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

James Honaker, Priyanka Nanayakkara, Salil Vadhan

School of Engineering & Applied Sciences
Harvard University

April 7, 2025

Housekeeping

- Detailed project descriptions due this Friday!
  - You can still change your topic, eg based on the feedback we gave.
  - Come to OH to discuss!

- No pset due this week, hw8b due Fri 4/18.

- Other project deadlines:
  - Full project paper: Wed 4/30
  - Revision of paper: Thu 5/8
  - Poster session: Thu 5/8, 9am-12pm in the SEC.
  - 3 late days per group on project deadlines.

Class-wide exercise

- Privately:
  - Write down your preference: vanilla (1) or chocolate (0)
  - Choose a random number from 1-4 using Google, www.random.org, or by tossing a coin twice.

- Class Poll: Salil will ask everyone to report their preference
  - If your random number is 1,2,3: report truthfully
  - If your random number is 4: report falsely

Group Exercise

1. Use the reported counts for vanilla and chocolate to compute an unbiased estimator $\hat{\mu}$ of the fraction of people in the class who prefer vanilla.

   Hint: write a formula for the expectation of the number V_{rep} of people who report vanilla in terms of the number v of people who actually prefer vanilla and $n - v$.

2. What is the standard deviation of your estimator?

3. For what ε is this method ε-DP? (Consider the release to be collection of everyone's "noisy" reports.)

Group Exercise: Solution

1. Use the reported counts for vanilla and chocolate to compute an unbiased estimator $\hat{\mu}$ of the fraction of people in the class who prefer vanilla.

$$\mathrm{E}[V_{\mathrm{rep}}] = \frac{3}{4} \cdot v + \frac{1}{4} \cdot (n - v) = \frac{v}{2} + \frac{n}{4}$$

$$\hat{\mu} = \frac{2}{n} \cdot V_{\mathrm{rep}} - \frac{1}{2}$$

2. What is the standard deviation of your estimator?

$$\sigma^2[\hat{\mu}] = \frac{4}{n^2} \cdot \sigma^2[V_{\mathrm{rep}}] = \frac{4}{n^2} \cdot n \cdot \frac{3}{4} \cdot \frac{1}{4}$$

$$\sigma[\hat{\mu}] = \sqrt{{}^{3}/_{4n}}$$

3. For what ε is this method ε-DP? (Consider the release to be collection of everyone's "noisy" reports.)

$$\varepsilon = \ln\left(\frac{3/4}{1/4}\right) = \ln 3 \approx 1.1$$

Individual Survey

Compare the method we just saw for doing a DP count to a standard noise-addition mechanism (e.g. the Laplace mechanism).

1. What is an advantage of the method we just used?

2. What is a disadvantage of the method we just used?

In either case, if you don’t think there’s an advantage or disadvantage, give your intuition.

