Privacy & Learning

A Story of an alliance between enemies

J. SAKETHA NATH (IITH)

JOINT WORK WITH ARUN IYER & SUNITA SARAWAGI

Private Data Release: Set-up

Private Data Release: Set-up

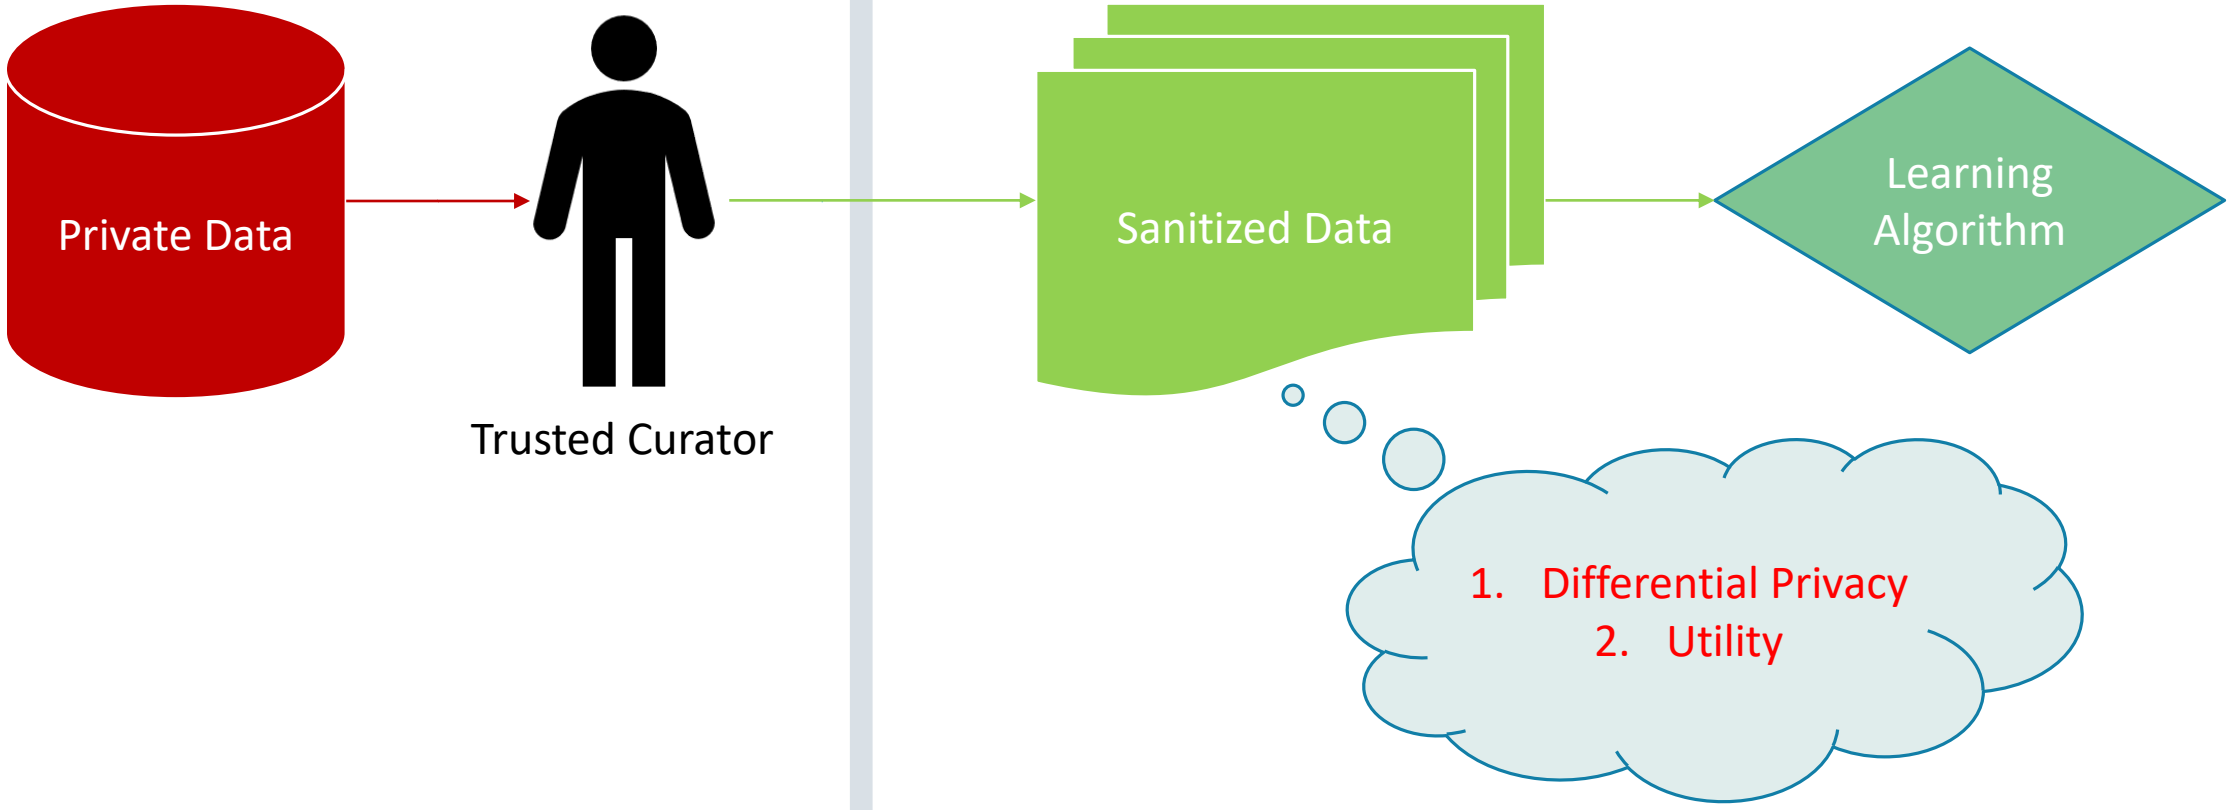

Class-Ratio (CR) Estimation

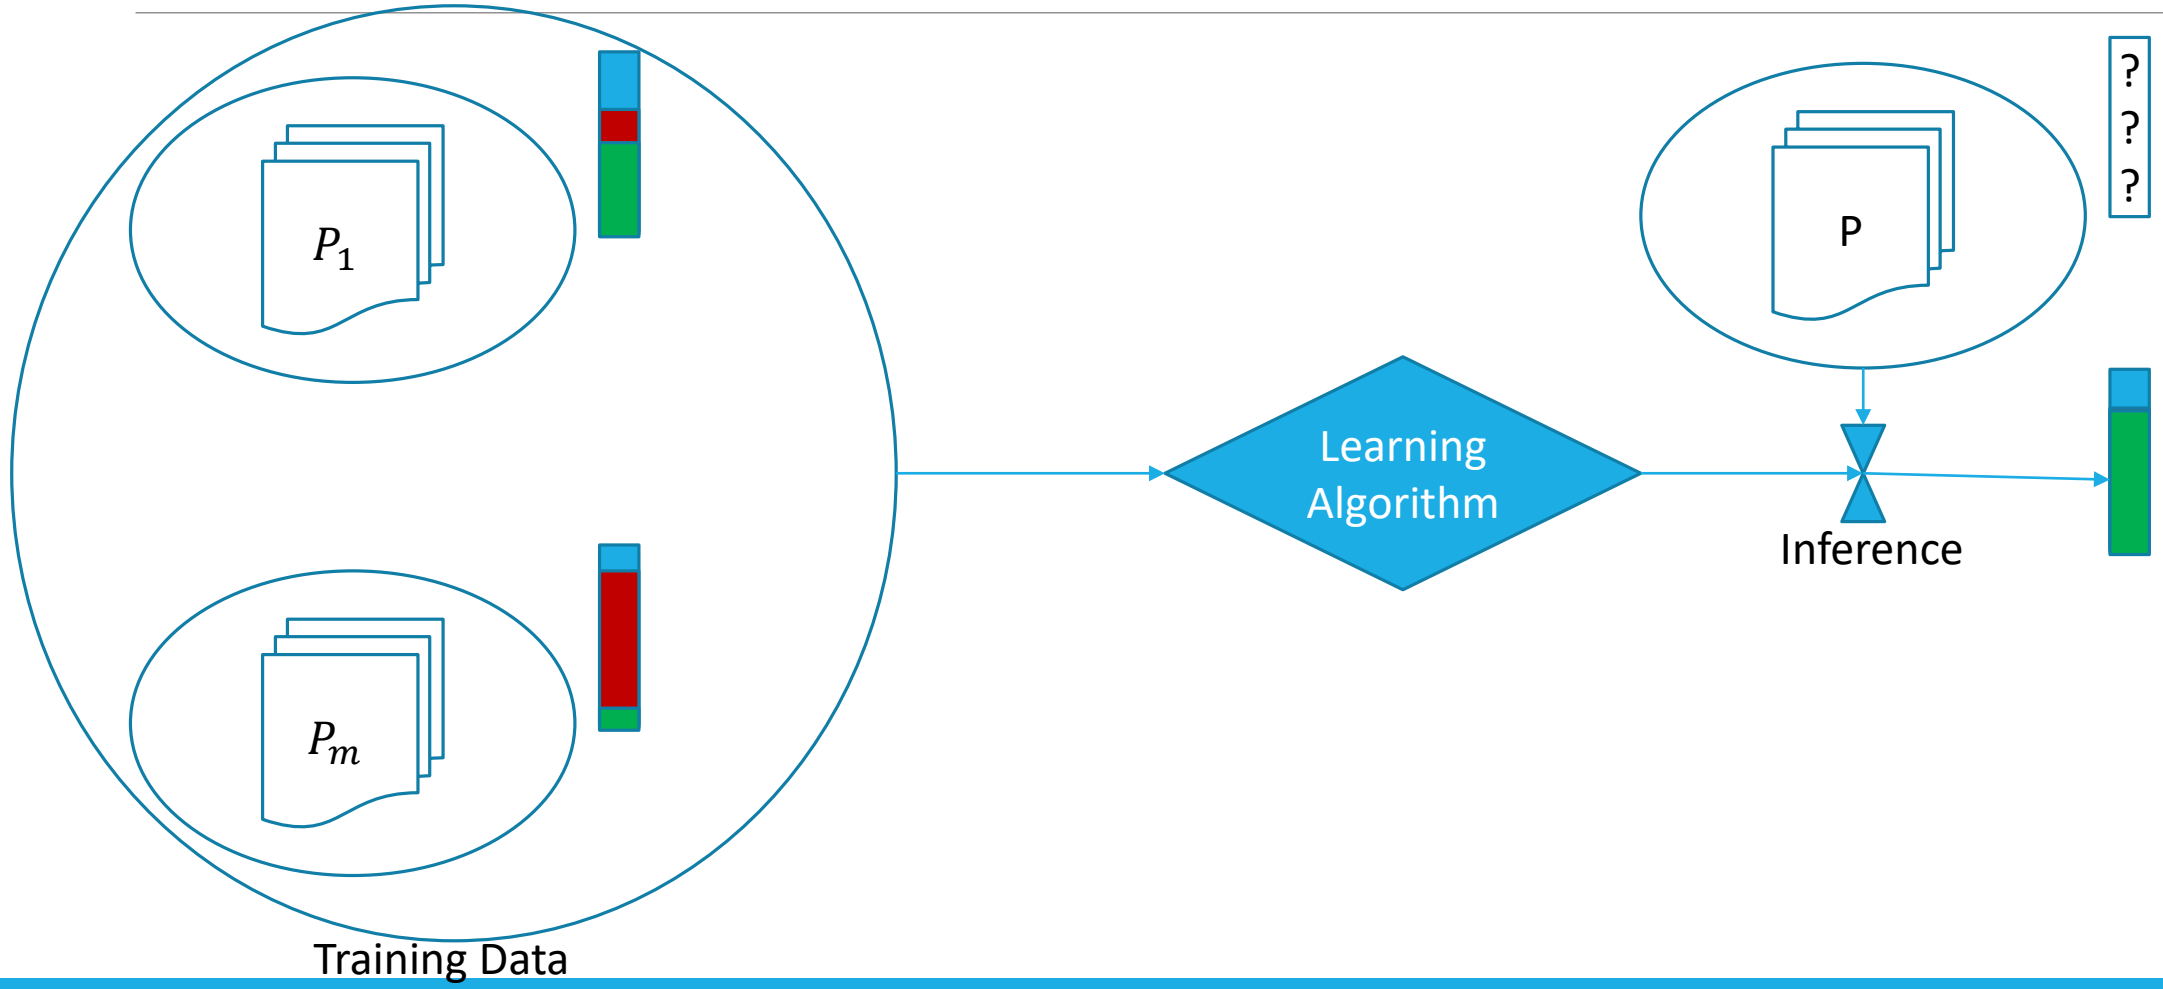

Label Sanitization

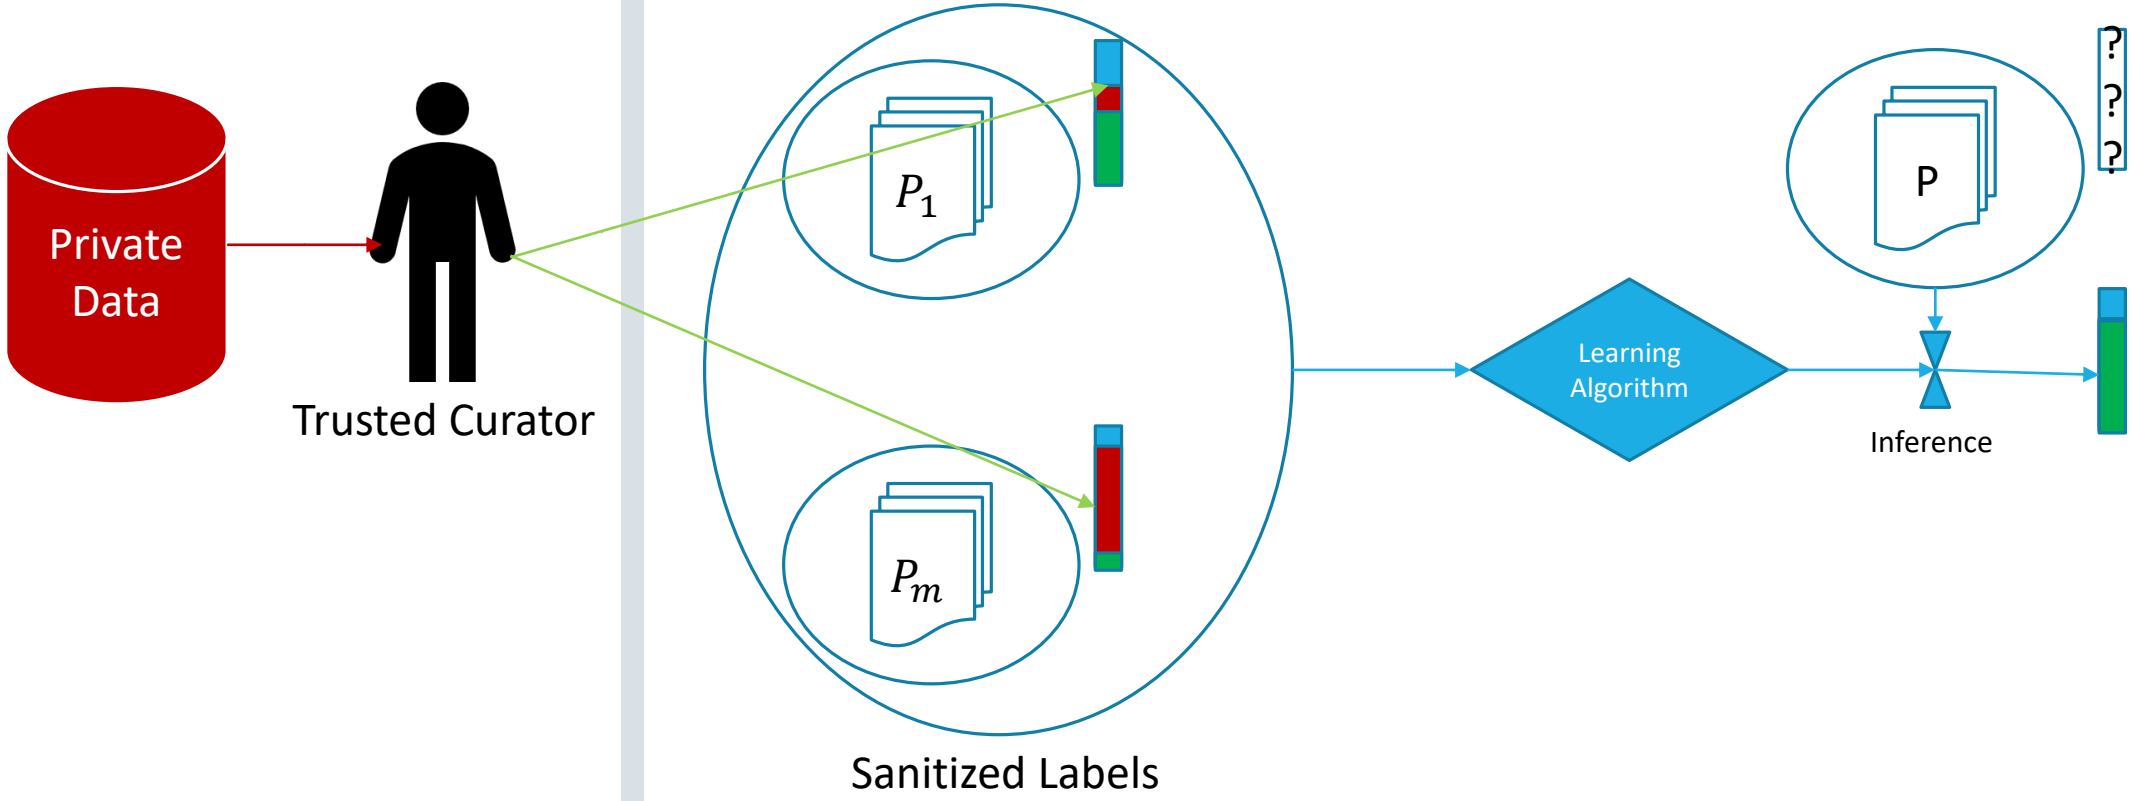

Related work

1. Ashwin M. et.al.'08 in DB setting

