Today

Probability:
Keep building it formally..
And our intuition.

Poll: blows my mind.

Flip 300 million coins.

Which is more likely?
(A) 300 million heads.
(B) 300 million tails.
(C) Alternating heads and tails.
(D) A tail every third spot.

Given the history of the universe up to right now.

What is the likelihood of our universe?
(A) The likelihood is 1. Cuz here it is.
(B) As likely as any other. Cuz of probability.
(C) Well. Quantum. IDK- TBH.

Perhaps a philosophical (“wastebasket”) question.

Also, “cuz” $==$ “because”

Probability Basics.

Probability Space.

1. **Sample Space:** Set of outcomes, Ω.
2. **Probability:** $Pr[\omega]$ for all $\omega \in \Omega$.
   2.1 $0 \le Pr[\omega] \le 1$.
   2.2 $\sum_{\omega \in \Omega} Pr[\omega] = 1$.

Example: Two coins.

1. $\Omega = \{HH, HT, TH, TT\}$
   (Note: Not $\Omega = \{H, T\}$ with two picks!)
2. $Pr[HH] = \cdots = Pr[TT] = 1/4$

Consequences of Additivity

Theorem

(a) Inclusion/Exclusion: $Pr[A \cup B] = Pr[A] + Pr[B] - Pr[A \cap B]$;

(b) Union Bound: $Pr[A_1 \cup \cdots \cup A_n] \le Pr[A_1] + \cdots + Pr[A_n]$;

(c) Law of Total Probability:

If $A_1, \dots A_N$ are a partition of Ω, i.e.,

pairwise disjoint and $\cup_{m=1}^{N} A_m = \Omega$, then

$$Pr[B] = Pr[B \cap A_1] + \cdots + Pr[B \cap A_N].$$

Proof Idea: Total probability.

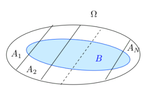

Add it up!

Add it up. Poll.

What does Rao mean by “Add it up.”

(A) Organize intuitions/proofs around $Pr[\omega]$.
(B) Organize intuition/proofs around $Pr[A]$.
(C) Some weird song whose refrain he heard in his youth.

(A), (B), and (C)

Conditional Probability.

Definition: The **conditional probability** of B given A is

$$Pr[B|A] = \frac{Pr[A \cap B]}{Pr[A]}$$

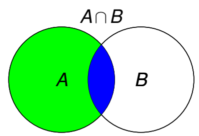

In A!
In B?
Must be in $A \cap B$.

$$Pr[B|A] = \frac{Pr[A \cap B]}{Pr[A]}.$$

Note also:

$Pr[A \cap B] = Pr[B|A]Pr[A]$

Product Rule

Def: $Pr[B|A] = \frac{Pr[A \cap B]}{Pr[A]}$.

Also: $Pr[A \cap B] = Pr[B|A]Pr[A]$

Theorem Product Rule
Let $A_1, A_2, \ldots, A_n$ be events. Then

$$Pr[A_1 \cap \cdots \cap A_n] = Pr[A_1]Pr[A_2|A_1] \cdots Pr[A_n|A_1 \cap \cdots \cap A_{n-1}].$$

Simple Bayes Rule.

$Pr[A|B] = \frac{Pr[A \cap B]}{Pr[B]}$, $Pr[B|A] = \frac{Pr[A \cap B]}{Pr[A]}$.

$Pr[A \cap B] = Pr[A|B]Pr[B] = Pr[B|A]Pr[A]$.

Bayes Rule: $Pr[A|B] = \frac{Pr[B|A]Pr[A]}{Pr[B]}$.

Is you coin loaded?

Your coin is fair w.p. $1/2$ or such that $Pr[H] = 0.6$, otherwise.

You flip your coin and it yields heads.

What is the probability that it is fair?

Analysis:

$$A = \text{ 'coin is fair'}, B = \text{ 'outcome is heads'}$$

We want to calculate $P[A|B]$.

