LEVERAGING REVIEWS: LEARNING TO PRICE WITH BUYER AND SELLER UNCERTAINTY

WSB SEMINAR, FEBRUARY 24 2023

KIRTHEVASAN KANDASAMY

DEPARTMENT OF COMPUTER SCIENCES, UW-MADISON

JOINT WORK WITH: WENSHUO GUO, NIKA HAGHTALAB, ELLEN VITERCIK

ONLINE MARKETPLACES ARE UBIQUITOUS

shopify

ebay

DOORDASH

amazon

ONLINE MARKETPLACES ARE UBIQUITOUS

- 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’

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'

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 '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:**

HOW CAN REVIEWS BE HELPFUL?

- **To buyers:**
  - Understand if the product is right for them.

HOW CAN REVIEWS BE HELPFUL?

- **To buyers:**
  - 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:**
  - 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:**

HOW CAN REVIEWS BE HELPFUL?

- **To buyers:**
  - 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.

HOW CAN REVIEWS BE HELPFUL?

- **To buyers:**
  - 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! We should increase the price.

GOAL

GOAL

- Study how reviews can help *both* sides of the market.

GOAL

- Study how reviews can help *both* sides of the market.
  - Sellers will **learn** optimal price.

GOAL

- Study how reviews can help *both* sides of the market.
  - Sellers will **learn** optimal price.
  - Buyers will **learn** their *value* for goods.

GOAL

- Study how reviews can help *both* sides of the market.
  - Sellers will **learn** optimal price.
  - Buyers will **learn** their *value* for goods.

- **Model:** Several market models …

GOAL

- Study how reviews can help *both* sides of the market.
  - Sellers will **learn** optimal price.
  - Buyers will **learn** their *value* for goods.

- **Model:** Several market models …
  - In this work: **posted-price mechanisms.**

GOAL

- Study how reviews can help *both* sides of the market.
  - Sellers will **learn** optimal price.
  - Buyers will **learn** their *value* for goods.

- **Model:** Several market models …
  - In this work: **posted-price mechanisms.**
  - **Prior work on feedback-driven market/auction design:** single-item auctions (*FPS '18, WPR'16, PPPR '22, ADG '16, DSS '19*), posted price mechanisms when buyers know values (*KL '03*), VCG mechanisms (*KGJS, JMLR '22*), matching markets *(LMJ, AISTATS '19)*, exchange economies (*GKGJS, AISTATS '22*), and several more …

POSTED PRICE MECHANISM

POSTED PRICE MECHANISM

- A single seller who has (an infinite amount) of a single item.

POSTED PRICE MECHANISM

- A single seller who has (an infinite amount) of a single item.
- Each buyer has a type i. All buyers of type i have **value** θ_i (in \$) for the item.

POSTED PRICE MECHANISM

- A single seller who has (an infinite amount) of a single item.
- Each buyer has a type i. All buyers of type i have **value** θ_i (in \$) for the item.
- There is a **distribution** $\mathscr{P}$ of buyer types, i.e $i \sim \mathscr{P}$.

POSTED PRICE MECHANISM

- A single seller who has (an infinite amount) of a single item.
- Each buyer has a type i. All buyers of type i have **value** θ_i (in \$) for the item.
- There is a **distribution** $\mathscr{P}$ of buyer types, i.e $i \sim \mathscr{P}$.

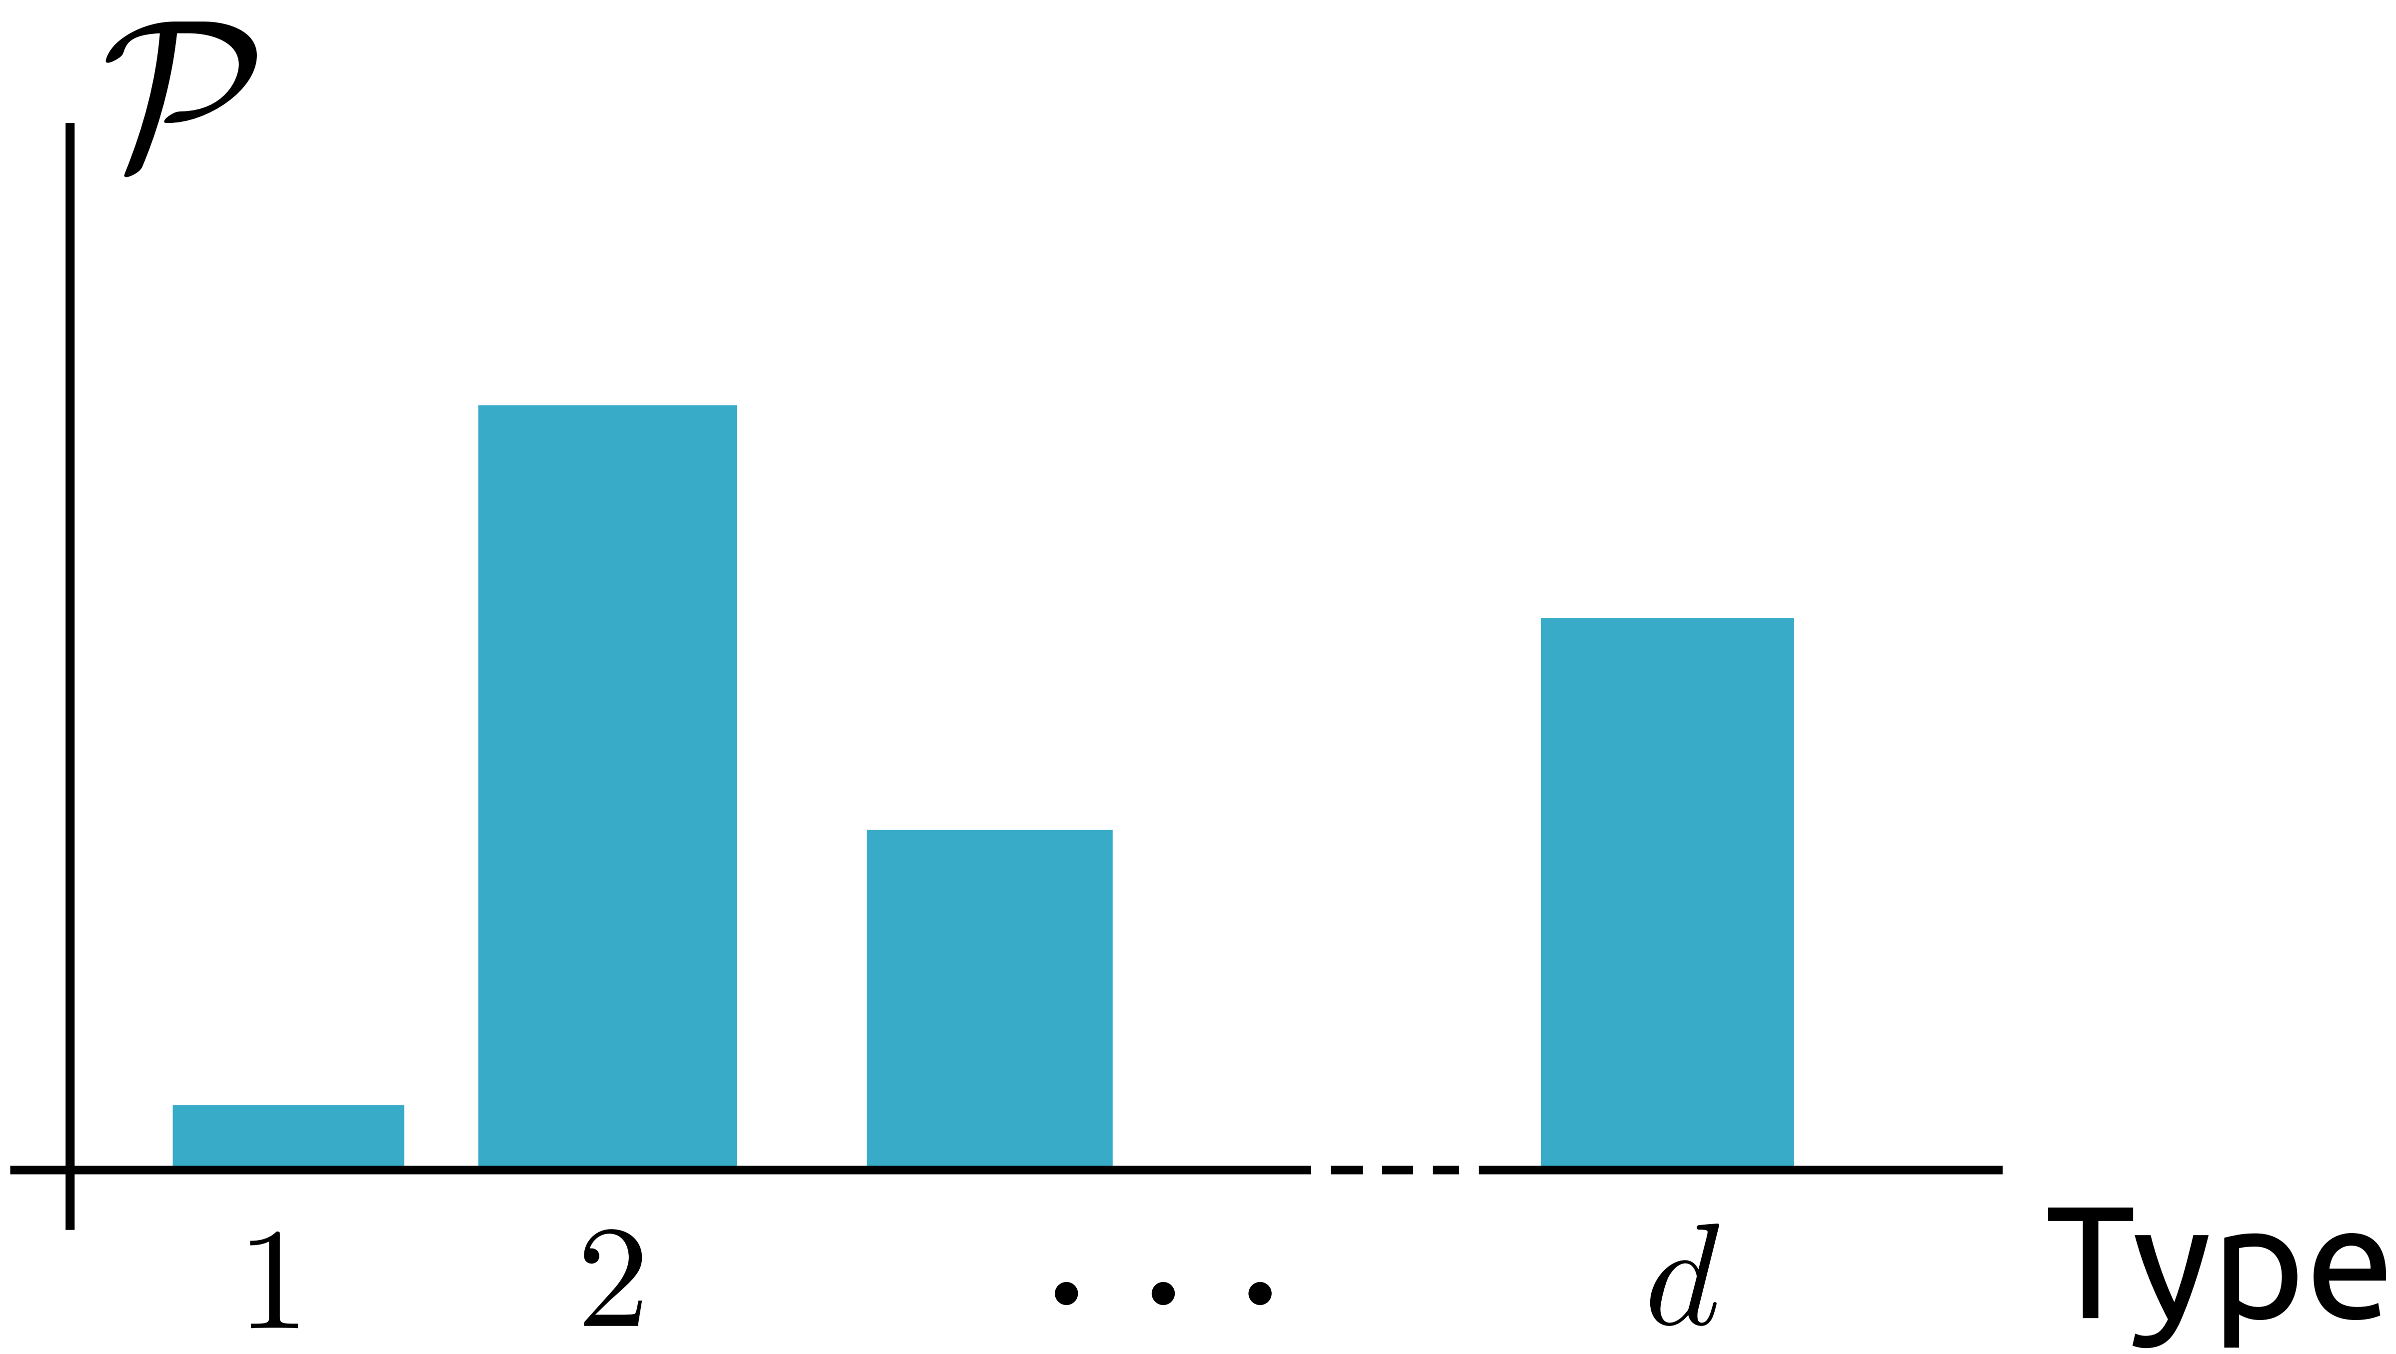

POSTED PRICE MECHANISM

- A single seller who has (an infinite amount) of a single item.
- Each buyer has a type i. All buyers of type i have **value** θ_i (in \$) for the item.
- There is a **distribution** $\mathscr{P}$ of buyer types, i.e $i \sim \mathscr{P}$.

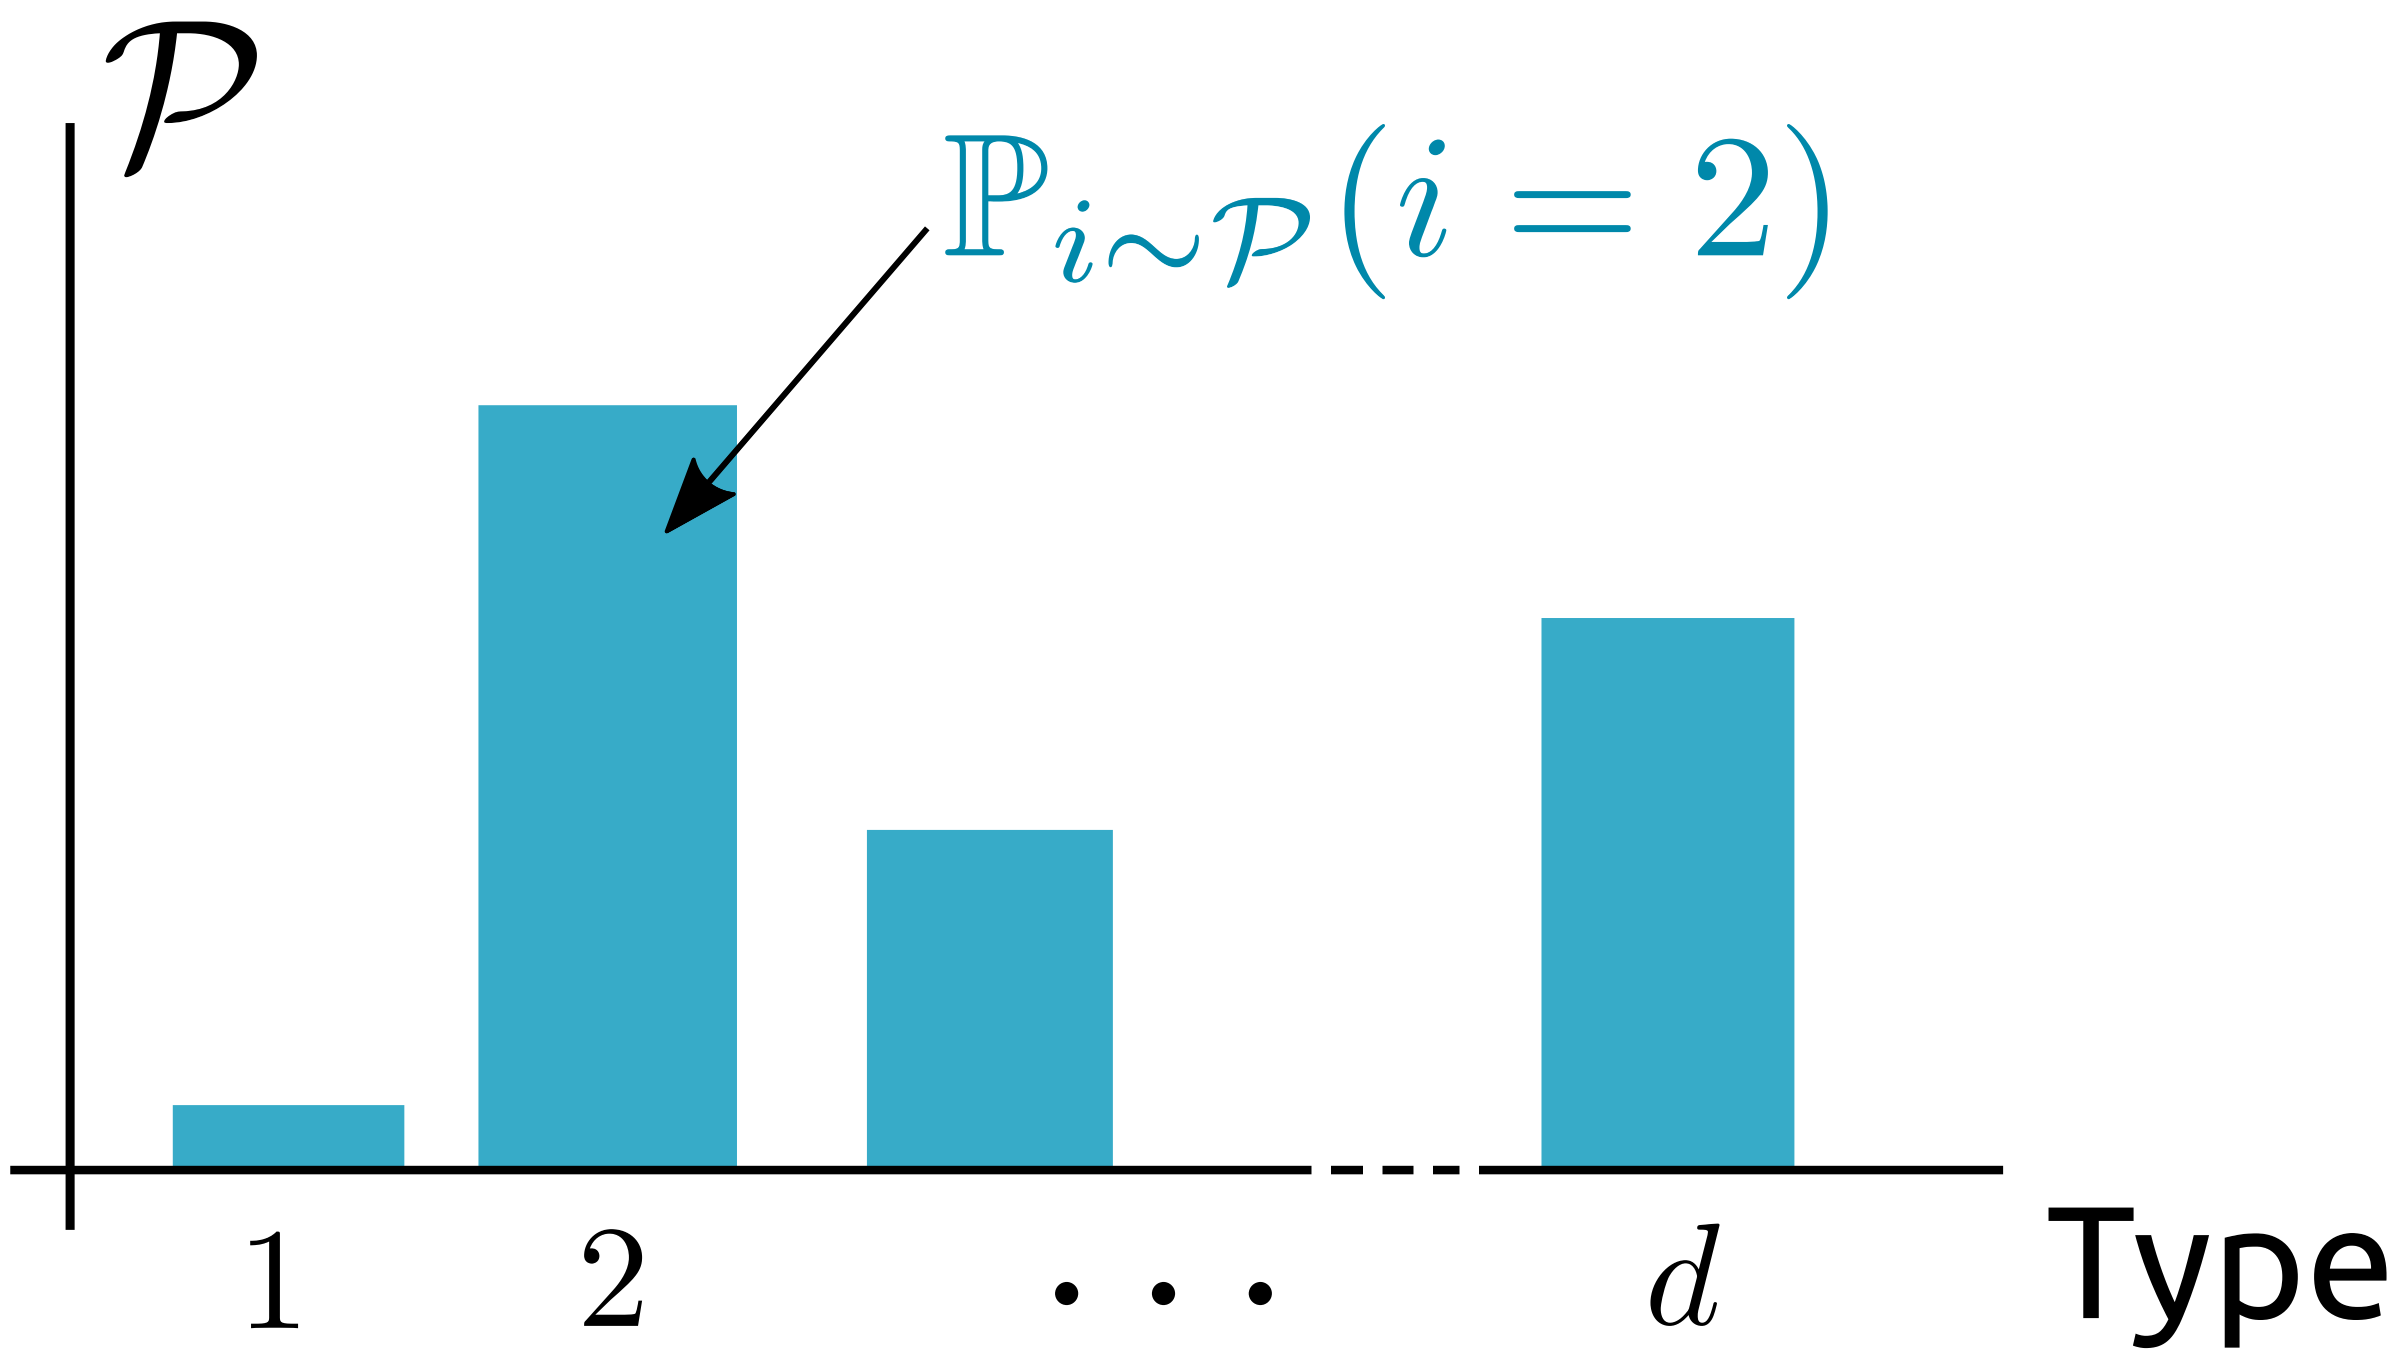

POSTED PRICE MECHANISM

- A single seller who has (an infinite amount) of a single item.
- Each buyer has a type i. All buyers of type i have **value** θ_i (in \$) for the item.
- There is a **distribution** $\mathscr{P}$ of buyer types, i.e $i \sim \mathscr{P}$.

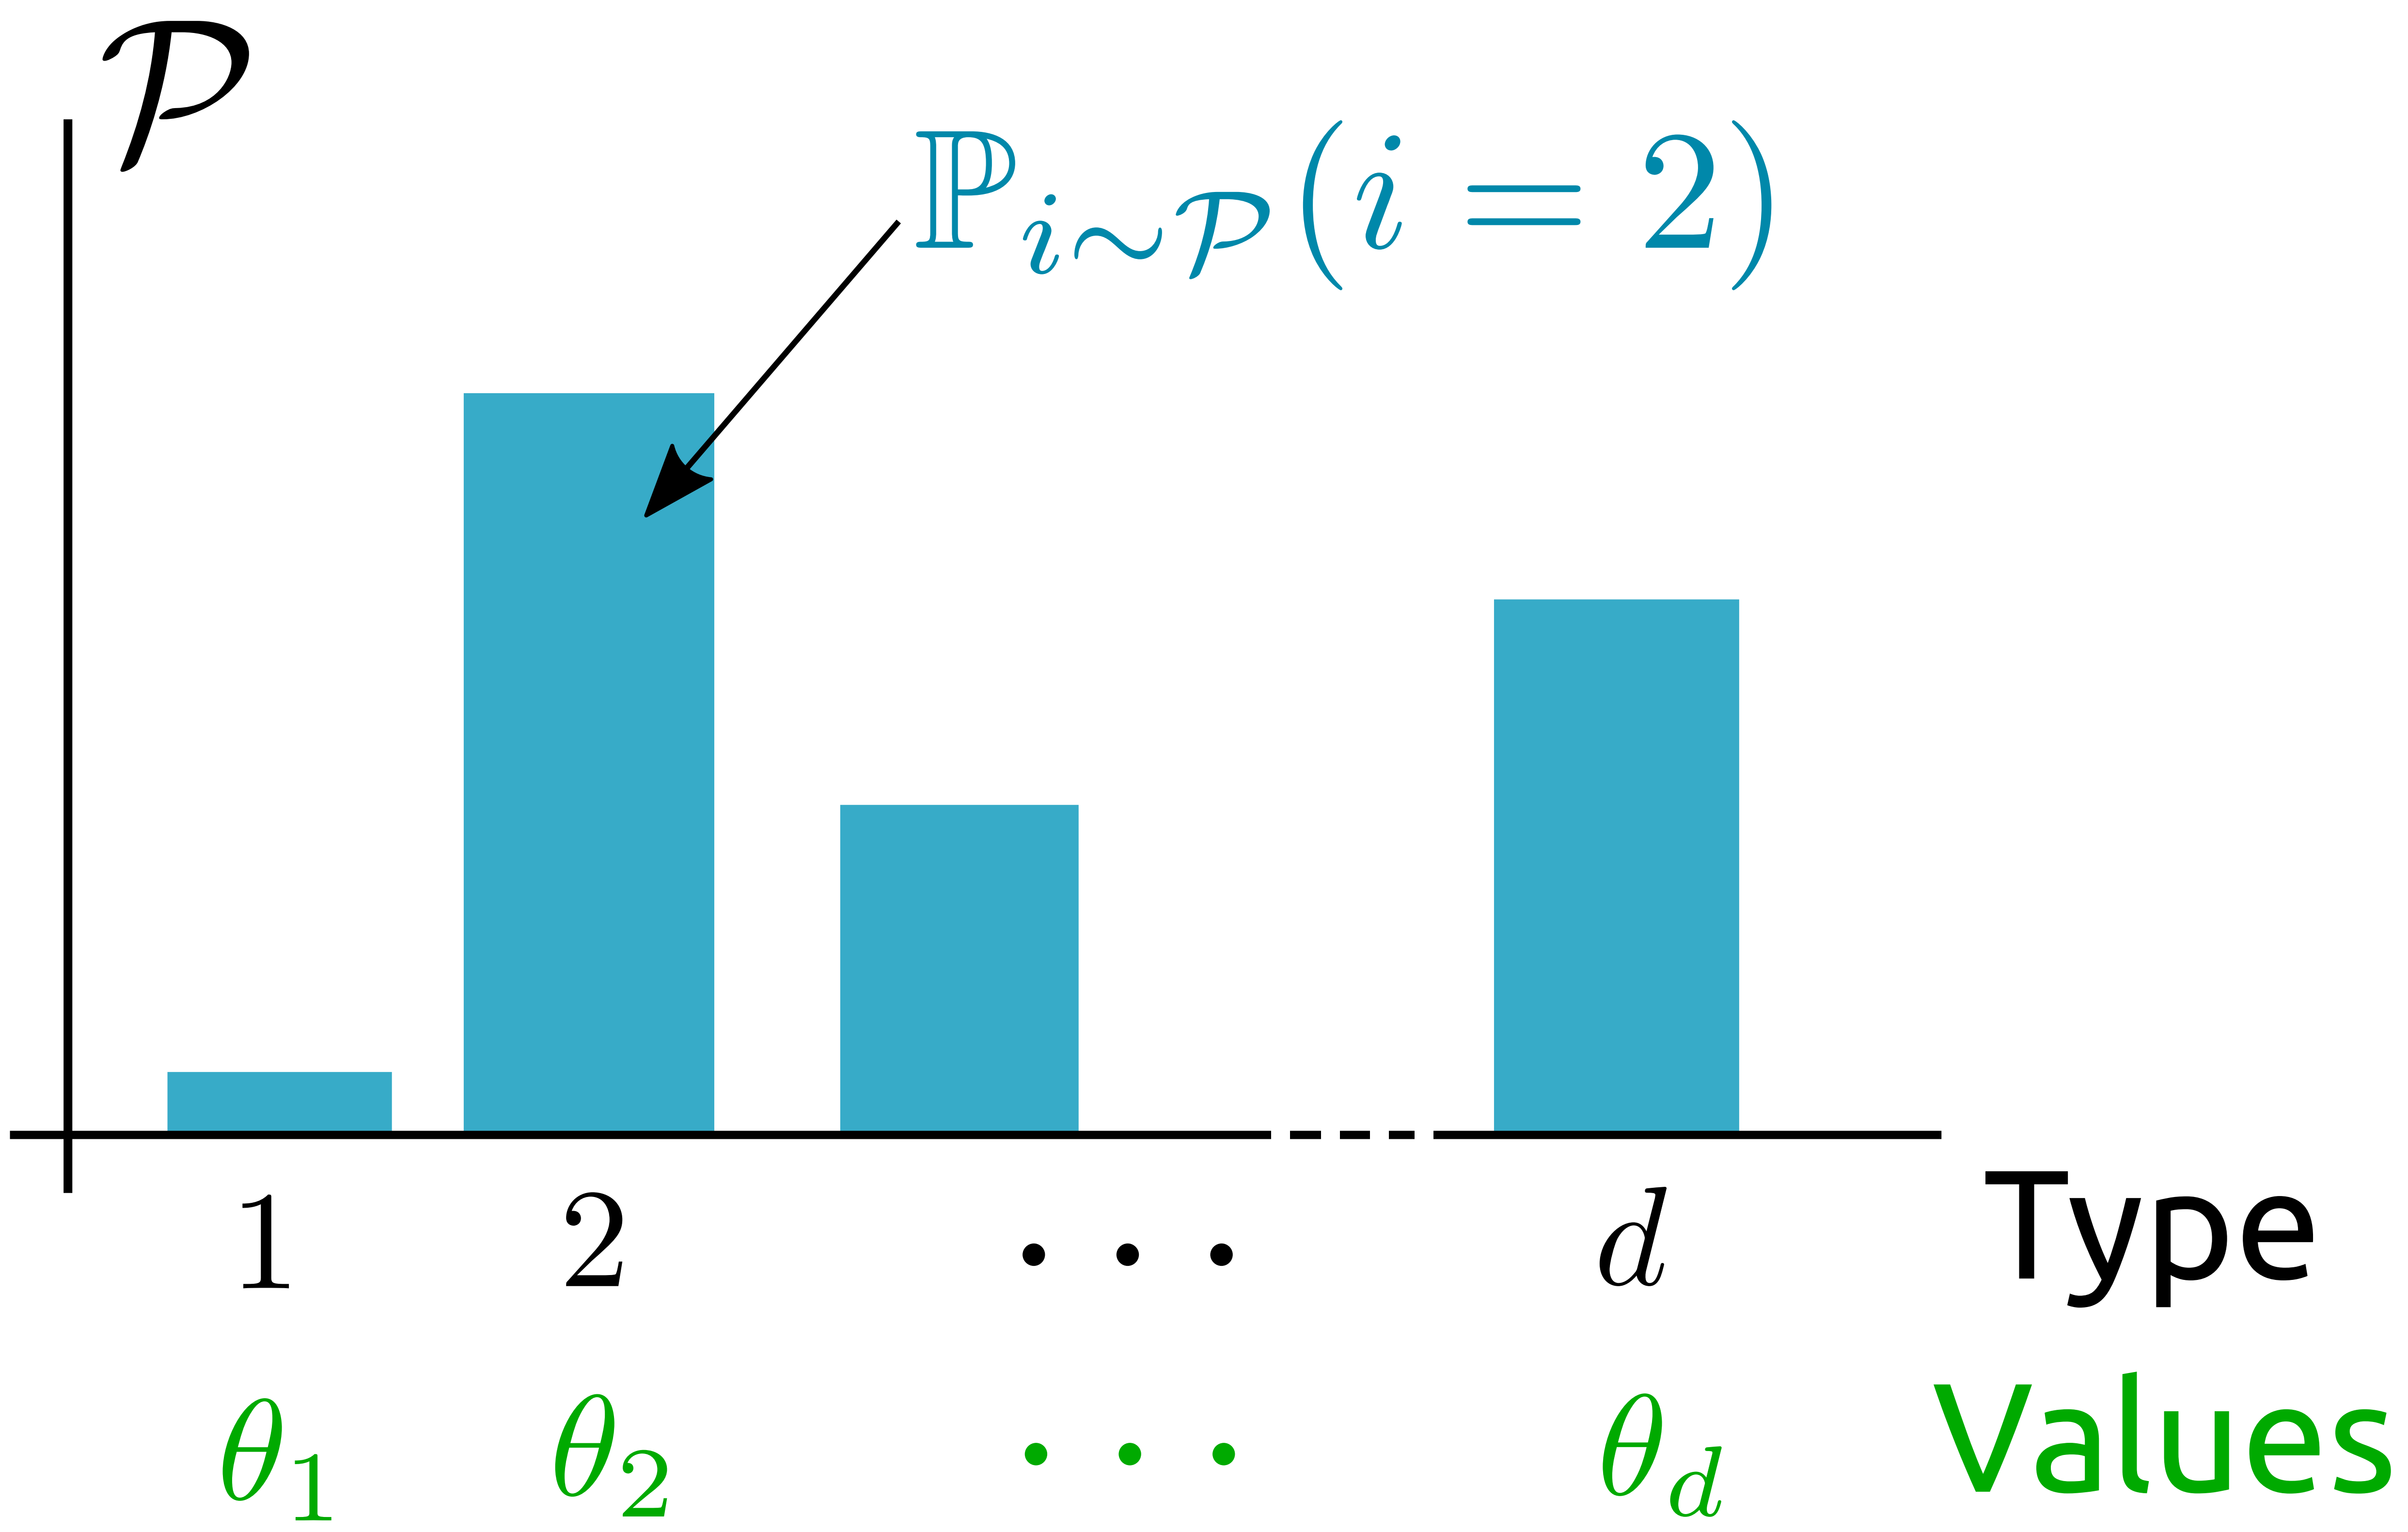

POSTED PRICE MECHANISM

- Seller posts a price p for the item. A buyer of type i will purchase if their value is larger than the price, i.e $\theta_i \geq p$.

POSTED PRICE MECHANISM

- Seller posts a price p for the item. A buyer of type i will purchase if their value is larger than the price, i.e $\theta_i \geq p$.

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

Visit the Cuisinart Store

14,945 ratings

4.7 out of 5

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 ›

POSTED PRICE MECHANISM

- Seller posts a price p for the item. A buyer of type i will purchase if their value is larger than the price, i.e $\theta_i \geq p$.

Price: $40

POSTED PRICE MECHANISM

- Seller posts a price p for the item. A buyer of type i will purchase if their value is larger than the price, i.e $\theta_i \geq p$.

Price: $40

Type 1

I will use it mostly
in the oven.

I value this pot at $20.

POSTED PRICE MECHANISM

- Seller posts a price p for the item. A buyer of type i will purchase if their value is larger than the price, i.e $\theta_i \geq p$.

Price: $40

Type 1

I will use it mostly in the oven.

I value this pot at $20.

Type 2

I will use it mostly for stove-top cooking.

I value this pot at $50.

MAXIMIZING REVENUE IN POSTED PRICE MECHANISMS

MAXIMIZING REVENUE IN POSTED PRICE MECHANISMS

- How does a seller choose a price, given the type values and type distribution $\mathscr{P}$?

MAXIMIZING REVENUE IN POSTED PRICE MECHANISMS

- How does a seller choose a price, given the type values and type distribution $\mathscr{P}$?

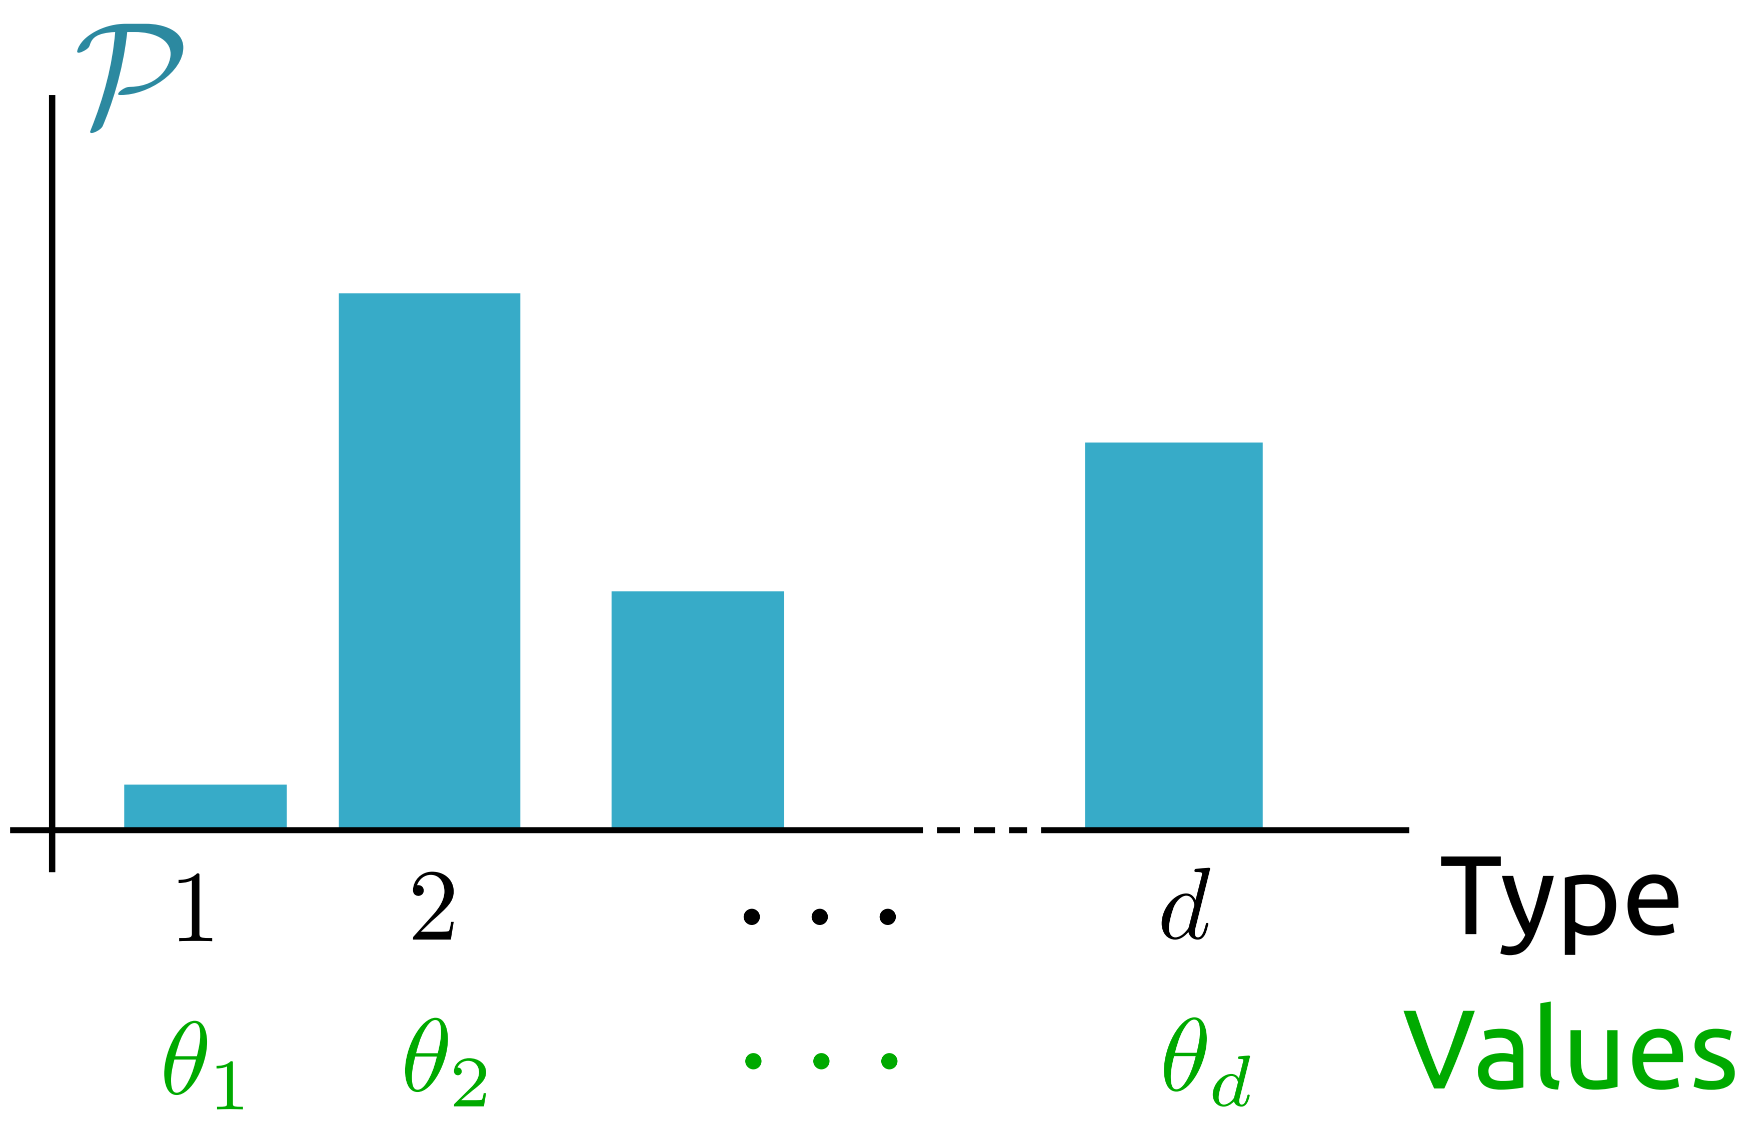

MAXIMIZING REVENUE IN POSTED PRICE MECHANISMS

- How does a seller choose a price, given the type values and type distribution $\mathscr{P}$?

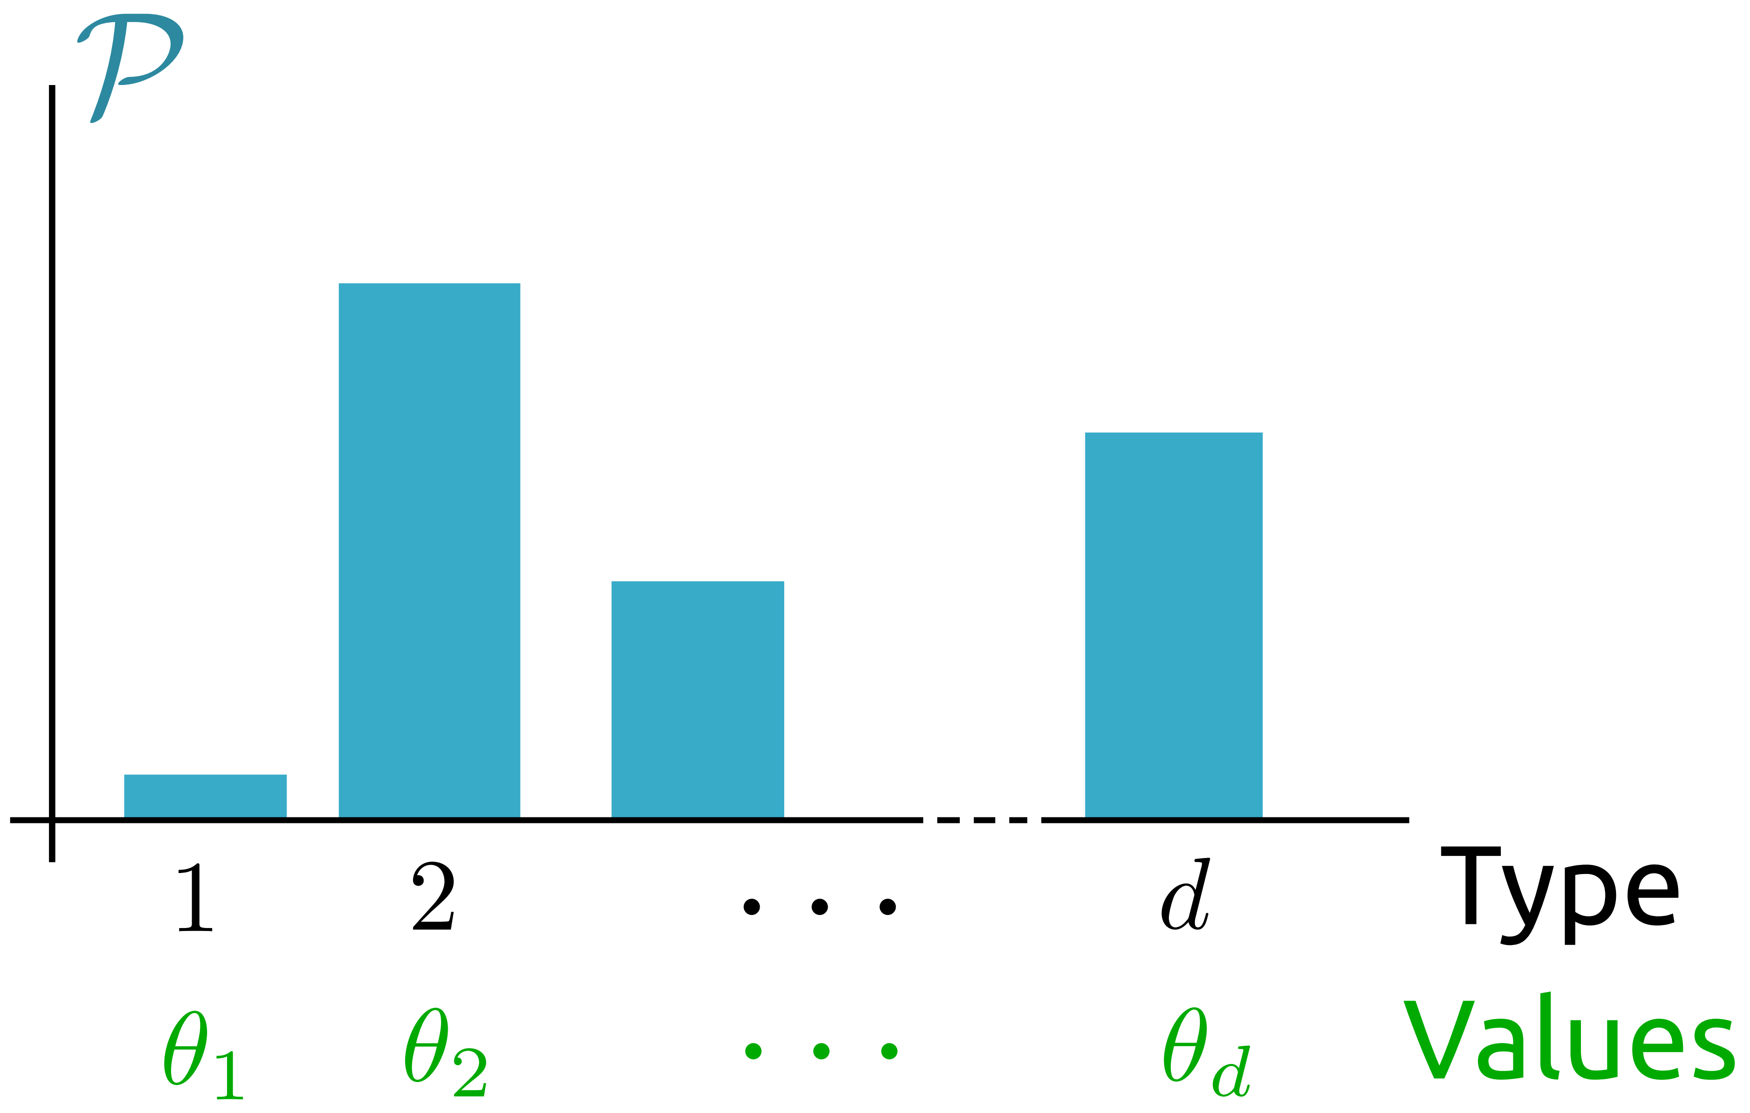

- Expected revenue *per purchase* if you set price p,

- How does a seller choose a price, given the type values and type distribution $\mathscr{P}$?

