Differential Privacy

Introduction

Fedor Fomin

Do you use torrents to download unsanctioned copyrighted material?

Don’t worry, when the results are published, we will remove any information that will allow your answers to be traced back to you. It will be totally anonymous…

Hiding some information cannot assure the protection of individual identity.

1990s: the Commonwealth of Massachusetts Group Insurance Commission (GIC) open access to the anonymous health record of its clients for research to benefit the society

Latanya Sweeney: Compared and co-relating GIC database with voter database, successfully identified a number of health records

Why hiding information does not assure the protection of individual identity?

There are not so many people in Norway who
have 12 children
were born in 1922
are 2.1 m tall, etc

Randomness could help

Flip a coin (and don’t tell us how it landed).

If heads -> tell honestly
If tails -> answer randomly

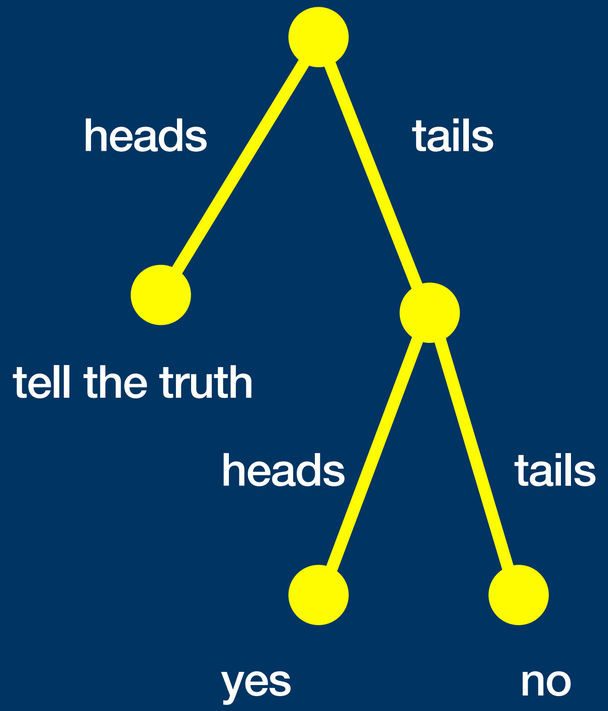

Randomness could help

If you say YES, it could be that you use torrents or because you had tails, then heads