We know $P[B|A] = 1/2, P[B|\bar{A}] = 0.6, Pr[A] = 1/2 = Pr[\bar{A}]$

Now,

$$\begin{aligned} Pr[B] &= Pr[A \cap B] + Pr[\bar{A} \cap B] = Pr[A]Pr[B|A] + Pr[\bar{A}]Pr[B|\bar{A}] \\ &= (1/2)(1/2) + (1/2)0.6 = 0.55. \end{aligned}$$

Thus,

$$Pr[A|B] = \frac{Pr[A]Pr[B|A]}{Pr[B]} = \frac{(1/2)(1/2)}{(1/2)(1/2) + (1/2)0.6} \approx 0.45.$$

Bayes and Biased Coin

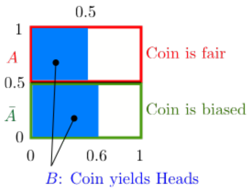

Pick a point uniformly at random in the unit square. Then

$Pr[A] = 0.5; Pr[\bar{A}] = 0.5$
$Pr[B|A] = 0.5; Pr[B|\bar{A}] = 0.6; Pr[A \cap B] = 0.5 \times 0.5$
$Pr[B] = 0.5 \times 0.5 + 0.5 \times 0.6 = Pr[A]Pr[B|A] + Pr[\bar{A}]Pr[B|\bar{A}]$

$$Pr[A|B] = \frac{0.5 \times 0.5}{0.5 \times 0.5 + 0.5 \times 0.6} = \frac{Pr[A]Pr[B|A]}{Pr[A]Pr[B|A] + Pr[\bar{A}]Pr[B|\bar{A}]}$$
$$\approx 0.46 = \text{fraction of } B \text{ that is inside } A$$

Bayes: General Case

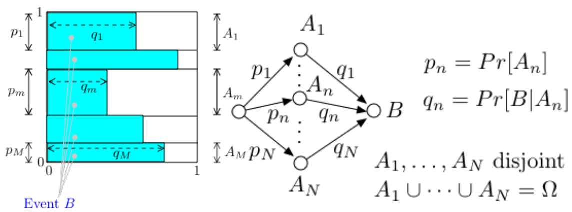

Pick a point uniformly at random in the unit square. Then

$Pr[A_n] = p_n, n = 1, \ldots, N$
$Pr[B|A_n] = q_n, n = 1, \ldots, N; Pr[A_n \cap B] = p_n q_n$
$Pr[B] = p_1 q_1 + \cdots p_N q_N$

$$Pr[A_n|B] = \frac{p_n q_n}{p_1 q_1 + \cdots p_N q_N} = \text{ fraction of } B \text{ inside } A_n.$$

Bayes Rule

A general picture: We imagine that there are N possible causes $A_1, \ldots, A_N$.

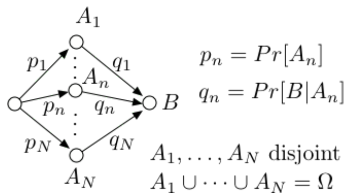

100 situations: $100 p_n q_n$ where A_n and B occur, for $n = 1, \ldots, N$.
In $100 \sum_m p_m q_m$ occurences of B, $100 p_n q_n$ occurrences of A_n.

Hence,

$Pr[A_n|B] = \frac{p_n q_n}{\sum_m p_m q_m}$.

But, $p_n = Pr[A_n], q_n = Pr[B|A_n], \sum_m p_m q - m = Pr[B]$, hence,

$Pr[A_n|B] = \frac{Pr[B|A_n]Pr[A_n]}{Pr[B]}$.

Why do you have a fever?

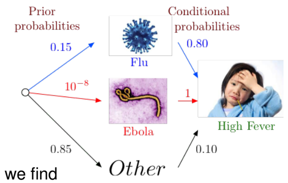

Using Bayes' rule, we find

$$Pr[\text{Flu}|\text{High Fever}] = \frac{0.15 \times 0.80}{0.15 \times 0.80 + 10^{-8} \times 1 + 0.85 \times 0.1} \approx 0.58$$

$$Pr[\text{Ebola}|\text{High Fever}] = \frac{10^{-8} \times 1}{0.15 \times 0.80 + 10^{-8} \times 1 + 0.85 \times 0.1} \approx 5 \times 10^{-8}$$

$$Pr[\text{Other}|\text{High Fever}] = \frac{0.85 \times 0.1}{0.15 \times 0.80 + 10^{-8} \times 1 + 0.85 \times 0.1} \approx 0.42$$

The values $0.58, 5 \times 10^{-8}, 0.42$ are the posterior probabilities.

Why do you have a fever?

