Online Appendix

To prove Theorem 1, we first prove the following combinatorial lemma.

LEMMA 6. *Let d be a multiple of 8 and $S_1, S_2, S_3, \ldots$ be a sequence of uniformly-drawn random subsets of $\{1, \ldots, d\}$ of size $d/4$. Then, $t_0 = \Omega(1.2^d)$:*

$$\Pr\left(|S_j \cap S_k| \le d/8, \forall 1 \le j < k \le t_0\right) \ge 1/2.$$

Proof of Lemma 6. Let S_j and S_k be two randomly sampled subsets of the set $\{1, \ldots, d\}$ of cardinality $d/4$. Consider a given feasible set S_k (a subset of $\{1, \ldots, d\}$ of size $d/4$). Then, there exists exactly $\binom{d}{d/4}$ possibilities for the set S_j. The number of possibilities for the set S_j that share exactly i elements with the set S_k is given by the product of binomials $\binom{d/4}{i}\binom{3d/4}{d/4-i}$. Therefore, the number of possible sets S_j with more than $d/8$ elements from S_k is given by $\sum_{i=d/8+1}^{d/4} \binom{d/4}{i}\binom{3d/4}{d/4-i}$. Thus,

$$\Pr\left(|S_j \cap S_k| > d/8\right) = \sum_{i=d/8+1}^{d/4} \frac{\binom{d/4}{i}\binom{3d/4}{d/4-i}}{\binom{d}{d/4}} \le O(1.69^{-d}), \tag{13}$$

where we will demonstrate the last inequality shortly. Now, consider a collection of $S_1, S_2, ..., S_{t_0}$ sets where $t_0 = 1.2^d$. Using the union bound, we can write:

$$\begin{aligned}
\Pr\left(|S_j \cap S_k| \le d/8, \forall 1 \le j < k \le t_0\right) &= 1 - \Pr\left(\exists\ j < k \in \{1, ..., t_0\} \mid |S_j \cap S_k| > d/8\right) \\
&\ge 1 - \sum_{1 \le j \le t_0} \sum_{1 \le k < j} \Pr\left(|S_j \cap S_k| > d/8\right) \\
&\ge 1 - (1.2^d)^2 \cdot O(1.69^{-d}) \\
&\ge 1/2 \qquad \text{for sufficiently high } d.
\end{aligned}$$

This proves the statement of the lemma.

We now prove the inequality in Eq. (13). From Stirling inequalities, $\sqrt{2\pi n}(n/e)^n \le n! \le e\sqrt{n}(n/e)^n$, we have

$$\binom{d}{d/4} \ge p_1(d) \frac{d^d}{(d/4)^{(d/4)}(3d/4)^{(3d/4)}} = p_1(d) 1.755^d, \tag{14}$$

where we use $p_1(d)$ to represent a polynomial function of d. Our next step is to use the same technique to show that the numerator from Eq. (13), $\sum_{i=d/8+1}^{d/4} \binom{d/4}{i}\binom{3d/4}{d/4-i}$, is bounded by $p_2(d)1.038^d$, where $p_2(d)$ is a polynomial function of d. We define:

$$h(d) = \sum_{i=d/8+1}^{d/4} \binom{d/4}{i}\binom{3d/4}{d/4-i}.$$

Management Science 00(0), pp. 000–000, © 0000 INFORMS

Using the Stirling inequalities once more, we can bound the quantity above by

$$\begin{aligned} h(d) &= \sum_{i=d/8+1}^{d/4} \frac{(d/4)!}{i!(d/4-i)!} \frac{(3d/4)!}{(d/4-i)!(d/2+i)!} \\ &\leq \tilde{p}_2(d) \sum_{i=d/8+1}^{d/4} \frac{(d/4)^{(d/4)}}{i^i(d/4-i)^{(d/4-i)}} \frac{(3d/4)^{(3d/4)}}{(d/4-i)^{(d/4-i)}(d/2+i)^{(d/2+i)}}, \end{aligned}$$

for some polynomial $\tilde{p}_2(d)$. Note that we can bound the sum by $(d/4-d/8)$ times the highest value of i and therefore

