CSE660 Differential Privacy

September 11, 2017

Marco Gaboardi

Room: 338-B

gaboardi@buffalo.edu

http://www.buffalo.edu/~gaboardi

Question: How can we make statistical queries private?

Private Statistical database

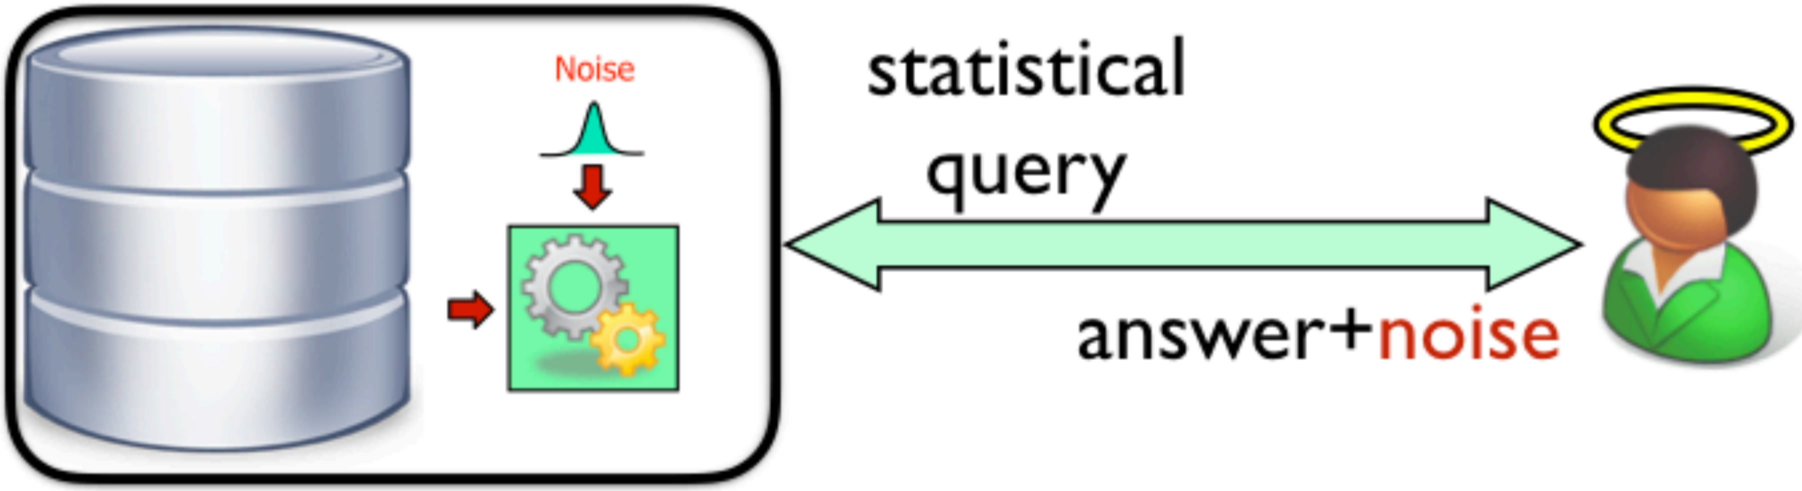

Question: What kind of noise?

