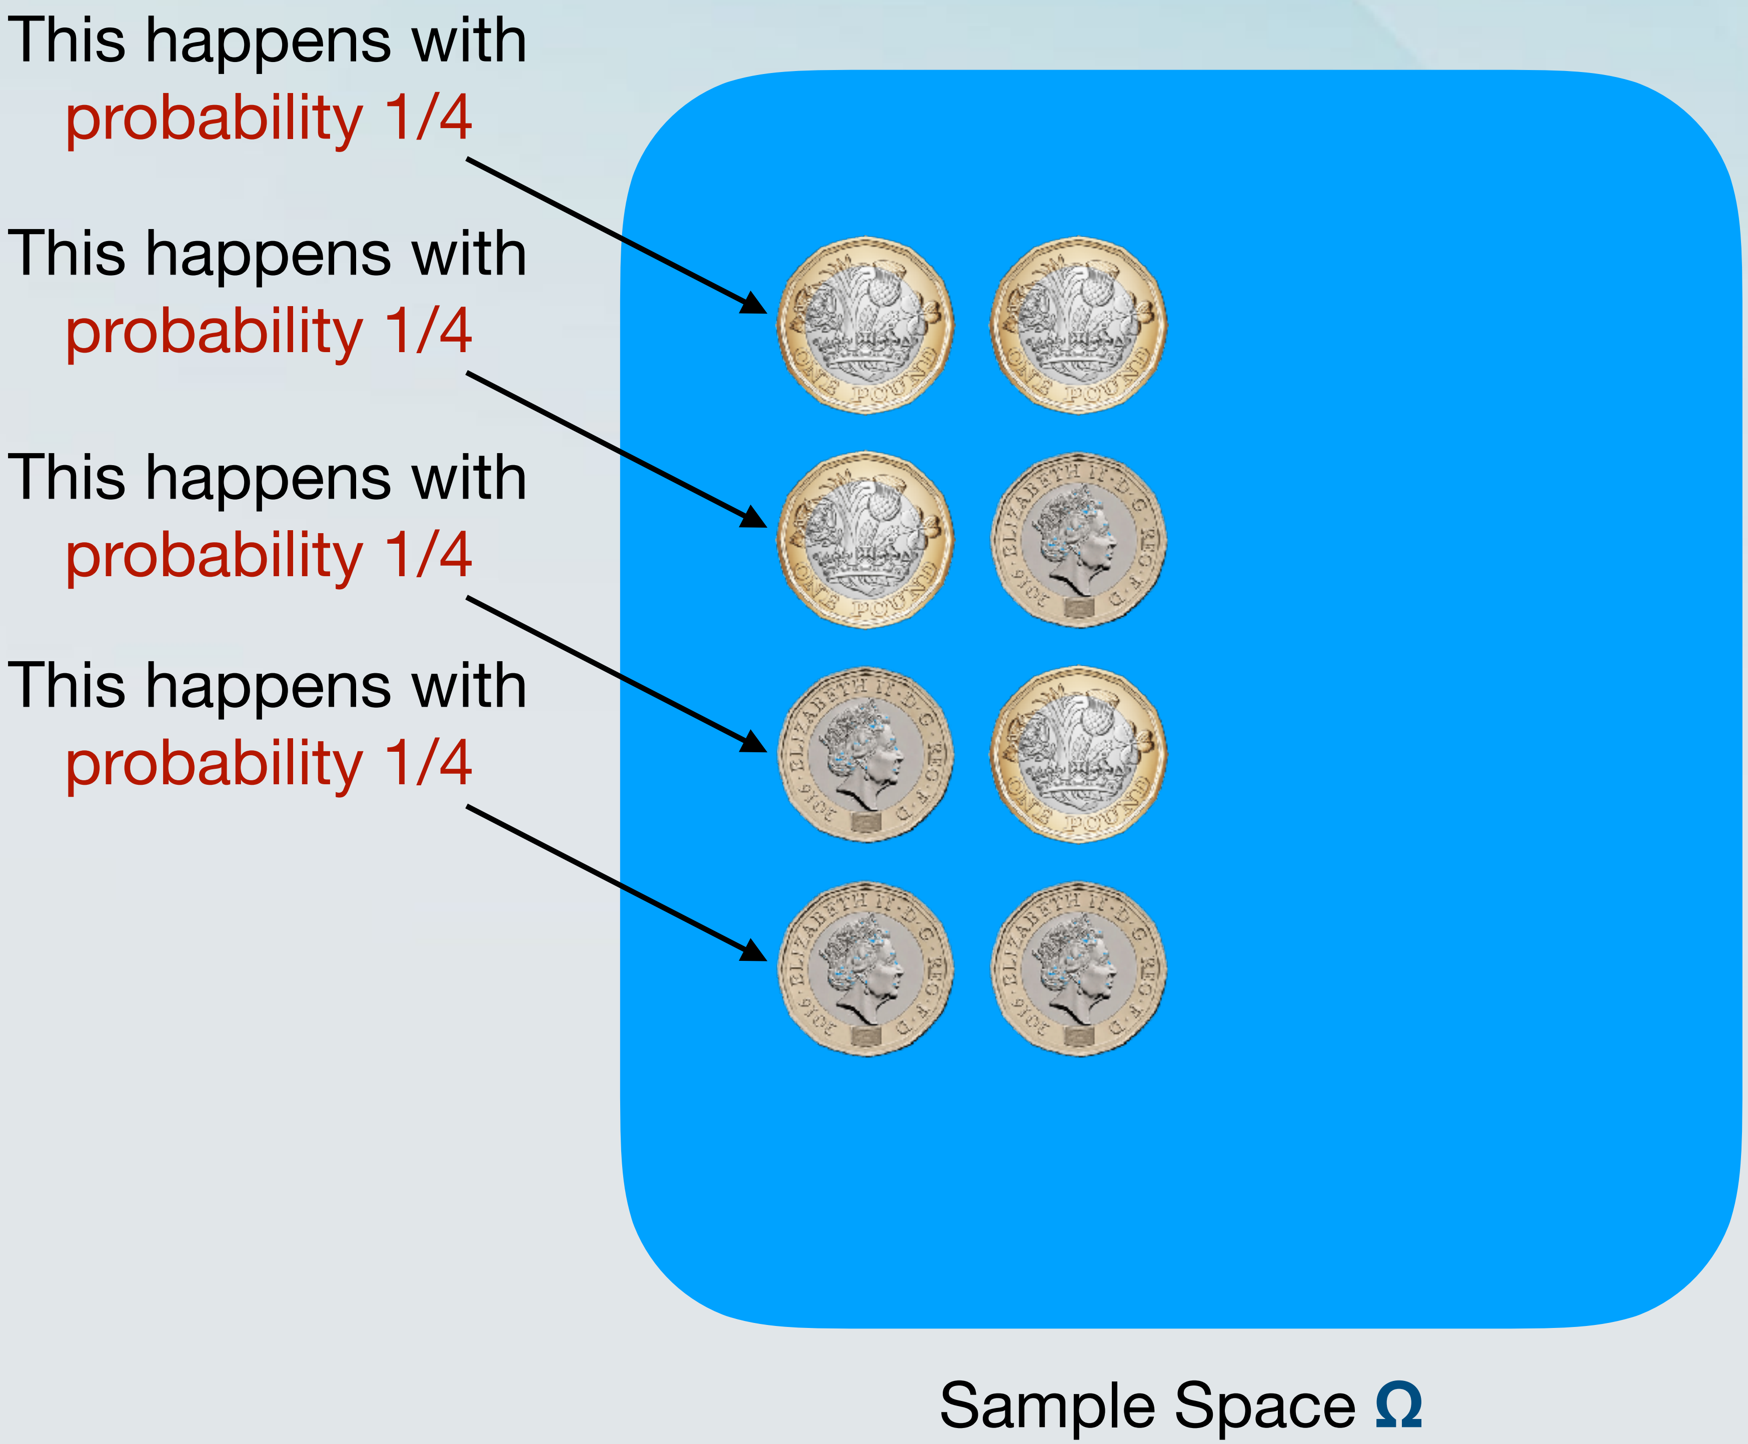

Heads or Tails (two fair coins)

