LEARNING COLLUSION-RESISTANT AUCTIONS VIA COALITION REGRET MINIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Auctions serve as efficient frameworks for resource allocation and have profoundly influenced digital economies, but deriving optimal mechanisms for multi-bidder, multi-item auctions remains analytically intractable. Recently, machine learning (ML) has partially addressed this by designing revenue-maximizing auctions through individual regret minimization. However, these approaches primarily ensure strategy-proofness against single bidders, leaving mechanisms vulnerable to collusion, which often leads to substantial revenue loss. To address this, we introduce **CoalitionNet**, an ML-powered auction framework designed to minimize coalition regret. Specifically, it leverages utility externalities to evaluate joint deviations and trains a strong collusion generator for mechanism adversarial learning. Extensive experiments demonstrate that CoalitionNet achieves near-zero coalition regret while maintaining competitive revenue performance. Furthermore, it also demonstrates strong generalization across unseen coalition sizes as well as valuation distributions.

1 INTRODUCTION

Auctions have long been used for allocating scarce resources and coordinating economic activities in a wide range of markets (Börgers et al., 2015). Classical auction theory derives revenue-optimal mechanisms under well-structured assumptions through closed-form analysis and careful reasoning about incentive constraints (Milgrom & Weber, 1982; Klemperer, 1999). However, these analytical approaches are typically confined to idealized environments with a small number of bidders and simple valuation structures. As the number of participants and items increases, the dimensionality of the strategy space grows rapidly, making purely analytical derivations increasingly difficult (Balcan et al., 2016; Duan et al., 2023).

Over the past few years, machine learning (ML) has emerged as a powerful tool for auction design, enabling data-driven discovery of mechanisms in complex, high-dimensional settings (Jordan & Mitchell, 2015; LeCun et al., 2015). By parameterizing mechanisms with flexible function approximators and optimizing them via gradient-based methods, ML approaches overcome the analytical intractability of classical designs and scale naturally to large problems. These methods have been successfully applied to real-world domains such as digital advertising (Duan et al., 2022; Li et al., 2025), blockchain ecosystems (Zhang et al., 2025), cloud resource markets (Xu et al., 2025; Yu et al., 2025) and large language models (Duetting et al., 2024).

Despite recent advances, existing learning-based approaches typically model bidders as independent agents, leaving coordinated strategic behavior largely unaddressed. However, the cost of such behavior is evident in real auction markets: bidder collusion in timber auctions was estimated to reduce seller revenue by 7.9% (Baldwin et al., 1997), while coordinated bidding in sponsored search auctions was found to notably reduce platform revenue (Decarolis et al., 2023). These findings motivate the development of a learning framework that explicitly accounts for coalition deviations.

To address the vulnerability of existing learning-based auctions to bidder collusion, this paper presents **CoalitionNet**, a robust mechanism design framework. Our main contributions are summarized as follows:

Adversarial Framework for Coalition Regret Minimization. We formulate coalition-regret minimization as an adversarial learning process between the auction mechanism and a collusion

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

generator. The generator searches for profitable joint misreports, while the mechanism is trained against these deviations. This process directly suppresses incentives for coordinated manipulation.

Externality-Guided Collusion Generation. We propose an externality-guided collusion generator that evaluates the marginal utility gains of pivot bidders and analytically identifies optimal coalitions, avoiding exhaustive search over exponentially many bidder subsets. This formulation also provides dense bidder-level signals, mitigating the near-zero gradients caused by coordinated bid reductions and enabling the detection of revenue-reducing joint deviations.

2 Related Work

Optimal Auction Design as a ML Problem With advances in numerical optimization and algorithmic game theory (Roughgarden, 2010; Vohra, 2011), the intersection of ML and auction design has emerged as a powerful paradigm for finding unknown optimal auctions, especially in high-dimensional settings (Sandholm, 2003; Guo & Conitzer, 2010). The seminal work of RegretNet by Dütting et al. (2019) pioneered the formulation of optimal auction design as a learning task, integrating incentive compatibility (IC) and individual rationality (IR) constraints as penalty terms within the loss function to enable the discovery of strategy-proof mechanisms through gradient-based optimization on neural network.

Subsequent research has significantly broadened this framework to accommodate increasingly complex configurations. Shen et al. (2018) adapted this paradigm to single-bidder settings, leveraging the neural approach to discover revenue-optimal mechanisms with provable optimality. To handle varying item sizes, Ivanov et al. (2022) replaced architectures with transformers, enabling flexibility for bid sequences of arbitrary length. In the realm of combinatorial auctions, Weissteiner & Seuken (2020) applied this learning paradigm to iterative formats, while Pham et al. (2025) tailored the network structure to specifically address bidders with combinatorial valuations and Ravindranath et al. (2024) integrated reinforcement learning to solve sequential combinatorial auctions. To ensure strict strategy-proofness, Wang et al. (2024) and Jiang et al. (2026) achieved this by incorporating a linear programming post-processing step. More recently, Liu & Zhang (2025) extended the methodology to double auctions, leveraging transformer-based architectures to tackle information asymmetry on both the demand and supply sides, and Wang et al. (2025) conceptualized optimal auction design as a diffusion process, demonstrating superior performance in large-scale settings.

However, a potential limitation is the exclusive focus on *individual* strategic behavior. Most existing ML methods, following the foundational RegretNet, assume that each bidder acts independently. The training objectives are designed to minimize the regret of a single agent acting unilaterally, ignoring the possibility of coordinated manipulation. As a result, there remains a significant gap in the design of mechanisms that are robust to collective strategic behavior and group collusion.

Collusions in Auction Design Collusion remains one of the main challenges in auction design, as joint deviations by a coalition of bidders can significantly undermine revenue (Marshall & Marx, 2009). In FCC spectrum auctions, Cramton & Schwartz (2002) document bid signaling and retaliation used to coordinate license allocation. Asker (2010) show that collusion can harm bidders outside the coalition and distort allocations in collectible-stamp auctions. Early studies on collusion-proofness focused on characterizing *group strategy-proof* (GSP) mechanisms (Laffont & Tirole, 1993), which guarantee that no coalition can jointly misreport to make every member better off. Subsequent studies examined collusion in standard auction formats such as second-price and VCG auctions (Graham & Marshall, 1987; Shoham & Leyton-Brown, 2008), proposing structural modifications such as reserve prices and bidder asymmetry to deter coordinated bidding. However, these analytical approaches are typically restricted to specific auction formats and limited coalition sizes (Che & Kim, 2009), limiting their scalability to large and heterogeneous markets where the space of joint deviations grows exponentially. These limitations motivate a learning-based approach that can account for bidder collusion without enumerating joint deviations.

3 PRELIMINARIES

In this section, we establish the formal framework for analyzing auctions under strategic uncertainty, extending classical primitives to account for bidder collusion.

3.1 OPTIMAL AUCTION DESIGN

We consider the design problem in an auction environment with a set of bidders $N = \{1, \dots, n\}$ and a set of items $M = \{1, \dots, m\}$. Each bidder i has a private valuation vector $\boldsymbol{v}_i = (v_{i,1}, \dots, v_{i,m}) \in \mathbb{R}^m_{\geq 0}$. $v_{i,j}$ denotes bidder i's value preference for item j. For simplicity, we consider the bidders have additive valuations, where bidder i's value for receiving a bundle $S \subseteq M$ is $\boldsymbol{v}_i(S) = \sum_{j \in S} v_{i,j}$.

The valuation vector $\boldsymbol{v}_i$ is private information, known only to the respective bidder. The auctioneer observes neither the realized individual valuations $\boldsymbol{v}_i$ nor the joint valuation profile $\boldsymbol{v} = (\boldsymbol{v}_1, \dots, \boldsymbol{v}_n)$. Instead, the auctioneer assumes that bidders' valuations are mutually independent, where each $\boldsymbol{v}_i$ follows a distribution $F_i = \prod_{j=1}^m F_{i,j}$ supported on the valuation space $V_i = \prod_{j=1}^m V_{i,j} \subseteq \mathbb{R}^m_{\geq 0}$. Consequently, the joint valuation profile $\boldsymbol{v}$ is drawn from the product distribution $\mathcal{F} = \prod_{i=1}^n F_i$ over the joint space $\mathcal{V} = \prod_{i=1}^n V_i$. Since valuations are private, each bidder will instead submit a bid vector $\boldsymbol{b_i} \in V_i$ reporting values for the m items during the auction. The complete bid profile is $\boldsymbol{b} = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_n)$, which may differ from the true valuation profile $\boldsymbol{v}$.

