Differential Privacy: From the Central Model to the Local Model and their Generalization

Catuscia Palamidessi

Inria informatiques mathématiques

Summary

- Privacy and motivation for probabilistic methods
- Privacy vs utility
- Differential privacy: central and local models
- Statistical utility
- Compositionality
- An hybrid mechanism for privacy in a distributed setting

Information age:
Data are very useful but they raise a risk for privacy

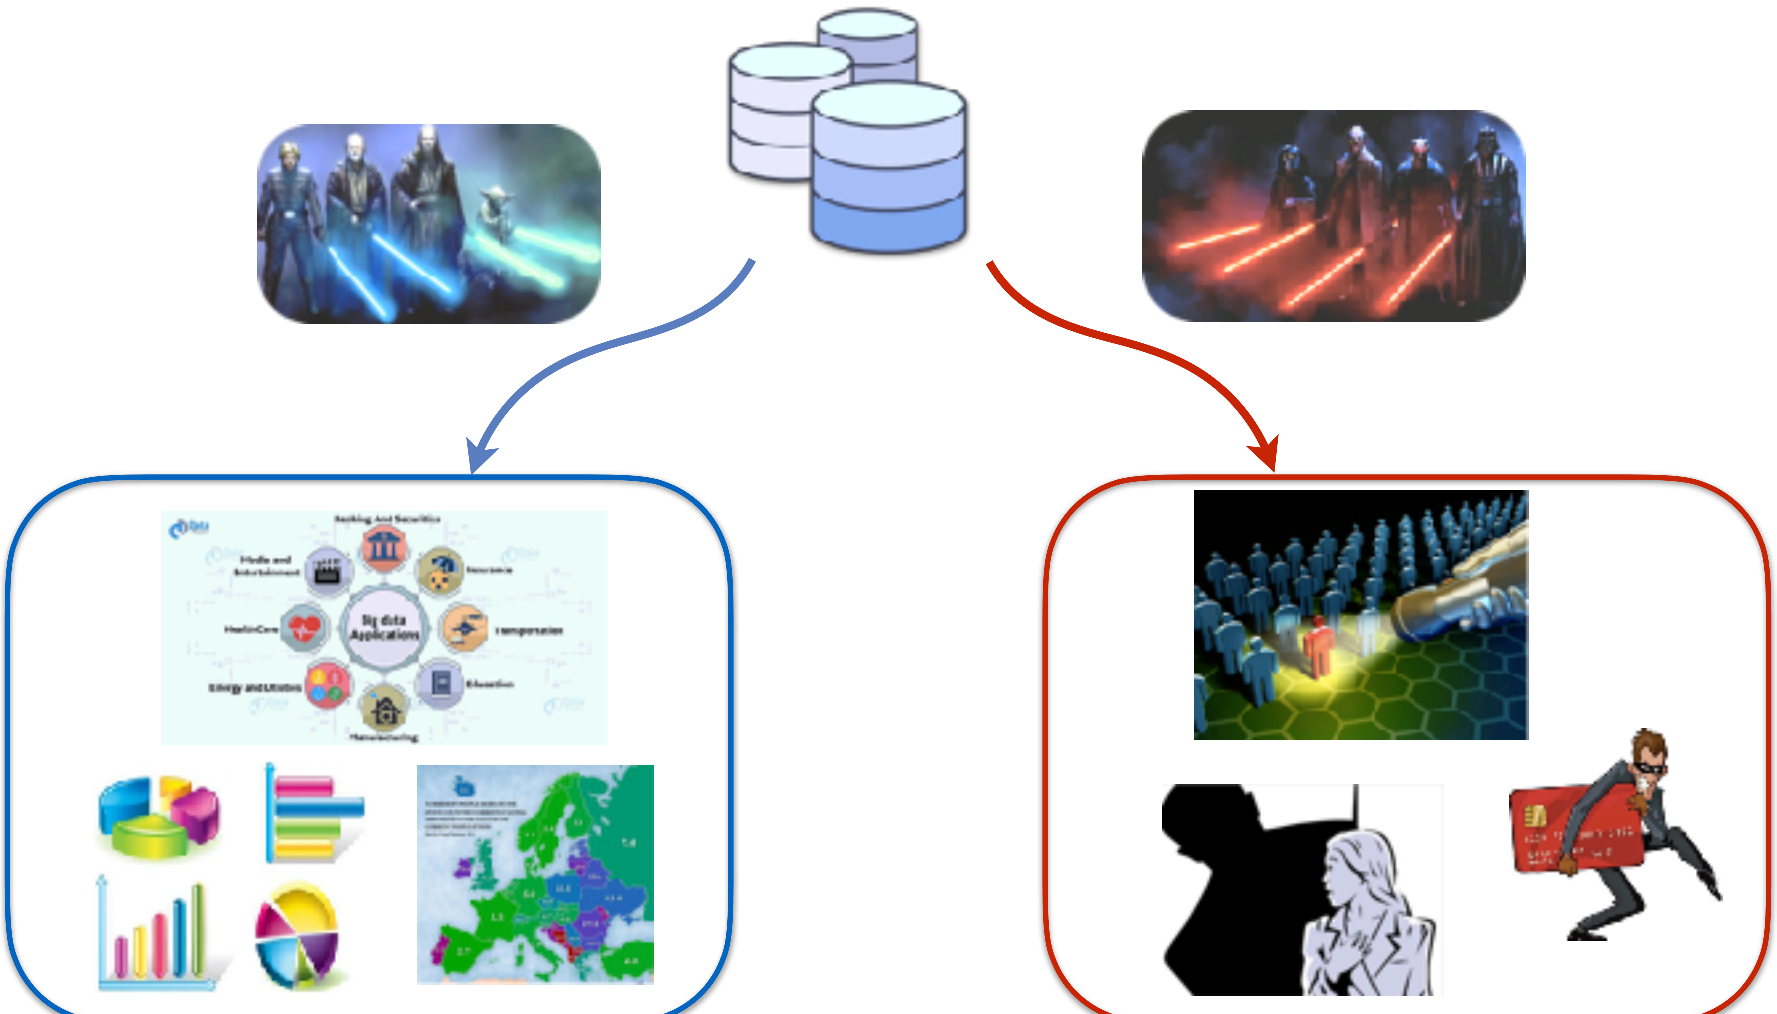

Privacy protection: Anonymization

In the past, most used technique for privacy protection was anonymization, i.e., removal of all personal identifiers: name, address, SSN, …

k-anonymity: every tuple of quasi-identifiers corresponds to at least k people

Unfortunately, **anonymization is not enough.**
Several famous attacks to anonymized datasets have shown the limitations of anonymity and k-anonymity

The Massachuttes Medical Database attack

The AOL attack

The Netflix prize attack

The Social Networks (Twitter) attack

De-anonymization attacks (I)

Robust De-anonymization of Large Sparse Datasets.
Narayanan and Shmatikov, 2008.

Showed the limitations of K-anonymity

De-anonymization of the **Netflix Prize dataset** (500,000 anonymous records of movie ratings), by linking it with the **IMDB dataset**.

They demonstrated that an adversary who knows just a few preferences about an individual subscriber can identify his record in the anonymous dataset.

De-anonymization attacks (II)

De-anonymizing Social Networks.
Narayanan and Shmatikov, 2009.

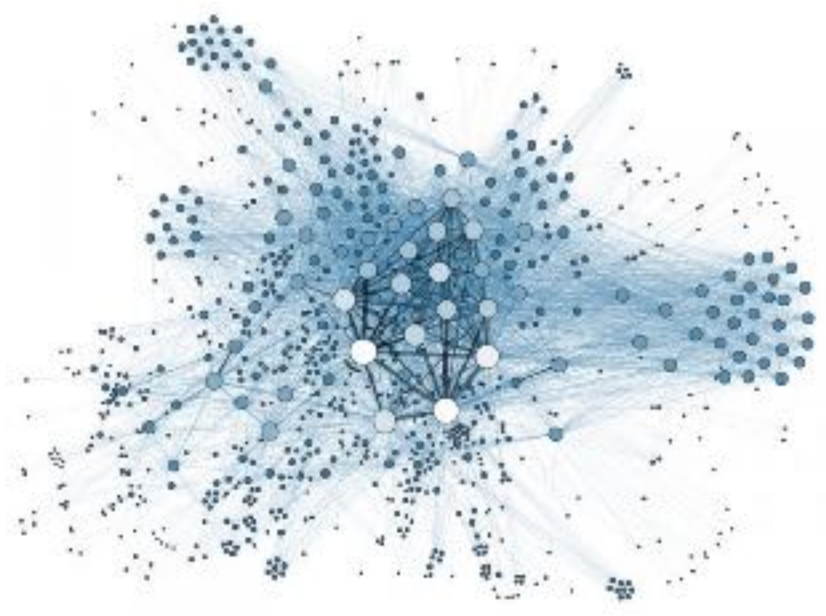

By using only the network topology, they were able to show that 33% of the users who had accounts on both **Twitter** and **Flickr** could be re-identified in the anonymous Twitter graph with only a 12% error rate.

General problem with deterministic methods

Deterministic methods for privacy are not robust wrt composition.

This is true even if the microdata are not accessible directly, and information can only be obtained by querying the dataset.

Deterministic methods are not robust wrt composition. Example

- A medical database D1 containing correlation between a certain disease and age.
- Query: "what is the minimal age of a person with the disease"

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

D1 is 2**-anonymous with respect to the age**. Namely, every possible answer partitions the records in groups of at least 2 elements

Alice	Bob
Carl	Don
Ellie	Frank

- A medical database D2 containing correlation between the disease and weight.
- Query: “what is the minimal weight of a person with the disease”

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Also D2 is 2**-anonymous wrt the weight**

Alice	Bob
Carl	Don
Ellie	Frank

k-anonymity is not compositional

Combine the two queries:

minimal weight and the minimal age of a person with the disease

Answers: 40, 100. **Unique!**

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

This is a general problem of **Deterministic approaches:** They are based on the principle that one observable corresponds to many possible values of the secret (group anonymity)

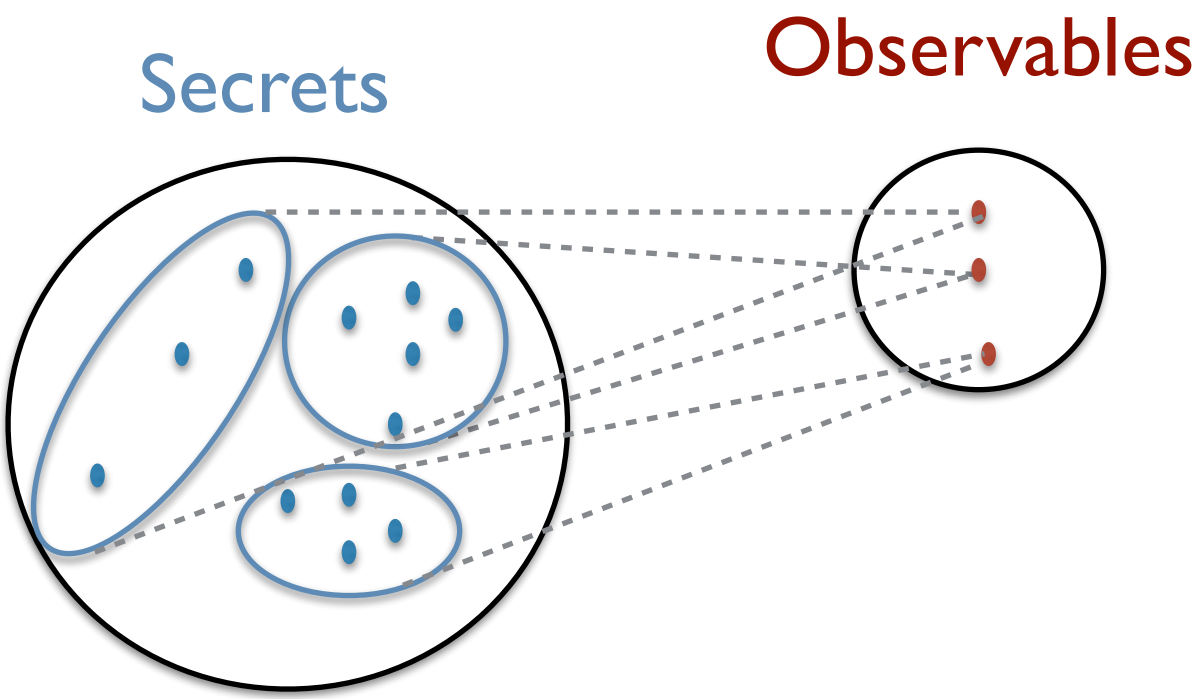

Problem of the deterministic approaches: the combination of observations determines smaller and smaller intersections on the domain of the secrets, and eventually result in singletones

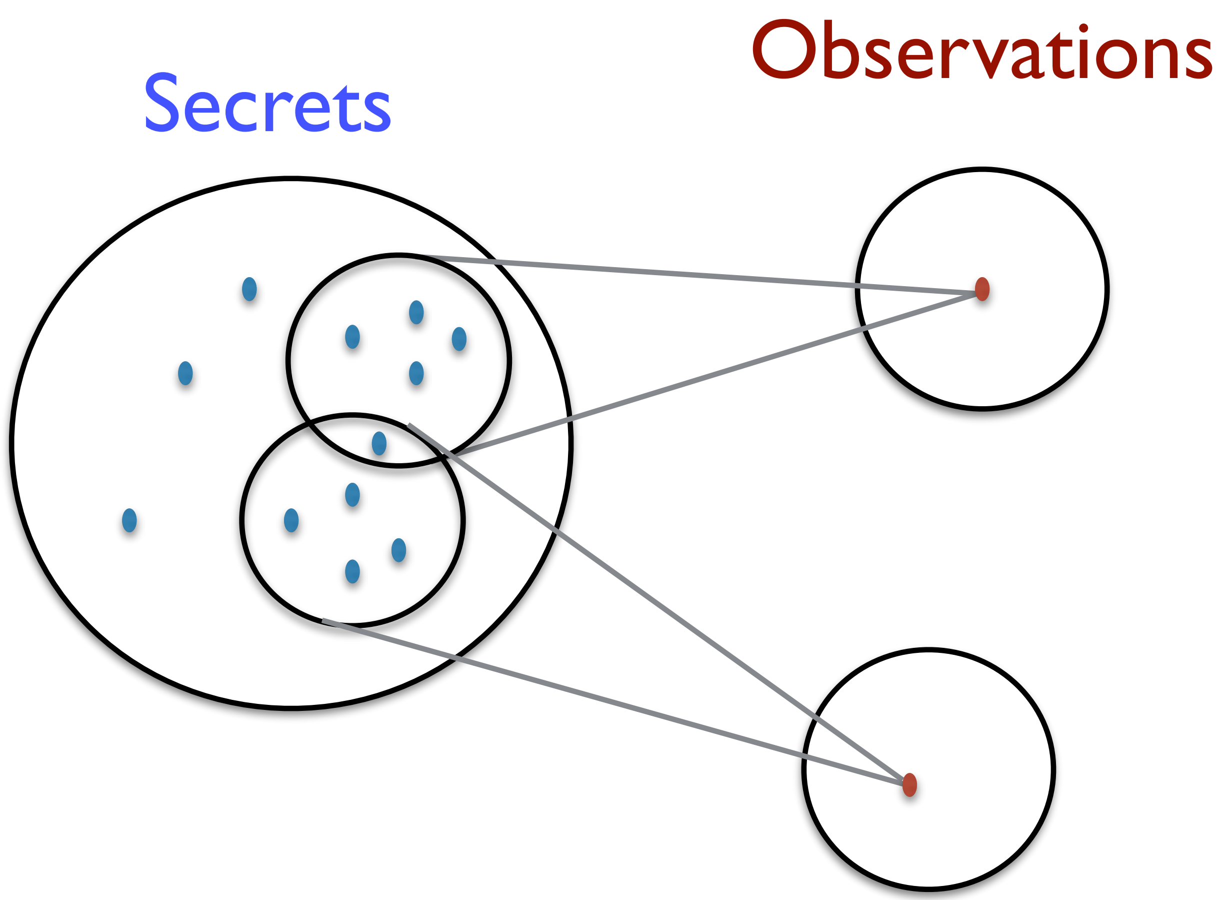

Problem of the deterministic approaches: the combination of observations determines smaller and smaller intersections on the domain of the secrets, and eventually result in singletones

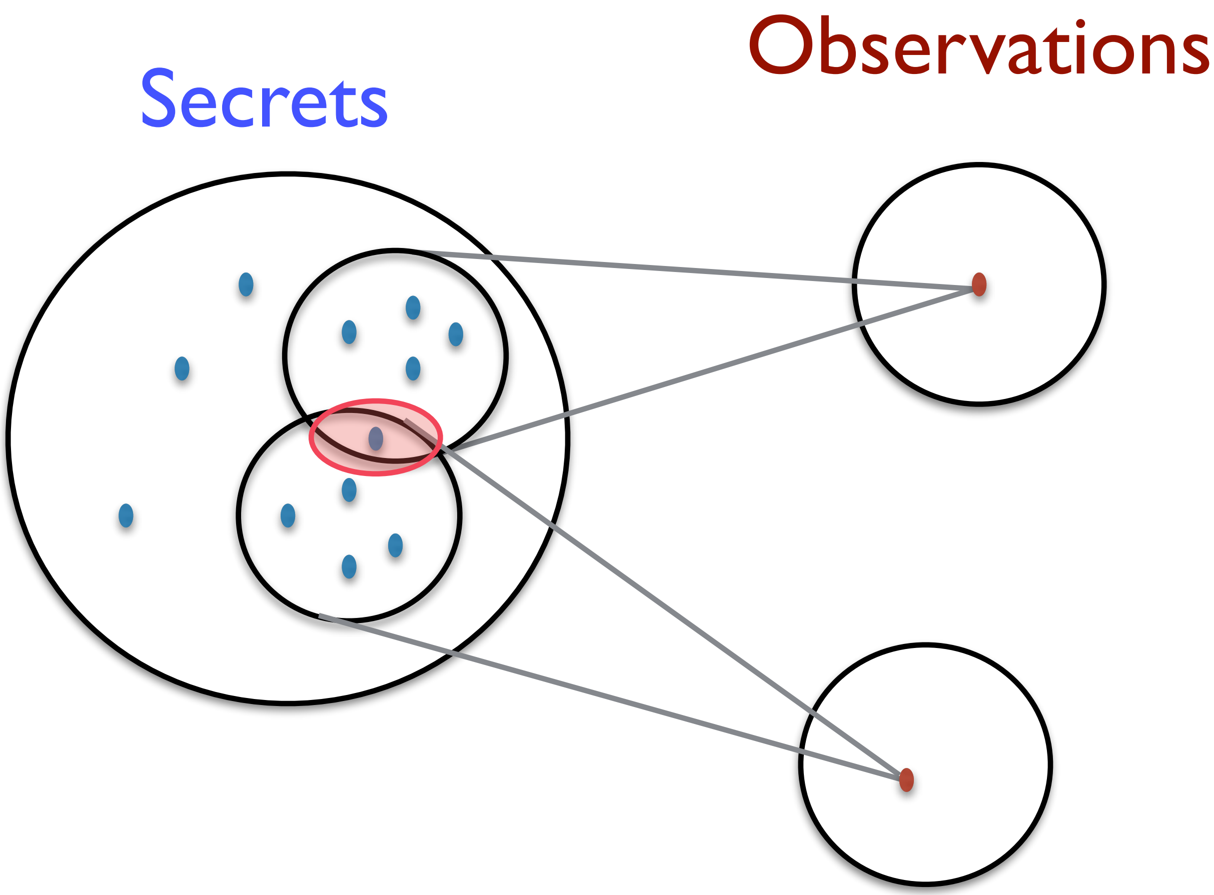

Too bad!!! What can we do?

Solution: use probabilistic method

Randomized approach for statistical databases

Introduce some probabilistic noise on the answer to obfuscate the link with any particular individual

Noisy answers

minimal age:

40 with probability 1/2

30 with probability 1/4

50 with probability 1/4

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy answers

minimal weight:

100 with prob. 4/7

90 with prob. 2/7

60 with prob. 1/7

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy answers

Even if he combines the answers, the adversary cannot tell for sure whether a certain person has the disease

name	weight	disease
Alice	60	no
Bob	90	no
Carl	90	no
Don	100	yes
Ellie	60	no
Frank	100	yes

name	age	disease
Alice	30	no
Bob	30	no
Carl	40	no
Don	40	yes
Ellie	50	no
Frank	50	yes

Alice	Bob
Carl	Don
Ellie	Frank

Noisy mechanisms

- The mechanisms reports an approximate answer, typically generated randomly on the basis of the true answer and of some probability distribution. This is the basic idea of differential privacy
- The probability distribution must be chosen carefully, in order to not destroy the utility of the answer
- A good mechanism should provide a good trade-off between privacy and utility. Note that, for the same level of privacy, different mechanisms may provide different levels of utility.

Summary

- Privacy vs utility
- Differential privacy: central and local models
- Statistical utility
- Compositionality
- An hybrid mechanism for privacy in a distributed setting

Utility

Various kinds of utility:

- Quality of service
- Precise statistical analyses
- Accuracy (machine learning)

Privacy, QoS, statistical estimation, accuracy are interrelated: The user often releases his data in exchange of a service, but it should not pose a threat to his privacy. In turn, the service provider offers the service because it's interested in collecting the user's data, which are often used to derive statistics or learning models.

It is important to find mechanisms that optimize the trade-off between utility and privacy

Utility

Various kinds of utility:

- Quality of service
- Precise statistical analyses
- Accuracy (machine learning)

Privacy, QoS, statistical analysis are interrelated: The user often releases his data in exchange of a service, but it should not pose a threat to his privacy. In turn, the service provider offers the service because it's interested in collecting the user's data, which are often used to derive statistics or learning models.

