Commitment to Sparse Strategies in Two-Player Games

Salam Afiouni[1], Jakub Černý[1], Chun Kai Ling[2], Christian Kroer[1]

[1]Columbia University, USA
[2]National University of Singapore, Singapore
sa4316@columbia.edu, chunkail@nus.edu.sg, jakub.cerny@columbia.edu, christian.kroer@columbia.edu

Abstract

While Nash equilibria are guaranteed to exist, they may exhibit dense support, making them difficult to understand and execute in some applications. In this paper, we study k-sparse commitments in games where one player is restricted to mixed strategies with support size at most k. Finding k-sparse commitments is known to be computationally hard. We start by showing several structural properties of k-sparse solutions, including that the optimal support may vary dramatically as k increases. These results suggest that naive greedy or double-oracle-based approaches are unlikely to yield practical algorithms. We then develop a simple approach based on mixed integer linear programs (MILPs) for zero-sum games, general-sum Stackelberg games, and various forms of structured sparsity. We also propose practical algorithms for cases where one or both players have large (i.e., practically innumerable) action sets, utilizing a combination of MILPs and incremental strategy generation. We evaluate our methods on synthetic and real-world scenarios based on security applications. In both settings, we observe that even for small support sizes, we can obtain more than 90% of the true Nash value while maintaining a reasonable runtime, demonstrating the significance of our formulation and algorithms.

Code — github.com/CoffeeAndConvexity/SparseEquilibria
Extended version — arxiv.org/abs/2412.14337

1 Introduction

Equilibrium finding in games has emerged as a key component of artificial intelligence, owing to applications in recreational game playing such as poker, Stratego, and Diplomacy, numerous applications in security, as well as logistics (Bowling et al. 2015; Brown and Sandholm 2018; Perolat et al. 2022; Bakhtin et al. 2022; Pita et al. 2008; Fang, Stone, and Tambe 2015; Tambe 2011; Jain et al. 2010; Černý et al. 2024). Central to the success of these applications are efficient approaches towards finding the optimal *randomized*, or *mixed* equilibria in their respective settings. Not only are randomized equilibria typically guaranteed to exist, they allow for the prescription of behavior that hedges over all possible opponent actions, making them particularly useful in zero-sum games. Nonetheless, a persistent difficulty faced by practitioners seeking to apply game theoretic solutions into the real world is that such solutions exhibit *dense support*, i.e., they could randomize over a large number of actions. Strategies with dense support are undesirable for many reasons; for instance, they are difficult to interpret and implement in practice, and are often faced with numerical issues if used for downstream processing.

In this paper, we propose overcoming these limitations by computing k-sparse solutions as opposed to Nash equilibria. Loosely speaking, a distinguished player termed the sparse player (representing the practitioner) is restricted to playing a mixed strategy with a support size at most k. The other player is free to play any mixed strategy they desire. We argue that in many practical use-cases, sparse strategies can almost match the performance of dense strategies, as long as the support of the sparse player is carefully chosen.

The problem of finding sparse solutions has been studied from various angles. Althöfer (1994) show that sparse approximations to strategies exist in zero-sum games without too much degradation, while seminal work by Lipton, Markakis, and Mehta (2003) show that *approximate* sparse Nash equilibria always exist via a sampling argument (Feder, Nazerzadeh, and Saberi (2007) show the optimality of these results); follow up work extends this argument to correlated and other equilibria (Babichenko, Barman, and Peretz 2014). Nonetheless, these arguments are nonconstructive in nature, and finding "good" sparse equilibria can be difficult. Indeed, finding the *optimal* k-sparse commitment even for zero-sum games is in general intractable (McCarthy et al. 2018). Foster, Golowich, and Kakade (2023) study a variant of sparsity in Markov games where correlated equilibria distributions are expressible as a mixture of k uniform product distributions, while Anagnostides et al. (2023) show hardness of computing sparse correlated equilibria for extensive form games. The line of work closest to ours is by McCarthy et al. (2018), who study k-sparse solutions under the name of *operationalizable* strategies, mainly from the perspective of security games. Our work differs from them in several ways. First, our algorithm is not restricted to finding k-uniform strategies, i.e., each of the k actions (possibly repeated) must be chosen with equal probability. Crucially, we significantly expand the scope of applications beyond that of security games, including large games with MILP representable strategy spaces.

For the majority of theoretical and algorithmic work above, the term "sparse" is formally defined by asymptotics in terms of the size of the game (e.g., the number of players, and the size of their actions). For our contributions, we adopt colloquial, but often times more practically useful, meanings for the terms sparse and dense, with sparse meaning "small enough" (e.g., say, in single digits), while dense implies the converse. Secondly, a significant amount of work (Lipton, Markakis, and Mehta 2003; Babichenko, Barman, and Peretz 2014; McCarthy et al. 2018) also requires sparse strategies to be k-uniform. This is a stronger restriction than our notion of k-sparsity, and it can adversely affect performance when we want a sparse strategy in the practical non-asymptotic sense where k is a very small constant; we show this both empirically and theoretically in Sections 3 and 5.

With these distinctions in mind, our key contributions are as follows. First, we formalize the notion of k-sparse commitments for two-player zero-sum games and Stackelberg equilibrium in general-sum games. Second, we study some of the properties of k-sparse commitments. We show that even for zero-sum games, supports of the optimal k-sparse equilibrium can be disjoint for almost all k, and the value function as a set function of allowable support is not submodular. We also demonstrate that *k-uniform* equilibria as computed by McCarthy et al. (2018) can yield a much worse performance than k-sparse equilibrium. Thirdly, we show how optimal k-sparse commitments can be computed via MILPs for a variety of settings, including zero-sum games, Stackelberg equilibrium, variants of structured action sets, as well as games where one or both players have a large, but MILP-representable strategy space. Finally, we demonstrate empirically the efficacy of our proposed method, using both randomly generated normal-form games, as well as settings based on security applications. Our results demonstrate scalability of our method which also often performs better than k-uniform strategies. Even for small k, we are often able to capture nearly 90% of the true optimal commitment, showing that k-sparse commitments are a practical and viable.