An auction mechanism is specified by an allocation rule $\boldsymbol{g} = (\boldsymbol{g}_1, \dots, \boldsymbol{g}_n) : \mathcal{V} \to [0,1]^{n \times m}$ and a payment rule $\boldsymbol{p} : \mathcal{V} \to \mathbb{R}^n_{\geq 0}$. Given a bid profile $\boldsymbol{b}$, the allocation $\boldsymbol{g}(\boldsymbol{b})$ represents the probability of assigns items to bidders, and the payment vector $\boldsymbol{p}(\boldsymbol{b})$ determines the amounts charged to each bidder. To ensure the feasibility of the allocation, the rule $\boldsymbol{g}$ must satisfy the constraint: $\sum_{i=1}^n g_{i,j}(\boldsymbol{b}) \leq 1, \quad \forall j \in M, \forall \boldsymbol{b} \in \mathcal{V}$, where $g_{i,j}(\boldsymbol{b})$ denotes the probability item j allocated to bidder i with the submitted bid profile $\boldsymbol{b}$.

Bidder i's utility under bid profile $\boldsymbol{b}$ and true valuation $\boldsymbol{v}_i$ is $u_i(\boldsymbol{v}_i; \boldsymbol{b}) = \boldsymbol{v}_i(\boldsymbol{g}_i(\boldsymbol{b})) - p_i(\boldsymbol{b}) = \sum_{j=1}^m v_{i,j} g_{i,j}(\boldsymbol{b}) - p_i(\boldsymbol{b})$, where $g_{i,j}(\boldsymbol{b})$ denotes the probability of item j being allocated to bidder i given the bid profile $\boldsymbol{b}$, and $p_i(\boldsymbol{b})$ is the corresponding payment charged to bidder i.

Definition 3.1 (Individual Regret)**.** For bidder i, given true valuation $\boldsymbol{v}_i$, the individual regret is defined as

$$R_I(\boldsymbol{v}_i) = \max_{\boldsymbol{b}'_i \in V_i} \Big[u_i(\boldsymbol{v}_i; (\boldsymbol{b}'_i, \boldsymbol{b}_{-i})) - u_i(\boldsymbol{v}_i; (\boldsymbol{v}_i, \boldsymbol{b}_{-i})) \Big]. \tag{1}$$

This quantity measures the maximum utility gain bidder i can obtain by unilaterally deviating from truthful reporting.

A mechanism is *dominant strategy incentive compatible* (DSIC) if truthful reporting maximizes each bidder's utility regardless of the bids submitted by others. Formally,

$$u_i(\boldsymbol{v}_i; (\boldsymbol{v}_i, \boldsymbol{b}_{-i})) \geq u_i(\boldsymbol{v}_i; (\boldsymbol{b}_i, \boldsymbol{b}_{-i})), \forall i \in N, \boldsymbol{v}_i, \boldsymbol{b}_i \in V_i. \tag{2}$$

A mechanism is *ex post individually rational* (IR) if truthful reporting guarantees bidders non-negative utility:

$$u_i(\boldsymbol{v}_i; (\boldsymbol{v}_i, \boldsymbol{b}_{-i})) \geq 0, \forall i \in N, \boldsymbol{v}_i \in V_i, \boldsymbol{b}_{-i} \in V_{-i}. \tag{3}$$

For a bid profile $\boldsymbol{b} = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_n)$, the payment rule $\boldsymbol{p}(\boldsymbol{b})$ determines the amount charged to each bidder, and the total revenue collected is $\text{Rev}(\boldsymbol{p}; \boldsymbol{b}) = \sum_{i=1}^n p_i(\boldsymbol{b})$.

The primary objective of optimal auction design is to identify an allocation rule $\boldsymbol{g}$ and a payment rule $\boldsymbol{p}$ that satisfy both DSIC and IR constraints while maximizing the expected revenue over the joint valuation distribution $\mathcal{F}$: $\max_{\boldsymbol{g}, \boldsymbol{p}} \mathbb{E}_{\boldsymbol{v} \sim \mathcal{F}} \left[\sum_{i=1}^n p_i(\boldsymbol{v}) \right]$, subject to the constraints of feasibility, DSIC, and IR.

3.2 COLLUSIONS IN AUCTION

Let $C \subseteq N$ denote a non-empty coalition of bidders, and let $\boldsymbol{b}_C$ and $\boldsymbol{v}_C$ denote the reported and true valuations of bidders in C, respectively. Similarly, write $\boldsymbol{b}_{-C}$ for the reported bids of bidders

outside the coalition. Given an allocation and payment rule $(\boldsymbol{g}, \boldsymbol{p})$ and a bid profile $(\boldsymbol{b}_C, \boldsymbol{b}_{-C})$, the total utility obtained by coalition C is $U_C(\boldsymbol{v}_C; \boldsymbol{b}) = \sum_{i \in C} \left[\boldsymbol{v}_i(\boldsymbol{g}_i(\boldsymbol{b}_C, \boldsymbol{b}_{-C})) - p_i(\boldsymbol{b}_C, \boldsymbol{b}_{-C}) \right]$.

Unlike classical GSP, we adopt a more general formulation to characterize coalition manipulation. Specifically, a coalition C engages in a coordinated deviation when its members jointly choose a misreport profile $\tilde{\boldsymbol{b}}_C$ that differs from $\boldsymbol{v}_C$ and is intended to increase the collective utility of the group. For a DSIC mechanism, individual bidders have no incentive to deviate on their own, but this guarantee does not extend to joint deviations. Formally, a mechanism is vulnerable to coalition manipulation if there exist $C \subseteq N$, true valuations $(\boldsymbol{v}_C, \boldsymbol{v}_{-C})$, and a coordinated misreport $\tilde{\boldsymbol{b}}_C$ such that

$$U_C(\boldsymbol{v}_C; (\tilde{\boldsymbol{b}}_C, \boldsymbol{b}_{-C})) > U_C(\boldsymbol{v}_C; (\boldsymbol{v}_C, \boldsymbol{b}_{-C})). \tag{4}$$

Definition 3.2 (Coalition Regret)**.** For a coalition $C \subseteq N$, given true valuations $\boldsymbol{v}_C$, the coalition regret is defined as

$$R_C(\boldsymbol{v}_C) = \max_{\tilde{\boldsymbol{b}}_C \in V_C} \Big[U_C(\boldsymbol{v}_C; (\tilde{\boldsymbol{b}}_C, \boldsymbol{b}_{-C})) \; - U_C(\boldsymbol{v}_C; (\boldsymbol{v}_C, \boldsymbol{b}_{-C})) \Big]. \tag{5}$$

The goal for collusion-resistant mechanism design is to design mechanisms that minimize $\bar{R}_C$ for all coalitions, thereby achieving robustness against coordinated strategic manipulation.

4 Collusion-Resistant Mechanism Design

4.1 Challenges in Extending RegretNet

To counter coordinated strategic behaviors among bidders, a natural starting point is to extend the adversarial learning framework of RegretNet (Dütting et al., 2019) from individual agents to bidder coalitions. In this formulation, an adversarial collusion network will be trained to discover joint misreports that maximize the coalition utility gain. However, directly optimizing this joint formulation encounters two critical challenges:

Combinatorial Search Space Searching over all $2^n - 1$ possible non-empty subsets of bidders coalitions induces an exponential combinatorial bottleneck, rendering discrete coalition enumeration computationally intractable in multi-bidder auctions.

Gradient Vanishing in Joint Space Effective collusion usually requires multiple bidders to reduce their bids simultaneously. Because of this, the partial derivative of collective surplus with respect to a single bidder's deviation is often close to zero, causing gradient-based training to converge slowly and fail to detect the revenue drop caused by collusion.

To overcome these challenges, we introduce **CoalitionNet**, a learning framework that exploits utility externalities to discover and suppress profitable coalitional deviations.

4.2 CoalitionNet

We instantiate CoalitionNet via an adversarial learning process between a mechanism network and a collusion generator, detailed formulations and derivations are provided in Appendix C.