Privacy vs Utility

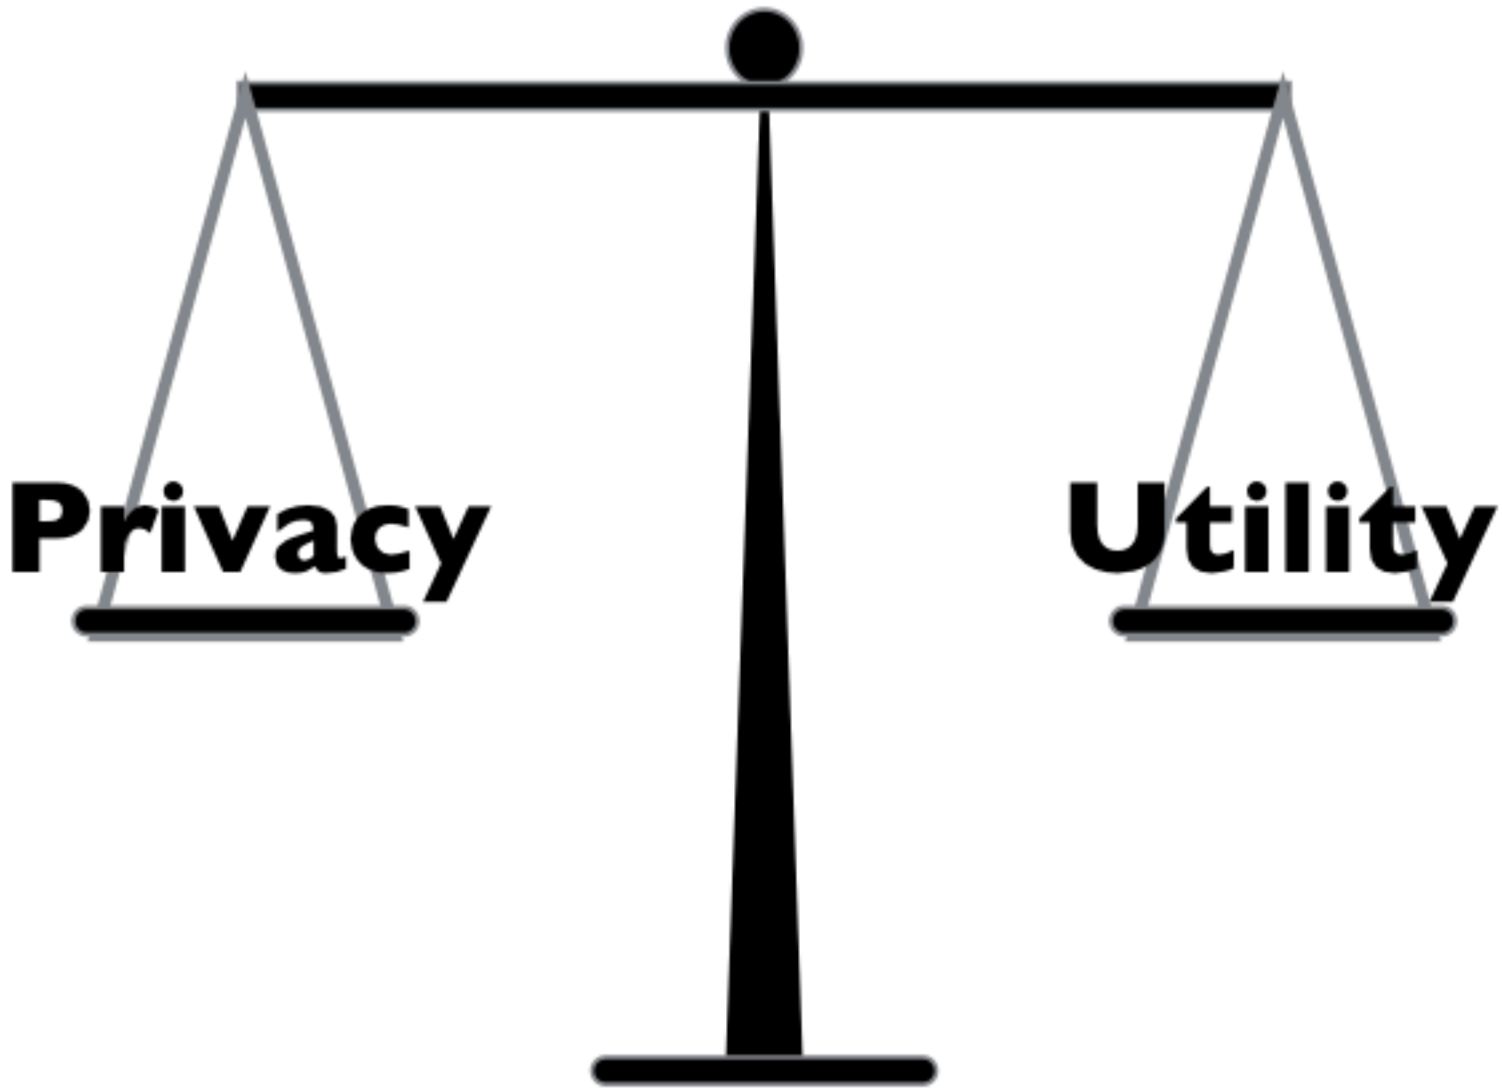

Question: Does this approach protect privacy?

Reconstruction attack

- Consider an adversary A (an algorithm) that has access to some data D through a privacy mechanism q*.
- The goal of the adversary is to output some data D' that is as similar as possible to D.
- To output D' the adversary can interact several times with q*.

