MOTIVATION

What market outcomes result from **competing firms** deploying **gradient-based learning** algorithms for pricing?

Can firms treat their market environment as a **black-box**?

- In theory, yes!
  - Under mild regularity conditions, gradient-based schemes achieve sublinear regret, **even in strategic environments**

Do there exist **meaningful models of customer behavior** for which theoretical guarantees cease to hold?

- Specifically interested in **service** platforms offering **identical goods**, e.g., Lyft vs. Uber

CONTRIBUTIONS

Satisficing Search over Ranking: a **bounded rationality** model explaining **profitable co-existence** of competing firms

Characterization of a **highly undesirable local Nash equilibrium** to which gradient-based learning dynamics converge

A mechanism which escapes the local equilibrium and **converges to the pure-strategy Nash equilibrium**

MODEL

- **Perfectly substitutable** items produced by **two firms**
- Non-atomic mass of **unit-demand** customers
  - Valuation $V \sim F(\cdot)$ for the item; tail $\bar{F}$
  - F has a **monotone hazard rate** $\Longrightarrow$ unique optimal **monopolist** price p^M
- **Two types** of customers, based on **exogenous preference** over firms

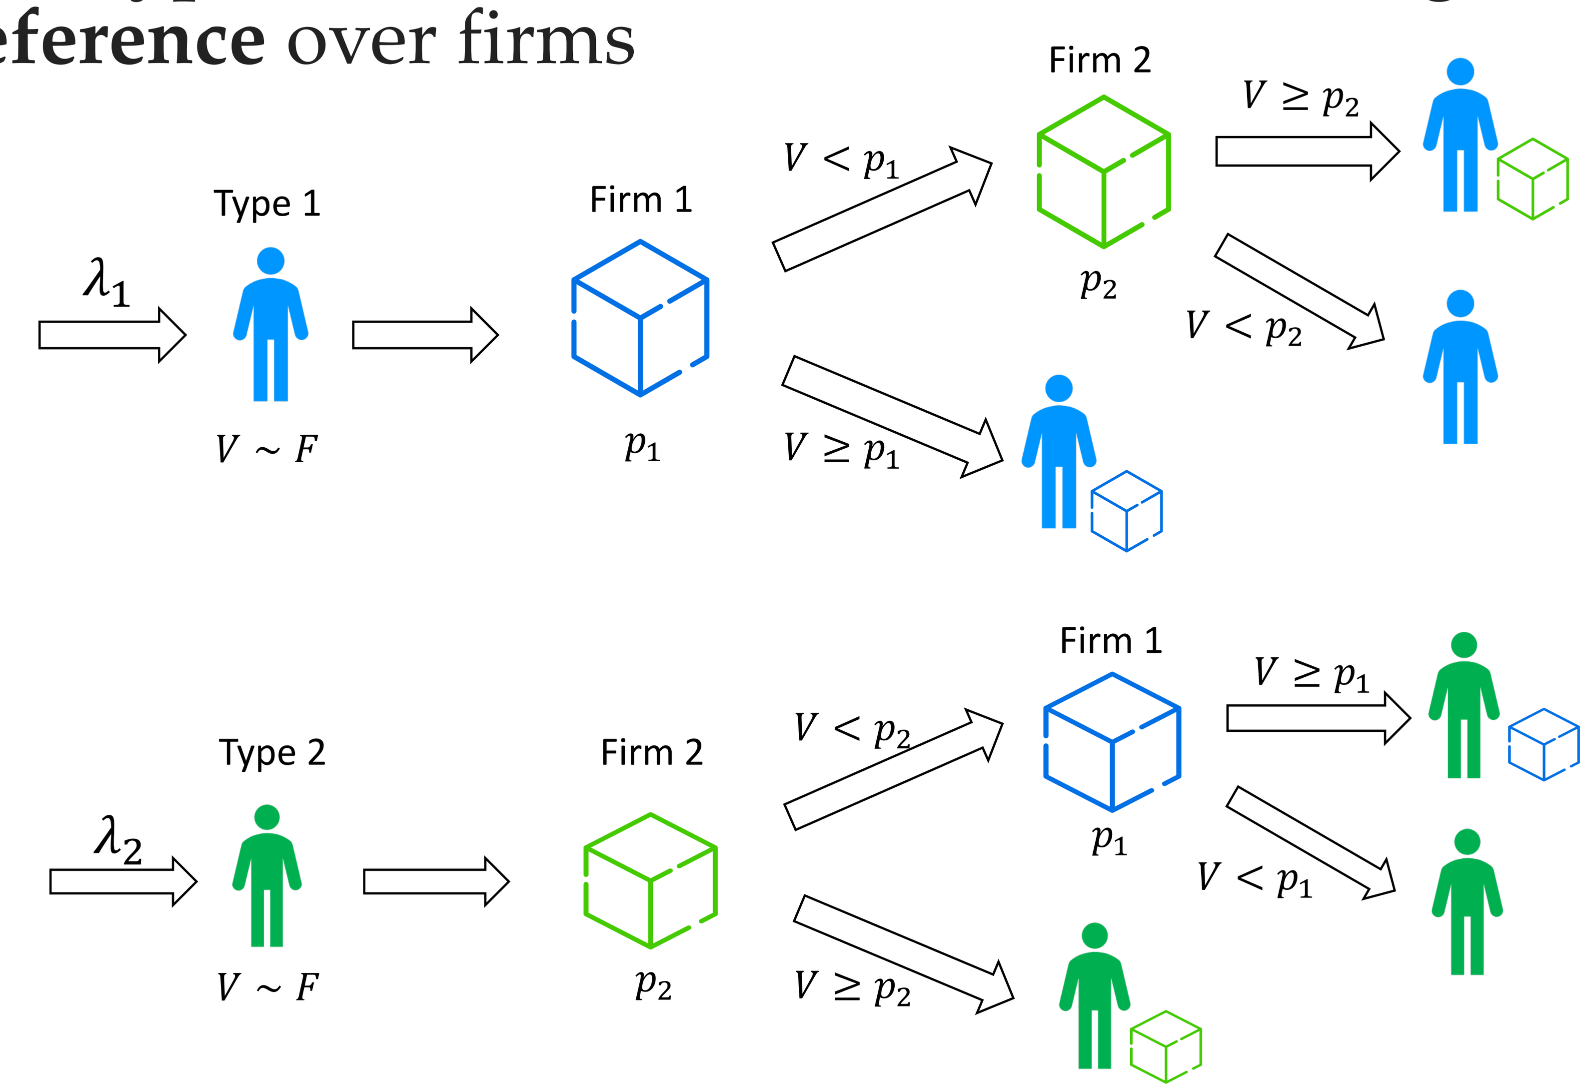

Monopolist over own loyal customer base vs. **Competitor** over entire market ?

