Robust channel ordering

- Given channels A and B on secret input X, the question of which leaks more will usually depend on the prior and the particular gain function used.
- Is there a **robust** ordering?
  - This could allow a stepwise refinement methodology.
  - This is arguably **indispensable** for security.
- For **deterministic** channels, a robust ordering has long been understood: the Lattice of Information [Landauer & Redmond '93].

The Lattice of Information

- A deterministic channel from X to Y induces a partition on $\mathcal{X}$: secrets are in the same block iff they map to the same output.
  - Example: $C_{country}$ maps a person x to the country of birth.

$C_{country}$'s partition:

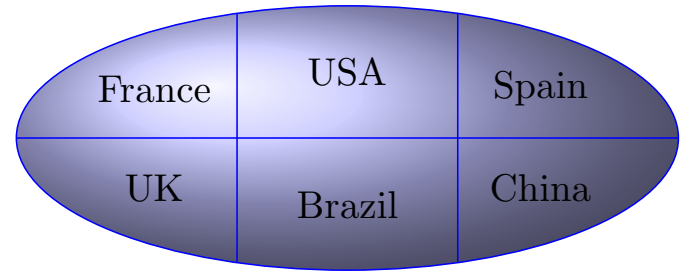

- Partition refinement $\sqsubseteq$: Subdivide zero or more of the blocks.
  - Example: C_{state} also includes the state of birth for Americans.

C_{state}'s partition:

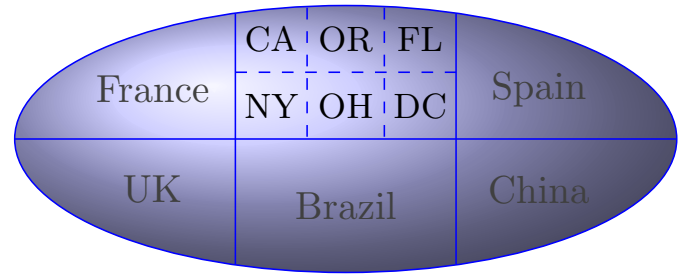

- $C_{country} \sqsubseteq C_{state}$

Partition refinement and leakage

- If A $\sqsubseteq$ B, then B leaks at least as much as A under **any** of the standard leakage measures (Shannon-, min-, and guessing entropy. The latter is the expected number of questions of the form "is S=s?" to figure out the secret entirely).
- Interestingly, the converse also holds:
  Theorem [Yasuoka &Terauchi '10, Malacaria '11]

  A $\sqsubseteq$ B

  iff

  A never leaks more than B on any prior, under **any** of the standard leakage measures
- Hence $\sqsubseteq$ is an ordering on deterministic channels with **both** a structural and a leakage-testing characterization.
- Can we generalize it to probabilistic channels?

Composition refinement

- Note that $C_{country}$ is the **composition** of C_{state} and C_{merge}, where C_{merge} post-processes by mapping all American states to USA.

$$C_{country} = C_{state}\ C_{merge}$$

- **Def:** $A \sqsubseteq_o B$ ("A is **composition refined** by B") if there exists a (post-processing) C such that A = BC.
- On deterministic channels, composition refinement $\sqsubseteq_o$ **coincides** with partition refinement $\sqsubseteq$.
  - So $\sqsubseteq_o$ **generalizes** $\sqsubseteq$ to probabilistic channels.

Strong leakage ordering

- **Def:** $A \leq_G B$ ("A never out-leaks B") if the g-leakage of A never exceeds that of B, for any prior π and any gain function g.

$A =$

	z_1	z_2
x_1	2/3	1/3
x_2	2/3	1/3
x_3	1/4	3/4

$B =$

	y_1	y_2	y_3
x_1	1/2	1/2	0
x_2	1/2	0	1/2
x_3	0	1/2	1/2

- **Def:** $A \leq_{\min} B$ if the min-entropy leakage of A never exceeds that of B, for any prior π.
- It turns out that $A \leq_{\min} B$, even though $A \not\leq_G B$

Relationship between $\sqsubseteq_o$ and $\leq_{\mathcal{G}}$

- **Theorem**: [Generalized data-processing inequality]

  If A $\sqsubseteq_o$ B then A $\leq_{\mathcal{G}}$ B.

  - Intuitively, the adversary should never prefer BC to B.

- **Theorem**: ["Coriaceous"]

  If A $\leq_{\mathcal{G}}$ B then A $\sqsubseteq_o$ B.

  - Conjectured for a long time. Proved by McIver et al. in 2014 using geometrical techniques (the Separating Hyperplane Lemma).

- So we have an ordering of probabilistic channels, with both **structural** and **leakage-testing** significance.

