Enhancing Affine Maximizer Auctions with Correlation-Aware Payment

Haoran Sun[1] Xuanzhi Xia[2] Xu Chu[1] Xiaotie Deng[1]
[1]CFCS, School of Computer Science, Peking University | [2]Department of Computer Science and Technology, Tsinghua University

Paper Homepage

ICML International Conference On Machine Learning
PEKING UNIVERSITY
Tsinghua University

Main Idea

Affine Maximizer Auctions (AMAs) are DSIC and IR by construction, but their VCG-style payment is too rigid under bidder-correlated valuations. CA-AMA adds a correlation-aware payment term that depends only on other bidders' valuations, so it can exploit correlation without changing the bidders' incentive to bid truthfully.

Motivated Example: Negative Correlation

Consider a single-item auction with two bidders, $v_1 = 1 - v_2$, and $v_1 \sim U[0,1]$. The optimal DSIC/IR mechanism extracts full surplus, but the required reserve has the opposite monotonicity from AMA.

Optimal reserve
$p_1^* = 1 - v_2$
Shape: decreases with value v_2.

Classic AMA
$p_1 = \max\{v_2 + \lambda_2, \lambda_1, \lambda_0\} - \lambda_1$
Shape: non-decreasing in v_2.

Mismatch. The optimal price moves in a direction that VCG-style AMA payments cannot express.

Why Classic AMA Struggles

For an AMA parameterized by $(\mathcal{A}, \boldsymbol{w}, \boldsymbol{\lambda})$, $\mathrm{asw}(k; V) = \sum_j w_j (\boldsymbol{v}_j \cdot (A_k)_j) + \lambda_k$ is the affine social welfare: linearly weighted social welfare plus an allocation boost.

$$k^* = \arg\max_{k \in [S]} \mathrm{asw}(k; V), \qquad p_i^{\mathrm{AMA}}(V) = \frac{1}{w_i}\left(\max_k \mathrm{asw}_{-i}(k; V) - \mathrm{asw}_{-i}(k^*; V)\right).$$

In negatively correlated environments, the revenue-optimal personalized reserve for bidder i may need to **decrease** as other bidders' valuations increase. Classic AMA payments cannot express this shape: bidder i's payment is induced only through the affine welfare of the other bidders.

Failure Mode
AMA payment has insufficient profile-specific flexibility

Goal
Keep DSIC while learning correlation-aware prices

CA-AMA Mechanism

CA-AMA keeps the AMA allocation and adds a peer-conditioned payment:

$$g^{\mathrm{CA}}(V; \mathcal{A}, \boldsymbol{w}, \boldsymbol{\lambda}) = g^{\mathrm{AMA}}(V; \mathcal{A}, \boldsymbol{w}, \boldsymbol{\lambda}),$$
$$p_i^{\mathrm{CA}}(V; \mathcal{A}, \boldsymbol{w}, \boldsymbol{\lambda}, p^{\mathrm{Cor}}) = p_i^{\mathrm{AMA}}(V; \mathcal{A}, \boldsymbol{w}, \boldsymbol{\lambda}) + p_i^{\mathrm{Cor}}(V_{-i}).$$

Why DSIC is preserved. For bidder i, the extra term $p_i^{\mathrm{Cor}}(V_{-i})$ is constant with respect to i's own report. Therefore the truthful bid that maximizes AMA utility also maximizes CA-AMA utility.

$$u_i^{\mathrm{CA}}(\boldsymbol{v}_i, (\boldsymbol{b}_i, V_{-i})) = u_i^{\mathrm{AMA}}(\boldsymbol{v}_i, (\boldsymbol{b}_i, V_{-i})) - p_i^{\mathrm{Cor}}(V_{-i}).$$

Theoretical Takeaways

- **Expressiveness.** For any distribution, CA-AMA can recover AMA by setting $p_i^{\mathrm{Cor}} = 0$.
- **No free gain under independence.** In single-item bidder-independent settings, optimal deterministic CA-AMA and deterministic AMA have equal revenue.
- **Large gain under correlation.** For any $\epsilon > 0$, there exists a correlated single-item distribution where $\mathrm{REV}_{\mathcal{F}}^{\mathrm{D\text{-}AMA}} \le \epsilon \cdot \mathrm{REV}_{\mathcal{F}}$, while $\mathrm{REV}_{\mathcal{F}}^{\mathrm{D\text{-}CA}} = \mathrm{REV}_{\mathcal{F}}$.
- **Learnability.** The optimal target
$$p_i^{\mathrm{OPT\text{-}core}}(V_{-i}) = \inf_{\boldsymbol{v}_i \in \mathrm{supp}(\mathcal{F}_i(V_{-i}))} u_i^{\mathrm{AMA}}((\boldsymbol{v}_i, V_{-i}))$$
is continuous under mild assumptions, supporting neural parameterization.