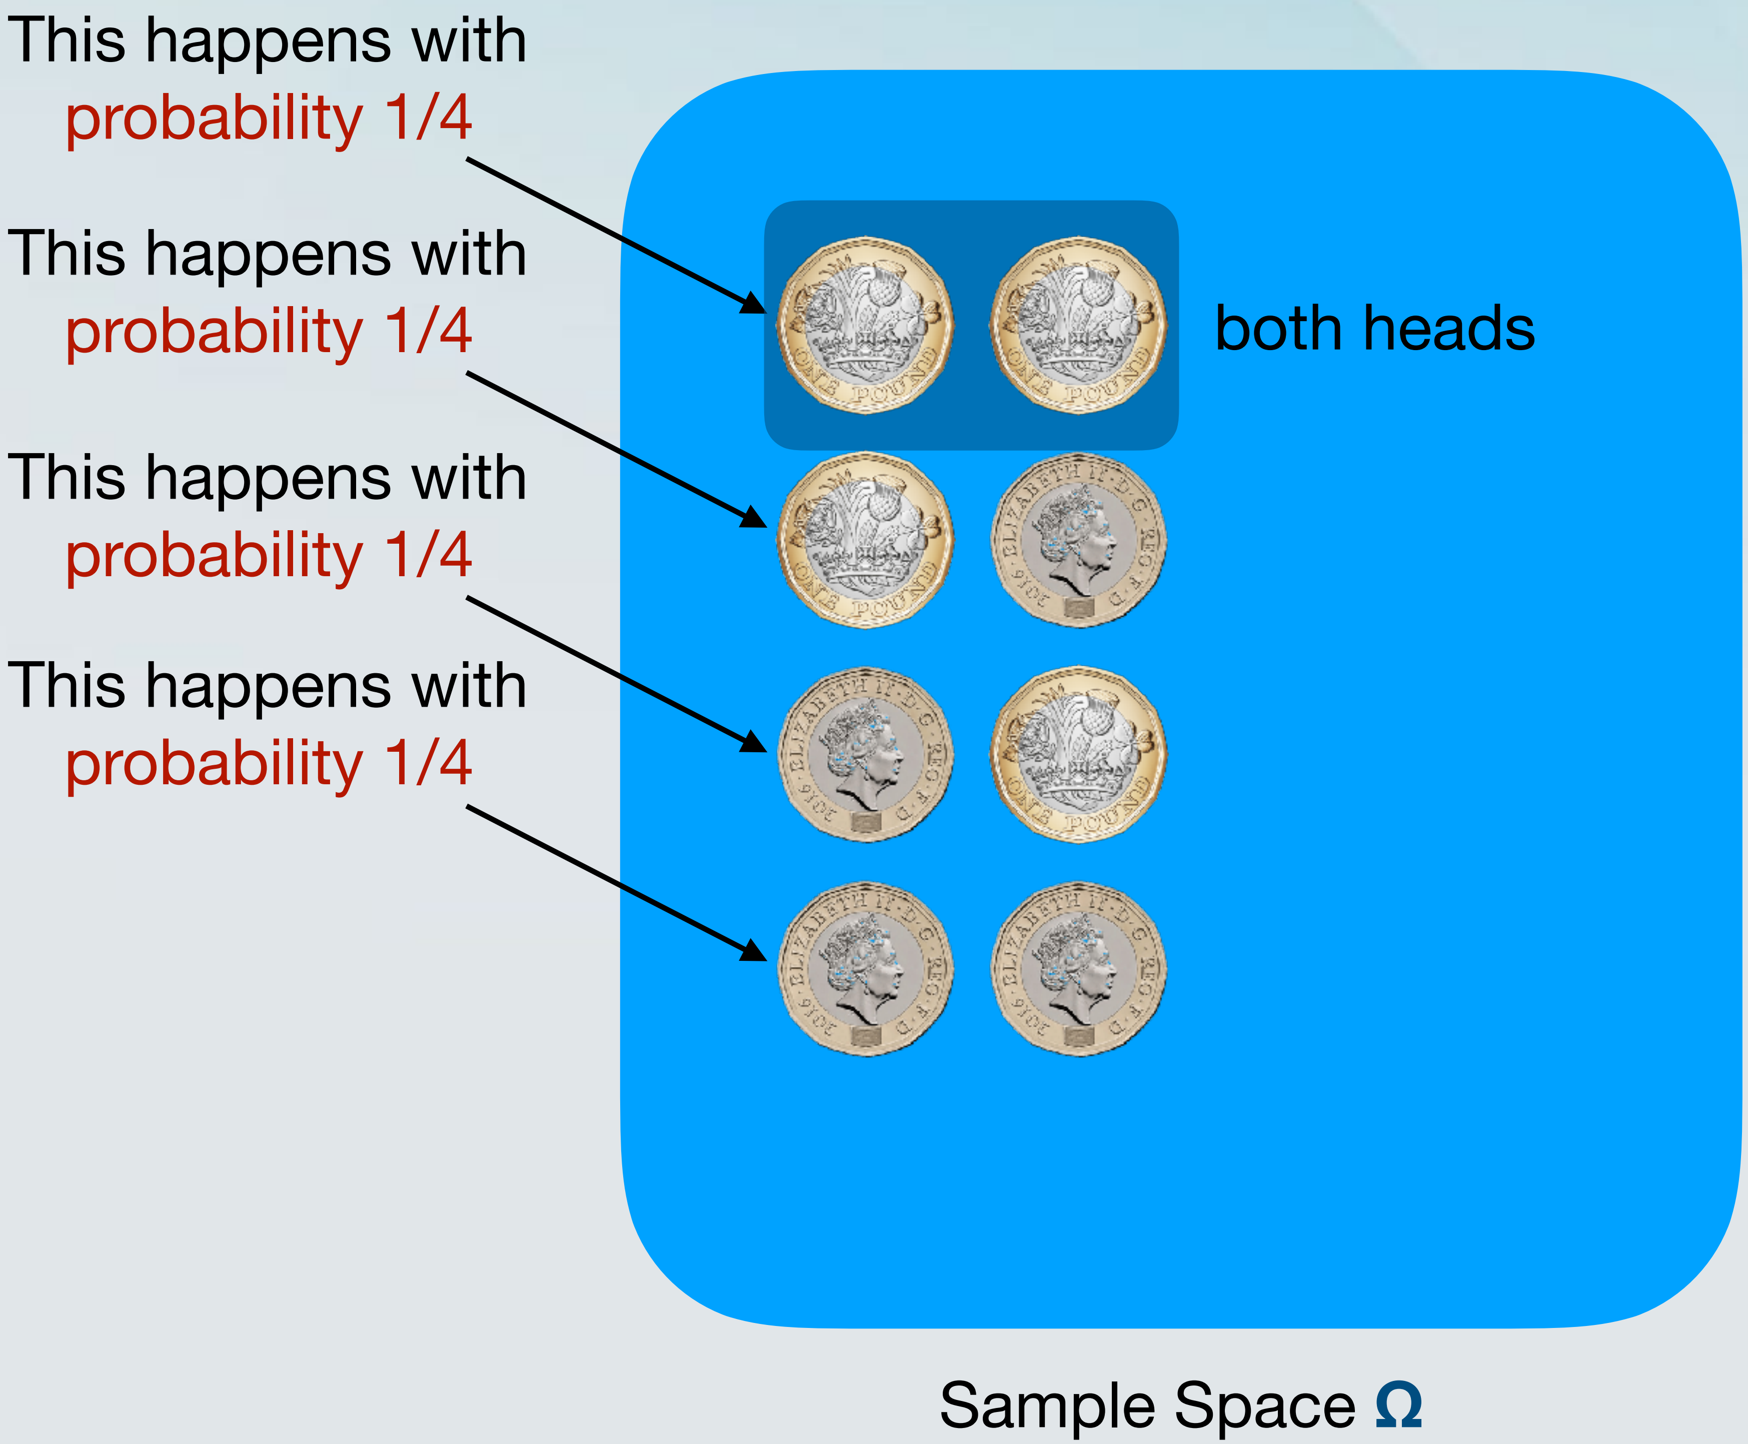

Heads or Tails (two fair coins)

