Sort, Partition, Randomize

Optimal Binary Hypothesis Testing under Local Differential Privacy

NeurIPS 2026 (Spotlight)

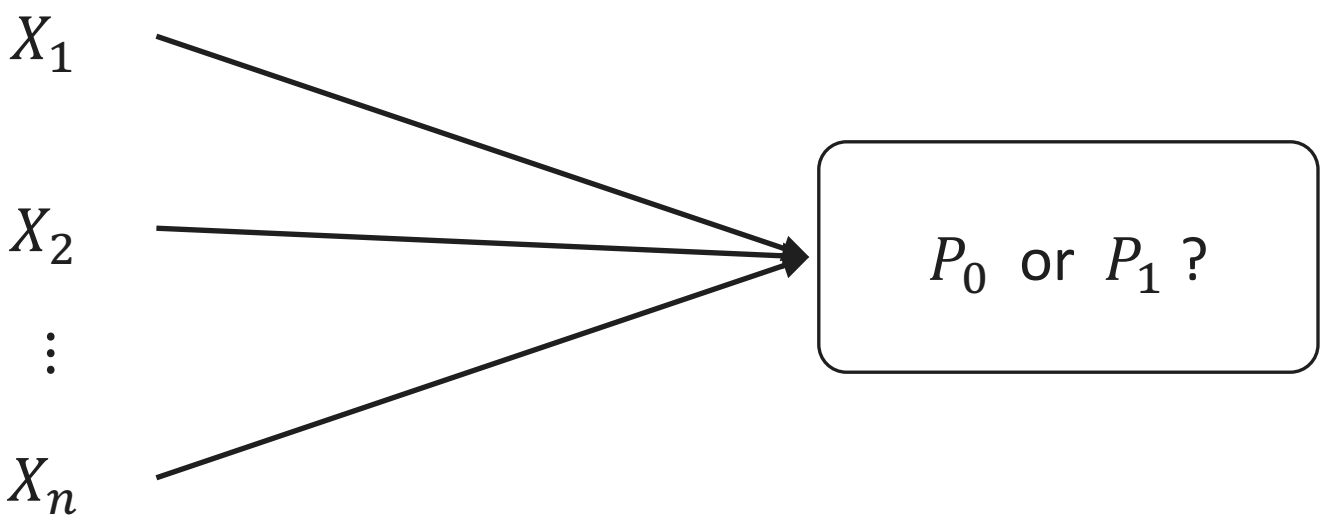
Elena Ghazi[1] Jawad Nasser[2] Flavio Calmon[1] Ibrahim Issa[2]

[1]Harvard University [2]American University of Beirut

Deciding from data

User locations

Analyst

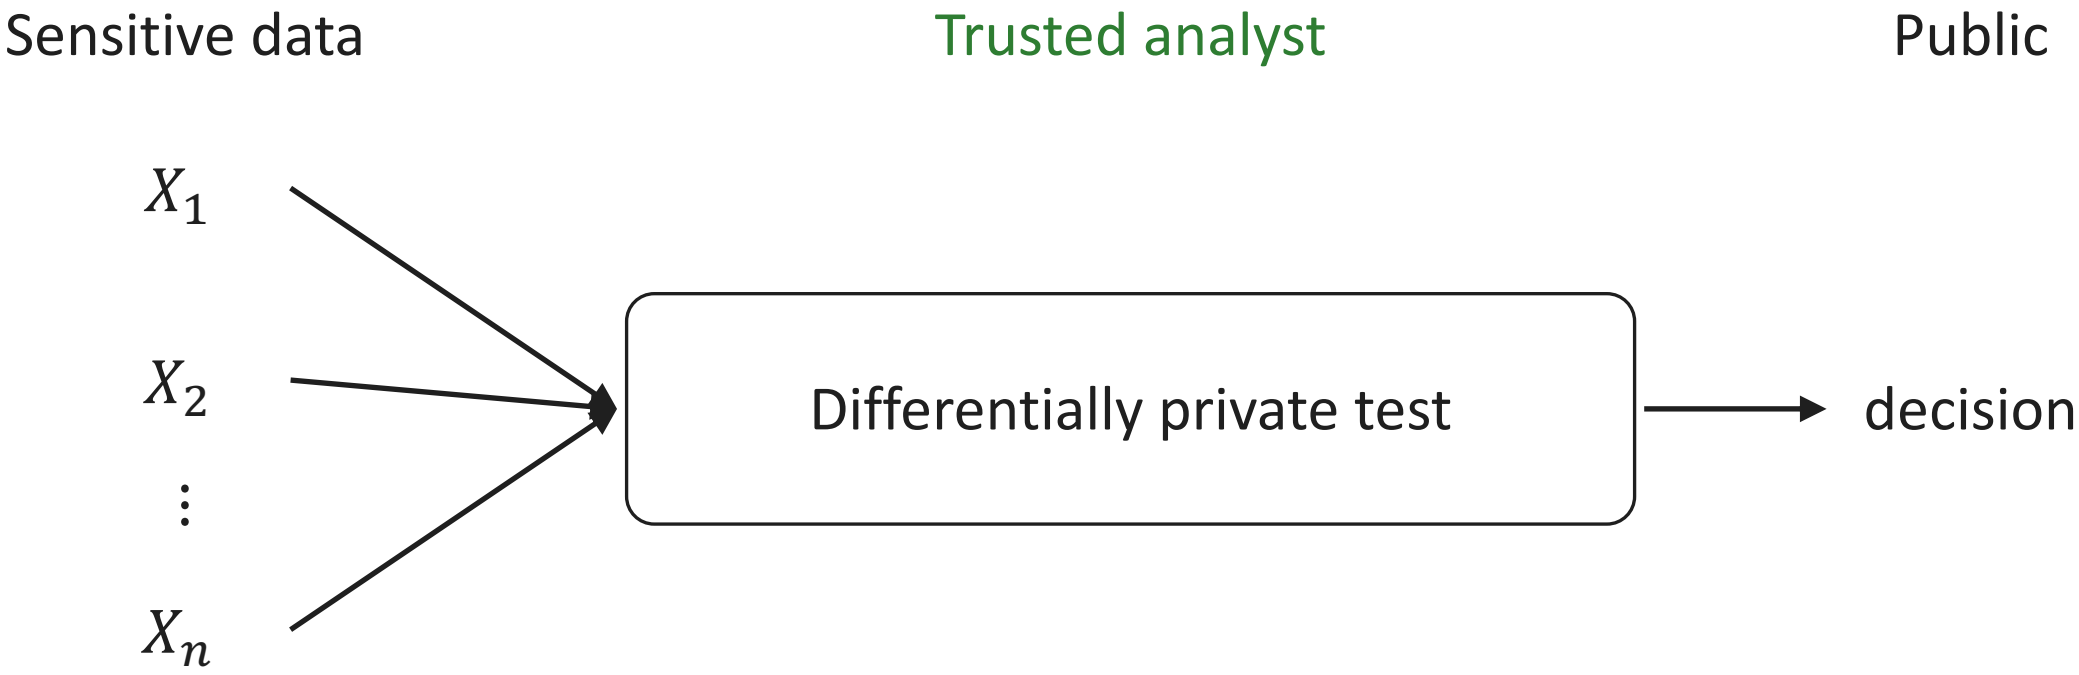


$X_i \in$ {Downtown, Cambridge, Back Bay, …}

X_1, …, X_n are drawn i.i.d. from one of two known distributions, P_0 or P_1

Normal conditions (P_0) or a transit shutdown (P_1)?

Deciding privately: trusted analyst



$X_1, \ldots, X_n$ drawn i.i.d. from P_0 or P_1

The analyst sees the raw data.

