LEARNING COMPETITIVE EQUILIBRIA IN NOISY COMBINATORIAL MARKETS

ENRIQUE AREYAN VIQUEIRA, CYRUS COUSINS & AMY GREENWALD

AAMAS 2021 — ONLINE, MAY 3–7

BROWN — Computer Science

IN A NUTSHELL

In a *combinatorial market*, buyers have preferences over bundles or packages of goods.

Computing the value of a bundle is often complicated, for example, the value of showing ads to users, the value of oil drilling sites, etc.

We present a market model with *noisy buyer valuations* and two *learning algorithms* that provably learn the *market's equilibria* in our model.

MODEL

A set of **goods**, for example: 🏀, 🛒, ⚽.

A set of **buyers**, let's call them Ana and Bob. Buyers have **values** for packages of goods.

For example,

Anna values 🏀 and 🛒 at 10 (basketball fan).

Bob values ⚽ at 100 (soccer fan).

In our example, there are $2^3 = 8$ many packages.

An **allocation** is a partition of goods for buyers.

For example, Anna gets 🏀 and 🛒, Bob gets ⚽.

A **competitive equilibrium** is an allocation and prices that make *all market participants* happy: the seller maximizes its *revenue* and the buyers maximize their *utility* (value minus payment) over all possible allocations at the fixed prices.

But!, buyers might not know their values *exactly*.

Instead, Ana and Bob can provide **confidence intervals** around their values.

For example, with 95% probability:

Bob values ⚽ at 100 ± 10 (soccer fan – mostly).

Goal: estimate a market's **competitive equilibrium** when buyers don't communicate their exact values, only confidence intervals around them.

WHY DO WE CARE?

We want markets to run *efficiently* even when buyers can't perfectly evaluate bundles. This is a common situation in practice.

For example, in *spectrum auctions*, telecommunication companies bid for spectrum licenses.

A company's bid reflects *future expected profit* and thus depends on unobservable factors: future customer demand, future competitors, etc.

ELICITATION ALGORITHMS

Baseline: estimate *all* values via some concentration inequality (e.g, Hoeffding's inequality).

Elicitation with *Pruning*.

Initially, all buyer, bundle pairs are *active*. Ask every buyer for *rough value estimates* (say up to "large" $\varepsilon_0 > 0$) for all bundles.

Repeat some number of times, $t = 0, 1, \ldots T$:

- Optimally allocate goods to buyers per current value estimates. Let V_t be this allocation's value.

- For each *active* buyer, bundle pair:

  (i) Suppose the buyer gets the bundle.
  (ii) Optimally allocate all other goods to buyers.
    – If V_t is higher than value accrued by (i) + (ii), then *deactivate* buyer, bundle pair.

- Increase value estimates accuracy: $\epsilon_{t+1} \leftarrow \epsilon_t/2$

We proved that *pruning criteria* (i) + (ii), preserves the *expected market's* competitive equilibria. Expectation is over all *buyers' values uncertainty*.

FUTURE RESEARCH

Our learning algorithm shows promise since it can learn markets well with far fewer reports from users compared to a baseline. Still, there is significant room for improvement.

In particular, our methodology is flexible enough to work in *any* combinatorial market. Minimal modifications will allow us to inject domain-dependent knowledge, an exciting future research direction.

EXPERIMENTAL RESULTS

The goal of our experiments is to robustly evaluate the *empirical* performance of our algorithms. We evaluate our algorithms on *unit-demand* valuations and the *local synergy value model*.

Unit-demand valuations. A unit-demand buyer has value for single-items packages. We constructed four different distributions over unit-demand valuations, each capturing distinct qualitative features of markets with unit-demand buyers (for details, please read our paper).

Figure 1 summarizes our unit-demand results. The plots show the number of buyers (y-axis) and the number of goods (x-axis) for each distribution. Elicitation with pruning *consistently* requires fewer confidence interval reports from buyers than the baseline, and darker colors indicate higher savings.


Figure 1: sample efficiency of elicitation with pruning on unit-demand markets results.

Local synergy value model (Scheffel et.al. (2012)). Telecommunication service providers might value different bundles of radio spectrum licenses differently, depending on whether the licenses in the bundle complement one another. For example, a bundle including New Jersey and Connecticut might not be very valuable unless it also contains New York City.

The local synergy value model is a simple model of regional complementarities. For various target errors ε, table 1's left column reports confidence interval report savings, the middle column reports the actual error achieved, and the right column a metric of buyers happiness (zero is maximal happiness).

Target Error	% Savings with Pruning (±4%)	Error guarantee (±0.01)	UM Loss (±0.0005)
1.25	18%	0.89	0.0011
2.50	11%	1.78	0.0018
5.00	-7%	3.59	0.0037
10.0	-35%	7.27	0.0072

Table 1: local synergy value model results. Elicitation with pruning is often more data efficient, achieves smaller errors, and results in outcomes where buyers are (almost) happy.

CONTACT INFORMATION

Lab Web: glamor.cs.brown.edu

Personal Web: www.enriqueareyan.com

Contact Email: eareyan@brown.edu