It is important to find mechanisms that optimize the trade-off between utility and privacy

Summary

- Privacy vs utility
- Differential privacy: central and local models
- Statistical utility
- Compositionality
- An hybrid mechanism for privacy in a distributed setting

Standard Differential Privacy (aka central model)

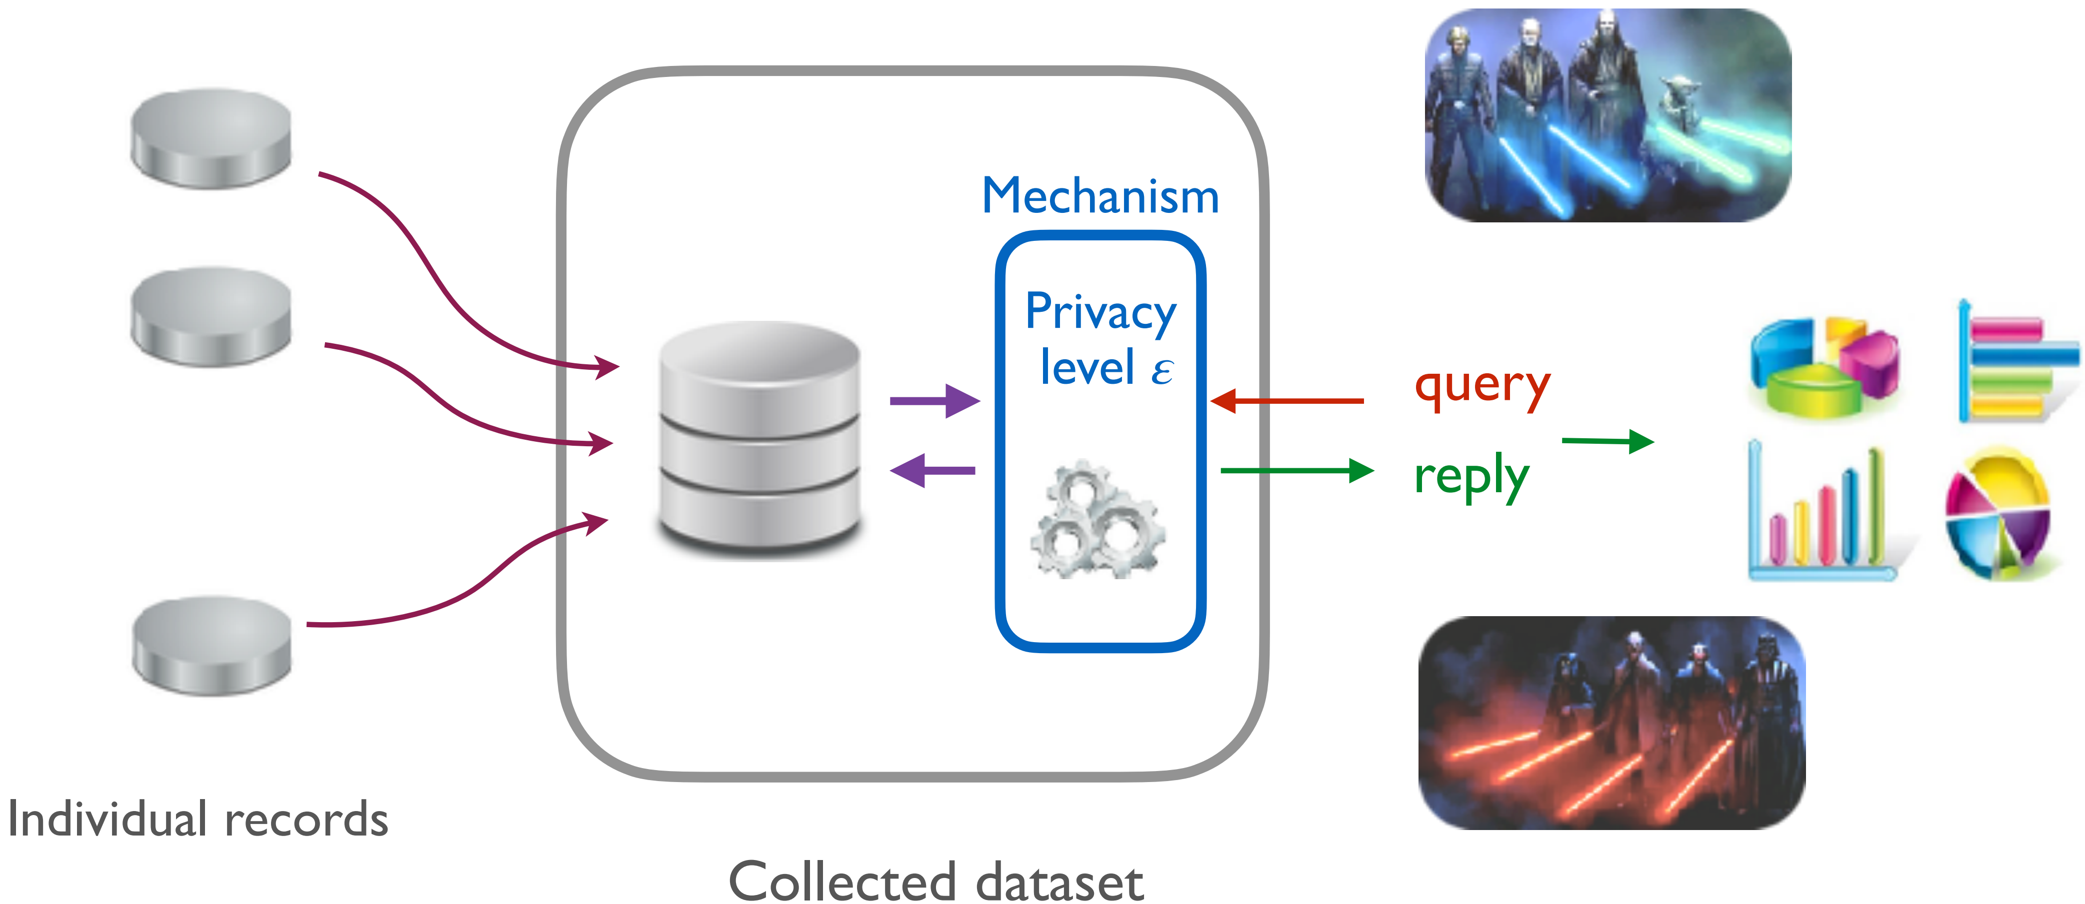

Privacy by randomization

Differential Privacy [Dwork et al., 2006]

A mechanism $\mathcal{K}$ (for a certain query) is ε-differentially private if for every pair of *adjacent* datasets x and x' and every possible answer y

$$P[\mathcal{K}(x) = y] \leq e^{\varepsilon}\ P[\mathcal{K}(x') = y]$$

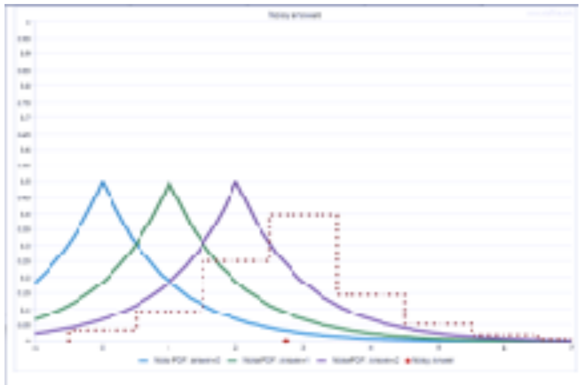

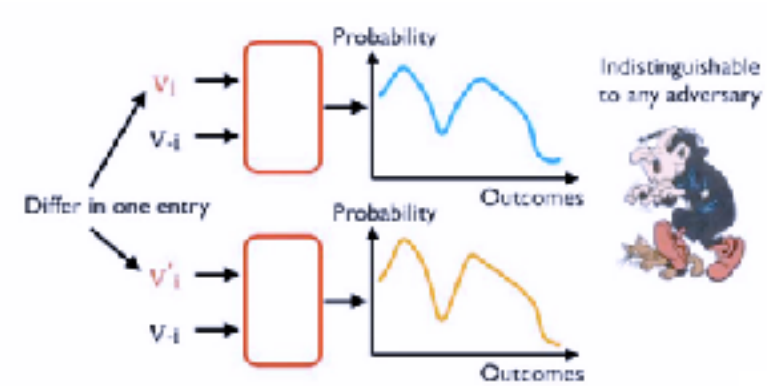

- **Compositionality**: the combination of two mechanisms which are e_1 and e_2 differentially private is $e_1 + e_2$ differentially private
- **Independent** from side knowledge

Typical DP mechanisms: Laplace, Geometric

Problem with Central Differential Privacy

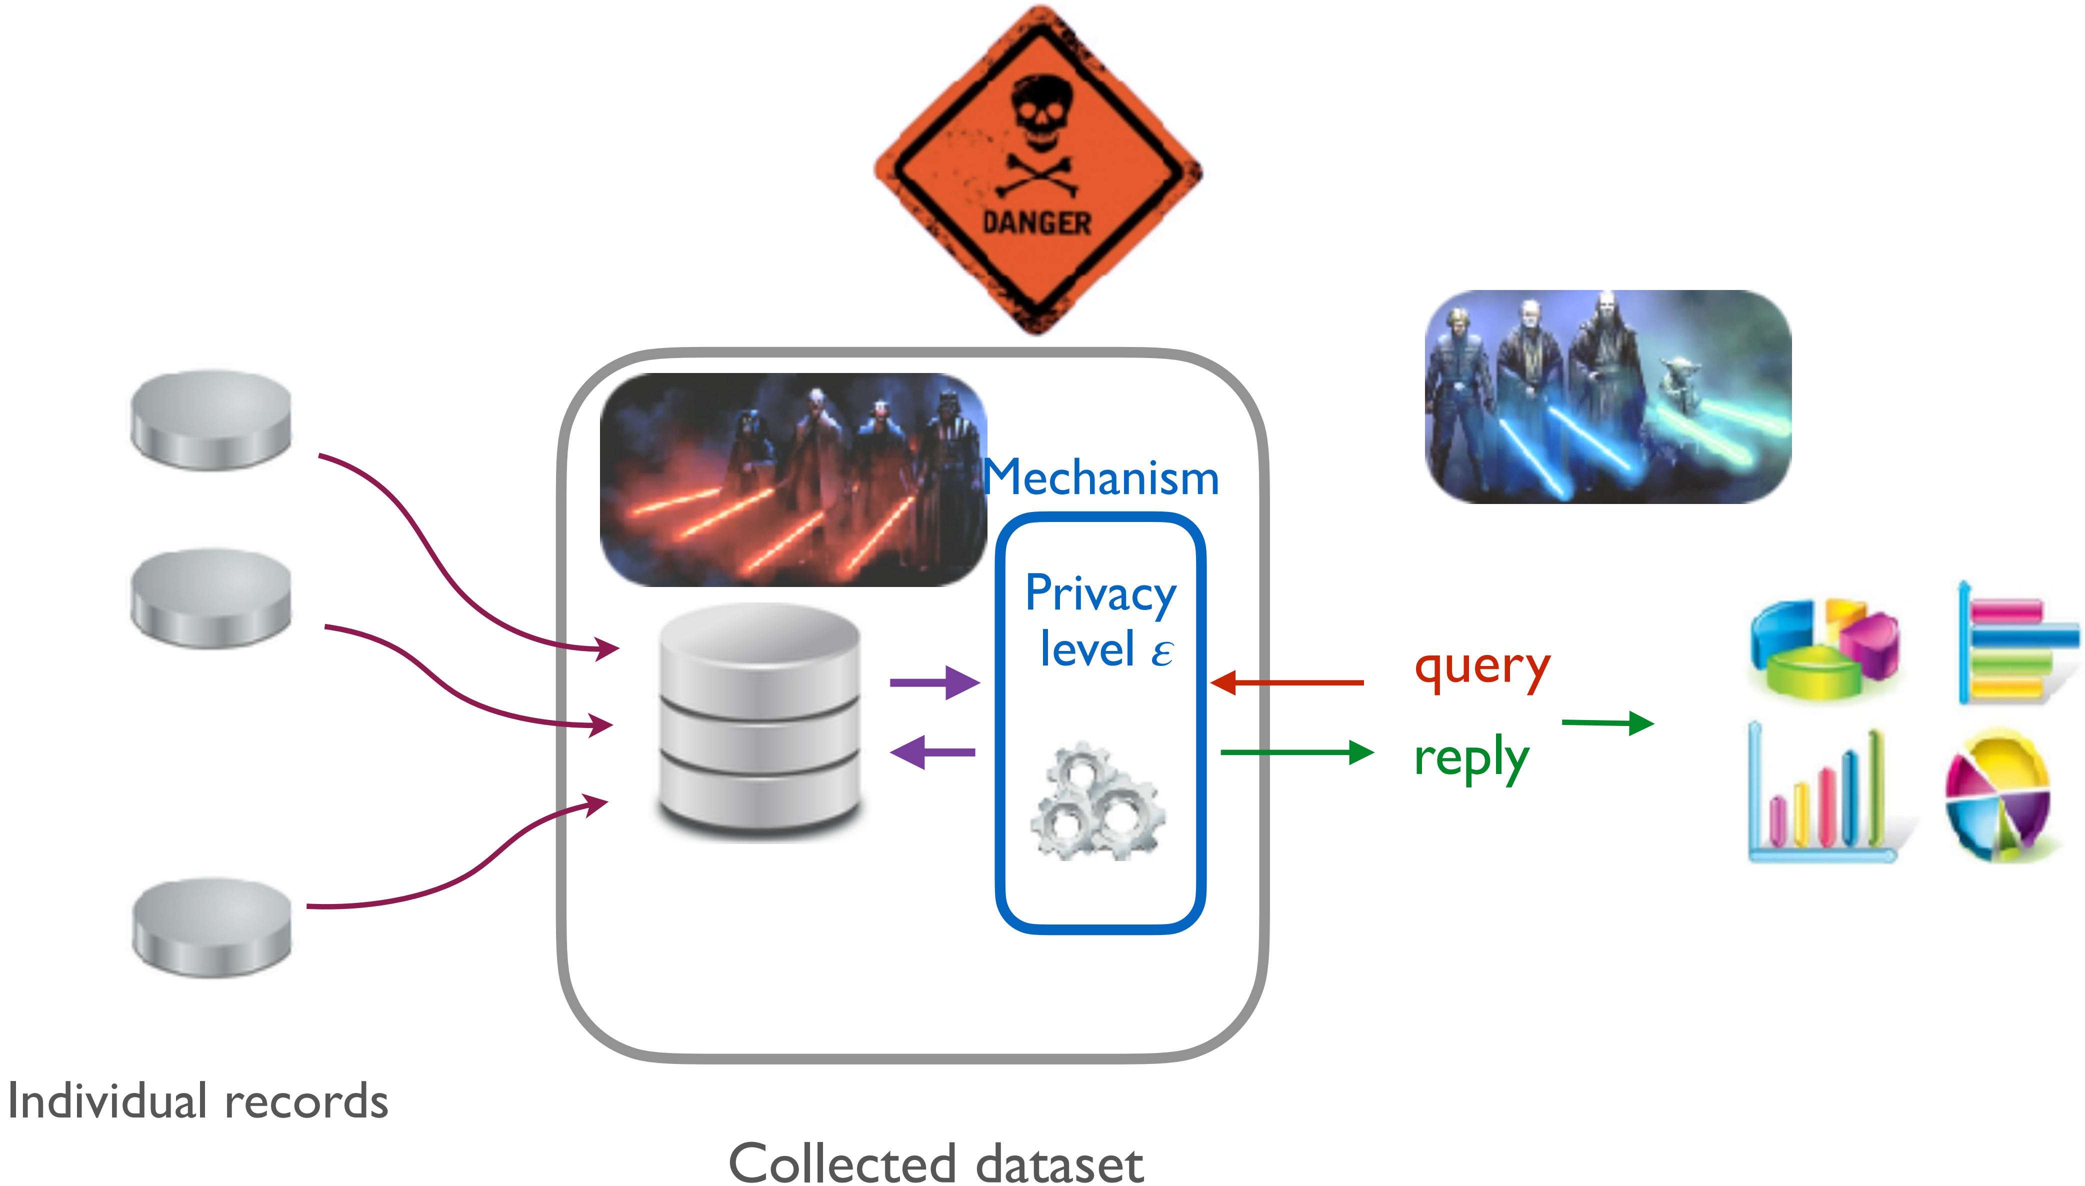

Local Differential Privacy

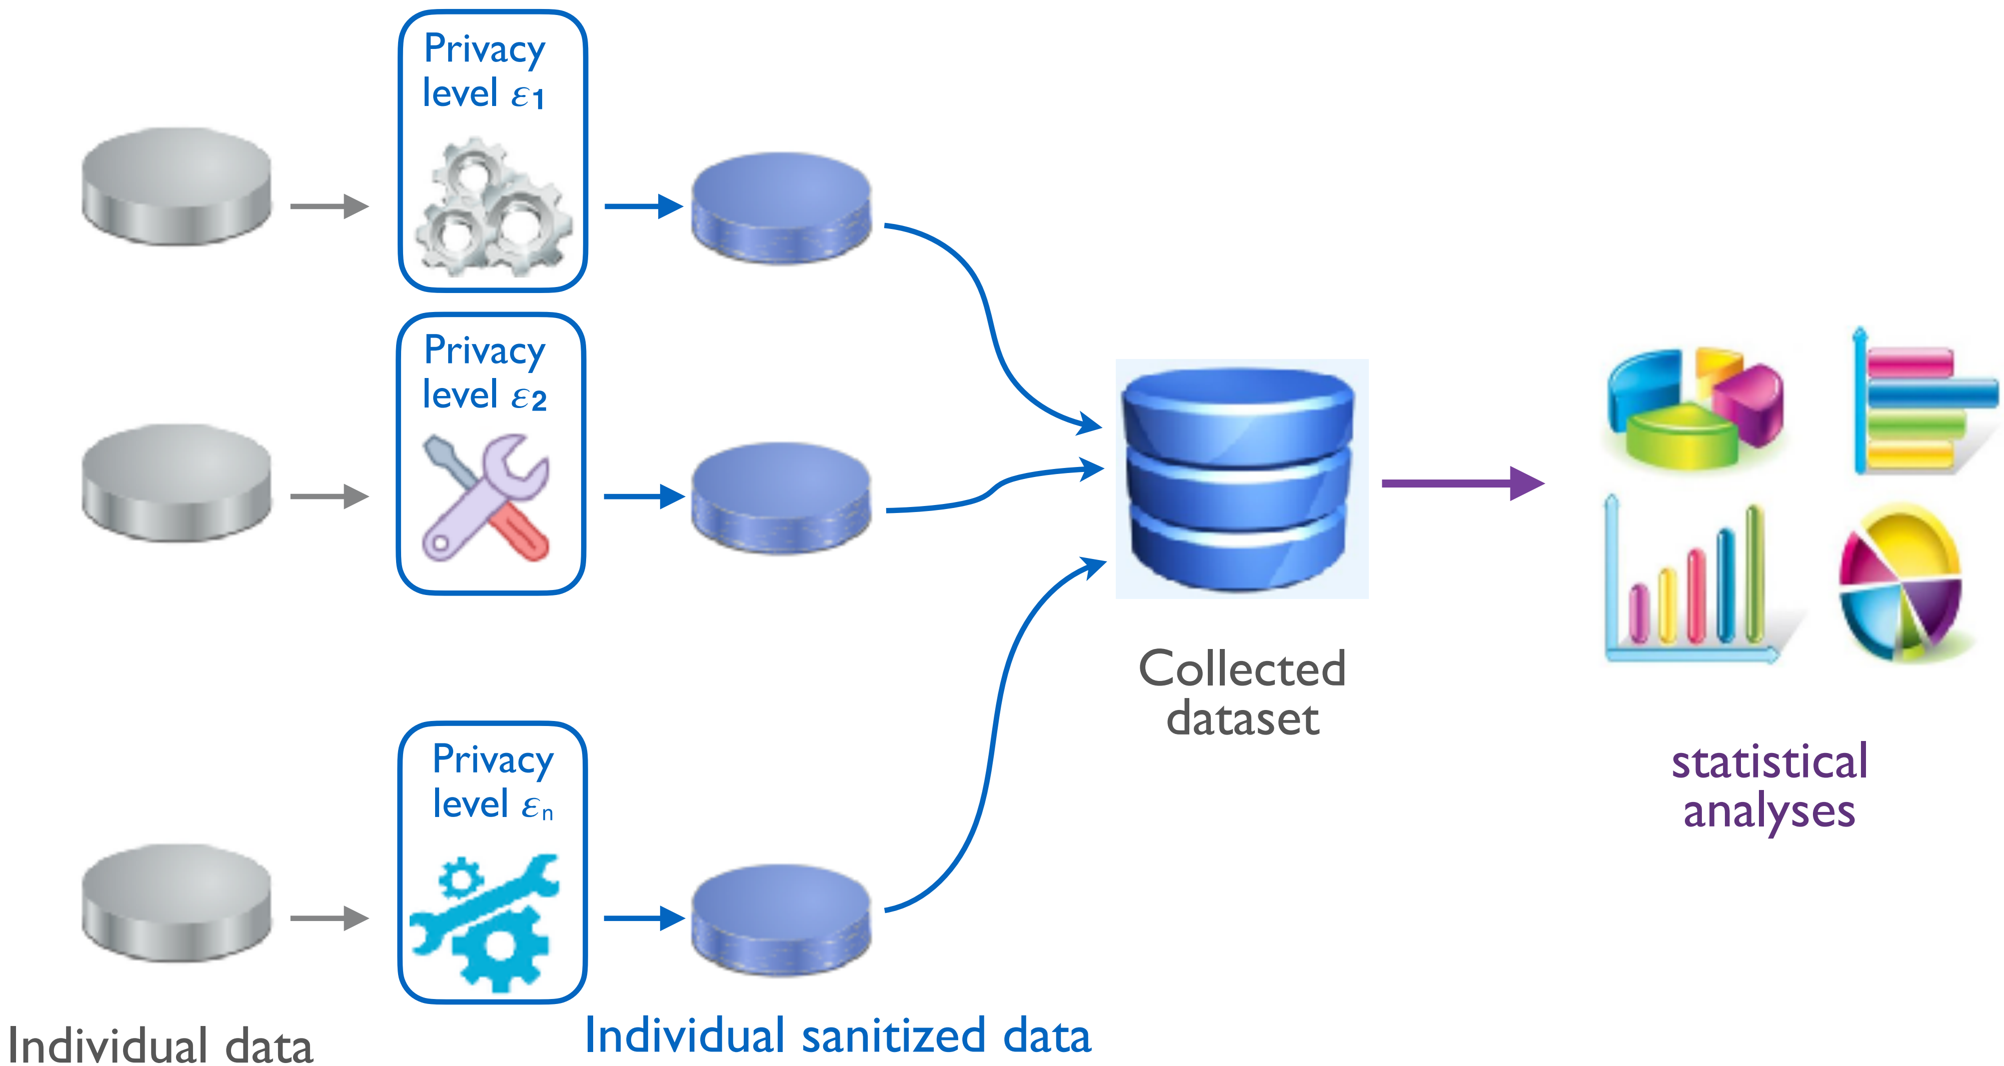

Local Differential Privacy

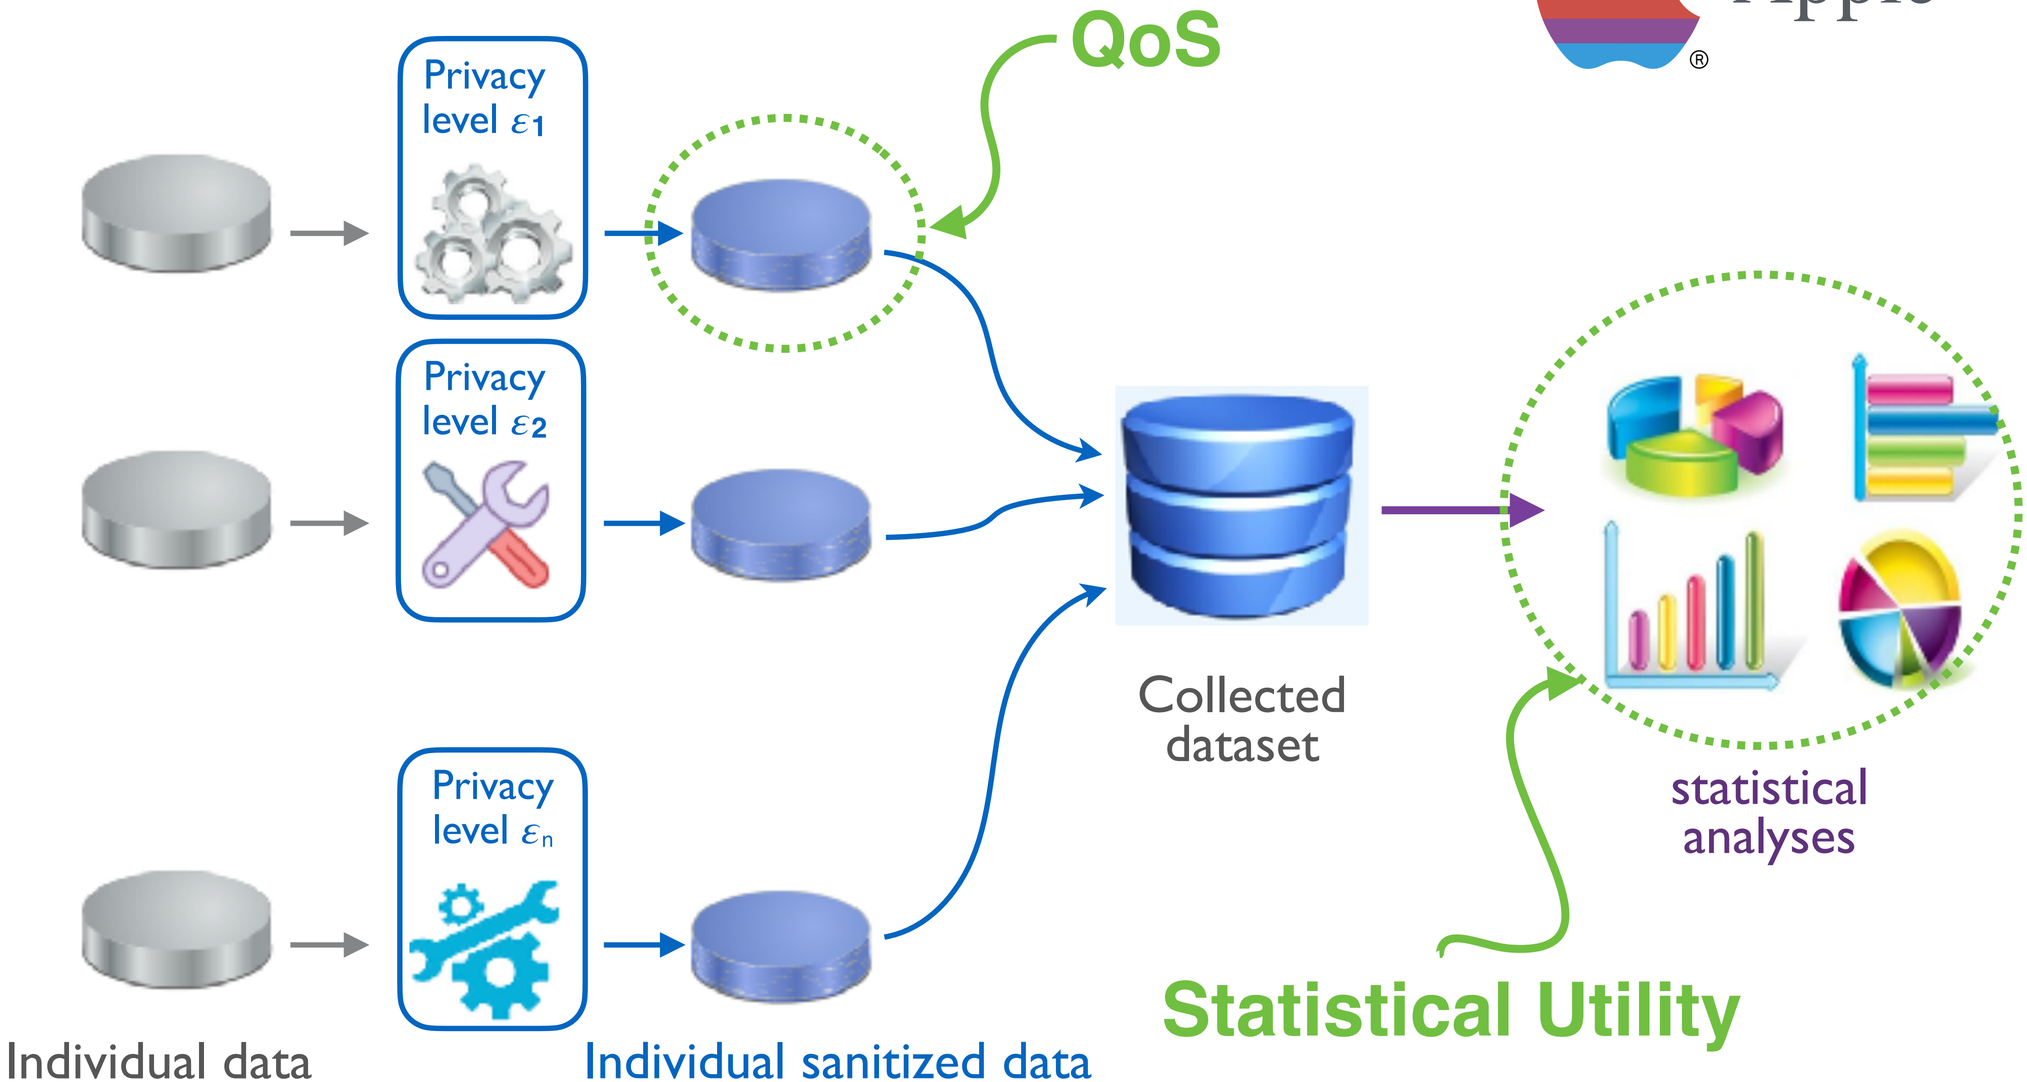

LDP versus Central DP

The trade-off utility-privacy is usually much worse in the local model than in the central model, especially when the collection of data is small. The is why so far is mainly used by large companies like Google and Apple