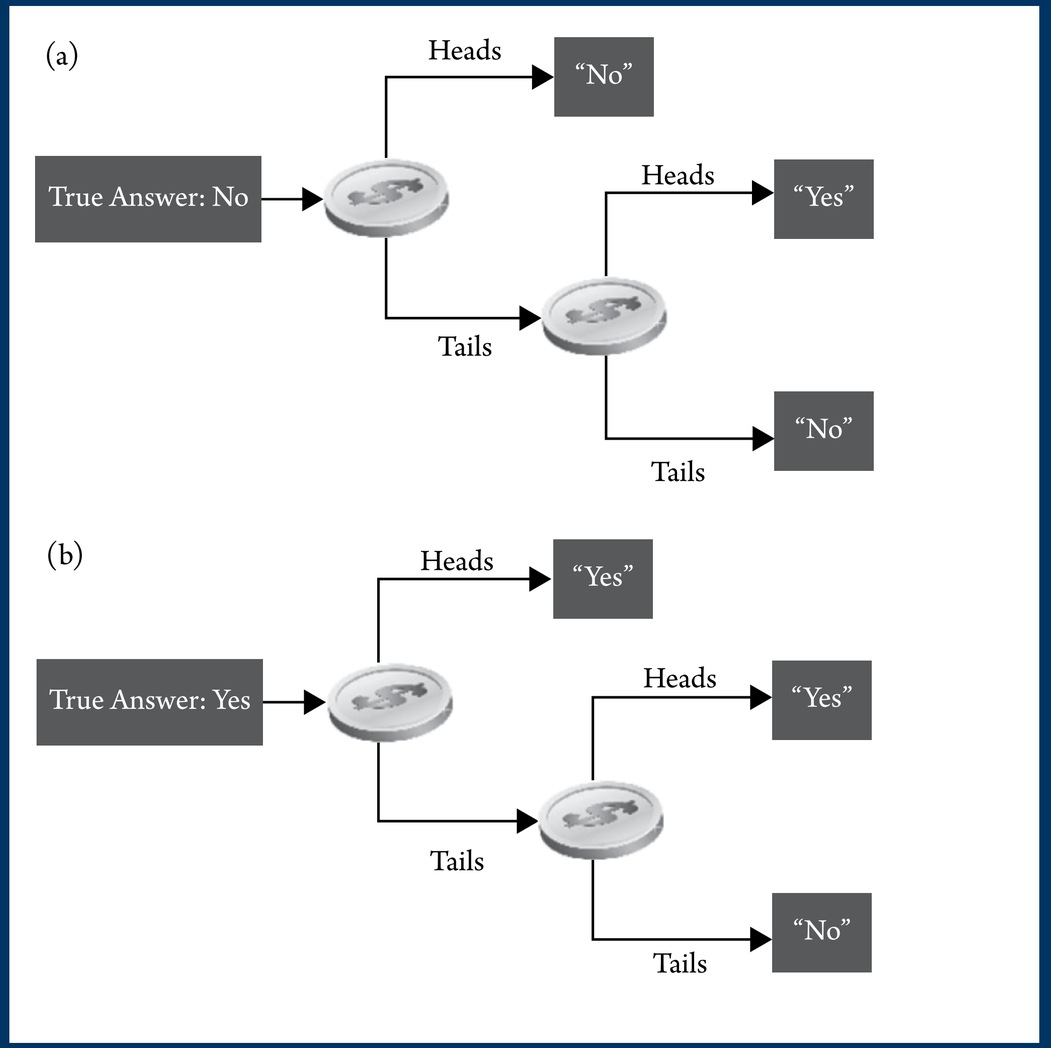

Randomness could help

If you say YES, it could be that you use torrents or because you had tails, then heads

3/4 times we say truth, 1/4 times we lie

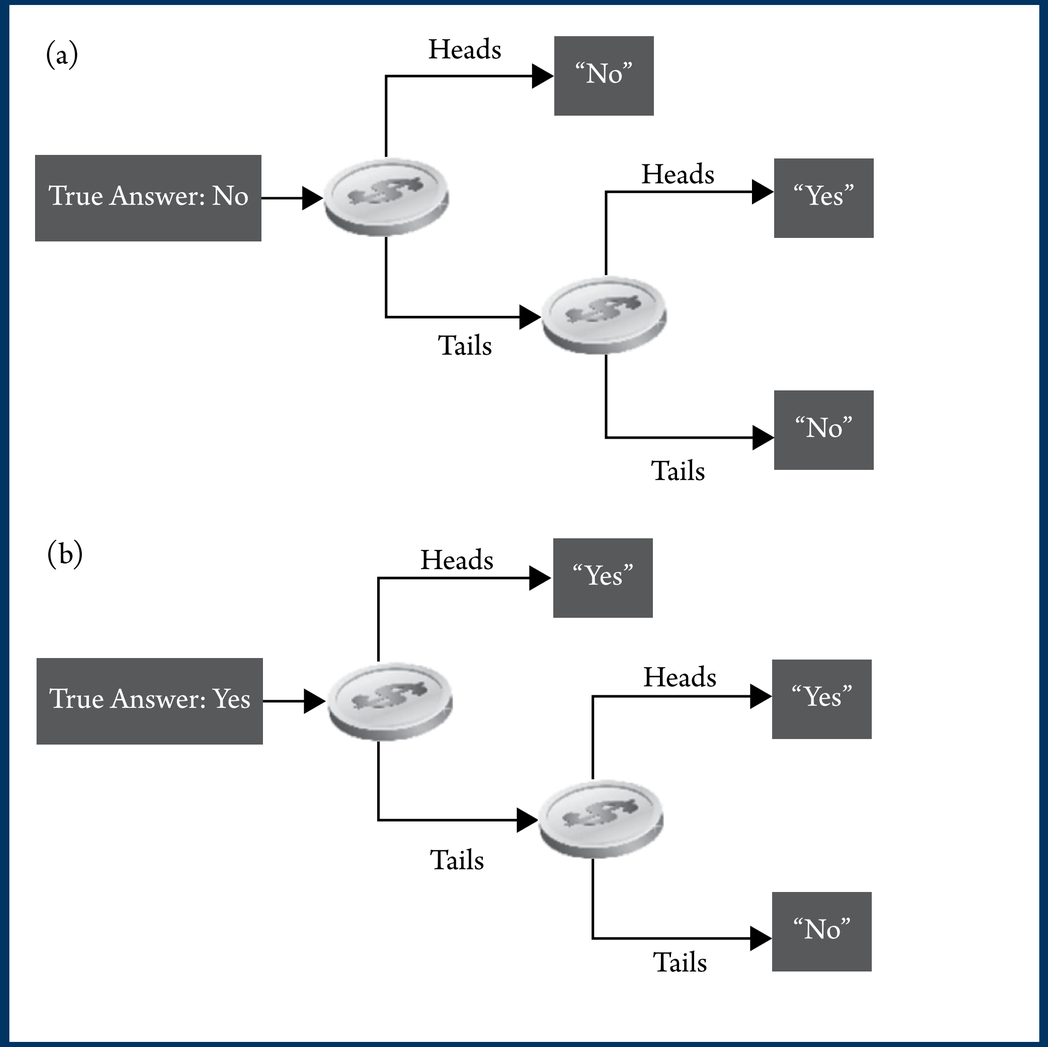

Randomness could help

If you say YES, it could be that you use torrents or because you had tails, then heads

