https://arxiv.org/abs/2009.12920

Privacy and Fairness in Data-driven Personalized Revenue Management

Yining Wang

Warrington College of Business,

University of Florida

Data-driven revenue management

- Using **data analytics** to help **revenue / profit** decisions.

Dynamic pricing

Best flights ⓘ

Total price includes taxes + fees for 1 adult. Additional bag fees and other fees may apply.

Sort by:

7:20 AM – 12:59 PM United · Operated by Air Wisconsin DBA United Express	5 hr 39 min MCO–YUL	1 stop 1 hr 38 min IAD	$141
10:40 AM – 4:58 PM American · Operated by PSA Airlines as American Eagle	6 hr 18 min MCO–YUL	1 stop 1 hr 54 min PHL	$141
1:10 PM – 7:16 PM Air Canada · Operated by Air Canada Rouge, Air Canada...	6 hr 6 min MCO–YUL	1 stop 2 hr 10 min YYZ	$174

Choose a ride, or swipe up for more

UberX 4 11:14am dropoff	$8-10
UberXL 11:16am	$11-13
Comfort New 11:14am	$10-11

Uber in California shows a range of prices instead of up-front prices

Data-driven revenue management

- Using **data analytics** to help **revenue / profit** decisions.

Dynamic Assortment Planning

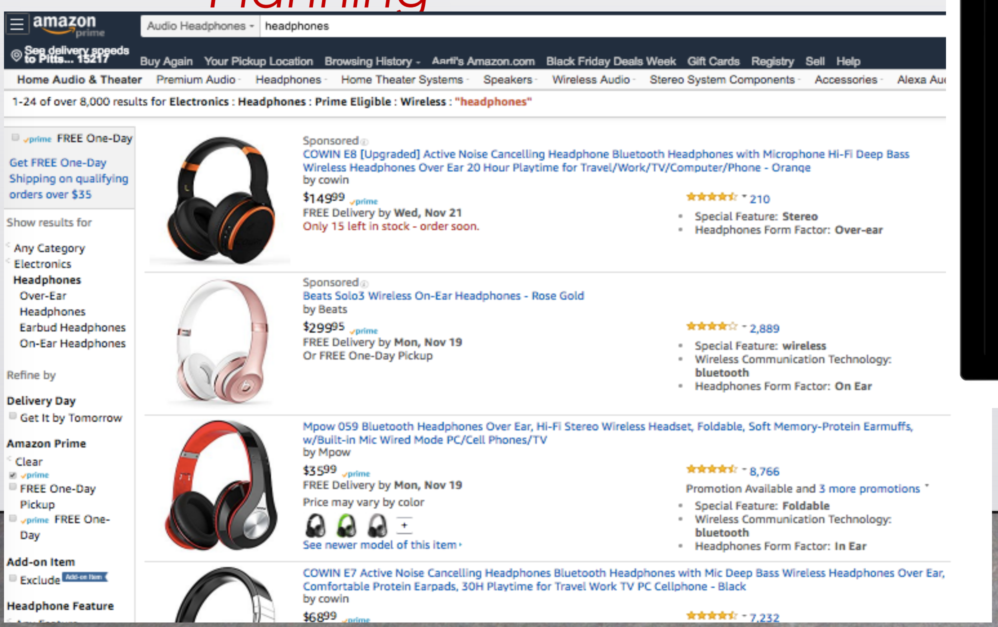

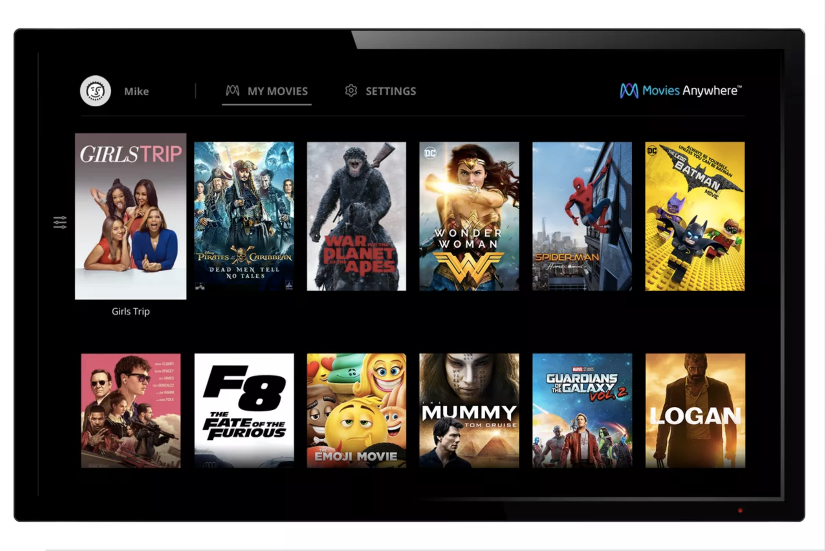

Data-driven revenue management

- The **population** approach:
  - Use *population data* such as the average demand or click-through rates over a region to make general price, promotion and inventory decisions.
- The **personalized** approach
  - Use *personalized data* to make individualized price/promotion/recommendation decisions.
  - More detailed, refined with higher profits

Personalized revenue management

- Example: **Yamibuy.com** (online retail)

User profiles

Personalized revenue management

- Example: **Yamibuy.com** (online retail)

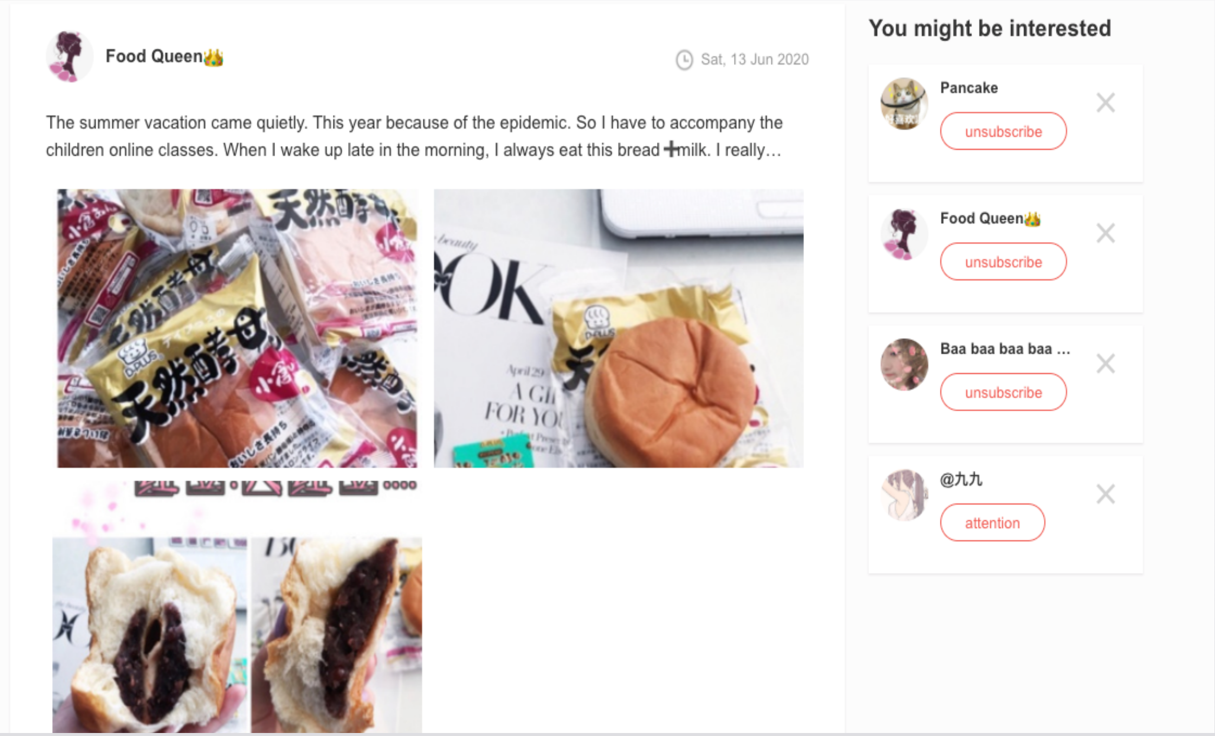

Social Network

Personalized revenue management

- Example: **Yamibuy.com** (online retail)
  - Personalized price decisions: set higher prices for those who target higher brands?
  - Personalized recommendation/promotion decisions: promote new/emerging items to social influencers (i.e. many posts / followers)?

