LEVERAGING REVIEWS: LEARNING TO PRICE WITH BUYER AND SELLER UNCERTAINTY

ECONOMICS & COMPUTATION 2023

WENSHUO GUO (UC BERKELEY)

NIKA HAGHTALAB (UC BERKELEY)

KIRTHI KANDASAMY (UW-MADISON)

ELLEN VITERCIK (STANFORD)

ONLINE MARKETPLACES ARE UBIQUITOUS

shopify

ebay

amazon

DOORDASH

ONLINE MARKETPLACES ARE UBIQUITOUS

shopify

ebay

amazon

DOORDASH

- Many reasons: convenience, many options, **reviews**.

CUSTOMERS USE REVIEWS TO MAKE AN INFORMED PURCHASE

Cuisinart 422-24 Contour Stainless 10-Inch Open Skillet

Visit the Cuisinart Store

3,625 ratings

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

Visit the Hartz Store

7,607 ratings

Paula's Choice Skin Perfecting 2% BHA Liquid Salicylic Acid Exfoliant, Gentle Facial Exfoliator for Blackheads, Large Pores, Wrinkles & Fine Lines, Travel Size, 1 Fluid Ounce - PACKAGING MAY VARY

Visit the Paula's Choice Store

79,839 ratings

(LCB '22 AI & Marketing, AMMO '22 Econometrica, MD '10 MIS Quarterly)

CUSTOMERS USE REVIEWS TO MAKE AN INFORMED PURCHASE

Cuisinart 422-24 Contour Stainless 10-Inch Open Skillet

Visit the Cuisinart Store

★★★★☆ 3,625 ratings

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

Visit the Hartz Store

★★★★☆ 7,607 ratings

Paula's Choice Skin Perfecting 2% BHA Liquid Salicylic Acid Exfoliant, Gentle Facial Exfoliator for Blackheads, Large Pores, Wrinkles & Fine Lines, Travel Size, 1 Fluid Ounce - PACKAGING MAY VARY

Visit the Paula's Choice Store

★★★★☆ 79,839 ratings

(LCB '22 AI & Marketing, AMMO '22 Econometrica, MD '10 MIS Quarterly)

- But customers do not look at just the average rating.

FILTERING REVIEWS BY 'CUSTOMER TYPE'

Cuisinart MCP22-24N MultiClad Pro Triple Ply 10-Inch, Open Skillet

Visit the Cuisinart Store

14,945 ratings

4.7 out of 5

FILTERING REVIEWS BY 'CUSTOMER TYPE'

Cuisinart MCP22-24N MultiClad Pro Triple Ply 10-Inch, Open Skillet

Visit the Cuisinart Store

14,945 ratings

4.7 out of 5

Looking for specific info?

oven

Customer Reviews

Warped

By Cheryl A. Jarrett in the United States on April 20, 2022

...They warp in the oven. see more

Warps

By Ricky K Workman in the United States on August 3, 2022

...Warps at 350 degrees see more

See 20 matching customer reviews ›

FILTERING REVIEWS BY ‘CUSTOMER 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

HOW CAN REVIEWS BE HELPFUL?

HOW CAN REVIEWS BE HELPFUL?

- **To buyers (customers):**
  - Understand if the product is right for them.
    - E.g. Several positive reviews for stovetop cooking, but warps too frequently in the oven.

HOW CAN REVIEWS BE HELPFUL?

- **To buyers (customers):**
  - Understand if the product is right for them.
    - E.g. Several positive reviews for stovetop cooking, but warps too frequently in the oven.
- **To sellers:**
  - Gauge the demand for the product $\Longrightarrow$ set prices to maximize revenue.
    - E.g. Several 5 star reviews! I should increase the price.

HOW CAN REVIEWS BE HELPFUL?

- **To buyers (customers):**
  - Understand if the product is right for them.
    - E.g. Several positive reviews for stovetop cooking, but warps too frequently in the oven.
- **To sellers:**
  - Gauge the demand for the product $\implies$ set prices to maximize revenue.
    - E.g. Several 5 star reviews! I should increase the price.
- **Goal:** Study how reviews can help *both* sides of the market.

HOW CAN REVIEWS BE HELPFUL?

- **To buyers (customers):**
  - Understand if the product is right for them.
    - E.g. Several positive reviews for stovetop cooking, but warps too frequently in the oven.
- **To sellers:**
  - Gauge the demand for the product $\implies$ set prices to maximize revenue.
    - E.g. Several 5 star reviews! I should increase the price.
- **Goal:** Study how reviews can help *both* sides of the market.
  - Sellers will **learn optimal price**.
  - Buyers will **learn their *value* for goods**.

POSTED PRICE MECHANISM

POSTED PRICE MECHANISM

- A single seller who wishes to sell an item.

POSTED PRICE MECHANISM

- A single seller who wishes to sell an item.
- Each buyer has a type i. There is a **distribution** $\mathscr{P}$ of buyer types.

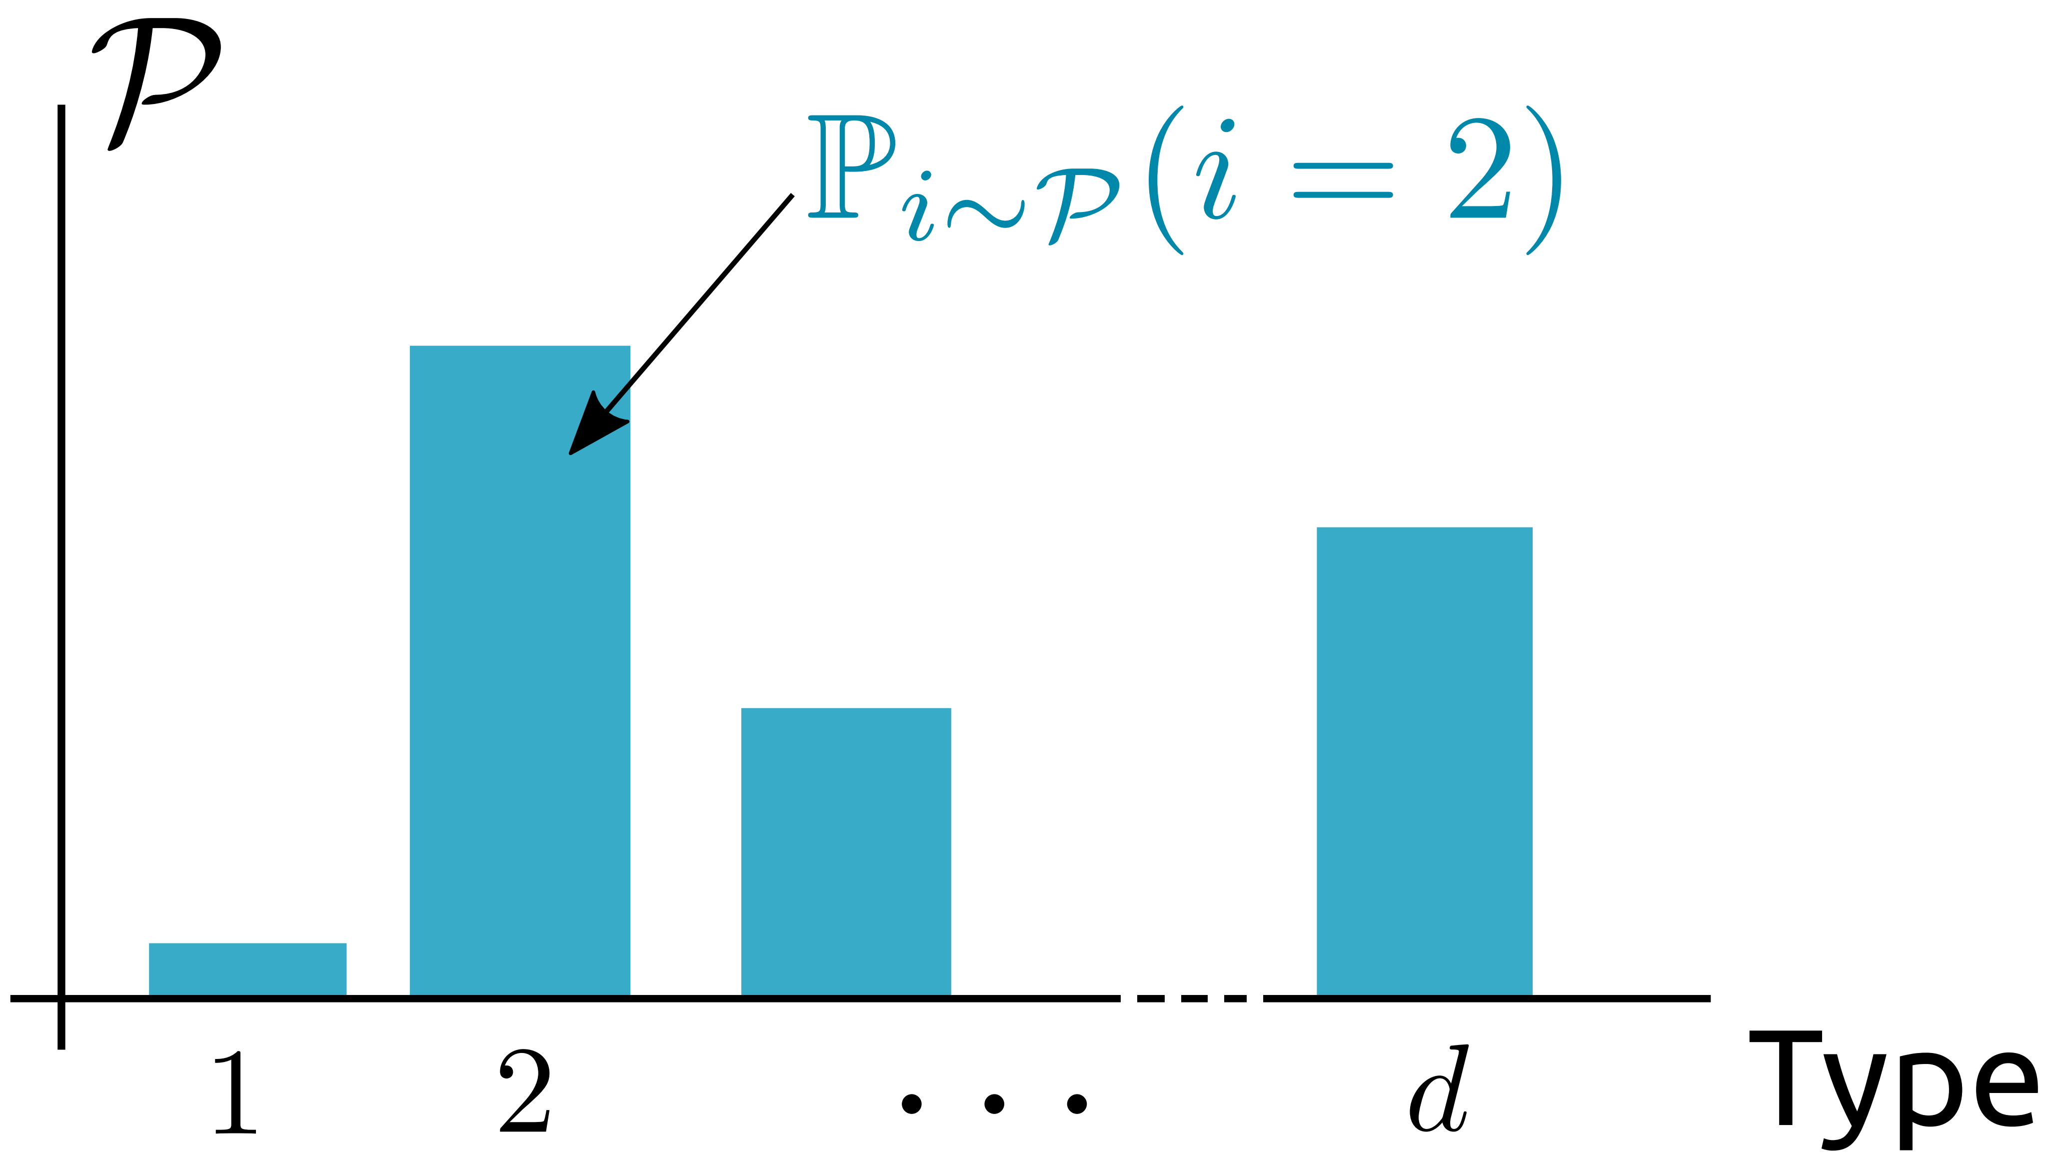

POSTED PRICE MECHANISM

- A single seller who wishes to sell an item.
- Each buyer has a type i. There is a **distribution** $\mathscr{P}$ of buyer types.
- All buyers of type i have **value** θ_i for the item.

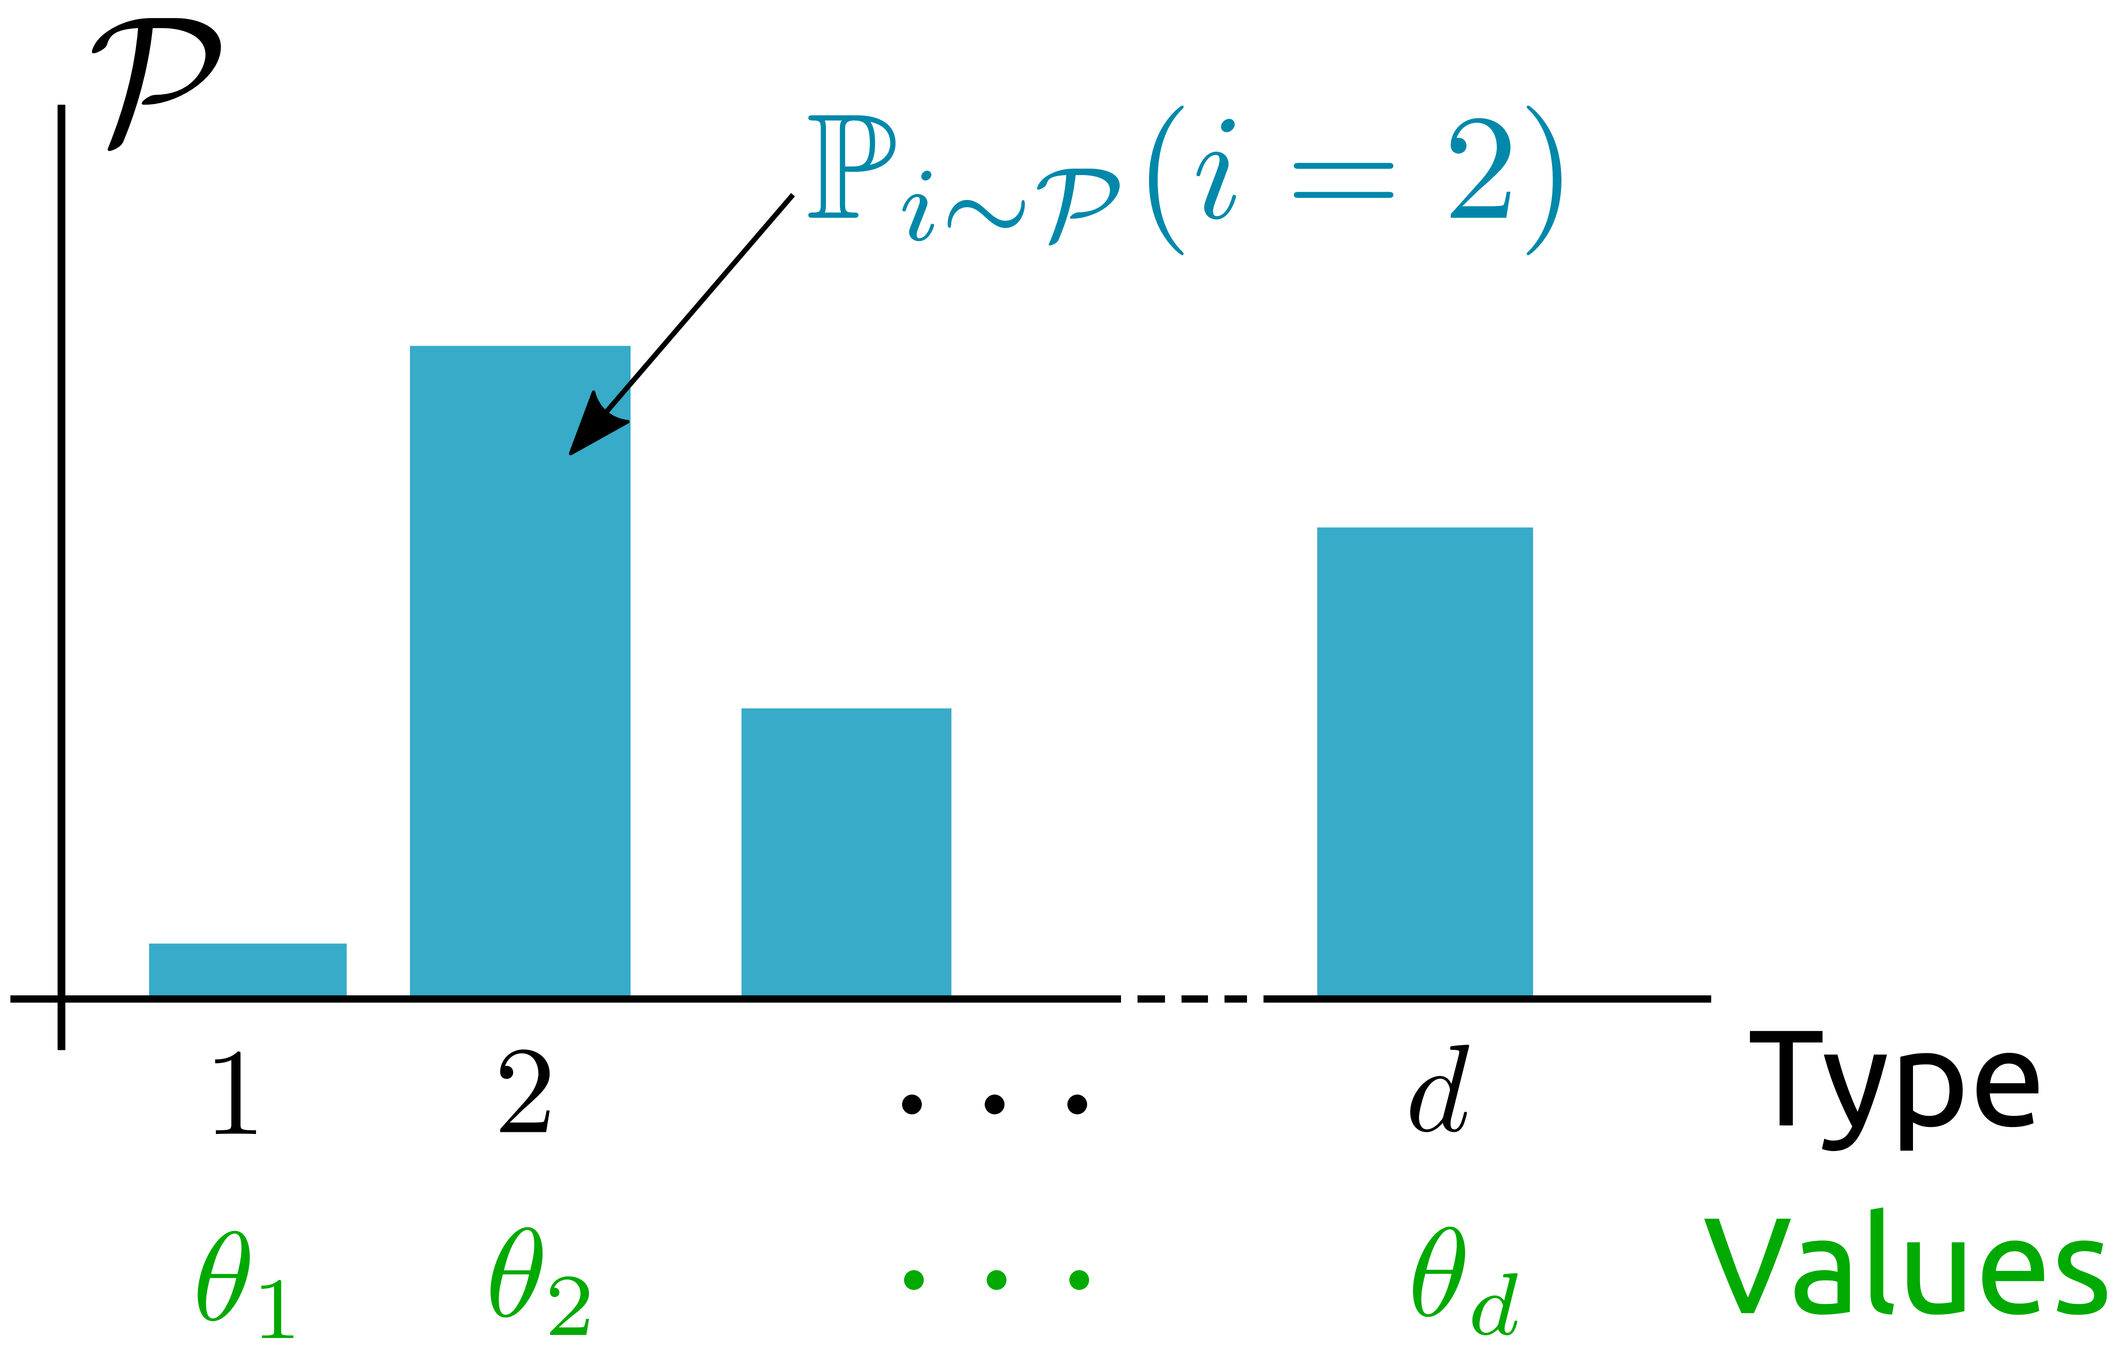

POSTED PRICE MECHANISM

- A single seller who wishes to sell an item.
- Each buyer has a type i. There is a **distribution** $\mathscr{P}$ of buyer types.
- All buyers of type i have **value** θ_i for the item.
- Seller posts a price p for the item. A buyer of type i will purchase if her value is larger than the price, i.e $\theta_i \geq p$.

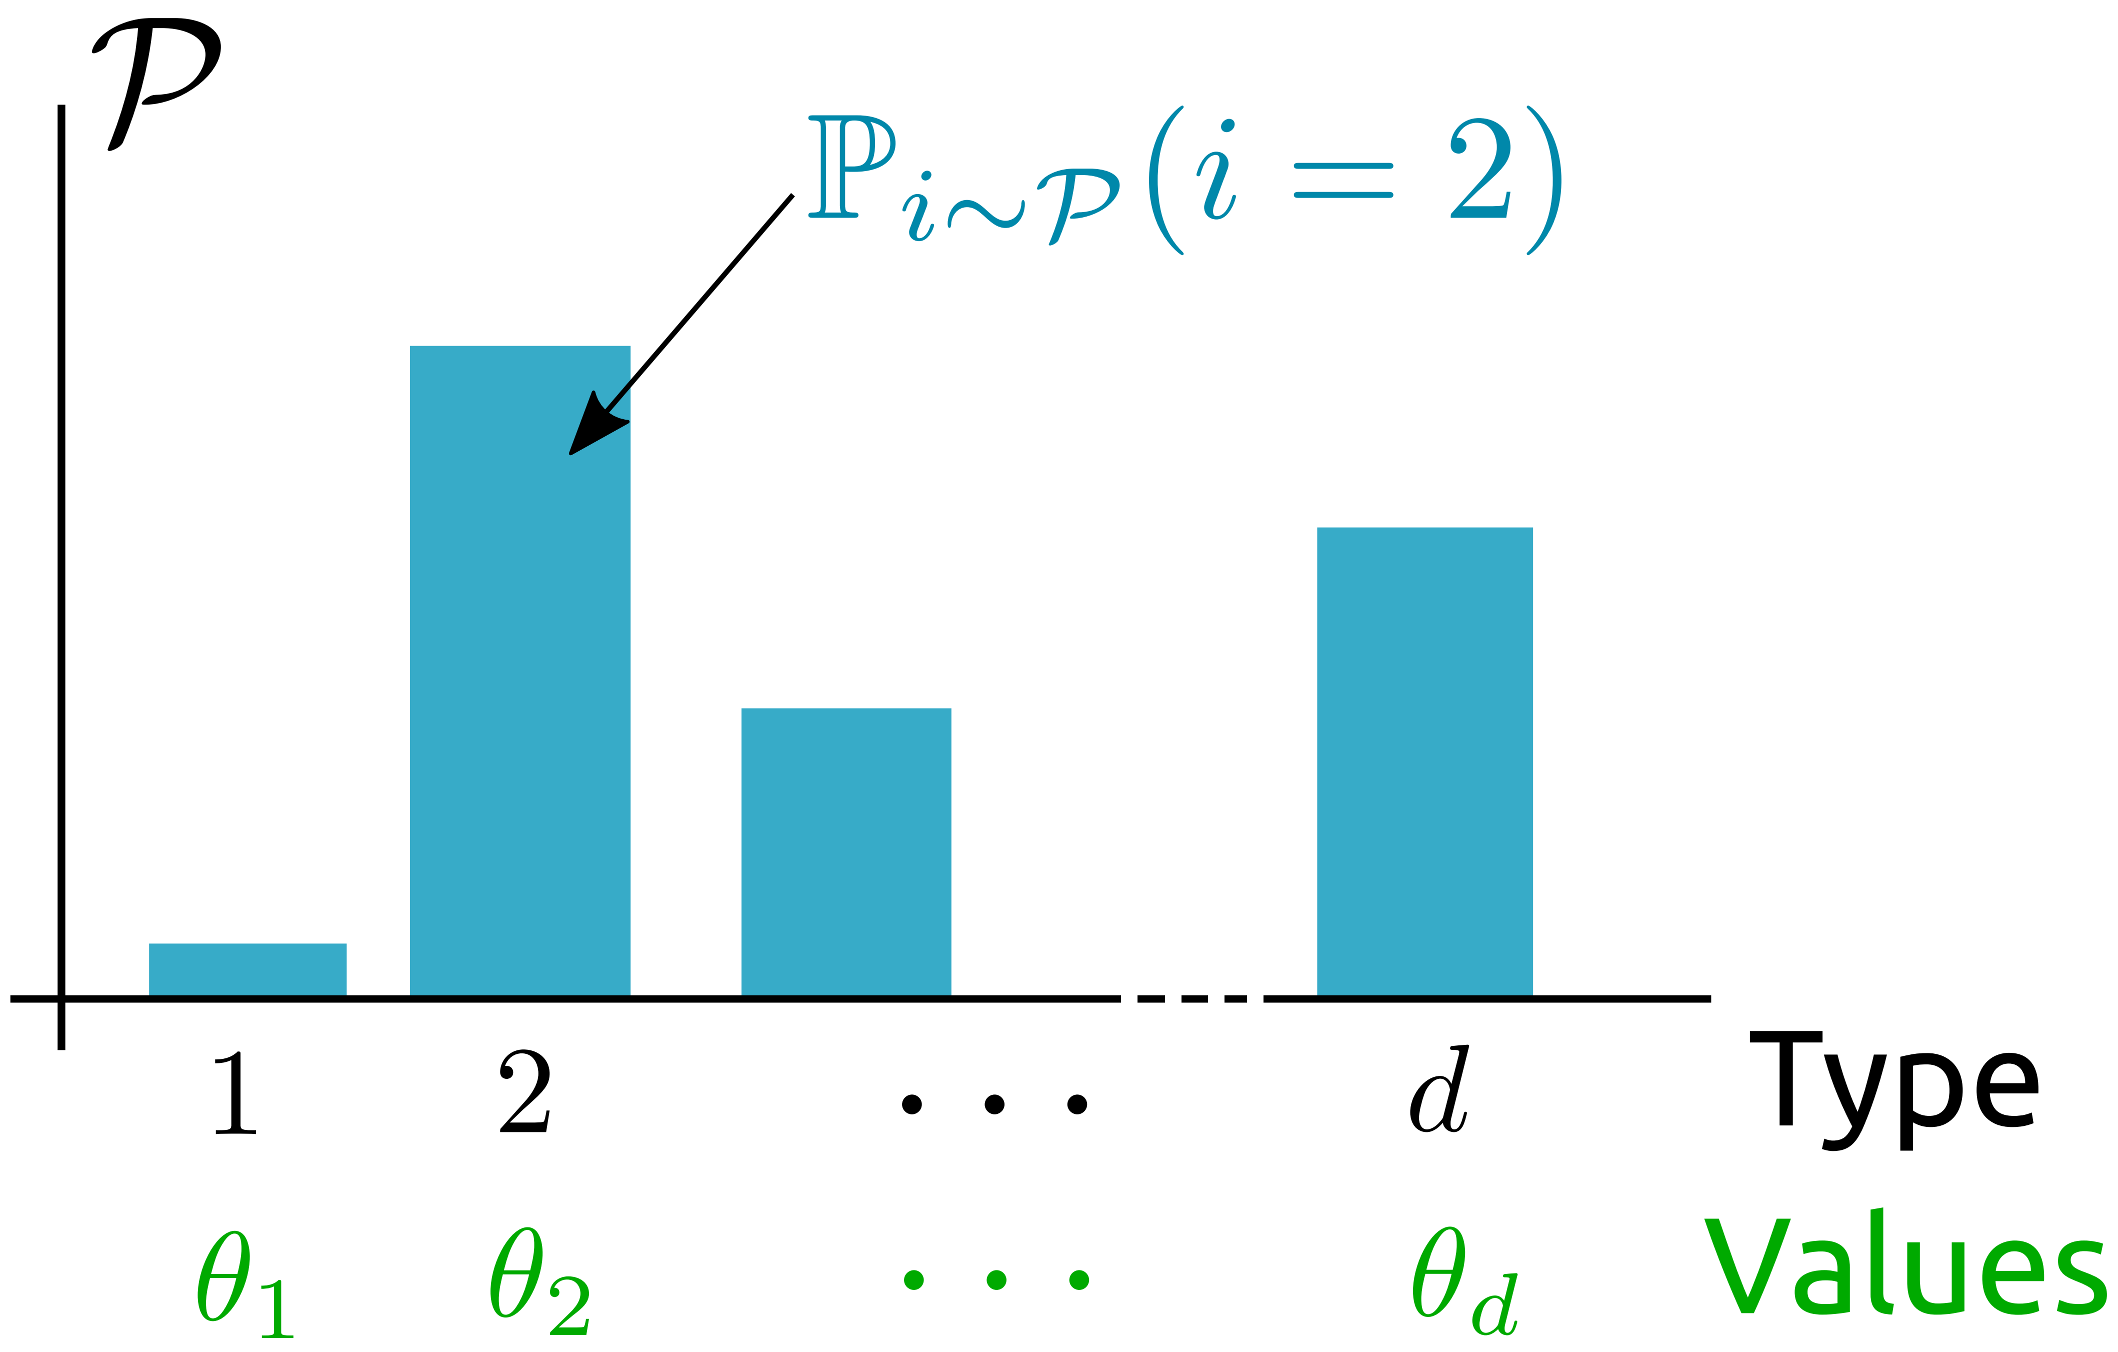

POSTED PRICE MECHANISM

- A single seller who wishes to sell an item.
- Each buyer has a type i. There is a **distribution** $\mathscr{P}$ of buyer types.
- All buyers of type i have **value** θ_i for the item.
- Seller posts a price p for the item. A buyer of type i will purchase if her value is larger than the price, i.e $\theta_i \geq p$.
- Optimal price:

$$p^\star = \arg\max_p p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

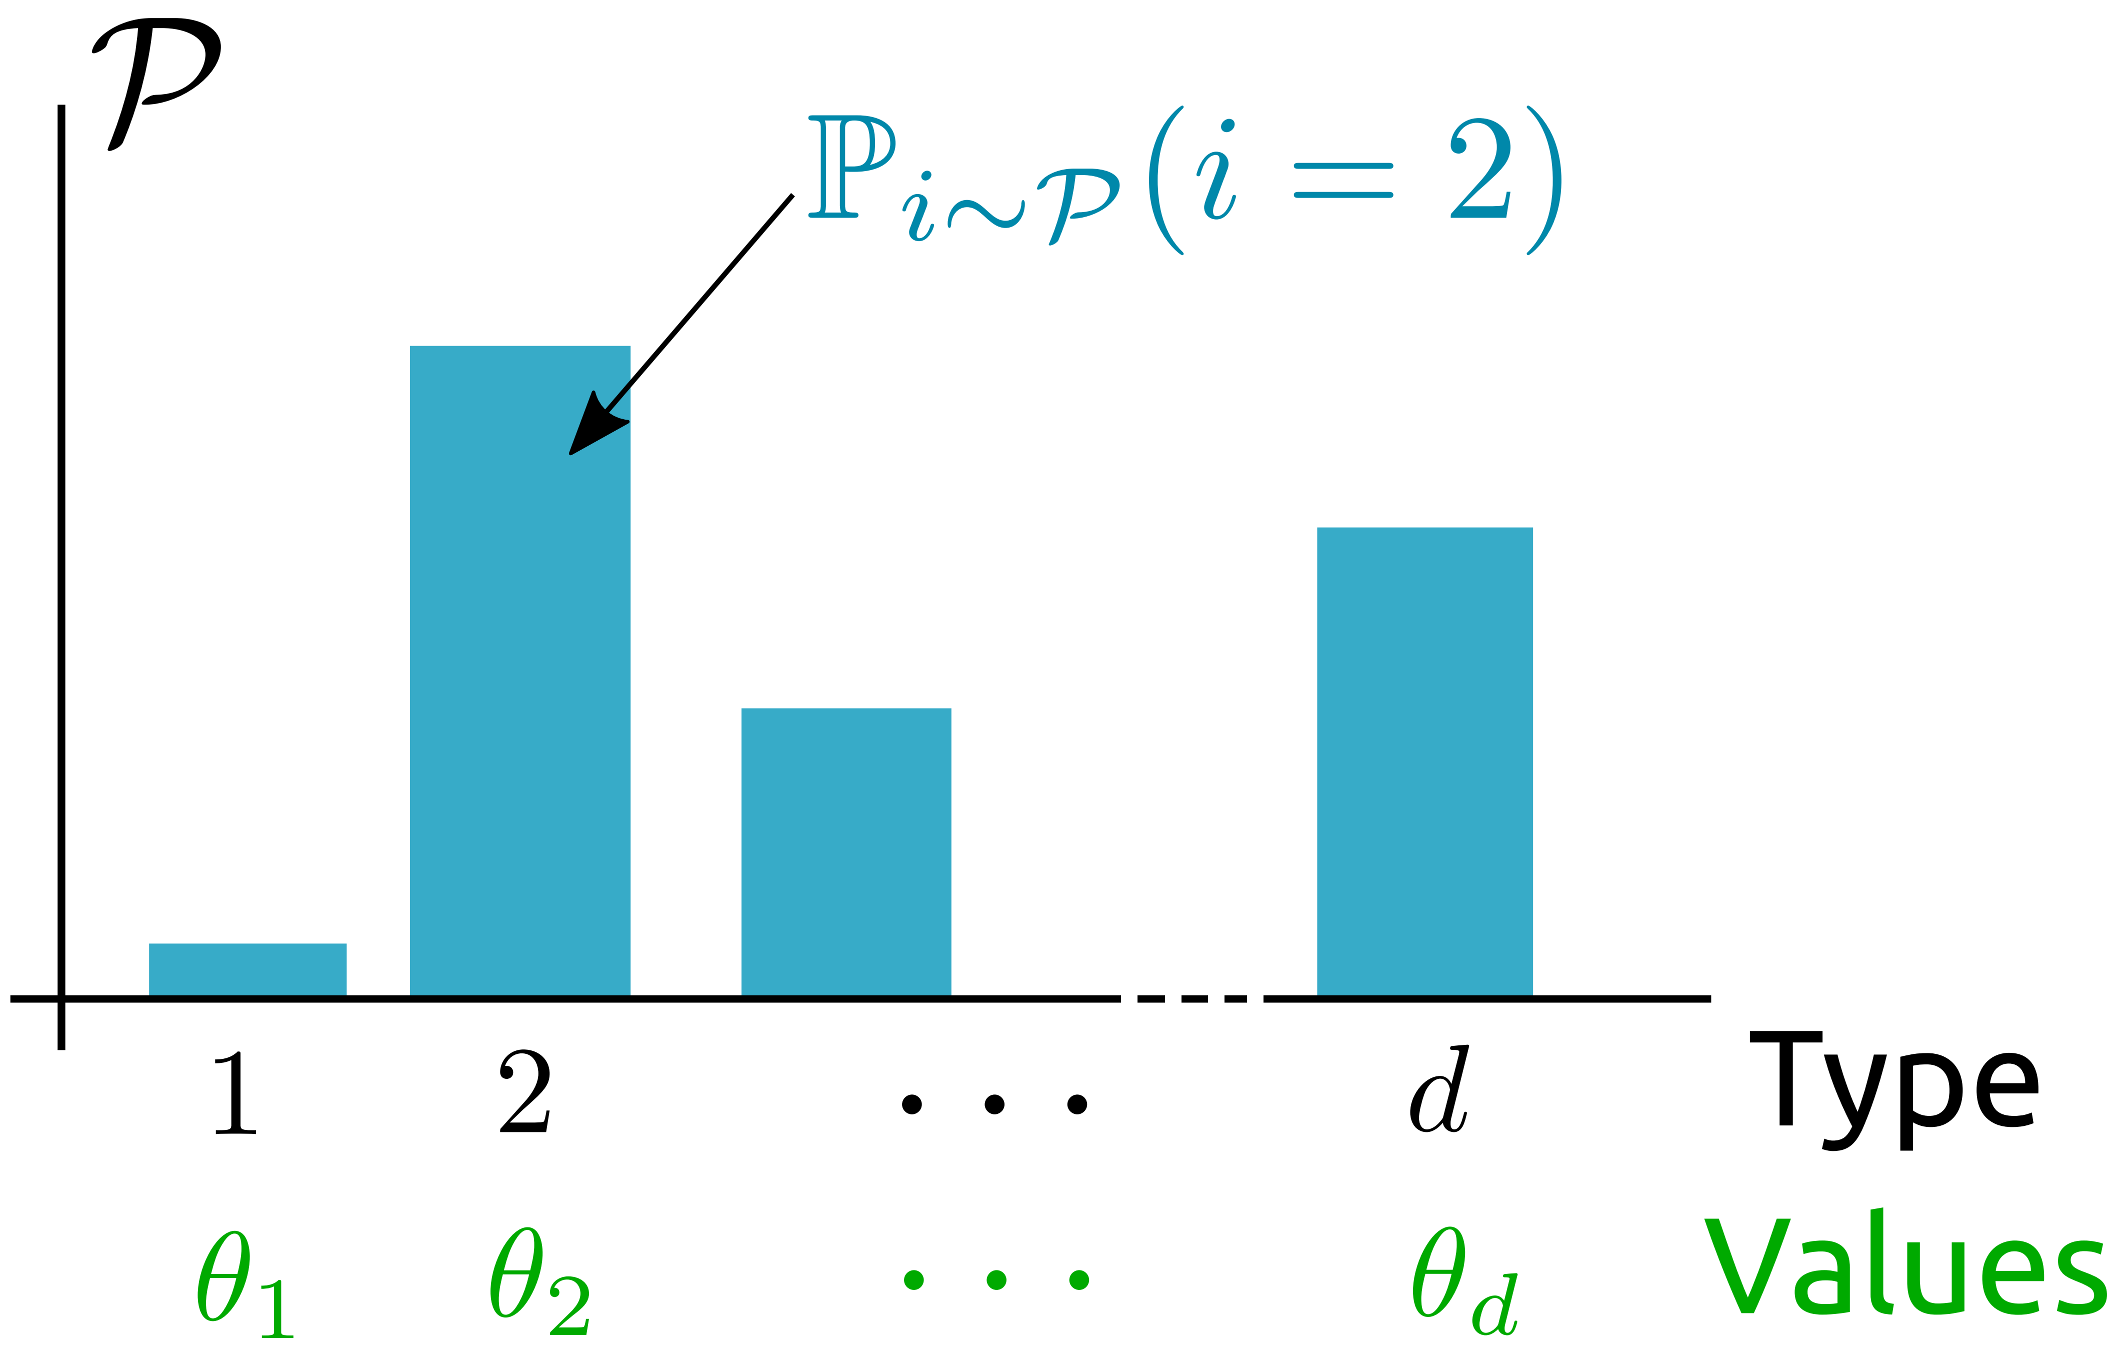

ISSUES

ISSUES

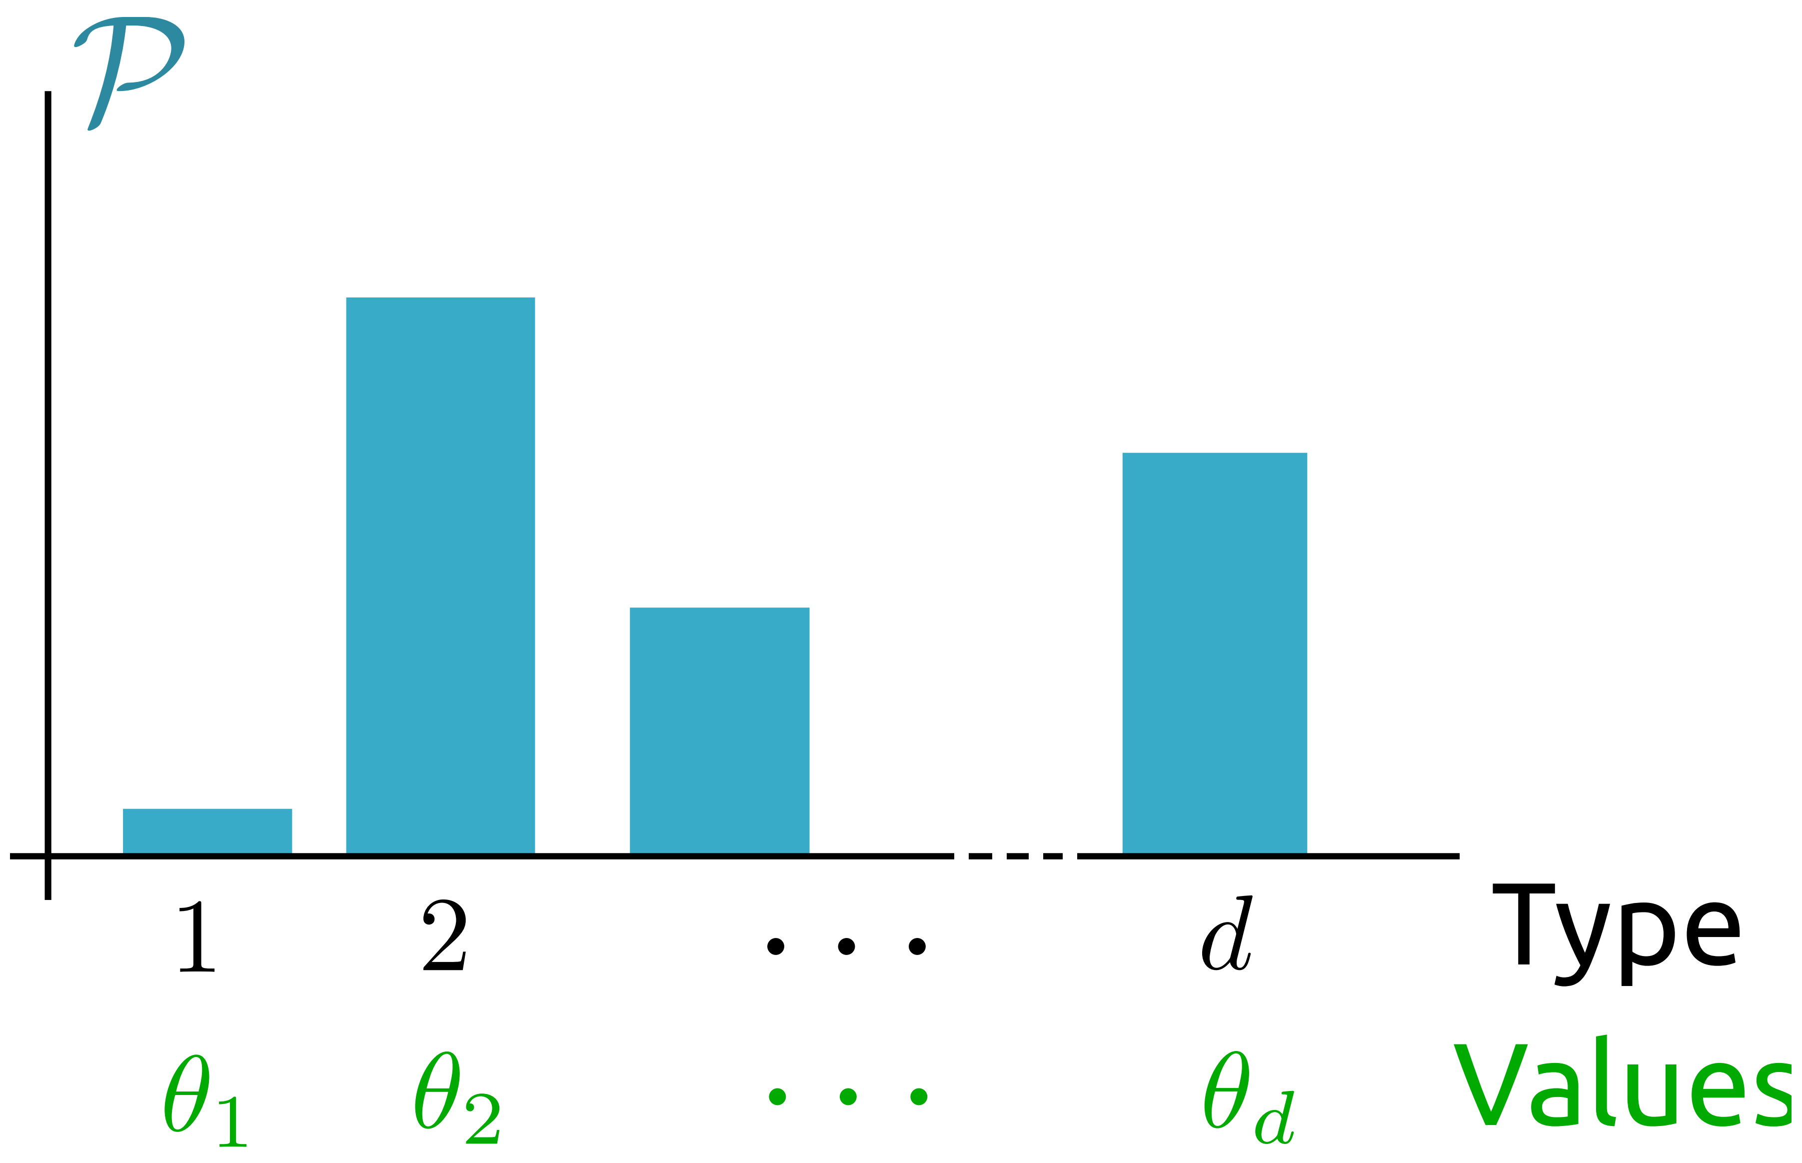

ISSUES

1. Seller wishes to maximize revenue, but may not know $\mathscr{P}$.

$$p^\star = \arg\max_p p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

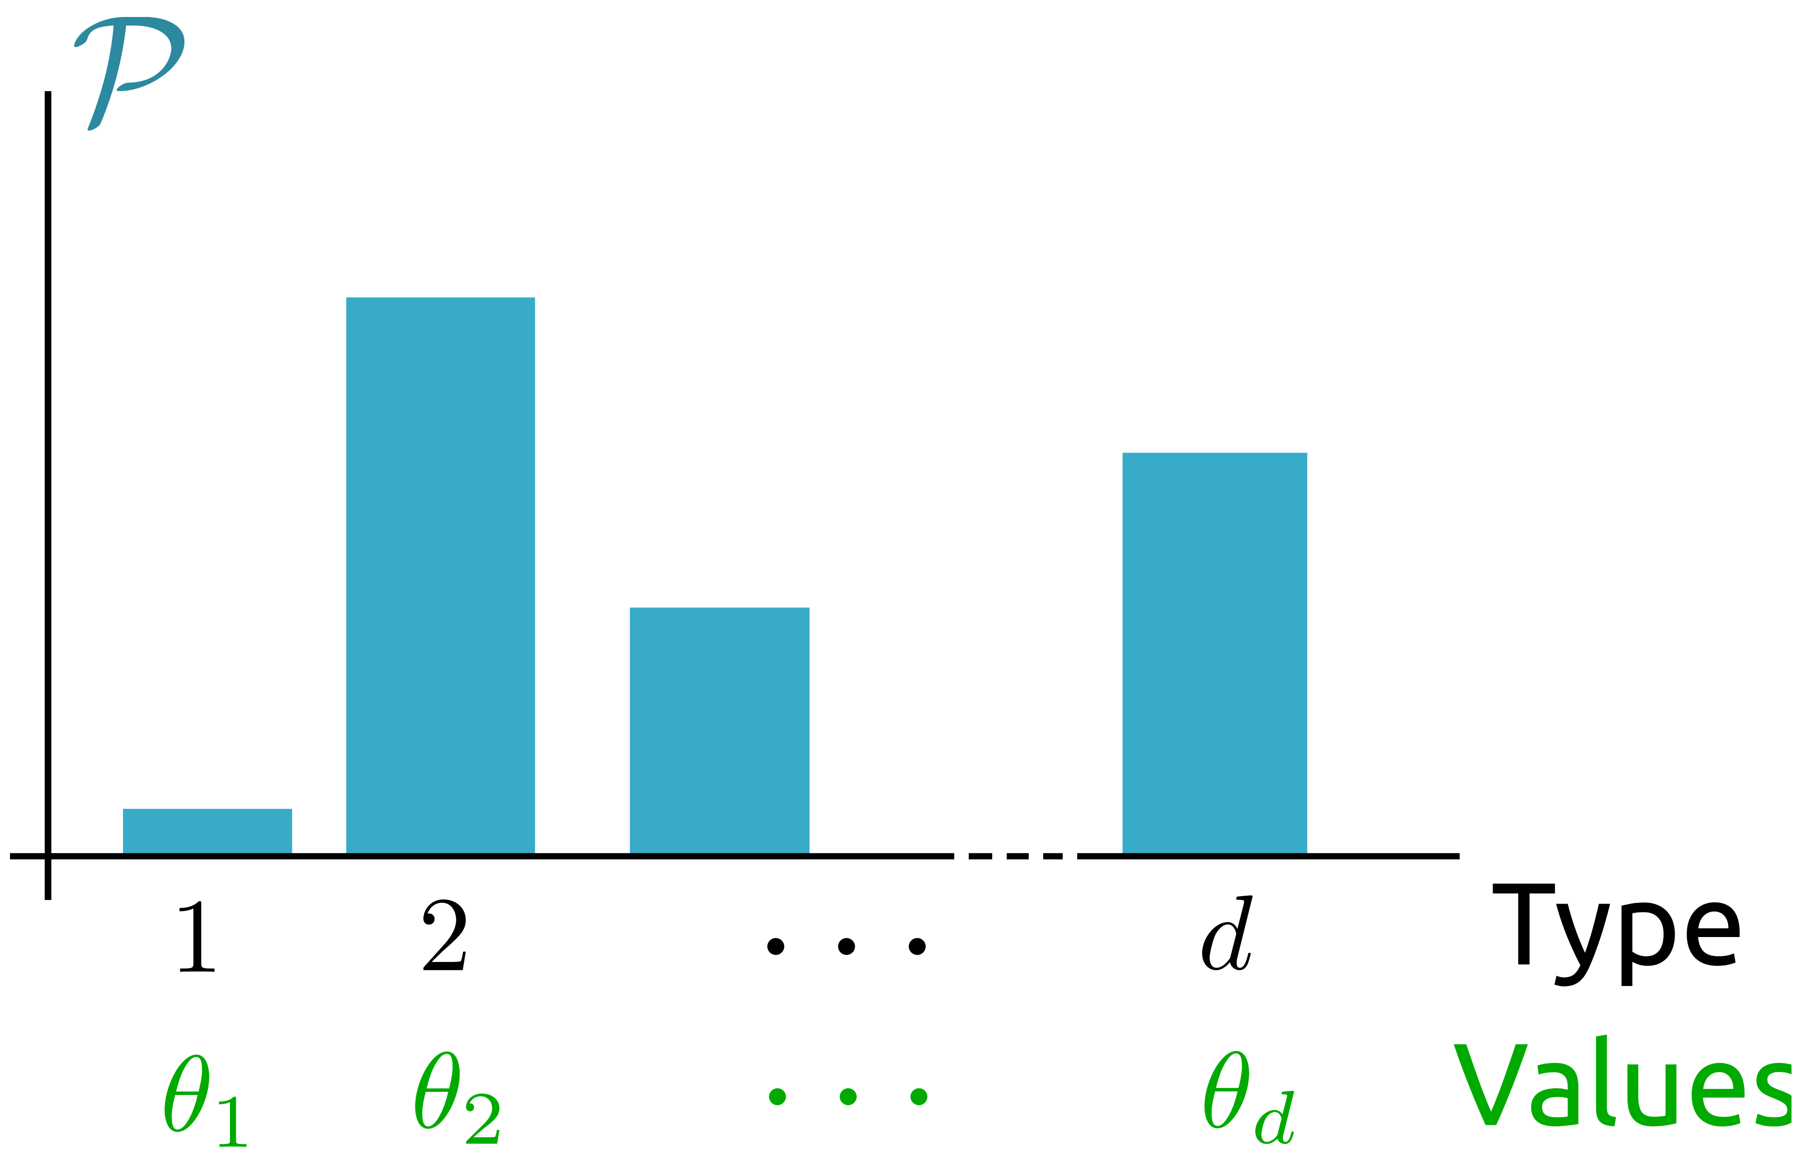

ISSUES

1. Seller wishes to maximize revenue, but may not know $\mathscr{P}$.

$$p^\star = \arg\max_p p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

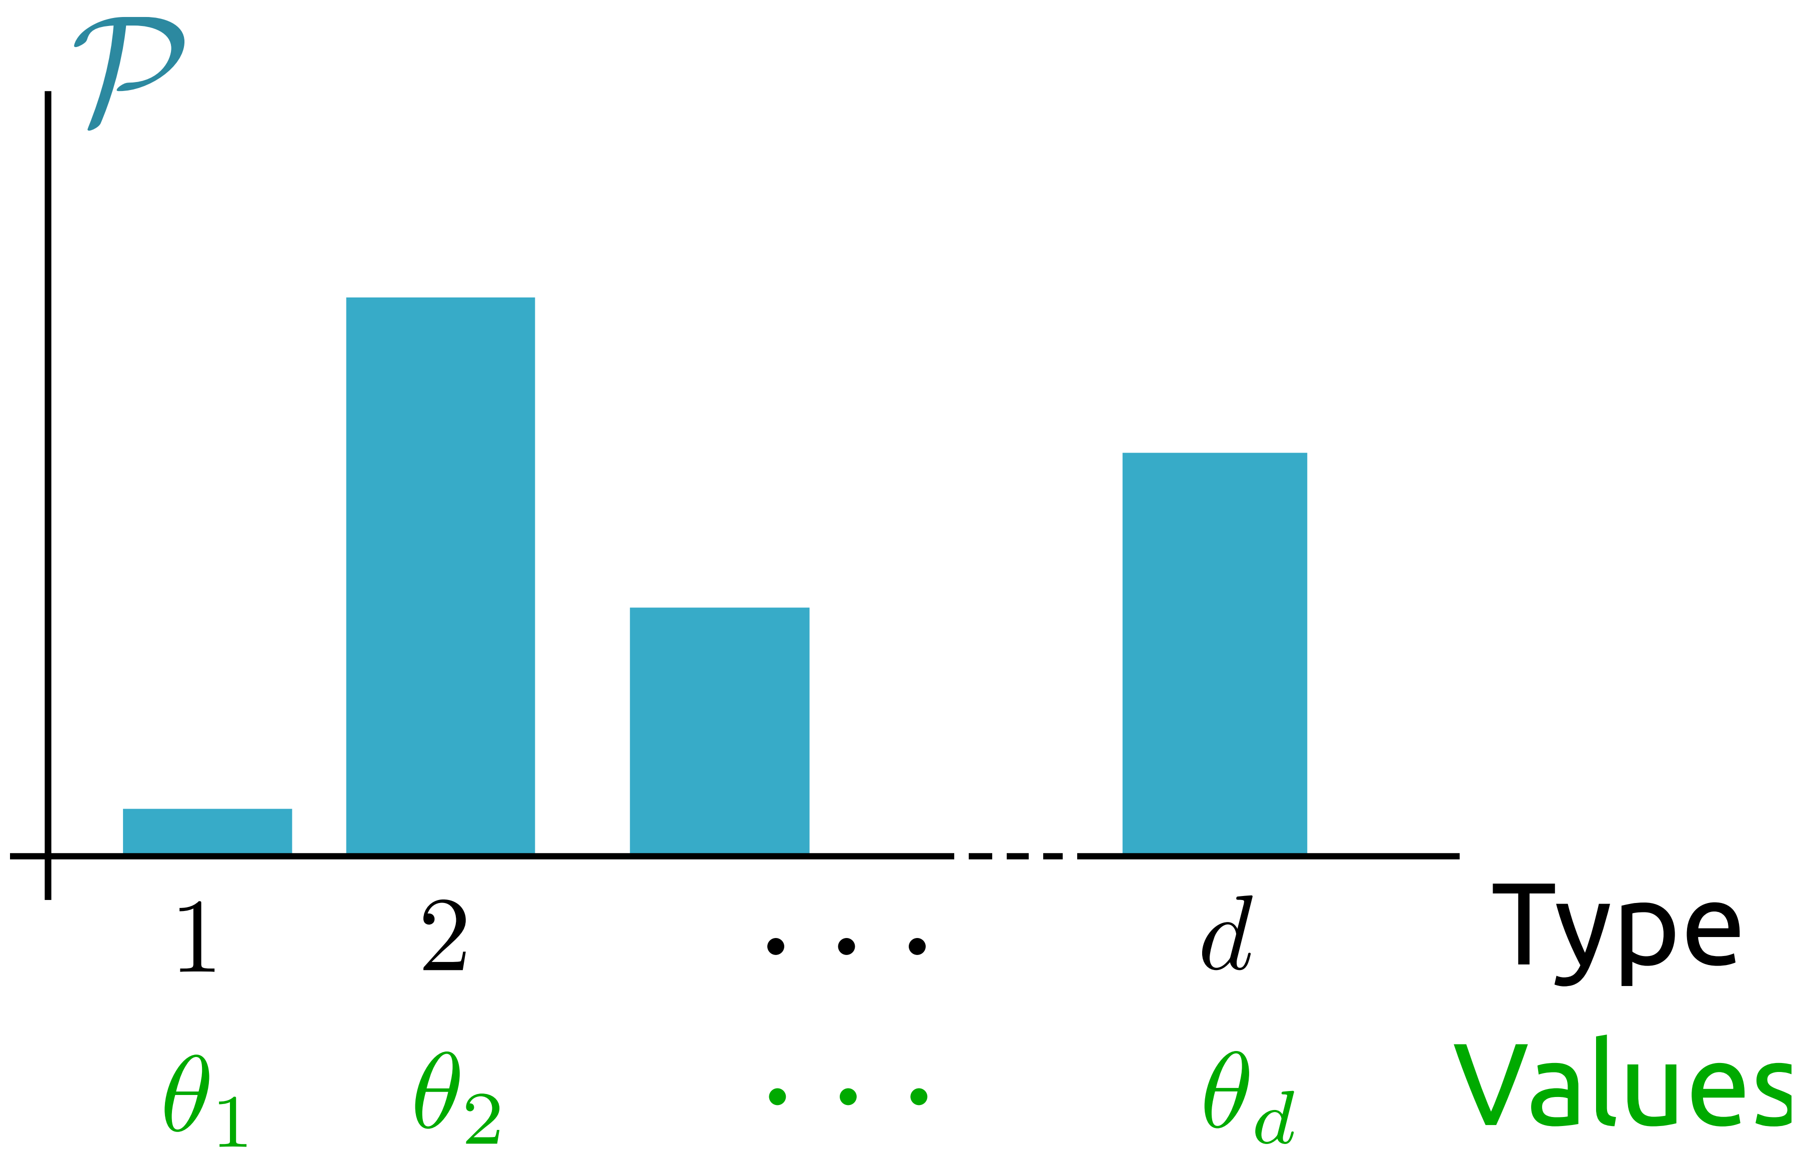

2. A buyer may know their type i, but not their value θ_i.

1. Seller wishes to maximize revenue, but may not know $\mathscr{P}$.

$$p^\star = \arg\max_p p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

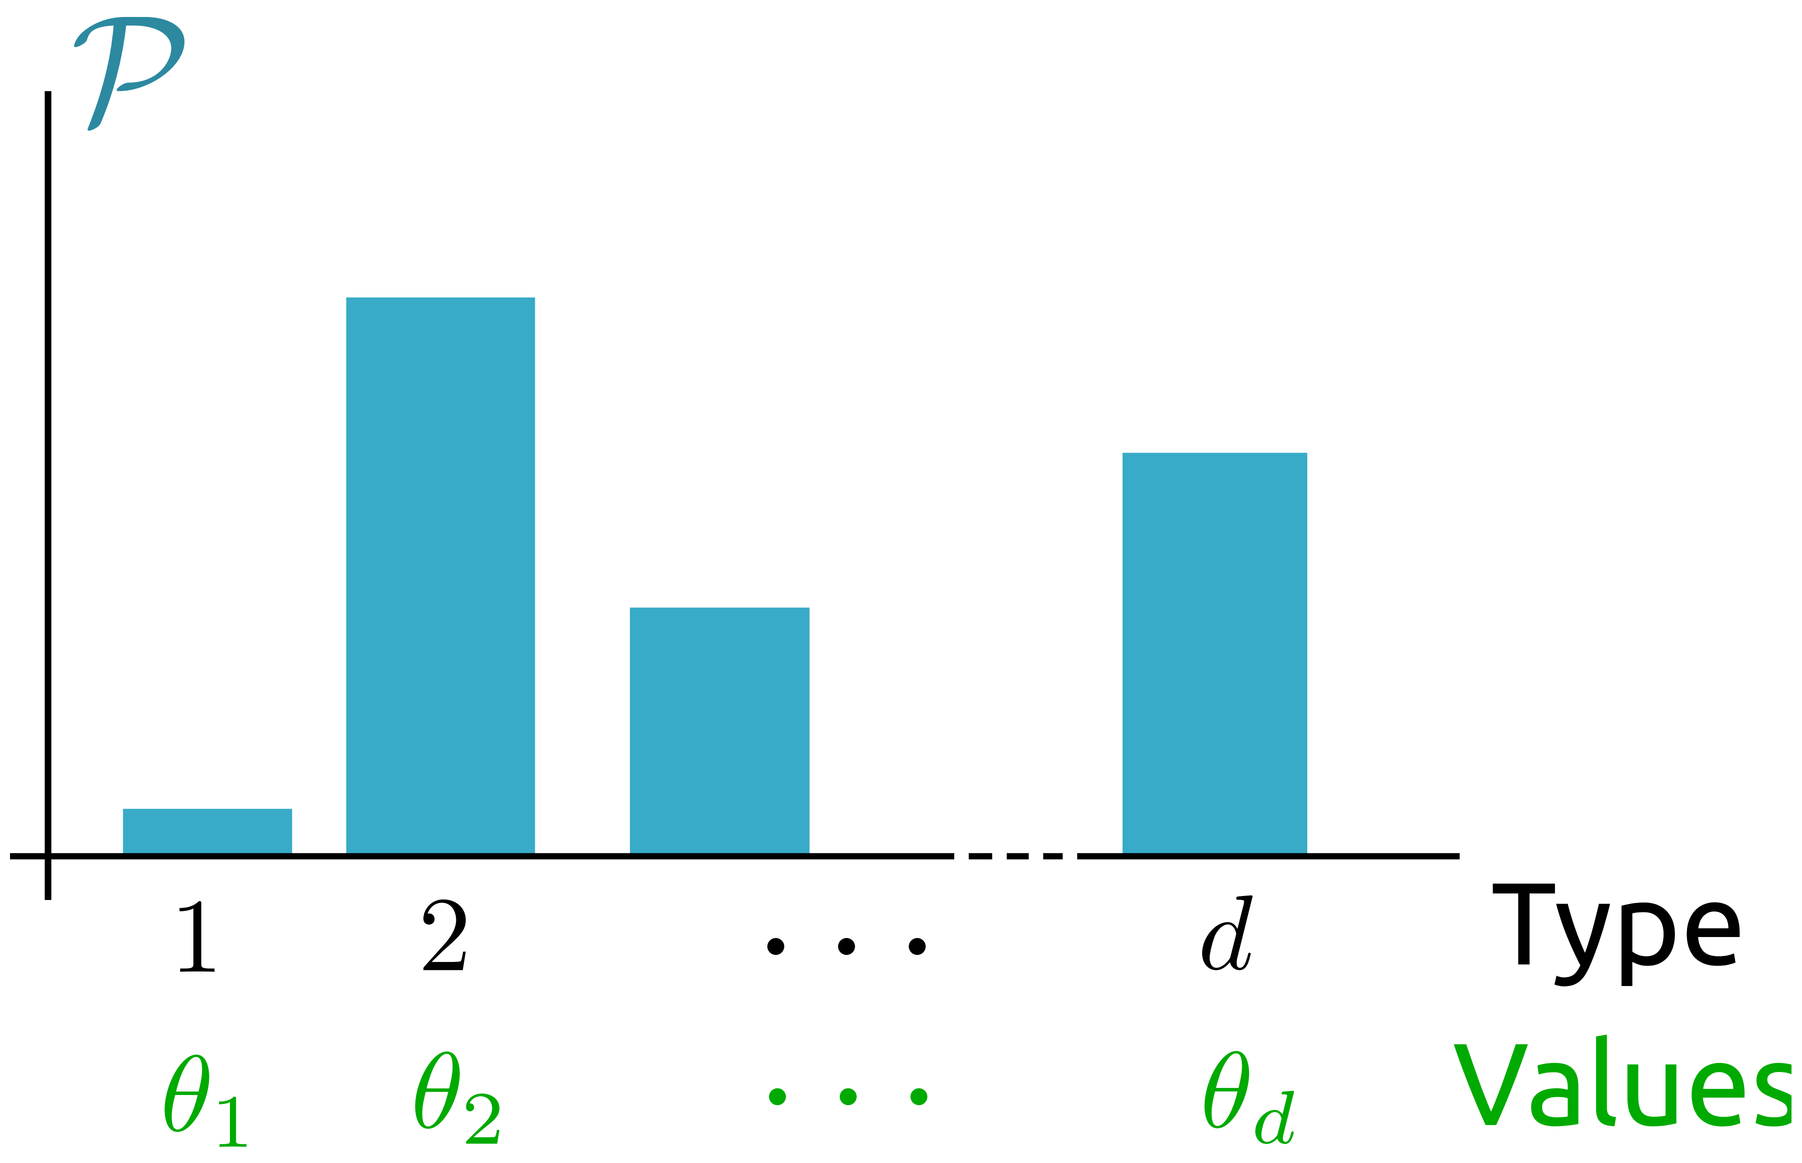

2. A buyer may know their type i, but not their value θ_i.
   - Due to uncertainty about their value, buyers may not be willing to buy an item except at a low price.

