

Learning Convex Partitions and Computing Game-theoretic Equilibria from Best Response Queries

Paul W. Goldberg *

Francisco J. Marmolejo-Cossío *†

Abstract

Suppose that an m -simplex is partitioned into n convex regions having disjoint interiors and distinct labels, and we may learn the label of any point by querying it. The learning objective is to know, for any point in the simplex, a label that occurs within some distance ε from that point. We present two algorithms for this task: Constant-Dimension Generalised Binary Search (CD-GBS), which for constant m uses $\text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries, and Constant-Region Generalised Binary Search (CR-GBS), which uses CD-GBS as a subroutine and for constant n uses $\text{poly}(m, \log(\frac{1}{\varepsilon}))$ queries.

We show via Kakutani’s fixed-point theorem that these algorithms provide bounds on the best-response query complexity of computing approximate well-supported equilibria of bimatrix games in which one of the players has a constant number of pure strategies. We also partially extend our results to games with multiple players, establishing further query complexity bounds for computing approximate well-supported equilibria in this setting.

Keywords: Query protocol, equilibrium computation, revealed preferences

1 Introduction

The computation of game-theoretic equilibria is a topic of long-standing interest in the algorithmic and AI communities. This includes computation in the “classical” setting of complete information about a game, as well as settings of partial information, communication-bounded settings, and distributed algorithms (for example, best-response dynamics). A recent line of research has studied computation of equilibria based on query access to players’ payoff functions. That work, along with the notion of revealed preferences in economics, inspires the new setting we study here.

We study algorithms that have query access to the players’ best-response behaviour: an algorithm may query a mixed-strategy profile (i.e. probability distributions constructed by the algorithm, over each player’s pure strategies) and learn the players’ best responses. Our focus is on standard bimatrix games, which is arguably the most natural starting-point for an investigation of this new query model. The solution concept of interest is ε -approximate Nash equilibria (exact equilibria are typically impossible to find using finitely many such queries, see Corollary 3). A basic challenge is to identify algorithms that achieve this goal with good bounds on their query complexity (and also, ideally, their runtime complexity).

*Department of Computer Science, University of Oxford. Emails: paul.goldberg@cs.ox.ac.uk, francisco.marmolejo@cs.ox.ac.uk

†Supported by the Mexican National Council of Science and Technology (CONACyT)

In more detail, we assume an $m \times n$ game G : a row player has m pure strategies and a column player has n pure strategies. G has two unknown $m \times n$ payoff matrices that represent payoffs to the players for all combinations of pure strategy choices they may make. A query consists of a probability distribution over the pure strategies of one of the players, and elicits an answer consisting of a best response for the other player (i.e. a pure strategy that maximises that player’s expected payoff). We seek an ε -well-supported Nash equilibrium (ε -WSNE): a pair of probability distributions over their pure strategies with the property that any strategy of player p whose expected payoff is more than ε below the value of p ’s best response, gets probability zero. The general question of interest is: how many queries are needed, as a function of m, n, ε .

Using Kakutani’s fixed point theorem, we reduce this question to a novel and more geometrical challenge in the design of query protocols. Suppose that the m -simplex Δ^m is partitioned into n convex regions having labels in $[n] = \{1, \dots, n\}$. When we query a point $x \in \Delta^m$ we are told the label of x . How many queries (in terms of m, n, ε) are needed in order to ensure that all points in Δ^m are within ε of a point whose label we know? We show how to achieve this using time and queries polynomial in $\log \varepsilon$ and $\max(m, n)$ provided that $\min(m, n)$ is constant. This leads to a polynomial query complexity algorithm for 2-player games, provided that one of the players has a constant number of strategies.

1.1 Further details

In essence, we consider partitions of the unit m -simplex Δ^m into n convex polytopes, P_1, \dots, P_n , with disjoint interiors, and aim to approximately learn the partition with access to a membership oracle that for a given $x \in \Delta^m$, returns a polytope to which x belongs. The notion of approximation we study is that of ε -close labellings, a collection of empirical polytopes, $\{\hat{P}_i\}_{i=1}^n$, such that $\hat{P}_i \subseteq P_i$ for $i = 1, \dots, n$ and $\cup_{i=1}^n \hat{P}_i$ is an ε -net of $\Delta^m \subset \mathbb{R}^m$ in the ℓ_2 norm.

Note that in one dimension ($m = 1$) we can use binary search to solve this problem using $n \log(1/\varepsilon)$ queries. We generalise to higher dimension, exploiting convexity of the regions to reduce query usage in computing ε -close labellings. We present two algorithms for this task: Constant-Dimension Generalised Binary Search (CD-GBS), which for constant m uses $\text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries, and Constant-Region Generalised Binary Search (CR-GBS), which uses CD-GBS as a subroutine and for constant n uses $\text{poly}(m, \log(\frac{1}{\varepsilon}))$ queries.

This problem derives from the question of how to compute approximate (well-supported) Nash equilibria (ε -WSNE) using only best response information, obtained via queries in which the algorithm selects a mixed strategy profile and a player, and receives a best response for that player to the mixed profile. Via Kakutani’s fixed-point theorem [20] we reduce this variant of equilibrium computation to finding ε -close labellings of polytope partitions. For $m \times n$ games where m is constant (or n equivalently, by symmetry), we show that an ε -WSNE can be computed using $\text{poly}(n, \log(\frac{1}{\varepsilon}))$ best response queries.

In addition, we briefly delve into the problem of computing ε -WSNE in multiplayer games with best response queries. Unfortunately, as soon as there are more than two players, the geometric connection between computing ε -WSNE and learning polytope partitions of the simplex breaks down. Nonetheless, fixed-point techniques from Section 8 can still be applied in this setting, and we present a simple algorithm that computes an ε -WSNE with a finite query complexity. To be more specific, in a game with n players each having k actions, our algorithm computes an ε -WSNE using $O\left(n \binom{nk}{\varepsilon}^{nk}\right)$ best response queries.

1.2 Related Work

Earlier work in computational learning theory has studied exact learning of geometrical regions over a discretised domain, where algorithms are sought with query complexity logarithmic in a resolution parameter and binary search is repeatedly applied in a systematic way [7]. Goldberg and Kwek [13] specifically study the learnability of polytopes in this context, deriving query efficient algorithms, and precisely classifying polytopes learnable in this setting. These algorithms can be adapted to approximately learn a single polytope with membership queries, but the obtained notion of approximation is not directly applicable to computing ε -close labellings.

The Nash equilibrium (NE) is a fundamental concept in game theory [23]. They are guaranteed to exist in finite games, yet computational challenges in finding one abound, most notably, the PPAD-completeness of computing an exact equilibrium even for two-player normal form games [8, 10]. For this reason, query complexity has been extensively used as a tool to differentiate hardness of equilibrium concepts in games. For payoff queries, some notable examples include: exponential lower bounds for randomised computation of exact Nash equilibria and exact correlated equilibria via communication complexity lower bounds in multiplayer games [17, 18]; exponential lower bounds for randomised computation of approximate well-supported equilibria and general approximate equilibria for a small enough approximation factor in multiplayer games [1]; upper and lower bounds for equilibrium computation in bimatrix games, congestion games [11] and anonymous games [16]; upper and lower bounds for randomised algorithms computing approximate correlated equilibria [14]. Babichenko et al. have also proved lower bounds in communication complexity for computing ε -WSNE for small enough ε in both bimatrix and multiplayer games [2].

Best response queries are a weaker but natural query model which is powerful enough to implement fictitious play, a dynamic first proposed by Brown [6], and proven to converge by Robinson [24] in two-player zero-sum games to an approximate NE. Fictitious play does not always converge for general games where both players have more than two strategies [12]. Furthermore, Daskalakis and Pan have proven that the rate of convergence of the dynamic is quite slow in the worst case (with arbitrary tie-breaking) [9]. Also, beyond non-convergence, the dynamic can have a poor approximation value for general games [15]. In addition, the relationship between best responses and convex partitions of simplices has been studied by Von Stengel [25] in the context of sequential games where one player has to commit to and announce a strategy before playing.

For a bimatrix game, simple ε -close labellings can be constructed by querying best responses at mixed strategies arising as uniform distributions over sufficiently large multisets of pure strategies. As a consequence of our main theorem, best responses to these multiset distributions contain enough information to compute approximate WSNE. This result is in the spirit of [3] and [22], who aim to quantify specific k such that some approximate equilibrium arises as a uniform mixture over multisets of size k . We note in our scenario that there is also a guaranteed existence of an approximate equilibrium using sufficiently large multisets, however *verifying* that a *specific* pair of mixed strategies is an approximate WSNE is not straightforward using only best response queries. This is in contrast to the verification of approximate equilibria via utility queries as studied in [3].

Separately, we note that the present paper is possibly relevant to the search for a price equilibrium in certain markets. Baldwin and Klemperer study markets consisting of *strong-substitutes* demand functions for N different goods available in multiple discrete units [4]. These markets are a generalisation of the *product-mix auction* of [21]; a basic task is to identify prices at which some desired bundle of the goods is demanded. Consider the space $(\mathbb{R}^+)^N$ of all price vectors. As analysed in [4], a strong-substitutes demand function partitions this price space into convex polytopes,




Figure 1: Polytope partition, cross-section and slices.

each of which comprises the prices at which some particular bundle of goods is demanded. So, the present paper relates to a setting where price vectors may be queried, and responses consist of demand bundles. The connection is imperfect, since the main objective in the context of [4] would be to learn a price at which some target bundle is demanded, rather than the entire demand function. The ideas here may be useful for learning the values that the market has for various bundles.

2 Preliminaries and Notation

Our main object of study will be families of polytopes that precisely cover the unit simplex, with the property that any two distinct polytopes from the family are either disjoint, meet at their boundary, or entirely coincide. Throughout, the polytopes we work with are convex.

Definition 1 ((m, n) -Polytope Partition). *A (m, n) -polytope partition consists of a set of n convex polytopes in \mathbb{R}^m , $\mathcal{P} = \{P_1, \dots, P_n\}$, with the following properties:*

- $\bigcup P_i = \Delta^m = \{x \in \mathbb{R}^m \mid \forall i, x_i \geq 0, \sum_i x_i \leq 1\}$.
- For each $i \neq j$, either $\text{relint}(P_i) \cap \text{relint}(P_j) = \emptyset$ or $P_i = P_j$, where $\text{relint}(H)$ means the relative interior of H .

Definition 2 (Cross-sections and Slices). *Let $P \subset \mathbb{R}^m$ be a polytope and $\pi : \mathbb{R}^m \rightarrow \mathbb{R}$ the projection function into the first coordinate. For $x \in \mathbb{R}$, we define the x -cross-section of P as $P^x = \pi^{-1}(x) \cap P$. For any $I = [x, y] \subset \mathbb{R}$ we define the $[x, y]$ -slice of P as $P^I = P^{x,y} = \pi^{-1}([x, y]) \cap P$. Suppose that $\mathcal{P} = \{P_i\}_i$ is an (m, n) -polytope partition. The definitions of cross-sections and slices extend to $\mathcal{P}^x = \{P_i^x\}_i$ and $\mathcal{P}^I = \mathcal{P}^{x,y} = \{P_i^{x,y}\}_i$.*

Figure 1 gives a visualisation of these two definitions. Notice that in the same figure, \mathcal{P}^x is essentially a lower-dimensional polytope partition linearly scaled by a factor of $(1 - x)$. This however, is not the case in general, as visible in Figure 2, where \mathcal{P}^x fails the second condition of Definition 1. We distinguish between these two scenarios with the following formal definition:

Definition 3 (Non-Degenerate and Degenerate cross-sections). *Let \mathcal{P} be an (m, n) -polytope partition. For $x \in [0, 1)$ let $f_x : \mathcal{P}^x \rightarrow \Delta^{m-1}$ be defined as $f_x(v_1, \dots, v_m) = \frac{1}{1-x}(v_2, \dots, v_m)$. If $f_x(\mathcal{P}^x)$ is an $(m-1, n)$ -polytope partition, we say that \mathcal{P}^x is a non-degenerate cross-section. Otherwise we say that \mathcal{P}^x is a degenerate cross-section.*

The recursive structure of polytope partitions on non-degenerate cross-sections will be crucial to our constructions. Luckily, for any polytope partition, there are only a finite number of points $x \in [0, 1)$ that give rise to degenerate cross-sections. Before showing this, we define an important discrete subset of $[0, 1]$ given by the projections of vertices of polytopes under π .

Definition 4 (Vertex Critical Coordinates). *For a given polytope $P \subset \mathbb{R}^m$ let $V_P \subset \mathbb{R}^m$ be the vertex set of P . Define the set of vertex critical coordinates as $C_P = \pi(V_P) \subset \mathbb{R}$. If $\mathcal{P} = \{P_i\}_{i=1}^n$ is an (m, n) -polytope partition, then we extend our definition to define $C_{\mathcal{P}} = \bigcup_{i=1}^n C_{P_i} \subset [0, 1]$ as the vertex critical coordinates of \mathcal{P} .*

Lemma 1. *Suppose that \mathcal{P} is an (m, n) -polytope partition and that $x \in [0, 1) \setminus C_{\mathcal{P}}$. Then \mathcal{P}^x is non-degenerate.*

Proof. First we show that if $P \subset \mathbb{R}^m$ is an arbitrary polytope and $x \in \mathbb{R} \setminus C_P$ then $\text{relint}(P^x) = \text{relint}(P) \cap \pi^{-1}(x)$.

First of all, we notice that $P \neq P^x$ since we have assumed that x is not the projection of a vertex of P . Suppose that the affine dimension of P is $k \leq m$ so that P is full dimensional in the affine subspace H of dimension k . Let $z \in \text{relint}(P^x) \subset P^x$. Clearly $\pi(z) = x$, hence we simply need to show that $z \in \text{relint}(P)$. Suppose that this is not the case, then z lies on some boundary hyperplane to P in H . Call this boundary hyperplane D . D cannot lie entirely in $\pi^{-1}(x)$ due to the fact that x is not a critical coordinate. It follows that $D \cap \pi^{-1}(x)$ is thus a boundary hyperplane to P^x . This contradicts the fact that $z \in \text{relint}(P^x)$, thus establishing the fact that $z \in \text{relint}(P) \cap \pi^{-1}(x)$.