Our "Bayes' Square" picture:

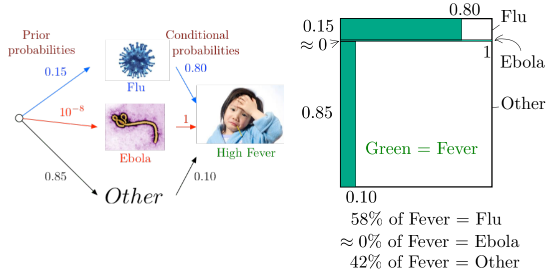

Note that even though $Pr[\text{Fever}|\text{Ebola}] = 1$, one has

$$Pr[\text{Ebola}|\text{Fever}] \approx 0.$$

This example shows the importance of the prior probabilities.

Why do you have a fever?

We found

$$Pr[\text{Flu}|\text{High Fever}] \approx 0.58,$$
$$Pr[\text{Ebola}|\text{High Fever}] \approx 5 \times 10^{-8},$$
$$Pr[\text{Other}|\text{High Fever}] \approx 0.42$$

'Flu' is Most Likely a Posteriori (MAP) cause of high fever.
'Ebola' is Maximum Likelihood Estimate (MLE) of cause:
causes fever with largest probability.
Recall that

$$p_m = Pr[A_m], q_m = Pr[B|A_m], Pr[A_m|B] = \frac{p_m q_m}{p_1 q_1 + \cdots + p_M q_M}.$$

Thus,

- MAP = value of m that maximizes $p_m q_m$.
- MLE = value of m that maximizes q_m.

Bayes' Rule Operations

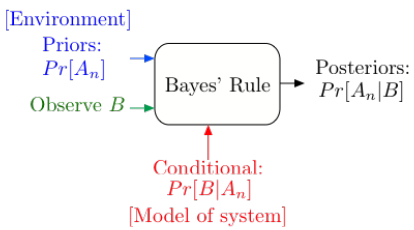

Bayes' Rule: canonical example of how information changes our opinions.

Thomas Bayes

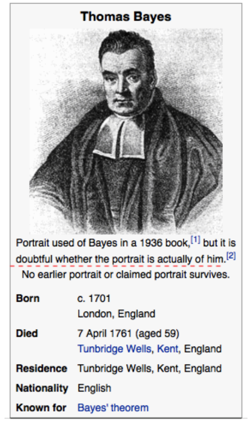
Thomas Bayes

Portrait used of Bayes in a 1936 book,[1] but it is doubtful whether the portrait is actually of him.[2] No earlier portrait or claimed portrait survives.

Born	c. 1701 London, England
Died	7 April 1761 (aged 59) Tunbridge Wells, Kent, England
Residence	Tunbridge Wells, Kent, England
Nationality	English
Known for	Bayes' theorem

Source: Wikipedia.

Thomas Bayes

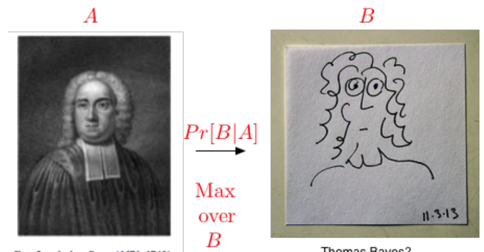

FIG. 3. *Joshua Bayes (1671–1746).*

Thomas Bayes?

A Bayesian picture of Thomas Bayes.

Testing for disease.

Random Experiment: Pick a random male.
Outcomes: $(test, disease)$
A - prostate cancer.
B - positive PSA test.

- $Pr[A] = 0.0016$, (.16 % of the male population is affected.)
- $Pr[B|A] = 0.80$ (80% chance of positive test with disease.)
- $Pr[B|\bar{A}] = 0.10$ (10% chance of positive test without disease.)

From http://www.cpcn.org/01_psa_tests.htm and http://seer.cancer.gov/statfacts/html/prost.html (10/12/2011.)

Positive PSA test (B). Do I have disease?

$$Pr[A|B]???$$

Bayes Rule.

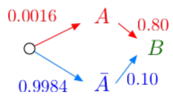

Using Bayes' rule, we find

$$P[A|B] = \frac{0.0016 \times 0.80}{0.0016 \times 0.80 + 0.9984 \times 0.10} = .013.$$

A 1.3% chance of prostate cancer with a positive PSA test.

Surgery anyone?

Impotence...

Incontinence..

Death.

Conditional Probability: Pictures/Poll.

Illustrations: Pick a point uniformly in the unit square

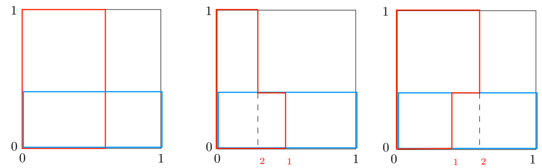

Which A and B are independent?