User profiles

Social Network

Data privacy in Personalized revenue management

- Personalized data involved in data-driven decision making are sensitive and private.
    - Example: age, gender, telephone number
    - More serious: medical history (drug stores), credit history (credit cards/loans)
- **Privacy breaches** of personalized data can have serious ethical and legal consequences!

Personalized revenue management

- **Question.** When using personalized data to make decisions, how to avoid inadvertently leaking private data of the users?

Data-driven personalized pricing

- The model.

- ○ *T* consumers, arriving sequentially.

x_1, p_1 x_2, p_2 x_3, p_3 … x_T, p_T

y_1 y_2 y_3

Personal info. (age, gender, etc.)
History (purchase, credit, medical, etc.)
Social network (e.g., page-rank)

Customer *t* profile: x_t
Posted price: p_t

Vector representation
$\phi_t = \phi(x_t, p_t)$

$$E[y_t | x_t, p_t] = f(\langle \phi_t, \theta^* \rangle)$$

Data-driven personalized pricing

- The "learning-while-doing" framework: learning the model θ^* while optimizing prices $\{p_t\}_{t=1}^T$
  - Many existing works in the literature. Zeevi & Besbes'09,15, Broder & Rusmevichientong'12, Chen & Gallego'19, Wang et al.'14, Keskin & Zeevi'14
  - The key principle: **"Optimism in the Face of Uncertainty"** (OFU), by Abbasi-Yadkori *et al.* in *NeurIPS, 2011.*

The predicted demand at *p*

$$\hat{p}_t = \arg\max_p p \times \left[f\left(\phi_t, \hat{\theta}_{t-1}\right) + \gamma \sqrt{\phi_t^T \Lambda_{t-1}^{-1} \phi_t} \right]$$

Confidence interval of the prediction

Concerns over privacy leakage

- The customer's profile x_t contains many sensitive information that shouldn't be published.
- The customer's purchase decision y_t is sometimes also sensitive information.
  - Whether the customer purchased certain medication
- **Concerns**: even if the pricing algorithm doesn't release x_t, y_t, could other people still **infer** these sensitive data, from the posted prices?

Concerns over privacy leakage

- **Example:** Privacy breach of purchase activity y_t.
  - Frequently, with active recent purchase activities the retailer spikes the price for larger profit margins.
  - A potential attack by a malicious agent: pretend as legitimate users before and after a customer of interest

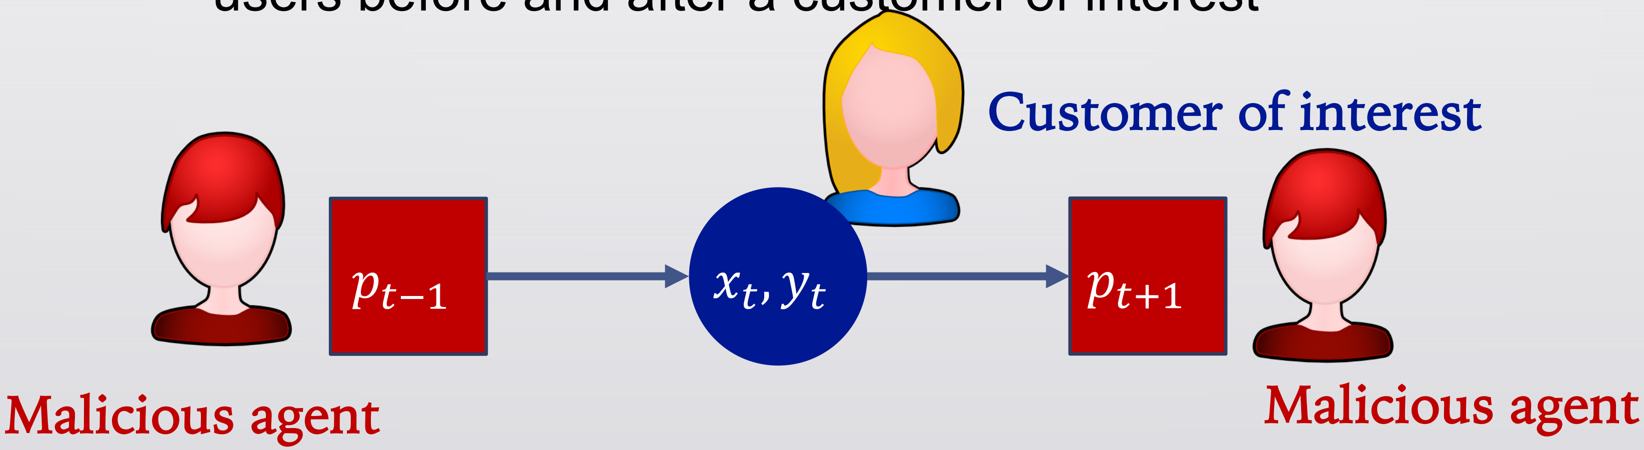

If the agents see a price increase $p_{t-1} < p_{t+1}$, it's more likely the person of interest made purchases.

Differentially private personalized pricing

- **Differential privacy**: a mathematically rigorous way to quantify privacy leakage. Dwork et al.'06

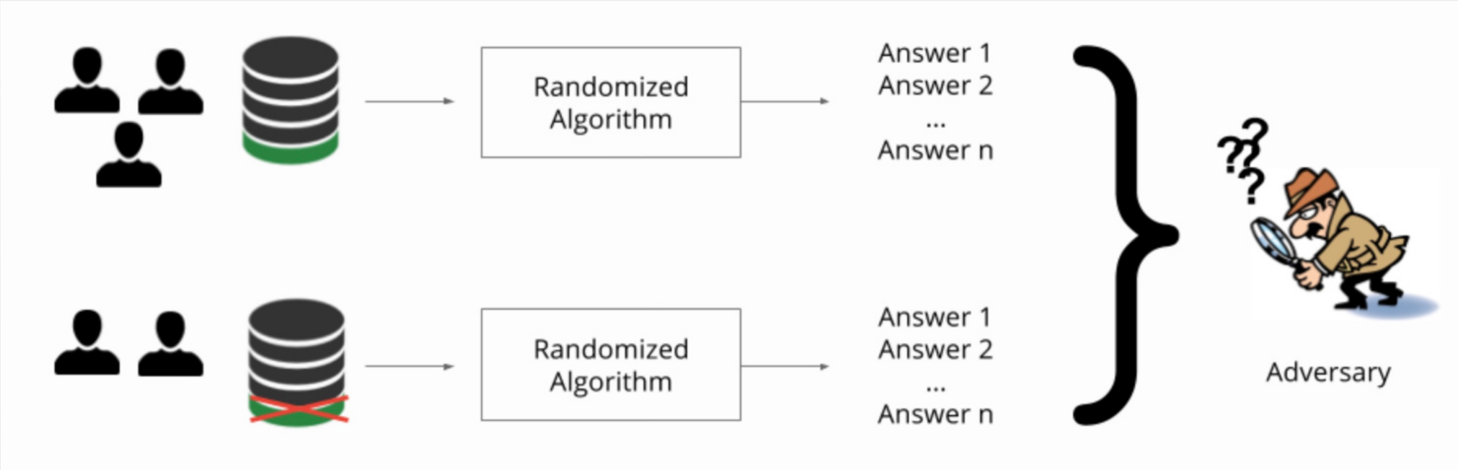

Differentially private personalized pricing

- **Differential privacy**: a mathematically rigorous way to quantify privacy leakage. Dwork et al.'06

$$\Pr[O|D] \leq e^{\varepsilon} \Pr[O|D'] + \delta$$

  - Interpretation: the probability of certain outcomes from the policy *O* does not change much, when a user's sensitive information changes ($D \rightarrow D'$).