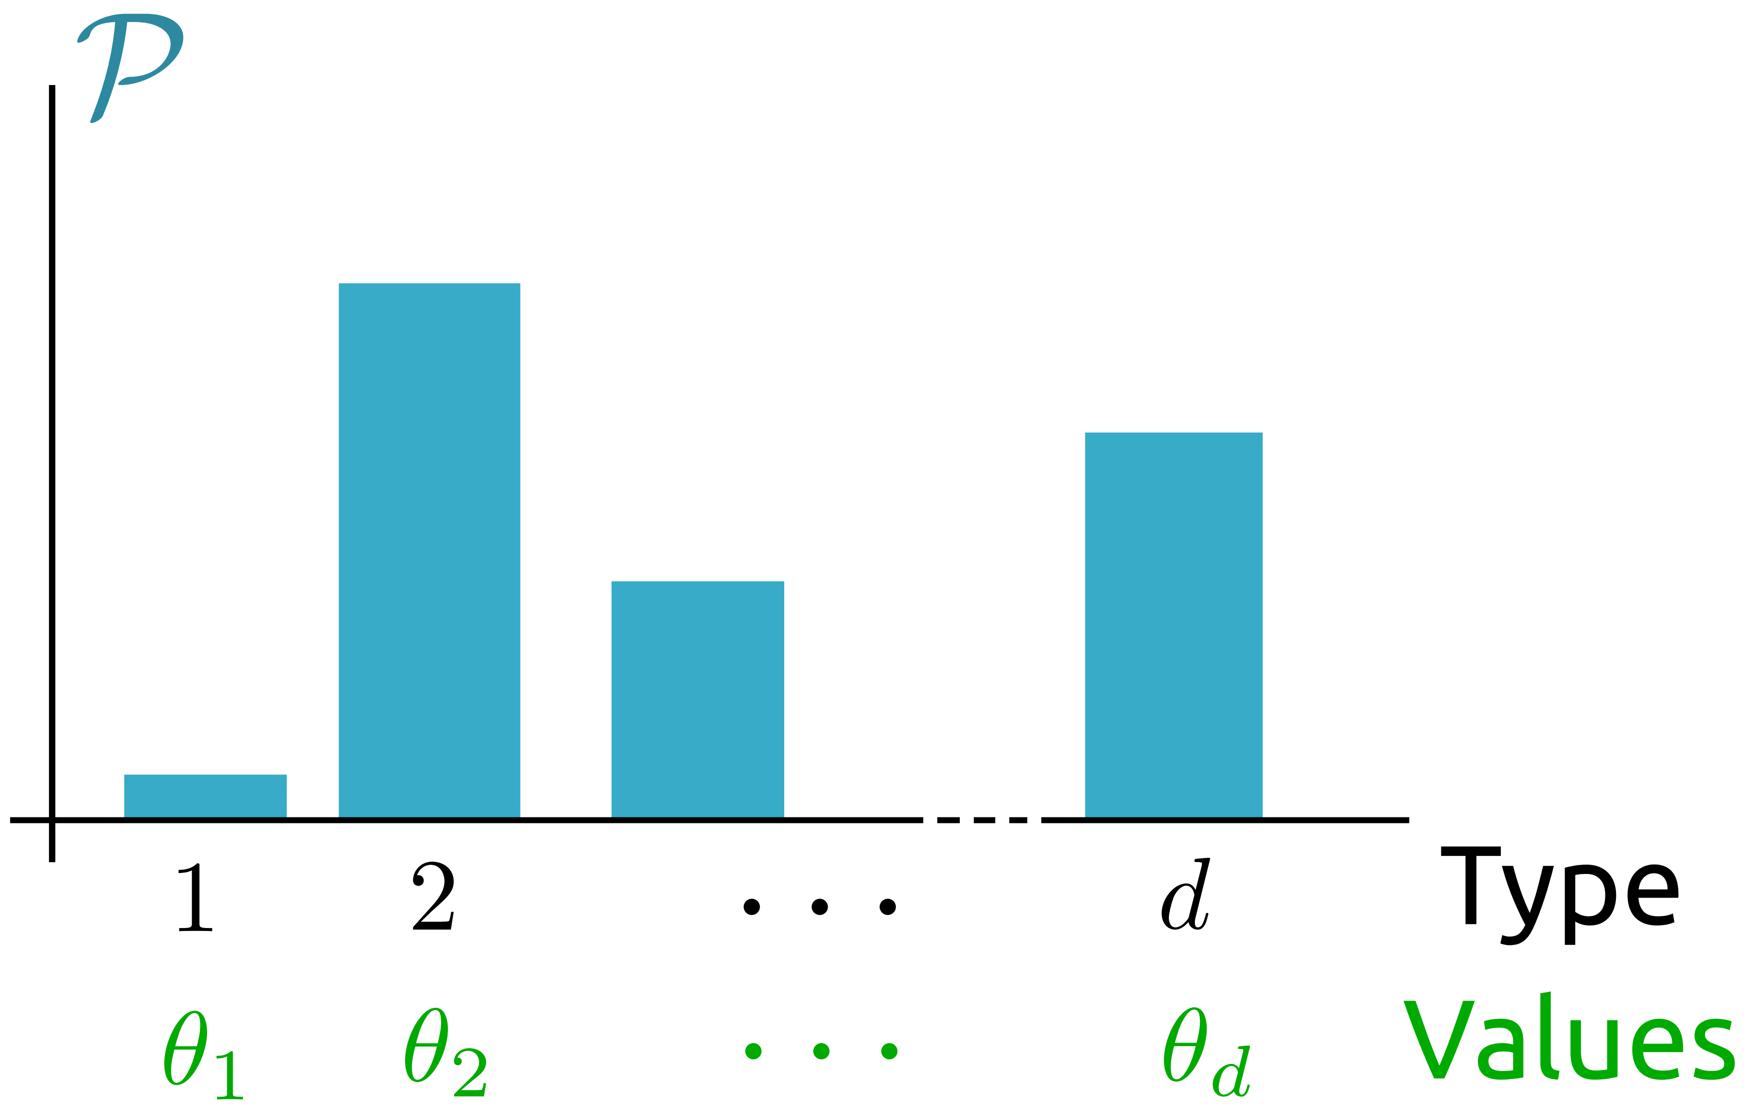

- Expected revenue *per purchase* if you set price p,

$$\mathrm{rev}(p) = p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

MAXIMIZING REVENUE IN POSTED PRICE MECHANISMS

- How does a seller choose a price, given the type values and type distribution $\mathscr{P}$?

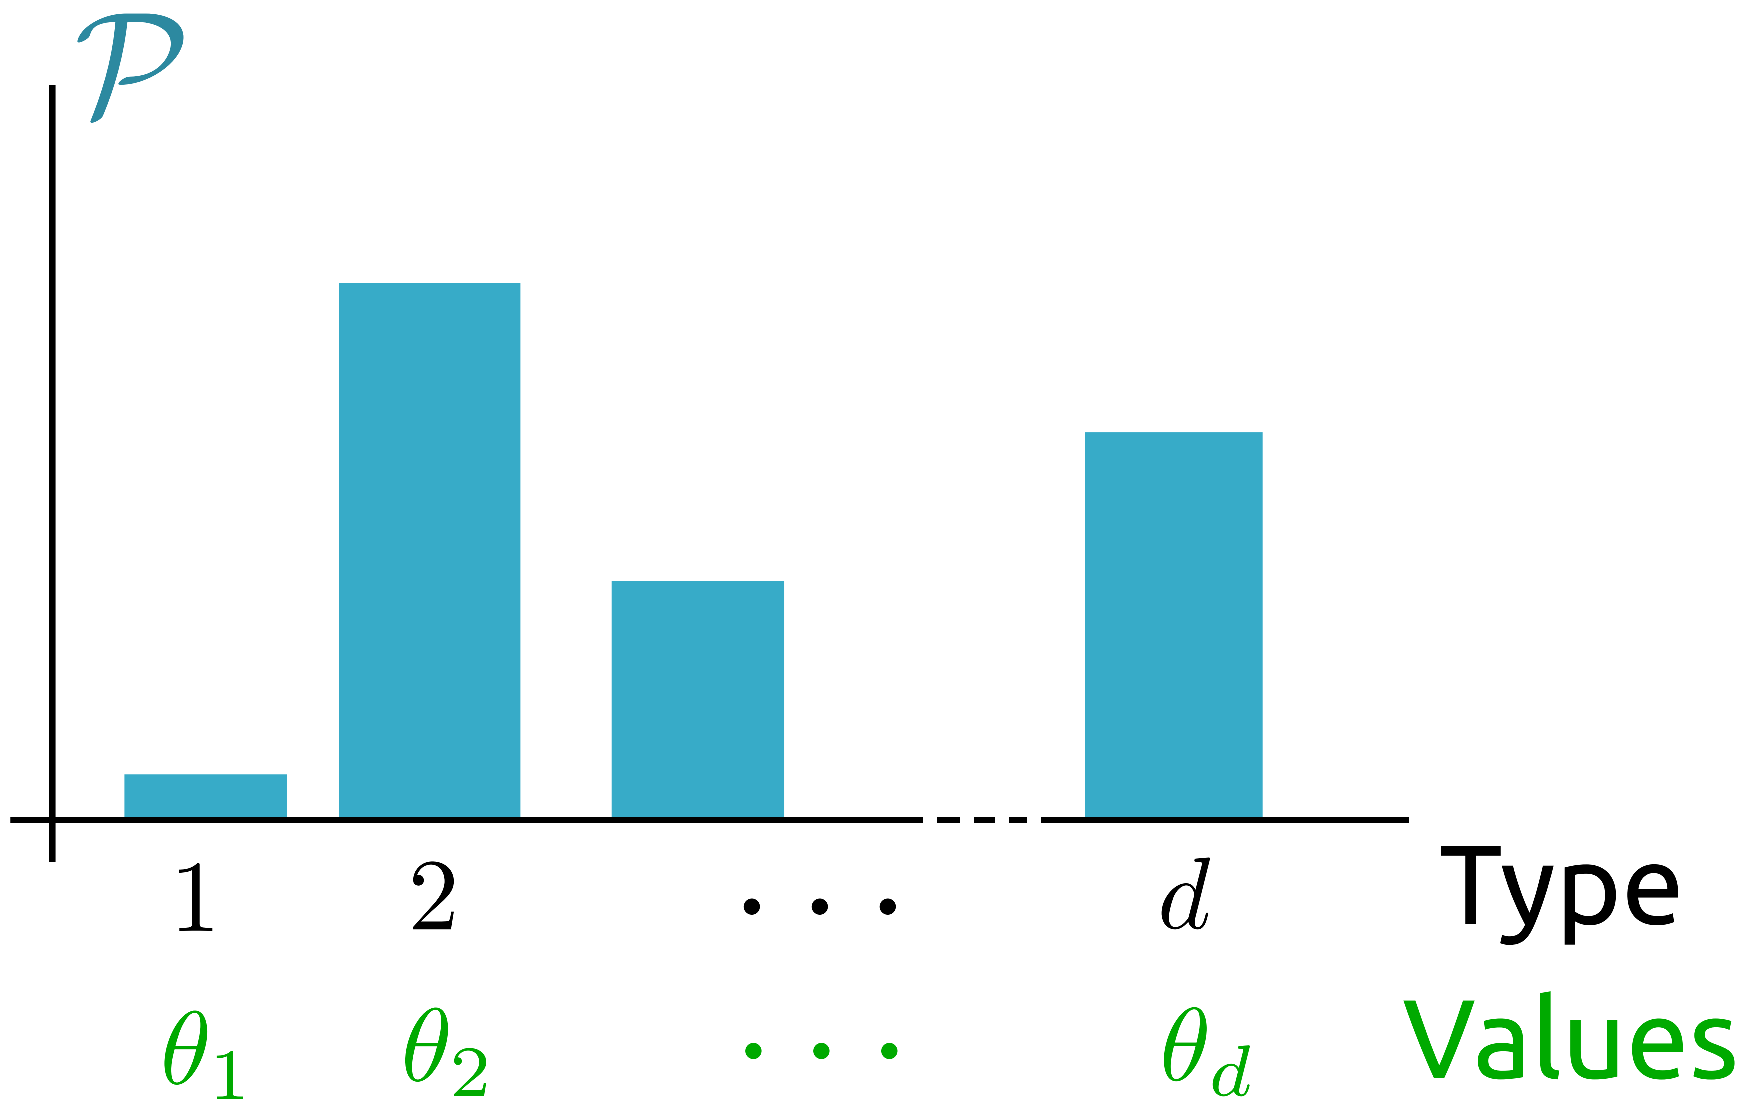

- Expected revenue *per purchase* if you set price p,

$$\text{rev}(p) = p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

- Optimal price

MAXIMIZING REVENUE IN POSTED PRICE MECHANISMS

- How does a seller choose a price, given the type values and type distribution $\mathscr{P}$?

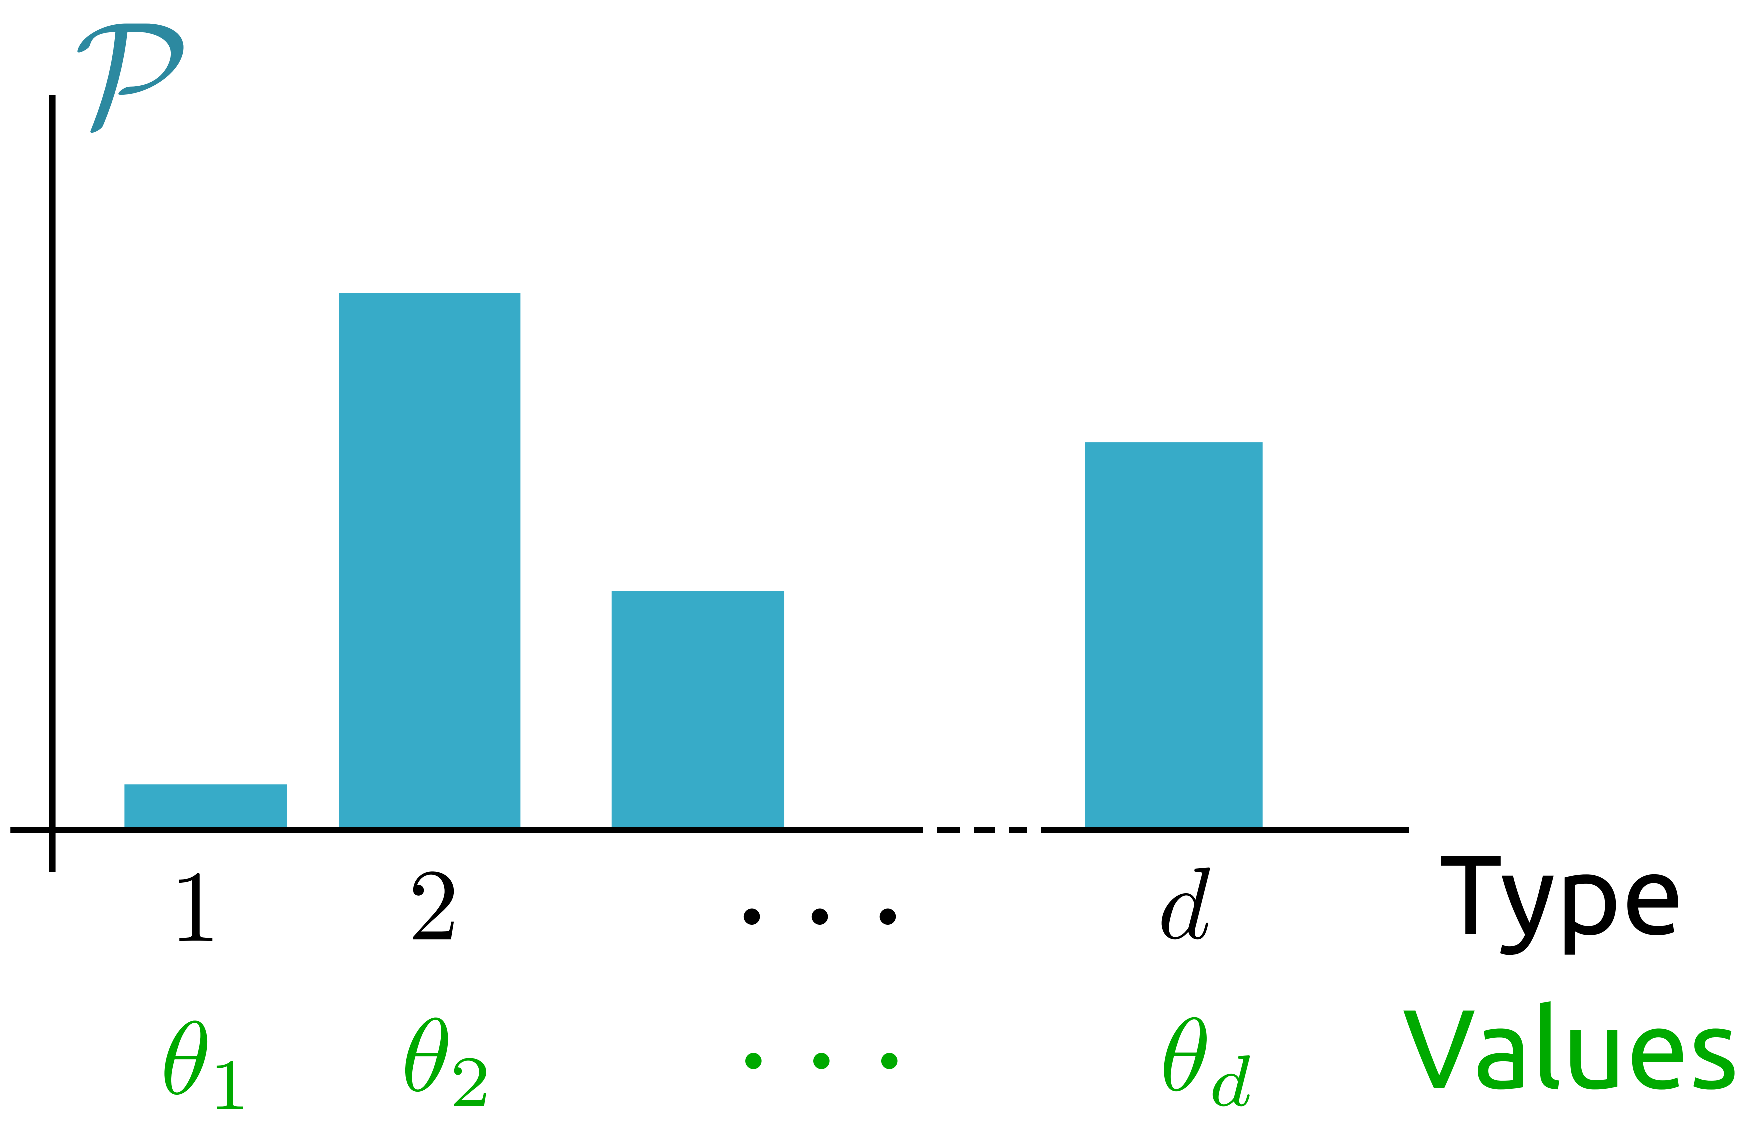

- Expected revenue *per purchase* if you set price p,

$$\mathrm{rev}(p) = p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p)$$

- Optimal price

$$p^\star = \arg\max_p \mathrm{rev}(p)$$

ISSUES

ISSUES

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

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

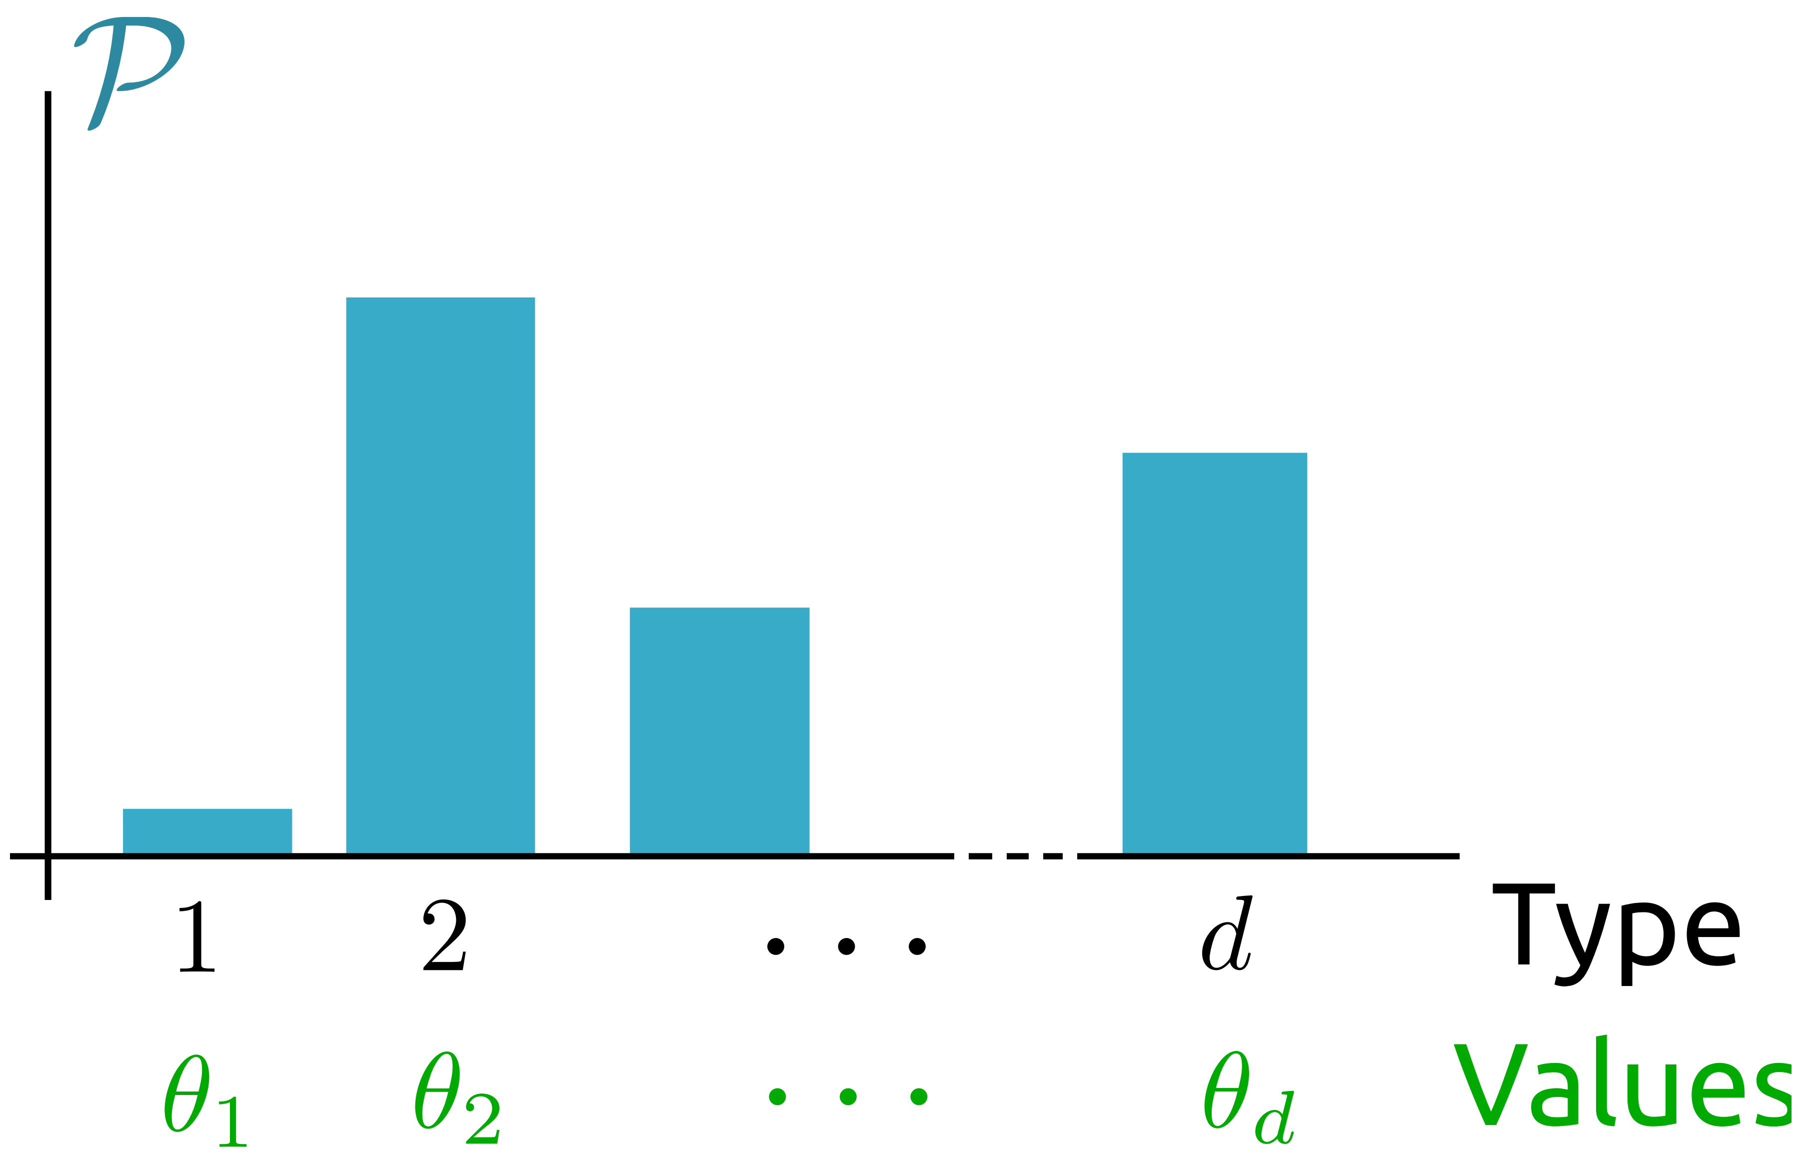

1. Seller wishes to maximize revenue, but may not know $\mathscr{P}$.

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

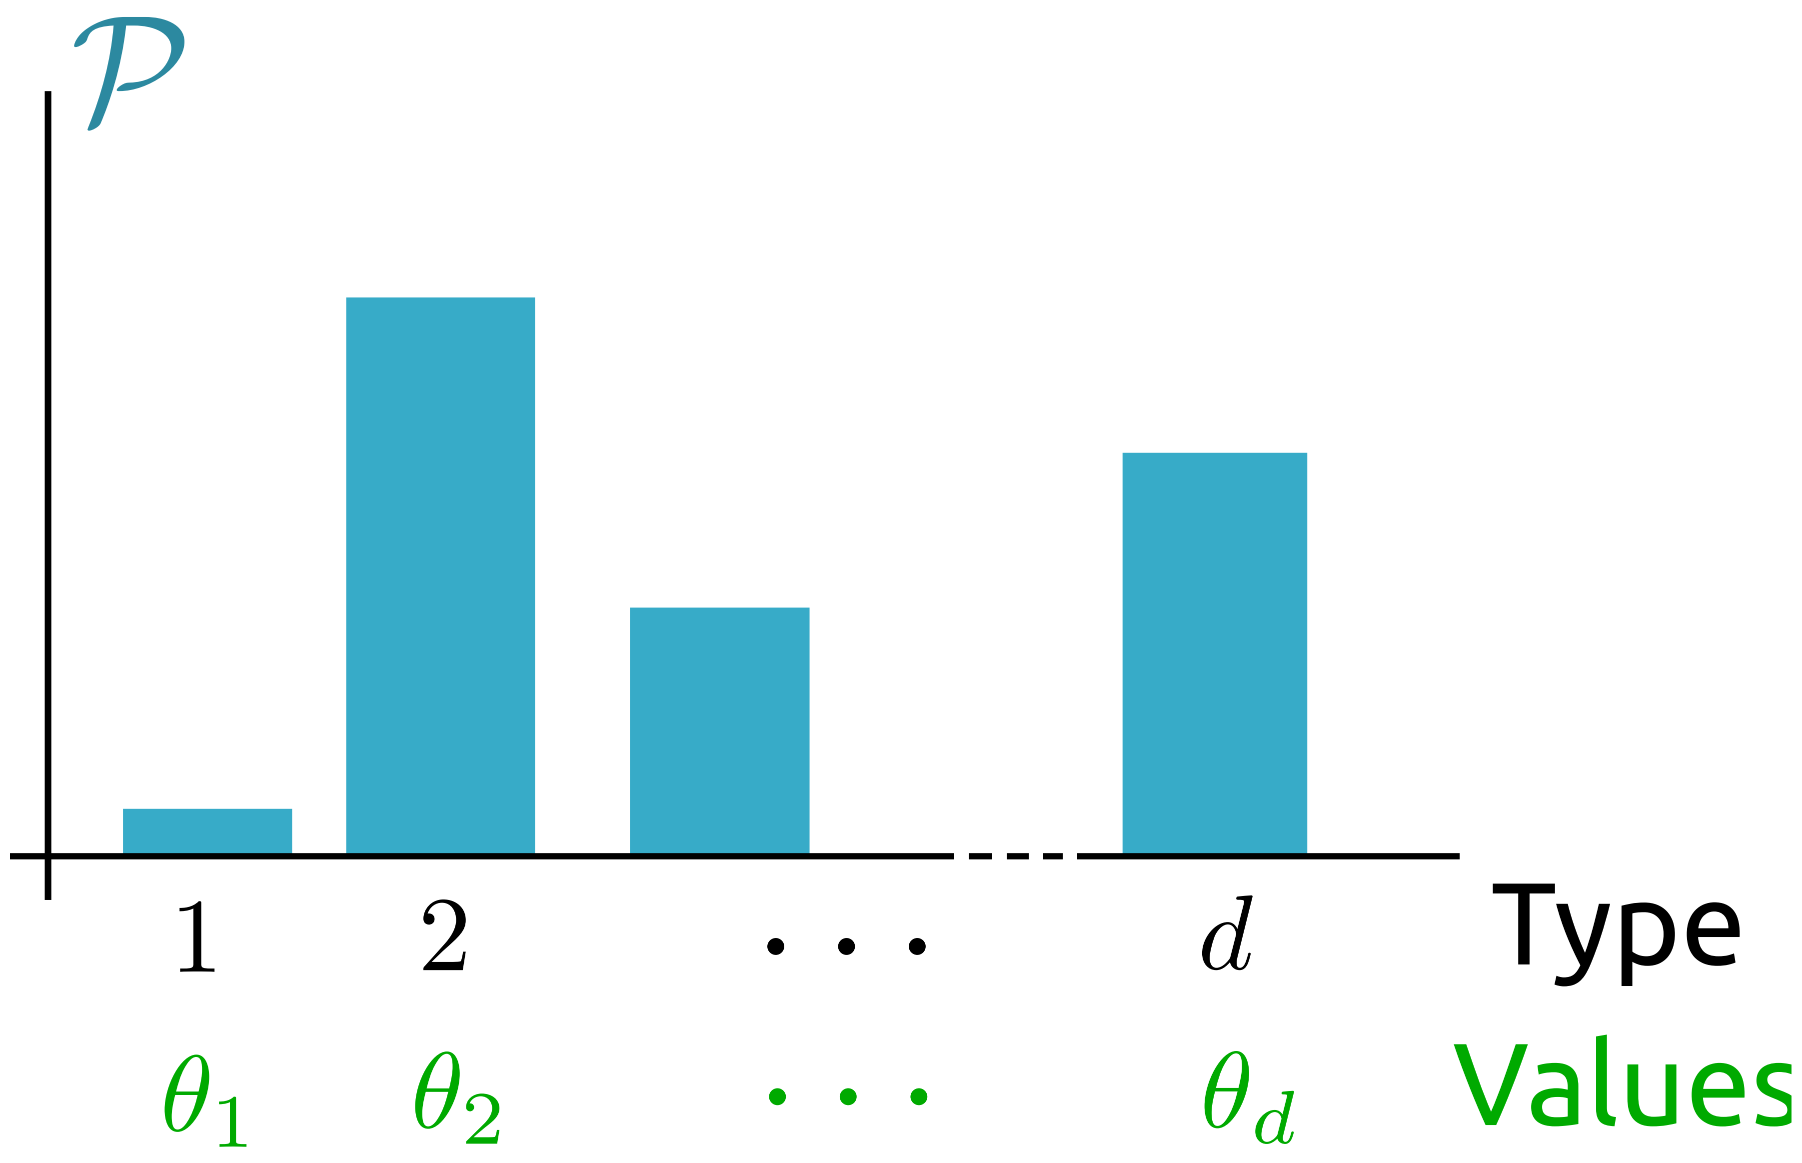

ISSUES

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

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

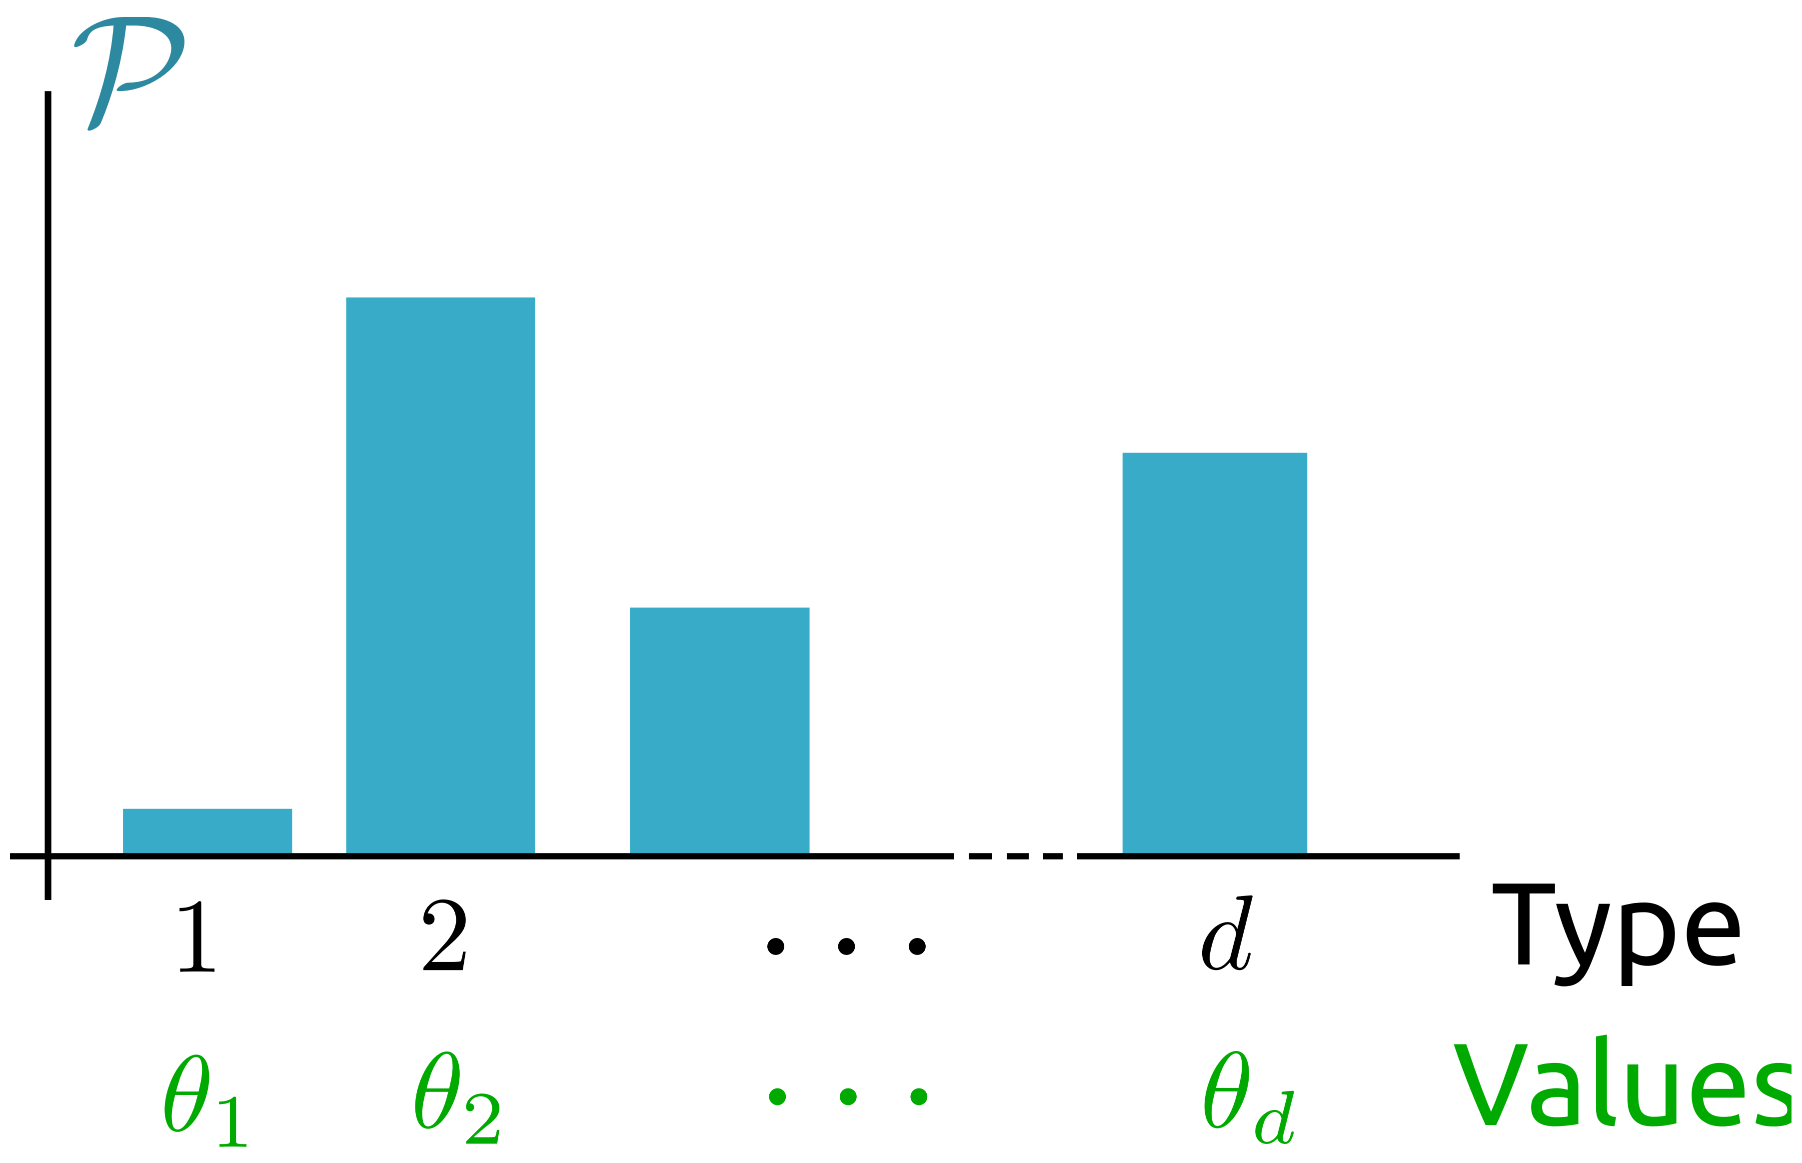

2. A buyer may know their type i, but not their value θ_i.

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

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

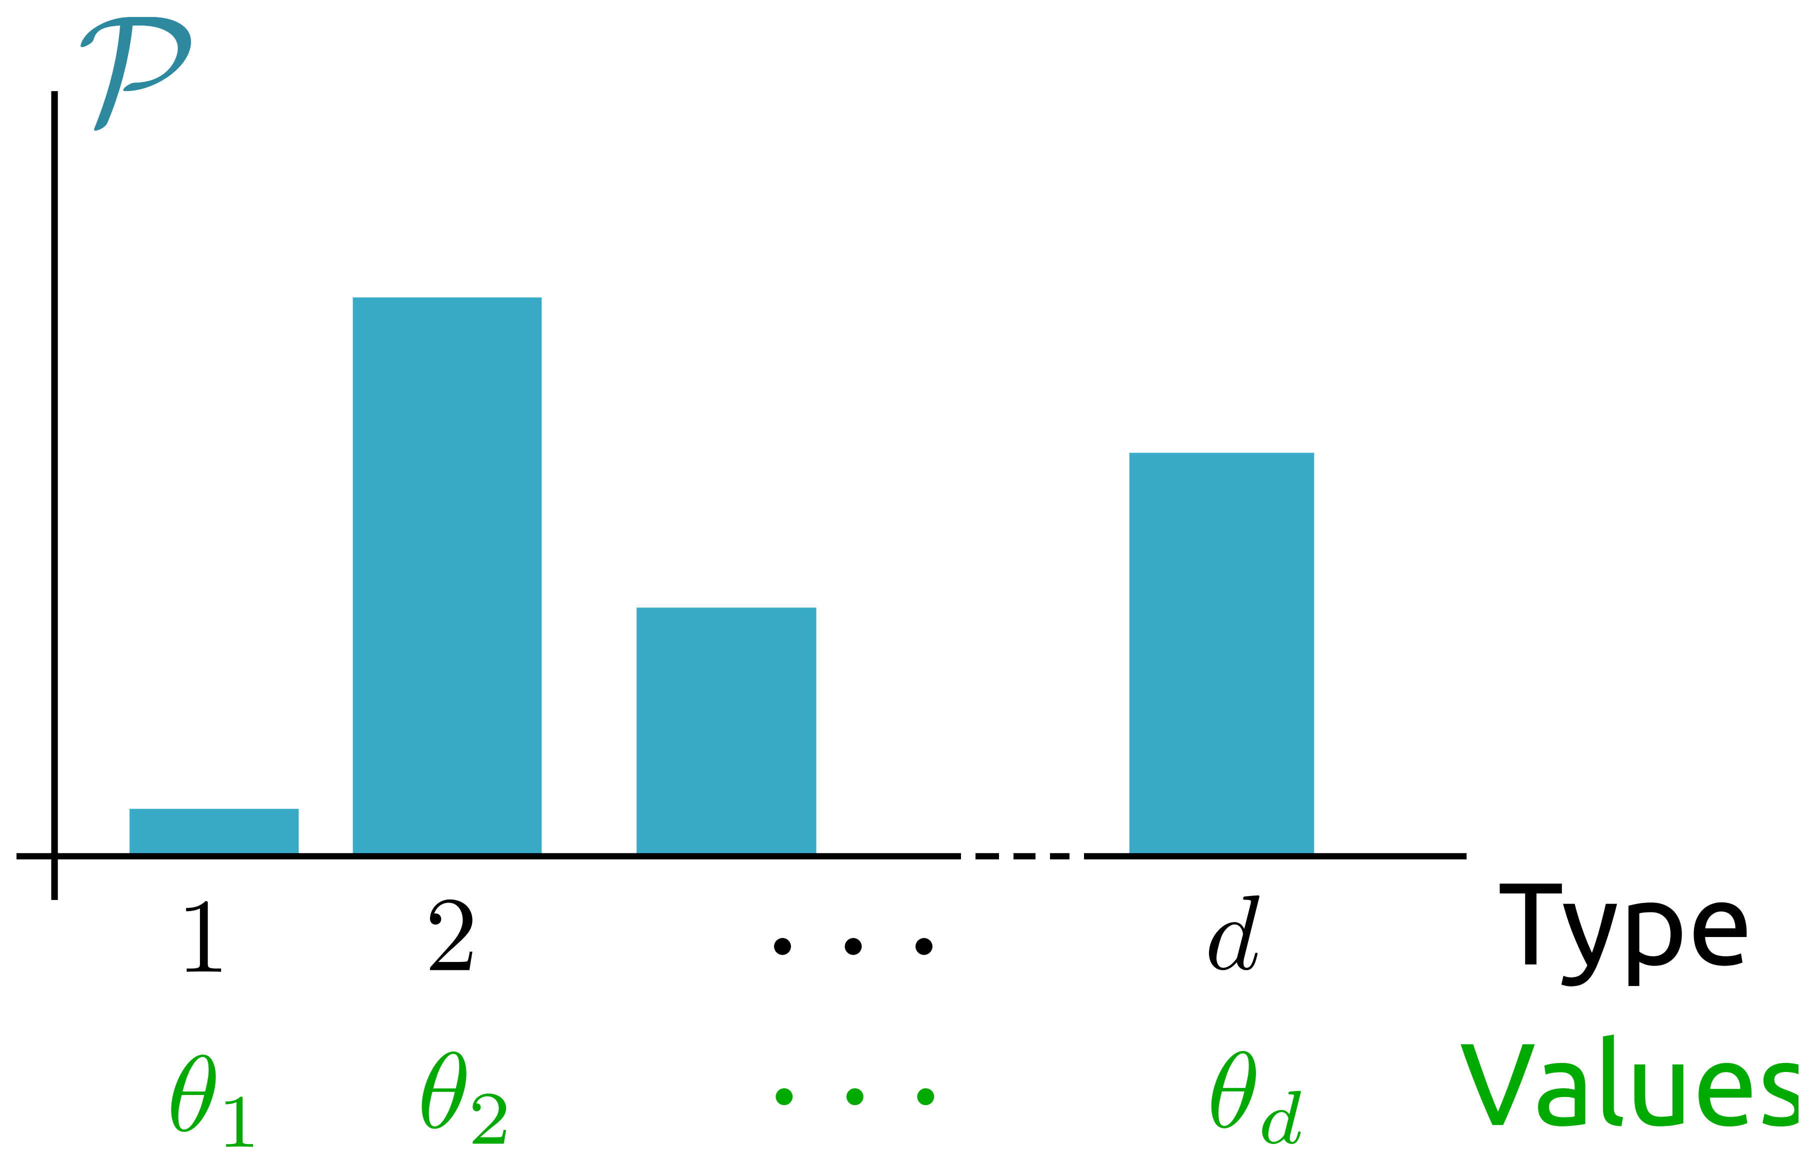

2. A buyer may know their type i, but not their value θ_i.
    - (In practice) due to uncertainty about their value, customers may not be willing to buy an item except at a low price.

