Foundations of Privacy

Lecture 4

Recalling the basic notions

Sex	Blood	···	HIV
F	B	···	Y
M	A	···	N
M	O	···	N
M	O	···	Y
F	A	···	N
M	B	···	Y

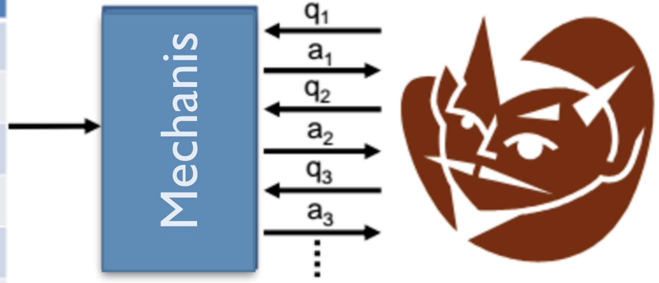

- A randomized mechanism for the query f is any probabilistic function $\mathcal{K}$ from $\mathcal{X}$ to a set of values $\mathcal{Z}$. Namely,

$$\mathcal{K} : \mathcal{X} \to \mathcal{DZ}$$

where $\mathcal{DZ}$ represents the set of probability distributions on $\mathcal{Z}$.

- **Definition (Differential Privacy)** $\mathcal{K}$ is ε-differentially private if for every pair of databases $x_1, x_2 \in \mathcal{X}$ such that $x_1 \sim x_2$, and for every $z \in Z$, we have:

$$p(Z = z|X = x_1) \leq e^{\varepsilon} p(Z = z|X = x_2)$$

where $p(Z = z|X = x)$ represents the conditional probability of z given x, namely the probability that on the database x the mechanism reports the answer z

Oblivious Mechanisms

- Given $f: \mathcal{X} \rightarrow \mathcal{Y}$ and $\mathcal{K}: \mathcal{X} \rightarrow \mathcal{Z}$, we say that $\mathcal{K}$ is oblivious if it can be seen as the composition of f and a randomized mechanism $\mathcal{H}$ (noise) defined on the exact answers $\mathcal{K} = \mathcal{H} \circ f$