Central Model vs Local Model

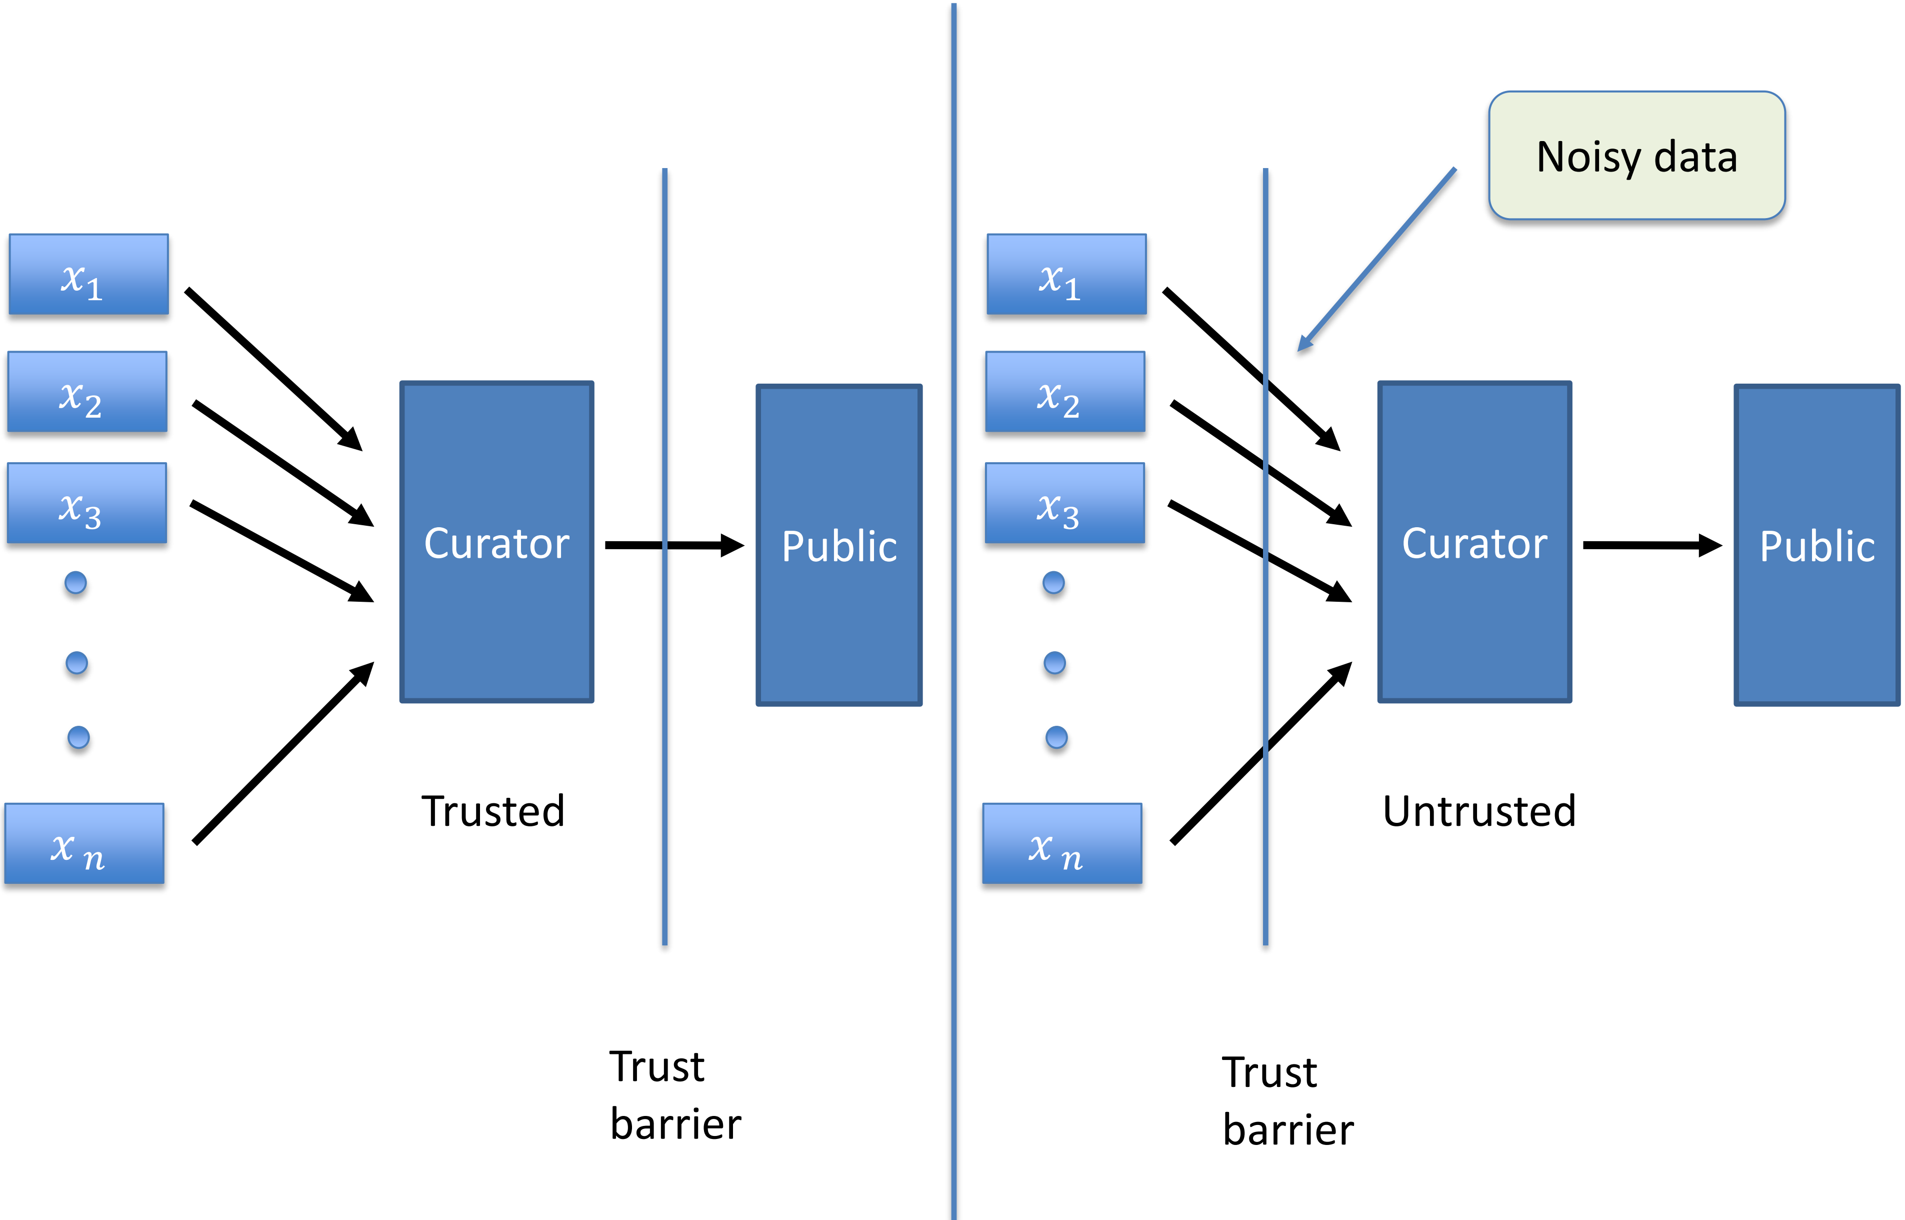

Local Differential Privacy

Sex	Blood	⋯	HIV?
F	B	⋯	Y

A single user!

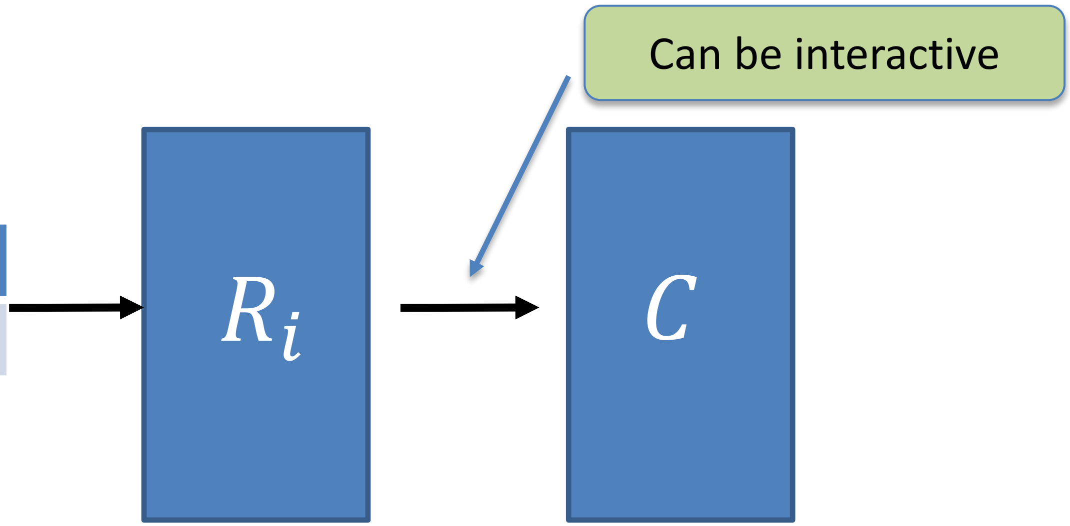

local randomizer

$R\colon \mathcal{X} \to \mathcal{Y}$ is (ε, δ)-locally differentially private (LDP) if for all $x, x' \in \mathcal{X}, S \subseteq \mathcal{Y}$

$$\Pr[R(x) \in S] \leq e^{\varepsilon} \cdot \Pr[R(x') \in S] + \delta$$

That is, a protocol is ε-LDP if each party's local randomizer R_i is an ε-DP mechanism for *1-row databases*.

