On the Complexity of Two-Party Differential Privacy

Noam Mazor

Joint work with

Iftach Haitner, Jad Silbak, Eliad Tsfadia

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

Differential Privacy (DP)

Dwork, McSherry, Nissim, Smith 2006

One record does not change the output distribution "too much"

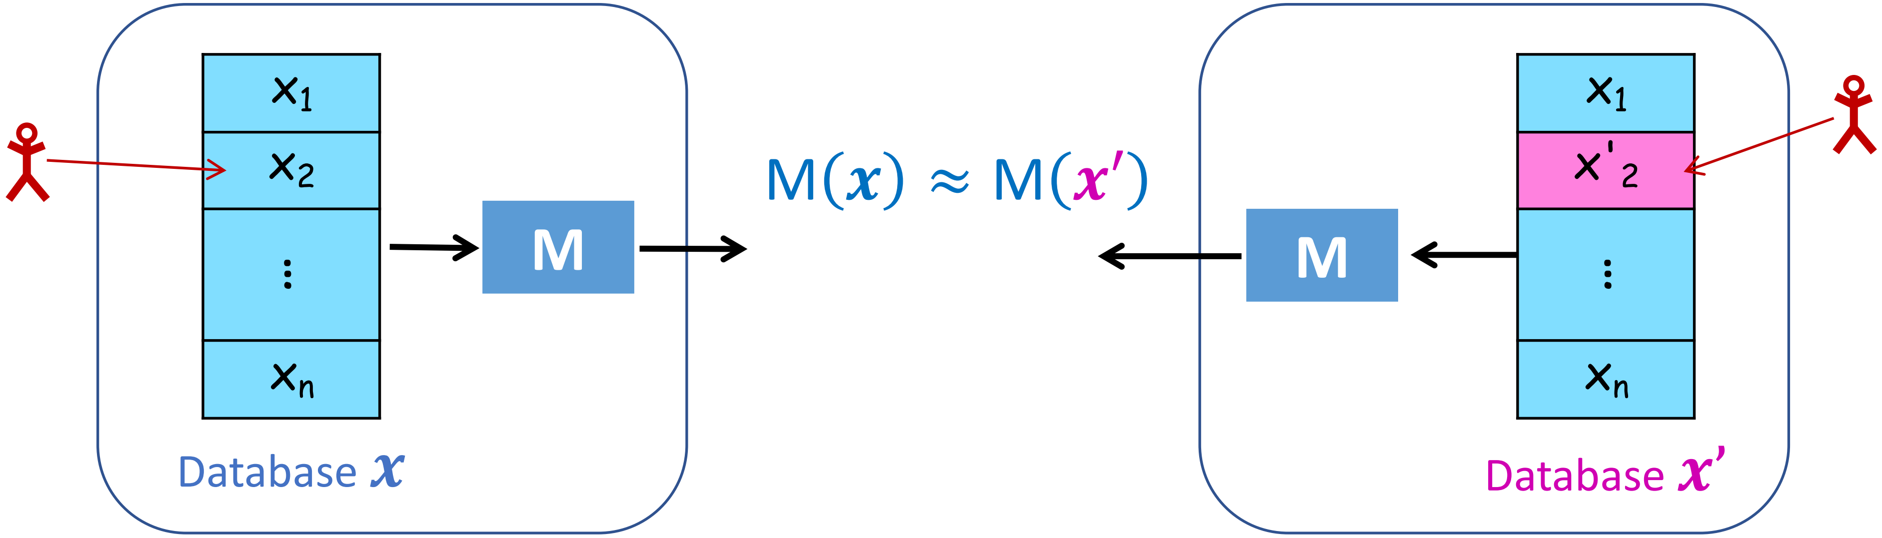

Alice participates

Alice provides fake info

Differential Privacy (DP)

Dwork, McSherry, Nissim, Smith 2006

One record does not change the output distribution "too much"