Reconstruction attack

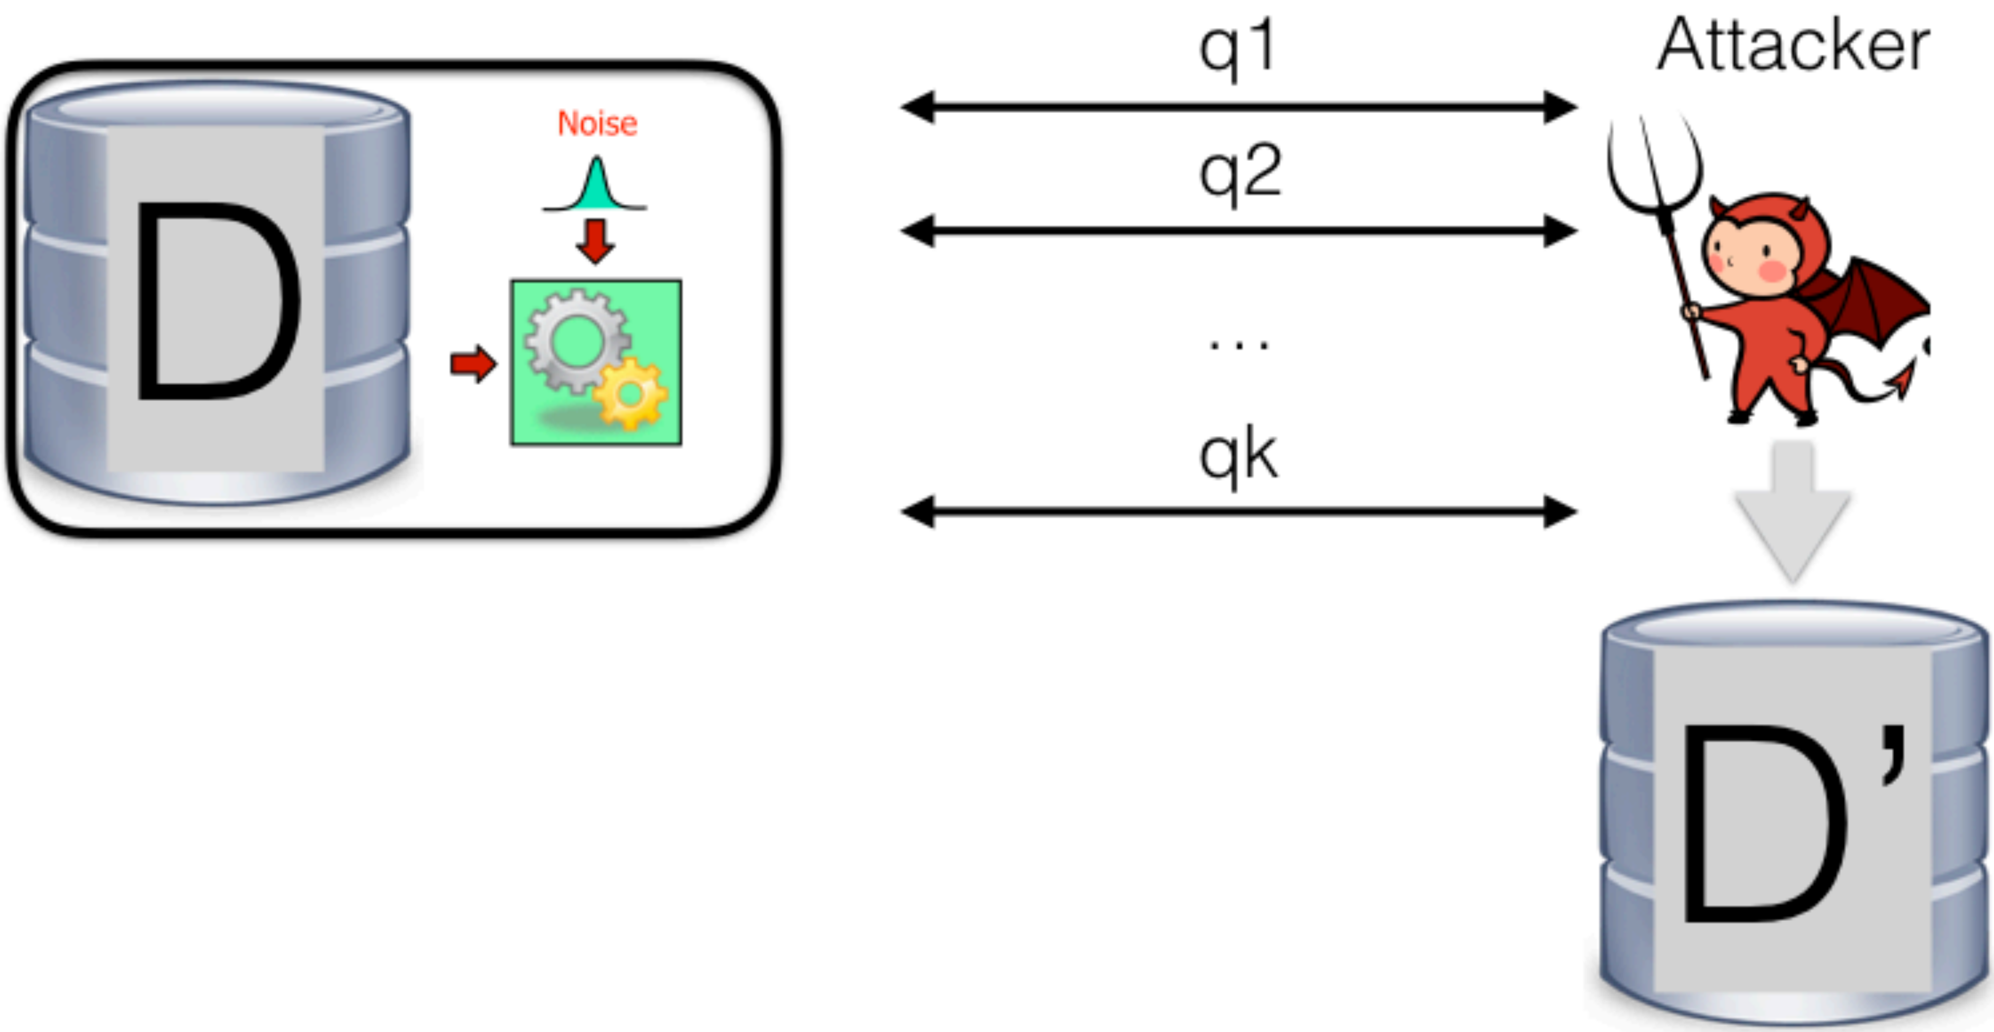

Reconstruction attack

We say that the attacker wins if

$$d(D, D') \sim 0$$

In our case we can use Hamming distance

Blatantly non-privacy

The privacy mechanism $M:\{0,1\}^n \rightarrow R$ is blatantly non-private if an adversary can build a candidate database $D' \in \{0,1\}^n$, that agrees with the real database D in all but $o(n)$ entries:

$$d_H(D,D') \in o(n)$$

Reconstruction attack with exponential adversary

Let $M:\{0,1\}^n \rightarrow R$ be a privacy mechanism adding noise within E perturbation. Then there is an adversary that can reconstruct the database within 4E positions.

[DinurNissim'02]

Reconstruction attack with exponential adversary

Let $M:\{0,1\}^n \rightarrow R$ be a privacy mechanism adding noise within **$E=o(n)$** perturbation. Then M is blatantly non-private against an adversary A running in exponential time.

[DinurNissim'02]

Reconstruction attack with polynomial adversary

Let $M:\{0,1\}^n \to R$ be a privacy mechanism adding noise within **$E=o(\sqrt{n})$** perturbation. Then we can show M blatantly non-private against an adversary A running in polynomial time and **answering n queries.**

[DinurNissim'02, DworkYekhanin'08]

Sample error

- Suppose that a database contains n individuals drawn uniformly at random from a population of size N>>n.
- Suppose we are interested in a medical condition that affects a fraction p of the population.
- Then we expect the number of individuals in the dataset with condition p is

$$np \pm \Theta(\sqrt{n})$$

- The sampling error is of the order of $\sqrt{n}$.

We would like the noise we introduce for privacy to be smaller than the sampling error.

Privacy

Preventing blatantly non-privacy is a low bar for a privacy mechanism.
However, we should expect that the previous results apply to other stronger notions, in particular Differential Privacy.

Assignment?

Quantitative notions of Privacy

- The impossibility results discussed above suggest a quantitative notion of privacy,
- A notion where the privacy loss depends on the number of queries that are allowed.

How much privacy loss shall we allow?

Privacy-preserving data analysis?

- The analyst knows no more about me after the analysis than what she knew before the analysis.

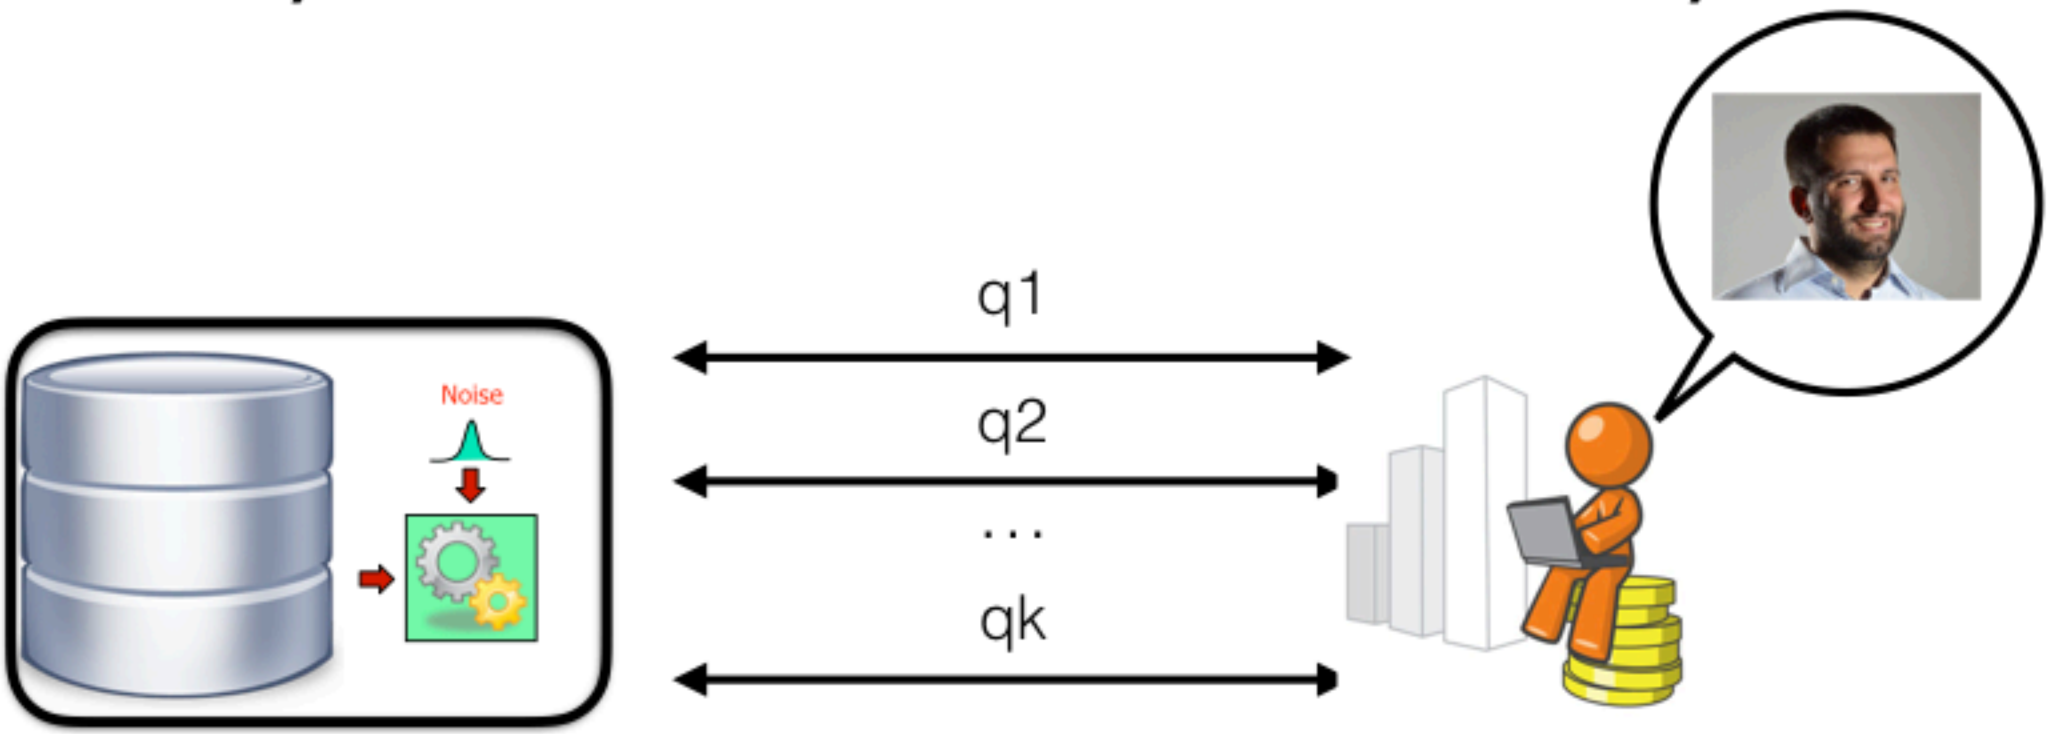

Privacy-preserving data analysis?

