Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty

Ellen Vitercik
Stanford University

Joint work with Wenshuo Guo, Nika Haghtalab, and Kirthevasan Kandasamy

Appeared in the Conference on Economics and Computation (EC) 2023

Learning from reviews

Online shopping accounts for 22% of global retail sales

Customers make far more informed decisions than ever before

Gain insights from ***hundreds of reviews*** *before making purchases*

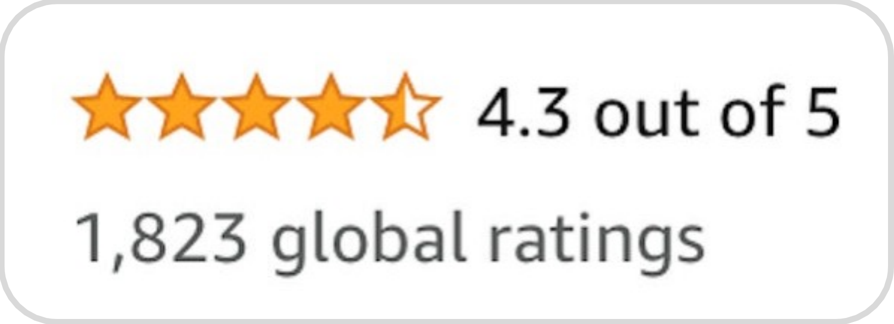

Learning from reviews

Often use reviews by buyers who share their "**type**," e.g.:

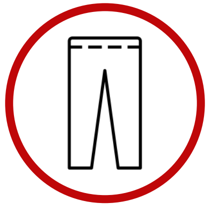
Body type for clothes


Skin type for skincare products

Use these reviews to **estimate** how much they will **value** items

Quantities they may be uncertain of before purchasing

Filtering reviews by type

Groomer's Best Small Combo Brush for Cats and Small Dogs

Visit the Hartz Store

7,579 ratings | 8 answered questions

Amazon's Choice for "hartz groomer's best combo dog brush"

Looking for specific info?

long-haired

Customer Reviews

Did not collect any hair off of my long haired cat

By Nazli Zeynep Turken on August 30, 2021

This brush/comb combo did not really collect any hair from my long-haired cat without a lot of pressure. The fur shedder work better.

Filtering reviews by type

Editor Mid Rise Bootcut Pant

★★★★☆ 4.2 (352) Write a review

Search topics and reviews

Rating | Body Type | Incentivized Review | Age

ATHLETIC × Clear All

1 – 7 of 7 Reviews

Disapprove

Nj

Review 1

Votes 12

★☆☆☆☆

Poor

4 months ago

Pockets flare out , not flattering I want my columnist pants back with the slit top pocket!

Filtering reviews by type

Paula's Choice
Skin Perfecting 2% BHA Liquid Exfoliant
★★★★☆ 1.1K | Ask a question | ♥ 254.6K

Sort | Rating | Verified Purchases | Non-Incentivized Reviews Only | Skin Type | Skin Concerns | Age Range

Oily ✕ Clear all

Viewing 1-6 of 189 reviews

★★★★★
6 d ago
✓ Recommended

LITERALLY NEED

I didn't notice a major difference until I ran out of it, then my forehead started to break out again and my skin just looked dull. It's the only thing that gets rid of pimples that are painful and under the skin.

Helpful? △ (3) | ▽ (1)

★★★★★

A MUST IN MY WEEKLY ROUTINE

Key challenge when pricing

For **rare** types of customers,

- May find only a few reviews from similar customers
- Due to uncertainty, may only be willing to buy at relatively **low prices**

