

Deep Learning for Revenue-Optimal Auctions with Budgets

Zhe Feng
Harvard University
Cambridge, MA
zhe_feng@g.harvard.edu

Harikrishna Narasimhan
Harvard University
Cambridge, MA
hnrarasimhan@seas.harvard.edu

David C. Parkes
Harvard University
Cambridge, MA
parkes@eecs.harvard.edu

ABSTRACT

The design of revenue-maximizing auctions for settings with private budgets is a hard task. Even the single-item case is not fully understood, and there are no analytical results for optimal, dominant-strategy incentive compatible, two-item auctions. In this work, we model the rules of an auction as a neural network, and use machine learning for the automated design of optimal auctions. We extend the *RegretNet* framework (Dütting et al.'17) to handle private budget constraints, as well as Bayesian incentive compatibility. We discover new auctions with high revenue for multi-unit auctions with private budgets, including problems with unit-demand bidders. For benchmarking purposes, we also demonstrate that *RegretNet* can obtain essentially optimal designs for simpler settings where analytical solutions are available [12, 24, 29].

KEYWORDS

Optimal auction design; Budget constraints; Deep Learning.

ACM Reference Format:

Zhe Feng, Harikrishna Narasimhan, and David C. Parkes. 2018. Deep Learning for Revenue-Optimal Auctions with Budgets. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, July 10–15, 2018, IFAAMAS, 9 pages.

1 INTRODUCTION

The design of revenue-optimal auctions in settings where bidders have private budget constraints is important yet poorly understood problem. Budget constraints arise when bidders have financial constraints that prevent them from making payments as large as their value for items. They are important in many economic settings, including spectrum auctions and land auctions, and are an integral part of the kinds of expressiveness provided to bidders in internet advertising [2, 13].

The design problem is not fully understood even for selling a single item. The technical challenge arises because this is a multi-dimensional mechanism design problem: a bidder's private information is her value for an item as well as her budget. This provides an obstacle to using Myerson's [26] characterization results. Even for selling a single item and with two bidders, the optimal dominant-strategy incentive compatible (DSIC) design with private budget constraints is not known. No revenue-optimal designs are known for selling two or more items to even a single bidder.

In this paper, we build upon the approach of Dütting et al. [16], and use deep neural networks for the automated design of optimal

auctions with budget constraints. We represent an auction as a feed-forward neural network, and optimize its parameters to maximize expected revenue. We need to include design constraints, namely *individual rationality* (IR), *budget constraints* (BC) and *incentive compatibility* (IC).¹ To the best of our knowledge, this is the first paper on automated mechanism design for settings with private budget constraints.

We design both DSIC and Bayesian Incentive Compatible (BIC) auctions. In DSIC auctions, reporting truthfully is the optimal strategy for a bidder no matter what the reports of others. In a BIC auction, truth-telling is the optimal strategy for a bidder in expectation with respect to the types of others, and given that the other bidders report truthfully. The literature has also considered two additional variations in the context of budget constraints: *conditional IC* and *unconditional IC* [12]. We can support both of these within our framework.

1.1 Main Contributions

Our main contributions are summarized below:

- We extend the *RegretNet* framework of Dütting et al. [16] to incorporate budget constraints, as well as, handle BIC and conditional IC constraints. A new aspect is that the utility of an agent can be unbounded in the presence of budgets (whenever an agent's payment exceeds her budget, her utility goes to negative infinity). To handle this, we refine the definition of regret to filter out misreports that would lead to budget violations.
- We show that our approach can be used to design new auctions with high revenue, including for the problem of selling multiple identical items to bidders with additive valuations and selling multiple distinct items to bidders with unit-demand valuations. In both cases, we consider continuous valuation distributions, which is a setting for which the problem cannot be solved through linear programming.
- We benchmark our approach in single-item settings for which analytical solutions exist, showing that neural networks can be used to learn essentially optimal auctions [12, 24, 29].

1.2 Related Work

The high-level approach that we follow is one of *automated mechanism design* (AMD) [14]. Early approaches to AMD involved the use of integer programs, and did not scale up to large settings, or heuristics to search over specialized classes of mechanisms known

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹We consider hard budget constraints for bidders, which means no bidder can pay more than her budget regardless of the bidder's value for the allocation. The literature also considers the case of soft budget constraints, where the bidders are allowed to gain additional funds from markets [22].

to be IC [30]. In recent years, efficient algorithms have been developed for BIC design, but they do not address problems with budget constraints or problems of DSIC design [7–9]

The use of machine learning for AMD was introduced by Dütting et al. [17], who use support vector machines for learning payment rules but not allocation rules, seeking payments that make the resulting mechanism maximally IC. Narasimhan et al. [27] also use support vector machines to learn social choice and matching rules from a restricted class of mechanisms. Narasimhan and Parkes [28] develop a statistical framework for learning assignment mechanisms without providing a computational procedure. Dütting et al. [16] first use deep neural networks for the automated design of optimal auctions. This approach, which we extend in the present paper, is more general, does not require specialized characterization results, and uses off-the-shelf deep learning tools. Very recently, [20] generalize the RegretNet in Dütting et al. [16] for the multi-facility location mechanisms.

Che and Gale [12] design the optimal single-item auction for a single bidder. Pai and Vohra [29] design the optimal BIC auction for a single item and multiple bidders.² Malakhov and Vohra [24] design the optimal auction for a single-item setting with two bidders, but consider a weaker, constrained form of DSIC. Che and Gale [11] develop a revenue ranking of three standard single-item auctions. Maskin [25] and Laffont and Robert [23] consider the problem of bidders with identical, known budgets.

In regard to approximation results: Borgs et al. [6] provide a multi-unit auction for private budget constraints with revenue that converges to the optimal, posted-price auction in the limit of a large population of bidders. Bhattacharya et al. [4] propose a constant approximation for revenue for selling multiple items to additive bidders with private budgets (BIC) and publicly known budgets (DSIC) respectively, adopting an approach that use linear programming relaxations. Chawla et al. [10] propose a multi-item auction with a constant approximation for revenue for bidders with identical, known budgets.

Budget constraints have been handled for the setting of allocative efficiency, with positive results for various multi-item settings, including for bidders with unit-demand valuations [1, 2, 15, 18, 19].³

2 PROBLEM SETUP

In this section, we describe the problem setup, starting with the simpler setting of single-item auctions.

2.1 Single-item auctions

There are n risk neutral bidders interested in a single indivisible good. Each bidder has a private (unknown to other bidders) *value* $v_i \in \mathbb{R}_{\geq 0}$ for the item, and a private *budget* $b_i \in \mathbb{R}_{\geq 0}$ on the amount she can pay. We let $t_i = (v_i, b_i)$ denote the *type* of bidder i and use $t = (t_1, t_2, \dots, t_n)$ to denote a *type profile*. Let \mathcal{T}_i denote the space of possible types for bidder i , and \mathcal{T} the space of type profiles. We assume that bidder i 's type is drawn from distribution F_i , and that F_i is known to both the auctioneer and, in the case of BIC, the other bidders. Let $F = \prod_{i=1}^n F_i$ and $F_{-i} = \prod_{j \neq i} F_j$. Further, let $v_{-i} =$

$(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ denote the valuation profile without v_i , $b_{-i} = (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ denote the budget profile without b_i , and $t_{-i} = (v_{-i}, b_{-i})$.