3/4 times we say truth, 1/4 times we lie

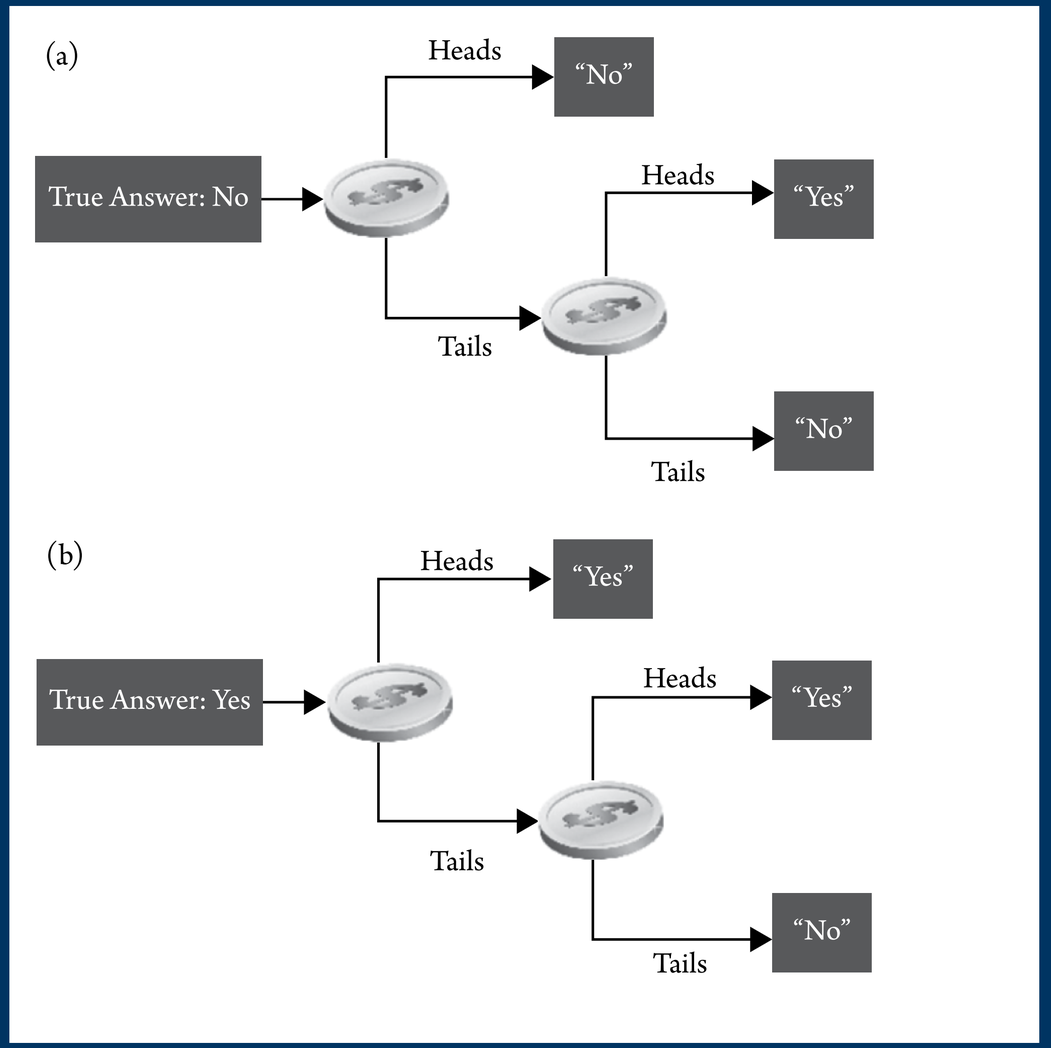

Nobody can form strong beliefs about the true data of any single individual!

How to extract the information?

Suppose 1/3 people use torrents. How many YES answers do we expect?