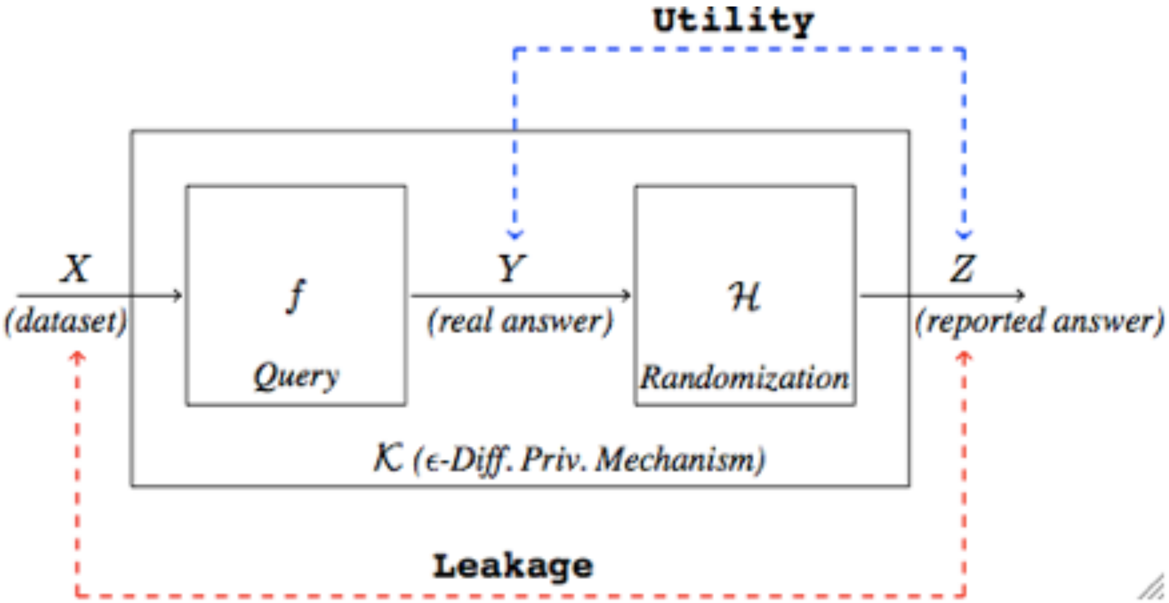

A typical oblivious mechanism: Laplacian noise

- Query $f\colon \mathcal{X} \rightarrow \mathcal{Y}$.
- **Laplacian noise.** If the exact answer is y, the reported answer is z, with a probability density function defined as:

$$dP_y(z) = c\, e^{-\frac{|z-y|}{\Delta f}\varepsilon}$$

where Δf is the *sensitivity* of f:

$$\Delta f = \max_{x \sim x' \in \mathcal{X}} |f(x) - f(x')|$$

($x \sim x'$ means x and x' are adjacent, i.e., they differ only for one record)

and c is a normalization factor:

$$c = \frac{\varepsilon}{2\,\Delta f}$$

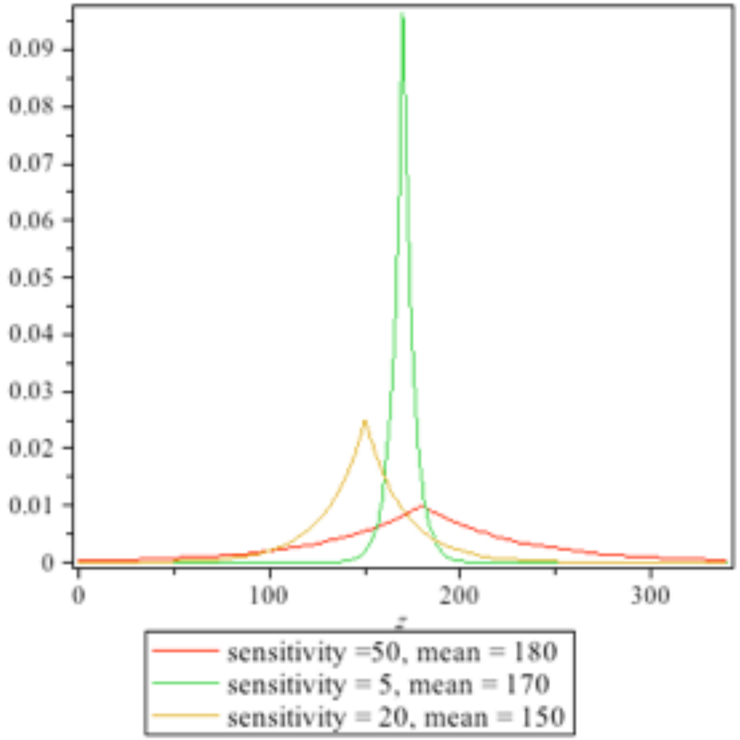

The geometric mechanism

- The Laplacian noise is typically used in the case that $\mathcal{Y}$ (the set of true answers of the query) is a **dense** numerical set, like the Reals or the Rationals.

- If $\mathcal{Y}$ is a **discrete** numerical set, like the Integers, then the typical mechanism used in this case is the **geometric mechanism**, which is a sort of discrete Laplacian.

