Dynamic Service Fee Pricing under Strategic Behavior: Actions as Instruments and Phase Transition

Rui Ai, David Simchi-Levi, Feng Zhu

MIT IDSS & LIDS

November 6, 2024

Motivations

- A large number of transactions nowadays take place on third-party platforms, e.g., Amazon, Uber and DoorDash.
- The platform earns profit by setting service fees $\Rightarrow$ Maximize.

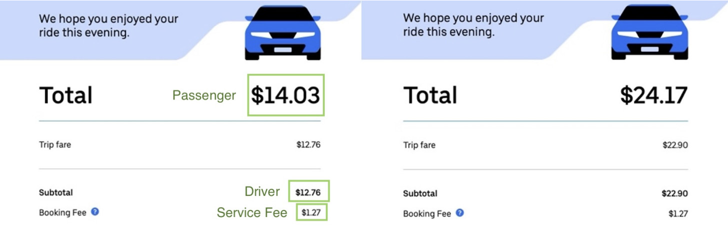

Figure: Uber charges a booking fee. From personal experience, Uber adjusts the value of the booking fee on a monthly basis.

Challenges

- Demand information needs to be learned.
- Only equilibria can be observed.
- Buyers may exhibit strategic behavior.
- Observe $P_t^e = P_{St}(Q_t^e) + a_t = P'_{Dt}(Q_t^e)$, but not P_{Dt}.

Question: Can non-i.i.d. actions (e.g. service fees) serve as instrumental variables in the problem of online pricing?

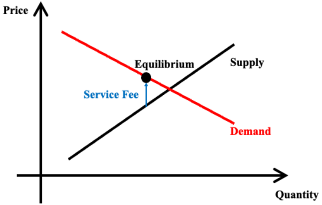

Figure: Market equilibrium.

Model and Assumptions

- Supply curve: $P_{St}(Q) = \alpha_0 + \alpha_1 Q + \epsilon_{St} \sim \mathcal{N}(0, \sigma_S^2)$.
- Demand curve: $P_{Dt}(Q) = \beta_0 + \beta_1 Q + \epsilon_{Dt}$.
- Utility-maximizing buyer: If the buyer has discount rate $\gamma = 1$, $\Omega(T)$-regret is unavoidable (Negative result).
- Time-sensitivity and user stickiness: Platform is patient that $\gamma = 1$ while buyer's discount rate is $\gamma \in [0, 1)$.
- $\text{Regret}(T) = \sum_{t=1}^{T} \mathbb{E}[a_t^* \cdot Q_t^e(P_{St}, P_{Dt}, a_t^*) - a_t \cdot Q_t^e(P_{St}, P'_{Dt}, a_t)]$.

Methods

- A carefully designed active randomness injection to balance exploration and exploitation effectively.
  - If market noise is smaller than $O(\frac{1}{\sqrt{T}})$, add artificial randomness to a_t.
- Using non-i.i.d. actions as instrumental variables to consistently estimate demand.
  - Martingale concentration analysis.
- A low-switching cost design that promotes nearly truthful buyer behavior.
  - Only update our policy $O(\log T)$ times.
  - In practice, Uber changes the booking fee at a low frequency.

The Algorithm

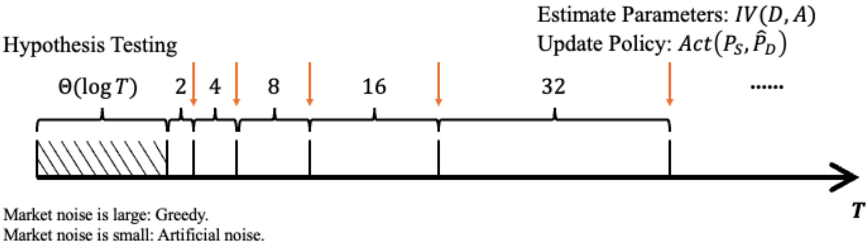

Figure: Timeline of the algorithm.

Results

- With probability at least $1-\iota$, the algorithm achieves Regret $\lesssim O(\sqrt{T}\log(\frac{\log T}{\iota})+\frac{\log T}{(1-\gamma)^2})$ without noise in the market.
- With probability at least $1-\iota$, the algorithm achieves Regret $\lesssim O(\frac{\log T\log(\frac{\log T}{\iota})}{\sigma_S^2}+\frac{\log T}{\sigma_S^2(1-\gamma)})$ with noise in the market.
- **Takehome: Noise Helps Learning!**
  - Explore the unknown environment (tail).
  - Bound the variance of the estimator.

Results

- Worst-case lower bound: $\Omega(\sqrt{T} \wedge \frac{\log T}{\sigma_S^2(1+\max\{0,\log(1/\sigma_S)\})})$.
  - New way to design hard-to-differentiate instances: matched noise magnitude.
- Our algorithm is optimal both in the number of rounds and the market noise level!
  - Regret $\lesssim \tilde{O}(\sqrt{T} \wedge \sigma_S^{-2}) \Rightarrow$ tight in both T and σ_S (no need to know in advance).
- **Takehome: Actions themselves can serve as instrumental variables in machine learning problems!**

Phase Transition

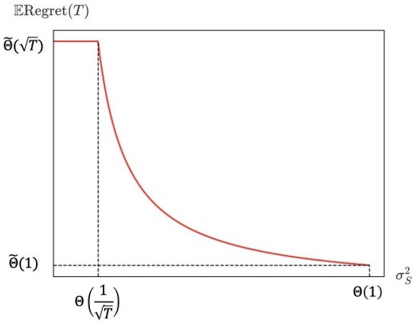

Figure: Theoretical phase transition and actual performance.

Thank You!

Contact me via `ruiai@mit.edu`**!**