Local Differential Privacy

[Jordan &Wainwright '13]

Definition Let $\mathcal{X}$ be a set of possible values and $\mathcal{Y}$ the set of noisy values. A mechanism $\mathcal{K}$ is ε-locally differentially private (ε-LDP) if for all $x_1, x_2 \in \mathcal{X}$ and for all $y \in \mathcal{Y}$

$$P[\mathcal{K}(x) = y] \leq e^{\varepsilon}\ P[\mathcal{K}(x') = y]$$

or equivalently, using the conditional probability notation:

$$p(y \mid x) \leq e^{\varepsilon}\ p(y \mid x')$$

Example: Randomized Response protocol

(log 3)-LPD

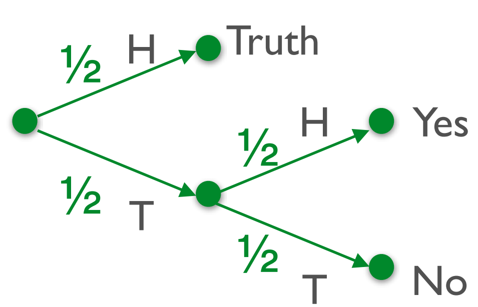

y

		yes	no
x	**yes**	$\frac{3}{4}$	$\frac{1}{4}$
	no	$\frac{1}{4}$	$\frac{3}{4}$

Mechanism's stochastic matrix

The k-RR mechanism (general RR for domains of size k)

[Kairouz et al, '16]

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

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

where c is a normalization constant.

namely $c = \dfrac{1}{k - 1 + e^{\varepsilon}}$ where k is the size of the domain

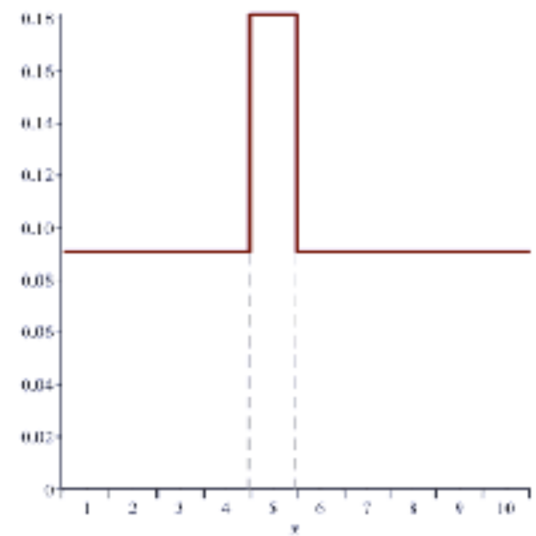

Privacy Properties:

- Compositionality
- Independence from the side knowledge of the adversary

Utility :

- Statistical Utility : ✓
- QoS : ✗

Our approach to LDP

d-privacy

[Chatzikokolakis & Palamidessi PETS'13]

d-privacy: a generalization of DP and LDP

d-privacy

On a generic domain $\mathcal{X}$ provided with a distance d:

$$\forall x, x' \in \mathcal{X}, \forall z \quad \frac{p(z \mid x)}{p(z \mid x')} \leq e^{\varepsilon\, d(x,x')}$$

generalizes

Differential Privacy

- x, x' are databases
- *d* is the Hamming distance

Local Differential Privacy

- *d* is the discrete distance

Properties

- Like LDP, it can be applied at the user side
- Like DP and LDP, it is compositional

QoS: we extensively studied *d*-privacy in the case of Location Privacy for Location Based Services

- Example of LBS: find the restaurants near the user
- Revealing the exact location may be dangerous: profiling, inference of sensitive information, etc.
- Revealing an approximate location is usually ok
- QoS: decreases with the expected distance between the real location and the noisy one.

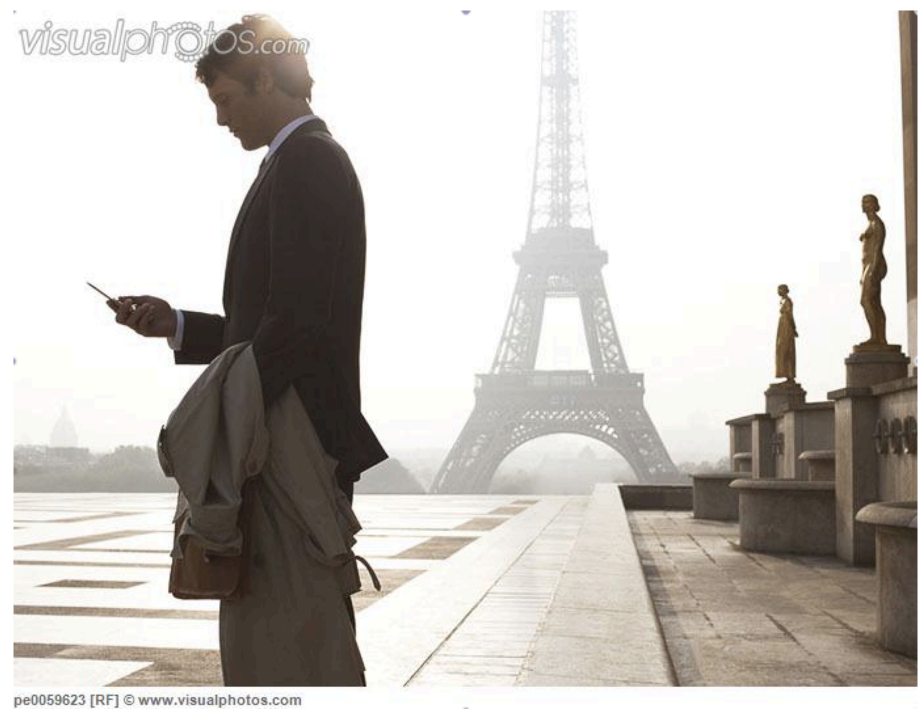

Location privacy: geo-indistinguishability

d : the Euclidean distance

x : the exact location

z : the reported location

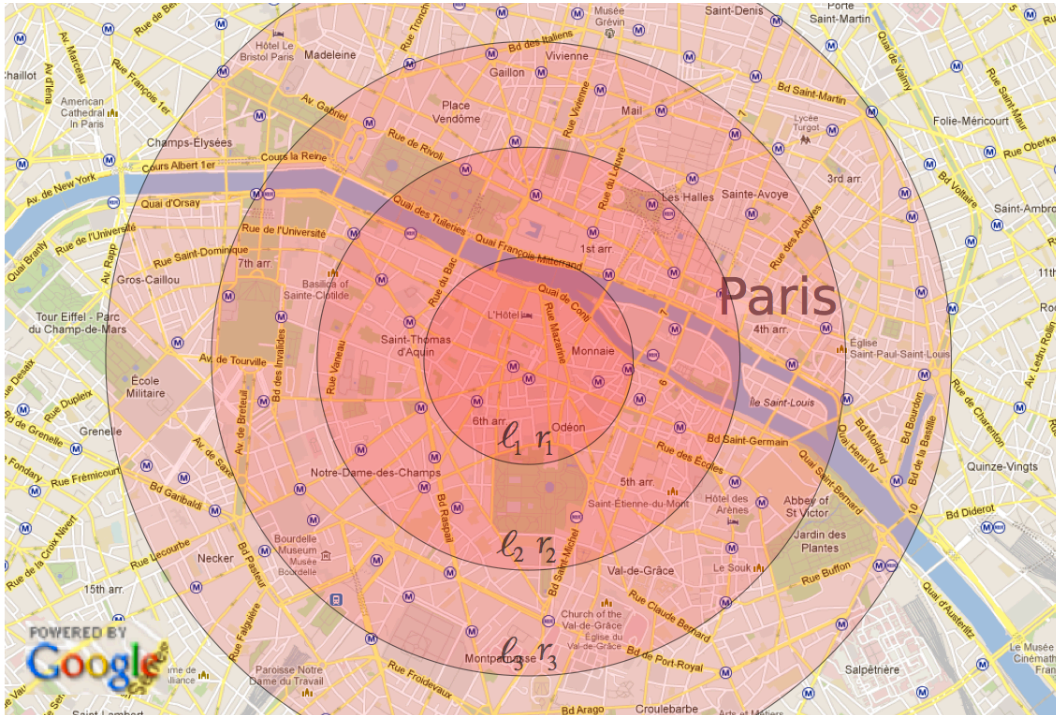

$d-$privacy

$$\frac{p(z|x)}{p(z|x')} \leq e^{\epsilon r}$$

where r is the distance between x and x'

We call this property **geo-indistinguishability.** Like DP, it is:

1) independent from the prior,
2) compositional

Typical *d*-private mechanisms:

Extended Laplace and Extended Geometric

Example: Location privacy

- Domain: points on a plane
- Distance: Euclidean

$$dp_x(z) \;=\; \frac{\epsilon^2}{2\pi}\, e^{\epsilon\, d(x,z)}$$

Efficient method to draw noisy locations based on polar coordinates

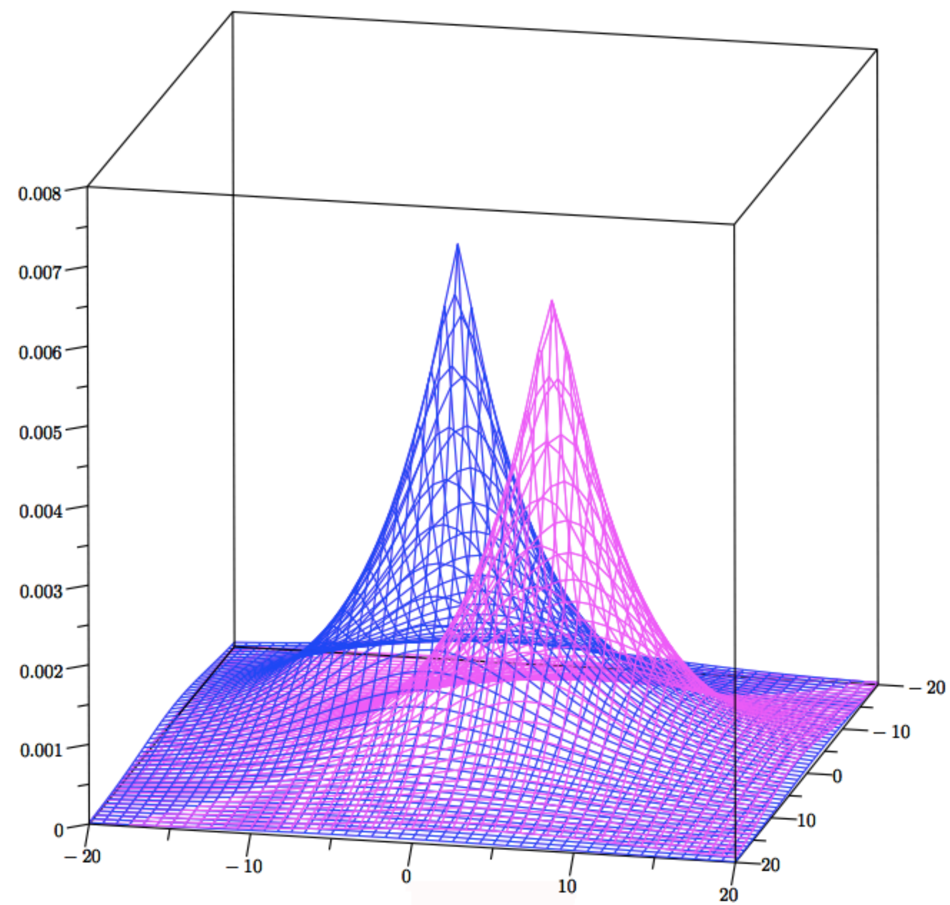

Tool: “Location Guard”

http://www.lix.polytechnique.fr/~kostas/software.html

Extension for Firefox, Chrome, and Opera. It has been released about two yeas ago, and nowadays it has about 60,000 active users.

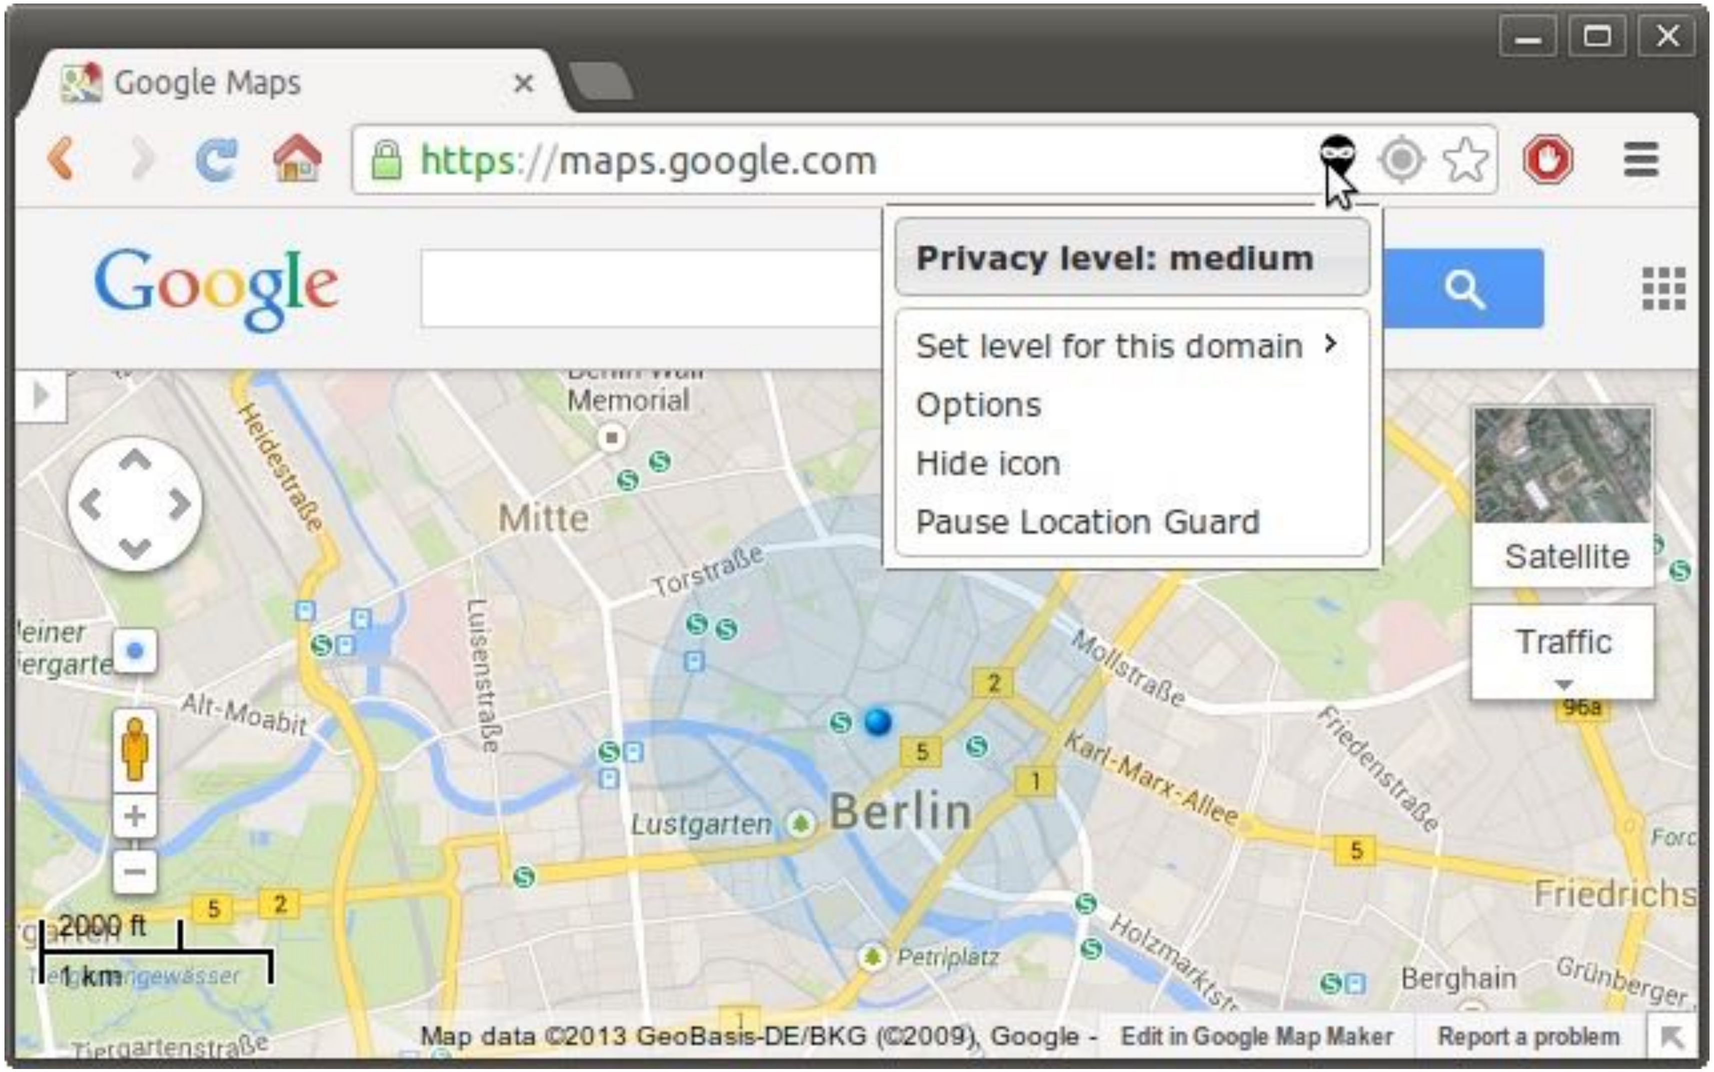

How it works

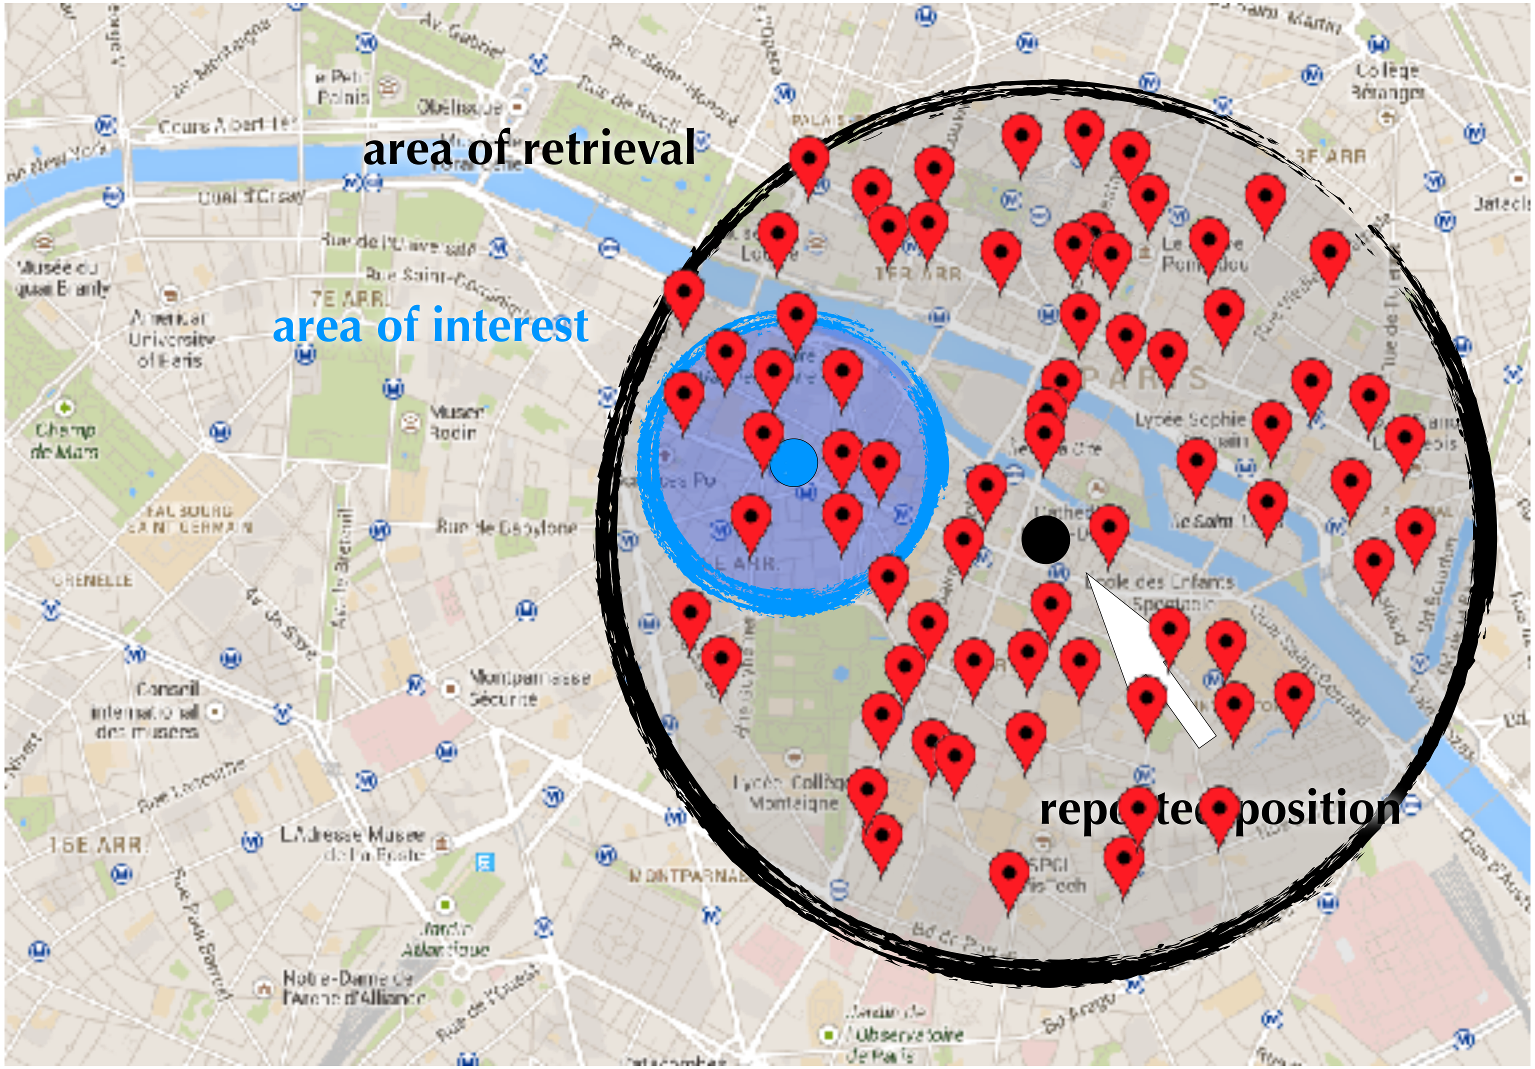

Summary

- Privacy vs utility
- Differential privacy: central and local models
- Statistical utility
- Compositionality
- An hybrid mechanism for privacy in a distributed setting

Statistical Utility:

Estimating the original distribution

i.e., the distribution from which the true data are sampled

Estimation mechanism

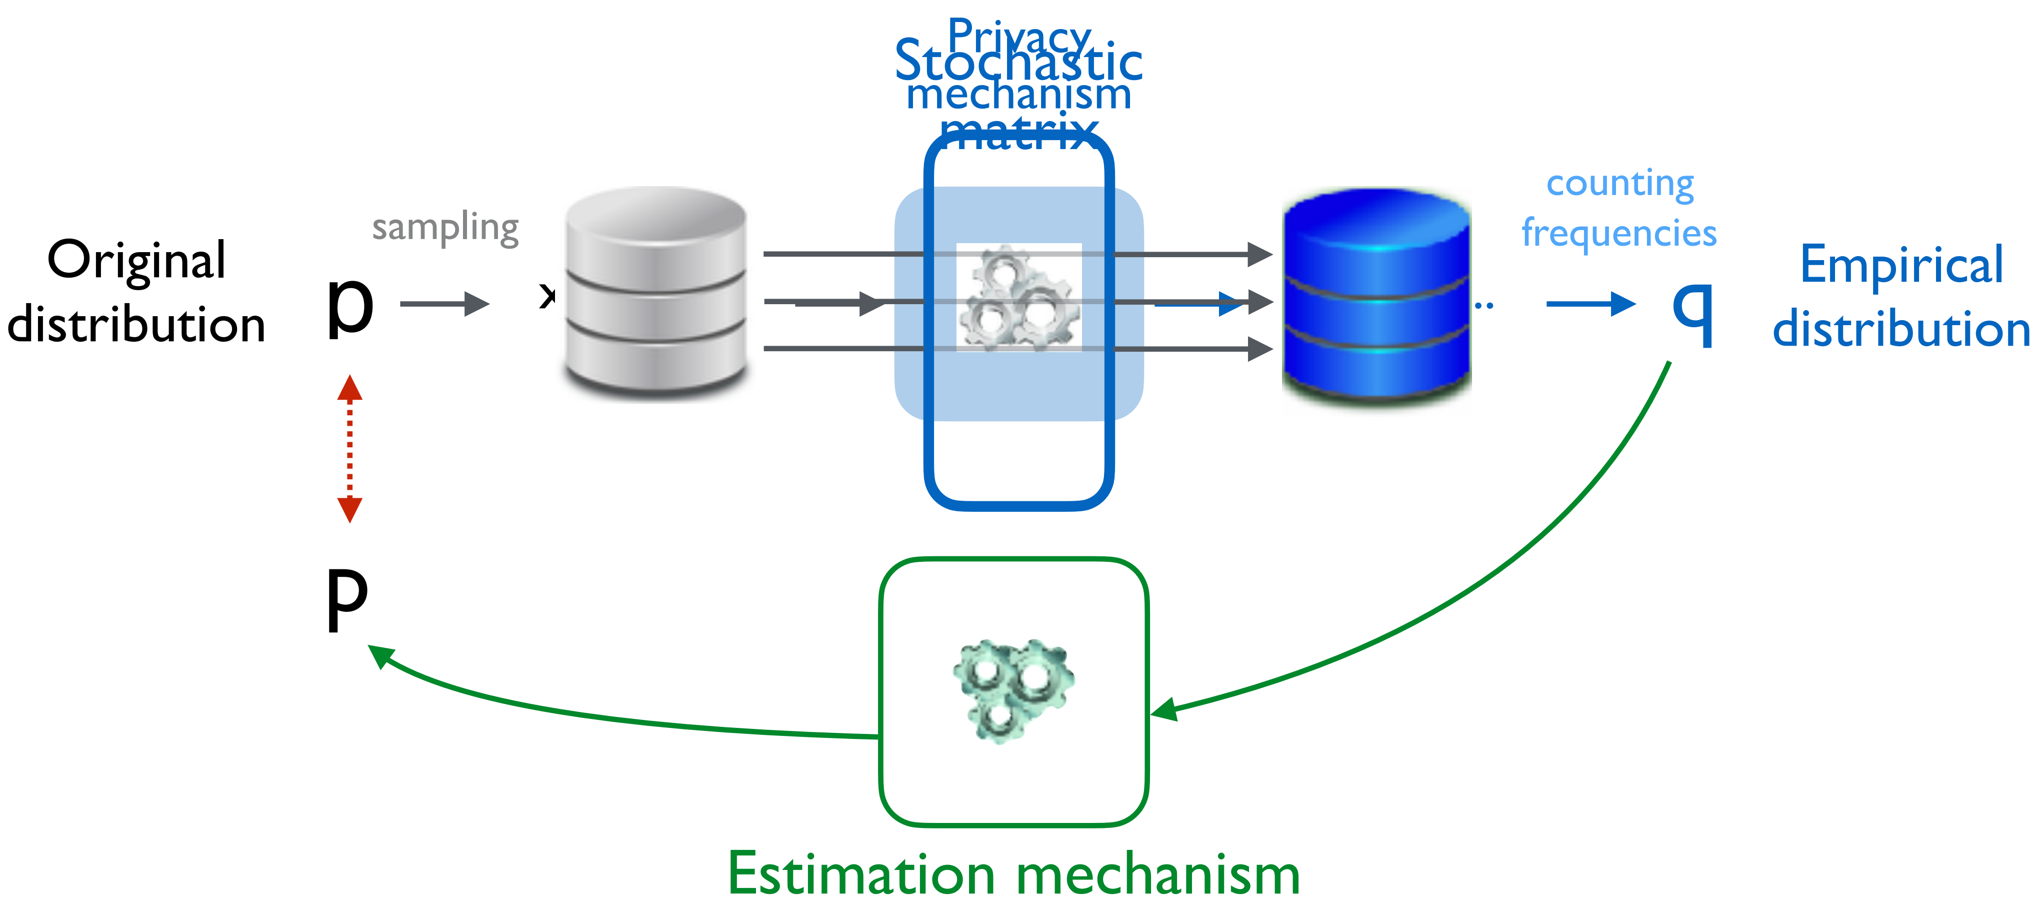

Estimation method

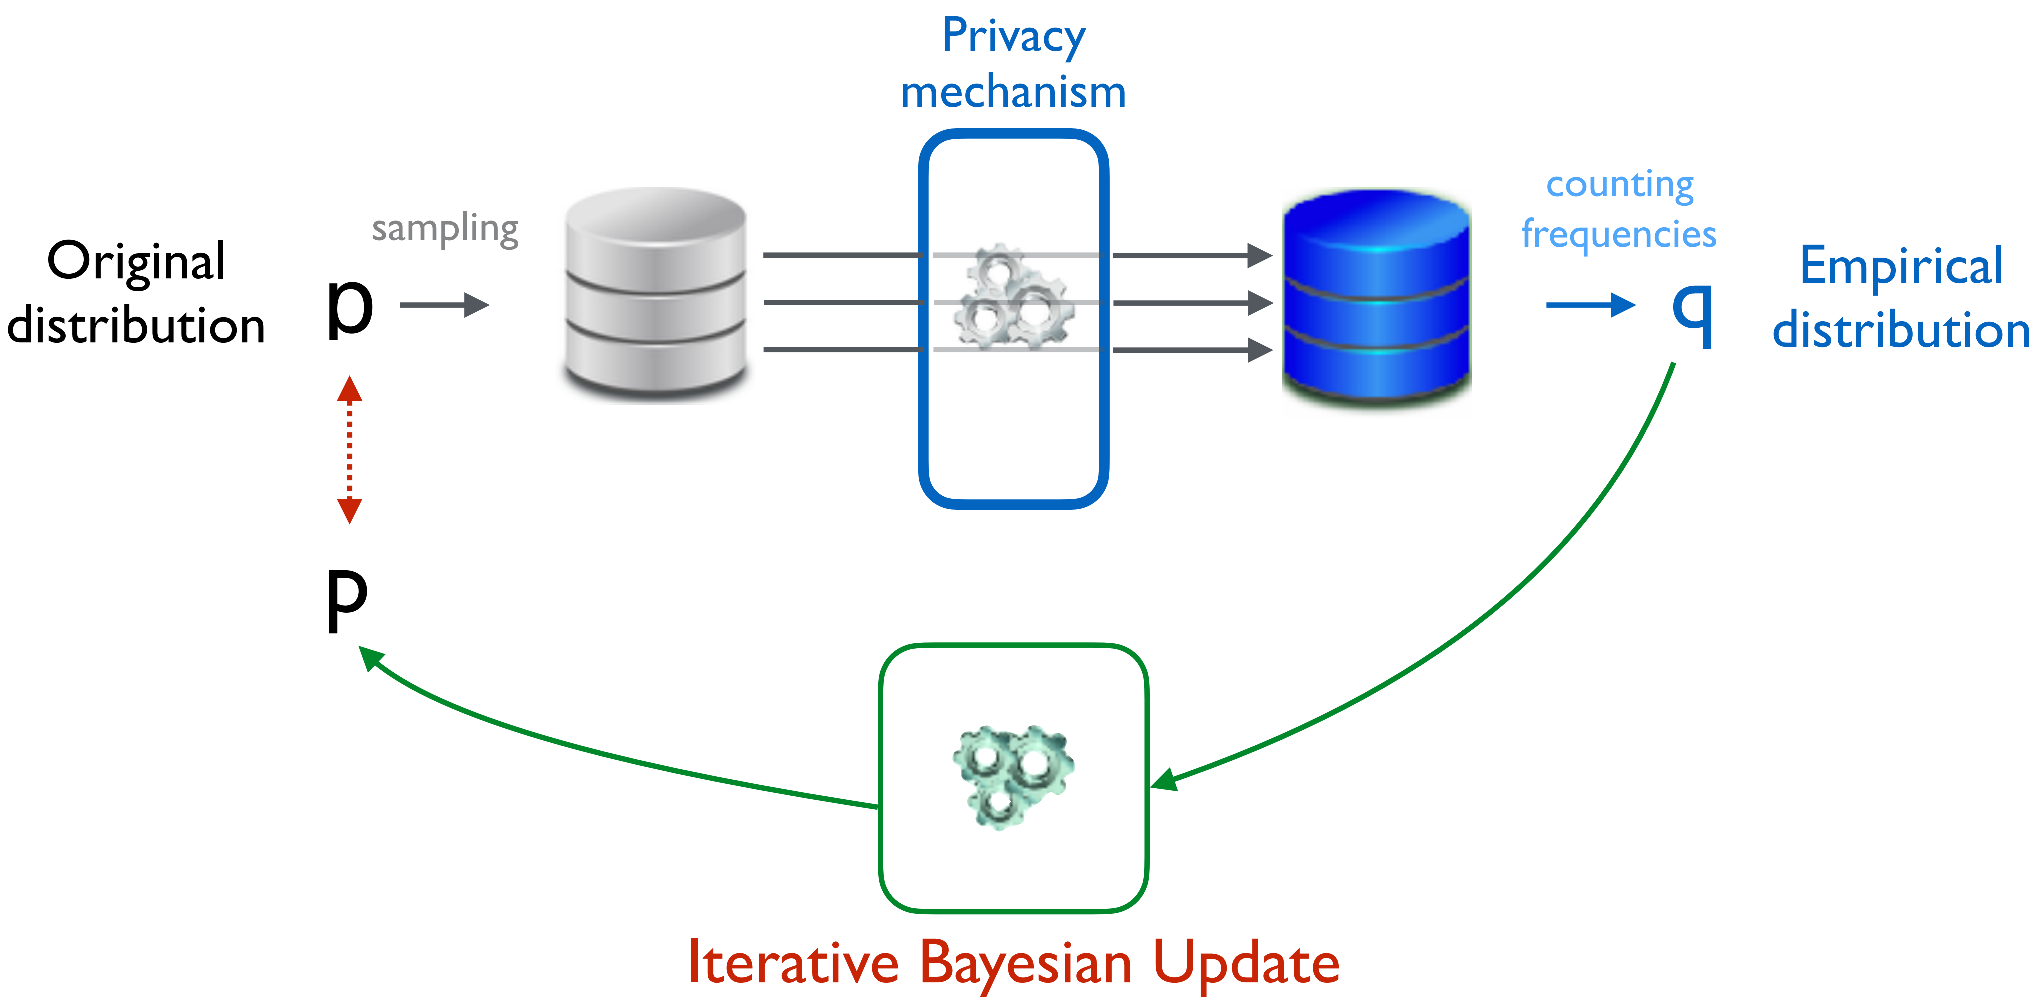

Estimation method

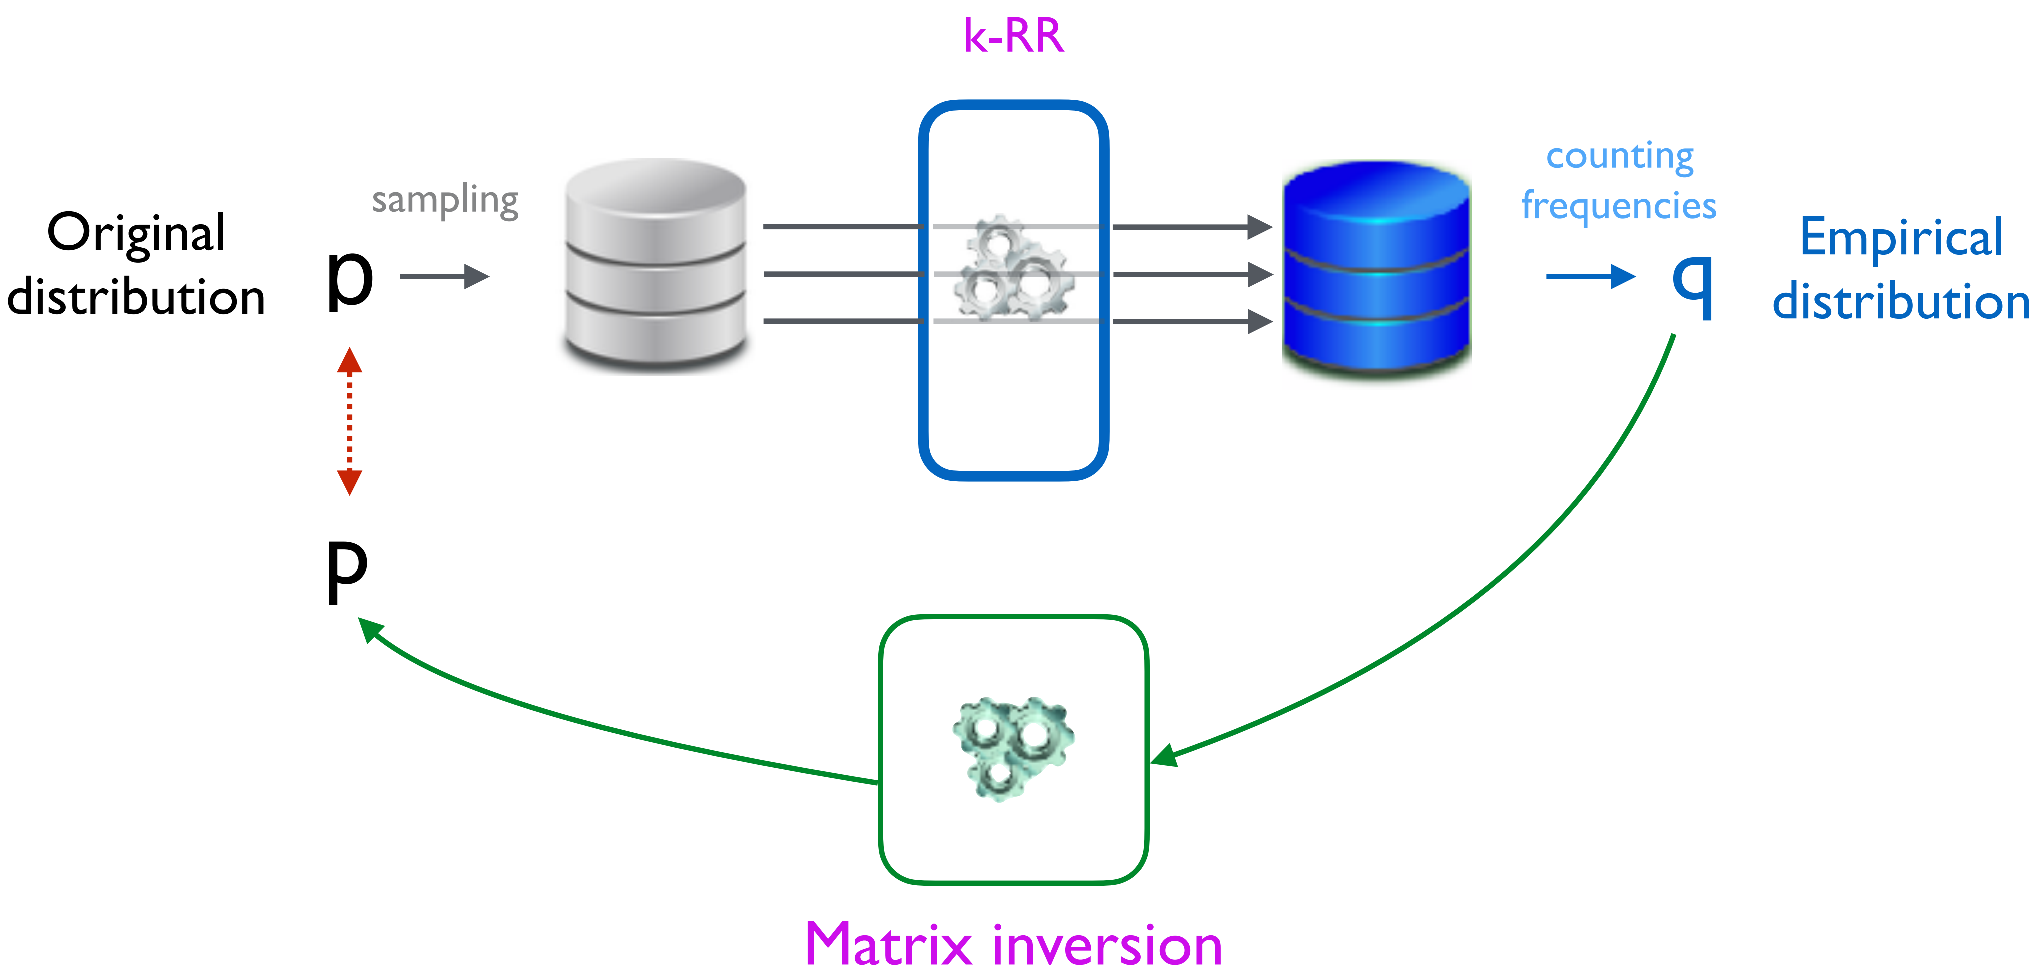

Estimation mechanism: The matrix inversion method
[Kairouz et al, '16]

- C : stochastic matrix associated to the privacy mechanism
- q : empirical distribution derived from the noisy data
- Estimation mechanism **Inv**: $p = q\,C^{-1}$

Example Assume $q(\mathit{Yes}) = \frac{6}{10}$ and $q(\mathit{No}) = \frac{4}{10}$. Then:

$$\frac{3}{4}\,p(\mathit{Yes}) + \frac{1}{4}\,p(\mathit{No}) = \frac{6}{10}$$

$$\frac{1}{4}\,p(\mathit{Yes}) + \frac{3}{4}\,p(\mathit{No}) = \frac{4}{10}$$

From which we derive $p(\mathit{Yes}) = \frac{7}{10}$ and $p(\mathit{No}) = \frac{3}{10}$

y

x	yes	no
yes	¾	¼
no	¼	¾

Estimation mechanism: The matrix inversion method

Problem 1: C must be invertible

Problem 2: $p = q\,C^{-1}$ may not be a distribution

Assume $q(Yes) = \frac{4}{5}$ and $q(No) = \frac{1}{5}$. Then:

$$\frac{3}{4}\,p(Yes) + \frac{1}{4}\,p(No) = \frac{4}{5}$$
$$\frac{1}{4}\,p(Yes) + \frac{3}{4}\,p(No) = \frac{1}{5}$$

x \ y	yes	no
yes	¾	¼
no	¼	¾

From which we derive $p(Yes) = \frac{11}{10}$ and $p(No) = -\frac{1}{10}$

Statistical utility: The matrix inversion method

$p = q\, C^{-1}$ may not be a distribution because it may contain negative elements.

In order to try to obtain the true distribution π we can either:

- set to 0 all the negative elements, and renormalize, or
- project p on the simplex.

Both these methods (especially the second one) give good estimate when used with the k-RR privacy mechanisms, but not with others. In particular, not with d-privacy.

An est. mech. that works well also for d-privacy: the Iterative Bayesian Update

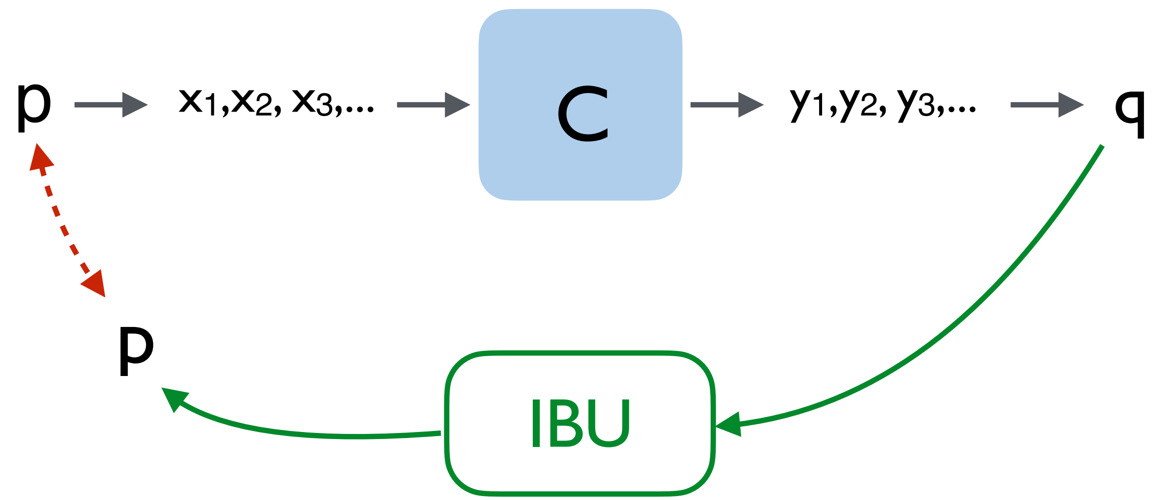

The IBU:

- is based on the **Maximization-Expectation** method
- produces a **Maximum Likelihood Estimator** $\hat{p}$ of the true distribution p
- If C is invertible, then the MLE is unique and IBU converges to p

An est. mech. that works well also for d-privacy: the Iterative Bayesian Update

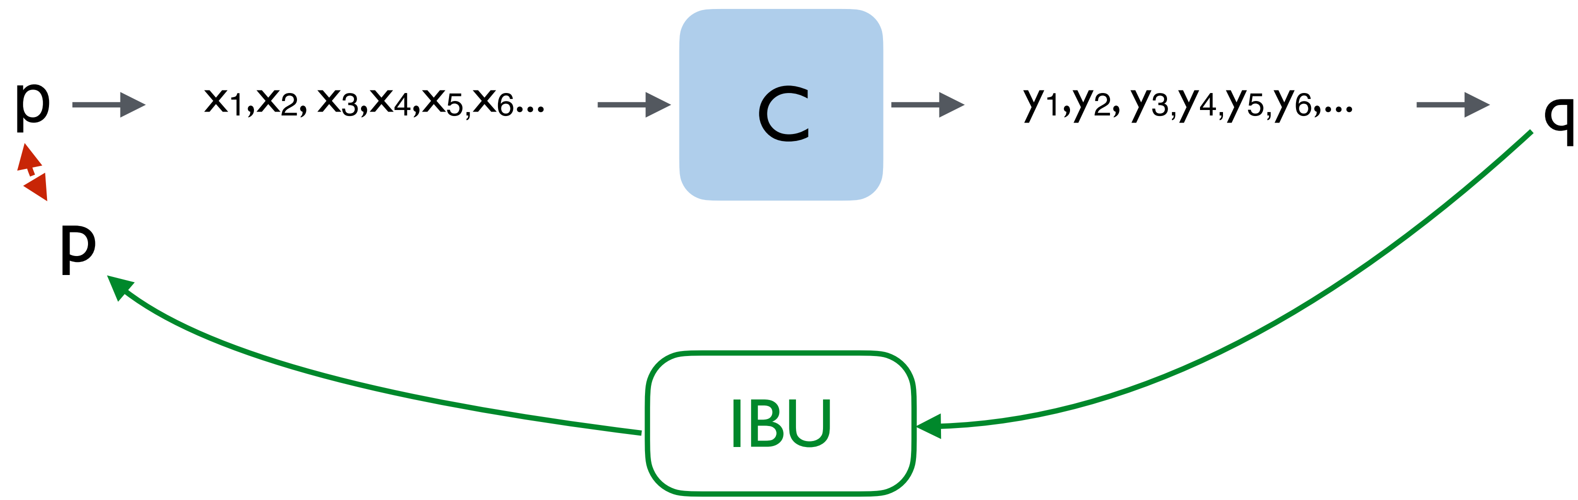

The IBU:

- is based on the **Maximization-Expectation** method
- produces a **Maximum Likelihood Estimator** *p* of the true distribution p
- If C is invertible, then the MLE is unique and IBU converges to p

The Iterative Bayesian Update

- Define $p^{(0)}$ = any fully supported distribution (e.g., the uniform distribution)
- Repeat: Define $p^{(n+1)}$ as the Bayesian update of $p^{(n)}$ weighted on the corresponding element of q, namely:

$$p_x^{(n+1)} = \sum_y q_y \frac{p_x^{(n)} C_{xy}}{\sum_z p_z^{(n)} C_{zy}}$$

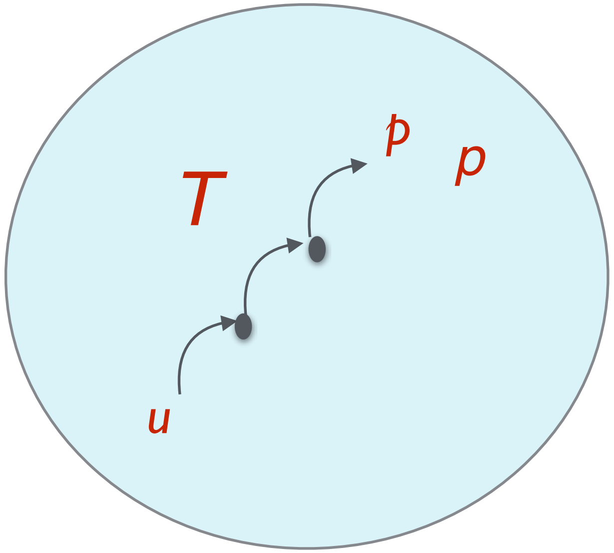

- Note that $p^{(n+1)} = T(p^{(n)})$
- If C is invertible then T is a contraction
- If T is a contraction then there is a unique fixed point p and it converges to p
- Stopping condition: when $p^{(n+1)}$ is close to $p^{(n)}$

Trade-off between privacy and statistical utility
d-privacy versus k-RR

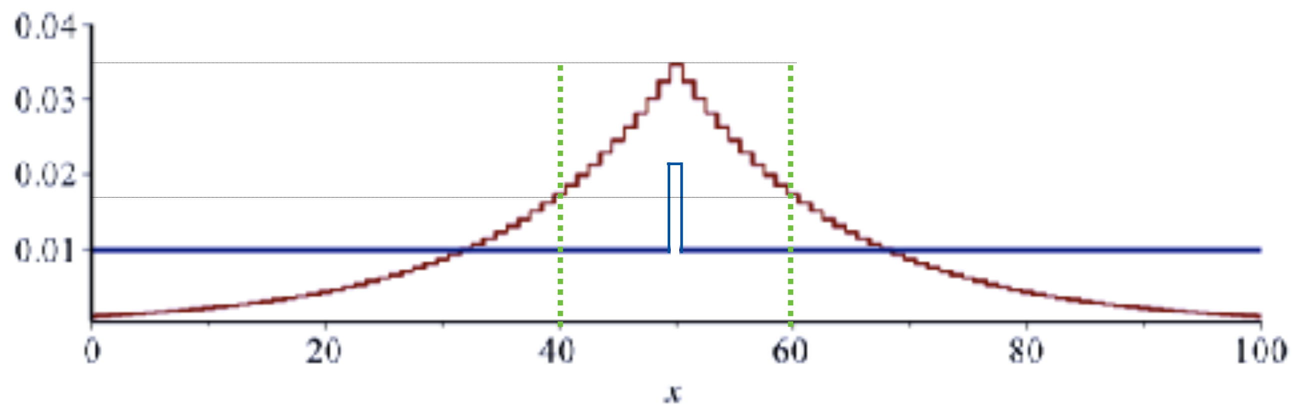

Both k-RR and the geometric / Laplace mechanisms are parametrized by ε, but it has a different meaning. To compare them wrt privacy, we need to calibrate ε, in such a way that the requested ratio is satisfied in the "area of interest" (area in which we want to be indistinguishable).