Each bidder reports (perhaps untruthfully) a value and budget. An auction (a, p) consists of a *randomized allocation rule* $a : \mathcal{T} \rightarrow [0, 1]^n$ and a *payment rule* $p : \mathcal{T} \rightarrow \mathbb{R}_{\geq 0}^n$. Given a reported type profile $t' \in \mathcal{T}$, $a_i(t')$ denotes the probability of bidder i being allocated the item and $\sum_{i=1}^n a_i(t') \leq 1$, and $p_i(t) \in \mathbb{R}_{\geq 0}$ denotes the expected payment by bidder i .⁴

The utility of bidder i with type $t_i = (v_i, b_i)$ for a reported type profile $t' \in \mathcal{T}$ is the difference between the value and payment if the payment is within the budget, and $-\infty$ otherwise:

$$u_i(t_i, t') = \begin{cases} v_i \cdot a_i(t') - p_i(t') & \text{if } p_i(t') \leq b_i, \\ -\infty & \text{if } p_i(t') > b_i. \end{cases} \quad (1)$$

We consider auctions (a, p) that satisfy the *budget constraints* (BC), i.e. charge each agent no more than her budget:

$$\forall i \in [n], t \in \mathcal{T} : p_i(t) \leq b_i \quad (\text{BC})$$

An auction that satisfies these budget constraints is *dominant strategy incentive compatible* (DSIC) if no bidder can strictly improve her utility by misreporting her type, i.e.⁵

$$\forall i \in [n], t \in \mathcal{T}, t'_i \in \mathcal{T}_i : u_i(t_i, (t_i, t_{-i})) \geq u_i(t_i, (t'_i, t_{-i})). \quad (\text{DSIC})$$

The *revenue* from an auction is $\sum_i p_i(t)$. We are interested in designing auctions that maximize expected revenue, while satisfying BC as well as ensuring *ex post individual rationality* (IR), i.e. that each bidder receives non-zero utility for participating:

$$\forall i \in [n], t \in \mathcal{T} : u_i(t_i, (t_i, t_{-i})) \geq 0. \quad (\text{IR})$$

We will also be interested in the design of BIC auctions because this will provide for benchmarking against some known analytical results. In practice, DSIC auctions are more preferred, at least when the effect on achievable revenue relative to BIC designs is small (and there are no other robustness concerns such as those that can arise in DSIC combinatorial auctions [3]), because they are more robust—the equilibrium does not rely on common knowledge of the type distribution or common knowledge of rationality.

For *Bayesian incentive compatibility* (BIC), define the *interim allocation* for bidder i and report t'_i as $\mathcal{A}_i(t'_i) = \mathbf{E}_{t_{-i} \sim F_{-i}}[a_i(t'_i, t_{-i})]$ and the *interim payment* as $\mathcal{P}_i(t'_i) = \mathbf{E}_{t_{-i} \sim F_{-i}}[p_i(t'_i, t_{-i})]$. Given this, we can define the *interim utility function* for a bidder with type t_i and reported type t'_i as:

$$\mathcal{U}_i(t_i, t'_i) = \begin{cases} v_i \mathcal{A}_i(t'_i) - \mathcal{P}_i(t'_i) & \text{if } \mathcal{P}_i(t'_i) \leq b_i, \\ -\infty & \text{if } \mathcal{P}_i(t'_i) > b_i. \end{cases} \quad (2)$$

An auction (a, p) satisfies *interim budget constraints* if

$$\forall i \in [n], t_i \in \mathcal{T}_i : \mathcal{P}_i(t_i) \leq b_i. \quad (\text{interim BC})$$

In addition, an auction satisfying *interim budget constraints* is BIC if:

$$\forall i \in [n], t_i, t'_i \in \mathcal{T}_i : \mathcal{U}_i(t_i, t_i) \geq \mathcal{U}_i(t_i, t'_i) \quad (\text{BIC})$$

Pai and Vohra [29] show that, for any BIC auction that satisfies *interim budget constraints* defined here, there exists an auction with

²They focus on the case of independent values and budgets, but mention that they can handle positive correlation in budget and value.

³The VCG mechanism is not incentive compatible for the budget-constrained case, even when modified in the natural way to truncate valuations by a bidder's budget.

⁴This is equivalent in expectation to charging each agent i a payment $p_i(t')/a_i(t')$ when she wins the auction and 0 otherwise.

⁵This inequality is well-defined for an auction that satisfies the budget constraints.

the same revenue that satisfies BIC for which the largest payment in the support of the interim payment rule is never greater than an agent's reported budget.

We will also insist that auctions that are BIC satisfy the property of *interim individual rationality*:

$$\forall i \in [n], t_i \in \mathcal{T}_i : \mathcal{U}_i(t_i, t_i) \geq 0 \quad (\text{interim IR})$$

There is also a weaker form of both DSIC and BIC, referred to as *conditional incentive compatibility* [12]. Conditional IC assumes that bidders can only underreport their budgets, and thus removes one direction of the incentive constraints. DSIC and BIC become, respectively,

$$\forall i \in [n], t \in \mathcal{T}, t'_i \in \mathcal{T}_i : u_i(t_i, (t_i, t_{-i})) \geq u_i(t_i, (t'_i, t_{-i})) \text{ if } b'_i \leq b_i \quad (\text{C-DSIC})$$

$$\forall i \in [n], t_i, t'_i \in \mathcal{T}_i : \mathcal{U}_i(t_i, t_i) \geq \mathcal{U}_i(t_i, t'_i) \text{ if } b'_i \leq b_i \quad (\text{C-BIC})$$

Conditional IC is motivated by settings in which the auctioneer can require each bidder to post a bond that is equal to her reported budget. Where this is not practical, the more typical, unconditional IC properties are required.

2.2 Multi-item auctions

We also consider a multi-item setting, with both additive and unit-demand valuations on items.

In the additive setting, there are m identical units of an item for sale, and each bidder i has a private value $v_i \in \mathbb{R}_{\geq 0}$ for each unit of an item, and a private budget $b_i \in \mathbb{R}_{\geq 0}$ on the payment. Here the valuation of bidder i for k units of the item is $k \cdot v_i$.