- In the geometric mechanism, the probability distribution of the noise is:

$$p(z|y) = c\, e^{-\frac{|z-y|}{\Delta f}\varepsilon}$$

$$c = \frac{1-\alpha}{1+\alpha} \qquad \text{where} \qquad \alpha = e^{-\frac{\varepsilon}{\Delta f}}$$

- Δf is the sensitivity of query f

Utility

- The utility $\mathcal{U}$ of a mechanism is the maximum expected gain over all possible databases. The maximum is over all possible remappings: It is assumed that the user is rational and therefore makes the guesses that are the most useful to him. Note that $\mathcal{U}$ depends also on the prior π over $\mathcal{X}$ Formally, let us denote by r a remapping function. For an oblivious mechanism we have:

$$\mathcal{U}(\mathcal{K},\pi,g) = \max_r \sum_x \pi(x) \sum_z p_{\mathcal{H}}(z|f(x))g(r(z),f(x))$$

For a general (possibly non-oblivious) mechanism, we have:

$$\mathcal{U}(\mathcal{K},\pi,g) = \max_r \sum_x \pi(x) \sum_z p_{\mathcal{K}}(z|x)g(r(z),f(x))$$

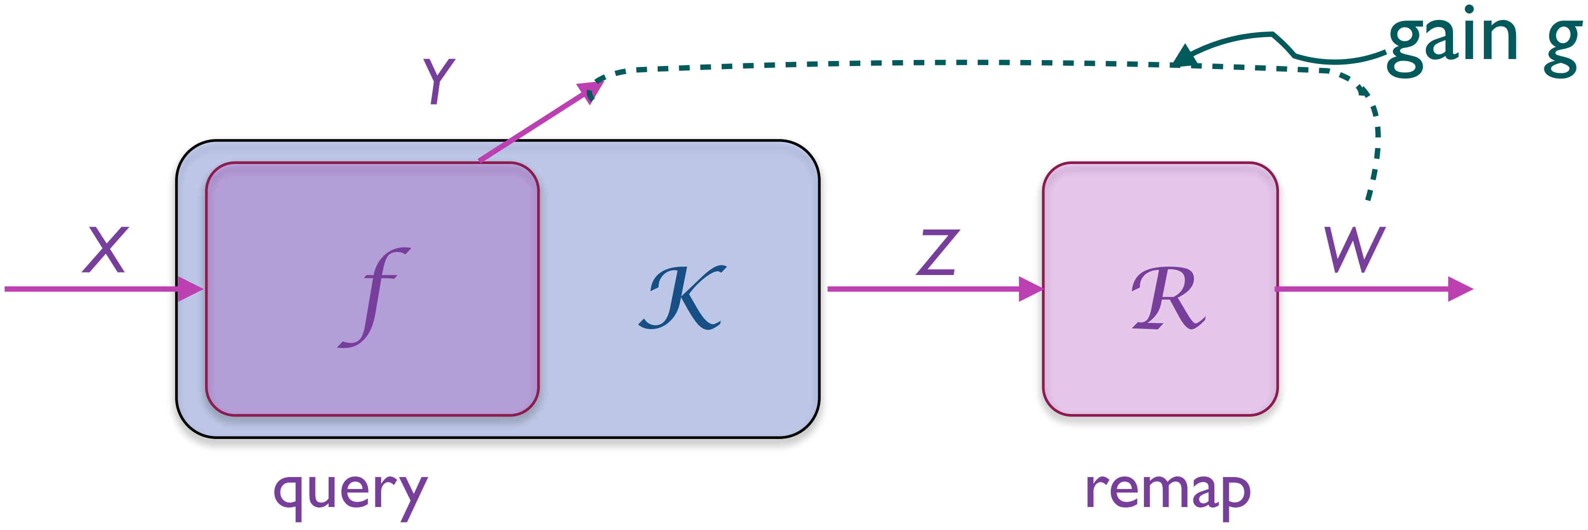

Exercises

1. Define the noise density function for the Laplacian mechanism for the query "What is the percentile of the people in the DB who earn more than 10K Euro a month", assuming that the database contains at least 1000 records.

2. Define the truncated geometric mechanism for the query ``how many people earn more than 10K Euro a month", assuming that the database contains exactly 100 records. Note that the true answers are in the interval [0,100].

