Differential Privacy

Amelia Andersdotter (ARTICLE19)[1]

Christoffer Långström (Uppsala University)[2]

IETF 104 Prague, 2019

[1] amelia@article19.org

[2] https://github.com/chrislangst

The Basic Problem

RFC6973 contains Privacy Considerations for IETF protocols. Differential Privacy is a way of remedying specific privacy threats, such as identifiability (RFC6973, 5.2.2) and secondary use (RFC6973, 5.2.3), while simultaneously providing a value to the degree of success with which these threats are remedied.

Differential Privacy aims to provide an individual with *plausible deniability*, in the sense that such an individual should be able to *deny* being part of a database (Dwork & Roth, 2009).

So formally...

$$\mathbb{P}[M(D)] \leq e^{\varepsilon}\mathbb{P}[M(D')] + \delta$$

where M is a *mechanism* applied to a database D and a database D', and D' differs from D only by a small amount, and $\mathbb{P}$ is a distribution over the possible outputs of the mechanism M as applied to D or D'.

We'll use the short-form M for $M(D)$ and M' for $M(D')$.

(ε, δ)-differential privacy: unpacking it

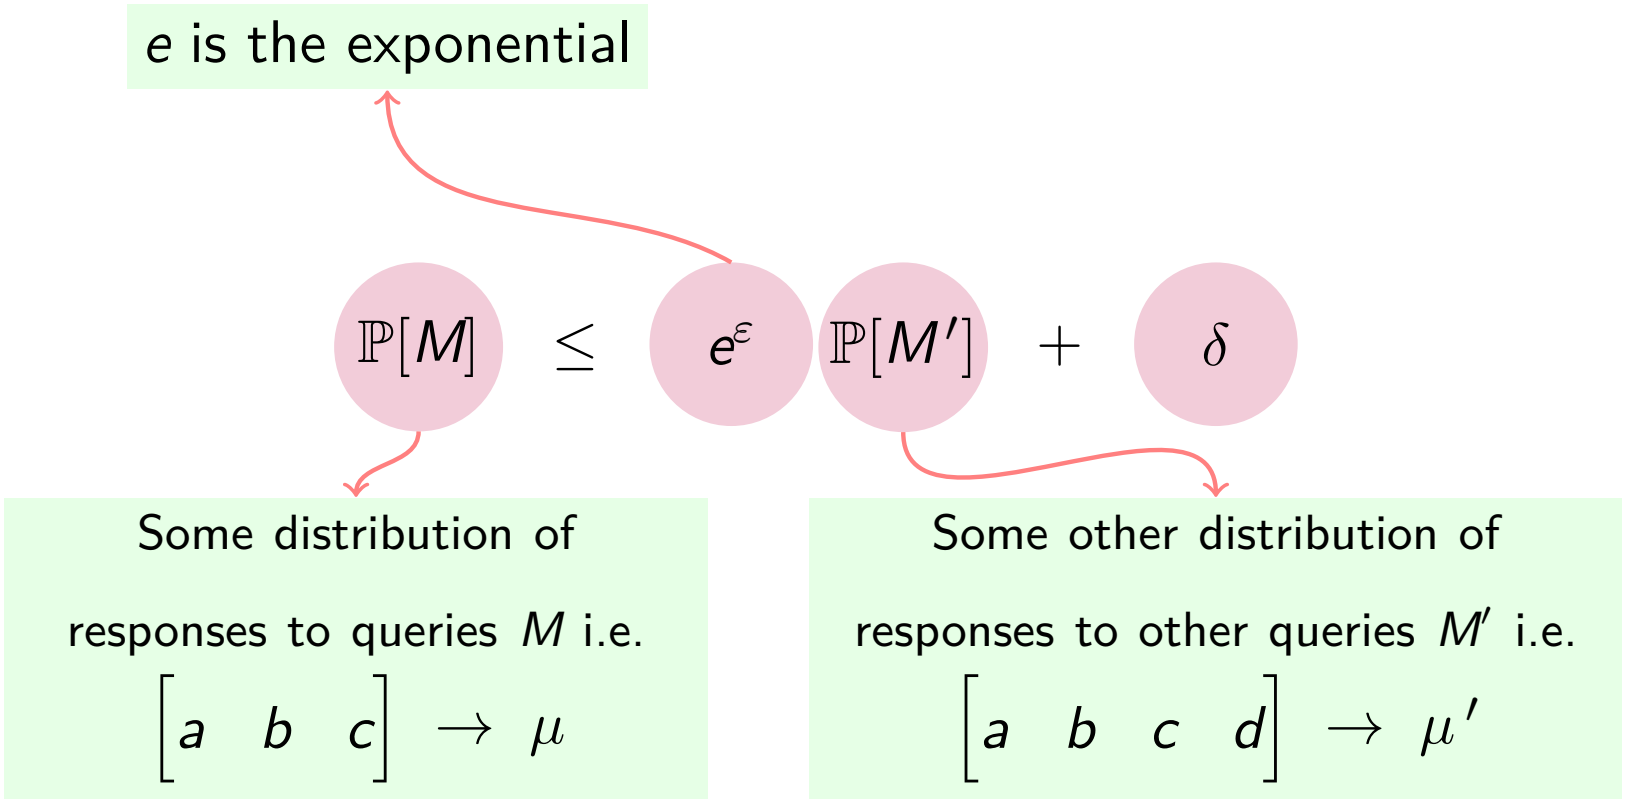

If M tries to find the mean (the sum of the values of all data points divided by the total number of data points), the idea is that M and M' should be *sufficiently* similar up to (ε, δ).