An allocation rule $a : \mathbb{R}_{\geq 0}^{2n} \rightarrow [0, 1]^{nm}$ maps a type profile $t' \in \mathbb{R}_{\geq 0}^{2n}$ to a matrix of allocation probabilities $a(t') \in [0, 1]^{nm}$, where $a_{ij}(t') \in [0, 1]$ denotes the probability of bidder i being allocated the j -th unit of the item, and $\sum_i a_{ij}(t') \leq 1, \forall j \in [m]$. The payment rule $p : \mathbb{R}_{\geq 0}^{2n} \rightarrow \mathbb{R}_{\geq 0}^n$ defines the expected payment $p_i(t')$ for each bidder.⁶ The utility of a bidder is given by:

$$u_i(t_i, t') = \begin{cases} \sum_{j=1}^m v_{ij} a_{ij}(t') - p_i(t) & \text{if } p_i(t) \leq b_i, \\ -\infty & \text{if } p_i(t) > b_i. \end{cases} \quad (3)$$

In the unit-demand setting, there are multiple distinct items $\{1, \dots, m\}$ for sale, and each bidder i has a private value $v_{ij} \in \mathbb{R}_{\geq 0}$ for each item j , and a private budget b_i . A bidder's valuation for a bundle of items T is the value of the most-valued item in the bundle: $v_i(T) = \max_{j \in T} v_{ij}$. Let $t_i = (v_{i1}, \dots, v_{im}, b_i)$ denote bidder i 's type. The allocation rule $a : \mathbb{R}_{\geq 0}^{n(m+1)} \rightarrow [0, 1]^{nm}$ maps a type profile $t' \in \mathbb{R}_{\geq 0}^{n(m+1)}$ to the probabilities $a_{ij}(t')$ that each bidder i is allocated item j probabilities, and the payment rule $p : \mathbb{R}_{\geq 0}^{n(m+1)} \rightarrow \mathbb{R}_{\geq 0}^n$ outputs the expected payments.

For revenue maximization with unit-demand bidders, it is sufficient to consider allocation rules that allocate at most one item to each bidder. Here we require the matrix of allocation probabilities to be doubly stochastic, i.e. to satisfy $\sum_j a_{ij}(t') \leq 1, \forall i \in [n]$ and $\sum_i a_{ij}(t') \leq 1, \forall j \in [m]$ for all t' . Such a randomized allocation

⁶If the payment rule p is *ex post* IR, for any reported type t' , there exists a set of payments $P_{ij}(t')$ on each outcome (i, j) s.t. each $P_{ij}(t') \leq v_{ij}$, which are equivalent in expectation to $p_i(t')$. These payments can be computed by solving a linear program.

can be decomposed into a lottery over deterministic, feasible assignments (the Birkhoff-von Neumann theorem [5, 31]). The utility of a unit-demand bidder under a doubly stochastic allocation a is again given by (3).

3 THE BUDGETED REGRETNET FRAMEWORK

In this section, we explain how to extend the *RegretNet framework* of Dütting et al. [16], which was developed and applied for settings without budget constraints, to a setting with budget constraints.

We represent an auction as a feed-forward neural network, and optimize the parameters to maximize revenue subject to regret, IR and budget constraints. While the framework of Dütting et al. enforces DSIC by requiring that the (empirical) *ex post* regret for the neural network be zero, we are able to handle more general forms of incentive compatibility by working with an appropriate notion of regret. For BIC, we constrain the (empirical) *interim* regret of the network to be zero; for conditional DSIC/BIC, we constrain the (empirical) conditional regret of the network to be zero. We additionally include budget constraints.

3.1 Network architecture

The allocation and payment rules are represented as separate feed-forward networks, but trained simultaneously, and connected through training loss function and constraints. The network architectures are shown in Figure 1 for the additive setting and in Figure 2 for the unit-demand setting.

Allocation network: The allocation rule for the additive setting takes a type profile t as input and outputs the probability $a_{ij}(t) \in [0, 1]$ of the j -th unit of the item being assigned to each bidder i . The neural network consists of R fully-connected hidden layers, with sigmoid activations and a fully-connected output layer. In the case of additive bidders, the output layer computes a real-valued score s_{ij} for each bidder-item pair (i, j) and converts these scores to allocation probabilities using a softmax function: $a_{ij}(t) = \frac{e^{s_{ij}}}{\sum_{k=1}^{n+1} e^{s_{kj}}}$, where $s_{n+1, j}$ is an additional "dummy score" computed for each item j . Through the inclusion of this dummy score, the softmax ensures that $\sum_{i=1}^n a_{ij}(t) \leq 1$ for each item j . The network can allocate multiple units to a single bidder.

For unit-demand bidders, we require the allocation probabilities to be doubly stochastic. For this, we modify the allocation network to generate a score s_{ij} and a score s'_{ij} for each bidder-item pair (i, j) , with the first set of scores normalized along the rows, and the second set of scores normalized along the columns using softmax functions. The final allocation is an element-wise minimum of the two sets of normalized scores, $a_{ij}(t) = \min \left\{ \frac{e^{s_{ij}}}{\sum_{k=1}^{n+1} e^{s_{kj}}}, \frac{e^{s'_{ij}}}{\sum_{k=1}^{m+1} e^{s'_{jk}}} \right\}$, and is guaranteed to be doubly stochastic.

Payment network: The payment rule is also defined through a feed-forward network, and consists of T fully-connected hidden layers, with sigmoid activations and a fully-connected output layer. Given an input type profile t , the neural network computes a payment $p_i(t)$ for each bidder i . In particular, the output layer computes a score $s'_i \in \mathbb{R}$ for each bidder, and applies the *ReLU activation function* to ensure that payments are non-negative: $p_i(t) = \max\{s'_i, 0\}$.

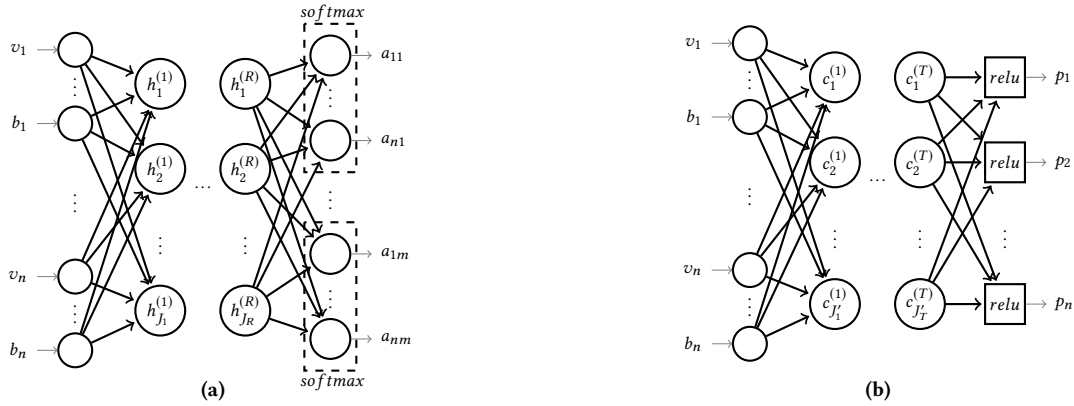


Figure 1: Budgeted RegretNet: (a) Allocation rule a and (b) Payment rule p for a setting with m identical items and n additive buyers.