(A) Left.
(B) Middle.
(B) Right.

See next slide.

Conditional Probability: Pictures

Illustrations: Pick a point uniformly in the unit square

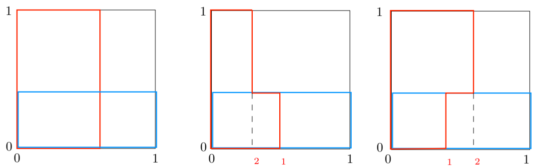

- Left: A and B are independent. $Pr[B] = b; Pr[B|A] = b$.
- Middle: A and B are positively correlated. $Pr[B|A] = b_1 > Pr[B|\bar{A}] = b_2$. Note: $Pr[B] \in (b_2, b_1)$.
- Right: A and B are negatively correlated. $Pr[B|A] = b_1 < Pr[B|\bar{A}] = b_2$. Note: $Pr[B] \in (b_1, b_2)$.

Quick Review

Events, Conditional Probability, Independence, Bayes' Rule

Key Ideas:

- Conditional Probability:

  $Pr[A|B] = \frac{Pr[A \cap B]}{Pr[B]}$
- Independence: $Pr[A \cap B] = Pr[A]Pr[B]$.
- Bayes' Rule:

  $$Pr[A_n|B] = \frac{Pr[A_n]Pr[B|A_n]}{\sum_m Pr[A_m]Pr[B|A_m]}.$$

  $Pr[A_n|B] =$ posterior probability; $Pr[A_n] =$ prior probability .
- All these are possible:

  $Pr[A|B] < Pr[A]; Pr[A|B] > Pr[A]; Pr[A|B] = Pr[A]$.

Independence

Recall :

A and B are independent
$\Leftrightarrow Pr[A \cap B] = Pr[A]Pr[B]$
$\Leftrightarrow Pr[A|B] = Pr[A]$.

In general: $Pr[A \cap B] = Pr[A|B]Pr[B]$.

If $Pr[A|B] = Pr[A]$, does $Pr[B|A] = Pr[B]$?

Yes. Independent: $Pr[A \cap B] = Pr[A]Pr[B] = Pr[A]Pr[B|A]$. Therefore $Pr[B|A] = Pr[B]$.
Consider the example below:

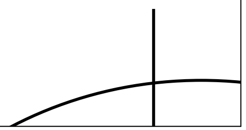

Mutually exclusive.

Events A and B are mutually exclusive if $A \cap B$ is empty.

Are A and B independent?

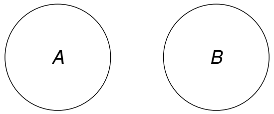

$P[A] = 1/3, Pr[B] = 1/3.$

$P[A|B]$? 0

Independent? $Pr[A] \neq Pr[A|B]$.

Pairwise Independence

Flip two fair coins. Let

- $A =$ 'first coin is H' $= \{HT, HH\}$;
- $B =$ 'second coin is H' $= \{TH, HH\}$;
- $C =$ 'the two coins are different' $= \{TH, HT\}$.

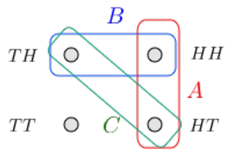

A, C are independent; B, C are independent;

$A \cap B, C$ are not independent. ($Pr[A \cap B \cap C] = 0 \neq Pr[A \cap B]Pr[C]$.)

False: If A did not say anything about C and B did not say anything about C, then $A \cap B$ would not say anything about C.

Example

Flip a fair coin 5 times. Let A_n = 'coin n is H', for $n = 1, \ldots, 5$.

Then,

$$A_m, A_n \text{ are independent for all } m \neq n.$$

Also,

$$A_1 \text{ and } A_3 \cap A_5 \text{ are independent.}$$

Indeed,

$$Pr[A_1 \cap (A_3 \cap A_5)] = \frac{1}{8} = Pr[A_1]Pr[A_3 \cap A_5].$$

Similarly,

$$A_1 \cap A_2 \text{ and } A_3 \cap A_4 \cap A_5 \text{ are independent.}$$

This leads to a definition ....

