CS208: Applied Privacy for Data Science
The Local Model: Foundations

James Honaker & Salil Vadhan

School of Engineering & Applied Sciences
Harvard University

March 25, 2019

Source Material

We will mostly follow slides from a talk of Adam Smith, "Local-ish Models for Statistical Differential Privacy," Workshop on Differential Privacy Meets Multiparty Computation (DPMPC), Boston University, June 5, 2018

[Change filename extension from .pptx to .pdf to open!]

Defining Privacy

- See Smith slides 4-6 (but ignore formal def of slide 5).

- Def: a protocol is ε-local DP if each party's local randomizer Q_i is an ε-DP interactive mechanism for 1-row databases.
  - stronger than Smith's formulation

- Q: What does it mean for an interactive mechanism to be DP?
  - let's return to this after seeing some local DP examples

Randomized Response

[Smith, slides 9-11]

ε-locally DP protocol that

- Estimates count/sum of a bounded function to $\pm O\left(\frac{\sqrt{n}}{\varepsilon}\right)$.
  - Q: how to prove this error bound?
- Estimates "statistical queries" (means/avgs) to $\pm O\left(\frac{1}{\varepsilon\sqrt{n}}\right)$.
  - Q: how to use RR for fractional-valued functions?

- Worse than centralized DP, but still useful.

- This is best possible for ε-local DP.

Histograms

[Smith, slides 12-13]

- Expected error on each bin is $\pm O\left(\frac{\sqrt{n}}{\varepsilon}\right)$.
- Expected max error over all D bins is $\pm O\left(\frac{\sqrt{n \cdot \log D}}{\varepsilon}\right)$.
- Sophisticated algorithmic ideas to get computational complexity sublinear in D.
- Note that Smith slide 13 is missing $\{0,1\} \leftrightarrow \{\pm 1\}$ translation.

Defining Privacy

- See Smith slides 4-6.

- Def: a protocol is ε-local DP if each party's local randomizer Q_i is an ε-DP interactive mechanism for 1-row databases.
  - stronger than Smith's formulation

- Q: What does it mean for an interactive mechanism to be DP?
  - let's return to this after seeing some local DP examples

