APPENDIX

The Appendix is structured as follows: Section A contains the missing proofs, Section B contains the result of the applicability of our techniques for Stackelberg games, Section C constains results about the sample complexity of standard SUQR, Section D contains the weaker sample complexity bound result for the generalized SUQR model derived using the approach of Haussler and Section E contains additional experiments.

A. PROOFS

Proof of Theorem 1

PROOF. First, Haussler uses the following pseudo metric ρ on A that is defined using the loss function l:

$$\rho(a,b) = \max_{y\in Y} |l(y,a) - l(y,b)|.$$

To start with, relying on Haussler's result, we show

$$Pr(\forall h \in \mathcal{H}. |\hat{r}_h(\vec{z}) - r_h(p)| < \frac{\alpha}{3}) \geq 1 - 4\mathcal{C}\Big(\frac{\alpha}{48}, \mathcal{H}, \rho\Big) e^{-\frac{\alpha^2 m}{576M^2}}$$

Choose $\alpha = \alpha'/4M$ and $\nu = 2M$ in Theorem 9 of [14]. Using property (3) (Section 2.2, [14]) of d_v we obtain $|r - s| \leq \epsilon$ whenever $d_v(r,s) \leq \alpha'$. Using this directly in Theorem 9 of Haussler [14] we obtain the desired result above.

Note the dependence of the above probability on m (the number of samples), and compare it to the first pre-condition in the PAC learning result. By equating $\delta/2$ to $4\mathcal{C}(\alpha/48, \mathcal{H}, \rho)e^{-\frac{\alpha^2 m}{576M^2}}$, we derive the sample complexity as

$$m \geq \frac{576M^2}{\alpha^2} \log \frac{8\mathcal{C}(\alpha/48, \mathcal{H}, \rho)}{\delta}$$

We wish to compute a bound on $\mathcal{C}(\epsilon, \mathcal{H}, \rho)$ in order to use the above result to obtain sample complexity. First, we prove that $\rho \leq 2T d_{\bar{l_1}}$ for the loss function we use. This result is used to bound $\mathcal{C}(\epsilon, \mathcal{H}, \rho)$, since, it is readily verified from definition that $\mathcal{C}(\epsilon, \mathcal{H}, \rho) \leq \mathcal{C}(\epsilon/2T, \mathcal{H}, d_{\bar{l_1}})$. Such a bounding directly gives

$$m \geq \frac{576M^2}{\alpha^2} \log \frac{8\mathcal{C}(\alpha/96T, \mathcal{H}, \rho)}{\delta}$$

Below we prove that $\rho \leq 2T d_{\bar{l_1}}$.

LEMMA 8. *Given the loss function defined above, we have* $\rho(a,b) \leq 2\max_i |a_i - b_i| \leq 2\sum_i |a_i - b_i| \leq 2T d_{\bar{l_1}}(a,b)$

PROOF. By definition, $\rho(a,b) = \max_i \Big| - a_i + b_i + \log \frac{1+\sum_{i=1}^{T-1} e^{a_i}}{1+\sum_{i=1}^{T-1} e^{b_i}} \Big| \leq \max_i |a_i - b_i| + \Big| \log \frac{1+\sum_{i=1}^{T-1} e^{a_i}}{1+\sum_{i=1}^{T-1} e^{b_i}} \Big|$. There is j and k such that $max_r = \frac{e^{a_j}}{e^{b_j}} \geq \frac{e^{a_i}}{e^{b_i}}$ for all i and $min_r = \frac{e^{a_k}}{e^{b_k}} \leq \frac{e^{a_i}}{e^{b_i}}$ for all i. Thus,

$$\log \frac{1 + min_r t}{1+t} \leq \log \frac{1+\sum_{i=1}^{T-1} e^{a_i}}{1+\sum_{i=1}^{T-1} e^{b_i}} \leq \log \frac{1 + max_r t}{1+t}$$

where $t = \sum_{i=1}^{T-1} e^{b_i}$. The greatest positive value of the RHS is $\log max_r \leq |a_j - b_j|$ and least negative value possible for LHS is $\log min_r \geq -|a_k - b_k|$. Thus,

$$\Big| \log \frac{1+\sum_{i=1}^{T-1} e^{a_i}}{1+\sum_{i=1}^{T-1} e^{b_i}} \Big| \leq \max_i \big|a_i - b_i\big|$$

Hence, we obtain $\rho(a,b) = \max_i |l(y_i,a) - l(y_i,b)| \leq 2\max_i |a_i - b_i|$, and the last inequality is trivial. □