Mechanism Network The auction mechanism $(\boldsymbol{g}, \boldsymbol{p})$ maps submitted bid profiles $\boldsymbol{b} \in [0, 1]^{n \times m}$ to allocation probabilities and payments via a Transformer, $\boldsymbol{H}_{\boldsymbol{\theta}} = \text{Transformer}(\boldsymbol{b}) \in \mathbb{R}^{n \times d}$. The bidder dimension is encoded as tokens, and the item dimension is encoded as features, which allows the network to take an unfixed number of bidders as input. The mechanism network consists of two components: an allocation network and a payment network, which are used to generate the allocation and payment rules, respectively.

The *allocation network* first produces unnormalized logits $\boldsymbol{Z} \in \mathbb{R}^{n \times m}$ and appends a learnable dummy-bidder logit $\boldsymbol{w}_{\text{dummy}} \in \mathbb{R}^{1 \times m}$ for unallocated items. The resulting logits then go through a column-wise softmax function, which strictly guarantees allocation feasibility and transforms the unnormalized logits into normalized allocation probabilities:

$$g_{i,j}(\boldsymbol{b}) = \frac{\exp(\tilde{Z}_{i,j})}{\sum_{k=1}^{n+1} \exp(\tilde{Z}_{k,j})}, \tilde{\boldsymbol{Z}} = [\boldsymbol{Z}; \boldsymbol{w}_{\text{dummy}}]. \tag{6}$$

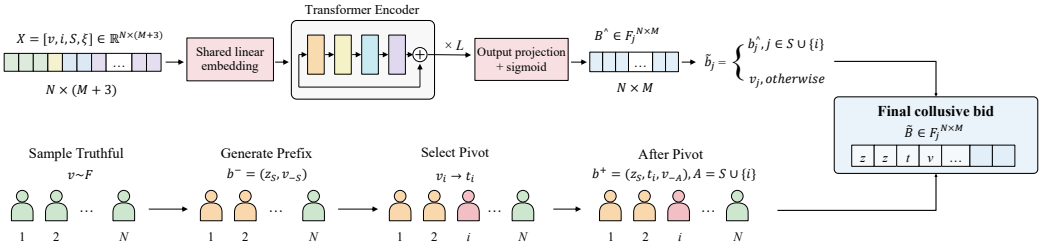

Figure 1: Overview of the collusion generator. Conditioned on truthful valuations $\boldsymbol{v}$, a deviating prefix S, and pivot bidder i, it generates collusive bids as adversary for the mechanism training.

The *payment network* first outputs an unnormalized payment vector $\tilde{\boldsymbol{p}} \in \mathbb{R}^{n\times 1}$, which is subsequently normalized into a fraction vector $\boldsymbol{p}_{\text{frac}}(\boldsymbol{b}) \in (0,1)^{n\times 1}$ via a sigmoid activation to guarantee IR. The final payment rule is

$$p_i(\boldsymbol{b}) = p_{\text{frac},i}(\boldsymbol{b}) \sum_{j=1}^{m} g_{i,j}(\boldsymbol{b}) b_{i,j}. \tag{7}$$

Collusion Generator Figure 1 illustrates the collusion generation process, where the coalition formation is modeled as a bidder transition. For each valuation profile $\boldsymbol{v}$, we sample a *pivot bidder* i and a *prefix group* $S \subseteq N \setminus \{i\}$, which is treated as an already deviated group. The active deviators are $A = S \cup \{i\}$. Conditioned on $(\boldsymbol{v}, i, S, \boldsymbol{\xi})$, where $\boldsymbol{\xi}$ is the exploration coefficient, the collusion generator f_ω proposes misreports $\boldsymbol{z}_S$ for the prefix group and $\boldsymbol{t}_i$ for the pivot bidder.

Generating a strong collusion adversary is equivalent to find a pivot bidder whose transition brings a high utility gain to the potential coalition groups. We measure the effect of the pivot's change by $\boldsymbol{d} = u(\boldsymbol{v}; \boldsymbol{b}^+) - u(\boldsymbol{v}; \boldsymbol{b}^-)$, and the total gain over truthful bidding by $\boldsymbol{w} = u(\boldsymbol{v}; \boldsymbol{b}^+) - u(\boldsymbol{v}; \boldsymbol{v})$, where $\boldsymbol{b}^- = (\boldsymbol{z}_S, \boldsymbol{v}_{-S})$ and $\boldsymbol{b}^+ = (\boldsymbol{z}_S, \boldsymbol{t}_i, \boldsymbol{v}_{-(S\cup\{i\})})$. We interpret $\boldsymbol{d}$ as the *utility-externality profile* induced by pivot bidder i. Each component d_j measures how bidder j's utility changes when i joins the deviation while the prefix S is held fixed; for $j \neq i$, it captures the utility externality that the pivot imposes on bidder j. Positive externalities therefore identify additional bidders who benefit from the pivot transition and can join the coalition. To measure the largest utility gain to a coalition containing the active deviators A, we define F_A as:

$$F_A(\Delta \boldsymbol{u}) = \sum_{j\in A} \Delta u_j + \sum_{j\notin A} [\Delta u_j]_+ = \max_{C \supseteq A} \sum_{j\in C} \Delta u_j, \tag{8}$$

where $[x]_+ = \max\{0, x\}$. In this way, a coalition with the largest aggregate utility gain is obtained analytically by adding every outside bidder with $\Delta u_j > 0$ to A, *without enumerating candidate coalitions*. However, a positive gain from the pivot's transition may only recover a loss incurred by the prefix. We therefore compare $F_A(\boldsymbol{d})$ with $F_A(\boldsymbol{w})$:

$$\Phi_A(\boldsymbol{d}, \boldsymbol{w}) = \min\{F_A(\boldsymbol{d}), F_A(\boldsymbol{w})\}. \tag{9}$$

Here, Φ_A measures the coalition gain from pivot bidder i joining the deviation after a sampled prefix S. By explicitly conditioning the deviation on this prefix, Φ_A captures local utility changes undetectable via unilateral deviations, thereby *alleviating gradient vanishing in the joint deviation space*. We use the Φ_A to search for collusive bids and to evaluate the resulting adversary. We train f_ω to maximize Φ_A:

$$\max_{\omega} \ \mathbb{E}_{\boldsymbol{v},S,i,\boldsymbol{\xi}} \left[\Phi_{S\cup\{i\}} \left(\boldsymbol{d}_\omega, \boldsymbol{w}_\omega \right) \right], \tag{10}$$

where $\boldsymbol{d}_\omega$ and $\boldsymbol{w}_\omega$ are induced by the generator's reports. At each adversarial update, we initialize the bid profile from the collusion generator's output and refine it with T steps of projected gradient ascent using Adam: $\tilde{\boldsymbol{b}}^{(t+1)} = \Pi_{\mathcal{B}_A} \left(\tilde{\boldsymbol{b}}^{(t)} + \alpha \, \text{Adam} \left(\nabla_{\tilde{\boldsymbol{b}}^{(t)}} \Phi_A(\tilde{\boldsymbol{b}}^{(t)}) \right) \right), t = 0, \ldots, T-1$, where $\mathcal{B}_A = \{\boldsymbol{b} \in [0,1]^{n\times m} : \boldsymbol{b}_{-A} = \boldsymbol{v}_{-A}\}$ is the set of valid bid profiles with truthful bids outside A. The refined profile $\tilde{\boldsymbol{b}}^{(T)}$ is then used in the adversarial update of the mechanism network.

4.3 Network Training and Updates

We alternately train the collusion generator and the mechanism network. In each iteration, we sample a batch of valuation profiles. For every profile and pivot bidder i, we consider an empty prefix $S_0 = \emptyset$ and a sampled prefix $S_1 \subseteq N \setminus \{i\}$; S_0 serves to test individual deviations, whereas S_1 exposes the mechanism to coordinated bids.

