**Optimal auctions through deep learning**

**LSE Research Online URL for this paper:** http://eprints.lse.ac.uk/100806/

Version: Published Version

**Book Section:**

Dütting, Paul, Feng, Zhe, Narasimham, Harikrishna, Parkes, David C. and Ravindranath, Sal S (2019) Optimal auctions through deep learning. In: Proceedings of ICML 2019. Proceedings of Machine Learning Research,97. UNSPECIFIED.

# Optimal Auctions through Deep Learning

## Appendix

## A. Omitted Proofs

### A.1. Proof of Lemma 1 and Proof of Lemma 2

***Proof of Lemma 1.*** First, given the property of Softmax function and the min operation, $\varphi^{DS}(s, s')$ ensures that the row sums and column sums for the resulting allocation matrix do not exceed 1. In fact, for any doubly stochastic allocation $z$, there exists scores $s$ and $s'$, for which the min of normalized scores recovers $z$ (e.g. $s_{ij} = s'_{ij} = log(z_{ij}) + c$ for any $c \in \mathbb{R}$). □

***Proof of Lemma 2.*** Similar to Lemma 1, $\varphi^{CF}(s, s^{(1)}, \ldots, s^{(m)})$ trivially satisfies the combinatorial feasibility (constraints (3)–(4)). For any allocation $z$ that satisfies the combinatorial feasibility, the following scores

$$\forall j = 1, \cdots, m, \quad s_{i,S} = s_{i,S}^{(j)} = \log(z_{i,S}) + c,$$

makes $\varphi^{CF}(s, s^{(1)}, \ldots, s^{(m)})$ recover $z$. □

### A.2. Proof of Theorem 1

We present the proof for auctions with general, randomized allocation rules. A randomized allocation rule $g_i : V \to [0, 1]^{2^M}$ maps valuation profiles to a vector of allocation probabilities for bidder $i$. Here $g_{i,S}(v) \in [0, 1]$ denote the probability that the allocation rule assigns subset of items $S \subseteq M$ to bidder $i$, and $\sum_{S \subseteq M} g_{i,S}(v) \le 1$. Note that this encompasses the allocation rules we consider for additive and unit-demand valuations, which only output allocation probabilities for individual items. The payment function $p : V \to R^n$ maps valuation profiles to a payment for each bidder $p_i(v) \in \mathbb{R}$. For ease of exposition, we omit the superscripts "$w$". As before, $\mathcal{M}$ is a class of auctions $(g, p)$.

We will assume that the allocation and payment rules in $\mathcal{M}$ are continuous and that the set of valuation profiles $V$ is a compact set.

**Notation.** For any vectors $a, b \in \mathbb{R}^d$, the inner product is denoted as $\langle a, b \rangle = \sum_{i=1}^{d} a_i b_i$. For any matrix $A \in \mathbb{R}^{k \times \ell}$, the $L_1$ norm is given by $\|A\|_1 = \max_{1 \le j \le \ell} \sum_{i=1}^{k} A_{ij}$.

Let $\mathcal{U}_i$ be the class of utility functions for bidder $i$ defined on auctions in $\mathcal{M}$, i.e.:

$$\mathcal{U}_i = \big\{ u_i : V_i \times V \to \mathbb{R} \,\big|\, u_i(v_i, b) = v_i(g(b)) - p_i(b) \text{ for some } (g, p) \in \mathcal{M} \big\}.$$

and let $\mathcal{U}$ be the class of profile of utility functions defined on $\mathcal{M}$, i.e. the class of tuples $(u_1, \ldots, u_n)$ where each $u_i : V_i \times V \to \mathbb{R}$ and $u_i(v_i, b) = v_i(g(b)) - p_i(b), \forall i \in N$ for some $(g, p) \in \mathcal{M}$. We will sometimes find it useful to represent the utility function as an inner product, i.e. treating $v_i$ as a real-valued vector of length $2^M$, we may write $u_i(v_i, b) = \langle v_i, g_i(b) \rangle - p_i(b)$.

Let $\text{rgt} \circ \mathcal{U}_i$ be the class of all regret functions for bidder $i$ defined on utility functions in $\mathcal{U}_i$:

$$\text{rgt} \circ \mathcal{U}_i = \Big\{ f_i : V \to \mathbb{R} \,\Big|\, f_i(v) = \max_{v'_i} u_i(v_i, (v'_i, v_{-i})) - u_i(v_i, v) \text{ for some } u_i \in \mathcal{U}_i \Big\}$$

and as before, let $\text{rgt} \circ \mathcal{U}$ be defined as the class of profiles of regret functions.

Define the $\ell_{\infty,1}$ distance between two utility functions $u$ and $u'$ as $\max_{v,v'} \sum_i |u_i(v_i, (v'_i, v_{-i})) - u_i(v_i, (v'_i, v_{-i}))|$ and $\mathcal{N}_\infty(\mathcal{U}, \epsilon)$ is the minimum number of balls of radius $\epsilon$ to cover $\mathcal{U}$ under this distance. Similarly, define the distance between $u_i$ and $u'_i$ as $\max_{v,v'_i} |u_i(v_i, (v'_i, v_{-i})) - u'_i(v_i, (v'_i, v_{-i}))|$, and let $\mathcal{N}_\infty(\mathcal{U}_i, \epsilon)$ denote the minimum number of balls of radius $\epsilon$ to cover $\mathcal{U}_i$ under this distance. Similarly, we define covering numbers $\mathcal{N}_\infty(\text{rgt} \circ \mathcal{U}_i, \epsilon)$ and $\mathcal{N}_\infty(\text{rgt} \circ \mathcal{U}, \epsilon)$ for the function classes $\text{rgt} \circ \mathcal{U}_i$ and $\text{rgt} \circ \mathcal{U}$ respectively.

Moreover, we denote the class of allocation functions as $\mathcal{G}$ and for each bidder $i$, $\mathcal{G}_i = \{g_i : V \to 2^M \,|\, g \in \mathcal{G}\}$. Similarly, we denote the class of payment functions by $\mathcal{P}$ and $\mathcal{P}_i = \{p_i : V \to \mathbb{R} \,|\, p \in \mathcal{P}\}$. We denote the covering number of $\mathcal{P}$ as $\mathcal{N}_\infty(\mathcal{P}, \epsilon)$ under the $\ell_{\infty,1}$ distance and the covering number for $\mathcal{P}_i$ using $\mathcal{N}_\infty(\mathcal{P}_i, \epsilon)$ under the $\ell_\infty$ distance.