Thus, using the above result we get

$$m \geq \frac{576M^2}{\alpha^2} \log \frac{8\mathcal{C}(\alpha/96T, \mathcal{H}, d_{\bar{l_1}})}{\delta}$$

Proof of Lemma 2

PROOF. First, note that $x_{iT} = x_i - x_T$ lies between $[-1,1]$ due to the constraints on x_i, x_T. Then, for any two functions $g, g' \in \mathcal{G}$ we have the following result:

$$\begin{aligned} & d_{L^1(P, d_{\bar{l_1}})}(g,g') = \\ &= \int_X \frac{1}{T-1} \sum_{i=1}^{T-1} d_{l_1}(w(x_i - x_T), w'(x_i - x_T))\, dP(x) \\ &= \int_X \frac{1}{T-1} \sum_{i=1}^{T-1} |(w - w')(x_i - x_T)|\, dP(x) \\ &\leq \int_X \frac{1}{T-1} \sum_{i=1}^{T-1} |(w - w')|\, dP(x) \;=\; |(w - w')| \end{aligned}$$

Also, note that since the range of any $g = w(x_i - x_T)$ is $[-\frac{M}{4}, \frac{M}{4}]$ and given $x_i - x_T$ lies between $[-1,1]$, we can claim that w lies between $[-\frac{M}{4}, \frac{M}{4}]$. Thus, given the distance between functions is bounded by the difference in weights, it enough to divide the $M/2$ range of the weights into intervals of size 2ϵ and consider functions at the boundaries. Hence the ϵ-cover has at most $M/4\epsilon$ functions.

The proof for constant valued functions $\mathcal{F}_i$ is similar, since its straightforward to see the distance between two functions in this space is the difference in the constant output. Also, the constants lie in $[-\frac{M}{4}, \frac{M}{4}]$, Then, the argument is same as the $\mathcal{G}$ case. □

Proof of Lemma 3

PROOF. First, the space of functions $\hat{\mathcal{H}} = \{h/\hat{K} \mid h \in \mathcal{H}_i\}$ is Lipschitz with Lipschitz constant ≤ 1 and $|h_i(x)| \leq M/2\hat{K}$. Clearly $\mathcal{N}(\epsilon, \mathcal{H}_i, d_{l_\infty}) \leq \mathcal{N}(\epsilon/\hat{K}, \hat{\mathcal{H}}, d_{l_\infty})$. Using the following result from [32]: for any Lipschitz real valued function space $\mathcal{H}$ with constant 1, any positive integer s and any distance d

$$\mathcal{N}(\epsilon, \mathcal{H}, d_{l_\infty}) \leq \Big(2\Big\lceil \frac{M(s+1)}{2\hat{K}\epsilon} \Big\rceil + 1\Big) \cdot (s+1)^{\mathcal{N}(\frac{s\epsilon}{s+1}, X, d)}$$

Then, we get the bound on $\mathcal{N}(\epsilon/\hat{K}, \hat{\mathcal{H}}, d_{l_\infty})$ by choosing $s = 1$ and $d = d_{l_\infty}$, and hence obtain the desired bound on $\mathcal{N}(\epsilon, \mathcal{H}_i, d_{l_\infty})$. □

Proof of Lemma 4

PROOF. For ease of notation, we do the proof with k standing for $K+1$. Let $Y_i = U_i - 0.5$, then $|Y| \leq 1/2$ and $S_T - 0.5T = \sum_i Y_i$. Using Bernstein's inequality with the fact that $E[Y_i^2] = 1/12$

$$P(\sum_i Y_i = S_T - 0.5T \leq -t) \leq e^{\frac{-0.5t^2}{T/12 + t/6}}$$

Thus, $P(S_T \leq 0.5T - t) \leq e^{\frac{0.5t^2}{T/12+t/6}}$. Take $k = 0.5T - t$, and hence $t = 0.5T - k = T(0.5 - k/T)$. Hence,

$$P(S_T \leq k) \leq e^{\frac{-3T(0.5-k/T)^2}{1-k/T}}$$

□

Proof of Theorem 3

PROOF. Given the results of Lemma 3, we get the sample complexity is of order

$$\frac{1}{\alpha^2}\Big(\log\frac{1}{\delta} + T\big(\mathcal{N}(\frac{\alpha}{T}, X, d_{l_1})\big)\Big)$$

Now, suing result of Lemma 4, we get the required order in the Theorem. We wish to note that if K/T is a constant then the $O(e^{-T})$ in Lemma 4 gets swamped by the T^T term. However, in practice for fixed T, this term does provide lower actual complexity bound than what is indicated by the order. $\square$

Proof of Lemma 5

PROOF. Observe that due to the definition of K^* any solution to MinLip will have Lipschitz constant $\geq K^*$. Thus, it suffices to show that the Lipschitz constant of h_i is K^*, to prove that h_i is a solution of MinLip. Take any two x, x'. If the min in the expression for h_i occurs for the same j for both x, x' then $|h_i(x) - h_i(x')|$ is given by $K^*|||x - x^j||_1 - ||x' - x^j||_1|$. By application of triangle inequality

$$-||x - x'||_1 \leq ||x - x^j||_1 - ||x' - x^j||_1 \leq ||x - x'||_1$$

Thus, $|h_i(x) - h_i(x')| \leq K^*||x - x'||_1$.