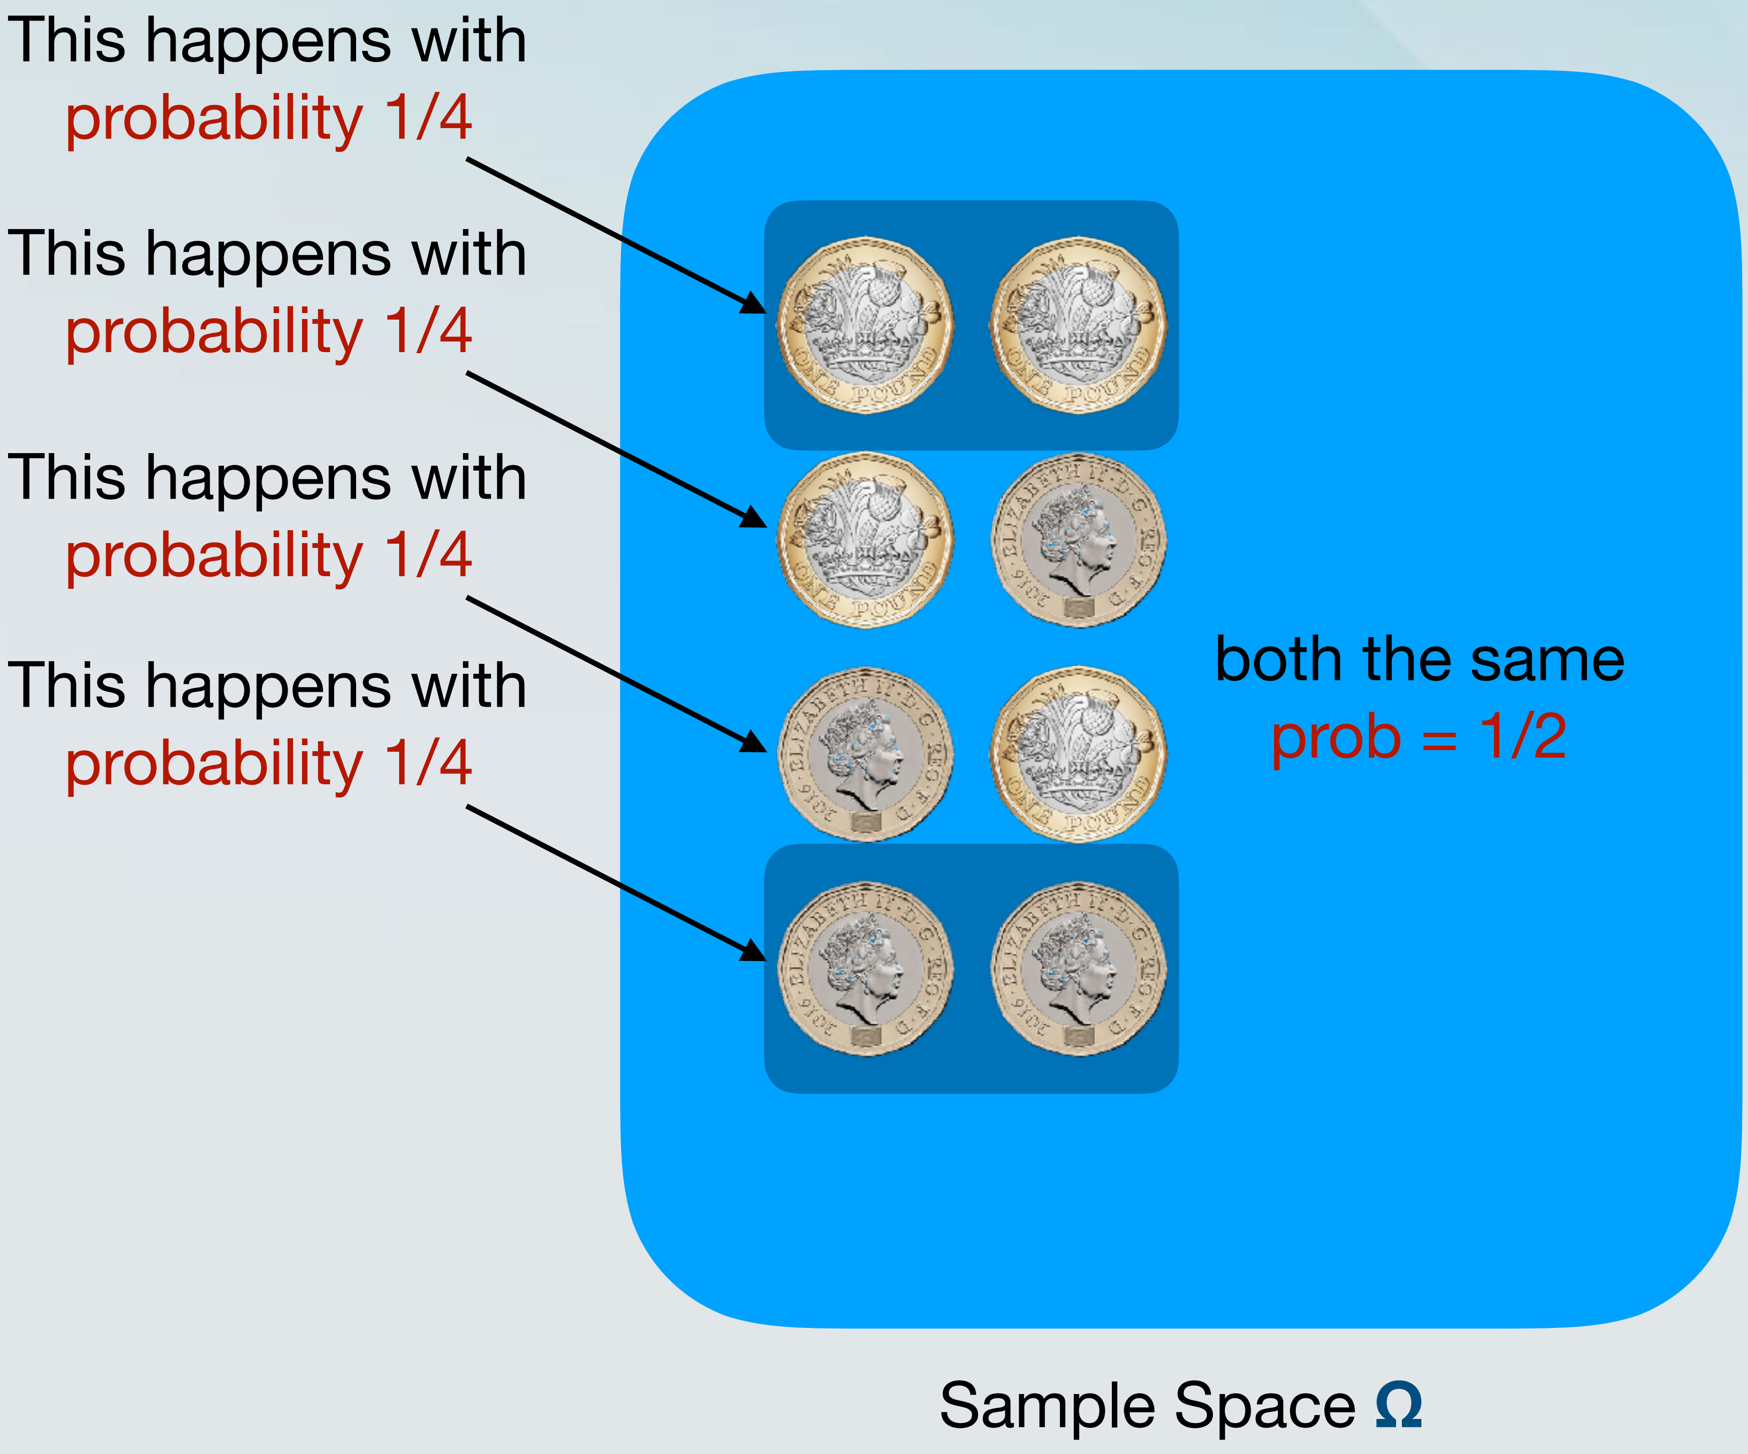

Identifier Selection

Identifier Selection

- There are n processes in a distributed system.

Identifier Selection

- There are *n* processes in a distributed system.
- The set of possible identifiers is the set of all *k*-bit strings.

Identifier Selection

- There are *n* processes in a distributed system.
- The set of possible identifiers is the set of all *k*-bit strings.
    - e.g., *1001001…01*

Identifier Selection

- There are *n* processes in a distributed system.
- The set of possible identifiers is the set of all *k*-bit strings.
  - e.g., *1001001…01*
- Each process chooses an identifier *uniformly at random.*

Identifier Selection

- There are *n* processes in a distributed system.
- The set of possible identifiers is the set of all *k*-bit strings.
  - e.g., *1001001…01*
- Each process chooses an identifier *uniformly at random.*
  - i.e., all strings have equal probability of being chosen.

Identifier Selection

- There are *n* processes in a distributed system.
- The set of possible identifiers is the set of all *k*-bit strings.
  - e.g., *1001001…01*
- Each process chooses an identifier *uniformly at random.*
  - i.e., all strings have equal probability of being chosen.
- What is the probability that processes *1* and *2* choose the same identifier?

