Online Market Intermediation

Yiannis Giannakopoulos[1], Elias Koutsoupias[2] and Philip Lazos[2]

[1]Department of Informatics, Technical University of Munich [2]Department of Computer Science, University of Oxford

to appear in ICALP 2017

The Online Market Intermediation Problem

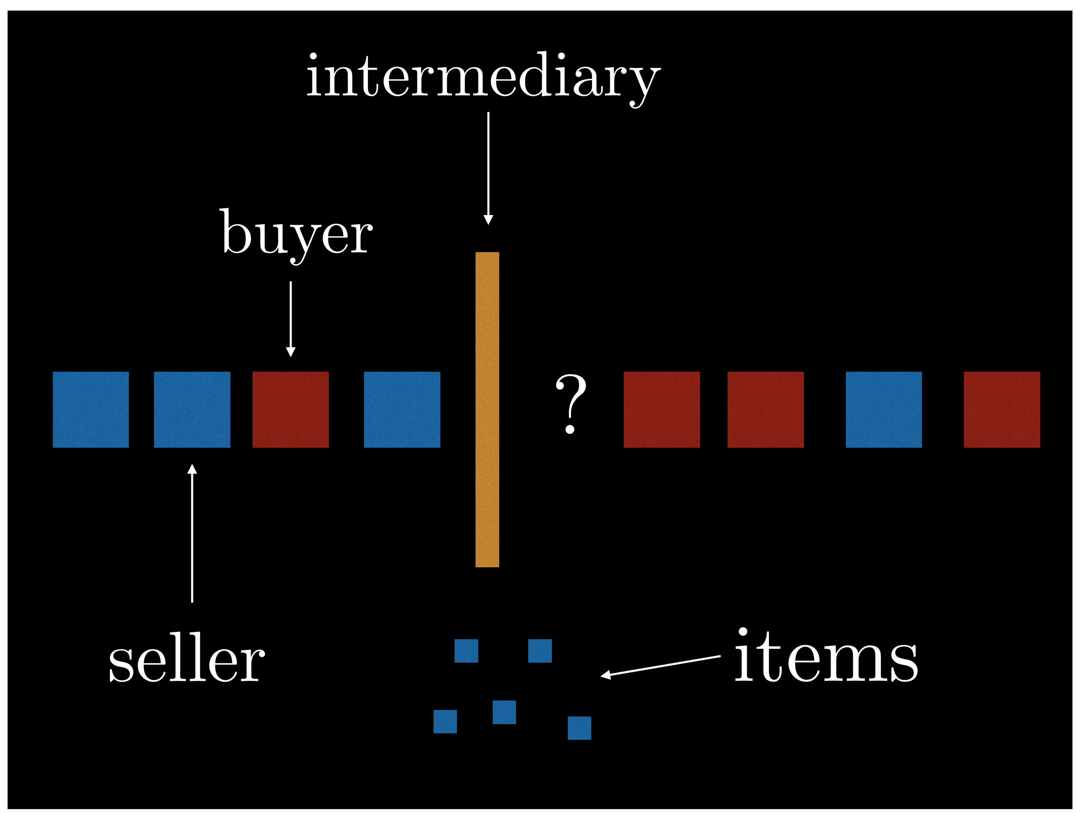

- A sequence σ of n agents, buyers and sellers
- Agents are interested in trading one item only and all items are identical.
  - Sellers enter the market with one item to sell and buyers want to buy one item.
  - Agents are strategic with quasilinear preferences. Their values follow distributions F_S and F_B for sellers and buyers respectively.

Intermediary: The intermediary interacts with the sequence σ in an online way. The number of agents is unknown and they are revealed one at a time. Interaction with agents is performed with **posted prices**. The intermediary starts with no items in stock.

Objectives

Welfare A natural objective is to maximize the social welfare $\mathcal{W}(\sigma)$: the sum of utilities of all agents, plus the intermediary. In this case payments cancel out, and the goal becomes transferring items to high value agents.

Profit Maximizing the intermediary's profit $\mathcal{R}(\sigma)$ is trickier: trades are only beneficial if performed at the right price and hoarding too many items can be easily penalized.

Variants We study three versions of the problem. The unrestricted, the K-item and α-balanced. In the K-item setting the intermediary is allowed to hold up to K items at most, while in the α-balanced the ratio between sellers and buyers is known.

Competitive Ratio

An algorithm is c-competitive for profit if for any σ, F_S and F_B we have:

$$\mathcal{R}_{OPT}(\sigma) \leq c\mathcal{R}(\sigma) + O(\mu_S),$$

where OPT is the optimal offline algorithm who knows the future, but not the result of random draws.

The additive term $O(\mu_S)$, where μ_S is the mean value of a seller, is required. Intuitively, it's the starting budget.

The definition for welfare is similar.

Distributional Assumptions

F_S and F_B have to follow stronger regularity assumptions than Myerson and Satterthwaite. In particular we need $\log(F_S(x))$ and $\log(1 - F_B(x))$ are concave (MHR). Just regularity would yield $\Omega(n)$ bounds.

The following properties are useful when dealing with such distributions. For $Y \sim F_B$:

1. $\Pr[Y \geq y] \geq \frac{1}{e}$ for any $y \geq \mu_B$ and $\Pr[Y \geq y] < \frac{1}{e}$ for any $y > 2\mu_B$
2. $\mathbb{E}\left[Y^{(m)}\right] \leq H_m \cdot \mu$ and $\Sigma_{i=m-k+1}^{m} \mathbb{E}[Y^{i:m}] \leq k\mu + 2\sqrt{kms}$
3. $x \leq e\mu F_S(x)$ for any $x \leq \mu_S$

These allow us to quantify relations between prices, probabilities and expectations.