Mutual Independence

Definition Mutual Independence

(a) The events $A_1, \ldots, A_5$ are mutually independent if

$$Pr[\cap_{k \in K} A_k] = \prod_{k \in K} Pr[A_k], \text{ for all } K \subseteq \{1, \ldots, 5\}.$$

(b) More generally, the events $\{A_j, j \in J\}$ are mutually independent if

$$Pr[\cap_{k \in K} A_k] = \prod_{k \in K} Pr[A_k], \text{ for all finite} K \subseteq J.$$

Example: Flip a fair coin forever. Let A_n = 'coin n is H.' Then the events A_n are mutually independent.

Mutual Independence

Theorem

(a) If the events $\{A_j, j \in J\}$ are mutually independent and if K_1 and K_2 are disjoint finite subsets of J, then

$$\cap_{k \in K_1} A_k \text{ and } \cap_{k \in K_2} A_k \text{ are independent.}$$

(b) More generally, if the K_n are pairwise disjoint finite subsets of J, then the events

$$\cap_{k \in K_n} A_k \text{ are mutually independent.}$$

(c) Also, the same is true if we replace some of the A_k by $\bar{A}_k$.

Balls in bins

One throws m balls into $n > m$ bins.

Balls in bins

One throws m balls into $n > m$ bins.

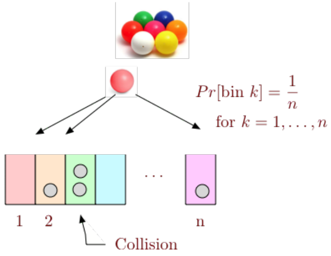

Theorem:
$Pr[\text{no collision}] \approx \exp\{-\frac{m^2}{2n}\}$, for large enough n.

Balls in bins

Theorem:
$Pr[\text{no collision}] \approx \exp\{-\frac{m^2}{2n}\}$, for large enough n.

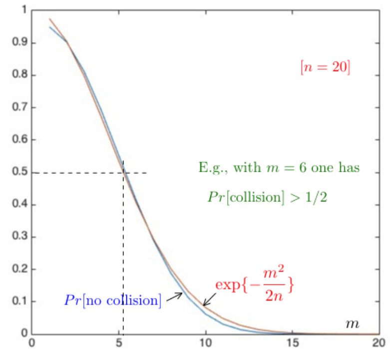

Balls in bins

Theorem:
$Pr[\text{no collision}] \approx \exp\{-\frac{m^2}{2n}\}$, for large enough n.

In particular, $Pr[\text{no collision}] \approx 1/2$ for $m^2/(2n) \approx \ln(2)$, i.e.,

$$m \approx \sqrt{2\ln(2)n} \approx 1.2\sqrt{n}.$$

E.g., $1.2\sqrt{20} \approx 5.4$.

Roughly, $Pr[\text{collision}] \approx 1/2$ for $m = \sqrt{n}$. ($e^{-0.5} \approx 0.6$.)

The Calculation.

A_i = no collision when ith ball is placed in a bin.

$Pr[A_i | A_{i-1} \cap \cdots \cap A_1] = (1 - \frac{i-1}{n})$.

no collision = $A_1 \cap \cdots \cap A_m$.

Product rule:
$Pr[A_1 \cap \cdots \cap A_m] = Pr[A_1]Pr[A_2|A_1] \cdots Pr[A_m|A_1 \cap \cdots \cap A_{m-1}]$

$$\Rightarrow Pr[\text{no collision}] = \left(1 - \frac{1}{n}\right) \cdots \left(1 - \frac{m-1}{n}\right).$$

Hence,

$$\ln(Pr[\text{no collision}]) = \sum_{k=1}^{m-1} \ln(1 - \frac{k}{n}) \approx \sum_{k=1}^{m-1} (-\frac{k}{n}) \ (*)$$
$$= -\frac{1}{n}\frac{m(m-1)}{2} \ (\dagger) \approx -\frac{m^2}{2n}$$

(*) We used $\ln(1-\varepsilon) \approx -\varepsilon$ for $|\varepsilon| \ll 1$.

(†) $1 + 2 + \cdots + m - 1 = (m-1)m/2$.

Approximation

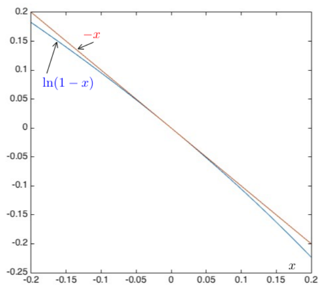

$$\exp\{-x\} = 1 - x + \frac{1}{2!}x^2 + \cdots \approx 1 - x, \text{ for } |x| \ll 1.$$

Hence, $-x \approx \ln(1-x)$ for $|x| \ll 1$.