(ε, δ)-differential privacy: unpacking it

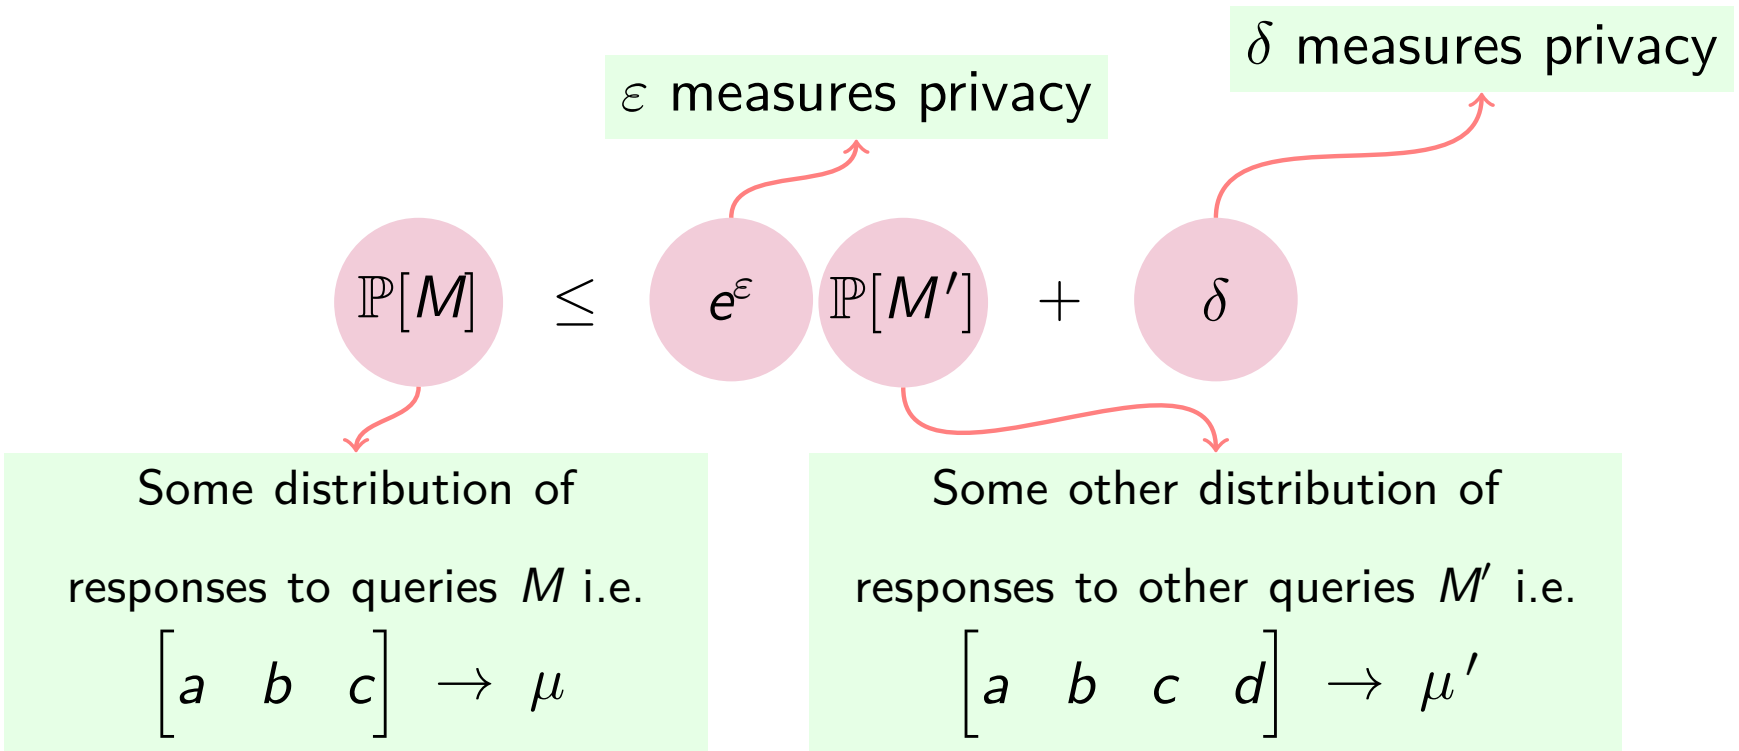

$$\mathbb{P}[M] \leq e^{\varepsilon} \mathbb{P}[M'] + \delta$$

When δ is small, and $\varepsilon \to 0$, then $\mathbb{P}[M] \approx \mathbb{P}[M']$.

So, differential privacy allows us to *quantify* the degree to which we are able to preserve the identity of the originating