2 Preliminaries

We are concerned with two-player games, where Player 1 and 2 have n and m actions respectively. We denote their payoff matrices by $A, B \in \mathbb{R}^{n\times m}$, such that A_{ij} (resp. B_{ij})) gives Player 1's payoff when Player 1 plays action i and Player 2 plays action j. The spaces of mixed strategies for each player are given by Δ^n and Δ^m respectively, where Δ^n is the probability simplex $\{x \in \mathbb{R}^n_+ | \sum_i^n x_i = 1\}$. For a vector $x \in \mathbb{R}^d$, we denote its support by $supp(x) = \{i \in [d] | x_i \neq 0\}$. The *best-response* of Player 2 against Player 1's (mixed) strategy $x \in \Delta^n$ is any pure strategy $\text{BR}_2(x) = \text{argmax}_{y\in[m]}\, x^T Ay$, where ties are broken arbitrarily. When y is a non-negative integer, we adopt the shorthand $x^T Ay = xAe_y$, where e_i is the i-th elementary basis vector. Best-responses of Player 1 are similarly denoted by $\text{BR}_1(y)$.

Nash equilibrium (NE) in zero-sum games Recall that in a zero-sum game, $A = -B$. Finding a Nash equilibrium (x^*, y^*) in a two-player zero-sum game can be formulated as a saddle-point problem using first player's payoff matrix.

In particular, von Neumann's *minimax theorem* shows that (x^*, y^*) is achieved by the saddle point of $x^T Ay$,

$$\max_{x\in\Delta^n} \min_{y\in\Delta^m} x^T Ay = \min_{y\in\Delta^m} \max_{x\in\Delta^n} x^T Ay = {x^*}^T Ay^*. \quad (1)$$

The minimax theorem shows that (x^*, y^*) is a NE if and only if each mixed strategy optimizes the players' payoff assuming that the other player best responds. Two-player zero-sum games can be solved efficiently using a variety of methods, including fictitious play (Brown 1951), linear programming (Shoham and Leyton-Brown 2008), and self-play with no-regret learners (Hart and Mas-Colell 2000).

Sparse commitments in zero-sum games We now consider a scenario where Player 1 is limited to playing strategies with support at most k. For the rest of the paper, we refer to Player 1 (resp. Player 2) as the sparse player (resp. the non-sparse player) and vice versa. Player 1's restricted strategy space is given by

$$\Delta^n_k = \{x \in \Delta^n \text{ such that } |supp(x)| \leq k\},$$

Then Player 1's set of optimal bounded-support strategies is

$$x^*_k = \operatorname*{argmax}_{x\in\Delta^n_k} \min_{y\in[m]} x^T Ay. \quad (2)$$

Note that Δ^n_S is not convex due to the cardinality constraint. Thus, the minimax theorem (1) does not hold. Instead, we have the weaker result $\max_{x\in\Delta^n_k} \min_{y\in\Delta^m} x^T Ay \leq \min_{y\in\Delta^m} \max_{x\in\Delta^n_k} x^T Ay$.

Remark 1. *In our definition, only one player is subjected to sparsity constraints; the other (i.e., inner) player is allowed to play any mixed strategy it desires, even one which is not sparse. This assumption has no bearing on our results, since at least one of Player 2's best responses will be pure.*

Strong Stackelberg equilibrium (SSE) in general-sum games In general-sum two-player games, where A and B may be arbitrary $n \times m$ matrices, a Stackelberg equilibrium (x^*, y^*) can be formulated as a bilevel problem (Von Stengel and Zamir 2004; Conitzer and Sandholm 2006):

$$x^* = \operatorname*{argmax}_{x\in\Delta^n} x^T Ay^*(x), \text{ where } \quad y^*(x) = \operatorname*{argmax}_{y\in[m]} x^T By.$$

The equilibrium is considered *strong* if, in addition, Player 2 breaks ties in favor of Player 1. In two player zero-sum games, NE and SSE coincide for Player 1, i.e., the optimal strategy x^* that Player 1 commits to is also a NE strategy. For this paper, we always assume that Player 1 is playing the role of the Stackelberg leader (outer optimization problem).

The SSE applies to general-sum games, always exists, enjoys a unique payoff, and can be computed in polynomial time, making it a popular choice of equilibrium for security applications (Sinha et al. 2018; Tambe 2011).

Sparse commitments in general-sum games Just like zero-sum games, we define optimal k-sparse commitments via the restricted strategy spaces Δ^n_k. The strategy x^*_k is then

$$x^*_k = \operatorname*{argmax}_{x\in\Delta^n_k} x^T Ay^*(x), \text{ where } \quad y^*(x) = \operatorname*{argmax}_{y\in\text{BR}_2(x)} x^T Ay.$$

It is important to differentiate between **natural sparsity** in game solutions, which refers to the *existence* of equilibria with small support, and the **enforced sparsity** in our context, which typically involves *trade-offs* in terms of strictly reduced performance for the sparse player. Although many structured games (e.g., extensive form games or other games played on graphs) exhibit naturally sparse equilibria, these are typically not sparse enough to be practically applicable.

3 Structure of Sparse Equilibria and the Limits of Naïve Sparsification Methods

We now explore pathological properties of sparse equilibria which will justify the necessity of our proposed method over more intuitive or heuristic approaches. For one, it is tempting to believe that the solutions x_k^* and $x_{k'}^*$ for $k < k'$ are "close" in some sense. Unfortunately, this is not the case: $supp(x_k^*)$ and $supp(x_{k'}^*)$ can be completely disjoint.