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

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

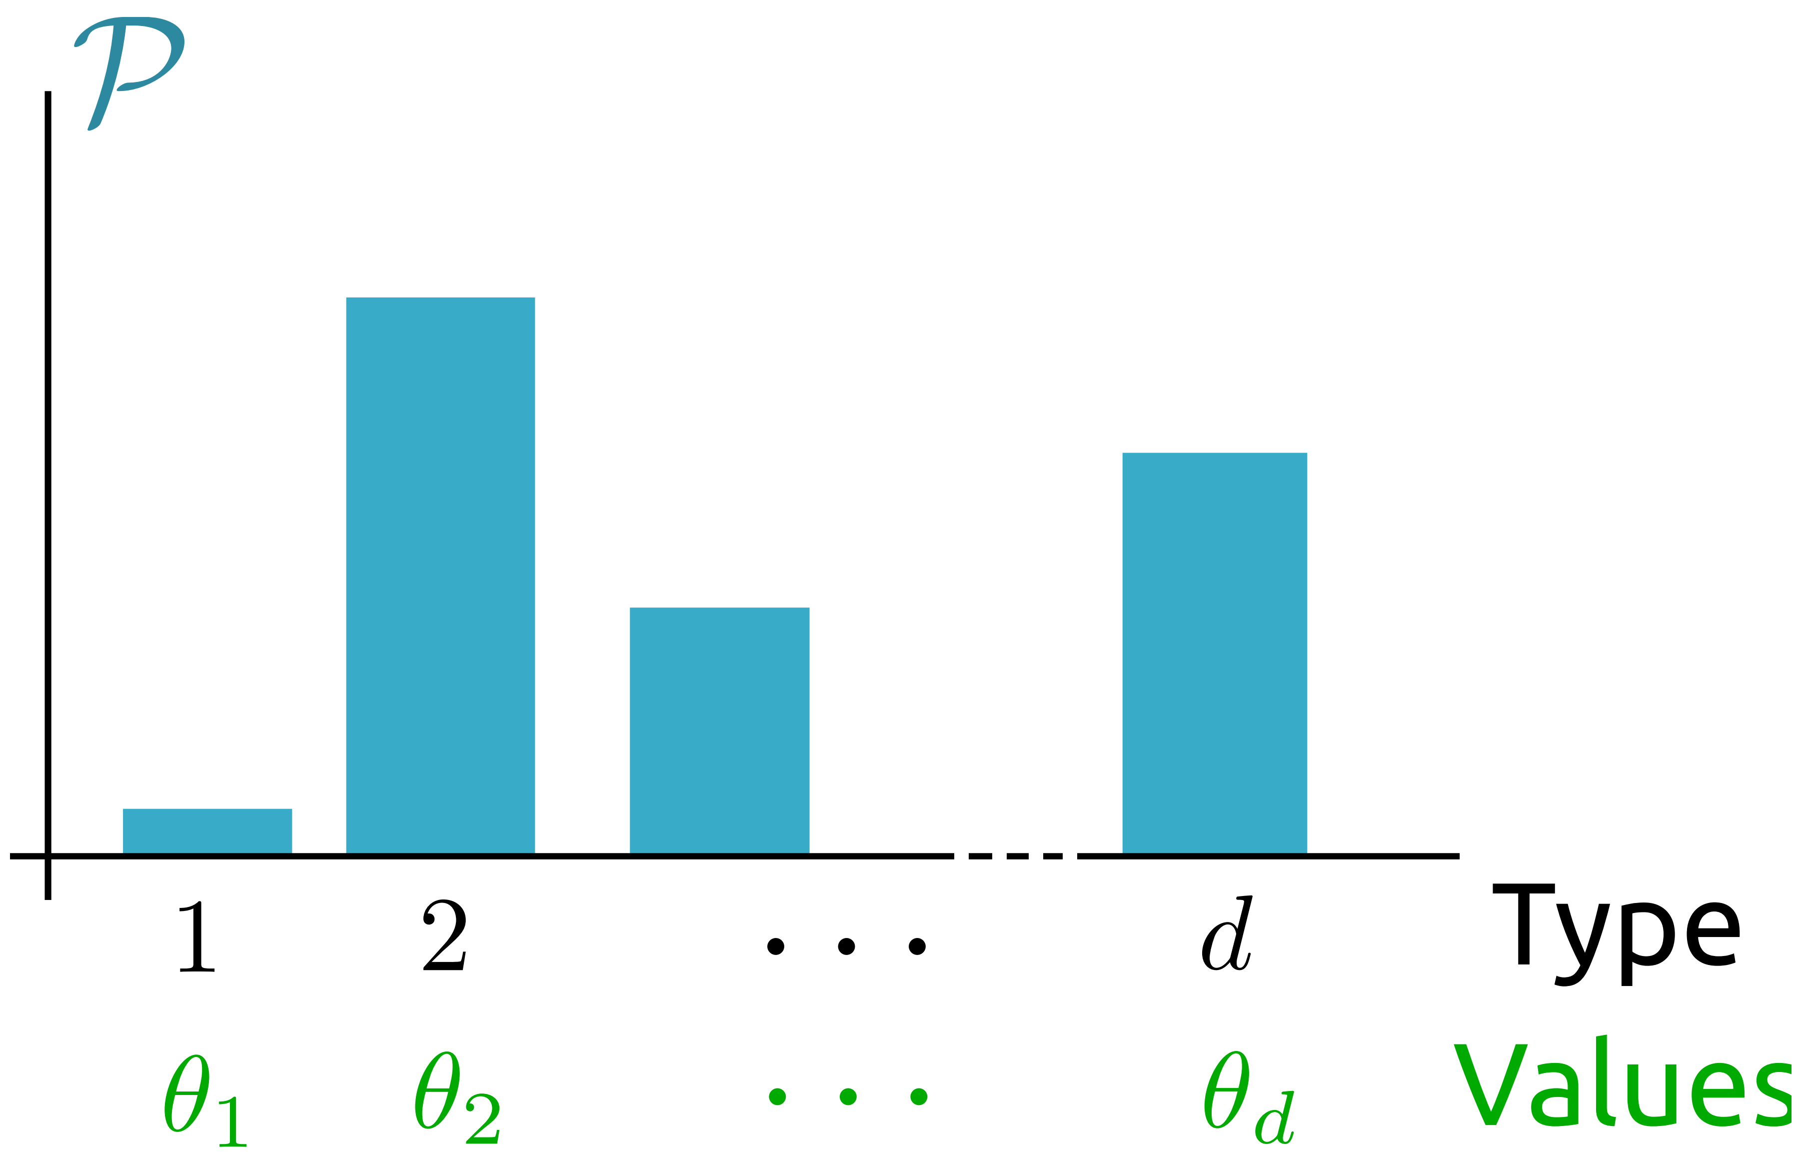

2. A buyer may know their type i, but not their value θ_i.
   - (In practice) due to uncertainty about their value, customers may not be willing to buy an item except at a low price.

- **In this work:** both customers and seller will **use reviews to learn.**

OUTLINE

1. **Problem set up**
   - Online learning framework, assumptions, challenges

2. **Algorithm**

3. **Theoretical results**
   - Upper bounds, lower bounds, proof sketches

OUTLINE

1. **Problem set up**
    - Online learning framework, assumptions, challenges

2. **Algorithm**

3. **Theoretical results**
    - Upper bounds, lower bounds, proof sketches

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 $i_t \sim \mathscr{P}$ arrives. She buys the item if she has reason to believe that $\theta_{i_t} \geq p$ (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 $\theta_{i_t} \geq p$ (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 $\theta_i \geq p$ (based on past reviews).
  - *If buyer buys,* she leaves a **review** based on her **experience** of using the item. Otherwise, no review.

EX-ANTE VS EX-POST VALUE

EX-ANTE VS EX-POST VALUE

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

EX-ANTE VS EX-POST VALUE

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

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

EX-ANTE VS EX-POST VALUE

- 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 $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.
  - Actual experience of the customer.

EX-ANTE VS EX-POST VALUE

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

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

EX-ANTE VS EX-POST VALUE

- 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 $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.
  - Actual experience of the customer.
  - Depends on exogenous factors that cannot be known at time of purchase.
    - E.g manufacturing defects, delivery quality.

EX-ANTE VS EX-POST VALUE

- 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 $\mathbb{E}_{\mathscr{D}_i}[v] = \theta_i$.
  - Actual experience of the customer.
  - Depends on exogenous factors that cannot be known at time of purchase.
    - E.g manufacturing defects, delivery quality.

- Customer reviews are based on *ex-post value* (actual experience).

WHAT IS A REVIEW?

WHAT IS A REVIEW?

- *If the buyers purchase*, they reveal their type i and ex-post value v to the seller and future buyers.

WHAT IS A REVIEW?

- *If the buyers purchase,* they reveal their type i and ex-post value v to the seller and future buyers.
    - Can extract type and value from written reviews, ratings, and buyer history *(AMMO '22 Econometrica)*

WHAT IS A REVIEW?

- *If the buyers purchase*, they reveal their type i and ex-post value v to the seller and future buyers.
  - Can extract type and value from written reviews, ratings, and buyer history *(AMMO '22 Econometrica)*
  - *'Revealing type'* is perhaps a new model for soliciting customer 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 $i \sim \mathscr{P}$ arrives. She buys the item if **she has reason to believe** that $\theta_i \geq p$ (based on past reviews).
  - *If buyer buys,* she leaves a **review** based on her **experience** of using the item. Otherwise, no review.

BUYER PURCHASE MODEL

BUYER PURCHASE MODEL

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

BUYER PURCHASE MODEL

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

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

BUYER PURCHASE MODEL

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

- **But** buyers cannot be overly conservative.
  - E.g: “I will only pay $0.01 since I do not know my value exactly”.

BUYER PURCHASE MODEL

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

- **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.

BUYER PURCHASE MODEL: η-RISKY CUSTOMERS

BUYER PURCHASE MODEL: η–RISKY CUSTOMERS

- Buyer on round t arrives with a threshold function τ_t.
  τ_t is an "estimate" of their value based on past reviews of their type.

BUYER PURCHASE MODEL: η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold function τ_t.
  τ_t is an "estimate" of their value based on past reviews of their type.
- Buyer purchases if $p_t \leq \tau_t$.

BUYER PURCHASE MODEL: η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold function τ_t.
  τ_t is an "estimate" of their value based on past reviews of their type.
- Buyer purchases if $p_t \leq \tau_t$.
- But this threshold has to be at least a η-lower confidence bound on the value.
  Let the buyer's type be i, and let $\Phi_{i,t}$ be reviews from past customers of type i. Then,

BUYER PURCHASE MODEL: η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold function τ_t.
  τ_t is an "estimate" of their value based on past reviews of their type.
- Buyer purchases if $p_t \leq \tau_t$.
- But this threshold has to be at least a η-lower confidence bound on the value.
  Let the buyer's type be i, and let $\Phi_{i,t}$ be reviews from past customers of type i. Then,

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

BUYER PURCHASE MODEL: η-RISKY CUSTOMERS

- Buyer on round t arrives with a threshold function τ_t.
  τ_t is an "estimate" of their value based on past reviews of their type.
- Buyer purchases if $p_t \leq \tau_t$.
- But this threshold has to be at least a η-lower confidence bound on the value.
  Let the buyer's type be i, and let $\Phi_{i,t}$ be 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:** The customer is willing to take at least a small risk. They may over-estimate their value (i.e $\tau_t > \theta_i$) with some small probability η.

ONLINE LEARNING FRAMEWORK

ONLINE LEARNING FRAMEWORK

- On each round t:

ONLINE LEARNING FRAMEWORK

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

ONLINE LEARNING FRAMEWORK

- On each round t:
  - Seller chooses a price p_t (based on past reviews)
  - A new buyer with type $i_t \sim \mathscr{P}$ and threshold τ_t arrives.

ONLINE LEARNING FRAMEWORK

- On each round t:
  - Seller chooses a price p_t (based on past reviews)
  - A new buyer with type $i_t \sim \mathscr{P}$ and threshold τ_t arrives.
  - Buyer buys if $p_t \leq \tau_t$.

ONLINE LEARNING FRAMEWORK

- On each round t:
  - Seller chooses a price p_t (based on past reviews)
  - A new buyer with type $i_t \sim \mathscr{P}$ and threshold τ_t arrives.
  - Buyer buys if $p_t \leq \tau_t$.
    - *If buyer buys,*

ONLINE LEARNING FRAMEWORK

- On each round t:
  - Seller chooses a price p_t (based on past reviews)
  - A new buyer with type $i_t \sim \mathscr{P}$ and threshold τ_t arrives.
  - Buyer buys if $p_t \leq \tau_t$.
    - *If buyer buys,*
      - Seller has revenue p_t.

ONLINE LEARNING FRAMEWORK

- On each round t:
  - Seller chooses a price p_t (based on past reviews)
  - A new buyer with type $i_t \sim \mathscr{P}$ and threshold τ_t arrives.
  - Buyer buys if $p_t \leq \tau_t$.
    - *If buyer buys,*
      - Seller has revenue p_t.
      - Buyer experiences ex-post value v_t, and reveals (i_t, v_t).

ONLINE LEARNING FRAMEWORK

- On each round t:
  - Seller chooses a price p_t (based on past reviews)
  - A new buyer with type $i_t \sim \mathscr{P}$ and threshold τ_t arrives.
  - Buyer buys if $p_t \leq \tau_t$.
    - *If buyer buys,*
      - Seller has revenue p_t.
      - Buyer experiences ex-post value v_t, and reveals (i_t, v_t).
    - *If buyer does not buy*, no revenue **and no review!**

REGRET

REGRET

- Compete against the best price $p^\star$ when sellers know $\mathscr{P}$ and customers know their type θ_{i_t}.

REGRET

- Compete against the best price $p^\star$ when sellers know $\mathscr{P}$ and customers know their type θ_{i_t}.

- Regret R_T after T rounds:

REGRET

- Compete against the best price $p^\star$ when sellers know $\mathscr{P}$ and customers know their type θ_{i_t}.

- Regret R_T after T rounds:

$$R_T = \sum_{t=1}^{T} p^\star \cdot b_t^\star - \sum_{t=1}^{T} p_t \cdot b_t$$

REGRET

- Compete against the best price $p^\star$ when sellers know $\mathscr{P}$ and customers know their type θ_{i_t}.

- Regret R_T after T rounds:

$$R_T = \sum_{t=1}^{T} p^\star \cdot b_t^\star - \sum_{t=1}^{T} p_t \cdot b_t$$

  - Here, $b_t = 1$ if there was a purchase on round t and $b_t^\star = 1$ if the customer would have purchased at $p^\star$ had they known their value.

REGRET

- Compete against the best price $p^\star$ when sellers know $\mathscr{P}$ and customers know their type θ_{i_t}.

- Regret R_T after T rounds:

$$R_T = \sum_{t=1}^{T} p^\star \cdot b_t^\star - \sum_{t=1}^{T} p_t \cdot b_t$$

  - Here, $b_t = 1$ if there was a purchase on round t and $b_t^\star = 1$ if the customer would have purchased at $p^\star$ had they known their value.
  - We want small R_T. Specifically $\mathbb{E}[R_T] \in o(T)$.

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 $\Longrightarrow$ no purchase $\Longrightarrow$ no review.

  1. *Seller learning:* Seller cannot gauge demand for the product.

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.