$$h(d) \leq p_2(d) \max_{i\in\{d/8,\dots,d/4\}} \frac{(d/4)^{(d/4)}}{i^i(d/4-i)^{(d/4-i)}} \frac{(3d/4)^{(3d/4)}}{(d/4-i)^{(d/4-i)}(d/2+i)^{(d/2+i)}},$$

with a slightly different polynomial $p_2(d)$ than before. We now replace i by yd for some $y \in [1/8, 1/4]$ to obtain the following bound:

$$\begin{aligned} h(d) &\leq p_2(d) \max_{y\in[1/8,1/4]} \frac{(d/4)^{(d/4)}}{(yd)^{(yd)}(d/4-yd)^{(d/4-yd)}} \frac{(3d/4)^{(3d/4)}}{(d/4-yd)^{(d/4-yd)}(d/2+yd)^{(d/2+yd)}} \\ &= p_2(d) \max_{y\in[1/8,1/4]} \left[\frac{(1/4)^{(1/4)}(3/4)^{(3/4)}}{y^y(1/4-y)^{(1/2-2y)}(1/2+y)^{(1/2+y)}}\right]^d. \end{aligned} \tag{15}$$

Now consider the function $g(y)$ defined as

$$g(y) = y^y(1/4-y)^{(1/2-2y)}(1/2+y)^{(1/2+y)}.$$

The second derivative of the logarithm of $g(y)$ is equal to

$$\frac{d^2 \ln g(y)}{dy^2} = \frac{-y^2+2.75y+.05}{(y-1)^2y(y+1)},$$

which is positive in the region $[1/8, 1/4]$. Therefore $\ln g(y)$ is strictly convex within this region and thus, has a unique minimizer. Using the first order condition, we can compute this minimizer to be $y = 0.169$. In addition, $g(0.169)$ evaluates to 0.549. Replacing this value in Eq. (15), we obtain:

$$h(d) \leq p_2(d)\left(\frac{.570}{.549}\right)^d = p_2(d)1.038^d. \tag{16}$$

The ratio between the two bounds from Eqs. (14) and (16) gives us

$$\frac{p_2(d)1.038^d}{p_1(d)1.755^d} = O(1.69^{-d}),$$

since the polynomials are absorbed by the exponential terms. This proves Eq. (13), completing our proof. □

Proof of Theorem 1. Consider a setting where $K_1 = [1,2]^d$. For $\Omega(a^d)$ steps, assume that nature draws a random subset S_t of $\{1,\ldots,\tilde{d}\}$ with size $\tilde{d}/4$, where $\tilde{d} = 8\lfloor d/8 \rfloor$. Nature chooses the feature vectors $x_t = \frac{1}{\sqrt{\tilde{d}/4}} I\{S_t\}$, i.e., the i-th coordinate of the vector x_t is $\frac{1}{\sqrt{\tilde{d}/4}}$ if $i \in S_t$, and 0 otherwise. We first show that with probability at least 1/2, the regret incurred over $\Omega(a^d)$ steps is at least $\Omega(Ra^d)$, where a is the constant 1.2 from Lemma 6. We will assume that d is a multiple of 8, and therefore use d instead of $\tilde{d}$.

We divide the proof into two cases depending on the value of ϵ.

1. Assume that $\epsilon < 0.5\sqrt{d}$ and consider the case $\theta = (1,1,\ldots,1)$. We now analyze the event where the pairwise intersection of sets S_t is at most $d/8$. In this case, we have:

$$\min_{\hat{\theta}\in K_1} \hat{\theta}' x_1 = \sqrt{d/4} = 0.5\sqrt{d} \qquad \text{and} \qquad \max_{\hat{\theta}\in K_1} \hat{\theta}' x_1 = \frac{2(d/4)}{\sqrt{d/4}} = \sqrt{d}.$$

The difference is equal to $0.5\sqrt{d}$ and is larger than ϵ, so that our algorithm will explore and set an explore price of $p_1 = 0.75\sqrt{d}$. Since $\theta' x_1 < p_1$, a sale does not occur, and the algorithm incurs a regret of $0.5\sqrt{d}$.