Prior Knowledge

~

Posterior Knowledge

Privacy-preserving data analysis?

Question: What is the problem with this requirement?

Privacy-preserving data analysis?

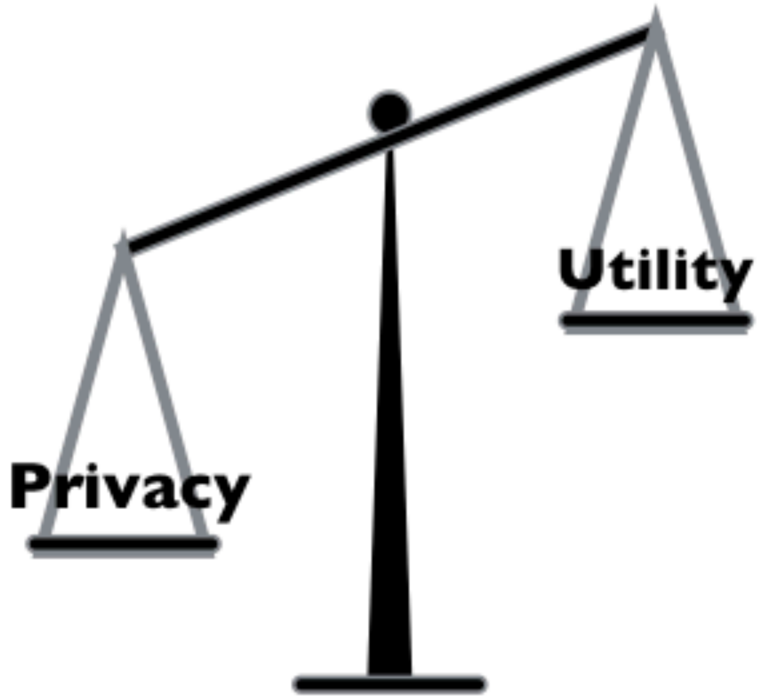

If nothing can be learned about an individual, then nothing at all can be learned at all!

[DworkNaor10]

Privacy-preserving data analysis?

- The analyst learn the same about me after the analysis as what she would have learnt if I didn't contribute my data.