We first state the following lemma from (Shalev-Shwartz & Ben-David, 2014). Let $\mathcal{F}$ be a class of functions $f : Z \to [-c, c]$ for some input space $Z$ and $c > 0$. Given a sample $\mathcal{S} = \{z_1, \ldots, z_L\}$ of points from $Z$, define the empirical Rademacher complexity of $\mathcal{F}$ as:

$$\hat{\mathcal{R}}_L(\mathcal{F}) := \frac{1}{L}\mathbf{E}_\sigma\left[\sup_{f \in \mathcal{F}} \sum_{z_i \in S} \sigma_i f(z_i)\right],$$

where $\sigma \in \{-1, 1\}^L$ and each $\sigma_i$ is drawn i.i.d from a uniform distribution on $\{-1, 1\}$.

**Lemma 3** (Generalization bound in terms of Rademacher complexity). *Let $\mathcal{S} = \{z_1, \ldots, z_L\}$ be a sample drawn i.i.d. from some distribution $D$ over $Z$. Then with probability of at least $1 - \delta$ over draw of $\mathcal{S}$ from $D$, for all $f \in \mathcal{F}$,*

$$\mathbf{E}_{z \in D}[f(z)] \leq \frac{1}{L}\sum_{i=1}^L f(z_i) + 2\hat{\mathcal{R}}_L(\mathcal{F}) + 4c\sqrt{\frac{2\log(4/\delta)}{L}}.$$

We are now ready to prove Theorem 1. We begin with the first part, namely a generalization bound for revenue.

*Proof of Theorem 1 (Part 1).* The proof involves a direct application of Lemma 3 to the class of revenue functions defined on $\mathcal{M}$:

$$\text{rev} \circ \mathcal{M} = \left\{f : V \to \mathbb{R} \,\middle|\, f(v) = \sum_{i=1}^n p_i(v), \text{ for some } (g, p) \in \mathcal{M}\right\},$$

and bounds the Rademacher complexity term for this class in terms of the covering number for the payment class $\mathcal{P}$, which in turn is bounded by the covering number for the auction class for $\mathcal{M}$.

Since we assume that the auctions in $\mathcal{M}$ satisfy individual rationality and the valuation functions are bounded in $[0, 1]$, we have for any $v$, $p_i(v) \leq 1$. By definition of the covering number $\mathcal{N}_\infty(\mathcal{P}, \epsilon)$ for the payment class, for any $p \in \mathcal{P}$, there exists a $f_p \in \hat{\mathcal{P}}$ where $|\hat{\mathcal{P}}| \leq \mathcal{N}_\infty(\mathcal{P}, \epsilon)$, such that $\max_v \sum_i |p_i(v) - f_{p_i}(v)| \leq \epsilon$. First we bound the Rademacher complexity, for a given $\epsilon \in (0, 1)$,

$$\hat{\mathcal{R}}_L(\text{rev} \circ \mathcal{M}) = \frac{1}{L}\mathbf{E}_\sigma\left[\sup_p \sum_{\ell=1}^L \sigma_\ell \cdot \sum_i p_i(v^{(\ell)})\right]$$

$$= \frac{1}{L}\mathbf{E}_\sigma\left[\sup_p \sum_{\ell=1}^L \sigma_\ell \cdot \sum_i f_{p_i}(v^{(\ell)})\right] + \frac{1}{L}\mathbf{E}_\sigma\left[\sup_p \sum_{\ell=1}^L \sigma_\ell \cdot \sum_i p_i(v^{(\ell)}) - f_{p_i}(v^{(\ell)})\right]$$

$$\leq \frac{1}{L}\mathbf{E}_\sigma\left[\sup_{\hat{p} \in \hat{\mathcal{P}}} \sum_{\ell=1}^L \sigma_\ell \cdot \sum_i \hat{p}_i(v^{(\ell)})\right] + \frac{1}{L}\mathbf{E}_\sigma\|\sigma\|_1\epsilon$$

$$\leq \sqrt{\sum_\ell(\sum_i \hat{p}_i(v^\ell))^2}\sqrt{\frac{2\log(\mathcal{N}_\infty(\mathcal{P}, \epsilon))}{L}} + \epsilon \quad \text{(By Massart's Lemma)}$$

$$\leq 2n\sqrt{\frac{2\log(\mathcal{N}_\infty(\mathcal{P}, \epsilon))}{L}} + \epsilon.$$

The last inequality is because

$$\sqrt{\sum_\ell\left(\sum_i \hat{p}_i(v^\ell)\right)^2} \leq \sqrt{\sum_\ell\left(\sum_i p_i(v^\ell) + n\epsilon\right)^2} \leq 2n\sqrt{L}.$$

Next we show $\mathcal{N}_\infty(\mathcal{P}, \epsilon) \leq \mathcal{N}_\infty(\mathcal{M}, \epsilon)$, for any $(g, p) \in \mathcal{M}$, take $(\hat{g}, \hat{p})$ s.t. for all $v$

$$\sum_{i,j}|g_{ij}(v) - \hat{g}_{ij}(v)| + \sum_i|p_i(v) - \hat{p}_i(v)| \leq \epsilon.$$

Thus for any $p \in \mathcal{P}$, for all $v$, $\sum_i |p_i(v) - \hat{p}_i(v)| \leq \epsilon$, which implies $\mathcal{N}_\infty(\mathcal{P}, \epsilon) \leq \mathcal{N}_\infty(\mathcal{M}, \epsilon)$. Applying Lemma 3 and $\sum_i p_i(v) \leq n$ for any $v$, with probability of at least $1 - \delta$,

$$\mathbf{E}_{v \sim F}\left[-\sum_{i \in N} p_i(v)\right] \leq -\frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{n} p_i(v^{(\ell)}) + 2 \cdot \inf_{\epsilon > 0}\left\{\epsilon + 2n\sqrt{\frac{2\log(\mathcal{N}_\infty(\mathcal{M}, \epsilon))}{L}}\right\} + Cn\sqrt{\frac{\log(1/\delta)}{L}}.$$

This completes the proof for the first part. $\qquad\square$

We move to the second part, namely a generalization bound for regret, which is the more challenging part of the proof.