PRIOR WORK ON SOCIAL LEARNING FROM REVIEWS

Chamley, 2004

Bose et al., 2006

Acemoglu et al., 2017

Crapis et al., 2017

Besbes and Scarsini, 2018

Ifrach et al., 2019

Boursier et al., 2020

Han and Anderson, 2020

OUTLINE

1. Introduction
2. **Problem set up & challenges**
3. Algorithm & theoretical results

FRAMEWORK FOR ONLINE LEARNING & PRICING

FRAMEWORK FOR ONLINE LEARNING & PRICING

- Proceeds over a sequence of rounds. On any given round there are reviews from previous rounds.

FRAMEWORK FOR ONLINE LEARNING & PRICING

- Proceeds over a sequence of rounds. On any given round there are reviews from previous rounds.

- On each round t:

FRAMEWORK FOR ONLINE LEARNING & PRICING

- Proceeds over a sequence of rounds. On any given round there are reviews from previous rounds.

- On each round t:
  - Seller chooses a price p_t (based on past reviews).

FRAMEWORK FOR ONLINE LEARNING & PRICING

- Proceeds over a sequence of rounds. On any given round there are reviews from previous rounds.

- On each round t:
  - Seller chooses a price p_t (based on past reviews).
  - A **new** buyer of type $i_t \sim \mathcal{P}$ arrives. She buys the item if she has reason to believe that her value is larger than the price p_t based on past reviews.

FRAMEWORK FOR ONLINE LEARNING & PRICING

- Proceeds over a sequence of rounds. On any given round there are reviews from previous rounds.

- On each round t:
  - Seller chooses a price p_t (based on past reviews).
  - A **new** buyer of type $i_t \sim \mathscr{P}$ arrives. She buys the item if she has reason to believe that her value is larger than the price p_t based on past reviews.
  - **If buyer buys**, she leaves a review based on her experience of using the item. **Otherwise, no review**.

FRAMEWORK FOR ONLINE LEARNING & PRICING

- Proceeds over a sequence of rounds. On any given round there are reviews from previous rounds

- On each round t:
  - Seller chooses a price p_t (based on past reviews).
  - A new buyer of type $i \sim \mathscr{P}$ arrives. She buys the item if **she has reason to believe that her value is larger than the price p_t** based on past reviews.
  - If buyer buys, she leaves a **review based on her experience** of using the item. Otherwise, no review.

LEAVING A REVIEW BASED ON EXPERIENCE

LEAVING A REVIEW BASED ON EXPERIENCE

- A buyer has ***ex-ante value*** θ_i based on their type .

LEAVING A REVIEW BASED ON EXPERIENCE

- A buyer has ***ex-ante value*** θ_i based on their type .

- Buyer's ***ex-post value*** v is drawn from distribution $\mathscr{D}_i$ such that $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.

LEAVING A REVIEW BASED ON EXPERIENCE

- A buyer has ***ex-ante value*** θ_i based on their type .

- Buyer's ***ex-post value*** v is drawn from distribution $\mathscr{D}_i$ such that $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.
  - Ex-post value is the actual experience of the customer.
  - Depends on exogenous factors that cannot be known at time of purchase.

LEAVING A REVIEW BASED ON EXPERIENCE

- A buyer has ***ex-ante value*** θ_i based on their type .

- Buyer's ***ex-post value*** v is drawn from distribution $\mathscr{D}_i$ such that $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.
  - Ex-post value is the actual experience of the customer.
  - Depends on exogenous factors that cannot be known at time of purchase.
    - E.g manufacturing defects, mismatch between description and product

LEAVING A REVIEW BASED ON EXPERIENCE

- A buyer has ***ex-ante value*** θ_i based on their type .

- Buyer's ***ex-post value*** v is drawn from distribution $\mathscr{D}_i$ such that $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.
  - Ex-post value is the actual experience of the customer.
  - Depends on exogenous factors that cannot be known at time of purchase.
    - E.g manufacturing defects, mismatch between description and product

- **Review: If the buyers buy**, they reveal their type and ex-post value (i, v) to the seller and future buyers. **Otherwise, no review.**

IS VALUE LARGER THAN PRICE BASED ON PAST REVIEWS?

IS VALUE LARGER THAN PRICE BASED ON PAST REVIEWS?

1. Buyers may not be willing to pay a high price if they are uncertain about their (ex-ante) value.

IS VALUE LARGER THAN PRICE BASED ON PAST REVIEWS?

1. Buyers may not be willing to pay a high price if they are uncertain about their (ex-ante) value.

2. **But** buyers cannot be overly conservative.

   - E.g: “I will only pay $0.01 since I do not know my value exactly”.

IS VALUE LARGER THAN PRICE BASED ON PAST REVIEWS?

1. Buyers may not be willing to pay a high price if they are uncertain about their (ex-ante) value.

2. **But** buyers cannot be overly conservative.
   - E.g: “I will only pay $0.01 since I do not know my value exactly”.
   - Revenue maximization would be hopeless with ultraconservative customers.

η-RISKY CUSTOMERS

η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold τ_t. Purchases if $p_t \le \tau_t$.

η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold τ_t. Purchases if $p_t \leq \tau_t$.
    - This threshold is a function of past reviews.

η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold τ_t. Purchases if $p_t \leq \tau_t$.
    - This threshold is a function of past reviews.

- τ_t has to be at least a η-lower confidence bound on the value.

η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold τ_t. Purchases if $p_t \leq \tau_t$.
  - This threshold is a function of past reviews.

- τ_t has to be at least a η-lower confidence bound on the value.
  If $\Phi_{i,t}$ are the reviews from past customers of type i, then,

$$\tau_t \geq \frac{1}{|\Phi_{i,t}|} \sum_{v \in \Phi_{i,t}} v - \sqrt{\frac{1}{|\Phi_{i,t}|} \log\left(\frac{t}{\eta}\right)}$$

η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold τ_t. Purchases if $p_t \leq \tau_t$.

  - This threshold is a function of past reviews.

- τ_t has to be at least a η-lower confidence bound on the value.

  If $\Phi_{i,t}$ are the reviews from past customers of type i, then,

$$\tau_t \geq \frac{1}{|\Phi_{i,t}|} \sum_{v \in \Phi_{i,t}} v - \sqrt{\frac{1}{|\Phi_{i,t}|} \log\left(\frac{t}{\eta}\right)}$$

  - **Bounded pessimism:** Customer is willing to take at least a small risk. She may over-estimate her value (i.e $\tau_t > \theta_i$) with some small probability η.

- Regret R_T after T rounds:

$$R_T = T\,\mathrm{rev}(p^\star) - \sum_{t=1}^{T} p_t \cdot 1(\text{purchase on round } t)$$

Optimal price when sellers know $\mathscr{P}$ and customers know their values θ_i.

OUTLINE

1. Introduction

2. Problem set up & challenges

3. Algorithm & theoretical results

OVERVIEW OF ALGORITHM & RESULTS

OVERVIEW OF ALGORITHM & RESULTS

- ***Algorithm:***
  - Choose low prices early, and increase them gradually.

OVERVIEW OF ALGORITHM & RESULTS

- ***Algorithm:***
  - Choose low prices early, and increase them gradually.
- ***Theoretical Results:***
  - Upper bound: $\tilde{\mathcal{O}}\left(d^{1/3}T^{2/3}\right)$ worst case regret, but $\tilde{\mathcal{O}}\left(T^{1/2}\right)$ regret when all types appear frequently.
  - Matching lower bounds.

CHALLENGES: PRICING AFFECTS LEARNING

CHALLENGES: PRICING AFFECTS LEARNING

- Seller wishes to set high prices on each round (to maximize current revenue).

CHALLENGES: PRICING AFFECTS LEARNING

- Seller wishes to set high prices on each round (to maximize current revenue).
- But higher prices $\implies$ no purchase $\implies$ no review.

CHALLENGES: PRICING AFFECTS LEARNING

- Seller wishes to set high prices on each round (to maximize current revenue).
- But higher prices $\implies$ no purchase $\implies$ no review.
    1. *Seller learning:* Seller cannot gauge demand for the product.
    2. *Buyer learning:* Future buyers cannot estimate their value.