3. Prove that ε-differential privacy can be equivalently defined as follows

   $\mathcal{K}$ is ε-differentially private if for every pair of databases $x_1, x_2 \in \mathcal{X}$ (not necessarily adjacents), and for every $z \in \mathcal{Z}$, we have:

   $$p(Z = z|X = x_1) \leq e^{\varepsilon h(x_1,x_2)} p(Z = z|X = x_2)$$

   where $h(x_1, x_2)$ represents the Hamming distance between x_1 and x_2

Exercises

4. Compute the utility of the geometric mechanism for the counting query "how many earn more than 10K Euro a month", with privacy degree ε, assuming a uniform prior distribution on the true answers, with the gain function defined as the identity relation.

5. We saw that post-processing cannot decrease privacy. Can it decrease the utility? Motivate your answer.

6. Find a mechanism for the same counting query, with the same degree of privacy, but lower utility.

7. Show that the graph on $\mathcal{Y}$ induced by the query "what is the average age of the people in the database" has cycles

8. Prove the result of Brenner and Nissim (hint: find two prior distributions on $\mathcal{Y}$ which have different optimal mechanisms)

Optimal mechanisms

- Given a prior π, and a privacy level ε, an ε-differentially private mechanism K is called optimal if it provides the best utility among all those which provide ε-differential privacy
- Note that the privacy does not depend on the prior, but the utility (in general) does.
- In the finite case the optimal mechanism can be computed with linear optimization techniques, where the variables are the conditional probabilities p(z | y)
  where y is the exact answer and z is the reported answer
- A mechanism is universally optimal if it is optimal for all priors π

Counting Queries

- Counting queries are typical examples of discrete queries. They are of the form: How many individuals in the database satisfy the property $\mathcal{P}$?
  - Examples:
    - How many individuals are affected by diabetes?
    - How many diabetic people are obese?
- Question: what is the sensitivity of a counting query?

Privacy vs utility: two fundamental results

1. [Ghosh et al., STOC 2009]
   The geometric mechanism and the truncated geometric mechanism are universally optimal for counting queries and any anti-monotonic gain function

Privacy vs utility: two fundamental results

2. [Brenner and Nissim, STOC 2010] The counting queries are the only kind of queries for which a universally optimal mechanism exists

- This means that for other kind of queries one the optimal mechanism is relative to a specific user.
- The precise characterization is given in terms of the graph $(\mathcal{Y},\sim)$ induced by $(\mathcal{X},\sim)$

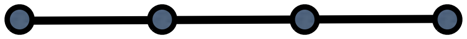

ok

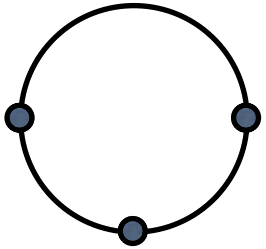

not ok

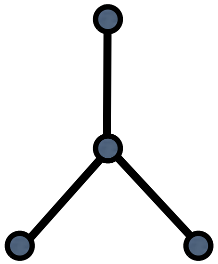

not ok

Exercises

4. Compute the utility of the geometric mechanism for the counting query "how many earn more than 10K Euro a month", with privacy degree ε, assuming a uniform prior distribution on the true answers, with the gain function defined as the identity relation.

5. We saw that post-processing cannot decrease privacy. Can it decrease the utility? Motivate your answer.

6. Find a mechanism for the same counting query, with the same degree of privacy, but lower utility.

7. Show that the graph on $\mathcal{Y}$ induced by the query "what is the average age of the people in the database" has cycles

8. Prove the result of Brenner and Nissim (hint: find two prior distributions on $\mathcal{Y}$ which have different optimal mechanisms)

Local Differential Privacy

- Differential Privacy assumes that we trust the party that collects the data.
- In the Local Differential Privacy Model, there is no such assumption. The individuals sanitize their data themselves, **before** they are collected
- Local Differential Privacy has become quite popular. For instance, it is the method that Google and Apple have adopted