Differentially private personalized pricing

- **Differential privacy**: a mathematically rigorous way to quantify privacy leakage. Dwork et al.'06

$$\Pr[O|D] \leq e^{\varepsilon} \Pr[O|D'] + \delta$$

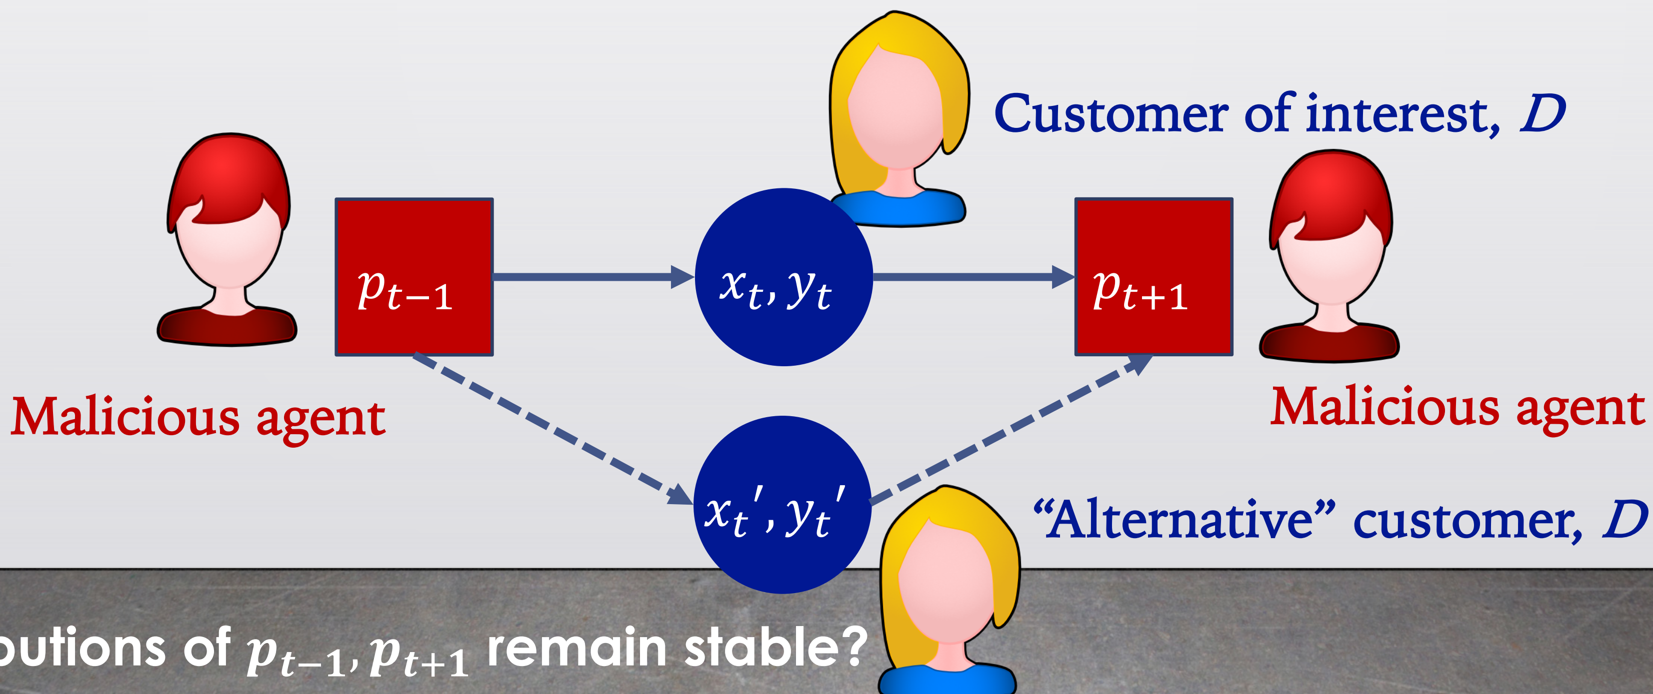

Distributions of p_{t-1}, p_{t+1} remain stable?

Differentially private personalized pricing

- **Differential privacy**: a mathematically rigorous way to quantify privacy leakage. Dwork et al.'06

$$\Pr[O|D] \leq e^{\varepsilon} \Pr[O|D'] + \delta$$

- **The (ε, δ)-differential privacy:** the smaller ε, δ are, the stronger privacy demands are requested by the firms/practitioners
- **Objective:** design differentially private algorithms without sacrificing too much profits.