Proposition 1. *There exist zero-sum games where the optimal sparse commitments x_k^*, $x_{k'}^*$ for $2 \leq k < k' \leq \sqrt{2n}$ have disjoint supports, i.e., $supp(x_k^*) \cap supp(x_{k'}^*) = \emptyset$.*

Despite Proposition 1, one may hope that even if $supp(x_k^*)$ and $supp(x_{k'}^*)$ differ greatly, player 1's expected utility under them may be somewhat nicely behaved. Define $v_k^* = \max_{x \in \Delta_k^n} \min_{y \in [m]} x^T Ay$, the reward to the first player under the optimal size k commitment. Clearly, v_k^* is non-decreasing in k. However, the utility of the first player is not necessarily "concave" in k, i.e., the marginal returns from increasing k is not necessarily diminishing.

Proposition 2 (non-diminishing marginal returns)**.** *There exist zero-sum games where $v_{k+1}^* - v_k^*$ is not non-increasing.*

Proposition 3 (non-submodularity)**.** *Let $S \subseteq [n]$ and $\Delta_S^n = \{x \in \Delta^n | x_i = 0 \quad \forall i \notin S\}$. Let $x_S^* = \text{argmax}_{x \in \Delta_S^n} \min_{y \in [m]} x^T Ay$ be the optimal commitment by Player 1 when restricted to playing only S, such that v_S^* is a set function (of S) denoting its corresponding utility. There exist zero-sum games where v_S^* is not submodular in S.*

We also remark that the counterexamples in Propositions 1 and 2 are not "over-engineered", we observe similar phenomena in our experiments. The above propositions have consequences for game solving. It is known that the problem of finding the optimal sparse strategy is NP-hard, even when the NE or the SSE can be found in polynomial time (McCarthy et al. 2018). Propositions 1, 2 and 3 go further, implying a combinatorial structure incompatible with some popular game solving approaches.

One such approach is that of *incremental strategy generation*, also known as *oracle-based methods*. These methods begin with a restricted (usually very small) set of actions, then an equilibrium is computed for the subgame restricted to those actions, and the subgames are iteratively expanded by querying *best-response oracles* for new actions to add. This is guaranteed to converge in finite though potentially exponential time (Zhang and Sandholm 2024). Oracle methods exploit the observation that in real-world games best responses can often be found or approximated efficiently, and usually the number of iterations of double oracle is low.

At first glance, incremental strategy generation seems appealing in our setting, as it naturally produces sparse strategies if the algorithm terminates before the restricted action set becomes larger than k. Unfortunately, it turns out that this approach, while guaranteed to produce k-sparse strategies, can yield low-quality commitments. For example, Proposition 1 implies that the restricted set must have size $\mathcal{O}(k^2)$ to include all optimal $1, 2, \ldots, k$-sparse strategies. In fact, Proposition 2 shows that for very large games (where the true NE may not be easily computed), one may be misled into prematurely terminating if increasing k does not improve v_k^*. Finally, Proposition 3 indicates that a greedy approach to expanding Player 1's commitment support does not ensure a good approximation, even in zero-sum games.

So far, our statements have been about the supports of the sparse equilibrium. Can we say something about the probabilities that each action is played? For instance, the algorithm of McCarthy et al. (2018) proposes finding "operationalizable" commitments where Player 1 is restricted to playing k-uniform strategies. Unfortunately, it turns out that k-uniform strategies can perform badly for small k.

Proposition 4. *Let $\bar{\Delta}_k^n = \{x \in \Delta^n | x_i = m/k, m \in \mathbb{Z}\}$ be the set of randomized strategies played uniformly over exactly k strategies (possibly repeated). Define the corresponding sparse equilibrium $\bar{x}_k^* = \text{argmax}_{x \in \bar{\Delta}_k^n} \min_{y \in [m]} x^T Ay$ and its value $\bar{v}_k^*$. Then, (i) $\bar{v}_k^*$ is no longer non-decreasing in k, and (ii) there exist games where finding an ϵ-NE operationalizable strategy $\bar{x}_k^*$ requires k to grow in a rate $\Omega(1/(\epsilon + c))$, where c is an adjustable game parameter.*

Proposition 4 suggests that finding optimal sparse equilibrium is not just about reasoning about which actions belong to the optimal support; one should also jointly reason about the probabilities with which these strategies are played. In the Appendix we show an instance where our definition of a sparse commitment gives a vastly different (and higher quality) solution compared to that of McCarthy et al. (2018). Notably, while McCarthy et al. (2018) argue that uniform strategies are easier to operationalize — our example shows that this comes at a price that could be large.

4 Finding Optimal k-sparse Commitments

The non-convex set Δ_k^n (a union of polytopes) can be represented via the following set of mixed-integer constraints:

$$1 = \sum_{a \in [n]} x(a), \quad k \geq \sum_{a \in [n]} z(a),$$

$$z(a) \geq x(a), \quad x(a) \in [0,1], \quad z(a) \in \{0,1\} \quad \forall a \in [n].$$

This gives a mixed integer linear program (MILP) for computing k-sparse equilibria, which we call the *basic method*:

$$\max_{g \in \mathbb{R}, x \in \Delta_k^n} g, \quad g \leq \sum_{a \in [n]} A(a,b)x(a) \quad \forall b \in [m]. \qquad \text{(B)}$$

Here, g represents the payoff which we are seeking to maximize, while the constraint upper bounds g over all possible opponent actions; such a formulation can be obtained by *dualizing* the inner minimization problem in (2). The basic method can be implemented easily via off the shelf MILP solvers. However, computation becomes prohibitively slow when n or m is large. We present methods to alleviate this.