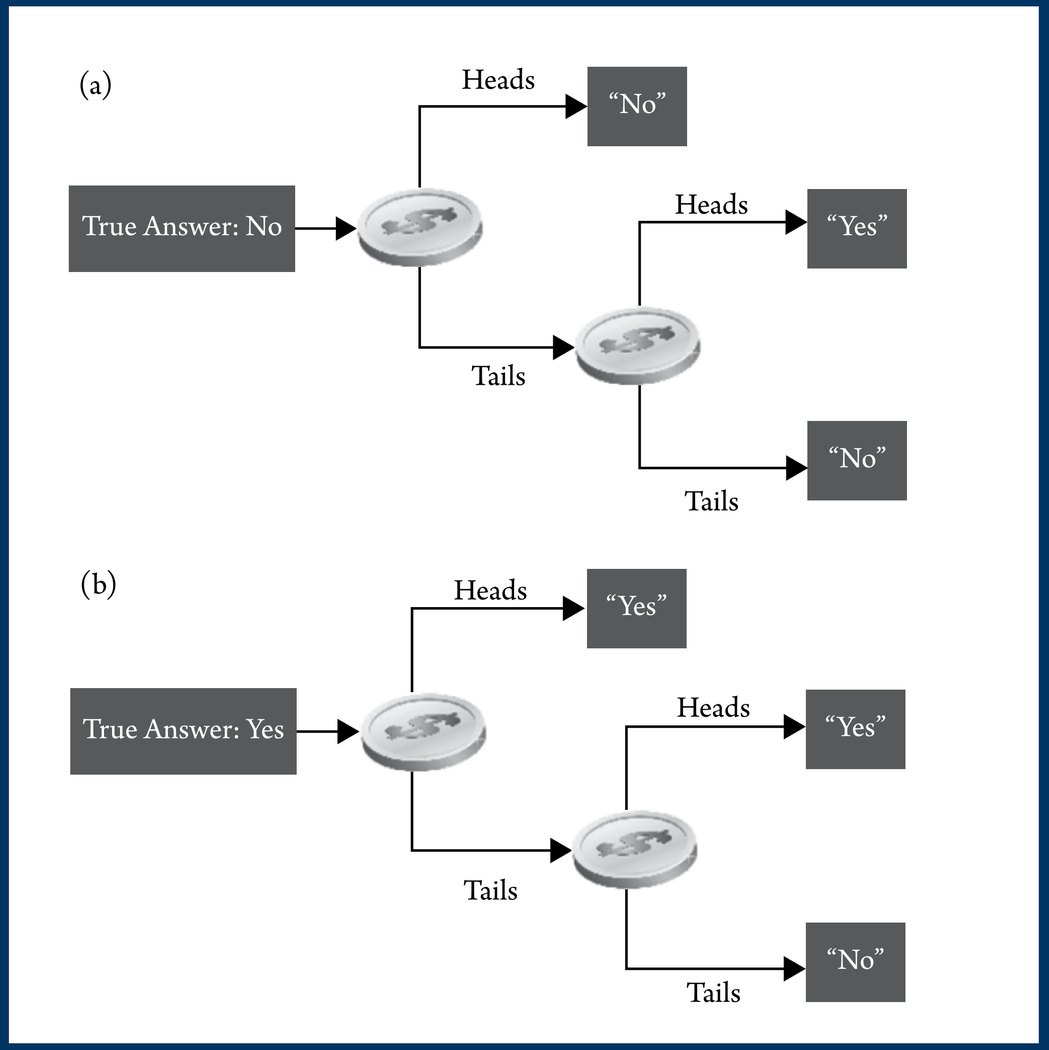

How to extract the information?

Suppose 1/3 people use torrents. How many YES answers do we expect?