Exercises

Consider again the two programs A and B on a uniformly distributed, 64-bit x:

A. y = (x or 00000… 0111);

B. if (x % 8 == 0) then y = x; else y = 0;

8. Show that they both have min-entropy leakage 61 bits.

9. Define g_8, which allows 8 tries, and show that it makes A worse than B.

10. Define g_{tiger}, which gives a penalty for a wrong guess (allowing guess "⊥" to mean "don't guess") and show that it makes B worse. For simplicity, allow g_{tiger} to range in [-1,1]

Differential Privacy

- Differential privacy [Dwork et al.,2006] is a notion of privacy originated from the area of **Statistical Databases**

- **The problem:** we want to use databases to get statistical information (aka aggregated information), but without violating the privacy of the people in the database

The problem

- Statistical queries should not reveal private information, but it is not so easy to prevent such privacy breach.

- Example: in a medical database, we may want to ask queries that help to figure the correlation between a disease and the age, but we want to keep private the info whether a certain person has the disease.

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

Query:
What is the youngest age of a person with the disease?

Answer:
40

Problem:
The adversary may know that Don is the only person in the database with age 40

The problem

- Statistical queries should not reveal private information, but it is not so easy to prevent such privacy breach.

- Example: in a medical database, we may want to ask queries that help to figure the correlation between a disease and the age, but we want to keep private the info whether a certain person has the disease.

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

k-anonymity: the answer always partitions the space in groups of at least k elements

Alice	Bob
Carl	Don
Ellie	Frank

Correlation: Many-to-one

- Principle: Ensure that there are **many** secret values that correspond to **one** observable
- This is the general principle of most deterministic approaches to protection of confidential information (group anonymity, k-anonymity, ℓ-anonymity, cloacking, etc.)

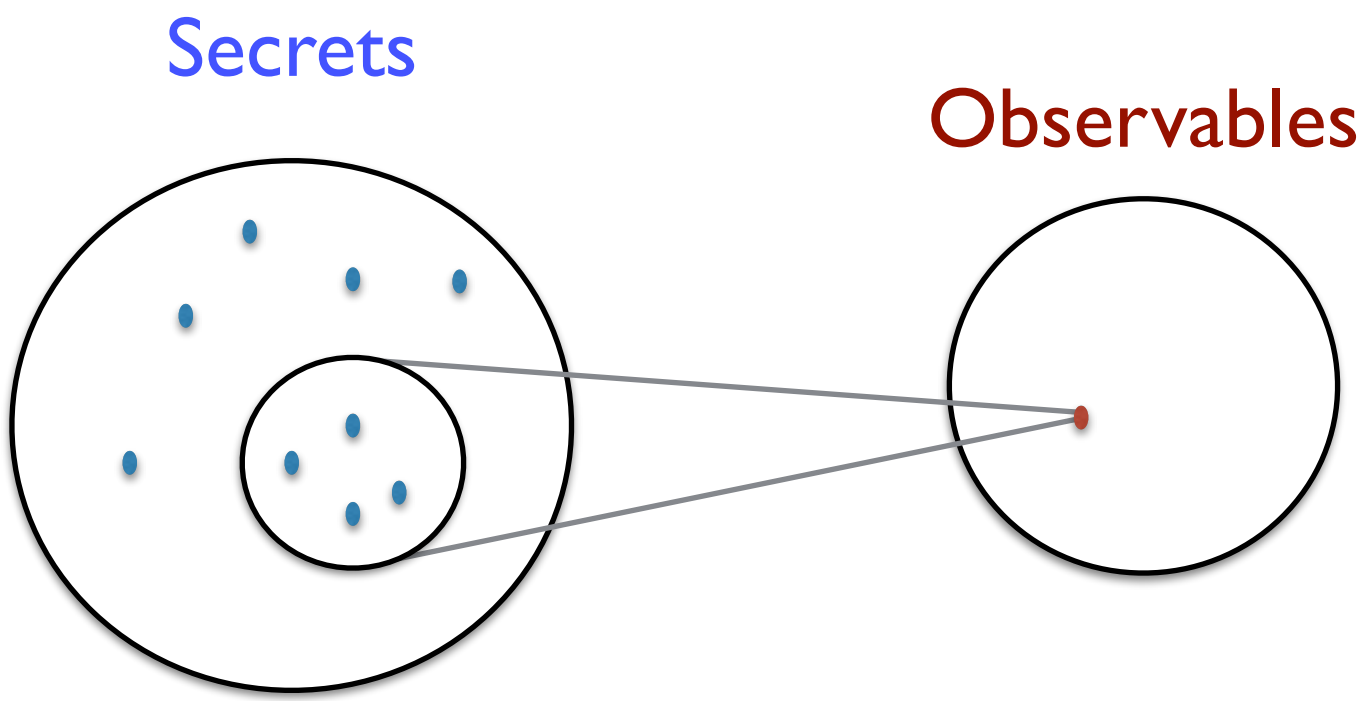

The problem

Unfortunately, the many-to-one approach is not robust under **composition**:

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

The problem of composition

Consider the query:

What is the minimal weight of a person with the disease?