The public sees only the randomized decision.

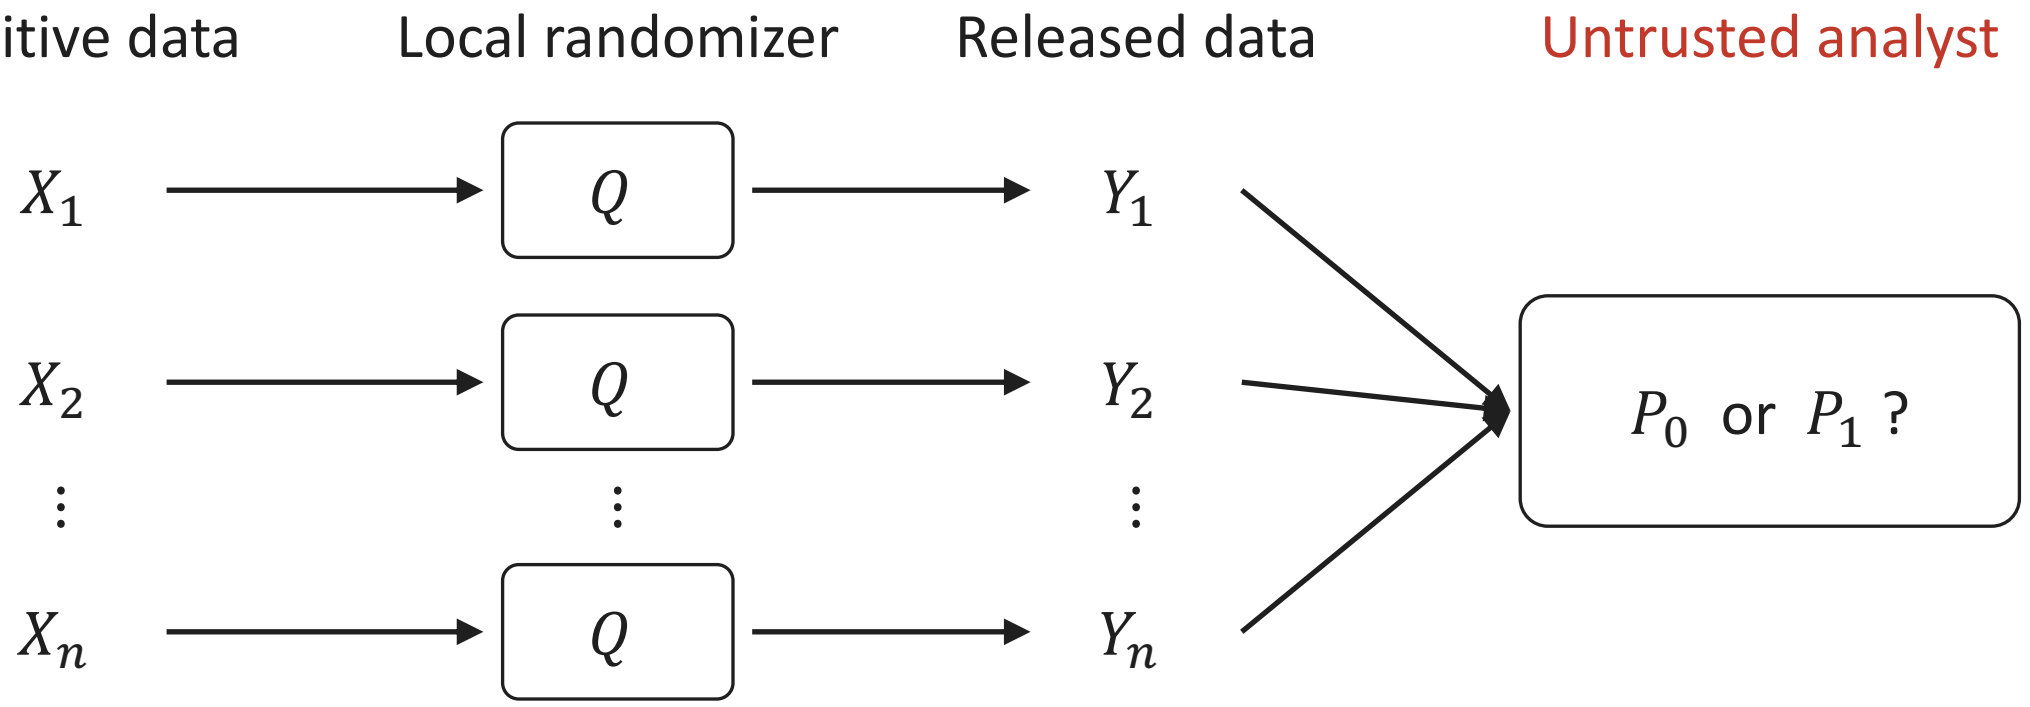
Differential privacy

Informally: changing one person's data has a limited effect on the distribution of the released result.

The privacy parameter ε controls the allowed change.
Smaller ε means stronger privacy.

Dwork, McSherry, Nissim & Smith, "Calibrating noise to sensitivity in private data analysis", TCC 2006

Deciding privately: untrusted analyst



$X_1, \ldots, X_n$ drawn i.i.d. from P_0 or P_1

The analyst sees only the randomized data.

Each user privatizes locally, with the same mechanism Q

Local differential privacy

A mechanism Q is ε-LDP if

$$Q(y \mid x) \leq e^{\varepsilon} Q(y \mid x') \quad \text{for all } x, x', y$$

Every output is about equally likely whatever the input: for any fixed output, its probability changes by at most a factor of e^{ε} across inputs
Small ε = strong privacy

Example: randomized response (Warner 1965), a yes/no answer

Report the true answer $\dfrac{e^{\varepsilon}}{e^{\varepsilon}+1}$

Report the opposite answer $\dfrac{1}{e^{\varepsilon}+1}$

Warner, “Randomized response: a survey technique for eliminating evasive answer bias”, JASA 1965

Local differential privacy

A mechanism Q is ε-LDP if

$$Q(y \mid x) \leq e^{\varepsilon} Q(y \mid x') \quad \text{for all } x, x', y$$

Every output is about equally likely whatever the input: for any fixed output, its probability changes by at most a factor of e^{ε} across inputs
Small ε = strong privacy

Example: m-ary randomized response: m possible inputs, m possible outputs

Report the true value $\dfrac{e^{\varepsilon}}{e^{\varepsilon} + m - 1}$

Report each other value $\dfrac{1}{e^{\varepsilon} + m - 1}$

Hypothesis testing under differential privacy

- **Canonne et al. (STOC 2019)**
  Two known distributions: optimal sample complexity up to constant factors, with a test achieving it.
- **Duchi, Jordan & Wainwright (FOCS 2013)**
  Minimax rates for estimation under local privacy.
- **Kairouz, Oh & Viswanath (JMLR 2016)**
  Structure of optimal local mechanisms.

This talk: finding the optimal LDP mechanism for testing two known distributions

What makes a mechanism good?

Let M_0, M_1 be the distributions of the released Y under P_0 and P_1:

$$M_i(y) = \sum_x Q(y \mid x) P_i(x)$$

A good mechanism keeps M_0 and M_1 "easy to tell apart"

A test can make two errors: $\alpha = P(\text{guess } P_1 \mid P_0)$ and $\beta = P(\text{guess } P_0 \mid P_1)$

Total variation (TV)

One sample, both errors weighted equally

$$\min_{\text{tests}} (\alpha + \beta) = 1 - \text{TV}(M_0, M_1)$$

(with a fair-coin prior: P(correct) = ½(1 + TV))

What makes a mechanism good?