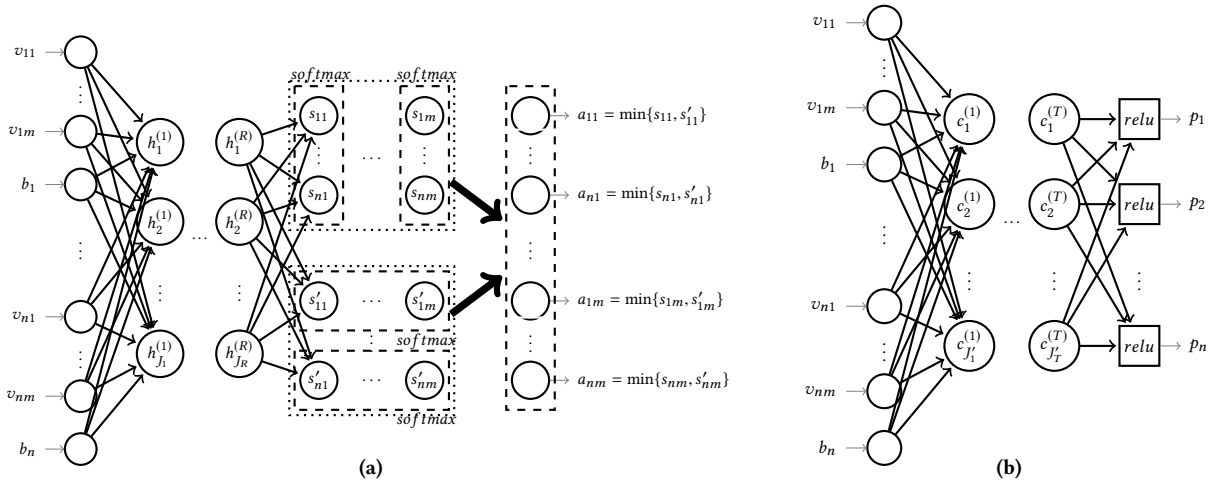


Figure 2: Budgeted RegretNet: (a) Allocation rule a and (b) Payment rule p for a setting with m distinct items and n unit-demand buyers.

3.2 Training problem

We use the following metrics to measure the degree to which an auction violates the BIC, IR and BC constraints.

Regret: We define the expected *interim* regret to bidder i , for an auction with rules (a, p) , as the maximum gain in *interim* utility by misreporting the bidder's type.

$$rgt_i(a, p) = \mathbf{E}_{t_i \sim F_i} \left[\max_{t'_i \in \mathcal{T}_i} \chi_{(\mathcal{P}_i(t'_i) \leq b_i)} (\mathcal{U}_i(t_i, t'_i) - \mathcal{U}_i(t_i, t_i)) \right], \quad (4)$$

where χ_A is an indicator function for whether predicate A is true. An auction is BIC if and only if it has zero *interim* regret. The indicator function in the above expression ensures that the first utility term does not go to $-\infty$. As long as the auction also satisfies interim BC, the second utility term is also finite for all type profiles, thus ensuring that the regret is always finite.

IR penalty: The penalty for violating IR for bidder i is given by:

$$irp_i(a, p) = \mathbf{E}_{t_i \sim F_i} [\max\{0, -\mathcal{U}_i(t_i, t_i)\}]. \quad (5)$$

BC penalty: The penalty for violating the budget constraint for bidder i is given by:

$$bcp_i(a, p) = \mathbf{E}_{t_i \sim F_i} [\max\{0, \mathcal{P}_i(t_i) - b_i\}]. \quad (6)$$

Further, we define the *loss function* as the negated expected revenue $\mathcal{L}(a, p) = -\mathbf{E}_{t \sim F} [\sum_{i=1}^n p_i(t)]$.

Let $\mathbf{w} \in \mathbb{R}^d$ denote the parameters of the allocation network, the induced allocation rule denoted by $a^{\mathbf{w}}$, and $\mathbf{w}' \in \mathbb{R}^{d'}$ denote the parameters of the payment network, the induced payment rule denoted by $p^{\mathbf{w}'}$.

The design objective is to minimize the loss function over the space of network parameters, such that the regret, IR penalty and BC penalty is zero for each bidder:

$$\begin{aligned} & \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{w}' \in \mathbb{R}^{d'}} \mathcal{L}(a^{\mathbf{w}}, p^{\mathbf{w}'}) \\ & \text{s.t. } rgt_i(a^{\mathbf{w}}, p^{\mathbf{w}'}) = 0, \forall i \in [n] \\ & \quad irp_i(a^{\mathbf{w}}, p^{\mathbf{w}'}) = 0, \forall i \in [n] \\ & \quad bcp_i(a^{\mathbf{w}}, p^{\mathbf{w}'}) = 0, \forall i \in [n]. \end{aligned} \quad (\text{OP1})$$

In practice, the loss, regret, IR penalty and BC penalty can be estimated from a sample of type profiles $S = \{t^{(1)}, t^{(2)}, \dots, t^{(L)}\}$ drawn i.i.d. from F . The loss for an auction with rules (a, p) can be estimated as $\widehat{\mathcal{L}}(a, p) = -\frac{1}{L} \sum_{\ell=1}^L \sum_{i=1}^n p_i(t^{(\ell)})$.

To estimate the *interim* regret, for each type profile $t^{(\ell)}$ in S , we draw additional samples $S_\ell = \{\tilde{t}^{(1)}, \dots, \tilde{t}^{(K)}\}$ from F , and $S'_\ell =$

$\{\tilde{t}^{(1)}, \dots, \tilde{t}^{(K')}\}$ from a uniform distribution over type space \mathcal{T} .⁷ Using sample S_ℓ , we define the empirical *interim* utility for bidder i with type t_i and report t'_i as:

$$\widehat{\mathcal{U}}_i(t_i, t'_i) = \frac{1}{K} \sum_{k=1}^K u_i(t_i, (t'_i, \tilde{t}_{-i}^{(k)}))$$

and the empirical *interim* payment as:

$$\widehat{\mathcal{P}}_i(t'_i) = \frac{1}{K} \sum_{k=1}^K p_i(t'_i, \tilde{t}_{-i}^{(k)})$$

Then the empirical *interim* regret is given by:

$$\widehat{rgt}_i(a, p) = \frac{1}{L} \sum_{\ell=1}^L \max_{t' \in S'_\ell} \left\{ \chi(\widehat{\mathcal{P}}_i(t'_i) \leq b_i^{(\ell)}) \cdot \left(\widehat{\mathcal{U}}_i(t_i^{(\ell)}, t'_i) - \widehat{\mathcal{U}}_i(t_i^{(\ell)}, t_i^{(\ell)}) \right) \right\}, \quad (7)$$

where the sample S'_ℓ provides a set of deviating type profiles to approximate the maximum over bidder misreports.