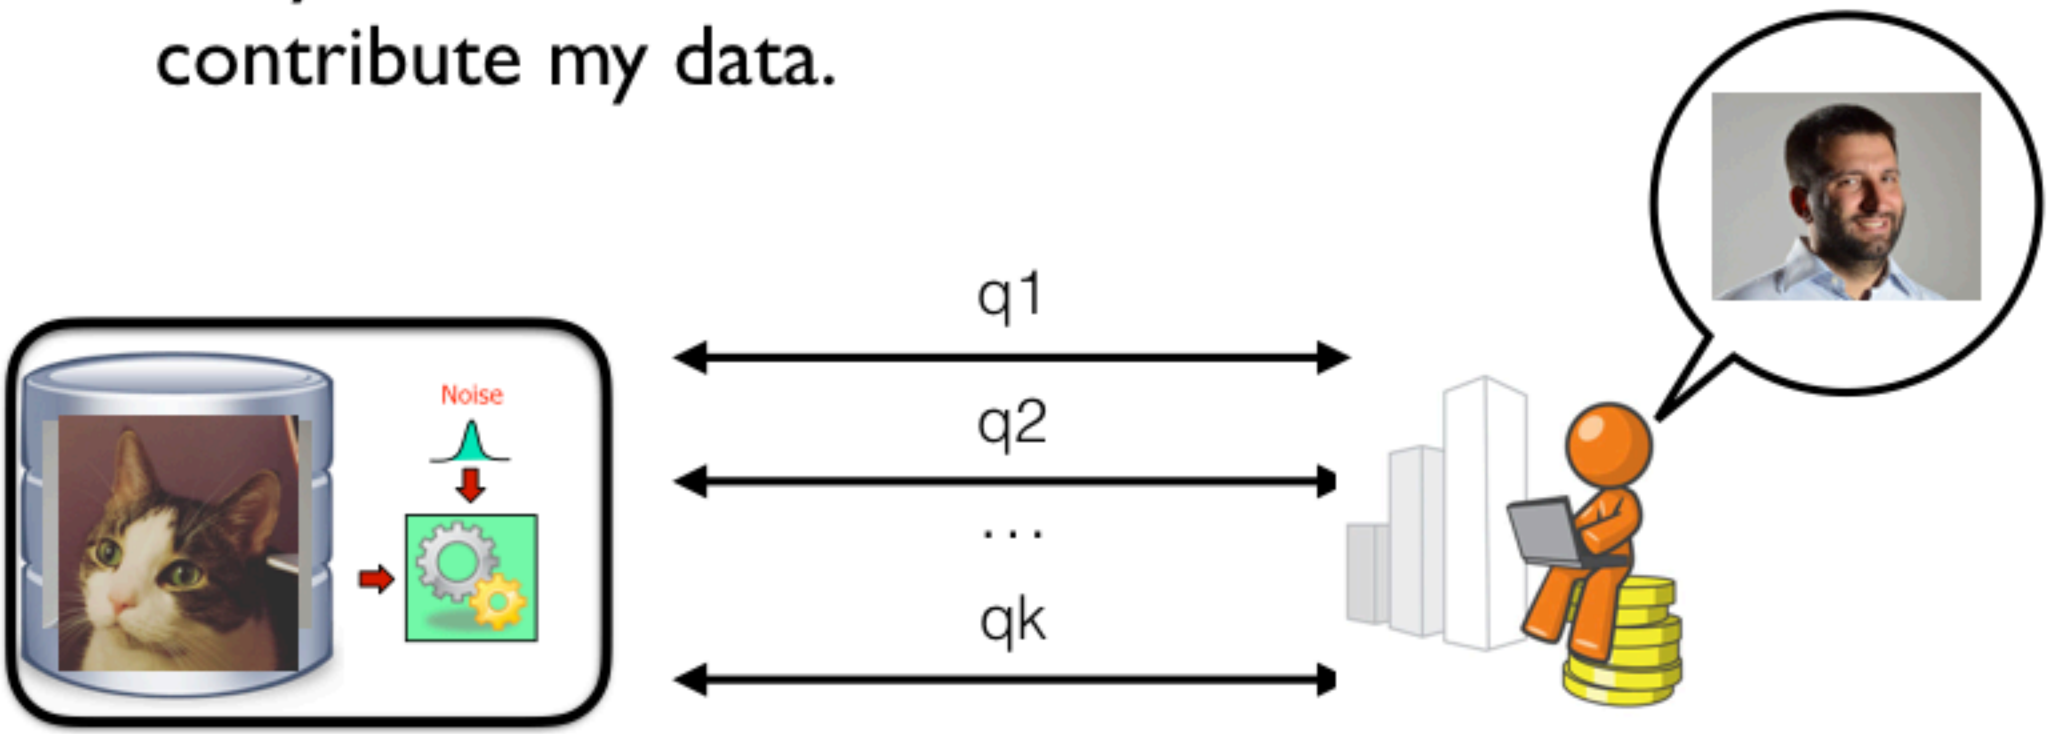

Adjacent databases

- We can formalize the concept of contributing my data or not in terms of a notion of distance between datasets.
- Given two datasets $D, D' \in \{0,1\}^n$, their distance is defined as:

$$D \Delta D' = |\{k \leq n \mid D(k) \neq D'(k)\}|$$

- We will call two datasets adjacent when $D \Delta D' = 1$ and we will write $D \sim D'$.

(ε,δ)-Differential Privacy

Definition

Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

(ε,δ)-Differential Privacy

A query returning a probability distribution

Definition

Given ε,δ ≥ 0, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

(ε,δ)-Differential Privacy

Privacy parameters

Definition

Given $\varepsilon, \delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

(ε,δ)-Differential Privacy

Definition

Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \to R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$Pr[Q(b_1) \in S] \leq \exp(\varepsilon) Pr[Q(b_2) \in S] + \delta$$

a quantification over all the databases

(ε,δ)-Differential Privacy

Definition

Given $\varepsilon, \delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

a notion of adjacency or distance

(ε,δ)-Differential Privacy

Definition

Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon) \Pr[Q(b_2) \in S] + \delta$$

and over all the possible outcomes

ε-Differential Privacy

Definition

Given $\varepsilon \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is ε-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon) \Pr[Q(b_2) \in S]$$

ε-Differential Privacy

Definition

Given $\varepsilon \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is ε-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon) \Pr[Q(b_2) \in S]$$

Let's substitute a concrete instance:

$$\Pr[Q(b \cup \{x\}) \in S] \leq \exp(\varepsilon) \Pr[Q(b \cup \{y\}) \in S]$$

ε-Differential Privacy

Definition

Given $\varepsilon \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is ε-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$Pr[Q(b_1) \in S] \leq \exp(\varepsilon) Pr[Q(b_2) \in S]$$

Let's substitute a concrete instance:

$$Pr[Q(b \cup \{x\}) \in S] \leq \exp(\varepsilon) Pr[Q(b \cup \{y\}) \in S]$$

Let's use the two quantifiers:

$$\exp(-\varepsilon) Pr[Q(b \cup \{y\}) \in S] \leq Pr[Q(b \cup \{x\}) \in S] \leq \exp(\varepsilon) Pr[Q(b \cup \{y\}) \in S]$$

ε-Differential Privacy

Definition

Given $\varepsilon \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is ε-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S]$$

Let's substitute a concrete instance:

$$\Pr[Q(b \cup \{x\}) \in S] \leq \exp(\varepsilon)\Pr[Q(b \cup \{y\}) \in S]$$

Let's use the two quantifiers:

$$\exp(-\varepsilon)\Pr[Q(b \cup \{y\}) \in S] \leq \Pr[Q(b \cup \{x\}) \in S] \leq \exp(\varepsilon)\Pr[Q(b \cup \{y\}) \in S]$$

And for $\varepsilon \rightarrow 0$

$$(1-\varepsilon)\Pr[Q(b \cup \{y\}) \in S] \leq \Pr[Q(b \cup \{x\}) \in S] \leq (1+\varepsilon)\Pr[Q(b \cup \{y\}) \in S]$$

ε-Differential Privacy

Definition

Given $\varepsilon \geq 0$, a probabilistic query $Q: X^n \rightarrow R$
is ε-differentially private iff
for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S]$$

Since we consider discrete distributions when $S=\{s_1,\ldots s_n\}$ we have:

$$\Pr[X \in S] = \Pr[X \in \{s_1\}] + \ldots + \Pr[X \in \{s_n\}]$$

So we can rewrite the condition above as for every $r \in R$:

$$\Pr[Q(b_1)=r] \leq \exp(\varepsilon)\Pr[Q(b_2)=r]$$

ε-Differential Privacy

In general we can think about the following quantity as the privacy loss incurred by observing r on the databases b and b'.