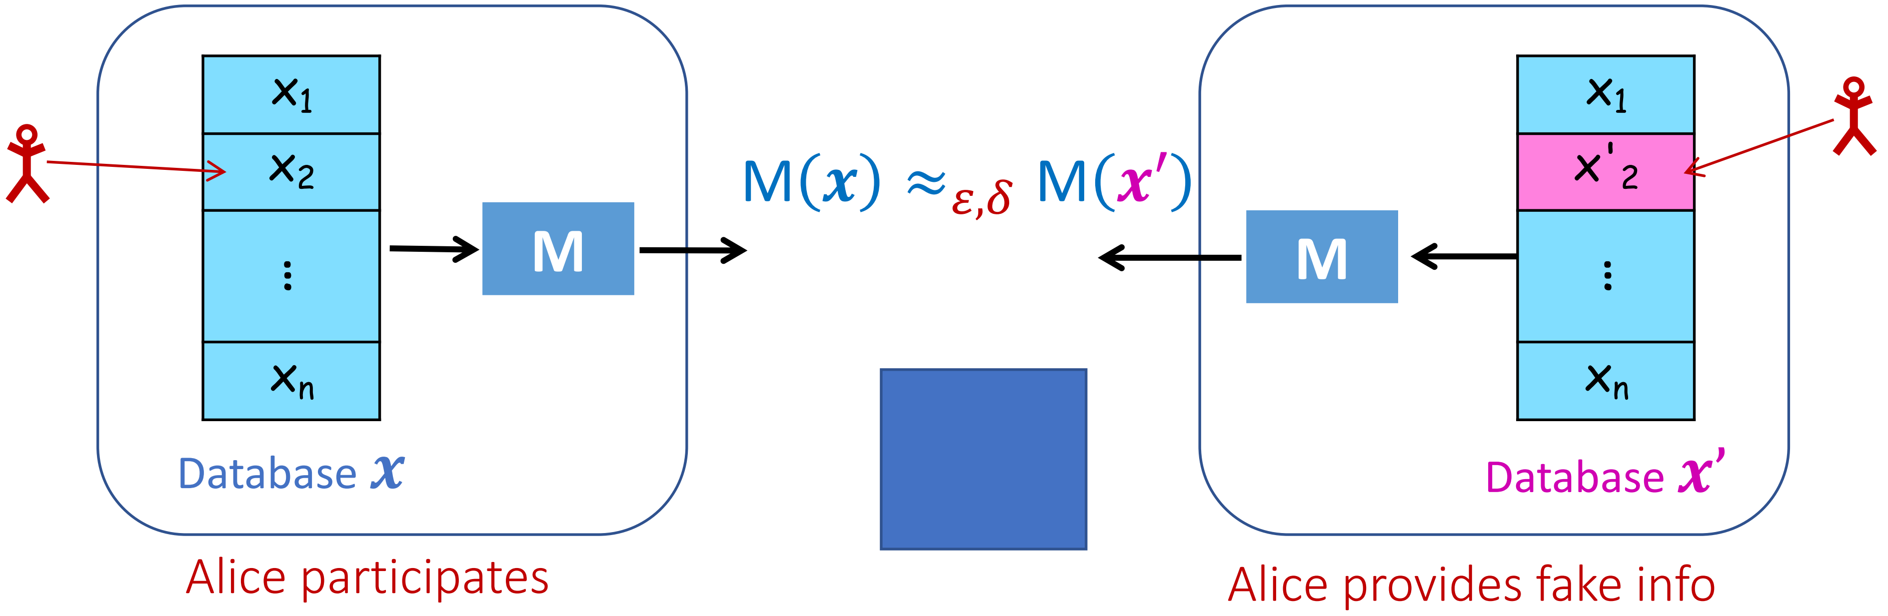

M is (ε, δ)-differentially private if

$\forall$ neighboring databases $\boldsymbol{x}, \boldsymbol{x}'$ and $\forall$(unbounded) distinguisher D:

$$\Pr[\mathrm{D}(\mathrm{M}(\boldsymbol{x})) = 1] \leq e^{\varepsilon} \cdot \Pr[\mathrm{D}(\mathrm{M}(\boldsymbol{x}')) = 1] + \delta$$

Centralized DP

$$\boldsymbol{x} = (x_1, \dots, x_n)$$

$$\mathrm{M}_1(\boldsymbol{x}) = \sum_{\boldsymbol{i}} \boldsymbol{x}_{\boldsymbol{i}}$$

Centralized DP

$\boldsymbol{x} = (x_1, \dots, x_n)$

$$\mathrm{M}_1(\boldsymbol{x}) = \sum_i \boldsymbol{x_i} + \boldsymbol{Noise}$$

Centralized DP

$\boldsymbol{x} = (x_1, \dots, x_n)$

$\mathrm{M}_1(\boldsymbol{x}) = \sum_i \boldsymbol{x_i} + \boldsymbol{Lap}(\frac{\mathbf{1}}{\boldsymbol{\epsilon}})$

Centralized model

- M has access to the **entire** database $\boldsymbol{x}$
- Goal: Estimate $f(\boldsymbol{x})$ while preserving privacy