*Proof of Theorem 1 (Part 2).* We first define the class of sum regret functions:

$$\overline{\mathrm{rgt}} \circ \mathcal{U} = \left\{f : V \to \mathbb{R} \,\middle|\, f(v) = \sum_{i=1}^{n} r_i(v) \text{ for some } (r_1, \ldots, r_n) \in \mathrm{rgt} \circ \mathcal{U}\right\}.$$

The proof then proceeds in three steps:

(1) bounding the covering number for each regret class $\mathrm{rgt} \circ \mathcal{U}_i$ in terms of the covering number for individual utility classes $\mathcal{U}_i$,

(2) bounding the covering number for the combined utility class $\mathcal{U}$ in terms of the covering number for $\mathcal{M}$, and

(3) bounding the covering number for the sum regret class $\overline{\mathrm{rgt}} \circ \mathcal{U}$ in terms of the covering number for the (combined) utility class $\mathcal{M}$.

An application of Lemma 3 then completes the proof. We prove each of the above steps below.

**Step 1.** $\mathcal{N}_\infty(\mathrm{rgt} \circ \mathcal{U}_i, \epsilon) \leq \mathcal{N}_\infty(\mathcal{U}_i, \epsilon/2)$.

By definition of covering number $\mathcal{N}_\infty(\mathcal{U}_i, \epsilon)$, there exists $\hat{\mathcal{U}}_i$ with size at most $\mathcal{N}_\infty(\mathcal{U}_i, \epsilon/2)$ such that for any $u_i \in \mathcal{U}_i$, there exists a $\hat{u}_i \in \hat{\mathcal{U}}_i$ with

$$\sup_{v, v'_i} |u_i(v_i, (v'_i, v_{-i})) - \hat{u}_i(v_i, (v'_i, v_{-i}))| \leq \epsilon/2.$$

For any $u_i \in \mathcal{U}_i$, taking $\hat{u}_i \in \hat{\mathcal{U}}_i$ satisfying the above condition, then for any $v$,

$$\left|\max_{v'_i \in V}\left(u_i(v_i, (v'_i, v_{-i})) - u_i(v_i, (v_i, v_{-i}))\right) - \max_{\bar{v}_i \in V}\left(\hat{u}_i(v_i, (\bar{v}_i, v_{-i})) - \hat{u}_i(v_i, (v_i, v_{-i}))\right)\right|$$

$$\leq \left|\max_{v'_i} u_i(v_i, (v'_i, v_{-i})) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i})) + \hat{u}_i(v_i, (v_i, v_{-i})) - u_i(v_i, (v_i, v_{-i}))\right|$$

$$\leq \left|\max_{v'_i} u_i(v_i, (v'_i, v_{-i})) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}))\right| + |\hat{u}_i(v_i, (v_i, v_{-i})) - u_i(v_i, (v_i, v_{-i}))|$$

$$\leq \left|\max_{v'_i} u_i(v_i, (v'_i, v_{-i})) - \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}))\right| + \epsilon/2.$$

Let $v_i^* \in \arg\max_{v'_i} u_i(v_i, (v'_i, v_{-i}))$ and $\hat{v}_i^* \in \arg\max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i}))$, then

$$\max_{v'_i} u_i(v_i, (v'_i, v_{-i})) = u_i(v_i^*, v_{-i}) \leq \hat{u}_i(v_i^*, v_{-i}) + \epsilon/2 \leq \hat{u}_i(\hat{v}_i^*, v_{-i}) + \epsilon/2 = \max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i})) + \epsilon/2, \text{ and}$$

$$\max_{\bar{v}_i} \hat{u}_i(v_i, (\bar{v}_i, v_{-i})) = \hat{u}_i(\hat{v}_i^*, v_{-i}) \leq u_i(\hat{v}_i^*, v_{-i}) + \epsilon/2 \leq u_i(v_i^*, v_{-i}) + \epsilon/2 = \max_{v'_i} u_i(v_i, (v'_i, v_{-i})) + \epsilon/2.$$

$$(6)$$

Thus, for all $u_i \in \mathcal{U}_i$, there exists $\hat{u}_i \in \hat{\mathcal{U}}_i$ such that for any valuation profile $v$,

$$\left|\max_{v'_i}\left(u_i(v_i, (v'_i, v_{-i})) - u_i(v_i, (v_i, v_{-i}))\right) - \max_{\bar{v}_i}\left(\hat{u}_i(v_i, (\bar{v}_i, v_{-i})) - \hat{u}_i(v_i, (v_i, v_{-i}))\right)\right| \leq \epsilon,$$

which implies $\mathcal{N}_\infty(\mathrm{rgt} \circ \mathcal{U}_i, \epsilon) \leq \mathcal{N}_\infty(\mathcal{U}_i, \epsilon/2)$.

This completes the proof for Step 1.

**Step 2.** $\mathcal{N}_\infty(\mathcal{U}, \epsilon) \leq \mathcal{N}_\infty(\mathcal{M}, \epsilon)$, for all $i \in N$.

Recall the utility function of bidder $i$ is $u_i(v_i, (v_i', v_{-i})) = \langle v_i, g_i(v_i', v_{-i}) \rangle - p_i(v_i', v_{-i})$. There exists a set $\hat{\mathcal{M}}$ with $|\hat{\mathcal{M}}| \leq \mathcal{N}_\infty(\mathcal{M}, \epsilon)$ such that there exists $(\hat{g}, \hat{p}) \in \hat{M}$ with

$$\sup_{v \in V} \sum_{i,j} |g_{ij}(v) - \hat{g}_{ij}(v)| + \|p(v) - \hat{p}(v)\|_1 \leq \epsilon.$$

We denote $\hat{u}_i(v_i, (v_i', v_{-i})) = \langle v_i, \hat{g}_i(v_i', v_{-i}) \rangle - \hat{p}_i(v_i', v_{-i})$, where we treat $v_i$ as a real-valued vector of length $2^M$.