Differential Privacy Schema

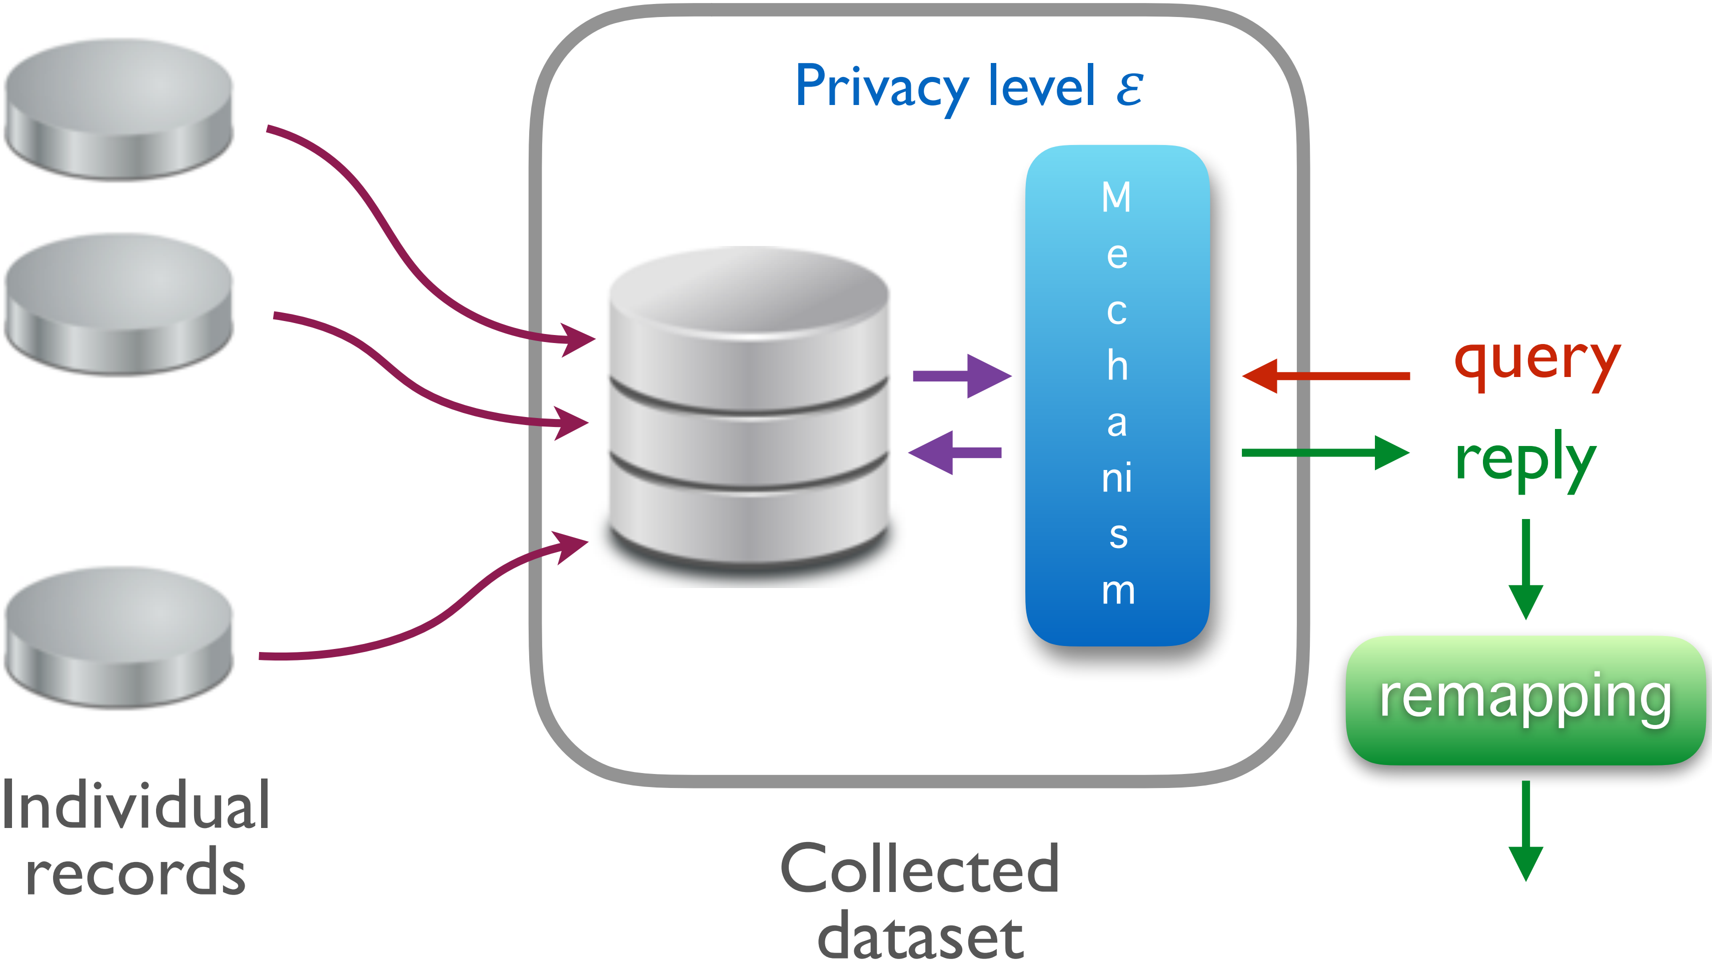

Local Differential Privacy Schema

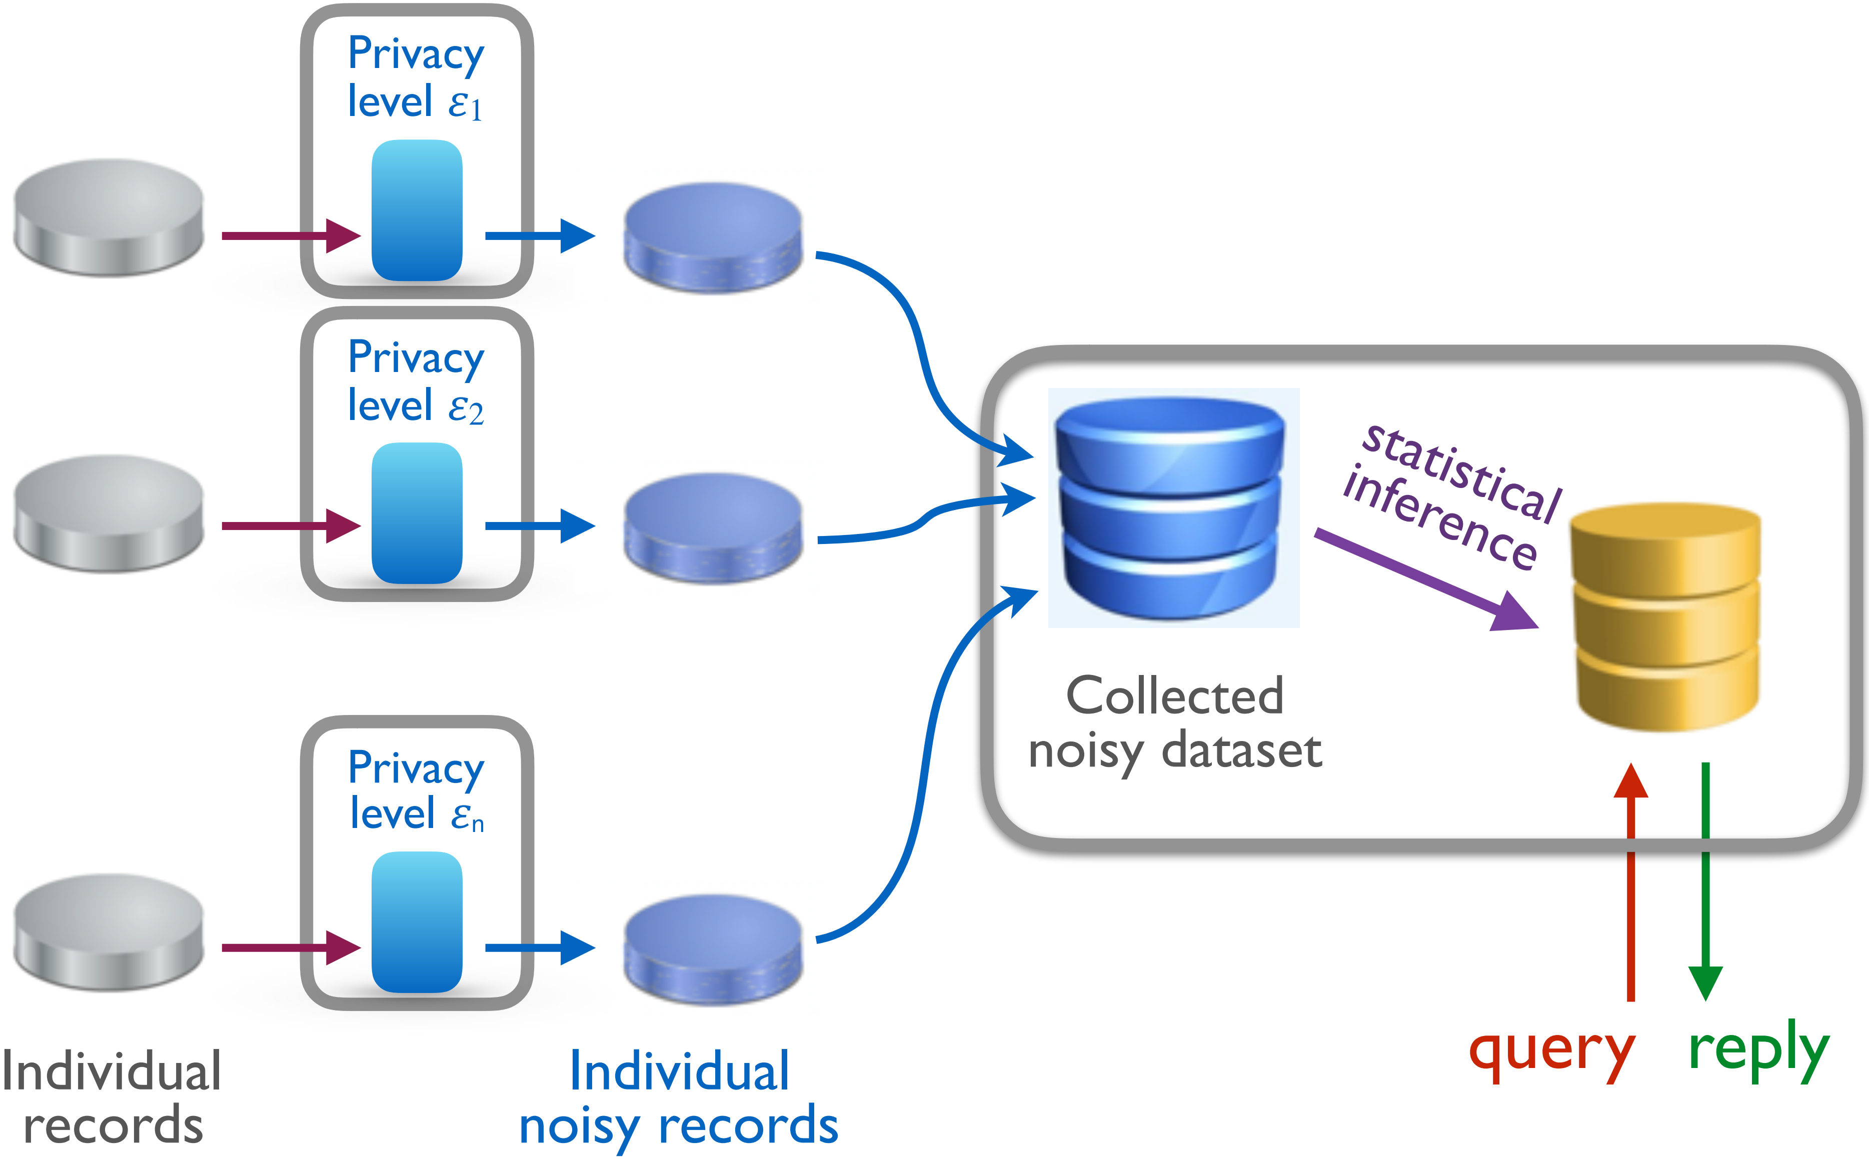

Example: embarrassing question

- We want to collect a statistic about a certain sensitive issue. For example: how many students cheat at exams.

- Students may not want to reveal that they cheated, but they may be willing to participate if they can give a noisy answer

- One possible method is the following: Each participant is posed the question "Have you cheated at the exam" and is has to reply according to the following protocol:
  - Flip a coin
  - If head, reply truthfully,
  - otherwise, flip the coin again
    - if head, reply "yes"
    - if tail, reply "no"

Example: embarrassing question