With the mechanism parameters θ fixed, we update the collusion generator parameters ω using Φ_A defined in Eq. 9:

$$\mathcal{L}_{\text{coll}}(\omega;\theta) = -\mathbb{E}_{\boldsymbol{v},i,S,\boldsymbol{\xi}}\left[\Phi_{S\cup\{i\}}\left(\tilde{\boldsymbol{b}}_\omega;\theta\right)\right], \tag{11}$$

where $\tilde{\boldsymbol{b}}_\omega$ is the generator's proposed bid profile with bids outside $S \cup \{i\}$ held truthful. We then refine the proposal via T steps of gradient ascent (Eq. 4.2), yielding $\tilde{\boldsymbol{b}}^{(T,k,i,S)}$.

For the mechanism update, we hold these bid profiles fixed. For each pivot bidder, we compute the maximum empirical coalition regret induced by $\tilde{\boldsymbol{b}}^{(T,k,i,S)}$ in both individual deviation and coalition manipulation, and average it over the batch:

$$\bar{R}_i(\theta) = \frac{1}{B}\sum_{k=1}^{B} \max_{S\in\{S_0^{(k,i)},S_1^{(k,i)}\}} \left[\Phi_{S\cup\{i\}}\left(\tilde{\boldsymbol{b}}^{(T,k,i,S)};\theta\right)\right]_+. \tag{12}$$

To maximize the revenue of the auctioneer while restricting the profitable space of potential collusions, the loss function for updating the mechanism network is designed as follows:

$$\mathcal{L}_{\text{mech}}(\theta) = -\frac{1}{B}\sum_{k=1}^{B}\sum_{j=1}^{n} p_j(\boldsymbol{v}^{(k)};\theta) + \lambda_t \sum_{i=1}^{n} \bar{R}_i(\theta). \tag{13}$$

Here, λ_t balances collusion resistance against revenue and controls the tradeoff between high revenue and low regret. Algorithm 1 summarizes the alternating updates.

Algorithm 1 CoalitionNet adversarial training

Require: Valuation distribution $\mathcal{F}$, batch size B, attack steps K, penalty weight λ_t.
1: Initialize mechanism parameters θ and generator parameters ω.
2: **for** $t = 1, \ldots, T_{\max}$ **do**
3: Sample $\boldsymbol{X} \sim \mathcal{F}^B$ and form $S_0 = \emptyset$ and a sampled S_1 for every profile and pivot.
4: With θ fixed, update ω using Eq. equation 10.
5: For each context, choose the better generator or random start, run T gradient ascent steps, and retain $\tilde{\boldsymbol{b}}^{(T,k,i,S)}$.
6: With the discovered bids fixed, compute $\bar{R}_i(\theta)$ using Eq. equation 12.
7: Update θ using Eq. equation 13.
8: Update λ_t.
9: **end for**

5 Experiments

The empirical study is designed to evaluate the capabilities of CoalitionNet. We investigate how effectively the mechanism minimizes coalition regret and protects seller revenue against collusive deviations, how it manages the revenue-regret trade-off, and whether this structural robustness holds across different coalition sizes, bidder counts, and heterogeneous valuation environments.

Experimental Setting and Data. We conduct our main experiments in a 5×5 auction environment with five bidders and five items. Bidder valuations are additive, with item values drawn independently from $\mathbf{Uniform(0, 1)}$. During training, we use batches of 512 valuation profiles sampled from this distribution, optimize the models for $10,000$ iterations with a learning rate of 10^{-3}, and employ 25 gradient-refinement steps to search for adversarial deviations.

Baselines. We compare the proposed CoalitionNet with three established baselines: **RegretNet** Dütting et al. (2019): A neural auction model that maximizes revenue while approximately enforcing individual incentive compatibility by penalizing bidder regret. **RegretFormer** Ivanov et al. (2022): An attention-based extension of RegretNet for learning high-revenue auctions with controlled bidder

regret. **AMenuNet** Duan et al. (2023): A neural affine-maximizer model that learns allocation menus and parameters while ensuring dominant-strategy incentive compatibility by construction. Details of model architectures and implementation can be found in Appendix C.

5.1 Regret And Revenue Comparison

To analyze the efficiency and robustness of CoalitionNet, we compare the coalition regret and revenue of four different models shows in Table 1. Among the four models, CoalitionNet achieves the lowest coalition regret, significantly reducing the value of the maximum coalition regret. In terms of revenue, CoalitionNet experiences a drop with truthful bidders, indicating that resistance to collusion comes at a cost to efficiency (see the further analysis in Section 5.3). However, with colluding bidders, CoalitionNet maintains stable revenue with a smaller drop compared with the other models, which are vulnerable to coalition manipulation and suffer larger revenue losses. These results support our adversarial learning framework, which introduces a collusion generator to simulate strong adversarial behavior and thus enhances mechanism robustness without hurting revenue performance too much.

Table 1: Regret and revenue comparison among different models. Results are reported as mean $\pm$ standard deviation across three training seeds.

Method	Avg. coal. regret ↓	Max coal. regret ↓	Revenue ↑	Attacked revenue ↑	Drop ↓
CoalitionNet	0.01 ± 0.00	0.01 ± 0.00	2.97 ± 0.02	2.75 ± 0.02	0.22 ± 0.00
RegretNet	0.26 ± 0.00	2.07 ± 0.03	3.18 ± 0.03	0.38 ± 0.01	2.80 ± 0.03
RegretFormer	0.25 ± 0.01	2.19 ± 0.06	3.26 ± 0.03	0.32 ± 0.15	2.93 ± 0.12
AMenuNet	0.23 ± 0.00	2.60 ± 0.04	3.36 ± 0.02	0.09 ± 0.02	3.27 ± 0.03

Additionally, Figure 2(a) shows the distribution of maximum coalition regret for all four models. CoalitionNet reduces this regret by nearly 100 times compared with the other models, which do not account for collusion during training. Figure 2(b) shows the seller revenue with truthful bidders. Interestingly, CoalitionNet exhibits a different distribution pattern from the other models. Its average revenue is close to that of the other models, and most revenue values are distributed similarly, while a few samples show relatively low revenue. The mechanism potentially treats these as high-risk valuation samples because they are more likely to involve collusion and harm allocation efficiency. In these cases, the mechanism restricts allocation by lowering the allocation probability, remaining aware of potential risks. This can be viewed as an *implicit reserve price* learned by the mechanism network during training. Reserve prices are a key method for avoiding collusion in practical auctions and auction research. These findings also suggest that CoalitionNet inherently learns strategies to deal with collusion, rather than merely overfitting to samples statistically.

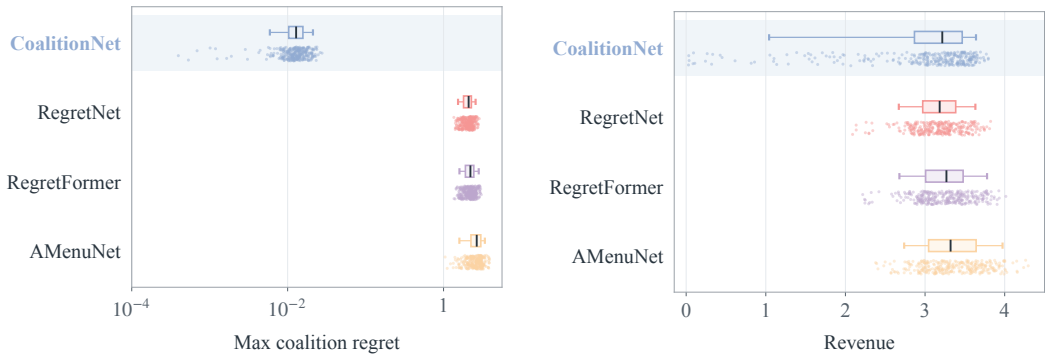

Figure 2: Regret and revenue distributions for randomly sampled valuation profiles.

5.2 Structural Generalization Across Coalitions and Bidders

In practical auction scenarios, the number of bidders and potential coalitions can be larger than those in our experiments. We thus evaluate the structural generalization ability of CoalitionNet against increased numbers of bidders and coalition sizes that were not encountered during training.

Figure 3(a) shows the average coalition regret across different training and testing coalition sizes. The coalition regret is very high when the training size is set to 1, but it reduces rapidly as the training coalition size increases. Notably, even when trained on small coalition sizes, CoalitionNet maintains a significantly low regret for unseen larger coalitions. Figure 3(b) investigates how a change in the number of bidders affects the coalition regret and seller revenue drop without retraining the model. Both the regret and revenue drop remain consistently low as the bidder count increases from 5 to 16. These two findings suggest that CoalitionNet is not confined to smaller training settings and retains its effectiveness and generalization ability for larger coalitions and bidder groups.