(i) Single oracle methods when m is large but n is small When only Player 2 (the non-sparse player) has a large strategy space, we propose a *single-oracle* approach. For this, we assume access to Player 2's best response oracle $\text{BR}_2(x)$. We start with a subset of strategies $\mathcal{M}_0 \subseteq [m]$. At each iteration i, we compute the optimal commitment $x_k^{\mathcal{M}_i}$ against the non-sparse player's strategy space $\mathcal{M}_i$. A best-response strategy $\text{BR}_2(x_k^{\mathcal{M}_i}))$ is then computed for the non-sparse player against the sparse player's strategy $x_k^{\mathcal{M}_i}$, using the best-response oracle. The best-response is added to the non-sparse player's strategy space, i.e. $\mathcal{M}_{i+1} \leftarrow \mathcal{M}_i \cup \text{BR}(x_k^{\mathcal{M}}))$, and the MIP (B) is resolved. This iterative process continues until the equilibrium gap, defined as $\nabla = g - u(x_k^{\mathcal{M}}, \text{BR}(x_k^{\mathcal{M}}))$, where g is the sparse player's optimal expected utility computed by (B), is below a pre-specified threshold $\epsilon > 0$. When the algorithm terminates, we are guaranteed an ϵ-optimal k-sparse commitment.

(ii) MILP-based method when n is large but m is small When only n is large, the basic method has a prohibitively large number of integer variables. However, in many cases, the huge action space arises because of some *combinatorial structure*, e.g., in security applications, the action space for patrollers is often the set all length L paths in a directed graph. While n is extremely large, we may exploit this combinatorial structure to compute k-sparse commitments. We term these generally as *MILP-representable* action spaces, which have actions given by the nonempty set $\mathcal{K} = \{z \in \{0,1\}^l | Fz = a\}$, for constants $F \in \mathbb{R}^{n \times l}, a \in \mathbb{R}^n$. Clearly $\mathcal{K}$ is finite, but extremely large. However, we can rewrite a k-mixture of strategies more concisely as: $\mathcal{S} = \{x^{(i)} \in \mathbb{R}^m, z^{(i)} \in \mathcal{K}, t \in \Delta^k | x^{(i)} \geq t^{(i)} - (1 - z^{(i)}), x^{(i)} \leq t^{(i)}, x^{(i)} \leq z^{(i)}, x^{(i)} \geq 0\}$. Essentially, $\mathcal{S}$ "duplicates" $\mathcal{K}$ up to k times, where $x^{(i)}$ is some $z^{(i)}$ "rescaled" by the probability $t^{(i)}$ with which it is played. The constraints enforce that $x^{(i)} = t^{(i)} \cdot z^{(i)}$. For a given pure strategy $z \in \mathcal{K}$, we define the payoff function for all opponent strategies $b \in [m]$:

$$A_{z,b} = \max\{\mathcal{C}_{z,b}\}, \text{where } \mathcal{C}_{z,b} = \{c \in \mathbb{R} | C_b z + d_b \geq 1c\}.$$

Here, $C_b \in \mathbb{R}^{r \times m}$, $d_b \in \mathbb{R}^r$ are game constants for every action $b \in [m]$. Rescaling $\mathcal{C}_{z,b}$ to account for $t^{(i)}$ yields:

$$\max_{g \in \mathbb{R}, (x,z,t) \in \mathcal{S}, c \in \mathbb{R}^{k \times m}} g$$
$$g \leq \sum_{i \in [k]} c(i,b), C_b x^{(i)} + d_b \cdot t^{(i)} \geq 1c(i,b) \quad \forall b \in [m].$$

While somewhat complicated, allowing the cost function to be *non-linear* function of z admits a much wider class of problems, including games involving security interdiction (Černý et al. 2024; Zhang et al. 2017). As a concrete example, we return to the patrolling example where Player 1's action is to choose a length L path. Suppose Player 2's action is to select a subset b of exactly w distinct edges to hide evidence of criminal activity, where w is a game constant. Player 1 obtains a payoff of 1 if and only if its chosen path intersects at least 1 of Player 2's chosen edges; no double counting is allowed. Then $z^{(i)}$ is a binary vector of length ℓ (the number of graph edges) such that $z^{(i)}(e) = 1$, $z^{(i)}(e) = t^{(i)}$ for all edges e included in path i. For a path z, $A_{z,b}$ is *not* a linear function of z, since path z can intersect the set of edges chosen in b *multiple times*. Yet, by using the above formulation we can set $r = 2$, and define the first row of C_b to contain 1 on every edge included in action b, $d_b(1) = 0$, second row of $C_b = 0$, and $d_b(2) = 1$. Therefore the constraint forces $\mathcal{A}_{z,b}$ to be upper bounded by 1 and the number of times path z intersects the edges in b, avoiding the issue of double counting.

(iii) Combined methods when n and m are large The case when both n and m are large can be handled by combining methods (i) and (ii). Naturally, we would require Player 2 to be able to compute best responses to any mixed strategy "efficiently" for any Player 1 strategy with constant-sized support. This approach also requires solving the k-sparse MILP from Section (ii) "somewhat-efficiently."

4.1 Extensions to gen-sum Stackelberg equilibria

The *multiple LPs* method (Conitzer and Sandholm 2006) is a standard approach for finding optimal commitments. Combined with our method, it yields k-sparse general-sum Stackelberg equilibria by solving *multiple MILPs*, one for each action $b \in [m]$ of Player 2. Specifically, we solve the following MILP for each action:

$$\max_{x \in \Delta_k^n} \sum_{a \in [n]} A(a,b)x(a) \tag{G}$$
$$0 \geq \sum_{a \in [n]} (B(a,b') - B(a,b))\, x(a) \quad \forall b' \in [m].$$

The constraint guarantees x is restricted such that b is a best response to x. Note that for some choices of b this MILP could be infeasible, though at least one is feasible. At the end, we select the maximum value achieved over all of the MILPs that are feasible. We describe practical speedups and an alternative single-MILP method in the appendix.