DP for Interactive Mechanisms

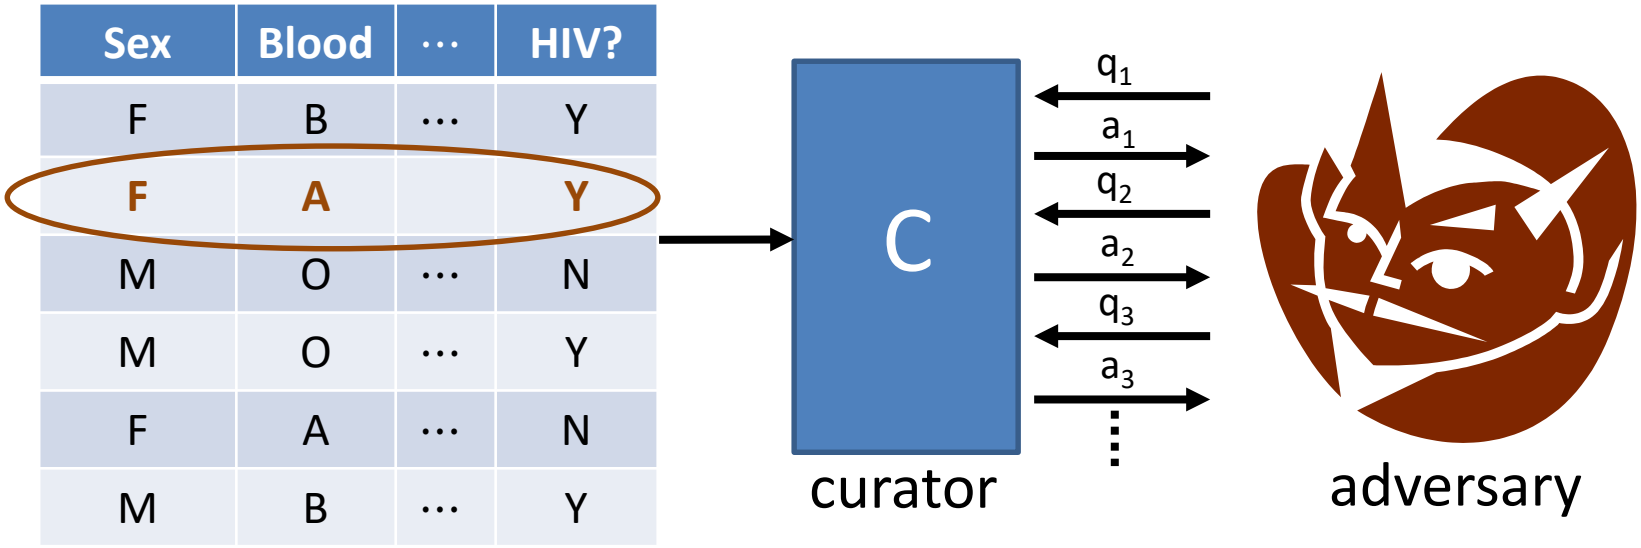

1st Attempt: for all D, D' differing on one row, all $q_1,\ldots,q_t$, all T

$$\Pr[C(D, q_1, \ldots, q_t) \in T] \leq e^{\varepsilon} \cdot \Pr[C(D', q_1, \ldots, q_t) \in T] + \delta$$