$$L_{b,b'}(r) = \log \frac{\Pr[Q(b)=r]}{\Pr[Q(b')=r]}$$

ε-Differential Privacy

Q : db => R   probabilistic

$Q(b \cup \{x\})$ $Q(b \cup \{y\})$

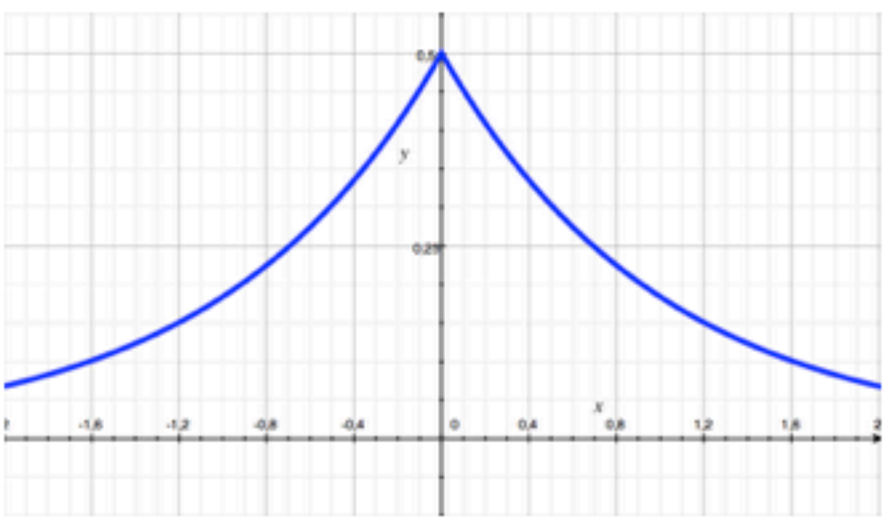

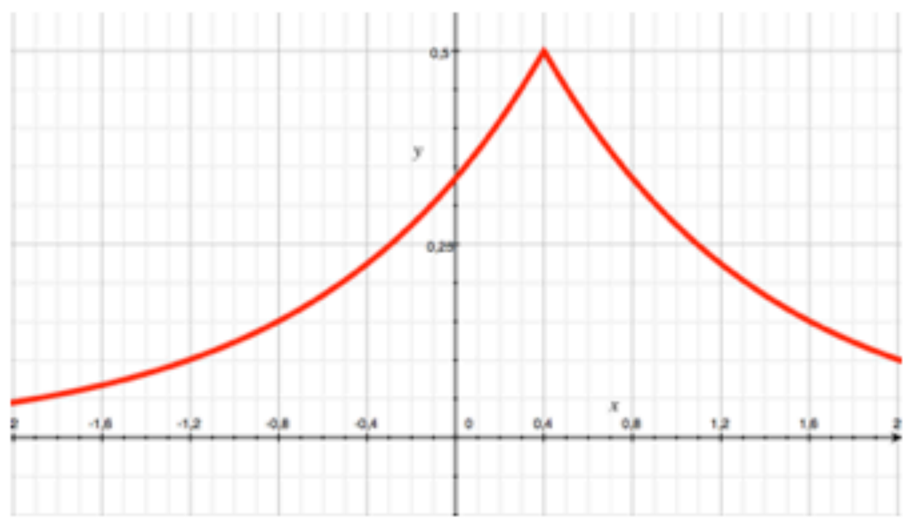

ε-Differential Privacy

$$d(Q(b \cup \{x\}), Q(b \cup \{y\})) \leq \varepsilon$$

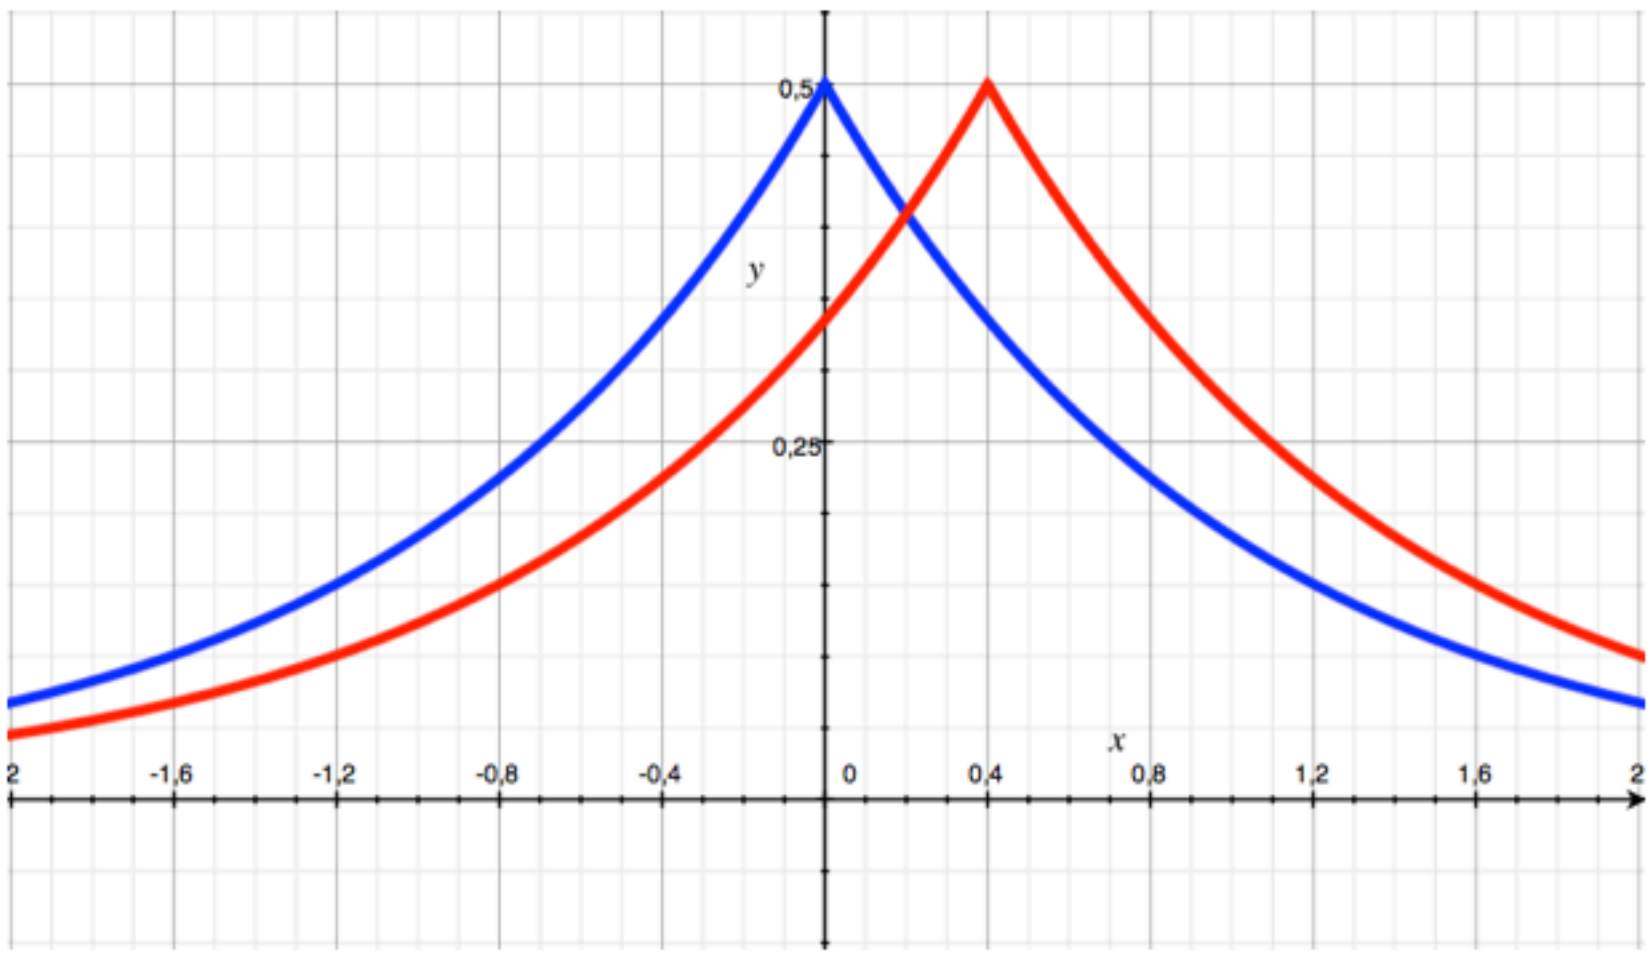

ε-Differential Privacy

$$\left|\log \frac{\Pr[Q(b \cup \{x\}) = r]}{\Pr[Q(b \cup \{y\}) = r]}\right| \leq \varepsilon$$