Algorithm framework

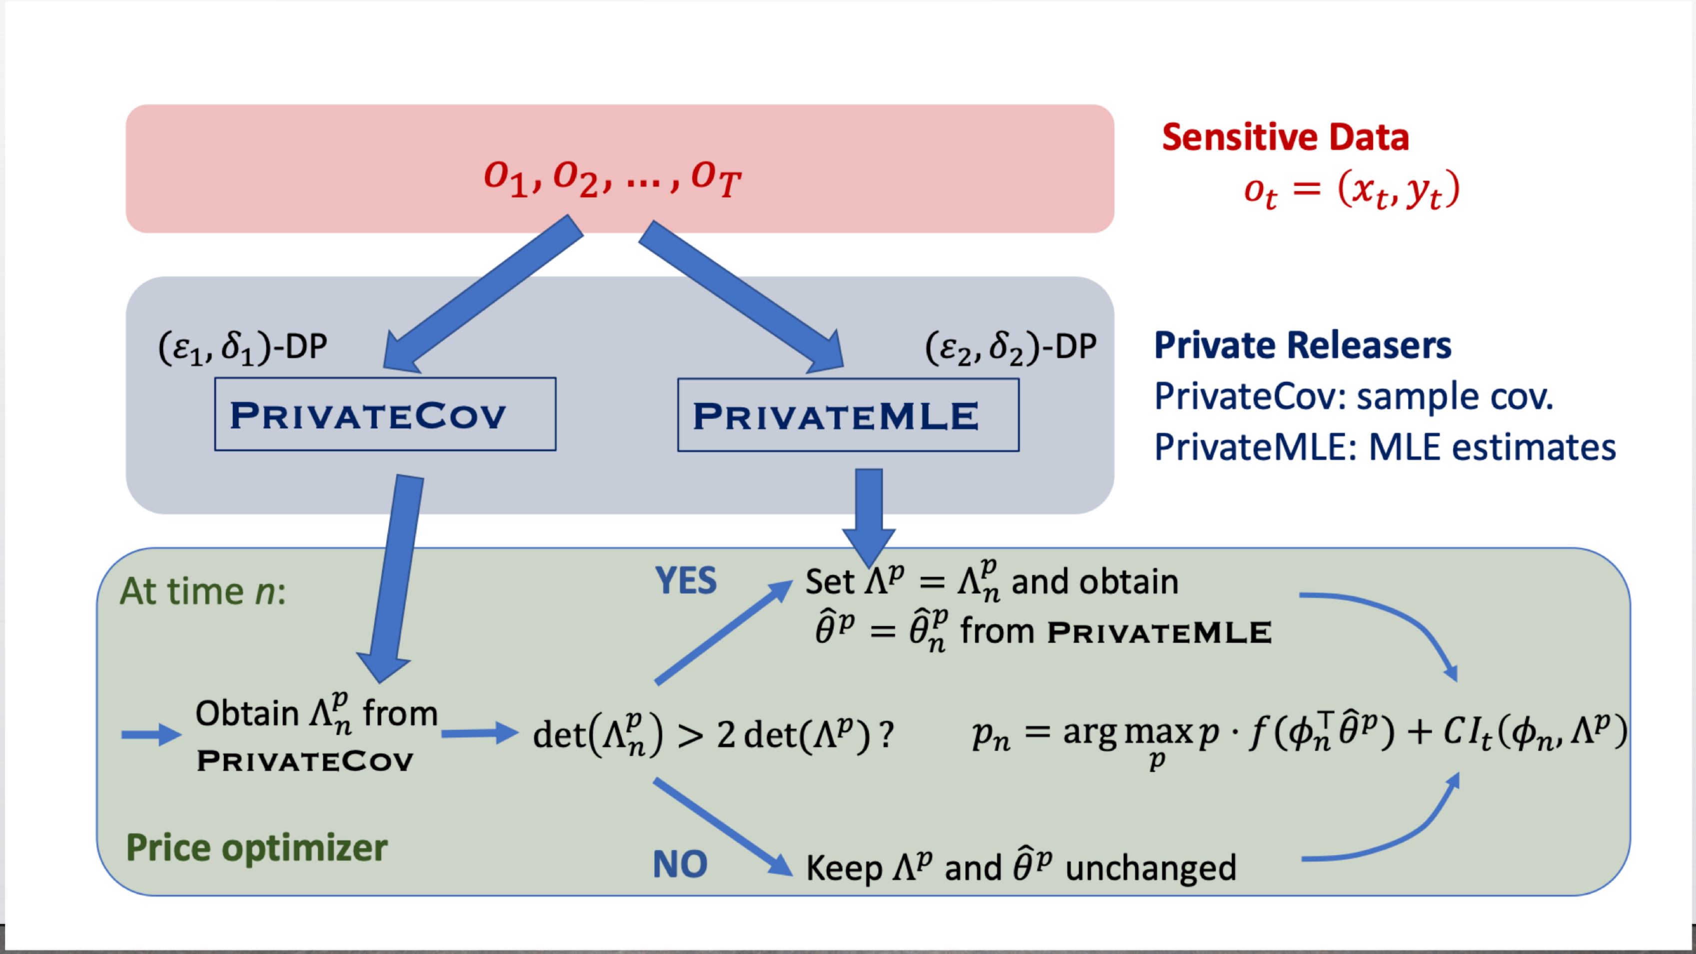

Algorithm details

- The **PrivateMLE** routine: produce privacy-aware model estimates using data prior to time *t*
- Key idea: "objective perturbation"

$$\max_{\theta} \sum_{\tau < t} \log P(y_\tau | x_\tau, p_\tau; \theta) - w_t^\top \theta$$

The calibrated noise
$w_t \sim N(0, \nu_{\varepsilon,\delta}^2)$

  - Privacy arguments in Kifer et al.'12, Chaudhuri et al.'11
  - Utility (error) analysis of $\hat{\theta}_t - \theta^*$ available by analyzing the first-order KKT condition of the perturbed objective.

Algorithm details

- The **PrivateCov** routine: give signals to invoke PrivateMLE for estimates, as few as possible.
- Approach: sequentially releasing differentially private sample covariance estimates.
    - "Tree-based" protocol in releasing consecutive sample covariances to facilitate frequent PrivateCov checks.
      Dwork et al.'10, 14, Chan et al.'11

Algorithm details

- At each time t, report **privatized** version $\tilde{\Lambda}_t$ of the sample covariance $\Lambda_t = \sum_{\tau<t} x_\tau x_\tau^T$ using **tree-based aggregation**

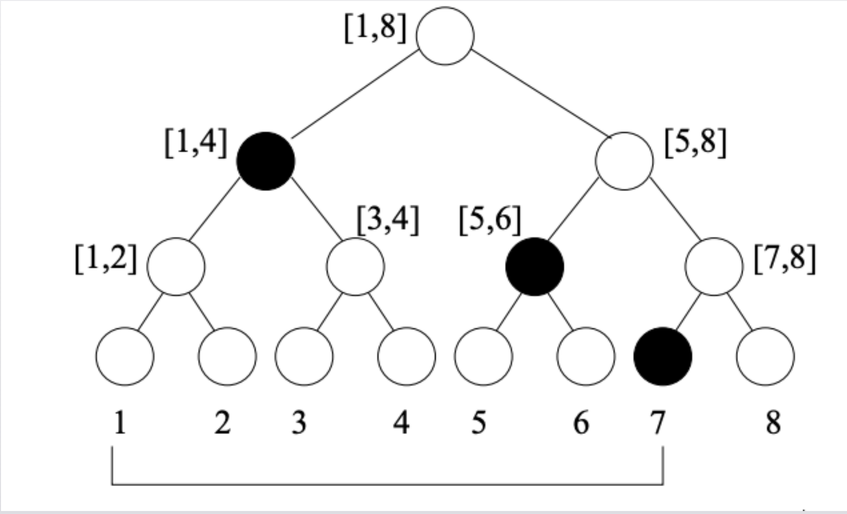

Example: $\sum_{\tau=1}^{7} x_\tau x_\tau^T$ is calculated

$$\sum_{\tau=1}^{4} x_\tau x_\tau^T + \text{noise}$$

$$\sum_{\tau=5}^{6} x_\tau x_\tau^T + \text{noise}$$

$$\sum_{\tau=7} x_\tau x_\tau^T + \text{noise}$$

Regret analysis

- **Regret measure**: performance of a (privacy-aware) policy π measured by $E^{\pi}[\sum_{t=1}^{T} r_t(p_t^*) - r_t(p_t)]$
  - p_t is the price offer by π and $r_t(p) = p \times E[y_t | p, x_t]$
  - p_t^* is the **optimal** price maximizing $r_t(.)$
- Without privacy concerns, the best algorithm has regret $\tilde{O}(d\sqrt{T})$. Filippi et al.'10, Abbasi-Yadkori et al.'11
- What does the regret look like for our proposed algorithm, subject to (ε, δ)-privacy constraints?

Regret analysis

- Without privacy concerns, the best algorithm has regret $\tilde{O}(d\sqrt{T})$. Filippi et al.'10, Abbasi-Yadkori et al.'11
- Subject to (ε, δ)-differential privacy constraints, our algorithm has regret $\tilde{O}(\varepsilon^{-1}\sqrt{d^3 T \ln^5(\delta^{-1})})$
  - Matches $\tilde{O}(\sqrt{T})$ regret, with slightly worse *d* dependency.
  - In practice *d* is usually small (few #. of covariates).
- If the contexts x_t are i.i.d. and non-degenerate, the regret can be improved to $\tilde{O}(d\sqrt{T} + \varepsilon^{-2} d^2 \ln^{10}(\delta^{-1}))$
  - Completely matches $\tilde{O}(d\sqrt{T})$ in the dominating term.