Suppose that $z \in \text{relint}(P) \cap \pi^{-1}(x)$. Since $\text{relint}(P) \subset P$, we know that $z \in P_x$. Furthermore, $z \in \text{relint}(P)$ means that for some $\varepsilon > 0$, the k -dimensional ball $B_\varepsilon(z) \cap H$ is entirely contained in P . Clearly this also holds for the $k-1$ dimensional ball $B_\varepsilon(z) \cap H \cap \pi^{-1}(x)$, establishing the fact that $z \in \text{relint}(P^x)$.

Let us return to the lemma statement which involves a polytope partition \mathcal{P} instead of a single polytope P . If $x \in [0, 1) \setminus C_{\mathcal{P}}$ then we have shown $\text{relint}(P_i^x) = \text{relint}(P_i) \cap \pi^{-1}(x)$ for all $P_i \in \mathcal{P}$, which from the fact that \mathcal{P} satisfies the second condition of Definition 1 establishes the fact that \mathcal{P}^x also satisfies this second condition. The fact that \mathcal{P}^x satisfies the first condition of Definition 1 trivially follows from the fact that \mathcal{P} covers Δ^m as per the first condition of Definition 1. □

2.1 Query Oracle Models

We study two natural query oracle models for polytope membership in any \mathcal{P} .

Definition 5 (Membership Query Oracles for Polytope Partitions). *Any (m, n) -polytope partition, \mathcal{P} has the following membership query oracles:*

- *Lexicographic query oracle: $Q_\ell : \Delta^m \rightarrow [n]$, which for a given y returns the smallest index of polytope to which y belongs, namely $Q_\ell(y) = \min\{i \in [n] \mid y \in P_i\}$.*




Figure 2: Degenerate cross-section at x

- *Adversarial query oracle(s):* $Q_A : \Delta^m \rightarrow [n]$, which can return any polytope to which y belongs. Namely Q_A is any function such that $Q_A(y) \in \{i \in [n] \mid y \in P_i\}$ for all $y \in \Delta^m$.

When we wish to refer to an arbitrary oracle from the above models, we use the notation Q . Before continuing, we also clarify that for $A, B \subseteq \mathbb{R}^m$, we denote $\text{Conv}(A, B) \subseteq \mathbb{R}^m$ as the convex combination of the two sets. In addition, if $A_i \subseteq \mathbb{R}^m$ is an indexed family of sets with $i = 1, \dots, r$, we denote $\text{Conv}(A_i \mid i = 1, \dots, r) \subseteq \mathbb{R}^m$ as the convex combination of all A_i .

2.2 ε -close Labellings

Upon making queries to Q , we can infer labels of $x \in \Delta^m$ by taking convex combinations. We abstract this notion in the following definition.

Definition 6 (Empirical Polytopes and Labellings). *Suppose that \mathcal{P} is an (m, n) -polytope partition and $S \subset \Delta^m$ is a finite set for which queries to Q have been made. Let $\hat{P}_i = \text{Conv}(\{x \in S \mid Q(x) = i\}) \subset P_i$. We say each \hat{P}_i is an empirical polytope of P_i and that $\hat{\mathcal{P}} = \{\hat{P}_i\}$ is an empirical labelling of \mathcal{P} . Furthermore, we use the notation $\hat{P}_\perp = \Delta^m \setminus \cup_{i=1}^n \hat{P}_i$ to refer to points in Δ^m unlabelled under $\hat{\mathcal{P}}$.*

An ε -net in the ℓ_2 norm for $\Delta^m \subset \mathbb{R}^m$ is a set $N_\varepsilon^m \subseteq \Delta^m$ with the property that for all $x \in \Delta^m$, there exists a $y \in N_\varepsilon^m$ such that $\|x - y\|_2 \leq \varepsilon$. Our learning goal is to use query access to an oracle, Q , to compute an empirical labelling $\hat{\mathcal{P}}$ such that $\cup_{i=1}^n \hat{P}_i$ is an ε -net of Δ^m .

Definition 7 (ε -close Labelling). *Suppose that $\varepsilon \geq 0$ and that $\hat{\mathcal{P}}$ is an empirical labelling for \mathcal{P} . If $\cup_{i=1}^n \hat{P}_i$ is an ε -net of $\Delta^m \subset \mathbb{R}^m$ in the ℓ_2 norm, we say that $\hat{\mathcal{P}}$ is an ε -close labelling of \mathcal{P} .*

Although ε -close labellings are defined for polytope partitions, we extend our terminology to also encompass slices of polytope partitions. As such, when we mention computing an ε -close labelling of $\mathcal{P}^{x,y}$, we mean an empirical labelling of $\mathcal{P}^{x,y}$ (in the same vein as Definition 6) with the property that the union of its empirical polytopes forms an ε -net of $(\Delta^m)^{x,y}$.

2.3 Learning in Thickness to Learning in Distance

Definition 8 (Thickness of Sets). *Suppose that $Z \subseteq \mathbb{R}^m$ is a set. We define the thickness of Z as the radius of the largest ℓ_2 ball fully contained in Z and we denote it by $\tau(Z) = \sup\{\delta \geq 0 \mid \exists x \in Z \text{ with } B_\delta(x) \subseteq Z\}$ where $B_\delta(x) = \{y \in \mathbb{R}^m \mid \|x - y\|_2 \leq \delta\}$. In the language of convex geometry, $\tau(Z)$ is the depth of the Chebyshev centre of Z .*

For a polytope partition \mathcal{P} , if $\widehat{\mathcal{P}}$ is an ε -close labelling, then $\tau(\widehat{P}_\perp) \leq \varepsilon$, but the converse does not hold in general. Even though \widehat{P}_\perp may be of small thickness, if it contains vertices of Δ^m , these vertices may be far from labelled points. The following results lead up to Lemma 4, a slightly weaker version of the converse. Lemma 4 shows that if we are able to learn an empirical labelling where the set of unlabelled points is of small enough thickness, then we will in fact have succeeded in learning an ε -close labelling, where any unlabelled point is close in distance to a labelled point.

Lemma 2. *Let $P \subset \mathbb{R}^m$ be a full-dimensional polytope with $\text{Diam}(P) = \sup_{x,y \in P} \|x - y\|_2$.*

- *If $A \subsetneq P$ and $\gamma > \left(\frac{\text{Diam}(P)}{\tau(P)}\right) \tau(A)$, then $B_\gamma(x) \cap (P \setminus A) \neq \emptyset$ for all $x \in A$.*
- *If $A \subseteq P$ is such that $\text{int}(P) \setminus A \neq \emptyset$ ($\text{int}(P)$ refers to the interior of P) and $\gamma > \left(\frac{\text{Diam}(P)}{\tau(P)}\right) \tau(A)$, then $B_\gamma(x) \cap (\text{int}(P) \setminus A) \neq \emptyset$ for all $x \in A$.*

Proof. The proof of the first claim follows by considering the picture given in Figure 3. Pick an arbitrary $x \in A$. Due to the definition of thickness, there exists some $v \in P$ such that $B_{\tau(P)}(v) \subset P$. Consider the convex combination, $\text{Conv}(x, B_{\tau(P)}(v)) \subset P$. The furthest point in this convex combination from x is at the other extreme of $B_{\tau(P)}(v)$ from x , and we denote the distance between these two points by $z = \sup_{a \in B_{\tau(P)}(v)} \|x - a\|_2 \leq \text{Diam}(P)$. By similarity however, it now follows that if we consider $B_\gamma(x)$, and the fact that $\text{Diam}(P) \geq z$, a similar inscribed sphere of radius strictly greater than $\tau(A)$ will exist within $F = B_\gamma(x) \cap \text{Conv}(x, B_{\tau(P)}(v)) \subset B_\gamma(x) \cap P$. By definition, $\tau(F) > \tau(A)$. It follows that $F \not\subseteq A$, which proves the claim as $F \subset B_\gamma(x)$.

As for a proof of the second claim, it follows by considering the same picture above and noticing that $\text{int}(\text{Conv}(x, B_{\tau(P)}(v))) \subset \text{int}(P)$ as well as the fact that the former set is non-empty since P is of full dimension. \square

Lemma 3. *$\text{Diam}(\Delta^m) = \sqrt{2}$ and $\tau(\Delta^m) \geq \frac{1}{m+\sqrt{m}}$.*

Proof. For the first part of the statement, let us fix an $x \in \Delta^m$. If we consider the function $f_x(z) = \|x - z\|_2^2$, then this function is differentiable and clearly achieves local maximum when z is a vertex of Δ^m . It thus follows that the distance between two points in Δ^m is maximised when both are vertices. This in turn is maximal when both points are vertices not equal to the zero vector, in which case they are at distance $\sqrt{2}$ from each other.

As for the second part of the claim, we explicitly construct an inscribed sphere of the desired radius. Let $\lambda = \frac{1}{m+\sqrt{m}}$, and define $x = \lambda \vec{1} \in \Delta^m$. Clearly $B_\lambda(x)$ is tangent to Δ^m on axis-aligned faces (defined by the set of all z such that $\pi_i(z) = 0$ in the positive orthant). The remaining face is given by the set of z in the positive orthant such that $\|z\|_1 = 1$. The most extremal point of $B_\lambda(x)$ in the direction of this face is given by $\frac{1}{m} \vec{1}$, hence the sphere is inscribed in Δ^m . \square




Figure 3: Proof of Lemma 2

Lemma 4. *Suppose that \mathcal{P} is an (m, n) -polytope partition. Furthermore suppose that $\widehat{\mathcal{P}}$ is an empirical labelling with $\tau(\widehat{P}_\perp) < \varepsilon$. For any $\gamma > \sqrt{2}(m + \sqrt{m})\varepsilon$, it follows that $\widehat{\mathcal{P}}$ is a γ -close labelling. In particular, if $\gamma > 4m\varepsilon$, the claim also holds.¹*

Proof. From Lemma 3, we know that $\tau(\Delta^m) \geq \frac{1}{m+\sqrt{m}}$ and $\text{Diam}(\Delta^m) = \sqrt{2}$. Suppose that $x \in \widehat{P}_\perp$. From Lemma 2 our choice of γ implies $B_\gamma(x) \cap (\Delta_m \setminus \widehat{P}_\perp) \neq \emptyset$. This in turn means that $\widehat{\mathcal{P}}$ is a γ -close labelling. As for the final claim, this holds since $m \geq 1$. \square

3 Constant-Dimension Generalised Binary Search for Q_ℓ

Let us first build some intuition for why generalisations of binary search lead to query efficient algorithms for computing ε -close labellings of (m, n) -polytope partitions.

Finding an ε -close labelling of a $(1, n)$ -polytope partition using a lexicographic oracle is the same as approximately learning n sub-intervals of $[0, 1]$. Using binary search techniques and an optimal $O(n \log(\frac{1}{\varepsilon}))$ queries, we can compute an ε -close labelling.

Query efficiency comes from the fact that if x, y have the same label, it becomes unnecessary to further query any point in $[x, y]$. To be more specific, unless $[x, y]$ contains the boundary of a sub-interval, all labels can be inferred within $[x, y]$. Boundary points of intervals thus serve as “critical points” with respect to the query oracle Q_ℓ , where the information it provides changes.

We will use a higher-dimensional analogue of this property at the core of CD-GBS. At a high level, suppose that we have an (m, n) -polytope partition that we want to learn via queries and an algorithm for computing arbitrary ε -close labellings of $(m - 1, n)$ -polytope partitions. We can use this algorithm as a subroutine on the cross-sections of two coordinates $x \neq y$ and ask whether the

¹An identical result which may be of separate interest holds if we consider partitions of arbitrary m -dimensional convex polytopes (not just Δ^m as per the definition of polytope partitions). As long as we can bound the thickness and diameter of the ambient convex polytope, learning in thickness translates to learning in distance.




Figure 4: Proof of Lemma 3

convex combination of these two ε -close labellings will itself result in a $g(\varepsilon)$ -close labelling of $\mathcal{P}^{x,y}$ (recall Definition 2) for a reasonable g .

Suppose that we could compute 0-close labellings (i.e. perfectly recover a polytope partition), it is clear that if we let \mathcal{P}_V be the set of all vertices of all P_i in the polytope partition, then $\pi(\mathcal{P}_V)$ is a suitable set of critical points (not necessarily the smallest one though) in the sense that if $[x, y] \cap \pi(\mathcal{P}_V) = \emptyset$, then the convex combination of both lower-dimensional 0-close labellings for \mathcal{P}^x and \mathcal{P}^y will result in a 0-close labelling for $\mathcal{P}^{x,y}$. Taking the contrapositive of this, if the convex combination does not result in a 0-close labelling —a condition which can be verified— then we know there is a critical point in $[x, y]$. Thus we recover a binary search mechanism, whereby we can isolate critical points up to a desired tolerance ε .

3.1 Warm-up: Learning Slices of Single Polytopes

We set up important groundwork by focusing on arbitrary polytopes $P \subset \mathbb{R}^m$. We let $\pi : \mathbb{R}^m \rightarrow \mathbb{R}$ be the projection function introduced in Definition 2, and we recall Definition 4 regarding the vertex critical coordinates of P denoted by C_P .

Lemma 5. *Suppose that $x, y \in \mathbb{R}$ are such that $[x, y] \cap C_P = \emptyset$. Then taking convex hulls of cross-sections we get $\text{Conv}(P^x, P^y) = P^{x,y}$.*

Proof. $[x, y] \cap C_P = \emptyset$ implies the vertices of the polytope $P^{x,y}$ lie in \mathcal{P}_x and \mathcal{P}_y . Since the convex hull of the set of all vertices of a bounded polytope is the polytope itself, the claim follows. \square

This property of polytopes whereby convex combinations give rise to complete information except when traversing a discrete set of critical points (visualised in Figure 5) is critical to CD-GBS. With query access to polytopes however, we no longer fully recover P_x perfectly, but instead an approximation given by an ε -close labelling, \widehat{P}_x . It becomes more subtle to show that by taking convex hulls of \widehat{P}_x and \widehat{P}_y , we recover the desired information along $[x, y]$.




Figure 5: $\text{Conv}(P^a, P^b) \neq P^{a,b}$ and $\text{Conv}(P^c, P^d) = P^{c,d}$

3.2 Necessary Machinery

We delve into the specifics of CD-GBS by defining some important machinery. We recall our notion of thickness in Definition 8, and see that it satisfies a sub-additivity property when the sets being considered are convex polytopes:

Lemma 6. *Let $P_1, \dots, P_k \subseteq \mathbb{R}^m$ be convex polytopes. Then $\tau(\cup_i P_i) < \frac{10}{3}(\sum_i \tau(P_i))(m+1)^{3/2}$.*

Proof. Let $R = \frac{10}{3}(\sum_i \tau(P_i))(m+1)^{3/2}$. Suppose that $x \in \cup_i P_i$. We will show that $B_R(x)$ cannot be a subset of $\cup_i P_i$ via a volume argument. For this proof, we will let $V(A)$ denote the volume of the set $A \subset \mathbb{R}^m$. We will also let $S(m, R)$ denote the volume of the hypersphere in m dimensions of radius R .