Let M_0, M_1 be the distributions of the released Y under P_0 and P_1:

$$M_i(y) = \sum_x Q(y \mid x) P_i(x)$$

A good mechanism keeps M_0 and M_1 "easy to tell apart"

A test can make two errors: $\alpha = P(\text{guess } P_1 \mid P_0)$ and $\beta = P(\text{guess } P_0 \mid P_1)$

Hockey-stick divergence (E_γ)

One sample, β weighted γ times more

$$\min_{\text{tests}} (\alpha + \gamma\beta) = 1 - E_\gamma(M_0 \| M_1)$$

($\gamma = 1$ gives back TV)

What makes a mechanism good?

Let M_0, M_1 be the distributions of the released Y under P_0 and P_1:

$$M_i(y) = \sum_x Q(y \mid x) P_i(x)$$

A good mechanism keeps M_0 and M_1 "easy to tell apart"

A test can make two errors: $\alpha = P(\text{guess } P_1 \mid P_0)$ and $\beta = P(\text{guess } P_0 \mid P_1)$

Kullback–Leibler divergence (KL)

n samples: keep α below a fixed level, make β vanish as fast as possible

$$\beta \approx e^{-n \cdot \mathrm{KL}(M_0 \| M_1)} \quad \text{(Stein's lemma)}$$

What makes a mechanism good?

Let M_0, M_1 be the distributions of the released Y under P_0 and P_1:

$$M_i(y) = \sum_x Q(y \mid x) P_i(x)$$

A good mechanism keeps M_0 and M_1 "easy to tell apart"

A test can make two errors: $\alpha = P(\text{guess } P_1 \mid P_0)$ and $\beta = P(\text{guess } P_0 \mid P_1)$

Hellinger distance (H^2)

How many samples until $\alpha + \beta \leq 0.1$?

$$\text{samples needed} \approx \frac{1}{H^2(M_0, M_1)} \quad \text{(up to constants)}$$

Classical testing bound; see Tsybakov, *Introduction to Nonparametric Estimation*, 2009

What makes a mechanism good?

Let M_0, M_1 be the distributions of the released Y under P_0 and P_1:

$$M_i(y) = \sum_x Q(y \mid x) P_i(x)$$

A good mechanism keeps M_0 and M_1 "easy to tell apart"

A test can make two errors: $\alpha = P(\text{guess } P_1 \mid P_0)$ and $\beta = P(\text{guess } P_0 \mid P_1)$

TV	$\boldsymbol{E_\gamma}$	**KL**	**Hellinger**
$f(t) = \frac{1}{2}\lvert t-1\rvert$	$f(t) = (t-\gamma)_+$	$f(t) = t \log t$	$f(t) = \left(\sqrt{t}-1\right)^2$

All four are f-divergences: $D_f(M_0 \| M_1) = \sum_y M_1(y)\, f\left(\frac{M_0(y)}{M_1(y)}\right)$

f-divergences

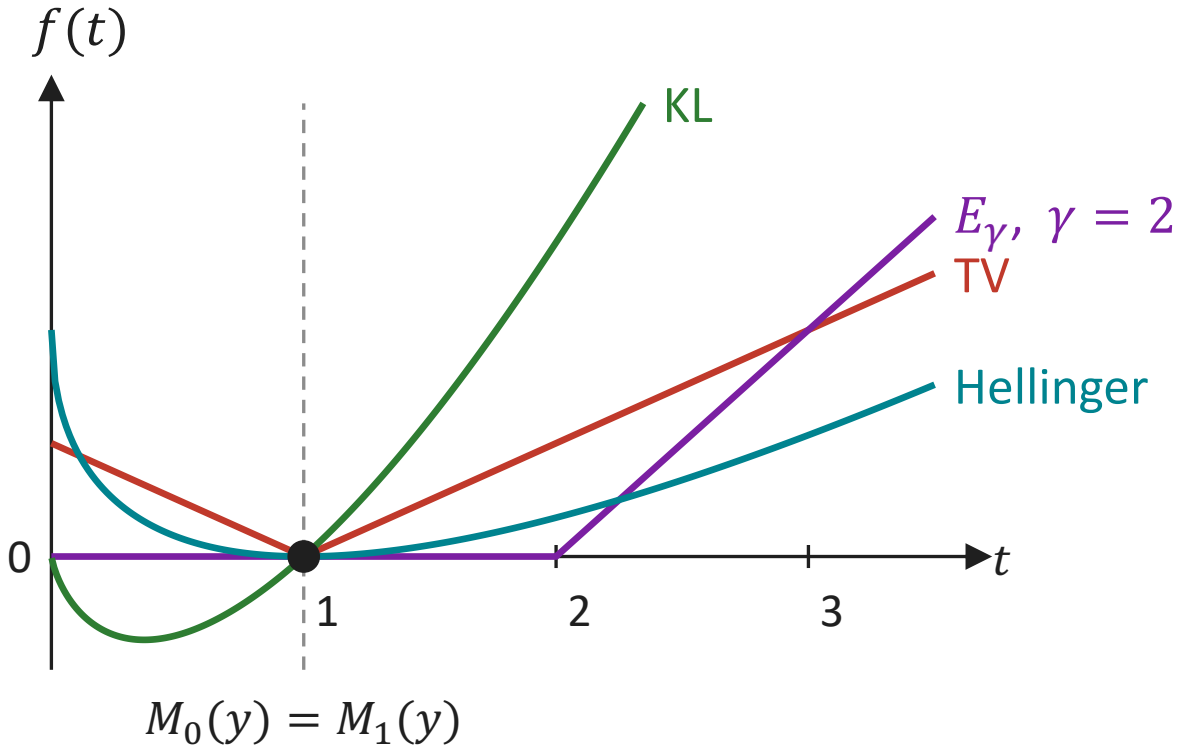
TV: $f(t) = \frac{1}{2}|t-1|$

KL: $f(t) = t \log t$

E_γ: $f(t) = (t-\gamma)_+$

Hellinger: $f(t) = \left(\sqrt{t}-1\right)^2$

$$D_f(M_0 \| M_1) = \sum_y M_1(y)\, f\left(\frac{M_0(y)}{M_1(y)}\right)$$



$t = M_0(y)/M_1(y)$: how much more likely output y is under P_0 than under P_1

$f(1) = 0$: an output equally likely under both tells nothing

f convex: $D_f(M_0 \| M_1)$ is jointly convex in (M_0, M_1)

Csiszár, "Information-type measures of difference of probability distributions and indirect observations", Studia Sci. Math. Hungar. 1967

The design problem

Given P_0, P_1, a privacy level ε and an f-divergence, find an ε-LDP mechanism Q keeping M_0, M_1 furthest apart (Kairouz, Oh & Viswanath 2016):

$$\max_{Q\ \varepsilon\text{-LDP}} D_f(M_0 \,\|\, M_1)$$

The best Q depends on P_0, P_1, ε and f

Variant: Q has at most ℓ outputs (Pensia, Asadi, Jog & Loh 2025)

Kairouz, Oh & Viswanath, "Extremal mechanisms for local differential privacy", JMLR 2016
Pensia, Asadi, Jog & Loh, "Simple binary hypothesis testing under local differential privacy and communication constraints", IEEE Trans. Inf. Theory 2025

Prior work: optimal mechanisms in different privacy regimes



small ε

Binary mechanism

optimal for every f
(for TV: at every ε)

$$b = \begin{cases} 0 & \text{if } P_0(x) \geq P_1(x) \\ 1 & \text{if } P_0(x) < P_1(x) \end{cases}$$

$$Q(y \mid x) = \begin{cases} e^{\varepsilon}/(e^{\varepsilon}+1) & \text{if } y = b \\ 1/(e^{\varepsilon}+1) & \text{if } y \neq b \end{cases}$$

