Learning to Clear the Market

International Conference on Machine Learning (ICML)
June 11, 2019

Weiran Shen, Tsinghua University
Sébastien Lahaie, Google Research
Renato Paes Leme, Google Research

Reserve Pricing

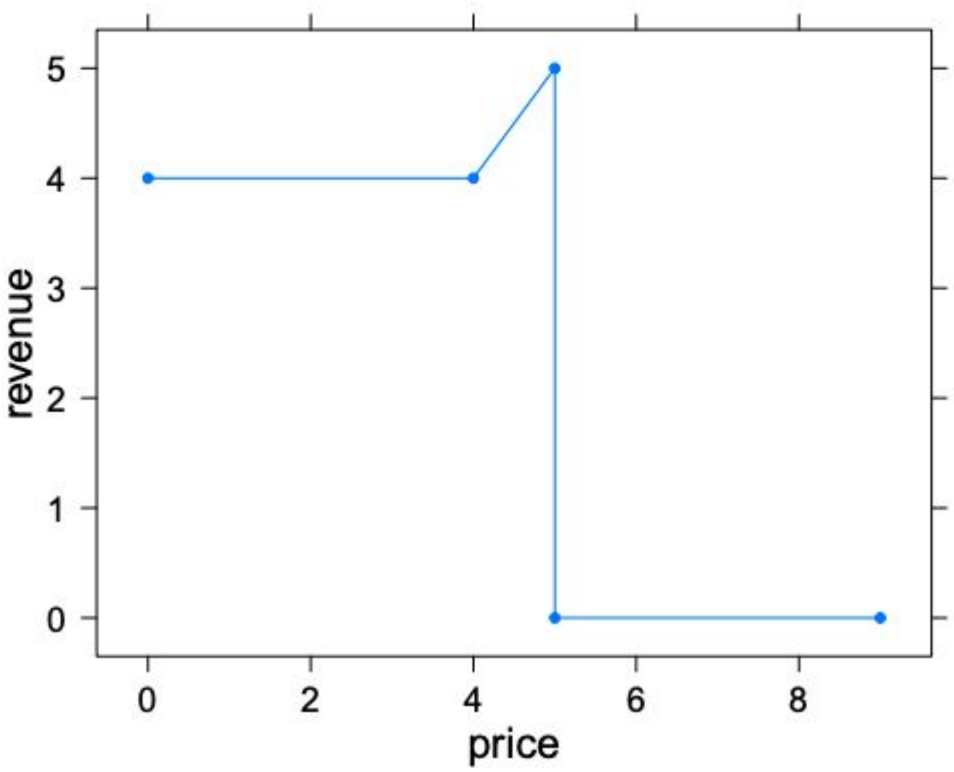

Bidders: $1, $2, $4, $5

- nonconvex, discontinuous [Medina-Mohri, 2014]
- gradient often zero
- optimal reserve is aggressive: high probability that the item is unsold

Market-Clearing Price

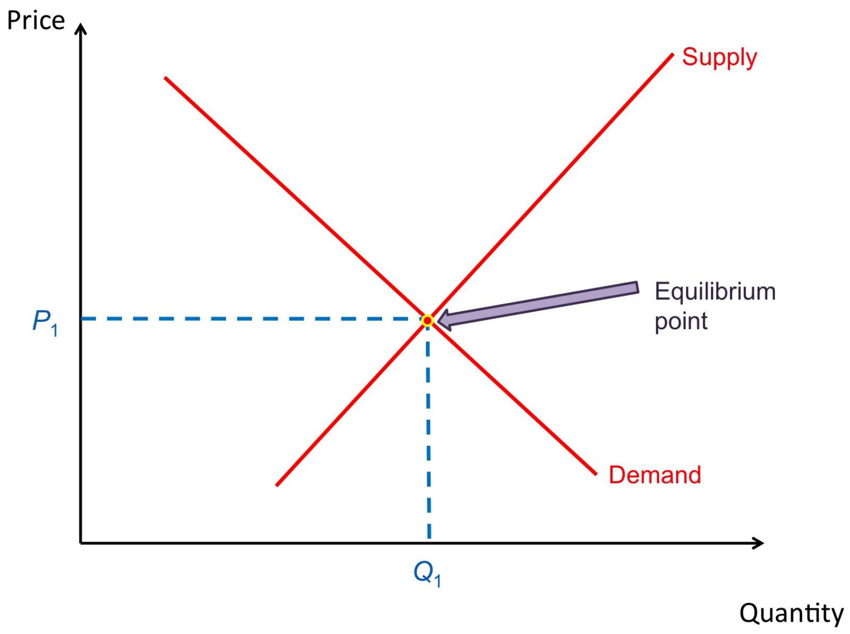

In a display ad auction:

- Supply = 1 single impression
- Want: Demand = 1 single bidder
- Set price between first- and second-highest bids.

Deriving the Loss Function

Formulate the (trivial) efficient allocation problem as an LP:

$$\max_{x \geq 0} \sum_i b_i x_i \quad \text{s.t. } \sum_i x_i = \lambda$$

Default choice is $\lambda = 1$. The dual of the allocation problem is a pricing problem:

$$\min_p \sum_i \max\{b_i - p, 0\} + \lambda p$$

Artificially **increasing** or **limiting** supply via λ controls how **conservative** or **aggressive** the resulting prices are.