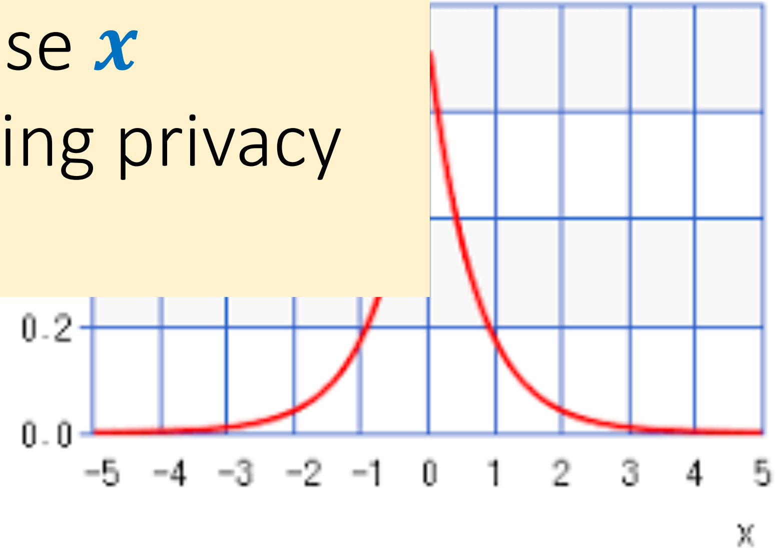

Two-Party DP

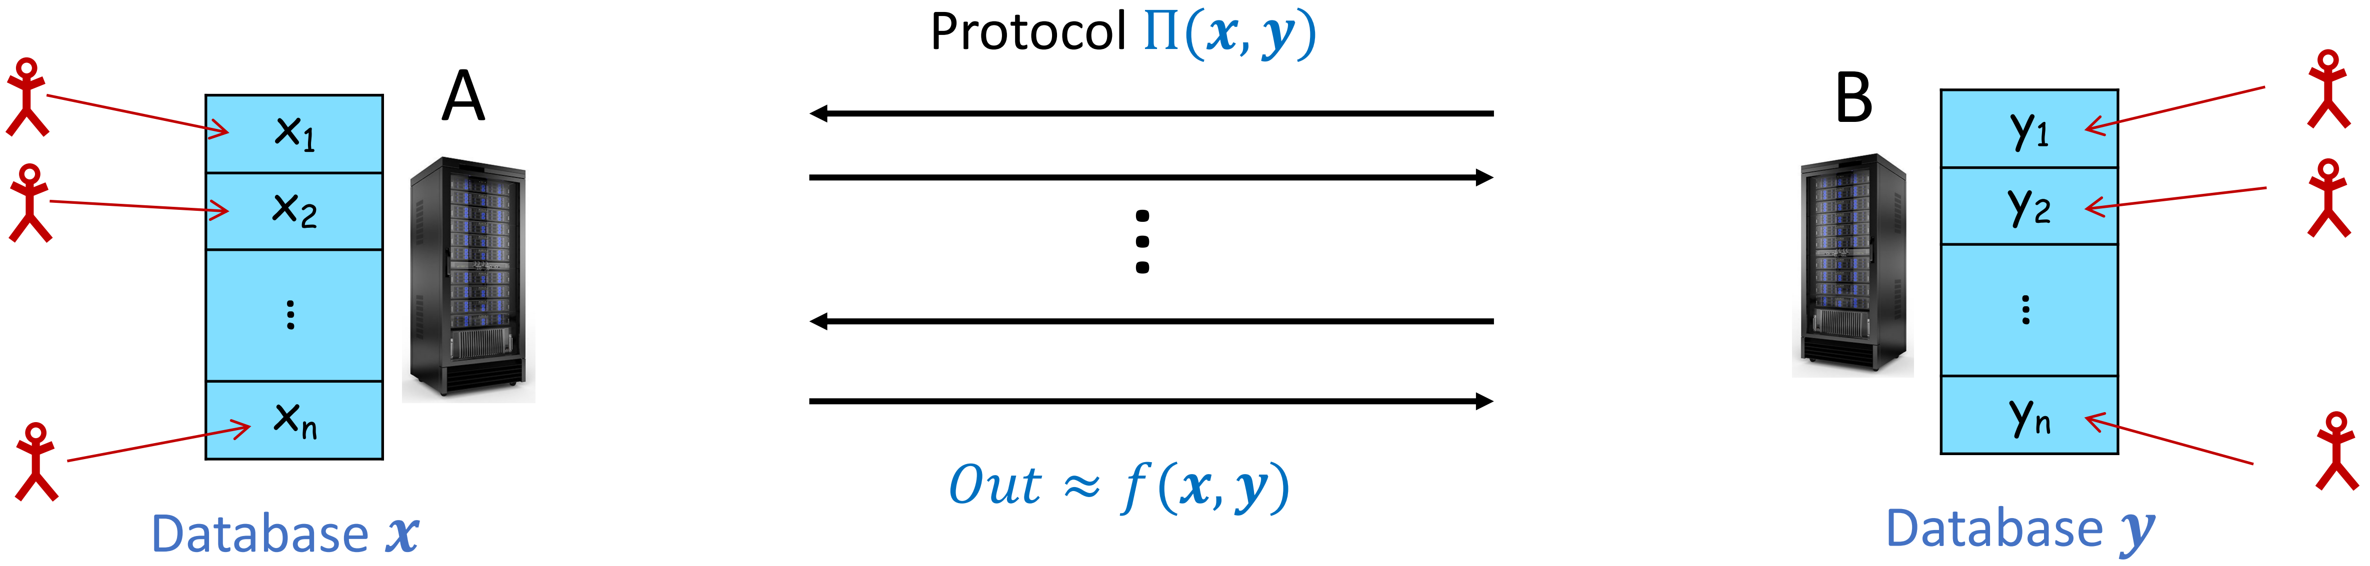

Goal: Estimate $f(\boldsymbol{x}, \boldsymbol{y})$ while preserving (ε, δ)-DP:

$$\forall \boldsymbol{x}, \forall \text{ neigh. } \boldsymbol{y}, \boldsymbol{y}': \quad view_A^{\Pi}(\boldsymbol{x}, \boldsymbol{y}) \approx_{\varepsilon,\delta} view_A^{\Pi}(\boldsymbol{x}, \boldsymbol{y}')$$

$view_A^{\Pi}(\boldsymbol{x}, \boldsymbol{y})$ – A's view in $\Pi(\boldsymbol{x}, \boldsymbol{y})$ (input, coins and transcript).

(and same for B)

Two-Party DP

Faculty of Exact Sciences

$\boldsymbol{x} = (x_1, \dots, x_n)$

$$out_1 = \sum_i x_i + Noise$$

$$out_2 = \sum_i y_i + Noise$$

$$out = out_1 + out_2$$

Faculty of Social Sciences

$\boldsymbol{y} = (y_1, \dots, y_n)$

Measure Correlation

A

$x \in \{-1,1\}^n$

?

B

$y \in \{-1,1\}^n$

$\langle x, y \rangle = \sum_{i=1}^{n} x_i y_i$ – measures **correlation** between databases

DP Inner Product

Centralized Model

Can achieve constant error.

Two-Party Protocol

For uniform inputs:

A: $\boldsymbol{x} \in \{-1,1\}^n$

B: $\boldsymbol{y} \in \{-1,1\}^n$

$$out = 0$$

- With prob. 0.99: $|out - \langle \boldsymbol{x}, \boldsymbol{y} \rangle| \approx \sqrt{n}$
- Can be generalized for every input distribution.

Using Crypto?

A: $\boldsymbol{x} \in \{-1,1\}^n$

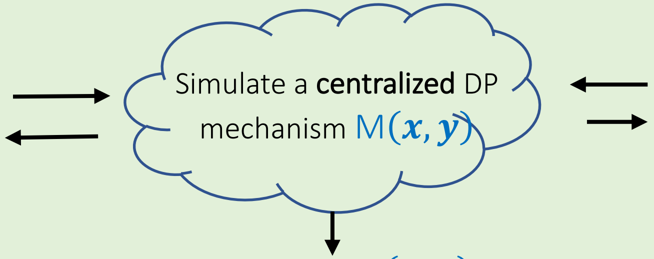

B: $\boldsymbol{y} \in \{-1,1\}^n$

$out = \mathrm{M}(\boldsymbol{x}, \boldsymbol{y})$

Need new definition of DP

Lower Bound

McGregor, Mironov, Pitassi, Reingold, Talwar and Vadhan 2010

For every DP protocol:

$$|out - \langle \boldsymbol{x}, \boldsymbol{y} \rangle| \approx \sqrt{n}$$

Computational DP

Beimel, Nissim, Omri 2008 Mironov, Pandey, Reingold, Vadhan 2009

- M is (ε, δ)-DP if:

$\forall$ neighboring databases $\boldsymbol{x}, \boldsymbol{x}'$ and $\forall$ **distinguisher** D:

$$\Pr[\mathrm{D}(\mathrm{M}(\boldsymbol{x})) = 1] \leq e^{\varepsilon} \cdot \Pr[\mathrm{D}(\mathrm{M}(\boldsymbol{x}')) = 1] + \delta$$

- M is (ε, δ)-**CDP** if:

the above only holds for any PPT D.