The IR and BC penalties can be similarly estimated as:

$$\begin{aligned} \widehat{irp}_i(a, p) &= \frac{1}{L} \sum_{\ell=1}^L \max \left\{ 0, -\widehat{\mathcal{U}}_i(t_i^{(\ell)}, t_i^{(\ell)}) \right\} \\ \widehat{bcp}_i(a, p) &= \frac{1}{L} \sum_{\ell=1}^L \max \left\{ 0, \widehat{\mathcal{P}}_i(t_i^{(\ell)}) - b_i^{(\ell)} \right\}. \end{aligned}$$

Following Dütting et al. [16], we use the *Augmented Lagrangian method* to solve the resulting sample-based optimization problem:

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{w}' \in \mathbb{R}^{d'}} \quad & \widehat{\mathcal{L}}(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) \\ \text{s.t.} \quad & \widehat{rgt}_i(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) = 0, \forall i \in [n] \\ & \widehat{irp}_i(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) = 0, \forall i \in [n] \\ & \widehat{bcp}_i(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) = 0, \forall i \in [n]. \end{aligned} \quad (\text{OP2})$$

Augmented Lagrangian Solver: The solver formulates a sequence of unconstrained optimization steps that combine the revenue, regret, IR penalty, and budget penalty terms into a single objective, with the relative weights on the regret, IR and budget penalty terms adjusted across iterations. More specifically, the solver constructs the following unconstrained, augmented Lagrangian objective:

$$\begin{aligned} \mathcal{F}_\rho(\mathbf{w}, \mathbf{w}'; \lambda_{rgt}, \lambda_{irp}, \lambda_{bcp}) \\ = \widehat{\mathcal{L}}(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) + \sum_{i \in [n]} \lambda_{rgt, i} \widehat{rgt}_i(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) + \frac{\rho}{2} \sum_{i \in [n]} \widehat{rgt}_i^2(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) \\ + \sum_{i \in [n]} \lambda_{irp, i} \widehat{irp}_i(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) + \frac{\rho}{2} \sum_{i \in [n]} \widehat{irp}_i^2(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) \\ + \sum_{i \in [n]} \lambda_{bcp, i} \widehat{bcp}_i(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) + \frac{\rho}{2} \sum_{i \in [n]} \widehat{bcp}_i^2(\mathbf{a}^{\mathbf{w}}, \mathbf{p}^{\mathbf{w}'}) \end{aligned}$$

where $\lambda_{rgt} \in \mathbb{R}^n$, $\lambda_{irp} \in \mathbb{R}^n$ and $\lambda_{bcp} \in \mathbb{R}^n$ are vectors of Lagrangian multipliers associated with the equality constraints in (OP2), and $\rho > 0$ is a fixed parameter that controls the weight on the augmented quadratic terms.

⁷The deviating types need not be sampled from the distributions of true types. We adopt a uniform sampling scheme, and find this to be effective in our experiments.

The solver operates across multiple iterations, and updates the Lagrange multipliers based on the violation of the constraints in each iteration t :

$$(\mathbf{w}^{t+1}, \mathbf{w}'^{t+1}) \in \operatorname{argmin}_{(\mathbf{w}, \mathbf{w}')} \mathcal{F}_\rho(\mathbf{w}, \mathbf{w}'; \lambda_{rgt}^t, \lambda_{irp}^t, \lambda_{bcp}^t) \quad (8)$$

$$\lambda_{rgt, i}^{t+1} = \lambda_{rgt, i}^t + \rho \widehat{rgt}_i(\mathbf{a}^{\mathbf{w}^{t+1}}, \mathbf{p}^{\mathbf{w}'^{t+1}}), \forall i \in [n], \quad (9)$$

$$\lambda_{irp, i}^{t+1} = \lambda_{irp, i}^t + \rho \widehat{irp}_i(\mathbf{a}^{\mathbf{w}^{t+1}}, \mathbf{p}^{\mathbf{w}'^{t+1}}), \forall i \in [n], \quad (10)$$

$$\lambda_{bcp, i}^{t+1} = \lambda_{bcp, i}^t + \rho \widehat{bcp}_i(\mathbf{a}^{\mathbf{w}^{t+1}}, \mathbf{p}^{\mathbf{w}'^{t+1}}), \forall i \in [n], \quad (11)$$

where the inner optimization in (8) is approximately solved through multiple iterations of the Adam solver [21]. Specifically, the gradient is pushed through the loss function as well as the empirical measures of violation of IC, IR and BC.⁸ In our experiments, the Lagrangian multipliers are initialized to zero.

3.3 Handling other kinds of IC constraints

The approach also extends to a design subject to conditional BIC, as well as DSIC and conditional DSIC. For conditional BIC, we replace the regret in (OP1) with the conditional regret, defined as:

$$crgt_i(a, p) = \mathbf{E}_{t_i \sim F_i} \left[\max_{t'_i \in \mathcal{T}_i} \chi(b'_i \leq b_i) \left(\mathcal{U}_i(t_i, t'_i) - \mathcal{U}_i(t_i, t_i) \right) \right], \quad (12)$$

and use the following estimate of the *conditional interim regret* in (OP2):

$$\begin{aligned} \widehat{crgt}_i(a, p) = \\ \frac{1}{L} \sum_{\ell=1}^L \max_{t' \in S'_\ell} \left\{ \chi(b'_i \leq b_i^{(\ell)}) \cdot \left(\widehat{\mathcal{U}}_i(t_i^{(\ell)}, t'_i) - \widehat{\mathcal{U}}_i(t_i^{(\ell)}, t_i^{(\ell)}) \right) \right\}. \end{aligned} \quad (13)$$

To handle DSIC and conditional DSIC, we replace the *interim* regret in the training problem with the *ex post* regret and a conditional version of the *ex post* regret, respectively. The expected *ex post* regret to bidder i in an auction (a, p) is defined as the maximum gain in *ex post* utility obtained by misreporting her type:

$$\begin{aligned} eprgt_i(a, p) = \\ \mathbf{E}_{t \sim F} \left[\max_{t'_i \in \mathcal{T}_i} \chi(p_i(t'_i, t_{-i}) \leq b_i) \left(u_i(t_i, (t'_i, t_{-i})) - u_i(t_i, (t_i, t_{-i})) \right) \right] \end{aligned} \quad (14)$$

Similarly, the *ex post* IR penalty and *ex post* BC penalty can be defined as:

$$epirp_i(a, p) = \mathbf{E}_{t \sim F} [\max\{0, -u_i(t_i, (t_i, t_{-i}))\}] \quad (15)$$

$$epbcp_i(a, p) = \mathbf{E}_{t \sim F} [\max\{0, p_i(t_i, (t_i, t_{-i})) - b_i\}] \quad (16)$$

To estimate the *ex post* regret, we use a set of deviating (misreport) samples $S'_\ell = \{\tilde{t}^{(1)}, \dots, \tilde{t}^{(K')}\}$, drawn from a uniform distribution over \mathcal{T} :

$$\begin{aligned} \widehat{eprgt}_i(a, p) = \frac{1}{L} \sum_{\ell=1}^L \max_{t' \in S'_\ell} \left\{ \chi(p_i(t'_i, t_{-i}^{(\ell)}) \leq b_i^{(\ell)}) \cdot \left(u_i(t_i^{(\ell)}, (t'_i, t_{-i}^{(\ell)})) - u_i(t_i^{(\ell)}, (t_i^{(\ell)}, t_{-i}^{(\ell)})) \right) \right\}. \end{aligned} \quad (17)$$

⁸The solver handles the indicator function in the regret definition by taking its gradient to be zero.