Identifier Selection (*1* process)

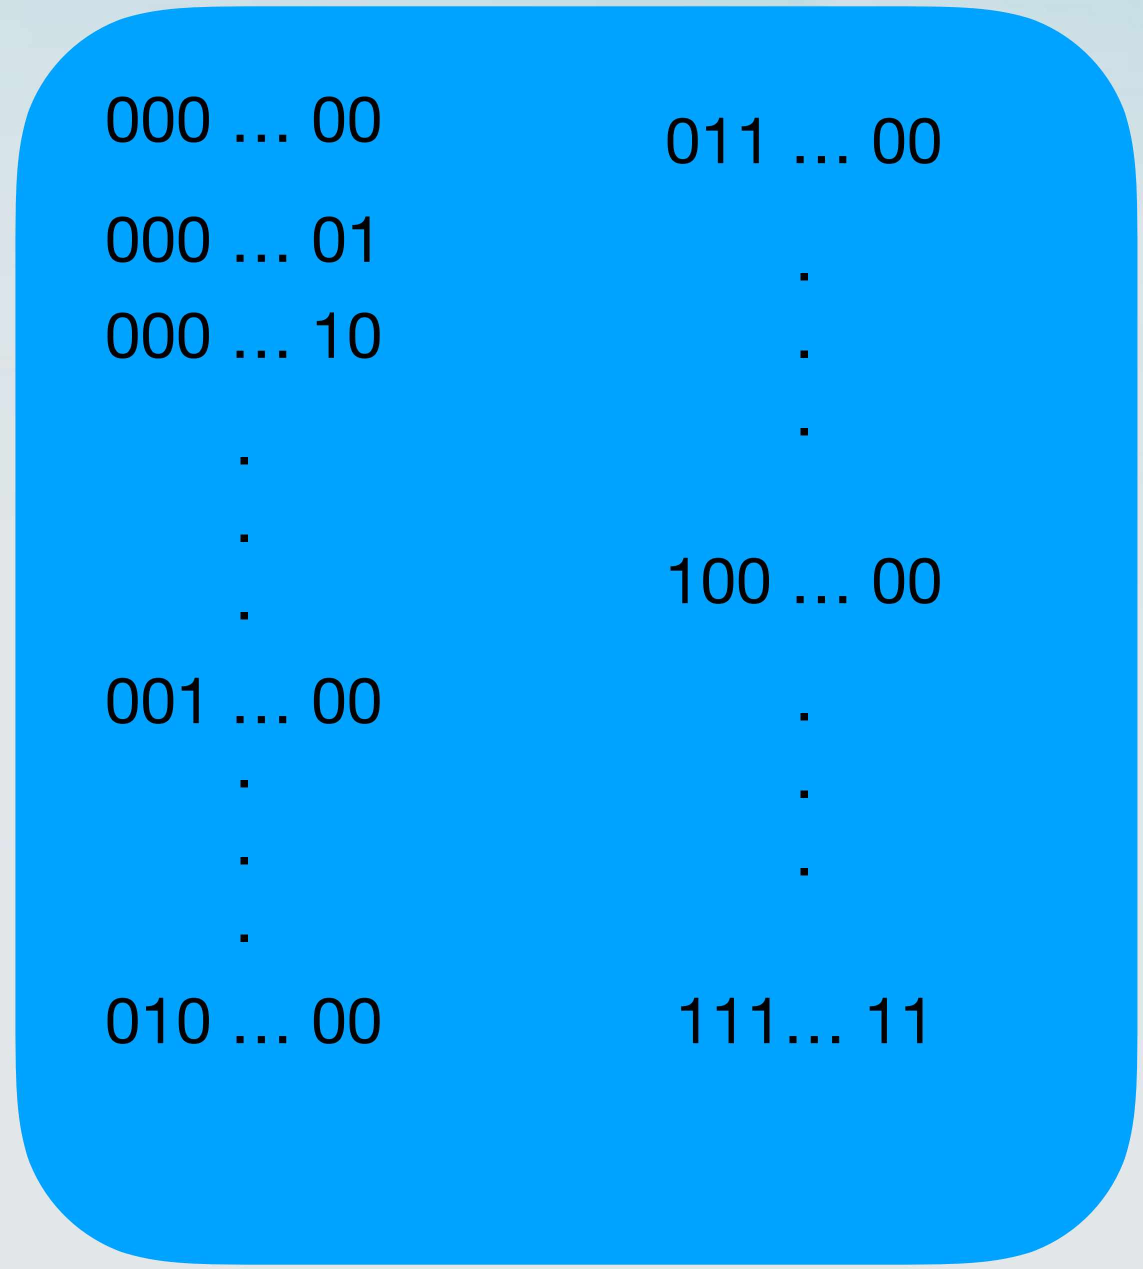

Sample Space **Ω**

Identifier Selection
(*n* processes)

000 … 00 000 … 00 000 … 00 …

.
.
.

111 … 11 011 … 10 011 … 11 …

.
.
.

111… 11 111 … 11 111 … 11 …

Sample Space Ω

Identifier Selection
(*n* processes)