CHALLENGE 1: PRICING VS SELLER LEARNING

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

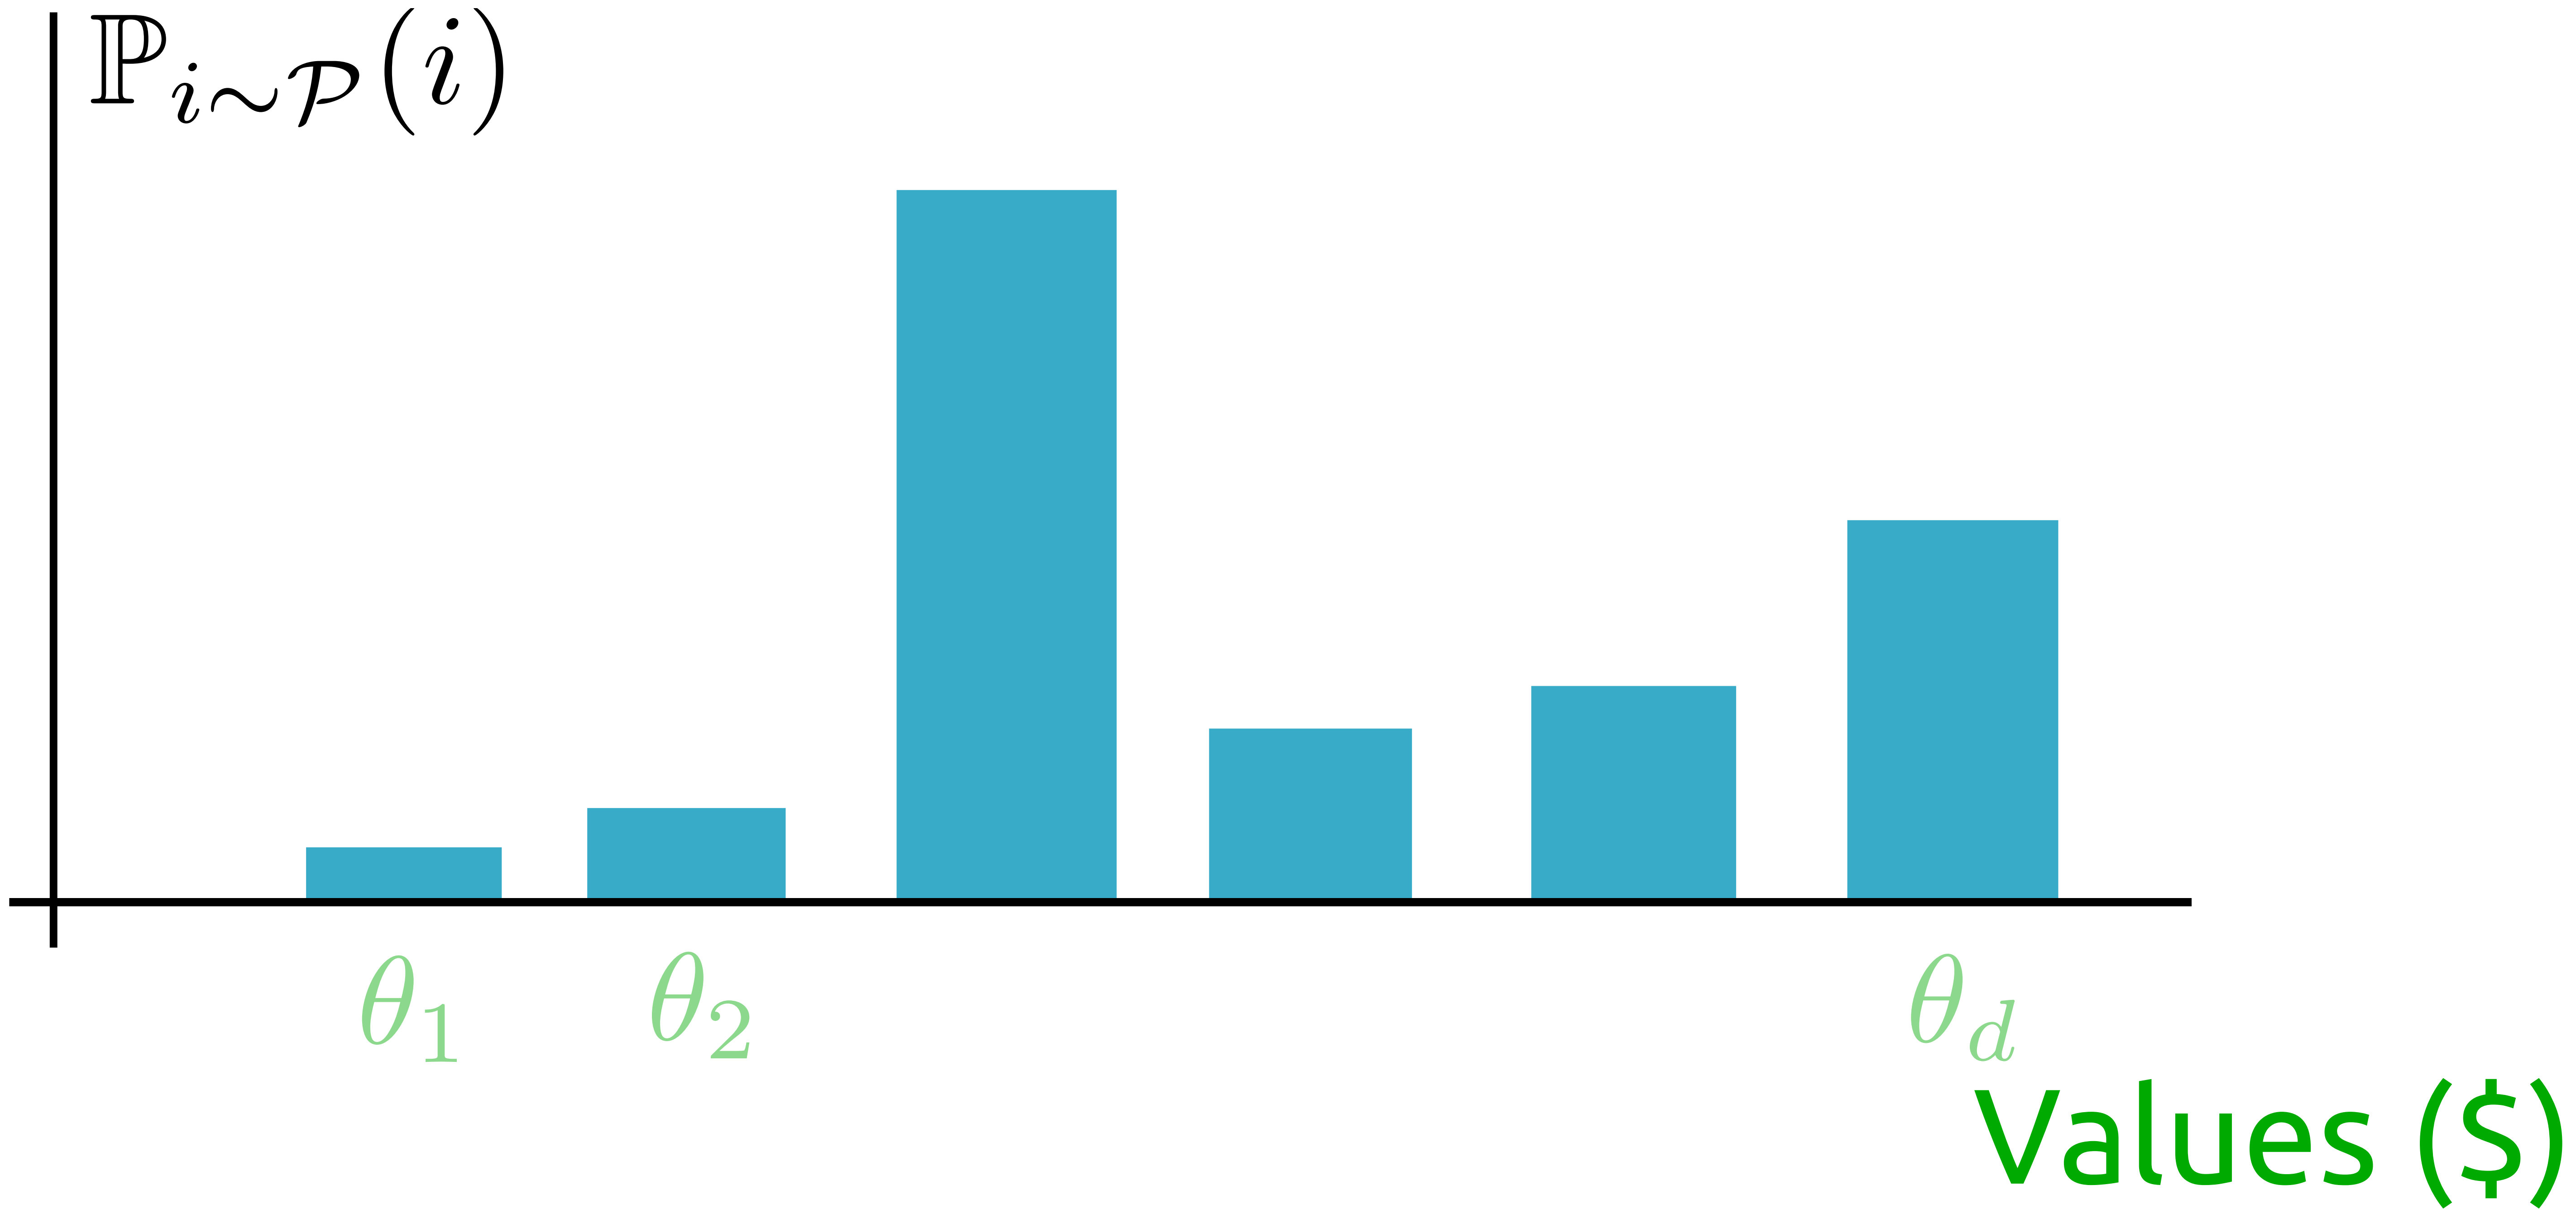

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

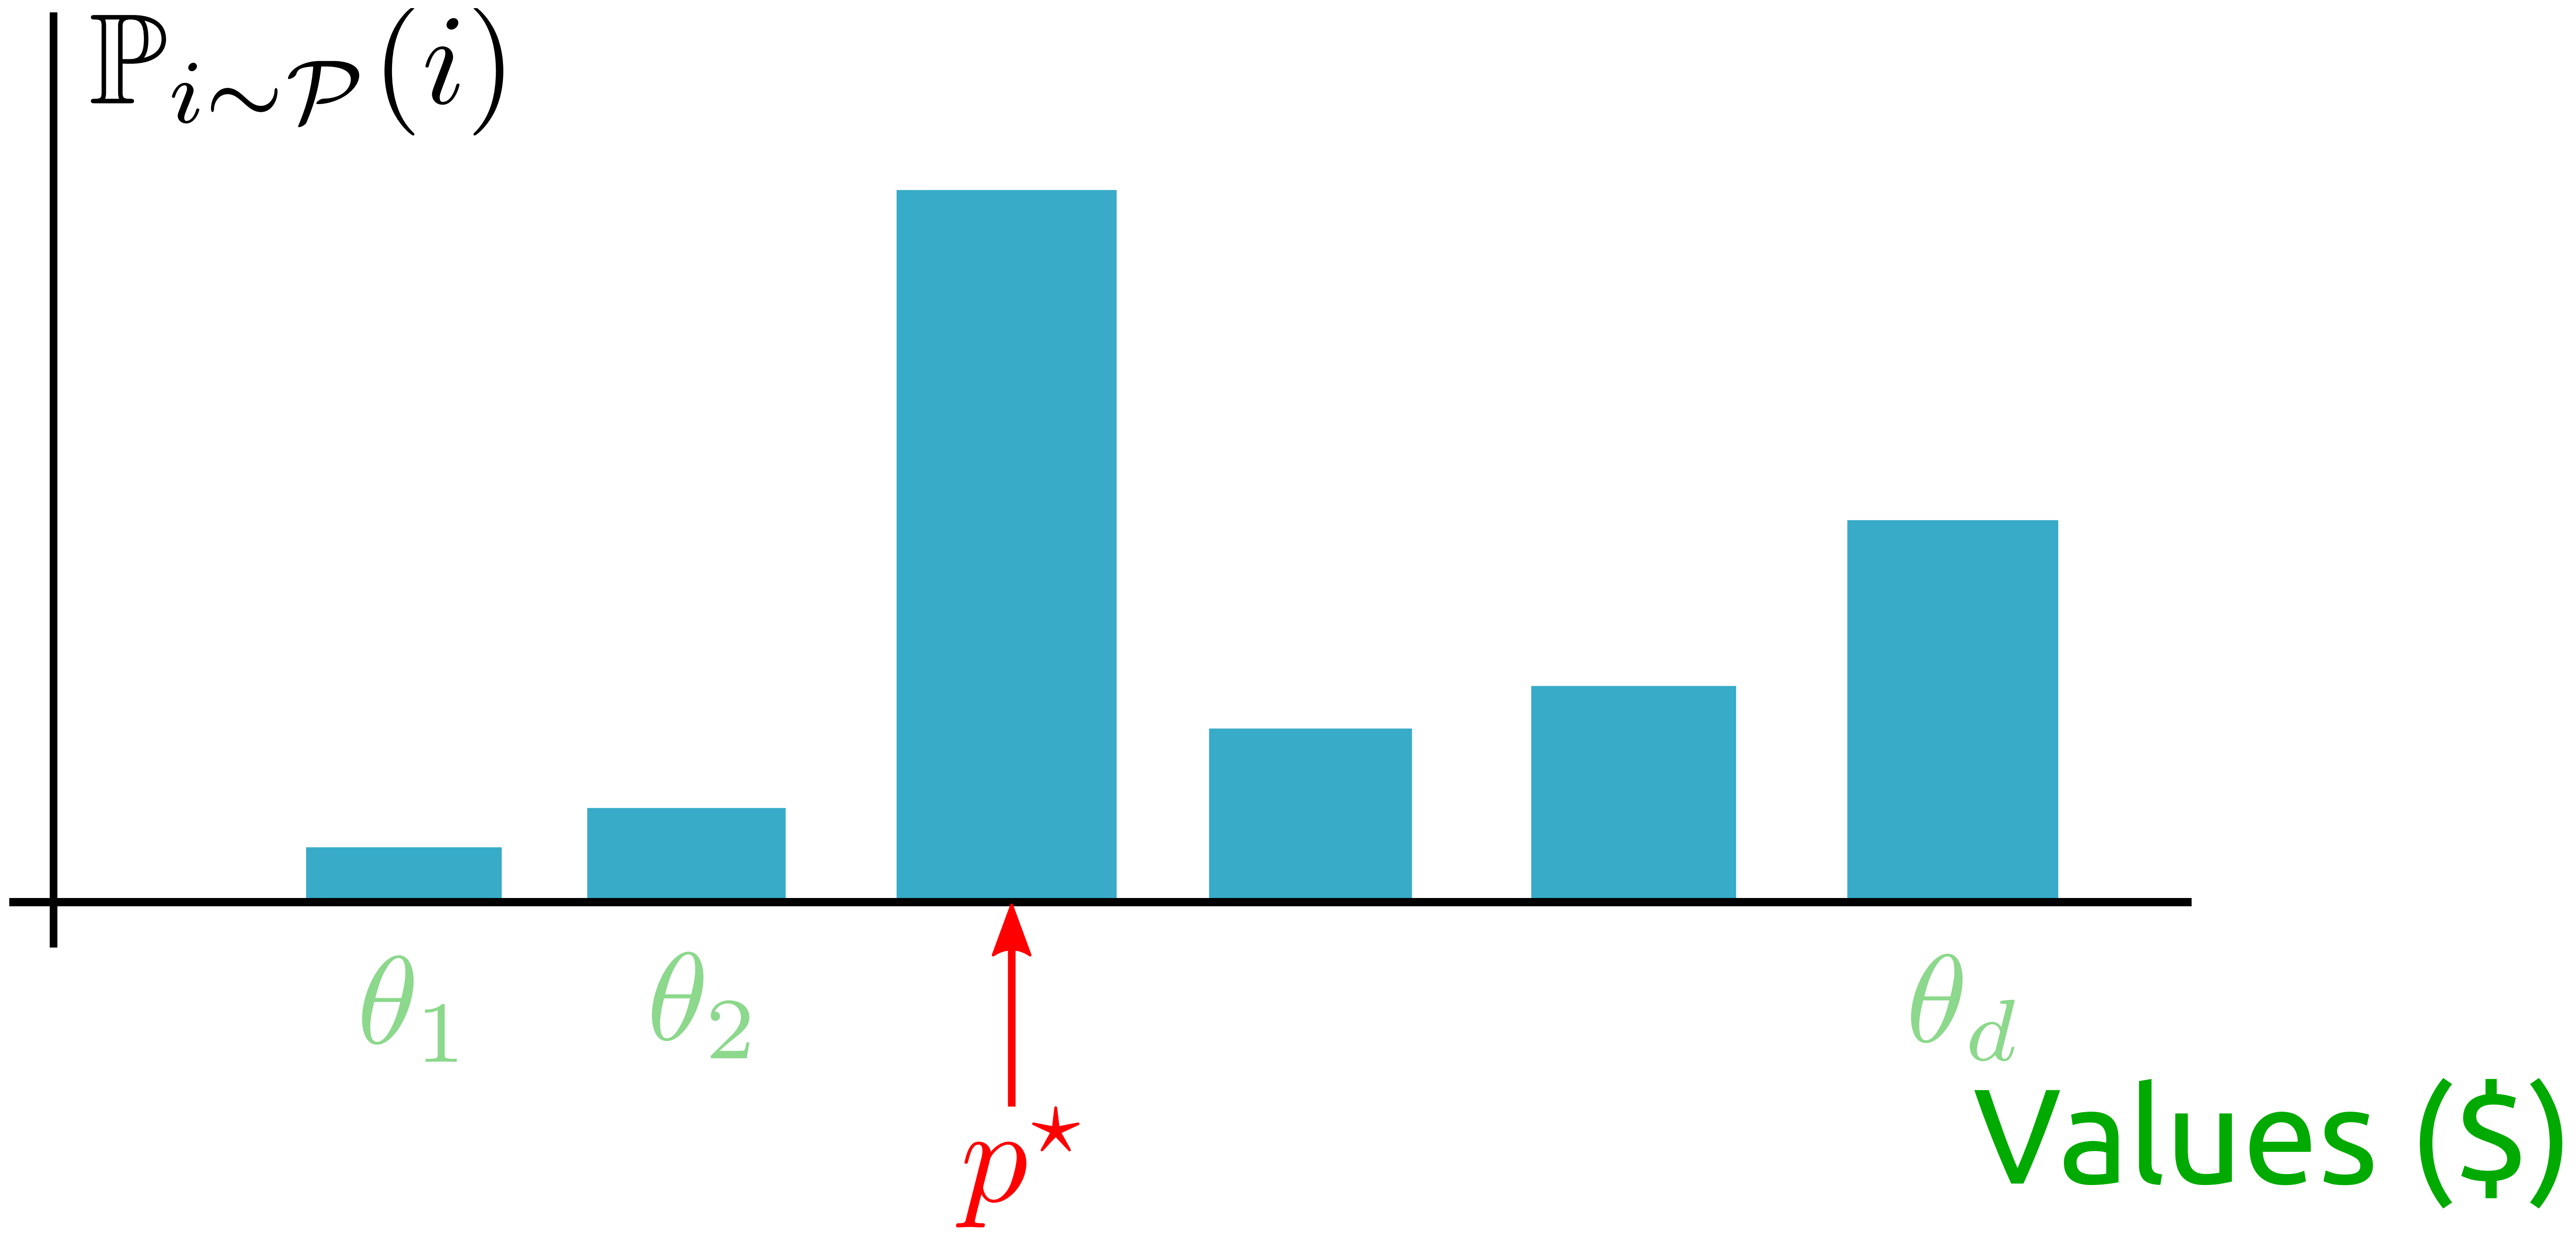

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

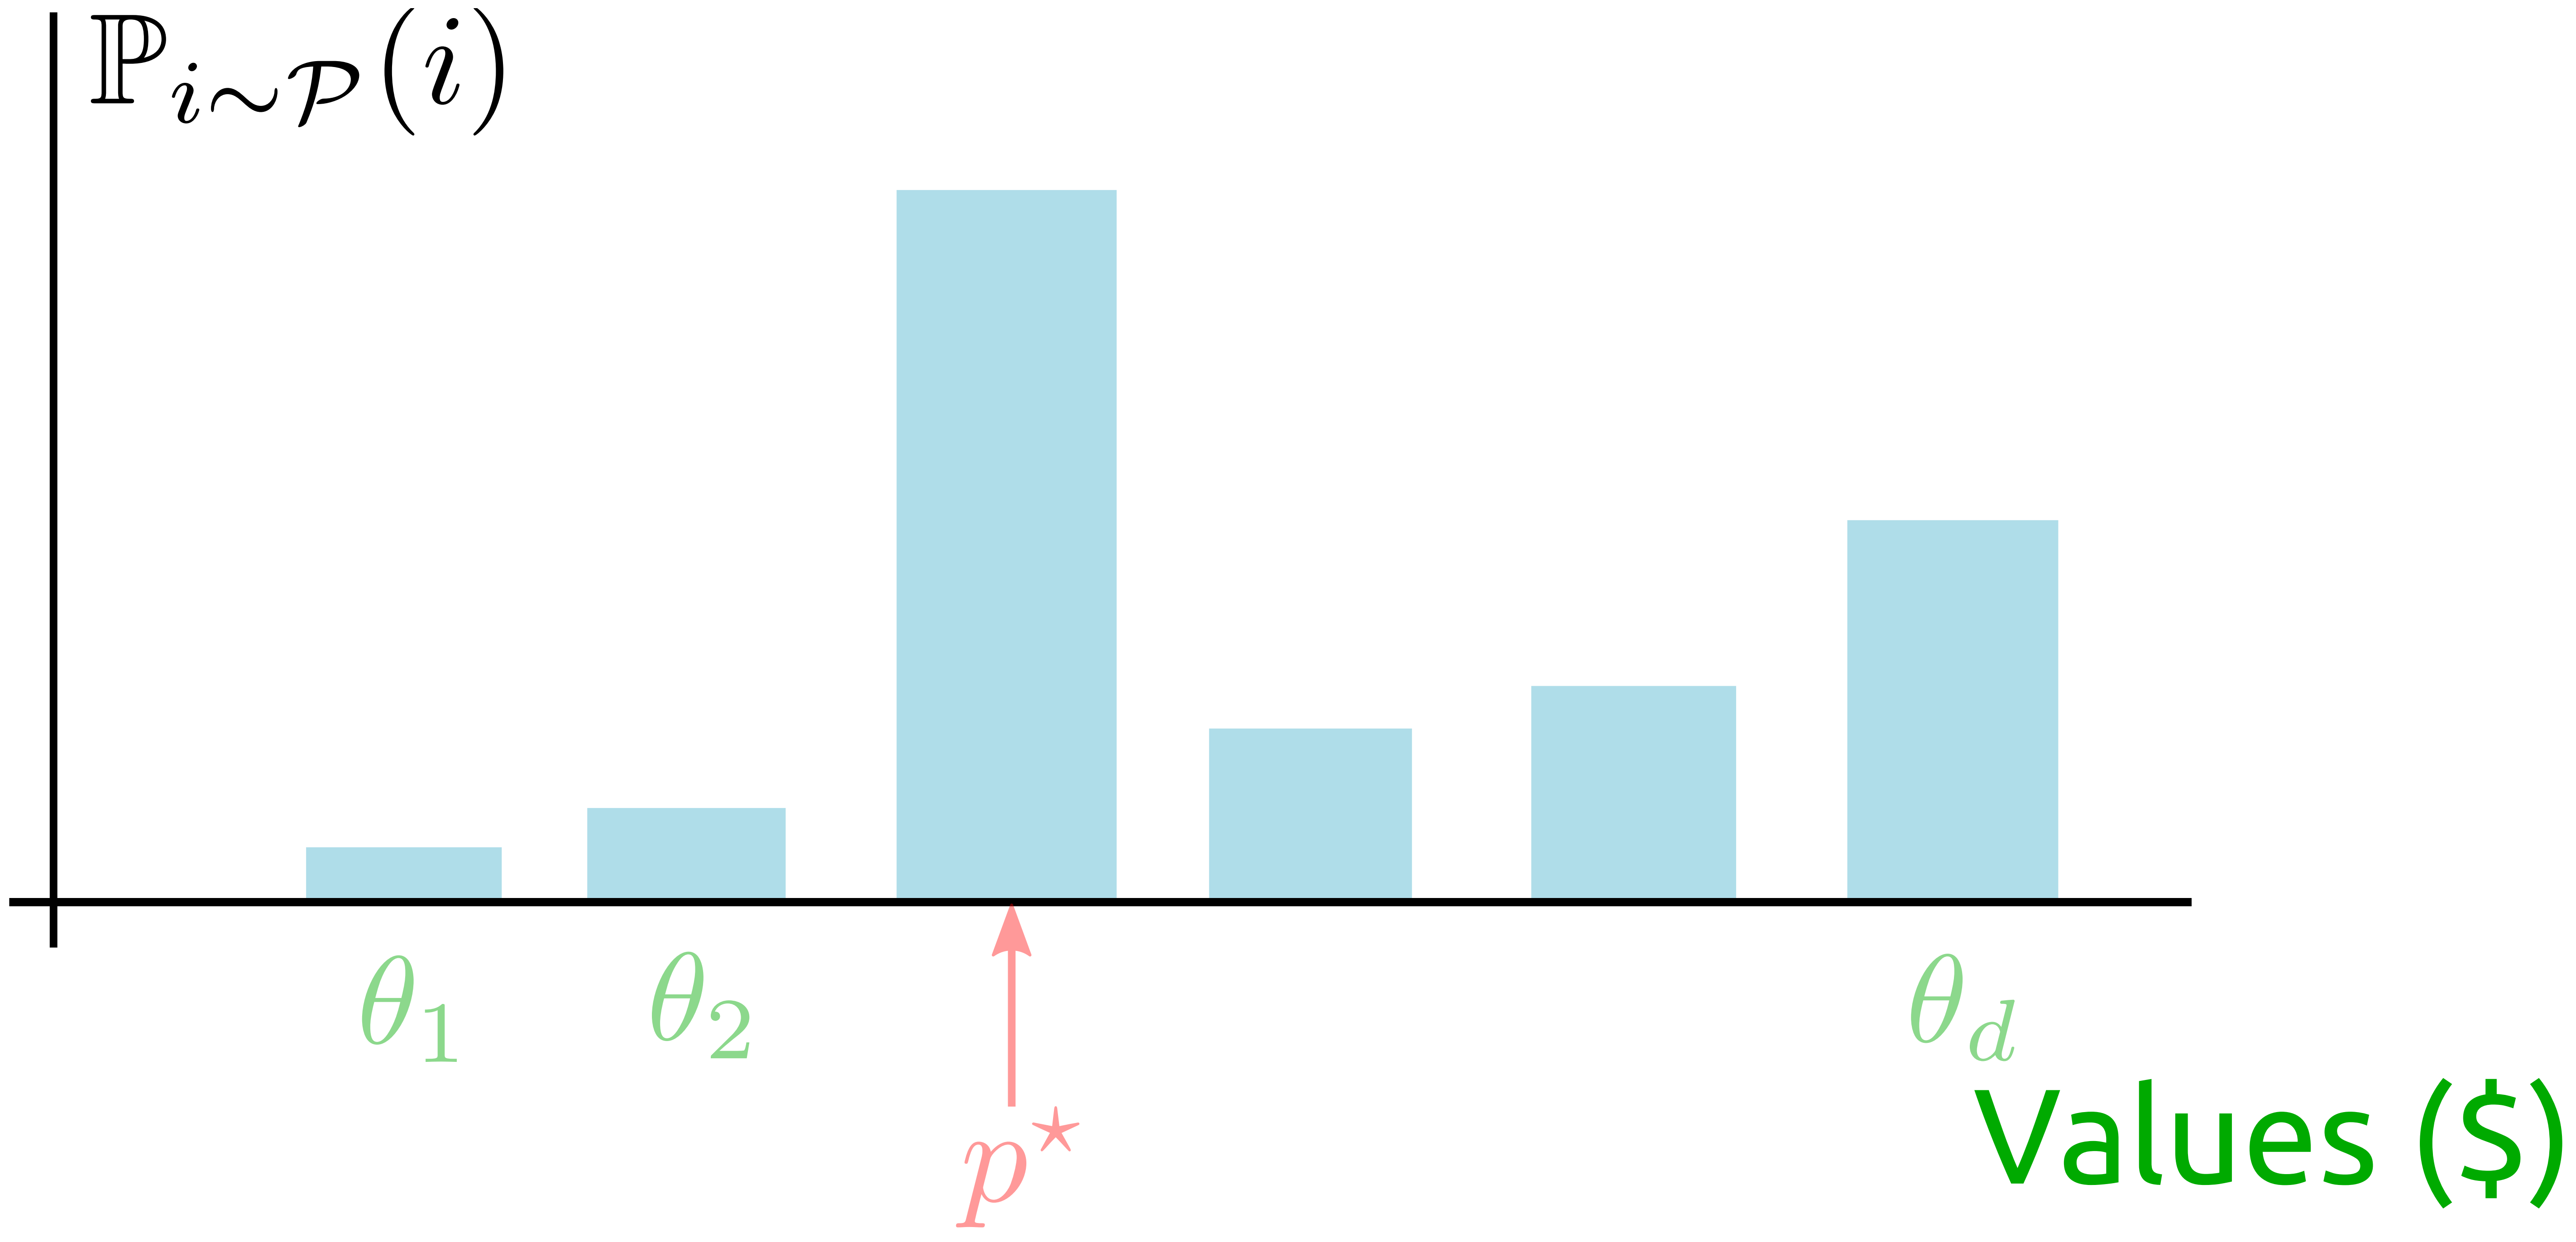

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

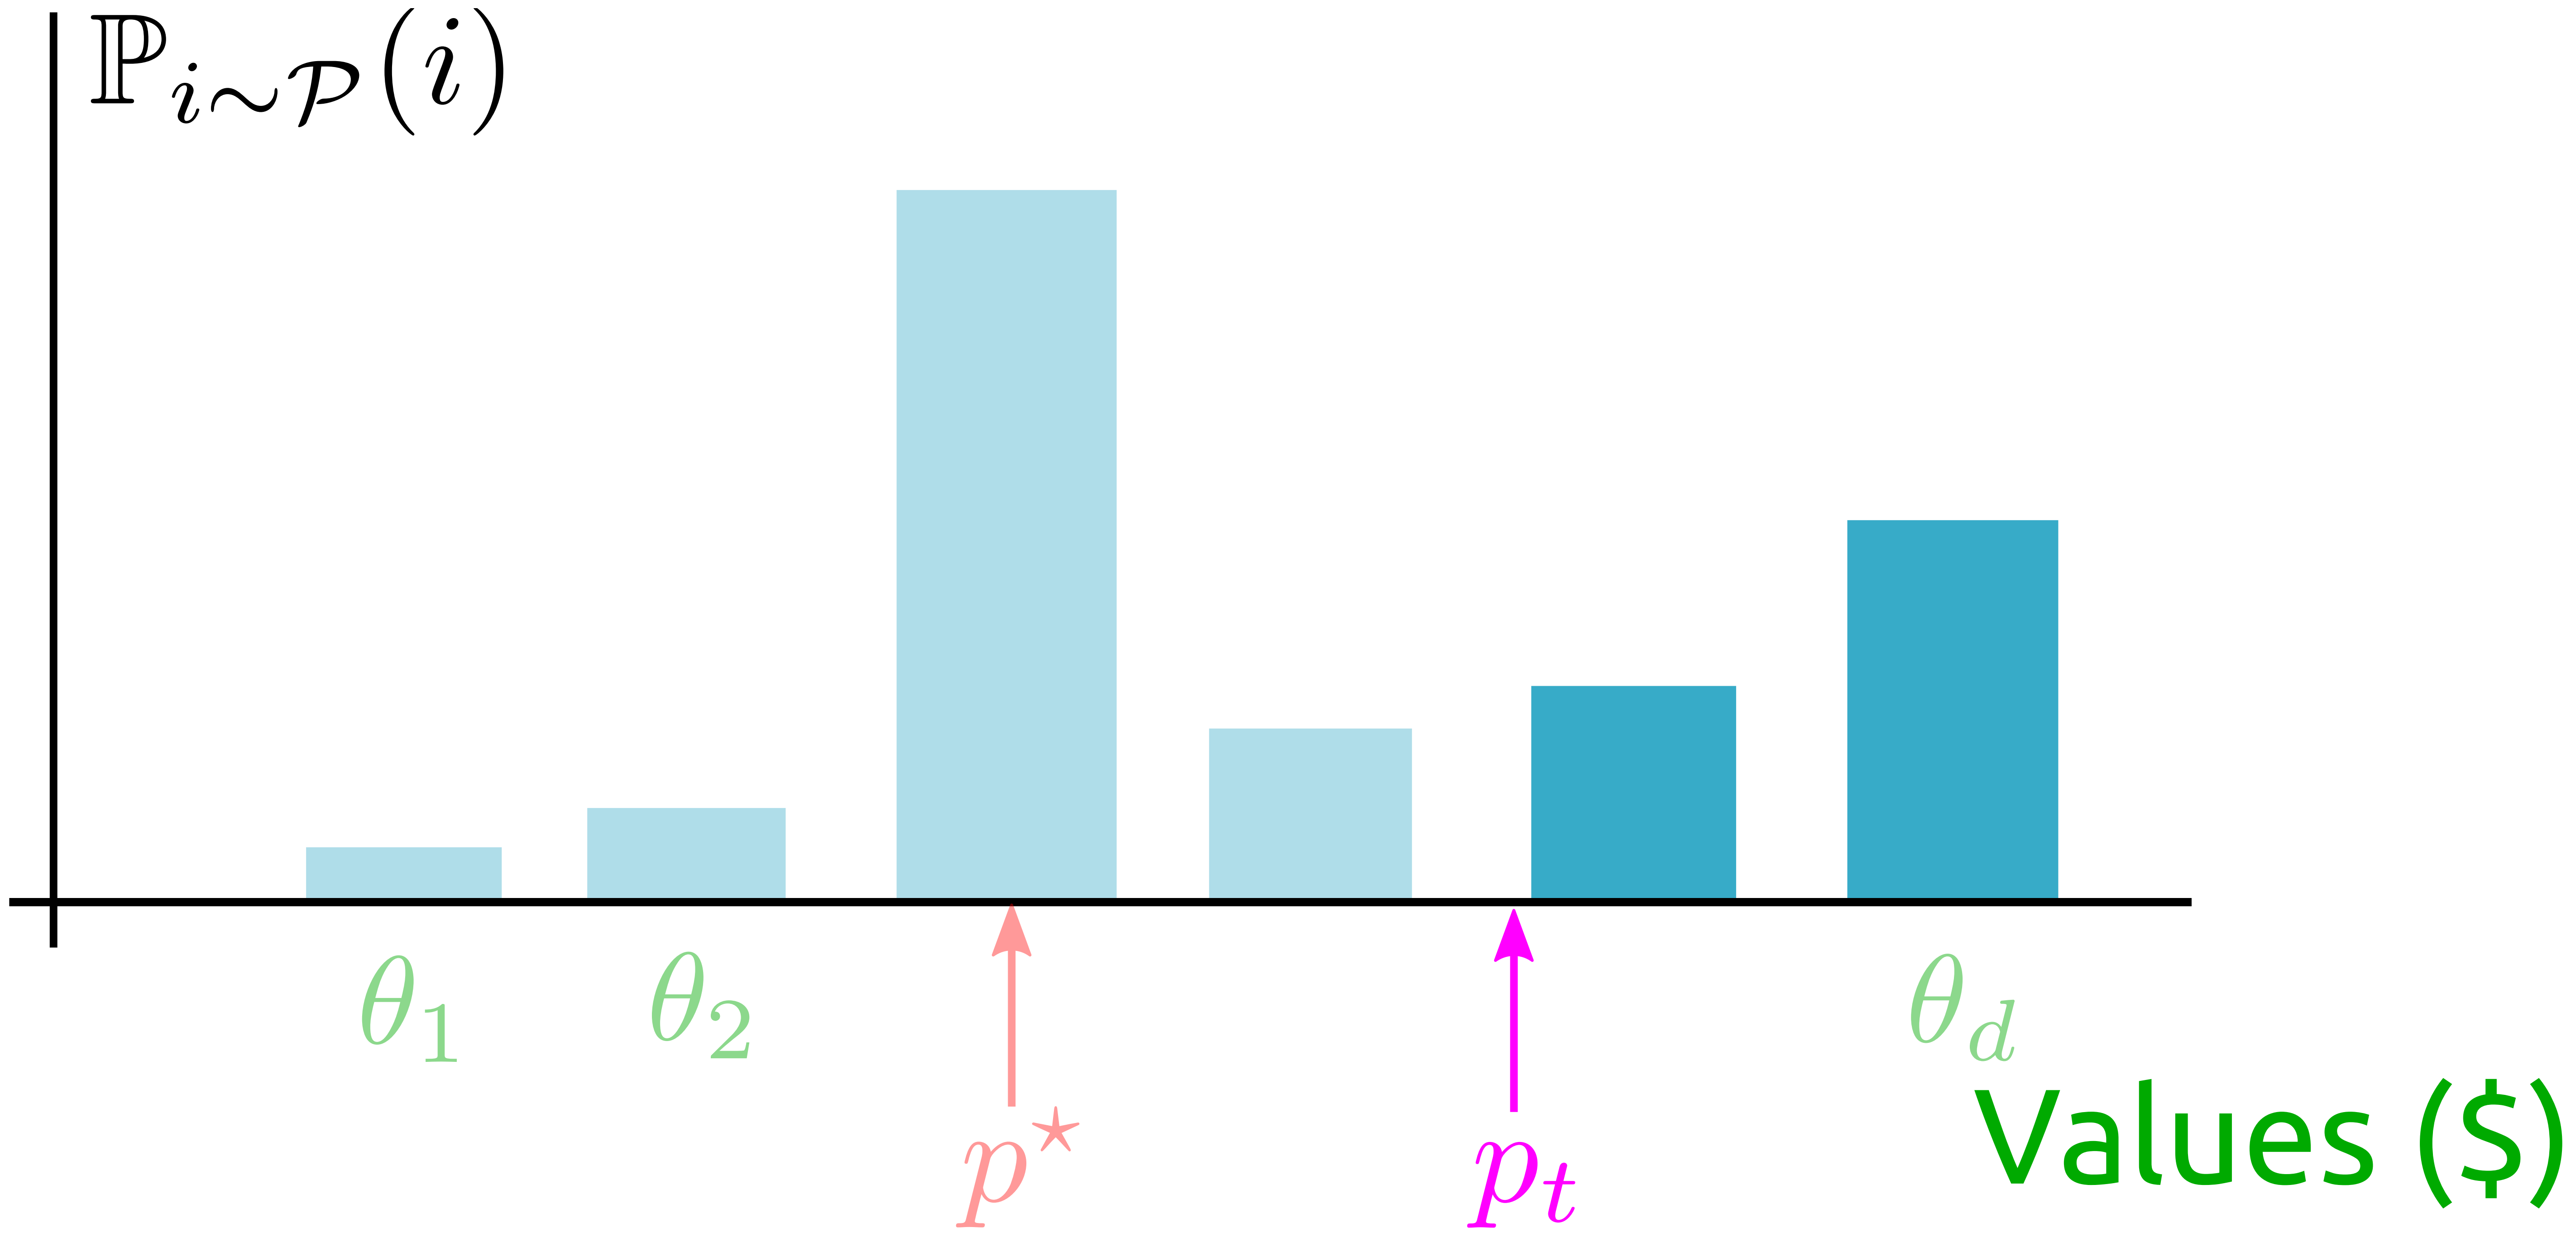

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

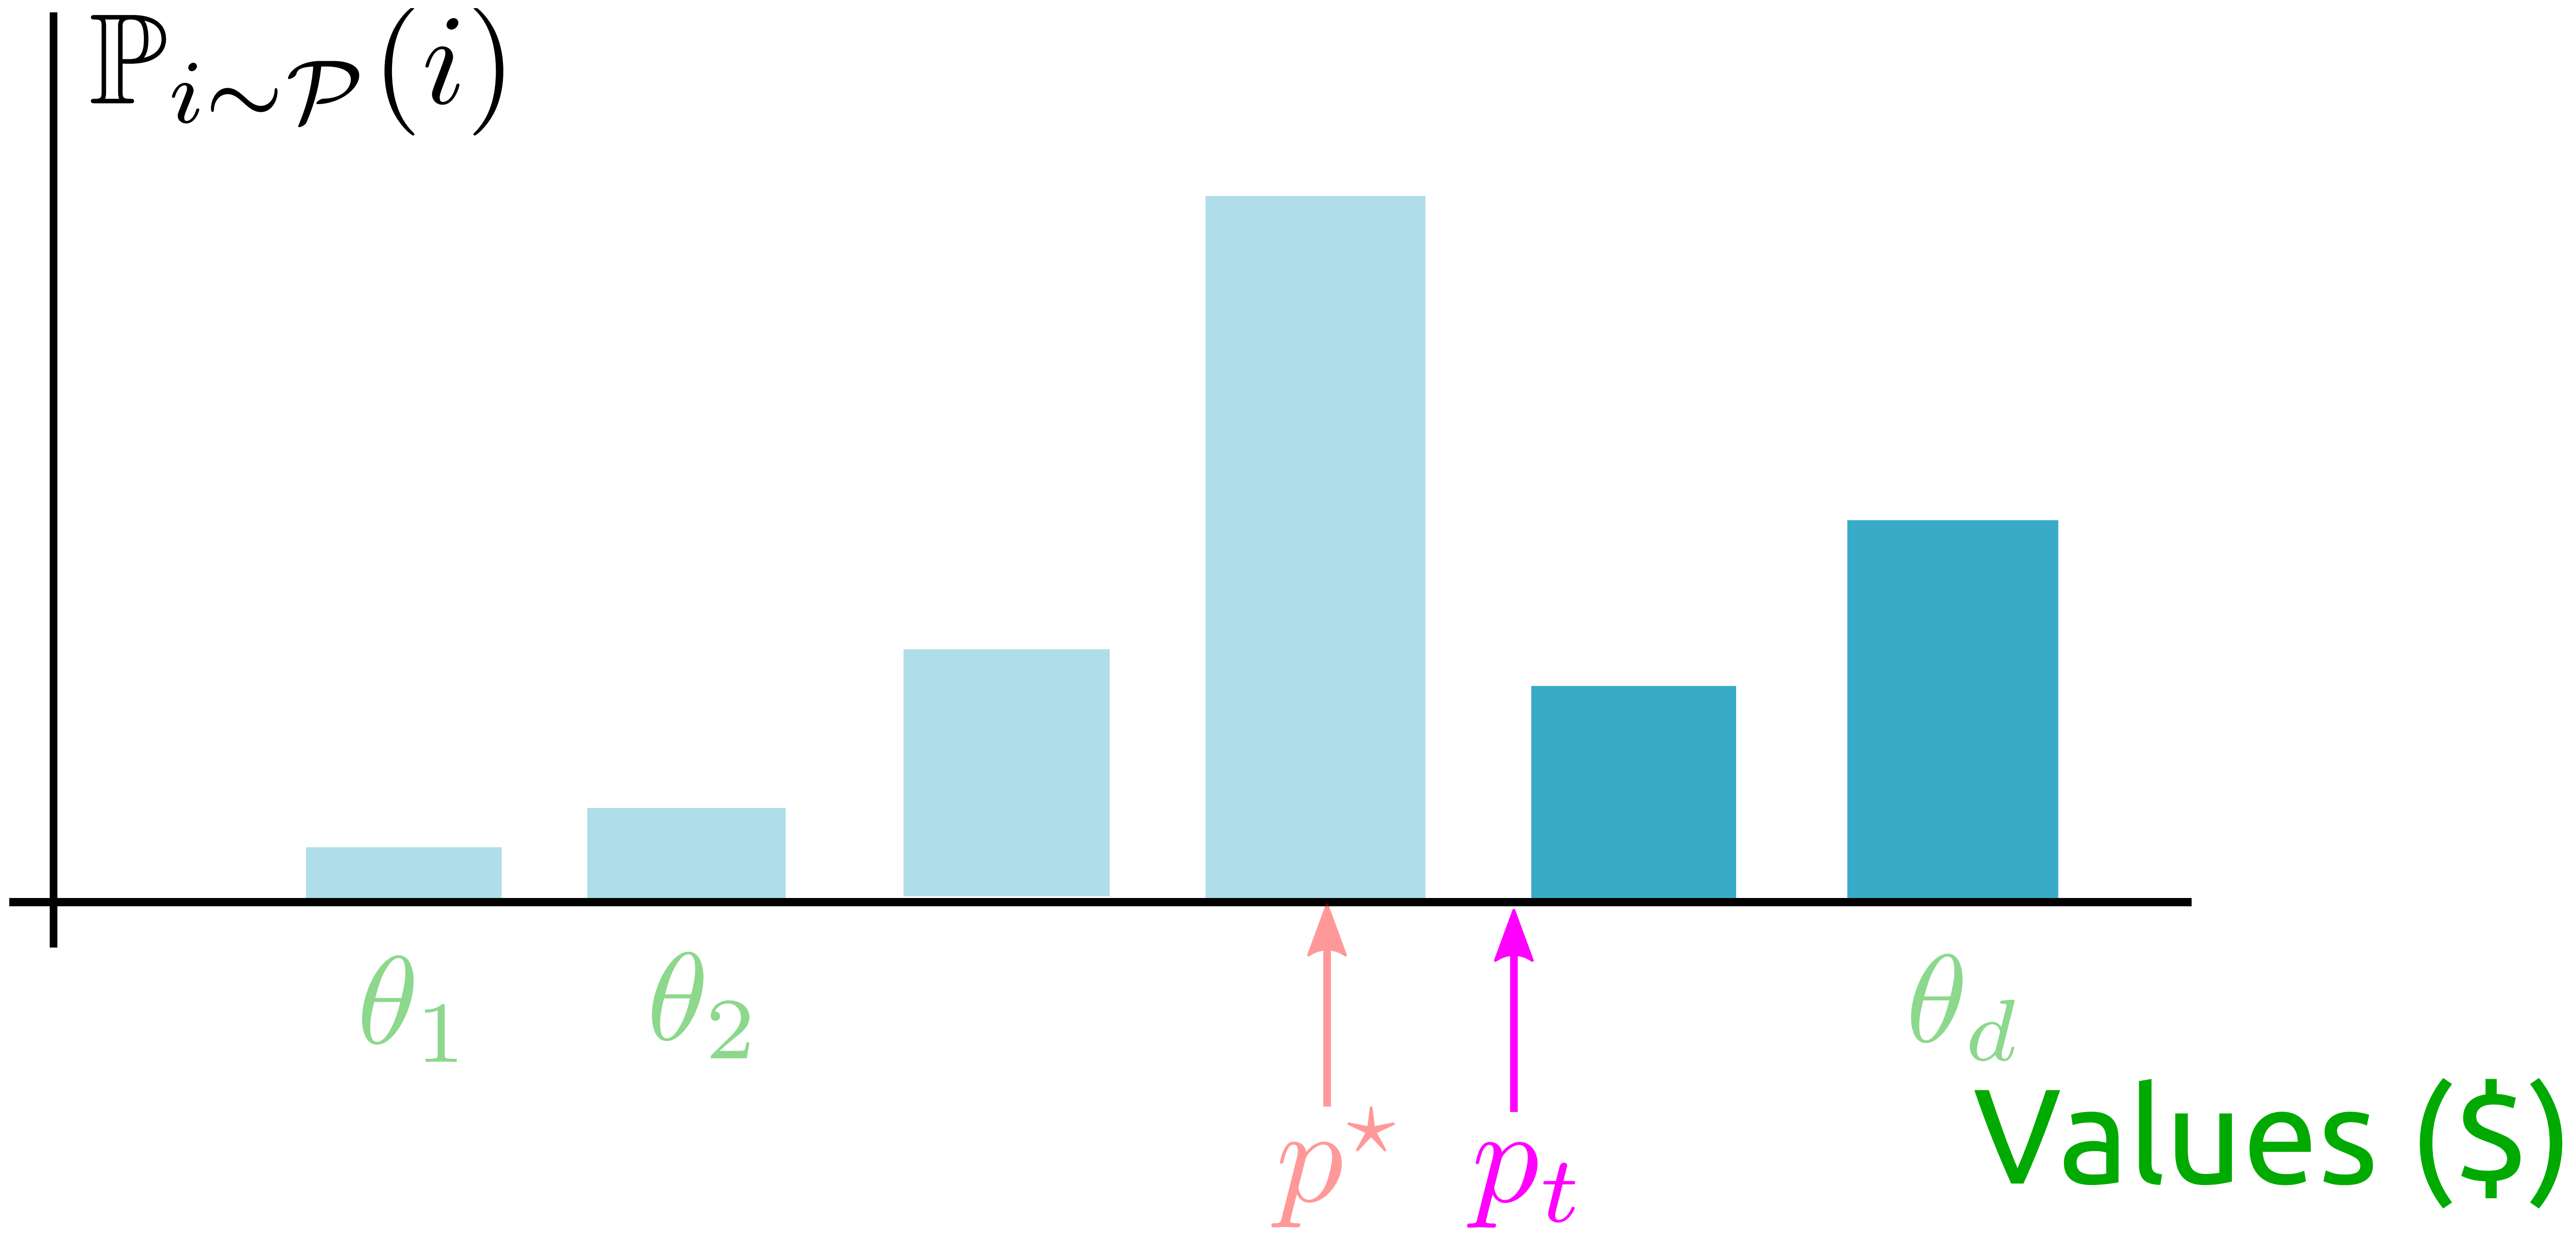

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

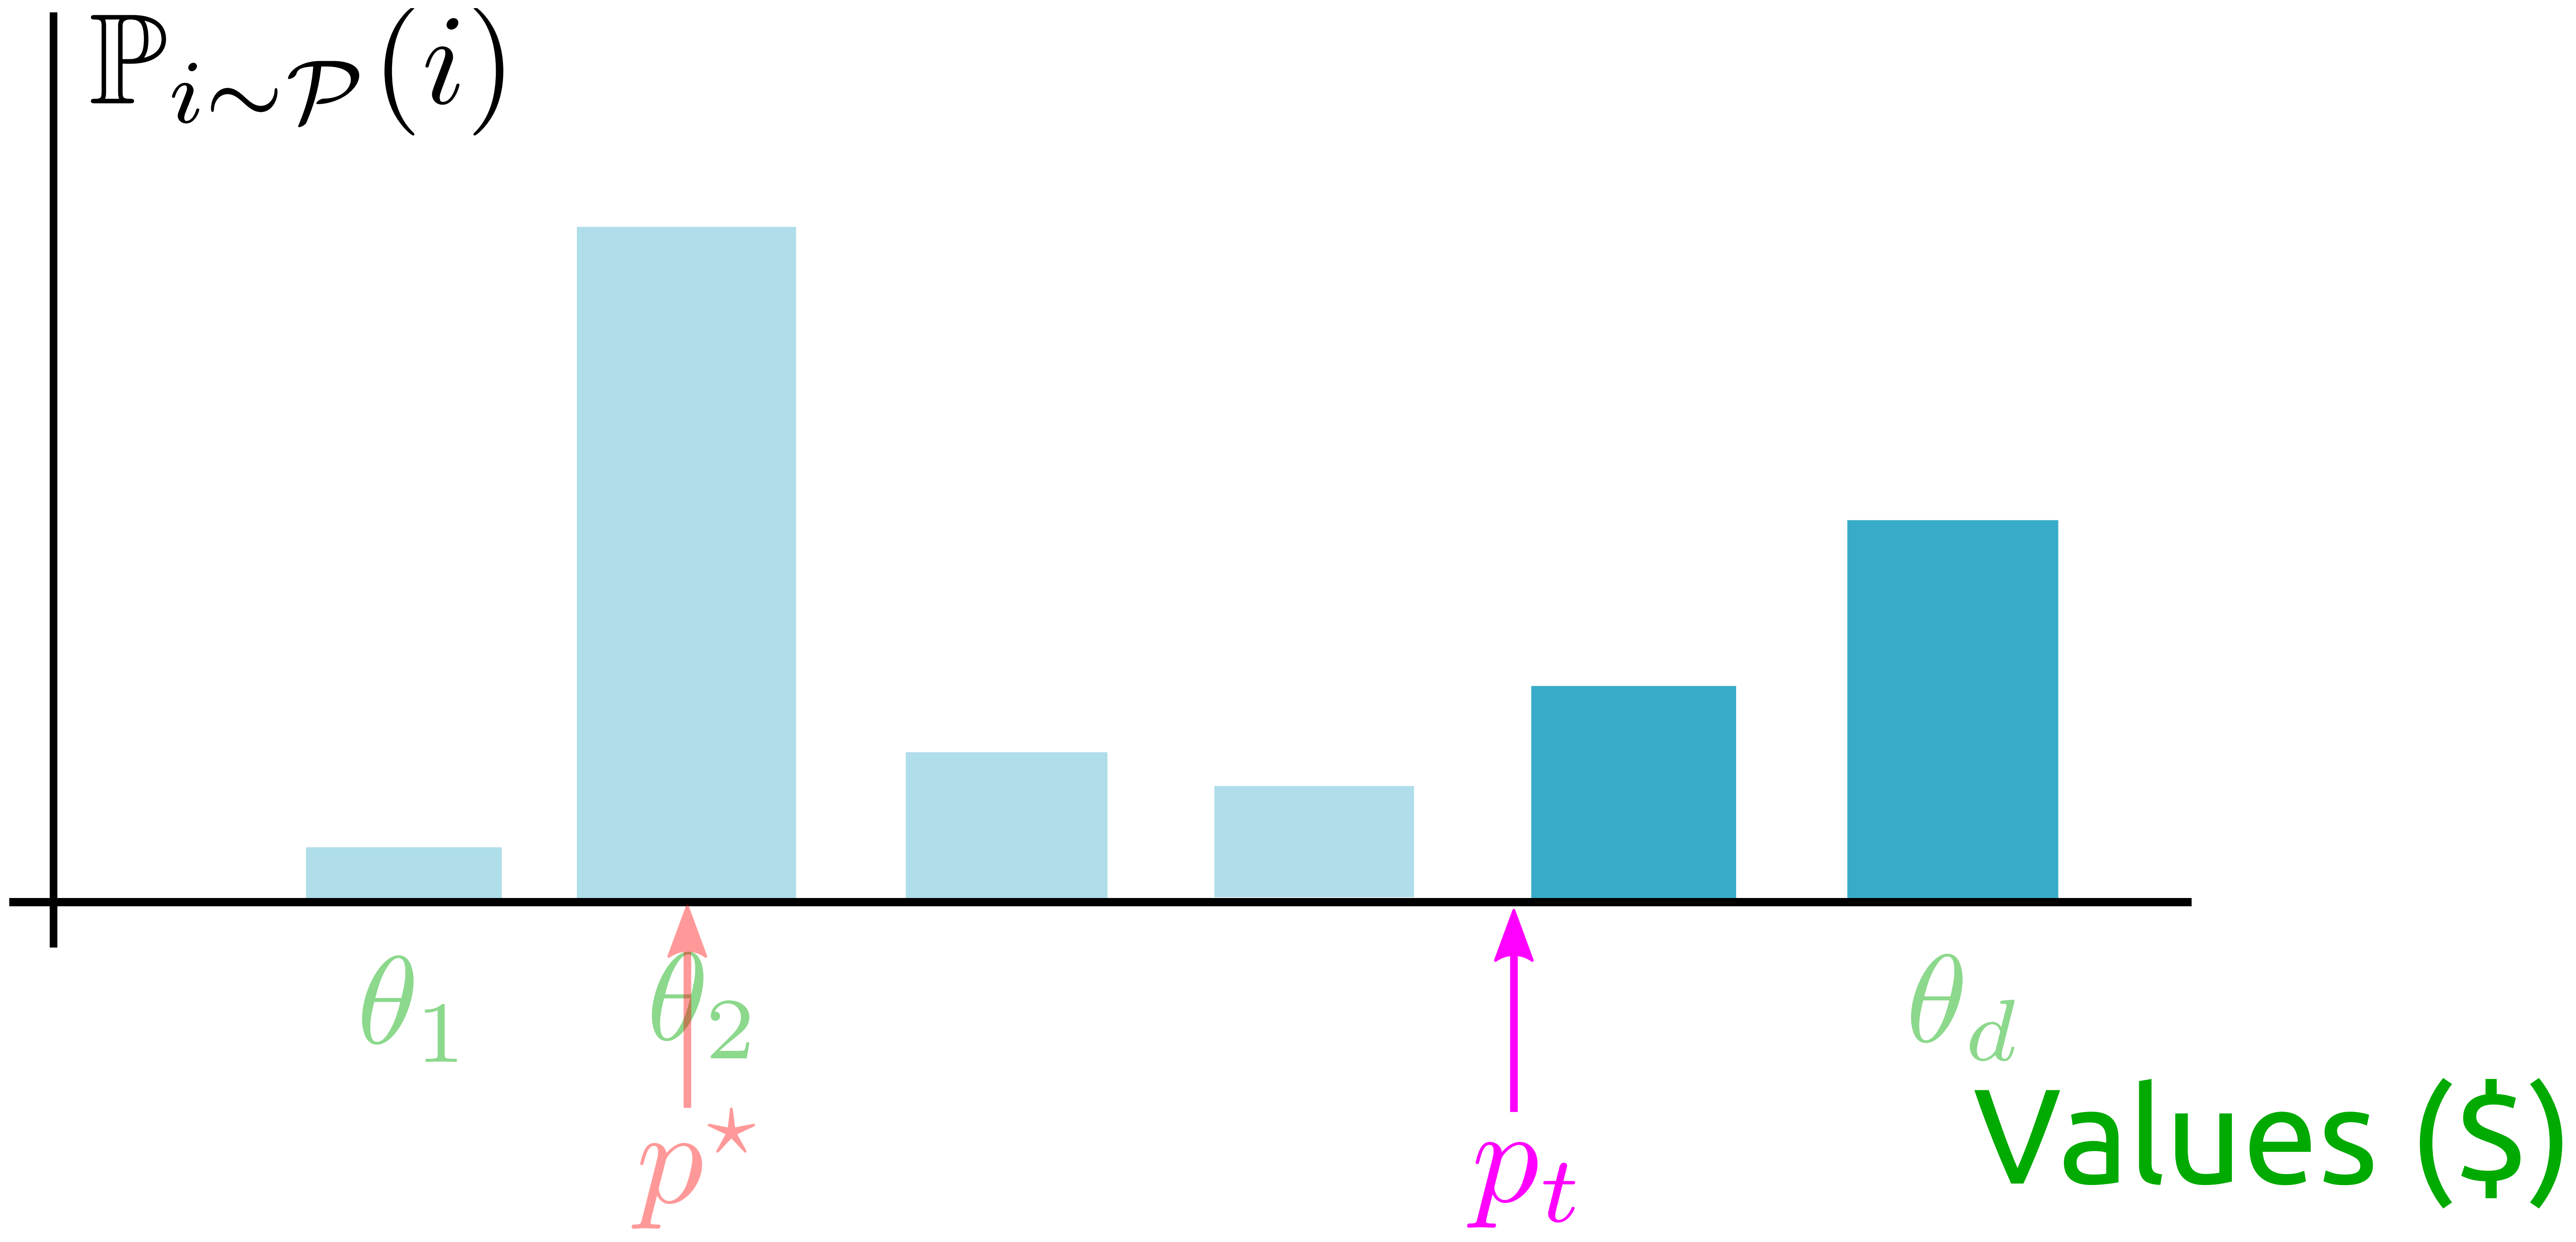

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