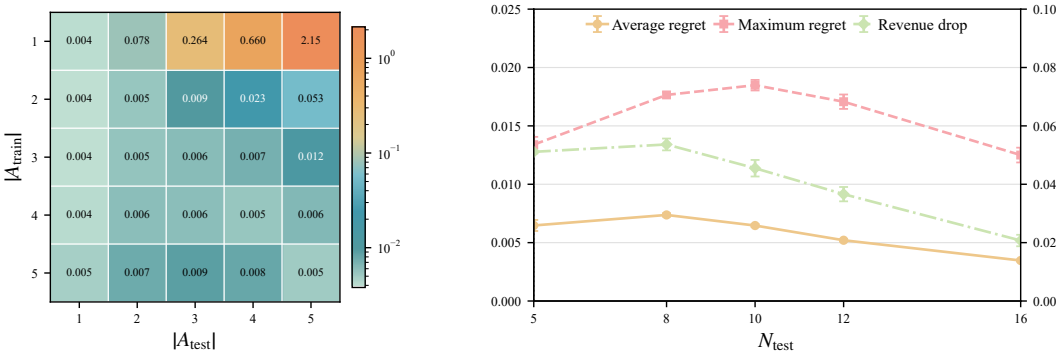

(a) Coalition-size generalization

(b) Zero-shot bidder generalization

Figure 3: Structural generalization of CoalitionNet. **(a)** Average coalition regret across maximum coalition sizes used in training and exact coalition sizes used in testing. **(b)** Coalition regret and seller revenue drop across different numbers of test bidders.

5.3 Empirical Revenue–Regret Frontier/Controllable Tradeoff

An inherent tension exists between maximizing seller revenue and ensuring robustness against bidder collusion. To deeply analyze this revenue-regret tradeoff, we trained each of the four models under five different configurations of the regret penalty multiplier λ_t and plotted their resulting performance. It is worth noting that RegretFormer exhibited significant instability during training under certain λ_t settings, easily degenerating into a trivial mechanism that allocates no items. While this trivial state naturally yields zero coalition regret, it also results in zero revenue. Therefore, we removed two such trivial data points for RegretFormer to ensure a fair comparison of meaningful mechanisms.

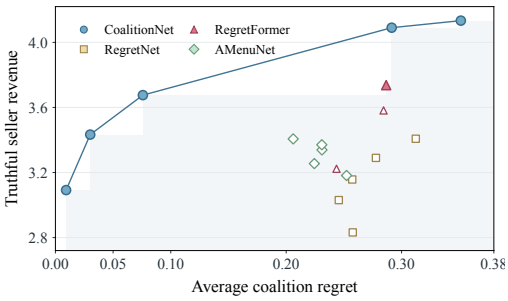

Figure 4: Empirical revenue–regret frontier.

A key observation from Figure 4 is that CoalitionNet forms a clear empirical revenue-regret frontier. By adjusting the hyperparameter λ_t, CoalitionNet provides a highly controllable tradeoff, empowering the mechanism designer to systematically balance efficiency and collusion resistance according to specific deployment needs. In contrast, baseline models fail to map out a predictable frontier; instead, their results appear as a random scatter of performance points. Furthermore, the plot demonstrates that the frontier established by CoalitionNet covers other baselines. As shown by the shaded region in Figure 4, CoalitionNet consistently achieves higher revenue at any given level of coalition regret, and incurs lower regret for any target revenue.

5.4 Distributional Heterogeneity

Bidder valuations distributions varies in real auction markets. Here we examines whether CoalitionNet can control coalition regret under distributional variation and asymmetric bidder valuations. We use the same auction setting as 5×5 and consider four different distributions: $\mathbf{Uniform(0, 1)}$, $\mathbf{Beta(2, 5)}$,

$\mathbf{Beta(5, 2)}$, and a bidder-asymmetric setting in which bidders 1–2 follow $\mathbf{Beta(2, 5)}$ while bidders 3–5 follow $\mathbf{Beta(5, 2)}$.

Table 2 shows that the average coalition regret remains low across all four valuation settings, while the revenue drop between truthful and collusive scenarios is moderate. However, the revenue drop is relatively larger in the $\mathbf{Beta(5, 2)}$ setting compared to the other settings. This likely arises because the distribution's concentration on high valuations creates greater payment exposure, allowing coordinated misreports to reduce seller revenue more substantially. Its higher truthful revenue also increases the space of the revenue loss. These results suggest that CoalitionNet can be trained effectively under different symmetric and asymmetric valuation distributions to handle bidder heterogeneity.

Table 2: Regret and revenue under different value distributions.

Distribution	Avg. coal. regret ↓	Max coal. regret ↓	Revenue ↑	Attacked revenue ↑	Drop ↓
$\mathrm{Uniform}(0, 1)$	0.01 ± 0.00	0.01 ± 0.00	2.96 ± 0.05	2.72 ± 0.05	0.25 ± 0.01
$\mathrm{Beta}(2, 5)$	0.01 ± 0.00	0.01 ± 0.00	1.65 ± 0.02	1.51 ± 0.03	0.15 ± 0.01
$\mathrm{Beta}(5, 2)$	0.01 ± 0.01	0.03 ± 0.00	3.75 ± 0.03	3.00 ± 0.06	0.75 ± 0.05
Bidder-asymmetric	0.01 ± 0.00	0.02 ± 0.00	3.55 ± 0.04	3.21 ± 0.06	0.33 ± 0.03

5.5 Learned Attack Efficacy for Collusion Generator

Low measured regret is only meaningful if the training adversary can expose strong collusive deviations. Otherwise, the apparent robustness of the learned mechanism could simply result from a weak collusion generator. We rule out this possibility by evaluating the collusion generator independently. We construct a set of feasible collusive bids through random sampling and independent gradient-based search over sampled valuation profiles, and rank the generator's proposal against these candidates. Figure 5 shows that the proposal alone outperforms 98.82% of the reference candidates on average, placing it close to the strongest attacks found by the search.

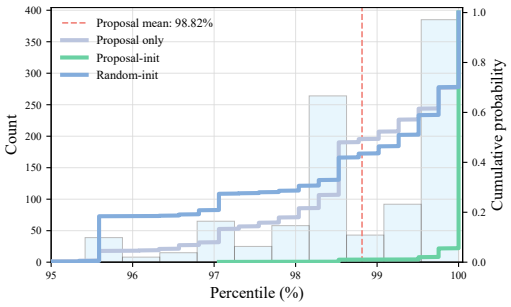

Figure 5: Learned attack efficacy.

We further compare three attack-generation strategies: the generator proposal without refinement, refinement from a random initialization, and refinement initialized from the generator proposal. As shown in Figure 5, the proposal-initialized strategy produces the strongest attacks, with its distribution concentrated near the top of the feasible collusive bids. This result indicates that the collusion generator does more than simply produce group manipulations: it identifies strong collusive bids and provides effective starting points for subsequent optimization.

6 Conclusion

In this paper, we introduce CoalitionNet, a collusion-resistant auction design framework based on coalition regret minimization. To overcome the challenges of gradient vanishing and combinatorial complexity in joint deviation spaces, our method introduces utility externalities to evaluate coalitional deviations. This externality-guided approach enables the training of a strong collusion generator, which effectively identifies worst-case joint misreports for mechanism adversarial learning. Empirical evaluations demonstrate that CoalitionNet minimizes coalition regret and protects seller revenue against collusion. Furthermore, the proposed framework exhibits decent generalization capabilities, maintaining its structural robustness across unseen larger coalition sizes and varying bidder counts, as well as adapting to heterogeneous valuation distributions. Finally, CoalitionNet establishes a clear empirical revenue-regret frontier, providing mechanism designers with a highly controllable tradeoff to systematically balance economic efficiency and collusion resistance.

AI USE STATEMENT

In this work, we used generative AI tools for polishing writing and assisting in the writing of proofs, and we manually checked all generated implementations for logical and mathematical correctness.