We next claim by induction that for every t, we have:

$$\min_{\hat{\theta}\in K_t} \hat{\theta}' x_t = 0.5\sqrt{d} \qquad \text{and} \qquad \max_{\hat{\theta}\in K_t} \hat{\theta}' x_t = \sqrt{d}.$$

As a result, the price is set to $p_t = 0.75\sqrt{d}$, no sale occurs, and the algorithm incurs a regret of $0.5\sqrt{d}$ in every period. The base case $(t=1)$ was shown above. Assume that the claim is true for t and we next show that it holds for $t+1$.

We have: $K_{t+1} = K_t \cap \{\theta' x_t \le 0.75\sqrt{d}\} = K_1 \cap_{s=1,2,\ldots t} \{\theta' x_s \le 0.75\sqrt{d}\}$. Note that for any $s = 1,2,\ldots,t$, we have: $(1,1,\ldots,1)' x_s = 0.5\sqrt{d}$ and hence, $\theta = (1,1,\ldots,1) \in K_{t+1}$. Therefore, we obtain:

$$\min_{\hat{\theta}\in K_{t+1}} \hat{\theta}' x_{t+1} = (1,1,\ldots,1)' x_{t+1} = 0.5\sqrt{d}.$$

Consider the vector $\widetilde{\theta}$ such that $\widetilde{\theta}_i = 2$ for $i \in S_{t+1}$, and $\widetilde{\theta}_i = 1$ otherwise. If we show that $\widetilde{\theta} \in K_{t+1}$, then we have:

$$\max_{\hat{\theta}\in K_{t+1}} \hat{\theta}' x_{t+1} \ge \widetilde{\theta}' x_{t+1} = \sqrt{d}.$$

Since the maximum over the initial set K_1 is also equal to $\sqrt{d}$, the above maximum cannot be larger than $\sqrt{d}$. The last step is to show that $\widetilde{\theta} \in K_{t+1}$. We know that $\widetilde{\theta} \in K_1$. In addition, we have for any $s = 1,2,\ldots,t$:

$$\begin{aligned}\widetilde{\theta}' x_s = \frac{1}{\sqrt{d/4}} \sum_{i\in S_s} [1 + I\{i \in S_{t+1}\}] &= \frac{1}{\sqrt{d/4}} \left[\frac{d}{4} + |S_s \cap S_{t+1}|\right] \\ &\le \frac{1}{\sqrt{d/4}} \left[\frac{d}{4} + \frac{d}{8}\right] = 0.75\sqrt{d},\end{aligned}$$

where the inequality follows from Lemma 6. Therefore, $\widetilde{\theta} \in K_{t+1}$.

For all $t = 1, 2, ..., k$, our algorithm incurs a regret of $0.5\sqrt{d}$. Recall that we have $k = \Omega(a^d)$ such steps, so that the total regret is given by: $0.5\sqrt{d} \cdot \Omega(a^d) = \Omega(\sqrt{d} \cdot a^d)$.

2. Assume that $\epsilon \geq 0.5\sqrt{d}$ and consider the case $\theta = (2, 2, \ldots, 2)$. In this scenario, our algorithm will always exploit (as the difference between the maximum and minimum is equal to $0.5\sqrt{d}$). The total regret is then: $(\sqrt{d} - 0.5\sqrt{d}) \cdot \Omega(a^d) = \Omega(\sqrt{d} \cdot a^d)$.

Note that in the argument above, we assumed $K_1 = [1,2]^d$, which is an uncertainty set where $R = \max_{\theta \in K_1} \|\theta\| = 2\sqrt{d}$. If we replace the uncertainty set with $K_1 = [\alpha, 2\alpha]^d$, for some $\alpha > 1$, R would increase to $2\alpha\sqrt{d}$ and the regret would scale with α. Consequently, the regret incurred over $\Omega(a^d)$ steps is $\Omega(Ra^d)$.