Kairouz, Oh & Viswanath, “Extremal mechanisms for local differential privacy”, JMLR 2016

Prior work: optimal mechanisms in different privacy regimes



large ε

k-ary RR

optimal for KL

$$Q(y \mid x) = \begin{cases} e^{\varepsilon}/(e^{\varepsilon} + k - 1) & \text{if } y = x \\ 1/(e^{\varepsilon} + k - 1) & \text{if } y \neq x \end{cases}$$

Kairouz, Oh & Viswanath, “Extremal mechanisms for local differential privacy”, JMLR 2016

Prior work: optimal mechanisms in different privacy regimes



small ε	intermediate ε	large ε
Binary mechanism	**No general closed-form solution known**	**k-ary RR**
optimal for every f (for TV: at every ε)		optimal for KL

Kairouz, Oh & Viswanath, "Extremal mechanisms for local differential privacy", JMLR 2016

Prior work: the cost of finding the optimal mechanism

	Prior work
No limit on outputs	**Kairouz et al. 2016** LP with 2^k variables (exponential in k)
At most ℓ outputs	**Pensia et al. 2025** $k^{O(\ell^2)} \cdot 2^{O(\ell^3 \log \ell)}$ (polynomial only for constant ℓ)

k = number of possible inputs

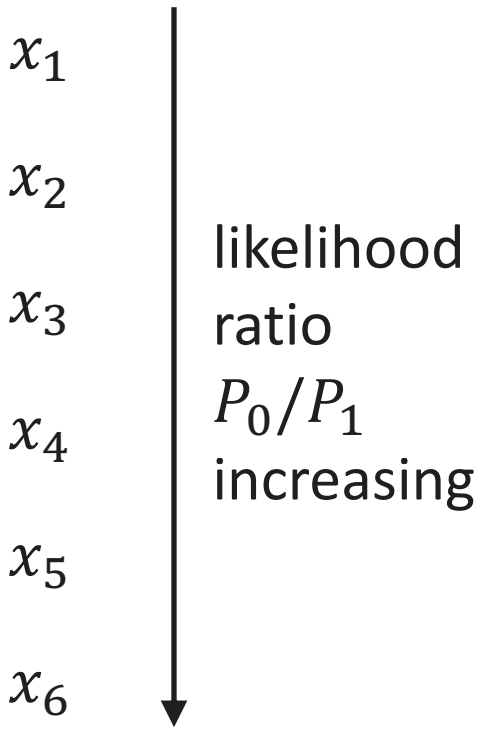
Our results

Structure	For every ε and every f, an optimal mechanism has a simple form: Sort, Partition, Randomize (SPR)
Algorithm	Find it in $O(k^3)$ time or $O(\ell k^2)$ with at most ℓ outputs
E_γ**-divergence**	Closed-form optimal mechanism

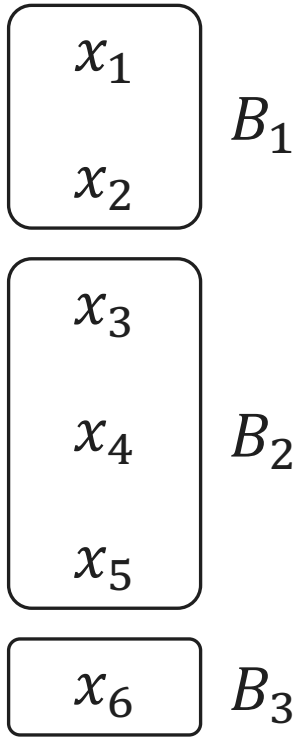
Optimal mechanisms: Sort, Partition, Randomize

1 Sort

the input alphabet



2 Partition



3 Randomize

Release the block label, privatized:

true block $\dfrac{e^{\varepsilon}}{e^{\varepsilon}+m-1}$

each other block $\dfrac{1}{e^{\varepsilon}+m-1}$

m = number of blocks

SPR is optimal

We want $\max_{Q\ \varepsilon\text{-LDP}} D_f(M_0 \,\|\, M_1)$

For every ε and every f-divergence, an optimal SPR mechanism exists.
With at most ℓ outputs, an SPR mechanism with at most ℓ blocks is optimal.

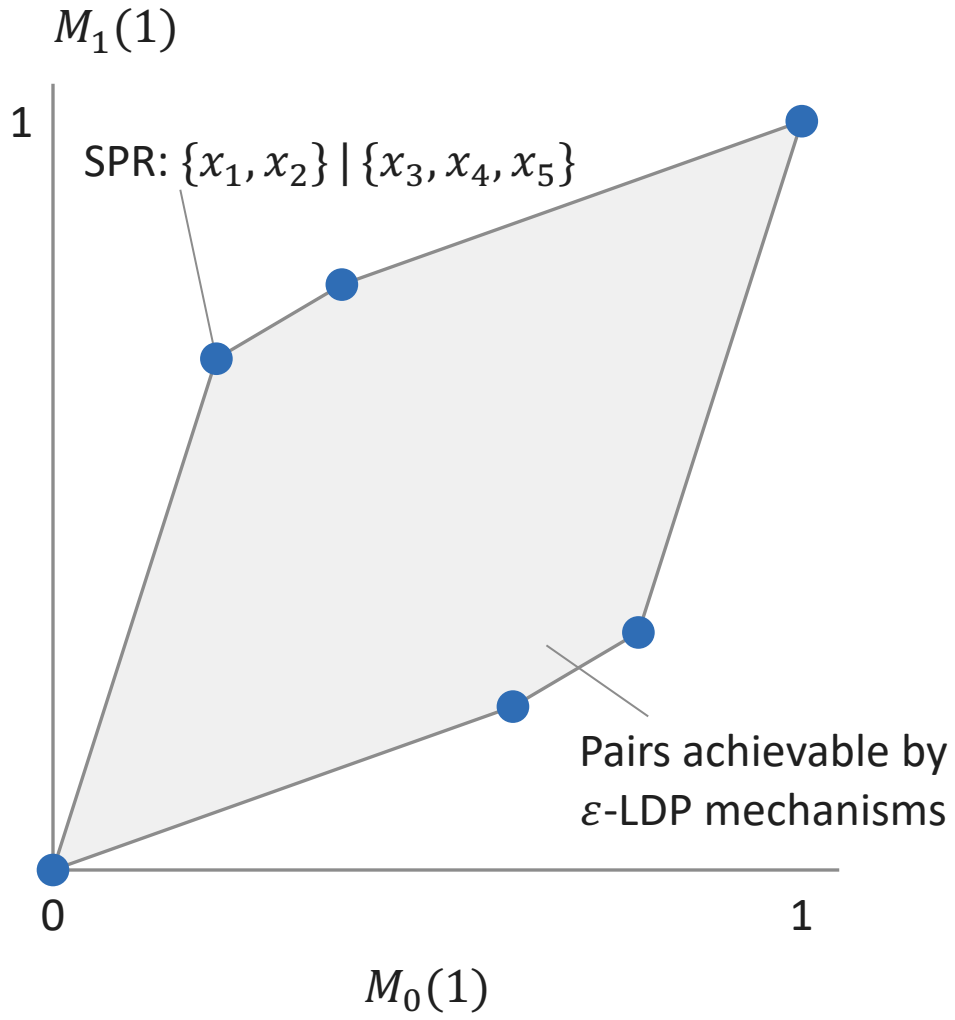
Special cases

- Binary mechanism = 2 blocks, cut at $P_0/P_1 = 1$
- k-ary RR = k singleton blocks