   a) $g(D) \sim Dir(m_1 + z_1, \ldots, m_c + z_c)$, where z_i is noise

Related work

1. Ashwin M. et.al.'08 in DB setting

   a) $g(D) \sim Dir(m_1 + z_1, \ldots, m_c + z_c)$, where z_i is noise

   b) Guarantees (ϵ, δ) differential privacy

$$P[g(D_1) \in B] \leq \delta + e^{\epsilon} P[g(D_2) \in B]$$

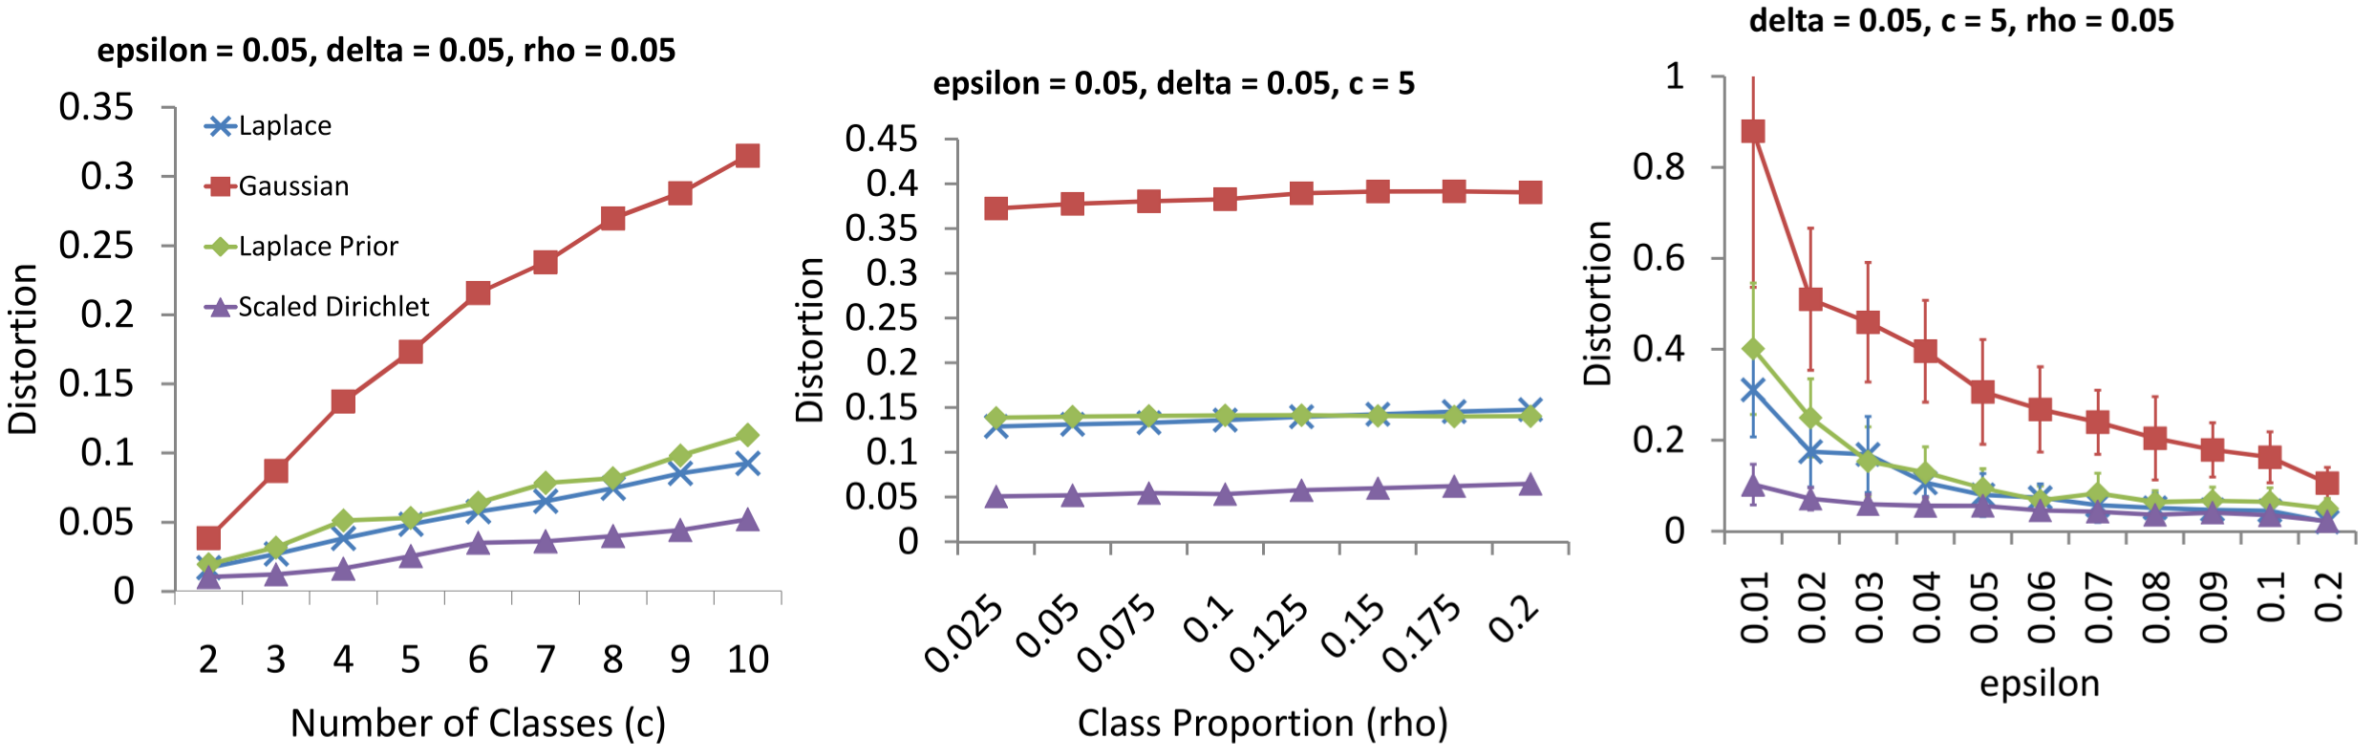

Figure 1: Comparing distortion of privacy mechanisms under varying settings of parameters. The first plot is for increasing number of classes, second plot for changing class proportions on five classes, third plot for increasing ϵ on five classes.