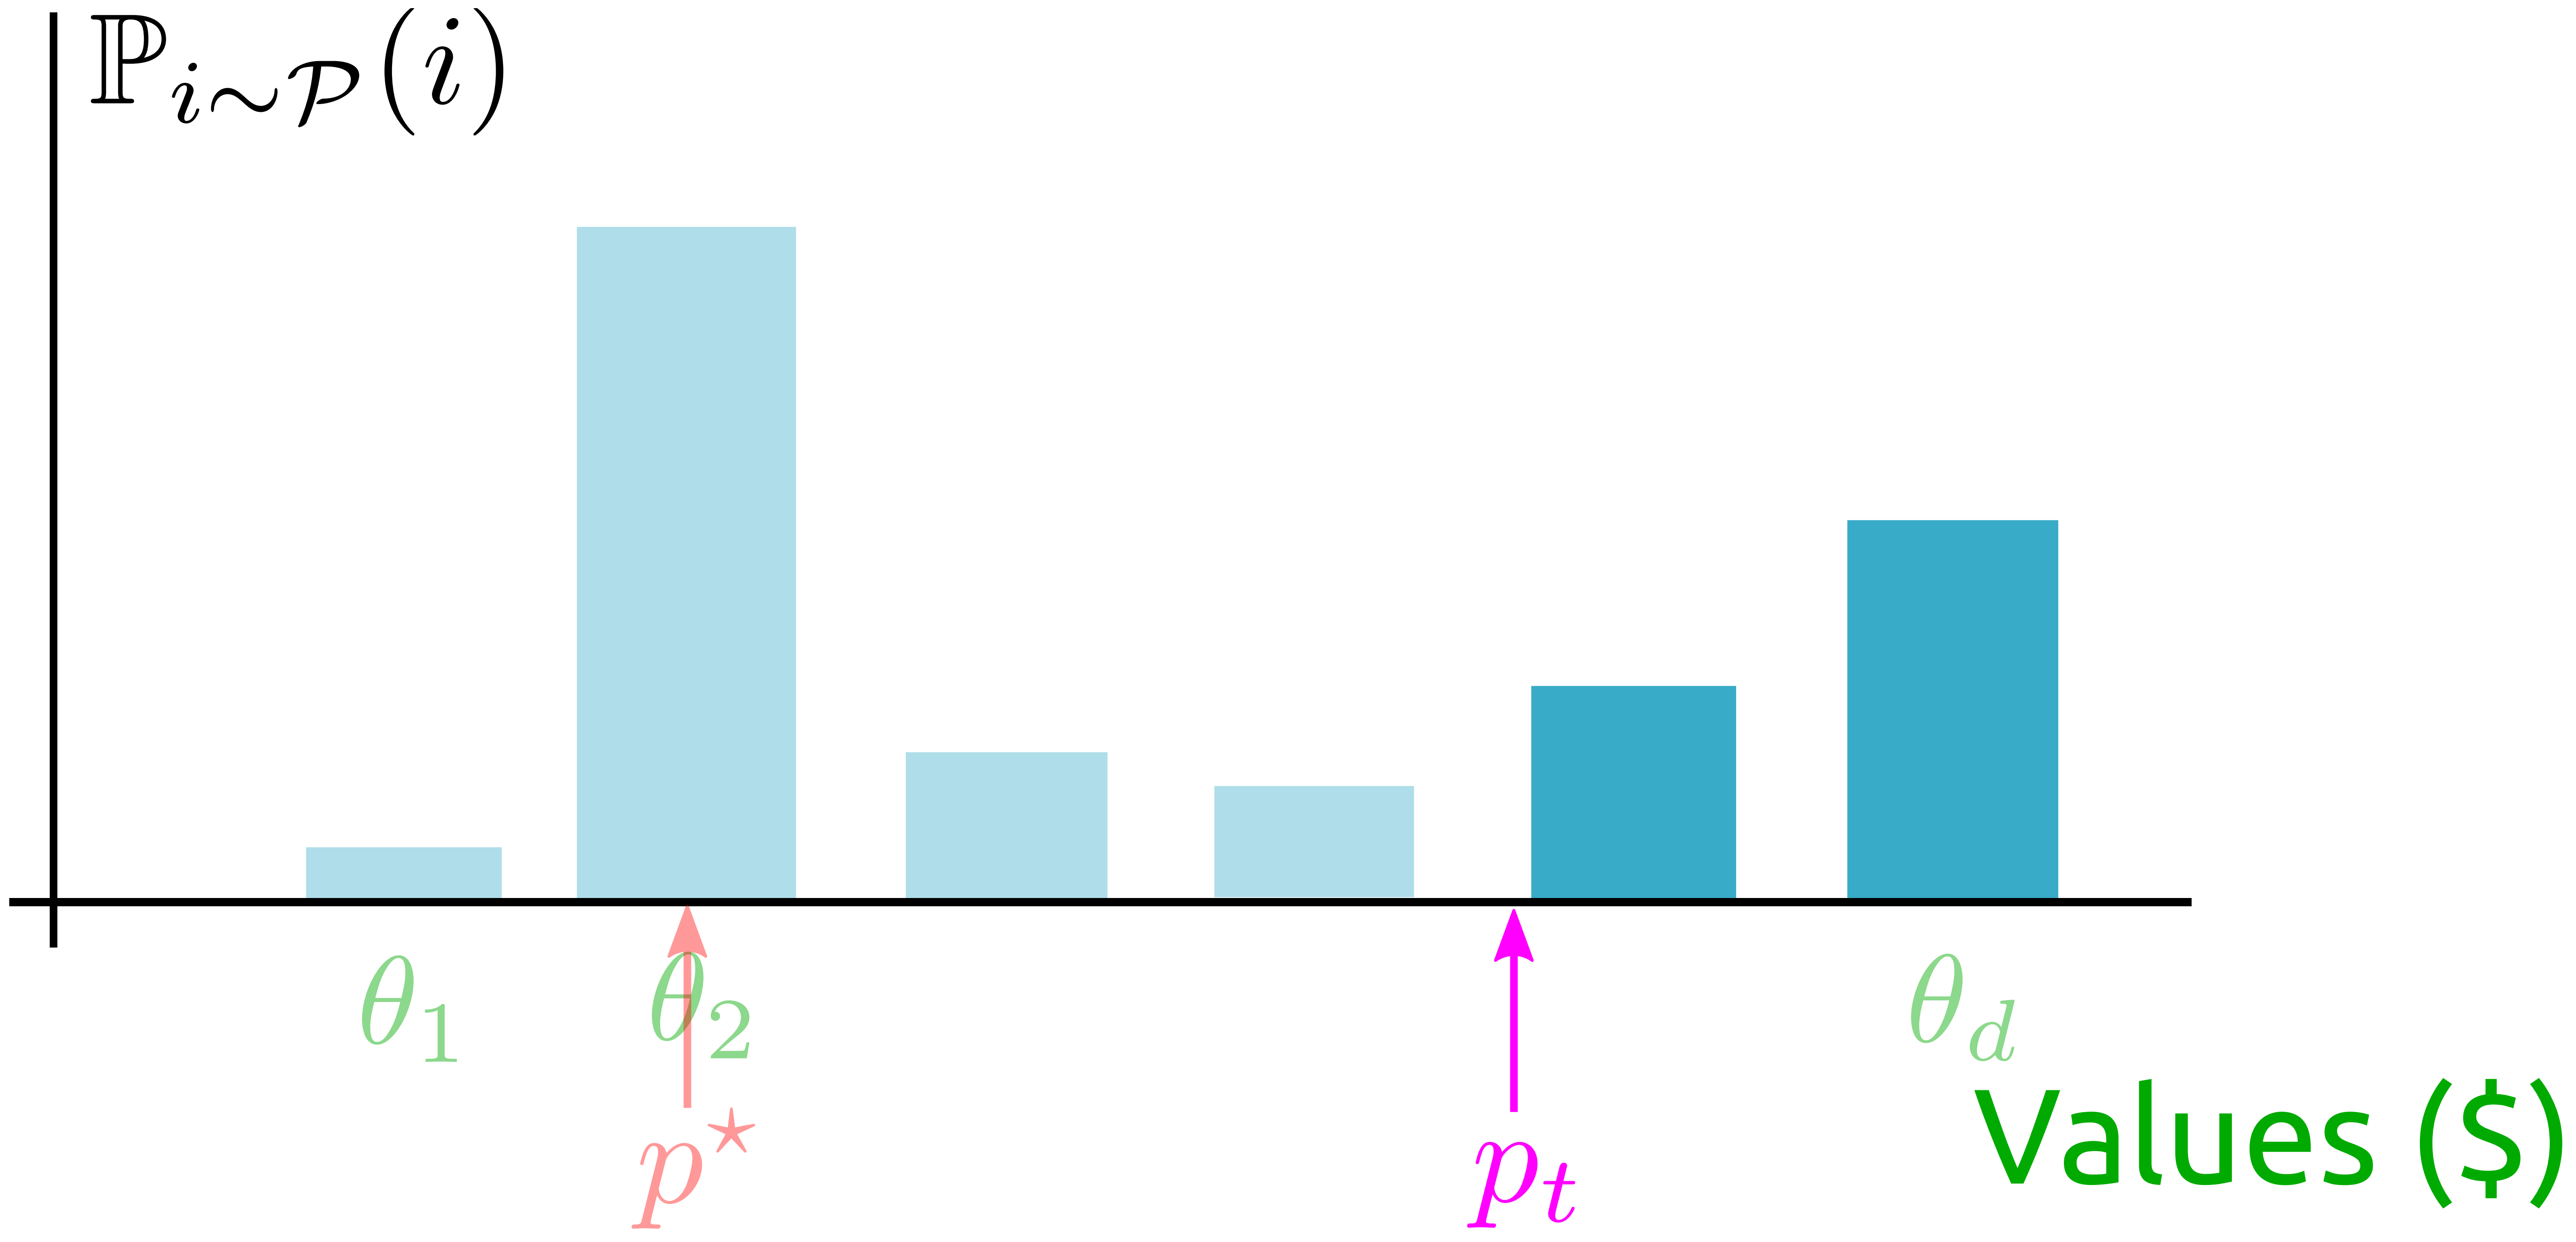

- Set prices too high $\implies$ no feedback about low value types.

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

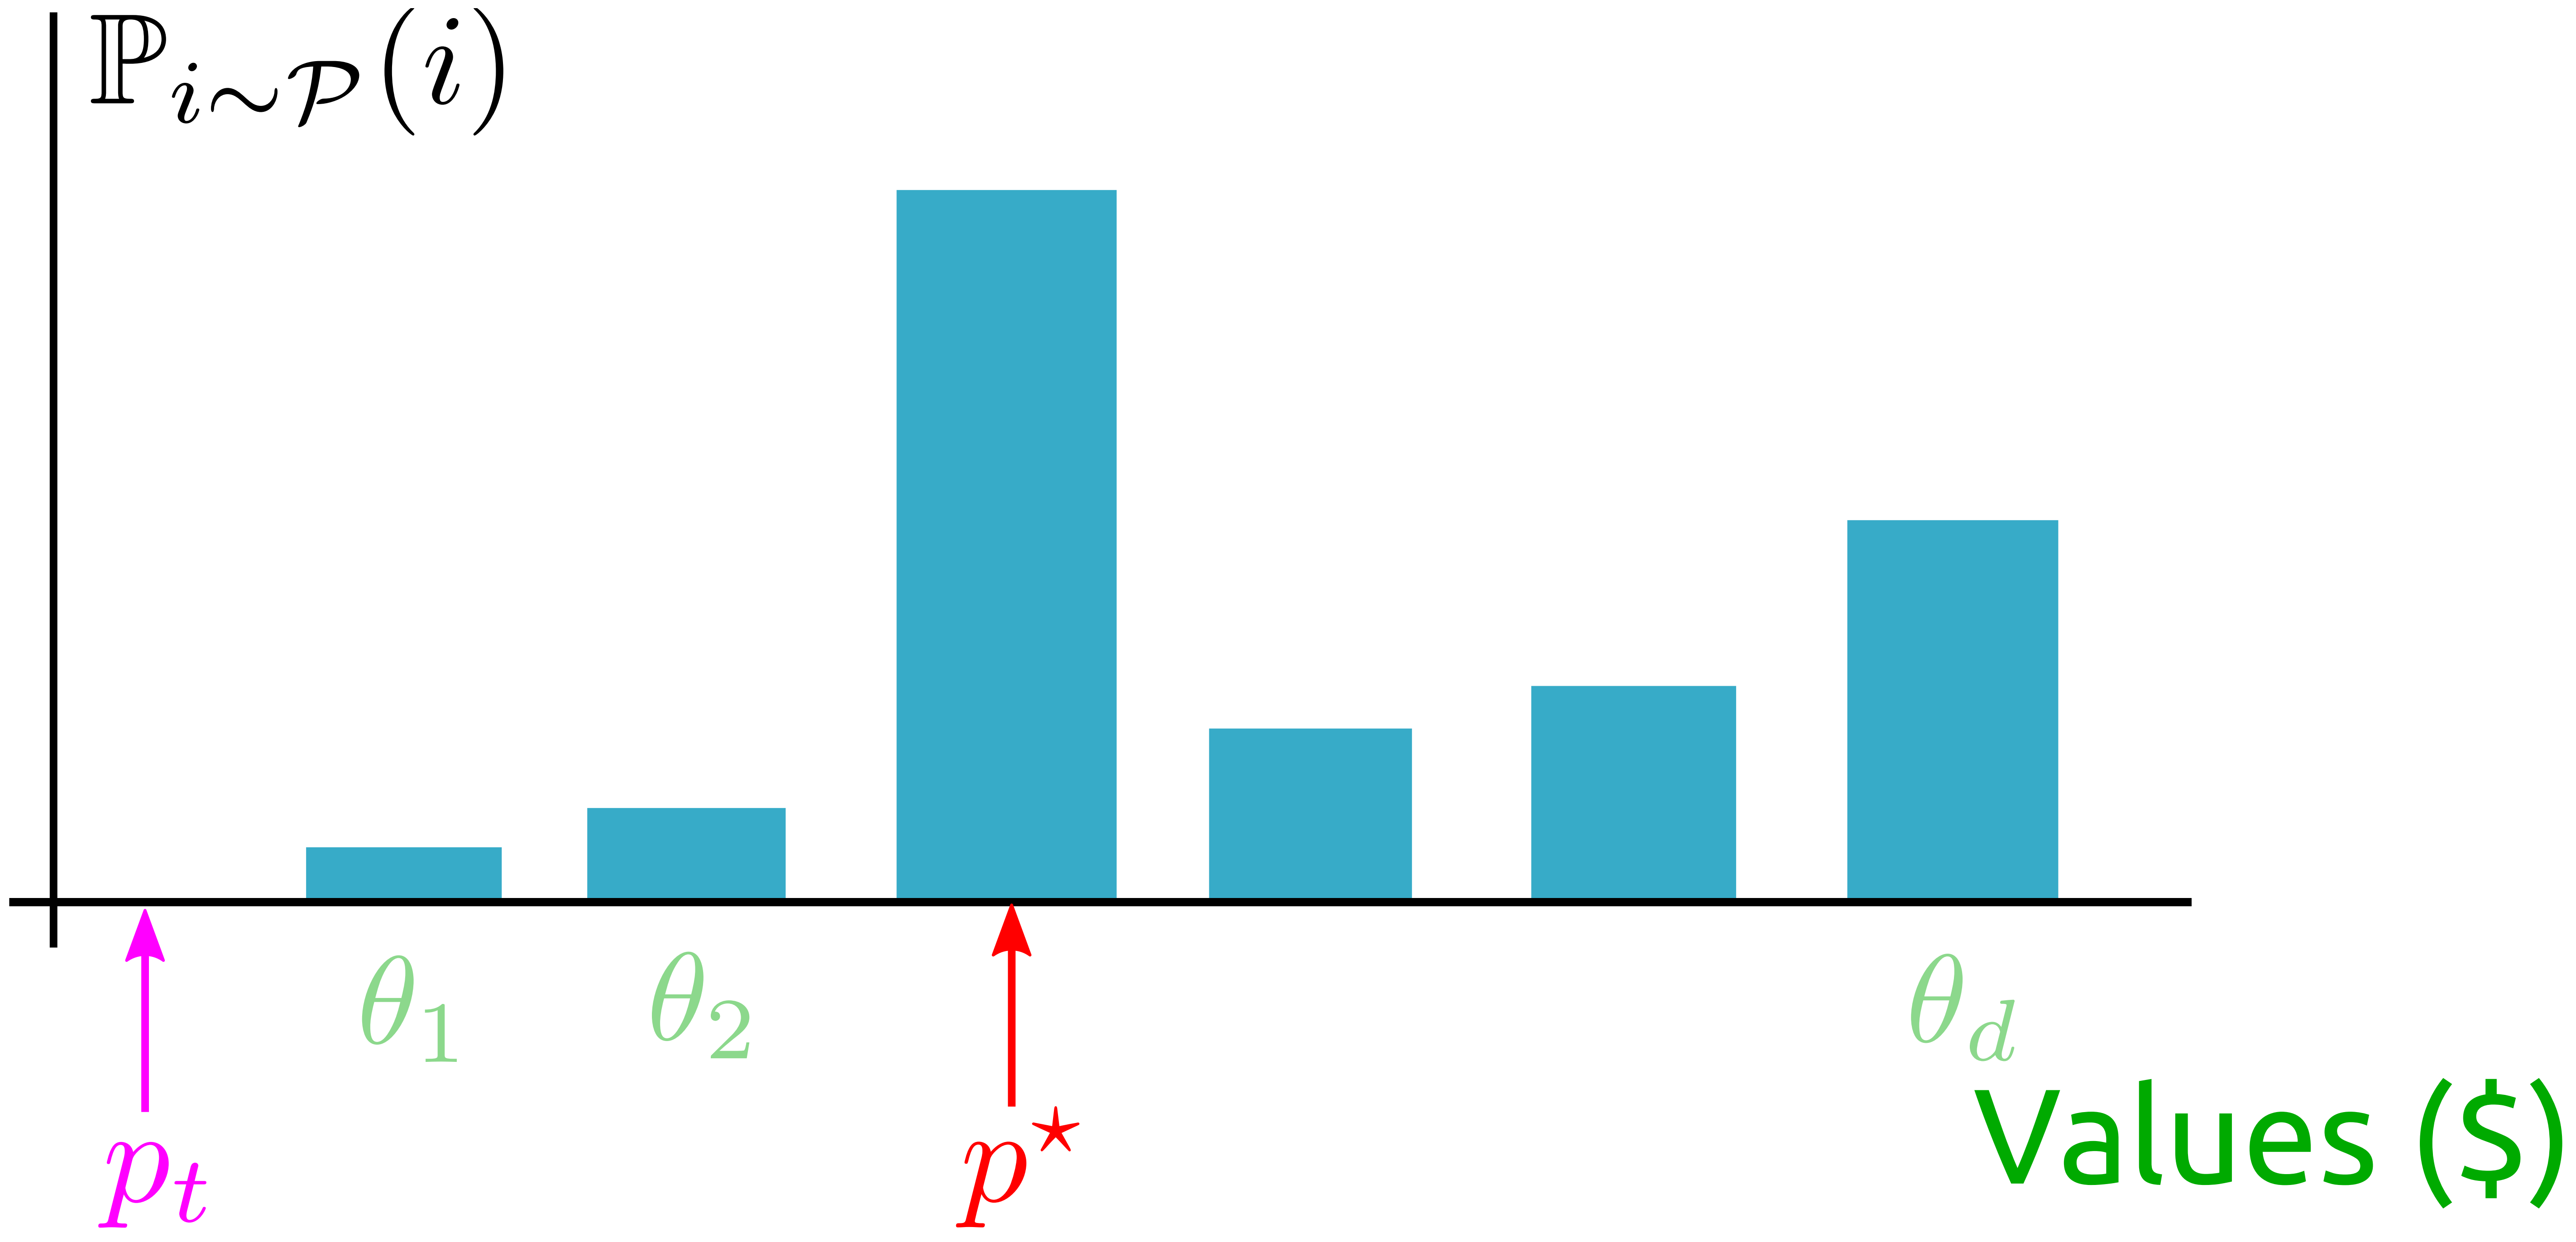

- Set prices too high $\implies$ no feedback about low value types.
- Set prices too low $\implies$ low revenue.

CHALLENGE 1: PRICING VS SELLER LEARNING

- *Even if buyers knew their values*, seller needs to be conservative with pricing.

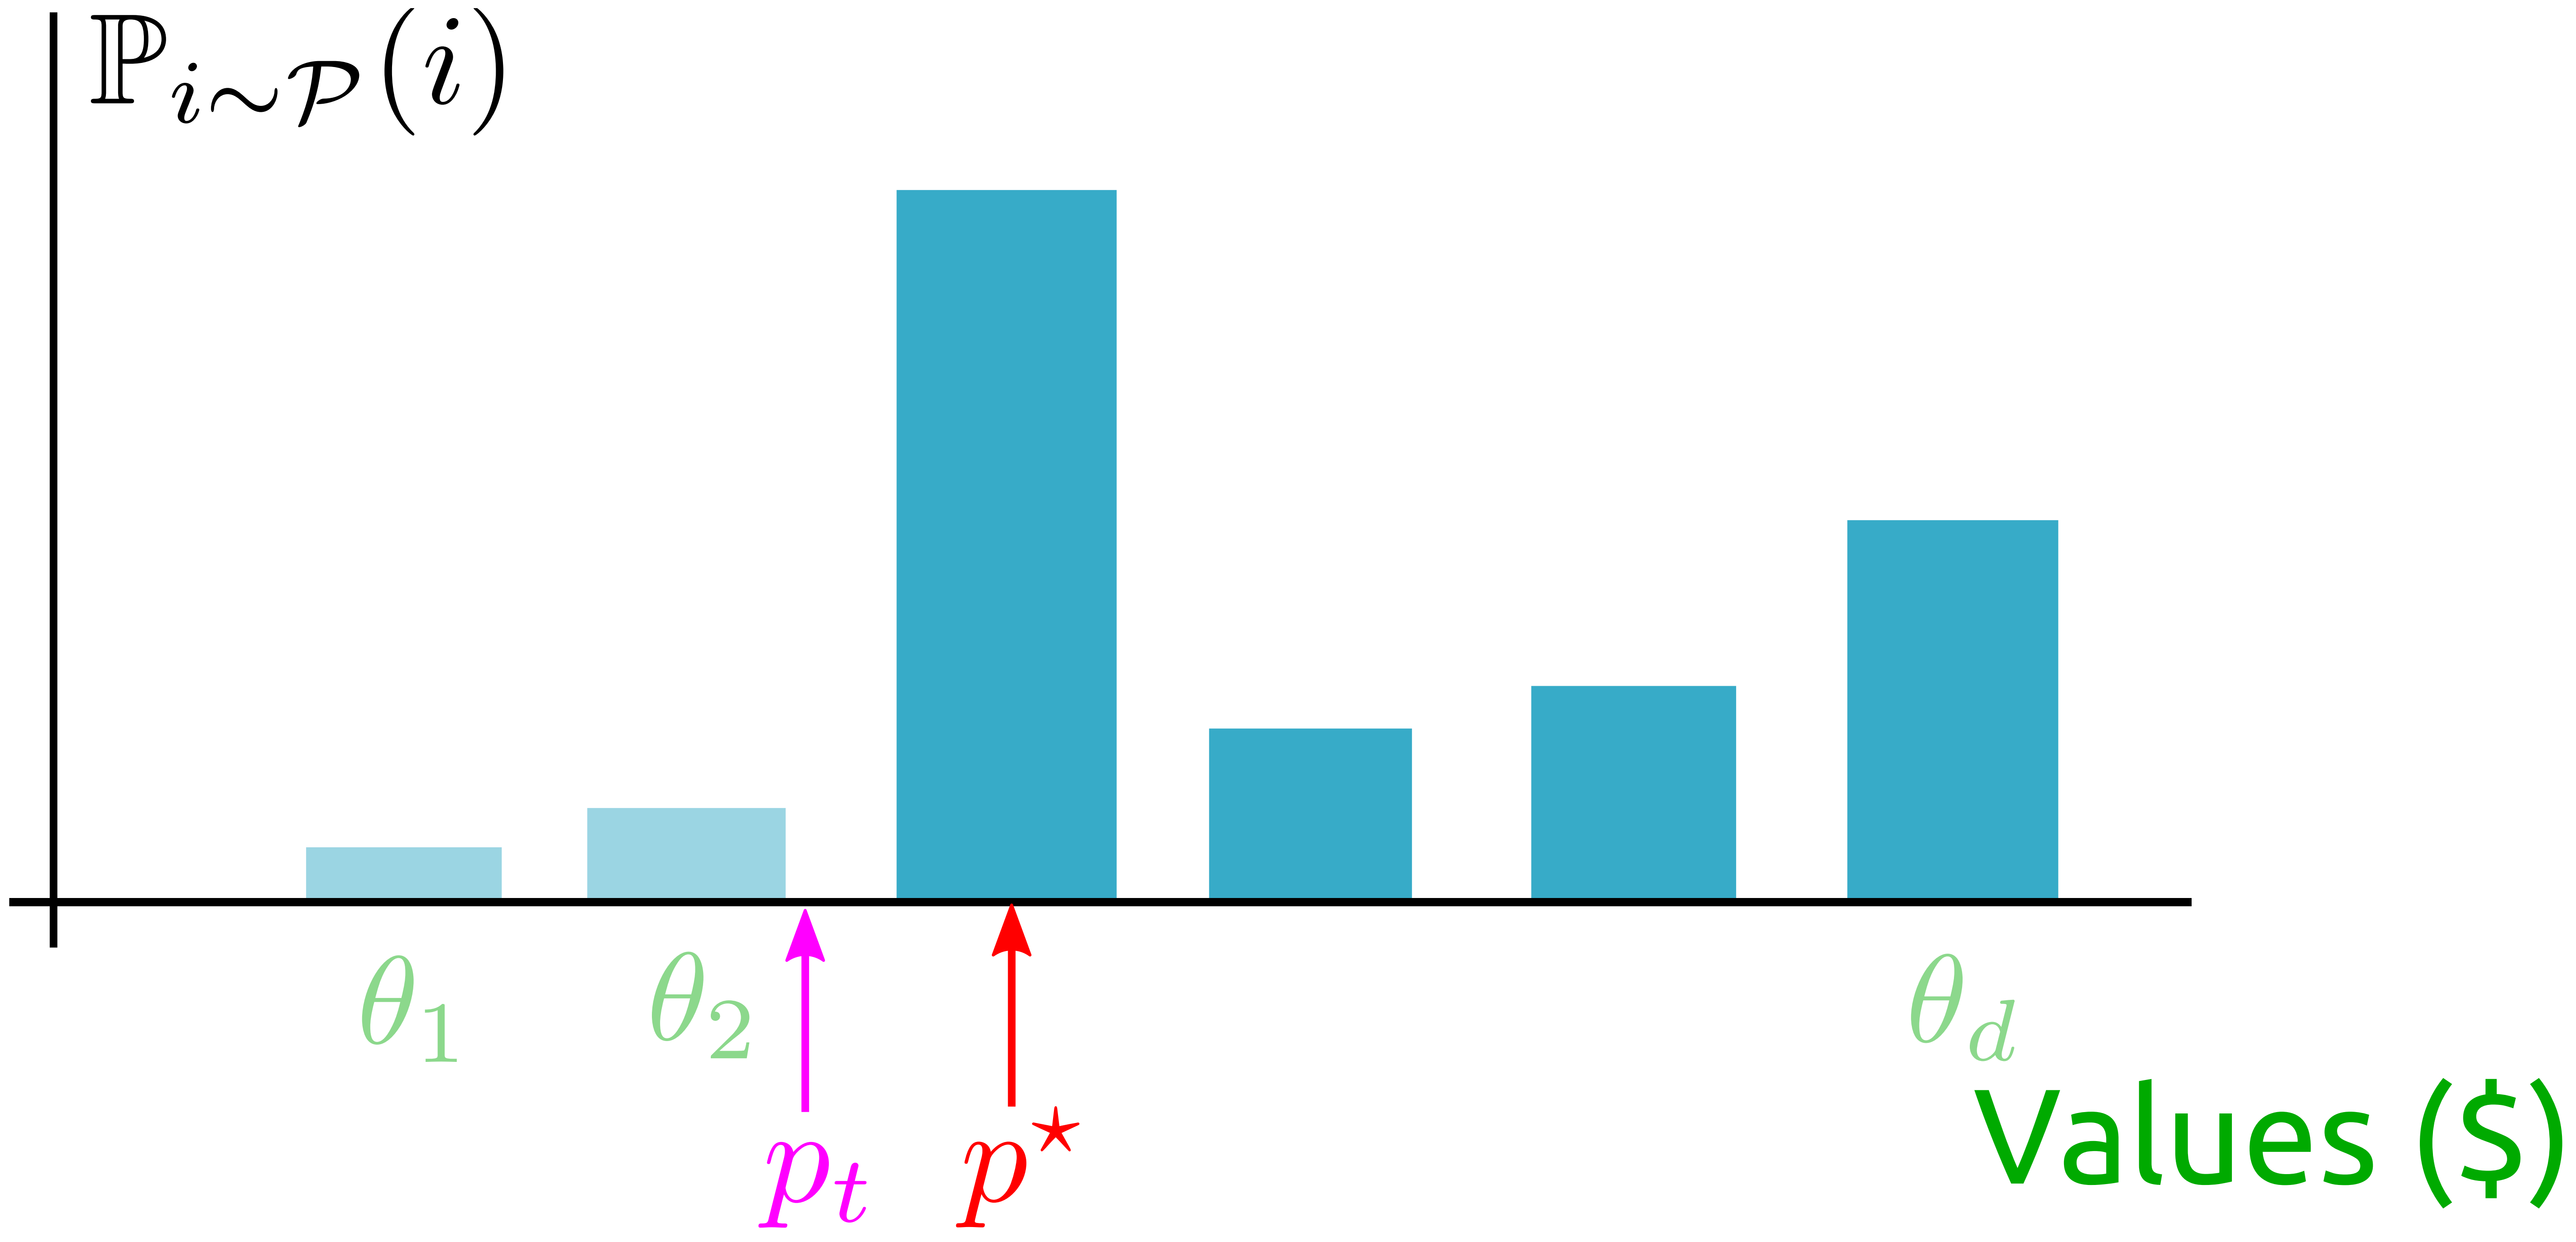

- Set prices too high $\implies$ no feedback about low value types.
- Set prices too low $\implies$ low revenue.
- *Property*: if $p_t \leq p^\star$, and buyers know values, sufficient feedback to learn $p^\star$.