As $\varepsilon \to \infty$, RR stops randomizing, and SPR becomes a likelihood-ratio quantizer, which is optimal without privacy (Tsitsiklis 1993)

Tsitsiklis, "Extremal properties of likelihood-ratio quantizers", IEEE Trans. Commun. 1993

Why SPR mechanisms are optimal



Example: $k = 5$ inputs, two outputs, $\varepsilon = 2$

M_0, M_1: distributions of the released Y under P_0, P_1

$$M_i(y) = \sum_x Q(y \mid x) P_i(x)$$

1 **Convexity**

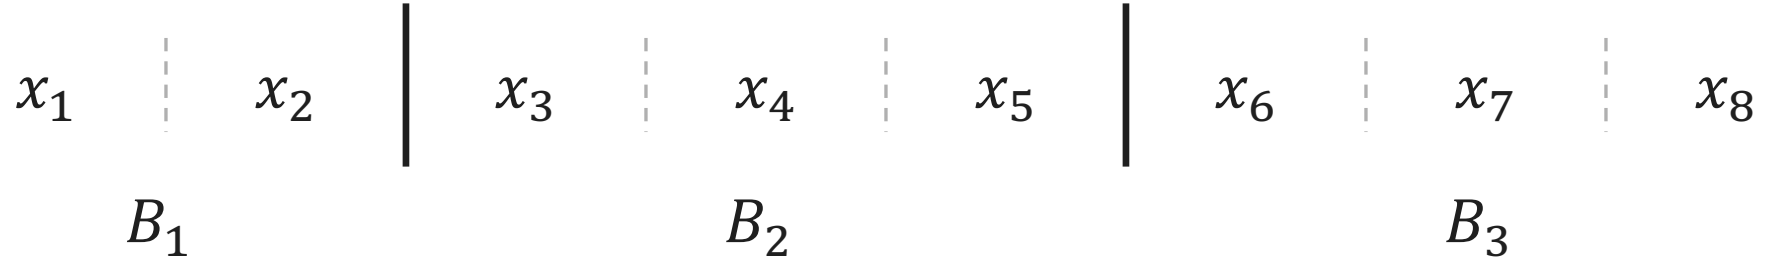
$D_f(M_0 \| M_1)$ is convex in (M_0, M_1), so an optimum is attained at a vertex

2 **Our structural result**

Every vertex is achieved by an SPR mechanism

Choosing the partition

Inputs sorted by $P_0(x)/P_1(x)$

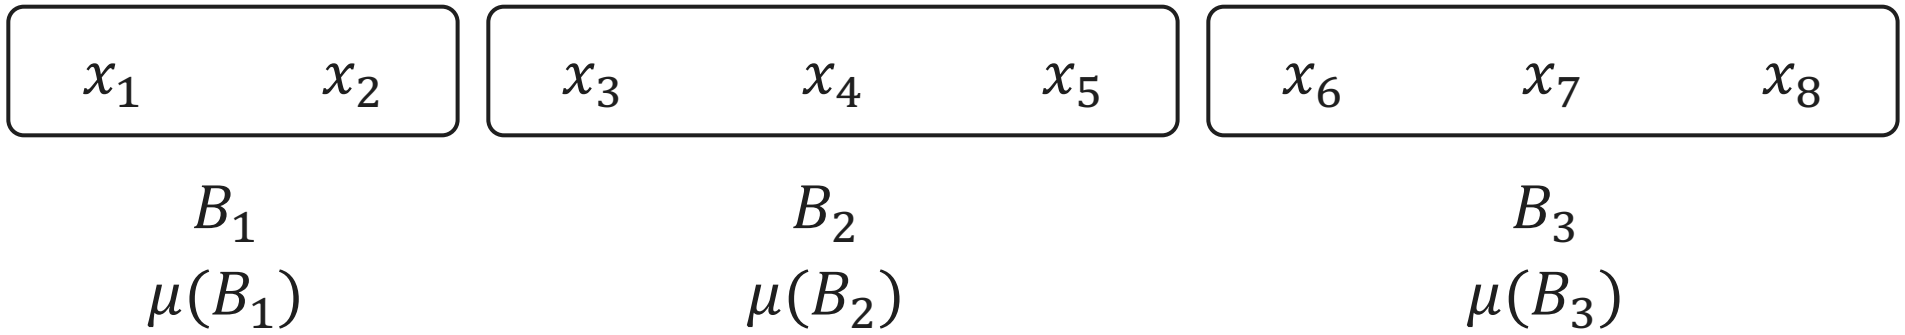


Choose where to cut: the cuts fix the blocks and their number m

$k-1$ gaps, each cut or not: 2^{k-1} partitions

Maximizing a sum of block scores

Inputs sorted by $P_0(x)/P_1(x)$

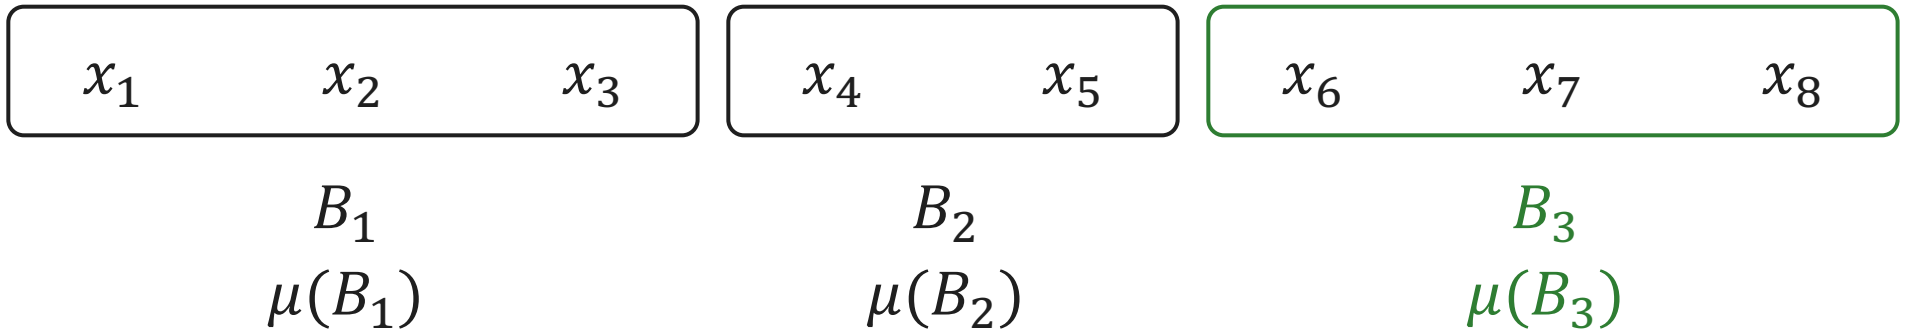


$$D_f(M_0 \| M_1) = \frac{\mu(B_1) + \cdots + \mu(B_m)}{e^{\varepsilon} + m - 1}$$

For fixed P_0, P_1, ε, f, each block has a score $\mu(B)$ determined by $P_0(B)$ and $P_1(B)$.