Two-Party CDP

Beimel, Nissim, Omri 2008 Mironov, Pandey, Reingold, Vadhan 2009

$\Pi(\boldsymbol{x}, \boldsymbol{y})$

A

$\boldsymbol{x} = (x_1, \dots, x_n)$

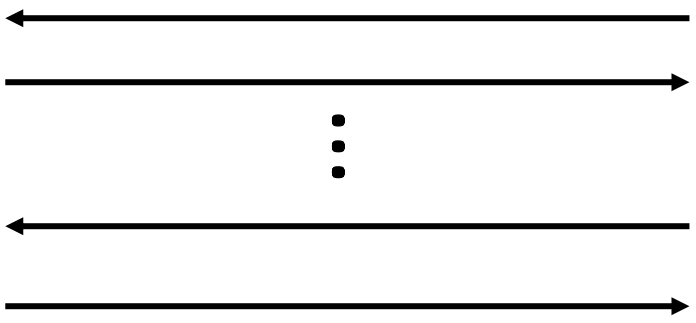

B

$\boldsymbol{y} = (y_1, \dots, y_n)$

$Out \approx f(\boldsymbol{x}, \boldsymbol{y})$

Relaxed Goal: Estimate $f(\boldsymbol{x}, \boldsymbol{y})$ while preserving (ε, δ)-CDP:

$\forall \boldsymbol{x}\ \forall$ neigh. $\boldsymbol{y}, \boldsymbol{y}'$: $\quad view_A^{\Pi}(\boldsymbol{x}, \boldsymbol{y}) \approx_{\varepsilon,\delta}^{c} view_A^{\Pi}(\boldsymbol{x}, \boldsymbol{y}')$

(and same for B)

CDP via Secure Multiparty Computation

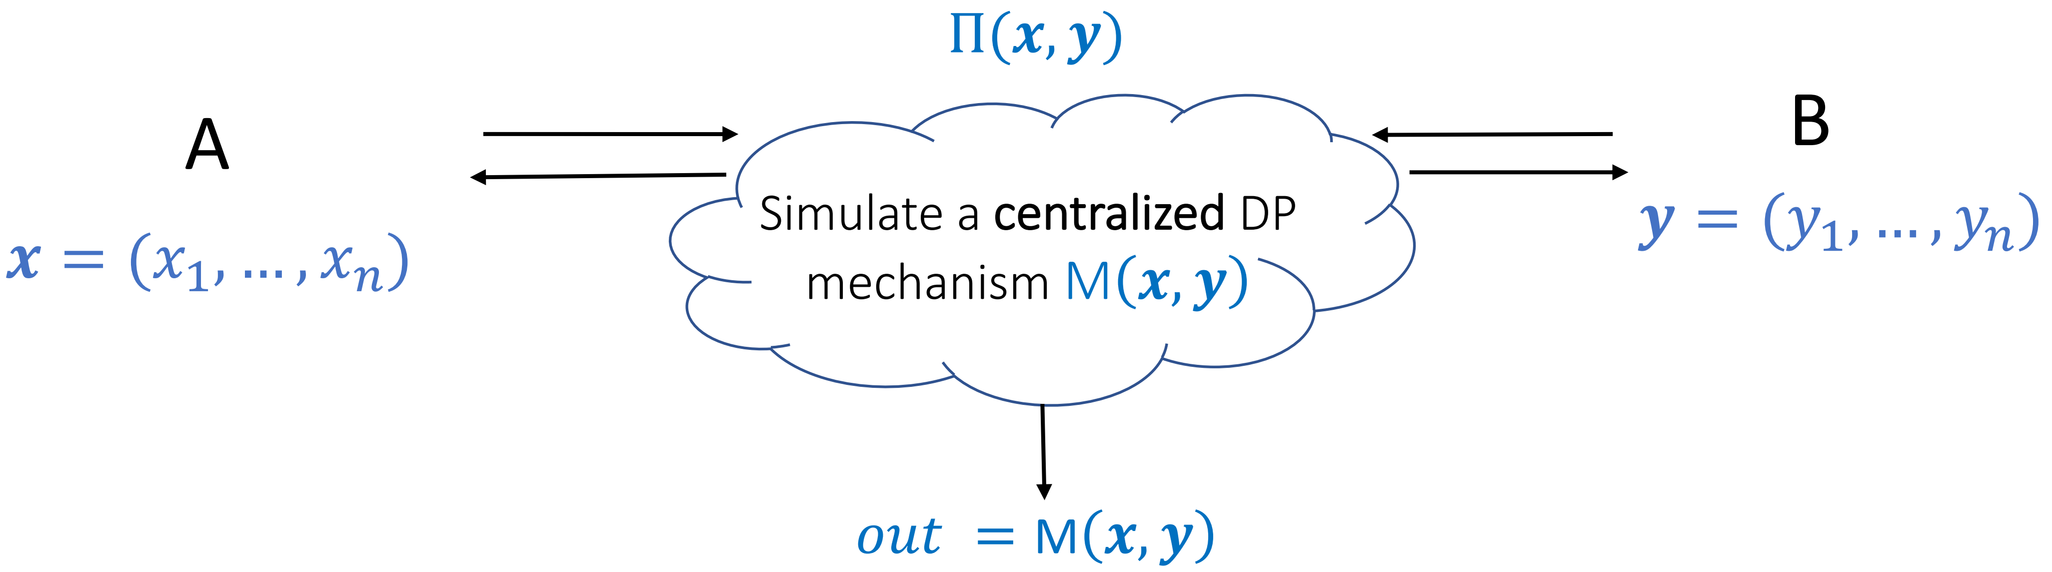

- M is (centralized) (ε, δ)-DP $\implies$ Π is (ε, δ)-CDP.
- Secure MPC via *Oblivious Transfer* (OT).
- For computing IP, take $\mathrm{M}(\boldsymbol{x}, \boldsymbol{y}) = \langle \boldsymbol{x}, \boldsymbol{y} \rangle + Lap(2/\varepsilon)$.

The Complexity of Two-Party CDP

Using OT, we can construct very accurate CDP protocols!

Main Questions:

- Are one-way functions sufficient?
- Is public-key cryptography necessary?
- Do we have to use (heavy) Secure MPC?