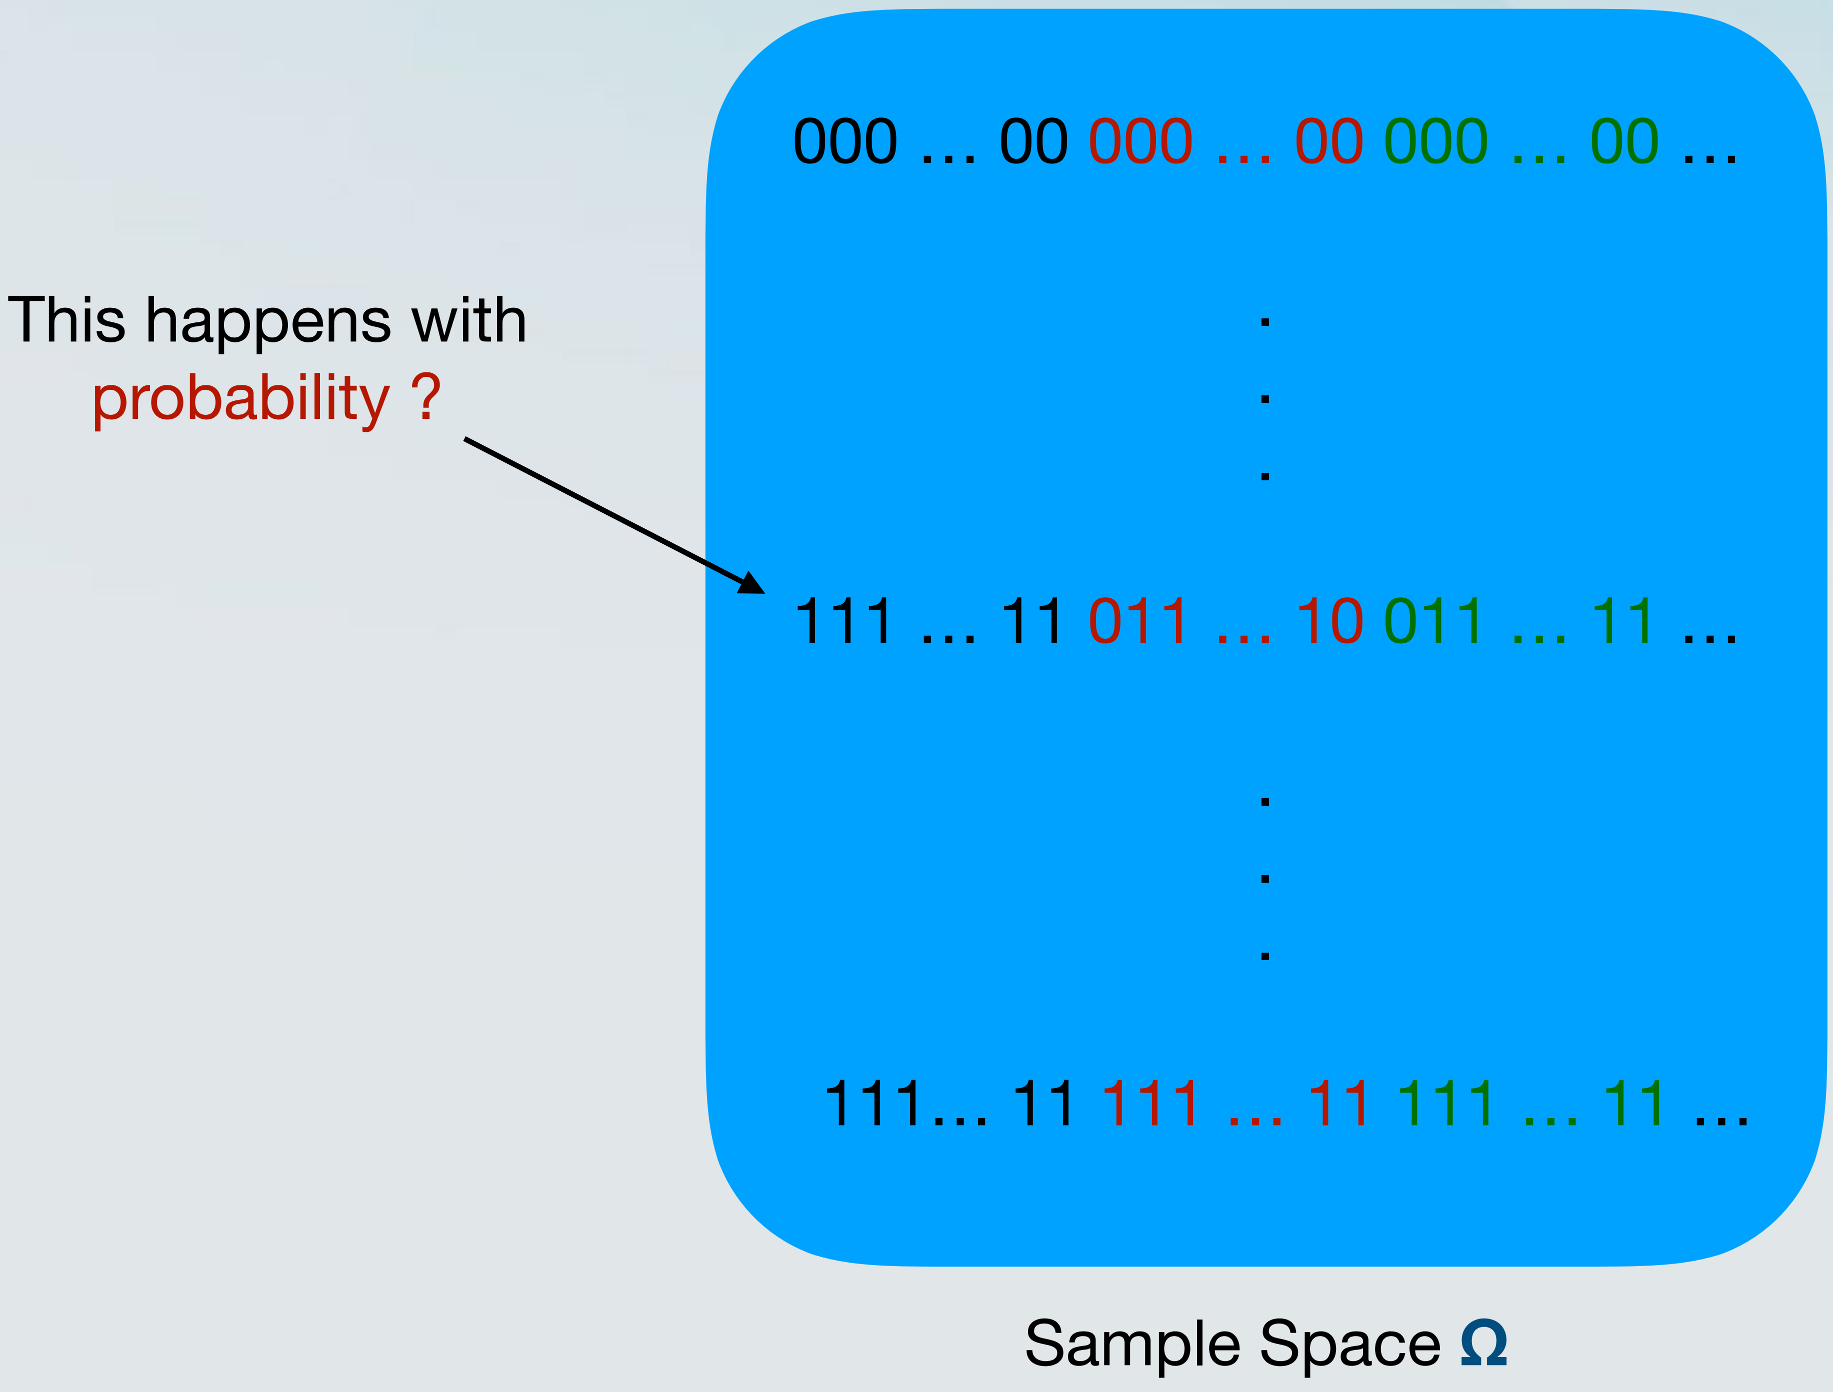

Identifier Selection
(n processes)

This happens with
probability ?

kn possible strings

000 … 00 000 … 00 000 … 00 …

.
.
.

111 … 11 011 … 10 011 … 11 …

.
.
.

111… 11 111 … 11 111 … 11 …

Sample Space **Ω**

Identifier Selection
(*n* processes)

000 … 00 000 … 00 000 … 00 …

.
.
.

111 … 11 011 … 10 011 … 11 …

.
.
.

111… 11 111 … 11 111 … 11 …

Sample Space **Ω**

This happens with probability ?

kn possible strings

2^{kn} possible choices

Identifier Selection
(*n* processes)

This happens with
probability ?

000 … 00 000 … 00 000 … 00 …

.
.
.

111 … 11 011 … 10 011 … 11 …

.
.
.

111… 11 111 … 11 111 … 11 …

Sample Space **Ω**

kn possible strings

2^{kn} possible choices

2^{kn} possible points

Identifier Selection
(*n* processes)

000 … 00 000 … 00 000 … 00 …

.
.
.

111 … 11 011 … 10 011 … 11 …

.
.
.

111… 11 111 … 11 111 … 11 …

Sample Space Ω

This happens with probability ?

kn possible strings

2^{kn} possible choices

2^{kn} possible points

same probability for all

Identifier Selection
(*n* processes)

This happens with
probability $1/2^{kn}$

kn-bit strings

2^{kn} possible choices

2^{kn} possible points

same probability for all

000 … 00 000 … 00 000 … 00 …

.
.
.

111 … 11 011 … 10 011 … 11 …

.
.
.

111… 11 111 … 11 111 … 11 …

Sample Space **Ω**

Identifier Selection

- What is the probability that processes *1* and *2* choose the same identifier?

Identifier Selection

- What is the probability that processes *1* and *2* choose the same identifier?
- The event E consists of all the points for which the first two coordinates (**black** and **red**) are the same.

Identifier Selection

- What is the probability that processes *1* and *2* choose the same identifier?

- The event E consists of all the points for which the first two coordinates (**black** and **red**) are the same.

  - All values possible for coordinates *3* to *n*, all values possible for coordinate 2 (**red**) and then coordinate *1* is fixed (**black**).

Identifier Selection

- What is the probability that processes *1* and *2* choose the same identifier?
- The event E consists of all the points for which the first two coordinates (**black** and **red**) are the same.
    - All values possible for coordinates *3* to *n*, all values possible for coordinate 2 (**red**) and then coordinate *1* is fixed (**black**).

$$\mathbf{Pr}[E] = \sum_{i \in E} p(i) = \frac{1}{2^{kn}} \cdot 2^{k(n-1)} = \frac{1}{2^k}$$

Conditional Probability

Conditional Probability

- Back to the Poker game.

Conditional Probability

- Back to the Poker game.
- What was my probability of winning?

Conditional Probability

- Back to the Poker game.
- What was my probability of winning?
- I could only win if the river was a King.

Conditional Probability

- Back to the Poker game.
- What was my probability of winning?
- I could only win if the river was a King.
- We already had drawn 8 cards, 2 of which were Kings.

Conditional Probability

- Back to the Poker game.
- What was my probability of winning?
- I could only win if the river was a King.
- We already had drawn 8 cards, 2 of which were Kings.
- What was the probability that another King would turn up?

Conditional Probability

- Back to the Poker game.
- What was my probability of winning?
- I could only win if the river was a King.
- We already had drawn 8 cards, 2 of which were Kings.
- What was the probability that another King would turn up?
    - 44 cards left, 2 Kings left, Probability $2/44 = 0.045$.