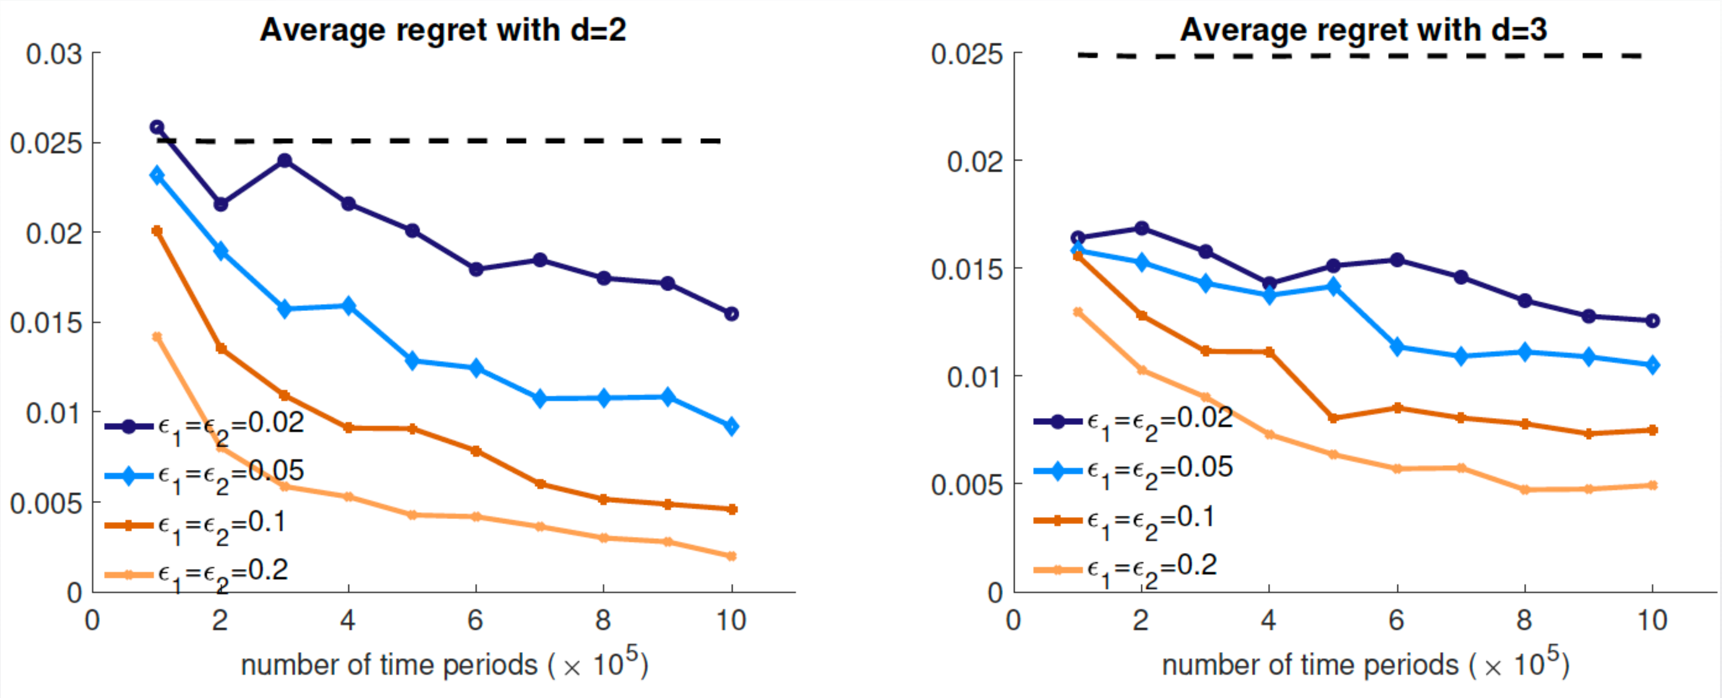

Numerical results

Average regret, with $\delta = 1/T^2$ and changing T

Additional results available in the paper

Future directions

- **Centralized (global) privacy** vs. **Local privacy**

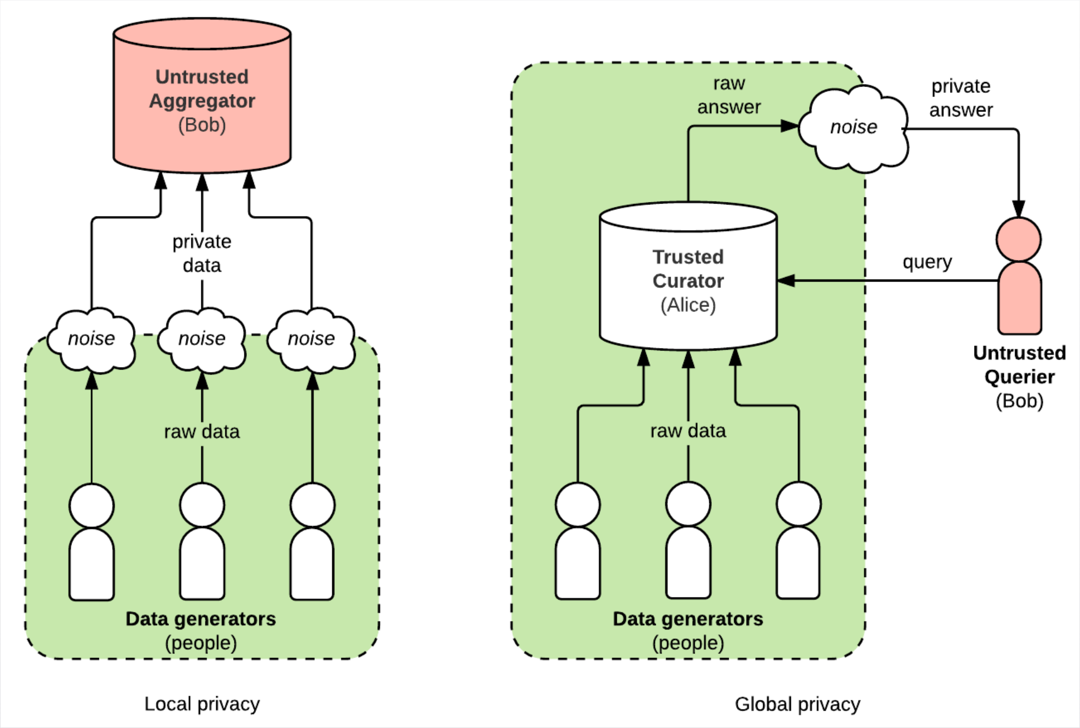

Key question:
Do ***I (as users)*** trust the *outside queriers (other users)*, or the *data curator (the company),* or *neither*?

Future directions

- **Centralized (global) privacy** vs. **Local privacy**
- For local privacy, the users do ***not*** trust the company and requires their profiles $\{x_t\}$ to be anonymized *first* before storing at the company's database.
- Idea: first-order methods with perturbed gradients

$$x_t, y_t \Longrightarrow g_t = \nabla_\theta \log P(y_t, x_t; \hat{\theta}_{t-1}) \Longrightarrow \tilde{g}_t = g_t + \xi$$

Future directions

- **Data privacy** vs. **Decision fairness**
- **Data privacy** requires the platform to avoid privacy leakage of users' data, through data storage or revenue decisions.
- **Decision fairness**, on the other hand, requires the firm to *not discriminate against* users inadvertently with their personalized data.