Complexity Hierarchy

“Non-trivial”:

Possible in two-party **CDP**

Impossible in two-party **DP**

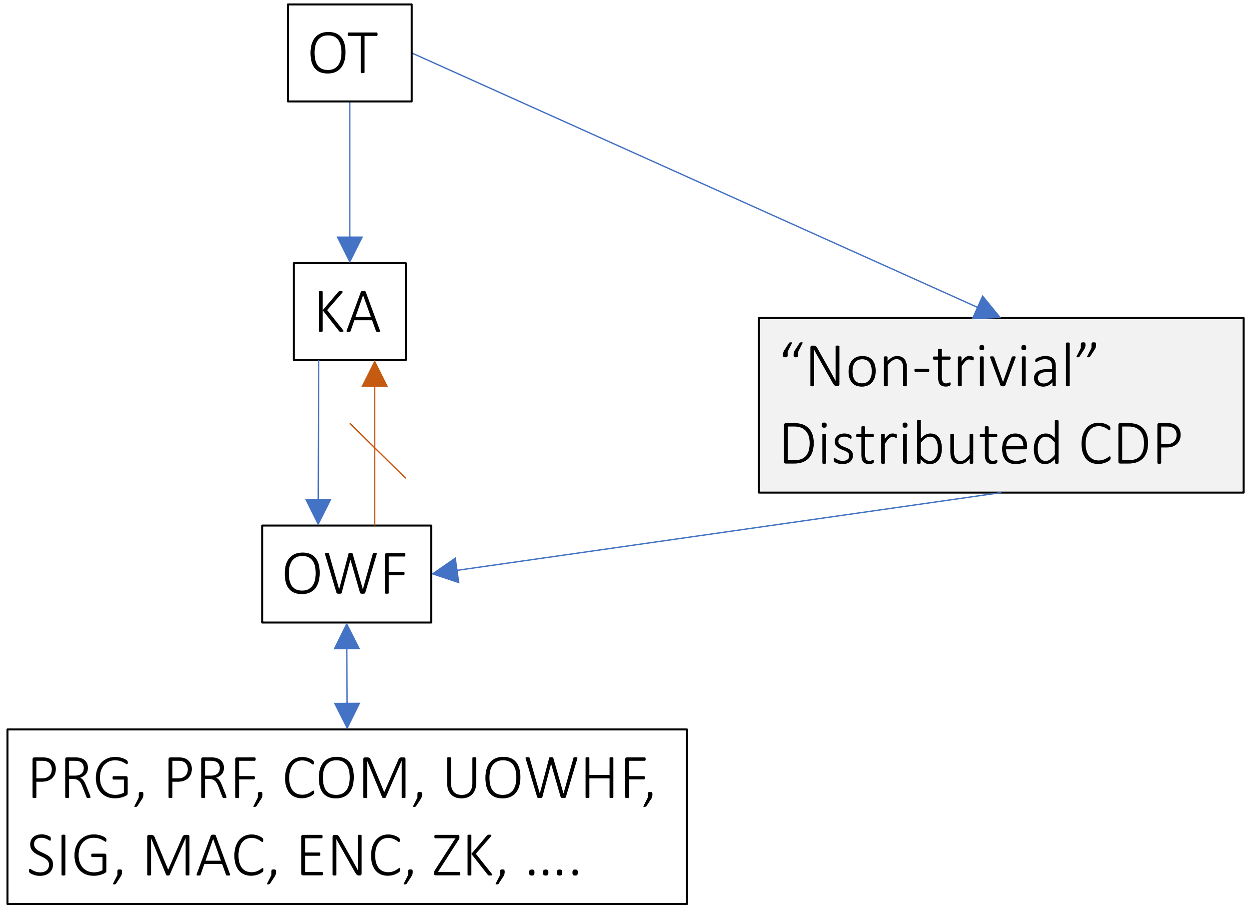

Complexity Hierarchy

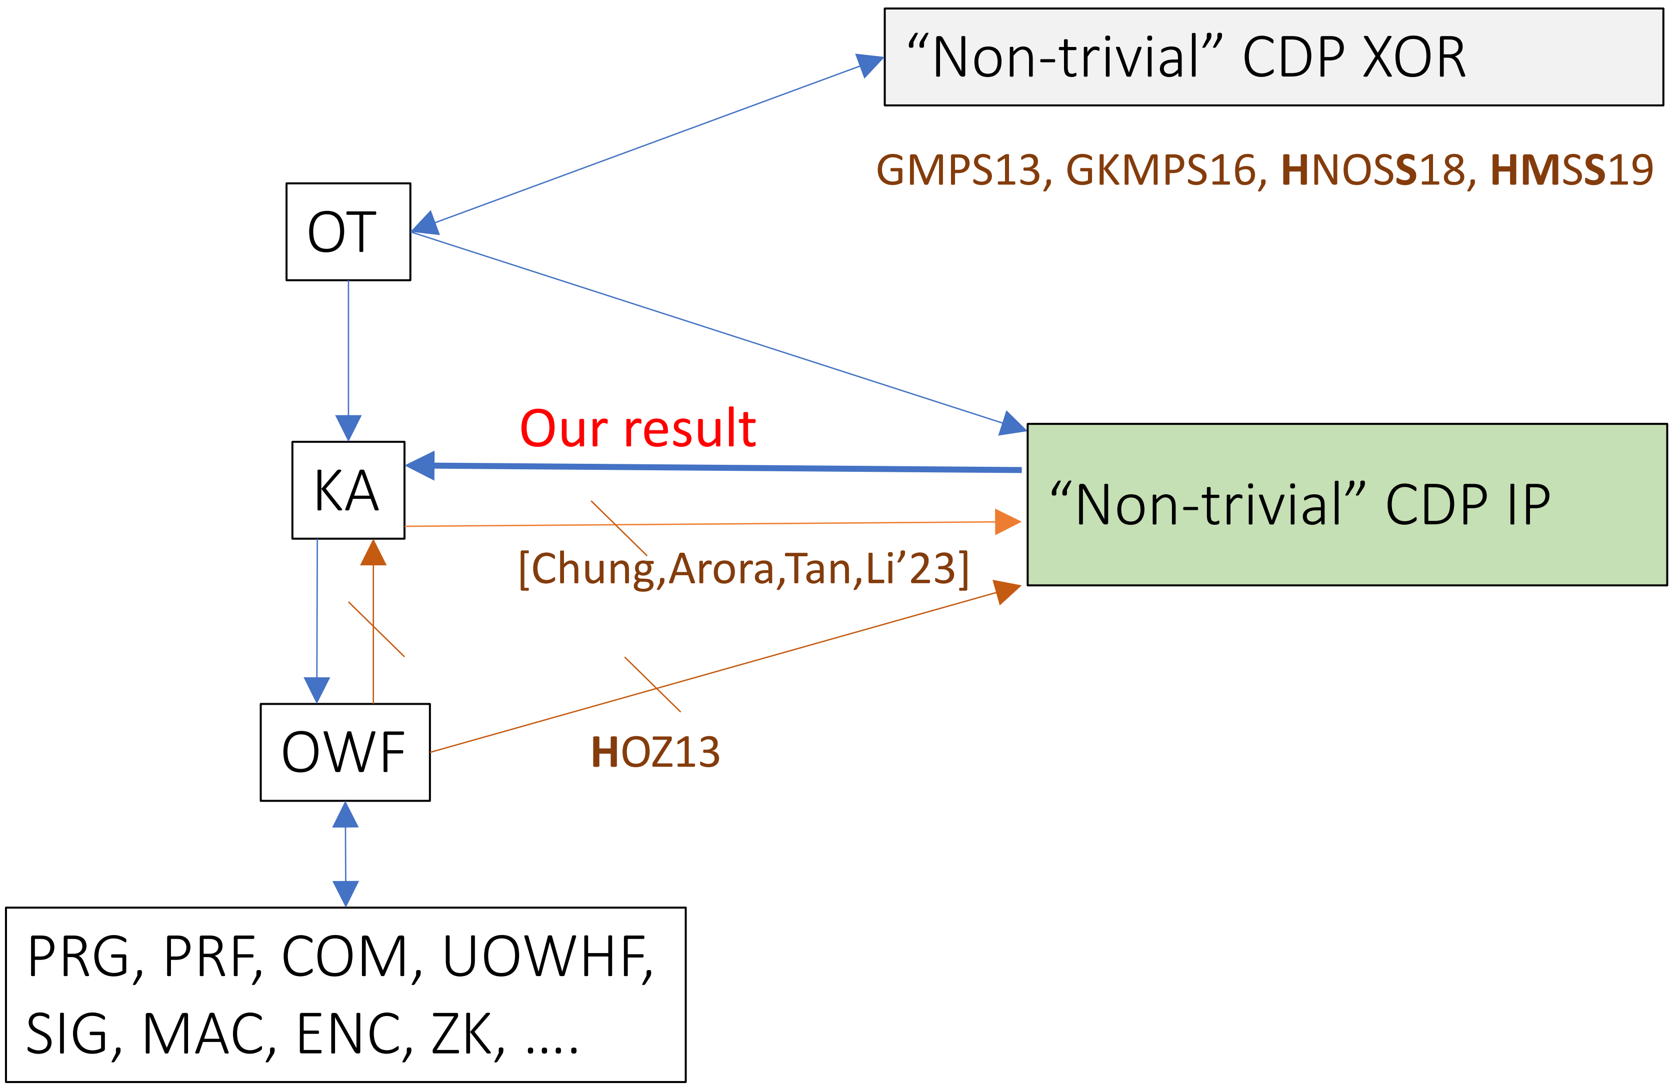

Our Main Result

<u>Thm 1 (informal)</u>: $(\varepsilon, \delta = 1/n^2)$-CDP two-party Π that, for some ℓ satisfies $\Pr_{\substack{x,y \leftarrow \{-1,1\}^n \\ out \leftarrow \Pi(x,y)}}[|out - \langle \boldsymbol{x}, \boldsymbol{y} \rangle| < \ell] > e^{\varepsilon} \cdot \ell / \sqrt{n}$, can be used to construct **Key Agreement.**

- For $\varepsilon = O(1)$ and $\ell = \sqrt{n}/c$ (for large enough constant c):

$$\Pr_{\substack{x,y \leftarrow \{-1,1\}^n \\ out \leftarrow \Pi(x,y)}}[|out - \langle \boldsymbol{x}, \boldsymbol{y} \rangle| \leq \sqrt{n}/c] \geq 0.01 \implies \text{Key Agreement}$$

  - ➢ Reproves the impossibility result of McGregor et al.

- Tight (up to a constant).
  - ➢ Protocol that outputs zero is w.p. $\Theta(\ell/\sqrt{n})$ at distance at most ℓ (for every ℓ).

The Information-Theoretic Lower Bound

McGregor, Mironov, Pitassi, Reingold, Talwar, and Vadhan 2010

Let Π be ε-DP, $X, Y \leftarrow \{-1,1\}^n$ and $T \leftarrow \Pi(X, Y)$ (transcript). Then:

1. $X|_T$ and $Y|_T$ are **independent.**

2. X_i is unpredictable given T, X_{-i} (*strong* $\boldsymbol{Santha\ Vazirani}$ Source)

IP is a good extractor for such sources.

$\Longrightarrow \langle X, Y \rangle|_T$ is **almost unifrom** modulo $\sqrt{n}$

Computational Setting

Let Π be ε-CDP, $X, Y \leftarrow \{-1,1\}^n$ and $T \leftarrow \Pi(X, Y)$ (transcript). Then:

- $X|_T$ and $Y|_T$ are independent, computationally *strong* SV Sources.

- IP is **not** a good extractor for such sources.

- Indeed, assuming OT, exists ε-CDP Π s.t. $\langle X, Y\rangle|_T$ is **predictable** (up to $\approx 1/\varepsilon$).

- $X|_T$ and $Y|_T$ are computationally **correlated**.

- Goal: Exploit the computational correlation into **Key Agreement**.