First of all, we need to show that for a given P_i , we have the following volume bound:

$$V(P_i \cap B_R(x)) \leq 2\tau(P_i)S(m-1, R).$$

This follows from Fritz John's Theorem [19] especially as referenced in [5]. The statement of this theorem says that if $K \subseteq \mathbb{R}^m$ is a convex body, then there exists a unique ellipsoid of maximal volume $\mathcal{E} \subseteq K$, with the property that $\mathcal{E} \subseteq P \subseteq m\mathcal{E}$. Any higher-dimensional ellipsoid has at most m axes of symmetry, and for \mathcal{E} , it must be the case that the smallest axis is at most the thickness of the convex body: $\tau(K)$ (Otherwise there would be a sphere of radius larger than $\tau(K)$ inscribed in \mathcal{E} , contradicting the definition of thickness). Furthermore, since $K \subseteq m\mathcal{E}$, the projection of K onto this axis must be contained in a segment of length at most $2m\tau(K)$. This means that if we take an arbitrary polytope P_i , there exist two parallel supporting hyperplanes to P_i , call them H_1 and H_2 that are at most $2m\tau(P_i)$ apart. Call the convex body between these halfspaces H . Since the majority of the mass of a hypersphere is contained around its centre, it follows that the volume of the intersection of H with $B_R(x)$ is maximised when x is equidistant from H_1 and H_2 . Furthermore, the volume of this cross-section is bounded by the distance between H_1 and H_2 multiplied by $S(m-1, R)$ which is at most $2\tau(P_i)S(m-1, R)$. Since $V(P_i \cap B_R(x)) \leq V(H \cap B_R(x))$, the claim holds.

Now if we take a union bound over all i , we get $V((\cup_i P_i) \cap B_R(x)) \leq 2m \sum_i \tau(P_i)S(m-1, R)$. If it were the case that the right hand side were strictly less than $S(m, R)$, we would have found an R such that $B_R(x)$ contains points not contained in any P_i . To this end, we use the following ratio:

$$\frac{S(m, R)}{S(m-1, R)} = R\sqrt{\pi} \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m+2}{2})} \geq 0.6R(m+1)^{-1/2}.$$

The inequality uses Stirling's formula for the gamma function. We can therefore see that with our value $R > \frac{10}{3}(\sum_i \tau(P_i))(m+1)^{3/2}$, we get the desired volume bound. \square

For a given polytope partition $\mathcal{P} = \{P_i\}_i$, it will be important to establish thickness bounds on P_i at specific cross-sections.

Definition 9 (α -Critical Coordinates). *Let $P \subset \mathbb{R}^m$ be a polytope. For $\alpha > 0$, we define $l_\alpha(P) = \inf\{x \in \mathbb{R} \mid \tau(P^x) \geq \alpha\}$ and $r_\alpha(P) = \sup\{x \in \mathbb{R} \mid \tau(P^x) \geq \alpha\}$ so that $\forall z \in \mathbb{R}$, $\tau(P^z) \geq \alpha$ if and only if $z \in [l_\alpha(P), r_\alpha(P)]$ (Here thickness is with respect to the natural embedding of P^x in \mathbb{R}^{m-1}). These are called α -critical coordinates for P .*

The previous definition allows us to associate to each polytope P_i a segment of $[0, 1]$ within which cross-sections of P_i are thick above a threshold. By combining this with Definition 4 we get the correct notion of critical coordinates mentioned at the beginning of Section 8.

Definition 10 (Critical Coordinates of a (m, n) -Polytope Partition). *Suppose that $\mathcal{P} = \{P_1, \dots, P_n\}$ is an (m, n) -polytope partition. For $\alpha > 0$, we let $C_{\mathcal{P}}^\alpha$ be the union of the sets of all vertex critical coordinates of all P_i as defined in Definition 4, and the set of all α -critical coordinates for all P_i as in Definition 9. Specifically, $C_{\mathcal{P}}^\alpha = (\cup_i C_{P_i}) \cup (\cup_i \{l_\alpha(P_i), r_\alpha(P_i)\})$.*

As mentioned in the beginning of this section, CD-GBS clusters queries around critical coordinates (up to a desired tolerance). For this reason it is important to bound the number of critical coordinates in a given (m, n) -Polytope partition.

Lemma 7. *If \mathcal{P} is a (m, n) -polytope partition $|C_{\mathcal{P}}^\alpha| \leq \binom{n+m}{m} + 2n$.*

Proof. For any given (m, n) -polytope partition, \mathcal{P} , if a vertex occurs, it must be the case that out of the n polytopes in \mathcal{P} and m boundary halfspaces of Δ^m , m of them meet. Furthermore, each collection of m polytopes and boundary halfspaces can give rise to only one vertex (which can be seen as a consequence of the fact that vertices are points in Δ^m). It follows that the set of all vertex critical coordinates is at most $\binom{n+m}{m}$ and the first part of the bound holds. As for the second half, there are at most two α -critical coordinates per P_i , which completes the expression above. \square

With this machinery in hand, we are in a position to prove the main result necessary to demonstrate correctness of CD-GBS. We show that if $x, y \in [0, 1]$ are such that $[x, y]$ contains no critical coordinates, then computing sufficiently fine empirical labellings of \mathcal{P}^x and \mathcal{P}^y with Q_ℓ will contain enough information to compute an ε -close labelling of $\mathcal{P}^{x,y}$ by simply taking convex combinations of the empirical labellings at both cross-sections.

Lemma 8. *Given $m, n, \varepsilon > 0$ let $\alpha = \frac{\varepsilon}{20nm^{5/2}}$ and $\beta = \frac{\varepsilon^2}{85nm^{5/2}}$. Suppose that \mathcal{P} is an (m, n) -polytope partition and that the following hold:*

- $x, y \in [0, 1]$ are such that $x < y \leq 1 - \frac{\varepsilon}{3}$.
- $[x, y] \cap C_{\mathcal{P}}^\alpha = \emptyset$.
- $\widehat{\mathcal{P}}^x$ and $\widehat{\mathcal{P}}^y$ are empirical labellings of \mathcal{P}^x and \mathcal{P}^y computed via Q_ℓ , such that $\cup_i \widehat{P}_i^x$ and $\cup_j \widehat{P}_j^y$ are β -nets for $(\Delta^m)^x$ and $(\Delta^m)^y$ respectively.

Then $\cup_i \text{Conv}(\widehat{P}_i^x, \widehat{P}_i^y)$ is an ε -net of $(\Delta^m)^{x,y}$.

Proof. Let us define the following:

$$U = \{i \in [n] \mid [l_\alpha(P_i), r_\alpha(P_i)] \cap [x, y] = \emptyset\}$$

$$V = \{i \in [n] \mid [x, y] \subsetneq [l_\alpha(P_i), r_\alpha(P_i)]\}$$

We call U the set of α -insignificant polytopes and V the set of α -significant polytopes. From the fact that $[x, y]$ contains no critical coordinates, we know that $U \cup V = [n]$ and from Lemma 1, we also know that for all $z \in [x, y]$, \mathcal{P}^z is non-degenerate. We proceed by proving the following claims:

1. $V \neq \emptyset$.
2. Any point in the cross-section of an α -insignificant polytope is $\frac{2\varepsilon}{3}$ close to an α -significant polytope (within that same cross-section).
3. If $e \in P_j^x \setminus \left(\bigcup_{i=1}^n \widehat{P}_i^x\right)$, and $j \in V$, then there exists a $e' \in \widehat{P}_j^x$ such that $\|e - e'\|_2 < \frac{\varepsilon}{3}$.
4. If $w \in P_j^z$ for some $j \in V$ and $z \in [x, y]$, then there exists a $w' \in \text{Conv}(\widehat{P}_j^x, \widehat{P}_j^y) \cap \pi^{-1}(z)$ such that $\|w - w'\|_2 < \frac{\varepsilon}{3}$.

(2) and (4) suffice to prove the theorem. To see this, suppose that $w \in (\Delta^m)^{x,y}$. This means that $w \in P_i^z$ for some $i \in [n]$ and $z \in [x, y]$. If $i \in V$, then from (4) $\exists w' \in \text{Conv}(\widehat{P}_i^x, \widehat{P}_i^y) \subset \bigcup_i \text{Conv}(\widehat{P}_i^x, \widehat{P}_i^y)$ such that $\|w - w'\|_2 < \frac{\varepsilon}{3}$. On the other hand, if $i \in U$, then by (2) $\exists w' \in P_j^z$ for some $j \in V$ such that $\|w - w'\|_2 < \frac{2\varepsilon}{3}$. In turn by (1) again, $\exists w'' \in \text{Conv}(\widehat{P}_j^x, \widehat{P}_j^y) \subset \bigcup_i \text{Conv}(\widehat{P}_i^x, \widehat{P}_i^y)$ such that $\|w' - w''\|_2 < \frac{\varepsilon}{3}$. Using the triangle inequality $\|w - w''\|_2 < \varepsilon$, and hence $\bigcup_i \text{Conv}(\widehat{P}_i^x, \widehat{P}_i^y)$ is an ε -net of $(\Delta^m)^{x,y}$ in the ℓ_2 norm as desired.

Let us prove statement (1). We know that if $i \in U$, for all $z \in [x, y]$ it holds that $\tau(P_i^z) \leq \alpha$. Using the union bound from Lemma 6 we see $\tau(\bigcup_{i \in U} P_i^z) \leq \frac{10n\alpha m^{3/2}}{3} = \frac{\varepsilon}{6m}$. On the other hand, we also know that $\bigcup_{i \in U} P_i^z \subset (\Delta^m)^z \cong (1-z)\Delta^{m-1}$. From Lemma 3, we know $\tau((\Delta^m)^z) \geq \frac{1-z}{((m-1)+\sqrt{m-1})}$. It follows that if $\frac{\varepsilon}{6m} < \frac{1-z}{((m-1)+\sqrt{m-1})}$, then $\bigcup_{i \in U} P_i^z \neq (\Delta^m)^z$. The condition $y \leq 1 - \frac{\varepsilon}{3}$ ensures that this happens for all $z \in [x, y]$. This in turn implies $V \neq \emptyset$.

Let us prove statement (2). From Lemma 3, we know $\frac{\text{Diam}((\Delta^m)^z)}{\tau((\Delta^m)^z)} \leq \sqrt{2}((m-1)+\sqrt{m-1})$. We can apply Lemma 2 in exactly the same fashion as Lemma 4 to get $\gamma_1 = \frac{2\varepsilon}{3} > \left(\frac{10n\alpha m^{3/2}}{3}\right) 4(m-1)$. We know that if $w \in \bigcup_{i \in U} P_i^z$, then $\exists w' \in \bigcup_{i \in V} P_i^z$ such that $\|w - w'\|_2 < \gamma_1 = \frac{2\varepsilon}{3}$, which is what we wanted to show.

Let us prove statement (3). Let us define $(\widehat{P}_j^x)_\perp = P_j^x \setminus \left(\bigcup_{i \in [n]} \widehat{P}_i^x\right)$. Note that these are the points in P_j^x that do not have any label whatsoever under the empirical labelling at x . Importantly, some points could have a label other than j if these points are on the boundary of another polytope with a label that has higher priority in the lexicographic ordering. By the fact that we have a β -close labelling of \mathcal{P}^x , it must hold that $\tau((\widehat{P}_j^x)_\perp) \leq \beta$. Also, since $P_j^x \subset (\Delta^m)^x$, we know $\text{Diam}(P_j^x) \leq \sqrt{2}$ from Lemma 3. Since $j \in V$, we also know that $\tau(P_j^x) \geq \alpha$, hence $\tau((\widehat{P}_j^x)_\perp) \leq \beta < \alpha \leq \tau(P_j^x)$ which in turn implies that $\text{int}(P_j^x) \setminus (\widehat{P}_j^x)_\perp \neq \emptyset$. Let $\eta^* = \frac{1}{2} \left(\frac{\varepsilon}{3} - \left(\frac{\sqrt{2}}{\alpha}\right) \beta\right) > 0$ and let $\gamma_2 = \frac{\varepsilon}{3} - \eta^* > \left(\frac{\sqrt{2}}{\alpha}\right) \beta$ (the addition of the η^* gap is to help with the proof of statement (4)). We can use the second part of Lemma 2 to see $B_{\gamma_2}(e) \cap \left(\text{int}(P_j^x) \setminus (\widehat{P}_j^x)_\perp\right) \neq \emptyset$. Since all points

in $\text{int}(P_j^x)$ only belong to P_j , it follows that under the lexicographic oracle one only sees the label j for those points. This implies that $B_{\gamma_2}(e) \cap \widehat{P}_j^x \neq \emptyset$, which in turn implies $\exists e' \in \widehat{P}_j^x$ such that $\|e - e'\|_2 < \gamma_2 = \frac{\varepsilon}{3} - \eta^* < \frac{\varepsilon}{3}$ as desired.