We have not used generative AI tools for generating synthetic data sets, helping develop theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, assisting in the writing of proofs, proposing or refining hypotheses, designing or providing feedback on research methodology or experiments, interpreting results, implementing methods, assisting with translation, cleaning and reformatting datasets, and supporting qualitative and thematic data analysis are not applicable to this work.

Additionally, we used generative AI tools for creating or editing software code and editing the research paper to improve readability.

We have reviewed all AI-assisted work. Specifically, all LLM-generated code was rigorously tested, debugged, and verified for mathematical and logical correctness by the authors. Furthermore, all AI-polished text was manually proofread to ensure it accurately preserves our original scientific intent, nuance, and meaning. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

Our research focuses on the theoretical and computational aspects of automated mechanism design. We propose **CoalitionNet**, a framework specifically designed to mitigate bidder collusion in auction environments. We do not foresee any significant ethical risks associated with this work.

REPRODUCIBILITY STATEMENT

We are highly committed to ensuring the reproducibility of our research. To facilitate this, we have provided comprehensive model training details for **CoalitionNet**. Specifically, the complete descriptions of our neural network architectures, optimization objectives, hyperparameter configurations, and the synthetic valuation generation process are thoroughly documented in the experimental section and the Appendix.

REFERENCES

John Asker. A study of the internal organization of a bidding cartel. *American Economic Review*, 100(3):724–762, 2010.

Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of automated mechanism design. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 2091–2099, Red Hook, NY, USA, 2016. Curran Associates Inc.

Laura H Baldwin, Robert C Marshall, and Jean-Francois Richard. Bidder collusion at forest service timber sales. *Journal of Political Economy*, 105(4):657–699, 1997.

Tilman Börgers, Daniel Krähmer, and Roland Strausz. *An Introduction to the Theory of Mechanism Design*. Oxford University Press, 2015.

Yeon-Koo Che and Jinwoo Kim. Optimal collusion-proof auctions. *Journal of Economic Theory*, 144(2):565–603, 2009.

Peter Cramton and Jesse A Schwartz. Collusive bidding in the fcc spectrum auctions. 2002.

Michael Curry, Tuomas Sandholm, and John Dickerson. Differentiable economics for randomized affine maximizer auctions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 2633–2641, 2023.

Michael Curry, Vinzenz Thoma, Darshan Chakrabarti, Stephen McAleer, Christian Kroer, Tuomas Sandholm, Niao He, and Sven Seuken. Automated design of affine maximizer mechanisms in dynamic settings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 9626–9635, 2024.

Francesco Decarolis, Maris Goldmanis, Antonio Penta, and Ksenia Shakhgildyan. Bid coordination in sponsored search auctions: detection methodology and empirical analysis. *The Journal of Industrial Economics*, 71(2):570–592, 2023.

Zhijian Duan, Jingwu Tang, Yutong Yin, Zhe Feng, Xiang Yan, Manzil Zaheer, and Xiaotie Deng. A context-integrated transformer-based neural network for auction design. In *International Conference on Machine Learning*, pp. 5609–5626. PMLR, 2022.

Zhijian Duan, Haoran Sun, Yurong Chen, and Xiaotie Deng. A scalable neural network for dsic affine maximizer auction design. *Advances in Neural Information Processing Systems*, 36:56169–56185, 2023.

Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. Mechanism design for large language models. In *Proceedings of the ACM Web Conference 2024*, pp. 144–155, 2024.

Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 1706–1715. PMLR, 2019.

Daniel A Graham and Robert C Marshall. Collusive bidder behavior at single-object second-price and english auctions. *Journal of Political economy*, 95(6):1217–1239, 1987.

Mingyu Guo and Vincent Conitzer. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, pp. 1676–1679, 2010.

Dmitry Ivanov, Iskander Safiulin, Igor Filippov, and Ksenia Balabaeva. Optimal-er auctions through attention. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2022. Curran Associates Inc.

Yanchen Jiang, David C Parkes, and Tonghan Wang. Duality for optimal multi-item, multi-bidder auction design: Revenue certificates through deep learning. *arXiv preprint arXiv:2606.10112*, 2026.

Michael I Jordan and Tom M Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.

Paul Klemperer. Auction theory: A guide to the literature. *Journal of economic surveys*, 13(3): 227–286, 1999.

Jean-Jacques Laffont and Jean Tirole. *A theory of incentives in procurement and regulation*. MIT press, 1993.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

Yang Li, Yuchao Ma, and Qi Qi. Optimal auction design in the joint advertising. *arXiv preprint arXiv:2507.07418*, 2025.

Jiayin Liu and Chenglong Zhang. Deep learning for double auction. *arXiv preprint arXiv:2504.05355*, 2025.

Robert C. Marshall and Leslie M. Marx. The vulnerability of auctions to bidder collusion. *The Quarterly Journal of Economics*, 124(2):883–910, 05 2009.

Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pp. 1089–1122, 1982.

Mai Pham, Vikrant Vaze, and Peter Chin. Advancing differentiable economics: A neural network framework for revenue-maximizing combinatorial auction mechanisms. *arXiv preprint arXiv:2501.19219*, 2025.

Sai Srivatsa Ravindranath, Zhe Feng, Di Wang, Manzil Zaheer, Aranyak Mehta, and David C Parkes. Deep reinforcement learning for sequential combinatorial auctions. *arXiv preprint arXiv:2407.08022*, 2024.

Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.

Tuomas Sandholm. Automated mechanism design: A new application area for search algorithms. In *International Conference on Principles and Practice of Constraint Programming*, pp. 19–36. Springer, 2003.

Weiran Shen, Pingzhong Tang, and Song Zuo. Automated mechanism design via neural networks. *arXiv preprint arXiv:1805.03382*, 2018.

Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.

Rakesh V Vohra. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press, 2011.

Tonghan Wang, Yanchen Jiang, and David C Parkes. Gemnet: Menu-based, strategy-proof multi-bidder auctions through deep learning. *arXiv preprint arXiv:2406.07428*, 2024.

Tonghan Wang, Yanchen Jiang, and David C Parkes. Bundleflow: Deep menus for combinatorial auctions by diffusion-based optimization. *arXiv preprint arXiv:2502.15283*, 2025.

Jakob Weissteiner and Sven Seuken. Deep learning—powered iterative combinatorial auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 2284–2293, 2020.

Xiaolong Xu, Yuhao Hu, Guangming Cui, Lianyong Qi, Wanchun Dou, and Zhipeng Cai. Cadec: A combinatorial auction for dynamic distributed dnn inference scheduling in edge-cloud networks. *IEEE Transactions on Mobile Computing*, 2025.

Jiguo Yu, Shun Liu, Yifei Zou, Guijuan Wang, and Chunqiang Hu. Auction theory and game theory based pricing of edge computing resources: A survey. *IEEE Internet of Things Journal*, 2025.

Mengqian Zhang, Yuhao Li, Xinyuan Sun, Elynn Chen, and Xi Chen. Maximal extractable value in batch auctions. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pp. 510–510, 2025.

A Extended Literature Review

A.1 Optimal Mechanism Design via Affine Maximizers

Beyond the line of research focused on *regret minimization* frameworks and deep learning-based automated mechanism design (e.g., RegretNet), another trajectory in the field centers on the structural characterization of incentive-compatible mechanisms. This alternative line of research seeks to identify the functional forms that optimal mechanisms must take, with **Affine Maximizers (AMs)** serving as the cornerstone of this approach.

The class of AMs represents a critical generalization of the VCG mechanism. A mechanism is an affine maximizer if it chooses an allocation x that maximizes a weighted sum of bidder valuations plus some constant per-allocation "bonus":

$$x^* \in \arg\max_{x \in \mathcal{X}} \left(\sum_{i=1}^{n} w_i \cdot v_i(x) + c_x \right) \tag{14}$$

where $w_i \geq 0$ are bidder weights and $c_x \in \mathbb{R}$ are constants for each possible allocation $x \in \mathcal{X}$. The importance of AMs is rooted in Roberts' Theorem, which states that if the bidders' valuation space is $V = \mathbb{R}^A$ and there are at least three possible outcomes, then every onto DSIC social choice function must be an affine maximizer. This theorem suggests that when searching for optimal mechanisms within the DSIC class, the space of AMs is a natural and often exhaustive candidate.