For all $v \in V, v_i' \in V_i$,

$$
\begin{aligned}
&|u_i(v_i, (v_i', v_{-i})) - \hat{u}_i(v_i, (v_i', v_{-i}))| \\
&\leq \quad |\langle v_i, g_i(v_i', v_{-i}) \rangle - \langle v_i, \hat{g}_i(v_i', v_{-i}) \rangle| + |p_i(v_i', v_{-i}) - \hat{p}_i(v_i', v_{-i})| \\
&\leq \quad \|v_i\|_\infty \cdot \|g_i(v_i', v_{-i}) - \hat{g}_i(v_i', v_{-i})\|_1 + |p_i(v_i', v_{-i}) - \hat{p}_i(v_i', v_{-i})| \\
&\leq \quad \sum_j |g_{ij}(v_i', v_{-i}) - \hat{g}_{ij}(v_i', v_{-i})| + |p_i(v_i', v_{-i}) - \hat{p}_i(v_i', v_{-i})|.
\end{aligned}
$$

Therefore, for any $u \in \mathcal{U}$, take $\hat{u} = (\hat{g}, \hat{p}) \in \hat{\mathcal{M}}$, for all $v, v'$,

$$
\sum_i |u_i(v_i, (v_i', v_{-i})) - \hat{u}_i(v_i, (v_i', v_{-i}))|
$$

$$
\leq \sum_{ij} |g_{ij}(v_i', v_{-i}) - \hat{g}_{ij}(v_i', v_{-i})| + \sum_i |p_i(v_i', v_{-i}) - \hat{p}_i(v_i', v_{-i})| \leq \epsilon.
$$

This completes the proof for Step 2.

**Step 3.** $\mathcal{N}_\infty(\overline{\mathrm{rgt}} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_\infty(\mathcal{M}, \epsilon/2)$

By definition of $\mathcal{N}_\infty(\mathcal{U}, \epsilon)$, there exists $\hat{\mathcal{U}}$ with size at most $\mathcal{N}_\infty(\mathcal{U}, \epsilon)$, such that, for any $u \in \mathcal{U}$, there exists $\hat{u}$ s.t. for all $v, v' \in V$, $\sum_i |u_i(v_i, (v_i', v_{-i})) - \hat{u}_i(v_i, (v_i', v_{-i}))| \leq \epsilon$. Therefore for all $v \in V$, $|\sum_i u_i(v_i, (v_i', v_{-i})) - \sum_i \hat{u}_i(v_i, (v_i', v_{-i}))| \leq \epsilon$, from which it follows that $\mathcal{N}_\infty(\overline{\mathrm{rgt}} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_\infty(\mathrm{rgt} \circ \mathcal{U}, \epsilon)$. Following Step 1, it is easy to show $\mathcal{N}_\infty(\mathrm{rgt} \circ \mathcal{U}, \epsilon) \leq \mathcal{N}_\infty(\mathcal{U}, \epsilon/2)$. This further with Step 2 completes the proof of Step 3.

Based on the same arguments as in the proof of Theorem 1 (Part 1) the empirical Rademacher complexity is bounded as:

$$
\hat{\mathcal{R}}_L(\overline{\mathrm{rgt}} \circ \mathcal{U}) \leq \inf_{\epsilon > 0} \left( \epsilon + 2n \sqrt{\frac{2 \log \mathcal{N}_\infty(\overline{\mathrm{rgt}} \circ \mathcal{U}, \epsilon)}{L}} \right) \leq \inf_{\epsilon > 0} \left( \epsilon + 2n \sqrt{\frac{2 \log \mathcal{N}_\infty(\mathcal{M}, \epsilon/2)}{L}} \right).
$$

Applying Lemma 3, completes the proof for generalization bound for regret. $\qquad \square$

### A.3. Proof of Theorem 2

We first bound the covering number for a general feed-forward neural network and specialize it to the three architectures we present in Section 3.

**Lemma 4.** *Let $\mathcal{F}_k$ be a class of feed-forward neural networks that maps an input vector $x \in \mathbb{R}^{d_0}$ to an output vector $y \in \mathbb{R}^{d_k}$, with each layer $\ell$ containing $T_\ell$ nodes and computing $z \mapsto \phi_\ell(w^\ell z)$, where each $w^\ell \in \mathbb{R}^{T_\ell \times T_{\ell-1}}$ and $\phi_\ell : \mathbb{R}^{T_\ell} \to [-B, +B]^{T_\ell}$. Further let, for each network in $\mathcal{F}_k$, let the parameter matrices $\|w^\ell\|_1 \leq W$ and $\|\phi_\ell(s) - \phi_\ell(s')\|_1 \leq \Phi \|s - s'\|_1$ for any $s, s' \in \mathbb{R}^{T_{\ell-1}}$.*

$$
\mathcal{N}_\infty(\mathcal{F}_k, \epsilon) \leq \left\lceil \frac{2Bd^2 W (2\Phi W)^k}{\epsilon} \right\rceil^d,
$$

*where $T = \max_{\ell \in [k]} T_\ell$ and $d$ is the total number of parameters in a network.*

*Proof.* We shall construct an $\ell_{1,\infty}$ cover for $\mathcal{F}_k$ by discretizing each of the $d$ parameters along $[-W, +W]$ at scale $\epsilon_0/d$, where we will choose $\epsilon_0 > 0$ at the end of the proof. We will use $\hat{\mathcal{F}}_k$ to denote the subset of neural networks in $\mathcal{F}_k$ whose parameters are in the range $\{-(\lceil Wd/\epsilon_0 \rceil - 1)\,\epsilon_0/d, \ldots, -\epsilon_0/d, 0, \epsilon_0/d, \ldots, \lceil Wd/\epsilon_0 \rceil \epsilon_0/d\}$. Note that size of $\hat{\mathcal{F}}_k$ is at most $\lceil 2dW/\epsilon_0 \rceil^d$. We shall now show that $\hat{\mathcal{F}}_k$ is an $\epsilon$-cover for $\mathcal{F}_k$.

We use mathematical induction on the number of layers $k$. We wish to show that for any $f \in \mathcal{F}_k$ there exists a $\hat{f} \in \hat{\mathcal{F}}_k$ such that:

$$\|f(x) - \hat{f}(x)\|_1 \leq Bd\epsilon_0(2\Phi W)^k.$$

Note that for $k = 0$, the statement holds trivially. Assume that the statement is true for $\mathcal{F}_k$. We now show that the statement holds for $\mathcal{F}_{k+1}$.