For the other case when the min for x occurs at some j and min for x' at some j' we have the following: $h_i(x') = h_{ij'} + K^*||x' - x^{j'}||_1$ and $h_i(x) = h_{ij} + K^*||x - x^{j'}||_1$. Also, due to the min, $h_i(x') \leq h_{ij} + K^*||x' - x^j||_1 = h_i(x) + K^*||x' - x^j||_1 - K^*||x - x^{j'}||_1$. Thus, we get

$$h_i(x') - h_i(x) \leq K^*(||x' - x^j||_1 - ||x - x^j||_1) \leq K^*||x' - x||_1$$

Using the symmetric case inequality for x we get

$$h_i(x) - h_i(x') \leq K^*(||x - x^j||_1 - ||x' - x^j||_1) \leq K^*||x - x'||_1$$

Combining both these we can claim that $|h_i(x) - h_i(x')| \leq K^*||x' - x||_1$. Thus, we have proved that h_i is K^* Lipschitz, and hence a solution of MinLip. $\square$

Proof of Lemma 6

PROOF. Let p_X be the marginal of $p(x, y)$ for space X. Define the expected entropy $E[H(x)] = \int p_X(x) \sum_{i=1}^T I_{y=\mathrm{t}_i} q_i^p(x) \log q_i^p(x)\, dx$. Given the loss function, we know that $r_h(p) = -\int p(x, y) \sum_{i=1}^T I_{y=\mathrm{t}_i} \log q_i^h(x)\, dx\, dy$. This is same as $-\int p_X(x) \sum_{i=1}^T I_{y=\mathrm{t}_i} q_i^p(x) \sum_{i=1}^T I_{y=\mathrm{t}_i} \log q_i^h(x)\, dx\, dy$. This reduces to $-\int p_X(x) \sum_{i=1}^T I_{y=\mathrm{t}_i} q_i^p(x) \log q_i^h(x)\, dx\, dy$. Thus, we have

$$E[H(x)] + r_h(p) = \int p_X(x) \sum_{i=1}^T I_{y=\mathrm{t}_i} q_i^p(x) \log \frac{q_i^p(x)}{q_i^h(x)}\, dx\, dy$$

Hence, we obtain

$$E[H(x)] + r_h(p) = E[\mathrm{KL}(q^p(x) \,||\, q^h(x))]$$

Hence, $|r_h(p) - r_{h^*}(p)|$ is equal to

$$|E[\mathrm{KL}(q^p(x) \,||\, q^h(x))] - E[\mathrm{KL}(q^p(x) \,||\, q^*(x))]|$$

Thus, from the assumptions, we get $E[\mathrm{KL}(q^p(x) \,||\, q^h(x))] \leq \alpha + \epsilon^*$ with probability $\geq 1 - \delta$. Next, using Markov inequality, with probability $\geq 1 - \delta$

$$Pr(\mathrm{KL}(q^p(x) \,||\, q^h(x)) \geq (\alpha + \epsilon^*)^{2/3}) \leq (\alpha + \epsilon^*)^{1/3}$$

that is using the notation $\Delta = (\alpha + \epsilon^*)^{1/3}$, with probability $\geq 1 - \delta$

$$Pr(\mathrm{KL}(q^p(x) \,||\, q^h(x)) \leq \Delta^{2/3}) \geq 1 - \Delta^{1/3}$$

Using Pinkser's inequality we get $(1/2)||q^p(x) - q^h(x)||_1^2 \leq \mathrm{KL}(q^p(x)||q^h(x))$. That is, the event $\mathrm{KL}(q^p(x) \,||\, q^h(x)) \leq \Delta^{2/3}$ implies the event $||q^p(x) - q^h(x)||_1 \leq \sqrt{2}\Delta$. Thus, $Pr(||q^p(x) - q^h(x)||_1 \leq \sqrt{2}\Delta) \geq Pr(\mathrm{KL}(q^p(x) \,||\, q^h(x)) \leq \Delta^{2/3})$. Thus, we obtain: with probability $\geq 1 - \delta$, $Pr(||q^p(x) - q^h(x)||_1 \leq \sqrt{2}\Delta) \geq 1 - \Delta$.