individual for a particular piece of data, when studying the results of a query over the population to which this individual pertains.

(ε, δ)-differential privacy in practice

Method 1: Perturb the data in response to the query.

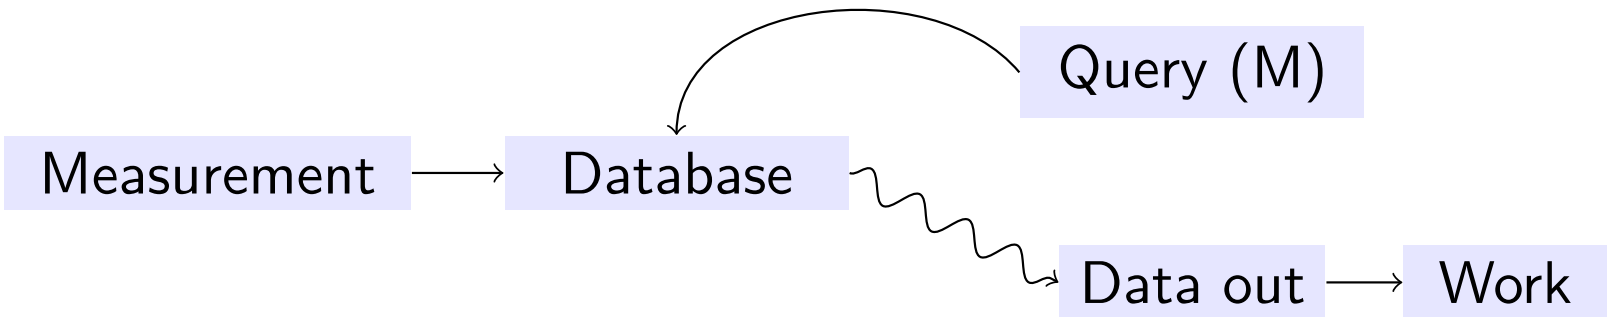

Method 1: continuation

Common methods

Adding noise to output (Gaussian or Laplacian).