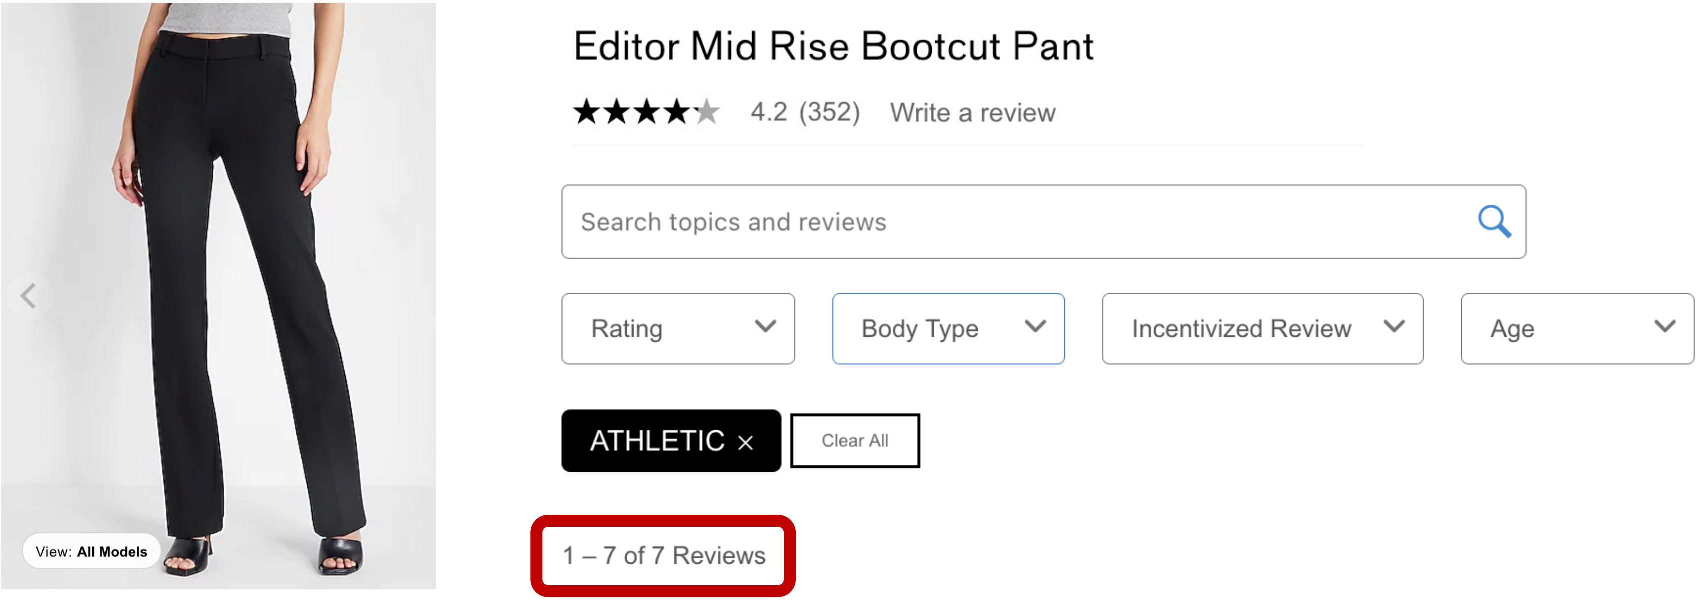

Key challenge when pricing

Customer's **purchase decision** isn't just a function of the price

- Depends on how **certain** the customer is about her valuation
- In turn, depends on the **earlier sales and reviews**

Leads to a **tension** between:

- Setting **revenue-optimal prices**, and
- Ensuring that buyers have **enough reviews to estimate** their values

Results overview

Introduce a model that simultaneously captures:

- The seller's pricing problem
- The buyers' learning problem
- The modus through which the buyers learn: reviews

We study how a seller can learn to set high-revenue prices

- Provide a no-regret learning algorithm
- Matching regret lower bounds

Outline

1. Introduction
2. **Mechanism design background**
3. Model
4. Main results
5. Conclusions and future directions

Mechanism design background

- Single item, single buyer
- Distribution $\mathcal{D}$ over buyer's value for item

  Seller knows $\mathcal{D}$

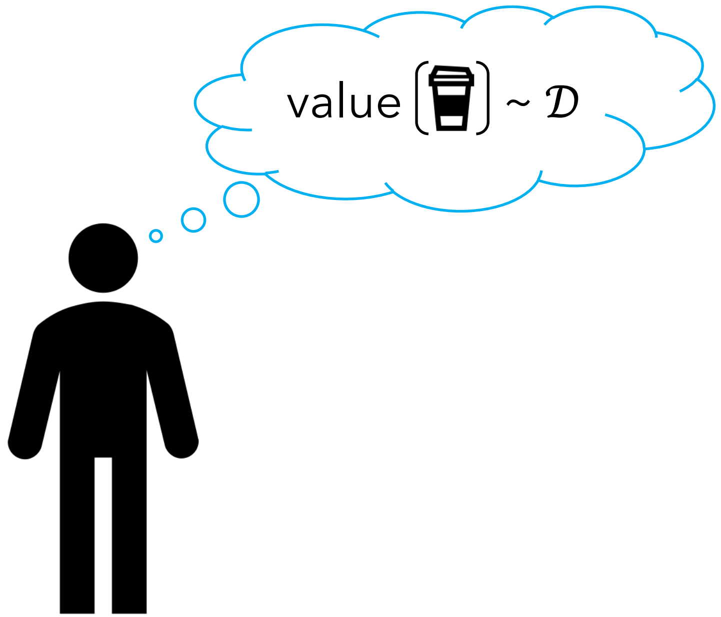

Mechanism design background

- Single item, single buyer
- Distribution $\mathcal{D}$ over buyer's value for item

  Seller knows $\mathcal{D}$
- Interaction between buyer and seller:
  1. Seller uses $\mathcal{D}$ to select choose **price p**
  2. Buyer draws value $v \sim \mathcal{D}$ and **purchases item if $v \geq p$**
- **Revenue-maximizing** price: $\text{argmax}\{p \cdot \mathbb{P}_{v\sim\mathcal{D}}[v \geq p]\}$
- Assumes seller knows $\mathcal{D}$ and buyer knows v

  We relax both these assumptions

Outline

1. Introduction
2. Mechanism design background
3. **Model**
4. Connections to prior research
5. Main results
6. Conclusions and future directions

Model

- Item sold repeatedly to sequence of buyers over T rounds
  - Buyers are distinct
- Each buyer has a **type** $i \in [d]$
  - E.g., height, weight, skin type, …
  - There's an unknown distribution $\mathcal{P}$ over types $[d]$

Model

- Buyer of type i's **value** for item drawn from distribution $\mathcal{D}_i$
  - support($\mathcal{D}_i$) $\subseteq [0,1]$
  - Has mean θ_i
- θ_i: buyer's *ex-ante* value
  - What buyer would **expect** their value to be before buying the item
- $v \sim \mathcal{D}_i$: buyer's *ex-post* value
  - What their **value actually is** after buying the item
- Seller knows $\theta_1, \ldots, \theta_d$ but not the distributions $\mathcal{P}, \mathcal{D}_1, \ldots, \mathcal{D}_d$

Distribution over types

Online learning model

At each timestep $t = 1, \ldots, T$:

1. Reviews σ_{t-1} describe past buyers' types & *ex-post* values

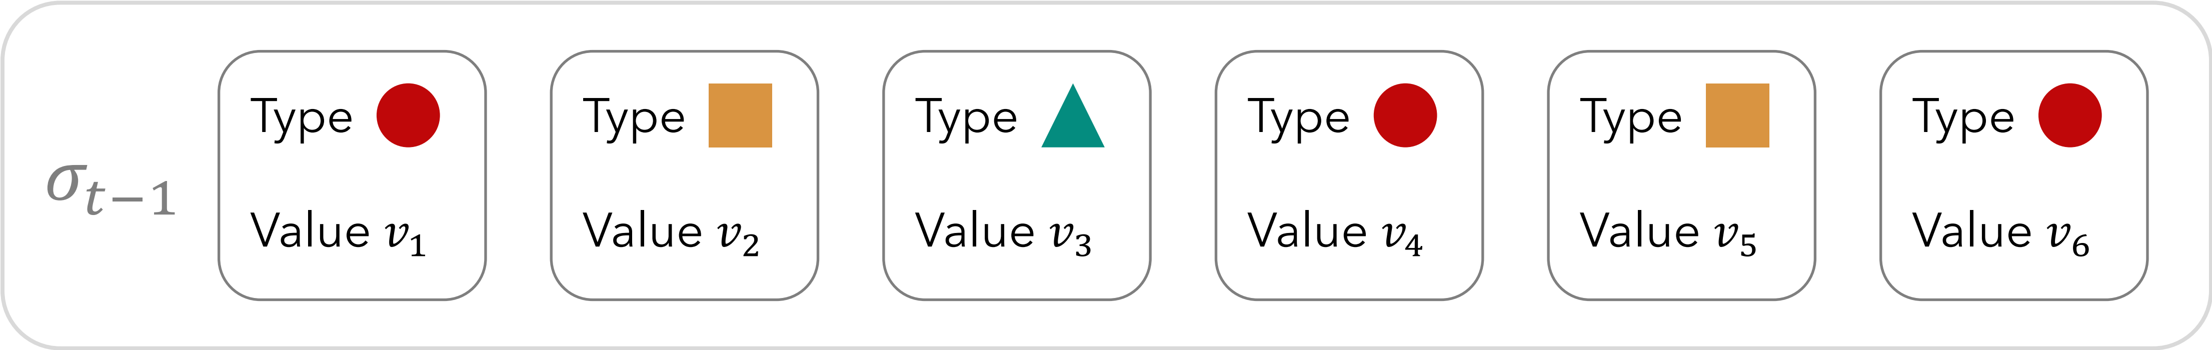

Online learning model

At each timestep $t = 1, \ldots, T$:

1. Reviews σ_{t-1} describe past buyers' types & *ex-post* values
2. Seller sets a **price** $p_t \in [0,1]$
3. Buyer arrives with type $i_t \sim \mathcal{P}$
   i. They observe the **past reviews** of buyers with type i_t
   ii. They decide **whether to purchase** the item

   Seller **doesn't know the type** i_t when they choose p_t
4. If the buyer purchases the item, they pay p_t
   i. If they buy, they **leave a review** of (i_t, v_t) with $v_t \sim \mathcal{D}_{i_t}$

Buyers' purchasing model

- Buyer's purchase decision defined by **threshold** $\tau_t(\sigma_{t-1}, i_t)$
  - $\tau_t(\sigma_{t-1}, i_t)$ represents the buyer's estimation of θ_{i_t} based on reviews
- Agent purchases the item if $p_t \leq \tau_t(\sigma_{t-1}, i_t)$
- **Conservative** agent would choose $\tau_t(\sigma_{t-1}, i_t)$ to be **low**
  - Extreme example would set $\tau_t(\sigma_{t-1}, i_t) = 0$

Agent only buys item if offered for free!

  - Optimizing revenue with such a conservative agent is **hopeless**

Bounded pessimism assumption

$\tau_t(\sigma_{t-1}, i_t)$ is at least a lower confidence bound LB_t that equals:

- The **average** of the reviews left by buyers with type i_t,
- Minus an **uncertainty term** that depends on # of reviews

Definition: *η-pessimistic agent*

- Φ_t = reviews left by previous buyers with type i_t
- $\mathrm{LB}_t = \frac{1}{|\Phi_t|}\sum_{v \in \Phi_t} v - \sqrt{\frac{1}{2|\Phi_t|}\ln\frac{t}{\eta}}$
- Agent is *η-pessimistic* if $\tau_t(\sigma_{t-1}, i_t) \geq \mathrm{LB}_t$
  - Will definitely buy if $\mathrm{LB}_t \geq p_t$

Bounded pessimism assumption

Definition: *η-pessimistic agent*

- Φ_t = reviews left by previous buyers with type i_t
- $\mathrm{LB}_t = \frac{1}{|\Phi_t|} \sum_{v \in \Phi_t} v - \sqrt{\frac{1}{2|\Phi_t|} \ln \frac{t}{\eta}}$
- Agent is *η-pessimistic* if $\tau_t(\sigma_{t-1}, i_t) \geq \mathrm{LB}_t$
  - Will definitely buy if $\mathrm{LB}_t \geq p_t$

With probability $1 - \eta$, for all t, $\theta_{i_t} \geq \mathrm{LB}_t$

- If $\mathrm{LB}_t \geq p_t$, agent's **expected utility** $\theta_{i_t} - p_t$ is likely positive
- Will buy if have **good reason to believe** their expected utility is ≥ 0

Key challenge

- Seller doesn't know i_t
  ⇒ Doesn't know # of reviews buyer will use to construct value estimate
- If i_t is a **rare type**, then LB_t will be low
  Would have to set a **low price** to ensure a purchase and a review
- If rare type's value is high, may be worth it to offer a low price
  Seller could "**win over**" these rare but high-value customers

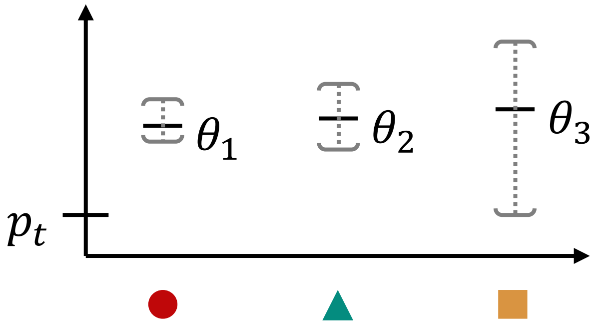

Key challenge

- Seller doesn't know i_t
  ⇒ Doesn't know # of reviews buyer will use to construct value estimate
- If i_t is a **rare type**, then LB_t will be low
  Would have to set a **low price** to ensure a purchase and a review
- If rare type's value is high, may be worth it to offer a low price
  Seller could "**win over**" these rare but high-value customers
- Seller has to decide **who to win over** without knowing i_t or $\mathcal{P}$
  May offer low price to a buyer who'd be willing to buy at a higher price