$\square$

Proof of Lemma 7

PROOF. We know that $q_i^h(x) = \frac{e^{h_i(x)}}{\sum_j e^{h_j(x)}}$ (assume $h_T(x) = 0$). Thus,

$$|q_i^h(x) - q_i^h(x')| = q_i^h(x')|e^{h_i(x) - h_i(x')} \frac{\sum_j e^{h_j(x')}}{\sum_j e^{h_j(x)}} - 1|$$

Let r denote $\frac{\sum_j e^{h_j(x')}}{\sum_j e^{h_j(x)}}$. There is l and k such that $max_r = \frac{e^{h_l(x')}}{e^{h_l(x)}} \geq \frac{e^{h_j(x')}}{e^{h_j(x)}}$ for all j and $min_r = \frac{e^{h_k(x')}}{e^{h_k(x)}} \leq \frac{e^{h_j(x')}}{e^{h_j(x)}}$ for all j. Then, $min_r \leq r \leq max_r$ First, note that due to our assumption that for each i $|h_i(x') - h_i(x)| \leq \hat{K}||x' - x||_1$, we have

$$e^{-\hat{K}||x'-x||_1} \leq min_r \leq r \leq max_r \leq e^{\hat{K}||x'-x||_1}$$

Using the Lipschitzness we can also claim that $e^{-\hat{K}||x'-x||_1} \leq e^{h_i(x) - h_i(x')} \leq e^{\hat{K}||x'-x||_1}$. Thus,

$$e^{-2\hat{K}||x'-x||_1} \leq e^{h_i(x) - h_i(x')} \cdot r \leq e^{2\hat{K}||x'-x||_1}$$

Since, $e^{-2\hat{K}||x'-x||_1} < 1$ and $e^{2\hat{K}||x'-x||_1} > 1$ we have

$$|e^{h_i(x) - h_i(x')} r - 1| \leq \max(|e^{-2\hat{K}||x'-x||_1} - 1|, |e^{2\hat{K}||x'-x||_1} - 1|)$$

Also, it is a fact that $|e^y - 1| \leq 1.5|y|$ for $|y| \leq 3/4$. Thus, we obtain

$$|e^{h_i(x) - h_i(x')} r - 1| \leq 3\hat{K}||x' - x||_1 \text{ for } 2\hat{K}||x' - x||_1 \leq 3/4$$

Thus, $||q^h(x') - q^h(x)||_1 = \sum_i |q_i^h(x) - q_i^h(x')| = \sum_i q_i^h(x')|e^{h_i(x) - h_i(x')} \frac{\sum_j e^{h_j(x')}}{\sum_j e^{h_j(x)}} - 1| \leq (\sum_i q_i^h(x'))3\hat{K}||x' - x||_1$ for $\hat{K}||x' - x||_1 \leq 3/8$. Since $\sum_i q_i^h(x') = 1$, we have

$$||q^h(x') - q^h(x)||_1 \leq 3\hat{K}||x' - x||_1 \text{ for } ||x' - x||_1 \leq 3/8\hat{K}$$

In other words q^h is *locally* $3\hat{K}$-Lipschitz for every l_1 norm ball of size $3/8\hat{K}$. The following allows us to prove global Lipschitzness.

LEMMA 9. *Any locally L-Lipschitz function f for every l_p ball of size δ_0 on a compact convex set $X \subset \mathbb{R}^n$ is Lipschitz on the set X. The Lipschitz constant is also L.*

PROOF. Take any two points $x, y \in X$, the straight line joining x, y lies in X (as X is convex). Also, a finite number of balls of size δ_0 cover X (due to compactness). Thus, there are finitely many points $x = z_1, \ldots, z_\mu = y$ on the line from x, y such that $d_{l_p}(z_i, z_{i+1}) \leq \delta_0$. Further, since these points lie on a straight line we have

$$d_{l_p}(x, y) = \textstyle\sum_1^{\mu-1} d_{l_p}(z_i, z_{i+1})$$

Then, let any metric d be used to measure distance in the range space of f, thus, we get

$$\begin{aligned} d(f(x), f(y)) &\leq \textstyle\sum_1^{\mu-1} d(f(z_i), f(z_{i+1})) \\ &\leq \textstyle\sum_1^{\mu-1} L d_{l_p}(z_i, z_{i+1}) \\ &= L d_{l_p}(x, y) \end{aligned}$$

$\square$

Since in our case the defender mixed strategy space is compact and convex and $q^h(x)$ satisfies the above lemma with $L = 3\hat{K}$ and $\delta_0 = 3/8\hat{K}$, $q^h(x)$ is $3\hat{K}$-Lipschitz.