CHALLENGE 2: PRICING VS BUYER LEARNING

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type "stovetop"

★★★★☆ $47	★★★★★ $51	★★★☆☆ $37
★★★☆☆ $39	★★★☆☆ $38	★★★★★ $52
★★★★☆ $42	★★★★☆ $44	★★★★☆ $46
★★★★★ $53	★★★★☆ $45	★★★★★ $56

Type "grill"

★★★★★ $53

★★★★★ $57

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type "stovetop"

★★★★☆ $47	★★★★★ $51	★★★☆☆ $37
★★★☆☆ $39	★★★☆☆ $38	★★★★★ $52
★★★★☆ $42	★★★★☆ $44	★★★★☆ $46
★★★★★ $53	★★★★☆ $45	★★★★★ $56

New stovetop user:

Type "grill"

★★★★★ $53

★★★★★ $57

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type "stovetop"

★★★★☆ $47	★★★★★ $51	★★★☆☆ $37
★★★☆☆ $39	★★★☆☆ $38	★★★★★ $52
★★★★☆ $42	★★★★☆ $44	★★★★☆ $46
★★★★★ $53	★★★★☆ $45	★★★★★ $56

New stovetop user:

"I will pay up to $42 for this pan"

Type "grill"

★★★★★ $53
★★★★★ $57

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type "stovetop"

★★★★☆ $47 ★★★★★ $51 ★★★☆☆ $37
★★★☆☆ $39 ★★★☆☆ $38 ★★★★★ $52
★★★★☆ $42 ★★★★☆ $44 ★★★★☆ $46
★★★★★ $53 ★★★★☆ $45 ★★★★★ $56

New stovetop user:

"I will pay up to $42 for this pan"

Type "grill"

★★★★★ $53
★★★★★ $57

New grill user:

"I will pay up to $33 for this pan"

CHALLENGE 2: PRICING VS BUYER LEARNING

- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

- Seller's dilemma: Only target "stovetop" buyers for high immediate revenue? Or also target "grill" customers for higher long term revenue?

ALGORITHM

- **Phase 1:**

- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)

- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
  - Offer item for *free*.
- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
  - Offer item for *free*.
  - Observe iid samples from type distribution $\mathscr{P}$.
- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
  - Offer item for *free*.
  - Observe iid samples from type distribution $\mathscr{P}$.
  - Eliminate types that did not appear often enough.

- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
  - Offer item for *free*.
  - Observe iid samples from type distribution $\mathscr{P}$.
  - Eliminate types that did not appear often enough.
    - $S_0 \leftarrow$ remaining types.
- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
  - Offer item for *free*.
  - Observe iid samples from type distribution $\mathscr{P}$.
  - Eliminate types that did not appear often enough.
    - $S_0 \leftarrow$ remaining types.
- **Phase 2:** (set $S_t = S_0$)

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
    - Offer item for *free*.
    - Observe iid samples from type distribution $\mathscr{P}$.
    - Eliminate types that did not appear often enough.
        - $S_0 \leftarrow$ remaining types.
- **Phase 2:** (set $S_t = S_0$)
    - On round t, set price p_t *low enough* that buyers in S_t will buy.

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(T^{1/3})$ rounds)
    - Offer item for *free*.
    - Observe iid samples from type distribution $\mathscr{P}$.
    - Eliminate types that did not appear often enough.
        - $S_0 \leftarrow$ remaining types.

- **Phase 2:** (set $S_t = S_0$)
    - On round t, set price p_t *low enough* that buyers in S_t will buy.
    - Update S_t: eliminate types which contribute too little to revenue.

PHASE 2

PHASE 2

- S_t: customers we are targeting in the current round.

PHASE 2

- S_t: customers we are targeting in the current round.
- Maintain η-confidence intervals for buyers' estimate of θ_i for each $i \in S_t$.

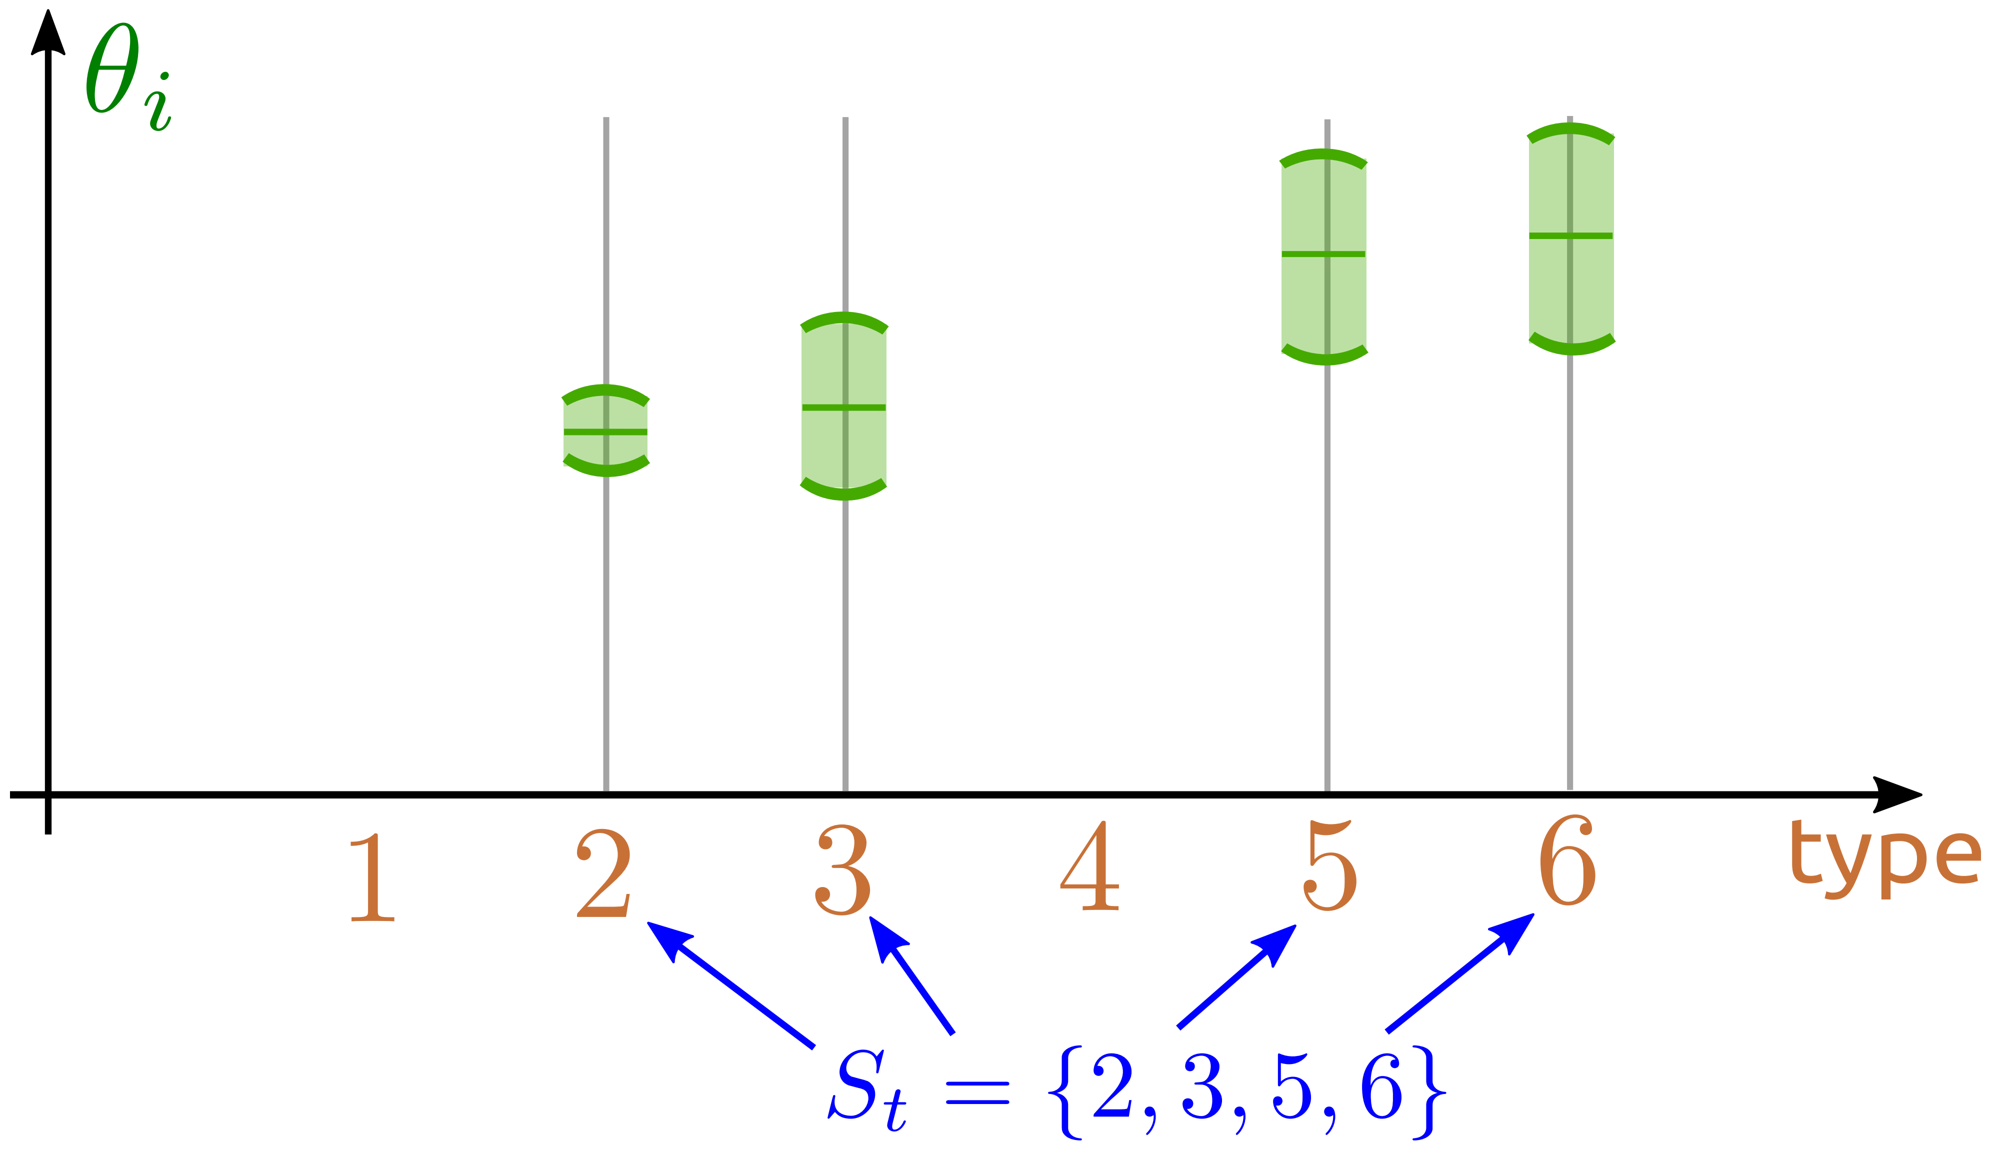

PHASE 2

- S_t: customers we are targeting in the current round.
- Maintain η-confidence intervals for buyers' estimate of θ_i for each $i \in S_t$.

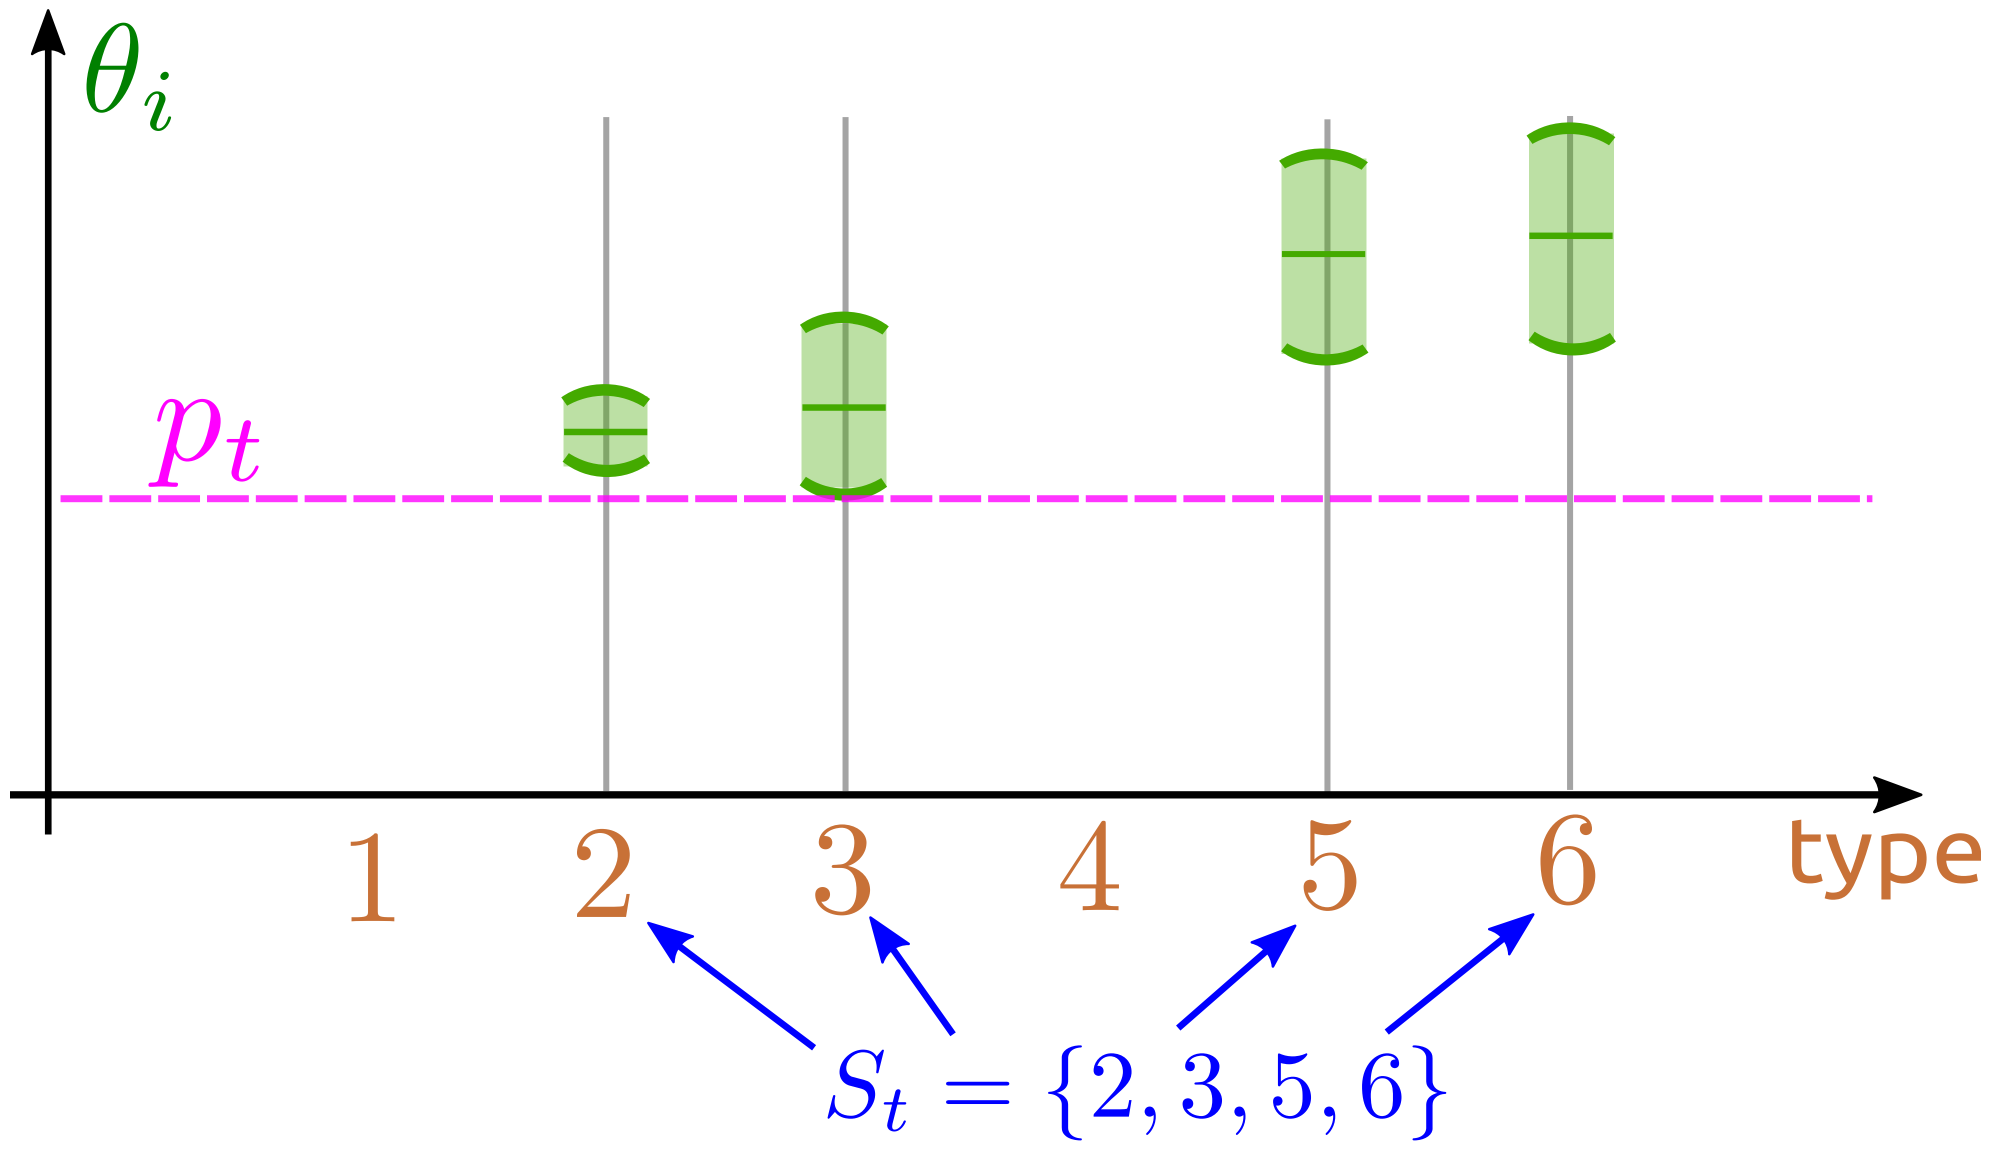

- Choose p_t to be the minimum LCB of these confidence intervals.

PHASE 2

- S_t: customers we are targeting in the current round.
- Maintain η-confidence intervals for buyers' estimate of θ_i for each $i \in S_t$.

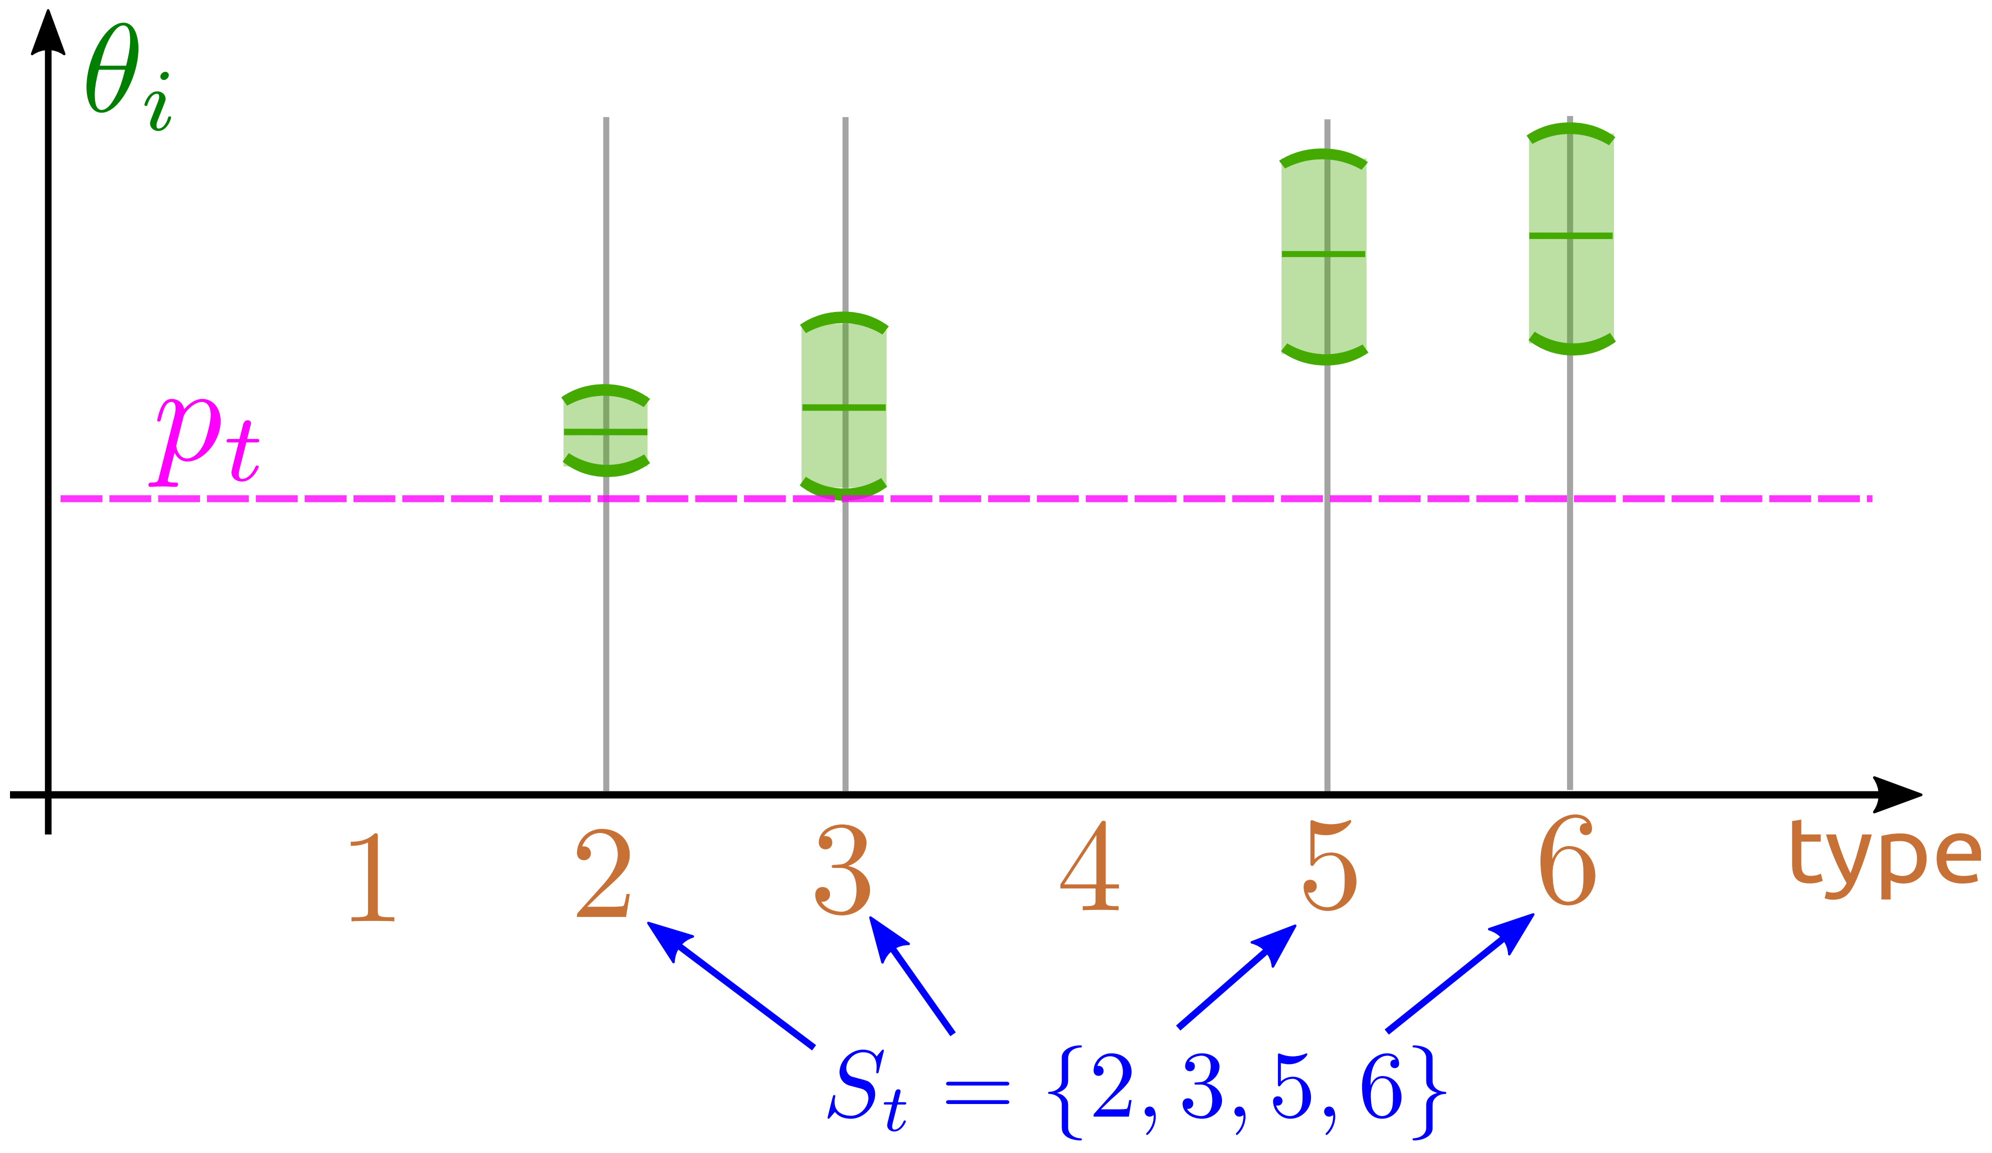

- Choose p_t to be the minimum LCB of these confidence intervals.
- Updating S_t: racing-like elimination used in Best Arm Identification.

WHY DO WE NEED A PHASE 1?

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for free, eliminate types that are infrequent.

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for free, eliminate types that are infrequent.