Finally, we prove statement (4). Since $[x, y]$ has no critical points, from Lemma 5 we know that $\text{Conv}(P_j^x, P_j^y) = P_j^{x,y}$, which in turn means that there exist $a \in P_j^x$ and $b \in P_j^y$ such that $w \in \text{Conv}(a, b)$. To be precise $w = \text{Conv}(a, b) \cap \pi^{-1}(z)$. Now if $a \in \widehat{P}_j^x$ and $b \in \widehat{P}_j^y$, then we are done. Let us suppose that this is not the case. We focus on a . If $a \in (\widehat{P}_j^x)_\perp$, the previously proved statement says there is some $a' \in \widehat{P}_j^x$ such that $\|a - a'\| < \frac{\varepsilon}{3}$. If $a \notin (\widehat{P}_j^x)_\perp \cup \widehat{P}_j^x$, then it must be the case that $a \in \widehat{P}_k^x \cap P_j^x$ for some other $k \in [n]$. This however only happens if $a \in P_j \cap P_k$ for some $P_k \neq P_j$, from the second property of polytope partitions from Definition 1 and the fact that using the lexicographic query oracle means that if $P_j = P_k$ and $j < k$ then $\widehat{P}_k = \emptyset$ always. Invoking the second property of Definition 1 again, we see that a lies on a bounding hyperplane of P_j . This in turn means that for every $\delta > 0$, $B_\delta(a) \cap \text{int}(P_j^x) \neq \emptyset$. Let us thus consider $\delta^* = \min\{\frac{\varepsilon}{3}, \frac{\eta^*}{2}\}$, where η^* is defined as in the previous paragraph. Let x^* be a point in $B_{\delta^*}(a) \cap \text{int}(P_j^x)$. Either $x^* \in \widehat{P}_j^x$ or $x^* \in (\widehat{P}_j^x)_\perp$. In the former case, since $\delta^* < \frac{\varepsilon}{3}$ we are done, we have succeeded in finding $a' = x^* \in \widehat{P}_j^x$ such that $\|a - a'\|_2 < \frac{\varepsilon}{3}$. In the latter case, from the previous paragraph, since $x^* \in (\widehat{P}_j^x)_\perp$ we know that $\exists a' \in \widehat{P}_j^x$ such that $\|x^* - a'\|_2 < \frac{\varepsilon}{3} - \eta^*$. Since $\delta^* < \frac{\eta^*}{2}$, we can use the triangle inequality to conclude that $\|a - a'\|_2 < \frac{\varepsilon}{3}$. In either case, we have proven what we wanted.

The same argumentation as the previous paragraph shows us that $\exists b' \in \widehat{P}_j^y$ such that $\|b - b'\|_2 < \frac{\varepsilon}{3}$. If we let $w' = \text{Conv}(a', b') \cap \pi^{-1}(z)$, then w' satisfies the requirements of statement (4) and we have finished our proof. \square

For the following corollary, suppose that \mathcal{P} is an (m, n) polytope partition and that $0 = t_0 < t_1, \dots, < t_k = 1$ are points in $[0, 1]$. Furthermore suppose that $\beta = \frac{\varepsilon^2}{85nm^{5/2}}$ as in Lemma 8. For each t_i , if $t_i \notin C_{\mathcal{P}}^\alpha$, let $\widehat{\mathcal{P}}^{t_i}$ be a β -close labelling of \mathcal{P}^x , otherwise let $\widehat{\mathcal{P}}^{t_i} = \emptyset$. Let $\widehat{\mathcal{P}} = \text{Conv}_i(\widehat{\mathcal{P}}^{t_i})$ and for $i = 1, \dots, k$, let $I_j = [t_{j-1}, t_j]$. If $\widehat{\mathcal{P}}^{t_{i-1}, t_i}$ is an ε -close labelling of $\mathcal{P}^{t_{i-1}, t_i}$, we say that I_j is covered, otherwise we say I_j is uncovered.

Corollary 1. *For any collection of $\{t_i\}_{i=1}^k$, there are no more than $2C_{\mathcal{P}}^\alpha$ intervals I_j that are uncovered.*

Proof. Suppose that I_j is uncovered, then one of the following holds:

- Either t_{j-1} or t_j are in $C_{\mathcal{P}}^\alpha$
- $t_{j-1}, t_j \notin C_{\mathcal{P}}^\alpha$ yet $\text{Conv}(\widehat{P}^{t_{j-1}}, \widehat{P}^{t_j})$ is not an ε -close labelling of $\mathcal{P}^{t_{j-1}, t_j}$.

From the contrapositive of Lemma 8, the latter case implies $I_j \cap C_{\mathcal{P}}^\alpha \neq \emptyset$, hence in either case there is a critical coordinate in I_j . In the worst case each $x \in C_{\mathcal{P}}^\alpha$ lies on a t_j , causing both I_j and I_{j+1} to be uncovered. This implies that there are at most $2C_{\mathcal{P}}^\alpha$ intervals I_j that are uncovered. \square

3.3 Specification of CD-GBS and Query Usage

Terms and Notation: The details of CD-GBS are presented in Algorithm 1. We recall our notation from Definition 3 where for $x \in [0, 1]$ we defined $f_x : (\Delta^m)^x \rightarrow \Delta^{m-1}$ given by $f_x(x, \dots, v_m) = \frac{1}{1-x}(v_2, \dots, v_m)$. We note that this is a bijection between both polytopes, hence it is well-defined to use f_x^{-1} . In addition, we let $\mathcal{D}^k = \{\frac{i}{2^k} \mid 1 \leq i \leq 2^k\}$ be the dyadic fractions of k -th power in the unit interval (excluding 0). For every $x \in \mathcal{D}^k$ we can associate the interval $I_x^k = [x - \frac{1}{2^k}, x]$. For each of these intervals $\text{midpoint}(I_x^k)$ denotes its midpoint. We also use the same language as Corollary 3.2 when we talk about whether I_x^k is covered or not (with respect to the current empirical labelling, $\widehat{\mathcal{P}}$, obtained from taking convex hulls of labels in Δ^m). We note that in order to have a well-defined base case of CD-GBS (which is equivalent to binary search), we let $\Delta^0 = \mathbb{R}^0 = \{0\}$. Finally, we say that a point $x \in [0, 1]$ is an uncovered critical point if $\widehat{\mathcal{P}}^x$ is computed via a recursive call to CD-GBS and for $(a, b) = B_{\varepsilon/2}(x) \cap [0, 1]$, it holds that $\widehat{\mathcal{P}}^{a,b}$ is not an ε -close labelling of $\mathcal{P}^{a,b}$.

Theorem 1. *If CD-GBS is given access to Q_ℓ for a (m, n) -polytope partition, it computes an ε -close labelling of \mathcal{P} using at most $\left(\prod_{i=1}^m \left(\binom{n+i}{i} + 2n\right)\right) 2^{2m^2} \log^m \left(\frac{170nm^{5/2}}{\varepsilon}\right)$ membership queries. For constant m this constitutes $O(n^{m^2} \log^m(\frac{n}{\varepsilon})) = \text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries².*

Proof. We first prove that CD-GBS indeed computes an ε -close labelling when given access to a valid Q_ℓ by inducting on m . It is straightforward to see that in the case $m = 1$, if CD-GBS is given access to a valid Q_ℓ for a $(1, n)$ polytope partition (a partition of the unit interval into connected subintervals), then it simply performs binary search on the interval $[0, 1] \cong \Delta^1$.

As for the inductive step, for $k = \lceil \log(2/\varepsilon) \rceil$, any two contiguous points of \mathcal{D}^k are less than $\varepsilon/2$ away from each other. For now suppose that every recursive call to CD-GBS was along a non-degenerate cross section \mathcal{P}^t . From the inductive assumption, this means that CD-GBS computes an $\varepsilon/2$ -close labellings of those cross-sections, using the triangle inequality, we know that $\widehat{\mathcal{P}}$ is an ε -close labelling of \mathcal{P} .

We note however that there is no guarantee for what a recursive call to CD-GBS does on a degenerate cross section $\widehat{\mathcal{P}}^t$. For this reason, it could be the case that at the end of the loop over \mathcal{D}^i , $\widehat{\mathcal{P}}$ is not an ε -close labelling. This can only happen if there is some $t \in C_{\mathcal{P}}^\alpha \cap \mathcal{D}^k$ which is an uncovered critical coordinate.

If t is an uncovered critical coordinate we can rectify the situation. If we find a $z \in B_{\varepsilon/2}$ that is not a critical coordinate, then \mathcal{P}^z is non-degenerate and computing CD-GBS along the cross-section gives us an $\frac{\varepsilon}{2}$ -close labelling of \mathcal{P}^z . Using the triangle inequality, we see that this in turn removes t from the set of uncovered critical coordinates, and we say that t is “fixed”. Thus the final while loop of the algorithm eliminates the set of uncovered critical coordinates so that $\widehat{\mathcal{P}}$ is indeed an ε -close labelling.

It thus remains to show that the final while loop terminates. However, there are at most $|C_{\mathcal{P}}^\alpha|$ uncovered critical coordinates, and over the course of fixing all uncovered critical coordinates, there are at most $|C_{\mathcal{P}}^\alpha|$ bad guesses for $z \in B_{\varepsilon/2}(x)$ where \mathcal{P}^z is degenerate. Therefore the final while

²CD-GBS runs in polynomial time for constant m . The time-intensive operation consists of identifying uncovered intervals, but since the dimension of the ambient simplex is constant, each empirical polytope \widehat{P}_i has at most a constant number of bounding hyperplanes. These hyperplanes can each be extruded by ε , and checking whether there exists a point outside all these extrusions can be done in time polynomial in n via brute force. In fact, all other algorithms in this paper have efficient runtimes (in their relevant parameters) due to similar reasoning.

loop makes at most $2|C_{\mathcal{P}}^\alpha|$ invocations to CD-GBS along cross-sections. This concludes the proof of correctness for CD-GBS.

Let us bound the total query usage of CD-GBS. For all values of k in the first for loop, we know from Corollary 3.2 that since Q_ℓ is a valid lexicographic oracle for \mathcal{P} , that the number of uncovered I_x^k will not exceed $2\left(\binom{n+m}{m} + 2n\right)$, and since CD-GBS is called once per uncovered interval, it follows that for each k there at most $2\left(\binom{n+m}{m} + 2n\right)$ recursive calls to CD-GBS. Furthermore, since Q_ℓ is a valid lexicographic oracle for \mathcal{P} , it will also never be the case that $\exists i, j \in [n]$, $z \in \Delta^m$ such that $\dim(\hat{P}_i) = m$ and $z \in \text{int}(\hat{P}_i)$.

In the worst case, k loops from 1 to $\lceil \log(2/\varepsilon) \rceil$ and makes an extra $2|C_{\mathcal{P}}^\alpha|$ recursive calls to CD-GBS to fix all uncovered critical coordinates. In total if we let $T(m, n, \varepsilon)$ denote the query cost of running CD-GBS on a valid lexicographic oracle, we get the following recursion:

$$T(m, n, \varepsilon) \leq 2|C_{\mathcal{P}}^\alpha| \log\left(\frac{2}{\varepsilon}\right) T\left(m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}\right) + 2|C_{\mathcal{P}}^\alpha|$$

In order to make this more amenable, we define $f(m) = \left(\binom{n+m}{m} + 2n\right)$ and use Lemma 7 to bound this expression as follows:

$$T(m, n, \varepsilon) \leq 3f(m) \log\left(\frac{2}{\varepsilon}\right) T\left(m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}\right)$$

Furthermore, from the fact that the base case is binary search, we know $T(1, n, \varepsilon) \leq n \log\left(\frac{2}{\varepsilon}\right)$.

To unpack the recursion. Let us define $\varepsilon_0 = \varepsilon$ and $\varepsilon_{k+1} = \frac{\varepsilon_k^2}{85n(m-k)^{5/2}}$ for $k = 1, \dots, m-1$. With this in hand, we can unroll the recursion to obtain:

$$T(m, n, \varepsilon) \leq \left(3^{m-1} \prod_{i=1}^{m-1} f(i)\right) \left(\prod_{k=1}^{m-1} \log\left(\frac{2}{\varepsilon_k}\right)\right)$$

Since each $\varepsilon_{k+1} < \varepsilon_k$, we can upper bound the right-hand product by bounding each term with ε_{m-1} . If we first solve for this value, we obtain:

$$\varepsilon_{m-1} = \frac{\varepsilon^{2^{m-1}}}{\prod_{j=1}^{m-1} (85nj^{5/2})^{2^j}} \geq \frac{\varepsilon^{2^{m-1}}}{\prod_{j=1}^{m-1} (85nm^{5/2})^{2^j}} \geq \left(\frac{\varepsilon}{85nm^{5/2}}\right)^{2^m}.$$

In the first inequality we bounded the denominator product in the base by $j \leq m$, as for the second inequality, we evaluated the geometric series in 2 for the exponent to bound the exponent by 2^m . With this in hand we obtain the desired bounds:

$$T(m, n, \varepsilon) \leq 3^m 2^{m^2} \prod_{i=1}^m f(i) \log^m\left(\frac{170nm^{5/2}}{\varepsilon}\right) \leq \left(\prod_{i=1}^m \left(\binom{n+i}{i} + 2n\right)\right) 2^{2m^2} \log^m\left(\frac{170nm^{5/2}}{\varepsilon}\right)$$

Finally, For large enough n , every term in $\prod_{i=1}^m \left(\binom{n+i}{i} + 2n\right)$ is bounded by $(n+m)^m + 2n$. It follows that this product is $O(n^{m^2})$, and thus for constant m , this constitutes $O(n^{m^2} \log^m\left(\frac{n}{\varepsilon}\right)) = \text{poly}(n, \log\left(\frac{1}{\varepsilon}\right))$ queries. \square

The previous results show that for constant dimension, m , CD-GBS is query efficient in n and $\frac{1}{\varepsilon}$. In the following section we use this algorithm as a building block to construct a method for computing efficient ε -close labellings when the number of regions, n , is held constant instead.

Algorithm 1 CD-GBS(m, n, ε, Q)

Input: $m \geq 0$, $n, \varepsilon > 0$, query access to function $Q : \Delta^m \rightarrow [n]$.

Output: $\widehat{\mathcal{P}}$: an ε -close labelling of \mathcal{P} .

if $m = 0$ **then**

 Query $Q(0)$

else

$\widehat{\mathcal{P}}^0 \leftarrow f_0^{-1} \left(\text{CD-GBS} \left(m - 1, n, \frac{\varepsilon^2}{85nm^{5/2}}, Q \circ f_0^{-1} \right) \right)$, $\widehat{\mathcal{P}}^1 \leftarrow Q(\vec{e}_1)$.

for $k = 1$ to $\lceil \log(2/\varepsilon) \rceil$ **do**

if Number of uncovered I_x^k exceeds $2 \left(\binom{n+m}{m} + 2n \right)$ **then**

 Halt

for $x \in \mathcal{D}^k$ **do**

if I_x^k is uncovered **then**

$t \leftarrow \text{midpoint}(I_x)$

$\widehat{\mathcal{P}}^t \leftarrow f_t^{-1} \left(\text{CD-GBS} \left(m - 1, n, \frac{\varepsilon^2}{85(1-t)nm^{5/2}}, Q \circ f_t^{-1} \right) \right)$

 Recompute $\widehat{\mathcal{P}}$ by taking convex hulls of labels

if $\exists i, j \in [n]$ such that $\text{int}(\widehat{P}_i) \cap \widehat{P}_j \neq \emptyset$ or $\widehat{\mathcal{P}}$ is an ε -close labelling **then**

 Halt

while $\exists x \in [0, 1]$ an uncovered critical point **do**

$t \leftarrow z$ for arbitrary $z \in B_{\varepsilon/2}(x)$

$\widehat{\mathcal{P}}^t \leftarrow f_t^{-1} \left(\text{CD-GBS} \left(m - 1, n, \frac{\varepsilon^2}{85(1-t)nm^{5/2}}, Q \circ f_t^{-1} \right) \right)$

 Recompute $\widehat{\mathcal{P}}$ by taking convex hulls of labels

return $\widehat{\mathcal{P}}$

4 Constant-Region Generalised Binary Search for Q_ℓ

In this section we introduce Constant-Region Generalised Binary Search, (CR-GBS), which as the name suggests, is a query-efficient algorithm for computing ε -close labellings of (m, n) -polytope partitions when n is constant and m and ε are allowed to vary.