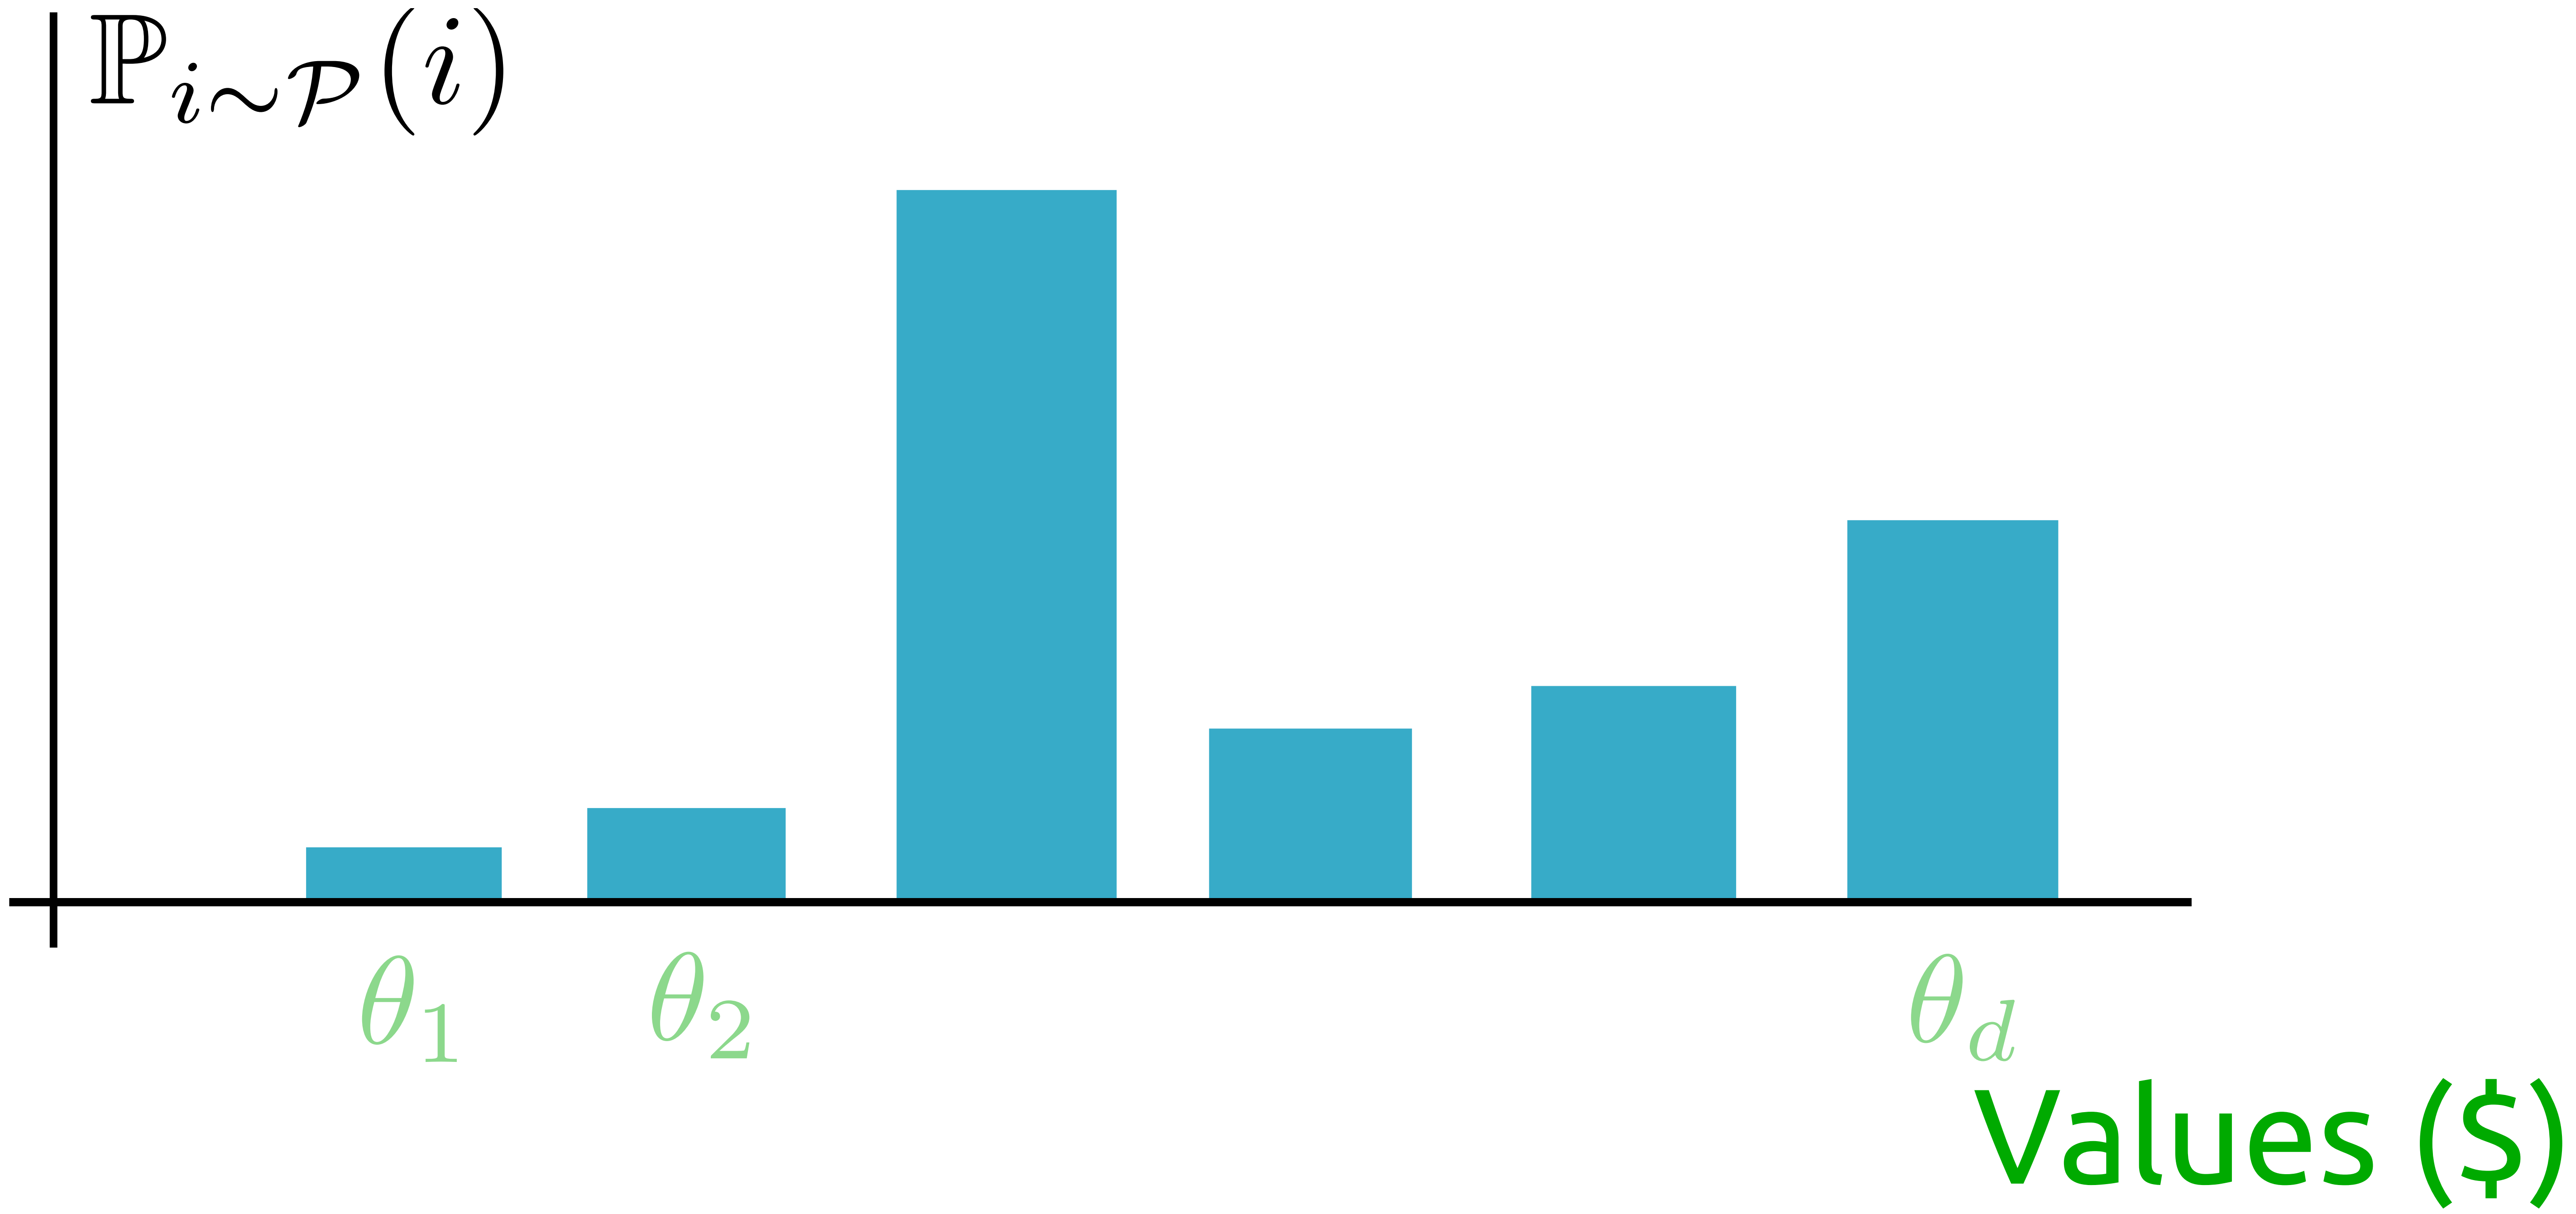

CHALLENGE 1: PRICING VS SELLER LEARNING

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

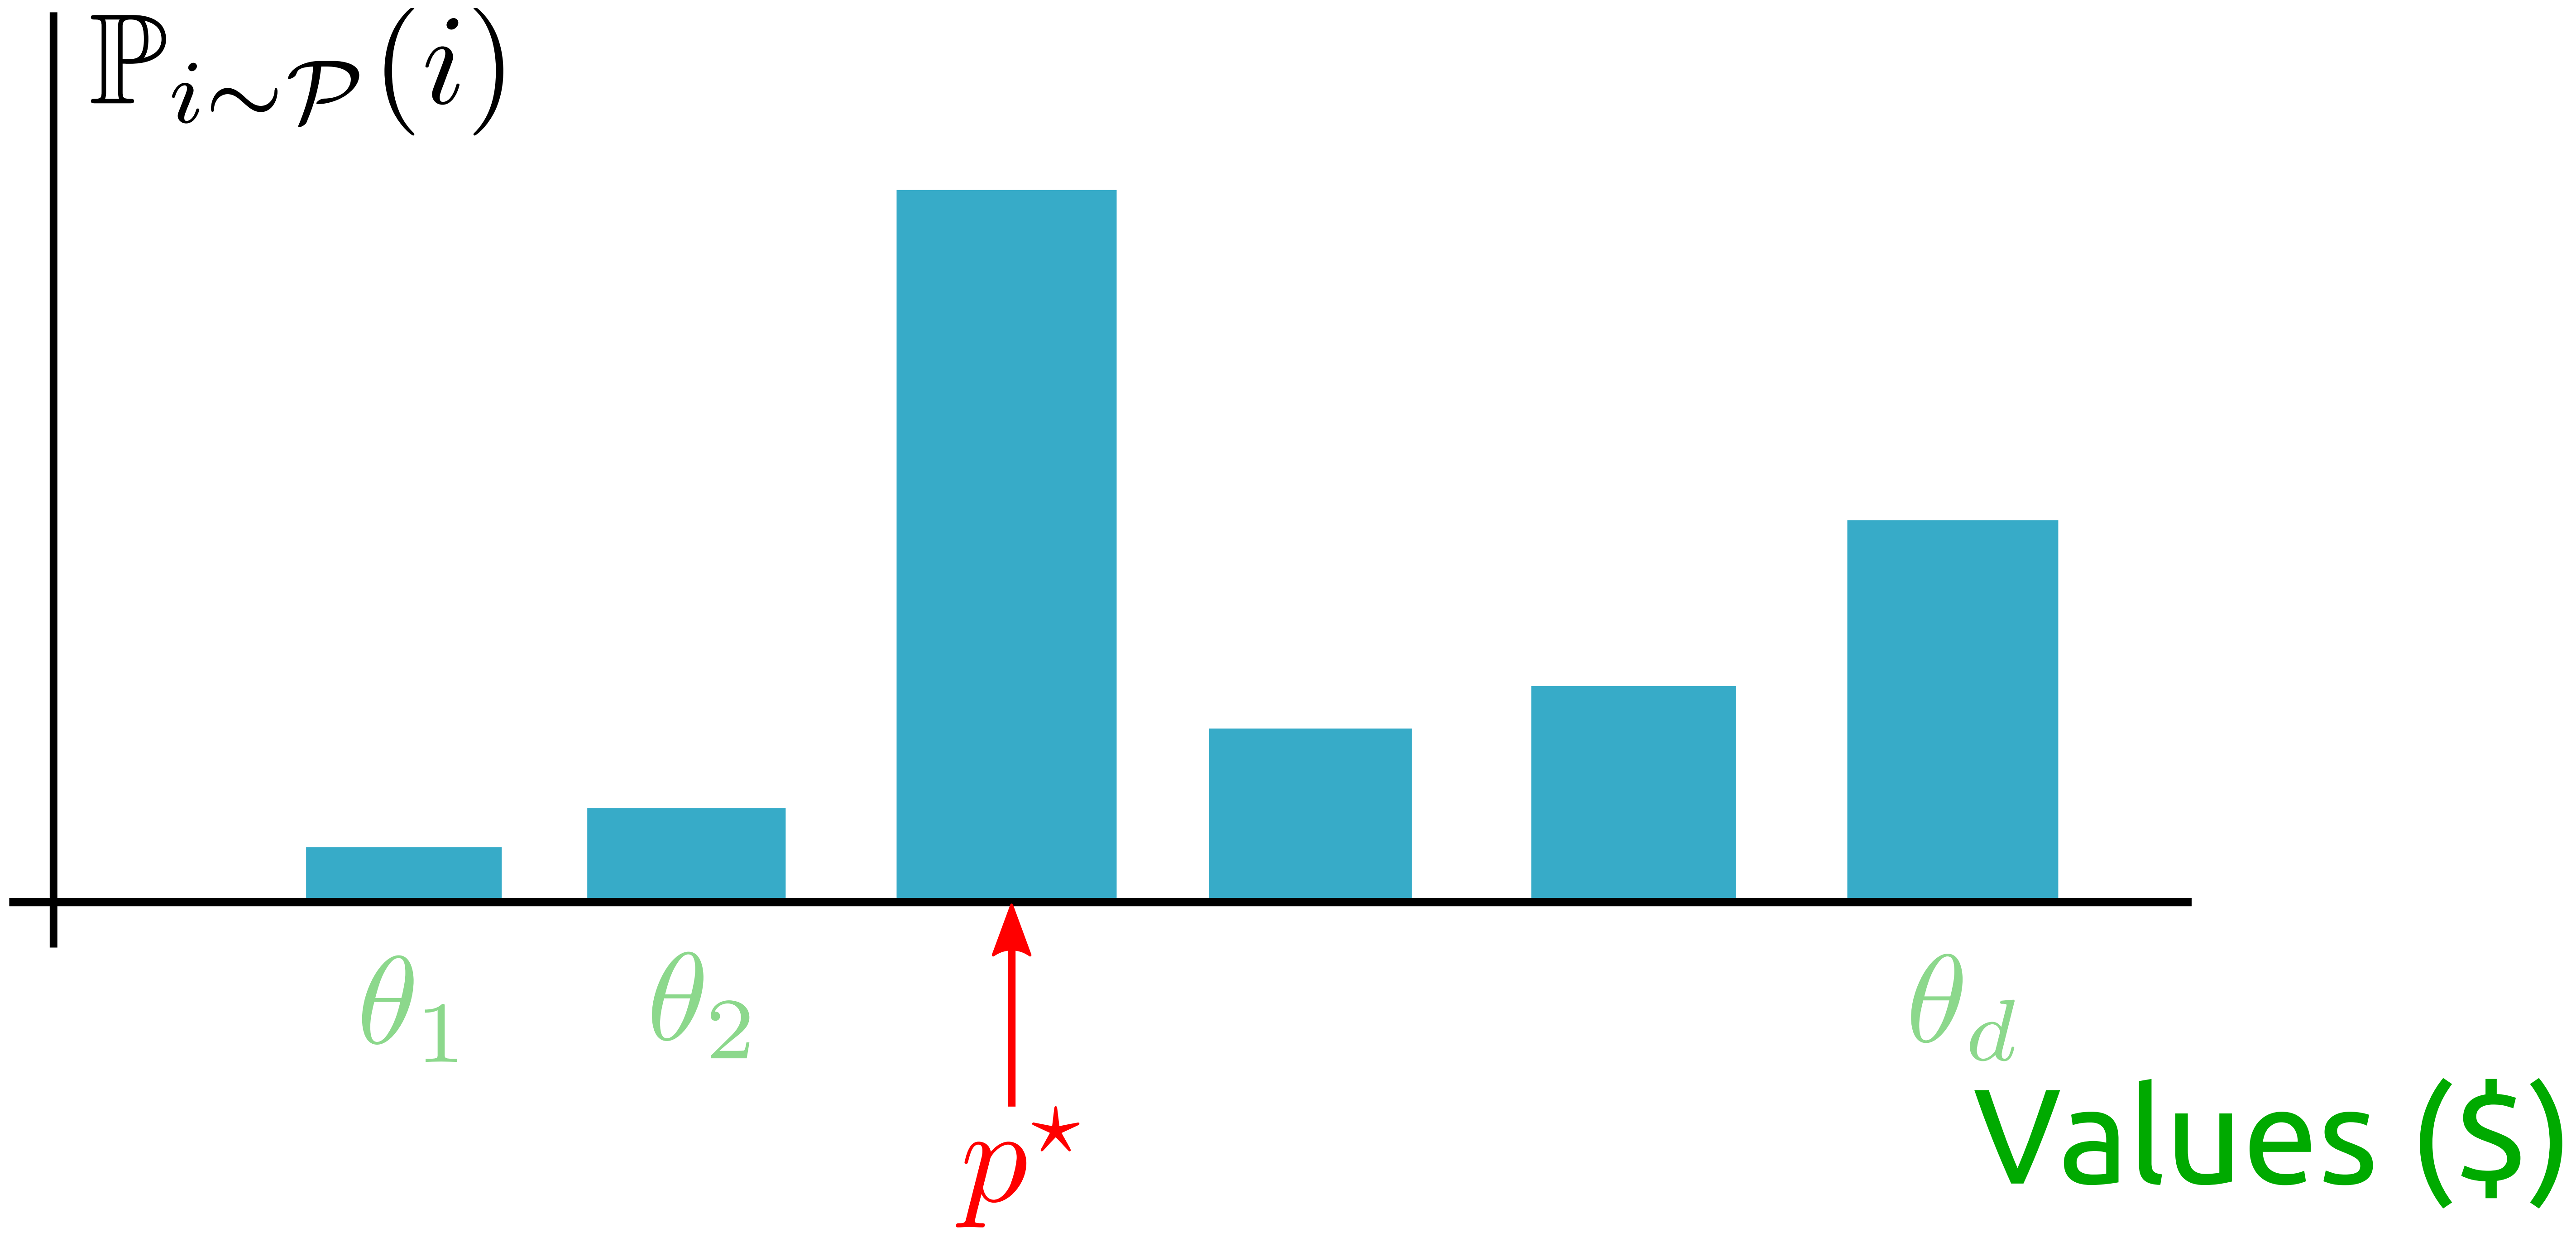

CHALLENGE 1: PRICING VS SELLER LEARNING

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

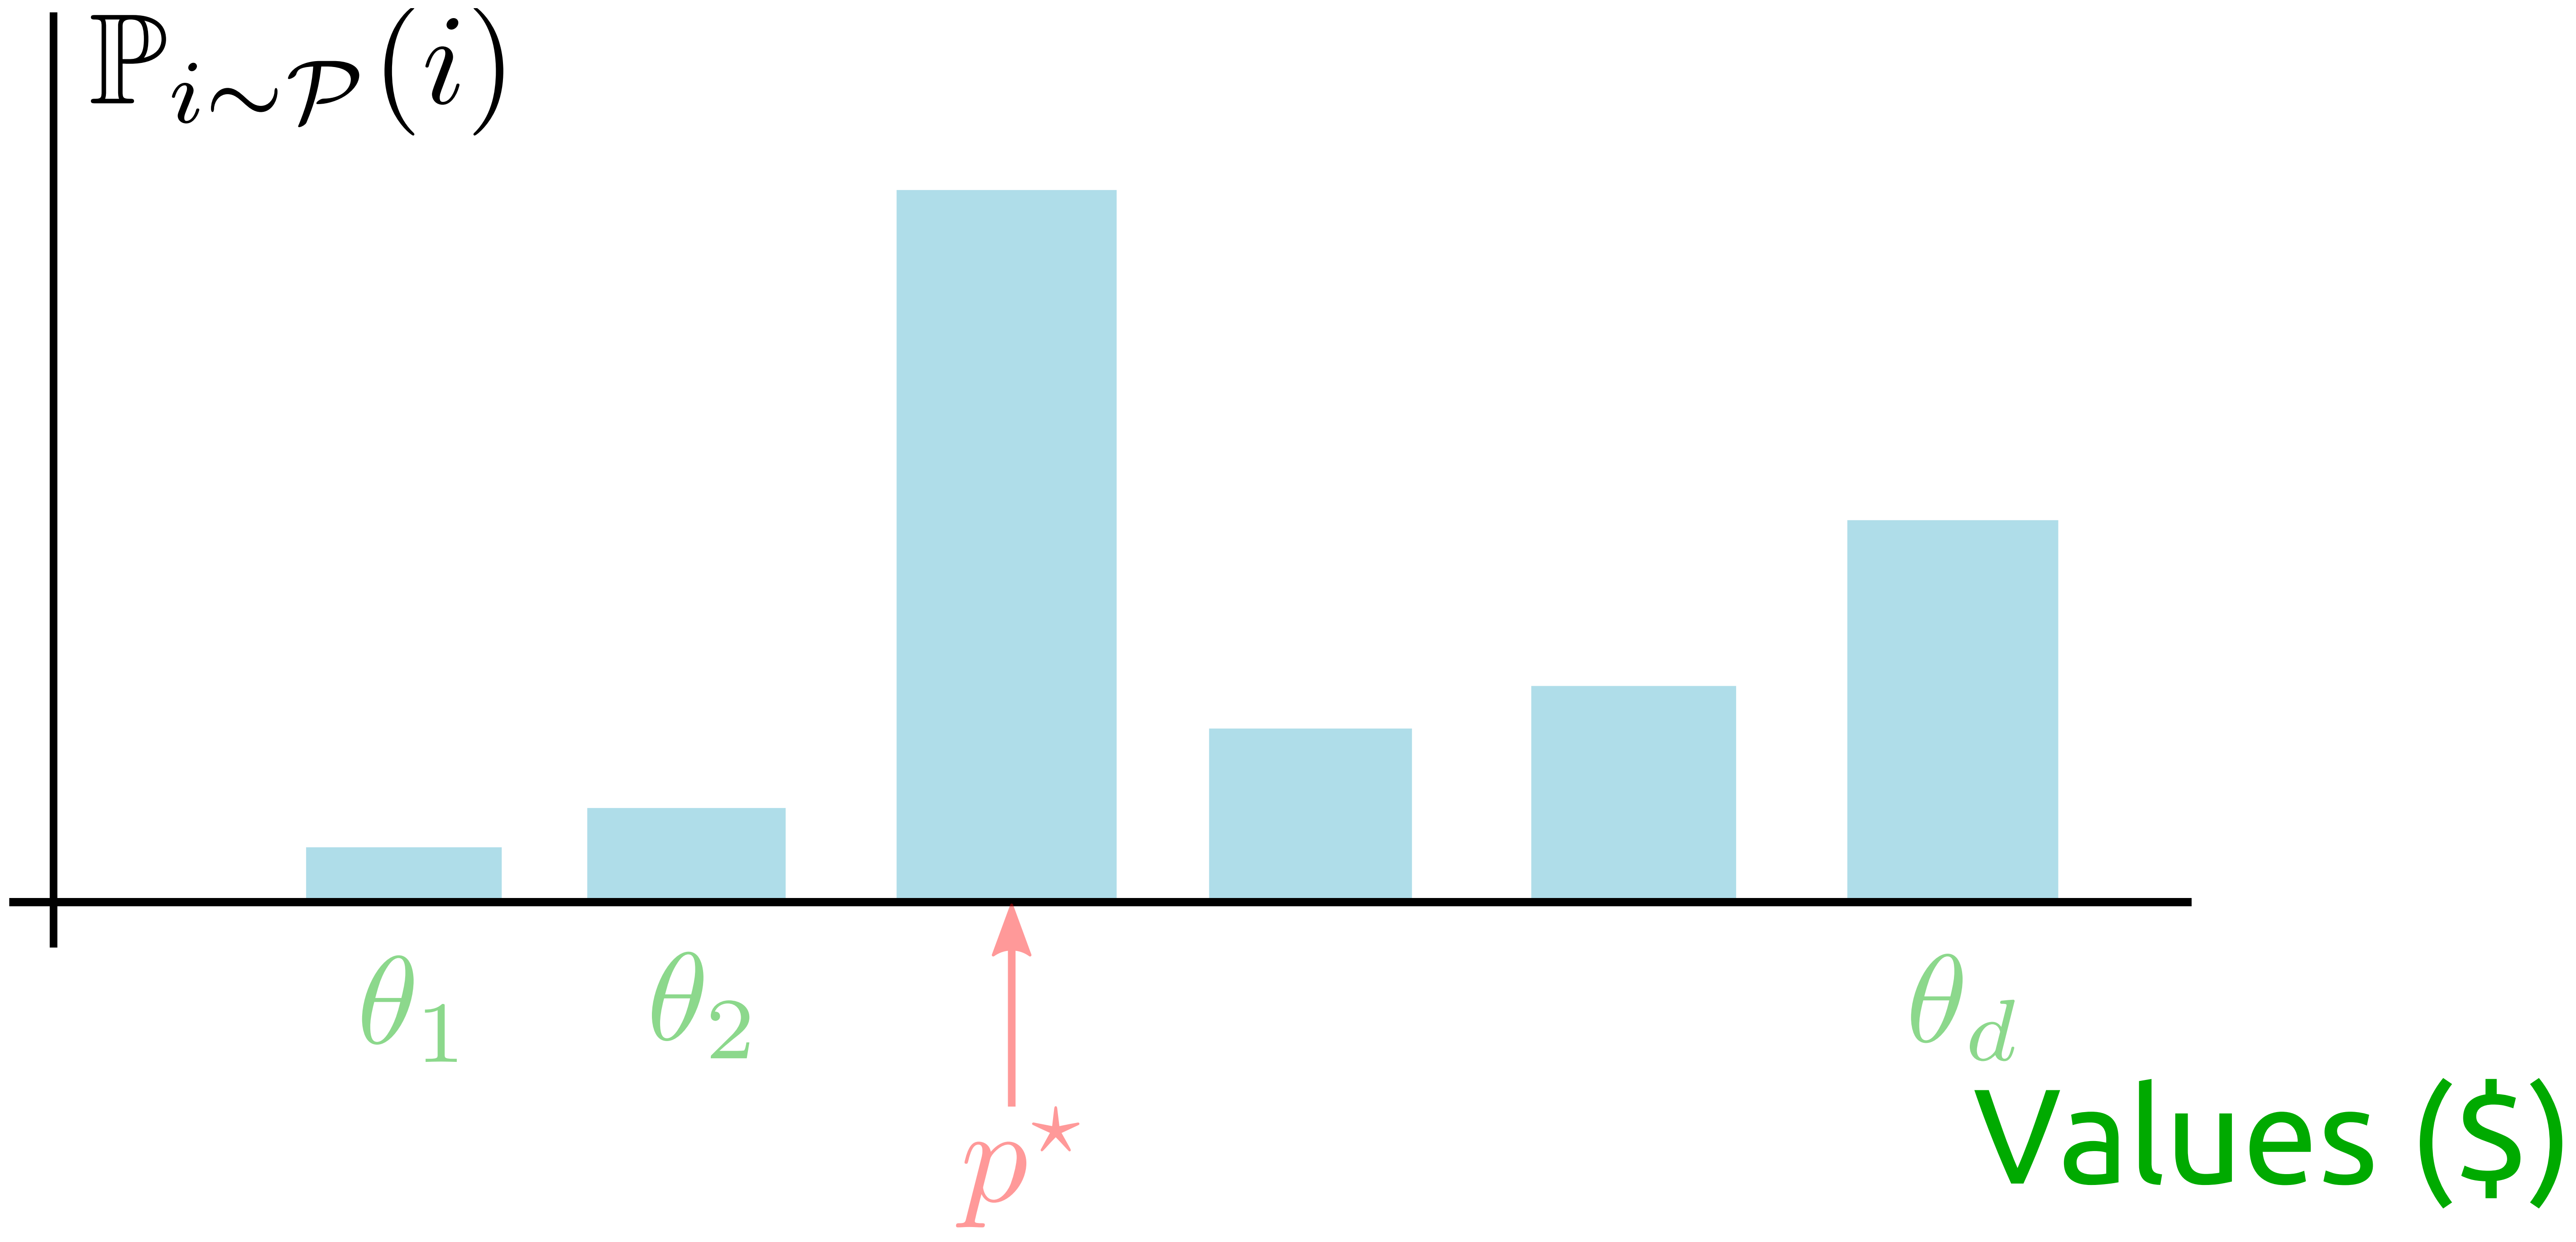

CHALLENGE 1: PRICING VS SELLER LEARNING

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

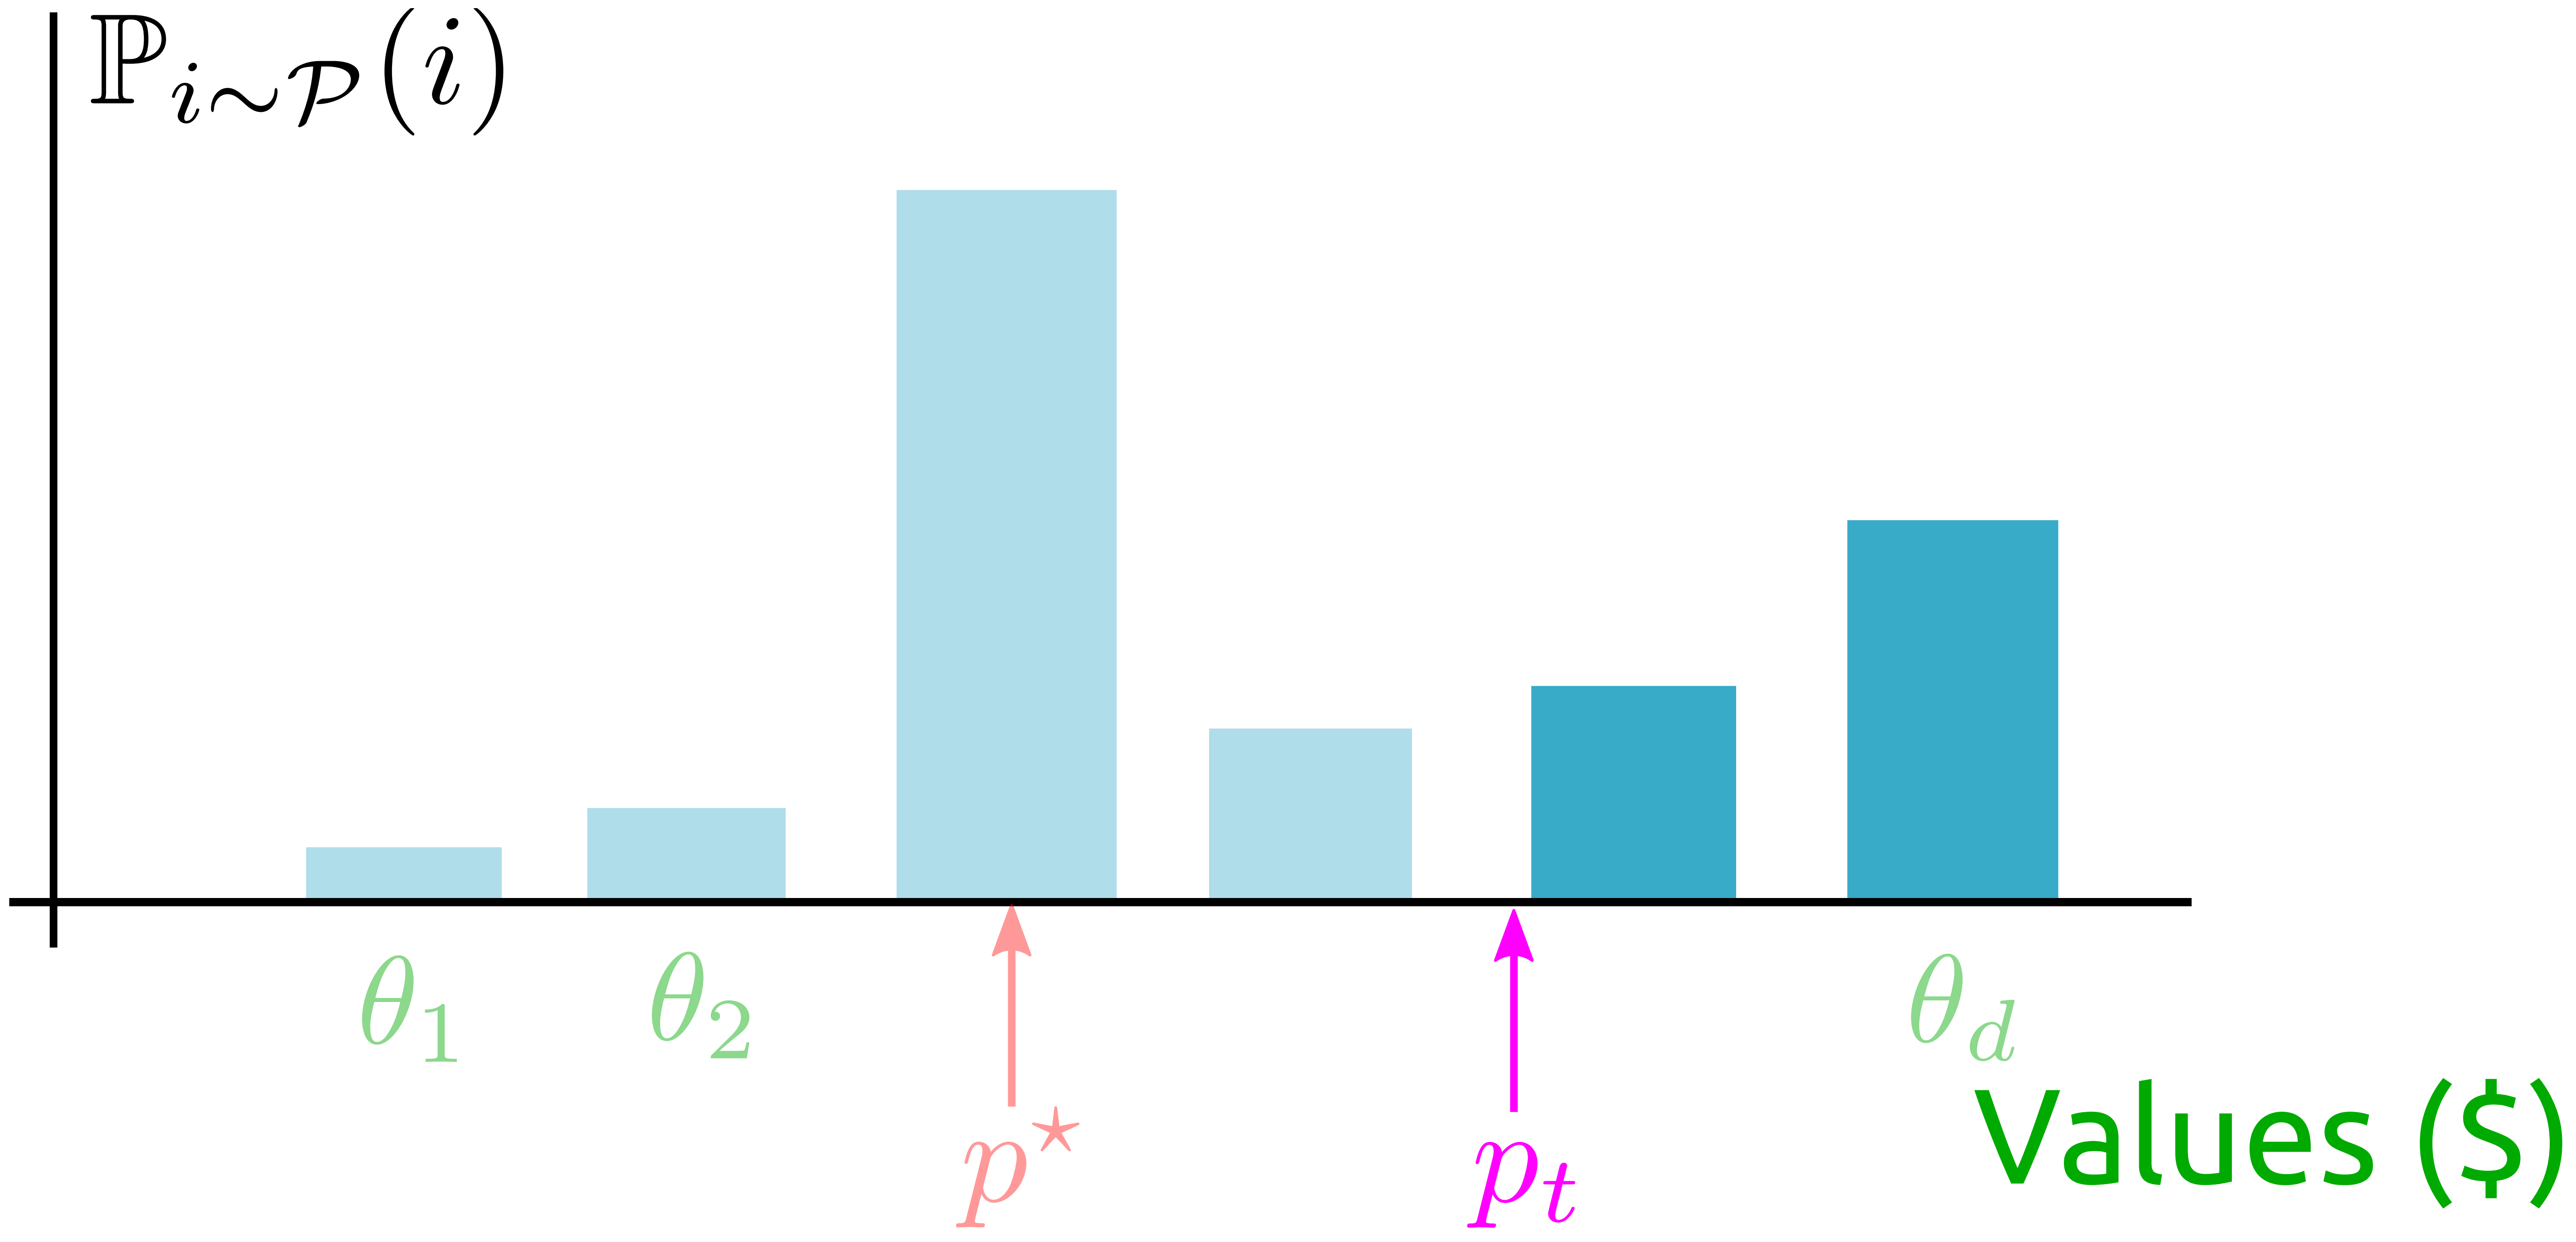

CHALLENGE 1: PRICING VS SELLER LEARNING

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

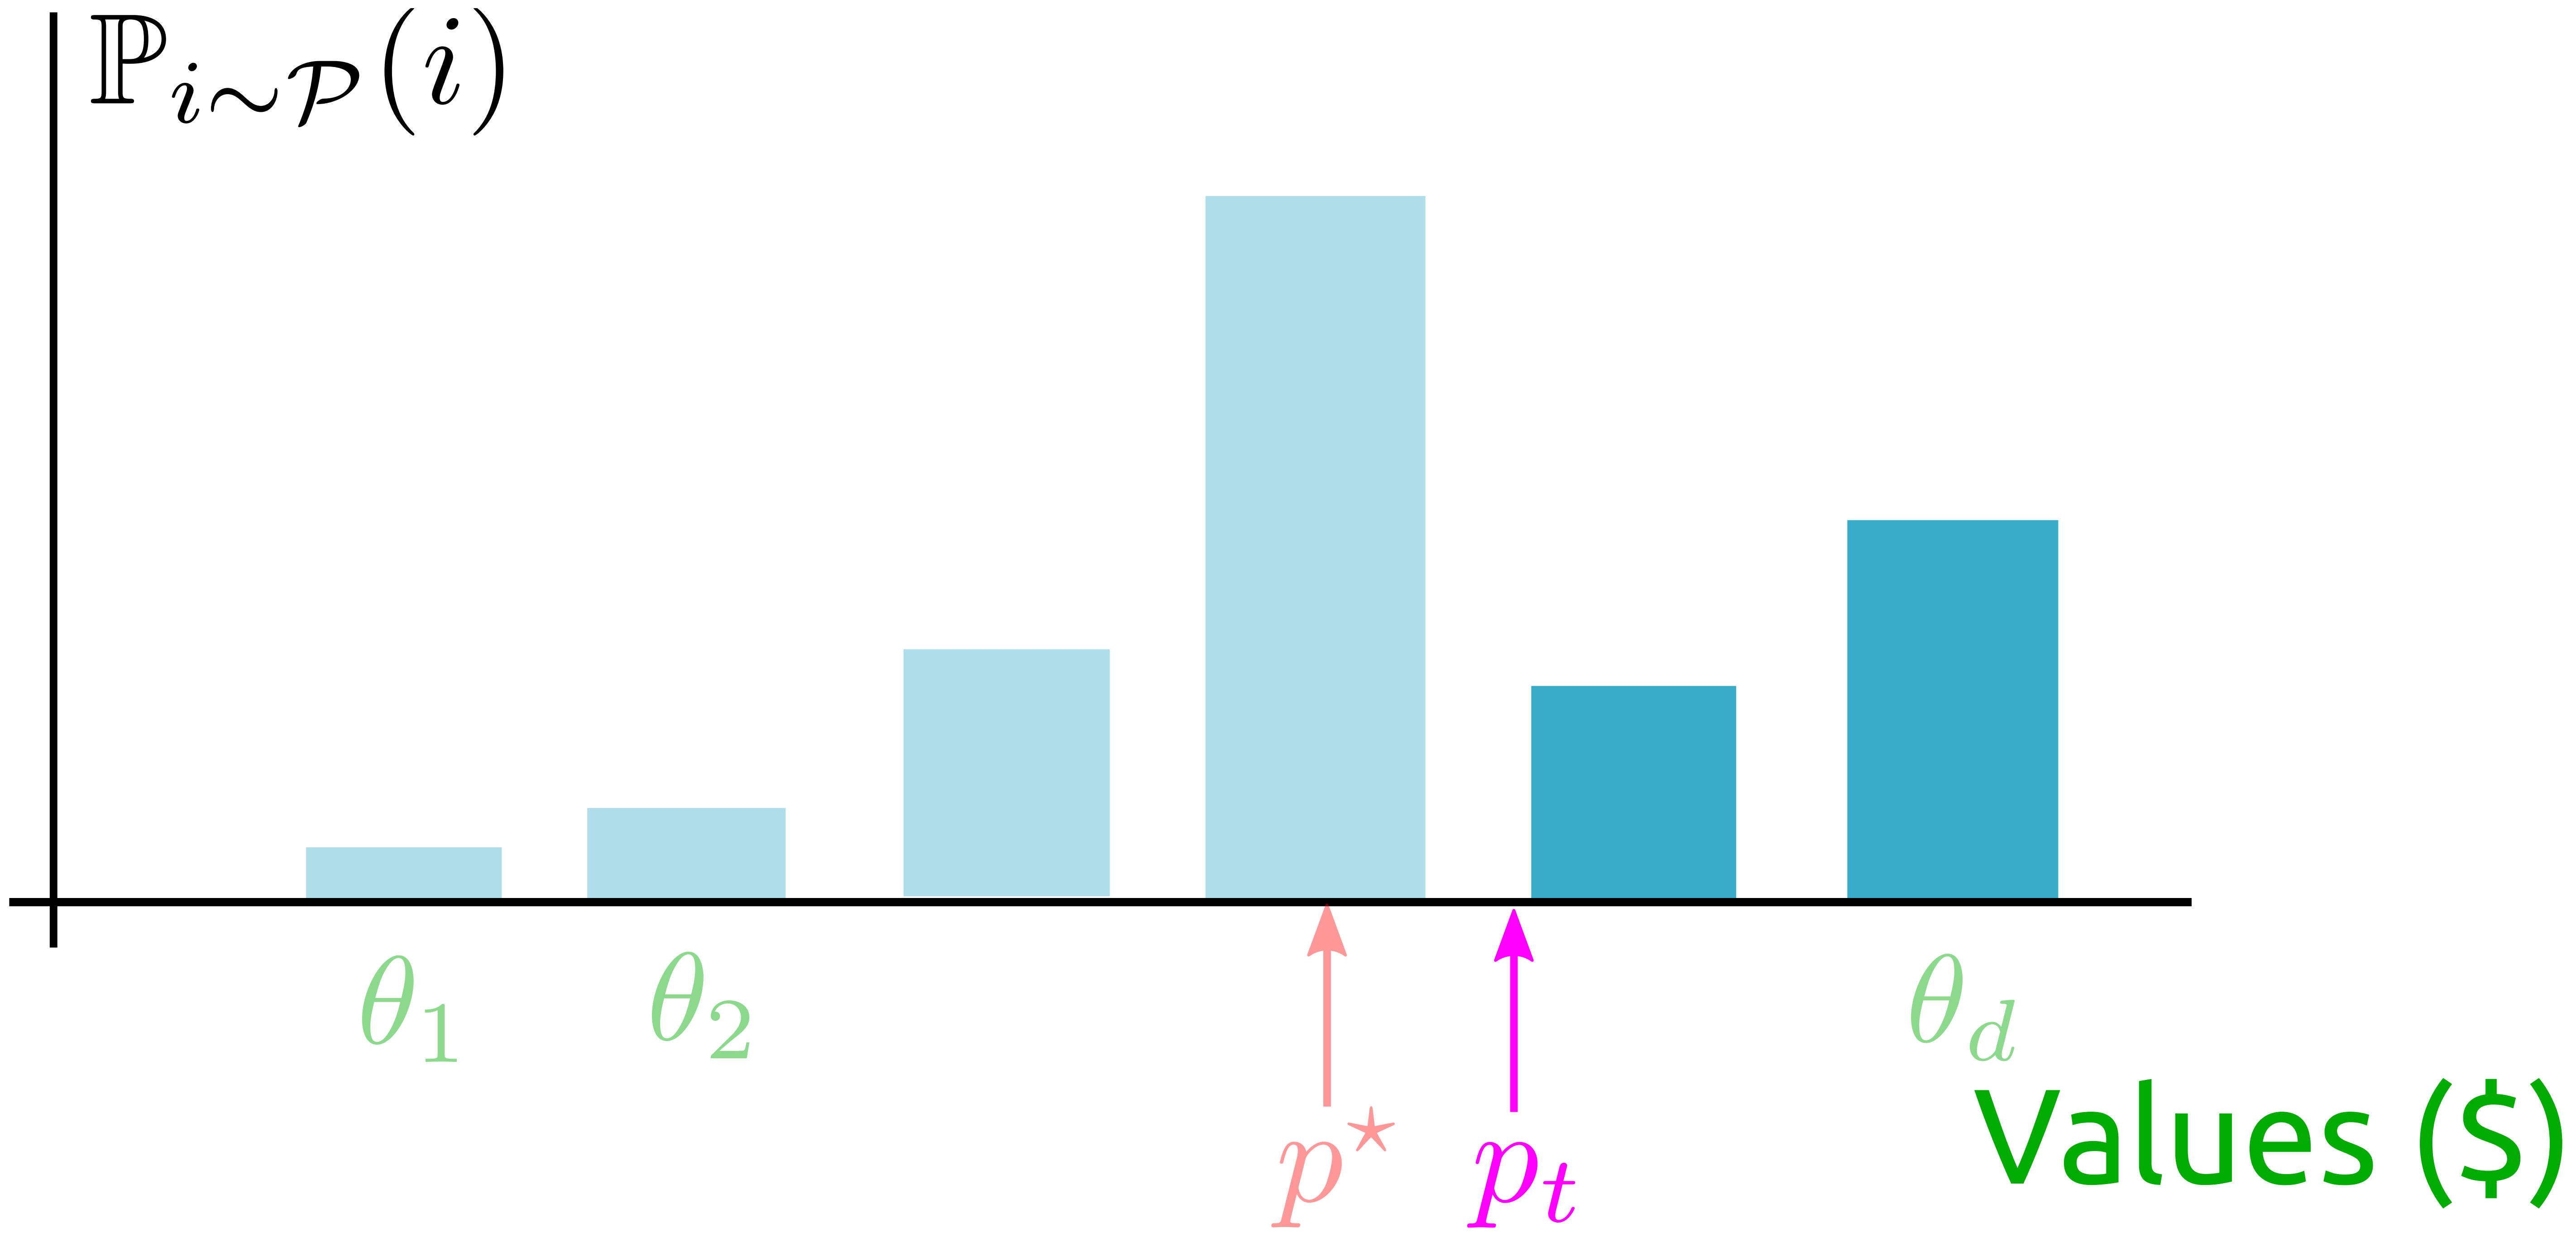

CHALLENGE 1: PRICING VS SELLER LEARNING

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

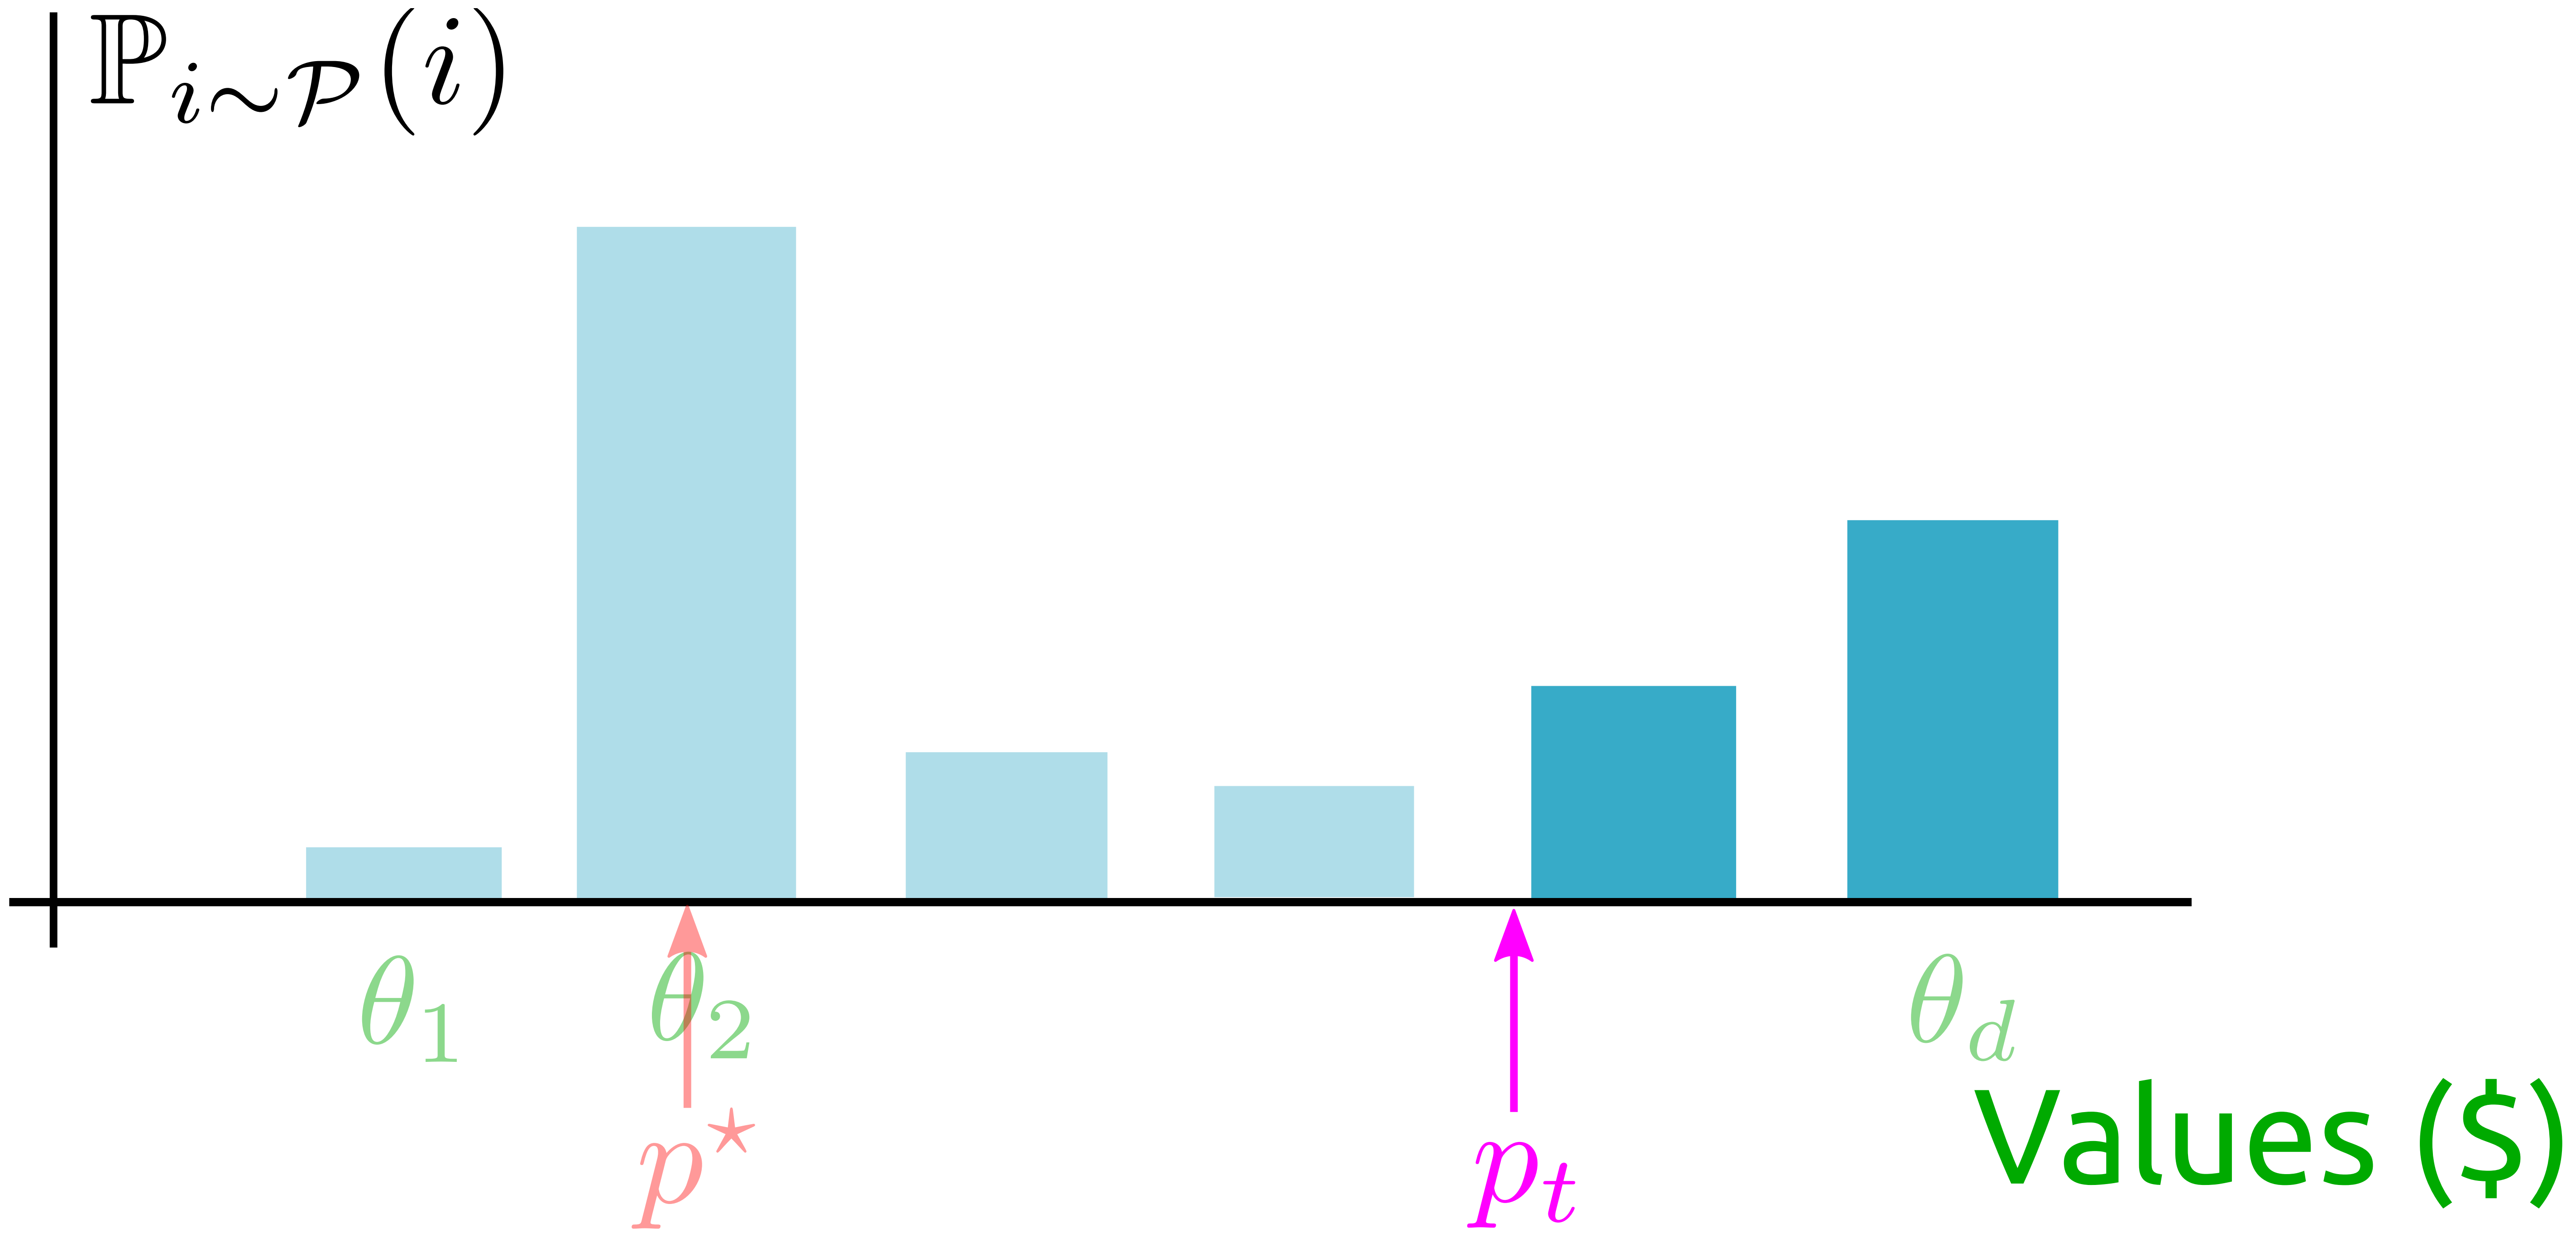

CHALLENGE 1: PRICING VS SELLER LEARNING

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

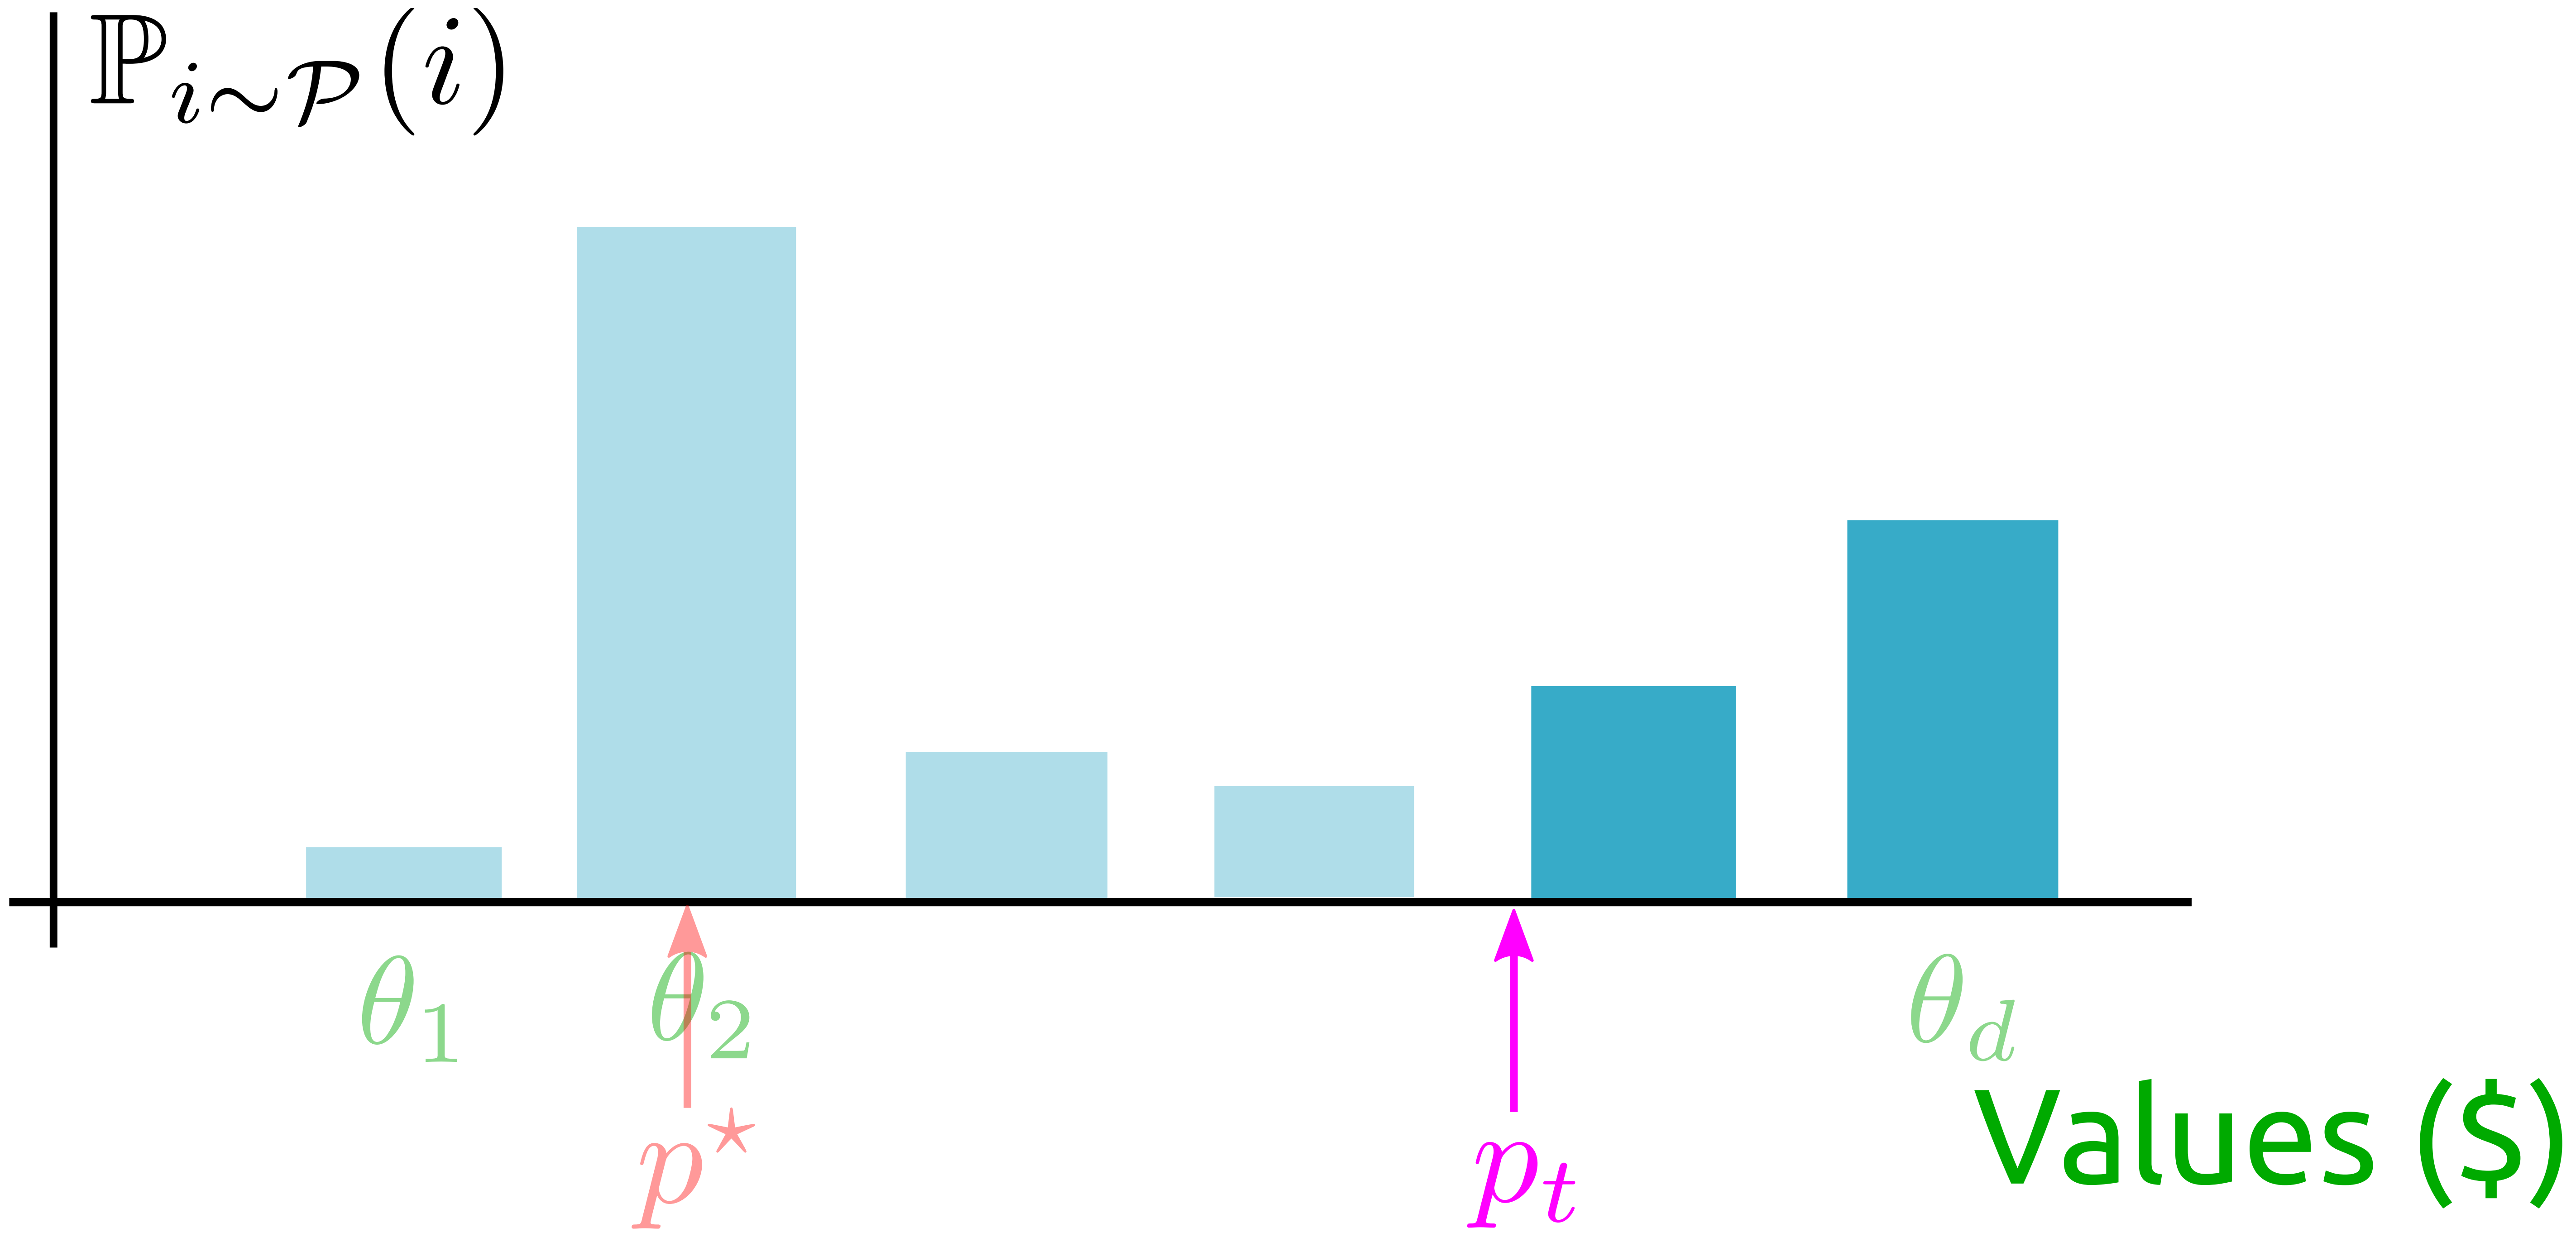

- Set prices too high $\implies$ no feedback about low value types.

CHALLENGE 1: PRICING VS SELLER LEARNING

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

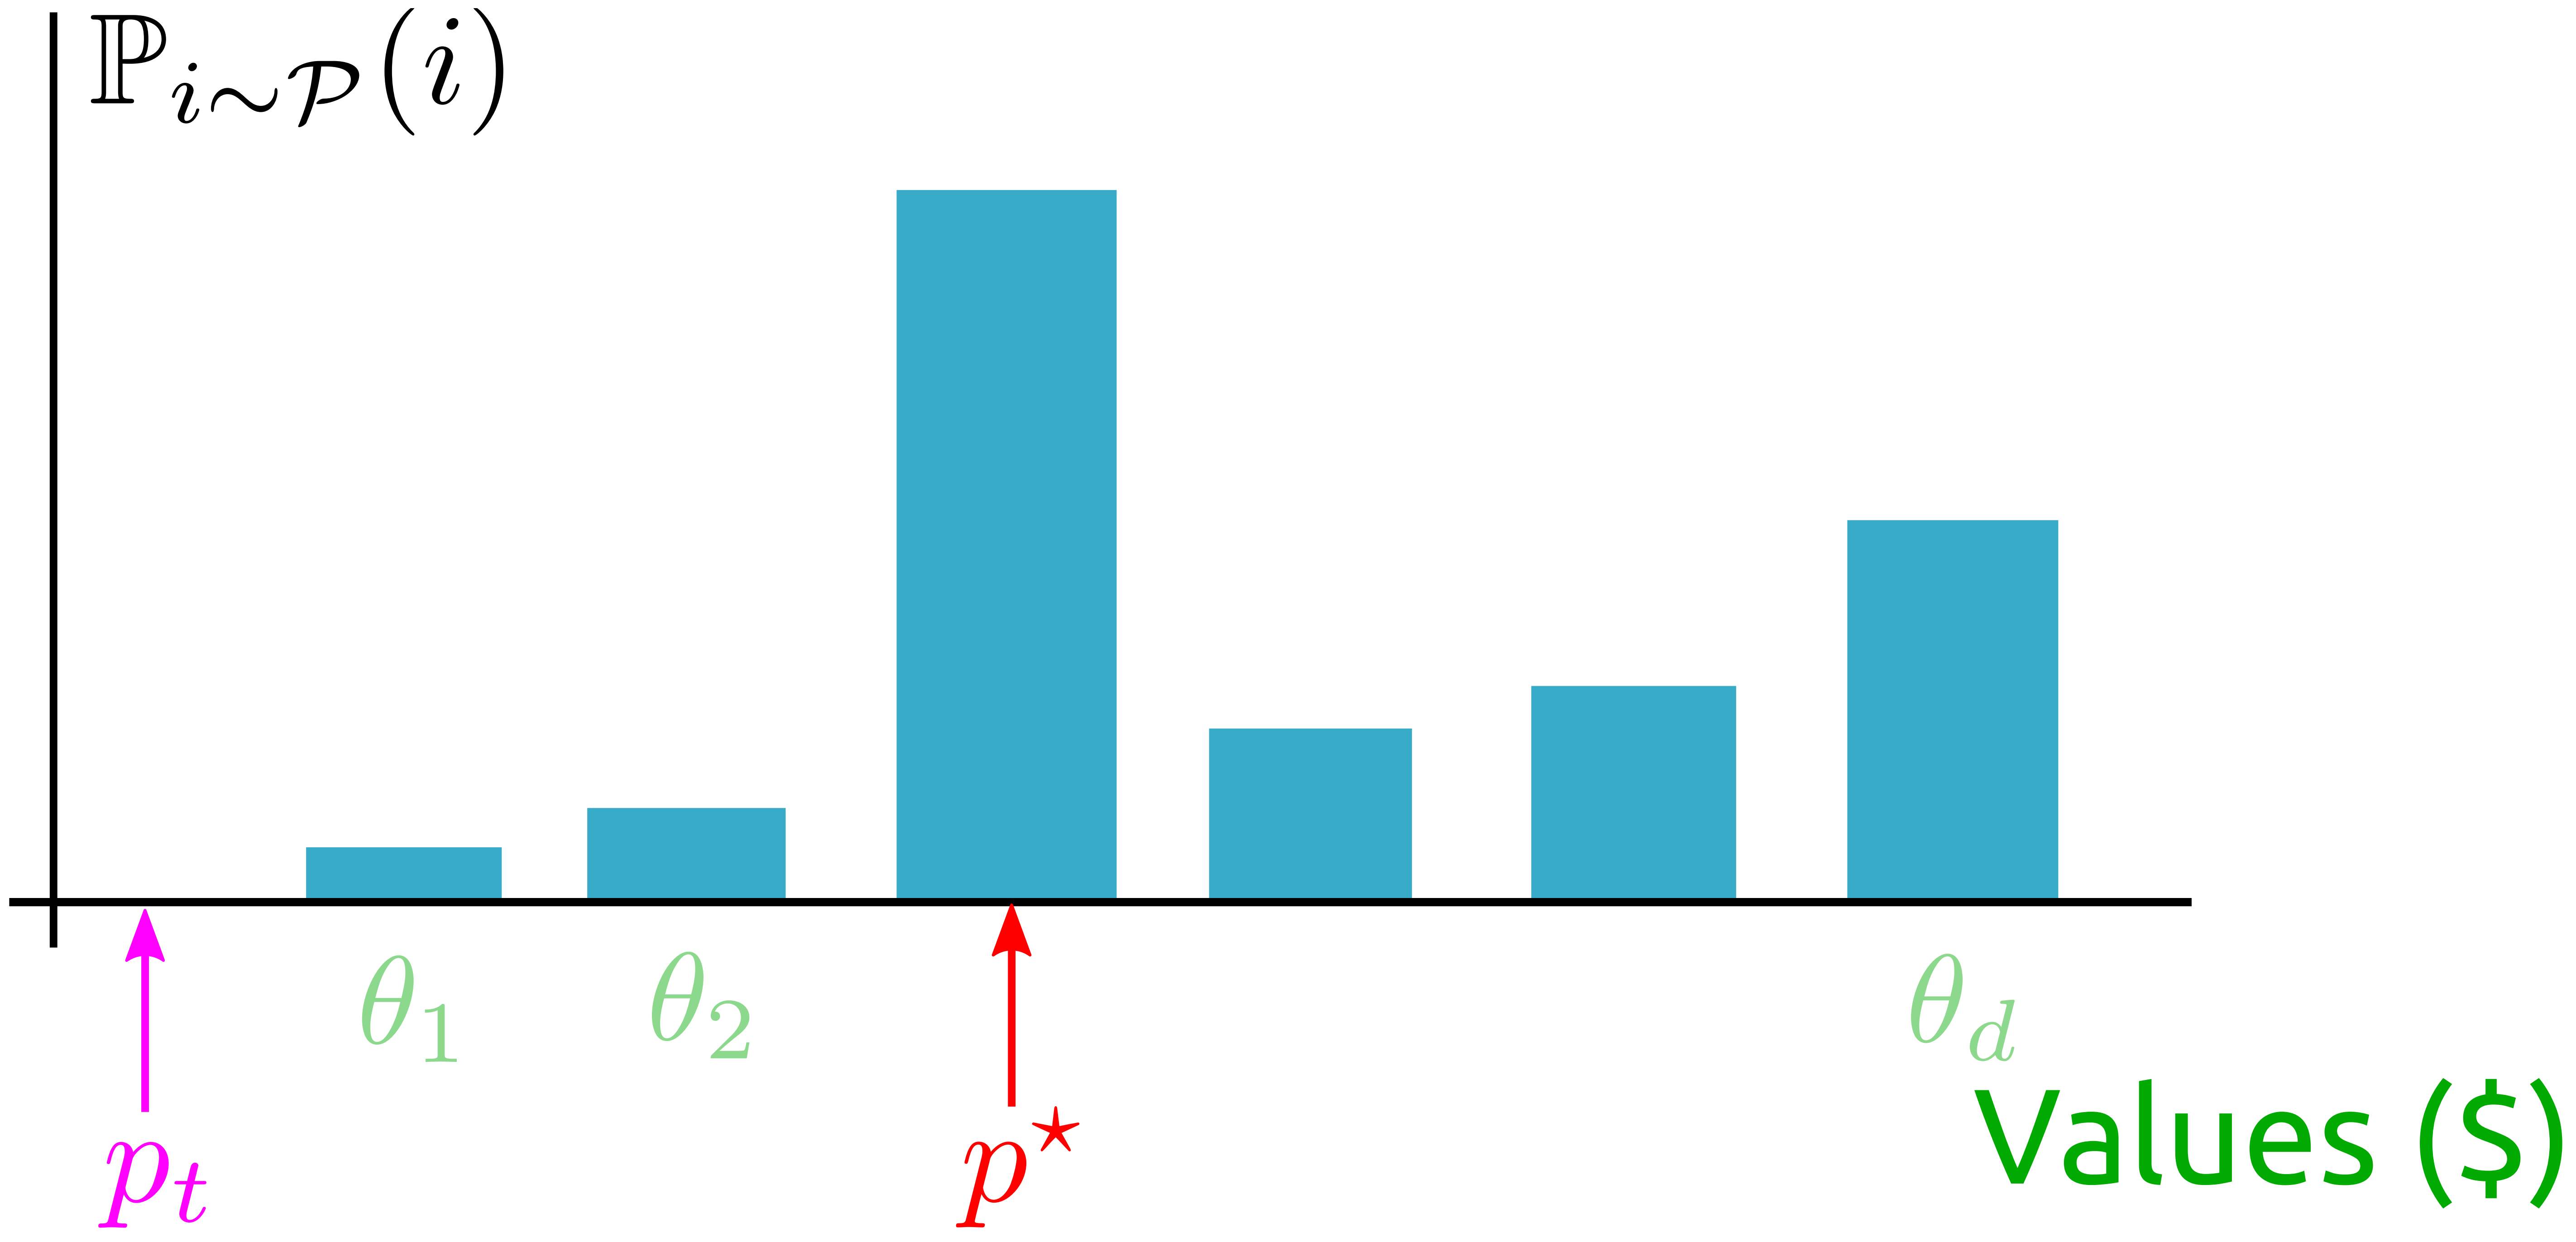

- Set prices too high $\implies$ no feedback about low value types.