Profit

Theorem 1. *The competitive ratio for profit is:*

- $\Theta(\sqrt{n})$ *in the unrestricted case.*
- $O(\log n)$ *in the K-item case.*
- $1 + o(1)$ *in the α-balanced case.*

Welfare

Theorem 2. *The competitive ratio for welfare is:*

- $\Theta(\log n)$ *in the unrestricted and K-item case.*
- 4 *in the α-balanced case.*

Algorithm

The posted price algorithm for the unrestricted setting:

- To the i-th seller post $q_i = F_S^{-1}\left(\frac{1}{e} \cdot \frac{1}{i^{1/2+\varepsilon}}\right)$
- Post to all buyers price $p = \mu_B$.

For welfare, the algorithm posts μ_S and μ_B as prices in all settings.

Lower Bound

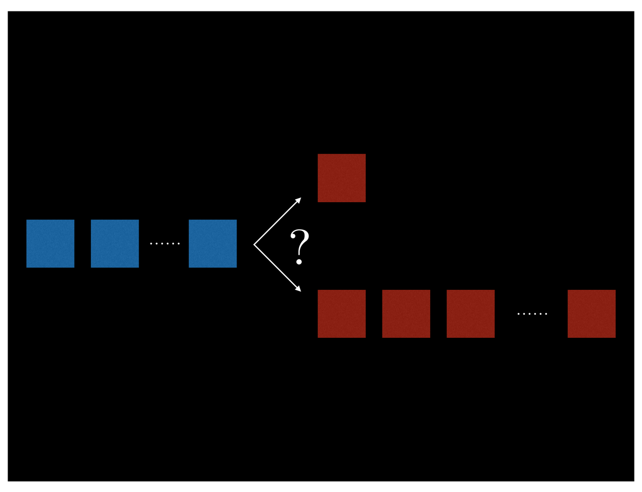

The lower bound of $\Omega(\sqrt{n})$ is achieved by a long sequence of sellers following either one or many buyers. The intermediary can only spend $O(\mu_S)$.

As such, the online can store at most $O(\sqrt{n})$ items from n consecutive sellers.

Upper Bound

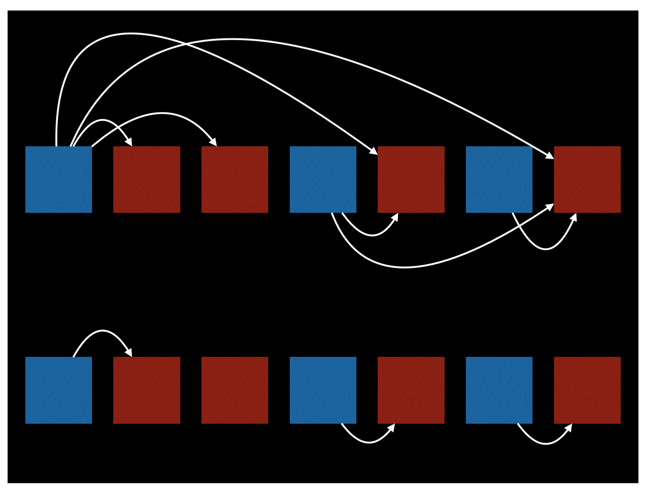

On top are the (potential) sales attempted by the offline. The online algorithm can always attempt a subset of those sales, by computing a FIFO matching between sellers and buyers.

The bound follows from the number of trades combined with Property 3 of the distribution.

K-Item

Since the intermediary can only hold at most K items, the trades generated by long runs of sellers are fewer. The pink edges are infeasible.

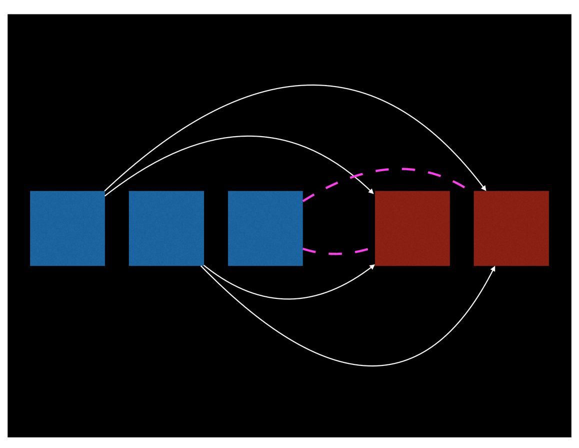

The online algorithm:

- Posts price $q = F_S^{-1}\left(\frac{1}{r} \cdot \frac{1}{2eK}\right)$ to sellers, if stock is not full.
- Posts price $p = \mu_B$ to all buyers.

Note that the potential losses are still $O(\mu_S)$. The online matching produced by this algorithm is the FIFO matching, rejecting sellers if the queue contains more than K elements.

α-Balanced

In this case, a ratio α between sellers and buyers is known. In particular the ratio must drop below α for any prefix of σ and should be tight at the end.

A fractional relaxation gives rise to the following constraint optimization, where m is the number of buyers.

$$\begin{aligned} &\max\ m\left(p(1 - F_B(p)) - \alpha \cdot qF_S(q)\right) \\ &\text{s.t.}\quad 1 - F_B(p) = \alpha F_S(q) \\ &\qquad p, q \in [0, \infty). \end{aligned}$$

Note that the prices do *not* depend on the length of σ.

These prices can then be used for any α-balanced sequence, with the expected profit converging to the optimal.