The intuition behind the algorithm lies in the fact that if m is much greater than n (it suffices for $m > \binom{n}{2}$), then any vertex of a given P_i cannot lie in the interior of the ambient simplex Δ^m . This is because a vertex in Δ^m must consist of the intersections of at least m half-spaces, all of which cannot arise from adjacencies between different P_i .

Not only do all vertices lie on the boundary of Δ^m , but one can easily show that they are all contained in faces of the boundary of Δ^m that have dimension $O(n^2)$ which is presumed to be constant. The number of such faces in the boundary of Δ^m is thus polynomial in m , and moreover if we could compute 0-close labellings of these faces we could take convex combinations and recover a 0-close labelling of the entire polytope partition.

We will demonstrate that for an appropriate value of ε' , if we compute ε' -close labellings of such faces in the boundary, we can recover an ε -close labelling of the entire polytope partition over all of Δ^m by taking convex combinations. CR-GBS computes the necessary ε' -close labellings of lower dimensional faces by using CD-GBS as a subroutine, which as we shall see results in our desired query efficiency for n constant.

We note however, that not all faces in the boundary of Δ^m are axis-aligned, which poses a problem if we are to use CD-GBS as a subroutine. As we show in the following section, this is not an issue since we can translate such simplices into axis-aligned simplices via a simple transformation.

4.1 Non-axis-aligned Simplices

So far we have focused on the case where $\Delta^m = \{x \in \mathbb{R}^m \mid \|x\|_1 \leq 1, x_i \geq 0\}$. In a straightforward fashion we transform our results to the equivalent simplex $\Lambda^m = \{x \in \mathbb{R}^{m+1} \mid \|x\|_1 = 1, x_i \geq 0\}$. To do so, we define the invertible linear map $\phi_m : \Delta^m \rightarrow \Lambda^{m+1}$ given by $\phi_m(x_1, \dots, x_m) = ((1 - \sum_{i=1}^m x_i), x_1, \dots, x_m)$. It is straightforward to see that ϕ_m is $\sqrt{m+1}$ -Lipschitz continuous. Via standard Lipschitz continuity arguments we get the following:

Lemma 9. *Suppose that \mathcal{P} is an (m, n) -polytope partition of Λ^{m+1} . If $\widehat{\mathcal{P}}$ is an $\frac{\varepsilon}{\sqrt{m+1}}$ -close labelling of $\phi_m^{-1}(\mathcal{P})$, then $\phi_m(\widehat{\mathcal{P}})$ is an ε -close labelling of \mathcal{P} .*

Proof. Suppose that $x \in \Lambda^{m+1}$ has no label under $\phi_m(\widehat{\mathcal{P}})$. Since $\widehat{\mathcal{P}}$ is an $\frac{\varepsilon}{\sqrt{m+1}}$ -close labelling, there must be some $y \in \Delta^m$ with the property that $\|\phi_m^{-1}(x) - y\|_2 < \frac{\varepsilon}{\sqrt{m+1}}$. If we consider $\phi_m^{-1}(x)$, since $\widehat{\mathcal{P}}$ is $\frac{\varepsilon}{\sqrt{m+1}}$ -close, there must be some y from an empirical polytope $\widehat{P}_i \subset \Delta^m$ with the property that $\|\phi_m(x) - y\|_2 < \frac{\varepsilon}{\sqrt{m+1}}$. It follows that $\phi_m(y) \in \phi_m(\widehat{\mathcal{P}})$ and by Lipschitz continuity of ϕ_m , $\|x - \phi_m(y)\|_2 < \varepsilon$ as desired. \square

4.2 Necessary Machinery for CR-GBS

Suppose that \mathcal{P} is an (m, n) -polytope partition with the property that $m > \binom{n}{2}$. Furthermore, let $k = \binom{n}{2}$ and let $\partial_k(\Delta^m)$ denote all k -dimensional faces of Δ^m . For each face F , let \mathcal{P}_F be the restriction of \mathcal{P} to F . If F is axis-aligned (equivalently, if F contains the origin), then it is an isometric embedding of Δ^k in Δ^m , so we let ϕ_F be a canonical isomorphism from F to Δ^k . If F is not axis-aligned, we let ϕ_F be any canonical isomorphism from F to Δ^k as per Section 4.1.

As mentioned previously, computing empirical labellings of every face in $\partial_k(\Delta^m)$ via CD-GBS will be enough to compute an empirical labelling for \mathcal{P} . The only issue with this strategy however, is that CD-GBS is only guaranteed to return an ε -close labelling if it is given access to a valid lexicographic membership oracle for a polytope partition, and for an arbitrary polytope partition, it is not always the case that $\phi_F(\mathcal{P}_F)$ is a (k, n) -polytope partition for all $F \in \partial_k(\Delta^m)$. As an example, consider a polytope partition with an arbitrary $m-1$ -dimensional polytope P_i contained in $F = \mathcal{P}^0$ (the 0-cross-section of \mathcal{P}). Any full-dimensional $P_j \in \mathcal{P}$ must have the property that $0 \notin \pi(\text{relint}(P_j))$, hence it still holds that $\text{relint}(P_i) \cap \text{relint}(P_j) = \emptyset$. However, when restricted to \mathcal{P}_F , relative interiors are with respect to \mathcal{P}^0 , and it can be the case that $\text{relint}((P_i)_F) \cap \text{relint}((P_j)_F) \neq \emptyset$. For this reason, we slightly refine our notion of polytope partition.

Definition 11. *Suppose that \mathcal{P} is an (m, n) -polytope partition such that for all $0 \leq k \leq m$ and $F \in \partial_k(\Delta^m)$, $\phi_F(\mathcal{P}_F)$ is a (k, n) -polytope partition. Then we say that \mathcal{P} is a proper polytope partition.*

For the remainder of this section, we focus on proper polytope partitions. In addition, in order to prove correctness of CR-GBS we define a robust approximation of any $P_i \in \mathcal{P}$.




Figure 6: γ -Interiors of Polytopes in a Partition

Definition 12. Suppose that $P \subset \Delta^m$ is a polytope. We define $\text{int}_\gamma(P)$ as

$$\text{int}_\gamma(P) = \{x \in P \mid B_\gamma(x) \cap \Delta^m \subset P\}$$

We call this the γ -interior of P .

Intuitively, the γ -interior of P consists of points that are “robustly” within P by a margin of γ relative to the interior of Δ^m , as visualised in Figure 6. In Lemma 10 we show that $\text{int}_\gamma(P)$ is a sub-polytope of P with certain supporting hyperplanes translated towards the interior of P by a margin of γ . We also show that if P_i is an element of an (m, n) -polytope partition where $m > k$, then the vertices of $\text{int}_\gamma(P_i)$ also lie in some $F \in \partial_k(\Delta^m)$.

Lemma 10. Suppose that \mathcal{P} is an (m, n) -polytope partition with $m > k = \binom{n}{2}$. For each $P_i \in \mathcal{P}$, and any $\gamma > 0$, $\text{int}_\gamma(P_i)$ is a sub-polytope of P_i . Furthermore, each vertex of $\text{int}_\gamma(P_i)$ lies in some $F \in \partial_k(\Delta^m)$.

Proof. Since $P_i \subset \mathbb{R}^m$ is a polytope, it can be expressed as the intersection of finitely many half-spaces: $P_i = \bigcap_{j=1}^q H_j$, such that $H_j = \{x \in \mathbb{R}^m \mid a_j \cdot x \geq b_j\}$, where $a_j \in \mathbb{R}^m$, $\|a_j\|_2 = 1$, $b_j \in \mathbb{R}$. As mentioned before, each half-space, H_j , can either arise as an adjacency of P_i with the boundary of Δ^m , or as an adjacency of P_i with some other $P_r \in \mathcal{P}$. Let us call the former set of half-spaces A and the latter B . We abuse notation slightly and also let A refer to the sets of indices $j \in [q]$ such that $H_j \in A$ (similarly for B).

For each $H_j \in B$, let $H'_j = \{x \in \mathbb{R}^m \mid a_j \cdot x \geq b_j + \gamma\}$, where $a_j \in \mathbb{R}^m$, $\|a_j\|_2 = 1$, $b_j \in \mathbb{R}$. Clearly $H'_j \subset H_j$, and in fact the boundary hyperplane of H'_j is parallel to that of H_j (and translated by a margin of γ towards the interior of H_j). We now define $C = \left(\bigcap_{j \in A} H_j\right) \cap \left(\bigcap_{j \in B} H'_j\right)$ and we show that $\text{int}_\gamma(P_i) = C$, which proves the first part of the lemma.

Suppose that $x \in C$. By virtue of the construction of all H'_j , it must be the case that $B_\gamma(x)$ does not intersect the boundary of any $H_j \in B$. Since all $H_j \in A$ are unchanged in C , we obtain $B_\gamma(x) \cap \Delta^m \subset P_i$, therefore $x \in \text{int}_\gamma(P_i)$.

Now suppose that $x \in \text{int}_\gamma(P_i)$. Since $\text{int}_\gamma(P_i) \subset P_i$, it is clear that $x \in H_j$ for all $H_j \in A$. As for $H_j \in B$, we know that $x \in H_j$ from the fact that $\text{int}_\gamma(P_i) \subset P_i$. If $x \notin H'_j$, then $B_\gamma(x) \not\subset H_j$, which in turn implies $B_\gamma(x) \not\subset P$, contradicting our assumption that $x \in \text{int}_\gamma(P_i)$. This proves the claim that $C = \text{int}_\gamma(P_i)$.

As for the final claim of the lemma, we note that since each $H_j \in A$ arises as an adjacency of P_i with the boundary of Δ^m , it must be the case that $|A| \leq m$. Furthermore, since each $H_j \in B$ arises

as an adjacency of two polytopes in \mathcal{P} , it follows that $|B| \leq \binom{n}{2} = k < m$. Since at least m half-spaces need to meet in \mathbb{R}^m to make a vertex, it must be the case that any vertex of $C = \text{int}_\gamma(P_i)$ lies on some $F \in \partial_k(\Delta^m)$. \square

Suppose that \mathcal{P} is a proper (m, n) -polytope partition with $m > k = \binom{n}{2}$. Furthermore, suppose that $P_i \in \mathcal{P}$ is of full affine dimension and consider a vertex, v , of $\text{int}_\gamma(P_i)$ which is “robustly” in the interior of P_i by definition. From the previous lemma we know that v lies in some $F \in \partial_k(\Delta^m)$. We now show that due to the margin γ with which v lies within P_i , we can recover a label of v by computing a suitable empirical-labelling of F .

Lemma 11. *Suppose that \mathcal{P} is a proper (m, n) -polytope partition with $m > k = \binom{n}{2}$. Furthermore, suppose that $P_i \in \mathcal{P}$ is of full affine dimension and that v is a vertex of $\text{int}_\gamma(P_i)$ that lies on some face $F \in \partial_k(\Delta^m)$. It follows that any $\frac{2\gamma}{5}$ -close labelling of F that correctly labels the vertices of F gives v the label i . Furthermore, suppose that for all $F \in \partial_k(\Delta^m)$ we compute a $\frac{2\gamma}{5}$ -close labelling. By taking convex combinations of these empirical labellings, we get $\tau(\widehat{P}_\perp) \leq \frac{10}{3}n^2\gamma(m+1)^{3/2}$.*

Proof. If v is a vertex as in the statement of the lemma, it must either be a vertex of the original simplex, or $B_\gamma(v) \cap P_i \cap F$ must contain a r -dimensional ℓ_2 ball of radius γ which we call A_2 (where r corresponds to the dimension of the sub-face of F that v lies on, implying $1 \leq r \leq k$). If v is a vertex of the original simplex, then it is correctly labelled by assumption, so we focus on the the latter case.

Let A_1 be any r -dimensional ℓ_1 ball of radius $\frac{3\gamma}{5}$ such that $A_1 \subsetneq A_2$, and denote the corners of A_1 by x_1, \dots, x_s . For $i = 1, \dots, s$, let $V_i = B_{2\gamma/5}(x_i) \cap A_2 \subset F$. We note that $V_i \cap V_j = \emptyset$ for all $i \neq j$.

By the conditions of empirical labellings and the fact that P_i is of full affine dimension, there must exist z_1, \dots, z_s such that $z_r \in V_r$ and z_r gets its correct label, i , under Q_ℓ . Furthermore, it is straightforward to see that $v \in \text{Conv}(z_1, \dots, z_s)$, hence v gets its correct label, i , as visualised in Figure 7 for $r = 2$.

Along with Lemma 10, this shows that if for all $F \in \partial(\Delta^m)^k$ we compute $\frac{2\gamma}{5}$ -close labellings that correctly label the vertices of F , then we will have correctly labelled all vertices of the polytope $\text{int}_\gamma(P_i)$. Consequently, by taking convex combinations of these labellings, the entirety of $\text{int}_\gamma(P_i)$ will be labelled correctly for an arbitrary full-dimensional P_i .