Maximizing a sum of block scores

Inputs sorted by $P_0(x)/P_1(x)$



$$D_f(M_0 \| M_1) = \frac{\mu(B_1) + \cdots + \mu(B_m)}{e^{\varepsilon} + m - 1}$$

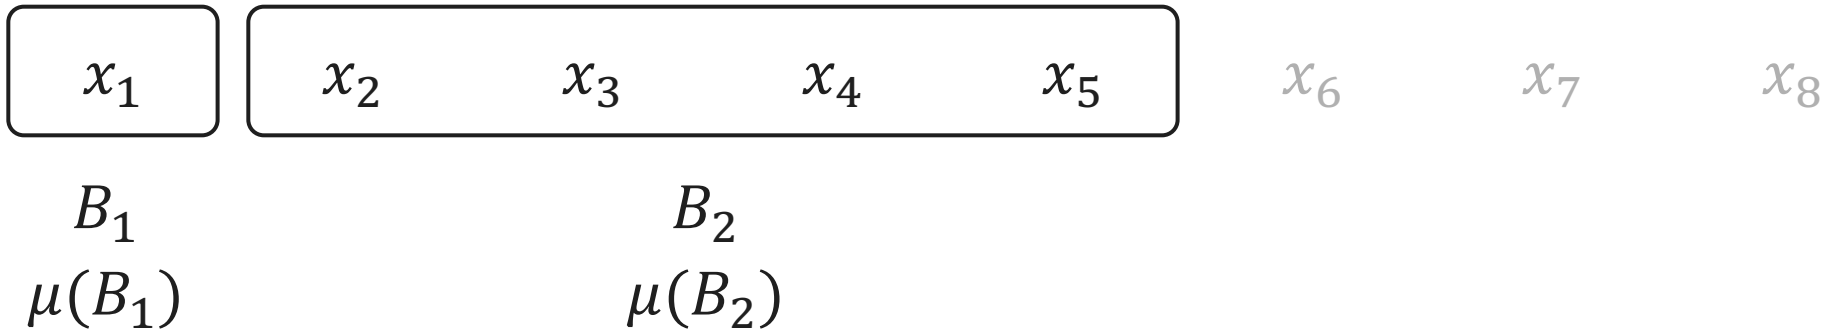
For fixed P_0, P_1, ε, f, each block has a score $\mu(B)$ determined by $P_0(B)$ and $P_1(B)$.

Moving the earlier cut leaves the last block and its score unchanged.

For a fixed number of blocks, maximize the sum of block scores.

A smaller problem: two blocks

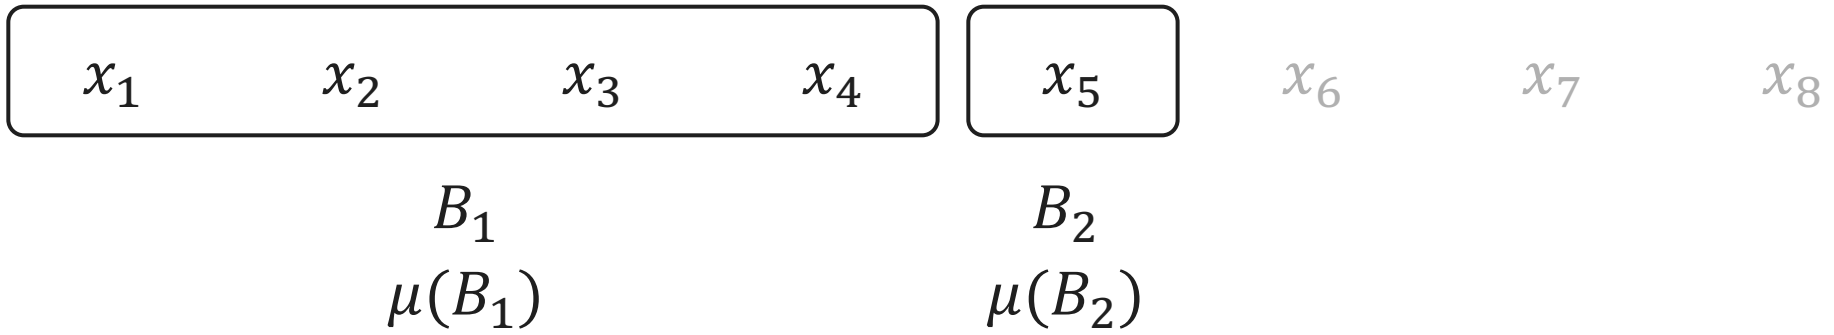
Partition the first five inputs into exactly two blocks



Try all four possible cut positions.
Add the two block scores and save the largest total.

A smaller problem: two blocks

Partition the first five inputs into exactly two blocks

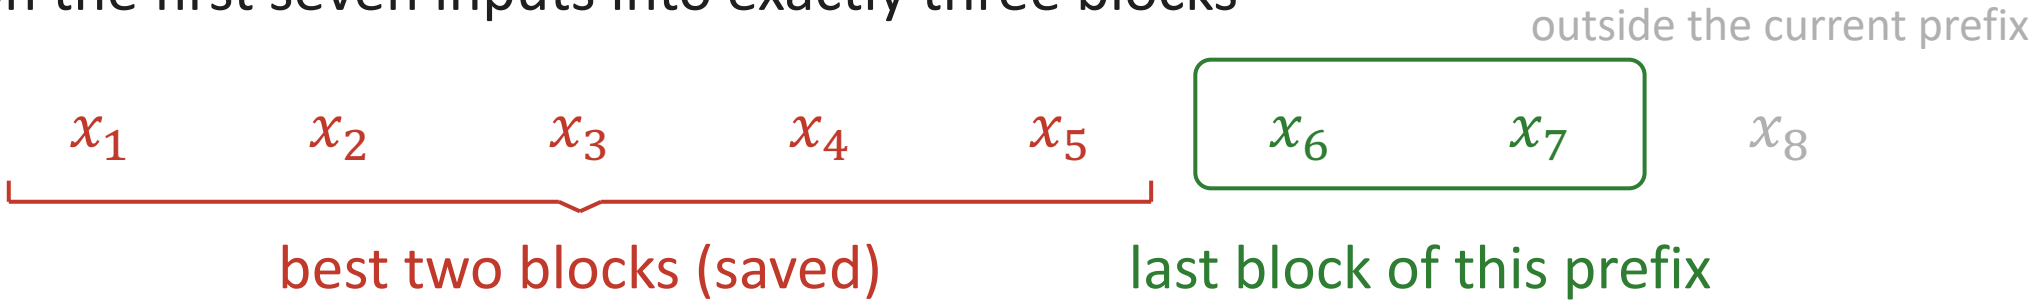


Try all four possible cut positions.
Add the two block scores and save the largest total.

Do the same for every prefix: this saves the best two-block score for the first i inputs, for each i.

Three blocks: fix the last block

Partition the first seven inputs into exactly three blocks

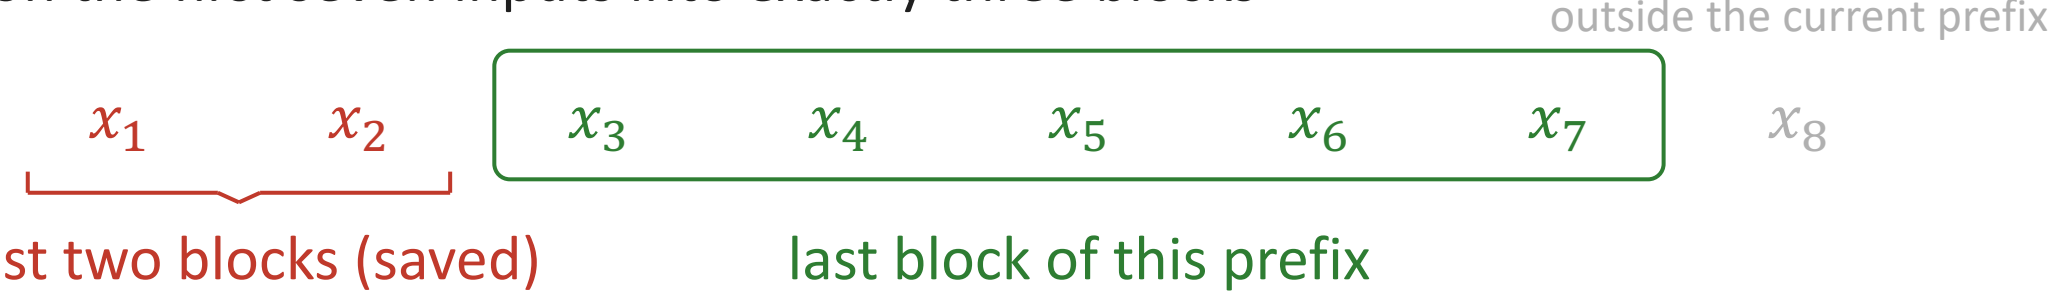


Best total with this last block:

saved best two-block score for the earlier inputs + score of the last block

Three blocks: fix the last block

Partition the first seven inputs into exactly three blocks



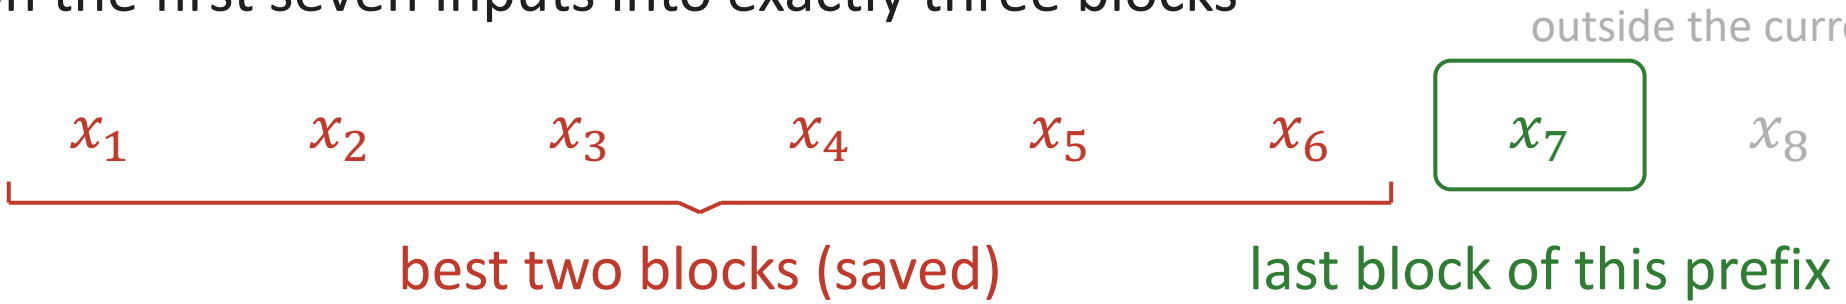
Best total with this last block:

saved best two-block score for the earlier inputs + score of the last block

Try every last block; keep the largest total.

Three blocks: fix the last block

Partition the first seven inputs into exactly three blocks



Best total with this last block:

saved best two-block score for the earlier inputs + score of the last block

Try every last block; keep the largest total.

Every three-block partition has a last block ending at the seventh input.
Trying every valid starting position therefore finds the best partition.

The general recurrence

$F[m, i]$: largest sum of block scores for the first i inputs in exactly m blocks

For $2 \leq m \leq i \leq k$:

$$F[m,i] = \max_{m-1 \leq t < i} \left\{ \underbrace{F[m-1,t]}_{\text{saved score for earlier inputs}} + \underbrace{\mu(\{x_{t+1}, \ldots, x_i\})}_{\text{last block's score}} \right\}$$

Base case: one block

$$F[1,i] = \mu(\{x_1, \ldots, x_i\}), \qquad 1 \leq i \leq k$$

With one block, there is no cut to choose: all i inputs belong to that block.

Choosing the number of blocks

For each m, $F[m, k]$ gives the largest sum of block scores for **all k inputs in exactly m blocks**.

Restore the denominator and compare:

$$\text{Optimal divergence} = \max_{1 \le m \le k} \frac{F[m, k]}{e^{\varepsilon} + m - 1}$$

Choose a maximizing m, then follow the saved cuts to recover its partition.

In our example, $k = 8$: compare block counts $m = 1, \ldots, 8$.

Comparison with prior work

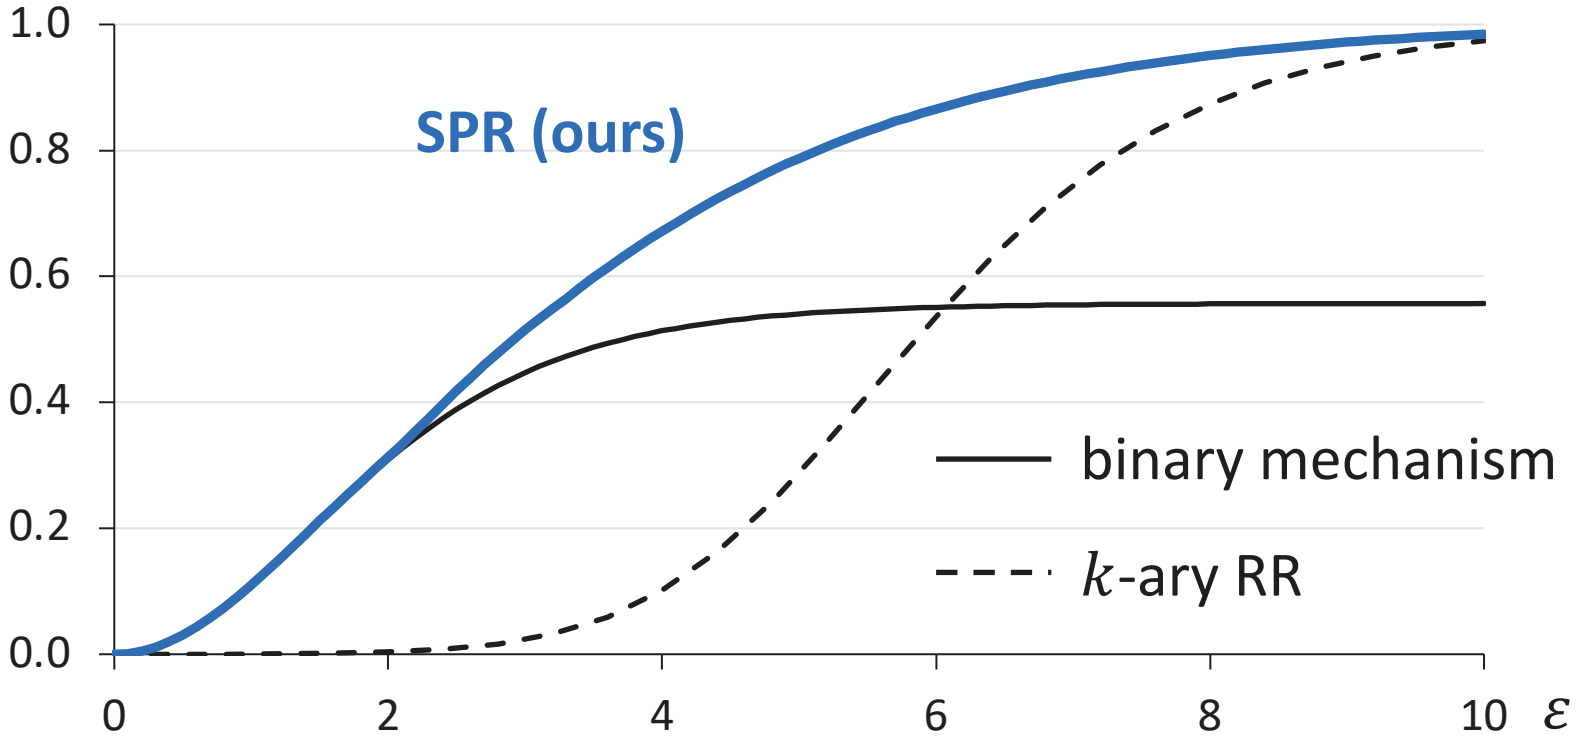
	Prior work	This work
No limit on outputs	**Kairouz et al. 2016** LP with 2^k variables (exponential in k)	$O(k^3)$
At most ℓ outputs	**Pensia et al. 2025** $k^{O(\ell^2)} \cdot 2^{O(\ell^3 \log \ell)}$ (polynomial only for constant ℓ)	$O(\ell k^2)$

k = number of possible inputs

Exact optimum at scales that were out of reach

$k = 100$ symbols, averaged over 100 random pairs (P_0, P_1)