The *ex post* IR and BC penalties can be estimated as:

$$\widehat{epir}p_i(a, p) = \frac{1}{L} \sum_{\ell=1}^L \max \left\{ 0, -u_i(t_i^{(\ell)}, (t_i^{(\ell)}, t_{-i}^{(\ell)})) \right\} \quad (18)$$

$$\widehat{epbc}p_i(a, p) = \frac{1}{L} \sum_{\ell=1}^L \max \left\{ 0, p_i(t_i^{(\ell)}, t_{-i}^{(\ell)}) - b_i^{(\ell)} \right\} \quad (19)$$

For the conditional *ex post* regret, we replace $\chi_{(p_i(t_i^{(\ell)}, t_{-i}^{(\ell)}) \leq b_i)}$ in $eprgt_i$ by $\chi_{(b_i' \leq b_i)}$. Similarly, in the empirical version of this quantity \widehat{eprgt}_i , we replace $\chi_{(p_i(t_i^{(\ell)}, t_{-i}^{(\ell)}) \leq b_i^{(\ell)})}$ by $\chi_{(b_i' \leq b_i^{(\ell)})}$.

4 EXPERIMENTAL RESULTS

We present experimental results to show that we can find new auctions for settings where the optimal design is unknown, and also recover essentially optimal DSIC and BIC auctions in a variety of simpler settings for which analytical results are available. Since DSIC is a stronger property than BIC, and preferred in practice, we give more focus to the automated design of DSIC auctions.

Experimental setup. We use the *TensorFlow* deep learning library for experiments. We solve the inner optimization in the augmented Lagrangian method using the ADAM solver [21], with a learning rate of 0.001 and a mini-batch size of 64. All the experiments are run on a compute cluster with NVIDIA GPU cores.

Evaluation. We generate training and test data from different type distributions, use the training set for fitting an auction network and evaluate performance of the learned auction on the test set. We use the following metrics for evaluation:

$$\begin{aligned} \text{Regret} &= \frac{1}{n} \sum_{i=1}^n \widehat{rgt}_i(a, p) \\ \text{Conditional Regret} &= \frac{1}{n} \sum_{i=1}^n \widehat{crgt}_i(a, p) \\ \text{IR penalty} &= \frac{1}{n} \sum_{i=1}^n \widehat{irp}_i(a, p) \\ \text{BC penalty} &= \frac{1}{n} \sum_{i=1}^n \widehat{bcp}_i(a, p). \end{aligned}$$

For experiments on DSIC auctions, the terms \widehat{rgt}_i , \widehat{crgt}_i , \widehat{irp}_i and \widehat{bcp}_i are *ex post* quantities. For experiments on BIC, these terms are *interim* quantities. The training and test set are large enough to avoid issues of overfitting. The specific sample sizes and network scale are provided in subsequent subsections.

4.1 Optimal DSIC auctions

We consider the design of DSIC auctions, adopting three different settings studied in the literature:

- Setting I: There is a single item and a single bidder, with the bidder's value $v_1 \sim \text{Unif}[0, 1]$ and budget $b_1 \sim \text{Unif}[0, 1]$. The optimal DSIC auction for this setting was derived by Che and Gale [12].
- Setting II: There is a single item and two bidders, where $v_1, v_2 \sim \text{Unif}\{1, 2, \dots, 10\}$. The first bidder is unconstrained while the second bidder has a budget of 4. The optimal auction under conditional DSIC for this setting was derived by Malakhov and Vohra [24].⁹

⁹ In this special case, the auctioneer knows the true budget of constrained bidder but allows her to misreport her budget. In effect, the budget of constrained bidder is publicly known.

Property	Setting	Opt rev	Budgeted RegretNet			
			rev	regret	irp	bcp
DSIC	I	0.192	0.196	0.002 (0.003)	0.002	0.001
	II (C)	4.664	4.638	0.002	0.005	0.002
	III	–	0.709	0.002 (0.004)	0.0	0.002
	IV	–	0.287	0.002 (0.003)	0.0	0.0
BIC	II (C)	4.847	4.788	0.0	0.0	0.0
	V	0.342	0.348	0.004 (0.005)	0.001	0.0

Table 1: Test metrics for Budgeted RegretNet auctions. Here (C) refers to conditional IC. For continuous valuation distributions, we also report within parenthesis the regret estimated using a larger misreport sample (i.e. with 1000 misreports for each type profile).

Setting	Misreport sample size $ S'_\ell $				
	100	200	400	800	1600
IV	0.0018	0.0021	0.0023	0.0026	0.0029

Table 2: Test regret for Budgeted RegretNet under Setting IV with misreport samples of different sizes for each type profile.

- Setting III: There are four identical items with two additive bidders where bidder i 's value for each item $v_i \sim \text{Unif}[0, 1]$ and the budget $b_i \sim \text{Unif}[0, 1]$. There is no analytical result.
- Setting IV: There are two items with two unit-demand bidders where bidder i 's value for the item j , $v_{ij} \sim \text{Unif}[0, 1]$ and the budget $b_i \sim \text{Unif}[0, 1]$. There is no analytical result.

We use allocation and payment networks with two hidden layers each, and with 100 hidden nodes in each layer. For all the experiments below, for each type profile $t^{(\ell)}$, we randomly generate a sample of 100 misreports S'_ℓ to evaluate the regret. We also report the regret estimated for continuous valuation distributions using a larger misreport sample (of size 1000 or more) for each type profile.¹⁰ A summary of our results is shown in Tables 1 and 2.

For setting I, we use a training and test sample of 5000 profiles each, with the parameter ρ in Augmented Lagrangian solver set to 0.01. Figure 3(a) presents plots of test revenue and test *ex post* regret for the learned auction as a function of solver iterations. Figure 3(b)-(c) show the allocation rule learned by the neural network, and compare this with the optimal rule of Che and Gale [12]. Not only does the learned auction yields revenue close to the optimal auctions and incur negligible regret, but the learned allocation rule closely matches the optimal rule. From Table 1, we see that the learned auction also incurs very small IR and budget violations.

For setting II, we use a smaller training and test sample of 1000 profiles, which are large enough for the discrete distribution considered here. We set ρ to 0.001. The optimal auction for this setting is given by Malakhov and Vohra [24]. We trained neural network for conditional DSIC. Figure 4(a) shows plots of the test revenue for the learned auction, as well as plots of the test *ex post* regret for the learned auction under conditional DSIC constraints. The learned auction yields revenue very close to the optimal revenue, while yielding negligible regret, IR violations, or budget violations. Furthermore, as seen in Figure 4(b)-(c), the learned allocation rule for conditional DSIC closely matches the analytical result in Malakhov and Vohra [24].

For setting III, we use a training and test sample of 5000 profiles, with ρ set to 0.01. Since the optimal auction is not provided by the

¹⁰For discrete valuation distributions in this paper, we find a sample of 100 misreports to be large enough to accurately estimate the regret.