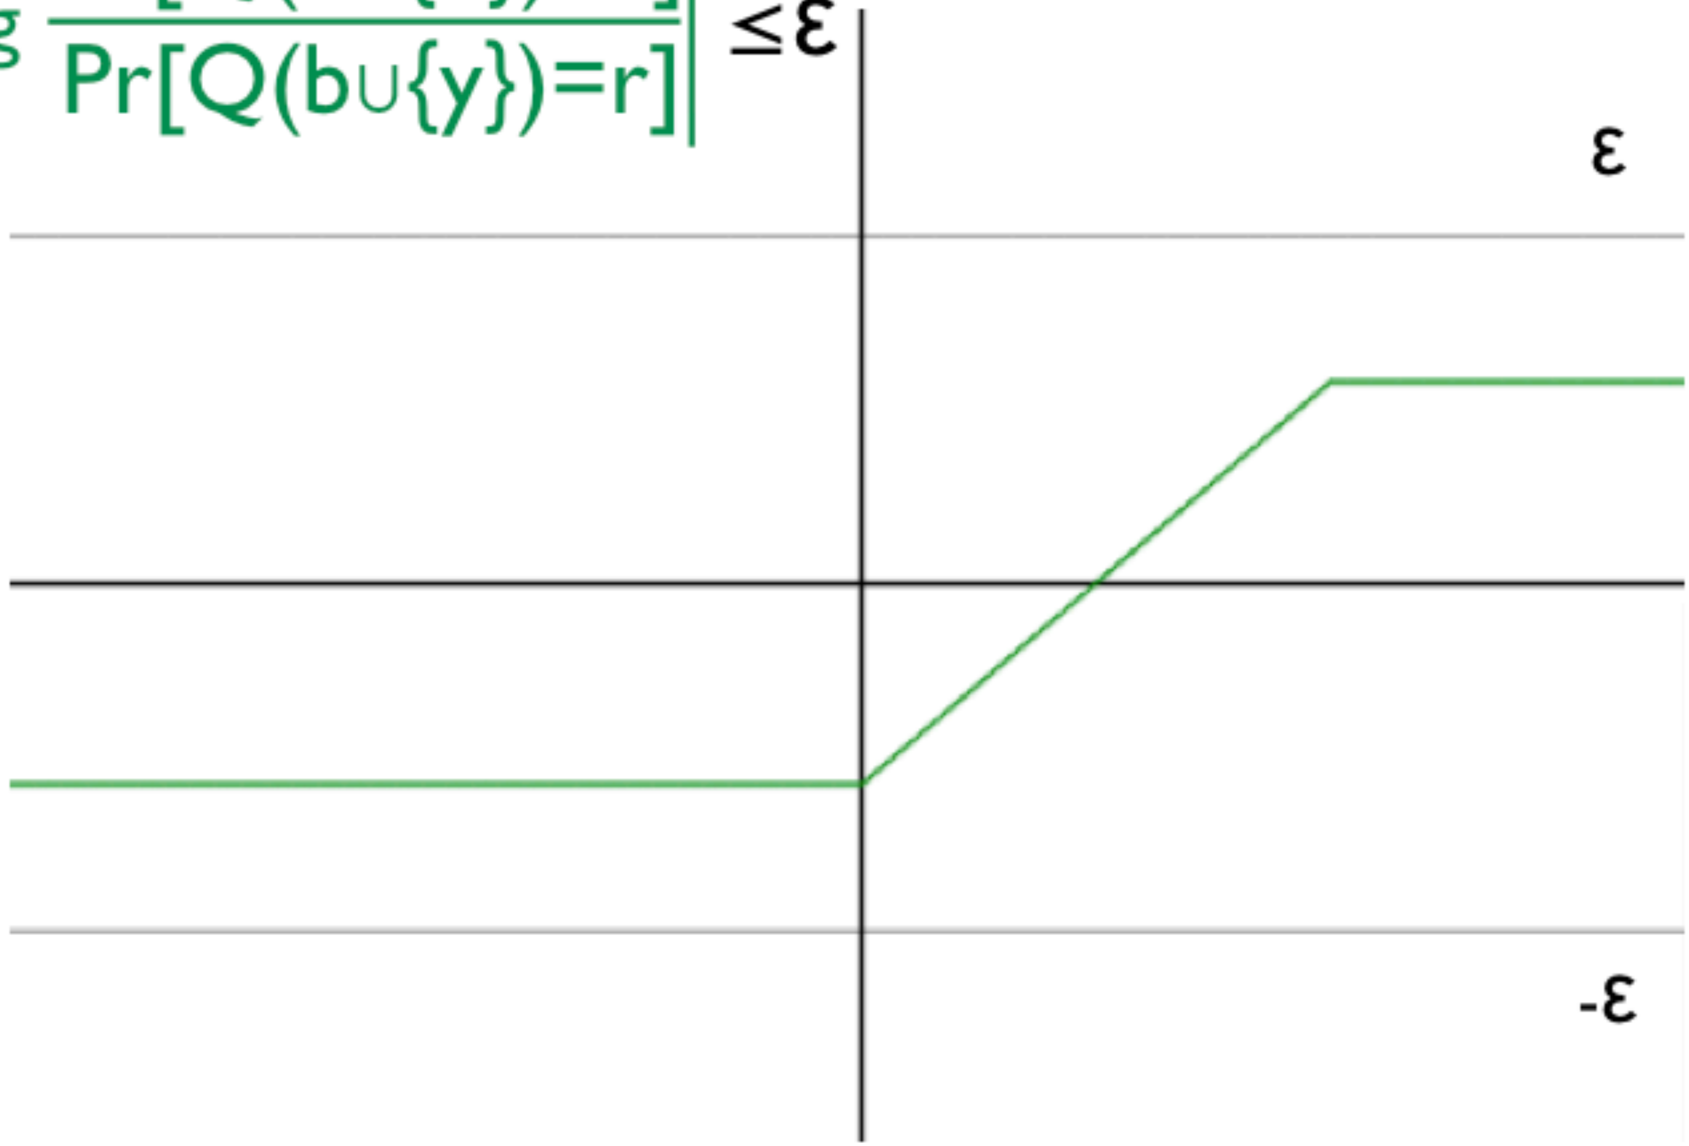

(ε,δ)-Differential Privacy

Definition

Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

(ε,δ)-Differential Privacy

Definition
Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

Similarly, we have

$$\log \frac{\Pr[Q(b_1) \in S] - \delta}{\Pr[Q(b_2) \in S]} \leq \varepsilon$$

$$-\varepsilon \leq \log \frac{\Pr[Q(b_1) \in S] + \delta}{\Pr[Q(b_2) \in S]}$$

(ε,δ)-Differential Privacy

Definition
Given $\varepsilon,\delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ε,δ)-differentially private iff for all adjacent database b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\varepsilon)\Pr[Q(b_2) \in S] + \delta$$

Similarly, we have

$$\log \frac{\Pr[Q(b_1) \in S] - \delta}{\Pr[Q(b_2) \in S]} \leq \varepsilon$$

$$-\varepsilon \leq \log \frac{\Pr[Q(b_1) \in S] + \delta}{\Pr[Q(b_2) \in S]}$$

Probability of failure

(ε,δ)-Differential Privacy

$$\left| \log \frac{\Pr[Q(b_1)=r]}{\Pr[Q(b_2)=r]} \right| \leq \varepsilon \quad \text{with probability } 1-\delta$$