CHALLENGE 1: PRICING VS SELLER LEARNING

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

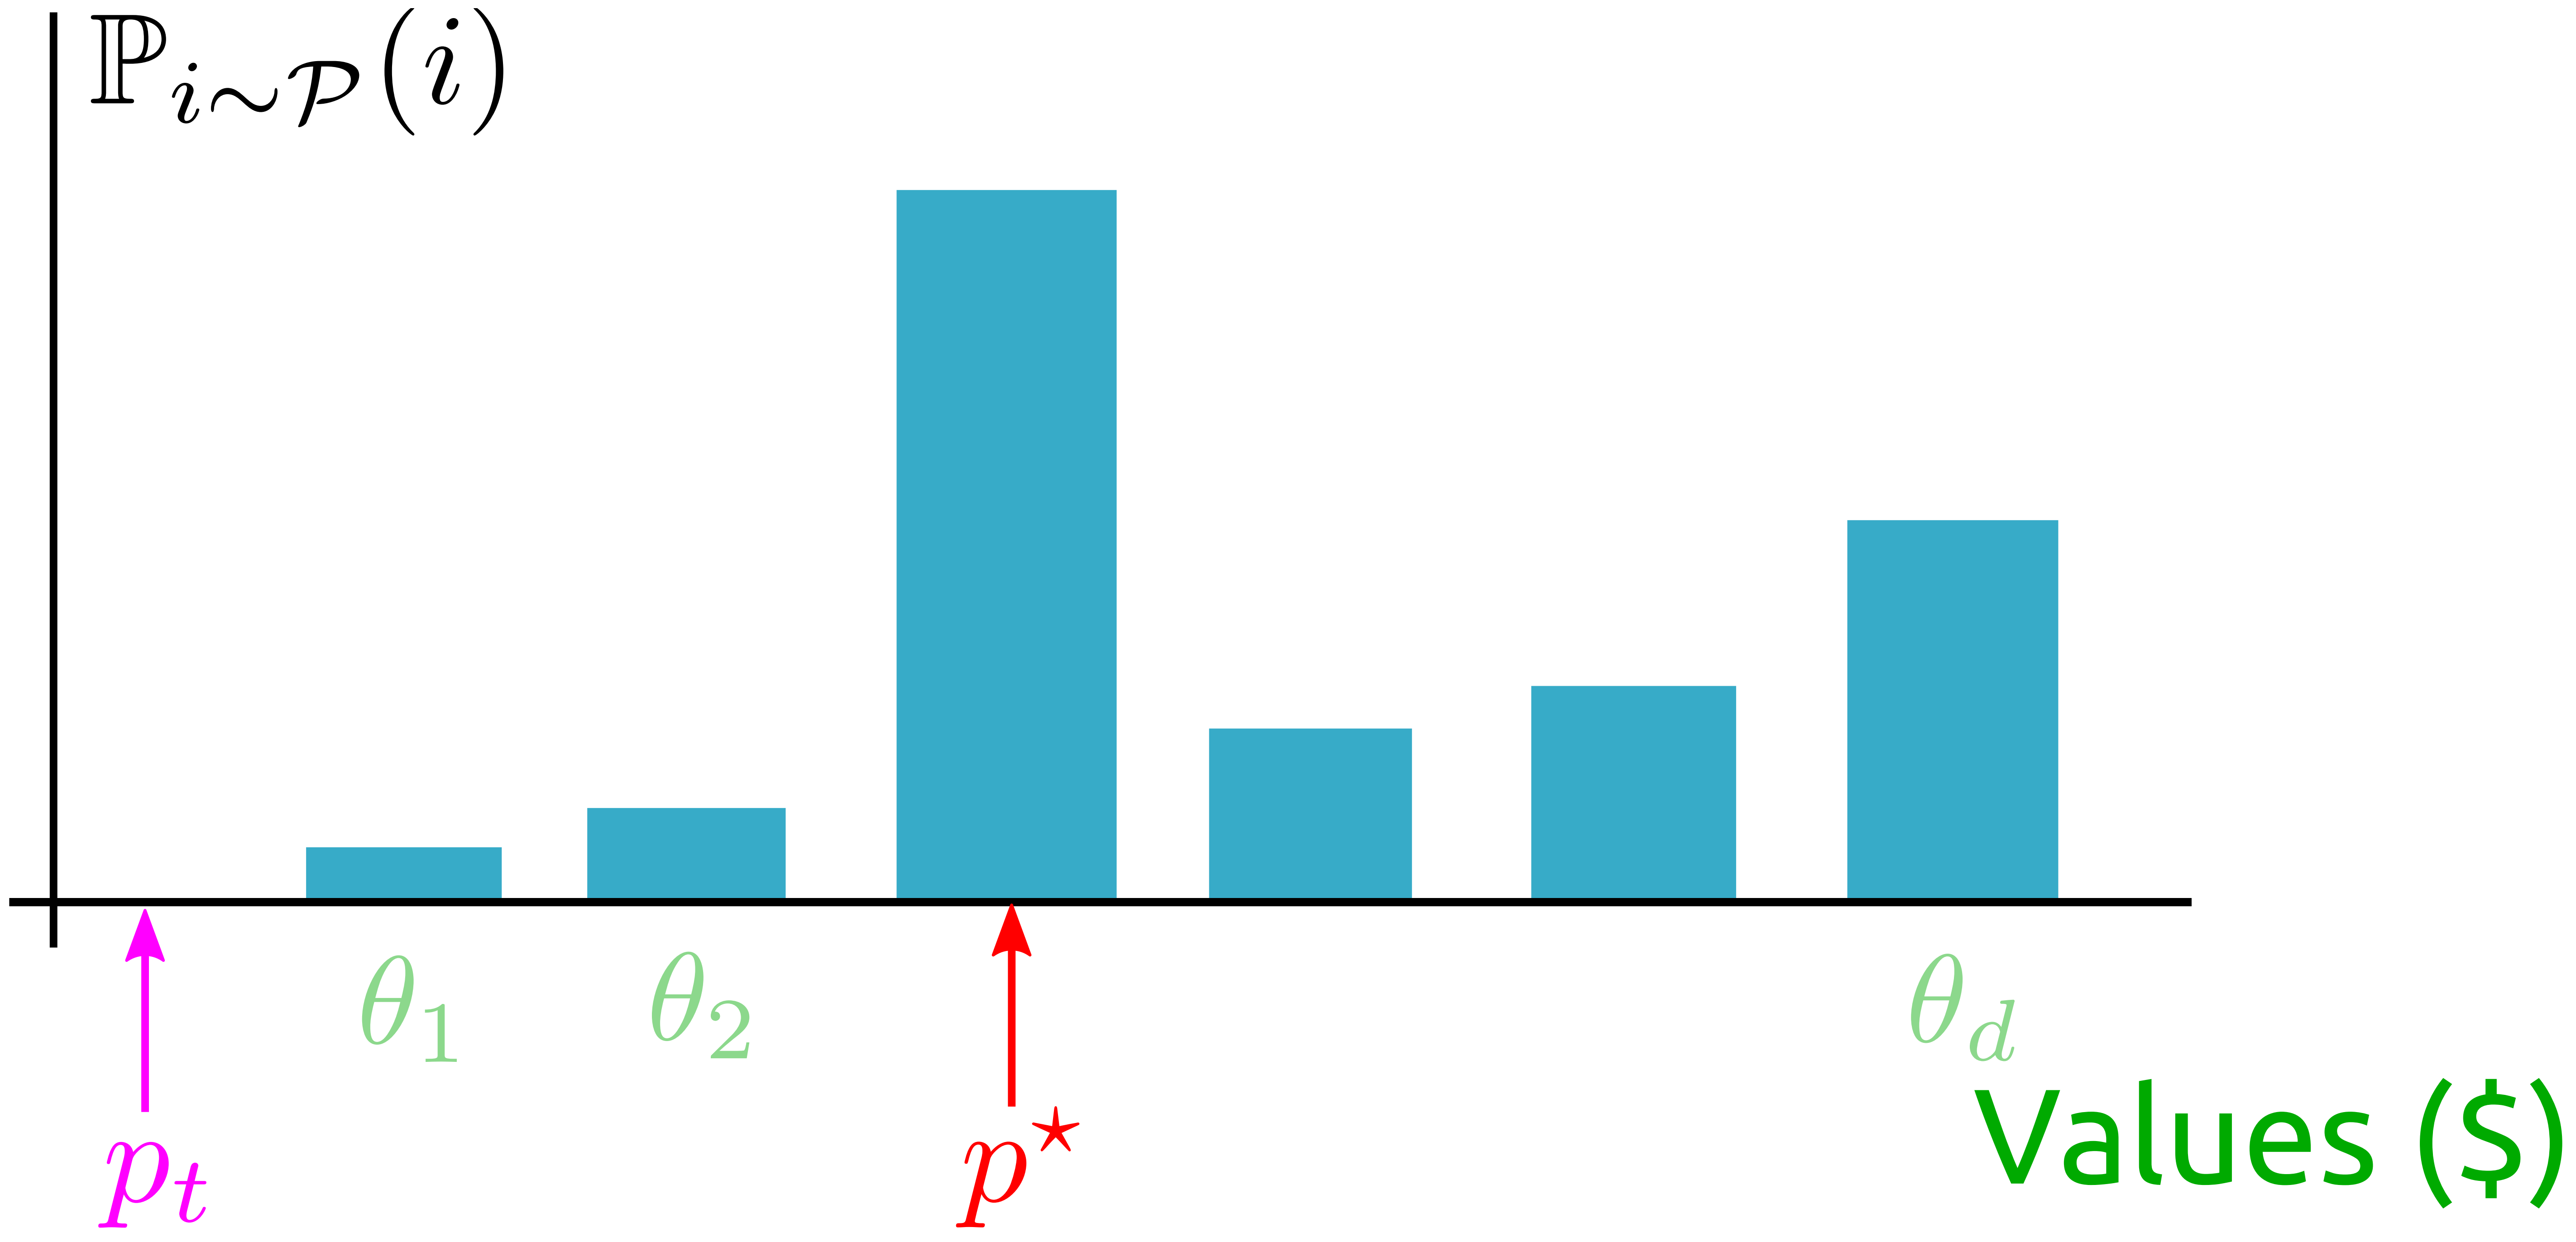

- Set prices too high $\Longrightarrow$ no feedback about low value types.
- Set prices too low $\Longrightarrow$ low revenue.

CHALLENGE 1: PRICING VS SELLER LEARNING

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

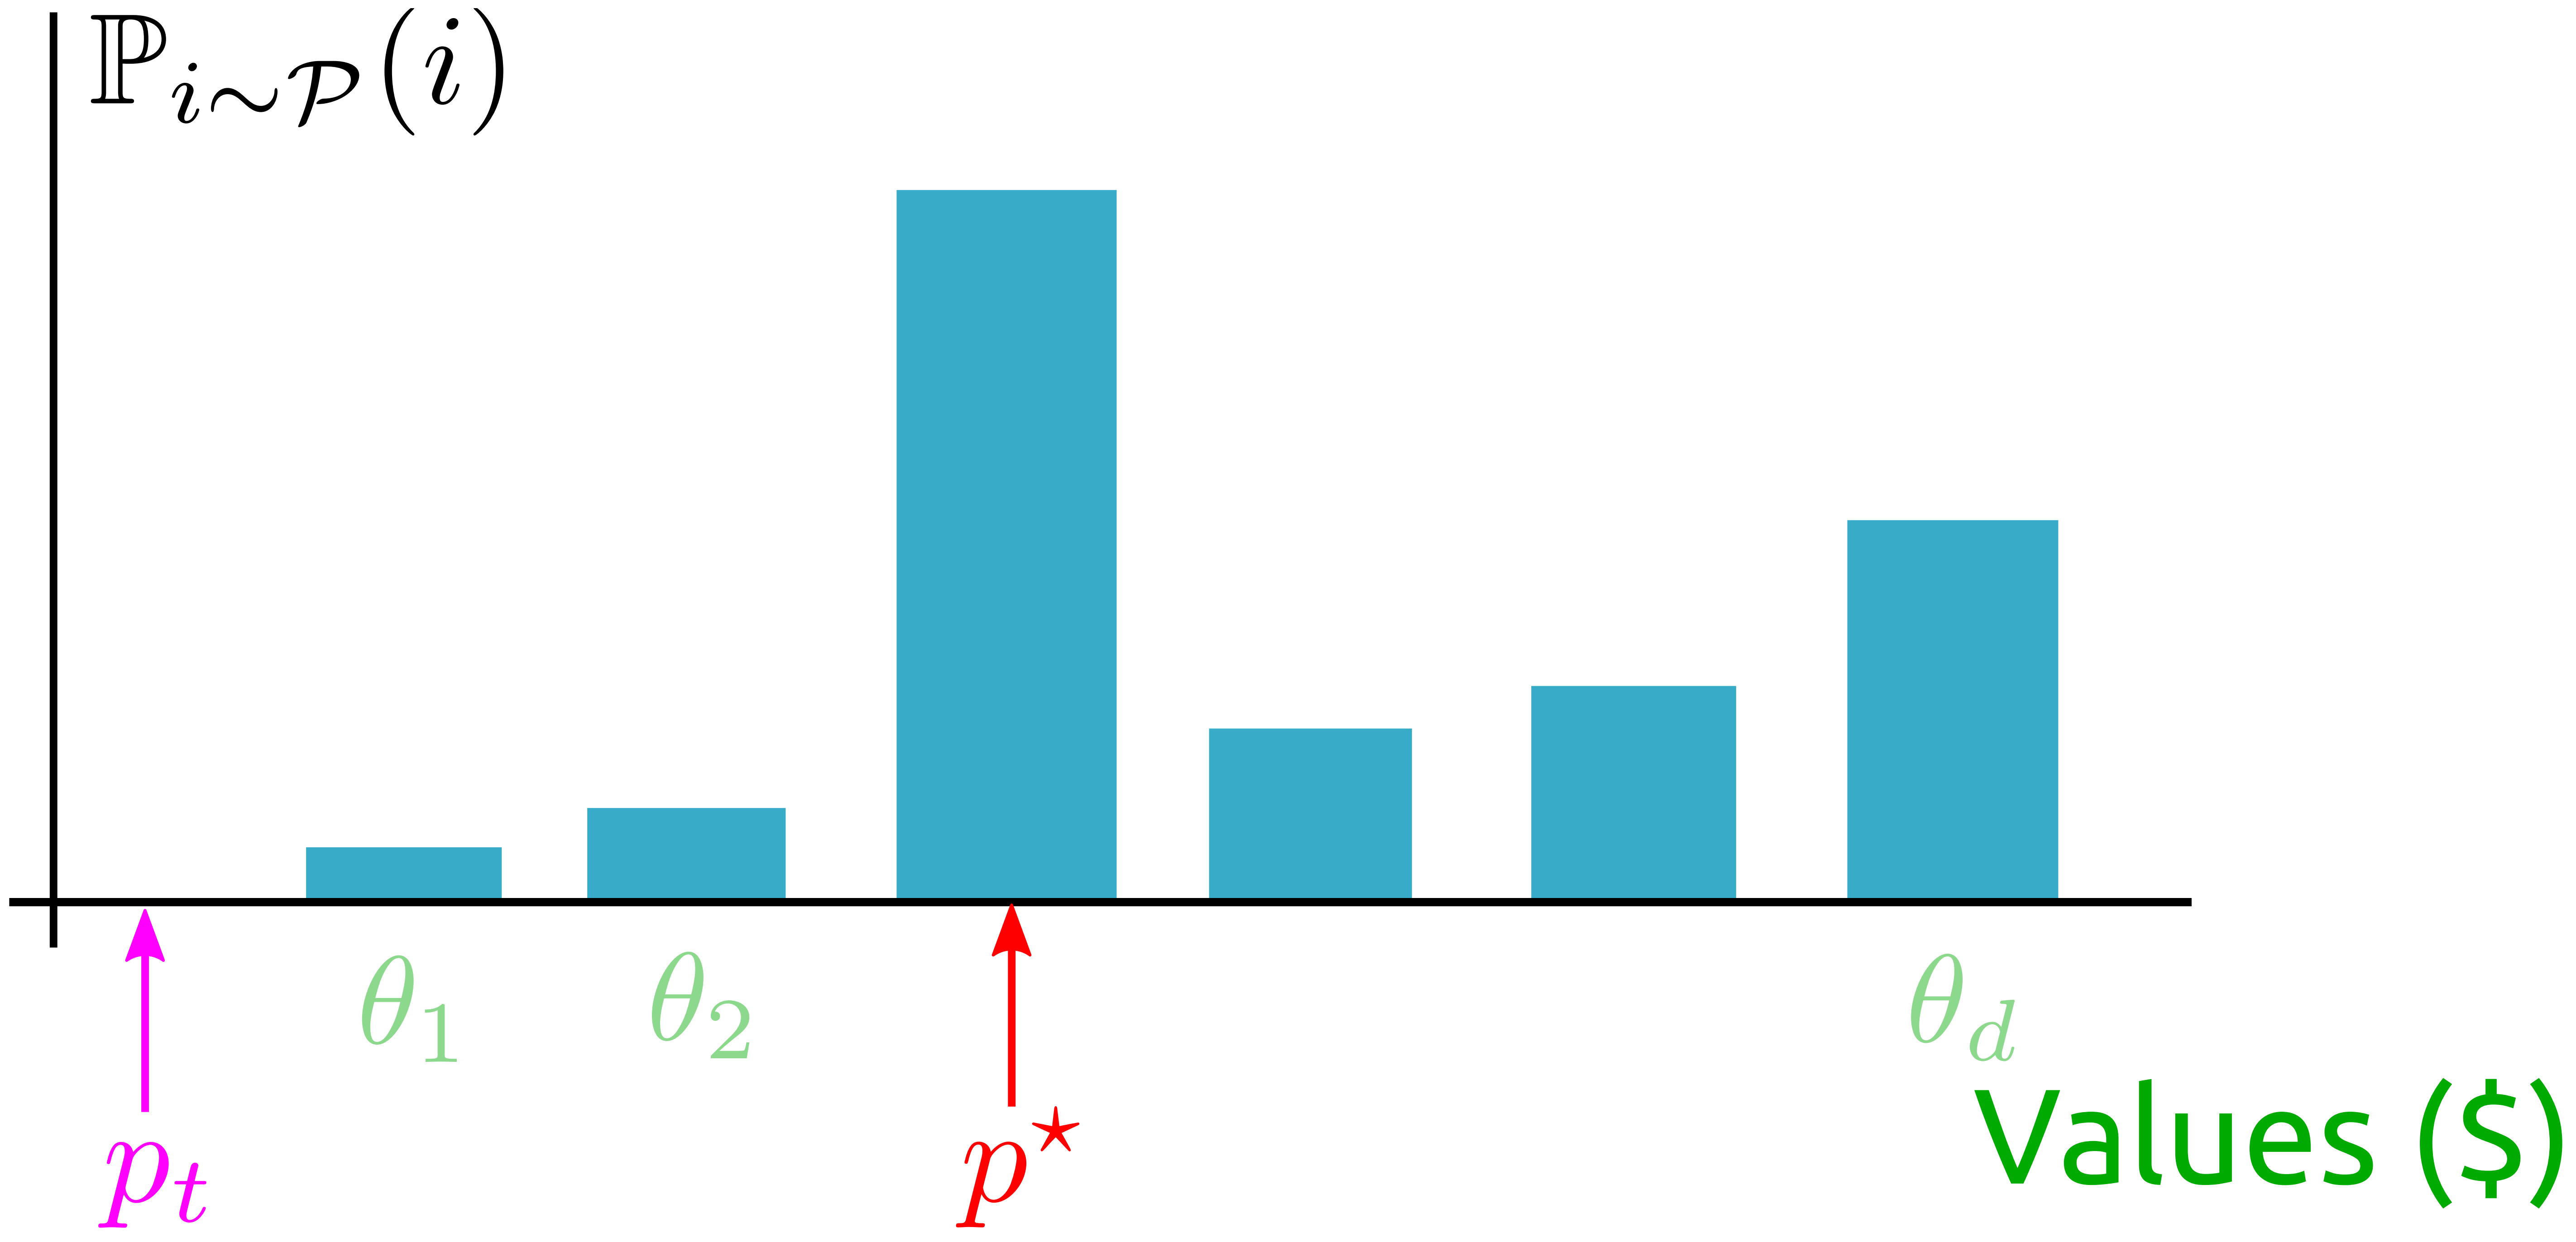

- Set prices too high $\implies$ no feedback about low value types.
- Set prices too low $\implies$ low revenue.
- *Property*: if $p_t \leq p^\star$, and buyers know values, sufficient feedback to learn $p^\star$.

CHALLENGE 1: PRICING VS SELLER LEARNING

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

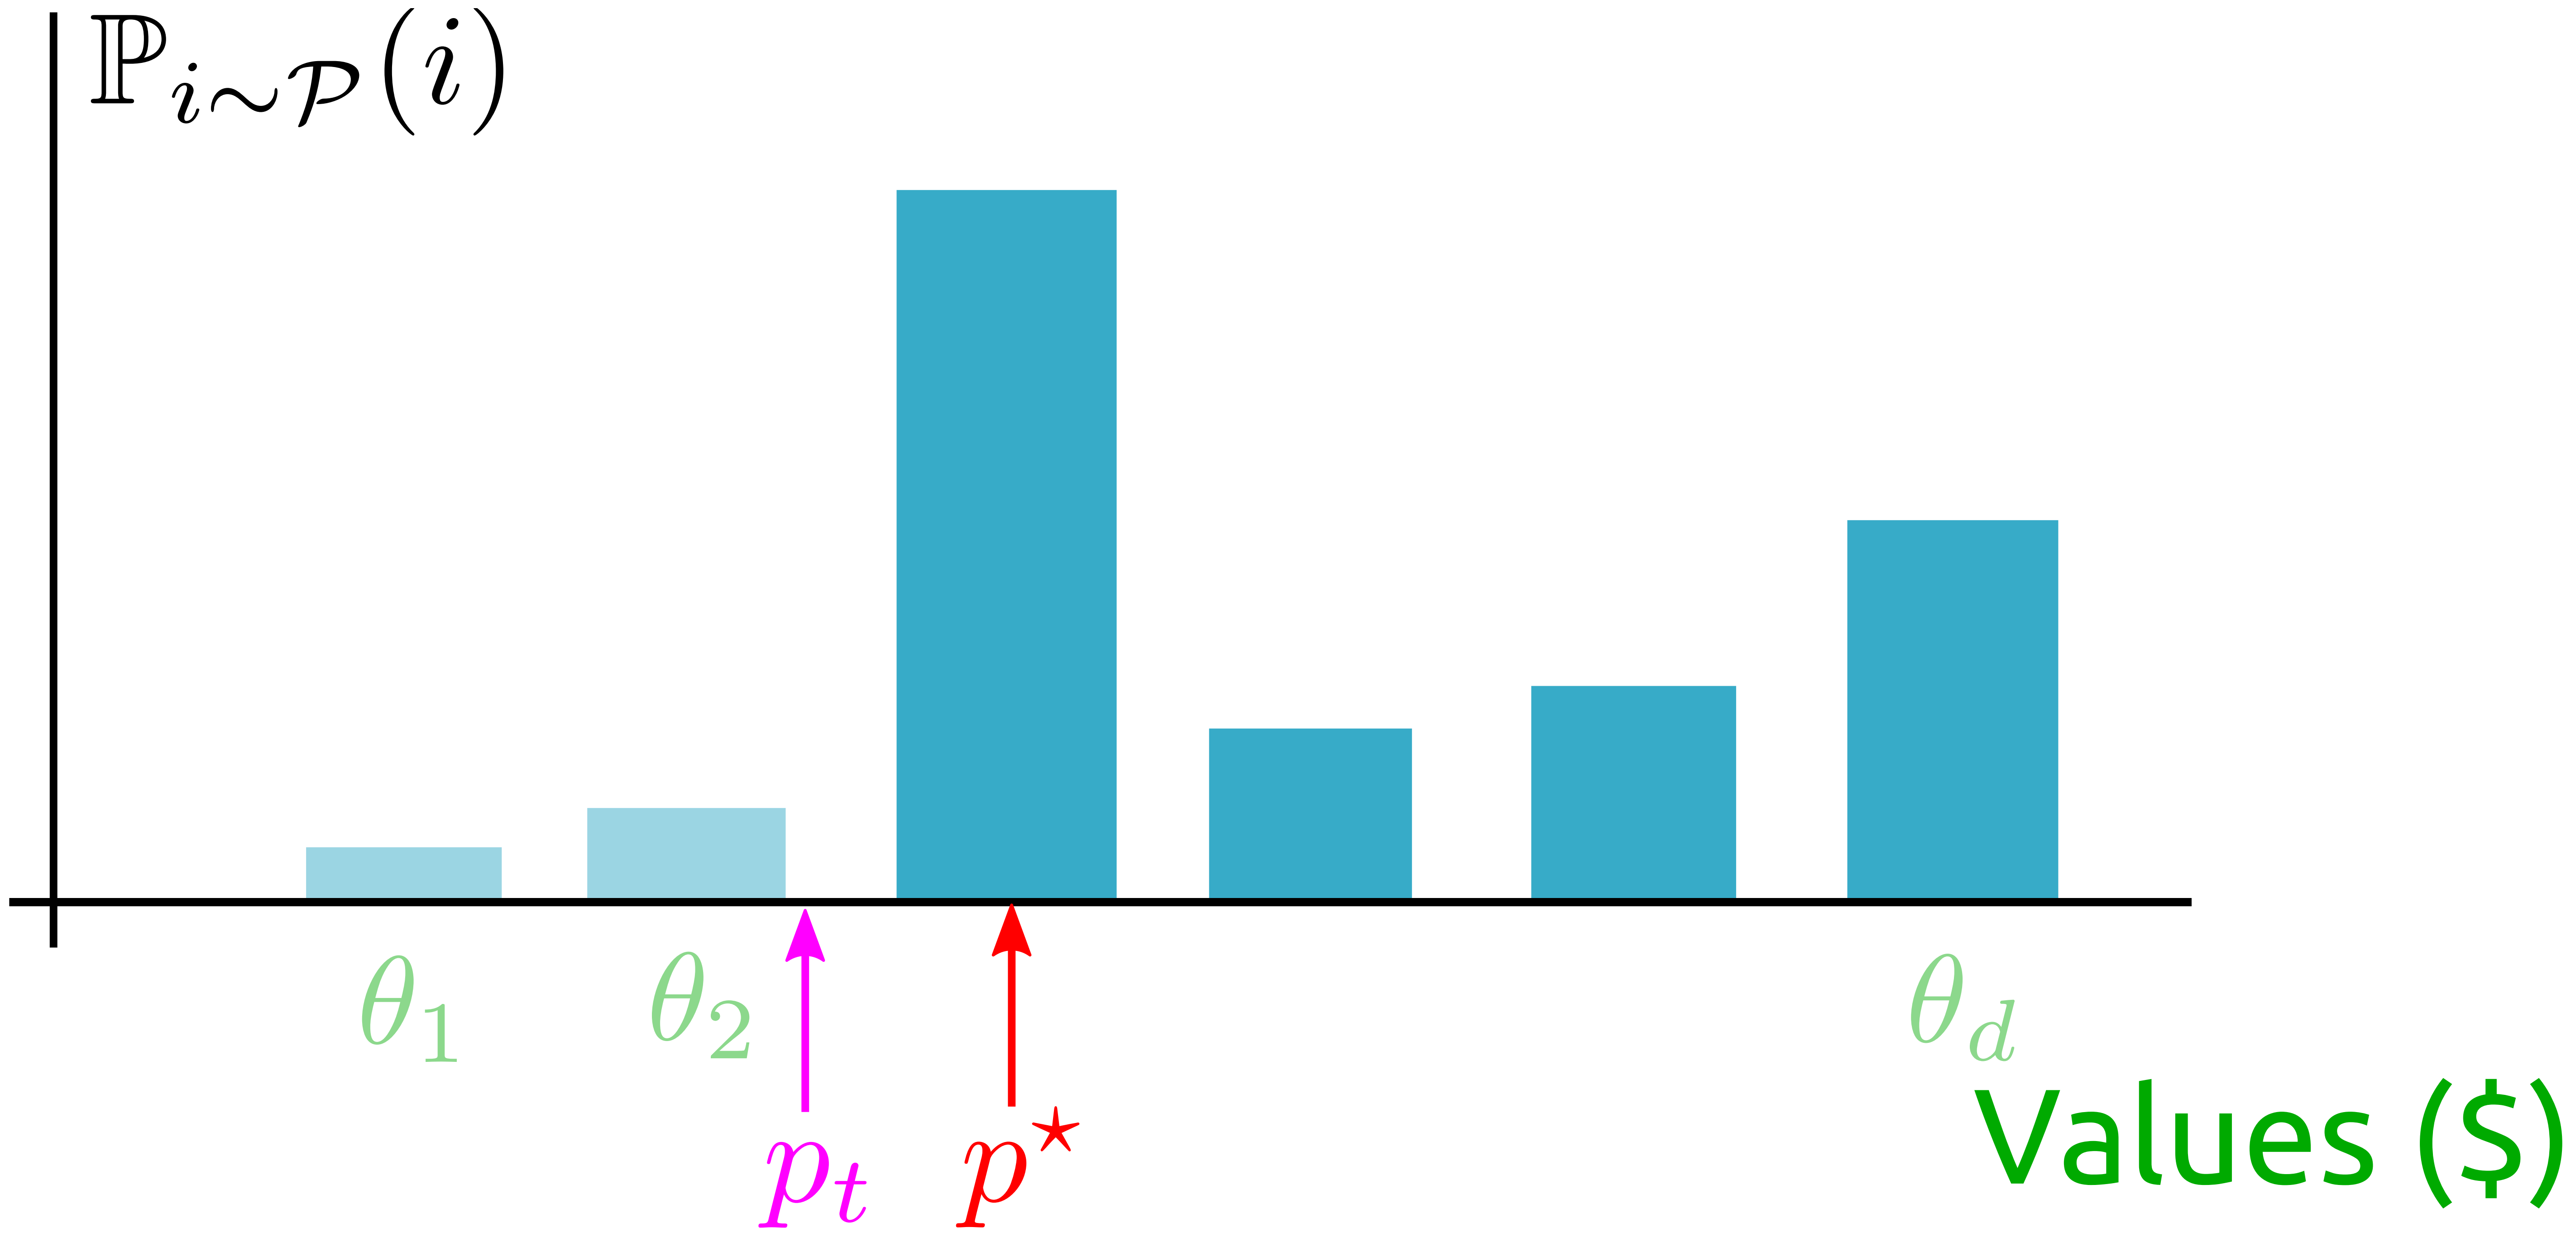

- Set prices too high $\implies$ no feedback about low value types.
- Set prices too low $\implies$ low revenue.
- *Property*: if $p_t \leq p^\star$, and buyers know values, sufficient feedback to learn $p^\star$.

CHALLENGE 2: PRICING VS BUYER LEARNING

CHALLENGE 2: PRICING VS BUYER LEARNING

- Seller chooses price *before* seeing the customer type.

CHALLENGE 2: PRICING VS BUYER LEARNING

- Seller chooses price *before* seeing the customer type.
- 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

- Seller chooses price *before* seeing the customer type.
- A buyer’s purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type 2	★★★★☆ $47	★★★★★ $51	★★★☆☆ $37
(stovetop)	★★★☆☆ $39	★★★☆☆ $38	★★★★★ $52
	★★★★☆ $42	★★★★☆ $44	★★★★☆ $46
	★★★★★ $53	★★★★☆ $45	★★★★★ $56

CHALLENGE 2: PRICING VS BUYER LEARNING

- Seller chooses price *before* seeing the customer type.
- 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

- Seller chooses price *before* seeing the customer type.
- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type 2

(stovetop)

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

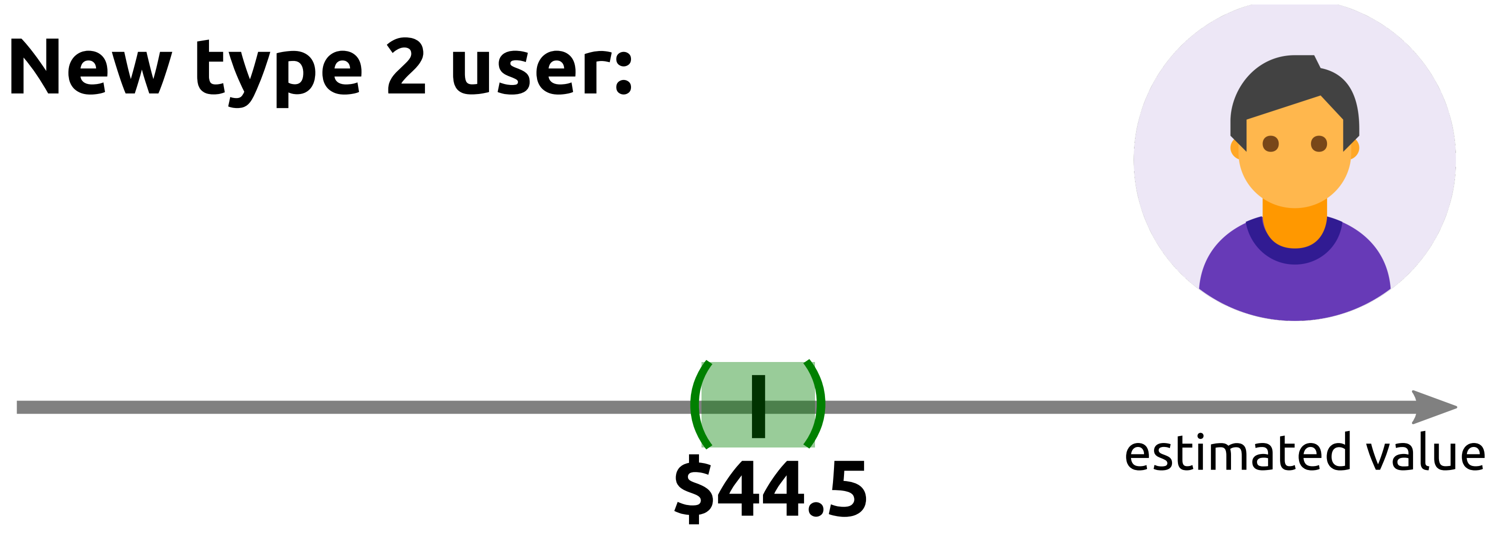

CHALLENGE 2: PRICING VS BUYER LEARNING

- Seller chooses price *before* seeing the customer type.
- 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

- Seller chooses price *before* seeing the customer type.
- 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

- Seller chooses price *before* seeing the customer type.
- 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

- Seller chooses price *before* seeing the customer type.
- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type 2 (stovetop)	★★★★☆ $47	★★★★★ $51	★★★☆☆ $37
	★★★☆☆ $39	★★★☆☆ $38	★★★★★ $52
	★★★★☆ $42	★★★★☆ $44	★★★★☆ $46
	★★★★★ $53	★★★★☆ $45	★★★★★ $56
Type 3 (grill)	★★★★★ $53		
	★★★★★ $57		

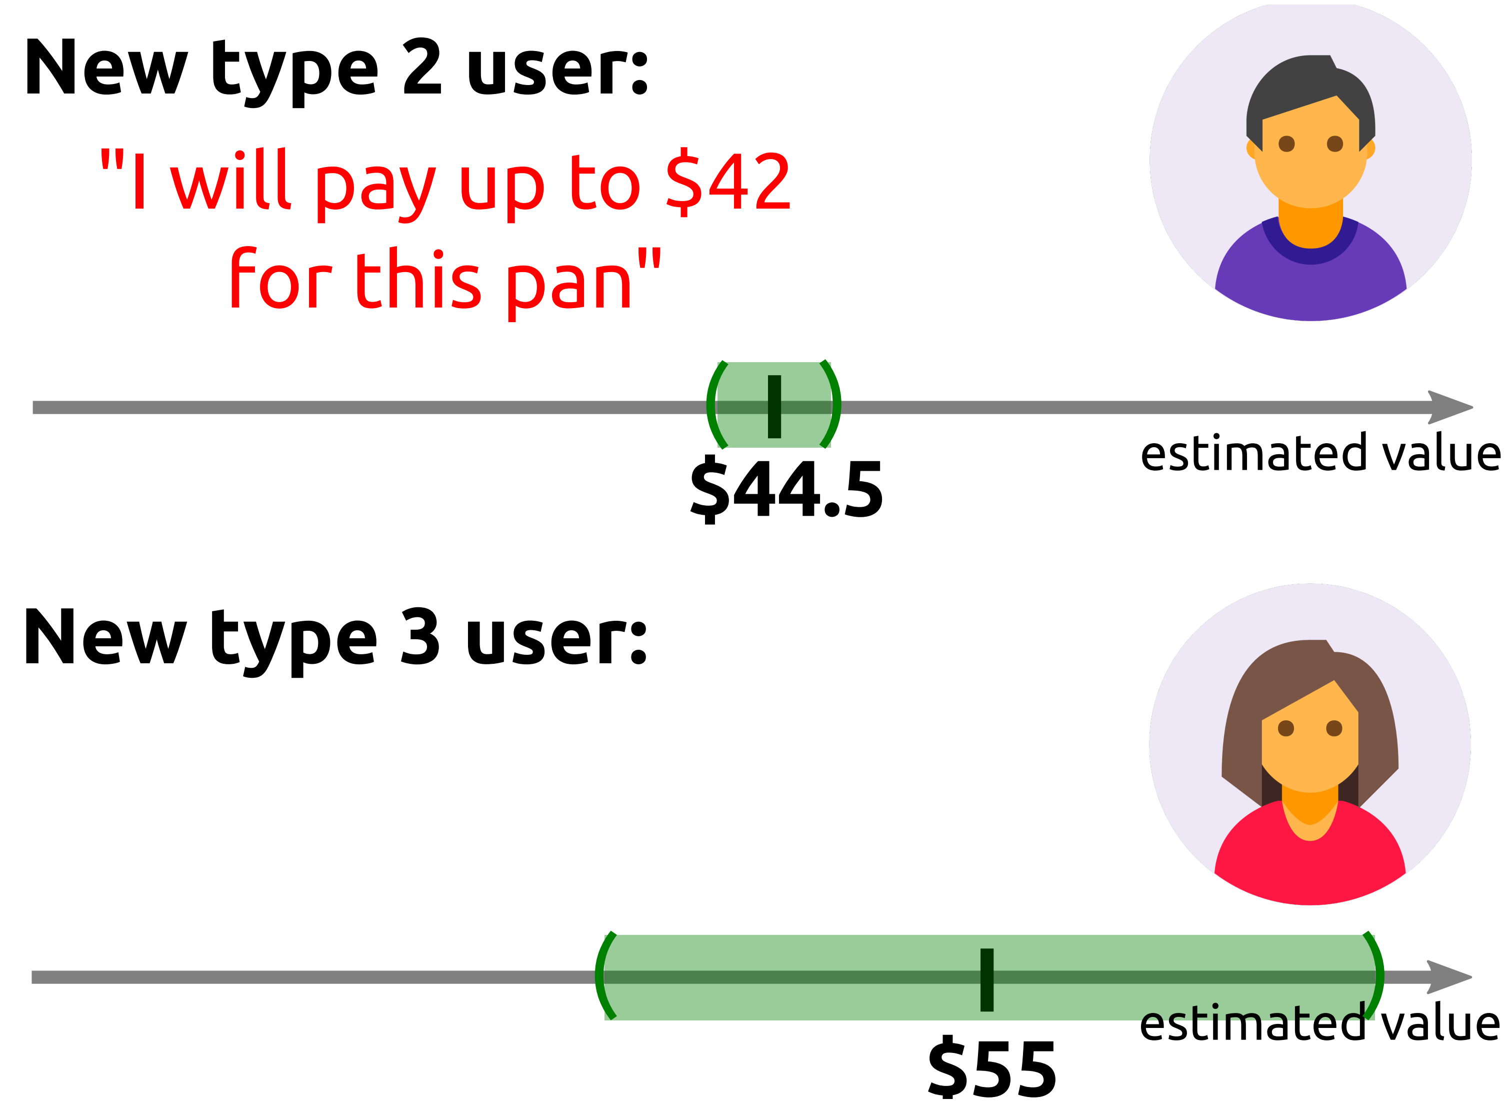

CHALLENGE 2: PRICING VS BUYER LEARNING

- Seller chooses price *before* seeing the customer type.
- 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

- Seller chooses price *before* seeing the customer type.
- A buyer's purchase decision depends on how certain she is of her value. This in turn depends on previous reviews.

Type 2 (stovetop)	★★★★☆ $47	★★★★★ $51	★★★☆☆ $37
	★★★☆☆ $39	★★★☆☆ $38	★★★★★ $52
	★★★★☆ $42	★★★★☆ $44	★★★★☆ $46
	★★★★★ $53	★★★★☆ $45	★★★★★ $56
Type 3 (grill)	★★★★★ $53		
	★★★★★ $57		

New type 2 user:

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

New type 3 user:

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

- Seller's dilemma: Only target type 1 buyers for high immediate revenue? Or also target type 3 customers for higher long term revenue?

OVERVIEW OF ALGORITHM & RESULTS

OVERVIEW OF ALGORITHM & RESULTS

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

OVERVIEW OF ALGORITHM & RESULTS

- ***Algorithmic insights:***
  - 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.

OUTLINE

1. **Problem set up**
   - Online learning framework, assumptions, challenges

2. **Algorithm**

3. **Theoretical results**
   - Upper bounds, lower bounds, proof sketches

ALGORITHM OVERVIEW

ALGORITHM OVERVIEW

- On each round t, maintain a set S_t of types

ALGORITHM OVERVIEW

- On each round t, maintain a set S_t of types
  1. Have sufficiently high value

ALGORITHM OVERVIEW

- On each round t, maintain a set S_t of types
  1. Have sufficiently high value
  2. Are not exceedingly rare

ALGORITHM OVERVIEW

- On each round t, maintain a set S_t of types
  1. Have sufficiently high value (high average ex-post value)
  2. Are not exceedingly rare (appeared frequently enough in the past)

- Both, based on past reviews.

ALGORITHM OVERVIEW

- On each round t, maintain a set S_t of types
  1. Have sufficiently high value (high average ex-post value)
  2. Are not exceedingly rare (appeared frequently enough in the past)
- Both, based on past reviews.
- On round t, set price so that all customers of types in S_t will buy.

ALGORITHM

- **Phase 1:**

- **Phase 2:**

ALGORITHM

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

- **Phase 2:**

ALGORITHM

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

ALGORITHM

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

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(d^{2/3}T^{1/3})$ rounds)
  - Offer item for a “very low price”
  - Observe iid samples from type distribution $\mathscr{P}$.
  - Set Q to be the set of types that appeared often enough
- **Phase 2:**

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(d^{2/3}T^{1/3})$ rounds)
    - Offer item for a "very low price"
    - Observe iid samples from type distribution $\mathscr{P}$.
    - Set Q to be the set of types that appeared often enough
- **Phase 2:** (set $S_t = Q$)

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(d^{2/3}T^{1/3})$ rounds)
    - Offer item for a "very low price"
    - Observe iid samples from type distribution $\mathscr{P}$.
    - Set Q to be the set of types that appeared often enough
- **Phase 2:** (set $S_t = Q$)
    - Set price p_t *low enough* that buyers in S_t will buy.

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(d^{2/3}T^{1/3})$ rounds)
  - Offer item for a "very low price"
  - Observe iid samples from type distribution $\mathscr{P}$.
  - Set Q to be the set of types that appeared often enough

- **Phase 2:** (set $S_t = Q$)
  - Set price p_t *low enough* that buyers in S_t will buy.
    - *low enough:* account for buyer uncertainty (η-risky).

ALGORITHM

- **Phase 1:** (for a small number of rounds, $\sim \tilde{\Theta}(d^{2/3}T^{1/3})$ rounds)
  - Offer item for a "very low price"
  - Observe iid samples from type distribution $\mathscr{P}$.
  - Set Q to be the set of types that appeared often enough
- **Phase 2:** (set $S_t = Q$)
  - Set price p_t *low enough* that buyers in S_t will buy.
    - *low enough:* account for buyer uncertainty (η-risky).
  - Update S_t: eliminate types which contribute too little to revenue.

PHASE 2: UPDATING S_t

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i\sim\mathscr{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

- Maintain confidence intervals for $\mathrm{rev}(\theta_i, Q)$ for each $i \in S_t$.

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i\sim\mathscr{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

- Maintain confidence intervals for $\mathrm{rev}(\theta_i, Q)$ for each $i \in S_t$.