3/4 times we say truth, 1/4 times we lie

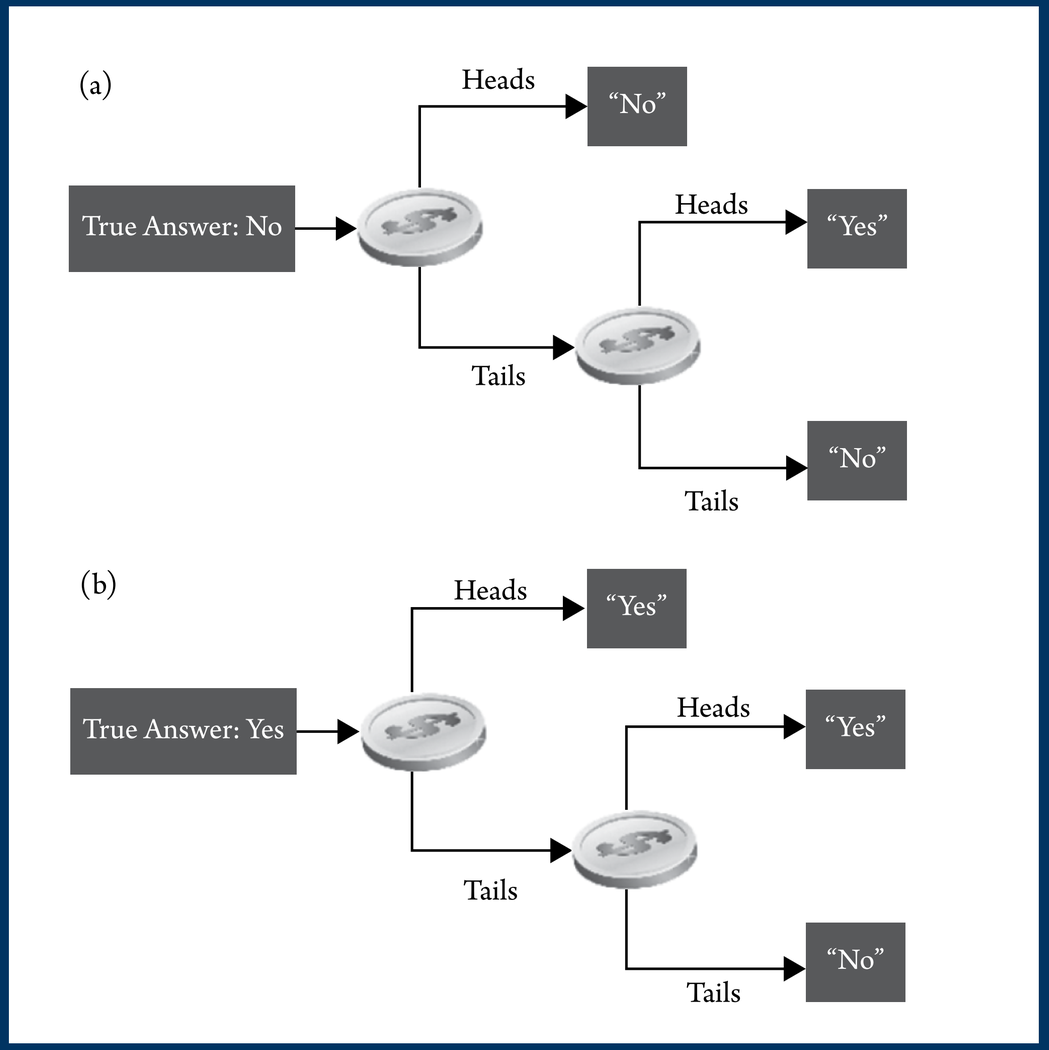

How to extract the information?

Suppose 1/3 people use torrents. How many YES answers do we expect?