As for utility, it depends on the metric used to compare distributions. If the metric takes into account the underlying distance (e.g., the Earth-mover's distance) then the trade-off utility-privacy of d-privacy its much better than that of k-RR.

Experiments on the Gowalla dataset

- Gowalla is a dataset of geographical checkins in several cities in the world
- We have used it to compare the statistical utility of kRR and Planar Laplacian with the respective e calibrated so to satisfy the same privacy constraint:
  same level of privacy within about 1 Km^2

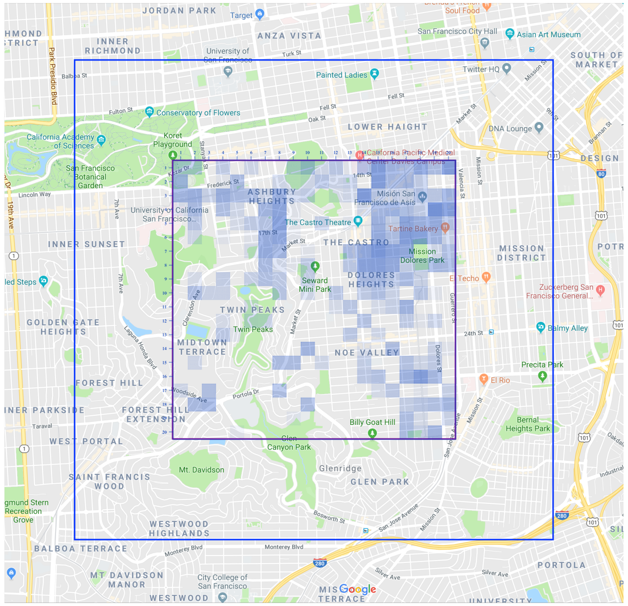

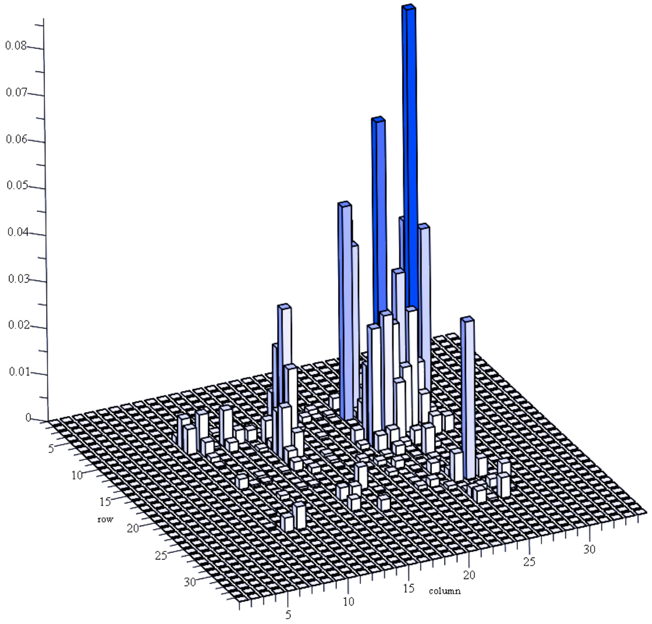

Gowalla checkins in an area of 3x3 km^2 in San Francisco downtown (about 10K checkins)

The Planar Laplace mechanism

$\varepsilon = \ln(2)$

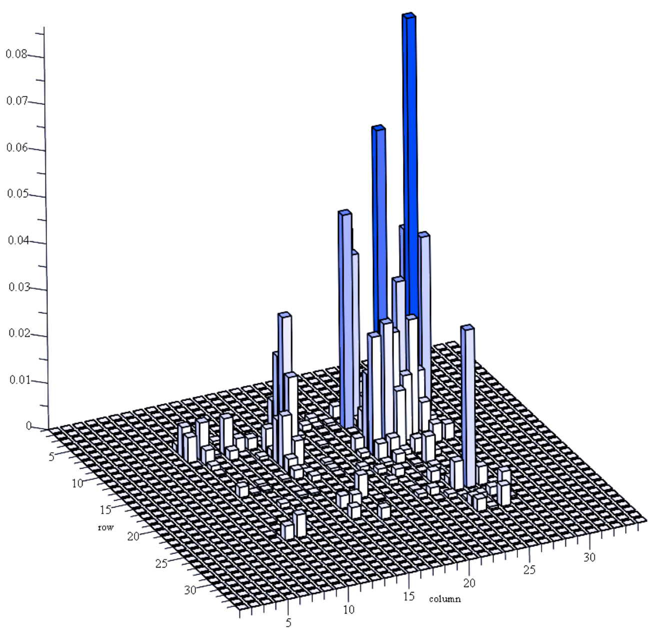

The real distribution

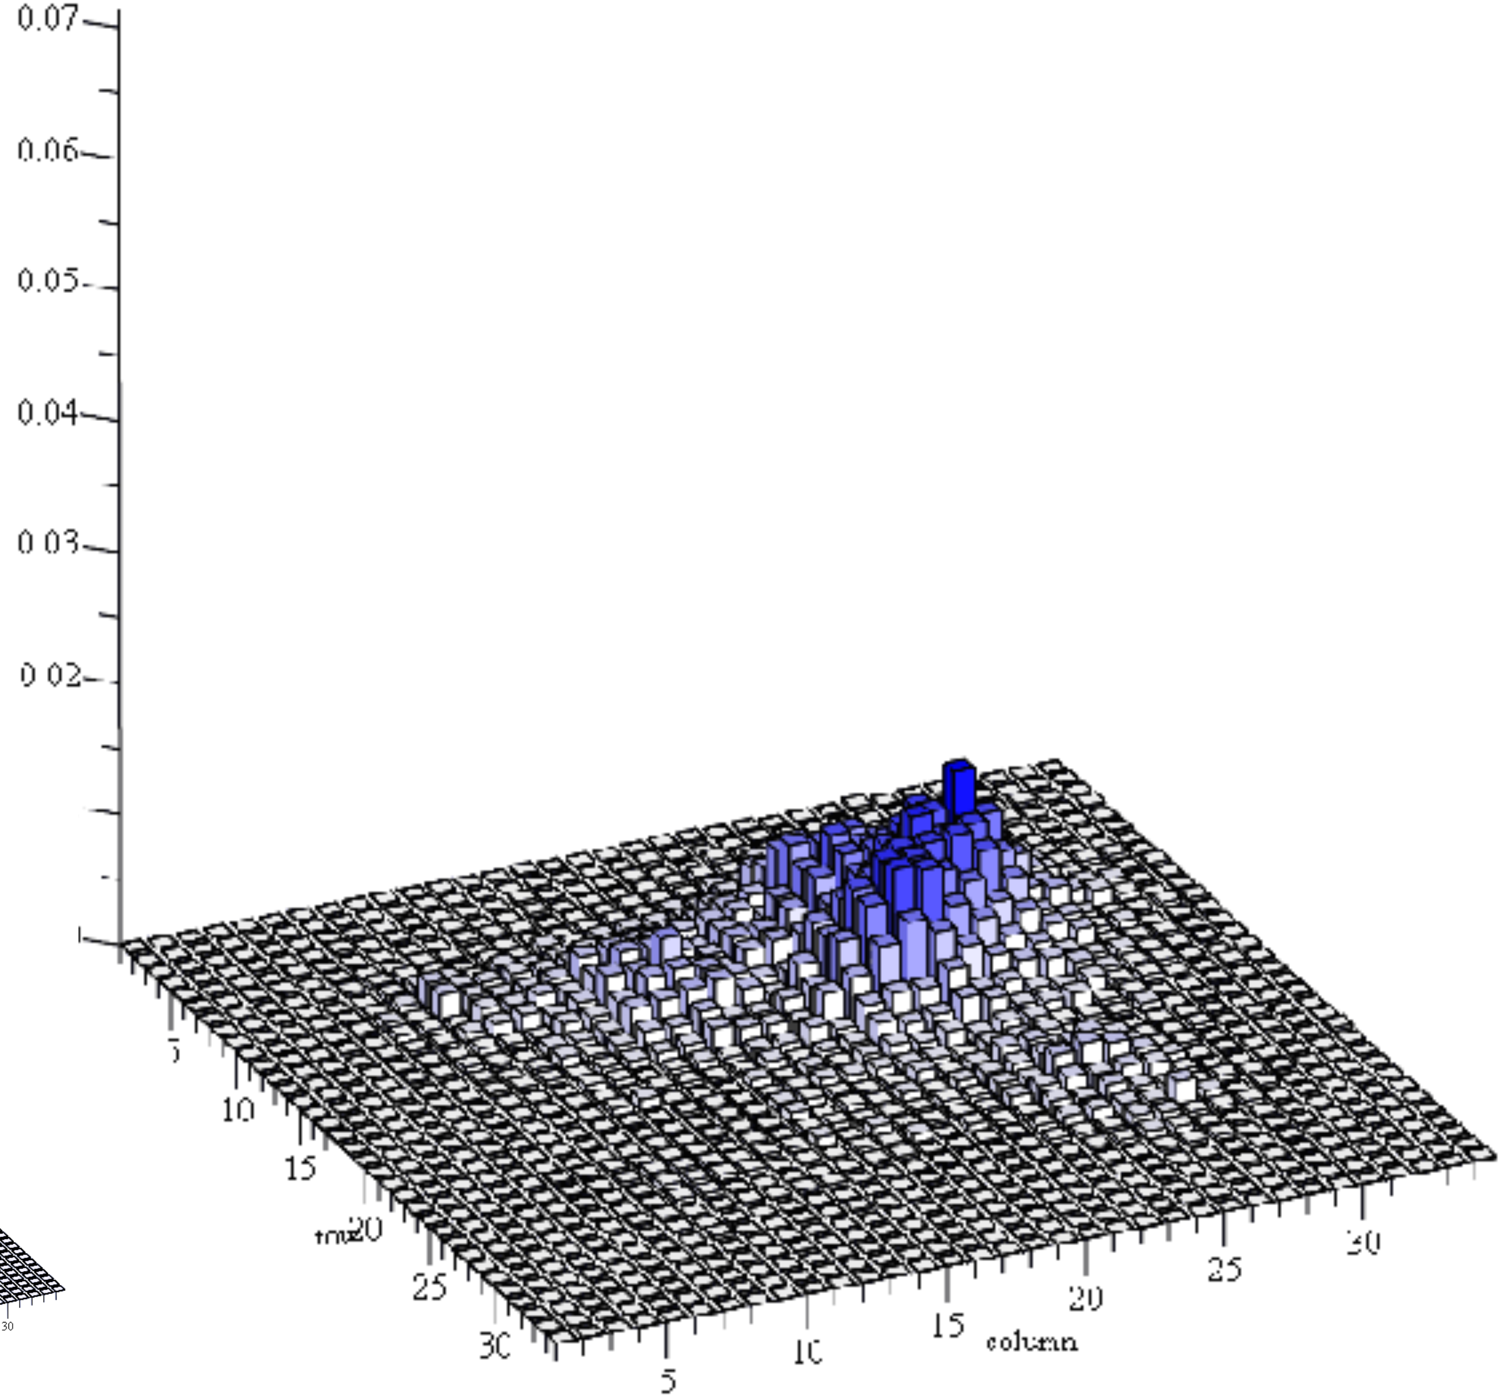

The noisy distribution and the result of the IBU (300 iterations)

The kRR mechanism

ε = ln(8)

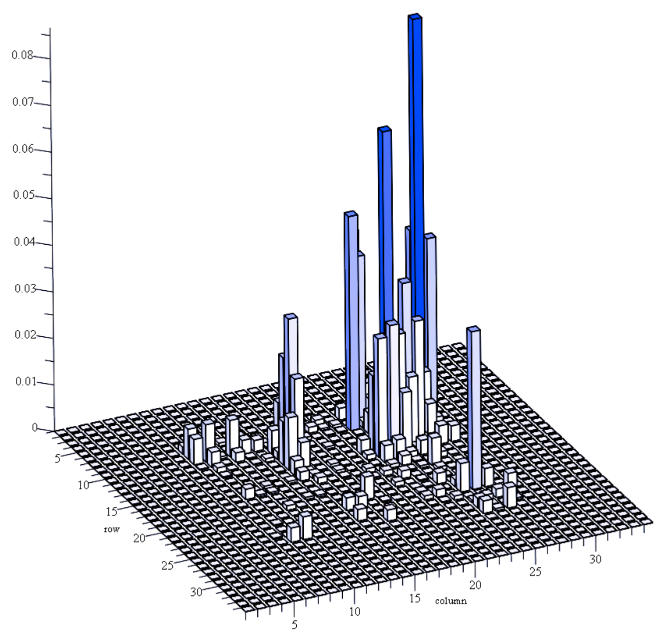

The real distribution

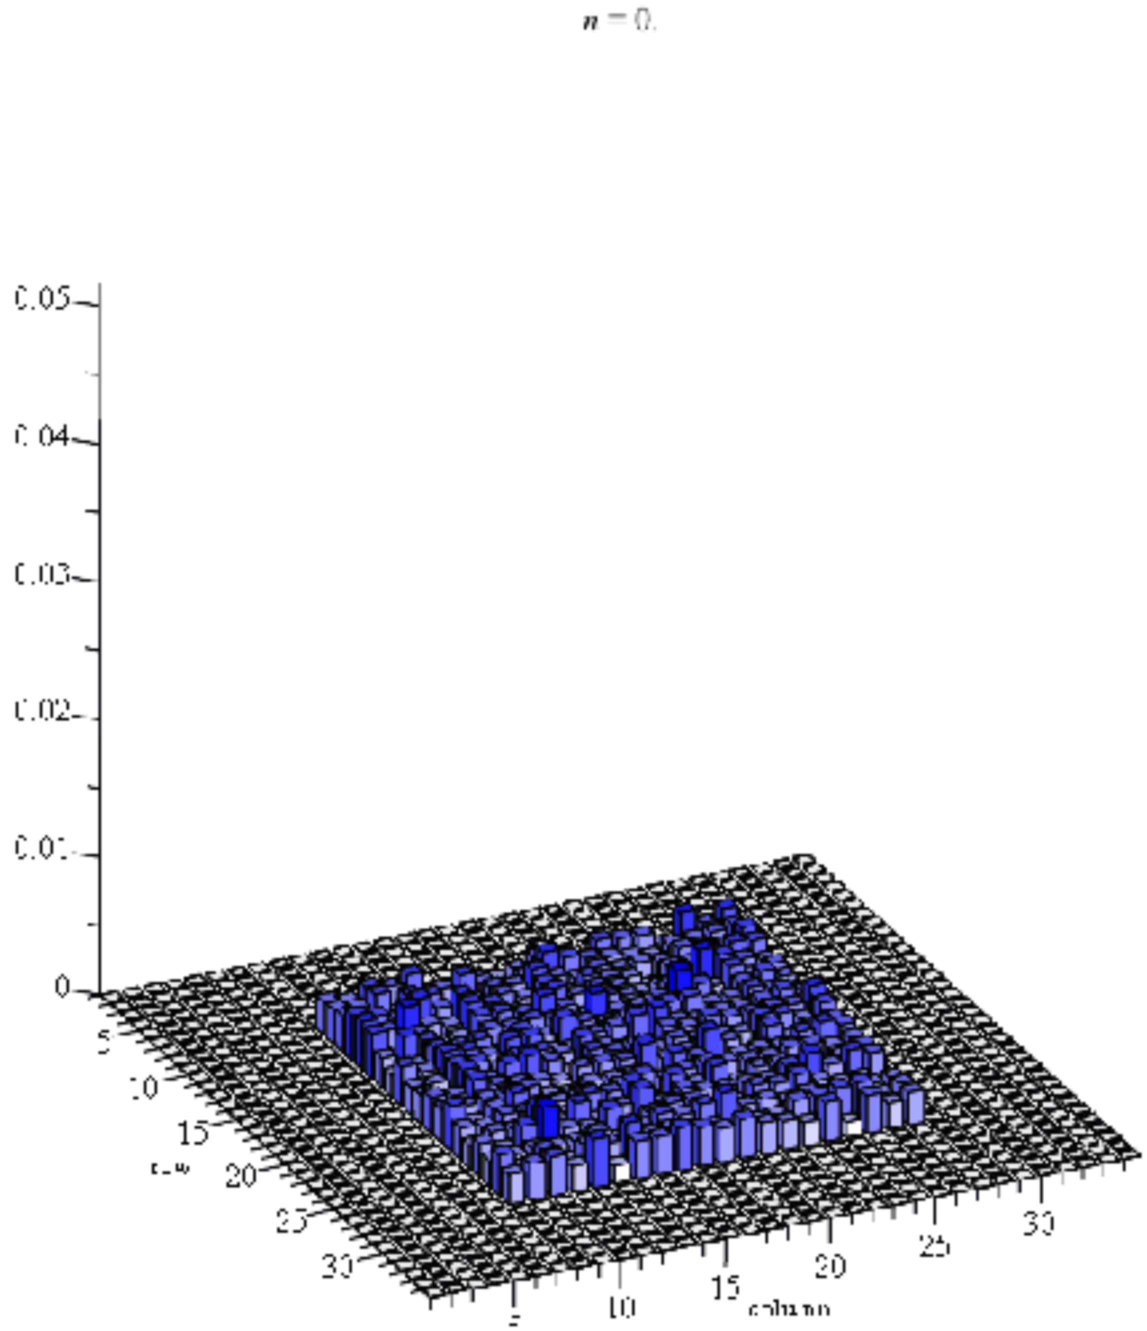

The noisy distribution and the result of the IBU (500 iterations)