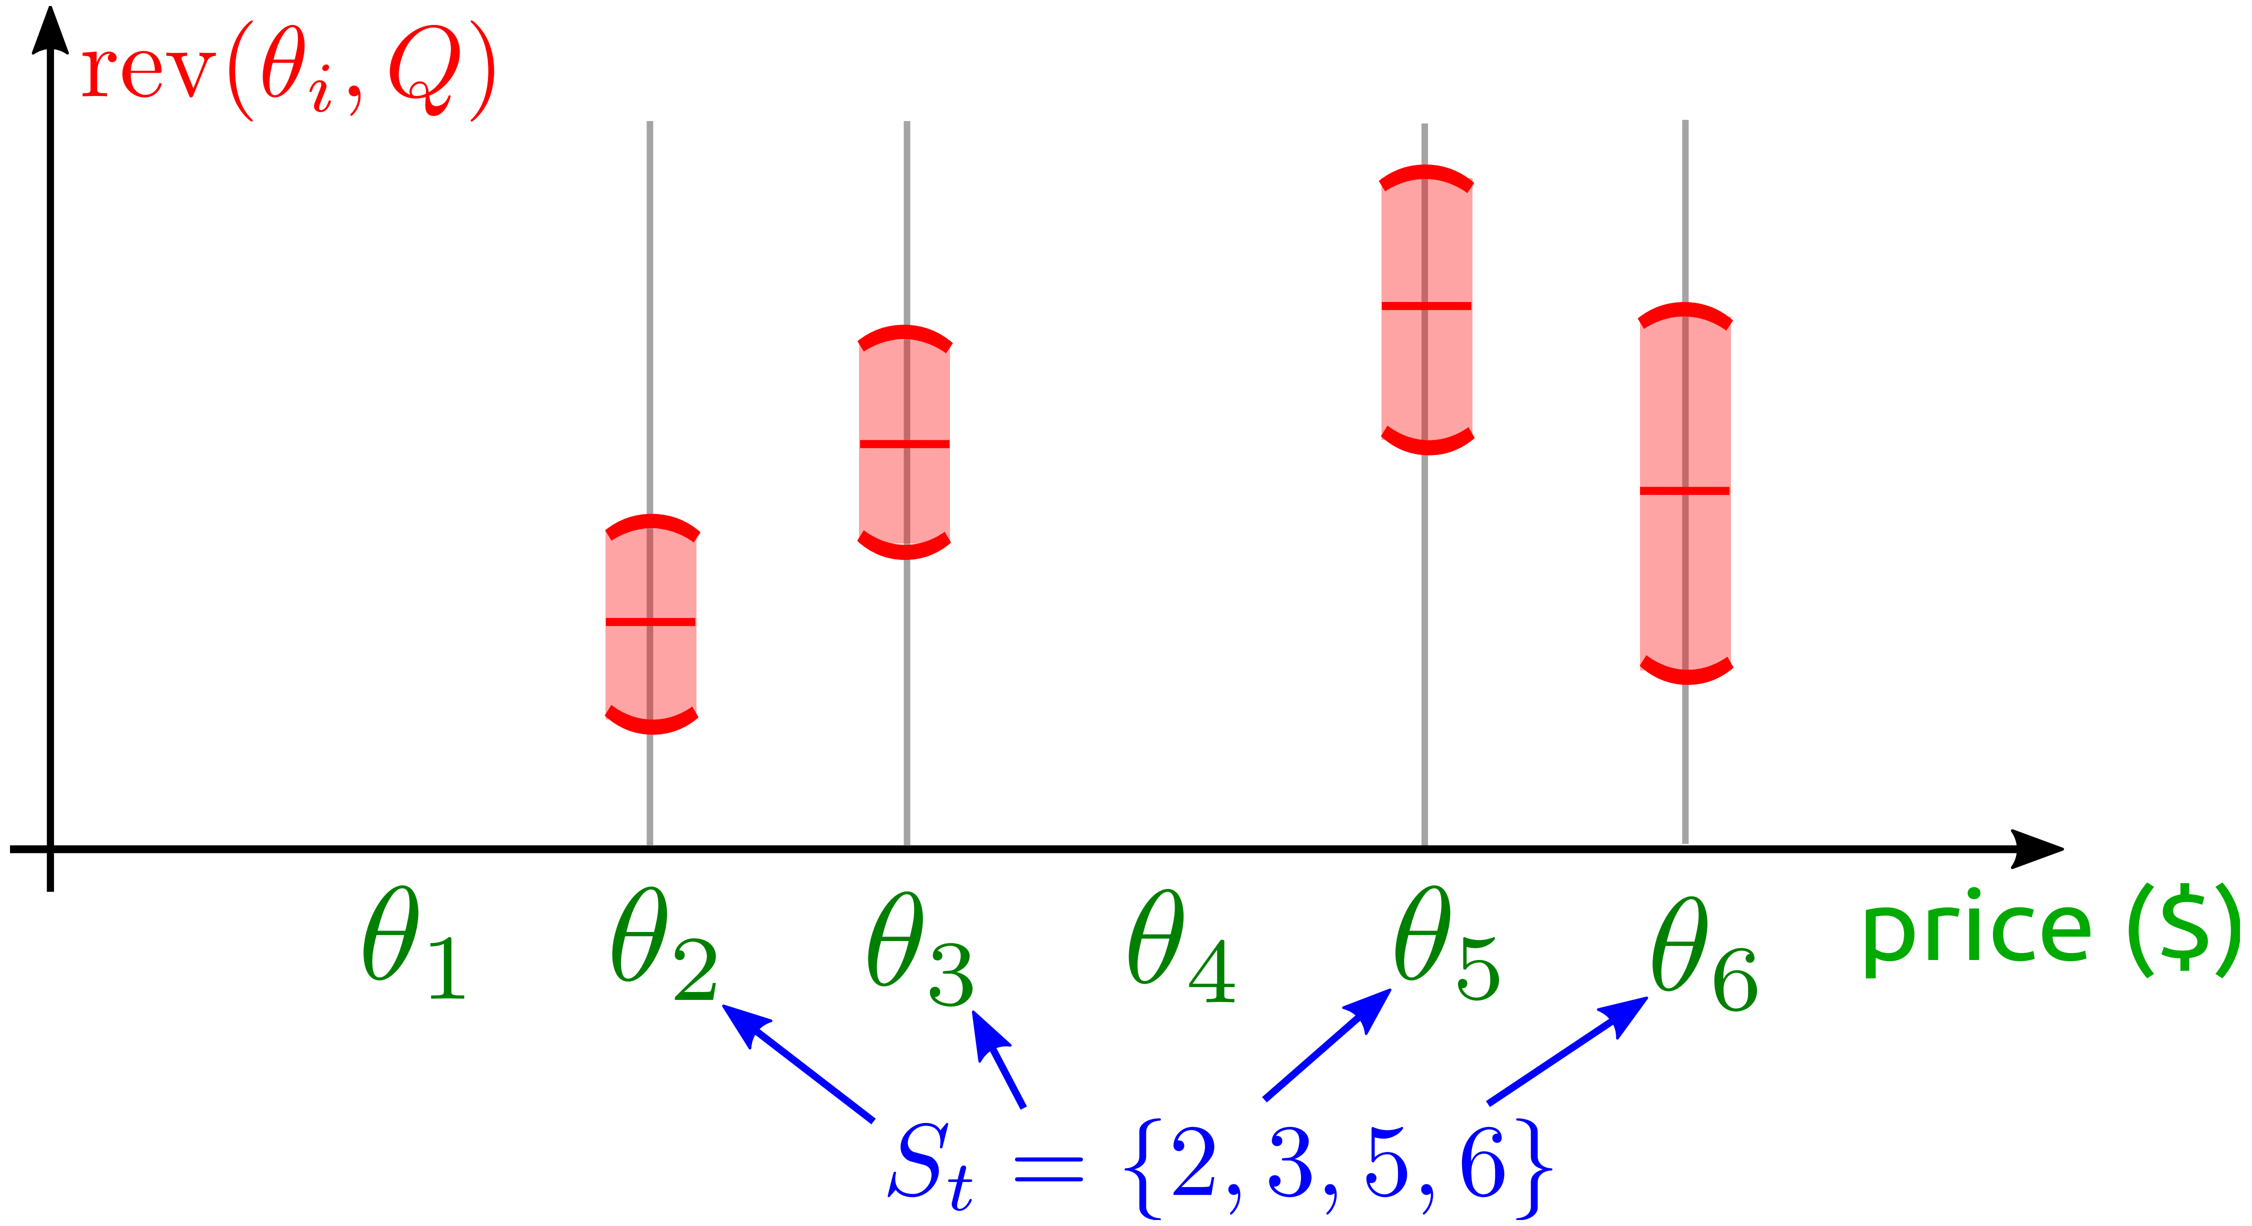

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i \sim \mathcal{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

- Maintain confidence intervals for $\mathrm{rev}(\theta_i, Q)$ for each $i \in S_t$.

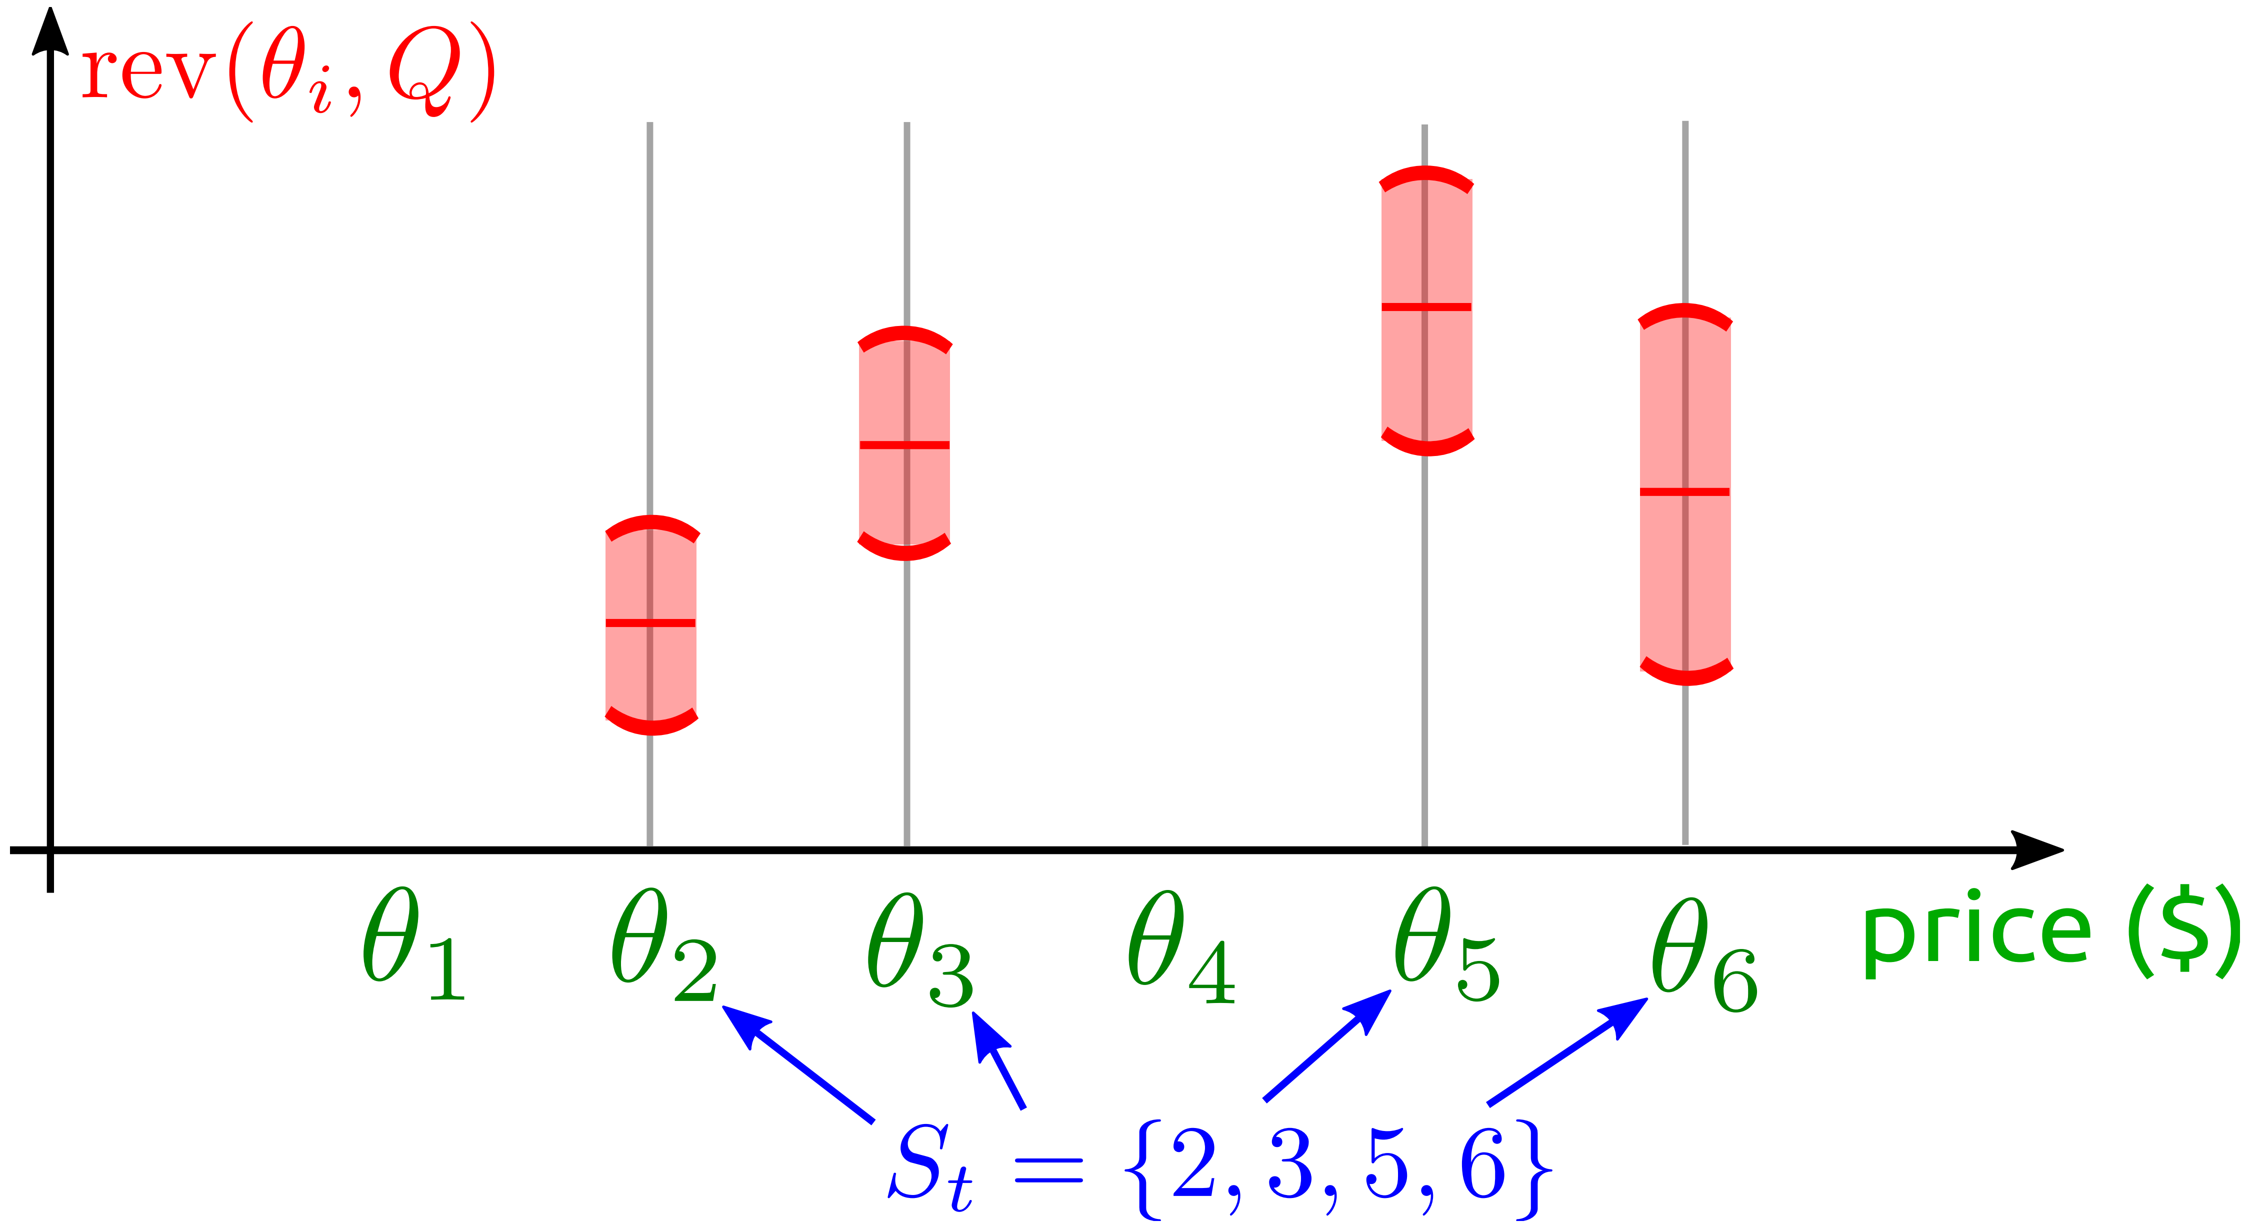

- Eliminate i from S_t if it's UCB is smaller than the highest LCB *(essentially)*.

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i \sim \mathscr{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

- Maintain confidence intervals for $\mathrm{rev}(\theta_i, Q)$ for each $i \in S_t$.

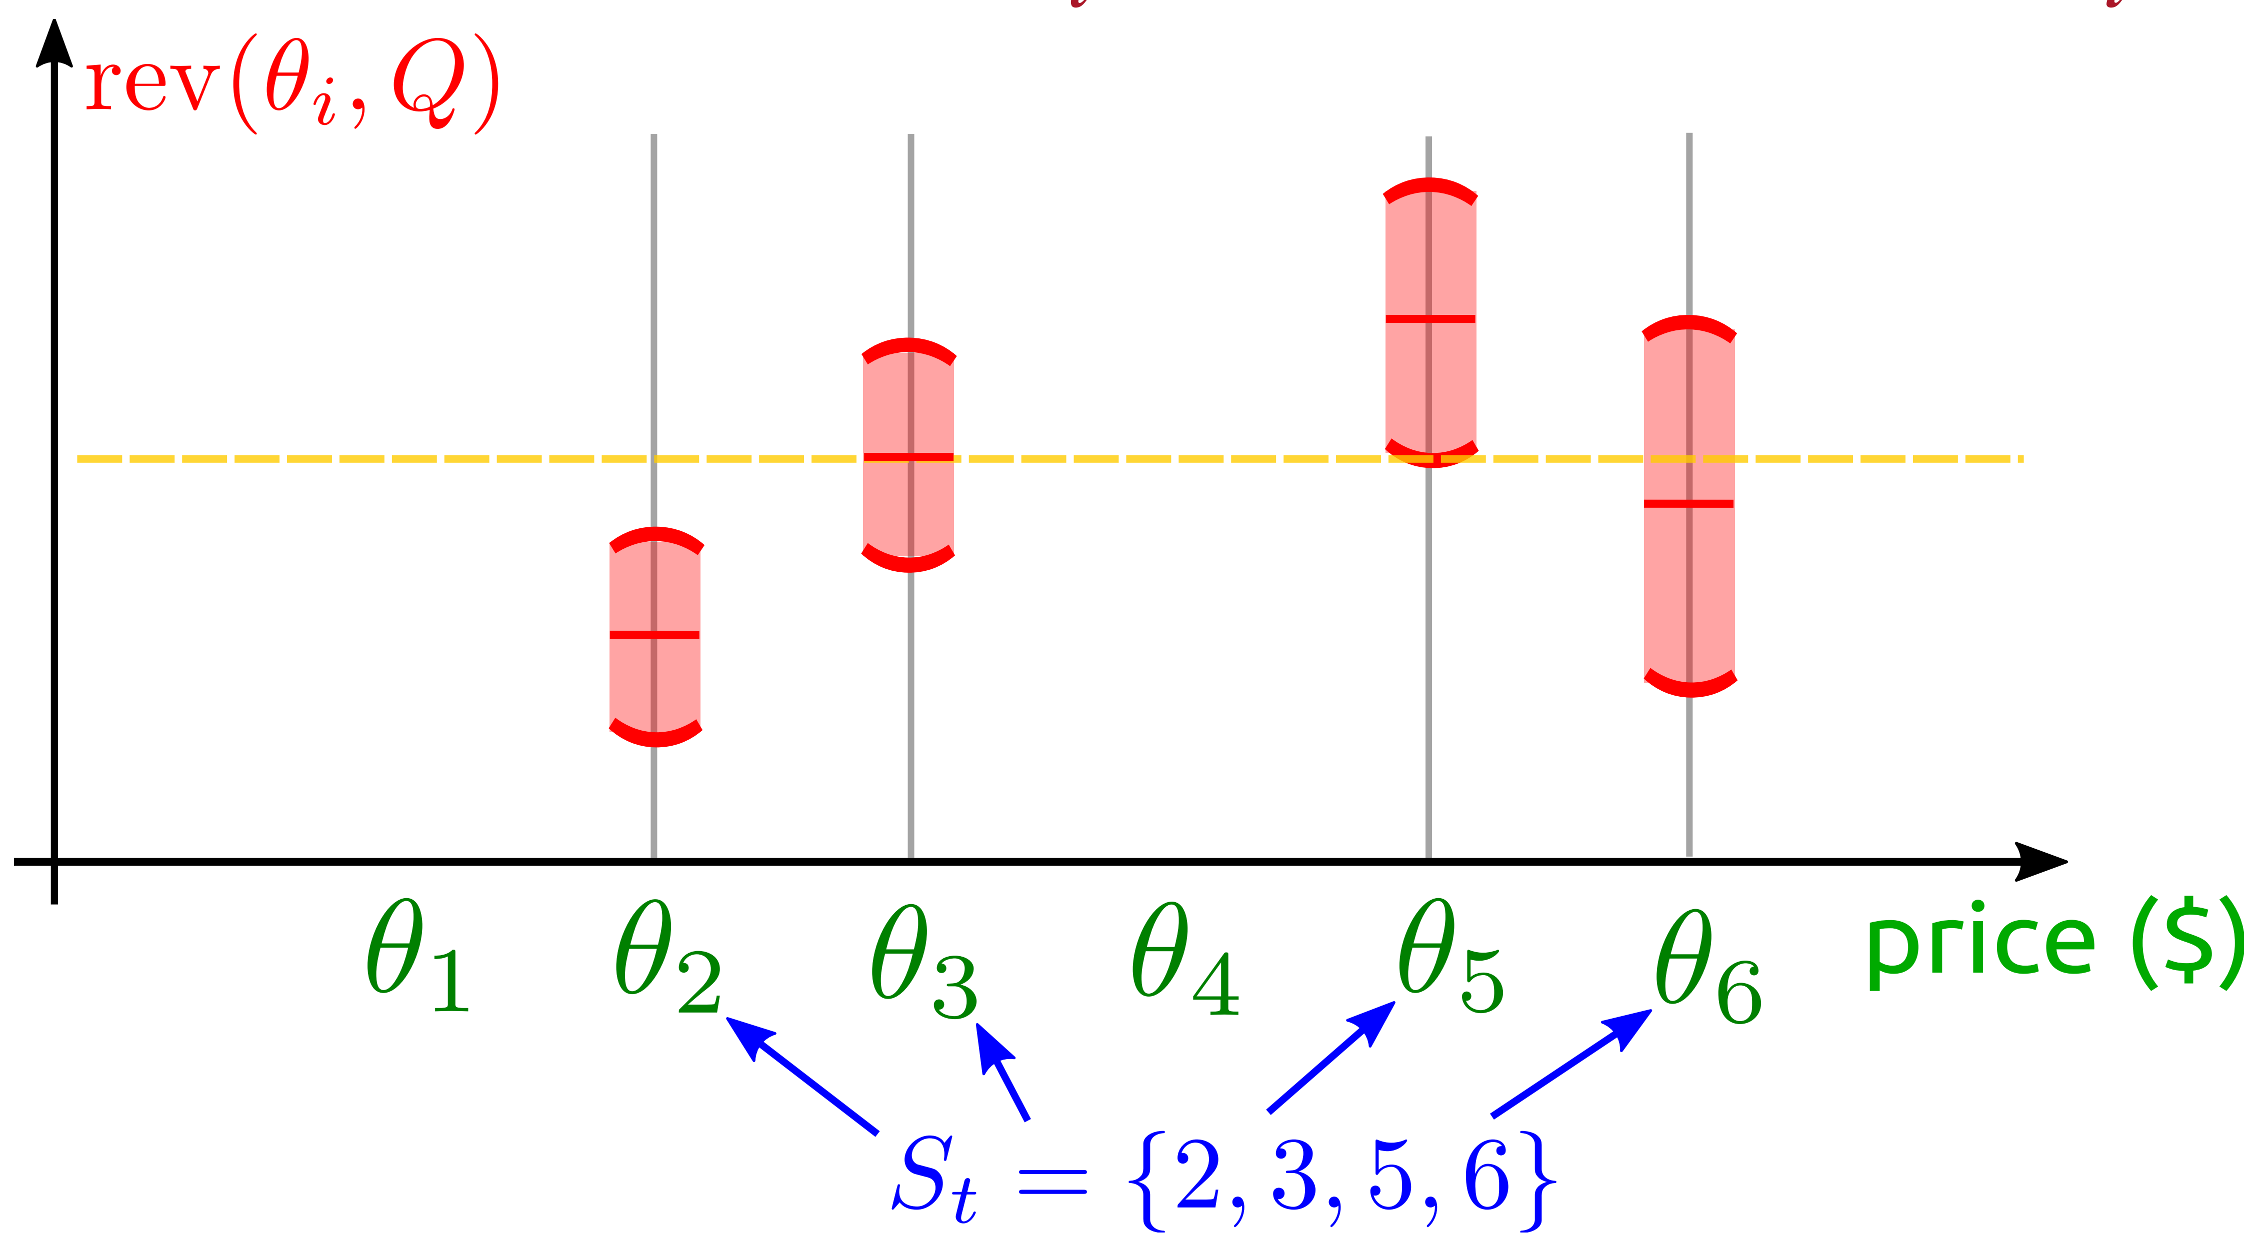

- Eliminate i from S_t if it's UCB is smaller than the highest LCB (*essentially*).

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i\sim\mathscr{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

- Maintain confidence intervals for $\mathrm{rev}(\theta_i, Q)$ for each $i \in S_t$.

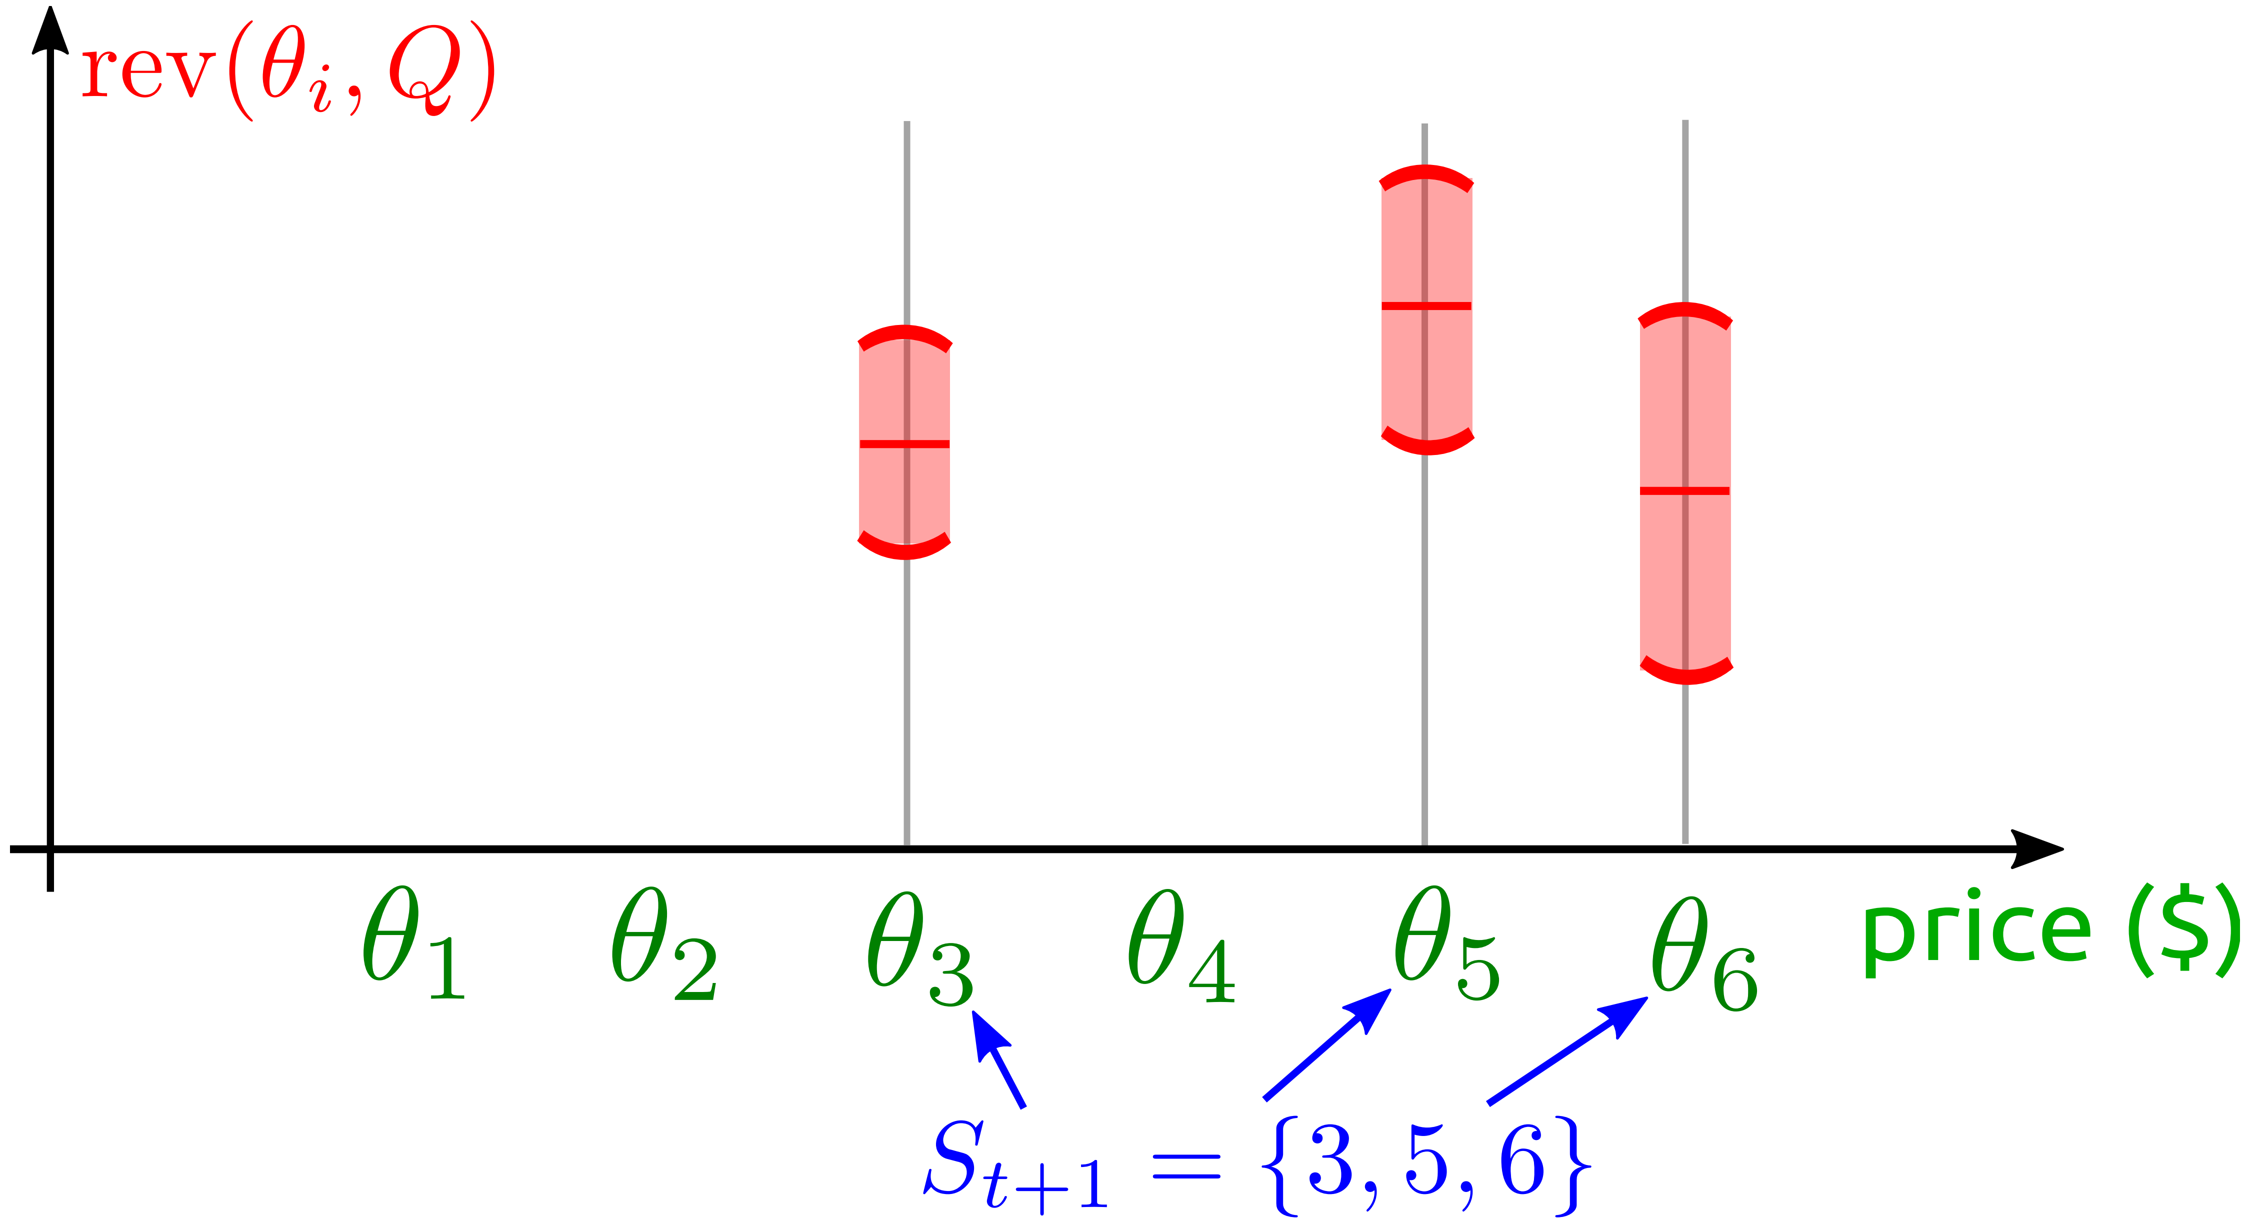

- Eliminate i from S_t if it's UCB is smaller than the highest LCB *(essentially)*.

PHASE 2: UPDATING S_t

- We will estimate the best price $p^\star(Q)$ for types in Q, instead of $p^\star$

$$p^\star(Q) = \arg\max_p \ \mathrm{rev}(p, Q) = \arg\max_p \ p \cdot \mathbb{P}_{i \sim \mathcal{P}}(\theta_i \geq p \ \text{ and } \ i \in Q)$$

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

- Maintain confidence intervals for $\mathrm{rev}(\theta_i, Q)$ for each $i \in S_t$.

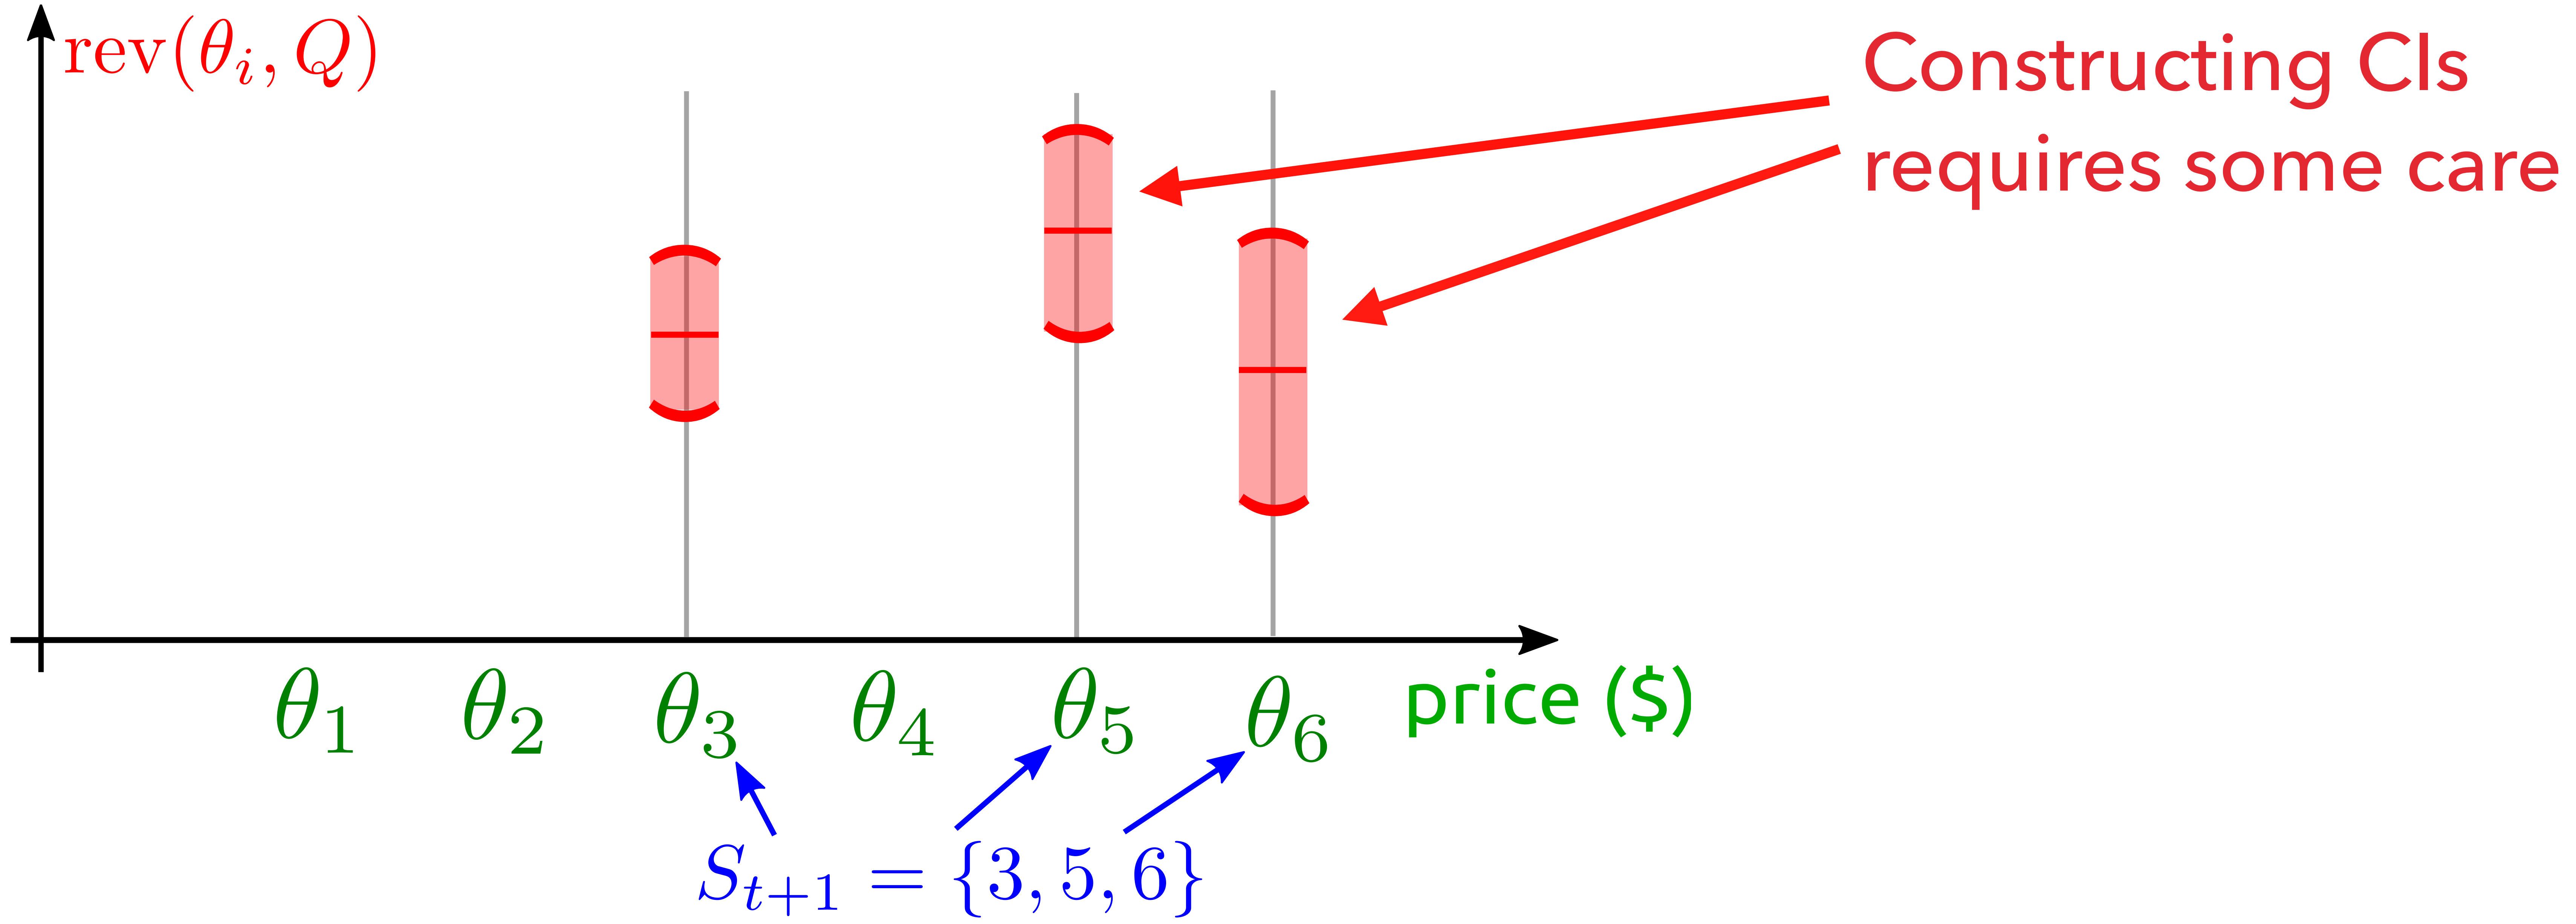

- Eliminate i from S_t if it's UCB is smaller than the highest LCB (*essentially*).