Decision fai

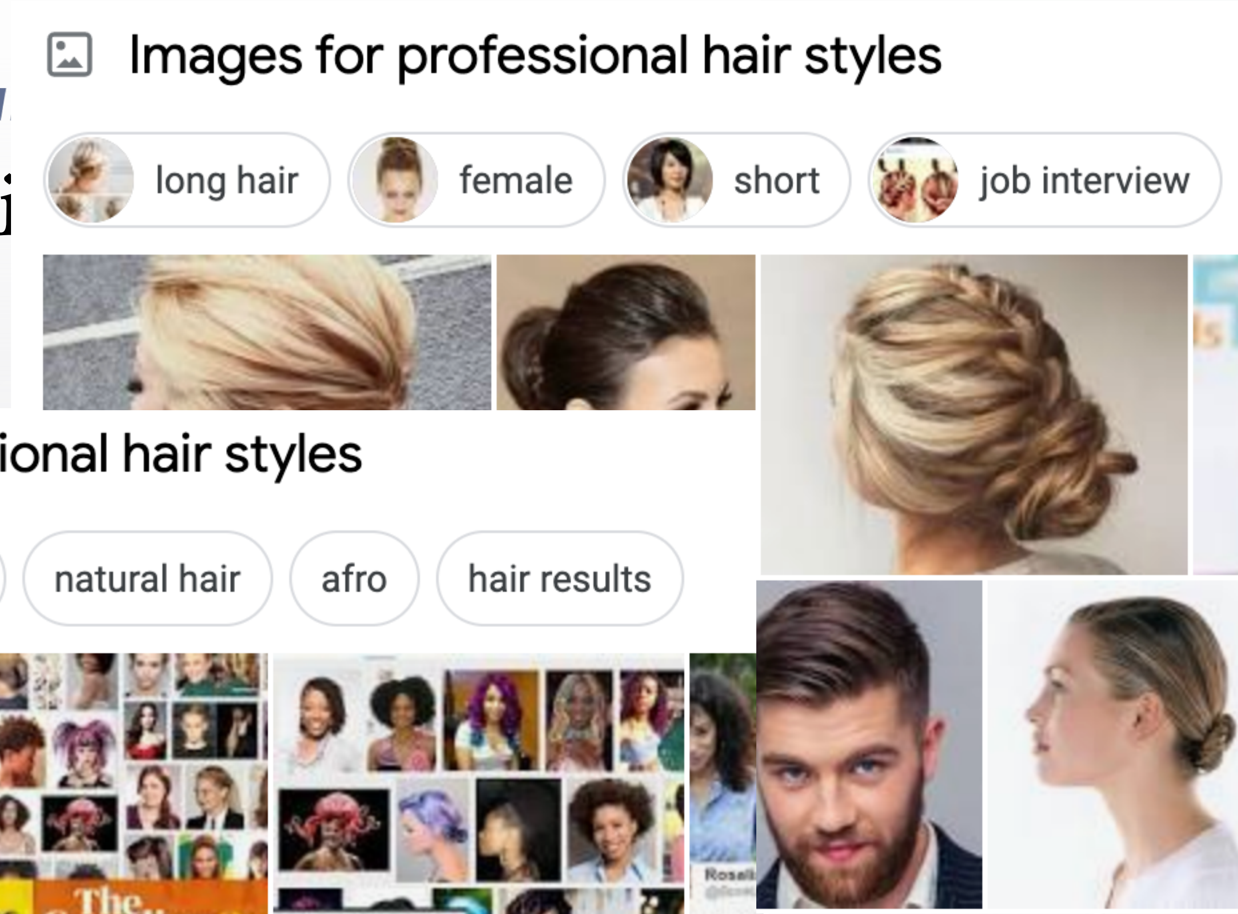

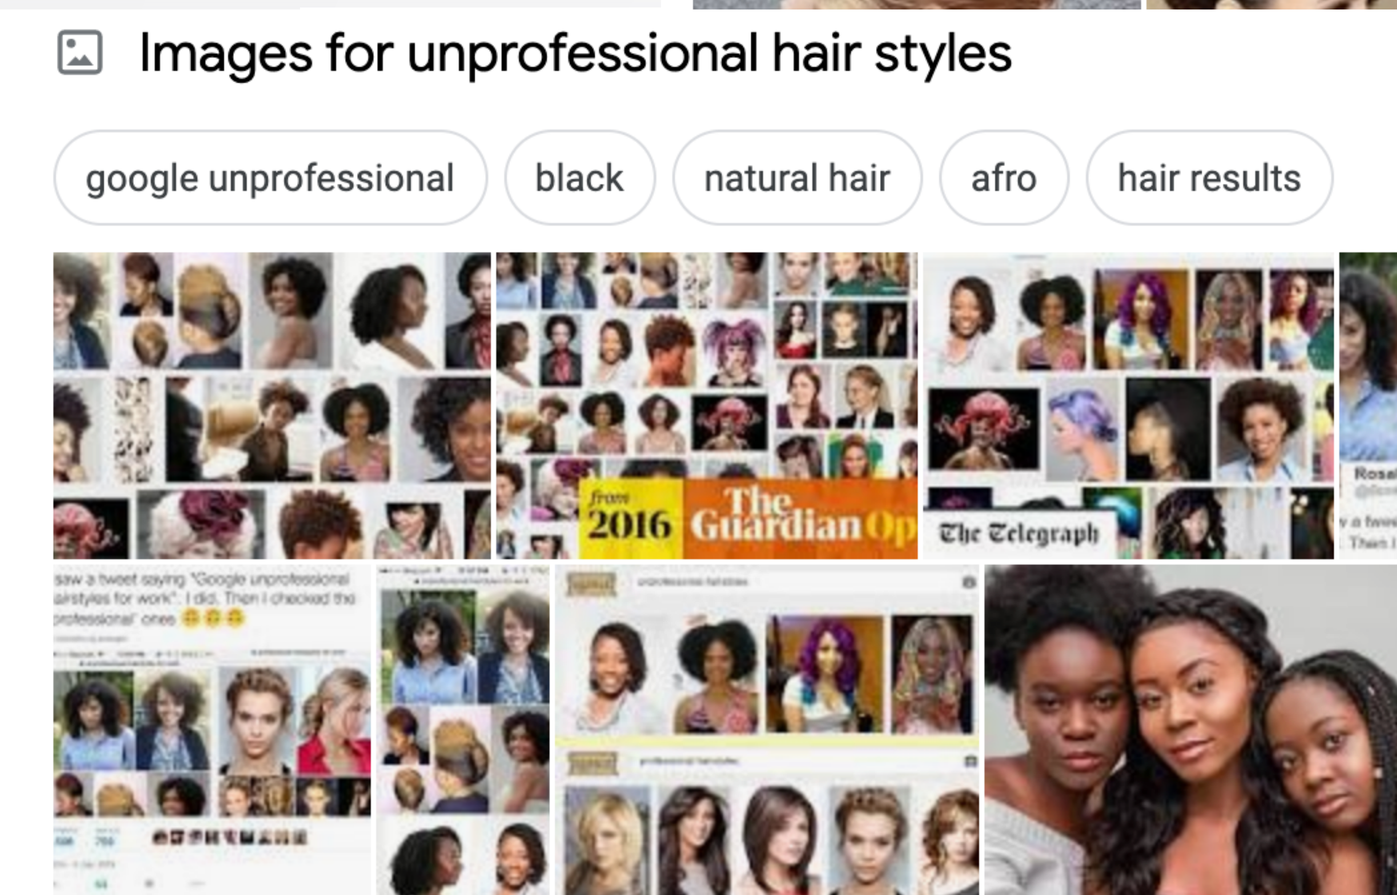

Decision fairness

- "**Individual fairness**" or "**Meritocratic fairness**".

Customers' profiles: x, x'

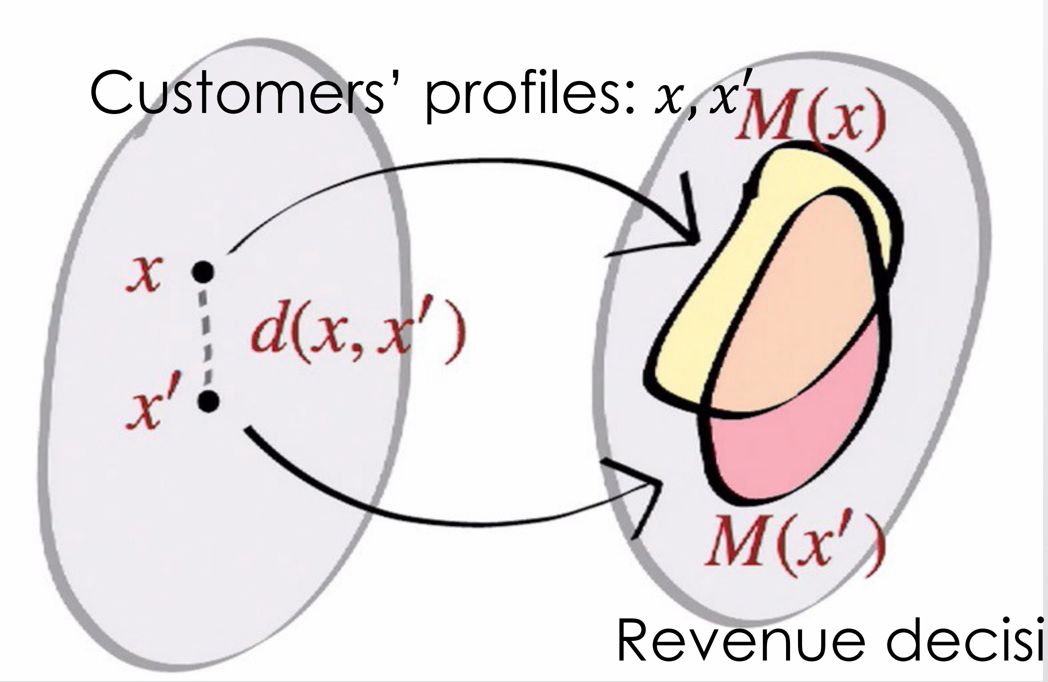

Revenue decisions: $M(x), M(x')$

Decision fairness

- "**Individual fairness**" or "**Meritocratic fairness**".
- "**Group fairness**": many times, fairness across user groups is more important/visible.