Outline

1. Introduction
2. Mechanism design background
3. Model
4. **Connections to prior research**
5. Main results
6. Conclusions and future directions

Prior research

Rich literature on **social learning from reviews**

[Chamley, '04; Bose et al., RAND J. Econ'06; Crapis et al., Manage Sci '17; Besbes, Scarsini, OR'18; Ifrach et al., OR'19; Kakhbod et al. SSRN'21; Boursier et al., ALT'22; Acemoglu et al. Econometrica'22]

Bayesian buyers:

Calculate item's posterior quality given the past reviews

[e.g., Ifrach et al., OR'19; Boursier et al., ALT'22; Acemoglu et al., Econometrica'22]

May be challenging to compute Bayesian updates

Several papers relax this assumption

[e.g., Crapis et al., Manage Sci '17; Besbes, Scarsini, OR'18]

Prior research

- E.g., Besbes and Scarsini [OR'18] study
  1. Fully **Baysian** buyers
  2. Buyers who can only observe the **average** of the past reviews
  - Conditions under which buyers can recover product's true quality
- Our model is situated between (1) and (2)
  - Purchase decisions depend on:
    - **Average** of the past reviews
    - **Number** of those reviews
- Besbes and Scarsini [OR'18] analyze **risk-neutral** buyers
  - We study **risk-averse** buyers:
    May not purchase even if the price is below the average reviews

Outline

1. Introduction
2. Mechanism design background
3. Model
4. Connections to prior research
5. **Main results**
6. Conclusions and future directions

Regret

Regret is the difference between:

1. The **algorithm**'s total expected revenue, and
2. The expected revenue of the **optimal fixed price** if:
   agents buy if their *ex-ante* value is larger than the price

Under 2, the buyers and seller know more than under 1:

- Seller knows all distributions $\mathcal{P}, \mathcal{D}_1, \ldots, \mathcal{D}_d$
  Knows which customers to target to maximize revenue
- Buyers know their *ex-ante* values $\theta_1, \ldots, \theta_d$
  Seller can extract more revenue than he could from uncertain buyers

Regret

Regret is the difference between:

1. The **algorithm**'s total expected revenue, and
2. The expected revenue of the **optimal fixed price** if:
   agents buy if their *ex-ante* value is larger than the price

$b_t = 1$ if buyer buys at round t at price p_t; $b_t = 0$ otherwise

1. $= \sum_{t=1}^{T} b_t \cdot p_t$

$p^* = \text{argmax}\{p \cdot \mathbb{P}_{i \sim \mathcal{P}}[\theta_i \geq p]\}$

2. $= T \cdot p^* \cdot \mathbb{P}_{i \sim \mathcal{P}}[\theta_i \geq p^*]$

Regret

Regret is the difference between:

1. The **algorithm**'s total expected revenue, and
2. The expected revenue of the **optimal fixed price** if:
   agents buy if their *ex-ante* value is larger than the price

In other words,

$$\mathbb{E}[R_T] = T \cdot p^* \cdot \mathbb{P}_{i \sim \mathcal{P}}[\theta_i \geq p^*] - \sum_{t=1}^{T} p_t \cdot b_t$$

Main result

$q_{\min}$= minimum probability of any type $\left(\min_{i\in[d]} \mathbb{P}_{j\sim\mathcal{D}}[j=i]\right)$

Theorem: We provide an algorithm such that

- If $q_{\min}$ not tiny $\left(q_{\min} > 2d^{-2/3}T^{-1/3}\right)$ then

$$\mathbb{E}[R_T] = O\left(\sqrt{\frac{T}{q_{\min}}} + T^{1/3}d^{2/3}\right)$$

- Otherwise,

$$\mathbb{E}[R_T] = O\left(T^{2/3}d^{1/3} + T^{1/3}d^{2/3}\right)$$

Also provide **lower bounds** that match up to lower order terms

Prior research

If seller only observes **purchase decisions** and not reviews:

- $\widetilde{\Theta}(T^{2/3})$ regret bound [Kleinberg & Leighton, '03]
- Can be improved to $\widetilde{\Theta}(T^{1/2})$ under distributional assumptions

If seller observes **purchase decisions** *and* **reviews**:

- Algorithm with $\tilde{O}(\sqrt{T})$ regret [Zhao & Chen, '20]
- Assumes buyers know their own values

Outline

1. Introduction
2. Mechanism design background
3. Model
4. Connections to prior research
5. Main results
    a. Regret bound overview
    b. **Algorithm**
    c. Regret bound proof sketch
    d. Lower bound
6. Conclusions and future directions

Algorithm overview

Algorithm maintains set S_t of buyer types which it estimates:

1. Have a sufficiently **high value**, and
2. Are **not exceedingly rare**

Intuitively, S_t is the set of buyers that the algorithm is targeting

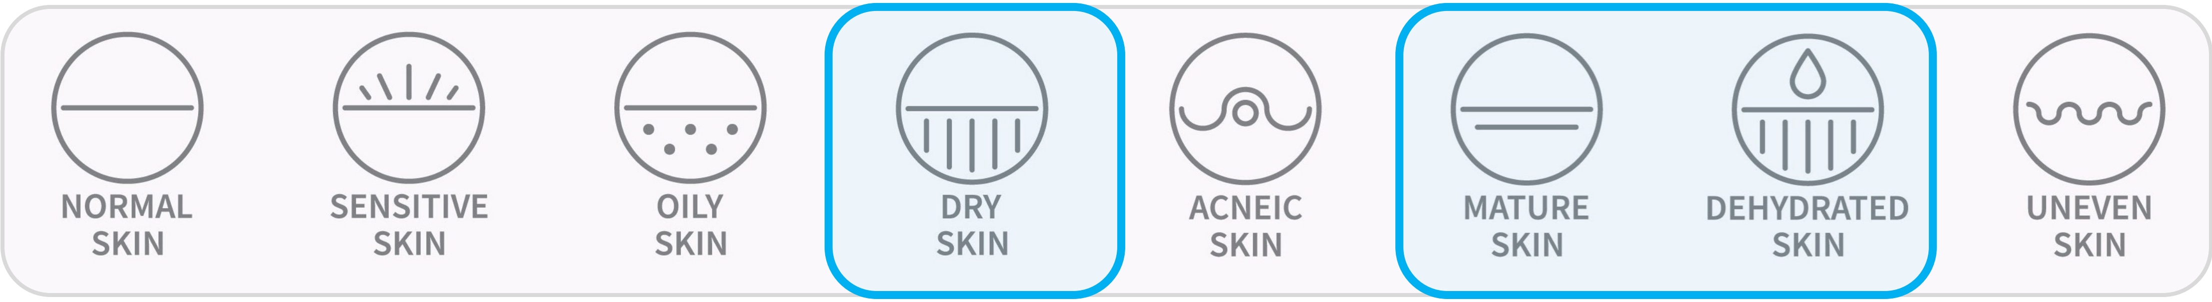

Algorithm overview

Algorithm has two phases

1st phase:

- Algorithm offers the item for **free**
- Observes **i.i.d. samples** from the type distribution
- Sets S_t to be the set of types that appeared **often enough**

2nd phase:

- Sets price low enough so that buyers in S_t **always buy**
- **Eliminates types** from S_t that contribute too little revenue

Algorithm: 1st phase

Offers item for free for $\widetilde{\Theta}(T^{1/3}d^{2/3})$ rounds

Q = set of buyer types that appeared frequently

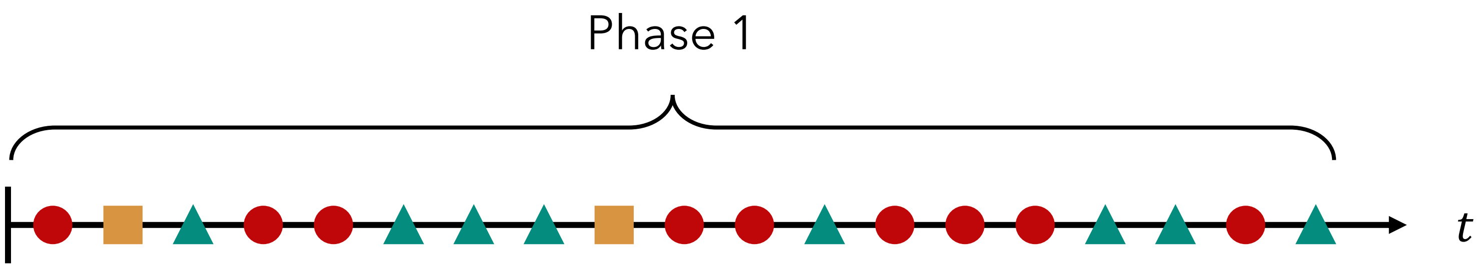

$Q = \{$●, ▲$\}$

Algorithm: 2nd phase overview

Algorithm will ignore types not in Q

- These customers are rare
- Will have more uncertainty about their value (low LB_t)
- Seller will have to set a low price to target these customers

⇒ Not worthwhile to target these customers

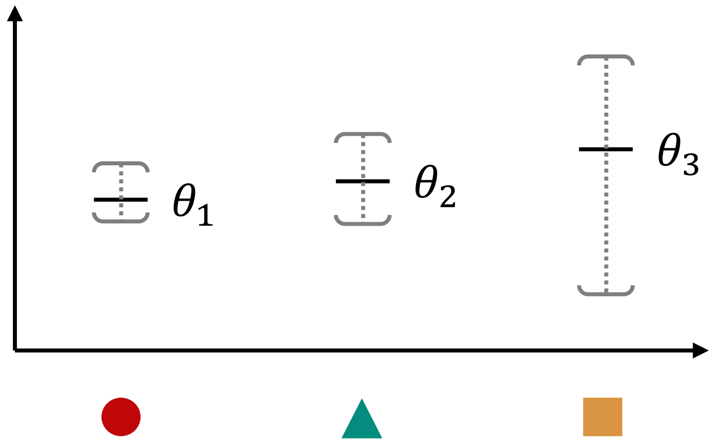

$Q = \{$●$,$▲$\}$

Algorithm: 2nd phase overview

In 2nd phase, only aims to maximize revenue WRT buyers in Q

- $\mathrm{rev}(p, Q) = p \cdot \mathbb{P}_{i \sim \mathcal{D}}[\theta_i \geq p \text{ and } i \in Q]$
- $p^*(Q) = \mathrm{argmax}_{p \in [0,1]} \mathrm{rev}(p, Q)$