Answer: 100

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

The problem of composition

Combine with the two queries:

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

Answers: 40, 100

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 the deterministic approaches (based on the principle of many-to-one): the combination of observations determines smaller and smaller intersections on the domain of the secrets, and eventually result in singletones

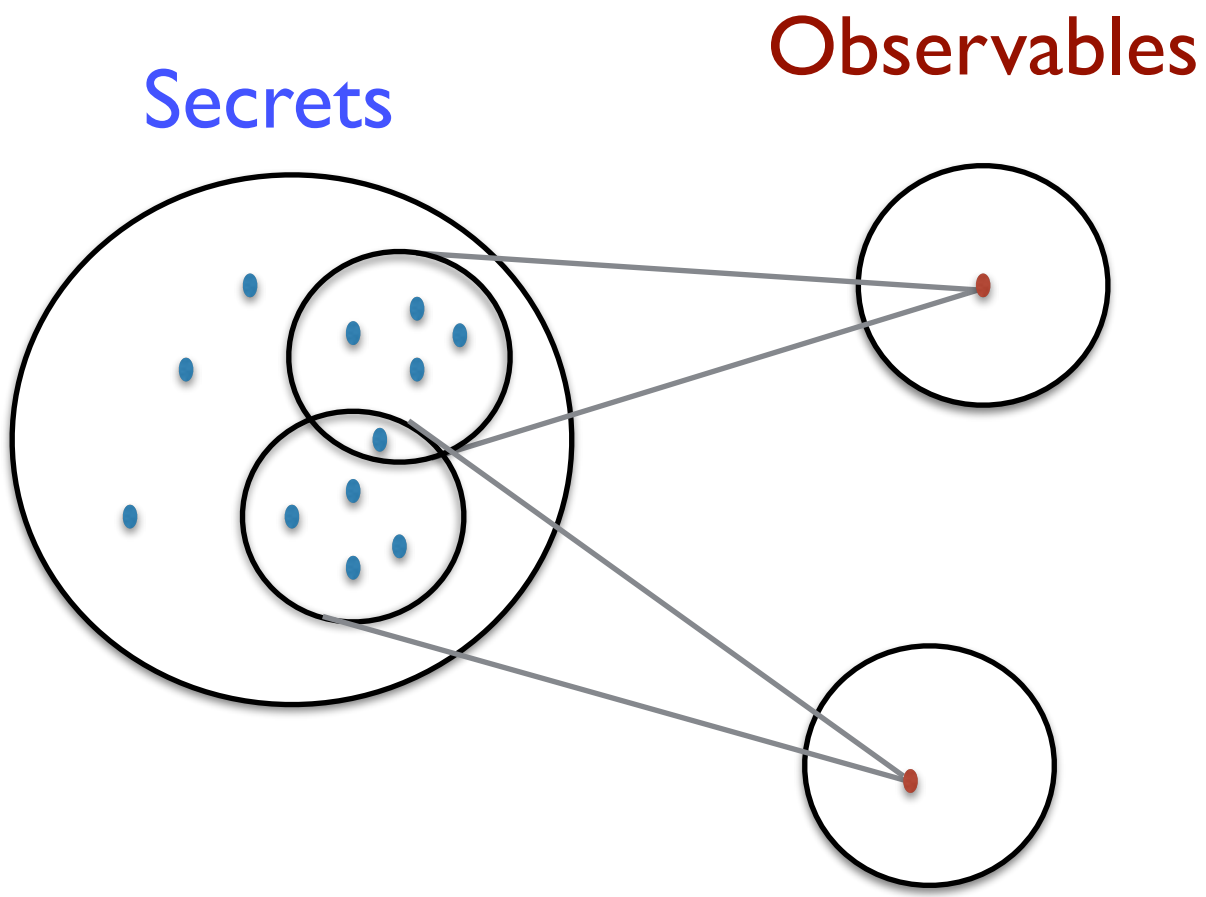

This is a general problem of the deterministic approaches (based on the principle of many-to-one): the combination of observations determines smaller and smaller intersections on the domain of the secrets, and eventually result in singletones