Despite their theoretical elegance, characterizing the specific weights $\{w_i\}$ and constants $\{c_x\}$ that yield the optimal revenue is analytically intractable for most complex distributions. Classical literature has explored "Virtual Valuations" to reduce revenue maximization to welfare maximization (e.g., Myerson's auction), but this approach does not scale easily to multi-item settings where the ironed virtual valuations become multi-dimensional.

In recent years, research has sought to combine the structural robustness of Affine Maximizer Auctions (AMAs) with the optimization power of deep learning, giving rise to the field of *differentiable economics*. Unlike RegretNet, which learns arbitrary allocation rules, these works restrict the mechanism's functional form to the AMA class to ensure *exact* DSIC and IR properties by construction, while using neural networks to optimize the previously "intractable" parameters.

Specifically, **AMenuNet** (Duan et al., 2023) and related architectures have been proposed to learn the weights w_i and boosts c_x directly from bidder and item representations. By parameterizing the set of candidate allocations (the "menu") and using smoothed maximization (e.g., via Softmax or smoothed linear programming), these models can be trained end-to-end using gradient-based methods. Another significant contribution is the exploration of **Randomized Affine Maximizers** (Curry et al., 2023), which extends the search space to include lottery allocations. This approach generalizes single-bidder architectures like RochetNet (Dütting et al., 2019) and MenuNet (Shen et al., 2018) to multi-bidder settings, demonstrating that even with a restricted AMA structure, learned mechanisms can outperform classical baselines and achieve revenue levels competitive with RegretNet, without sacrificing strategyproofness. Furthermore, recent extensions have moved into **Dynamic AMA** settings (Curry et al., 2024), applying reinforcement learning and zeroth-order optimization to design truthful mechanisms for sequential auctions where agents arrive and depart over time.

While the aforementioned advancements successfully integrate individual IC into the AMA learning framework, extending these models to accommodate coalition-level IC remains a formidable and promising frontier. In many real-world auction environments, bidders may strategically coordinate their reports to improve their collective utility, a challenge that standard DSIC mechanisms often fail to address. Incorporating CIC constraints into affine maximizers, however, introduces a significant leap in *combinatorial dimension complexity*. Specifically, preventing coalitional deviations requires enforcing constraints across the power set of bidders, which exacerbates the already demanding task of optimizing over an exponential allocation space $\mathcal{X}$. Although recent work has begun to explore group strategyproofness in simpler settings, a scalable approach that maintains the revenue-optimizing power of neural AMAs while being robust to coalitional behavior is still lacking. Developing sophisticated neural architectures or relaxation techniques that can navigate this dimensionality bottleneck represents a critical direction for future research in automated mechanism design.

B GENERALIZATION BOUND FOR COALITION REGRET

The collusion generator optimizes a pivot transition, whereas coalition regret measures the gain from an entire joint deviation. We derive a bound connecting these two quantities to clarify whether the externality-guided objective controls the incentive for a coalition to deviate.

Proposition B.1 (Externality objective and coalition regret)**.** *Fix a valuation profile $\boldsymbol{v}$, and define worst-case coalition regret, with bidders outside the coalition reporting truthfully, by*

$$\mathcal{R}^*(\boldsymbol{v}) = \sup_{\varnothing \neq C \subseteq N} \sup_{\tilde{\boldsymbol{b}}_C \in V_C} \Big[U_C\big(\boldsymbol{v}_C; (\tilde{\boldsymbol{b}}_C, \boldsymbol{v}_{-C})\big) - U_C(\boldsymbol{v}_C; \boldsymbol{v}) \Big]_+ .$$

For each pivot bidder $i \in N$, let

$$\epsilon_i(\boldsymbol{v}) = \sup_{\substack{S \subseteq N \setminus \{i\} \\ \boldsymbol{z}_S, \boldsymbol{t}_i \text{ feasible}}} \big[\Phi_{S\cup\{i\}}(\boldsymbol{d}, \boldsymbol{w})\big]_+ , \qquad B(\boldsymbol{v}) = \sum_{i \in N} \epsilon_i(\boldsymbol{v}),$$

where $\mathbf{d}$ and $\mathbf{w}$ are induced by the pre-transition and post-transition profiles defined in Section 4.2. Then

$$\mathcal{R}^*(\boldsymbol{v}) \le B(\boldsymbol{v}) \le n\, \mathcal{R}^*(\boldsymbol{v}), \qquad n = |N|. \tag{15}$$

Proof. For the upper bound, fix i, a prefix S, and feasible reports, and write $A = S \cup \{i\}$. From the definition of Φ_A,

$$[\Phi_A(\boldsymbol{d}, \boldsymbol{w})]_+ \le [F_A(\boldsymbol{w})]_+ = \max_{C \supseteq A} \left[\sum_{j \in C} w_j \right]_+ .$$

The post-transition profile $\boldsymbol{b}^+$ changes reports only in A. Hence, for every $C \supseteq A$, it is a feasible joint report for C, with members of $C \setminus A$ reporting truthfully. Each term in the maximum is therefore at most $\mathcal{R}^*(\boldsymbol{v})$. Taking the supremum over prefixes and reports gives $\epsilon_i(\boldsymbol{v}) \le \mathcal{R}^*(\boldsymbol{v})$. Summing over $i \in N$ proves $B(\boldsymbol{v}) \le n\mathcal{R}^*(\boldsymbol{v})$.

For the lower bound, fix any nonempty coalition C and any feasible joint report $\tilde{\boldsymbol{b}}_C$. Order its members as $i_1, \dots, i_m$. Set $\boldsymbol{b}^0 = \boldsymbol{v}$, and let $\boldsymbol{b}^t$ use the reports in $\tilde{\boldsymbol{b}}_C$ for $i_1, \dots, i_t$, with all other bidders reporting truthfully. Define

$$A_t = \{i_1, \dots, i_t\}, \qquad G_t = U_C(\boldsymbol{v}_C; \boldsymbol{b}^t) - U_C(\boldsymbol{v}_C; \boldsymbol{v}).$$

Then $G_0 = 0$. At transition t, take pivot i_t and prefix $S_t = A_{t-1}$, so that the pre-transition and post-transition profiles are $\boldsymbol{b}^{t-1}$ and $\boldsymbol{b}^t$. Since $C \supseteq A_t$, the definition of F_{A_t} implies

$$F_{A_t}(\boldsymbol{d}_t) \ge G_t - G_{t-1}, \qquad F_{A_t}(\boldsymbol{w}_t) \ge G_t,$$

where

$$\boldsymbol{d}_t = u(\boldsymbol{v}; \boldsymbol{b}^t) - u(\boldsymbol{v}; \boldsymbol{b}^{t-1}), \qquad \boldsymbol{w}_t = u(\boldsymbol{v}; \boldsymbol{b}^t) - u(\boldsymbol{v}; \boldsymbol{v}).$$

Consequently,

$$\begin{aligned} [\Phi_{A_t}(\boldsymbol{d}_t, \boldsymbol{w}_t)]_+ &\ge [\min\{G_t - G_{t-1}, G_t\}]_+ \\ &\ge [G_t]_+ - [G_{t-1}]_+ . \end{aligned}$$

This transition is feasible in the supremum defining $\epsilon_{i_t}(\boldsymbol{v})$. Summing over the transitions yields

$$\begin{aligned} B(\boldsymbol{v}) &\ge \sum_{t=1}^{m} \epsilon_{i_t}(\boldsymbol{v}) \\ &\ge \sum_{t=1}^{m} [\Phi_{A_t}(\boldsymbol{d}_t, \boldsymbol{w}_t)]_+ \\ &\ge \sum_{t=1}^{m} \big([G_t]_+ - [G_{t-1}]_+\big) = [G_m]_+ . \end{aligned}$$

Taking the supremum over C and $\tilde{\boldsymbol{b}}_C$ proves $\mathcal{R}^*(\boldsymbol{v}) \le B(\boldsymbol{v})$. □

The proposition shows that, when maximized over all feasible prefixes and reports, the positive part of Φ_A yields an upper bound on worst-case coalition regret. It therefore provides a theoretical interpretation of the objective used to search for collusive deviations.

C Implementation Details

To ensure the reproducibility of our results and the fairness of the comparative analysis, this section provides a comprehensive description of the experimental setup, including computing resources, network architectures, baseline configurations, and hyperparameter settings.