Evaluation: San Francisco

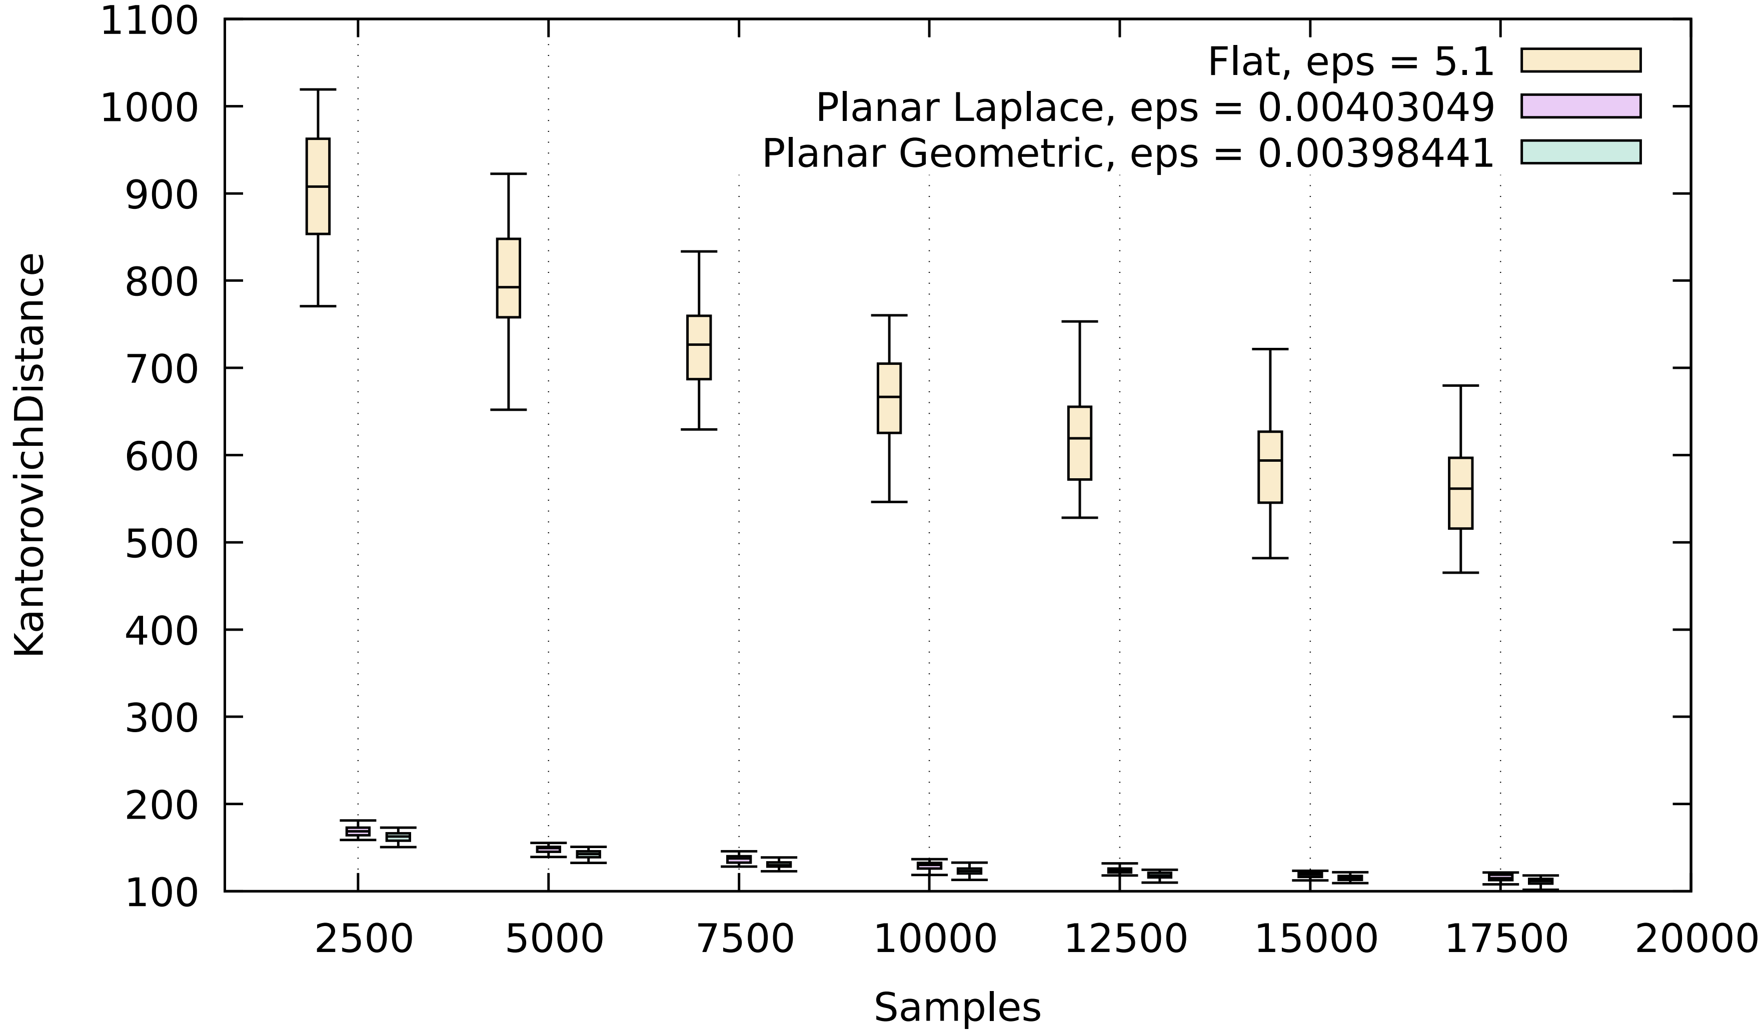

Evaluation: Paris

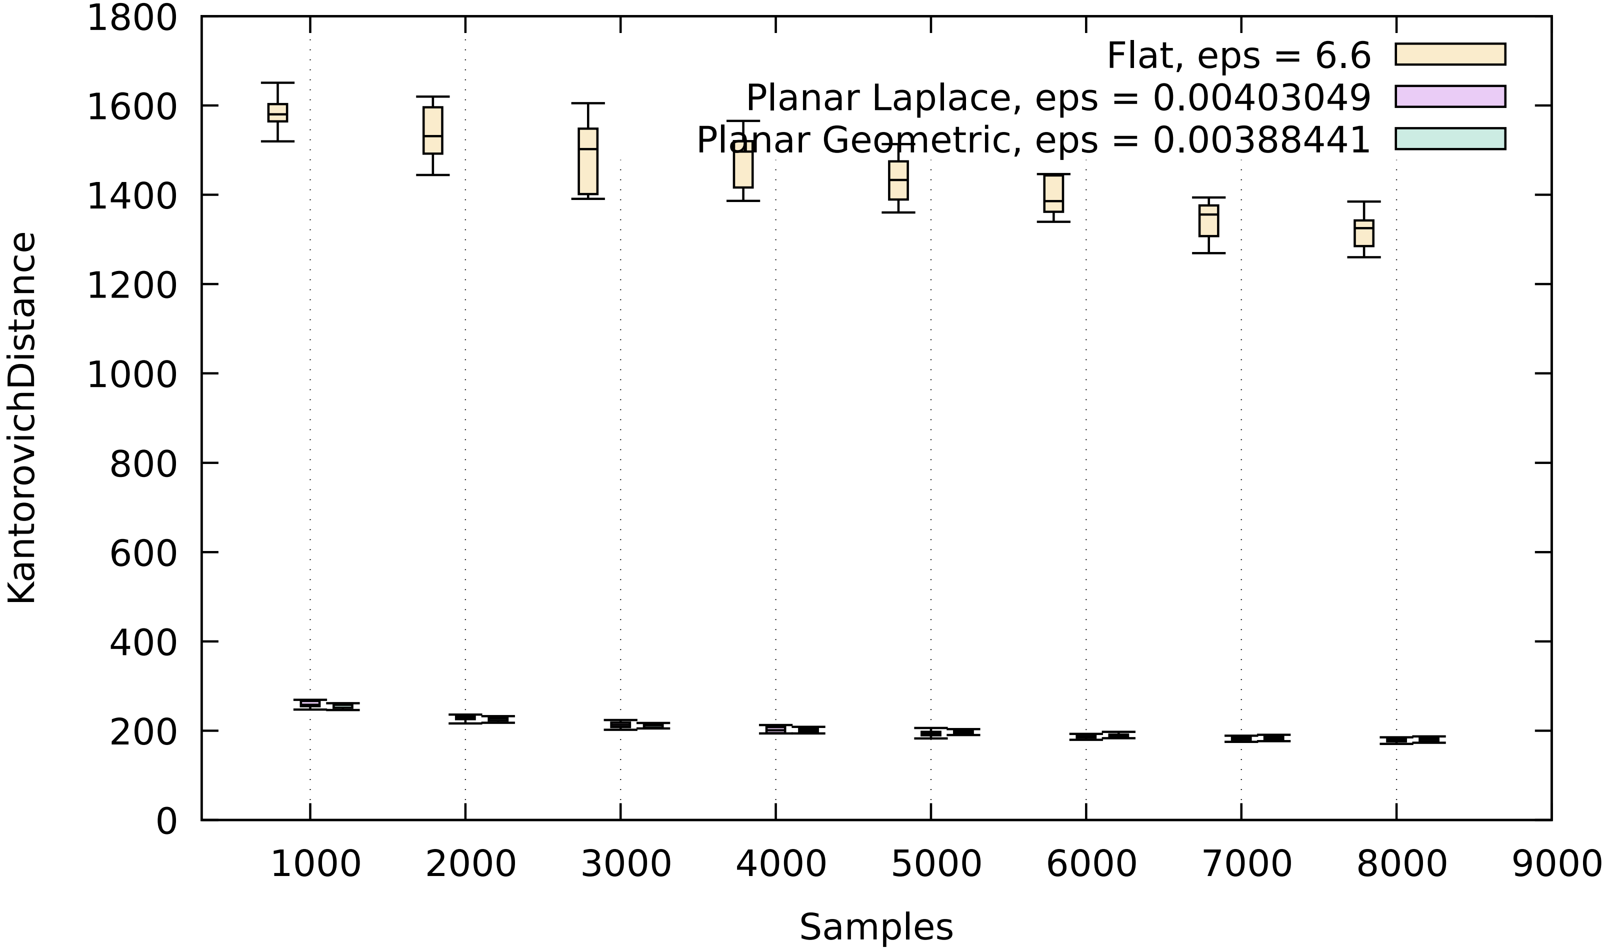

Summary

- Privacy vs utility
- Differential privacy: central and local models
- Statistical utility
- Compositionality
- An hybrid mechanism for privacy in a distributed setting

Compositionality of k-RR & Inv

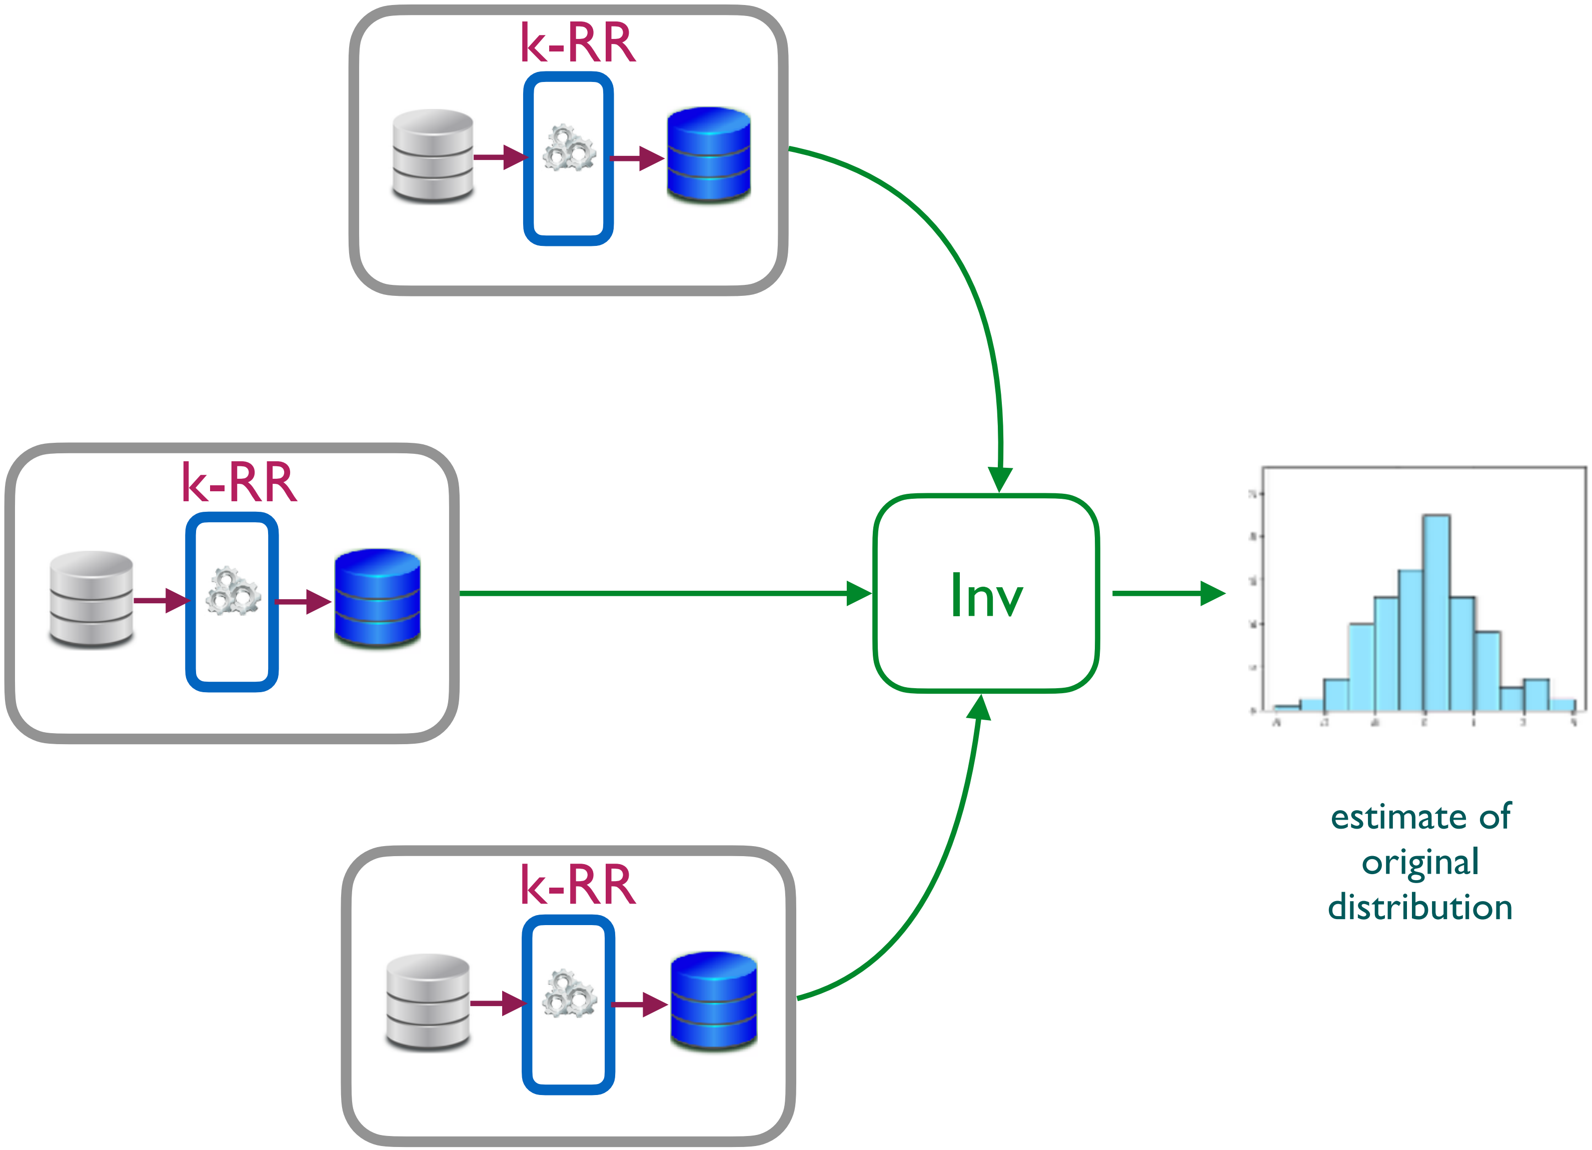

Compositionality of k-RR & Inv:

Exactly the same estimation accuracy as if the dataset was centralized

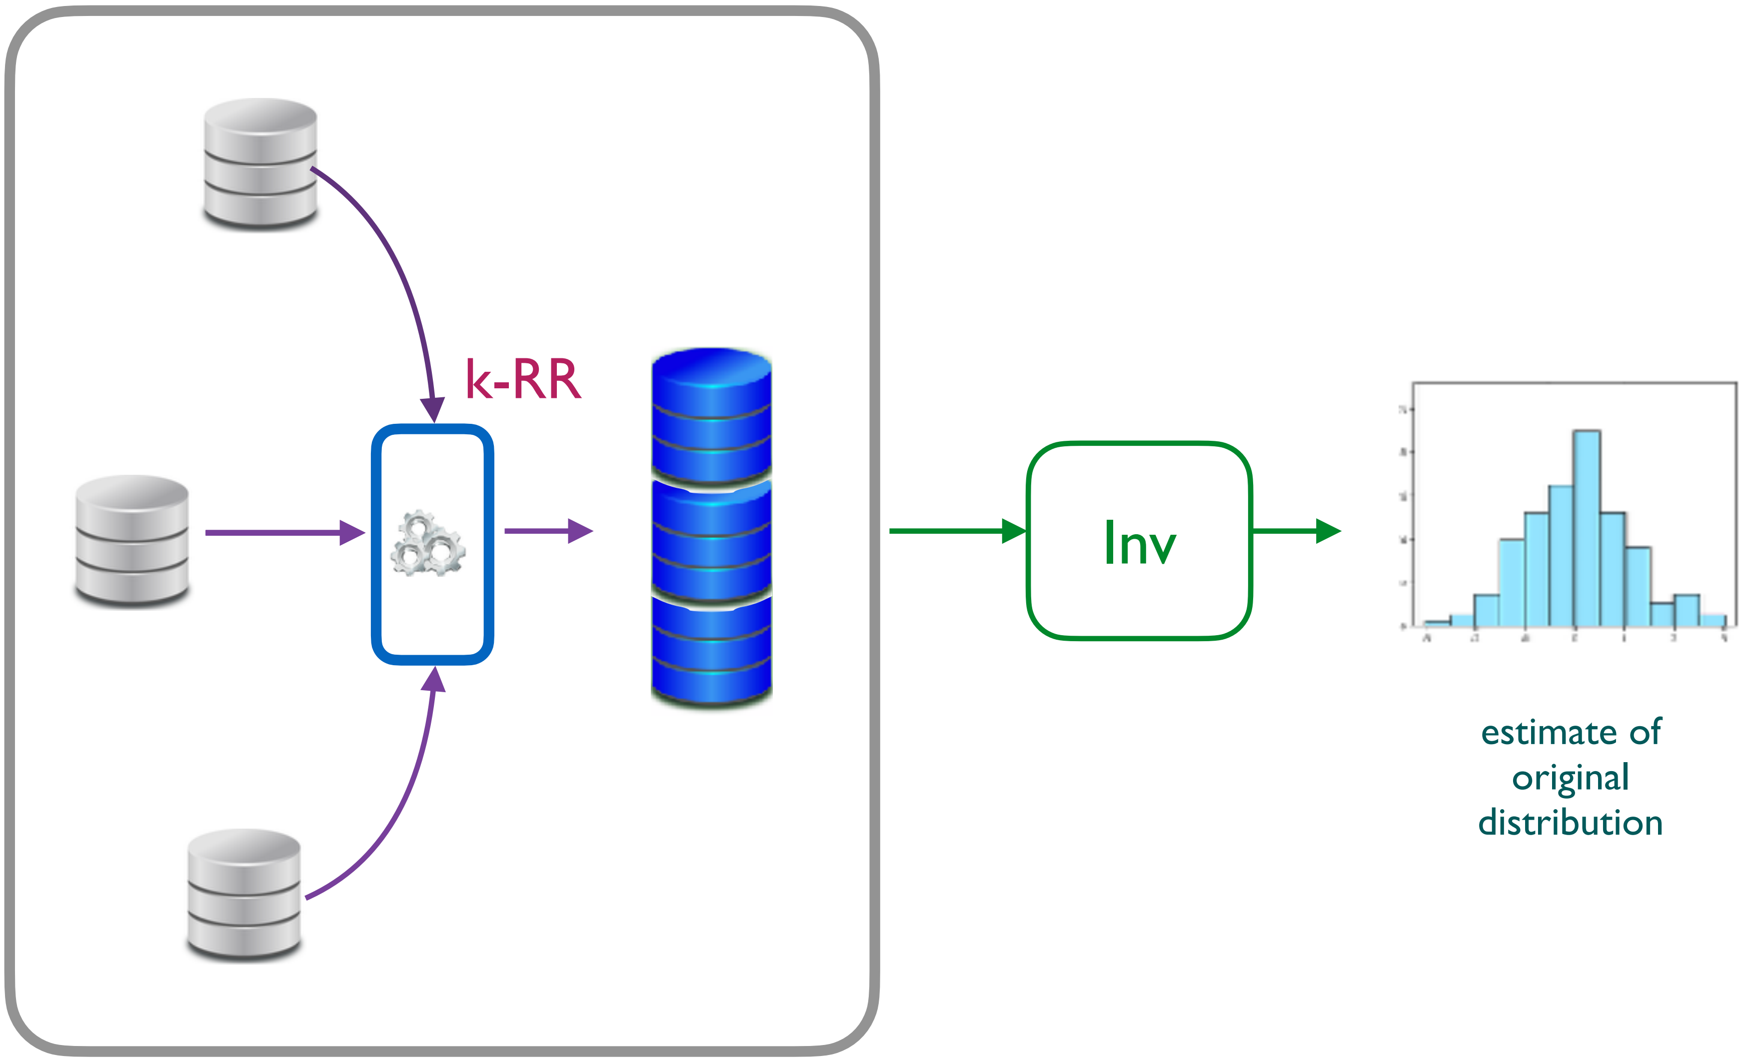

Compositionality
of k-RR & Inv:

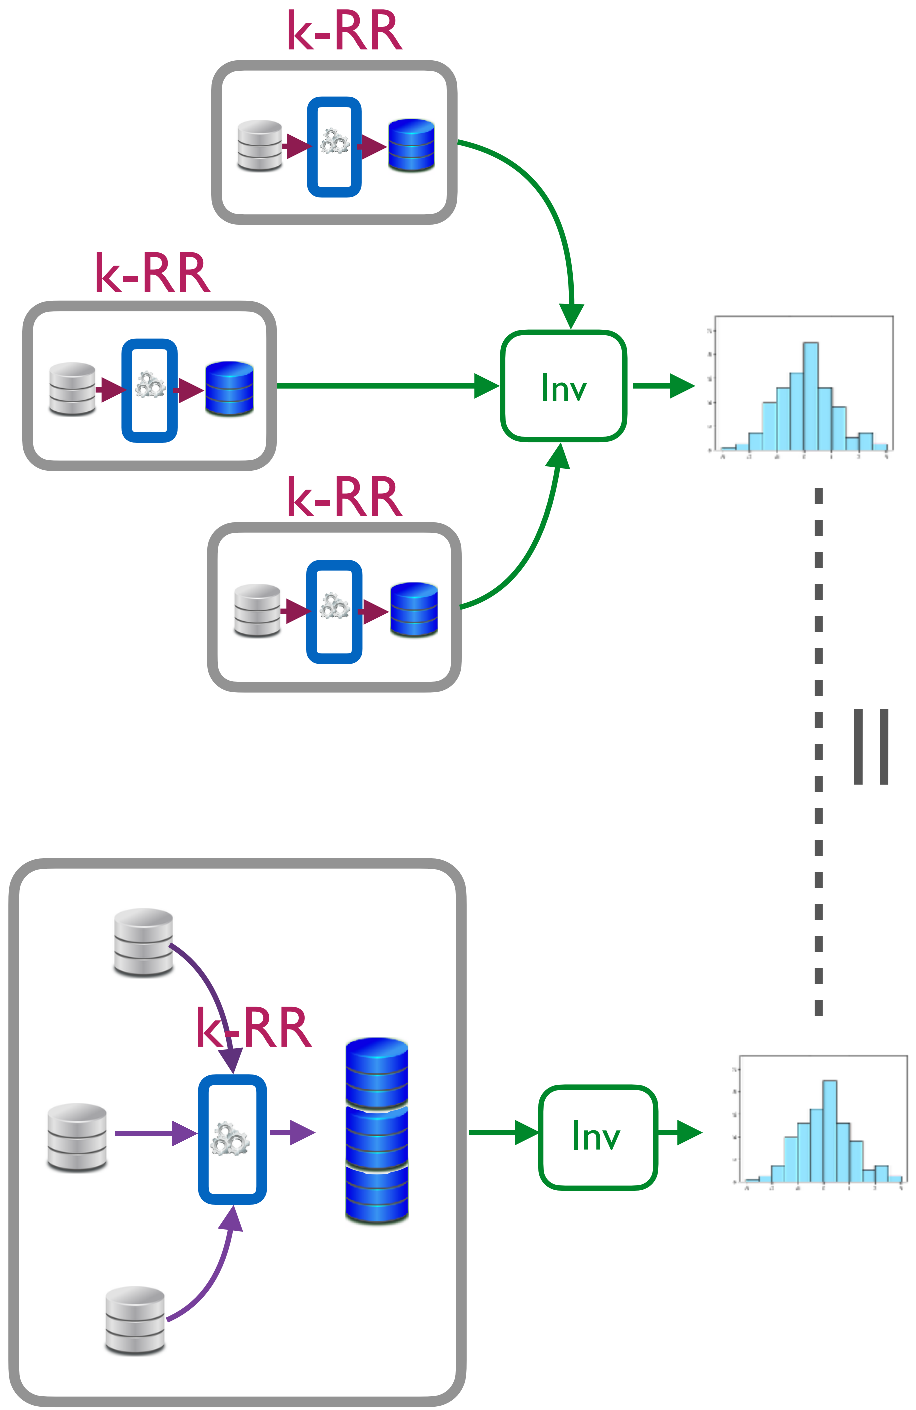

Compositionality of k-RR & Inv:

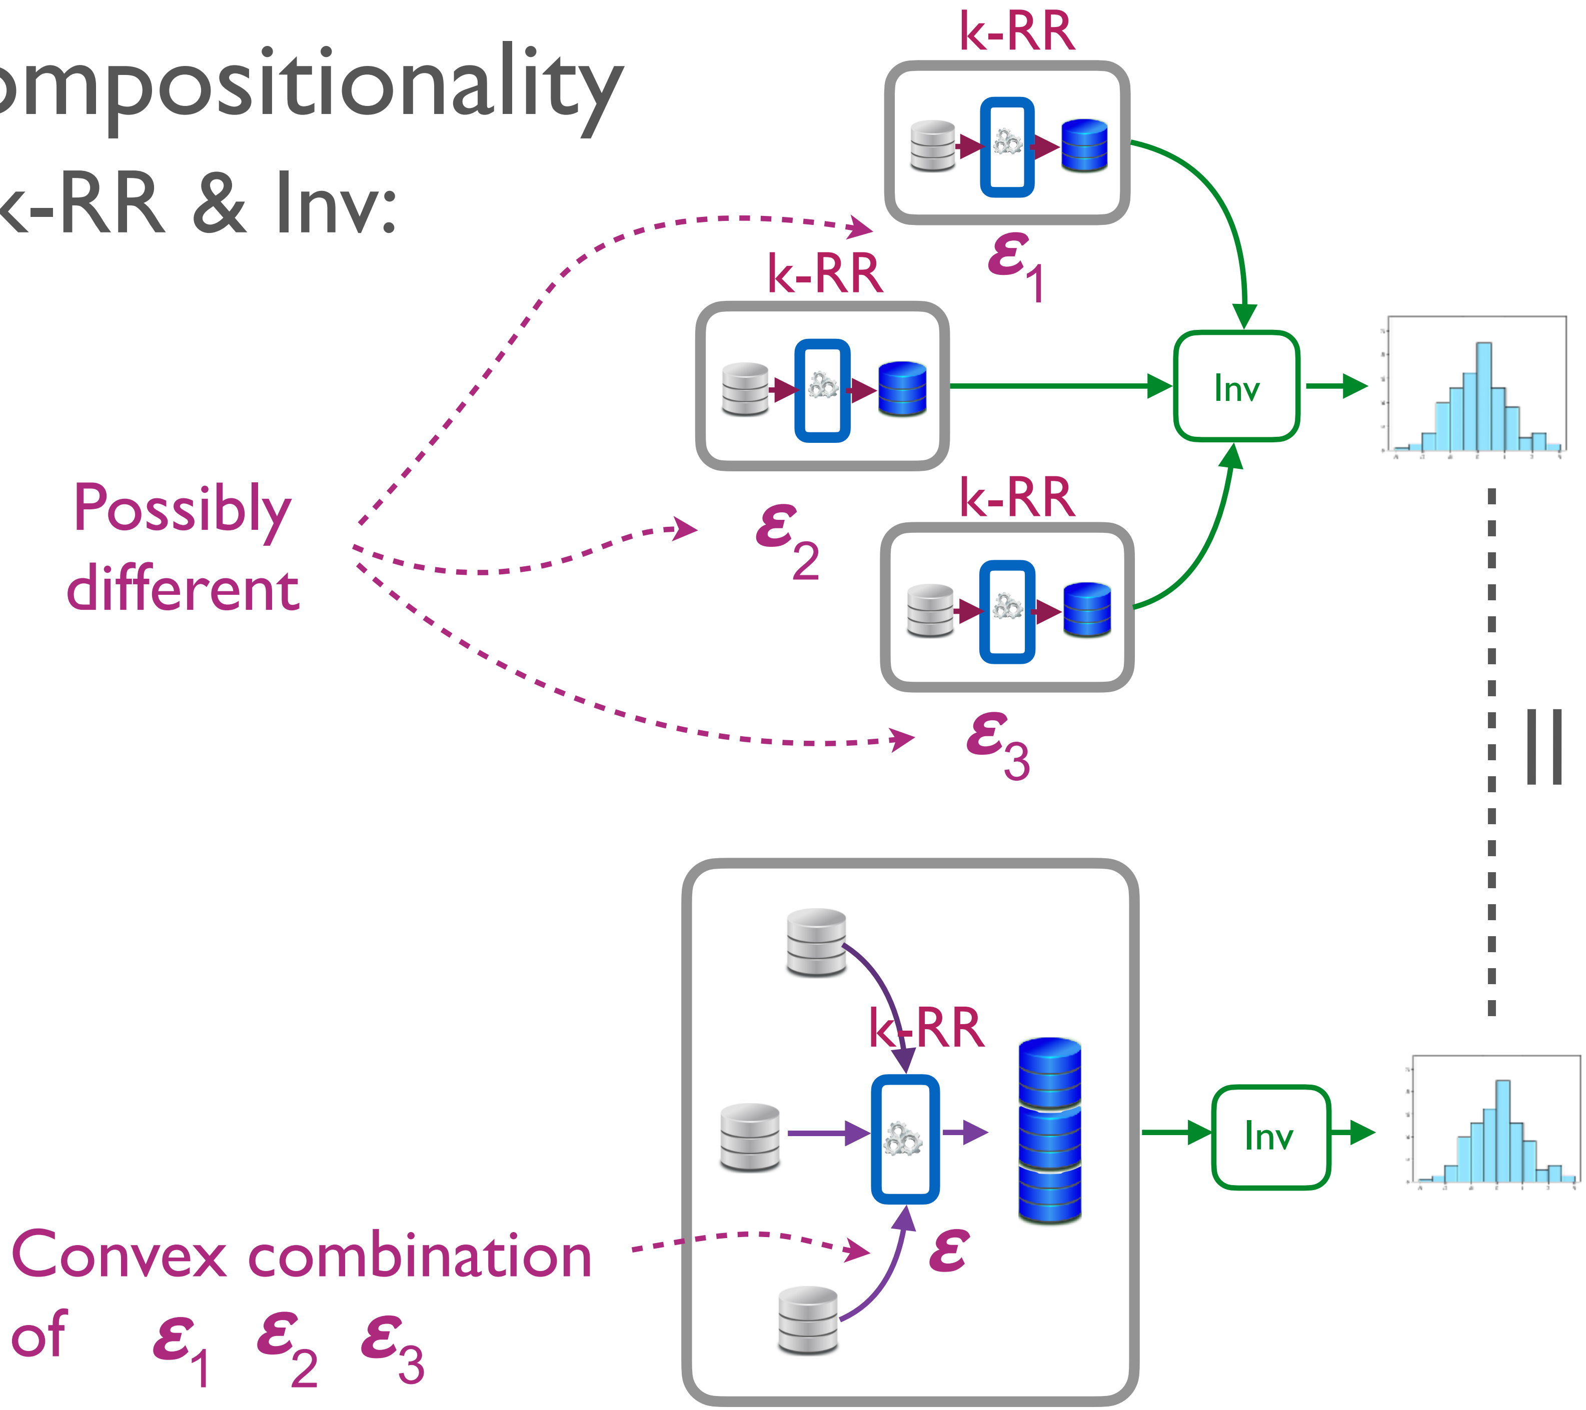

Generalized IBU (GIBU)