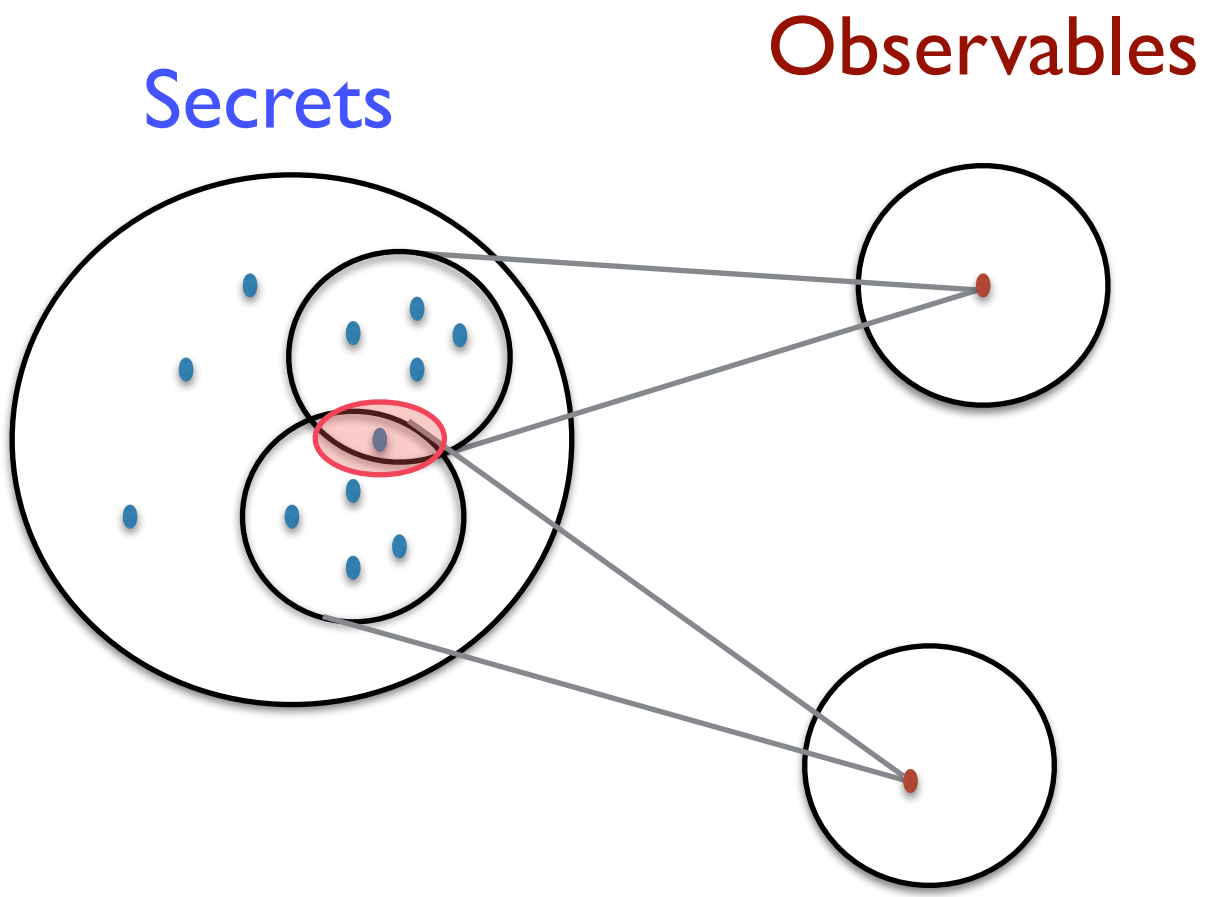

Composition attacks

Composition attacks are real!
For instance, in a recent paper, Narayanan et Smatikov showed that by combining the information of two popular social network (Twitter and Flickr) they were able to de-anonymize a large percentage of the users (about 80%) and retrieve their private information with only a small probability of error (12%).

De-anonymizing Social Networks, Arvind Narayanan and Vitaly Shmatikov. Security & Privacy '09.

Solution

Introduce some probabilistic noise on the answer, so that the answers of minimal age and minimal weight can be given also by other people with different age and weight

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

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

Combination of the answers

The adversary cannot tell for sure whether a certain person has the disease

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

Noisy mechanisms

- The mechanisms reports an approximate answer, typically generated randomly on the basis of the true answer and of some probability distribution
- 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 mechanism may provide different levels of utility.
- First of all, we need to formalize the notions of privacy and utility

Differential Privacy

- There have been various attempts to formalize the notion of privacy, but the most successful one is the notion of Differential Privacy, recently introduced by Dwork
- **Differential Privacy** [Dwork 2006]: a randomized function $\mathcal{K}$ provides ε-differential privacy if for all databases x, x' **which are adjacent (**i.e., differ for only one record), and for all $z \in \mathcal{Z}$, we have

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

- The idea is that the likelihoods of x and x' are not too far apart, for every S
- Differential privacy is robust with respect to composition of queries
- The definition of differential privacy is independent from the prior (but this does not mean that the prior doesn't help in breaching privacy!)