- Equilibrium concept: **local Nash equilibrium (LNE)**

(p_1, p_2) form a **LNE** if $\exists$ open neighborhoods $\mathcal{N}_1, \mathcal{N}_2$ such that
1. p_1 is a best response (BR) to p_2 in $\bar{\mathcal{N}}_1$.
2. p_2 is a BR to p_1 in $\bar{\mathcal{N}}_2$.

- Pure-strategy NE = **global** Nash equilibrium **(GNE)**

MAIN RESULTS

Theorem 1. In a duopoly setting with with loyal customer bases of size $\lambda_1 \geq \lambda_2$, there are **exactly two LNE**:
1. **Major** LNE: $p_1 = p^M, p_2 = BR_2(p^M) < p^M$
2. **Minor** LNE: $p_2 = p^M, p_1 = BR_1(p^M) < p^M$

$\Longrightarrow$ in all equilibria, **both firms are profitable**

Theorem 2 (Existence of pure-strategy NE). The major LNE is a GNE for all $\lambda_1 \geq \lambda_2$.

Proposition. For any online algorithms used by the firms, if the resulting sequence of prices converges to the minor LNE, then Firm 2 incurs **linear** regret.

NUMERICAL RESULTS

- Learning dynamics under demand uncertainty
- Feedback assumptions:
  - First-order/bandit feedback (OGD/FKM)
- Dependence on initial prices
  1. $p_1(0) > p_2(0)$
  2. $p_1(0) = p_2(0)$
  3. $p_1(0) < p_2(0)$
- **"Surreptitious exploration"**: model testing mechanism; escapes the minor LNE in 100% of replications

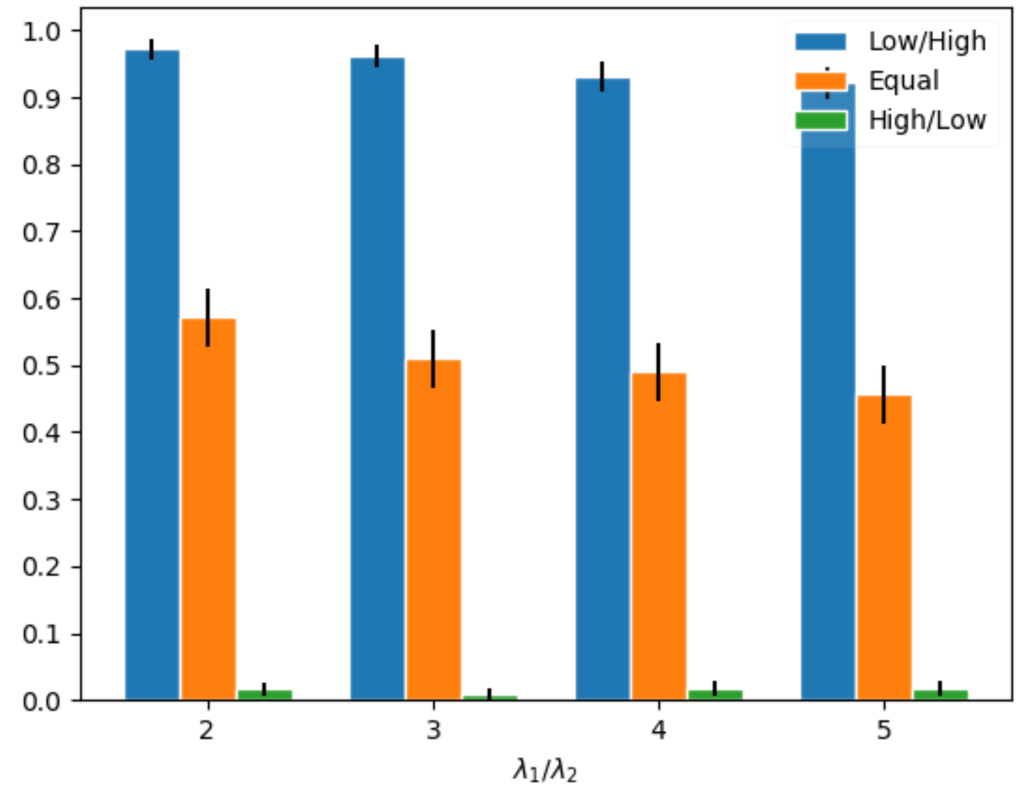

Fraction of time prices converge to minor LNE (OGD)