4.2 Extensions to structured sparsity constraints

For some security applications, we may want Player 1's sparsity to be imposed in a more structured manner. For example, suppose we are planning patrols in a graph, such that actions are paths of a given length L. Suppose further that we are also required to decide the locations of k_1 and k_2 command posts at which paths are allowed to begin and end. Thus, we do not require sparsity over the path distribution, but rather sparsity in the *number of starting and ending vertices* over length L paths chosen with non-zero probability. Consequently, utilizing Δ_k^n is overly restrictive, as there may be paths sharing a starting or ending (or both) vertex.

More generally, we will define R *constraint sets* $\mathcal{S} = \{S^1, \ldots, S^R\}$. For each constraint set indexed by $i \in [R]$, we have *action sets* $S^i = \{S_1^i, \ldots, S_{q_i}^i\}$, where each $S_j^i \subseteq \mathcal{A}$. Each constraint set has a specified sparsity constraint k_i.

In the above patrolling example, $R = 2$, one for starting points, and another for ending points. Then q_1 (resp. q_2) is equal to the number of vertices, and S^1_j (resp. S^2_j) contains all length L paths starting from (resp. ending at) vertex j.

Note that the sets in S^i do not necessarily form a partition of $\mathcal{A}$; for a given constraint set i, an action $a \in \mathcal{A}$ can either be in one, multiple, or no action sets. For every action set, we incur a sparsity cost as long as at least one action it contains is played with positive probability. This yields the MILP:

$$
\begin{aligned}
1 &= \sum_{a\in[n]} x(a), && x(a) \in [0,1] \quad \forall a \in [n] \\
k_i &\geq \sum_{S^i_j \in S^i} z(S^i_j) && \forall i \in [R] \qquad \text{(S)} \\
z(S^i_j) &\geq x(a) && \forall a \in S^i_j,\ S^i_j \in S^i,\ i \in [R] \\
z(S^i_j) &\in \{0,1\} && \forall S^i_j \in S^i,\ i \in [R].
\end{aligned}
$$

5 Empirical Evaluation

We empirically analyze the performance and scalability of our proposed methods, focusing on (i) wall-clock computational time, (ii) the utility that Player 1 achieves for different support sizes k, and (iii) comparisons to k-uniform strategies. For purposes of comparison, we normalize running times using $t_{\text{normalized}} = \frac{t - t_{\min}}{t_{\max} - t_{\min}}$, where $t_{\min}$ and $t_{\max}$ are the minimum and maximum runtimes respectively. We normalize the utility value by the Nash value for zero-sum games and the Stackelberg equilibrium (SE) for general-sum games. Specifically, $u_{\text{normalized}} = \frac{u - u_{\min}}{u_{\text{equi}} - u_{\min}}$ where $u_{\min}$ is the utility at $k = 1$ and u_{equi} is the true value, i.e. $k = \infty$. We use Gurobi (Gurobi Optimization, LLC 2023) for all MILP solvers. Unless otherwise stated, for experiments with randomness, 30 instances were generated and solved. We report the standard errors in plots (usually negligible). Other experimental details are deferred to the Appendix.

5.1 Fully enumerable strategy spaces

We begin by analyzing enforced sparsity on games where n and m are small, under two different experiment setups.

Randomly generated normal-form games We consider both zero-sum and general-sum games and solve them using the proposed MIP (B) and multiple MILP (G) respectively. We generate payoff matrices A, B with sizes $n = m \in \{20, 30, \ldots, 80\}$. For zero-sum games, each A_{ij} is uniformly chosen in $[10, 100]$. In general-sum games, A_{ij} and B_{ij} are randomly chosen in $[-50, 50]$, under the condition that A_{ij} and B_{ij} have opposite signs.

Figure 1 shows the normalized runtime and relative utility as a function of k, each averaged over all instances per game size. For visualization purposes, we normalized the values of k with respect to the support size of NE (resp. SE) for zero-sum (resp. general-sum) games, denoted by k_{NE} (resp. k_{SE}). Thus, the x-axis represents the enforced support size w.r.t. the optimal strategy. Interestingly, the runtime curves exhibit a notable phase transition, i.e., the difficulty is concentrated in a specific range of k values. For example, for zero-sum games we find a "hard zone" when k lies between $0.2k_{NE}$ and $0.8k_{NE}$. For general-sum games, this hard zone was observed for values of $k \leq 0.2k_{SE}$. This pattern is reminiscent of the behavior observed when solving random SAT instances (Gent and Walsh 1994). Furthermore, the **utility curves** demonstrate that as the game size increases, we are still able to capture nearly 90% of the unrestricted optimal commitment, even with a fraction of the support size. When $n = m = 80$, about 90% of the game value is attained at around $0.25k_{NE}$ and $0.45k_{SE}$.

Patrolling games We evaluate our MIP (B) on a security motivated example using the layered graph framework of Černý et al. (2024). We design a pursuit evasion (PE) game on a university campus: a professor is trying to apprehend a student who stole exam questions. We constructed the game using a real university campus path network. Each player starts at one of several potential initial locations on campus corresponding to vertices, and selects a path of length $T = 5$ to traverse (also known as *depth*); the student is apprehended if and only if the chosen paths share an edge. We call a set of fixed starting points and desirable final locations a *scenario*. Individual games are then generated by assigning random (uniform from $[6, 10]$ for desirable and $[1, 5]$ for all others) rewards for escaped students based on their final location.

Figure 2 shows results across 30 instances when we impose sparsity on the distribution over the professor's paths. Just like random games, we again observe a phase transition pattern, with a pronounced "hard zone" when k ranges from 5 to 10. Furthermore, the professor is able to achieve more than 90% of the optimal even for k as low as 3.