Privacy Trade-offs

Our Mechanism

$$g(D) \sim Dir(\sigma m_1, \ldots, \sigma m_c)$$

1. Unbiased for CR estimation
2.

Our Mechanism

$$g(D) \sim Dir(\sigma m_1, \ldots, \sigma m_c)$$

1. Unbiased for CR estimation
2. σ controls privacy vs. distortion

Our Mechanism

$$g(D) \sim Dir(\sigma m_1, \ldots, \sigma m_c)$$

1. Unbiased for CR estimation
2. σ controls privacy vs. distortion
3. Guaranteed (ϵ, δ) Privacy unless m_i are too low

Asymptotic Behaviour

CR Estimation	Differential Privacy
1. Asymptotic consistency. 2. Skewed class ratios are good $err \propto maxsing(\hat{P})$	1. Asymptotic case has guaranteed privacy for any ϵ, δ. $\delta = \frac{\left(\max_{ij} \frac{1-\hat{\rho}_i}{\hat{\rho}_i} + e^{2\frac{\epsilon}{\sigma}} \frac{1-\hat{\rho}_j}{\hat{\rho}_j}\right) + 2e^{\frac{\epsilon}{\sigma}}}{(\sigma m + 1)\left(1 - e^{\frac{\epsilon}{\sigma}}\right)^2}$ 2. Uniform class-ratios are good.

Allies or Enemies?

LEMMA 4. *With probability* $1-\zeta$ *we have:*

$$\|\tilde{\mathcal{P}}-\hat{\mathcal{P}}\|_F \leq \frac{Mc}{\sqrt{\zeta}} \max_{i\in\{1,\ldots,M\};y\in\mathcal{Y}} \sqrt{\frac{\hat{\rho}_{iy}\left(1-\hat{\rho}_{iy}\right)}{\sigma_i\left(\epsilon,\delta\right)m_i+1}}$$

