


Extending Myerson's Optimal Auctions to Correlated Bidders via Neural Network Interpolation

Mingyu Guo¹Jiayuan Liu²Vincent Conitzer²¹Adelaide University²Carnegie Mellon University


The problem

- ▶ **Goal:** given a **correlated** prior over bids, a **deterministic, strategy-proof, ex post individually rational** single-item auction maximizing expected revenue.
- ▶ **Myerson (1981)** solved this for **independent** bidders: greedily allocate by (ironed) virtual valuation. Monotone allocation $\Rightarrow$ strategy-proofness.
- ▶ **Real bidders are correlated.** *Marginal profit* (Papadimitriou & Pierrakos 2011) is the correlated version of virtual valuation, but greedily allocating by it is usually **non-monotone**, hence not strategy-proof.
- ▶ Optimal design for ≥ 3 correlated bidders is **NP-hard to approximate**, so we take a **computational approach**.


Key idea: tiny network + verification




(a) greedy on marginal profit: **non-monotone**



(b) tiny-network interpolation: **monotone**



(c) revenue fix: C and D become no-sale



(d) final strategy-proof allocation

- ▶ Axes are the two bids. **Blue**: bidder 1 wins; **red**: bidder 2 wins; white: no sale.
- ▶ **Monotone** = blue region *rightward-closed* and red region *upward-closed*. Greedy allocation on **marginal profit** (a) violates this: blue points enclosed in red, red points enclosed in blue.
- ▶ A **tiny MLP+ReLU** trained to imitate greedy lacks the capacity to reproduce these violations: it **“glosses over” the noise** and recovers the macro-trend.
- ▶ The result is **formally verified** monotone by MIP, so it is exactly strategy-proof. Payments follow from Myerson's characterization.

Pipeline

1. Interpolate. Train a tiny MLP+ReLU on sampled bid profiles to imitate the greedy allocation on marginal profit

2. Verify. Each ReLU can be modeled with a binary variable, so the network's worst case can be analyzed via MIP. We find the **worst-case violation**: the largest bid increase that turns a winner into a loser. Objective zero $\Rightarrow$ monotone

3. Counterexample-guided training. Each violation found by the MIP is a counterexample: penalize any drop in the bidder's win share along that bid increase, and keep training

4. Monotonicity fix (post-processing). For the tentative winner, a second MIP finds the highest bid at which she would still lose; the sale price is raised to at least this **cutoff**, pricing out every remaining violation

Guarantee. Exact strategy-proofness for the given distribution; near-optimal revenue in practice

Verification needs **small** networks, which also filter out non-monotone detail. Steps 3–4 are rarely needed on natural distributions, but **essential on adversarial ones** (method 1: evolution; method 2: MIP construction).

Why marginal profit?

- ▶ **Virtual valuation:** allocate to the highest; no sale if all negative.
- ▶ **Marginal profit** (Papadimitriou & Pierrakos 2011) is **always** ≥ 0 : the “when not to sell” decision is deferred to a post-hoc **revenue fix**. If the network’s winner i faces offer o , raise it to

$$\arg \max_{x \geq o} x (1 - F_i(x \mid b_{-i})),$$

where $F_i(\cdot \mid b_{-i})$ is i ’s conditional CDF given the others’ bids.

- ▶ Raising offers never breaks monotonicity, so the network can defer the “no sale” decision and only learn who wins. This is the step from panel (c) to panel (d).

Benchmark: the greedy upper bound

- ▶ Expected greedy revenue on marginal profit is an **unattainable upper bound** on the optimal revenue (it ignores monotonicity). **First such benchmark**; we systematically test against it on synthetic, adversarial, and real distributions.
- ▶ We report

$$\text{GAP}(\text{METHOD}) = \frac{\text{Rev}(\text{GREEDY}) - \text{Rev}(\text{METHOD})}{\text{Rev}(\text{GREEDY})},$$

i.e. **distance to optimality**.

- ▶ **Baselines**: second price with optimal reserve (2ND), Myerson on marginals (MYER), **fully strategy-proof RegretNet** (RNET, built with our verification), and Automated Mechanism Design (AMD) for 2 bidders, as an alternative to Papadimitriou & Pierrakos's FPTAS.

Results: 65 correlated distributions

Average revenue gap from the greedy upper bound (random grid distributions)

MP = our marginal-profit interpolation

	MP (OURS)	RNET	2ND	MYER
$n = 2$	0.26%	2.7%	4.2%	4.3%
$n = 3$	0.36%	3.6%	3.9%	3.8%
$n = 4$	0.78%	3.5%	3.2%	3.2%
$n = 5$	0.87%	2.8%	2.4%	2.4%

- ▶ **Tenfold** smaller gap than the best baseline for $n = 2$. MP is the best scalable method on every distribution; for $n = 2$ it often beats AMD, the discrete-bid optimum.
- ▶ Test bed: 40 random, 7 hand-crafted, 12 adversarial grids; 3 continuous; 3 real-world.

Gap from the greedy upper bound on empirical bid distributions

Domain	Bidders	2ND	MYER	MP (OURS)
Consumer (eBay)	layperson	22.7%	34.5%	1.3%
Financial (Ethereum)	algorithmic	7.7%	14.5%	0.6%
Electricity (PJM)	professional	37.1%	37.2%	5.3%

- ▶ Three domains, three bidder types: laypeople bidding for electronics, algorithms bidding for block-building rights, professionals bidding for electricity generation rights.

- ▶ **Multi-unit, unit demand** (m items): greedy serves the m highest marginal profits; a winner must stay a winner when raising her bid. Same pipeline.
- ▶ **RegretNet made exactly strategy-proof**: keep a tiny MLP allocation net, verify and fix it with our MIPs; the payment net is only a training surrogate.
- ▶ **A general mechanism design property verification tool**: interpolate any mechanism (formula or code) with a tiny network, then verify it via MIP.

Take-aways

- ▶ Where theory hits NP-hardness, **computational rigor** can still deliver: verifiably strategy-proof and near-optimal auctions for correlated bidders.
- ▶ Small networks are a feature: limited capacity acts as a monotonicity-seeking regularizer.
- ▶ **Simple is best:** of the three monotonicity-seeking measures (counterexample-guided training, post-processing monotonicity fix, and simply re-initializing, training again and re-verifying), **try again** usually works best.

Thank you. Poster session: come and see the allocations up close.