Air defense battery placements Our task is to select which and where to place air defenses to maximize defensive coverage. Player 1's actions involve choosing (a) which, out of 4 different types of air defense batteries; those which provide stronger coverage have lower range and vice versa, and (b) where to place it's chosen battery. These locations are based off maps from the real world. Player 2 chooses a location to attack; the probability of success depends on how the defensive converges there. k-sparsity is sensible since batteries and ammunition dumps are expensive to procure and construct. Due to space constraints, we report detailed results in the Appendix. Nonetheless, our main observations are similar: while natural sparsity is quite low (e.g., < 20), we can still do well even with a small k (around 5). We also remark that the pathological behavior in Propositions 1 and 2 are seen in practice in these experiments.

5.2 Fully enumerable structured strategy spaces

We evaluate MILP (S) on experiments requiring structured sparsity, using the patrolling game (with $T = 5$) as the basis for two tests. In the first experiment, the professor chooses paths from 8 starting points but is limited to using k starting points, potentially with multiple paths per starting point. The results are reported in Figure 2. Interestingly, we no longer observe a clear phase transition. Again, we obtain almost maximum utility at just $k = 4$, i.e., around half of the possible starting points. Secondly, we experimented on a more

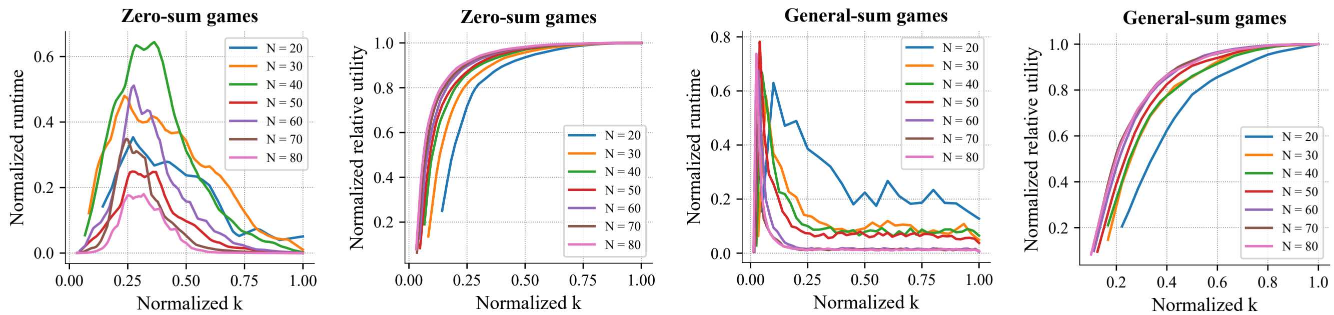

Figure 1: Average normalized runtime and relative utility for solving random zero-sum (**left**) and general-sum (**right**) games.

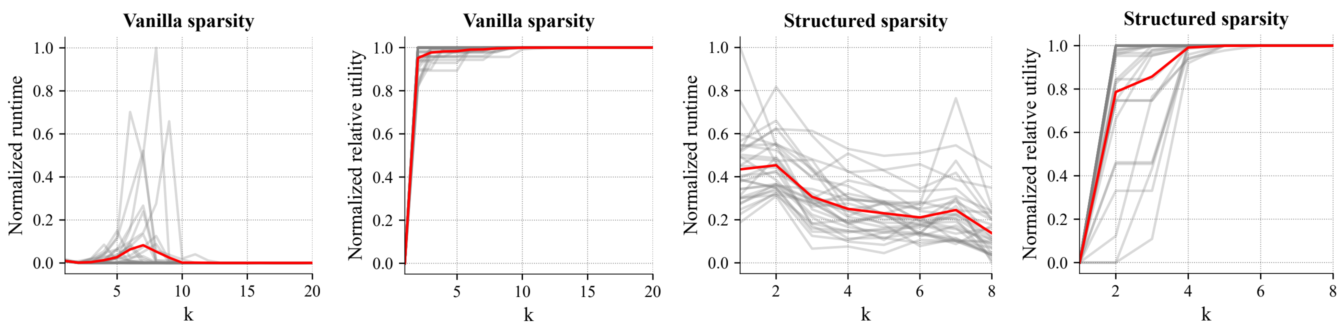

Figure 2: Average (red) normalized runtime and relative utility over 30 instances (gray) of a PE on a university campus. We impose *vanilla sparsity* on path distribution (**left**) and *structured sparsity* on the starting point distribution of Player 1 (**right**).

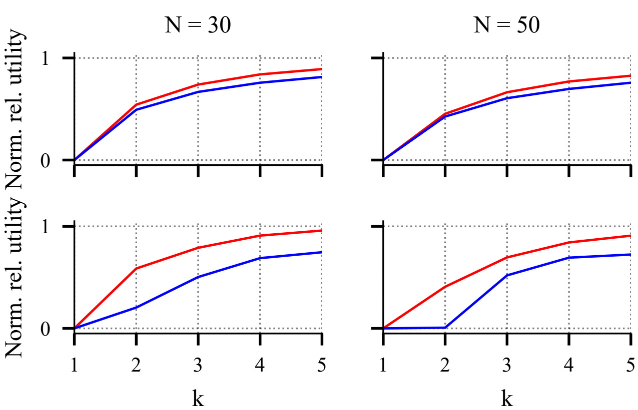

Figure 3: Average normalized relative utility for solving randomly generated zero-sum and general-sum games using optimal k-sparse (red) and k-uniform (blue) commitments. **Top:** zero-sum games. **Bottom:** general-sum games.