PHASE 2: CHOOSING PRICE p_t

PHASE 2: CHOOSING PRICE p_t

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

PHASE 2: CHOOSING PRICE p_t

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

PHASE 2: CHOOSING PRICE p_t

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

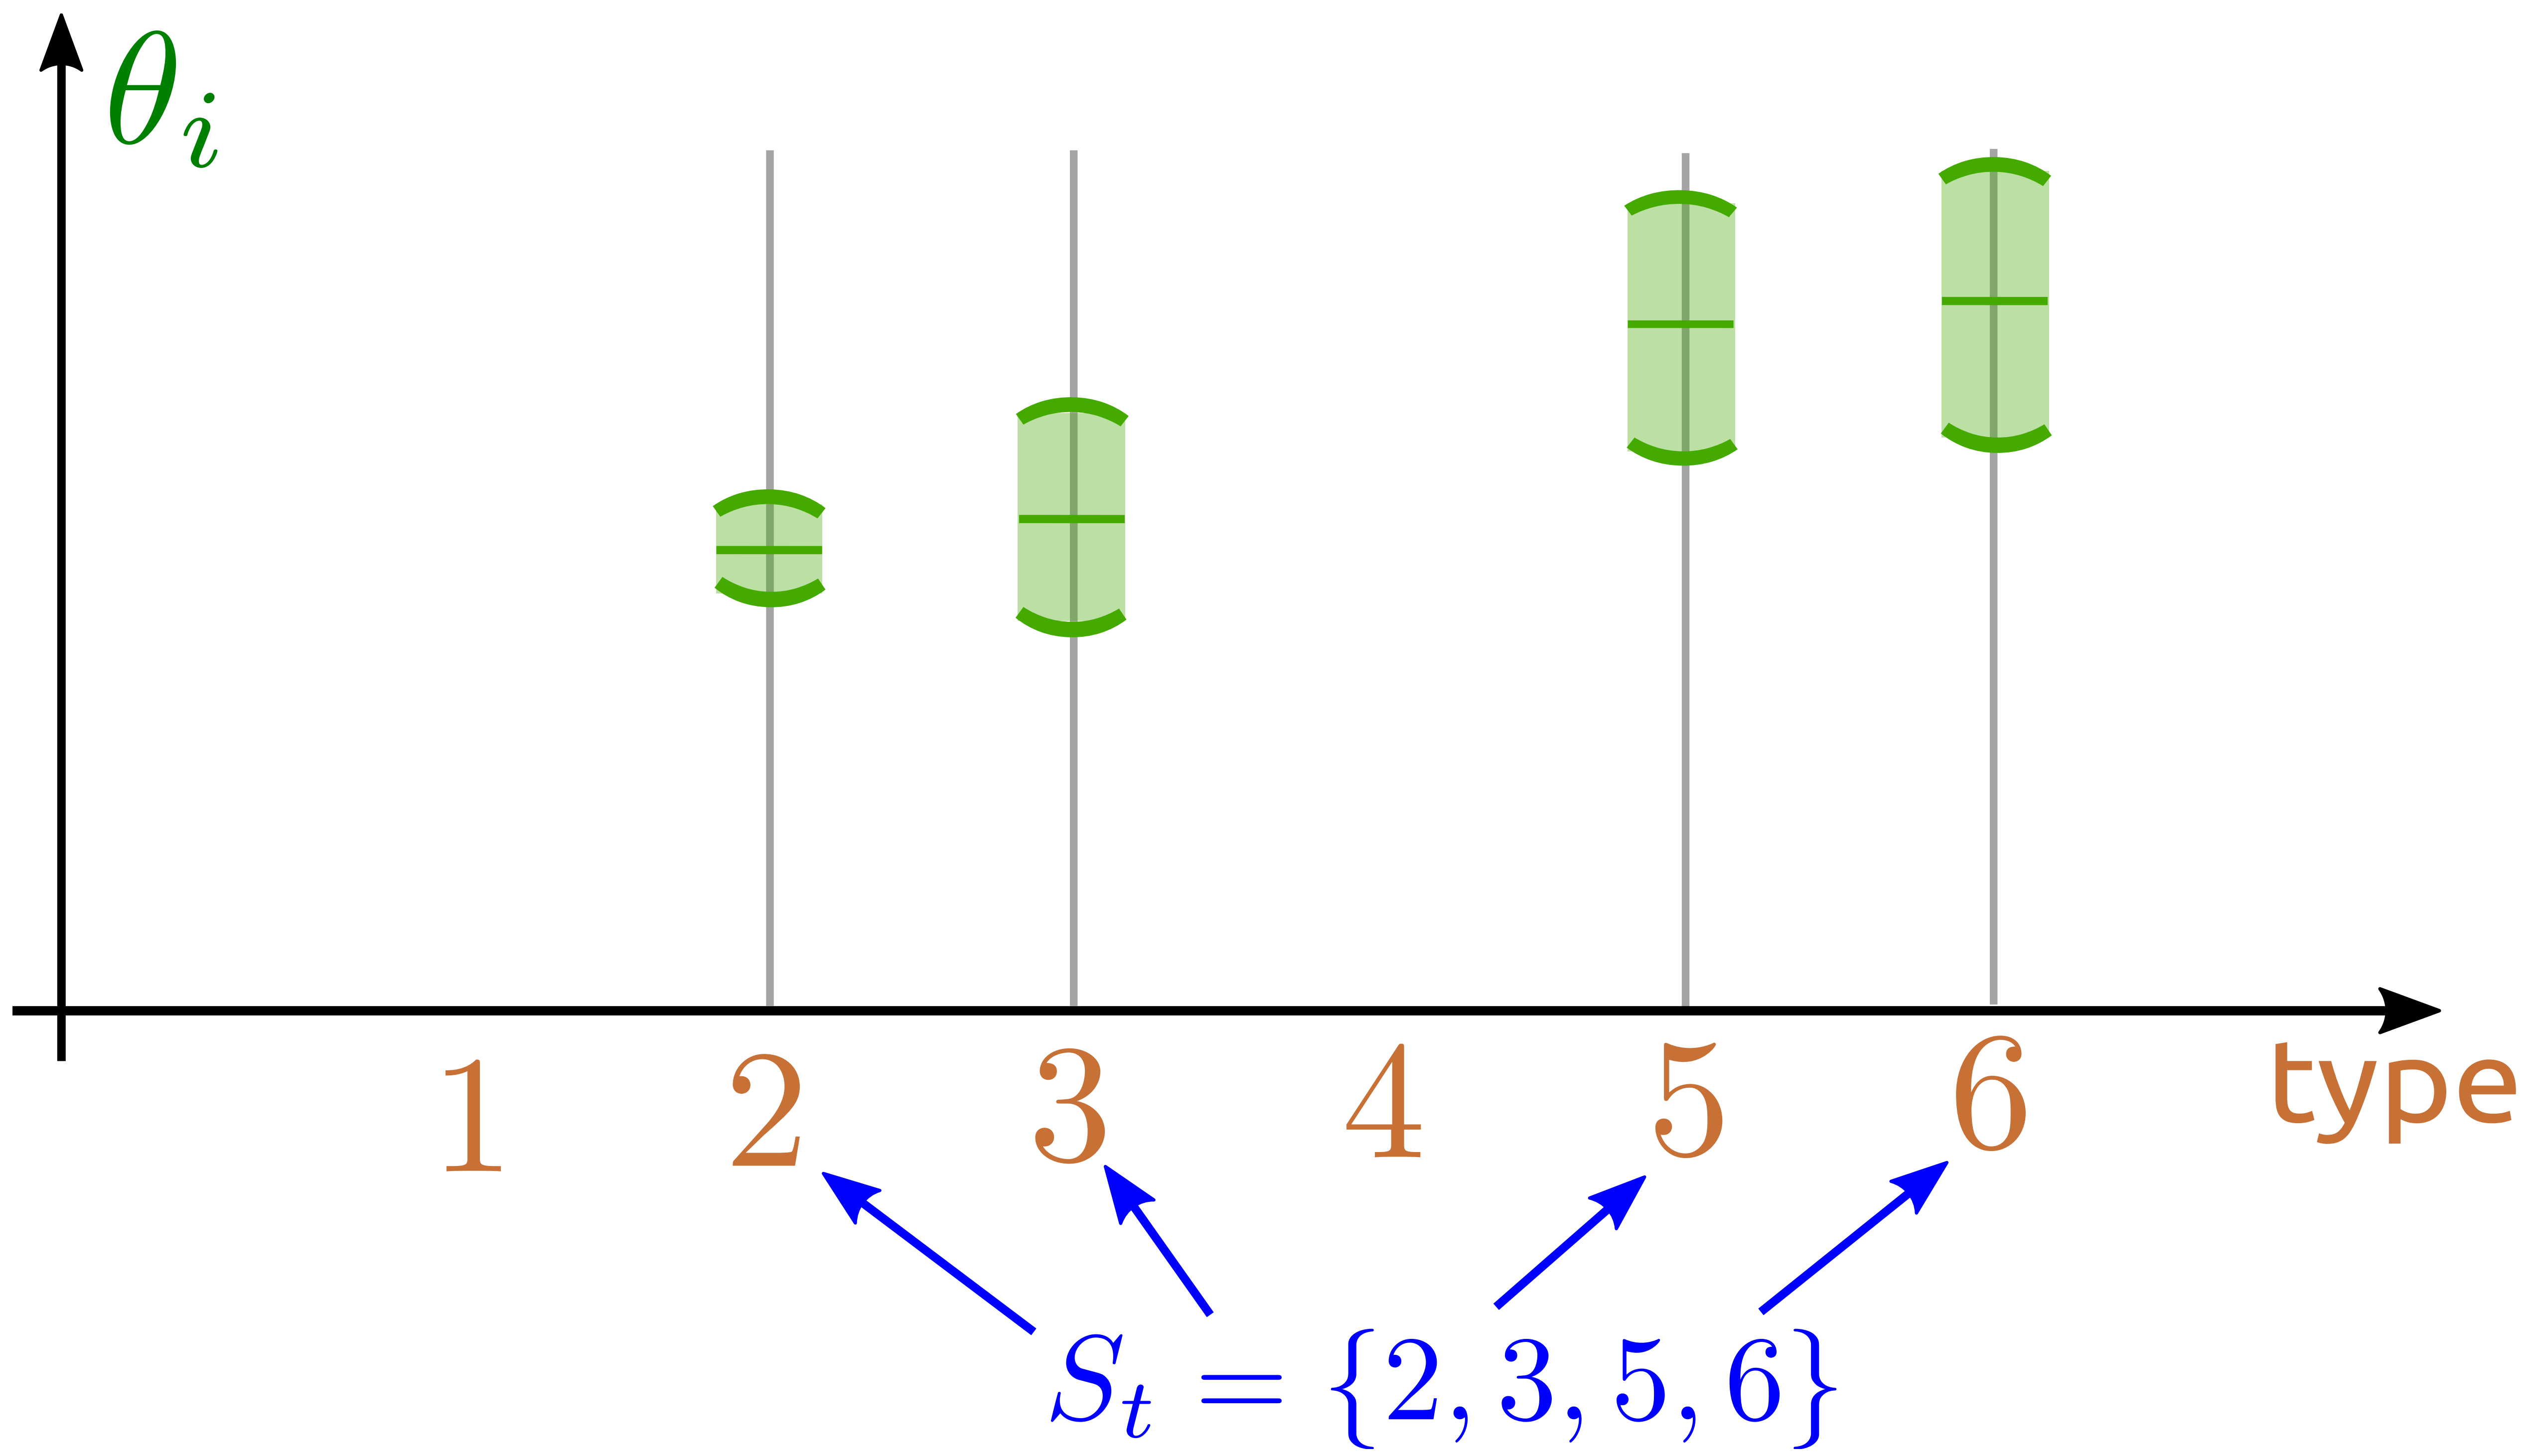

PHASE 2: CHOOSING PRICE p_t

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

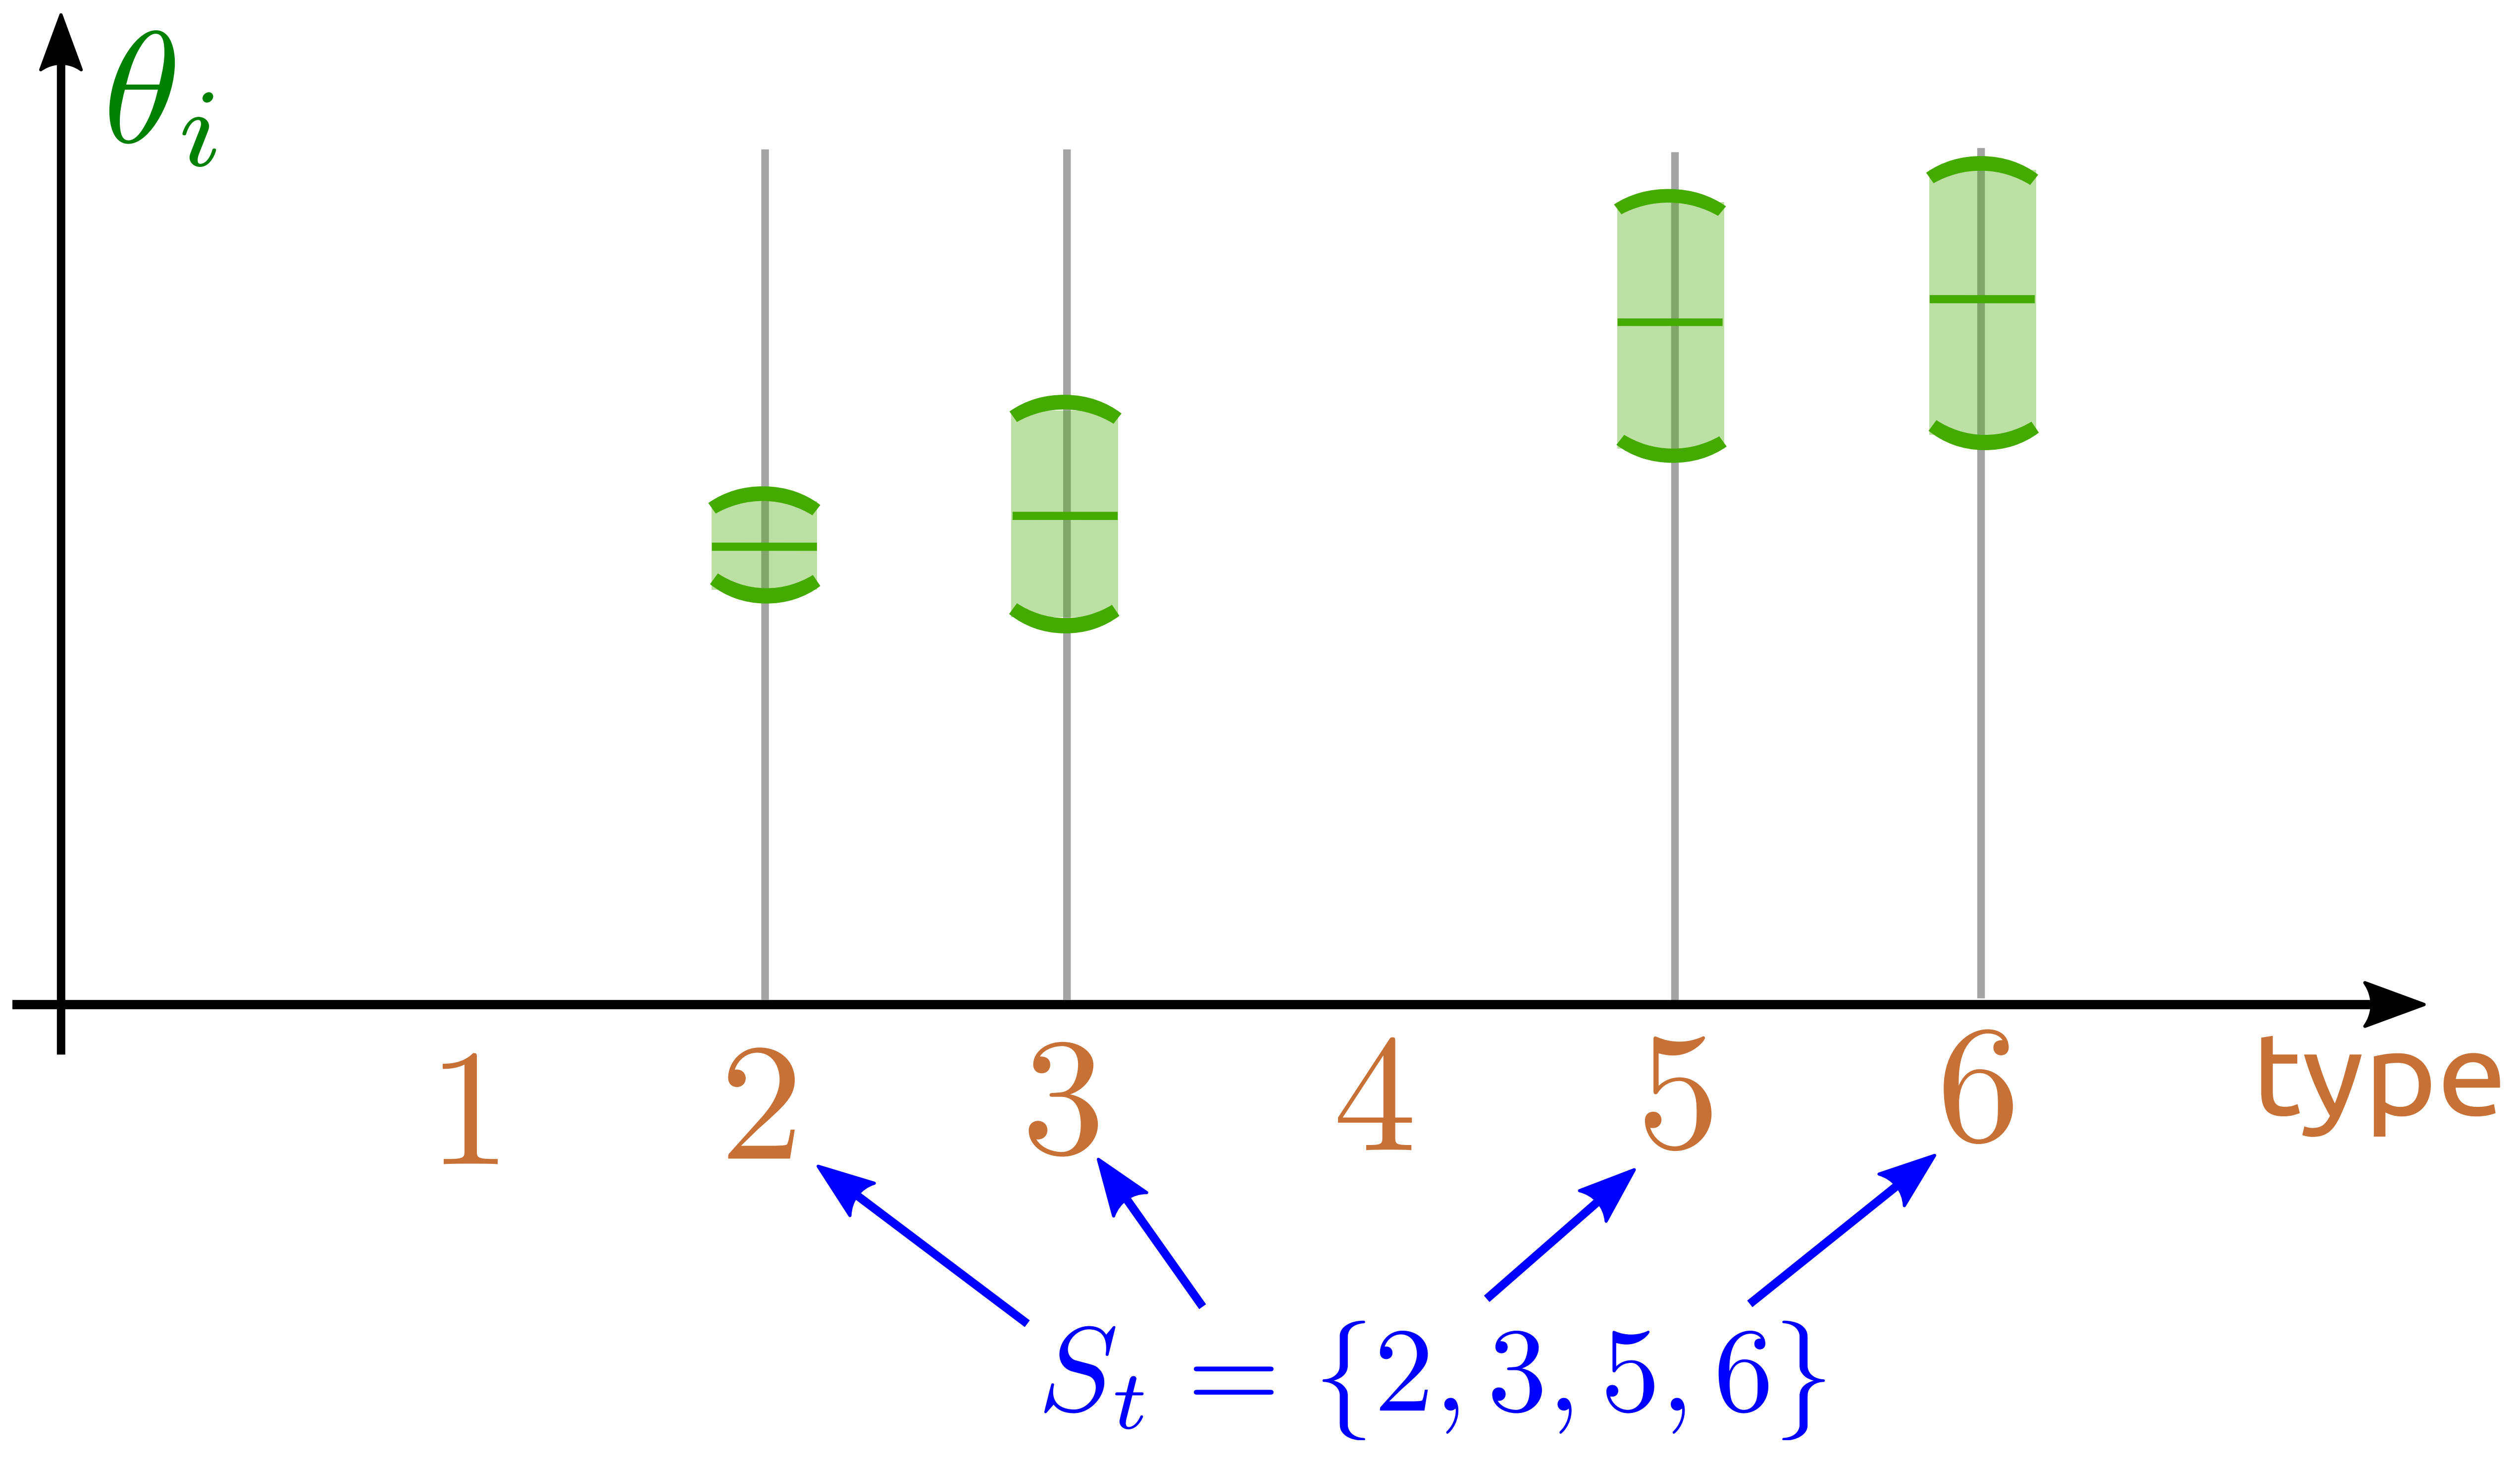

- Choose p_t to be the minimum LCB of these confidence intervals.

PHASE 2: CHOOSING PRICE p_t

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

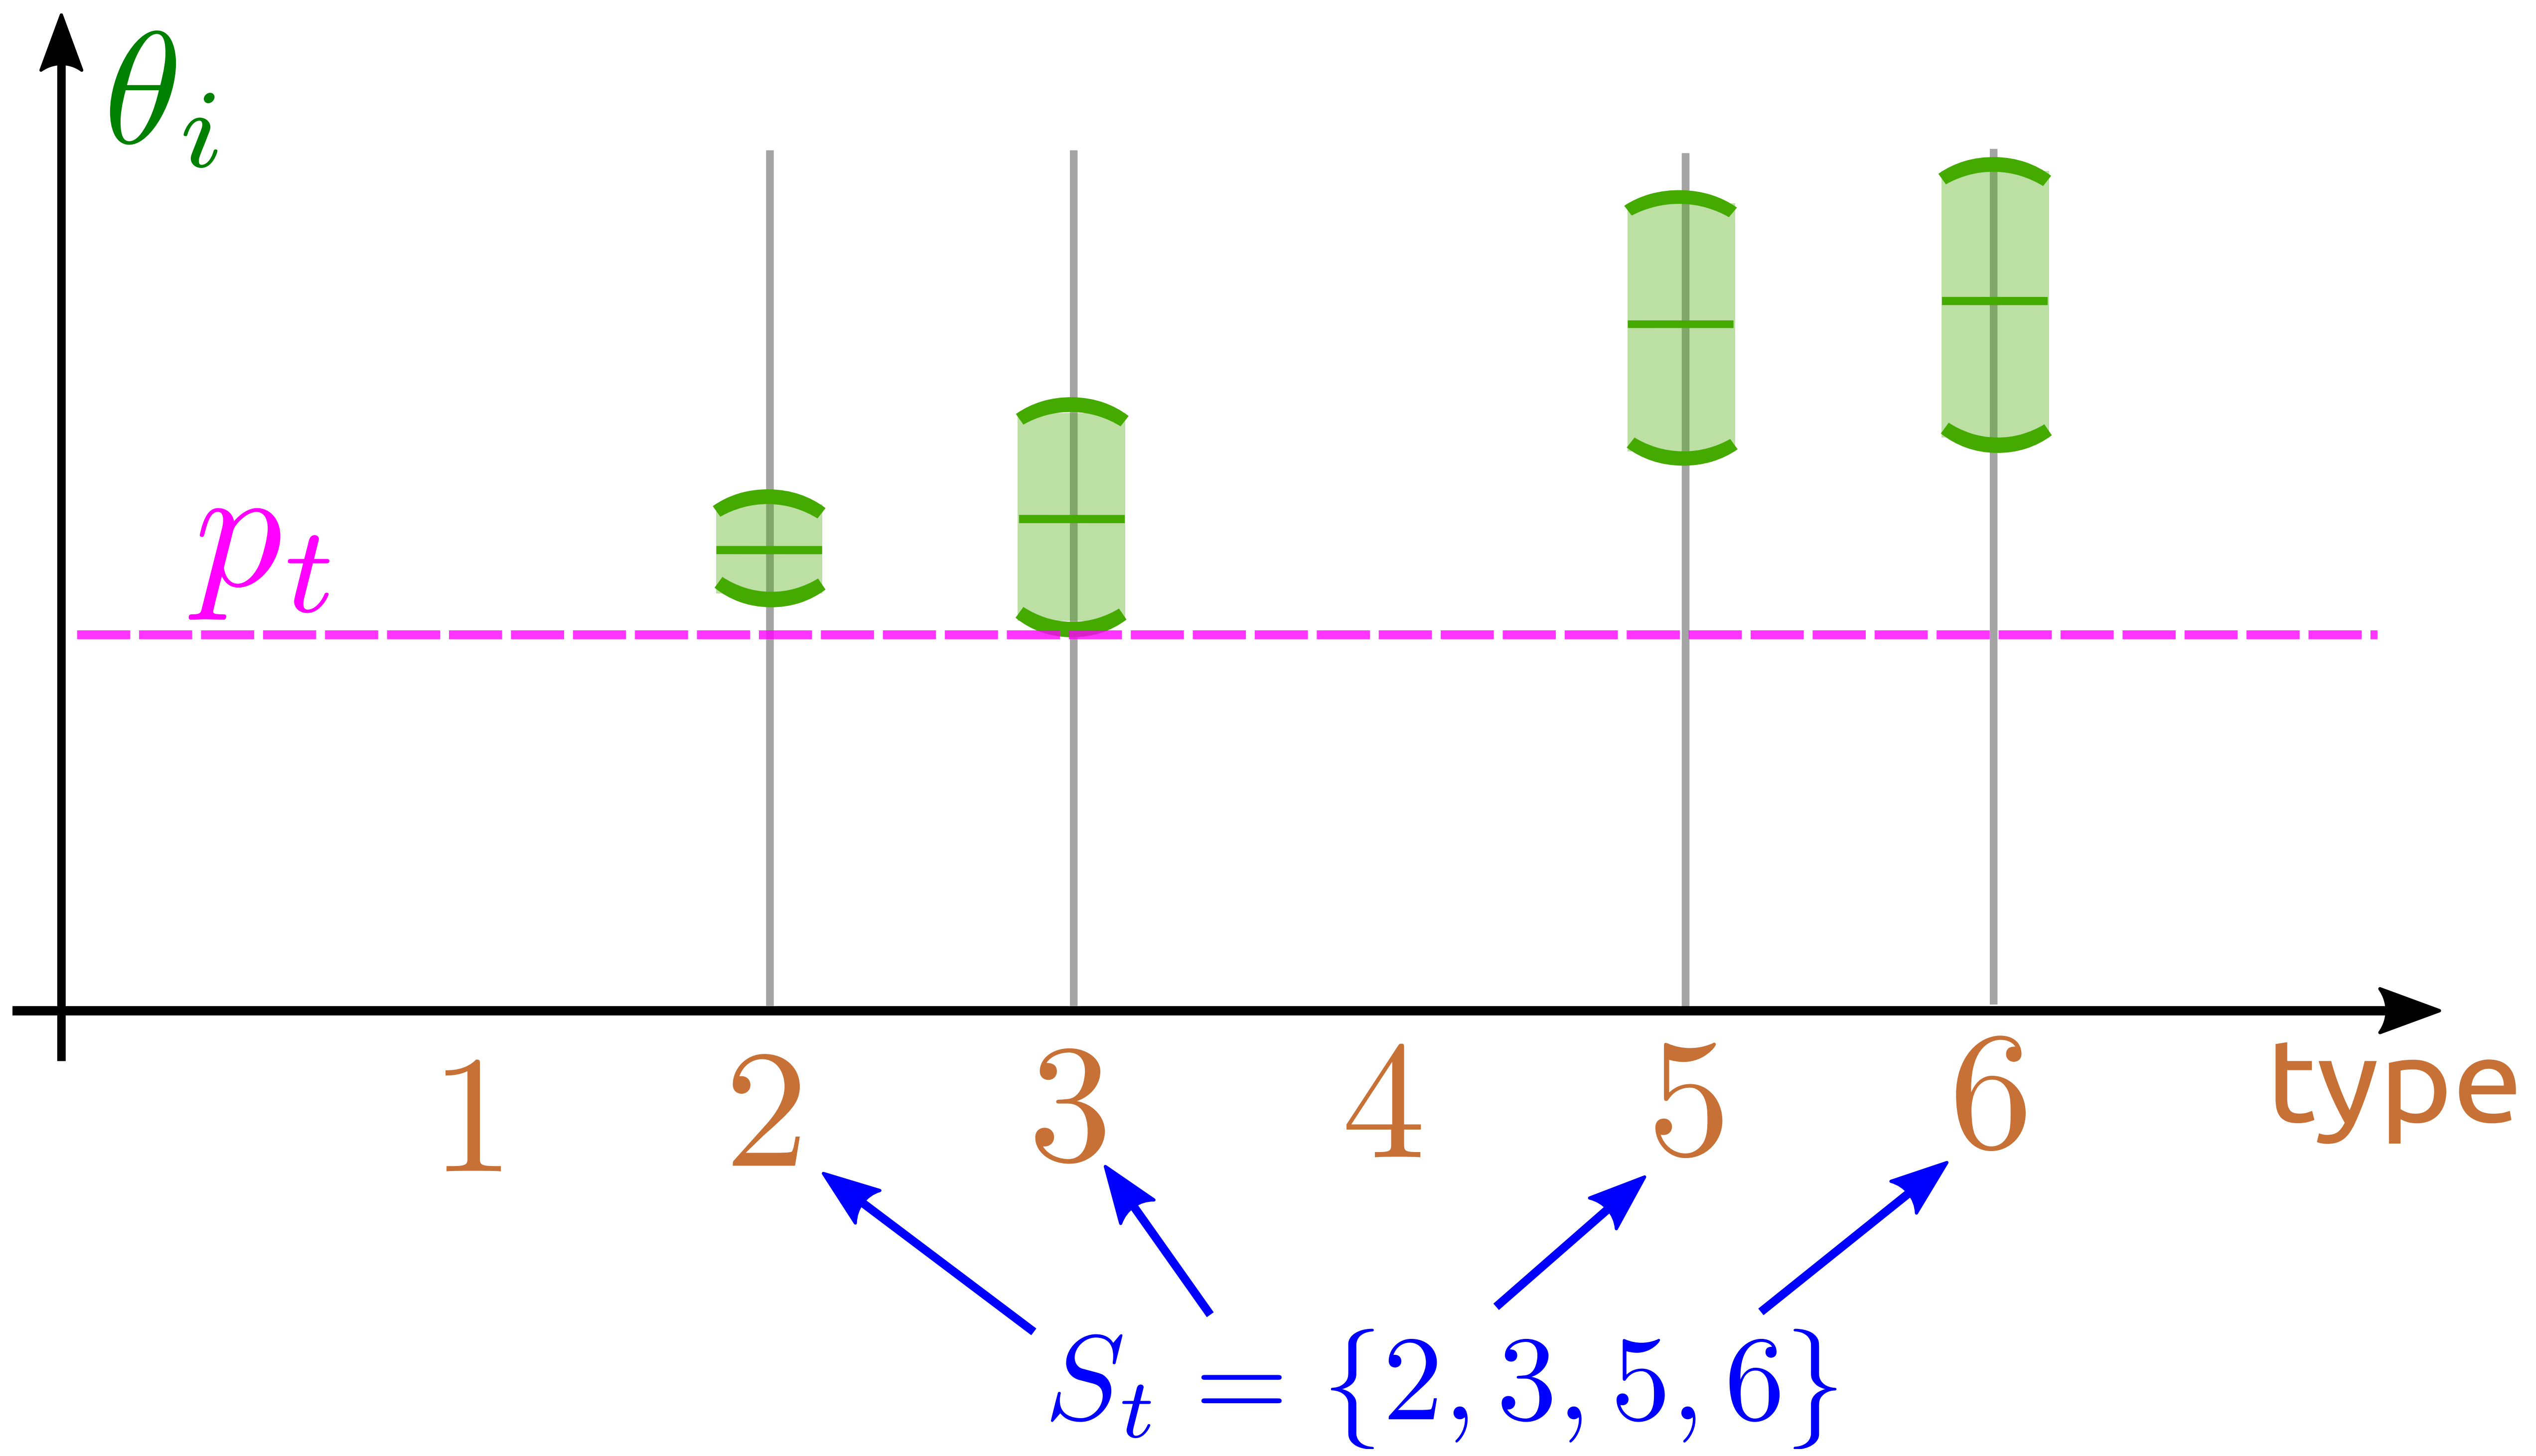

- Choose p_t to be the minimum LCB of these confidence intervals.

WHY DO WE NEED A PHASE 1?

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for a “very low price”, eliminate types that are infrequent.

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for a "very low price", eliminate types that are infrequent.

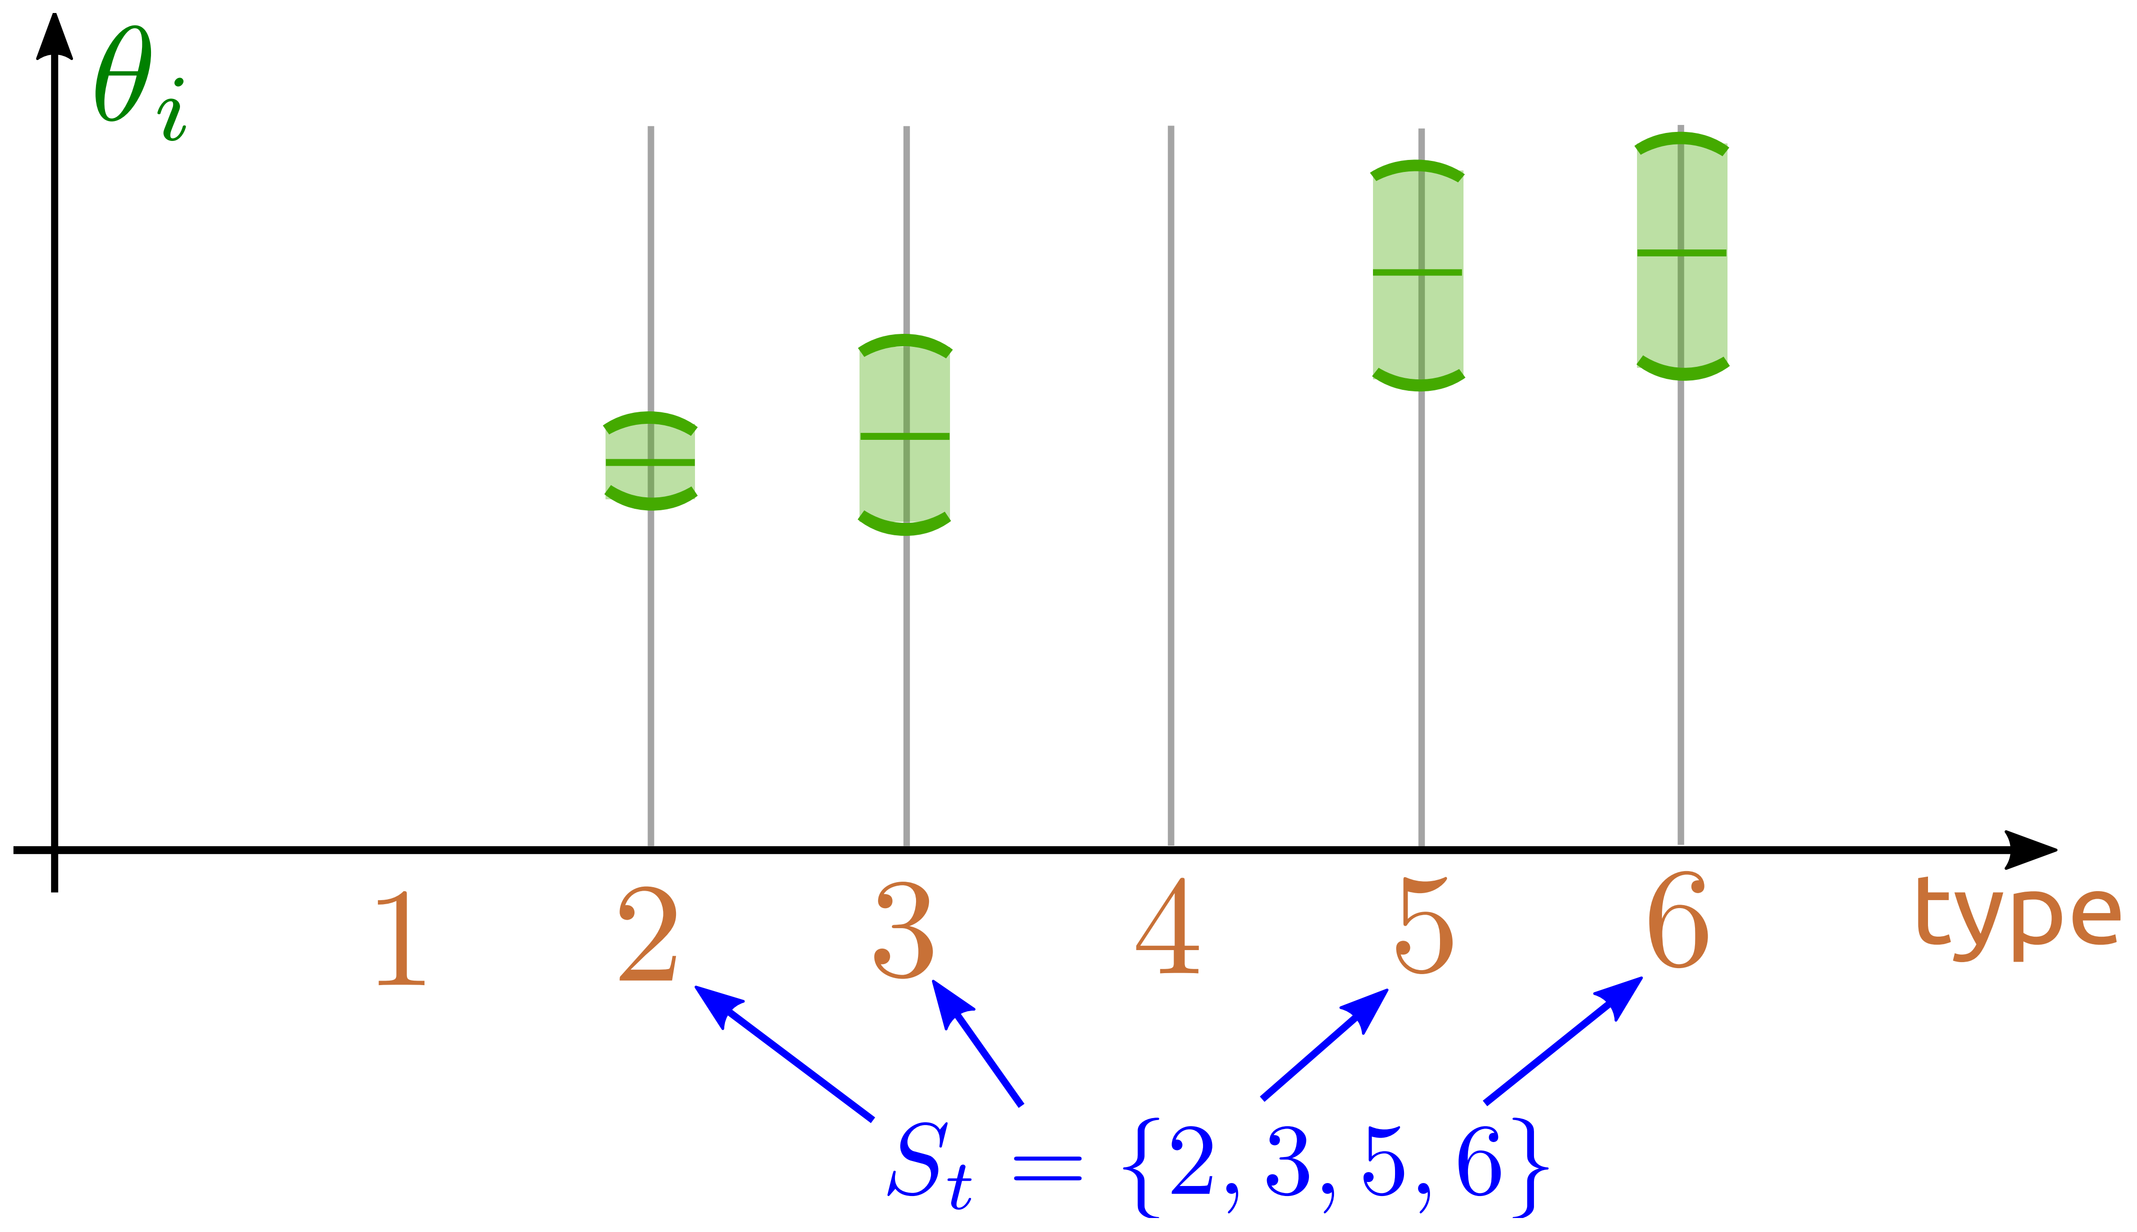

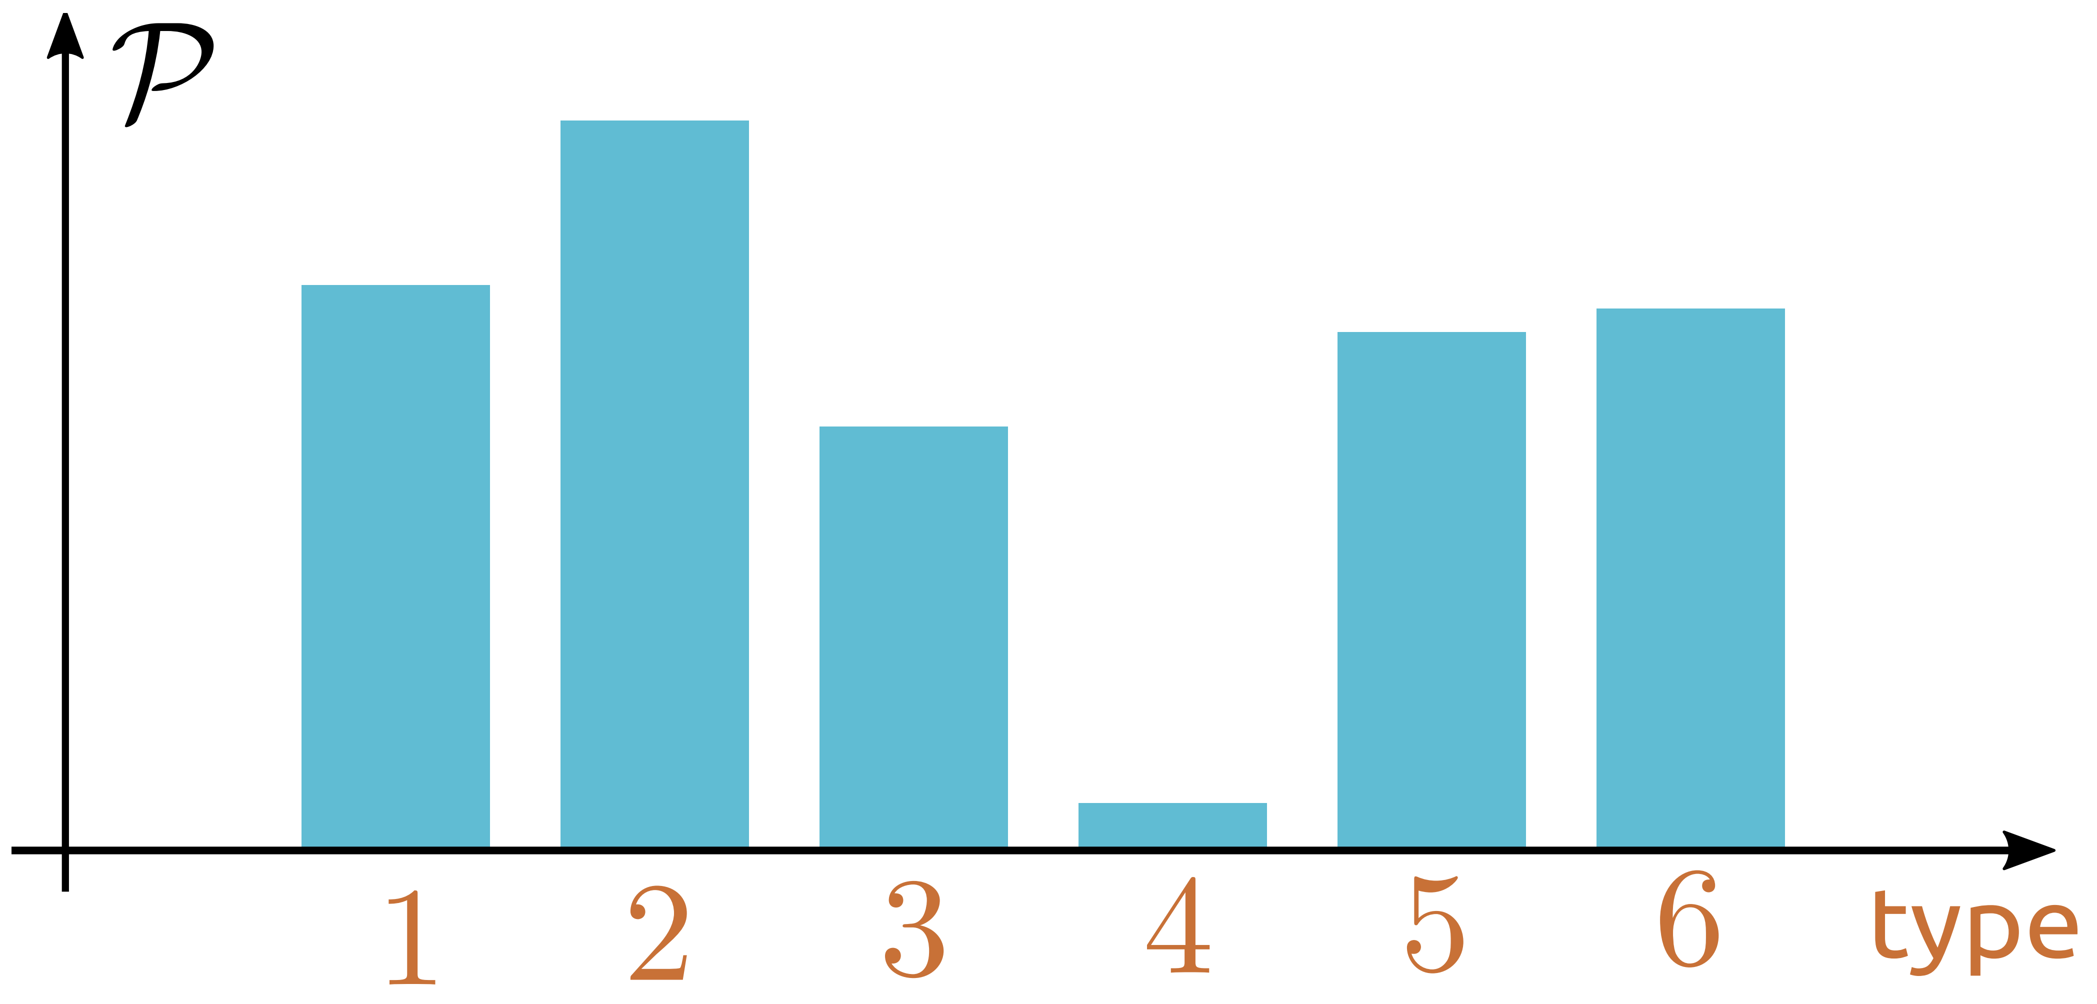

WHY DO WE NEED A PHASE 1?

- Phase 1: offer the item for a “very low price”, eliminate types that are infrequent.
  - Low probability of appearance $\Longrightarrow$ fewer reviews.