Interactive Local DP

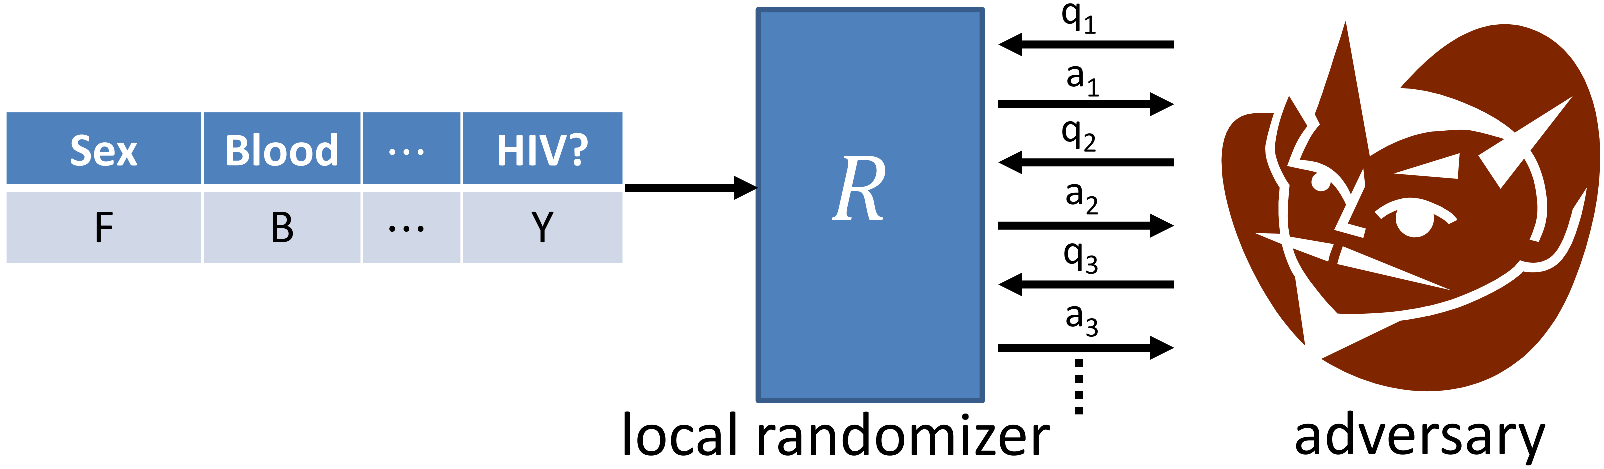

Require: for all x, x', all adversarial strategies A

$$\underbrace{\text{View}_A\big(A \leftrightarrow M(x)\big)}\approx_\varepsilon \underbrace{\text{View}_A(A \leftrightarrow M(x'))}$$

Everything A sees (its internal randomness & query answers)

Equivalently: $\forall\, A \;\; \Pr[A \text{ outputs "In" after interacting w}/M(x)]$
$\leq e^{\varepsilon} \cdot \Pr[A \text{ outputs "In" after interacting w}/M(x')]$

Randomized Response

[Warner'65]

For $x_i \in \{0,1\}$, $\mathrm{RR}_\varepsilon(x_i) = \begin{cases} x_i & \text{w.p. } \frac{e^\varepsilon}{1+e^\varepsilon} \\ 1 - x_i & \text{w.p. } \frac{1}{1+e^\varepsilon} \end{cases}$

Theorem: RR_ε is ε-LDP.

Unbiased estimator of the mean μ given $y_i = \mathrm{RR}_\varepsilon(x_i)$ for $i = 1, \ldots, n$:

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{(1 + e^\varepsilon) \cdot y_i - 1}{e^\varepsilon - 1} \right).$$

Standard deviation: $O\left(\frac{1}{\varepsilon\sqrt{n}}\right)$ for $\varepsilon \leq 1$.