Market Clearing Loss

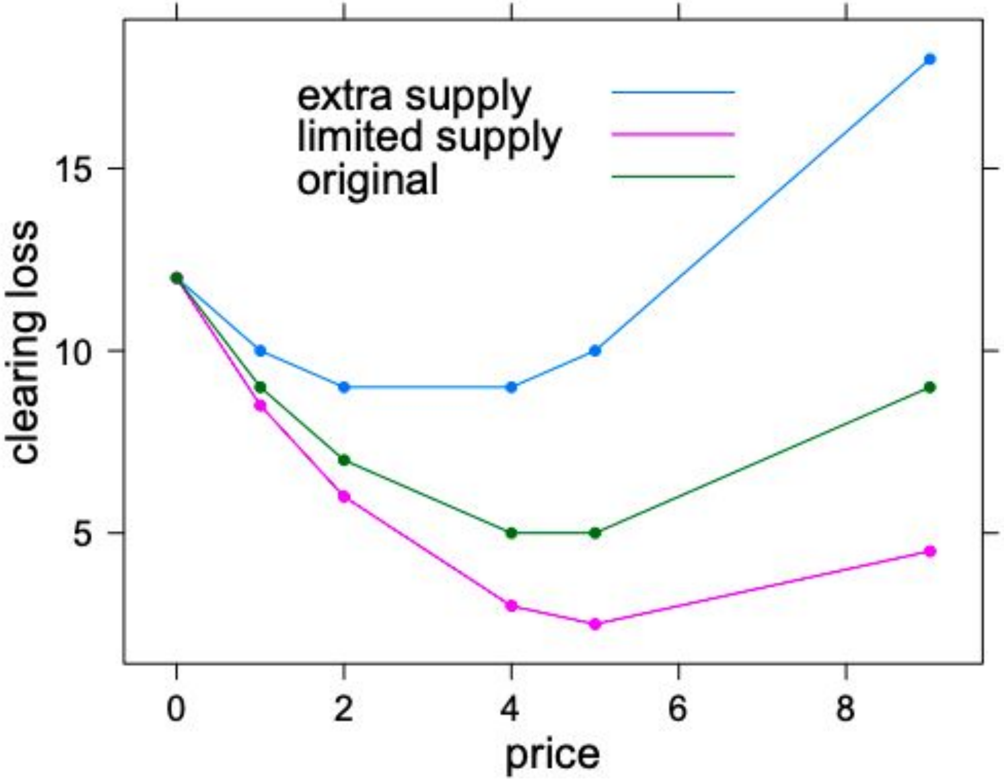

Bidders: $1, $2, $4, $5

- piecewise linear, convex
- robust to outliers
- all bids shape the loss

Revenue vs. Match Rate Trade-off

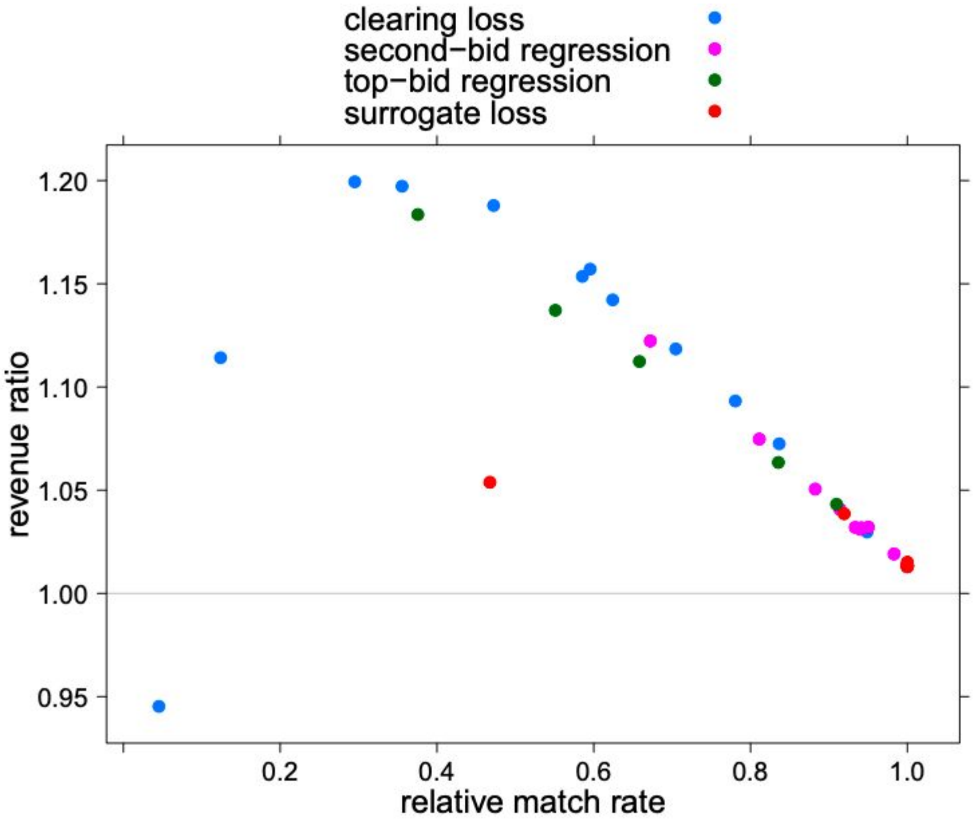

Summary

- Loss that captures the “market value” of an item (e.g., an ad impression).
- Allows fine-grained control of the revenue vs. match rate trade-off.
- Outperforms regression and surrogate loss benchmarks in terms of trade-offs and convergence rates.

More details at Poster #156.