For a given full-dimensional $P_i \subset \mathcal{P}$, it is the case that $P_i \setminus \text{int}_\gamma(P_i)$ can be expressed as a disjoint union of at most $k \leq n^2$ polytopes of thickness bounded by γ (using the notation from the proof of Lemma 10, these polytopes are all of the form $(H_j \setminus H'_j) \cap P_i$, of which there are at most $|B| = k$). For a given P_j that is not full-dimensional, it trivially holds that $\tau(P_j) = 0$. Thus we can use Lemma 6 to see that $\tau(\widehat{P}_\perp) = \tau(\cup_i (P_i \setminus \text{int}_\gamma(P_i))) \leq \frac{10}{3}n^2\gamma(m+1)^{3/2}$. \square

Corollary 2. *Suppose that \mathcal{P} is a proper (m, n) -polytope partition. Let $\gamma = \frac{3\varepsilon}{40n^2(m+1)^{5/2}}$, and suppose that for all $F \in \partial_k(\Delta^m)$, a $\frac{2\gamma}{5}$ -close labelling that correctly labels the vertices of F is computed with Q_ℓ . Taking a convex combination of these empirical labellings results in an ε -close labelling of \mathcal{P} .*

Proof. This follows from the fact that $\tau(\widehat{P}_\perp) \leq \frac{10}{3}n^2\gamma(m+1)^{3/2}$ from the previous theorem. We can therefore use Lemma 4 and obtain the desired result. \square

The previous result gives us precisely what we need to prove the correctness of CR-GBS. In fact, it shows that CR-GBS can use any algorithm as a sub-routine (not just CD-GBS) as long as




Figure 7: Proof of Lemma 11

it computes empirical labellings of polytope partitions along all faces $F \in \partial(\Delta^m)^k$ while correctly labelling the vertices of Δ^m .

4.3 Specification of CR-GBS and Query Usage

Terms and Notation: For $F \in \partial_k(\Delta^m)$, we let ϕ_F denote a canonical isomorphism from F to Δ^k as per Section 4.2. Furthermore, for each such F , we let $\hat{\mathcal{P}}_F$ empirical labelling returned by CD-GBS on a given face, F .

Algorithm 2 CR-GBS(m, n, ε, Q)

Input: $m, n, \varepsilon > 0$, query access to membership oracle Q for (m, n) -polytope partition \mathcal{P} .

Output: ε -close labelling of \mathcal{P} .

```

 $k \leftarrow \binom{n}{2}$ 
for  $F \in \partial_k(\Delta^m)$  do
     $\hat{\mathcal{P}}^F \leftarrow \phi_F^{-1} \left( \text{CD-GBS} \left( k, n, \frac{3\varepsilon}{100n^2\sqrt{k+1}(m+1)^{5/2}}, Q \circ \phi_F^{-1} \right) \right)$ 
 $\hat{\mathcal{P}} \leftarrow \text{Conv}_F(\hat{\mathcal{P}}^F)$ 
return  $\hat{\mathcal{P}}$ 

```

Theorem 2. Let \mathcal{P} be a proper (m, n) -polytope partition where n is constant and $m > k = \binom{n}{2}$. CR-GBS computes an ε -close labelling of \mathcal{P} and uses $O(m^k \log^k(\frac{m}{\varepsilon})) = \text{poly}(m, \log(\frac{1}{\varepsilon}))$ queries.

Proof. The correctness follows from Corollary 2. In the worst case faces are of the form Λ^k , which incur an extra cost of $\sqrt{k+1}$ in the approximation factor of empirical labellings. We use this as a worst case bound.

For simplicity in notation, we define $m_0 = k$, $\varepsilon_0 = \frac{3\varepsilon}{100n^2\sqrt{k+1}(m+1)^{5/2}}$. From Theorem 1, the CD-GBS subroutine uses at most $\left(\prod_{i=1}^{m_0} \left(\binom{n+i}{i} + 2n\right)\right) 2^{2m_0^2} \log^{m_0} \left(\frac{170nm_0^{5/2}}{\varepsilon_0}\right)$ queries. Since $k = \binom{n}{2}$ is constant, this expression can be written as $O\left(\log^k\left(\frac{m^{5/2}}{\varepsilon}\right)\right) = O\left(\log^k\left(\frac{m}{\varepsilon}\right)\right)$. Finally, there are $\binom{m}{k}$ possible faces upon which CD-GBS can be called as a subroutine, hence the total query usage is indeed $O\left(m^k \log^k\left(\frac{m}{\varepsilon}\right)\right) = \text{poly}(m, \log\left(\frac{1}{\varepsilon}\right))$. \square

5 Upper Envelope Polytope Partitions

Up until now we have focused completely on the lexicographic query oracle Q_ℓ , creating algorithms CD-GBS and CR-GBS that compute ε -close labellings of (m, n) -polytope partitions when given access to Q_ℓ . If these algorithms are given access to an adversarial oracle Q_A however, they may fail. It suffices to see this for CD-GBS since CR-GBS uses it as a subroutine.

To see why CD-GBS may fail under Q_A we recall that the algorithm recursively computes ε -close labellings of cross-sections \mathcal{P}^t for different values of $t \in [0, 1]$. If ever CD-GBS is called on a degenerate cross-section \mathcal{P}^t , it has conditions to either tell that it is being called on a degenerate cross-section (when it notices that there exist $i, j \in [n]$ and $z \in \Delta^m$ such that $z \in \text{int}(\hat{P}_i) \cap \hat{P}_j$), or in the worst case, prevent it from exceeding its query balance. In both cases however, the algorithm returns a valid empirical labelling, i.e., $\hat{P} = \{\hat{P}_i\}_{i=1}^n$ such that $\hat{P}_i \subseteq P_i$.

When an adversarial oracle is used however, we may see $i, j \in [n]$ and $z \in \Delta^m$ such that $z \in \text{int}(\hat{P}_i) \cap \hat{P}_j$. Indeed this can occur if $P_i = P_j$ and both are full-dimensional. The natural solution seems to merge P_i and P_j (since the second condition of the definition of polytope partitions tells us that $P_i = P_j$ in this case). The main problem however, is that there is no way of telling when the condition above is an artifice of the adversarial oracle, or simply due to the fact that \mathcal{P}^t is degenerate. If we blindly merge labels, we may in fact be performing an incorrect merge on a degenerate cross-section! This of course may return inconsistent polytope partitions.

Since the key problem is the existence of degenerate cross-sections, we consider a slightly stronger variant of polytope partitions with the key property that cross-sections are never degenerate. Furthermore, this special type of polytope partition is expressive enough for our game theoretic applications, and best of all, it allows us to prove results in the adversarial query oracle model.

Definition 13 (Upper Envelope Polytope Partition). *Suppose that $A \in \mathbb{R}^{n \times m}$ is an $n \times m$ real-valued matrix and that $b \in \mathbb{R}^n$. Let $P_i = y \in \Delta^m$ such that $(Ay + b)_i \geq (Ay + b)_j$ for all $j \neq i$. We denote the collection $\mathcal{P}(A, b) = P_1, \dots, P_n$, as the upper envelope polytope partition (UEPP) arising from (A, b) .*

It is straightforward to see that for any (A, b) , $\mathcal{P}(A, b)$ is itself an (m, n) -polytope partition. Crucially however, it satisfies more properties than the previous definition of polytope partitions.

Lemma 12. *Suppose that A is an $n \times m$ real valued matrix and that $b \in \mathbb{R}^n$. Then $\mathcal{P}(A, b) = \{P_1, \dots, P_n\}$ has the following properties:*

- For any $x \in [0, 1)$ let f_x be the canonical affine transformation that maps $(\Delta^m)^x$ to Δ^{m-1} . There exists an $n \times (m-1)$ real matrix A^x and $b^x \in \mathbb{R}^n$ such that $\mathcal{P}(A^x, b^x) = f_x(\mathcal{P}(A, b)^x)$.
- $\mathcal{P}(A, b)$ is a proper polytope partition (Definition 11).

- If $A_{i,\bullet} = A_{j,\bullet}$ and $b_i = b_j$ then $P_i = P_j$. Conversely if P_i is of full affine dimension and $\text{relint}(P_i) \cap P_j \neq \emptyset$, then $A_{i,\bullet} = A_{j,\bullet}$ and $b_i = b_j$; consequently, $P_i = P_j$.
- Suppose that $a_1, \dots, a_k \in \mathbb{R}$ are such that $\sum_{i=1}^k a_i < 1$ with $k < m$. Let $H = \{(z_1, \dots, z_m) \in \Delta^m \mid z_i = a_i, i = 1, \dots, k\}$ where H has affine codimension k . If $x_1, \dots, x_{m-k} \in \Delta^m$ are affinely independent points of $P_i \cap H$ and $y \in \text{Conv}(x_1, \dots, x_{m-k})$ belongs to P_j , then P_i and P_j coincide in H .

Proof. The first bullet point follows from two facts: affine transformations restricted to affine subspaces are themselves affine transformations, and compositions of affine transformations are themselves affine transformations.

To be rigorous, define the affine transformation $g : \mathbb{R}^m \rightarrow \mathbb{R}^m$ to be $g(x) = Ax + b$. Let $g' = g \upharpoonright_{(\Delta^m)^x}$ be the restriction of g to the affine subspace $(\Delta^m)^x \subset \mathbb{R}^m$ of codimension 1. As we mentioned before, g' is itself an affine transformation.

Now let us recall that f_x is the canonical affine transformation that maps $(\Delta^m)^x$ to Δ^{m-1} . It is straightforward to see that f_x^{-1} exists (Δ^{m-1} and $(\Delta^m)^x$ are clearly isomorphic) and is itself an affine transformation. Consequently $g' \circ f_x^{-1} : \Delta^{m-1} \rightarrow \mathbb{R}^n$ is itself an affine transformation, which can be identified with a matrix A^x and vector b^x such that $(g' \circ f_x^{-1})z = A^x z + b^x$. It is straightforward to see that (A^x, b^x) are such that $\mathcal{P}(A^x, b^x) = f_x(\mathcal{P}(A, b)^x)$ as desired.

As for the second bullet point, let $F \in \partial_{m-1}(\Delta^m)$ be an arbitrary face of Δ^m of codimension 1. In addition, we use the notation ϕ_F as before to denote the canonical isomorphism from F to Δ^{m-1} . ϕ_F is itself an affine transformation, hence we can use identical argumentation from before to show that by restricting the original affine functions arising from (A, b) to $F \in \partial_{m-1}(\Delta^m)$, we can find equivalent affine functions that render $\phi_F(\mathcal{P}_F)$ an upper-envelope polytope partition. For arbitrary $0 \leq k \leq m-1$, we can use the previous statement inductively to show that for any $F \in \partial_k(\Delta^m)$, \mathcal{P}_F is a (k, n) -polytope partition. This concludes the proof that \mathcal{P} is a proper polytope partition.

As for the third bullet point, the fact that $A_{i,\bullet} = A_{j,\bullet}$ and $b_i = b_j$ implies $P_i = P_j$ is trivial. Let us focus on the case when P_i is of full affine dimension and $\text{relint}(P_i) \cap P_j \neq \emptyset$. For the sake of contradiction, let us suppose that $A_{i,\bullet} = A_{j,\bullet}$ and $b_i \neq b_j$. If this holds, then $(Ay + b)_i \neq (Ay + b)_j$ for all y , which contradicts our assumption that $\text{relint}(P_i) \cap P_j \neq \emptyset$. Let us therefore suppose that $A_{i,\bullet} \neq A_{j,\bullet}$. Let H be the set of y such that $(Ay + b)_i = (Ay + b)_j$. Since $A_{i,\bullet} \neq A_{j,\bullet}$, H has codimension of at least 1. By assumption, there exists a $z \in \text{relint}(P_i) \cap P_j$. It must be the case that $z \in H$ as well. However, using the fact that $z \in \text{relint}(P_i)$ and that P_i is of full affine dimension, for some $\varepsilon > 0$, the $B_\varepsilon(z) \subsetneq P_i$. However, since $z \in H$, which is of codimension 1, then half of $B_\varepsilon(z)$ must not belong to P_i , which is a contradiction.

The final bullet point follows from putting the first and third bullet points together and inducting on k . The base case follows from the fact that for $w \in [0, 1)$, we know that \mathcal{P}^w is itself a scaled upper envelope polytope partition (from the first bullet point). Now suppose that $x_1, \dots, x_{m-1} \in P_i$ are affinely independent in $\mathcal{P}(A, b)^w$. Furthermore suppose that $\text{Conv}(x_1, \dots, x_{m-1})$ contains a point $y \in P_j$. Since the x_i are affinely independent, it follows that P_i^w is full-dimensional in $\mathcal{P}(A, b)^w$, hence we can apply the third bullet point to show that P_i^w and P_j^w coincide in $\mathcal{P}(A, b)^w$ (which is in fact what we desired).

Let us suppose that the claim holds for a given $k-1 < m-1$ and that we are given a_1, \dots, a_k . From the first bullet point, $\mathcal{P}(A, b)^{a_1}$ is a scaled lower-dimensional upper envelope polytope partition. Let us define $H = \{(z_1, \dots, z_m) \in \Delta^m \mid z_i = a_i, i = 1, \dots, k\}$ and $H_2 = \{(z_1, \dots, z_m) \in \Delta^m \mid z_i = a_i, i = 2, \dots, k\}$. It follows that $\mathcal{P}(A, b) \cap H = \mathcal{P}(A, b)^{a_1} \cap H_2$, and in the later we can use the

inductive assumption (since $(m-1) - (k-1) = m-k$) to show that if x_1, \dots, x_{m-k} are affinely independent points in $P_i^{a_1} \cap H_2 = P_i \cap H$, and $y \in \text{Conv}(x_1, \dots, x_{m-k})$ belongs to P_j , then $P_i^{a_1}$ and $P_j^{a_1}$ coincide in H_2 , which is the same as saying P_i and P_j coincide in H as desired. \square

5.1 Adversarial CD-GBS

Suppose that \mathcal{P} is an UEPP. Since it is also a proper (m, n) -polytope partition, it inherits all the properties from before. Along with Lemma 12 we have the necessary tools to show that Algorithm 3 is a query efficient way of computing ε -close labellings of \mathcal{P} with an adversarial query oracle. In the specification of CD-GBS, we use identical terms and notation from Algorithm 1.

Algorithm 3 Adversarial CD-GBS(m, n, ε, Q_A)

Input: $m \geq 0$, $n, \varepsilon > 0$, query access to oracle $Q_A : \Delta^m \rightarrow [n]$.

Require: Recursive calls to CD-GBS $\left(m-1, n, \frac{\varepsilon^2}{85(1-x)nm^{5/2}}, Q_A \circ f_x^{-1}\right)$.