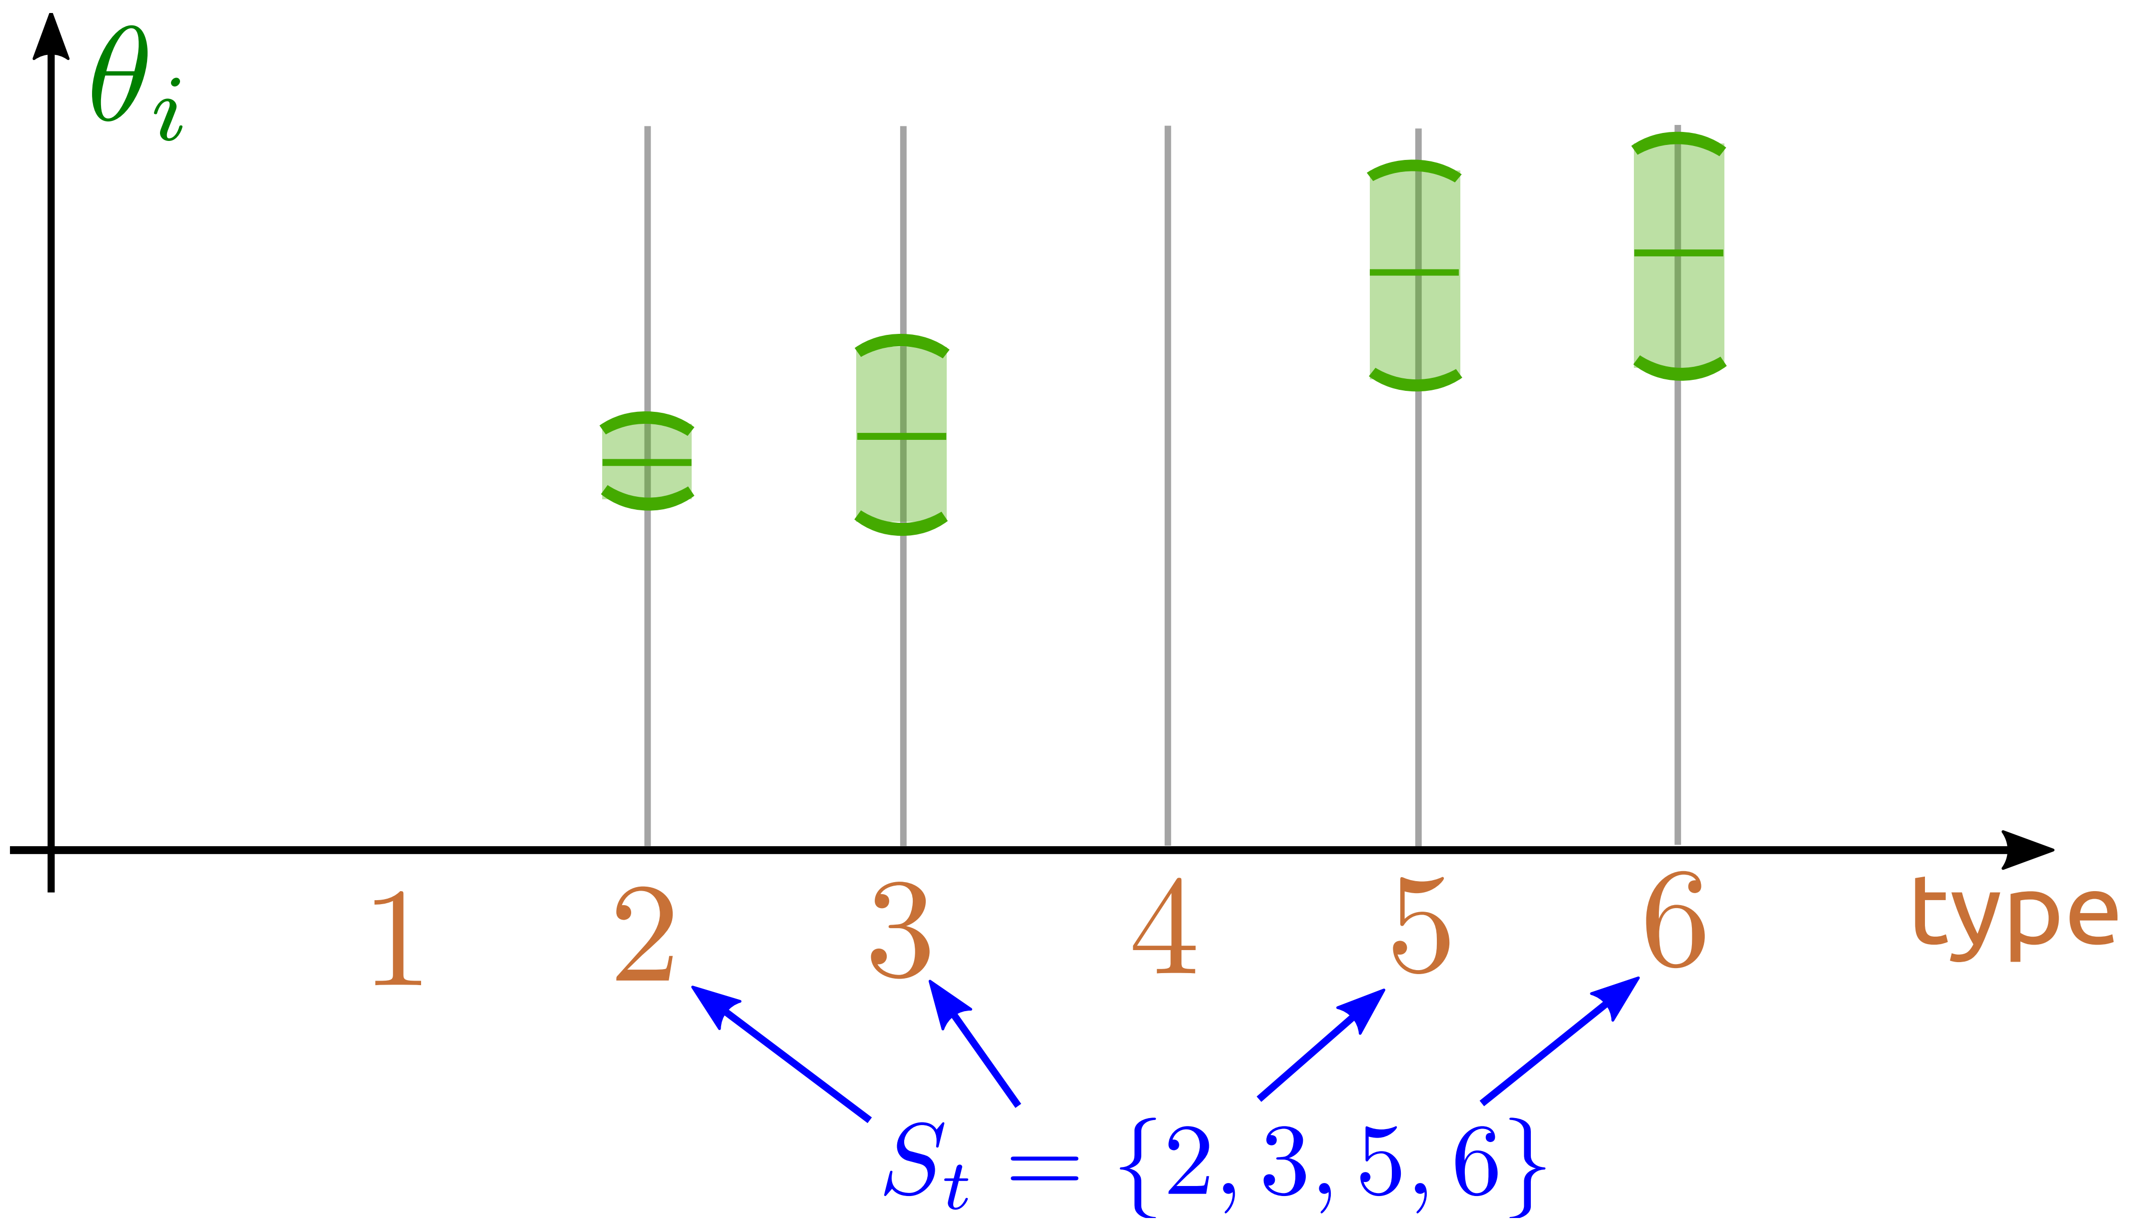

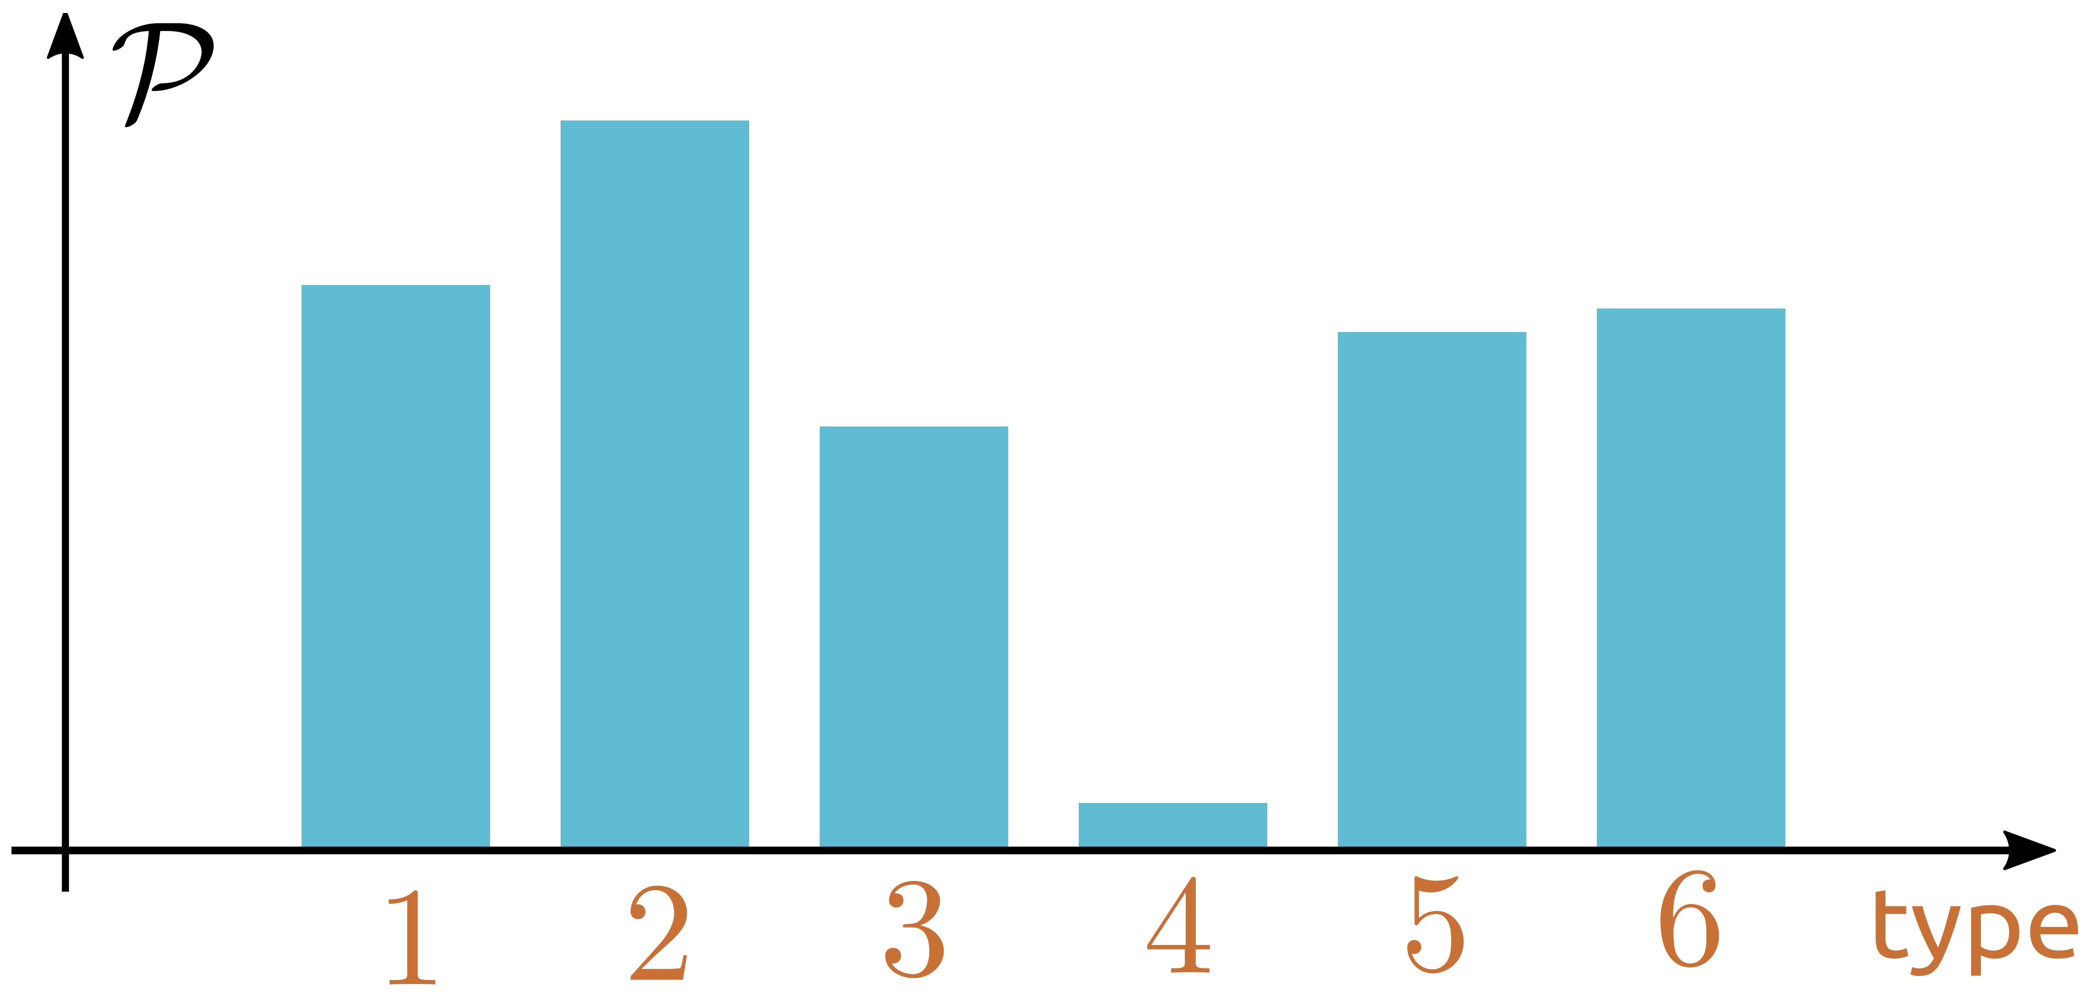

WHY DO WE NEED A PHASE 1?

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

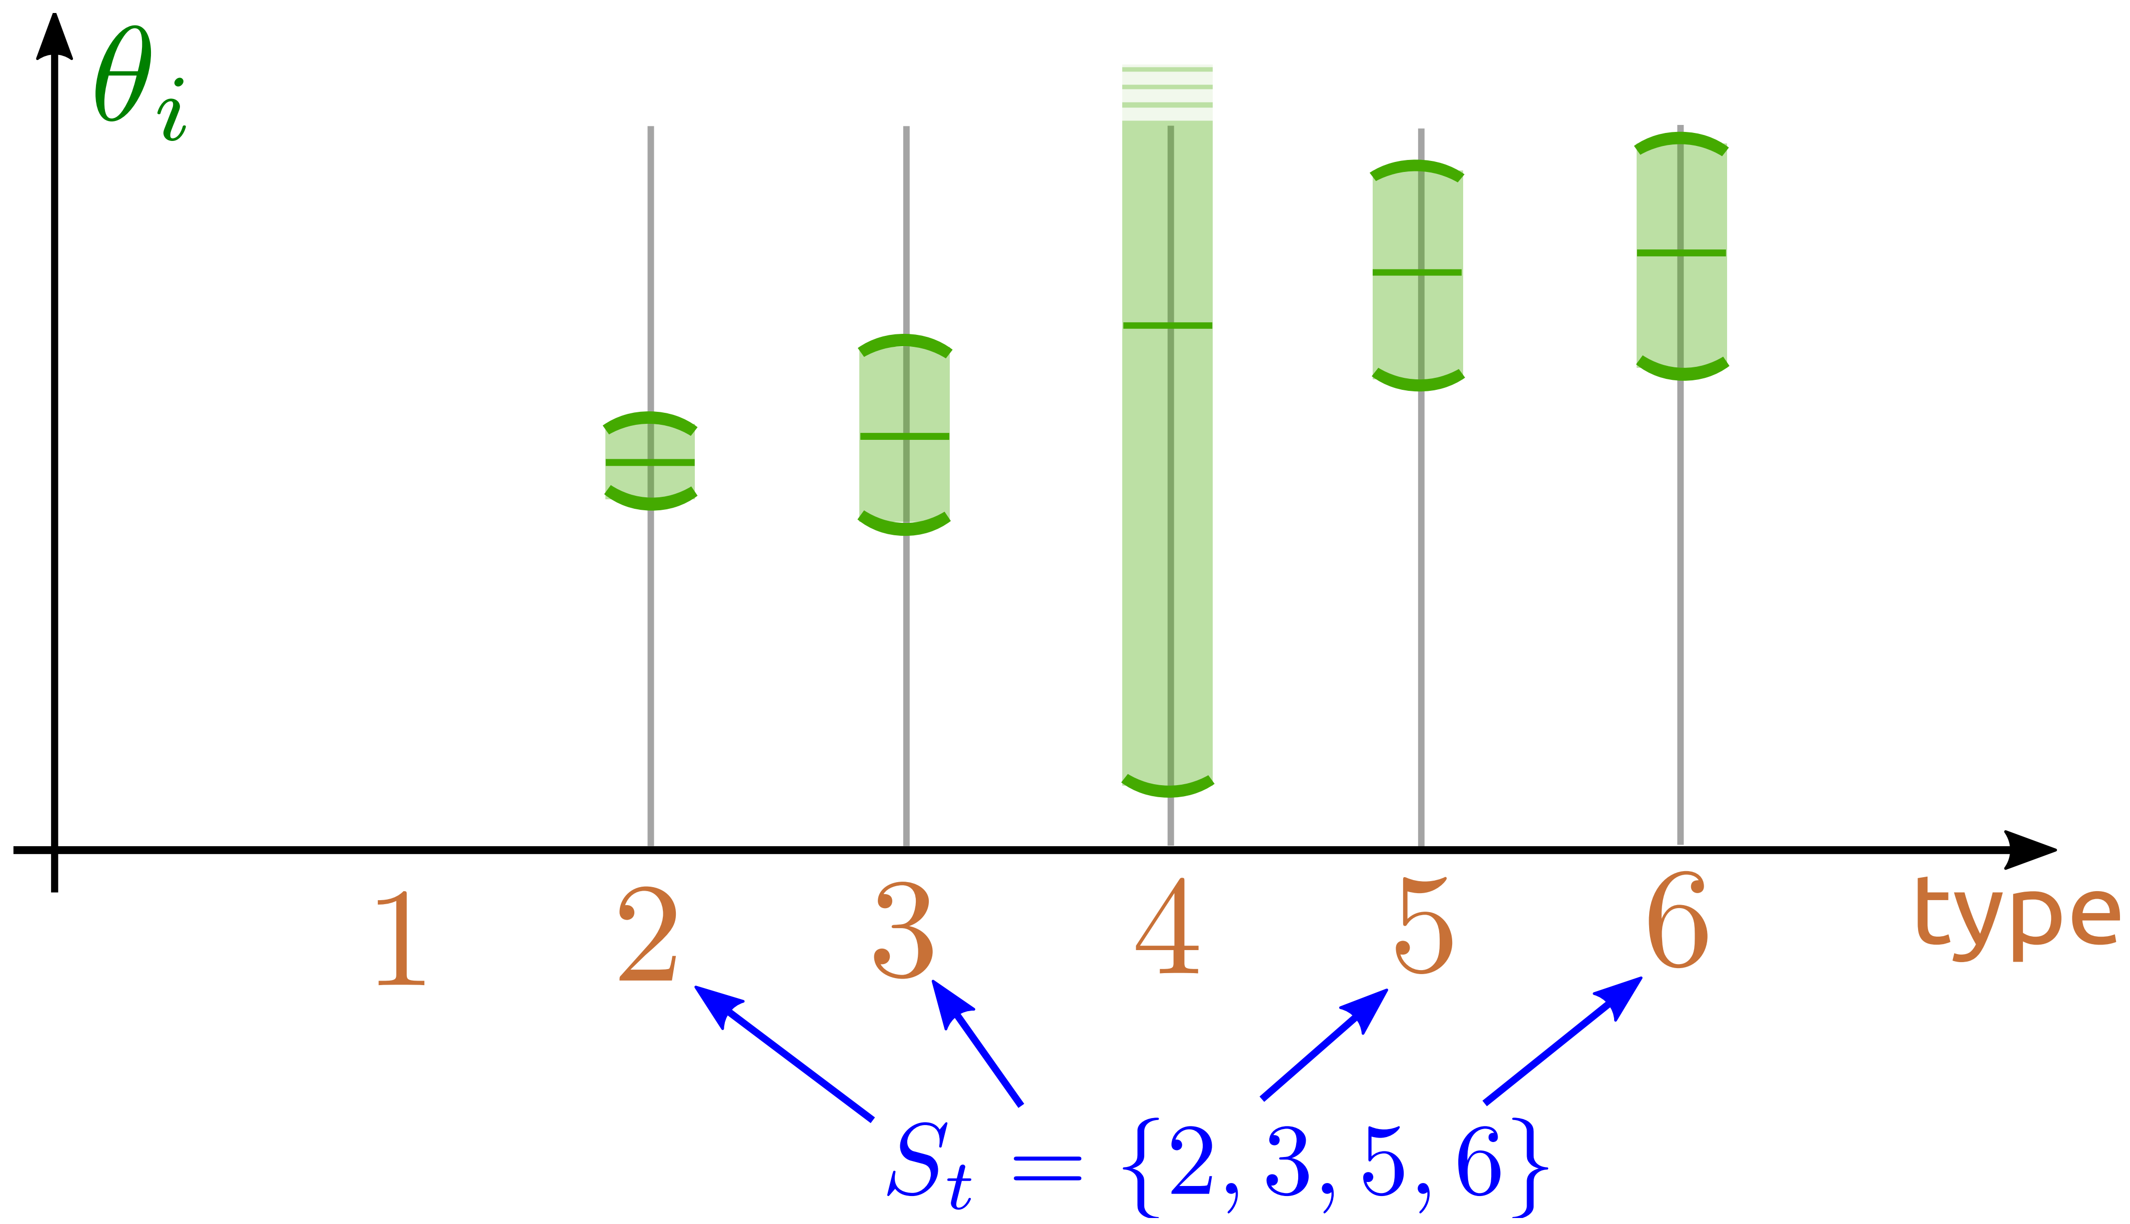

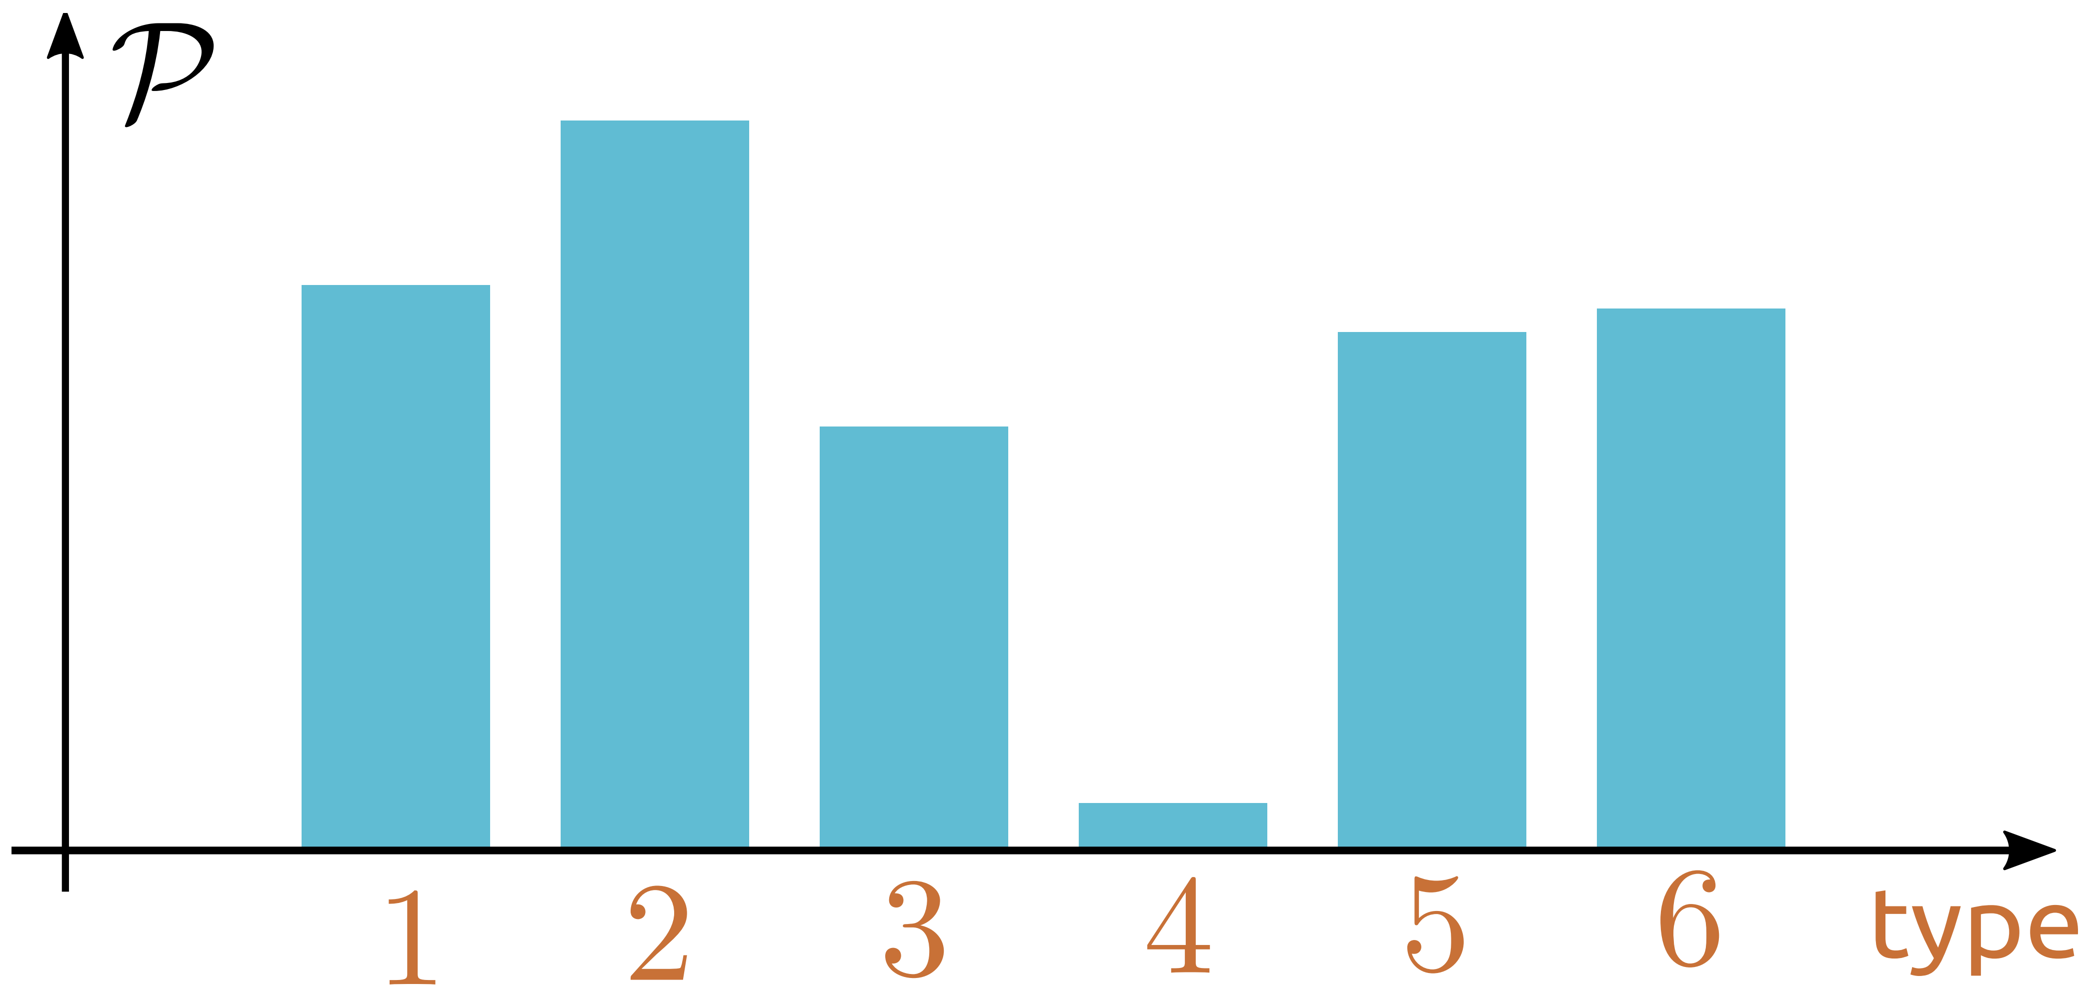

WHY DO WE NEED A PHASE 1?

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

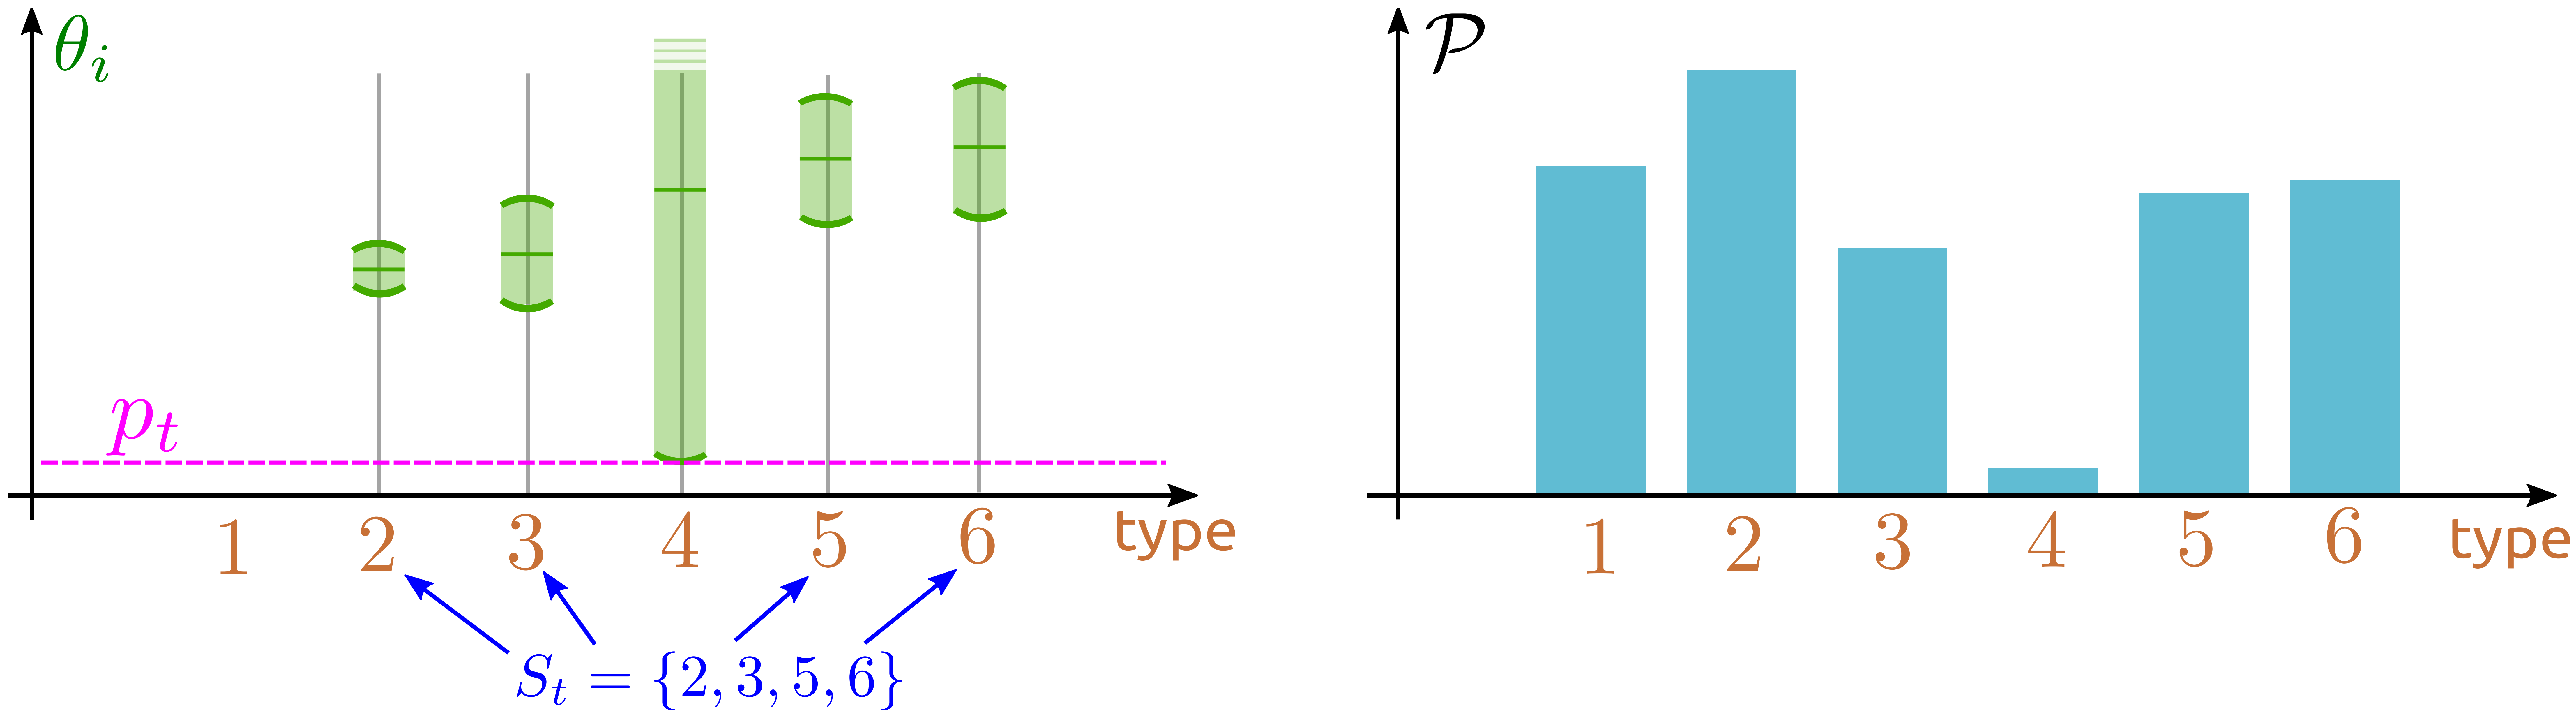

OUTLINE

1. **Problem set up**
   - Online learning framework, assumptions, challenges

2. **Algorithm**

3. **Theoretical results**
   - Upper bounds, lower bounds, proof sketches

MAIN UPPER BOUND

Theorem: In the worst case,

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

Theorem: In the worst case,

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

But if the smallest probability of appearance for any type is large, i.e. $q_{\min} \geq d^{-2/3}T^{-1/3}$

Theorem: In the worst case,

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

But if the smallest probability of appearance for any type is large, i.e. $q_{\min} \geq d^{-2/3}T^{-1/3}$

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

UPPER BOUND PROOF SKETCH

UPPER BOUND PROOF SKETCH

Worst case bound: $\mathbb{E}[R_T] \in \tilde{\mathcal{O}}\left(d^{1/3}T^{2/3} + d^{2/3}T^{1/3}\right)$

UPPER BOUND PROOF SKETCH

Worst case bound: $\mathbb{E}[R_T] \in \tilde{\mathcal{O}}\left(d^{1/3}T^{2/3} + d^{2/3}T^{1/3}\right)$

Four sources of regret:

UPPER BOUND PROOF SKETCH

Worst case bound: $\mathbb{E}[R_T] \in \tilde{\mathcal{O}}\left(d^{1/3}T^{2/3} + d^{2/3}T^{1/3}\right)$

Four sources of regret:

1. *Low price in phase 1:* $d^{2/3}T^{1/3}$ rounds.

UPPER BOUND PROOF SKETCH

Worst case bound: $\mathbb{E}[R_T] \in \tilde{\mathcal{O}}\left(d^{1/3}T^{2/3} + d^{2/3}T^{1/3}\right)$

Four sources of regret:

1. *Low price in phase 1:* $d^{2/3}T^{1/3}$ rounds.
2. *Eliminating some types after Phase 1 due to low probability of appearance:*
   we are competing with $p^\star(Q)$ instead of Q: $d^{1/3}T^{2/3}$ regret.

UPPER BOUND PROOF SKETCH

Worst case bound: $\mathbb{E}[R_T] \in \tilde{\mathcal{O}}\left(d^{1/3}T^{2/3} + d^{2/3}T^{1/3}\right)$

Four sources of regret:

1. *Low price in phase 1:* $d^{2/3}T^{1/3}$ rounds.
2. *Eliminating some types after Phase 1 due to low probability of appearance:*
   we are competing with $p^\star(Q)$ instead of Q: $d^{1/3}T^{2/3}$ regret.
3. *Error in estimating the revenue* $\mathrm{rev}(\theta_i, Q)$ *in Phase 2:* $\sqrt{T}$ regret.

UPPER BOUND PROOF SKETCH

Worst case bound: $\mathbb{E}[R_T] \in \tilde{\mathcal{O}}\left(d^{1/3}T^{2/3} + d^{2/3}T^{1/3}\right)$

Four sources of regret:

1. *Low price in phase 1:* $d^{2/3}T^{1/3}$ rounds.
2. *Eliminating some types after Phase 1 due to low probability of appearance:*
   we are competing with $p^\star(Q)$ instead of Q: $d^{1/3}T^{2/3}$ regret.
3. *Error in estimating the revenue* $\mathrm{rev}(\theta_i, Q)$ *in Phase 2:* $\sqrt{T}$ regret.
4. *Agents learning their values:* $d^{1/3}T^{2/3}$ regret.

MAIN LOWER BOUND

Theorem: In the worst case,

$$\inf_{\text{algorithms}} \sup_{\text{problems}} \mathbb{E}[R_T] \in \tilde{\Omega}\left(d^{1/3}T^{2/3}\right)$$

MAIN LOWER BOUND

Theorem: In the worst case,

$$\inf_{\text{algorithms}} \sup_{\text{problems}} \mathbb{E}[R_T] \in \tilde{\Omega}\left(d^{1/3}T^{2/3}\right)$$

But if the smallest probability of appearance for any type is large, i.e. $q_{\min} \geq d^{-2/3}T^{-1/3}$

MAIN LOWER BOUND

Theorem: In the worst case,

$$\inf_{\text{algorithms}} \sup_{\text{problems}} \mathbb{E}[R_T] \in \tilde{\Omega}\left(d^{1/3}T^{2/3}\right)$$

But if the smallest probability of appearance for any type is large, i.e. $q_{\min} \geq d^{-2/3}T^{-1/3}$

$$\inf_{\text{algorithms}} \sup_{q_{\min} \geq d^{-2/3}T^{-1/3}} \mathbb{E}[R_T] \in \tilde{\Omega}\left(\sqrt{\frac{T}{q_{\min}}}\right)$$

LOWER BOUND PROOF SKETCH

LOWER BOUND PROOF SKETCH

- Consider difficult problem instances where

  1. different types have similar ex-ante values,

  2. but, large variance in type appearance probabilities $\mathscr{P}$.

LOWER BOUND PROOF SKETCH

- Consider difficult problem instances where
  1. different types have similar ex-ante values,
  2. but, large variance in type appearance probabilities $\mathscr{P}$.
- Algorithm must decide if
  - it will target low probability types (low price due to high uncertainty)
  - or, ignore low probability types (foregoing potential future revenue)

LOWER BOUND PROOF SKETCH

- Consider difficult problem instances where
  1. different types have similar ex-ante values,
  2. but, large variance in type appearance probabilities $\mathscr{P}$.

- Algorithm must decide if
  - it will target low probability types (low price due to high uncertainty)
  - or, ignore low probability types (foregoing potential future revenue)

- Either way, seller suffers high regret.

LOWER BOUND PROOF SKETCH

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.
    - who contribute too little revenue.

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.
    - who contribute too little revenue.
  - On each round, set price low enough to target all other customer types (while accounting for buyer uncertainty).

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.
    - who contribute too little revenue.
  - On each round, set price low enough to target all other customer types (while accounting for buyer uncertainty).

- No algorithm can do significantly better than the best algorithm in this class.

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.
    - who contribute too little revenue.
  - On each round, set price low enough to target all other customer types (while accounting for buyer uncertainty).

- No algorithm can do significantly better than the best algorithm in this class.
- **Proof supports the structure of our algorithm** (even without knowledge of $\mathscr{P}$)**!**

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.
    - who contribute too little revenue.
  - On each round, set price low enough to target all other customer types (while accounting for buyer uncertainty).

Use a few rounds (Phase 1) to eliminate low prob types

- No algorithm can do significantly better than the best algorithm in this class.
- **Proof supports the structure of our algorithm** (even without knowledge of $\mathscr{P}$)**!**

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.
    - who contribute too little revenue.
  - On each round, set price low enough to target all other customer types (while accounting for buyer uncertainty).

Use a few rounds (Phase 1) to eliminate low prob types

Phase 2, eliminating types from S_t

- No algorithm can do significantly better than the best algorithm in this class.
- **Proof supports the structure of our algorithm** (even without knowledge of $\mathscr{P}$)**!**

LOWER BOUND PROOF SKETCH

- A class of algorithms when $\mathscr{P}$ is known to the seller:
  - Ahead of time (before round 1), ignore types
    - whose probability of appearance is smaller than a chosen threshold.

      Use a few rounds (Phase 1) to eliminate low prob types
    - who contribute too little revenue.

      Phase 2, eliminating types from S_t
  - On each round, set price low enough to target all other customer types (while accounting for buyer uncertainty).

    Phase 2, pricing strategy

- No algorithm can do significantly better than the best algorithm in this class.
- **Proof supports the structure of our algorithm** (even without knowledge of $\mathscr{P}$)**!**

SUMMARY

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

- ***Algorithmic insight:*** 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.

Wenshuo Guo
UC Berkeley

Nika Haghtalab
UC Berkeley

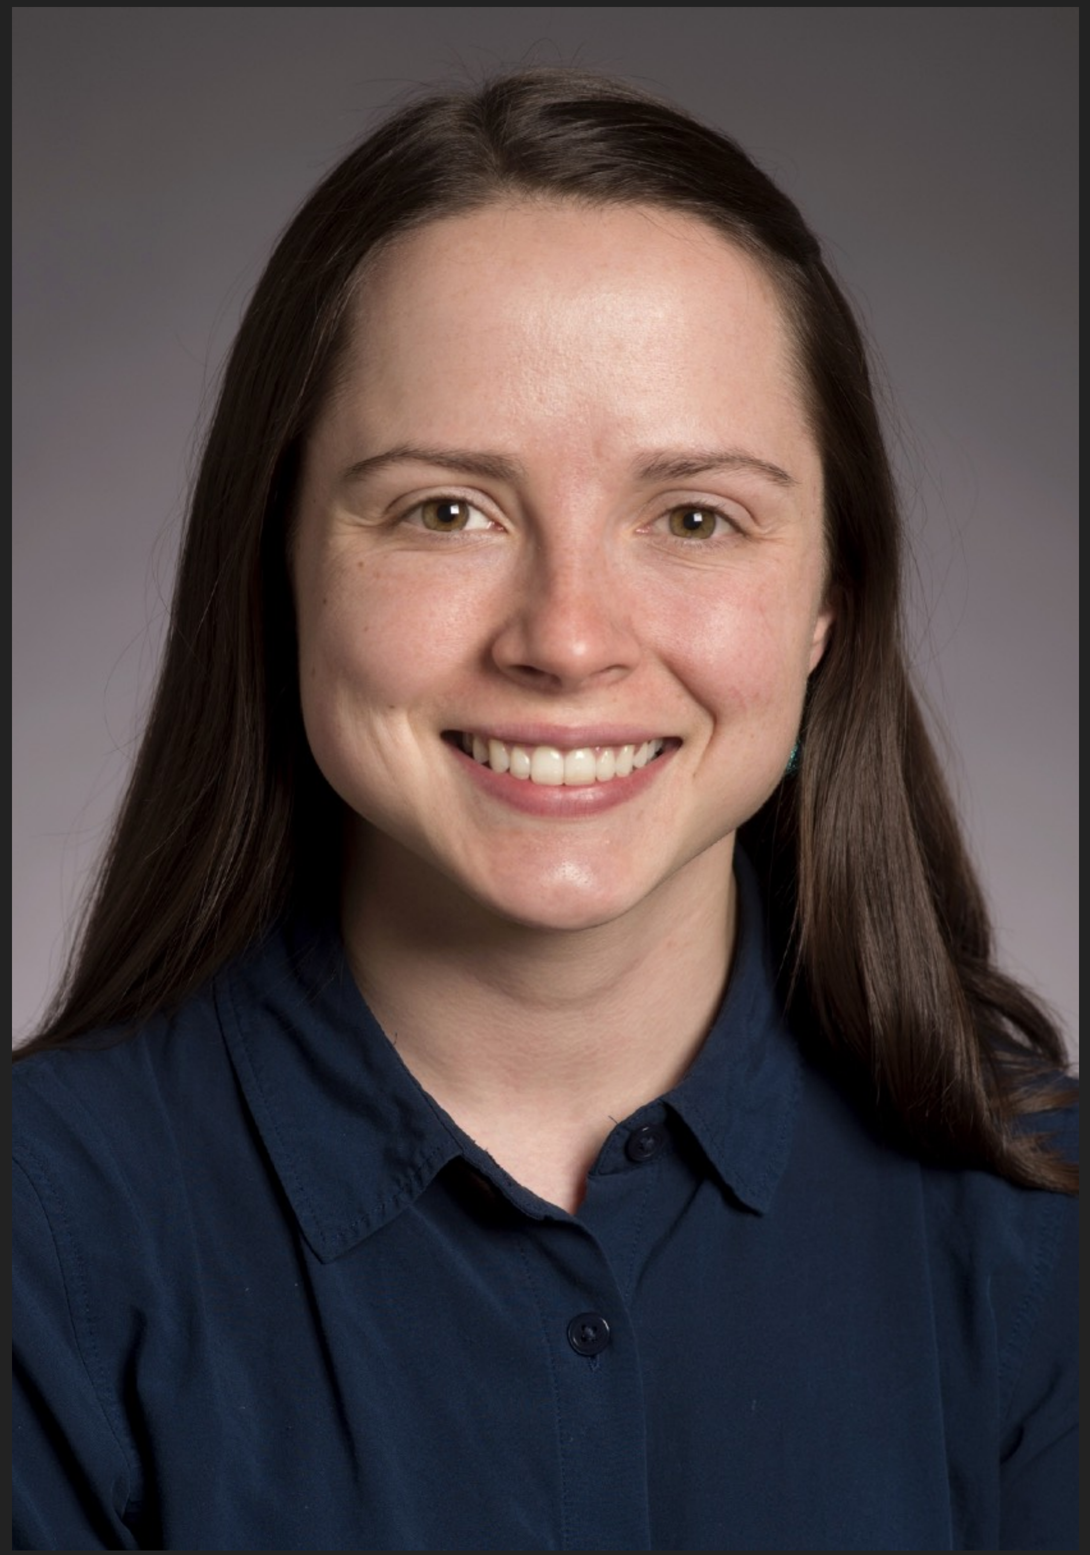

Ellen Vitercik
Stanford

THANK YOU!