Optimization Framework

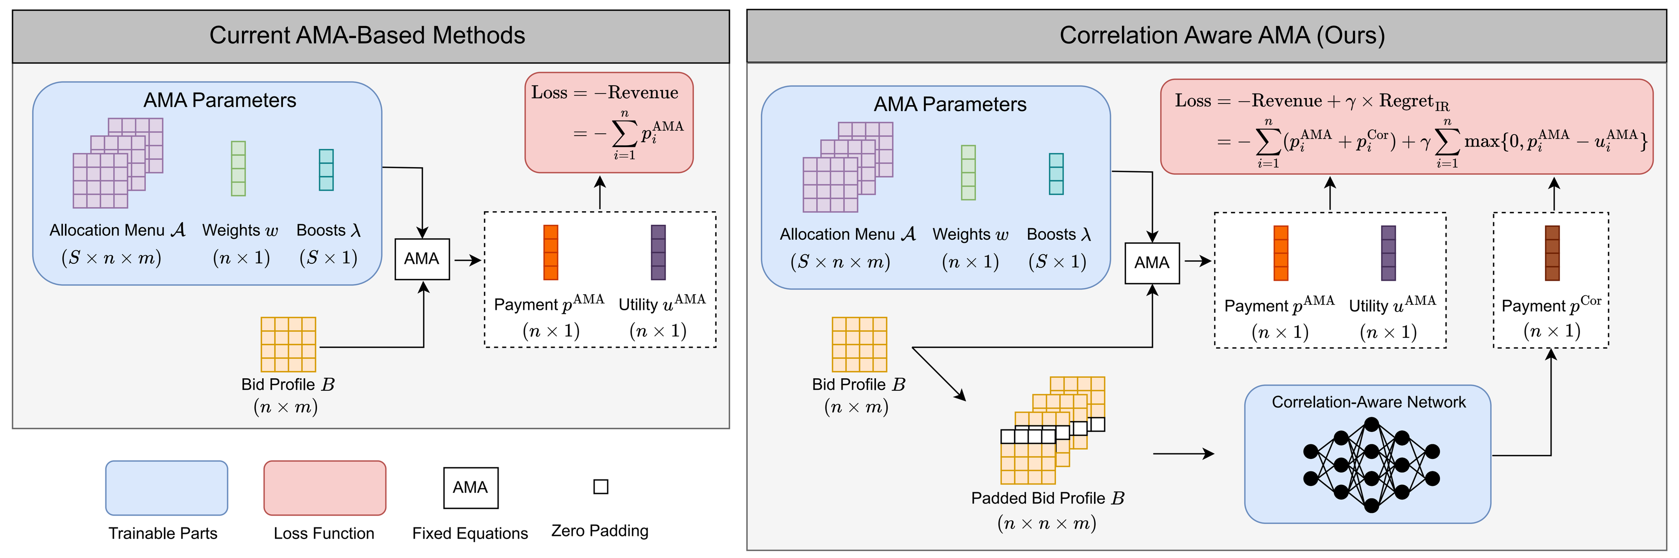

— Classic AMA optimizes allocation weights and boosts. CA-AMA adds p^{Cor} and uses IR-regret to control the only constraint that can be affected by the new payment term.

Two-Stage Training

We parameterize AMA components by a network with parameters θ and correlation-aware payments by a three-layer ReLU MLP with parameters ϕ.

$$\mathcal{L}(\theta, \phi) = \sum_{V^{(k)} \in D} \left(-\mathrm{Revenue}(V^{(k)}) + \gamma \cdot \mathrm{Regret}_{\mathrm{IR}}(V^{(k)})\right),$$

$$\mathrm{Regret}_{\mathrm{IR}}(V) = \sum_{i=1}^{n} \max\{0, -u_i^{\mathrm{CA}}(\boldsymbol{v}_i, V)\}.$$

Sample valuation profiles $D = \{V^{(k)}\}_{k=1}^{K}$ → Mutual training update θ, ϕ softmax AMA → Post-training freeze θ exact AMA payment

- **Mutual training:** jointly finds AMA structure and core prices.
- **Post-training:** freezes AMA parameters and refines p^{Cor} using exact payments.
- **Strict ex-post IR option:** let bidders with negative utility opt out, replacing u_i by $\max\{u_i, 0\}$ while preserving DSIC.