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Overall, Northpointe's assessment tool correctly predicts recidivism 61 percent of the time. But blacks are almost twice as likely as whites to be labeled a higher risk but not actually re-offend. It makes the opposite mistake among whites: They are much more likely than blacks to be labeled lower risk but go on to commit other crimes. (Source: ProPublica analysis of data from Broward County, Fla.)

Decision fairness

- Suppose users come from K **sensitive groups**, which are observable to the firm (racial, financial, demographical, etc.)
- The revenue decisions are required to solicit similar **average demands** across all **sensitive groups.**

Group fairness in personalized pricing

- Customer has profile $x \in X$, belongs to group k
  - Finite profile set $|X| < \infty$;
  - Personalized price decision $p_t: X \to [\underline{p}, \overline{p}]$;
- Revenue maximization with fairness constraints:

Distrb. Of profiles for ALL consumers

$\bar{G} = \frac{1}{K}\sum_{k=1}^{K} \pi_k G_k$

$$\max E_{x\sim\bar{G}}[p_t(x)D(p_t(x))]$$

Discrepancy between sensitive groups $\forall k \neq k'$

$$s.t. \quad \left| E_{x\sim G_k}\left[D\left(p_t(x)\right)\right] - E_{x\sim G_{k'}}\left[D\left(p_t(x)\right)\right] \right| \leq \varepsilon$$

Group fairness in personalized pricing

- Customer has profile $x \in X$, belongs to group *k*
  - Finite profile set $|X| < \infty$;
  - Personalized price decision $p_t: X \to [\underline{p}, \overline{p}]$;
- Revenue maximization with fairness constraints:

$$\max E_{x\sim\bar{G}}[p_t(x)D(p_t(x))]$$
$$s.t. \quad \left|E_{x\sim G_k}\left[D\left(p_t(x)\right)\right] - E_{x\sim G_{k'}}\left[D\left(p_t(x)\right)\right]\right| \le \varepsilon$$

- **Learning-While-Doing:**
  - Replace $D(.)$ with $\overline{D}_t$ (UCB) or $\widehat{D}_t \sim Q(\cdot\,|y_{<t})$ (TS)

Thank you!
Questions?

Personalized revenue management

- Example: **Booking.com** (hotel reservations)

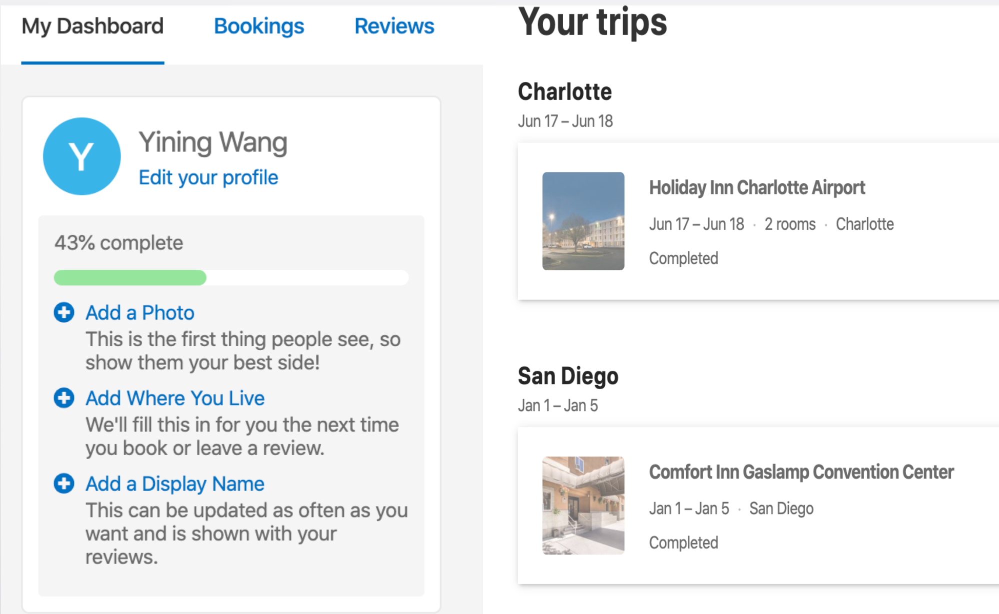

Can we use the user's
- *home address*, or
- *past booking history,* to
1. **promote** certain hotels (destinations closer to the user's home address), or
2. **price** stays at differently (set high prices for high-end or frequent business travelers)

Technical comment

- Why not perturb the user profiles x_t directly?
- Imagine a simple task of releasing the **sample average** of $x_1, \ldots, x_n$, $\bar{x} = (x_1 + \cdots + x_n)/n$
  - If I add noise first: $\tilde{x}_i = x_i + \xi_i$, and then report the average $\hat{x}^1 = (\tilde{x}_1 + \cdots + \tilde{x}_n)/n$, we have that
    $$\left|\hat{x}^1 - \bar{x}\right| = \widetilde{O}(1/\varepsilon\sqrt{n})$$
  - If I compute $\bar{x} = (x_1 + \cdots + x_n)/n$ first and then report $\hat{x}^2 = \bar{x} + \bar{\xi}$, then we have that
    $$\left|\hat{x}^2 - \bar{x}\right| = \widetilde{O}(1/\varepsilon n)$$

Machine learning for revenue management

- Machine learning and big-data analytics
  - Supervised, unsupervised and semi-supervised learning
  - Active learning, online learning, design of experiments
  - Reinforcement learning and multi-agent learning
  - Deep learning and learning representations
  - Resource-constrained learning (communications, computations, privacy, fairness, etc.)
- Many of the above techniques can be adapted to solve challenges in data-driven revenue management!

Machine learning for revenue management

- Question 1. How to systematically incorporate personalized data to maximize revenue/profit performances as much as possible?
- *Applicable ML techniques:* **online and bandit learning**

Machine learning for revenue management

- Question 2. When using personalized data to make decisions, how to avoid inadvertently leaking private data of the users?
- *Applicable ML techniques:* **differential privacy**

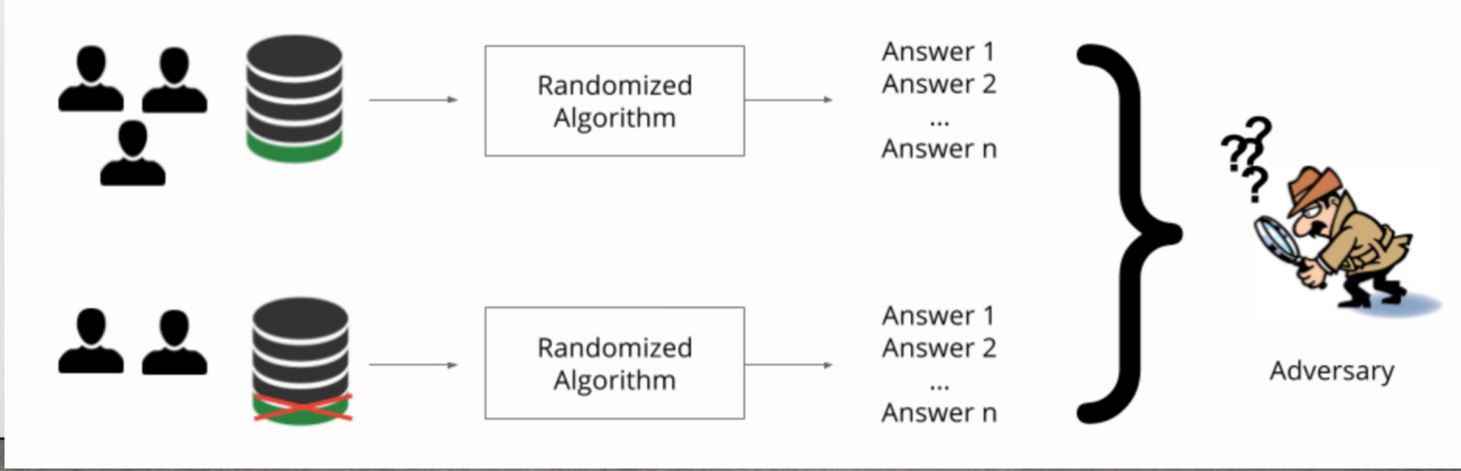

Concerns over privacy leakage

- **Example:** Privacy breach of customer profile x_t.
  - Most pricing systems post similar prices to consumers with similar profiles in the future (i,e., similar x_t)
  - A potential attack by a malicious agent: pretend as consumers before and after a customer of interest.