Proof of Theorem 4

PROOF. Coupled with the guarantee that with prob. $\geq 1-\delta$, $Pr(||q^p(x)-q^h(x)||_1 \leq \sqrt{2}\Delta) \geq 1-\Delta$, the assumptions guarantee that with prob. $\geq 1-\delta$ for the learned hypothesis h there must exist a $x' \in B(x^*, \epsilon)$ such that $||q^p(x') - q^h(x')||_1 \leq \sqrt{2}\Delta$ and there must exist $x'' \in B(\tilde{x}, \epsilon)$ such that $||q^p(x'') - q^h(x'')||_1 \leq \sqrt{2}\Delta$.

First, for notational ease let γ denote $\sqrt{2}\Delta$. The following are immediate using triangle inequality, with the results $||q^p(x') - q^h(x')||_1 \leq \gamma$ and $||q^p(x'')-q^h(x'')||_1 \leq \gamma$ and the Lipschitzness assumptions

$$||q^p(x^*) - q^h(x')||_1 \leq K\epsilon + \gamma \quad (\text{opt}x^*)$$
$$||q^p(\tilde{x}) - q^h(x'')||_1 \leq 3\hat{K}\epsilon + \gamma \quad (\text{opt}\tilde{x})$$

We call $\tilde{x}^T U q^h(\tilde{x}) \geq x'^T U q^h(x')$ as equation opt$h$. Thus, we bound the utility loss as following

$$\begin{aligned}
&x^{*T} U q^p(x^*) - \tilde{x}^T U q^p(\tilde{x}) \\
&= x^{*T} U q^p(x^*) - \tilde{x}^T U q^h(\tilde{x}) + \tilde{x}^T U q^h(\tilde{x}) - \tilde{x}^T U p(y/\tilde{x}) \\
&\leq x^{*T} U q^p(x^*) - x'^T U q^h(x') + \tilde{x}^T U q^h(\tilde{x}) - \tilde{x}^T U p(y/\tilde{x}) \\
&\qquad \text{using opt}h \\
&= (x^* - x')^T U q^p(x^*) + x'^T U (q^p(x^*) - q^h(x')) + \\
&\quad \tilde{x}^T U q^h(\tilde{x}) - \tilde{x}^T U q^p(\tilde{x}) \\
&\leq \epsilon + (K\epsilon + \gamma) + \tilde{x}^T U q^h(\tilde{x}) - \tilde{x}^T U q^p(\tilde{x}) \\
&\qquad \text{using } x' \in B(x^*, \epsilon), \text{opt}x^* \\
&= ((K+1)\epsilon + \gamma) + \tilde{x}^T U (q^h(\tilde{x}) - q^h(x'')) + \\
&\quad \tilde{x}^T U (q^h(x'') - q^p(\tilde{x})) \\
&\leq (K+1)\epsilon + \gamma + 6\hat{K}\epsilon + \gamma \\
&\qquad \text{using } x'' \in B(\tilde{x}, \epsilon) \text{ with Lipschitz } q^h, \text{opt}\tilde{x}
\end{aligned}$$

□

B. EXTENSION TO STACKELBERG GAMES

Our technique extends to Stackelberg games by noting that the single resource case $K = 1$ with $T-1$ targets gives $\sum_{i=1}^{T-1} x_i \leq 1$. This directly maps to a probability distribution over T actions. The x_i's with $x_T = 1 - \sum_{i=1}^{T-1} x_i$ is the probability of playing an action. With this set-up now the security game is a standard Stackelberg game, but where the leader has T actions and follower has $T-1$ actions.

Thus, in order to capture the general Stakelberg game, for the adversary, we assume N actions for the adversary (instead of $T-1$ above). Then, similar to security games $q_1, \ldots, q_N$ denotes the adversary's probability of playing an action. Thus, the function h now outputs vectors of size $N-1$ (instead of $O(T)$), i.e., A is a subset of $N-1$ dimensional Euclidean space. The model of security game in the PAC framework extends as is to this Stackelberg setup, just with $h(x)$ and A being $N-1$ dimensional. The rest of the analysis proceeds exactly as for security games for both parametric and non-parametric case, by replacing the T corresponding to the adversary's action space by N. Since, the proof technique is exactly same, we just state the final results. Thus, for a Stackelberg game with T leader actions and N follower actions, the bound for Theorem 1 becomes

$$\frac{576M^2}{\alpha^2} \log \frac{8\mathcal{C}(\alpha/96N, \mathcal{H}, d_{l_1^-})}{\delta}$$