3/4 times we say truth, 1/4 times we lie

1/3 x 3/4 + 1/4 x 2/3=1/4+1/6=5/12

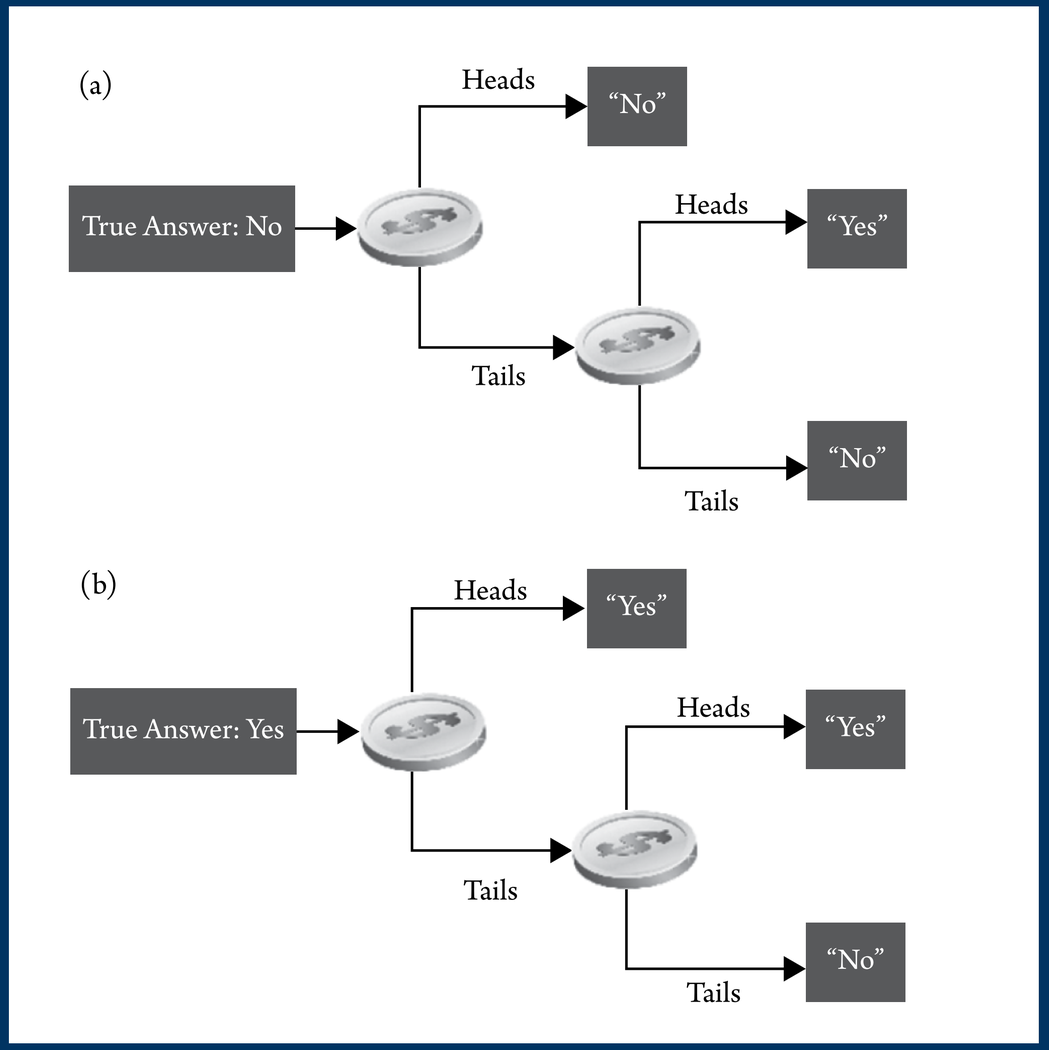

Remark

1/3 x 3/4 + 1/4 x 2/3=1/4+1/6=5/12

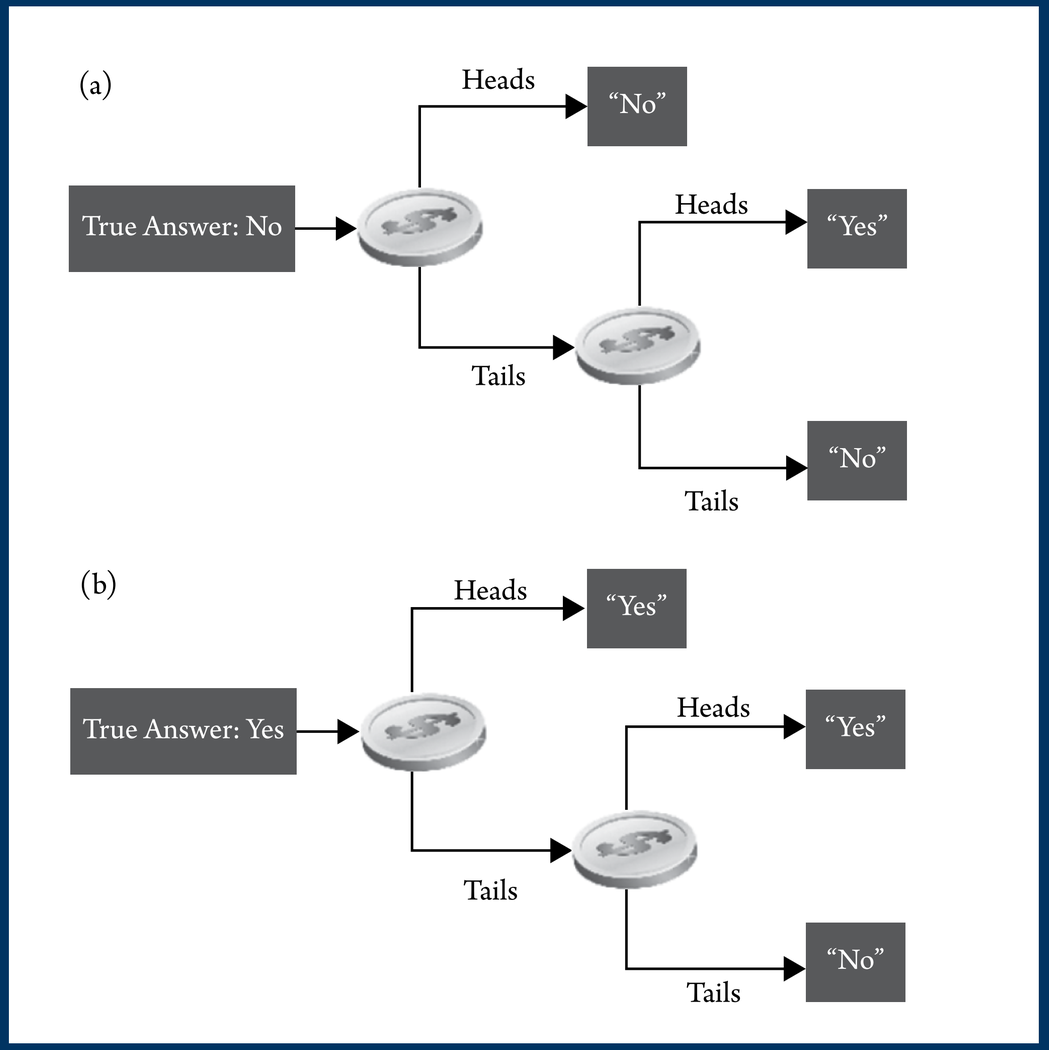

Remark