[Elsalamouny and Palamidessi, EuroS&P'20]

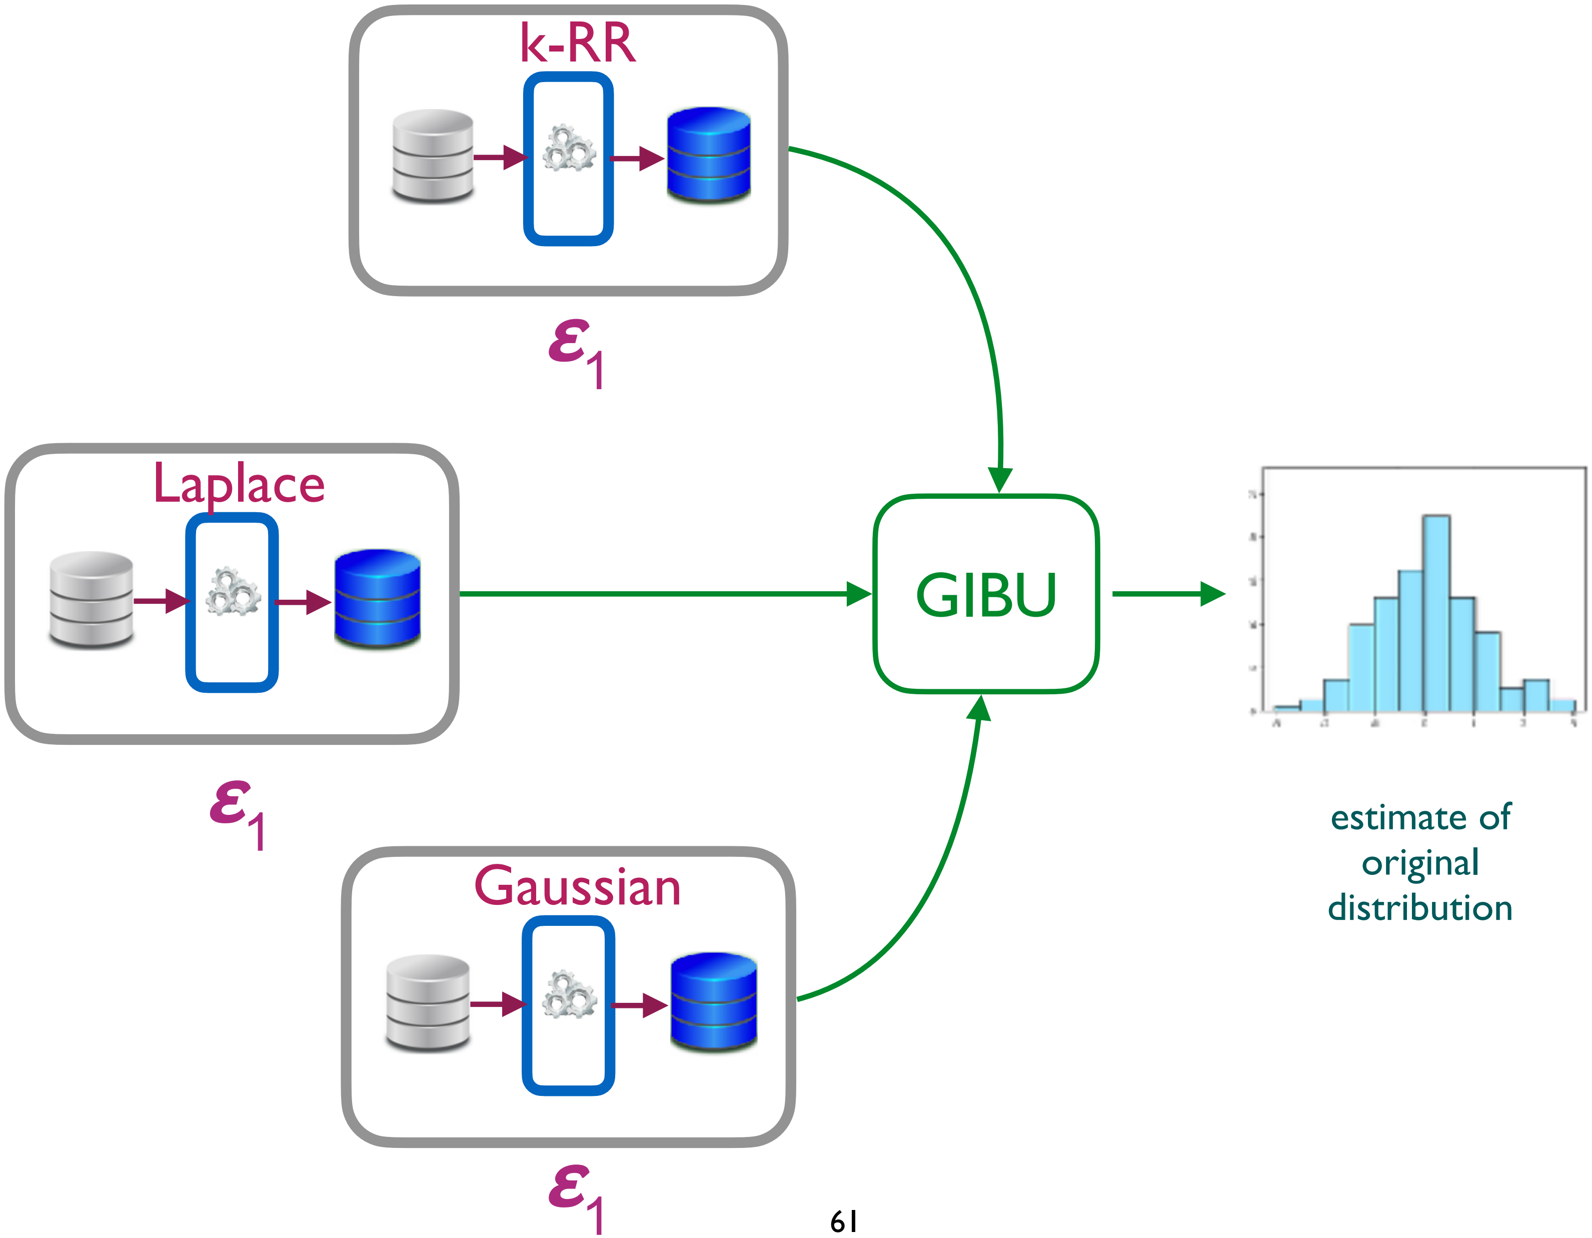

Compositionality of Generalized IBU

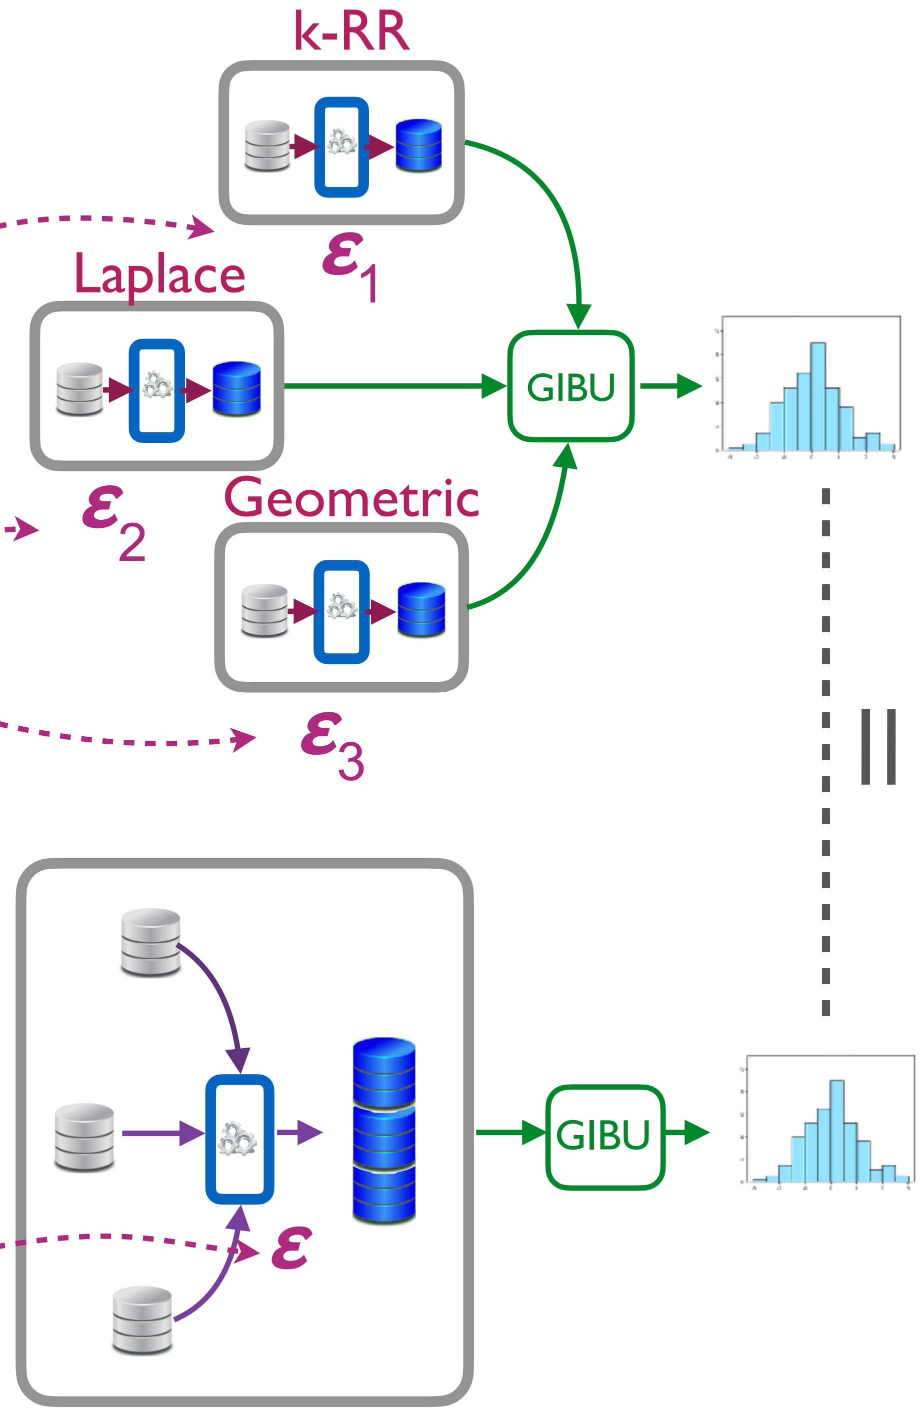

Possibly different

Convex combination of ε_1 ε_2 ε_3

Summary

- Privacy vs utility
- Differential privacy: central and local models
- Statistical utility
- Compositionality
- An hybrid mechanism for privacy in a distributed setting

Distributed setting

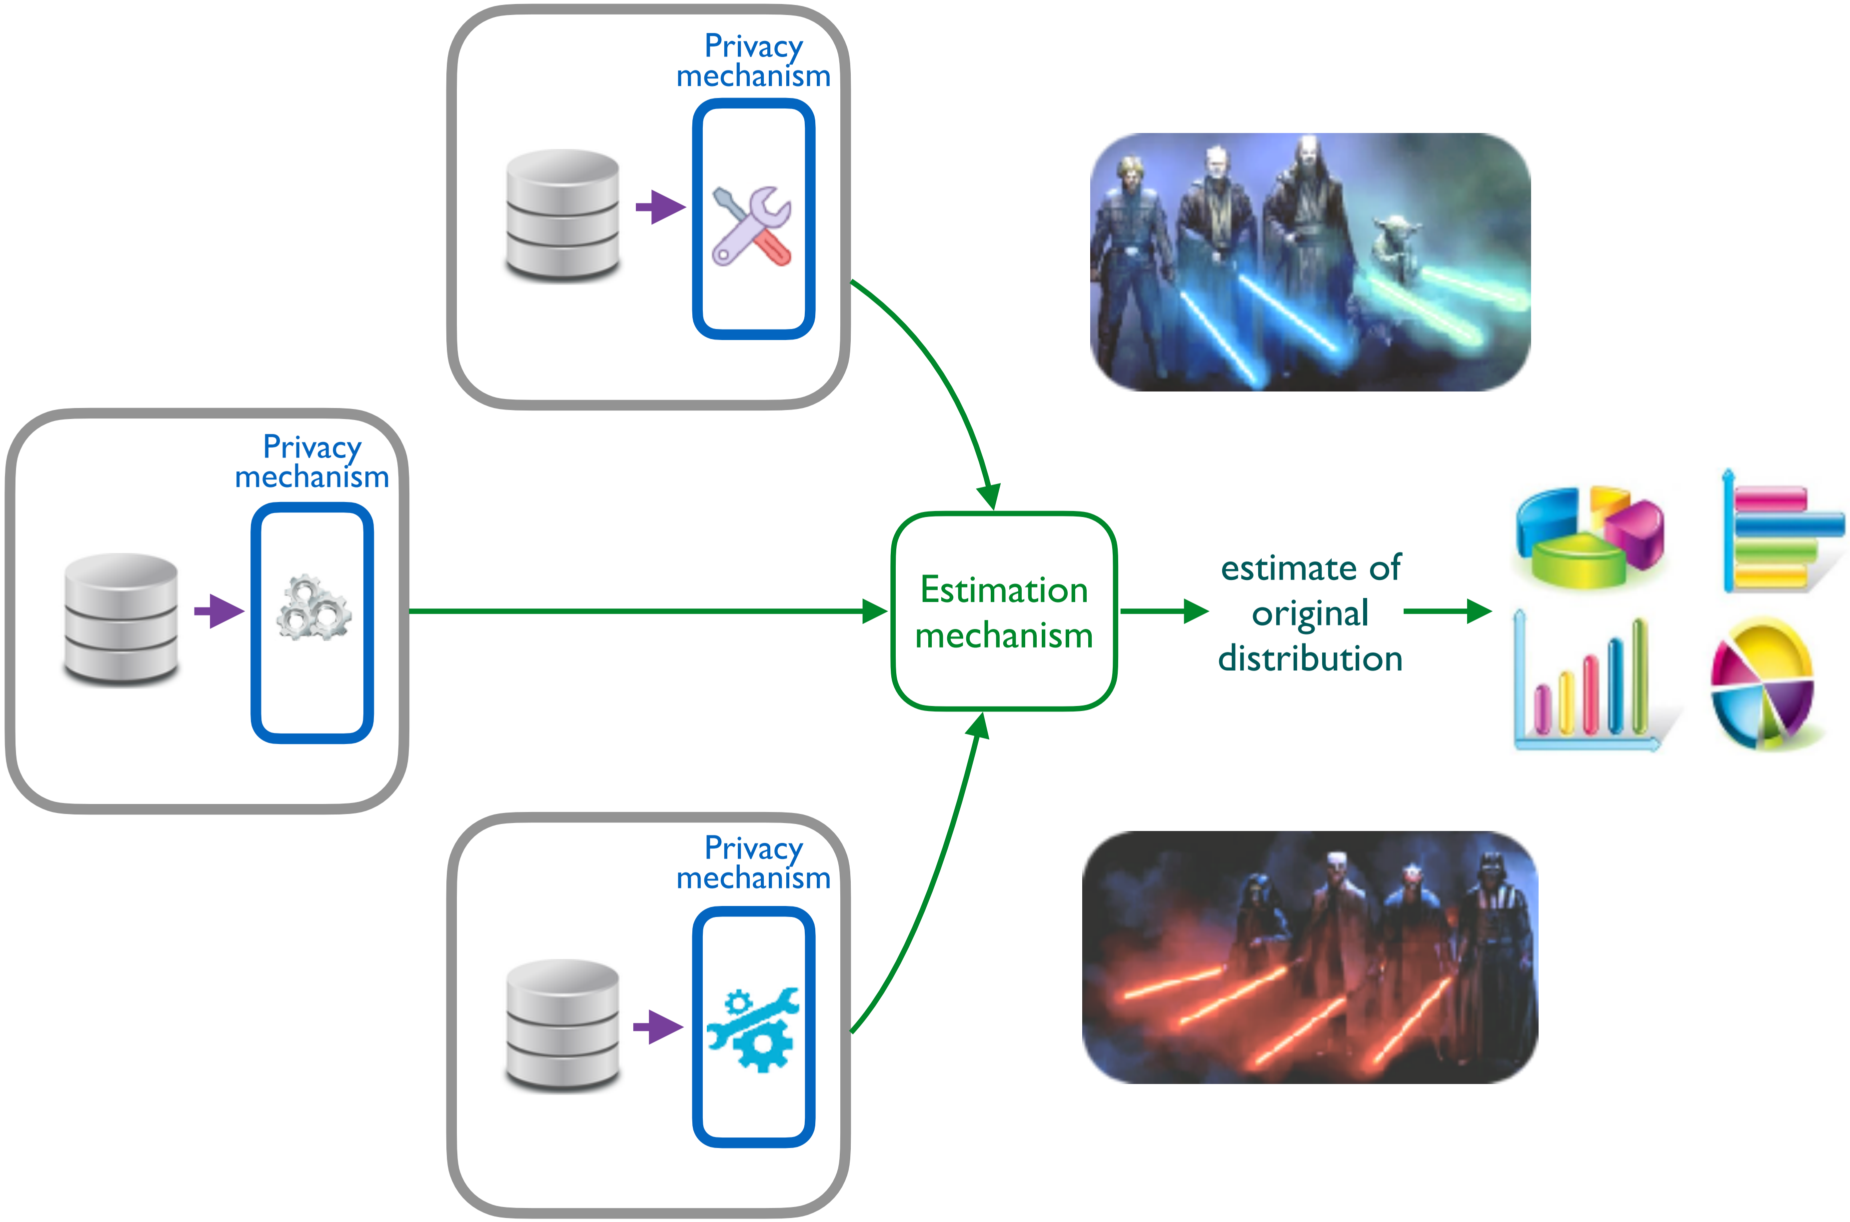

The hybrid approach

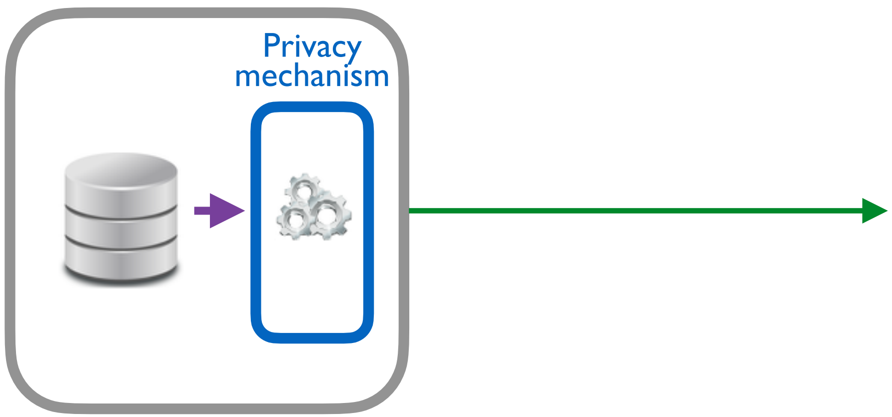

The hybrid approach

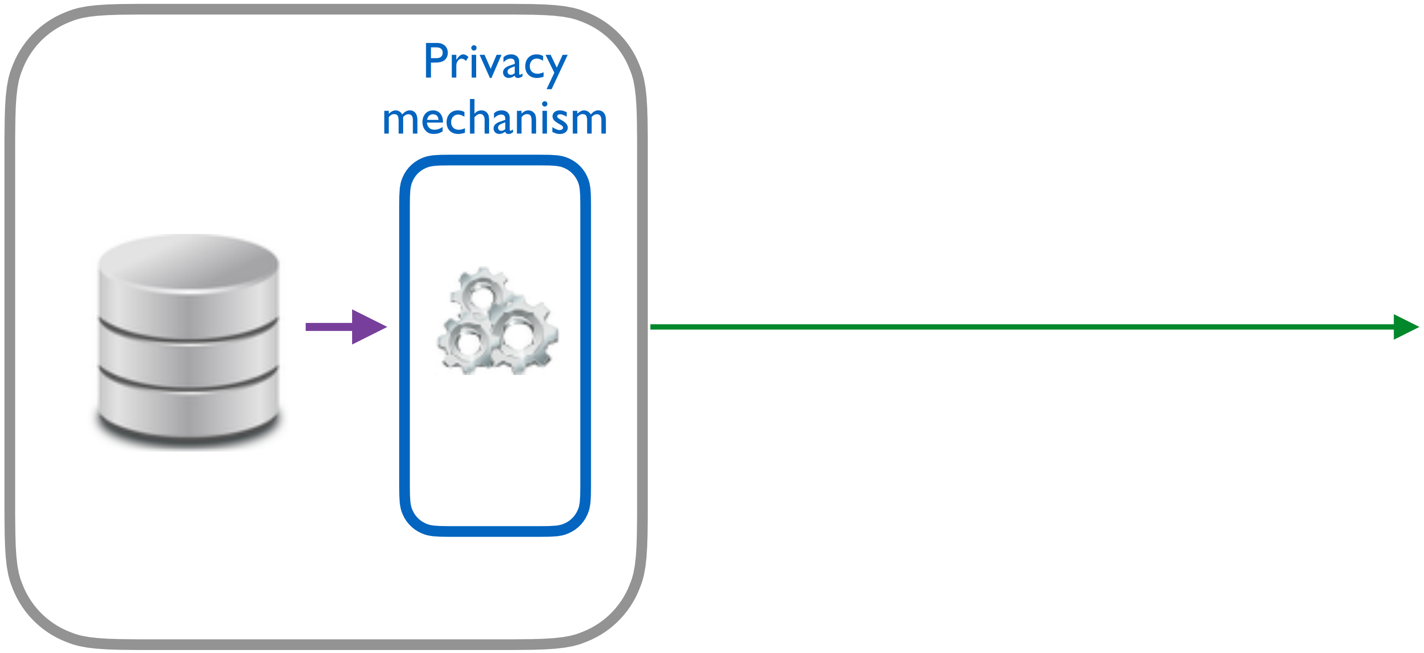

The hybrid approach

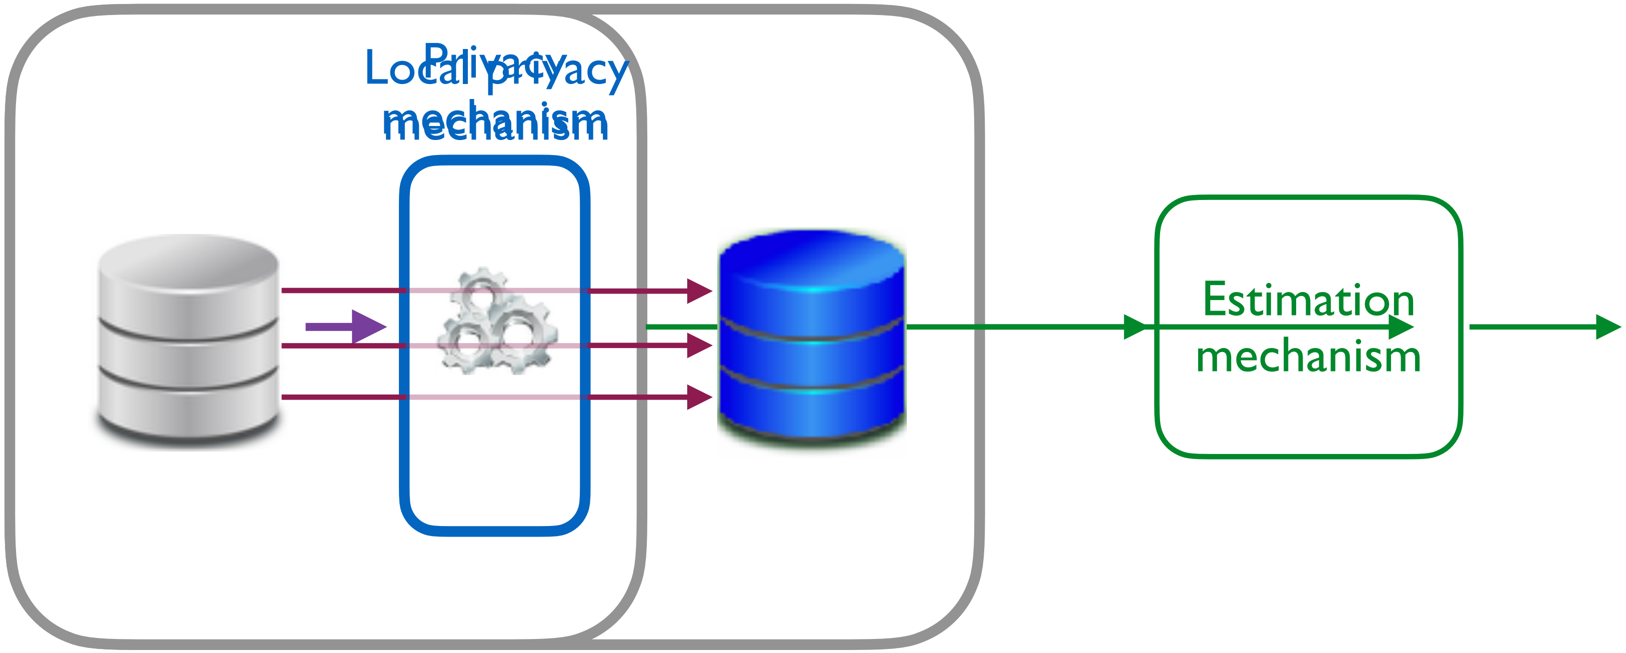

Apply a LDP mechanism to each record individually

Estimate the original distribution like in LDP

Advantages of hybrid wrt local

- The trade-off utility-privacy is usually much worse in the local model than in the central model

- However, in the hybrid model, the trade-off of certain mechanisms (kRR + Inv and d-privacy + GIBU) is as good as in the central model.