Randomized Response

RR gives an ε-locally DP protocol that

- Estimates "statistical queries" (means/avgs) to $\pm O\left(\frac{1}{\varepsilon\sqrt{n}}\right)$.
  - Q: how to use RR for fractional-valued functions?
  - A: first randomly *round* $x_i \in [0,1]$ to 1 w.p. x_i, 0 w.p. $1 - x_i$.
- Estimates count/sum of a bounded function to $\pm O\left(\frac{\sqrt{n}}{\varepsilon}\right)$.

- Worse than centralized DP by a factor of $\sqrt{n}$, but still useful.

- Fact: The above privacy-accuracy tradeoff is the best possible for ε-local DP.

Local DP Histograms

$x_1, \ldots, x_n \in [D]$ (D bins). Use a 1-hot encoding:

x_i = 0 0 0 … 1 0 … 0 …. 0 Length D

$\mathrm{RR}_{\varepsilon/2}$ on every coordinate

y_i = 1 0 1 … 1 0 … 0…. 1

$$\hat{h} = \sum_{i=1}^{n} \left(\frac{\left(1 + e^{\varepsilon/2}\right) \cdot y_i - \vec{1}}{e^{\varepsilon/2} - 1} \right).$$

Local DP Histograms

- Expected error on each bin is $\pm O\left(\frac{\sqrt{n}}{\varepsilon}\right)$.
- Expected max error over all D bins is $\pm O\left(\frac{\sqrt{n \cdot \log D}}{\varepsilon}\right)$.

- We need to communicate D bits from each user.
  There exist protocols that use sophisticated algorithmic ideas to get communication complexity sublinear in D.

Local vs. Centralized DP

Central Model

- Central curator collects the data from all users, then performs privatization
- Requires the users to trust the curator with their private data
- Most DP algorithms are in this model

Local Model

- Each user privatizes their own data then sends it to a central curator
- Requires less trust from users
- Worse accuracy

Slide based on one from Brendan Avent's presentation

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).