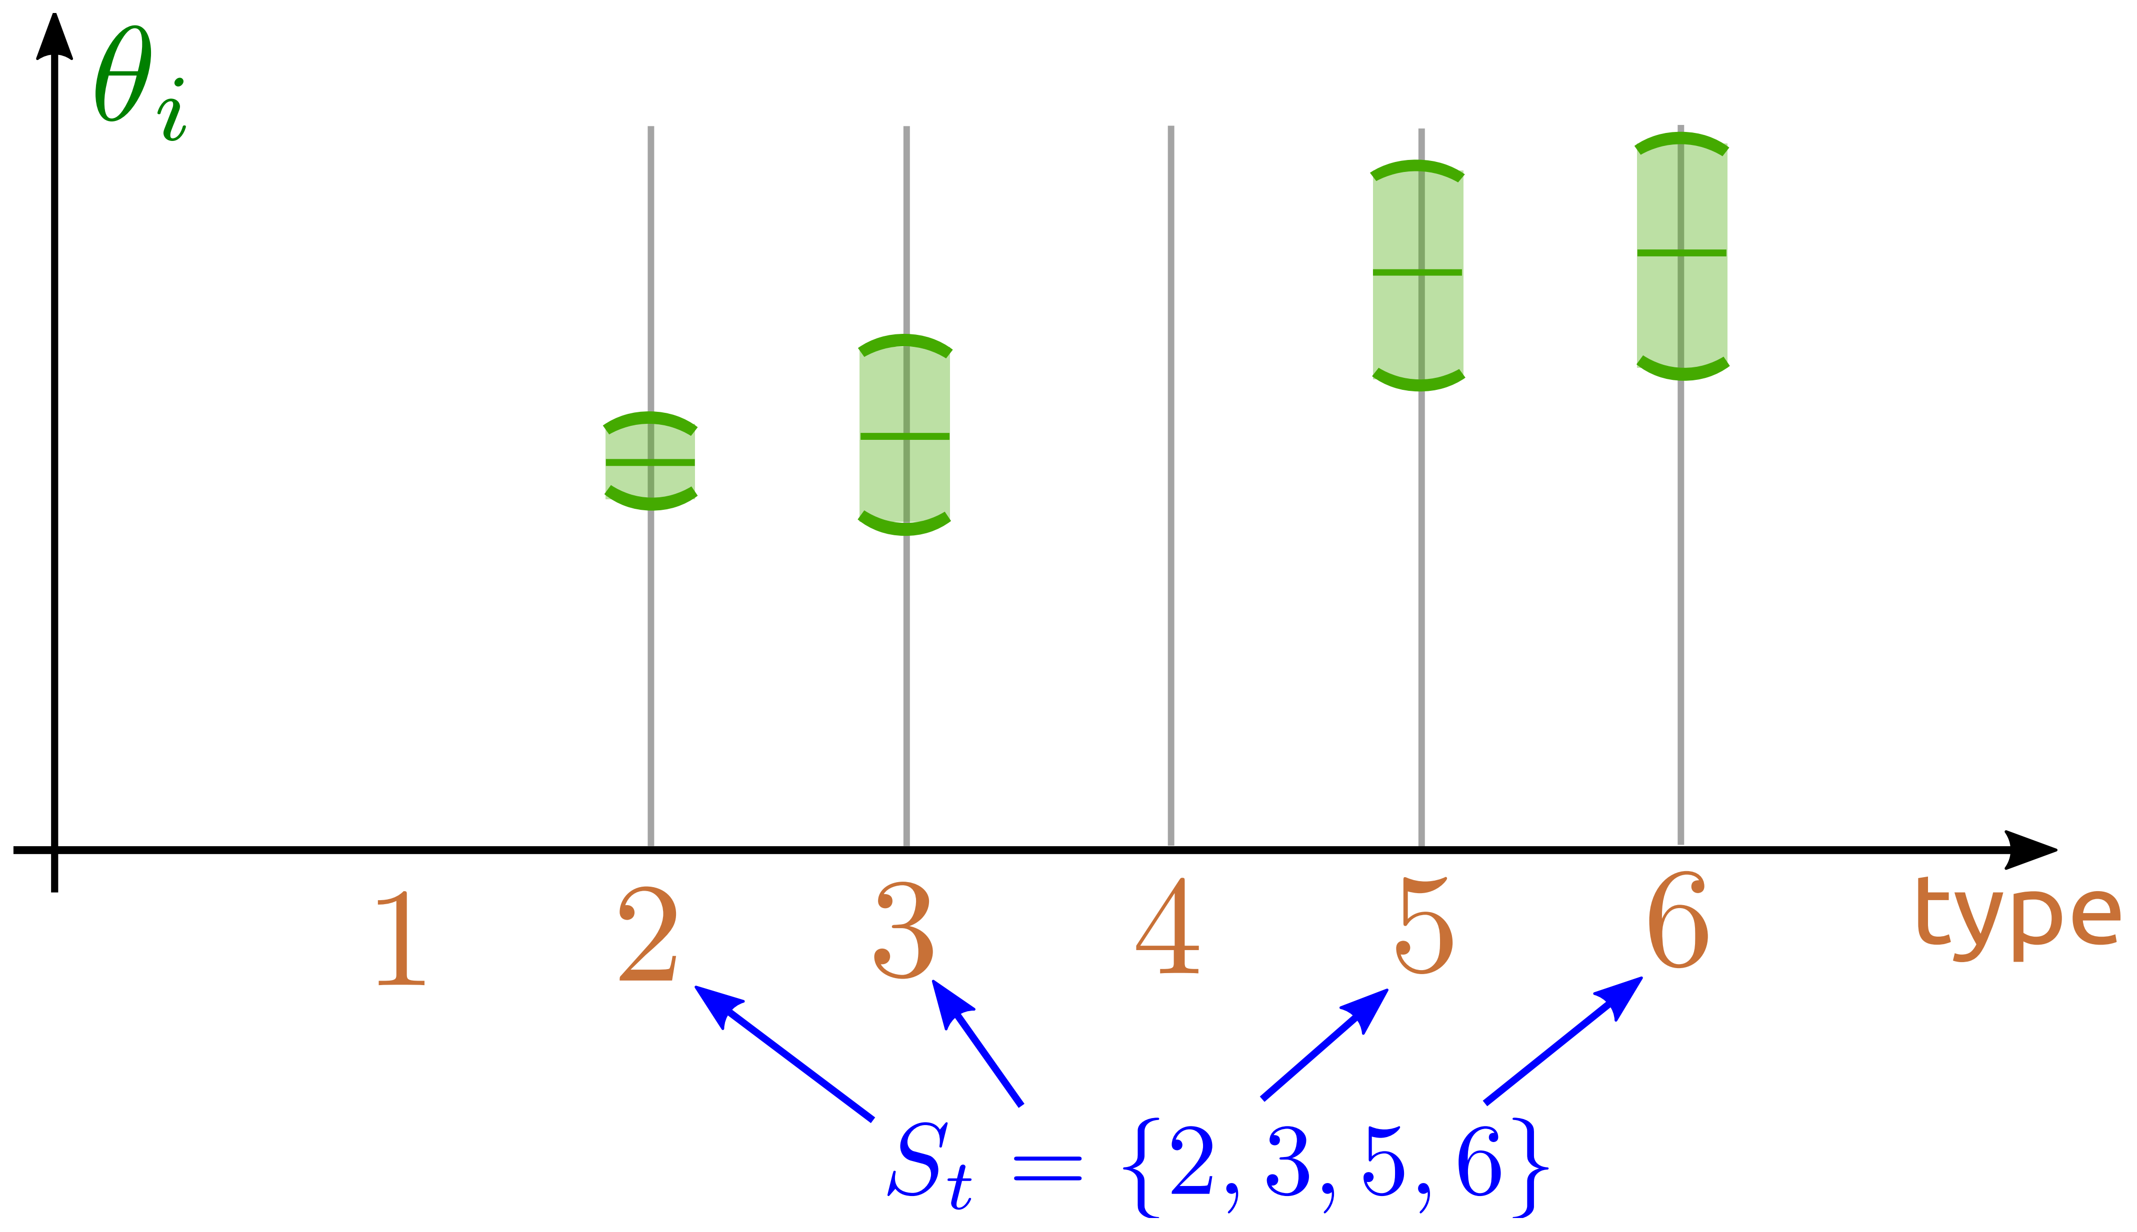

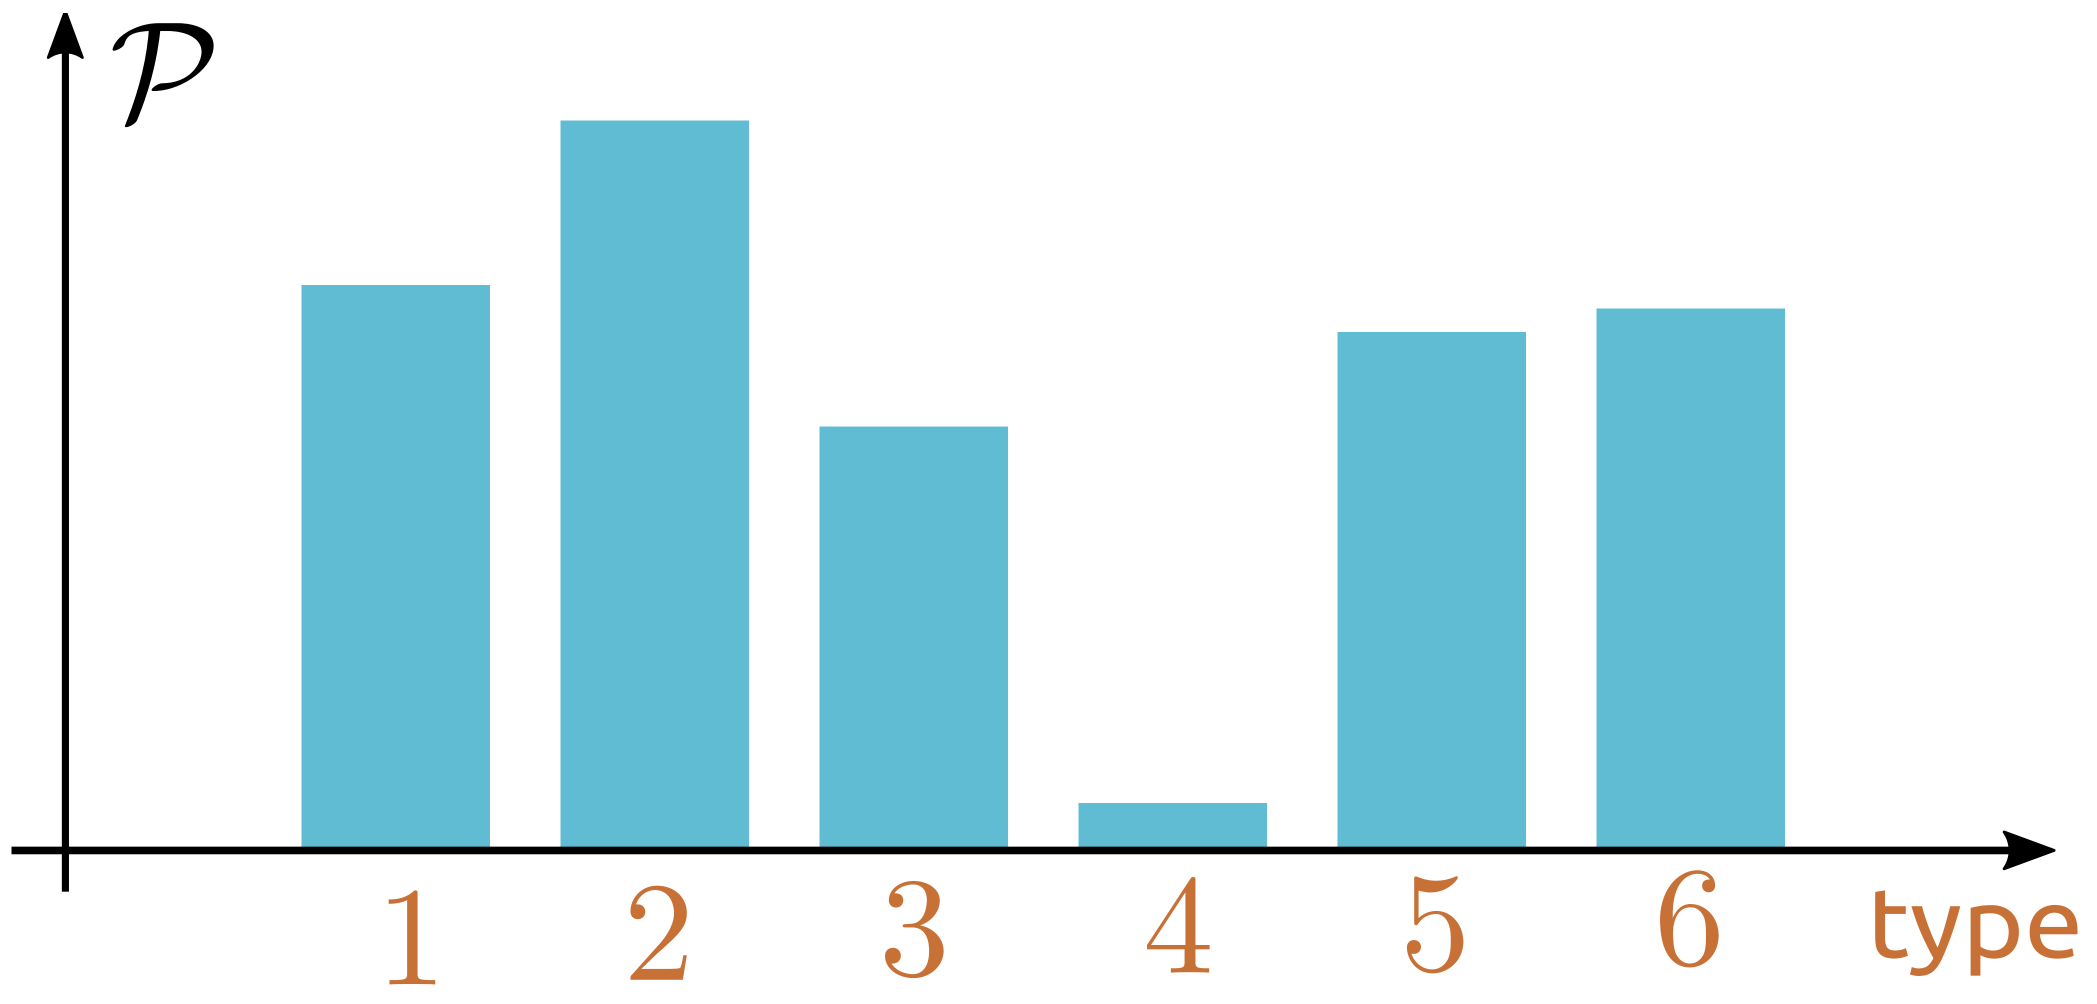

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for free, eliminate types that are infrequent.
  - Low probability of appearance $\Longrightarrow$ fewer reviews.

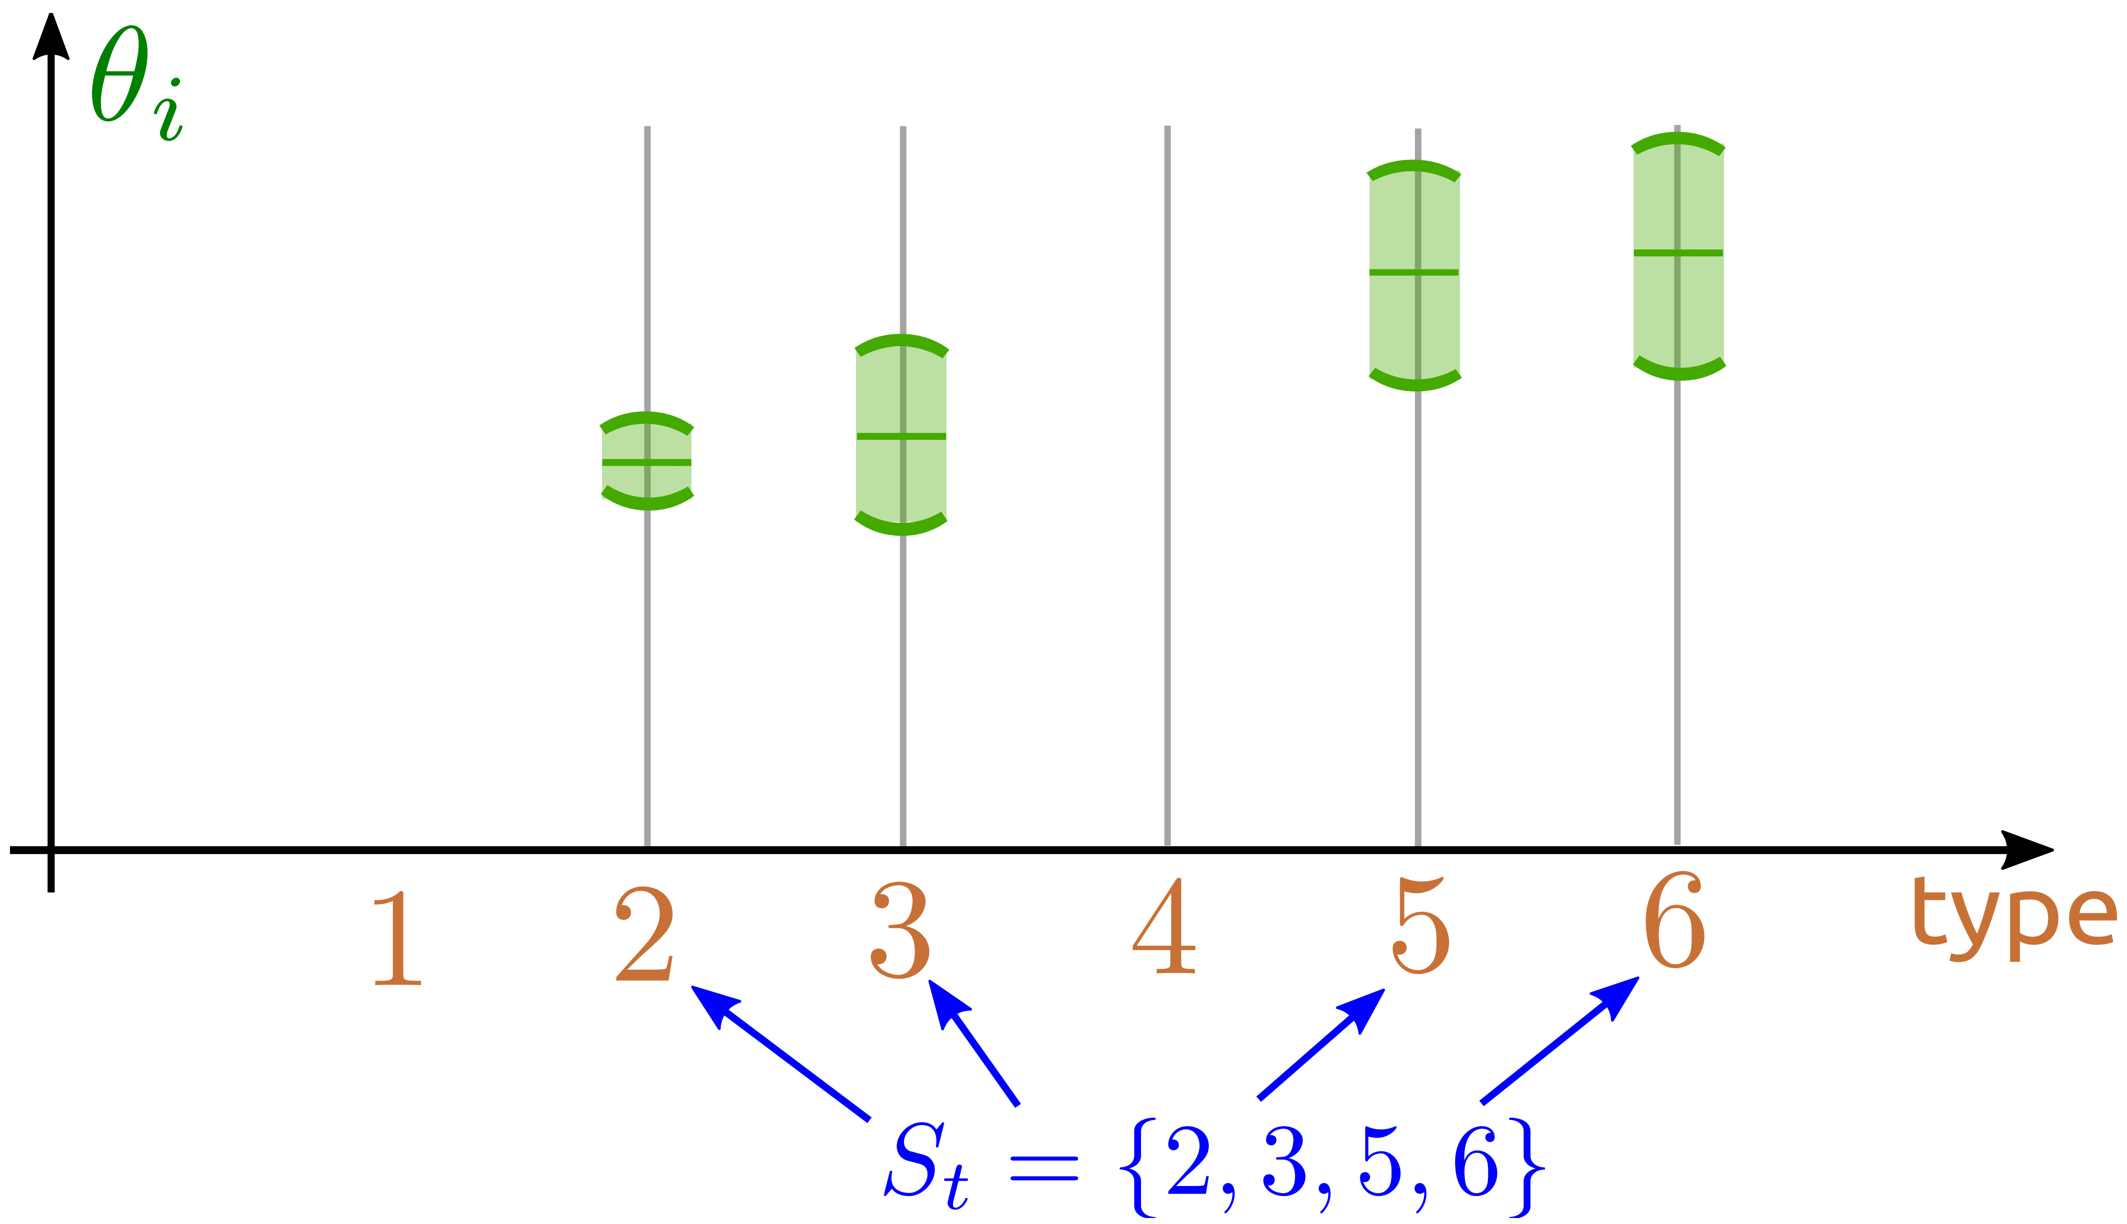

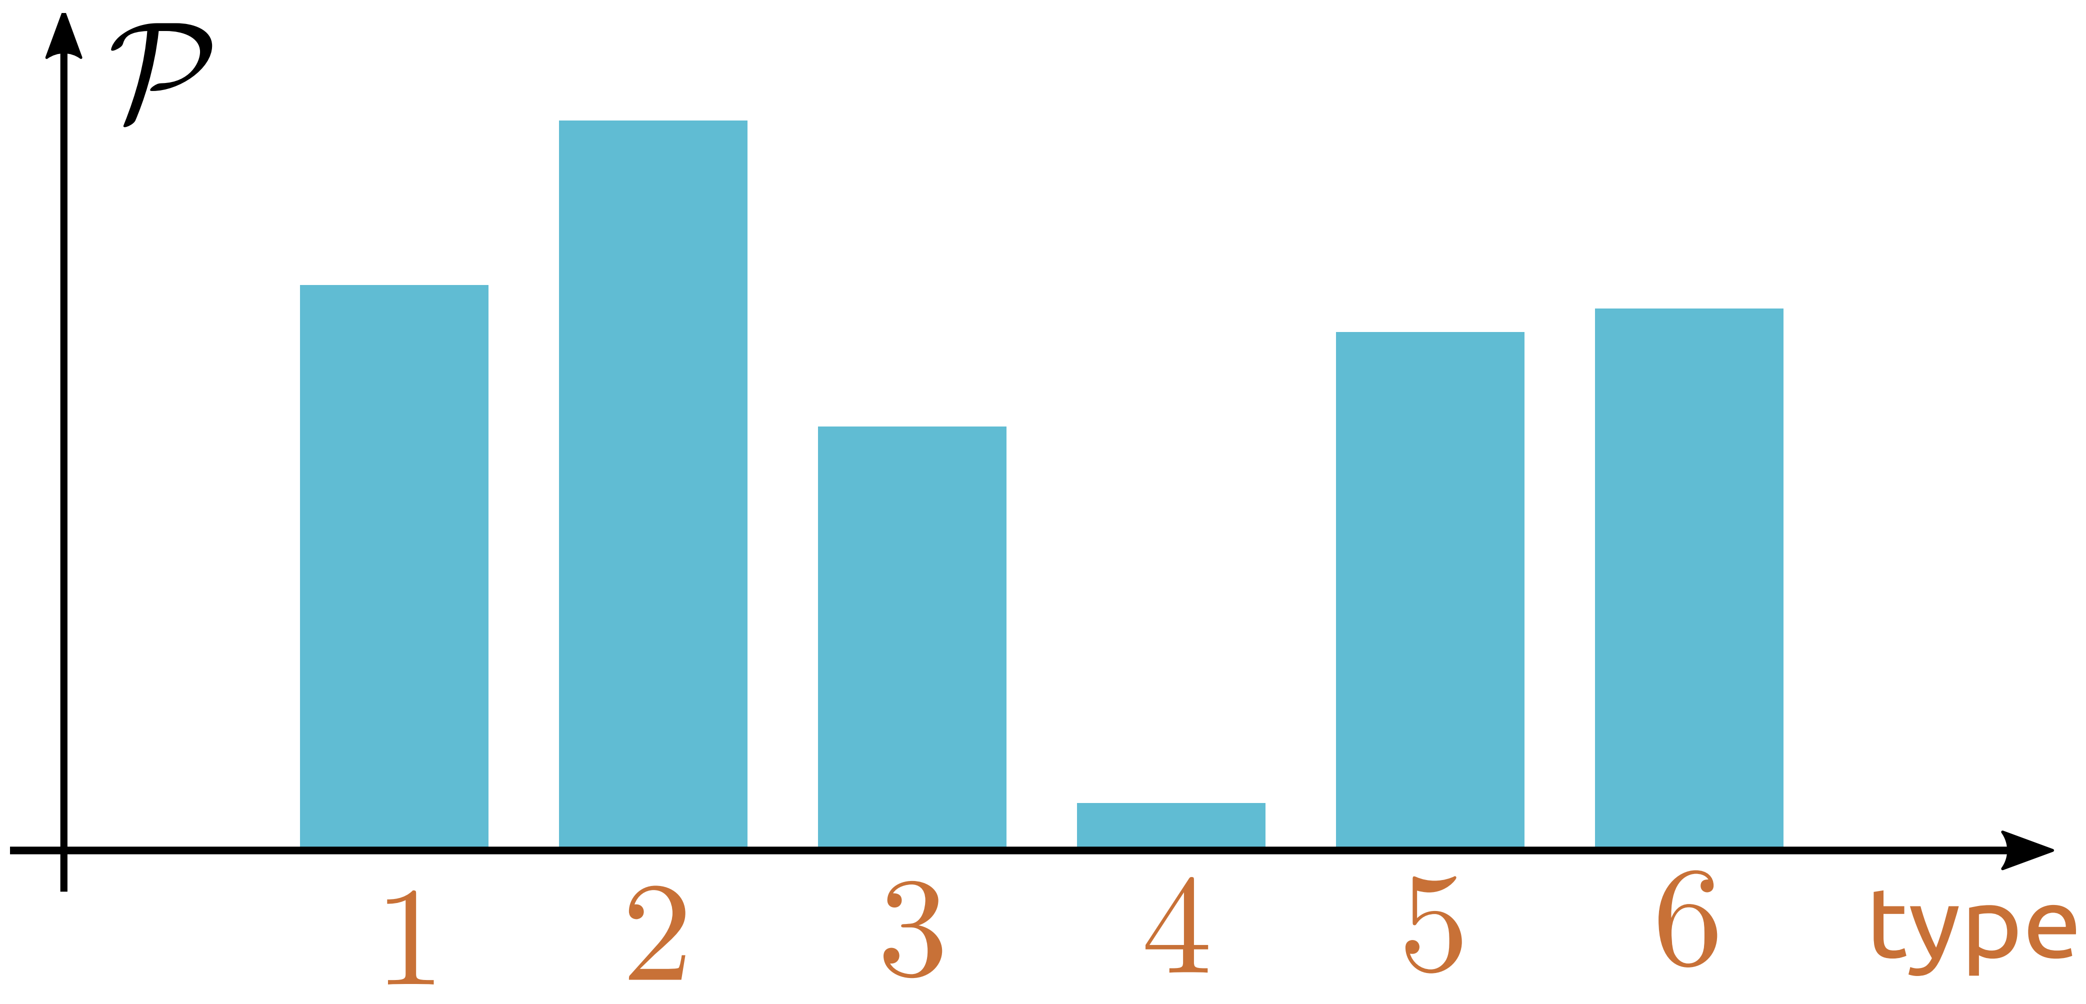

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for free, eliminate types that are infrequent.
  - Low probability of appearance $\Longrightarrow$ fewer reviews.
  - More uncertainty about their value.