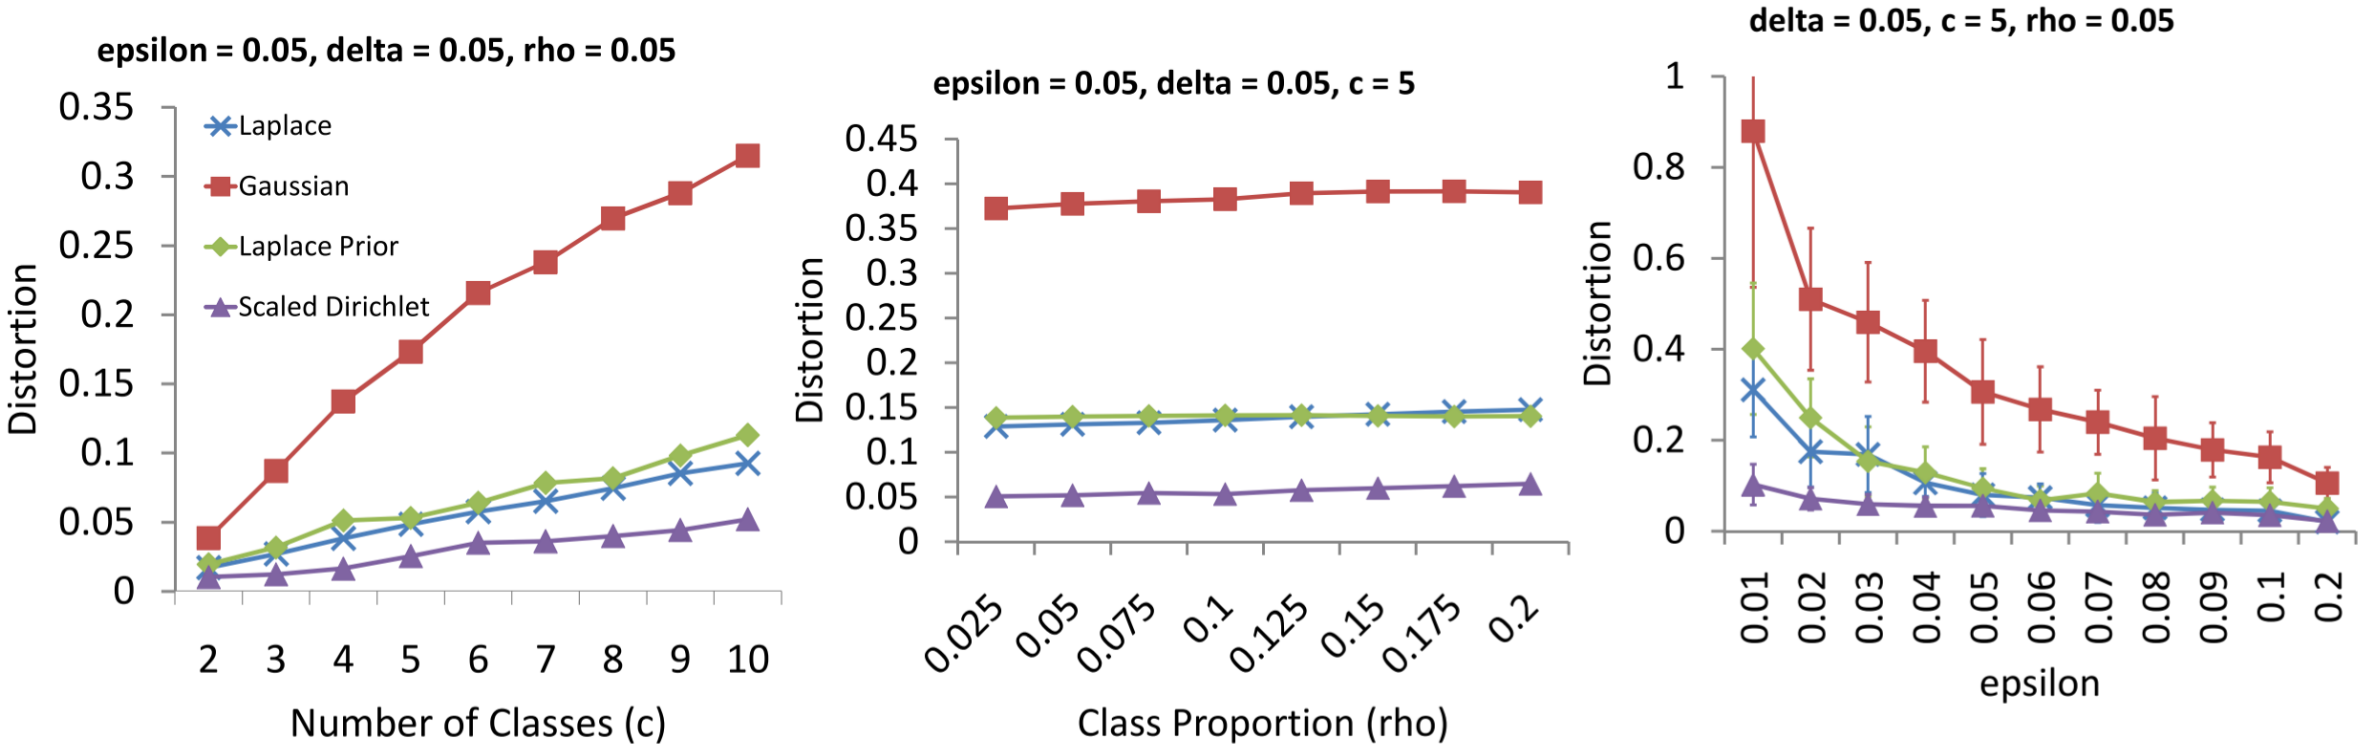

Figure 1: Comparing distortion of privacy mechanisms under varying settings of parameters. The first plot is for increasing number of classes, second plot for changing class proportions on five classes, third plot for increasing ϵ on five classes.