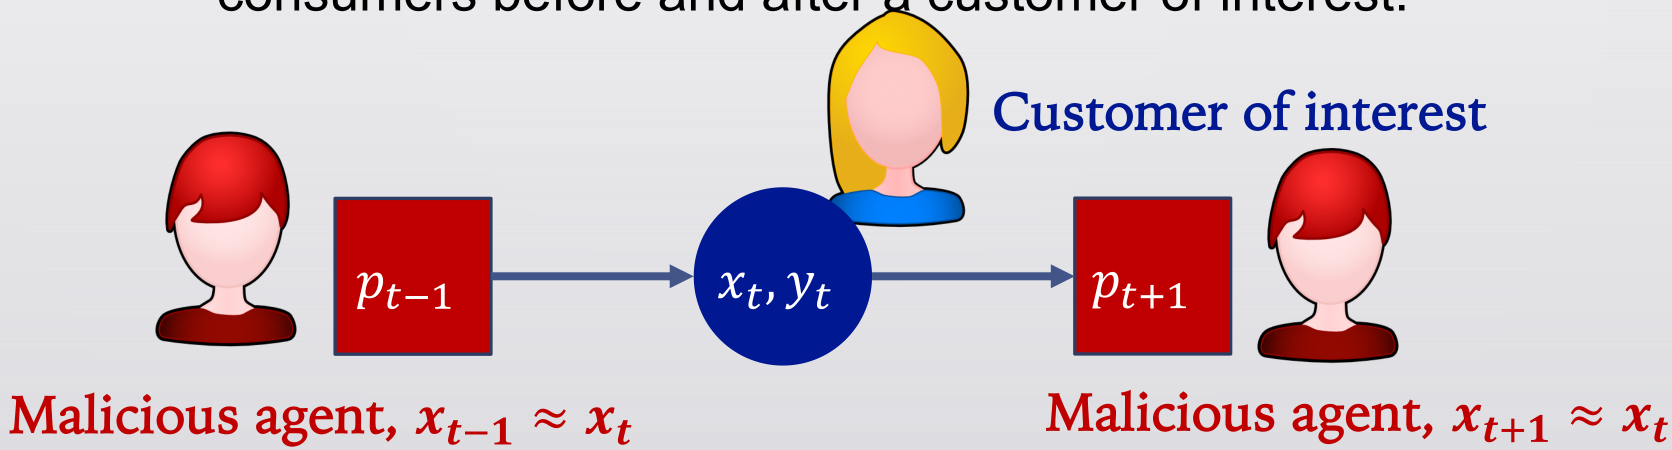

If the agents see similar prices $p_{t-1} \approx p_{t+1}$, it is more likely that the customer of interest has similar profiles.

Technical challenges

- **<u>Challenge 1.</u>** General demand models do not admit *sufficient statistics* like the linear regression.
  - Cannot directly apply Shariff & Sheffet'18 which simply perturbs the sample covariance and average demand.
  - **Solution: privacy-aware maximum likelihood estimation with a concave/convex formulation**
  - Privacy analysis comes from Kifer et al.'12, Chaudhuri et al.'11, but utility/error analysis is re-done and novel.

Technical challenges

- **<u>Challenge 2.</u>** the "curse of composition": releasing too many statistics in DP formulation.
  - Cannot update demand model **after every customer**. That leaks top much privacy through composition.
  - **Solution: infrequent private model updates, with private protocols signaling updates as well.**
  - Ideas drawn from non-private low-switching policies Abbasi-Yadkori et al.'11 and private protocols for sample covariance and sequence releases. Dwork et al.'10, 14, Chan et al.'11

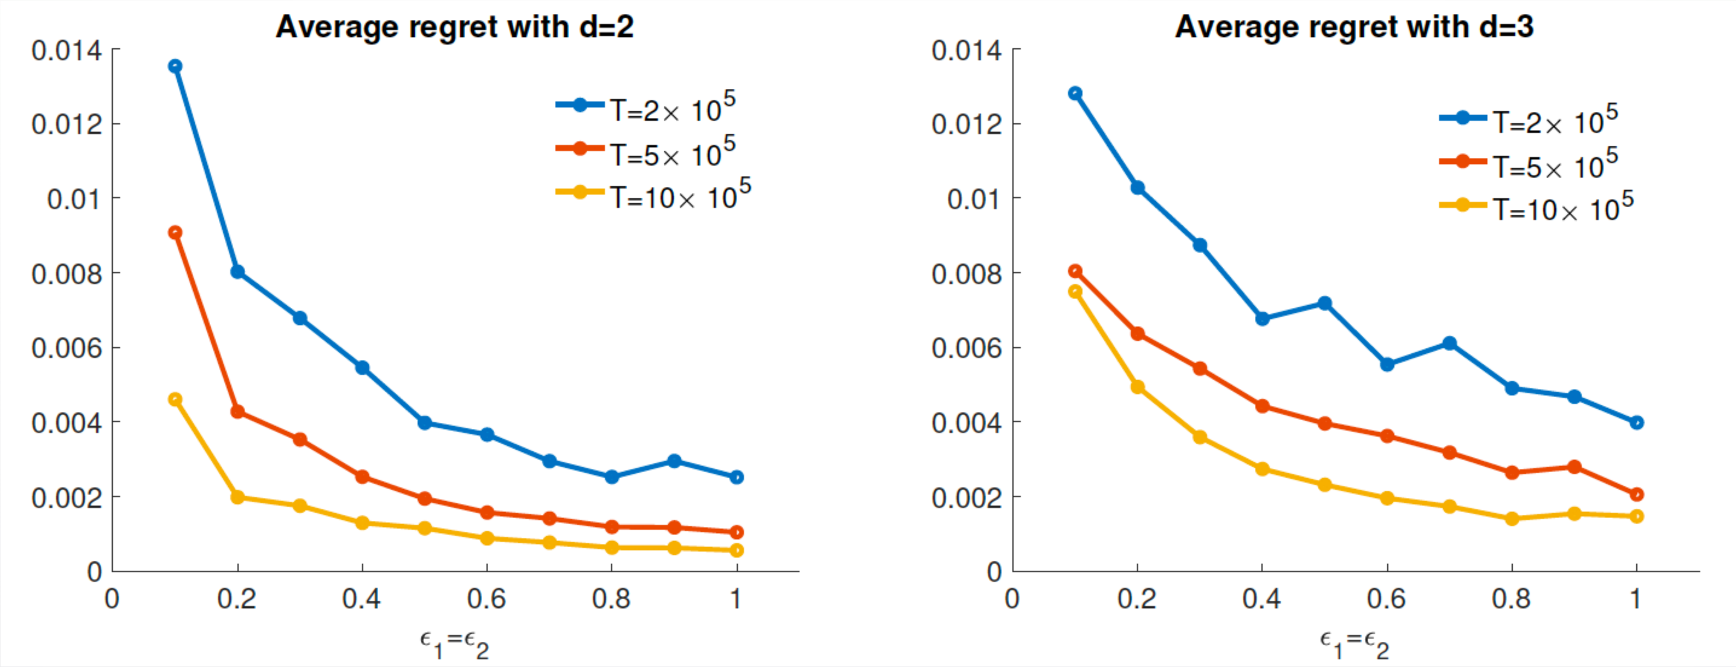

Numerical results

Average regret, with $\delta = 1/T^2$ and changing ε