It can be seen from the proof for the parametric part that the sample complexity does not depend on the dimensionality of X, but only on the dimensionality of A. Hence, the sample complexity results from generalized SUQR parametric case is

$$O\big(\frac{1}{\alpha^2}(\log\frac{1}{\delta} + N\log\frac{N}{\alpha})\big)$$

and for the non-parametric case, which depends on both dimensionality of X and T, the sample complexity is

$$O\big(\frac{1}{\alpha^2}(\log\frac{1}{\delta} + \frac{N^{T+1}}{\alpha^T})\big)$$

C. ANALYSIS OF STANDARD SUQR FORM

For SUQR the rewards and penalties are given and fixed. Let the rewards be given and fixed $r = \langle r_1, \ldots, r_T \rangle$ (each $r_i \in [0, r_{\max}], r_{\max} > 0$), and the penalty values are $p = \langle p_1, \ldots, p_T \rangle$ (each $p_i \in [0, p_{\min}], p_{\min} < 0$). Thus, the output of h is

$$\begin{aligned}
&h(x) = \\
&\langle w_1 x_{1T} + w_2 r_{1T} + w_3 p_{1T}, \ldots, \\
&w_1 x_{T-1T} + w_2 r_{T-1T} + w_3 p_{T-1T} \rangle
\end{aligned}$$

where $r_{iT} = r_i - r_T$ and same for p_{iT}. Note that in the above formulation all the component functions $h_i(x)$ have same weights. We can consider the function space $\mathcal{H}$ as the following direct-sum semi-free product $\mathcal{G} \oplus \mathcal{F} \oplus \mathcal{E} = \{\langle g_1 + f_1 + e_1, \ldots, g_{T-1} + f_{T-1} + e_{T-1} \rangle \mid \langle g_1, \ldots, g_{T-1} \rangle \in \mathcal{G}, \langle f_1, \ldots, f_{T-1} \rangle \in \mathcal{F}, \langle e_1, \ldots, e_{T-1} \rangle \in \mathcal{E}\}$, where each of $\mathcal{G}, \mathcal{F}, \mathcal{E}$ is defined below. $\mathcal{G} = \{\langle g_1, \ldots, g_{T-1} \rangle \mid \langle g_1, \ldots, g_{T-1} \rangle \in \times_i \mathcal{G}_i$, all g_i have same weight$\}$ where $\mathcal{G}_i$ has functions of the form wx_{iT}. $\mathcal{F} = \{\langle f_1, \ldots, f_{T-1} \rangle \mid \langle f_1, \ldots, f_{T-1} \rangle \in \times_i \mathcal{F}_i$, all f_i have same weight$\}$ where $\mathcal{F}_i$ has constant valued functions of the form wr_{iT}. $\mathcal{E} = \{\langle e_1, \ldots, e_{T-1} \rangle \mid \langle e_1, \ldots, e_{T-1} \rangle \in \times_i \mathcal{E}_i$, all e_i have same weight$\}$ where $\mathcal{E}_i$ has constant valued functions of the form wp_{iT}.

Consider an $\epsilon/3$-cover U_e for $\mathcal{E}$, an $\epsilon/3$-cover U_f for $\mathcal{F}$ and $\epsilon/3$-cover U_g for $\mathcal{G}$. We claim that $U_e \times U_f \times U_g$ is an ϵ-cover for $\mathcal{E} \oplus \mathcal{F} \oplus \mathcal{G}$. Thus, the size of the ϵ-cover for $\mathcal{E} \oplus \mathcal{F} \oplus \mathcal{G}$ is bounded by $|U_e||U_f||U_g|$. Thus,

$$\mathcal{N}(\epsilon, \mathcal{H}, d_{l_1^-}) \leq \mathcal{N}(\epsilon/3, \mathcal{G}, d_{l_1^-}) \mathcal{N}(\epsilon/3, \mathcal{F}, d_{l_1^-}) \mathcal{N}(\epsilon/3, \mathcal{E}, d_{l_1^-})$$

Taking sup over P we get

$$\mathcal{C}(\epsilon, \mathcal{H}, d_{l_1^-}) \leq \mathcal{C}(\epsilon/3, \mathcal{G}, d_{l_1^-}) \mathcal{C}(\epsilon/3, \mathcal{F}, d_{l_1^-}) \mathcal{C}(\epsilon/3, \mathcal{E}, d_{l_1^-})$$

Now, we show that $U_e \times U_f \times U_g$ is an ϵ-cover for $\mathcal{H} = \mathcal{E} \oplus \mathcal{F} \oplus \mathcal{G}$ Fix any $h \in \mathcal{H} = \mathcal{E} \oplus \mathcal{F} \oplus \mathcal{G}$. Then, $h = e + f + g$ for some $e \in \mathcal{E}, f \in \mathcal{F}, g \in \mathcal{G}$. Let $e' \in U_e$ be $\epsilon/3$ close to e, $f' \in U_f$ be $\epsilon/3$ close to f and $g' \in U_g$ be $\epsilon/3$ close to g.