Today's your birthday, it's my birthday too..

Probability that m people all have different birthdays?
With $n = 365$, one finds

$Pr[\text{collision}] \approx 1/2$ if $m \approx 1.2\sqrt{365} \approx 23$.

If $m = 60$, we find that

$$Pr[\text{no collision}] \approx \exp\{-\frac{m^2}{2n}\} = \exp\{-\frac{60^2}{2 \times 365}\} \approx 0.007.$$

If $m = 366$, then $Pr[\text{no collision}] = 0$. (No approximation here!)

Checksums!

Consider a set of m files.
Each file has a checksum of b bits.
How large should b be for $Pr[\text{share a checksum}] \leq 10^{-3}$?

Claim: $b \geq 2.9\ln(m) + 9$.

Proof:

Let $n = 2^b$ be the number of checksums.
We know $Pr[\text{no collision}] \approx \exp\{-m^2/(2n)\} \approx 1 - m^2/(2n)$. Hence,

$$Pr[\text{no collision}] \approx 1 - 10^{-3} \Leftrightarrow m^2/(2n) \approx 10^{-3}$$
$$\Leftrightarrow 2n \approx m^2 10^3 \Leftrightarrow 2^{b+1} \approx m^2 2^{10}$$
$$\Leftrightarrow b + 1 \approx 10 + 2\log_2(m) \approx 10 + 2.9\ln(m).$$

Note: $\log_2(x) = \log_2(e)\ln(x) \approx 1.44\ln(x)$.

Coupon Collector Problem.

There are n different baseball cards.
(Brian Wilson, Jackie Robinson, Roger Hornsby, ...)

One random baseball card in each cereal box.

Theorem: If you buy m boxes,

(a) $Pr[\text{miss one specific item}] \approx e^{-\frac{m}{n}}$

(b) $Pr[\text{miss any one of the items}] \leq ne^{-\frac{m}{n}}$.

Coupon Collector Problem: Analysis.

Event A_m = 'fail to get Brian Wilson in m cereal boxes'

Fail the first time: $(1 - \frac{1}{n})$
Fail the second time: $(1 - \frac{1}{n})$
And so on ... for m times. Hence,

$$\begin{aligned} Pr[A_m] &= (1 - \frac{1}{n}) \times \cdots \times (1 - \frac{1}{n}) \\ &= (1 - \frac{1}{n})^m \\ ln(Pr[A_m]) &= m\ln(1 - \frac{1}{n}) \approx m \times (-\frac{1}{n}) \\ Pr[A_m] &\approx \exp\{-\frac{m}{n}\}. \end{aligned}$$

For $p_m = \frac{1}{2}$, we need around $n\ln 2 \approx 0.69n$ boxes.

Collect all cards?

Experiment: Choose m cards at random with replacement.

Events: E_k = 'fail to get player k', for k = 1, ..., n

Probability of failing to get at least one of these n players:

$$p := Pr[E_1 \cup E_2 \cdots \cup E_n]$$

How does one estimate p? Union Bound:

$$p = Pr[E_1 \cup E_2 \cdots \cup E_n] \leq Pr[E_1] + Pr[E_2] \cdots Pr[E_n].$$

$$Pr[E_k] \approx e^{-\frac{m}{n}}, k = 1, \ldots, n.$$

Plug in and get

$$p \leq ne^{-\frac{m}{n}}.$$

Collect all cards?

Thus,

$$Pr[\text{missing at least one card}] \leq ne^{-\frac{m}{n}}.$$

Hence,

$$Pr[\text{missing at least one card}] \leq p \text{ when } m \geq n\ln(\frac{n}{p}).$$

To get $p = 1/2$, set $m = n\ln(2n)$.

$(p \leq ne^{-\frac{m}{n}} \leq ne^{-\ln(n/p)} \leq n(\frac{p}{n}) \leq p.)$

E.g., $n = 10^2 \Rightarrow m = 530; n = 10^3 \Rightarrow m = 7600$.

Quick Review.

Bayes' Rule, Mutual Independence, Collisions and Collecting

Main results:

- Bayes' Rule: $Pr[A_m|B] = p_m q_m/(p_1 q_1 + \cdots + p_M q_M)$.
- Product Rule: $Pr[A_1 \cap \cdots \cap A_n] = Pr[A_1]Pr[A_2|A_1] \cdots Pr[A_n|A_1 \cap \cdots \cap A_{n-1}]$.