1/3 x 3/4 + 1/4 x 2/3=1/4+1/6=5/12

5/12 is on average

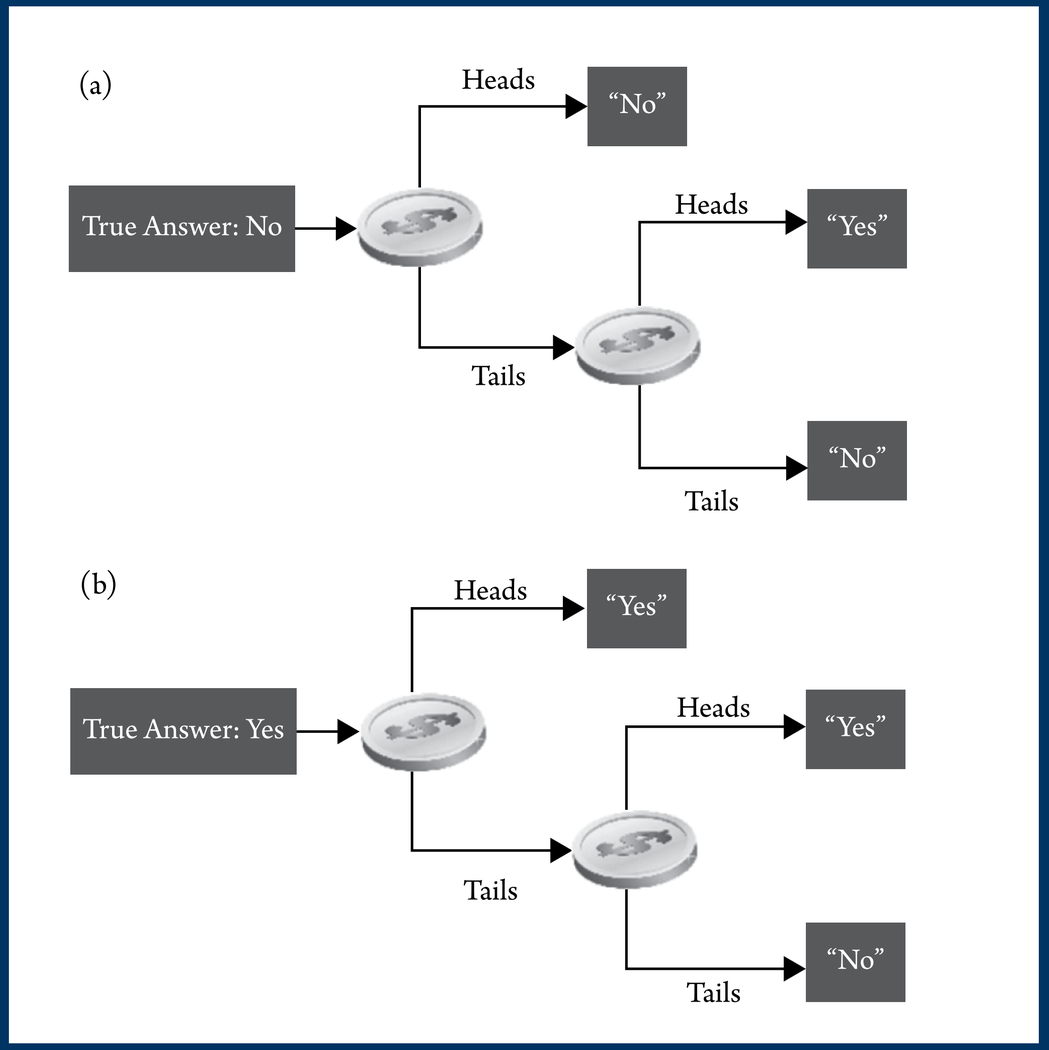

Remark

1/3 x 3/4 + 1/4 x 2/3=1/4+1/6=5/12

5/12 is on average

The law of big numbers: the more people we have, the smaller is the error.