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

Federated DP

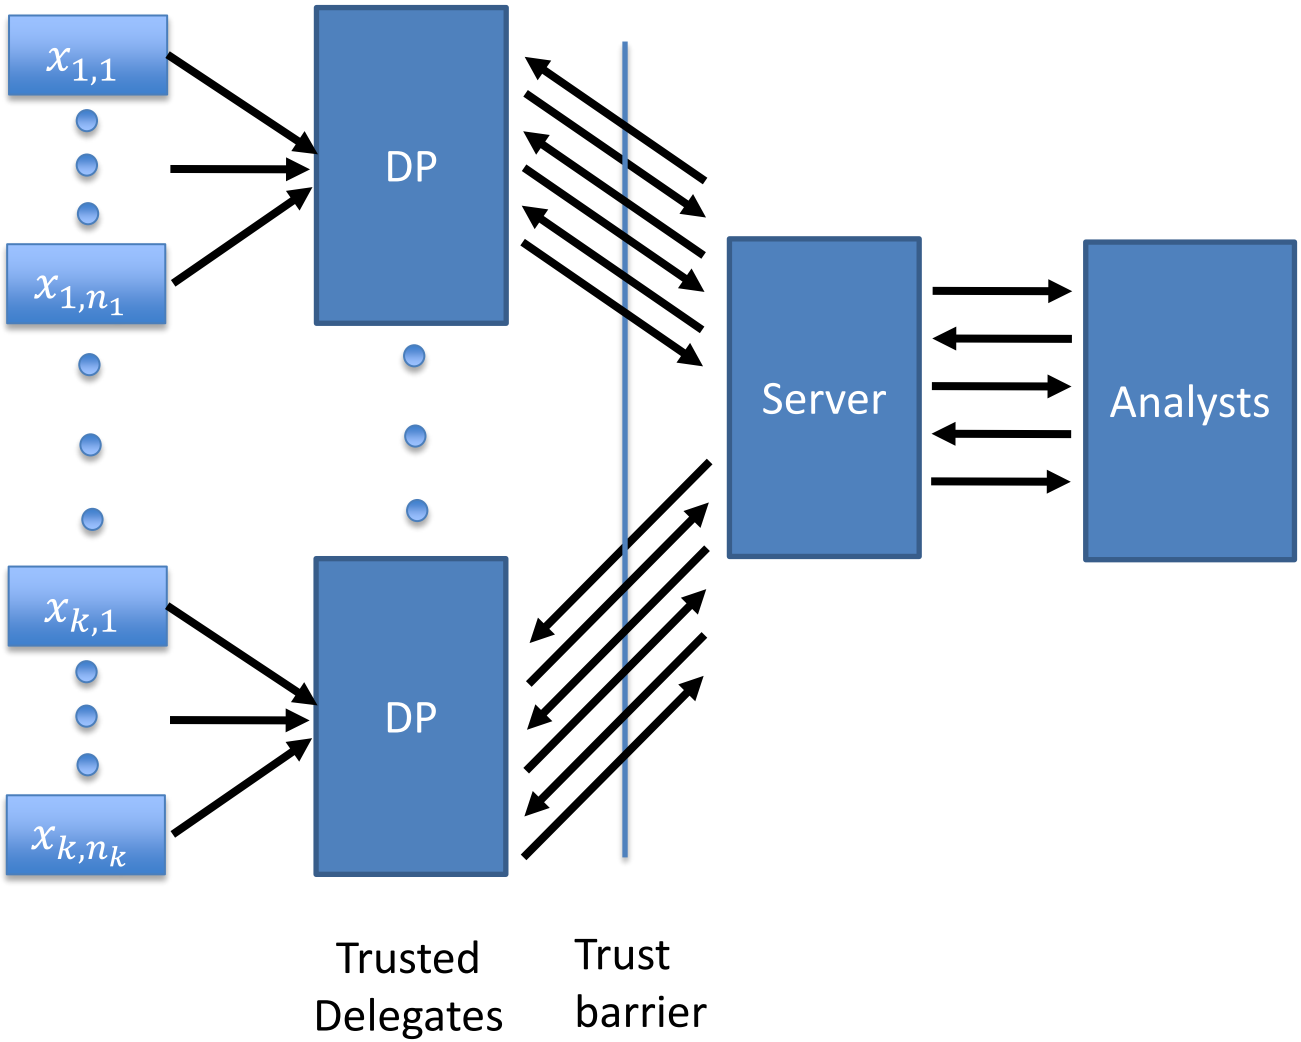

Comparing the Models

- Federated DP with k delegates, $n = n_1 + \cdots + n_k$
  - "horizontally partitioned" data
  - $k = 1$: central DP
  - $k = n$: local DP

- Error for sum of bounded values (like in DP-SGD) = $\Theta\left(\frac{\sqrt{k}}{\varepsilon}\right)$.
  - Interpolates between local & central model

- Error for set intersection when $k = 2$: $\Theta\left(\frac{\sqrt{n}}{\varepsilon}\right)$
  - No better than local model!

Other Models

- Can we get the “best of both worlds”?
  - Privacy protections like the local model
  - Accuracy like the central model

- Two approaches
  - The shuffle model
  - Using cryptography (secure multiparty computation)