Then,

$$\begin{aligned}
&d_{L^1(P, d_{l_1^-})}(h, h') \\
&\quad = \int_X \frac{1}{k} \sum_{i=1}^{k} d_{l_1}(h_i(x), h'_i(x))\, dP(x) \\
&\quad \leq \int_X \frac{1}{k} \sum_{i=1}^{k} d_{l_1}(g_i(x), g'_i(x)) \\
&\qquad + d_{l_1}(f_i(x), f'_i(x)) + d_{l_1}(e_i(x), e'_i(x))\, dP(x) \\
&\quad = d_{L^1(P, d_{l_1^-})}(g, g') + d_{L^1(P, d_{l_1^-})}(f, f') + d_{L^1(P, d_{l_1^-})}(e, e') \\
&\quad \leq \epsilon
\end{aligned}$$

Similar to Lemma 2, it is possible to show that for any probability distribution P, for any function g, g' $d_{l_1^-}(g, g') \leq |w - w'|$

and f, f' $d_{l_1^-}(f, f') \leq |w - w'|r_{max}$ and e, e' $d_{l_1^-}(e, e') \leq |w - w'||p_{min}|$. Assume each of the functions have a range $[-M/6, M/6]$ (this does not affect the order in terms of M). Given, these ranges w for g can take values in $[-M/6, M/6]$, w for g can take values in $[-M/6r_{max}, M/6r_{max}]$ and w for g can take values in $[-M/6|p_{min}|, M/6|p_{min}|]$. To get a capacity of $\epsilon/3$ it is enough to divide the respective w range into intervals of $2\epsilon/3$, and consider the boundaries. This yields an $\epsilon/3$-capacity of $M/2\epsilon$, $M/2\epsilon r_{max}$ and $M/2\epsilon|p_{min}|$ for $\mathcal{G}$, $\mathcal{F}$ and $\mathcal{E}$ respectively.

Thus,

$$\mathcal{C}(\epsilon, \mathcal{H}, d_{l_1^-}) \leq (M/2\epsilon)^3 \frac{1}{r_{max}|p_{min}|}$$

Plugging this in sample complexity from Theorem 1 we get the results that the sample complexity is

$$O\big(\frac{1}{\alpha^2}(\log\frac{1}{\delta} + \log\frac{T}{\alpha})\big)$$

D. ALTERNATE PROOF FOR GENERALIZED SUQR SAMPLE COMPLEXITY

As discussed in the main paper we use the function space $\mathcal{H}'$ with each component function space $\mathcal{H}'_i$ given by $w_i x_{iT} + c_{iT}$. Then, we can directly use Equation 2. We still need to bound $\mathcal{C}(\epsilon, \mathcal{H}'_i, d_{l_1})$. For this, we note the set of functions $w_i x_{iT} + c_{iT}$ has two free parameters w_i and c_i, thus, this function space is a subset of the vector space of functions of dimension two (two values needs to represent each function). Using the pseudo-dimension technique [14] we know that for psuedo-dimension d of function space $\mathcal{H}_i$ we get

$$\mathcal{C}(\epsilon, \mathcal{H}'_i, d_{l_1}) \leq 2(\frac{eM}{\epsilon}\log\frac{eM}{\epsilon})^d$$

Also, we know [14] that pseudo-dimension is equal to the vector space dimension if the function class is a subset of a vector space. Therefore, for our case $d = 2$. Therefore, using Equation 2 we get

$$\mathcal{C}(\epsilon, \mathcal{H}', d_{l_1}) \leq 2^T(\frac{eM}{\epsilon}\log\frac{eM}{\epsilon})^{2T}$$

Plugging this result in Theorem 1 we get the sample complexity of

$$O\Big((\frac{1}{\alpha^2})\big(\log(\frac{1}{\delta}) + T\log(\frac{T}{\alpha}\log\frac{T}{\alpha})\big)\Big)$$

E. EXPERIMENTAL RESULTS

Here we provide additional experimental results on the Uganda, AMT and simulated datasets. The AMT dataset consisted of 32 unique mixed strategies, 16 of which were deployed for one payoff structure and the remaining 16 for another. In the main paper, we provided results on AMT data for payoff structure 1. Here, in Figs. 2(a) and 2(b), we show results on the AMT data for both the parametric (SUQR) and NPL learning settings on payoff structure 2.