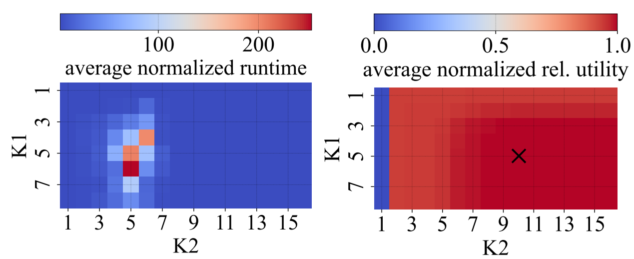

Figure 4: Average normalized runtime (left) and relative utility (right) for solving a pursuit-evasion game on a university campus, imposing *structured sparsity* on the distribution over starting points (k_1) and paths (k_2) of Player 1. The smallest values of k_1 and k_2 at which the Nash value is attained is marked by "$\times$".

complex setting with sparsity constraints k_1 and k_2 on starting points and paths respectively. The runtime heat map in Figure 4 reveals a "2 dimensional hard zone", with the hardest instances occurring at k_1 between $[4, 7]$ and k_2 in $[4, 6]$.

5.3 Scaling up to large strategy spaces

Next, we examine settings where the basic MIP approach fails due to the large strategy space of one or both players. Using patrolling game variants, we test three setups: (i) large n, small m; (ii) small n, large m; and (iii) large n and m. For each setup, we run 30 instances and report average relative utility results for instances solved within a 12-hour limit.

(i) Sparse player has a large action space Here, only Player 1 has a large action space, which we increase by varying T from 5 to 8 (note that n grows exponentially with the path length T). We observe in Figure 5 (top) that for $T \in \{5, 6, 7\}$, Player 1 achieves more than 90% of the optimal game value for k as low as 3. Even at $T = 8$, we obtain a respectable 80% of the true value when $k = 3$.

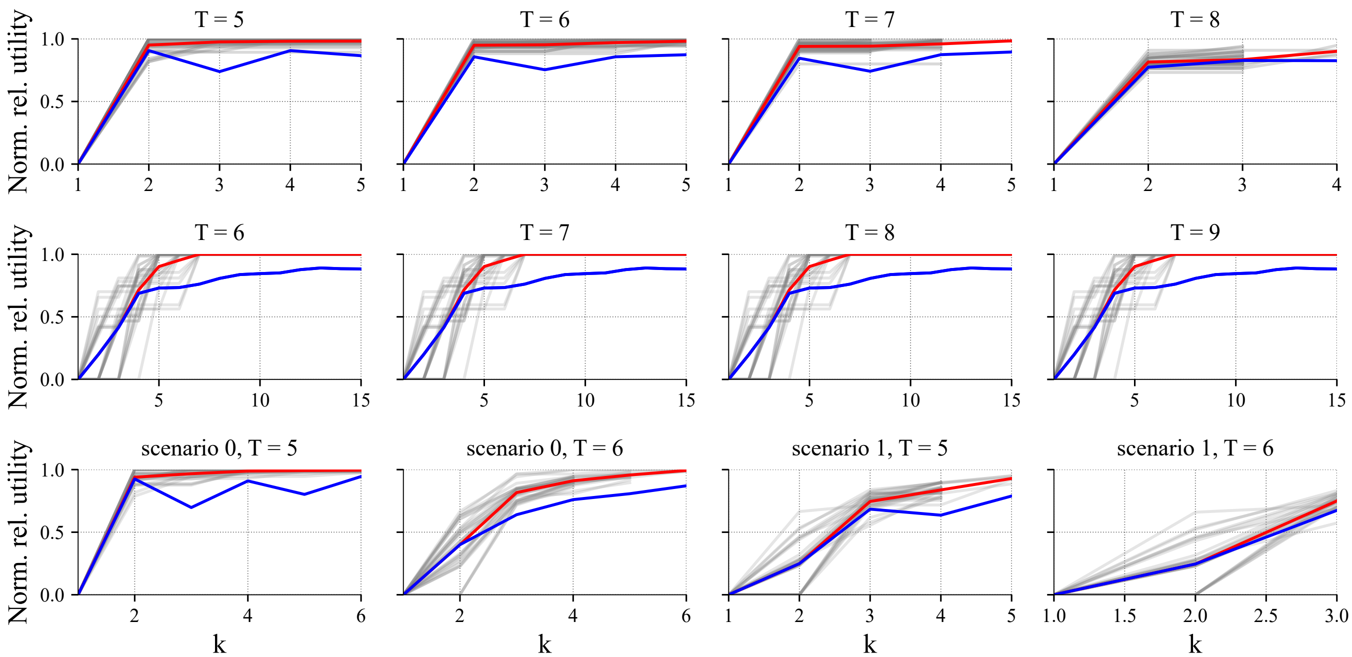

Figure 5: Average normalized relative utility curves for optimal k-sparse (red) and k-uniform (blue) commitments. Results for individual instances using optimal k-sparse commitments are in gray. **Top row:** sparse player has a large strategy space. **Middle row:** non-sparse player has a large strategy space. **Bottom row:** both players have large strategy spaces.

(ii) Non-sparse player has a large action space In this setting, Player 2 has a large action space whereas Player 1 places only a single patrol on a single intersection in the graph, which then stays there. We vary the depth of the graph from 6 to 9. The relative utility plots reported in Figure 5 (middle) show that Player 1 captures around 90% of the true Nash value at $k = 3$ for all depths. Notice however that the plots are almost identical. The reason behind this is that increasing depth does seem to provide Player 2 with no strategic advantage, Player 1 hence positions the patrol on the same intersections as with smaller T.