Where the Revenue Comes From

In correlated settings, $p_i^{\mathrm{Cor}}(V_{-i})$ acts like a learned personalized reserve inferred from peer information.

Component	Role in CA-AMA
AMA allocation	Keeps feasibility and interpretable allocation search.
p^{AMA}	Retains the VCG-style externality payment.
p^{Cor}	Extracts correlation-dependent surplus while independent of bidder i's own report.
IR-regret	Soft training signal for the only new feasibility risk.

Generalization of IR-Regret

For a 3-layer ReLU p_i^{Cor} network with bounded spectral norms, the empirical and expected IR-regret are uniformly close:

$$\left|\frac{1}{K}\sum_{k=1}^{K} \mathrm{Regret}_{\mathrm{IR}}(V^{(k)}) - \mathbb{E}_{V\sim\mathcal{F}}[\mathrm{Regret}_{\mathrm{IR}}(V)]\right| \le O\left(\sqrt{\frac{\log(1/\delta)}{K}}\right).$$

This makes the training objective a meaningful proxy for unseen valuation profiles.

Experimental Summary

Dist.	Setting	AMA	CA	CA ex-IR	Gain
D-0.5	2×2	0.7480	0.8532	**0.8261**	+10.4%
D-0.5	2×5	1.8808	2.3663	**2.2694**	+20.7%
D-0.5	3×10	3.1363	3.6205	**3.5623**	+13.6%
D-0.5	5×5	1.2123	1.3209	**1.3084**	+7.9%
D-2.0	2×2	0.6516	0.7131	**0.6995**	+7.4%
D-2.0	2×5	1.7362	1.9437	**1.8934**	+9.1%
D-2.0	3×10	2.8706	2.9448	**2.9292**	+2.0%
D-2.0	5×5	1.0239	1.0406	**1.0330**	+0.9%
L-0.6	2×5 Sym	2.2923	2.6305	**2.5628**	+11.8%
L-0.6	2×5 Asym	1.7135	1.9359	**1.8553**	+8.3%
L-0.8	2×5 Sym	2.4239	3.0837	**2.9643**	+22.3%
L-0.8	2×5 Asym	1.6677	2.2028	**2.1124**	+26.7%

— AMA = Randomized AMA; CA = IR-regret CA-AMA; CA ex-IR = strict ex-post IR opt-out.