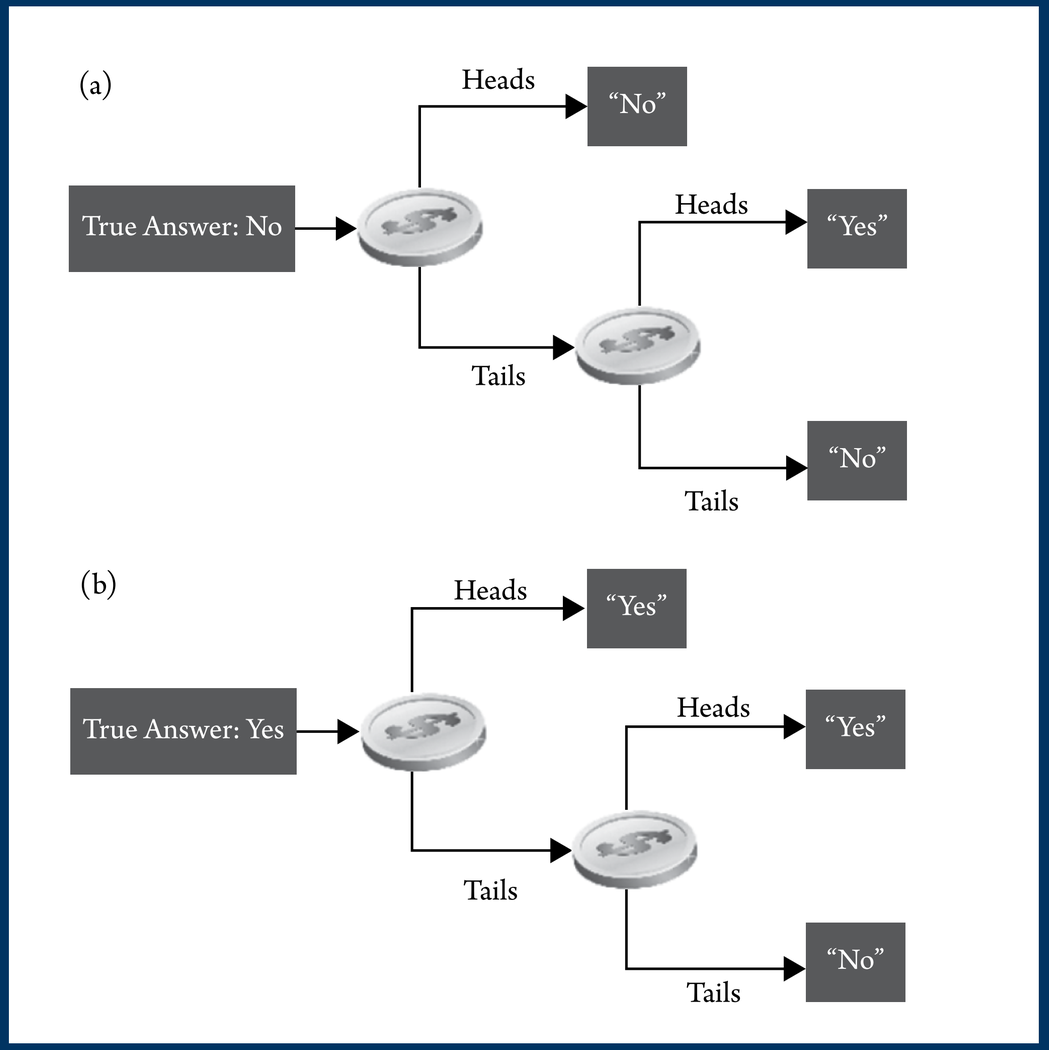

Differential privacy: informal idea

Results of something (a dataset, a model, a statistic, etc.) do not change when any individual is removed from or added to the data.

It is not possible to discover any one person’s contribution because the data you’re looking for would be the same regardless of whether or not their data was present.

Differential privacy

$\mathbf{X}$: The data *universe*

$\mathbf{D \subset X}$: The dataset (one element per person)

Definition: Two datasets $\mathrm{D}, \mathrm{D}' \subset \mathrm{X}$ are *neighbors* if they differ in the data of a single individual.

Differential privacy

X: The data *universe*

$D \subset X$: The dataset (one element per person)

Definition: An algorithm M is ϵ-differentially private if for all pairs of neighboring datasets D, D', and for all outputs x:

$$\Pr[M(D)=x] \leq (1+\epsilon)\Pr[M(D')=x]$$

Differential privacy

Definition: An algorithm M is ϵ-differentially private if for all pairs of neighboring datasets D, D', and for all outputs x:

$$\Pr[M(D)=x] \leq (1+\epsilon)\Pr[M(D')=x]$$

This is a mathematical concept.
You cannot run DP on something.

Differential privacy

The algorithmic questions: what level of noise is required, what distribution, how to apply it correctly to complicated datasets, are often very difficult.

Differential privacy in practice

2008: U.S. Census Bureau, for showing commuting patterns.
2014: Google's RAPPOR, for telemetry such as learning statistics about unwanted software hijacking users' settings.
2015: Google, for sharing historical traffic statistics.
2016: Apple announced its intention to use differential privacy in iOS 10 to improve its Intelligent personal assistant technology.
2017: Microsoft, for telemetry in Windows.
2020: LinkedIn, for advertiser queries.

Is DP a silver bullet?

"Don't worry, whatever you tell us will be DP-protected."

Is DP a silver bullet?

"Don't worry, whatever you tell us will be DP-protected."

A meaningful privacy guarantee, the analytical utility of DP outputs is likely to be very poor

A meaningful analytical utility could result in a poor privacy guarantee.

Further reading

- The Ethical Algorithm: The Science of Socially Aware Algorithm Design by Michael Kearns and Aaron Roth
- A Gentle Introduction to Differential Privacy by Tom Titcombe
- The Limits of Differential Privacy (and Its Misuse in Data Release and Machine Learning) By Josep Domingo-Ferrer, David Sánchez, Alberto Blanco-Justicia Communications of the ACM, July 2021, Vol. 64 No. 7, Pages 33-35 10.1145/3433638
- The Algorithmic Foundations of Differential Privacy by Cynthia Dwork and Aaron Roth