A function $f \in \mathcal{F}_{k+1}$ can be written as $f(z) = \phi_{k+1}(w_{k+1}H(z))$ for some $H \in \mathcal{F}_k$. Similarly, a function $\hat{f} \in \hat{\mathcal{F}}_{k+1}$ can be written as $\hat{f}(z) = \phi_{k+1}(\hat{w}_{k+1}\hat{H}(z))$ for some $\hat{H} \in \hat{\mathcal{F}}_k$ and $\hat{w}_{k+1}$ is a matrix of entries in $\{-(\lceil Wd/\epsilon_0 \rceil - 1)\,\epsilon_0/d, \ldots, -\epsilon_0/d, 0, \epsilon_0/d, \ldots, \lceil Wd/\epsilon_0 \rceil \epsilon_0/d\}$. Also note that for any parameter matrix $w^\ell \in \mathbb{R}^{T_\ell \times T_{\ell-1}}$, there is a matrix $\hat{w}^\ell$ with discrete entries s.t.

$$\|w_\ell - \hat{w}_\ell\|_1 = \max_{1 \leq j \leq T_{\ell-1}} \sum_{i=1}^{T_\ell} |w^\ell_{\ell,i,j} - \hat{w}_{\ell,i,j}| \leq T_\ell \epsilon_0/d \leq \epsilon_0. \tag{7}$$

We then have:

$$
\begin{aligned}
\|f(x) - \hat{f}(x)\|_1 &= \|\phi_{k+1}(w_{k+1}H(x)) - \phi_{k+1}(\hat{w}_{k+1}\hat{H}(x))\|_1 \\
&\leq \Phi\|w_{k+1}H(x) - \hat{w}_{k+1}\hat{H}(x)\|_1 \\
&\leq \Phi\|w_{k+1}H(x) - w_{k+1}\hat{H}(x)\|_1 + \Phi\|w_{k+1}\hat{H}(x) - \hat{w}_{k+1}\hat{H}(x)\|_1 \\
&\leq \Phi\|w_{k+1}\|_1 \cdot \|H(x) - \hat{H}(x)\|_1 + \Phi\|w_{k+1} - \hat{w}_{k+1}\|_1 \cdot \|\hat{H}(x)\|_1 \\
&\leq \Phi W\|H(x) - \hat{H}(x)\|_1 + \Phi B\|w_{k+1} - \hat{w}_{k+1}\|_1 \\
&\leq Bd\epsilon_0 \Phi W(2\Phi W)^k + \Phi Bd\epsilon_0 \\
&\leq Bd\epsilon_0(2\Phi W)^{k+1},
\end{aligned}
$$

where the second line follows from our assumption on $\phi_{k+1}$, and the sixth line follows from our inductive hypothesis and from (7). By choosing $\epsilon_0 = \frac{\epsilon}{B(2\Phi W)^k}$, we complete the proof. $\qquad\square$

We next bound the covering number of the mechanism class in terms of the covering number for the class of allocation networks and for the class of payment networks. Recall that the payment networks computes a fraction $\alpha : \mathbb{R}^{m(n+1)} \to [0,1]^n$ and computes a payment $p_i(b) = \alpha_i(b) \cdot \langle v_i, g_i(b) \rangle$ for each bidder $i$. Let $\mathcal{G}$ be the class of allocation networks and $\mathcal{A}$ be the class of fractional payment functions used to construct auctions in $\mathcal{M}$. Let $\mathcal{N}_\infty(\mathcal{G}, \epsilon)$ and $\mathcal{N}_\infty(\mathcal{A}, \epsilon)$ be the corresponding covering numbers w.r.t. the $\ell_\infty$ norm. Then:

**Lemma 5.** $\mathcal{N}_\infty(\mathcal{M}, \epsilon) \leq \mathcal{N}_\infty(\mathcal{G}, \epsilon/3) \cdot \mathcal{N}_\infty(\mathcal{A}, \epsilon/3)$.

*Proof.* Let $\hat{\mathcal{G}} \subseteq \mathcal{G}$, $\hat{\mathcal{A}} \subseteq \mathcal{A}$ be $\ell_\infty$ covers for $\mathcal{G}$ and $\mathcal{A}$, i.e. for any $g \in \mathcal{G}$ and $\alpha \in \mathcal{A}$, there exists $\hat{g} \in \hat{\mathcal{G}}$ and $\hat{\alpha} \in \hat{\mathcal{A}}$ with

$$\sup_b \sum_{i,j} |g_{ij}(b) - \hat{g}_{ij}(b)| \leq \epsilon/3, \text{ and} \tag{8}$$

$$\sup_b \sum_i |\alpha_i(b) - \hat{\alpha}_i(b)| \leq \epsilon/3. \tag{9}$$

We now show that the class of mechanism $\hat{\mathcal{M}} = \{(\hat{g}, \hat{\alpha}) \mid \hat{g} \in \hat{\mathcal{G}}, \text{ and } \hat{p}(b) = \hat{\alpha}_i(b) \cdot \langle v_i, \hat{g}_i(b) \rangle\}$ is an $\epsilon$-cover for $\mathcal{M}$ under the $\ell_{1,\infty}$ distance. For any mechanism in $(g, p) \in \mathcal{M}$, let $(\hat{g}, \hat{p}) \in \hat{\mathcal{M}}$ be a mechanism in $\hat{\mathcal{M}}$ that satisfies (9). We have:

$$\sum_{i,j} |g_{ij}(b) - \hat{g}_{ij}(b)| + \sum_i |p_i(b) - \hat{p}_i(b)|$$

$$\leq \epsilon/3 + \sum_i |\alpha_i(b) \cdot \langle b_i, g_{i,\cdot}(b) \rangle - \hat{\alpha}_i(b) \cdot \langle b_i, \hat{g}_i(b) \rangle|$$

$$\leq \epsilon/3 + \sum_i |(\alpha_i(b) - \hat{\alpha}_i(b)) \cdot \langle b_i, g_i(b) \rangle| + |\hat{\alpha}_i(b) \cdot (\langle b_i, g_i(b) \rangle - \langle b_i, \hat{g}_{i,\cdot}(b) \rangle)|$$

$$\leq \epsilon/3 + \sum_i |\alpha_i(b) - \hat{\alpha}_i(b)| + \sum_i \|b_i\|_\infty \cdot \|g_i(b) - \hat{g}_i(b)\|_1$$

$$\leq 2\epsilon/3 + \sum_{i,j} |g_{ij}(b) - \hat{g}_{ij}(b)| \leq \epsilon,$$

where in the third inequality we use $\langle b_i, g_i(b) \rangle \leq 1$. The size of the cover $\hat{\mathcal{M}}$ is $|\hat{\mathcal{G}}||\hat{\mathcal{A}}|$, which completes the proof. $\square$

We are now ready to prove covering number bounds for the three architectures in Section 3.