vectors of answers $a_1, \ldots, a_t$

DP for Interactive Mechanisms

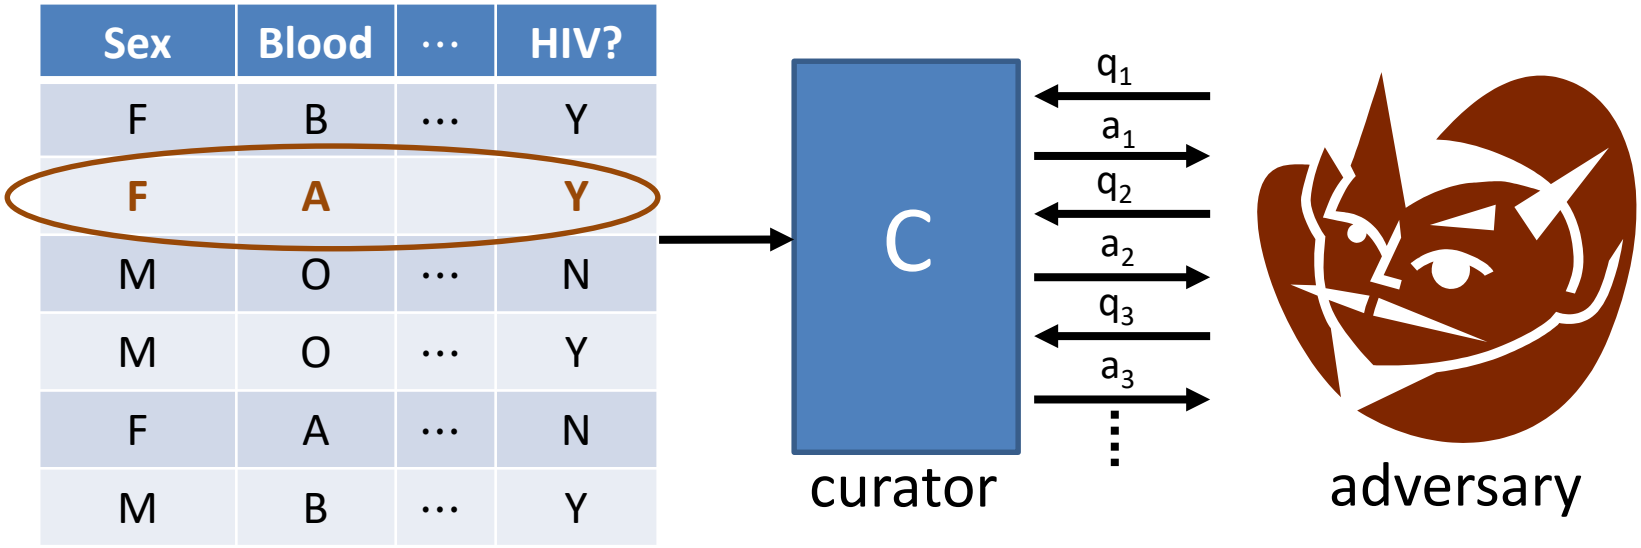

Better: for all D, D' differing on one row, all adversarial strategies A

$$\Pr[A \text{ outputs YES after interacting w}/C(D)] \leq e^{\varepsilon} \cdot \Pr[A \text{ outputs YES after interacting w}/C(D')] + \delta$$

Fact: composition thms for DP yield interactive DP in this sense.

(advanced/optimal comp. requires privacy params to be non-adaptive [Rogers et al. `16].)

Local DP

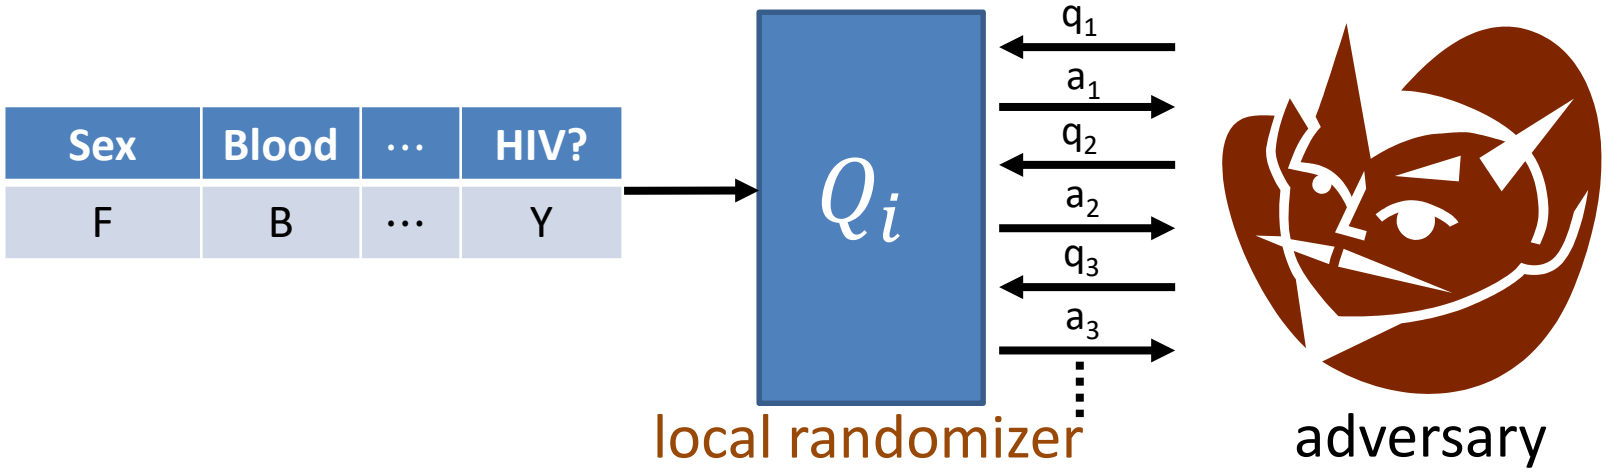

Require: for all i, x_i, x_i' ~~differing on one row,~~ all strategies A

$$\Pr[A \text{ outputs YES after interacting w}/Q_i(x_i)]$$
$$\leq e^{\varepsilon} \cdot \Pr[A \text{ outputs YES after interacting w}/Q_i(x_i')] + \delta$$