$\mathcal{K}$ can be seen as a noisy channel, in the information-theoretic sense from the domain $\mathcal{X}$ of databases to the domain $\mathcal{Z}$ of reported answers

Channel matrix

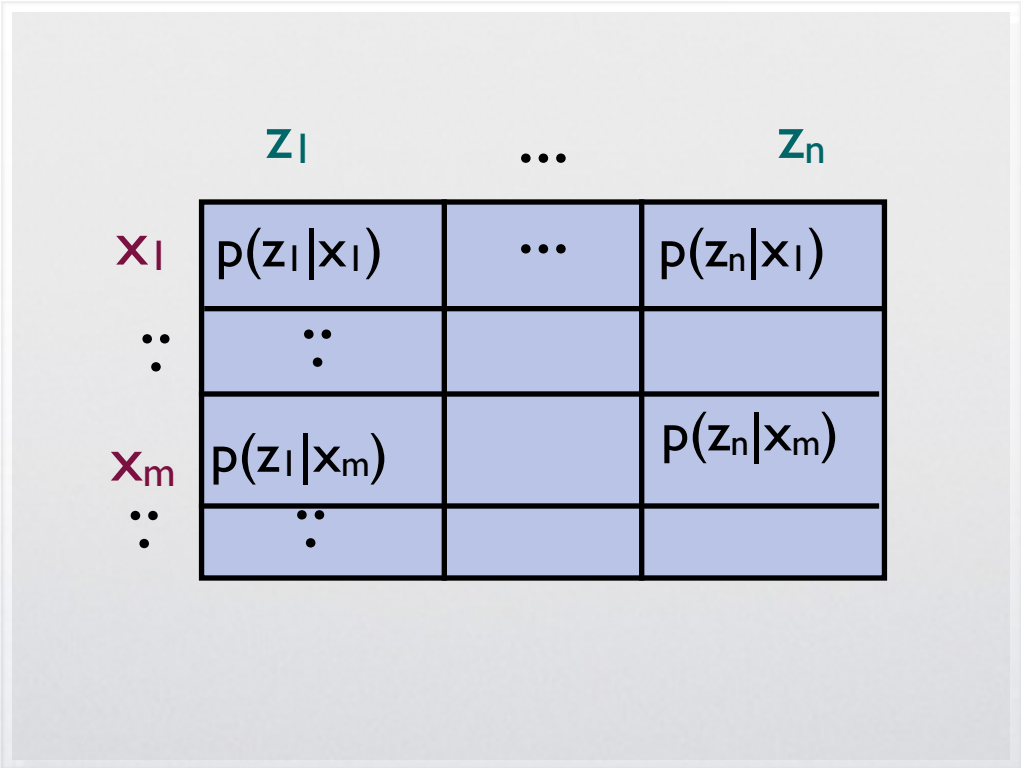

Differential privacy on the channel matrix

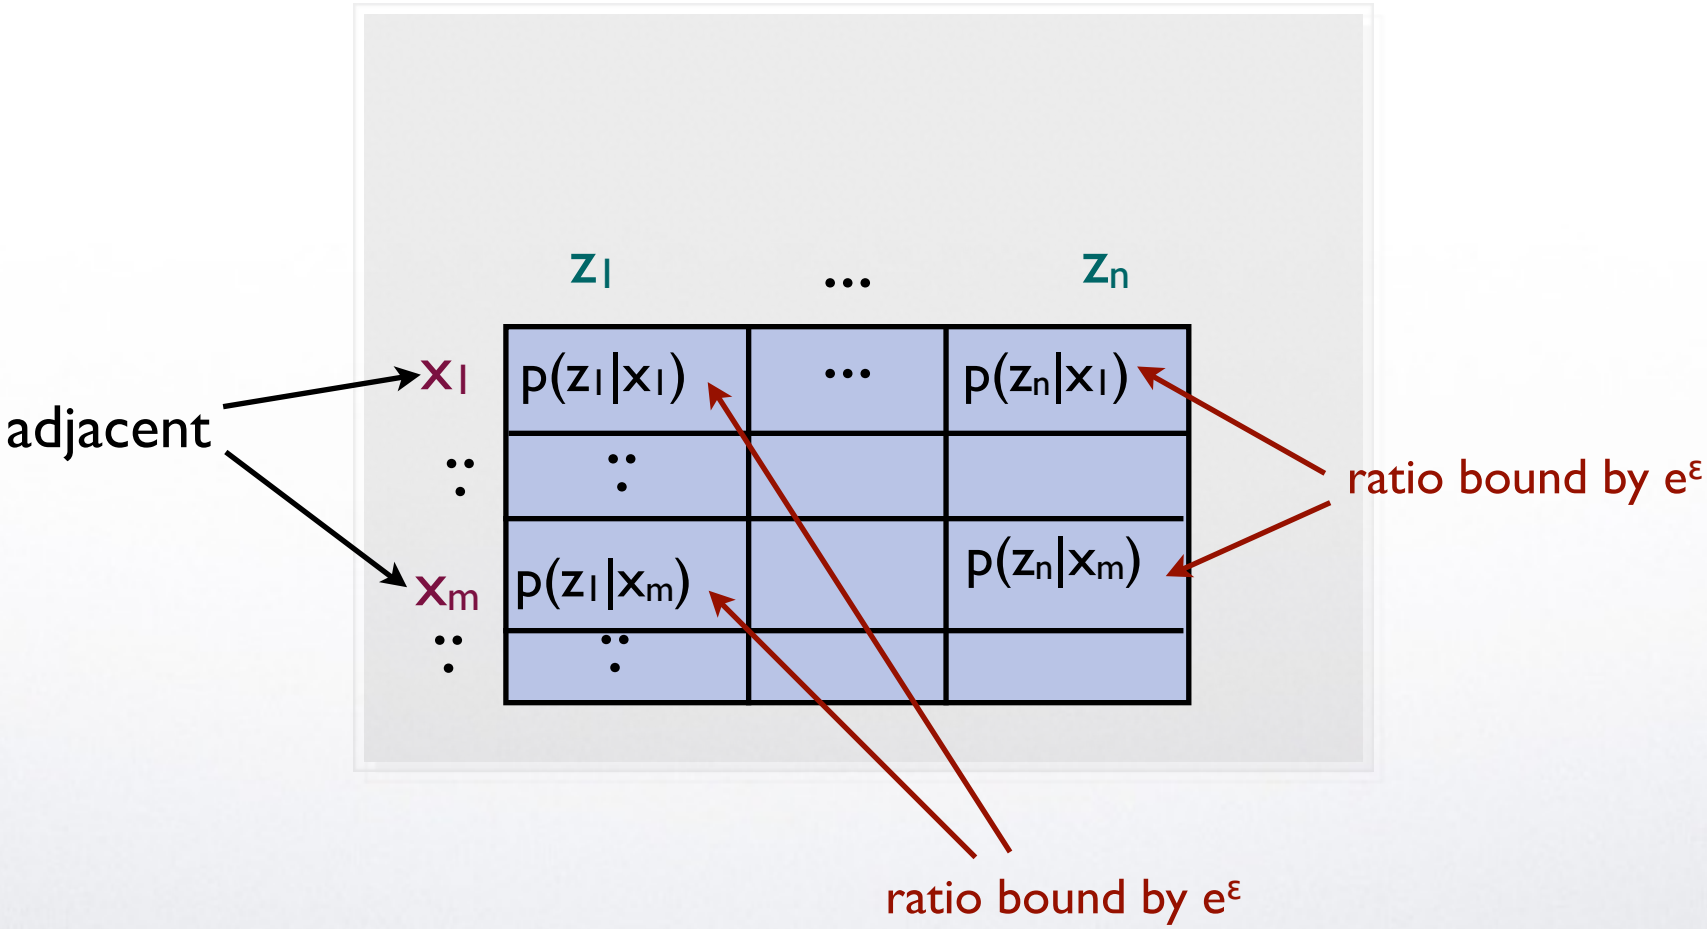

Differential Privacy: alternative definition

- Perhaps the notion of differential privacy is easier to understand under the following equivalent characterization.
- In the following, X_i is the random variable representing the value of the individual i, and $X_{\neq i}$ is the random variable representing the value of all the other individuals in the database

- **Differential Privacy, alternative characterization:** a randomized function $\mathcal{K}$ provides ε-differential privacy if:

$$\text{for all } x \in \mathcal{X}, z \in \mathcal{Z}, p_i(\cdot)$$

$$\frac{1}{e^\epsilon} \leq \frac{p(X_i = x_i | X_{\neq i} = x_{\neq i})}{p(X_i = x_i | X_{\neq i} = x_{\neq i} \wedge K = z)} \leq e^\epsilon$$

Question

- What is the basic difference between the protection guarantees offered by differential privacy, and those of Information flow?