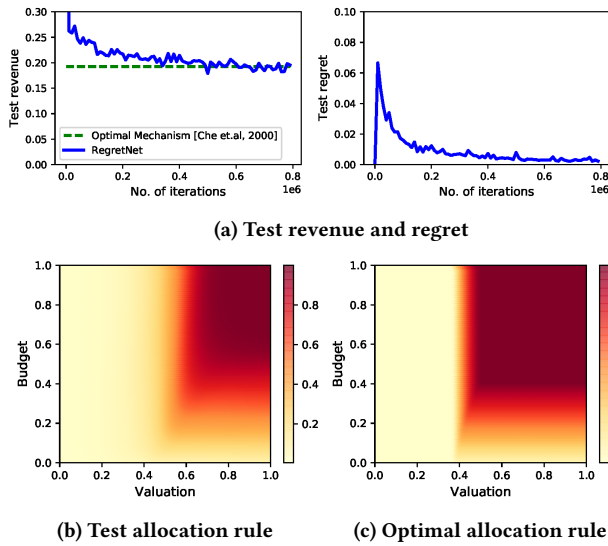


Figure 3: The auction learned under DSIC for Setting I with a single item and single bidder, where $v_1 \sim Unif[0, 1]$ and $b_1 \sim Unif[0, 1]$. The solid regions in (b) and (c) depict the probability of the item being allocated to the bidder.

theoretical literature, we compare the learned auction rule against the optimal posted pricing auction, as well as the auction proposed by Borgs et al. [6]. Figure 5 shows test revenue and *ex post* regret as functions of solver iterations. In this case, the neural network is able to discover an auction with a higher revenue than the baseline, while incurring a very small regret, as well as, very small IR and budget violations.

For setting IV, we use a training and test sample of 5000 profiles, with ρ set to 0.03. Since there is no analytical result for this setting, we compare the learned auction rules against the ascending auction of Ashlagi and Braverman [1]. Figure 6 shows the test revenue and *ex post* regret as functions of the number of solver iterations. The auction learned by RegretNet has a higher revenue than the baseline, while incurring very small regret, IR, and budget violations.

Since the regret is estimated using a sample of misreports, for this experiment, we also evaluate the regret using misreport samples S'_ℓ of different sizes. The results are summarized in Table 2. Figure 7 shows the test *ex post* regret as functions of solver iterations for different sizes of misreport samples. As seen, even with larger number of misreport samples, the regret is still very small, implying that the learned auction is indeed essentially IC.

4.2 Optimal BIC auctions

Next, we consider the automated design of BIC auctions. Here we focus on settings for which analytical results are available. This serves to provide a validation that we are able to use *RegretNet* to learn BIC designs. We are less interested in optimal BIC for new settings because we consider DSIC of more practical interest. We consider the following settings:

- Setting II from Section 4.1. The optimal BIC auction for this setting was derived by Malakhov and Vohra [24].
- Setting V: There is a single item and two symmetric budget constrained bidders. Each bidder draws a value $v_i \sim$

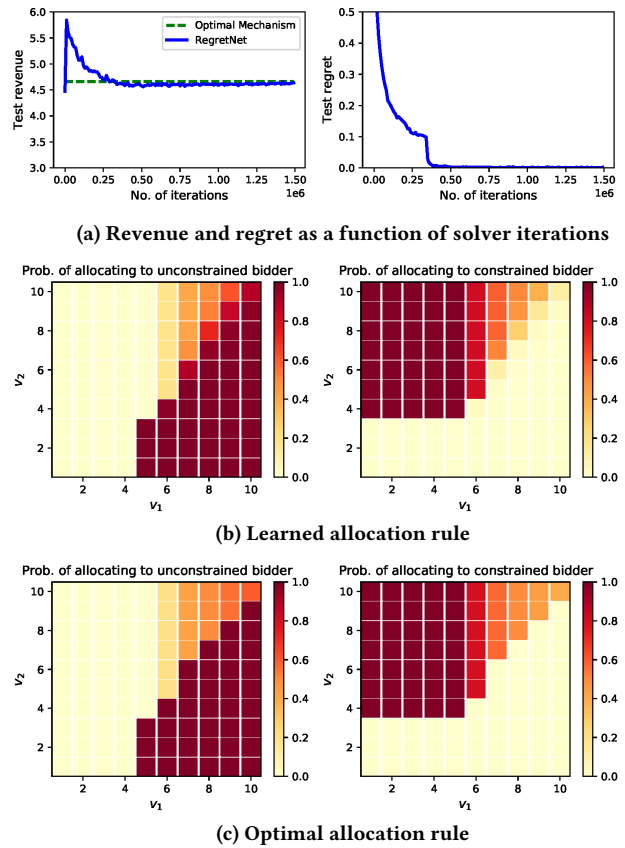


Figure 4: The auction learned under conditional DSIC for Setting II with a single item and two bidders, where $v_1, v_2 \sim Unif\{1, 2, \dots, 10\}$, bidder 1 is unconstrained, and bidder 2 has a budget of 4.




Figure 5: Revenue and regret for the DSIC auction learned under Setting III with four identical items and two additive bidders, where bidder i 's value for each item $v_i \sim Unif[0, 1]$ and $b_i \sim Unif[0, 1]$.

$Unif[0, 1]$ and budget $b_i \sim Unif\{0.22, 0.42\}$. The optimal auction for this setting was derived by Pai and Vohra [29].

For these experiments, we use allocation and payment networks with two hidden layers with 50 nodes each.¹¹ A summary of the results is provided in Table 1. The training and test set have 1000 type profiles each and ρ was set to 0.05. To learn the BIC auctions, we need additional samples S_ℓ from known distribution F for each type profile $t^{(\ell)}$, which makes the training of RegretNet more costly than for the case of DSIC auctions.

¹¹Unlike DSIC settings, we reduce the size of neural networks in BIC settings to trade-off the cost of more computation for estimating interim rules.

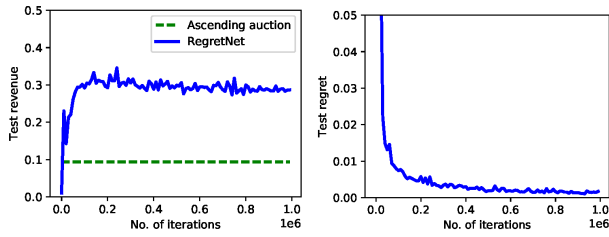


Figure 6: Revenue and regret for the DSIC auction learned under Setting IV with two items and two unit-demand bidders, where bidder i 's value for item j $v_{ij} \sim \text{Unif}[0, 1]$ and $b_i \sim \text{Unif}[0, 1]$.

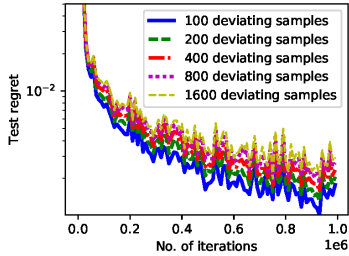


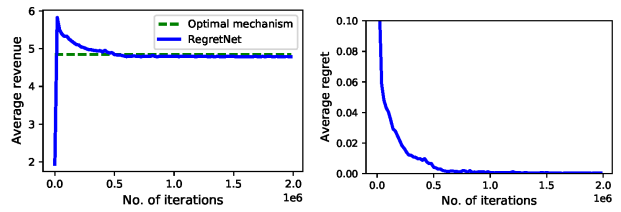
Figure 7: A semi-logarithmic plot of test regret as a function of the number of iterations for different misreport sample sizes for the DSIC auction learned under Setting IV.

Figure 8 presents the results of learning a BIC auction for setting II, providing the test revenue and test interim regret as a function of the number of solver iterations. We also illustrate the learned allocation rule, and compare it with the optimal allocation rule of Malakhov and Vohra [24]. Not only does the auction that is derived through machine learning achieve near-optimal revenue with essentially zero regret, IR and budget violations, but we closely recover the design of the optimal allocation rule. Figure 9 shows the test revenue and interim regret of the learned auction for setting V. Again, we are able to achieve almost-optimal revenue, while incurring very small regret, IR, and budget violations.