Multiparty DP

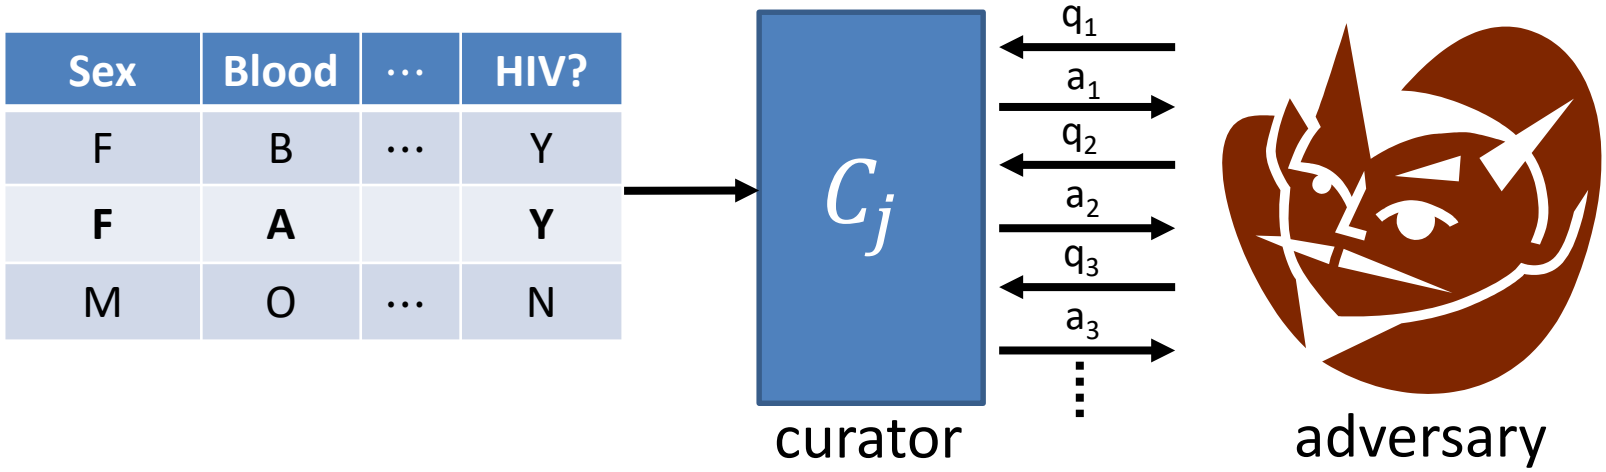

Intermediate model where there are k curators C_j, each with a dataset $D^{(j)}$ on n_j subjects, interacting to carry out a joint analysis.

Require: for all $D^{(j)}$, $D^{(j)\prime}$ differing on one row, all adversarial A

$$\Pr[A \text{ outputs YES after interacting w/} C_j(D^{(j)})]$$
$$\leq e^{\varepsilon} \cdot \Pr[A \text{ outputs YES after interacting w/} C_j(D^{(j)\prime})] + \delta$$

Adversarial Models for DP

Can quantify over restricted sets of adversaries for DP protocols

- Computationally bounded adversaries
  - E.g. "all polynomial-time A".
  - Allows for using cryptography, secure multiparty computation
- "Honest-but-curious" (a.k.a. "semi-honest") adversaries
  - Follows the protocol properly but carries out arbitrary computation afterwards.
  - May be reasonable for Google, Apple as local DP aggregators trying to protect against subpoenas.

Adversarial Models for DP (cont.)

- Anonymous participants
  - Aggregator doesn't know which data is from which participants (e.g. due to use of a "mix-net" like Tor)
  - Can boost privacy protections of local DP!

- Threshold adversaries
  - Adversary that controls up to t parties (e.g. "honest majority").

- Another choice: should correctness/accuracy hold even when parties deviate from the protocol or drop out?

Local vs. Centralized DP

- Local DP protocols provably have lower accuracy for counts/averages than centralized DP protocols.
  - $\Theta(1/\varepsilon\sqrt{n})$ error vs. $\Theta(1/\varepsilon n)$.
  - Successful deployments have very large n (Google, Apple).

- For some tasks there is an exponential separation.
  - E.g. “learning parities”
  - See Smith slides 15-21.

- Gap can be closed by relaxing adversarial model (e.g. anonymous participants, computationally bounded adversaries) and using crypto/infrastructure (secure MPC, mix-nets).