Proof Overview

CDP IP to KA

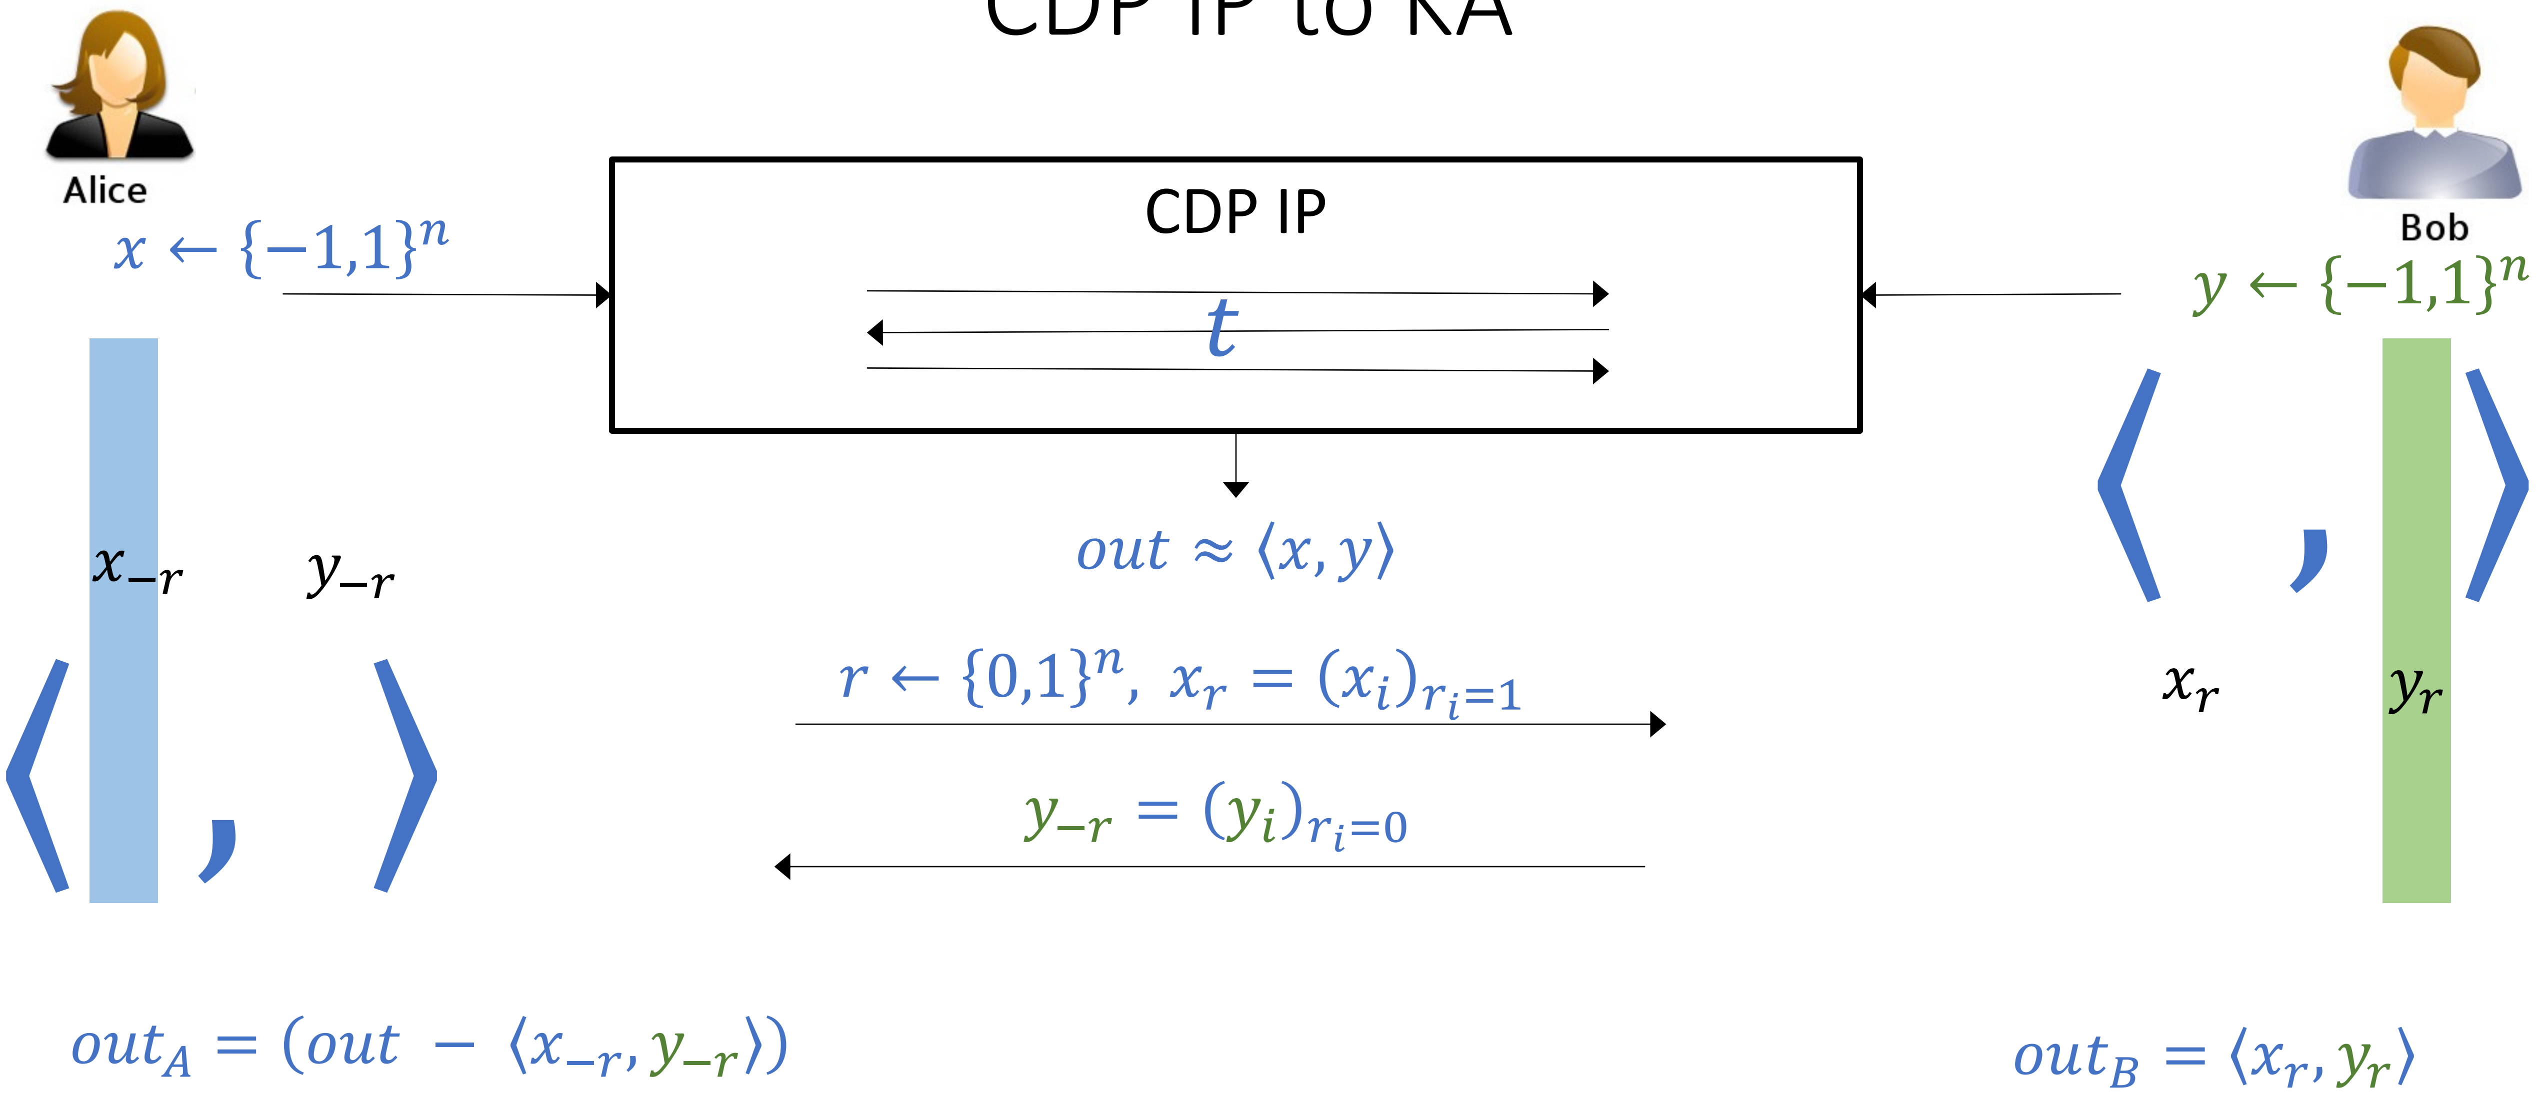

Alice
$x \leftarrow \{-1,1\}^n$
CDP IP
t
Bob
$y \leftarrow \{-1,1\}^n$
$out \approx \langle x, y \rangle$
x_{-r}
y_{-r}
$r \leftarrow \{0,1\}^n,\ x_r = (x_i)_{r_i=1}$
$y_{-r} = (y_i)_{r_i=0}$
x_r
y_r
$out_A = (out - \langle x_{-r}, y_{-r} \rangle)$
$out_B = \langle x_r, y_r \rangle$

Analysis

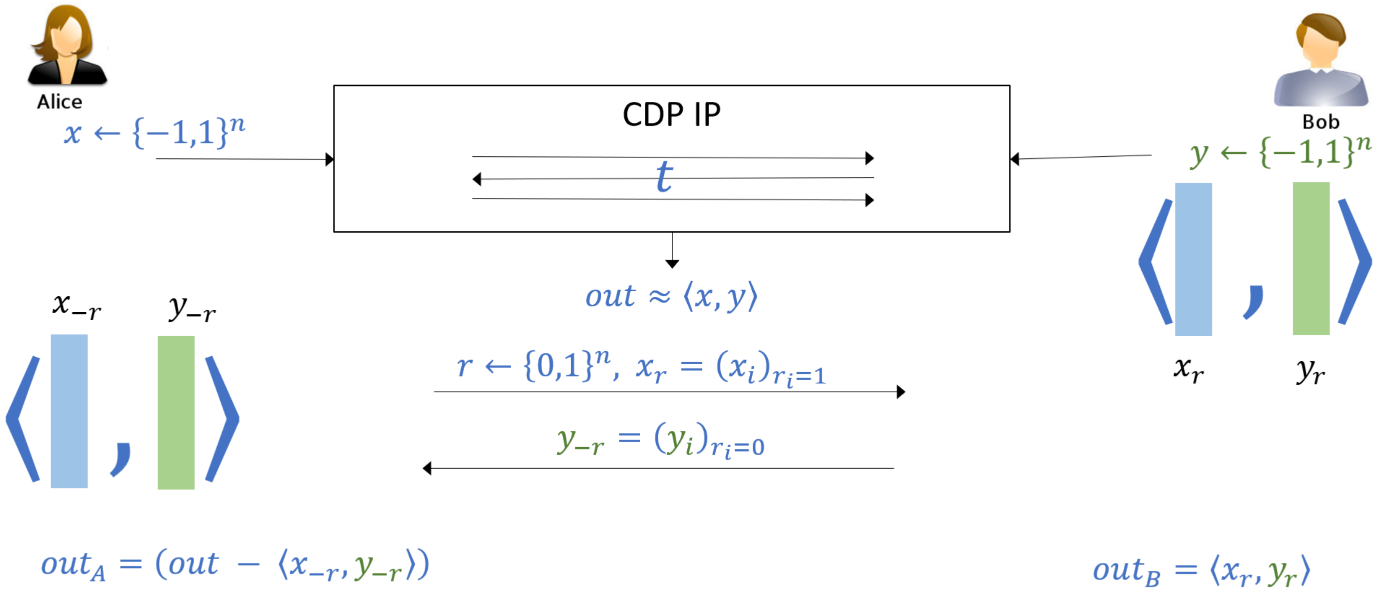

- Agreement:

  $$Out \approx \langle X, Y \rangle \implies Out_A \approx Out_B$$

- Secrecy:

  Goal: showing that $\forall$ PPT Eve, Eve(T, R, X_R, Y_{-R}) is far from Out_B.
  - Should hold since $(X,Y)|_T$ is highly unpredictable by privacy (computationally strong SV).
    - The proof is not trivial.
    - Done via a new theorem about *strong SV sources.*

- Simple proof for the case $E[|out - \langle x, y \rangle|] \le \frac{\sqrt{n}}{\log^c(n)}$.

Seed-dependent condenser for strong SV

Thm 2 (informal): Let (X, Y) be $e^{-\varepsilon}$-strong SV. Then whp over $R \leftarrow \{0,1\}^n$:

$$H_\infty(\langle X_R, Y_R \rangle \mid R, X_R, Y_{-R}) \geq \log\left(\frac{\sqrt{n}}{e^{\varepsilon} \cdot \log n}\right)$$

Constructive proof.

- High min-entropy **conditioned** on the *seed-dependent* leakage (X_R, Y_{-R}).

- Constructive proof: $\exists$ PPT Rec and $i \in [n]$ such that:
  $\forall$ PPT E(R, X_R, Y_{-R}) that predicts $\langle X_R, Y_R \rangle$ ``too well'',
  Rec$^{\mathrm{E}}(X_{-i}, Y)$ reconstructs X_i ``too well''.
  - Applicable for *computational* SV sources.

Conclusions & Open Problems

Non-trivial CDP-IP $\Rightarrow$ Key Agreement

Open Questions:

- Finding a more general characterization that capture more functionalities.
- Determine whether OT is the minimal required assumption for CDP IP.
    - Our result is tight for *DP against external observer.*

Thanks!