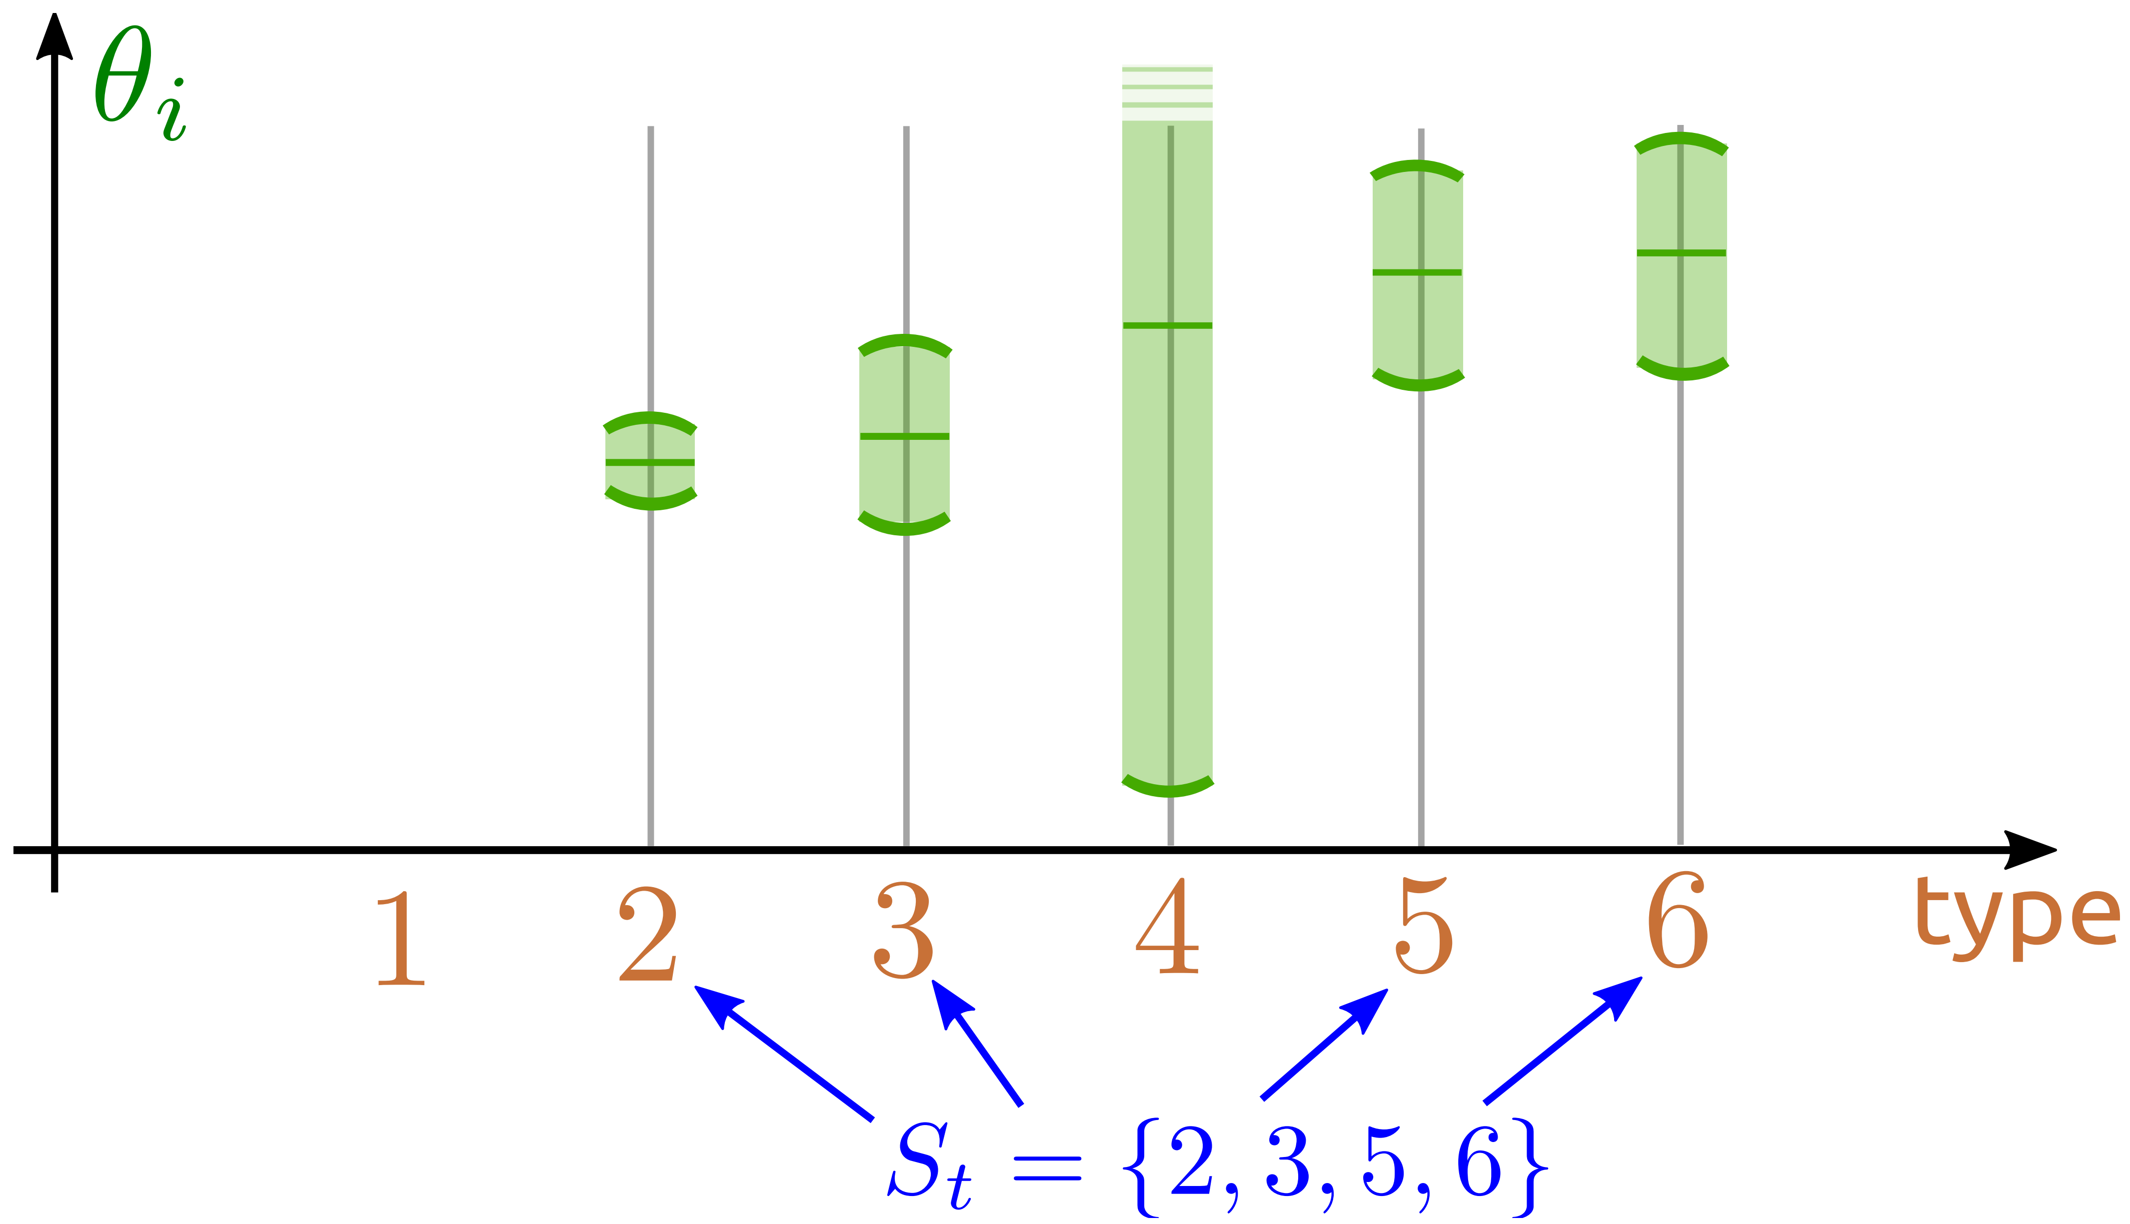

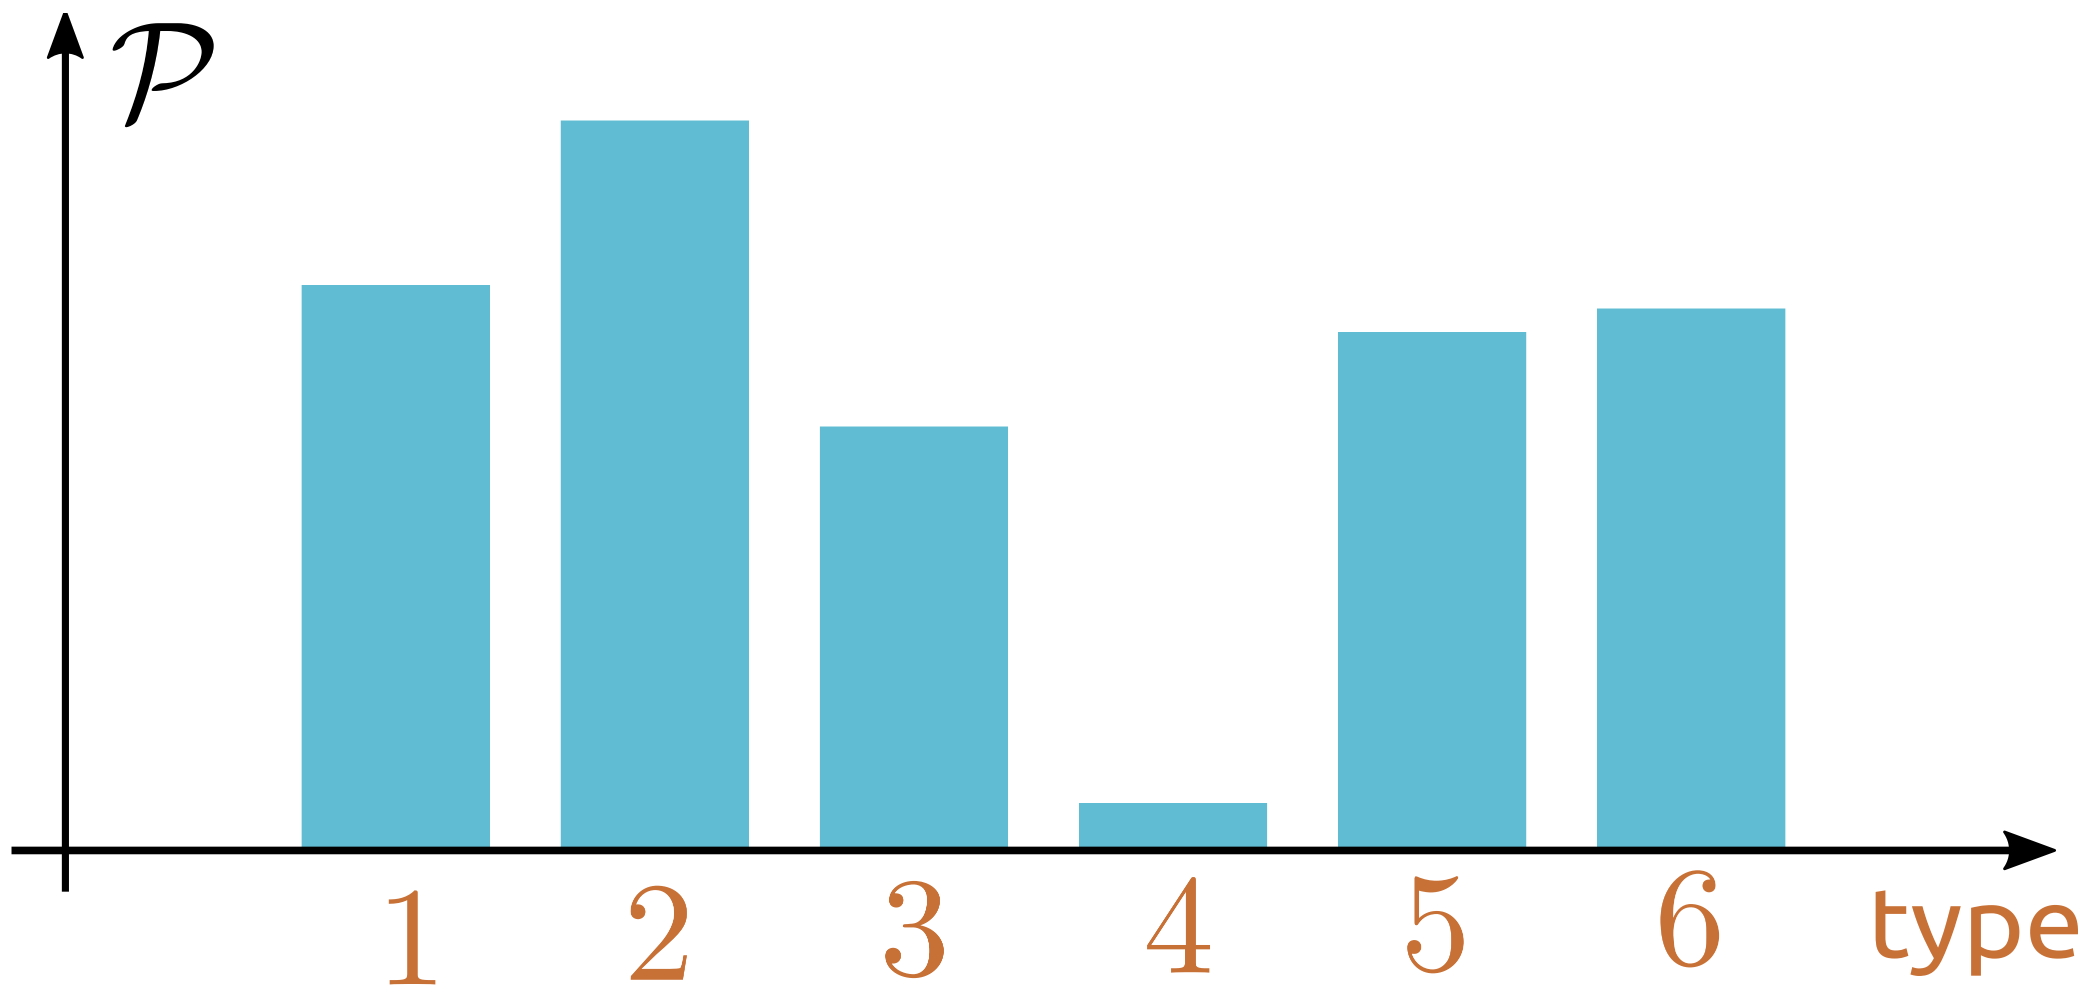

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for free, eliminate types that are infrequent.
  - Low probability of appearance $\implies$ fewer reviews.
  - More uncertainty about their value.
  - Need to set a low price to target these buyers $\implies$ low revenue.

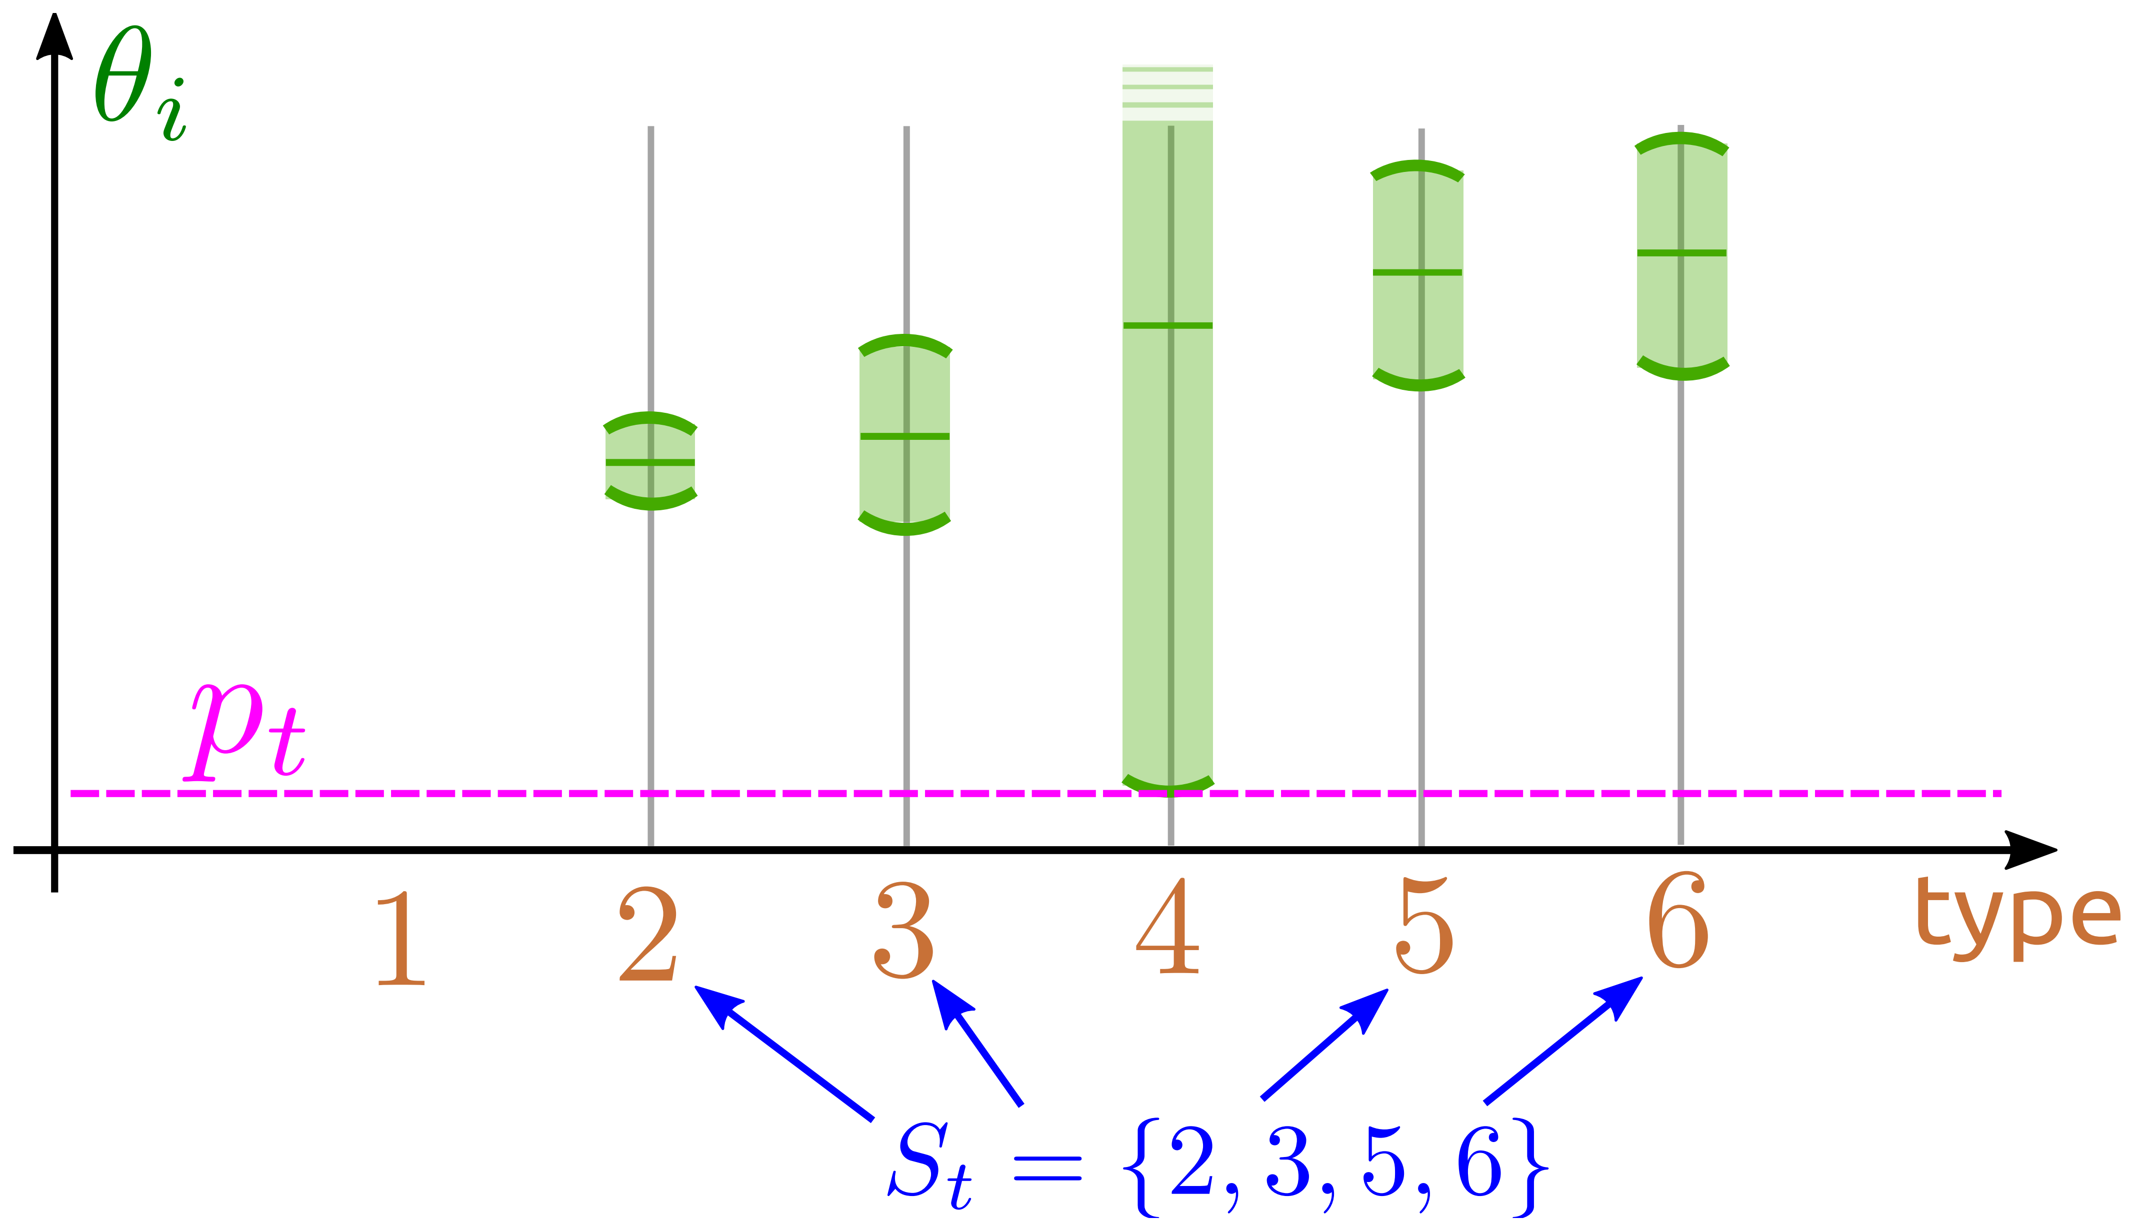

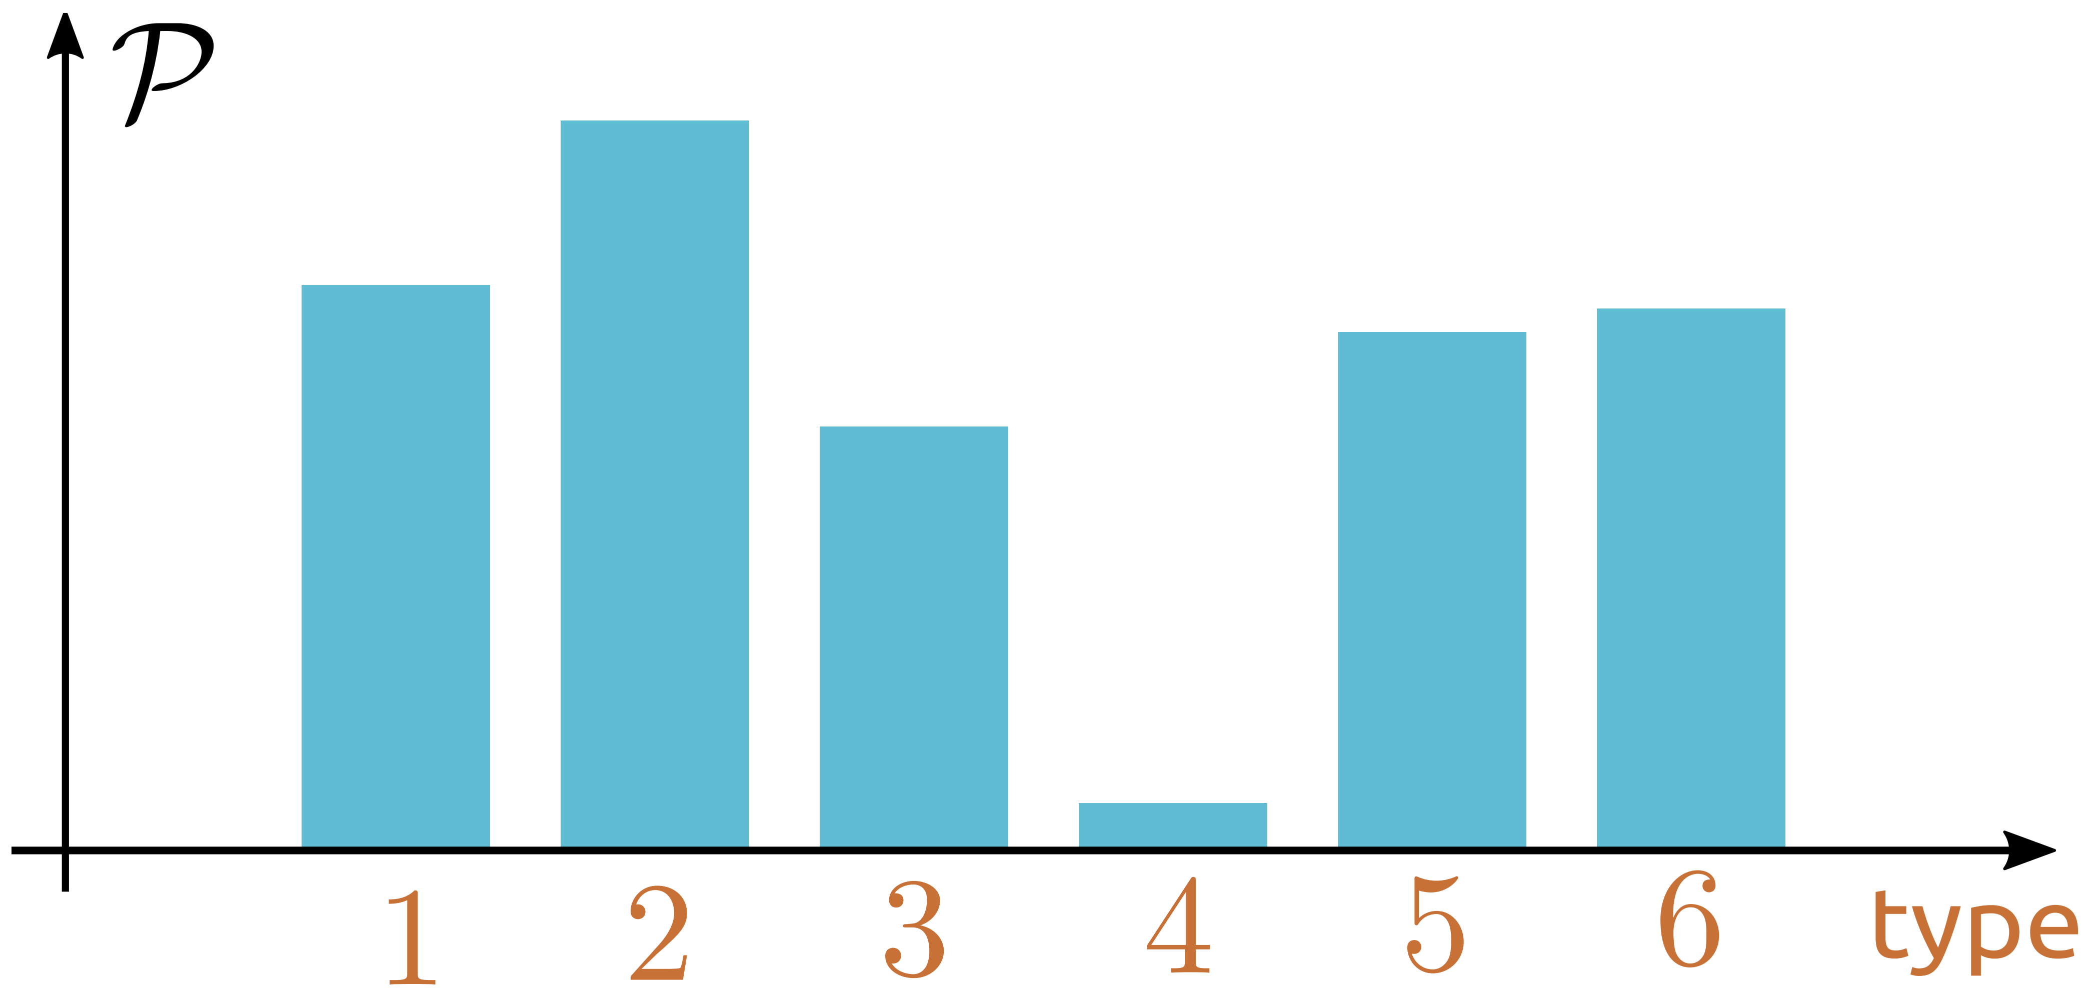

THEORETICAL RESULTS

THEORETICAL RESULTS

Upper bound: After T rounds with d types,

THEORETICAL RESULTS

Upper bound: After T rounds with d types,

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

$q_{\min}$: lowest type probability

THEORETICAL RESULTS

Upper bound: After T rounds with d types,

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

$q_{\min}$: lowest type probability

Lower bound:

THEORETICAL RESULTS

Upper bound: After T rounds with d types,

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

$q_{\min}$: lowest type probability

Lower bound:

$$\inf_{\text{algorithms}} \ \sup_{\text{problems}} \ \mathbb{E}[R_T] \ \gtrsim \ d^{1/3}T^{2/3}$$

THEORETICAL RESULTS

Upper bound: After T rounds with d types,

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

$q_{\min}$: lowest type probability

Lower bound:

$$\inf_{\text{algorithms}} \sup_{\text{problems}} \mathbb{E}[R_T] \gtrsim d^{1/3}T^{2/3}$$

$$\inf_{\text{algorithms}} \sup_{q_{\min} \geq d^{-2/3}T^{-1/3}} \mathbb{E}[R_T] \gtrsim \sqrt{\frac{T}{q_{\min}}}$$

Upper bound: After T rounds with d types,

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

$q_{\min}$: lowest type probability

Lower bound:

$$\inf_{\text{algorithms}} \sup_{\text{problems}} \mathbb{E}[R_T] \gtrsim d^{1/3}T^{2/3}$$

$$\inf_{\text{algorithms}} \sup_{q_{\min} \geq d^{-2/3}T^{-1/3}} \mathbb{E}[R_T] \gtrsim \sqrt{\frac{T}{q_{\min}}}$$

Lower bound proof: No algorithm can do significantly better than an algorithm which eliminates low probability types and then focuses on the rest.

Wenshuo Guo

Nika Haghtalab

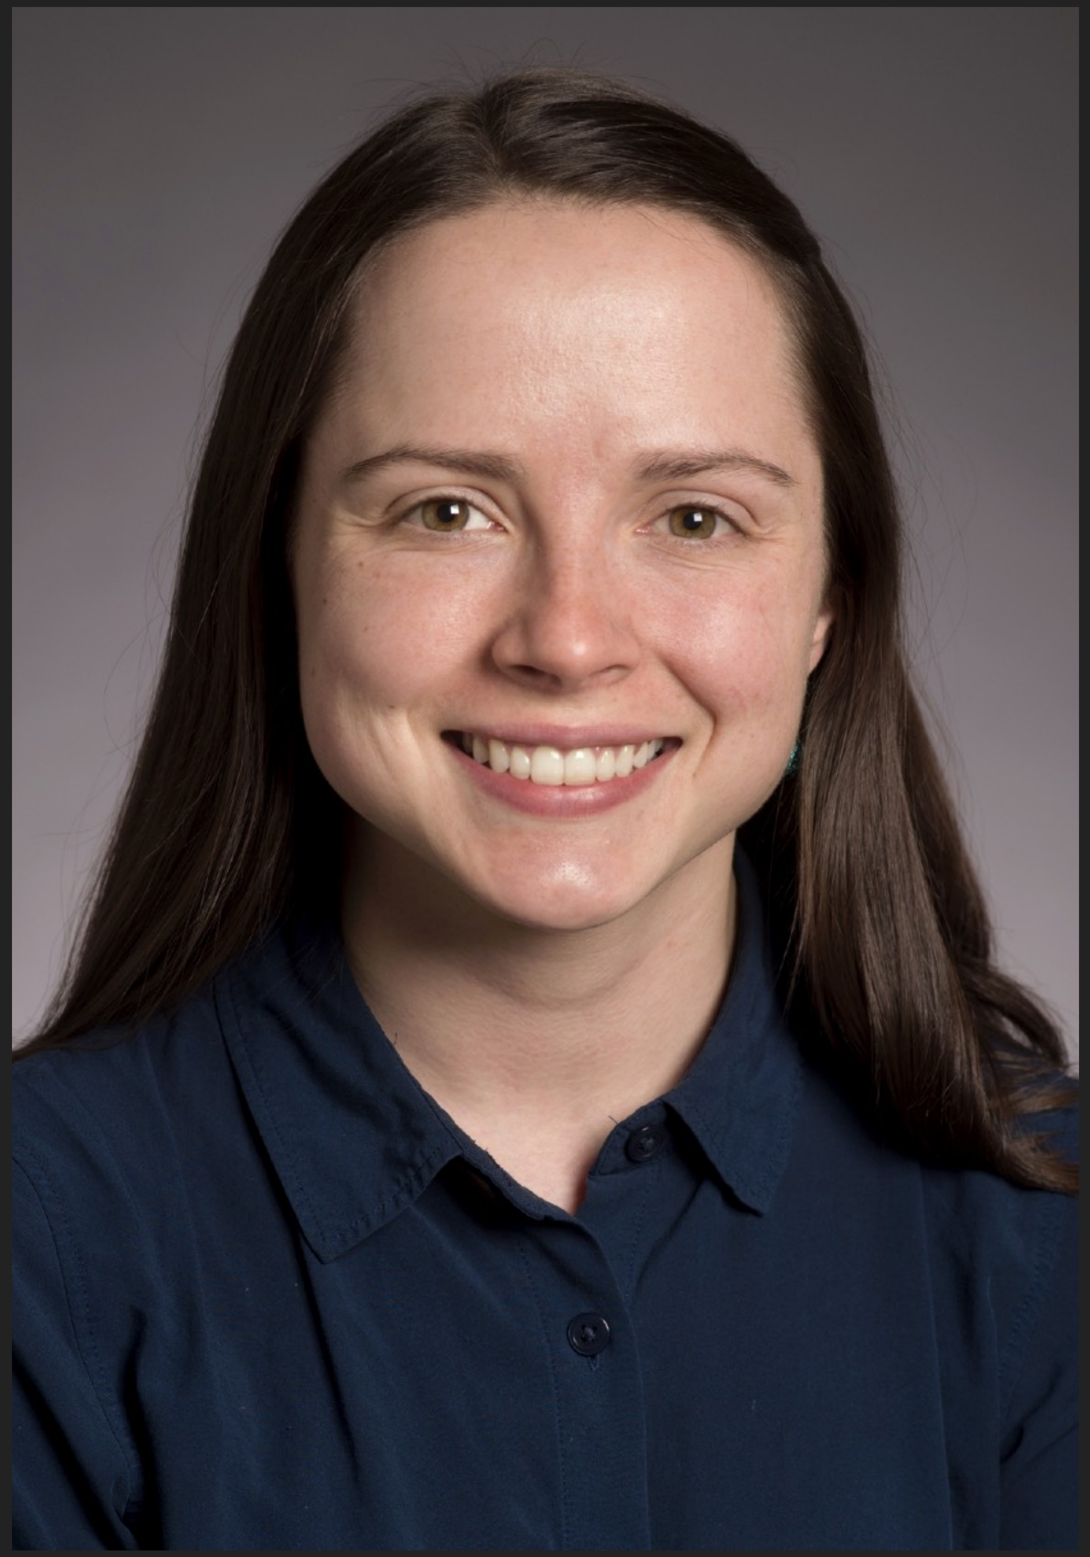

Ellen Vitercik

THANK YOU!

SUMMARY

- ***Challenges:*** Setting high prices for high instantaneous revenue
  $\implies$ Both buyer and seller cannot learn
  $\implies$ Poor revenue in the long run

- ***Algorithm:*** Choose low prices early, and increase them gradually.

- ***Theoretical Results:***
  - Upper bound: $\tilde{\mathcal{O}}\left(d^{1/3}T^{2/3}\right)$ worst case regret, but $\tilde{\mathcal{O}}\left(T^{1/2}\right)$ regret when all types appear frequently.
  - Matching lower bounds.