*Proof of Theorem 2.* All three architectures use the same feed-forward architecture for computing fractional payments, consisting of $K$ hidden layers with tanh activation functions. We also have by our assumption that the $L_1$ norm of the vector of all model parameters is at most $W$, for each $\ell = 1, \ldots, R+1$, $\|w_\ell\|_1 \leq W$. Using that fact that the tanh activation functions are 1-Lipschitz and bounded in $[-1, 1]$, and there are at most $\max\{K, n\}$ number of nodes in any layer of the payment network, we have by an application of Lemma 4 the following bound on the covering number of the fractional payment networks $\mathcal{A}$ used in each case:

$$\mathcal{N}_\infty(\mathcal{A}, \epsilon) \leq \left\lceil \frac{\max(K, n)^2 (2W)^{R+1}}{\epsilon} \right\rceil^{d_p},$$

where $d_p$ is the number of parameters in payment networks.

For the covering number of allocation networks $\mathcal{G}$, we consider each architecture separately. In each case, we bound the Lipschitz constant for the activation functions used in the layers of the allocation network and followed by an application of Lemma 4. For ease of exposition, we omit the dummy scores used in the final layer of neural network architectures.

**Additive bidders.** The output layer computes $n$ allocation probabilities for each item $j$ using a softmax function. The activation function $\phi_{R+1} : \mathbb{R}^n \to \mathbb{R}^n$ for the final layer for input $s \in \mathbb{R}^{n \times m}$ can be described as: $\phi_{R+1}(s) = [\text{softmax}(s_{1,1}, \ldots, s_{n,1}), \ldots, \text{softmax}(s_{1,m}, \ldots, s_{n,m})]$, where $\text{softmax} : \mathbb{R}^n \to [0, 1]^n$ is defined for any $u \in \mathbb{R}^n$ as $\text{softmax}_i(u) = \exp(u_i)/\sum_{k=1}^n \exp(u_k)$.

We then have for any $s, s' \in \mathbb{R}^{n \times m}$,

$$\|\phi_{R+1}(s) - \phi_{R+1}(s')\|_1 = \sum_j \left\| \text{softmax}(s_{1,j}, \ldots, s_{n,j}) - \text{softmax}(s'_{1,j}, \ldots, s'_{n,j}) \right\|_1$$

$$\leq \sqrt{n} \sum_j \left\| \text{softmax}(s_{1,j}, \ldots, s_{n,j}) - \text{softmax}(s'_{1,j}, \ldots, s'_{n,j}) \right\|_2$$

$$\leq \sqrt{n} \frac{\sqrt{n-1}}{n} \sum_j \sqrt{\sum_i \|s_{ij} - s'_{ij}\|^2}$$

$$\leq \sum_j \sum_i |s_{ij} - s'_{ij}|, \tag{10}$$

where the third step follows by bounding the Frobenius norm of the Jacobian of the softmax function.

The hidden layers $\ell = 1, \ldots, R$ are standard feed-forward layers with tanh activations. Since the tanh activation function is 1-Lipschitz, $\|\phi_\ell(s) - \phi_\ell(s')\|_1 \leq \|s - s'\|_1$. We also have by our assumption that the $L_1$ norm of the vector of all model parameters is at most $W$, for each $\ell = 1, \ldots, R+1$, $\|w_\ell\|_1 \leq W$. Moreover, the output of each hidden layer node is in $[-1, 1]$, the output layer nodes is in $[0, 1]$, and the maximum number of nodes in any layer (including the output layer) is at most $\max\{K, mn\}$.

By an application of Lemma 4 with $\Phi = 1$, $B = 1$ and $d = \max K, mn$, we have

$$\mathcal{N}_\infty(\mathcal{G}, \epsilon) \leq \left\lceil \frac{\max\{K, mn\}^2 (2W)^{R+1}}{\epsilon} \right\rceil^{d_a},$$

where $d_a$ is the number of parameters in allocation networks.

**Unit-demand bidders.** The output layer $n$ allocation probabilities for each item $j$ as an element-wise minimum of two softmax functions. The activation function $\phi_{R+1} : \mathbb{R}^2 n \to \mathbb{R}^n$ for the final layer for two sets of scores $s, \bar{s} \in \mathbb{R}^{n \times m}$ can be described as:

$$\phi_{R+1,i,j}(s, s') = \min\{\text{softmax}_j(s_{i,1}, \ldots, s_{i,m}), \text{softmax}_i(s'_{1,j}, \ldots, s'_{n,j})\}.$$

We then have for any $s, \tilde{s}, s', \tilde{s}' \in \mathbb{R}^{n \times m}$,

$$
\begin{aligned}
\|\phi_{R+1}(s, \tilde{s}) - \phi_{R+1}(s', \tilde{s}')\|_1 &= \sum_{i,j} \Big| \min\{\text{softmax}_j(s_{i,1}, \ldots, s_{i,m}), \text{softmax}_i(\tilde{s}_{1,j}, \ldots, \tilde{s}_{n,j})\} \\
&\qquad - \min\{\text{softmax}_j(s'_{i,1}, \ldots, s'_{i,m}), \text{softmax}_i(\tilde{s}'_{1,j}, \ldots, \tilde{s}'_{n,j})\} \Big| \\
&\leq \sum_{i,j} \Big| \max\{\text{softmax}_j(s_{i,1}, \ldots, s_{i,m}) - \text{softmax}_j(s'_{i,1}, \ldots, s'_{i,m}), \\
&\qquad\qquad \text{softmax}_i(\tilde{s}_{1,j}, \ldots, \tilde{s}_{n,j}) - \text{softmax}_i(\tilde{s}'_{1,j}, \ldots, \tilde{s}'_{n,j})\} \Big| \\
&\leq \sum_i \big\| \text{softmax}(s_{i,1}, \ldots, s_{i,m}) - \text{softmax}(s'_{i,1}, \ldots, s'_{i,m}) \big\|_1 \\
&\qquad + \sum_j \big\| \text{softmax}(\tilde{s}_{1,j}, \ldots, \tilde{s}_{n,j}) - \text{softmax}(\tilde{s}'_{1,j}, \ldots, \tilde{s}'_{n,j})\} \big\|_1 \\
&\leq \sum_{i,j} |s_{ij} - s'_{ij}| + \sum_{i,j} |\tilde{s}_{ij} - \tilde{s}'_{ij}|,
\end{aligned}
$$

where the last step can be derived in the same way as (10).