Observation: $p^*(Q) = \theta_{i_Q}$ for some $i_Q \in Q$

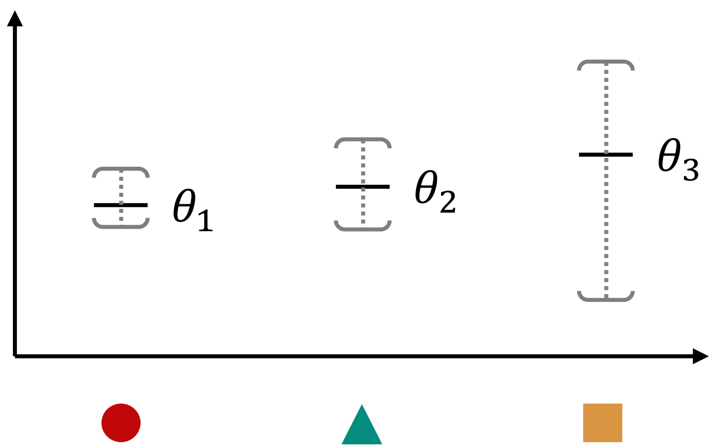

$Q = \{$●, ▲$\}$

Algorithm: 2nd phase price selection

Maintains set S_t of "active types" such that i_Q is likely in S_t

S_t initially set to Q

Sets p_t low enough to ensure if $i_t \in S_t$, then the buyer will buy

Ensures a review if $i_t \in S_t$

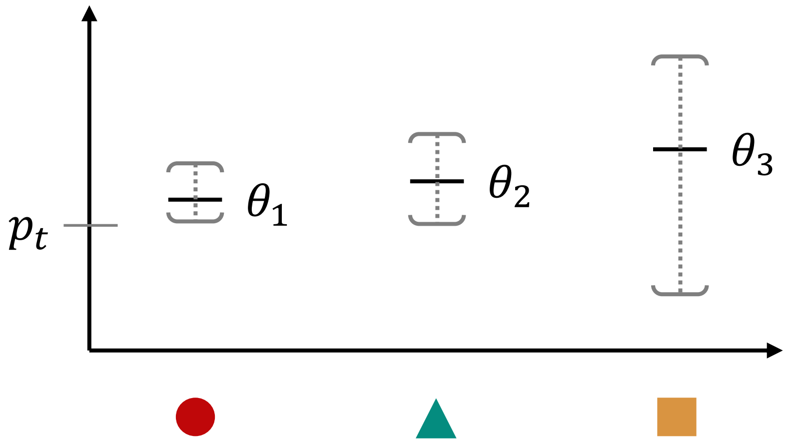

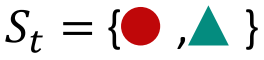

Algorithm: 2nd phase type elimination

For each active type $i \in S_t$, algorithm estimates $\mathrm{rev}(\theta_i, Q)$

$$\mathrm{rev}(\theta_i, Q) = \theta_i \cdot \mathbb{P}_{j \sim \mathcal{D}}\left[\theta_j \geq \theta_i \text{ and } j \in Q\right]$$

Requires care because at each round:

- Don't observe $\mathbf{1}_{\{\theta_{i_t} \geq \theta_i \text{ and } i_t \in Q\}}$
- Only observe $\mathbf{1}_{\{b_t = 1,\, \theta_{i_t} \geq \theta_i, \text{ and } i_t \in Q\}}$

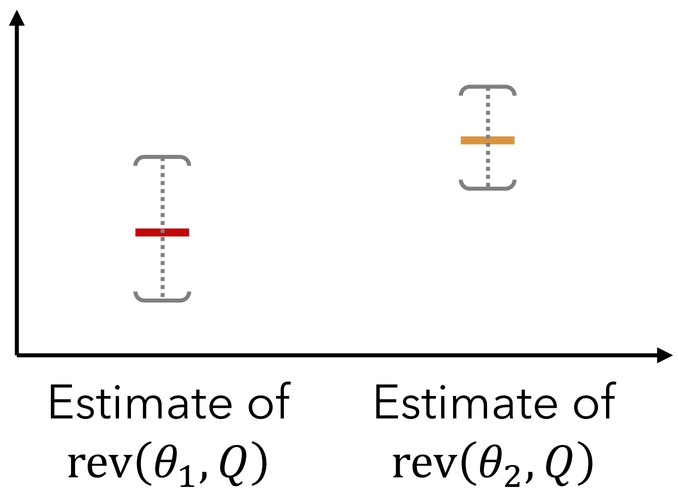

Algorithm: 2nd phase type elimination

For each active type $i \in S_t$, algorithm estimates $\text{rev}(\theta_i, Q)$

$$\text{rev}(\theta_i, Q) = \theta_i \cdot \mathbb{P}_{j \sim \mathcal{D}}[\theta_j \geq \theta_i \text{ and } j \in Q]$$

Removes types from S_t if estimate is too small

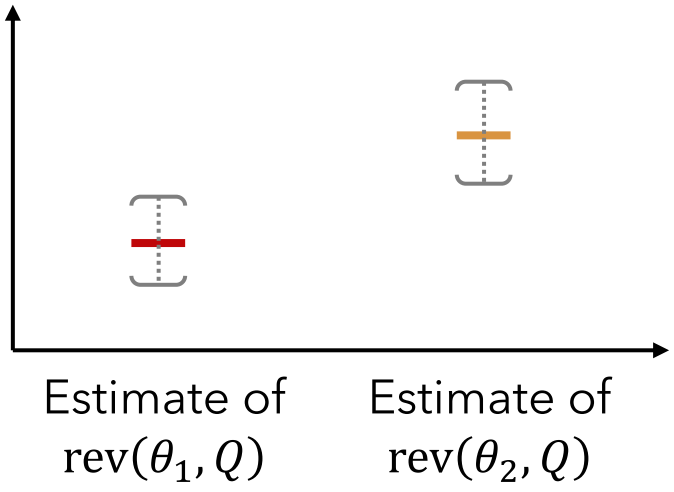

Algorithm summary

Phase 1:

- Offer item for free to get samples from type distribution
- Set Q to be set of types that appeared sufficiently often

Phase 2:

- Only aim to compete with $p^*(Q) = \theta_{i_Q}$ for some $i_Q \in Q$
- Maintain set S_t such that $i_Q \in S_t$
- Set price low enough so that buyers in S_t always buy
- Eliminate types from S_t that contribute too little revenue

Distinctions from explore-then-commit

"Explore" phase of ETC is much **longer** (often $O\left(T^{2/3}\right)$ rounds)

ETC algorithms focus on learning **all unknowns** in 1st phase

We only focus on eliminating low probability types

"Commit" phase of ETC often **doesn't include any learning**

In the 2nd phase, our algorithm is still learning the optimal price

Unlike our algorithm, ETC **can't obtain $O\left(T^{1/2}\right)$ regret**

Outline

1. Introduction
2. Mechanism design background
3. Model
4. Connections to prior research
5. Main results
    1. Regret bound overview
    2. Algorithm
    3. **Regret bound proof sketch**
    4. Lower bound
6. Conclusions and future directions