(iii) Both players have large action spaces Finally, we examine the case where both Player 1 and 2 have large action spaces. We study two scenarios differing in the locations of the players' starting points and desirable exits. We present in Figure 5 the plots for $T = \{5, 6\}$ across the two scenarios. In scenario 0, Player 1 gains high utility value at small k, notably for depth 5. Solving for the optimal k-sparse strategy at depth 6 in scenario 1 proves to be challenging. While some instances achieve approximately 80% of the Nash value, a support size of 3 remains insufficient for attaining higher utility. We point out that here, there are $\sim$ 16k paths per player, accounting for the relatively poor performance.

5.4 Comparison with k-uniform strategies

We compare the average relative utility achieved by Player 1 when using the optimal k-sparse commitment against the k-uniform mixed strategy proposed by McCarthy et al. (2018). Again, we consider random normal-form games and patrolling games with large action spaces. For the former, we generate zero-sum and general-sum games as described earlier, considering games where $n = m \in 30, 50$. We limit the x-axis to $k = 5$ to emphasize small support sizes, which are of primary interest. As shown in Figure 3, the results indicate that k-sparse commitments generally outperform k-uniform strategies, especially when k is small. This difference is more pronounced in general-sum games.

Figure 5 depicts the results in patrolling games, where the blue curves represent the average relative utility normalized by the Nash value when the sparse player employs k-uniform strategies. Consistent with the findings from random games, we observe that optimal k-sparse strategies provide higher utility for the sparse player, particularly when the support size is small. For instance, for $T = 5, 6, 7$, the sparse player secures over 95% of the Nash value with $k = 3$ when both players have large strategy spaces, whereas a uniform strategy yields less than 80%.

6 Conclusion

This paper studies properties and algorithms for finding k-sparse commitments for two player games. Our proposed method performs competitively with non-sparse commitments even for very small values of k. Furthermore, we have proposed extensions to handle Stackelberg commitments in general-sum games, structured sparsity, as well as cases with a large number of player actions. Our notion of k-sparse commitments often leads to significantly better performance when compared to k-uniform strategies. Future work include extensions to extensive-form or multiplayer games, exploring approximate sparsity using sparsity-inducing regularization techniques, or online learning of sparse strategies.

Acknowledgments

This research was supported by the Office of Naval Research awards N00014-22-1-2530 and N00014-23-1-2374, and the National Science Foundation awards IIS-2147361 and IIS-2238960. Part of this work was carried out during a time when CKL was still a postdoctoral research scientist at Columbia University, before transitioning to NUS.

References

Althöfer, I. 1994. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199: 339–355.

Anagnostides, I.; Kalavasis, A.; Sandholm, T.; and Zampetakis, M. 2023. On the Complexity of Computing Sparse Equilibria and Lower Bounds for No-Regret Learning in Games. *arXiv preprint arXiv:2311.14869*.

Babichenko, Y.; Barman, S.; and Peretz, R. 2014. Simple approximate equilibria in large games. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 753–770.

Bakhtin, A.; Brown, N.; Dinan, E.; Farina, G.; Flaherty, C.; Fried, D.; Goff, A.; Gray, J.; Hu, H.; et al. 2022. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624): 1067–1074.

Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2015. Heads-up limit hold'em poker is solved. *Science*, 347(6218): 145–149.

Brown, G. W. 1951. Iterative Solution of Games by Fictitious Play. In Koopmans, T. C., ed., *Activity Analysis of Production and Allocation*. New York: Wiley.

Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.

Černý, J.; Ling, C. K.; Chakrabarti, D.; Zhang, J.; Farina, G.; Kroer, C.; and Iyengar, G. 2024. Contested Logistics: A Game Theoretic Approach. In *International Conference on Decision and Game Theory for Security*. Springer.

Černý, J.; Ling, C. K.; Kroer, C.; and Iyengar, G. 2024. Layered Graph Security Games. In Larson, K., ed., *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 2695–2703. International Joint Conferences on Artificial Intelligence Organization.

Conitzer, V.; and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*, 82–90.

Fang, F.; Stone, P.; and Tambe, M. 2015. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

Feder, T.; Nazerzadeh, H.; and Saberi, A. 2007. Approximating Nash equilibria using small-support strategies. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, 352–354.

Foster, D. J.; Golowich, N.; and Kakade, S. M. 2023. Hardness of independent learning and sparse equilibrium computation in markov games. In *International Conference on Machine Learning*, 10188–10221. PMLR.

Gent, I. P.; and Walsh, T. 1994. The SAT phase transition. In *ECAI*, volume 94, 105–109. PITMAN.

Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual.

Hart, S.; and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5): 1127–1150.

Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordónez, F. 2010. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4): 267–290.

Lipton, R. J.; Markakis, E.; and Mehta, A. 2003. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, 36–41.

McCarthy, S. M.; Laan, C. M.; Wang, K.; Vayanos, P.; Sinha, A.; and Tambe, M. 2018. The price of usability: Designing operationalizable strategies for security games. IJCAI.

Perolat, J.; De Vylder, B.; Hennes, D.; Tarassov, E.; Strub, F.; de Boer, V.; Muller, P.; Connor, J. T.; Burch, N.; Anthony, T.; et al. 2022. Mastering the game of Stratego with model-free multiagent reinforcement learning. *Science*, 378(6623): 990–996.

Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, 125–132.

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

Sinha, A.; Fang, F.; An, B.; Kiekintveld, C.; and Tambe, M. 2018. Stackelberg security games: Looking beyond a decade of success. IJCAI.

Tambe, M. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge university press.

Von Stengel, B.; and Zamir, S. 2004. Leadership with commitment to mixed strategies. Technical report, Citeseer.

Zhang, B. H.; and Sandholm, T. 2024. Exponential Lower Bounds on the Double Oracle Algorithm in Zero-Sum Games. *arXiv preprint arXiv:2405.06797*.

Zhang, Y.; An, B.; Tran-Thanh, L.; Wang, Z.; Gan, J.; and Jennings, N. R. 2017. Optimal escape interdiction on transportation networks.