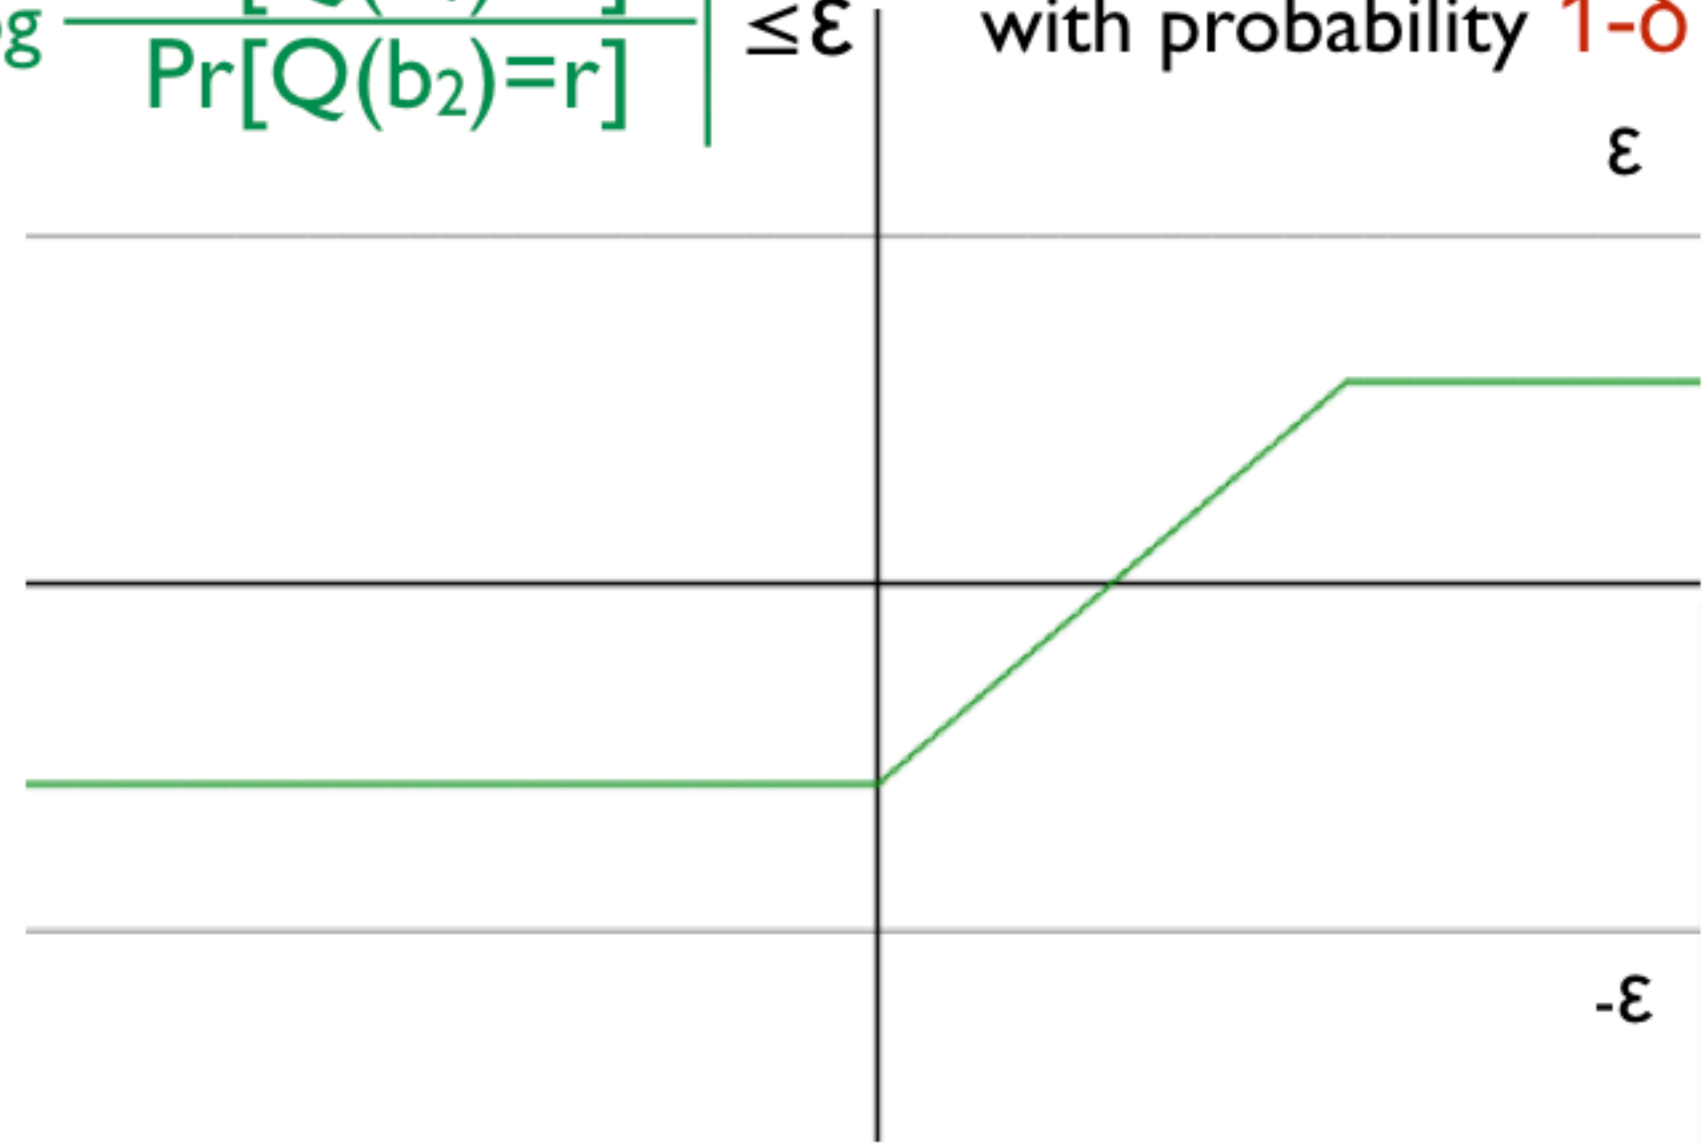

Randomized Response

[Warner65]

Suppose I ask a yes/no question.

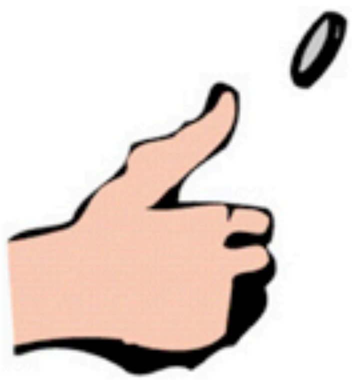

biased
coin

True
answer

Opposite
answer

The value of the bias is what determines the epsilon

An example

```
AlmostRandom (b : bool) : bool {
   if coinToss()
   then
      return b;
   else
      return coinToss();
}
```

An example

AlmostRandom is (ln 3,0)-differentially private

Let's consider the case we have two adjacent data `b` and `b'`. By the fact that they are adjacent we know that one of them is 1 and one of them is 0. Without loss of generality let's assume b=1 and b'=0.

We have:

$$\Pr[\mathbf{AR(b)} = 1] = 3/4 \qquad \Pr[\mathbf{AR(b')} = 1] = 1/4$$

$$\Pr[\mathbf{AR(b)} = 0] = 1/4 \qquad \Pr[\mathbf{AR(b')} = 0] = 3/4$$

So:

$$\left| \frac{\Pr[\mathbf{AR(b)} = R]}{\Pr[\mathbf{AR(b')} = R]} \right| \leq 3$$