Regret bound

Theorem: If $q_{\min}$ is tiny $\left(q_{\min} < 2d^{-2/3}T^{-1/3}\right)$

$$\mathbb{E}[R_T] = O\left(T^{1/3}d^{2/3} + T^{2/3}d^{1/3}\right)$$

Otherwise,

$$\mathbb{E}[R_T] = O\left(\sqrt{\frac{T}{q_{\min}}} + T^{1/3}d^{2/3}\right)$$

$q_{\min}$= minimum probability of any type $\left(\min_{i \in [d]} \mathbb{P}_{j \sim \mathcal{D}}[j = i]\right)$

Proof sketch

Theorem: If $q_{\min}$ is tiny $\left(q_{\min} < 2d^{-2/3}T^{-1/3}\right)$

$$\mathbb{E}[R_T] = O\left(T^{1/3}d^{2/3} + T^{2/3}d^{1/3}\right)$$

Proof sketch:

- In 1st phase, item offered for free
- Phase lasts $\widetilde{\Theta}\left(T^{1/3}d^{2/3}\right)$ rounds

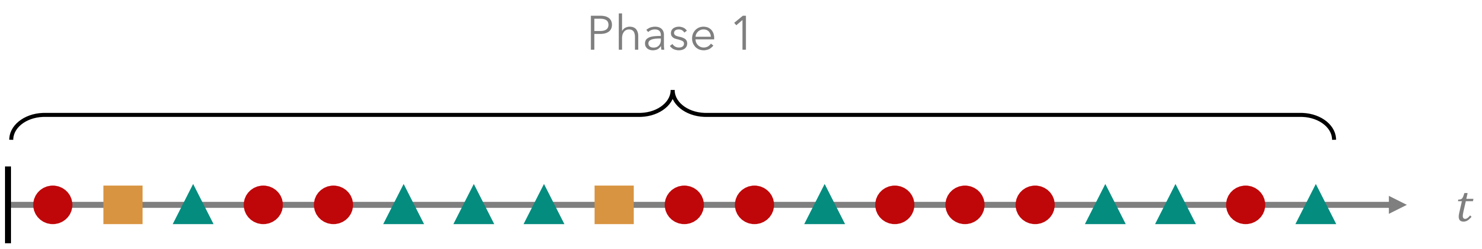

Proof sketch

Theorem: If $q_{\min}$ is tiny $\left(q_{\min} < 2d^{-2/3}T^{-1/3}\right)$

$$\mathbb{E}[R_T] = O\left(T^{1/3}d^{2/3} + T^{2/3}d^{1/3}\right)$$

Proof sketch:

- In 2nd phase, alg competes with $p^*(Q)$
- Competing with $p^*(Q)$ instead of optimal price adds $O\left(T^{2/3}d^{1/3}\right)$ regret

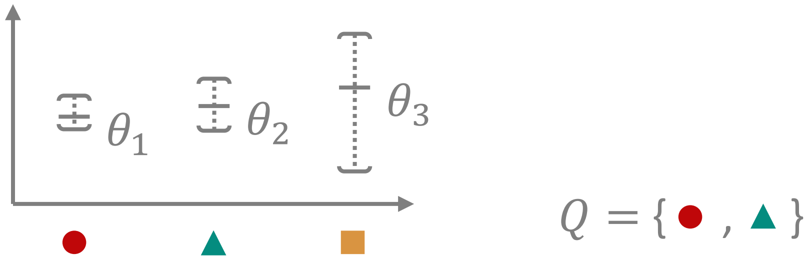

Proof sketch

Theorem: If $q_{\min}$ is tiny $\left(q_{\min} < 2d^{-2/3}T^{-1/3}\right)$

$$\mathbb{E}[R_T] = O\left(T^{1/3}d^{2/3} + T^{2/3}d^{1/3}\right)$$

Proof sketch:

- In 2nd phase, maintains estimates of $\mathrm{rev}(\theta_i, Q)$ for all $i \in S_t$
- Error of estimates contributes $\tilde{O}(\sqrt{T})$ to regret

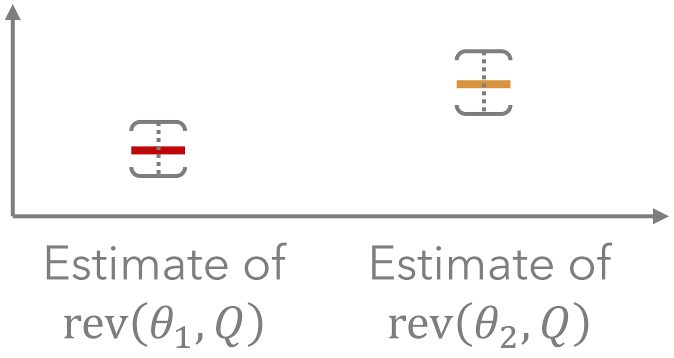

Proof sketch

Theorem: If $q_{\min}$ is tiny $\left(q_{\min} < 2d^{-2/3}T^{-1/3}\right)$

$$\mathbb{E}[R_T] = O\left(T^{1/3}d^{2/3} + T^{2/3}d^{1/3}\right)$$

Proof sketch:

- Agents themselves are learning
- Increases the regret by $O\left(d^{1/3}T^{1/3}\sqrt{\ln\frac{1}{\eta}}\right)$

What changes when $q_{\min}$ isn't tiny?