As with additive bidders, using additionally hidden layers $\ell = 1, \ldots, R$ are standard feed-forward layers with tanh activations, we have from Lemma 4 with $\Phi = 1$, $B = 1$ and $d = \max\{K, mn\}$,

$$\mathcal{N}_\infty(\mathcal{G}, \epsilon) \leq \left\lceil \frac{\max\{K, mn\}^2 (2W)^{R+1}}{\epsilon} \right\rceil^{d_a}.$$

**Combinatorial bidders.** The output layer outputs an allocation probability for each bidder $i$ and bundle of items $S \subseteq M$. The activation function $\phi_{R+1} : \mathbb{R}^{(m+1)n2^m} \to \mathbb{R}^{n2^m}$ for this layer for $m + 1$ sets of scores $s, s^{(1)}, \ldots, s^{(m)} \in \mathbb{R}^{n \times 2^m}$ is given by:

$$
\begin{aligned}
&\phi_{R+1,i,S}(s, s^{(1)}, \ldots, s^{(m)}) \\
&= \min\Big\{ \text{softmax}_S(s_{i,S'} : S' \subseteq M), \text{softmax}_S(s^{(1)}_{i,S'} : S' \subseteq M), \ldots, \text{softmax}_S(s^{(m)}_{i,S'} : S' \subseteq M) \Big\},
\end{aligned}
$$

where $\text{softmax}_S(a_{S'} : S' \subseteq M) = \exp(a_S) / \sum_{S' \subseteq M} \exp(a_{S'})$.

We then have for any $s, s^{(1)}, \ldots, s^{(m)}, s', s'^{(1)}, \ldots, s'^{(m)} \in \mathbb{R}^{n \times 2^m}$,

$$
\begin{aligned}
&\|\phi_{R+1}(s, s^{(1)}, \ldots, s^{(m)}) - \phi_{R+1}(s', s'^{(1)}, \ldots, s'^{(m)})\|_1 \\
&= \sum_{i,S} \Big| \min\Big\{ \text{softmax}_S(s_{i,S'} : S' \subseteq M), \text{softmax}_S(s^{(1)}_{i,S'} : S' \subseteq M), \ldots, \text{softmax}_S(s^{(m)}_{i,S'} : S' \subseteq M) \Big\} \\
&\qquad - \min\Big\{ \text{softmax}_S(s'_{i,S'} : S' \subseteq M), \text{softmax}_S(s'^{(1)}_{i,S'} : S' \subseteq M), \ldots, \text{softmax}_S(s'^{(m)}_{i,S'} : S' \subseteq M) \Big\} \Big| \\
&\leq \sum_{i,S} \max\Big\{ \big| \text{softmax}_S(s_{i,S'} : S' \subseteq M) - \text{softmax}_S(s'_{i,S'} : S' \subseteq M) \big|, \\
&\qquad\qquad \big| \text{softmax}_S(s^{(1)}_{i,S'} : S' \subseteq M) - \text{softmax}_S(s'^{(1)}_{i,S'} : S' \subseteq M) \big|, \ldots \\
&\qquad\qquad \big| \text{softmax}_S(s^{(m)}_{i,S'} : S' \subseteq M) - \text{softmax}_S(s'^{(m)}_{i,S'} : S' \subseteq M) \big| \Big\} \\
&\leq \sum_i \big\| \text{softmax}(s_{i,S'} : S' \subseteq M) - \text{softmax}(s'_{i,S'} : S' \subseteq M) \big\|_1 \\
&\qquad + \sum_{i,j} \big\| \text{softmax}(s^{(j)}_{i,S'} : S' \subseteq M) - \text{softmax}(s'^{(j)}_{i,S'} : S' \subseteq M) \big\|_1
\end{aligned}
$$

| Distretization | Number of decision variables | Number of constraints |
|---|---|---|
| 5 bins/value | $1.25 \times 10^5$ | $3.91 \times 10^6$ |
| 6 bins/value | $3.73 \times 10^5$ | $2.02 \times 10^7$ |
| 7 bins/value | $9.41 \times 10^5$ | $8.07 \times 10^7$ |

**Table 2:** Number of decision variables and constraints of LP with different discretizations for a 2 bidder, 3 items setting with uniform valuations.

$$\leq \sum_{i,S} |s_{i,S} - s'_{i,S}| + \sum_{i,j,S} |s_{i,S}^{(j)} - s_{i,S}'^{(j)}|,$$

where the last step can be derived in the same way as (10).

As with additive bidders, using additionally hidden layers $\ell = 1, \ldots, R$ are standard feed-forward layers with tanh activations, we have from Lemma 4 with $\Phi = 1$, $B = 1$ and $d = \max\{K, n \cdot 2^m\}$

$$\mathcal{N}_\infty(\mathcal{G}, \epsilon) \leq \left\lceil \frac{\max\{K, n \cdot 2^m\}^2 (2W)^{R+1}}{\epsilon} \right\rceil^{d_a},$$

where $d_a$ is the number of parameters in allocation networks. $\qquad\square$

We now bound $\Delta_L$ for the three architectures using the covering number bounds we derived above. In particular, we upper bound the the 'inf' over $\epsilon > 0$ by substituting a specific value of $\epsilon$:

(a) For additive bidders, choosing $\epsilon = \frac{1}{\sqrt{L}}$, we get $\Delta_L \leq O\left(\sqrt{R(d_p + d_a)\frac{\log(W \max\{K, mn\}L)}{L}}\right)$.

(b) For unit-demand bidders, choosing $\epsilon = \frac{1}{\sqrt{L}}$, we get $\Delta_L \leq O\left(\sqrt{R(d_p + d_a)\frac{\log((W \max\{K, mn\}L)}{L}}\right)$.

(c) For combinatorial bidders, choosing $\epsilon = \frac{1}{\sqrt{L}}$, we get $\Delta_L \leq O\left(\sqrt{R(d_p + d_a)\frac{\log(W \max\{K, n \cdot 2^m\}L)}{L}}\right)$.