Shuffle DP

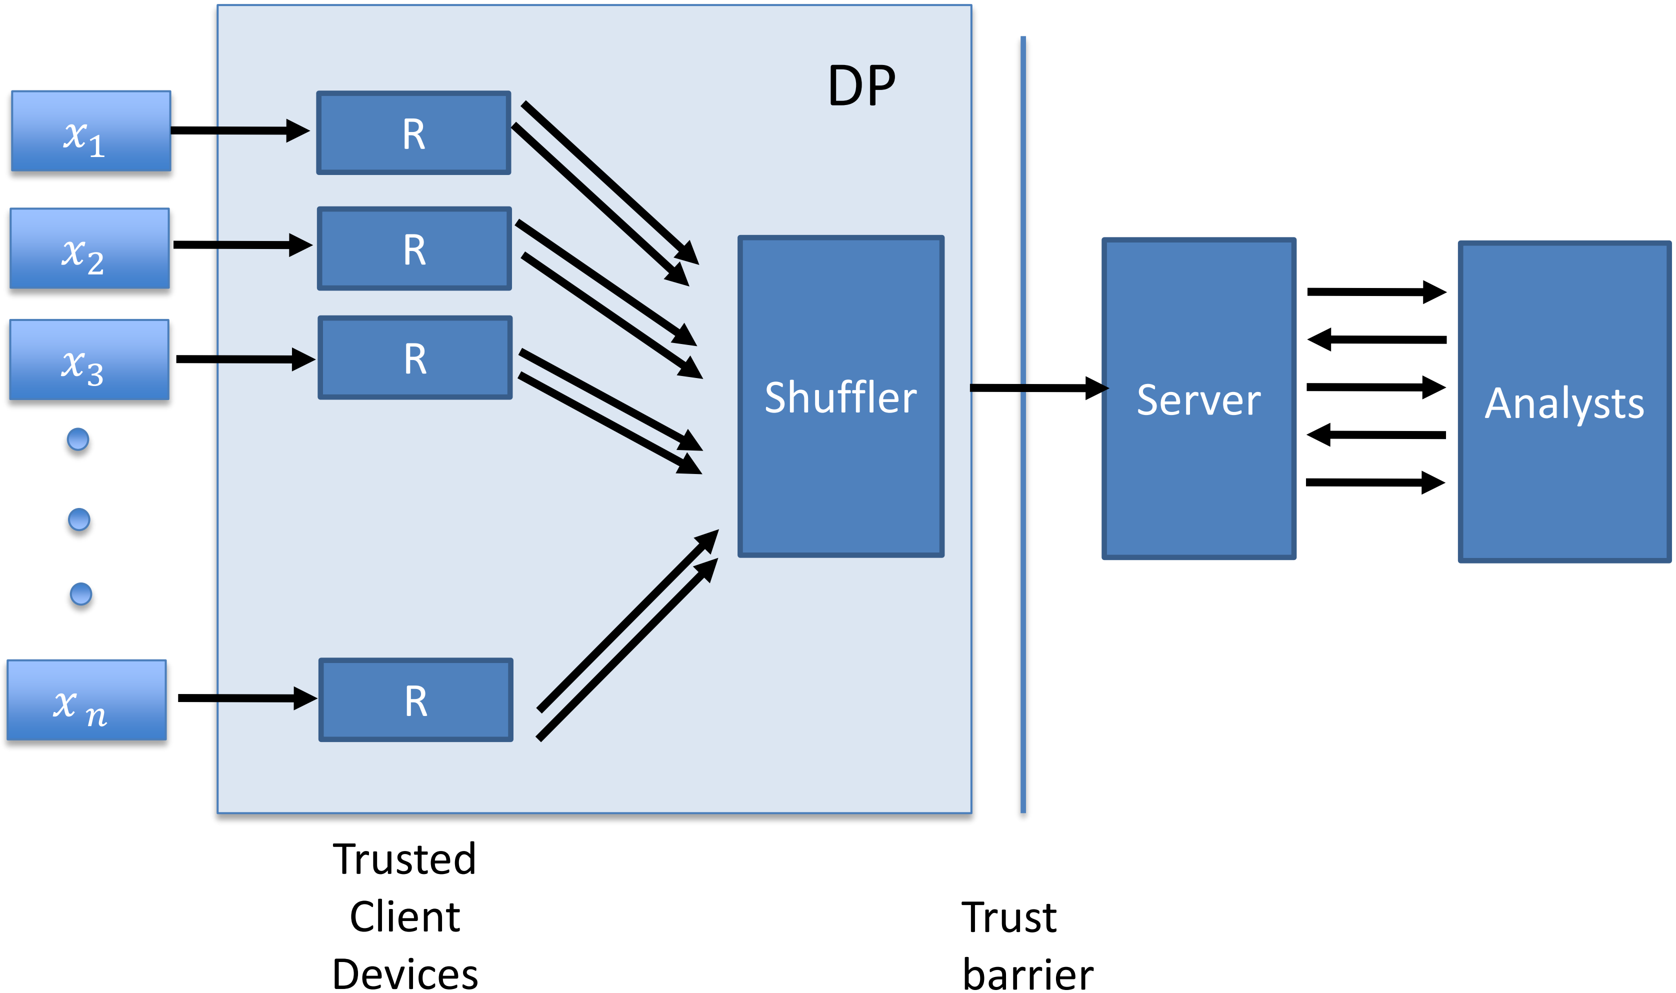

DP
x_1
x_2
x_3
x_n
R
R
R
R
Shuffler
Server
Analysts
Trusted Client Devices
Trust barrier