Theorem: If $q_{\min}$ isn't tiny $\left(q_{\min} > 2d^{-2/3}T^{-1/3}\right)$

$$\mathbb{E}[R_T] = O\left(\sqrt{\frac{T}{q_{\min}}} + T^{1/3}d^{2/3}\right)$$

Proof sketch:

- Same analysis structure, but we prove that WHP, $Q = [d]$
- Significantly reduces the sources of regret

Outline

1. Introduction
2. Mechanism design background
3. Model
4. Main results
    1. Regret bound overview
    2. Algorithm
    3. Regret bound proof sketch
    4. **Lower bound**
5. Conclusions and future directions

Nearly-matching regret lower bound

Theorem:

- $q_{\min}$-independent lower bound of $\Omega\left(T^{2/3} d^{1/3}\right)$
- If $q_{\min} > T^{-1/3} d^{-2/3}$, lower bound of $\Omega\left(\sqrt{\frac{T}{q_{\min}}}\right)$

Lower bound proof intuition

Different types have similar value distributions

- But large variation in appearance probabilities

Intuitively, any algorithm must decide if:

- It will target low-probability buyers (large confidence intervals)

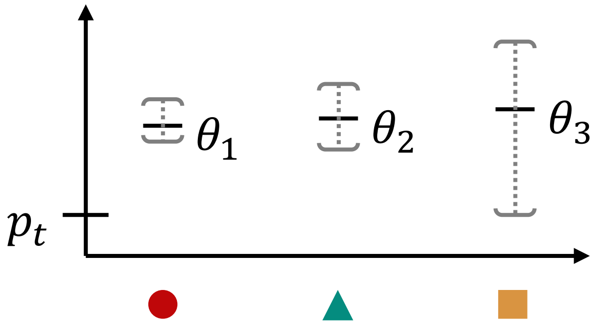

Lower bound proof intuition

Different types have similar value distributions

- But large variation in appearance probabilities

Intuitively, any algorithm must decide if:

- It will target low-probability buyers (large confidence intervals)
- Or ignore low-probability buyers

Either way, any algorithm suffers high regret

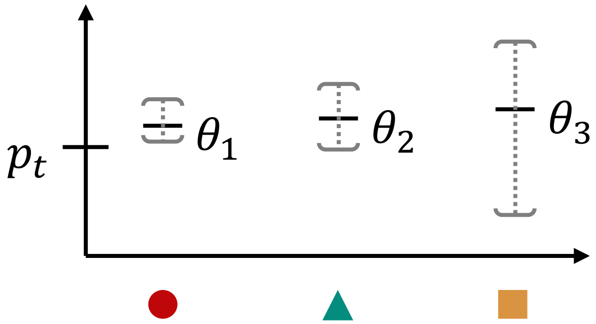

Lower bound proof intuition ($d = 2$)

- $\theta_1 = \theta_2 = \frac{1}{2} \quad \Rightarrow \quad$ baseline's price is $\frac{1}{2}$
- $\mathbb{P}_{i \sim \mathcal{P}}[i = 1] = q, \quad \mathbb{P}_{i \sim \mathcal{P}}[i = 2] = 1 - q, \quad q = T^{-1/3}$
- Type 1's lower bound will always be $\lesssim \frac{1}{2} - \sqrt{\frac{1}{qT}} = \frac{1}{2} - T^{-1/3}$
- If target Type 1 and 2, must set $p_t \lesssim \frac{1}{2} - T^{-1/3}$
  - Means regret is at least $T \cdot T^{-1/3} = T^{2/3}$
- If only target Type 2, will lose $\approx qT = T^{2/3}$ rev. from Type 1
- No algorithm can do better than these two extremes

Lower bound proof insight

Proof indicates that any policy can't do better than one that

- Chooses ahead of time to **target all customer types**, or
- Only focus on the **high probability** types

⇒ Doesn't help to dynamically change which types to target

This **mirrors the behavior** of our algorithm as well:

- Uses a short initial phase to eliminate low probability types
- Thereon, it only targets the remaining high probability types

Outline

1. Introduction
2. Mechanism design background
3. Model
4. Connections to prior research
5. Main results
6. **Conclusions and future directions**

Conclusions

No-regret pricing strategies

- Both sides of the market are **learning from reviews**

Algorithm strategically sets **low prices early on**

- Boosts sales from customers who have **rare types** and **high values**

Algorithm trades off:

- Revenue loss due to **discounts** from the initial phase, and
- Future **revenue gains**

Lower bound: algorithm is optimal up to lower order terms

Future directions

Pricing strategies when buyers don't always leave reviews

Mimics real-world buyer behaviors

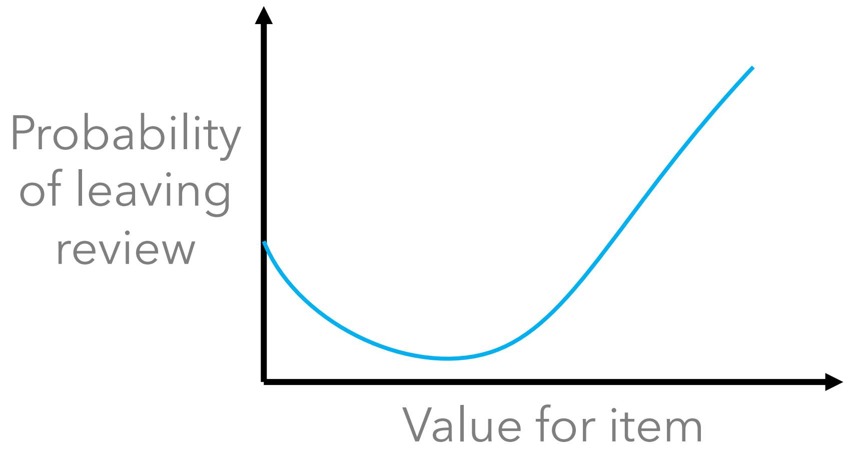

Future directions

What if the buyers **appear over several rounds**?

May **behave strategically** in order to purchase at lower future prices

Prior research: Buyers bid strategically over many interactions

- Key difference: buyers know their own values
- [Braverman et al., ‘18, Deng et al., ‘19, Nekipelov et al., ‘15, Devanur et al., ‘14]

Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty

Ellen Vitercik
Stanford University

Joint work with Wenshuo Guo, Nika Haghtalab, and Kirthevasan Kandasamy

Appeared in the Conference on Economics and Computation (EC) 2023