Privacy Trade-offs

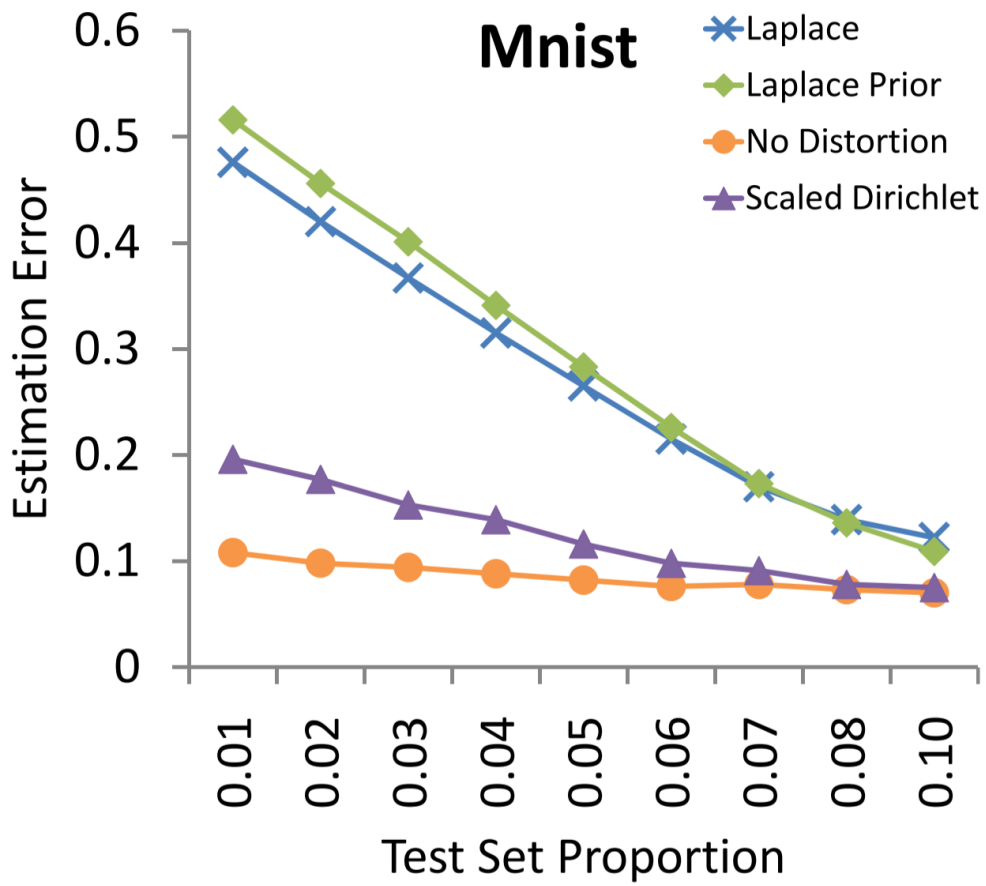

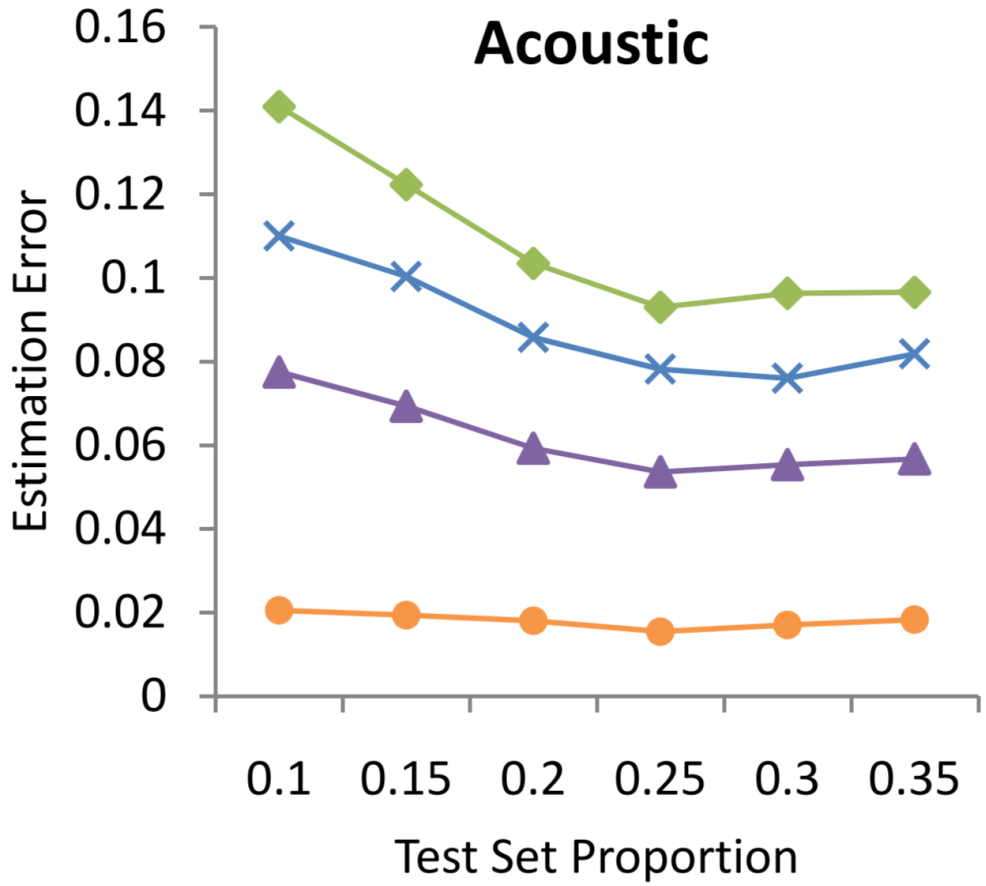

Learning with Privacy

MNIST has higher number of classes

Related Work

1. Series of negative results by Dinur&Nissim’03, Dwork et.al’07’08 in DB settings

Related Work

1. Series of negative results by Dinur&Nissim'03, Dwork et.al'07'08 in DB settings

2. Seminal work by Blum et.al.'08 gives positive result for learning settings
   a) Perturbations maintaining fractional counts
   b) Mechanism computationally infeasible

Related Work

1. Series of negative results by Dinur&Nissim'03, Dwork et.al'07'08 in DB settings

2. Seminal work by Blum et.al.'08 gives positive result for learning settings
   a) Perturbations maintaining fractional counts
   b) Mechanism computationally infeasible

3. Dwork et.al.'09 formalizes the trade-offs between privacy, utility, computational feasibility

Summary

Learning bounds highlighting trade-offs

1. Showed in what sense privacy and learning may aid or may compete with each other

2. Explicit bounds involving the trade-offs

3. Motivates learning theoretic formalisms for other issues:
   a) Fairness
   b) Compactness
   c) Robustness
   d) Explainability
   e) Causality
   f) Tractability