Output: ε -close labelling of \mathcal{P} .

if $m = 0$ **then**

 Query $Q_A(0)$

else

$\widehat{\mathcal{P}}^0 \leftarrow f_0^{-1} \left(\text{CD-GBS} \left(m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}, Q_A \circ f_0^{-1} \right) \right)$

$\widehat{\mathcal{P}}^1 \leftarrow Q(\vec{e}_1)$.

for $k = 1$ to $\lceil \log(2/\varepsilon) \rceil$ **do**

for $x \in \mathcal{D}^k$ **do**

if I_x^k is uncovered **then**

$t \leftarrow \text{midpoint}(I_x)$

$\widehat{\mathcal{P}}^t \leftarrow f_t^{-1} \left(\text{CD-GBS} \left(m-1, n, \frac{\varepsilon^2}{85(1-t)nm^{5/2}}, Q_A \circ f_t^{-1} \right) \right)$

 Recompute $\widehat{\mathcal{P}}$ by taking convex hulls of labels

while $\exists i, j \in [n]$, $z \in \Delta^m$ such that $\dim(\widehat{P}_i) = m$ and $z \in \text{int}(\widehat{P}_i)$ **do**

 Merge label i with label j

 Recompute $\widehat{\mathcal{P}}$ by taking convex hulls of labels

if $\widehat{\mathcal{P}}$ is an ε -close labelling **then**

 Break

return $\widehat{\mathcal{P}}$

Theorem 3. *If CD-GBS is given access to an adversarial query oracle Q_A of an (m, n) -polytope partition based on a UEPP, it computes an ε -close labelling of \mathcal{P} using at most $\left(\prod_{i=1}^m \left(\binom{n+i}{i} + 2n \right) \right) 2^{2m^2} \log^m \left(\frac{170nm^{5/2}}{\varepsilon} \right)$ membership queries. For constant m this constitutes $O(n^{m^2} \log^m \left(\frac{n}{\varepsilon} \right)) = \text{poly}(n, \log \left(\frac{1}{\varepsilon} \right))$ queries.*

Proof. As in the proof of correctness of CD-GBS, we begin by noting that when $m = 1$ the algorithm runs identical to binary search. We thus focus on the case where $m > 1$.

The key observation of the proof of correctness is the following: At any given k in the first for loop there are at most $2|C_{\mathcal{P}}^a|$ values of x such that I_x^k is uncovered. Let us consider the empirical polytope $\widehat{\mathcal{P}}$ that has been constructed at the time of the execution of the k -th loop. Due to the fact that we have merged any labels from the execution of the loop at value $k-1$, it follows that

for all i, j , $\widehat{P}_i \cap \widehat{P}_j$ is not of full affine dimension. In turn this means that there exists a hyperplane $H_{i,j}$ that separates the interiors of \widehat{P}_i and \widehat{P}_j . Furthermore, denote $H_{i,j}^+$ as the halfspace defined by $H_{i,j}$ in which \widehat{P}_i is contained. This means that in turn we can define $\overline{P}_i = \cap_j H_{i,j}^+$ so that $\widehat{P}_i \subset \overline{P}_i$. Furthermore, it is straightforward to see that we can define $\overline{\mathcal{P}} = \{\overline{P}_i\}$ as a valid (m, n) -polytope partition that is consistent with $\widehat{\mathcal{P}}$. Since $\overline{\mathcal{P}}$ is consistent with our current observations from Q_A , we can actually simulate CD-GBS on \overline{P}_i for the first $k - 1$ iterations of the algorithm (ordering polytopes accordingly to simulate a lexicographic query oracle). The empirical polytope returned when doing so will in fact be $\widehat{\mathcal{P}}$, and thus we can apply corollary 3.2 to tell us that the number of uncovered I_x^k is in fact bounded by $2|C_{\overline{\mathcal{P}}}^\alpha|$.

Returning to our proof of correctness of the algorithm. It is not hard to see that upon termination it is correct if we assume that calling CD-GBS as a subroutine works correctly as per the inductive assumption and crucially the fact that from Lemma 12 each cross-section of \mathcal{P} is non-degenerate. Therefore we focus on the query cost of the algorithm. At each $k = 1$ to $\lceil \log(2/\varepsilon) \rceil$, from our previous result there can only be at most $2|C_{\mathcal{P}}^\alpha|$ uncovered I_x^k , which are precisely the I_x^k that result in queries. By simple multiplication we thus get that the number of cross-section queries is at most $2\lceil \log(2/\varepsilon) \rceil |C_{\mathcal{P}}^\alpha|$, hence we get the following recursion for bounding the query cost of adversarial CD-GBS:

$$T(m, n, \varepsilon) \leq 2 \left(\binom{n+m}{m} + 2n \right) \log \left(\frac{2}{\varepsilon} \right) T \left(m-1, n, \frac{\varepsilon^2}{85nm^{5/2}} \right)$$

with base case $T(1, n, \varepsilon) \leq n \log \left(\frac{2}{\varepsilon} \right)$. If we unpack the recursion in the same way as Theorem 1, we get the desired result. \square

5.2 Adversarial CR-GBS

In this section we formalize an adversarial variant of CR-GBS. We note that most of the notation is identical to lexicographic CR-GBS.

Algorithm 4 CR-GBS($m, n, \varepsilon, \mathcal{P}$)

Input: $m, n, \varepsilon > 0$, query access to Q_A for (m, n) -polytope partition \mathcal{P} .

Output: ε -close labelling of \mathcal{P} .

$k \leftarrow \binom{n}{2}$

for $F \in \partial(\Delta^m)^k$ **do**

$\widehat{\mathcal{P}}^F \leftarrow \phi_F^{-1} \left(\text{CD-GBS} \left(k, n, \frac{3\varepsilon}{100n^2\sqrt{k+1}(m+1)^{5/2}}, Q \circ \phi_F^{-1} \right) \right)$.

$\widehat{\mathcal{P}} \leftarrow \text{Conv}_F(\widehat{\mathcal{P}}^F)$

while $\exists i, j \in [n]$, $z \in \Delta^m$ such that $\dim(\widehat{P}_i) = m$ and $z \in \text{int}(\widehat{P}_i)$ **do**

Merge label i with label j

Recompute convex hulls of labels

return \widehat{Q}

Theorem 4. *Let \mathcal{P} be an (m, n) -polytope partition where n is constant. Furthermore, let $k = \binom{n}{2}$. CR-GBS computes an ε -close labelling of \mathcal{P} and uses $O \left(m^k \log^k \left(\frac{m}{\varepsilon} \right) \right) = \text{poly} \left(m, \log \left(\frac{1}{\varepsilon} \right) \right)$ queries.*

Proof. As in the proof of Theorem 3, we know that there exists a polytope partition $\overline{\mathcal{P}}$ that is consistent with $\widehat{\mathcal{P}}$. Once again, we notice that this invocation of adversarial CR-GBS is identical

to running lexicographic CR-GBS on $\overline{\mathcal{P}}$, which would in turn return $\widehat{\mathcal{P}}$ an ε -close labelling of $\overline{\mathcal{P}}$. However, it is straightforward to see from the definition of ε -close labellings that this also makes $\widehat{\mathcal{P}}$ an ε -close labelling of \mathcal{P} as desired. Finally, the query usage of adversarial CR-GBS is identical to lexicographic CR-GBS, hence the rest of the theorem follows. \square

6 Games and Best Responses

Now that we have established query-efficient algorithms for learning ε -close labellings of polytope partitions, we turn our attention to game theory to prove the connection between learning these labellings and computing approximate well-supported Nash equilibria.

Suppose that $G = (A, B)$ is an $m \times n$ bi-matrix game where $A, B \in [0, 1]^{m \times n}$ are the row player and column player payoff matrices respectively with payoffs normalised to $[0, 1]$. We wish to identify an ε -well-supported Nash equilibrium (ε -WSNE) using only limited information on G . The set of row player pure strategies is $[m] = \{1, \dots, m\}$ and similarly that of the column player pure strategies is $[n] = \{1, \dots, n\}$. Furthermore, the set of all row player mixed strategies can be associated with the axis-aligned $(m - 1)$ -simplex: $\Delta^{m-1} = \{\vec{x} \in \mathbb{R}^{m-1} \mid \sum_{i=1}^{m-1} x_i \leq 1 \text{ and } x_i \geq 0\}$. Similarly, column player mixed strategies are identified with Δ^{n-1} .

Definition 14 (Utility Functions). *Suppose that $u \in \Delta^{m-1}$, and $v \in \Delta^{n-1}$ are row and column player mixed strategies. Let $u' = (1 - \sum u_i, u_1, \dots, u_{n-1})$ and $v' = (1 - \sum v_i, v_1, \dots, v_{n-1})$. Then for strategy profile (u, v) , row player utility is $U_r(u, v) = u'^T A v'$ and column player utility is $U_c(u, v) = u'^T B v'$.*

It will also be useful to have shorthand for the following functions: $U_r^i(y) = U_r(e_i, y)$ as the row player utility for playing pure strategy i , and $E_R(y) = \max_{i \in [m]} U_r^i(y)$ as the maximal utility the row player can achieve against mixed strategies. In an identical fashion we can define U_c^j and E_C as the column player utility in playing strategy j and the maximal column player utility. With this notation in hand, we can define the best response oracles algorithms will have access to when computing approximate Nash equilibria.

Definition 15 (Best Response Query Oracles). *Any bimatrix game has the following best response query oracles:*

- *Strong query oracles: for the column player, $BR_s^C(u) = \{j \in [n] \mid U_c^j(u) = E_C(u)\}$ and for the row player, $BR_s^R(v) = \{i \in [m] \mid U_r^i(v) = E_R(v)\}$*
- *Lexicographic query oracles: for the column player, $BR_\ell^C(u) = \min_{j \in [n]} j \in BR_s^C(u)$ and for the row player, $BR_\ell^R(v) = \min_{i \in [m]} i \in BR_s^R(v)$*
- *Adversarial query oracles: for the column player, any function BR_A^C such that $BR_A^C(u) \in BR_s^C$ and for the row player, any function such that $BR_A^R(v) \in BR_s^R(v)$*

For a given mixed strategy, $u \in \Delta^{m-1}$, we say the support of u is the set of pure strategies that are played in u with non-zero probability. It will be useful to formulate this as a function in order to define Nash equilibria combinatorially in the following section.

Definition 16 (Support Functions). *Let $S^R : \Delta^{m-1} \rightarrow \mathcal{P}([m])$ be the function which returns the support of a row player mixed strategy. Similarly let $S^C : \Delta^n \rightarrow (\mathcal{P}[n])$ return the support of column player mixed strategies.*

7 Nash Equilibria and Lipschitz Continuity of Utility Functions

Definition 17 (Nash Equilibrium). *Suppose that u and v are row and column player strategies respectively. We say that the pair (u, v) is a Nash Equilibrium (NE) if for all $u' \in \Delta^{m-1}$ and $v' \in \Delta^{n-1}$: $U_r(u, v) \geq U_r(u', v)$ and $U_c(u, v) \geq U_c(u, v')$.*

Though the definition of a Nash equilibrium involves utility values of both players at their mixed strategy profiles, there is an equivalent combinatorial formulation of the above definition:

Proposition 5. *(u, v) is a NE if and only if $S^R(u) \subseteq BR_s^R(v)$ and $S^C(v) \subseteq BR_s^C(u)$. In other words u is supported by best responses to v and vice versa.*

When using best response queries only, one does not have access to utility values (as emphasised in the first definition), however this second equivalent definition of Nash equilibria can be verified by using best response oracles and support functions alone.³

We also note that for utility queries, the complexity of an exact NE is finite: we can exhaustively query the game. On the other hand, Corollary 3 shows that this is not the case for best response queries. As a consequence, we relax the notion of a NE when using best response queries. the relaxation of NE which we study is that of approximate well-supported equilibria. Before proceeding with the formal definition, we say that a row player mixed strategy $u \in \Delta^m$ is an ε best response against a column player mixed strategy $v \in \Delta^n$ if $U_r(u, v) \geq U_r(u', v) - \varepsilon$ for all $u' \in \Delta^m$. An identical notion holds for when a column player mixed strategy $v \in \Delta^n$ is an ε best response against a row player mixed strategy $u \in \Delta^m$. Intuitively, an ε best response is a mixed strategy where a player has only an ε incentive to deviate.

Definition 18 (ε -Well-Supported Nash Equilibrium). *Suppose that u and v are row and column player strategies respectively. We say that the pair (u, v) is an ε -well-supported Nash equilibrium (ε -WSNE) if and only if u is supported by ε -best responses to v and vice versa.⁴*

Theorem 6. *The query complexity of computing an ε -WSNE with best response queries is $\Omega(\log(\frac{1}{\varepsilon}))$, even when given access to strong query oracles.*

Proof. Suppose that $x, y \in (0, 1)$ are arbitrary, and let us consider a two-player binary action game, $G_{x,y}$, with the following row and column player payoff matrices:

$$A_x = \begin{pmatrix} x & x \\ 0 & 1 \end{pmatrix} \quad B_y = \begin{pmatrix} 0 & y \\ 1 & y \end{pmatrix}$$

Since $G_{x,y}$ is a binary action game, we can express any mixed strategy profile of both players by a tuple $(p_r, p_c) \in [0, 1]^2$. p_r represents the probability the row player plays the second row and p_c represents the probability that the column player plays the second column.

It is clear that this game has no pure equilibria, but its unique NE is the mixture: $(y, x) \in [0, 1]^2$. Upon close inspection, one can also see that the set of ε -WSNE of $G_{x,y}$ lie in $B_\varepsilon(y) \times B_\varepsilon(x) \cap [0, 1]^2$, where $B_\varepsilon(z)$ the set of points at a distance ε from z in \mathbb{R} . Suppose that an algorithm, \mathcal{A} , is given a game $G_{x,y}$ from the family above and access to the game's strong best response query oracles. In order for \mathcal{A} to compute an ε -WSNE, the previous observation tells us that the it has to at least

³In general one is unable to recover utility values from best responses, even up to affine transformations.

⁴Note that the conditions for an ε -WSNE imply that no player has more than an ε incentive to deviate from the approximate equilibrium.

find a point $z \in B_\varepsilon(x)$ by querying the row player's best response oracle. From the structure of A_x however, we know that for $p_c \in [0, x]$ the first row is a best response for the row player, and for $p_c \in [x, 1]$, the second row is a best response for the row player. This problem formulation however is identical to binary search, hence finding a $z \in B_\varepsilon(x)$ takes at least $\Omega(\log(\frac{1}{\varepsilon}))$ queries. \square

Corollary 3. *The query complexity of computing a NE with best response queries is infinite, even when given access to strong query oracles.*

7.1 Algebraic Properties of Utility Functions

Definition 18 mentions approximate best responses, yet we only have access to the best response oracle in our model. In order to resolve this, we delve into the algebraic properties of utility functions of both the column and row player.