- We can represent the protocol as a stochastic matrix C, where the rows are the true answers x_i and the columns are the noisy answers y_j. The cell C_{ij} contains the conditional probability to get the noisy answer y_j if the true answer was x_i , namely $C_{ij} = P(y_j \mid x_i)$
- For example, the protocol of the previous page has the following matrix:

	yes	no
yes	¾	¼
no	¼	¾

The conditional probabilities of the matrix are determined by the bias of the coin used in the protocol

The more uniform the probabilities are, the higher is the protection of the participants' privacy

Local Differential Privacy

Definition Let $\mathcal{X}$ be the set of possible values. We say that a mechanism $\mathcal{K} : \mathcal{X} \to \mathcal{Y}$ is ε-LPD (locally differentially private) if

$$\forall x_1, x_2 \in \mathcal{X}, \forall y \in \mathcal{Y} \quad p(y|x_1) \leq e^{\varepsilon}\, p(y|x_2)$$

- For instance, the protocol in previous page is log 3 - LPD

- Note that a collection of data sanitized with LPD is DP with respect to any possible query. We leave for exercise to determine the level of privacy of the resulting DP mechanism, as a function of the level(s) of privacy of the LPD mechanism(s).

The flat mechanism

The flat mechanism is the simplest way to implement LPD. It is defined as follows:

$$p(y|x) = \begin{cases} c\,e^{\varepsilon} & \text{if } x = y \\ c & otherwise \end{cases}$$

where c is a normalization constant.

- Exercise: determine *c*

Statistical inference

- The noisy data determine a distribution $\underline{p}$ on the output domain $\mathcal{Y}$ of the LPD mechanism(s).

- From $\underline{q}$, we want to reconstruct a distribution $\underline{p}$ which should approximates as much as possible the real distribution p on the real data. The utility of the collected dataset depends on how good this approximation is.

- There are two methods: the Matrix Inversion method, and the Iterative Bayesian update.

The matrix inversion method

- Define $\underline{r}$ as the product $C^{-1}\,\underline{q}$

- $\underline{r}$ may not be a distribution because it may contain negative elements. In order to obtain a distribution (which we will take as $\underline{p}$) we can either:

- set to 0 all the negative elements, and renormalize, or

- project $\underline{r}$ on the simplex.

The Iterative Bayesian method

- Define $\underline{p}_0$ as any distribution, for instance $\underline{q}$ if they have the same domain, or the uniform distribution.

- Repeat the following procedure: Define $\underline{p}_{n+1}$ as the Bayesian update of $\underline{p}_n$ weighted on the corresponding element of $\underline{q}$

- Stop when $\underline{p}_{n+1}$ “is almost the same” as $\underline{p}_n$

- Define $\underline{p}$ as $\underline{p}_{n+1}$