Answer

- Information flow is an average measure (in the domain of privacy, it measures the "common good"). Differencial privacy is a worst-case measure, it protects every individual.
- Differencial privacy induces a bound on the information flow. The vice versa is not true
- In a sense, information flow represents the point of view of a company (e.g., ensurance company). Differential privacy represents the point of view of the individual.

Utility in Oblivious Mechanisms

- Given $f: \mathcal{X} \to \mathcal{Y}$ and $\mathcal{K}: \mathcal{X} \to \mathcal{Z}$, we say that $\mathcal{K}$ is oblivious if it depends only on $\mathcal{Y}$ (not on $\mathcal{X}$)

- If $\mathcal{K}$ is oblivious, it can be seen as the composition of f and a randomized mechanism $\mathcal{H}$ (noise) defined on the exact answers $\mathcal{K} = f \times \mathcal{H}$

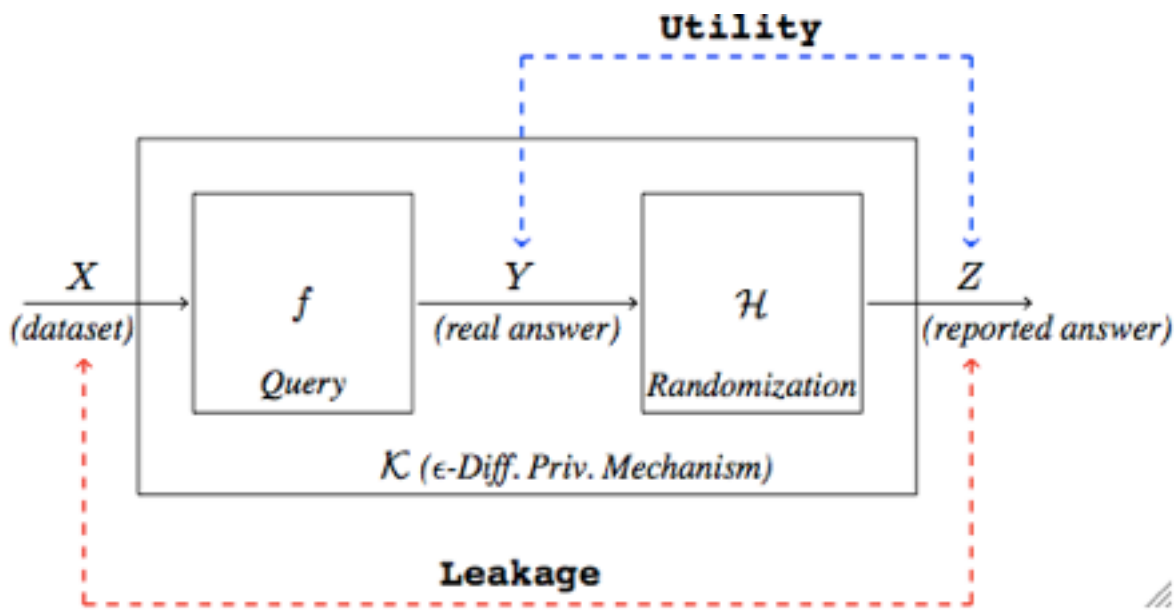

- Privacy concerns the information flow between the databases and the reported answers, while utility concerns the information flow between the correct answer and the reported answer

Utility

The reported answer, i.e. the answer given by the randomized function, should allow to approximate the true (i.e. the exact) answer to some extent

Z = reported answer; *Y* = exact answer

Utility:

$$\mathcal{U}(Y,Z) = \sum_{y,z} p(y,z)\, gain(y, remap(z))$$

In this formula, the remap is chosen so to maximize the result. The remapping allows the user to use side information (i.e. a the priori pb) to maximize utility.

For instance, if the reported answer is 20, but I know that the minimum possible answer is 21, then I will remap the answer to 21

Example: **binary gain function**:

$$gain(y_1, y_2) = \begin{cases} 1 & y_1 = y_2 \\ 0 & y_1 \neq y_2 \end{cases}$$

In the binary case the utility is **the expected value of the probability of success** to obtain the true answer (note the correspondence with the min-vulnerability)

In general, the gain function is anti-monotonic with the distance between the real value and the reported (and remapped) value.