Perfect Negative Correlation

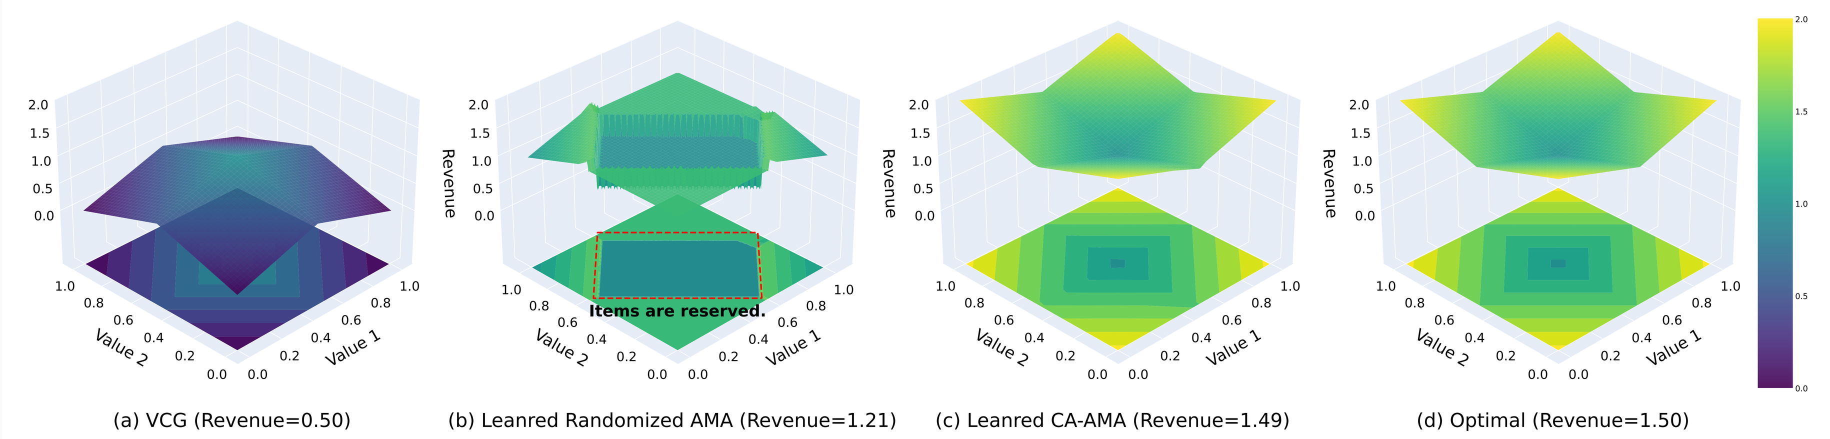

— CA-AMA tracks the optimal surface in a 2-bidder, 2-item negative-correlation case; Randomized AMA often loses revenue by reserving items.

Equal-Revenue Stress Test

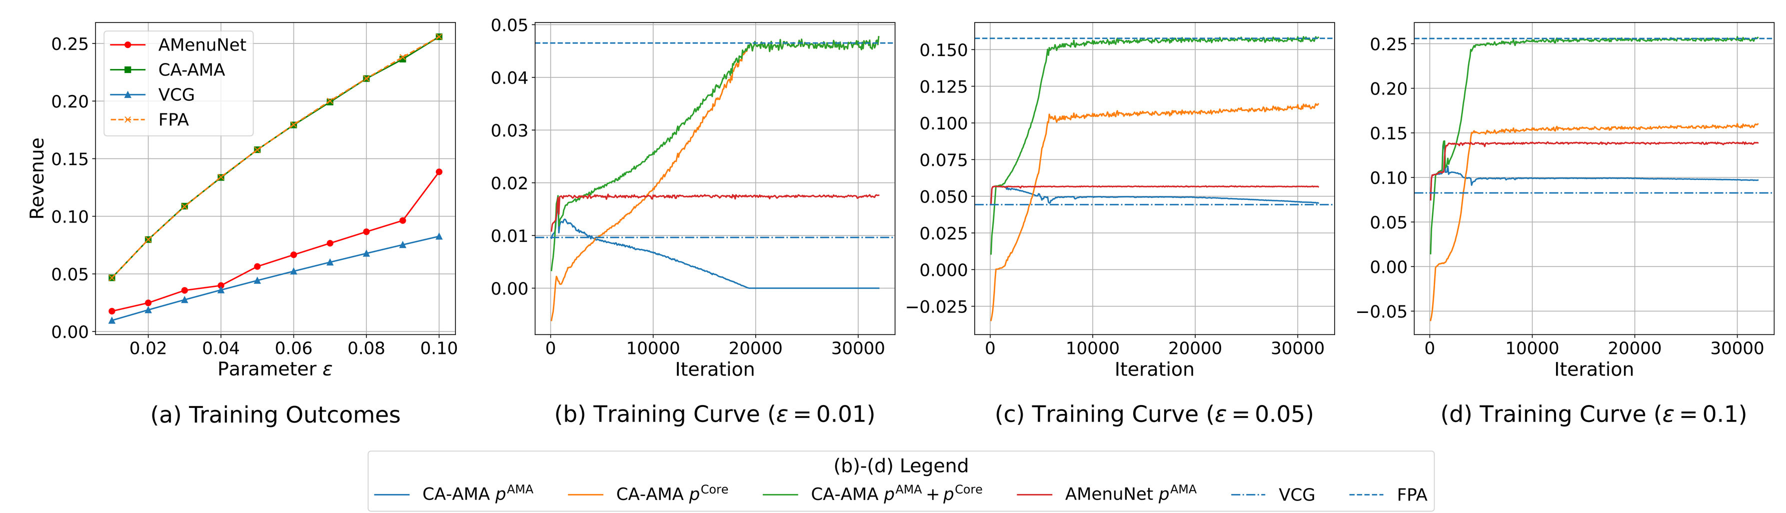

— CA-AMA converges near optimal revenue; the learned p^{Cor} explains most of the improvement.

Practical Cost

Setting	2×2	2×5	3×10	5×5
Randomized AMA	14 min	26 min	52 min	1 h 50 min
CA-AMA	15 min	27 min	54 min	1 h 55 min

CA-AMA: use bidder correlations to repair AMA payments while preserving DSIC by construction

Code: github.com/Haoran0301/CA-AMA