Conditional Probability

- Given that event F has occurred, what is the probability that even E will occur?

$$\mathbf{Pr}[E \mid F] = \frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]}$$

$$\mathbf{Pr}[E] = \sum_{j=1}^{k} \mathbf{Pr}[E \mid F_j] \cdot \mathbf{Pr}[F_j]$$

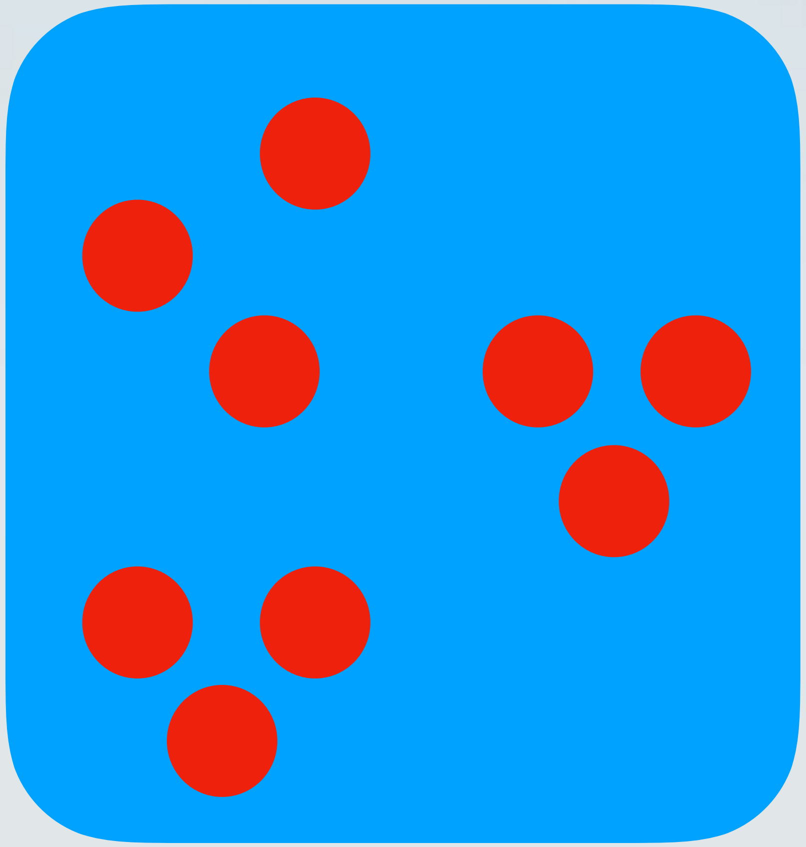

Sample Space **Ω**

Conditional Probability

- Given that event F has occurred, what is the probability that even E will occur?

$$\mathbf{Pr}[E \mid F] = \frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]}$$

$$\mathbf{Pr}[E] = \sum_{j=1}^{k} \mathbf{Pr}[E \mid F_j] \cdot \mathbf{Pr}[F_j]$$

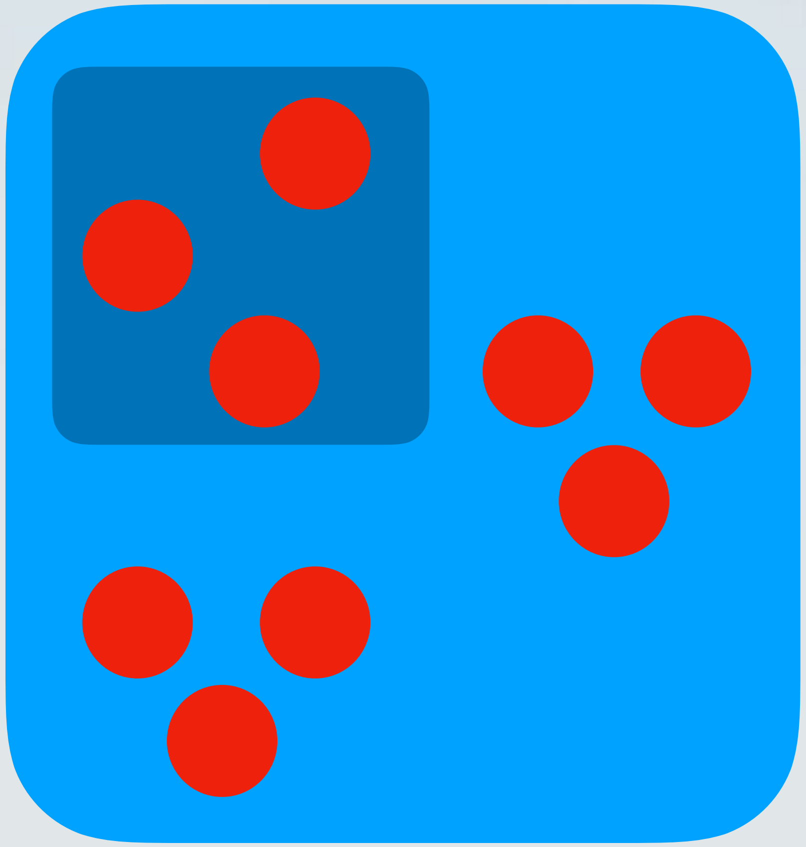

Sample Space **Ω**

Conditional Probability

- Given that event F has occurred, what is the probability that even E will occur?

$$\mathbf{Pr}[E \mid F] = \frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]}$$

$$\mathbf{Pr}[E] = \sum_{j=1}^{k} \mathbf{Pr}[E \mid F_j] \cdot \mathbf{Pr}[F_j]$$

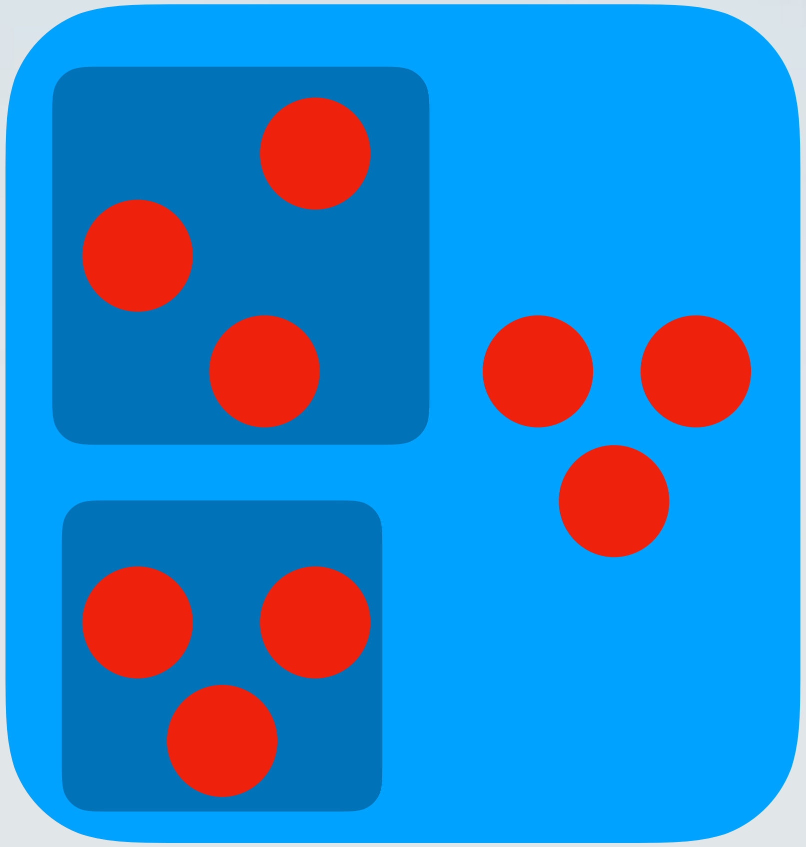

Sample Space Ω

Conditional Probability

- Given that event F has occurred, what is the probability that even E will occur?

$$\mathbf{Pr}[E \mid F] = \frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]}$$

$$\mathbf{Pr}[E] = \sum_{j=1}^{k} \mathbf{Pr}[E \mid F_j] \cdot \mathbf{Pr}[F_j]$$