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
- A mechanism is universally optimal if it is optimal for all priors π
- Note that the level of privacy does not depend on the prior, but the utility (in general) does.
- The optimal mechansim can be computed with linear optimization techniques

A typical ε-differentially-private mechanism: Laplacian noise

- Randomized mechanism for a query $f: \mathcal{X} \to \mathcal{Y}$.
- A typical randomized method: **add 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:

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

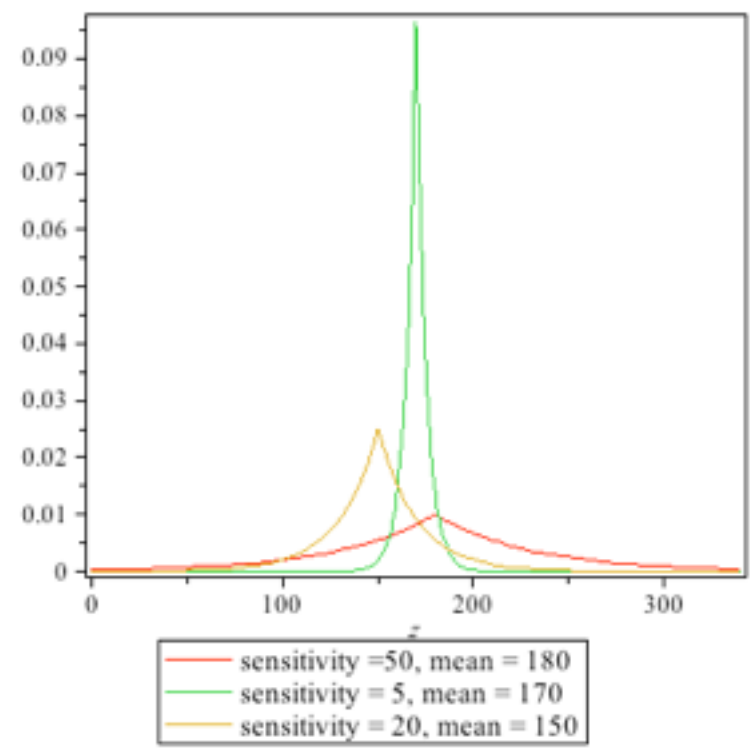

The geometric mechanism

- The geometric mechanism is a sort of discrete Laplacian.
- Assume that Y and Z are sets of integers. In the geometric mechanism, the probability distribution of the noise is:

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

- where c is a normalization factor, defined so to obtain a probability distribution, and Δf is the sensitivity of query f

- Note that it does not make much sense to report answers outside Y. If Y is an interval [a,b], we can truncate the mechanism, i.e., set Z = Y, and transfer on the extremes a and b all the probability that (according to the formula above) would fall outside the interval, to the left or to the

Counting Queries

- A counting query is a query of the form: How many individuals (tuples) in the database satisfy the property $\mathcal{P}$?

- The sensitivity of a counting query is 1

Exercise

- Define the noise probability distribution for the geometric mechanism for a counting query
- Truncate the above mechanism to the left of 0 (because for a counting query it does not make sense to report negative answers)

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 utility 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 $(\mathcal{X}, \sim)$

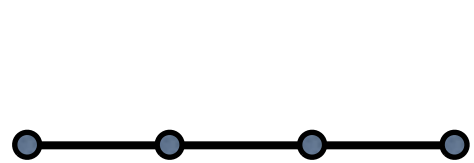

ok

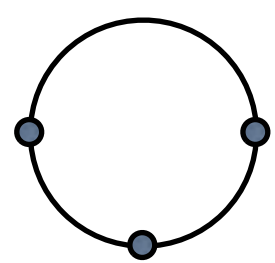

not ok

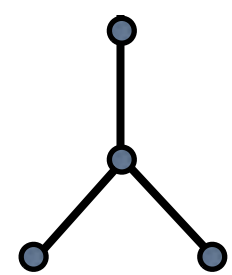

not ok

Further reading

TUTORIAL

Mário S. Alvim, Miguel E. Andrés, Konstantinos Chatzikokolakis, and Catuscia Palamidessi.
Quantitative Information Flow and Applications to Differential Privacy.
In A. Aldini and R. Gorrieri, editors, Foundations of Security Analysis and Design VI – FOSAD Tutorial Lectures, LNCS 6858, pages 211–230. Springer, 2011.

RESEARCH PAPERS

Boris Köpf and David Basin.
An information-theoretic model for adaptive side-channel attacks.
Proc. of CCS, pp. 286–296, ACM, 2007.

Geoffrey Smith.
On the Foundations of Quantitative Information Flow
Proc. of FOSSACS, LNCS 5504, pp. 288–302, Springer, 2009.

Mário S. Alvim and Konstantinos Chatzikokolakis and Catuscia Palamidessi and Geoffrey Smith.
Measuring Information Leakage Using Generalized Gain Functions.
Proc. of CSF, pp. 265-279, IEEE, 2012.

Cynthia Dwork. A firm foundation for private data analysis.
Communications of the ACM, 54(1):86–96, 2011.

Thank you !