Drawbacks

Worsens estimator quality (Duchi et al, 2016).

Repeated queries cancel out noise ("privacy budget").

Need to trust the database holder.

(ε, δ)-differential privacy in practice

Method 2: Perturb the measurement.

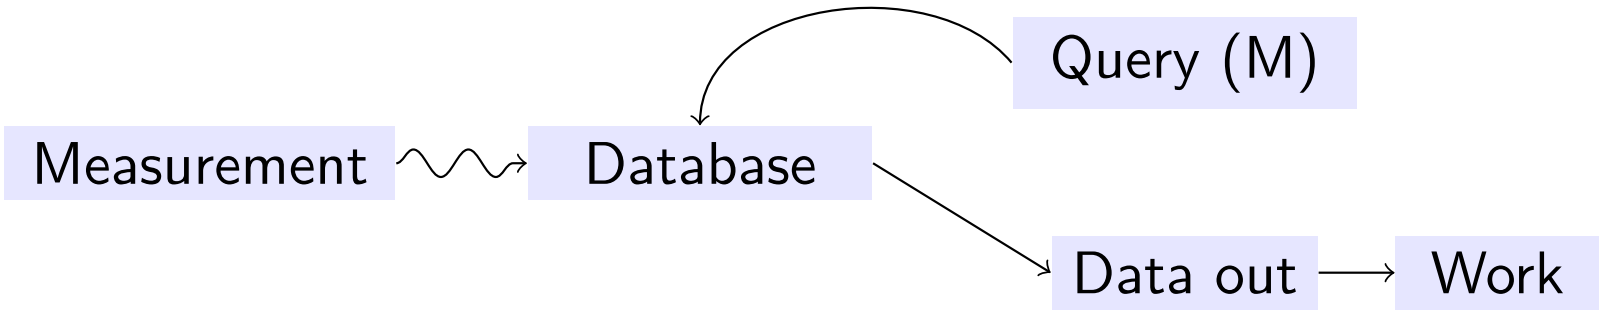

Method 2: continuation

Common methods

Removing/encrypting identifiers.

Swapping data between different identifiable flows.

Randomizing responses.

Drawbacks

Worsens estimator quality (think filtering problem).

Need to trust the entity making measurements.

(ε, δ)-differential privacy is not the only privacy needed!

(ε, δ)-differential privacy specifically deals with the case when we are trying to protect the identity of an originating individual for a particular piece of a data in a data set.

Data sanitization and security still important!

(ε, δ)-differential privacy is not the only measure!

A survey in 2015 found hundreds of metrics, adapted for a range of different threats (Wagner & Eckhoff, 2015).

Another study demonstrates that the identity of any particular individual in a sufficiently big data set is already statistically (ε, δ)-protected (Duan, 2009).

(ε, δ)-differential privacy is for API:s

Nevertheless, some ideas:

- Protocols which provide predictably false data (i.e. for security analytics, or similar) according to some random distribution?
- Example: QUIC Spin bit? Other protocol data which is convenient for feature indication or security analytics, but not for establishing the communication flow?

Questions?

References / Further Reading

Duan, Y., “Privacy without noise”, CIKM ’09 Proceedings of the 18th ACM conference on Information and knowledge management (2009).

Duchi, J., et al, “Minimax Optimal Procedures for Locally Private Estimation”, arXiv:1604.02390 [math.ST]

Dwork, C., Roth, A., “The Algorithmic Foundations of Differential Privacy” (2014)

RFC6973, Privacy Considerations for Internet Protocols, A. Cooper et al (July 2013)

Wagner, I., Eckhoff, D., “Technical Privacy Metrics: a Systematic Survey”, arXiv:1512.00327 [cs.CR]