- **Hardware:** A workstation equipped with eight NVIDIA GeForce RTX 4090 GPUs.
- **Software:** Algorithms were implemented using PyTorch.
- **Statistical Significance:** We report the mean performance and standard deviation across **3 independent runs** with different random seeds for every task and algorithm.

C.1 CoalitionNet

We implement CoalitionNet using a dual-component architecture comprising a permutation-equivariant Transformer mechanism and a distinct collusion generator. In the mechanism network, each bidder's bid vector is processed as a token through a four-layer Transformer encoder with a hidden dimension of 128, four attention heads, and a feed-forward dimension of 256. To determine the outcomes, an allocation head computes one logit per bidder and item. By appending a learned dummy-bidder logit prior to a bidder-wise softmax, the architecture inherently ensures that each item's allocation probabilities validly sum to at most one across the real bidders. Concurrently, a sigmoid payment head outputs a fractional multiplier $q_i(\boldsymbol{b}) \in (0, 1)$, defining the final payment as $p_i(\boldsymbol{b}) = q_i(\boldsymbol{b}) \sum_j g_{ij}(\boldsymbol{b}) b_{ij}$.

The collusion generator, designed to propose strategic misreports, employs a two-layer Transformer with a hidden dimension of 64, two heads, and a feed-forward dimension of 128. Its input tokens concatenate the true valuation vector, a pivot indicator, a prefix indicator, and a noise input. Based on these, it proposes reports specifically for the pivot and prefix bidders, while assuming all other bids remain truthful.

During training, the process iterates dynamically between the mechanism and the generator. For each valuation profile and pivot bidder, the objective considers both an empty prefix and one sampled prefix. The generator takes a single optimizer update on the signed objective Φ_A per mechanism update. To ensure robust attacks, a random report is also evaluated; the better of the generator's or the random proposal is then used to initialize 25 steps of projected Adam ascent. For the mechanism loss, we take the positive part $[\Phi_A]_+$, computing the maximum across the two sampled contexts for each pivot before summing over pivots and averaging across the batch. In the 5×5 comparison checkpoints, we regulate this using a scalar dynamic penalty λ_t. The penalty controller initializes at 10 (restricted to the range $[0.1, 100]$) and updates every 100 iterations using a separate calibration set of 64 profiles. Optimization for the mechanism is driven by Adam with a cosine learning-rate schedule, whereas the generator relies on Adam with an initial learning rate of 10^{-3}. Figure 6 illustrates the workflow of CoalitionNet.

C.2 RegretNet

In contrast to the attention-based approach, RegretNet implementation relies on two independent multilayer perceptrons (MLPs) for allocations and payments, each featuring two tanh hidden layers of width 100. The allocation network outputs $(n + 1)m$ logits and, similar to CoalitionNet, applies a bidder-wise softmax that incorporates a dummy bidder. Meanwhile, the payment network produces a sigmoid fraction that is multiplied by each bidder's reported allocated value.

The training regimen is structured to maximize truthful revenue subject to approximate individual incentive compatibility. To enforce this, a 25-step projected Adam attack actively searches for each bidder's optimal unilateral misreport. The resulting per-bidder regret is then penalized via an augmented-Lagrangian framework. Within this setup, dual variables are updated every 100 mechanism updates, and the quadratic penalty coefficient is progressively increased every 1,000 updates. Figure 7 illustrates the workflow of RegretNet.

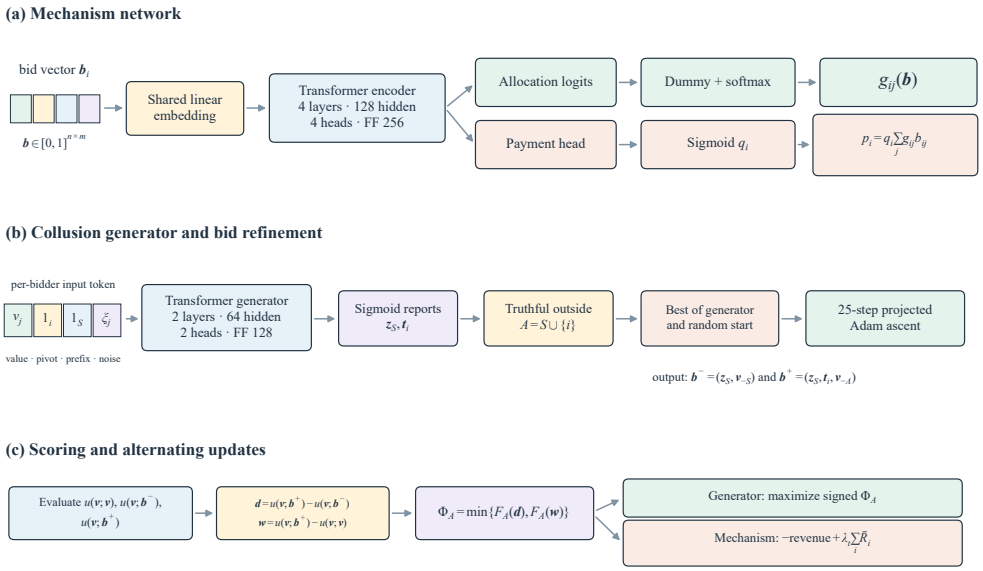

Figure 6: Workflow for CoalitionNet.

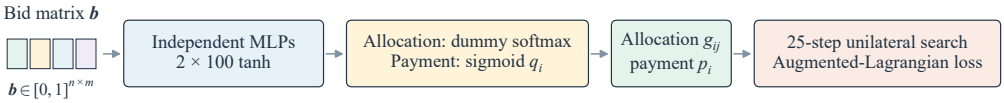

Figure 7: Workflow for RegretNet.

C.3 REGRETFORMER

Compare with RegretNet, RegretFormer bridges the architectural gap by incorporating attention mechanisms into the regret-based framework. It applies an exchangeable input layer to the bid matrix, followed by two bidder/item attention blocks characterized by a hidden width of 128, four heads, and a head width of 32. These attention heads directly output feasible allocation probabilities and payment fractions, with the final payments properly scaled by each bidder's reported allocated value.

While RegretFormer utilizes the same individual-misreport search strategy described for RegretNet, its training objective is distinct. It combines negative truthful revenue with total individual regret, guided by a budget-based weight update. We set the reference budget parameter to 10^{-3}. Figure 8 illustrates the workflow of RegretFormer.

C.4 AMENUNET

Taking a fundamentally different approach, AMenuNet generates an affine-maximizer menu based solely on fixed bidder and item identifiers, making it entirely independent of the current bid profile. The generator constructs this menu using bidder and item embeddings of dimension 10, processed through five two-dimensional Transformer layers with a hidden width of 64 and four attention heads.

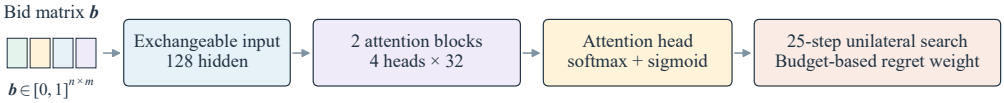

Figure 8: Workflow for RegretFormer.

Table 3: Training settings for the 5×5 uniform-valuation comparison. Attack steps denote projected misreport-ascent steps during training.

Model	Batch size	Adam LR	Attack steps	Updates
CoalitionNet	512	10^{-3}	25	10,000
RegretNet	512	10^{-3}	25	10,000
RegretFormer	512	10^{-3}	25	10,000
AMenuNet	512	10^{-3}	0	10,000

For the 5×5 comparison setting, this network outputs 512 candidate allocations, bidder weights, and allocation boosts, alongside an additional null allocation.

At evaluation, the mechanism acts as a hard selector: it chooses the candidate that maximizes the weighted reported welfare plus its associated boost, and then computes exact affine-maximizer payments by optimizing the score with each bidder systematically removed. During training, however, the model employs differentiable softmax relaxations—using a temperature of 50 for the allocation menu and 500 for the menu selection.

C.5 Training and Codes

Training parameters are listed in Table 3. Codes are available on `https://anonymous.4open.science/r/ICLR_2027_LEARNING-COLLUSION-RESISTANT-AUCTIONS-VIA-COALITION-REGRET-MINIMIZATION_CODE-FB56/`.