5 CONCLUSION

We have used deep learning to design essentially optimal, multi-item auctions under private budget constraints. Whereas the state-of-the-art analytical results for the design of optimal, DSIC auctions cannot handle more than two bidders, or more than one item (to even a single bidder), *RegretNet* can discover new, essentially incentive-compatible designs with high revenue in these settings (consider Setting III and Setting IV). We also validate the approach by demonstrating that *RegretNet* can recover essentially optimal designs in settings for which optimal analytical results do exist, including the case of BIC auction design.


In the future, it will be interesting to study the robustness of the learned auctions to perturbations in the type distributions, develop methods that allow a single network to handle different number of bidders or items, improve the efficiency with which we can train *RegretNet* in the case of BIC design, and use our approach to estimate both upper- and lower-bounds on the revenue from exactly IC designs. It will also be interesting to explore the effect of allowing for correlation between value and budget and across bidders, soft



(a) Revenue and regret as a function of solver iterations



(b) Learned allocation rule



(c) Optimal allocation rule

Figure 8: Auction learned under BIC for Setting II with a single item and two bidders, where $v_1, v_2 \sim \text{Unif}\{1, 2, \dots, 10\}$, bidder 1 is unconstrained and bidder 2 has a budget of 4.




Figure 9: Revenue and regret of auction learned under BIC for Setting V with a single item and two bidders, where $v_1, v_2 \sim \text{Unif}[0, 1]$ and $b_1, b_2 \sim \text{Unif}\{0.22, 0.42\}$.

budget constraints, and budgets that depend on a bidder's allocation. All of these seem within reach of automated methods, but are extremely challenging to handle through theoretical analysis.

6 ACKNOWLEDGMENTS

We would like to thank Paul Dütting for his contributions to the broader project on using deep learning for optimal economic design. Thanks also to Mallesh Pai for a helpful discussion on budget-constrained auctions, and to the anonymous reviewers.

REFERENCES

- [1] Itai Ashlagi and Mark Braverman. 2009. Ascending Unit Demand Auctions with Budget Limits. *Technical Report* (2009).
- [2] Itai Ashlagi, Mark Braverman, Avinatan Hassidim, Ron Lavi, and Moshe Tennenholtz. 2010. Position Auctions with Budgets: Existence and Uniqueness. *The B.E. Journal of Theoretical Economics* 10, 1 (2010).
- [3] Lawrence M. Ausubel and Paul Milgrom. 2006. The Lovely but Lonely Vickrey Auction. In *Combinatorial Auctions, chapter 1*. MIT Press.
- [4] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. 2012. Budget-Constrained Auctions with Heterogeneous Items. *Theory of Computing* 8, 20 (2012), 429–460.
- [5] Garrett Birkhoff. 1946. Three observations on linear algebra. *Univ. Nac. Tucuman. Revista A* 5 (1946), 147–151.
- [6] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. 2005. Multi-unit Auctions with Budget-constrained Bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce*. 44–51.
- [7] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th ACM Symposium on Theory of Computing*. 459–478.
- [8] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. Optimal Multi-dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*. 130–139.
- [9] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2013. Understanding Incentives: Mechanism Design Becomes Algorithm Design. In *Proceedings of the 54th IEEE Symposium on Foundations of Computer Science*. 618–627.
- [10] Shuchi Chawla, David L. Malec, and Azarakhsh Malekian. 2011. Bayesian Mechanism Design for Budget-constrained Agents. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. 253–262.
- [11] Yeon-Koo Che and Ian Gale. 1998. Standard Auctions with Financially Constrained Bidders. *The Review of Economic Studies* 65, 1 (1998), 1–21.
- [12] Yeon-Koo Che and Ian Gale. 2000. The Optimal Mechanism for Selling to a Budget-Constrained Buyer. *Journal of Economic Theory* 92, 2 (2000), 198–233.
- [13] Riccardo Colini-Baldeschi, Stefano Leonardi, Monika Henzinger, and Martin Starnberger. 2015. On Multiple Keyword Sponsored Search Auctions with Budgets. *ACM Trans. Econ. Comput.* 4, 1, Article 2 (Dec. 2015), 2:1–2:34 pages.
- [14] Vincent Conitzer and Tuomas Sandholm. 2002. Complexity of Mechanism Design. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*. San Francisco, CA, USA, 103–110.
- [15] Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2008. Multi-unit Auctions with Budget Limits. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*. Washington, DC, USA, 260–269.
- [16] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, and David C. Parkes. 2017. Optimal Auctions through Deep Learning. *CoRR* abs/1706.03459 (2017).
- [17] Paul Dütting, Felix A. Fischer, Pichayut Jirapinyo, John K. Lai, Benjamin Lubin, and David C. Parkes. 2012. Payment rules through discriminant-based classifiers. In *Proceedings of the 13th ACM Conference on Electronic Commerce*.
- [18] Paul Dütting, Monika Henzinger, and Martin Starnberger. 2015. Auctions for Heterogeneous Items and Budget Limits. *ACM Trans. Econ. Comput.* 4, 1, Article 4 (2015), 4:1–4:17 pages.
- [19] Paul Dütting, Monika Henzinger, and Ingmar Weber. 2011. An Expressive Mechanism for Auctions on the Web. In *Proceedings of the 20th International Conference on World Wide Web*. 127–136.
- [20] Noah Golowich, Harikrishna Narasimhan, and David C. Parkes. 2018. Deep Learning for Multi-Facility Location Mechanism Design. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*. To appear.
- [21] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014).
- [22] János Kornai, Eric Maskin, and Gérard Roland. 2003. Understanding the Soft Budget Constraint. *Journal of Economic Literature* 41, 4 (2003), 1095–1136.
- [23] Jean-Jacques Laffont and Jacques Robert. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2 (1996), 181–186.
- [24] Alexey Malakhov and Rakesh V. Vohra. 2008. Optimal auctions for asymmetrically budget constrained bidders. *Review of Economic Design* 12, 4 (2008), 245.
- [25] Eric S. Maskin. 2000. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review* 44, 4 (2000), 667–681.
- [26] Roger Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6 (1981), 58–73.
- [27] Harikrishna Narasimhan, Shivani Agarwal, and David C. Parkes. 2016. Automated Mechanism Design without Money via Machine Learning. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*. 433–439.
- [28] Harikrishna Narasimhan and David C Parkes. 2016. A general statistical framework for designing strategy-proof assignment mechanisms. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*. 527–536.
- [29] Malleh M. Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150 (2014), 383–425.
- [30] Tuomas Sandholm and Anton Likhodedov. 2015. Automated Design of Revenue-Maximizing Combinatorial Auctions. *Operations Research* 63, 5 (2015), 1000–1025.
- [31] John von Neumann. 1953. A Certain Zero-sum Two-person Game equivalent to the Optimal Assignment Problem. *Contributions to the Theory of Games (AM-28), Volume II* (1953), 5–12.