For running experiments on simulated data, we used the same mixed strategies and features as for the AMT data, but simulated attacks, first using the actual SUQR model and then using a modified form of the SUQR model. Figs. 2(c) and 2(d) show results on simulated data on payoff structures 1 and 2 for the parametric cases, when the data is generated by an adversary with an SUQR model with true weight vector reported in Nguyen et. al [26] ($(w_1, w_2, w_3) = (-9.85, 0.37, 0.15)$ ($c_i = w_2 R_i + w_3 P_i$)). Similar results for the NPL model are shown in Figs. 2(e) and 2(f) respectively. We can see that the NPL approach performs poorly with only one or five samples as expectied but improves significantly as more samples are added. To further show its potential, we modified the true adversary model of generating attacks from SUQR to the following: $q_i \propto e^{w_1 x_i^2 + c_i}$, i.e., instead of x_i, the adversary reasons based on x_i^2. We considered the same true weight vector to simulate attacks. Then, we observe in Figs. 2(g) (for payoff structure 1) and 2(h) (for payoff structure 2 data), that α approaches a value closer to zero for 500 or more sample. Also, the NPL model performs better than the parametric model with 500 or more samples. This shows that the NPL approach is more accurate when the true adversary does not satisfy the simple parametric logistic form, indicating that when we don't know the true function of the adversary's decision making process, adopting a non-parametric method to learn the adversary's behavior is more effective.

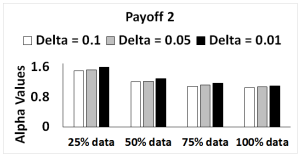

(a) AMT Parametric Results Payoff 2

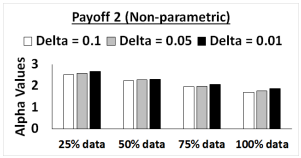

(b) AMT Nonparametric Results Payoff 2

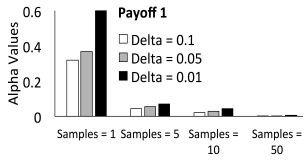

(c) Simulated Data Payoff 1 - Parametric results

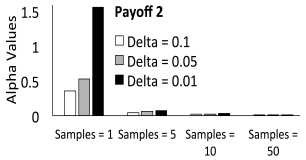

(d) Simulated Data Payoff 2 - Parametric results

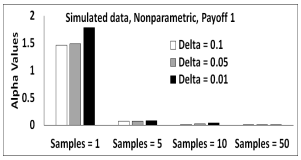

(e) Simulated Data Payoff 1 - Nonparametric results

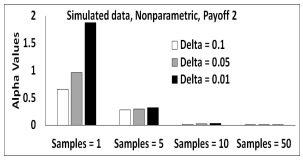

(f) Simulated Data Payoff 2 - Nonparametric results

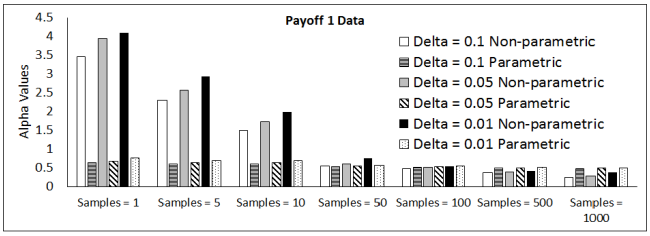

(g) Parametric vs Non-parametric results on Simulated (for various sample sizes) data from payoff 1 when the true adversary model is different from the parametric learned function

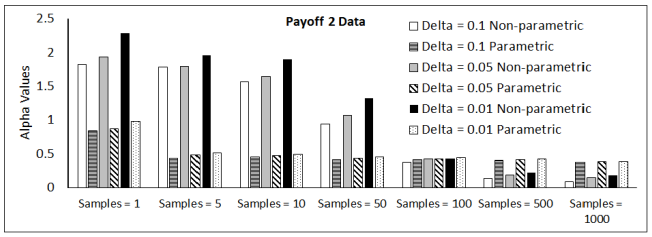

(h) Parametric vs Non-parametric results on Simulated (for various sample sizes) data from payoff 2 when the true adversary model is different from the parametric learned function

Figure 2: Results on Uganda, AMT and simulated datasets for the parametric and non-parametric cases respectively.