We next show the $\ln T$ part of the regret. Recall that $K_1 = [1,2]^d$ and $\theta = (1, 1, ..., 1)$. Let $\hat{T} = \Omega(a^d)$ be the period in which the sequence of steps outlined above ends. We first argue that the vectors z_i belong to the set $K_{\hat{T}}$, for all $i = 1, ..., d$, where z_i is a vector with 2 in dimension i and 1 in all other dimensions. To see this, we show that z_i satisfies $(z_i)'x_t \leq p_t$ for all $t = 1, ..., \hat{T}$. We have: $(z_i)'x_t = 2x_i + \sum_{j=2}^d x_j \leq \frac{4}{\sqrt{d}} + \frac{d}{4}\frac{2}{\sqrt{d}}$ using the way x_is are chosen in this construction. Recall that $p_t = 0.75\sqrt{d}$, and hence we obtain $(z_i)'x_t \leq \frac{4}{\sqrt{d}} + \frac{d}{4}\frac{2}{\sqrt{d}} \leq p_t$, for $d \geq 16$.

For periods $k = \hat{T} + i$ for $i = 1, ..., d$, we will assume that nature chooses the vectors $x_{\hat{T}+i} = e_i$. For $k = \hat{T} + 1$, we have

$$\min_{\hat{\theta} \in K_{\hat{T}+1}} \hat{\theta}' x_{\hat{T}+1} = \theta' e_1 = 1 \qquad \text{and} \qquad \max_{\hat{\theta} \in K_{\hat{T}+1}} \hat{\theta}' x_{\hat{T}+1} = (z_1)' e_1 = 2.$$

Therefore, $p_{\hat{T}+1} = 3/2$. Note that this does not eliminate any of the vectors z_i from the uncertainty set for $i > 1$. Repeating this argument d times, we will have added the inequalities $\theta_i \leq 3/2$ for all $i = 1, ..., d$ by step $k = \hat{T} + d$.

It remains to argue that $K_{\hat{T}+d+1} = [1, 3/2]^d$, which requires proving that no point in $[1, 3/2]^d$ was removed from the set during steps $k = 1, ..., \hat{T}$. This is equivalent to showing that for all $\hat{\theta} \in [1, 3/2]^d$, $\hat{\theta}' x_t \leq p_t = \frac{3}{4}\sqrt{d}$. It is sufficient to show this inequality for $\hat{\theta} = (3/2, 3/2, ...3/2)$, which is the worst value of $\hat{\theta}$. By the construction of x_t, $(3/2, 3/2, ...3/2)' x_t = \frac{3}{2}\sum_{i=1}^{d/4} \frac{2}{\sqrt{d}} = \frac{3}{4}\sqrt{d}$, satisfying the desired inequality.

With this construction, we have returned in period $\hat{T} + d + 1$ to the same stage we were in period 1, except that the uncertainty set has been reduced by half in all dimensions. We can now repeat the same argument by using the identical sequence of contexts, $x_t = \frac{1}{\sqrt{\tilde{d}/4}} I\{S_t\}$, and argue that the price will be $p_t = (3/8)\sqrt{d}$, resulting in no sales for $\Omega(a^d)$ steps, followed by the contexts $x_t = e_i$ for d steps, where we again reduce by half each dimension of the uncertainty set.

After repeating this entire argument k times, we are left with an uncertainty set equals to $[1, 1 + 2^{-k}]^d$. After $k = \log_2(1/\epsilon)$ iterations, the uncertainty set becomes $[1, 1 + 1/\epsilon]^d$. The regret is bounded by (i) the loss due to explore periods, which is proportional to the number of explore steps, $a^d \log_2(1/\epsilon)$ plus (ii) the loss due to exploit periods, which corresponds to the maximum loss under an exploit price, $\epsilon\sqrt{d}$, multiplied by the number of exploit steps, $(T - a^d \log_2(1/\epsilon))$:

$$a^d \log_2(1/\epsilon) + \epsilon\sqrt{d} \cdot (T - a^d \log_2(1/\epsilon)).$$

The value of ϵ that minimizes this regret is $\epsilon = a^d/(T\sqrt{d})$, yielding a regret bound of $\Omega(a^d \ln(T\sqrt{d}/a^d))$. Replacing $a = 1.2$ with any number strictly between 1 and 1.2 yields the desired $\Omega(a^d \ln(T))$. $\square$