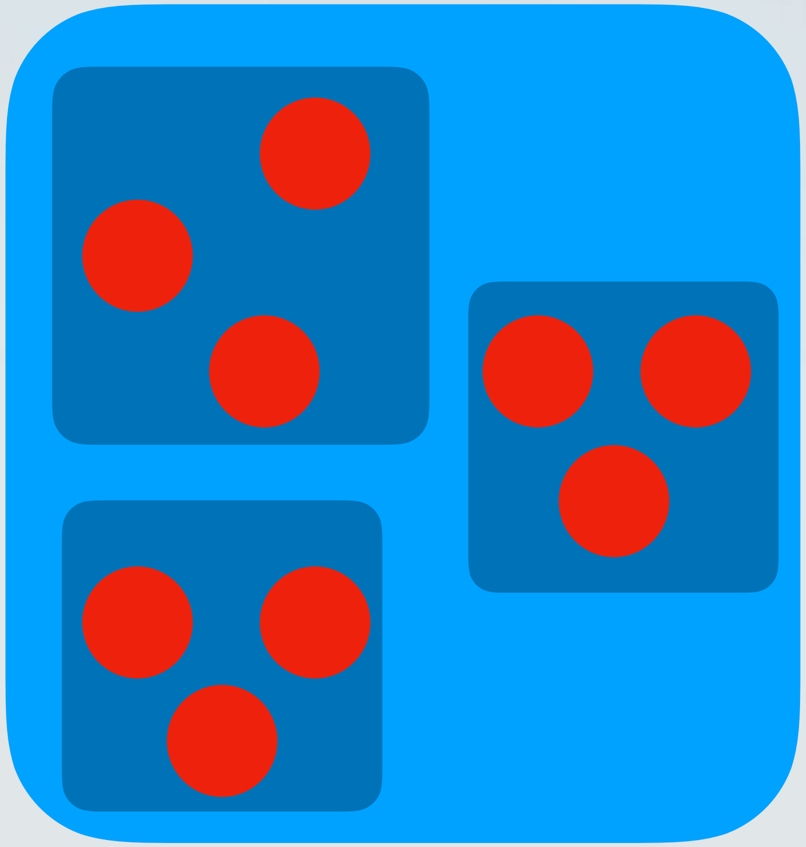

Sample Space $\boldsymbol{\Omega}$

Conditional Probability

- If we toss two fair coins, what is the probability that we get 2 "heads", given that the first toss was "heads"?

- E = 2 heads, F = first toss heads

$$\mathbf{Pr}[F] = 1/2 \qquad \mathbf{Pr}[E \cap F] = 1/4$$

$$\mathbf{Pr}[E \mid F] = \frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]} = \frac{1}{2}$$

Birthday Problem

- We have a room of 25 people.
- Assume that one’s birthday is drawn uniformly at random from all the days of the year.
- What is the probability that there exist two of them that have the same birthday?

Birthday Problem

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
    - 1

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 1
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 1
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 364/365

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 1
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 364/365
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 1
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 364/365
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 363/365?

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 1
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 364/365
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 363/365?
  - What if the first two people actually share a birthday?

Birthday Problem

- Pick a person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 1
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen is …
  - 364/365
- Pick another person. The probability that he/she does not share a birthday with anyone previously chosen, *given that all the previously chosen people do not share a birthday* is …
  - 363/365

Birthday Problem

Birthday Problem

- The probability that there do not exist any two people that share a birthday is equal to:

Birthday Problem

- The probability that there do not exist any two people that share a birthday is equal to:

    - Pr[*1st and 2nd don't share*] x
      Pr[*3rd does not match the previous*] x
      Pr[*4th does not match the previous*] x … x
      Pr[*last does not match the previous*]

Birthday Problem

- The probability that there do not exist any two people that share a birthday is equal to:

  - Pr[*1st and 2nd don't share*] x
    Pr[*3rd does not match the previous*] x
    Pr[*4th does not match the previous*] x … x
    Pr[*last does not match the previous*]

  - *1* x *364/365* x *363/365* x … x *341/365 = 0.431*

Birthday Problem

- The probability that there do not exist any two people that share a birthday is equal to:

  - Pr[*1st and 2nd don't share*] x
    Pr[*3rd does not match the previous*] x
    Pr[*4th does not match the previous*] x … x
    Pr[*last does not match the previous*]

  - *1* x *364/365* x *363/365* x … x *341/365 = 0.431*

- The probability that there exist two people that share a birthday is then equal to: *1-0.431 = 0.569*

Independent Events

- Two events are *independent*, if information about the outcome of one of them does not affect our estimate of the likelihood of the other.

Independent Events

- Two events are *independent*, if information about the outcome of one of them does not affect our estimate of the likelihood of the other.

- Formally: Pr[E | F] = Pr[E] and Pr[F | E] = Pr[E].

Independent Events

- Two events are *independent*, if information about the outcome of one of them does not affect our estimate of the likelihood of the other.

- Formally: Pr[E | F] = Pr[E] and Pr[F | E] = Pr[E].

- This implies: $\frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]} = \mathbf{Pr}[E] \Rightarrow \mathbf{Pr}[E \cap F] = \mathbf{Pr}[E] \cdot \mathbf{Pr}[F]$

Independent Events

- Two events are *independent*, if information about the outcome of one of them does not affect our estimate of the likelihood of the other.

- Formally: Pr[E | F] = Pr[E] and Pr[F | E] = Pr[E].

- This implies: $\frac{\mathbf{Pr}[E \cap F]}{\mathbf{Pr}[F]} = \mathbf{Pr}[E] \Rightarrow \mathbf{Pr}[E \cap F] = \mathbf{Pr}[E] \cdot \mathbf{Pr}[F]$

- In other words, *the probability that two independent events happen is the product of the probabilities that each one of them happens.*

Independent Events

- Generalising:

$$\mathbf{Pr}\left[\bigcap_{i\in I} E_i\right] = \prod_{i\in I} \mathbf{Pr}[E_i]$$

Independent Events

- Generalising:

$$\mathbf{Pr}\left[\bigcap_{i \in I} E_i\right] = \prod_{i \in I} \mathbf{Pr}[E_i]$$

The probability A **and** B **and** C happens is the product of their probabilities.

The Union Bound

independent events

$$\mathbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] = \sum_{i=1}^{n} \mathbf{Pr}[E_i]$$

generally

$$\mathbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] \leq \sum_{i=1}^{n} \mathbf{Pr}[E_i]$$

The Union Bound

independent events

$$\textbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] = \sum_{i=1}^{n} \textbf{Pr}[E_i]$$

The probability that A **or** B **or** C happens is the sum of their probabilities.

generally

$$\textbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] \leq \sum_{i=1}^{n} \textbf{Pr}[E_i]$$

The Union Bound

independent events

$$\textbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] = \sum_{i=1}^{n} \textbf{Pr}[E_i]$$

The probability that A **or** B **or** C happens is the sum of their probabilities.

generally

$$\textbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] \leq \sum_{i=1}^{n} \textbf{Pr}[E_i]$$

The probability that A **or** B **or** C happens is at most the sum of their probabilities.

How to use the Union Bound

- Suppose that we design an algorithm which will produce the correct outcome *with high probability*.

How to use the Union Bound

- Suppose that we design an algorithm which will produce the correct outcome *with high probability*.
- We first formulate a set of "bad events" E_1, E_2, … , E_n, which will force our algorithm to produce an incorrect outcome.

How to use the Union Bound

- Suppose that we design an algorithm which will produce the correct outcome *with high probability*.
- We first formulate a set of "bad events" E_1, E_2, … , E_n, which will force our algorithm to produce an incorrect outcome.
- If none of these "bad events" happens, our algorithm will produce the correct outcome.

How to use the Union Bound

- Suppose that we design an algorithm which will produce the correct outcome *with high probability*.
- We first formulate a set of "bad events" E_1, E_2, … , E_n, which will force our algorithm to produce an incorrect outcome.
- If none of these "bad events" happens, our algorithm will produce the correct outcome.
- Suppose that F is the event that the algorithm fails.

How to use the Union Bound

- Suppose that we design an algorithm which will produce the correct outcome *with high probability*.
- We first formulate a set of "bad events" E_1, E_2, … , E_n, which will force our algorithm to produce an incorrect outcome.
- If none of these "bad events" happens, our algorithm will produce the correct outcome.
- Suppose that F is the event that the algorithm fails.
  - We have:

$$\mathbf{Pr}[F] \leq \mathbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] \leq \sum_{i=1}^{n} \mathbf{Pr}[E_i]$$

How to use the Union Bound

- We have: $$\mathbf{Pr}[F] \leq \mathbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] \leq \sum_{i=1}^{n} \mathbf{Pr}[E_i]$$

How to use the Union Bound

- We have: $$\mathbf{Pr}[F] \leq \mathbf{Pr}\left[\bigcup_{i=1}^{n} E_i\right] \leq \sum_{i=1}^{n} \mathbf{Pr}[E_i]$$

- If we can prove that the sum of probabilities of these events is small, then we can prove that our algorithm succeeds with high probability.

Identifier Selection

- There are *n* processes in a distributed system.
- The set of possible identifiers is the set of all *k*-bit strings.
  - e.g., *1001001…01*
- Each process chooses an identifier *uniformly at random.*
  - i.e., all strings have equal probability of being chosen.
- What is the probability that processes *1* and *2* choose the same identifier?