$D_{\mathrm{KL}}(M_0 \| M_1)$ / $D_{\mathrm{KL}}(P_0 \| P_1)$ (1 = no loss)



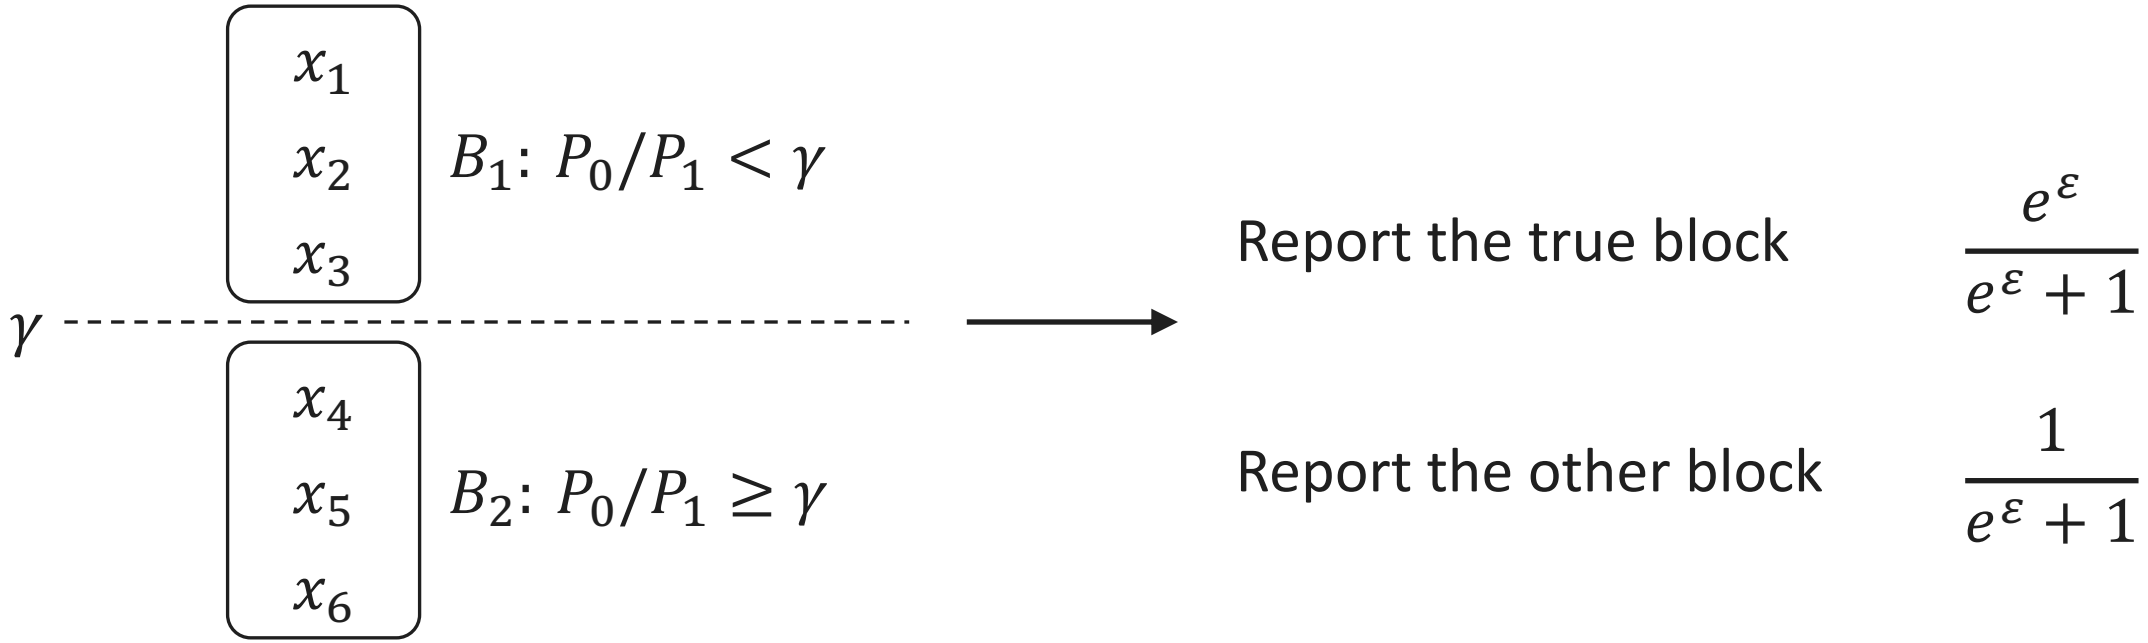
Exact LP
$2^{100} \approx 10^{30}$ variables

SPR dynamic program
≈ 2 ms per optimal mechanism

Hockey-stick divergence E_γ

Maximizing E_γ = minimizing $\alpha + \gamma\beta$

Optimal: the **generalized binary mechanism**



Optimal for every ε and every number of inputs (closed form)

With $\gamma = 1$, this recovers the TV result of Kairouz, Oh & Viswanath (2016)

Takeaways

1	**Structure**	For every ε and every f-divergence, optimal mechanisms are Sort, Partition, Randomize
2	**Algorithm**	Exact optimum in $O(k^3)$ instead of an LP with 2^k variables
3	**E_γ-divergence**	Two blocks suffice: cut once at $P_0/P_1 = \gamma$

Open questions

Mechanism design with imperfect knowledge of the distributions

Our optimal mechanism depends on P_0, P_1.

- How well does it perform when those distributions are only approximately known?
- Can we choose mechanisms that perform well despite that uncertainty?

More than two hypotheses

- How should we design a locally private mechanism to minimize testing error among several known distributions?

References

Awan & Slavković, “Differentially private uniformly most powerful tests for binomial data”, NeurIPS 2018

Canonne, Kamath, McMillan, Smith & Ullman, “The structure of optimal private tests for simple hypotheses”, STOC 2019

Csiszár, “Information-type measures of difference of probability distributions and indirect observations”, Studia Sci. Math. Hungar. 1967

Duchi, Jordan & Wainwright, “Local privacy and statistical minimax rates”, FOCS 2013

Dwork, McSherry, Nissim & Smith, “Calibrating noise to sensitivity in private data analysis”, TCC 2006

Kairouz, Oh & Viswanath, “Extremal mechanisms for local differential privacy”, JMLR 2016

Pensia, Asadi, Jog & Loh, “Simple binary hypothesis testing under local differential privacy and communication constraints”, IEEE Trans. Inf. Theory 2025

Tsitsiklis, “Extremal properties of likelihood-ratio quantizers”, IEEE Trans. Commun. 1993

Warner, “Randomized response: a survey technique for eliminating evasive answer bias”, JASA 1965