## B. Omitted Details in Experiments

In this section, we show more details of the experiments in this paper.

**Discussion on size of LP.** First, we provide more evidence about the efficiency of our RegretNet compared with LP. As mentioned in (Conitzer & Sandholm, 2002), the number of decision variables and constraints are exponential in the number of bidders and items. We consider the setting with $n$ additive bidders and $m$ items and the value is divided into $D$ bins per item. There are $D^{mn}$ valuation profiles in total, each involving $(n + nm)$ variables ($n$ payments and $nm$ allocation probabilities). For the constraints, there are $n$ IR constraints (for $n$ bidders) and $n \cdot (D^m - 1)$ IC constraints (for each bidder, there are $(D^m - 1)$ constraints) for each valuation profile. In addition, there are $n$ bidder-wise and $m$ item-wise allocation constraints. In Table 2, we show the explosion of decision variables and constraints with finer discretization of the valuations for 2 bidders, 3 items setting. As we can see, the decision variables and constraints blow up extremely fast, even for a small setting with a coarse discretization over value.

**Additional discussion of experiments.** For small settings (I)–(V), we get similar performance as in Figure 3 with smaller training samples (around 5000). ReLU activations yield comparable results for smaller settings (I)–(V), but tanh works better for larger settings (VI)–(VII). Our RegretNet is scalable for auctions with more bidders and items. A single iteration of augmented Lagrangian took on an average 1–17 seconds across experiments. Even for the larger settings (VI)–(VII), the running time of our algorithm was less than 13 hours. For the settings (VI)–(VII) for which the optimal auction is *not* known, we also compare with a Myerson auction to sell the entire bundle of items as one unit, which is optimal in the limit of number of items (Palfrey, 1983).

| Distribution | Opt | RegretNet | |
|---|---|---|---|
| | rev | rev | rgt |
| Setting (a): $v_1 \sim [4, 16], v_2 \sim U[4, 7]$ | 9.781 | 9.734 | < 0.001 |
| Setting (b): $v_1, v_2$ drawn uniformly from a unit triangle | 0.388 | 0.392 | < 0.001 |
| Setting (c): $v_1, v_2 \sim U[0, 1]$ | 0.384 | 0.384 | < 0.001 |

**Table 3:** Revenue of auctions for single additive bidder, two items obtained with RegretNet.



**Figure 6:** Allocation rule learned by RegretNet for (a) the single additive bidder, two items setting with values $v_1 \sim U[4, 16]$ and $v_2 \sim U[4, 7]$, and for (b) the single additive bidder, two items setting with values $v_1, v_2$ drawn jointly, uniformly from a triangle with vertices $(0, 0)$, $(0, 1)$ and $(1, 0)$, The optimal mechanisms due to (Daskalakis et al., 2017) for (a) and (Haghpanah & Hartline, 2015) for (b) are described by the regions separated by the dashed orange lines. The numbers in orange are the probability the item is allocated in a region.



**Figure 7:** Allocation rule learned by RegretNet for (a) the single unit-demand bidder, two items setting with values $v_1, v_2 \sim U[0, 1]$ (optimal mechanism due to (Pavlov, 2011)), and for (b) the single additive bidder, two items setting with values $v_1 \sim U[0, 4], v_2 \sim U[0, 3]$. The subset of valuations $(v_1, v_2)$ where the bidder receives neither item looks like a pentagonal shape.

| Distribution | Item-wise Myerson | Bundled Myerson | RegretNet | |
|---|---|---|---|---|
| | rev | rev | rev | rgt |
| Setting (d): $v_i \sim U[0, 1]$ | 2.495 | 3.457 | **3.461** | < 0.003 |
| Setting (e): $v_1 \sim U[0, 4], v_2 \sim U[0, 3]$ | 1.877 | 1.749 | **1.911** | < 0.001 |

**Table 4:** Revenue of auctions for single additive bidder, 10 items obtained with RegretNet and single additive bidder, 2 items with $v_1 \sim U[0, 4], v_2 \sim U[0, 3]$.

| Distribution | Ascending auction | RegretNet | |
|---|---|---|---|
| | rev | rev | rgt |
| Setting (f): $v_1, v_2 \sim U[0, 1]$ | 0.179 | **0.706** | < 0.001 |

**Table 5:** Revenue of auctions for 2 unit-demand bidders, 2 items obtained with RegretNet. For the ascending auction, the price were raised in units of 0.3 (which was empirically tuned using a grid search.)

# C. Additional Experiments

In this section, we show the additional experiments for both the single bidder case and the mulitple bidders case. We consider the following settings:

(a) Single additive bidder with preferences over two non-identically distributed items, where $v_1 \sim U[4, 16]$ and $v_2 \sim U[4, 7]$.

(b) Single additive bidder with preferences over two items, where $(v_1, v_2)$ are drawn jointly and uniformly from a unit triangle with vertices $(0, 0), (0, 1)$ and $(1, 0)$.

(c) Single unit-demand bidder with preferences over two items, where the item values $v_1, v_2 \sim U[0, 1]$,

(d) Single additive bidder with preferences over ten items, where each $v_i \sim U[0, 1]$.

(e) Single additive bidder with preferences over two items, where the item values $v_1 \sim U[0, 4], v_2 \sim U[0, 3]$,

(f) Two unit-demand bidders and two items, where the bidders draw their value for each item from identical uniform distributions over $[0, 1]$.

For setting (a), we show our *RegretNet* almost exactly recovers the optimal mechanism of (Daskalakis et al., 2017). For setting (b), we show that the approach almost exactly recovers the optimal mechanism of (Haghpanah & Hartline, 2015). For setting (c), we show that the approach almost exactly recovers the optimal mechanisms of (Pavlov, 2011). For settings (a), (b), (c), we show our results in Table 3, and we show the allocation plots for the three settings above in Figure 6 and Figure 7. To our knowledge, an analytical solution for the optimal mechanism for setting (d) is not available (Daskalakis, 2015). Here our approach finds a new mechanism that has higher revenue than both a Myerson auction on each item and a Myerson on the entire bundle, we show it in Table 4. For setting (e), we plot the allocation figures in Figure 7 and test the performance of our *RegretNet* compared with Myerson auction on each item and Myerson auction on the entire bundle in Table 4. For setting (f), the optimal auction is again not known; we show in Table 5 that the learned auctions beat reasonable baseline mechanisms.