Identifier Selection

- There are *1000* processes in a distributed system.
- The set of possible identifiers is the set of all *32*-bit strings.
  - e.g., *1001001…01*
- Each process chooses an identifier *uniformly at random.*
  - i.e., all strings have equal probability of being chosen.
- What is the probability that any two of them choose the same identifier?

Identifier Selection

- What is the probability that any two of them choose the same identifier?

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.
- Simple “failure” events: E_{ij}

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.
- Simple “failure” events: E_{ij}
  - Processes *i* and *j* choose the same identifier.

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.
- Simple "failure" events: E_{ij}
  - Processes *i* and *j* choose the same identifier.
- How many such events?

$$\binom{1000}{2}$$

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.
- Simple "failure" events: E_{ij}
  - Processes *i* and *j* choose the same identifier.
- How many such events?

$$\binom{1000}{2}$$

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.
- Simple "failure" events: E_{ij}
  - Processes *i* and *j* choose the same identifier.
- How many such events?

$$\binom{1000}{2}$$

- What is the probability of each happening?

Identifier Selection

- What is the probability that any two of them choose the same identifier?
  - Call this event F.
- Simple "failure" events: E_{ij}
  - Processes *i* and *j* choose the same identifier.
- How many such events?

$$\binom{1000}{2}$$

- What is the probability of each happening?
  - $1/2^{32}$

Identifier Selection

Identifier Selection

- What is the probability of failure?

Identifier Selection

- What is the probability of failure?

$$\mathbf{Pr}[F] \leq \sum_{i,j} \mathbf{Pr}[E_{ij}] = \binom{1000}{2} \cdot \frac{1}{2^{32}} \leq 0.000125$$

Identifier Selection

- What is the probability of failure?

$$\mathbf{Pr}[F] \le \sum_{i,j} \mathbf{Pr}[E_{ij}] = \binom{1000}{2} \cdot \frac{1}{2^{32}} \le 0.000125$$

- What is the probability of success?

at least 1-0.000125 = 0.999875

Random Variables and Expectations

- Random Variable: (Informally) A variable X whose values depend on outcomes of a random phenomenon.
- Pr[X = *j*] : *the probability that the value of* X *is j*.
- Expectation ("average value") of X:

$$\mathbb{E}[x] = \sum_{j=1}^{n} \mathbf{Pr}[X = j]$$

Expectation

- Simple example:
  - Assume that X takes a value in $\{1, 2, \ldots, n\}$ with probability $1/n$.
  - $E[X] = 1(1/n) + 2(1/n) + \ldots + n(1/n) = (1+2+\ldots+n)/n = (n+1)/2$

Waiting for the first success

- We flip a *biased* coin, where
  Pr[H] = p and
  Pr[T] = 1-p

- We flip repeatedly until we get one "*heads*" result.

- What is the expected number of times that we need to flip for that to happen?

Waiting for the first success

- We flip a *biased* coin, where
  Pr[H] = p and
  Pr[T] = 1-p

- We flip repeatedly until we get one “*heads*” result.

- What is the expected number of times that we need to flip for that to happen?

- Let X be the random variable of the number of flips.

Waiting for the first success

- We flip a *biased* coin, where
  Pr[H] = p and
  Pr[T] = 1-p

- We flip repeatedly until we get one "*heads*" result.

- What is the expected number of times that we need to flip for that to happen?

- Let X be the random variable of the number of flips.

  - We are looking for E[X].

Waiting for the first success

Waiting for the first success

- Suppose that we get “*heads*” on the j-th flip.

Waiting for the first success

- Suppose that we get "*heads*" on the j-th flip.
- We have: $\mathbf{Pr}[X = j] = (1 - p)^{j-1} \cdot p$

Waiting for the first success

- Suppose that we get "*heads*" on the *j*-th flip.
- We have: $\mathbf{Pr}[X = j] = (1-p)^{j-1} \cdot p$
- The expectation then becomes:

$$\mathbb{E}[X] = \sum_{j=0}^{\infty} j \cdot \mathbf{Pr}[X = j] = \sum_{j=1}^{\infty} j(1-p)^{j-1}p = \frac{p}{1-p} \sum_{j=1}^{\infty} j(1-p)^{j}$$

$$= \frac{p}{1-p} \cdot \frac{(1-p)}{p^2} = \frac{1}{p}$$

Application

- Suppose that we repeat an experiment multiple times, and each time the probability of success is $p > 0$.
    - e.g., compute a minimum cut in a graph.
- The expected number of repetitions that we need until the experiment succeeds is $1/p$.

Linearity of Expectation

- Let X and Y be random variables defined over the same space.
- Let X+Y be the random variable equal to X(ω) + Y(ω) on a point ω of the sample space.
- It holds that E[X+Y] = E[X] + E[Y]
- Generally:

$$\mathbb{E}[X_1 + X_2 + \ldots + X_n] = \sum_{i=1}^{n} X_i$$

Guessing a card

- A deck with *n* cards.
- We draw a card, and before we see it, we guess what it is.
  - We pick one of the cards *uniformly at random from the whole deck*.
- How many of our predictions do we expect to be correct?

Guessing a card

- How many of our predictions do we expect to be correct?

Guessing a card

- How many of our predictions do we expect to be correct?
    - Let X be the random variable denoting the number of correct predictions.

Guessing a card

- How many of our predictions do we expect to be correct?
  - Let X be the random variable denoting the number of correct predictions.
  - To help us compute X, we will defined variables X_i:

Guessing a card

- How many of our predictions do we expect to be correct?
  - Let X be the random variable denoting the number of correct predictions.
  - To help us compute X, we will defined variables X_i:
    - X_i is *1* if the ith prediction is correct and *0* otherwise.

Guessing a card

- How many of our predictions do we expect to be correct?
  - Let X be the random variable denoting the number of correct predictions.
  - To help us compute X, we will defined variables X_i:
    - X_i is *1* if the ith prediction is correct and *0* otherwise.
    - Obviously: $X = X_1 + X_2 + \ldots + X_n$

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$
  - We have $\mathbb{E}[X_i] = 0 \cdot \mathbf{Pr}[X_i = 0] + 1 \cdot \mathbf{Pr}[X_i = 1] = \frac{1}{n}$

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$
  - We have $\mathbb{E}[X_i] = 0 \cdot \mathbf{Pr}[X_i = 0] + 1 \cdot \mathbf{Pr}[X_i = 1] = \frac{1}{n}$
  - By linearity of expectation: $\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = 1$

Guessing a card

- A deck with *n* cards.
- We draw a card, and before we see it, we guess what it is.
  - We pick one of the cards *uniformly at random from the whole deck*.
- How many of our predictions do we expect to be correct?

Guessing a card

- A deck with *n* cards.
- We draw a card, and before we see it, we guess what it is.
    - We pick one of the cards *uniformly at random from the cards you have not seen*.
- How many of our predictions do we expect to be correct?

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$
  - We have $\mathbb{E}[X_i] = 0 \cdot \mathbf{Pr}[X_i = 0] + 1 \cdot \mathbf{Pr}[X_i = 1] = \frac{1}{n}$
  - By linearity of expectation: $\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = 1$

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$
  - We have $\mathbb{E}[X_i] = 0 \cdot \mathbf{Pr}[X_i = 0] + 1 \cdot \mathbf{Pr}[X_i = 1] = \frac{1}{n-i+1}$
  - By linearity of expectation:

$$\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = \sum_{i=1}^{n} \frac{1}{n-i+1} = \sum_{i=1}^{n} \frac{1}{i} = \mathbf{H}(n)$$

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$
  - We have $\mathbb{E}[X_i] = 0 \cdot \mathbf{Pr}[X_i = 0] + 1 \cdot \mathbf{Pr}[X_i = 1] = \frac{1}{n - i + 1}$
  - By linearity of expectation:

The *n*-th harmonic number

$$\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = \sum_{i=1}^{n} \frac{1}{n - i + 1} = \sum_{i=1}^{n} \frac{1}{i} = \mathbf{H}(n)$$

Guessing a card

- To help us compute X, we will defined variables X_i:
  - X_i is *1* if the ith prediction is correct and *0* otherwise.
  - Obviously: $X = X_1 + X_2 + \ldots + X_n$
  - We have $\mathbb{E}[X_i] = 0 \cdot \mathbf{Pr}[X_i = 0] + 1 \cdot \mathbf{Pr}[X_i = 1] = \frac{1}{n-i+1}$
  - By linearity of expectation:

The *n*-th harmonic number

$$\mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = \sum_{i=1}^{n} \frac{1}{n-i+1} = \sum_{i=1}^{n} \frac{1}{i} = \mathbf{H}(n)$$

$H(n) = \Theta(\log n)$