- Hybrid approach: combination of the local and central model. The **mechanism is local**, while the **attacker is like in the central model**, which is **weaker** than the one of the local model

- The hybrid mechanism for k-RR is also known as the Shuffle Model [Balle et al]

Privacy in the hybrid model

Attacker in the local model

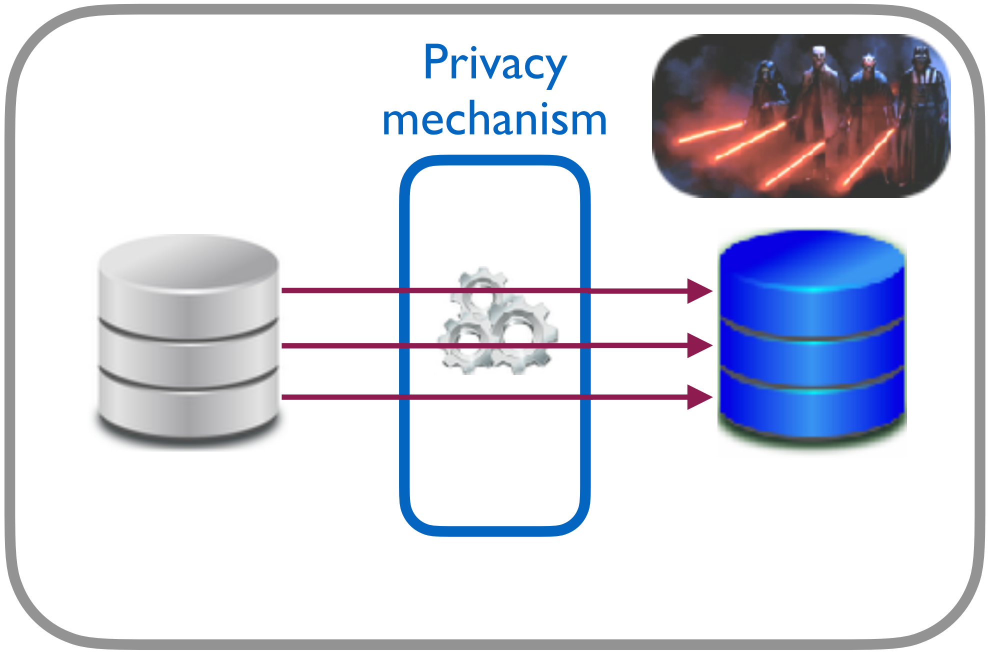

In the local model the attacker can see the obfuscated version of each record

Attacker in the hybrid model

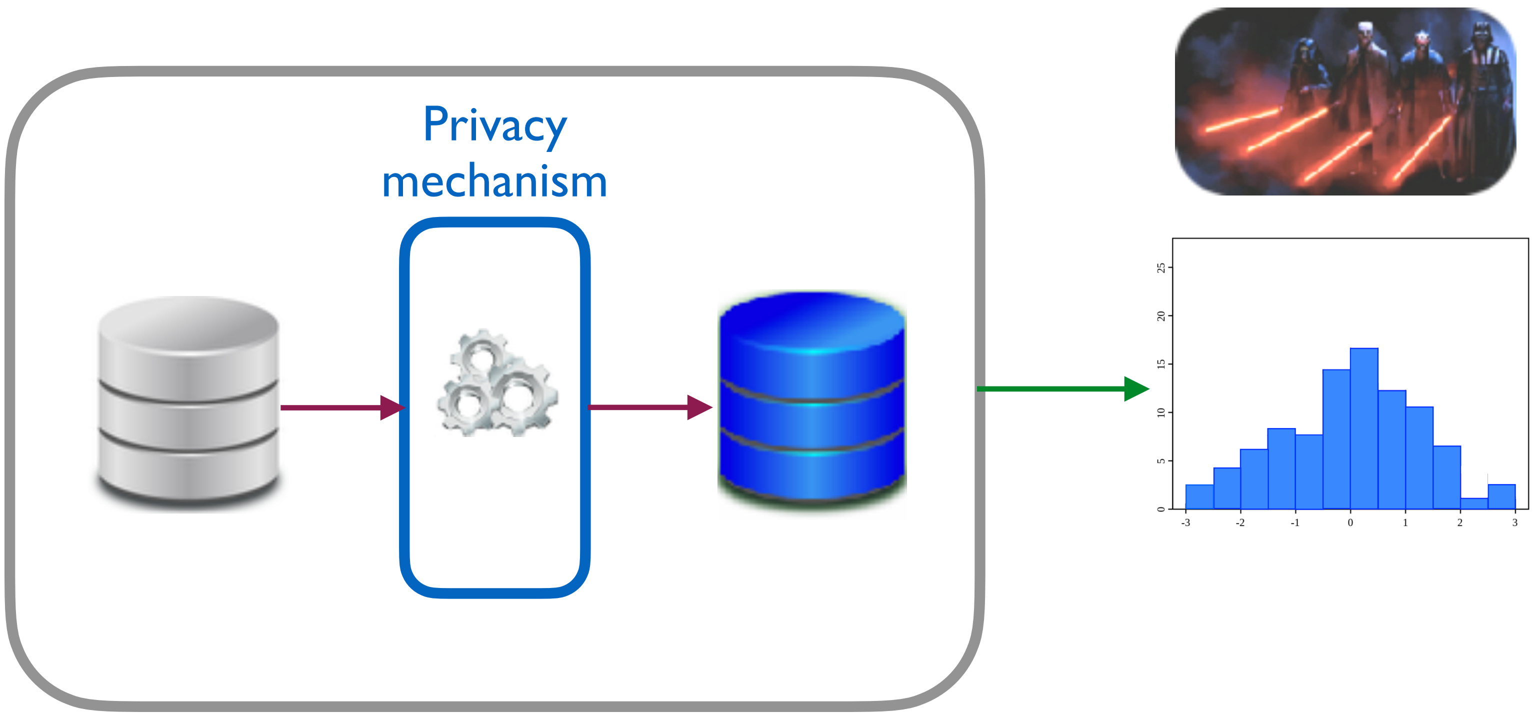

In the hybrid model the attacker only see the aggregated result of the obfuscation. This is achieved thanks to:

- users trusting their local data collector, or
- data collection is done via a technique called ``shuffling'' that re-orders the records, so that the relation with the owner is lost (anonymization)

Utility in the hybrid model

Utility in the hybrid model

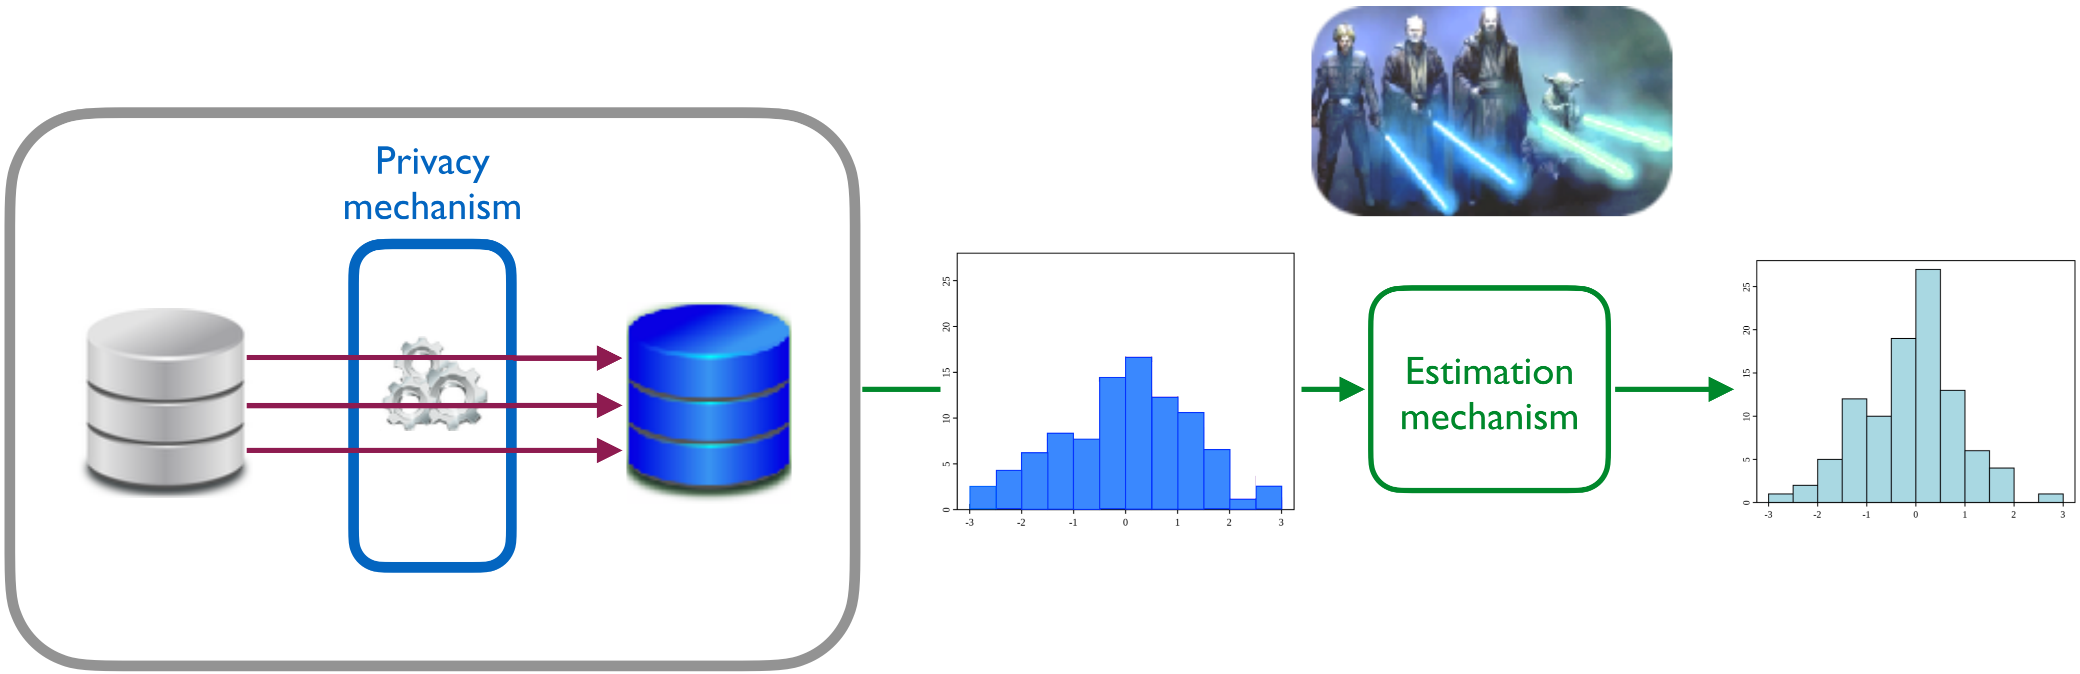

In the hybrid model, we can use the estimation mechanisms of LDP (e.g., the GIBU)

Utility in the central model

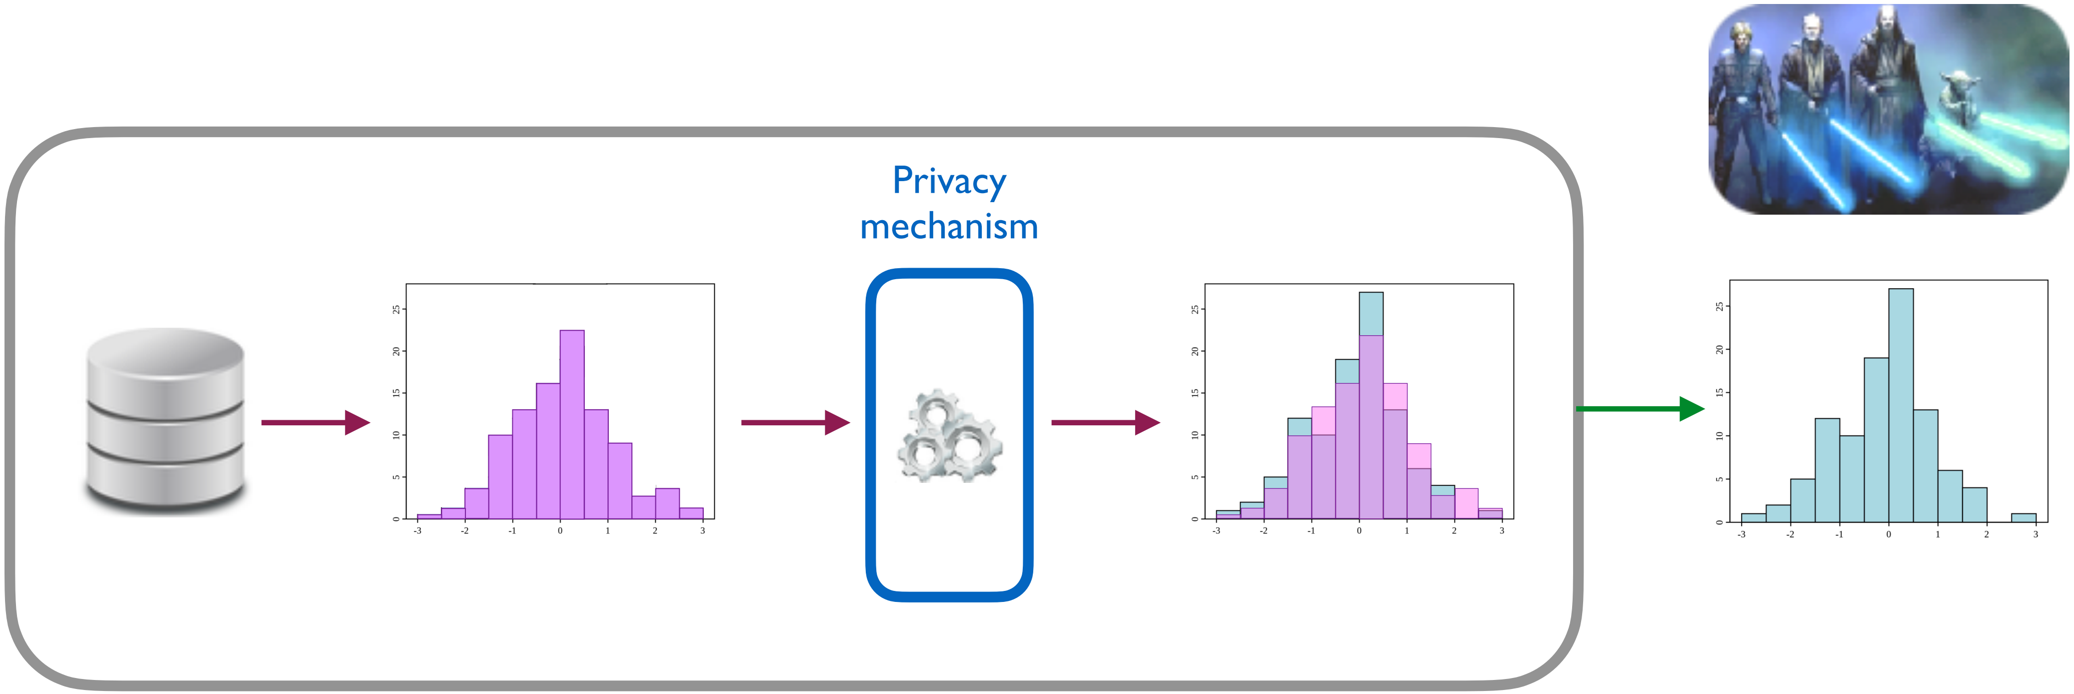

Utility: hybrid vs central

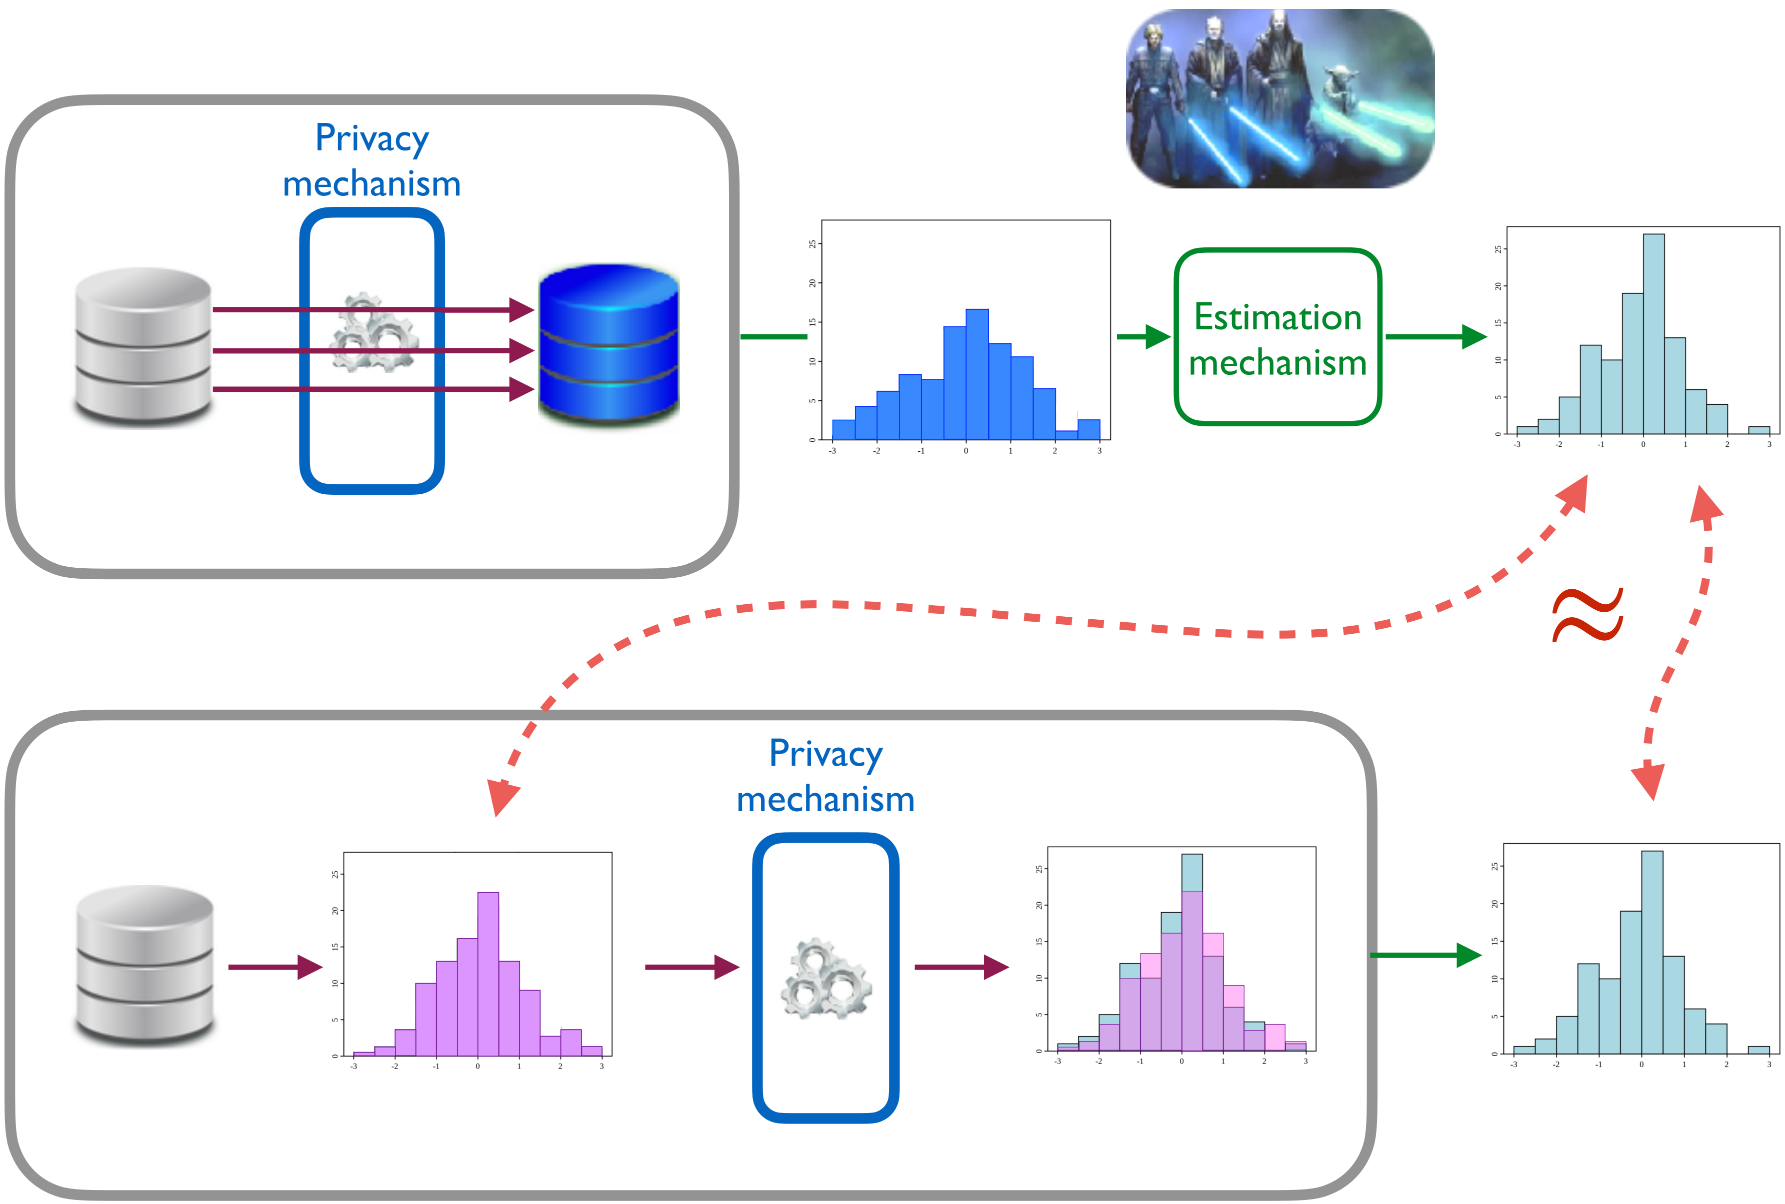

Advantage of hybrid wrt central: Compositionality

- GIBU is compositional (on any local mechanism)
- Inv applied to k-RR is compositional

Advantage of hybrid wrt central: Compositionality

- GIBU is compositional (on any local mechanism)
- Inv applied to k-RR is compositional

We could also compose the results of standard DP obfuscation (noise added to histogram), but when the mechanisms have different levels of privacy, we have observed experimentally that we not get the same estimation accuracy

d-privacy + IBU vs kRR + Inv

- IBU is more general: it can be applied to any privacy mechanism (and MLE is unique if the mechanism is invertible)
- d-privacy + IBU: better estimation accuracy if the distance between distributions takes into account the ground distance (e.g. , the Earth Movers' distance)
- kRR + Inv: more efficient

Thanks!

Questions ?