Lemma 13. *If the domains of U_r^i and U_c^j are endowed with the ℓ_2 norm, then the functions are λ_R and λ_C Lipschitz continuous respectively, for some $0 \leq \lambda_R \leq \sqrt{n-1}$ and $0 \leq \lambda_C \leq \sqrt{m-1}$. If the domains are endowed with the ℓ_1 norm, then both functions are 1-Lipschitz continuous.*

Proof. We focus on U_r^i , the case for the column player is identical. Let $c = [A^T]_i = (a_1, \dots, a_n)$ be the i -th row vector of the row player's payoff matrix, and suppose that $v = (v_1, \dots, v_{n-1}) \in \Delta^{n-1}$ is a column player mixed strategy, where $v_n = 1 - \sum_{i=1}^{n-1} v_i$ is implicit. Let $z = (z_i)_{i=1}^{n-1}$ with $z_i = (a_i - a_0)$ for $i = 1, \dots, n-1$. Then it is clear that $U_r^i(v) = a_0 + \sum_{i=1}^{n-1} z_i \cdot v_i$. This function is linear, and trivially $\|z\|_2$ -Lipschitz continuous. Since the game is normalised, $\|z\|_2 \leq \|\vec{1}\|_2 = \sqrt{n-1}$.

As for the second part of the claim, the domain of U_r^i can be equivalently represented as $\Lambda^n = \{x \in \mathbb{R}^n \mid \|x\|_1 = 1, x_i \geq 0\}$ by using the invertible linear map $\phi_{n-1} : \Delta^{n-1} \rightarrow \Lambda^n$ given by $\phi_{n-1}(x_1, \dots, x_{n-1}) = (x_1, \dots, x_{n-1}, (1 - \sum_{i=1}^{n-1} x_i))$. This space can be endowed with total variation distance as a metric, which for two distributions, $x, y \in \Lambda^n$ is defined as $TV(x, y) = \max_{S \subseteq [n]} |\mathbb{P}_x(S) - \mathbb{P}_y(S)| = \frac{1}{2} \|x - y\|_1$. Since utilities are bounded to be in the interval $[0, 1]$, it follows that U_r^i is 1-Lipschitz as a function with domain Λ^n in the total variation metric.

Now suppose that $x, y \in \Delta^{n-1}$. We wish to show that $TV(\phi_{n-1}(x), \phi_{n-1}(y)) \leq \|x - y\|_1$. To see this, let $x_n = 1 - \sum_{i=1}^{n-1} x_i$ and $y_n = \sum_{i=1}^{n-1} y_i$. Then $\|\phi_{n-1}(x) - \phi_{n-1}(y)\|_1 = \|x - y\|_1 + |x_n - y_n|$. From this we see that $|x_n - y_n| = |\sum_{i=1}^{n-1} (y_i - x_i)| \leq \|x - y\|_1$, which in turn implies $\|\phi_{n-1}(x) - \phi_{n-1}(y)\|_1 \leq 2\|x - y\|_1$. Dividing the expression by 2 and applying the fact that $TV(x, y) = \frac{1}{2} \|x - y\|_1$ proves our desired inequality. 1-Lipschitz continuity in the ℓ_1 norm for domain Δ^{n-1} follows immediately. \square

Corollary 4. *Since E_C and E_R are defined as a pointwise maximum, it follows that they are also Lipschitz continuous with constant $\lambda_C \leq \sqrt{m-1}$ and $\lambda_R \leq \sqrt{n-1}$ in the ℓ_2 norm.*

With bounded Lipschitz continuity, we have guarantees on how much utilities can deviate between ‘‘close’’ mixed strategy profiles. This has interesting implications even for the best response query oracle, for this means that if u and u' are close in the ℓ_2 norm with $c_i \in BR_s^C(u)$, $c_j \in BR_s^C(u')$ and $c_i \neq c_j$, then we can say that c_i and c_j are both approximate best responses in the vicinity of u and u' . We formalise this as follows.

Lemma 14. *Fix $\varepsilon > 0$ and let $\delta_C = \frac{\varepsilon}{2\sqrt{m-1}}$. Suppose that $u \in \Delta^{m-1}$ is a row player mixed strategy with $c_j \in BR_s^C(u)$. For any u' such that $\|u - u'\|_2 \leq \delta_C$, if $c_i \in BR_s^C(u')$, then $|U_c^i(u) - U_c^j(u)| \leq \varepsilon$.*

In other words, c_i is an ε -best response to u . Similarly, let $\delta_R = \frac{\varepsilon}{2\sqrt{n-1}}$. Suppose that $v \in \Delta^{n-1}$ is a column player mixed strategy with $r_j \in BR_s^R(u)$. For any v' such that $\|v - v'\|_2 \leq \delta_R$, if $r_i \in BR_s^R(v')$, then $|U_r^i(v) - U_r^i(v')| \leq \varepsilon$. In other words, r_i is an ε -best response to v .

Proof. Suppose that u' is such that $\|u - u'\| \leq \delta_C$ and $c_i \in BR_s^C(u')$. By definition, $E_c(u') = U_c^i(u') \geq U_c^j(u')$ and by Lemma 13, $|U_j^c(u) - U_j^c(u')| \leq \lambda_C \|u - u'\|_2$, and $\|U_i^c(u) - U_i^c(u')\| \leq \lambda_C \|u - u'\|_2$. With these expressions we obtain the following inequalities:

$$|E_c(u) - U_c^i(u)| \leq 2\lambda_C \|u - u'\|_2 \leq 2\lambda_C \frac{\varepsilon}{2\lambda_C} = \varepsilon$$

The proof of the second half of the lemma is identical. \square

The previous Lemma establishes the important idea that we can obtain some information regarding approximate best responses using only the best response oracle and “nearby” queries. With some thought one can see that this in general does not reveal *all* approximate best response information. For example, if a strategy were strictly dominated, a best response oracle would never see it, and hence never be able to tell if it was an approximate best response.

8 Nash’s Theorem with Discrete Approximations

We are now in a position to prove the intimate connection between computing ε -close labellings of upper envelope polytope partitions and computing ε -WSNE for bimatrix games using best response queries.

Definition 19 (Best Response Sets). *Let $G = (A, B)$ be a bimatrix game. We define column best response sets as the collection of $C_i = \{x \in \Delta^{m-1} \mid BR_s^C(x) = c_i\}$. Similarly we define row player best response sets as the collection of $R_j = \{y \in \Delta^{n-1} \mid BR_s^R(y) = r_j\}$. We denote the collections by $\mathcal{C} = \{C_i\}_{i=1}^n$ and $\mathcal{R} = \{R_j\}_{j=1}^m$.*

Since utilities are affine functions, it is immediately clear that \mathcal{C} and \mathcal{R} are upper envelope polytope partitions. Now the best response oracles play the same role as membership oracles, Q , from before. Since adversarial oracles are the weakest of the three membership oracles (in the sense that they are a valid lexicographic oracle and they can be simulated with access to a strong oracle), we focus on using adversarial best response oracles. Furthermore, with our language of empirical labellings we can now define a key object used in the computation of approximate equilibria. Before doing so, we clarify some notation: $d(x, S)$ denotes the infimum distance of a point, x to a set S .

Definition 20 (Voronoi Best Response Sets). *Suppose that $\widehat{\mathcal{C}} = \{\widehat{C}_i\}$ and $\widehat{\mathcal{R}} = \{\widehat{R}_j\}$ are empirical labellings of \mathcal{C} and \mathcal{R} as in Definition 6. The Voronoi Best Response Sets of the row and column player are $VR_j = \{y \in \Delta^{n-1} \mid \operatorname{argmin}_j d(y, \widehat{R}_j) = r_j\}$ and $VC_i = \{x \in \Delta^{m-1} \mid \operatorname{argmin}_i d(x, \widehat{C}_i) = c_i\}$, defined for any $j \in [m]$ and $i \in [n]$. Furthermore, we let $V^R(v) = \{i \mid VR_i \ni v\}$ and $V^C(u) = \{j \mid VC_j \ni u\}$ be the row and column player Voronoi Best Responses.*

Lemma 15. *Voronoi best response sets partition Δ^{m-1} and Δ^{n-1} into closed connected regions with non-empty interior and piecewise linear boundaries.*

Proof. Without loss of generality, let us focus on a given column player Voronoi best response set; i.e. some $VC_i \subset \Delta^{m-1}$ that is non-empty and arises from the empirical labelling $\widehat{\mathcal{C}} = \{\widehat{C}_i\}$ of column-player best responses. First we note that $\widehat{C}_i \subset VC_i$ by definition, and the former is a convex, closed, and connected polytope of Δ^{m-1} . Therefore it remains to show that if we pick an arbitrary $x \in VC_i \setminus \widehat{C}_i$ it is connected to C_i . To do so, suppose that $p \subset \Delta^{m-1}$ is the unique shortest path from x to C_i . It is clear that all points along p must also lie in VC_i , therefore the claim holds.

As for closedness, note that if $\{x_n\}$ is a sequence in VC_i that converges to some $x \in \Delta^{m-1}$ then x must also be in VC_i . This follows from the continuity of Euclidian distance for $\Delta^{m-1} \subset \mathbb{R}^m$. Now suppose that x is a limit point of VC_i , then we can construct a sequence as above, and thus x must also be in VC_i , rendering the set closed.

Finally, the piecewise linear boundary arises from the fact that \widehat{C}_i is itself a closed convex polytope which has a piecewise linear boundary. Decision boundaries between different \widehat{C}_i and \widehat{C}_j are composed of nearest neighbour decision boundaries between piecewise linear boundaries, which in turn results in piecewise linear decision boundaries between VC_i and VC_j . \square

Although the previous lemma proves that Voronoi best response sets partition Δ^{m-1} and Δ^{n-1} into closed connected regions with non-empty interior and piecewise linear boundaries, they need not be convex. This ends up not being an issue for our subsequent results. The reason we deal with these objects however is due to the following Lemma. We recall that $\lambda_R \leq \sqrt{m-1}$ and $\lambda_C \leq \sqrt{n-1}$ are the relevant Lipschitz continuity constants for row player and column player expected utility functions. The following is a straightforward consequence of Lemma 14.

Lemma 16. *Suppose that $\widehat{\mathcal{C}}$ is a $\frac{\varepsilon}{2\sqrt{m-1}}$ -close labelling and $\widehat{\mathcal{R}}$ is a $\frac{\varepsilon}{2\sqrt{n-1}}$ -close labelling. Then Voronoi best responses are ε best-responses in G*

We recall that the combinatorial formulation of Nash's theorem implies that with full information of best response sets in all of Δ^m and Δ^n , one is able to compute and verify an exact Nash equilibrium. Best responses only partially recover this information in convex patches of Δ^m and Δ^n . Furthermore, it is not clear how a game G' with best responses sets consistent with empirical best response sets of G can be used to compute an approximate equilibrium of G . Voronoi best response sets however allow us to take the partial information provided by empirical best response sets and extend it to approximate best response information *across the entire domains* Δ^m and Δ^n (Voronoi best response sets cover Δ^m and Δ^n after all). This hints at the fact that Voronoi best response sets hold enough information to compute ε -WSNE. In fact we can prove this in the same way as Nash's theorem: via Kakutani's fixed point theorem. In order to do so, we define a Voronoi best response correspondence (which as we have shown before is an approximate best response correspondence), and show that it satisfies the properties of Kakutani's fixed point theorem. The guaranteed fixed point of this correspondence will in turn be an ε -WSNE.

Definition 21 (Voronoi Approximate Best Response Correspondence). *For a given mixed strategy profile of both the row and column player, $(u, v) \in \Delta^{m-1} \times \Delta^{n-1}$, we define $B^*(u, v)$ to be the set of all possible mixtures over Voronoi best response profiles both players may have to the other player's strategy. $B^* : \Delta^{m-1} \times \Delta^{n-1} \rightarrow \mathcal{P}(\Delta^{m-1} \times \Delta^{n-1})$ is defined as follows:*

$$B^*(u, v) = (\text{conv}(V^R(v)), \text{conv}(V^C(u))) \subseteq \Delta^{m-1} \times \Delta^{n-1}.$$

Theorem 7 (Kakutani's Fixed Point Theorem [20]). *Let A be a non-empty subset of a finite-dimensional Euclidian space and $f : A \rightarrow \mathcal{P}(A)$ be a set-valued function satisfying the following conditions:*

- A is a compact and convex set.
- $f(x)$ is non-empty for all $x \in A$.
- $f(x)$ is a convex-valued correspondence: for all $x \in A$, $f(x)$ is a convex set.
- $f(x)$ has a closed graph: that is, if $\{x_n, y_n\} \rightarrow \{x, y\}$ with $y_n \in f(x_n)$ for all n , then $y \in f(x)$.

Then f has a fixed point, that is there exists some $x \in A$ such that $x \in f(x)$.

Theorem 8. B^* satisfies all the conditions of Kakutani's fixed point Theorem, and hence there exists a strategy profile (u^*, v^*) such that $(u^*, v^*) \in B^*(u^*, v^*)$. In particular, if the Voronoi best responses for B^* arise from $\widehat{\mathcal{C}}$, a $\frac{\varepsilon}{2\sqrt{m-1}}$ -close labelling and $\widehat{\mathcal{R}}$, a $\frac{\varepsilon}{2\sqrt{n-1}}$ -close labelling, then this in turn implies that (u^*, v^*) is an ε -WSNE of G .

Proof. We need to prove the following conditions for Kakutani's fixed point Theorem:

- B^* has a compact and convex domain.
- $B^*(u, v)$ is non-empty and convex for all $(u, v) \in \Delta^{m-1} \times \Delta^{n-1}$.
- (Graph Closedness) Suppose that $\{\sigma_n\}$ and $\{\sigma'_n\}$ are sequences in $\Delta^{m-1} \times \Delta^{n-1}$ that converge to σ and σ' respectively. Furthermore, suppose that $\sigma'_n \in B^*(\sigma_n)$ for all n . Then $\sigma' \in B^*(\sigma)$.

For the first item, the domain of B^* is $\Delta^{m-1} \times \Delta^{n-1}$ which clearly satisfies the desired condition.

As for the second and third item, from the definition of B^* the image of any (u, v) consists of convex combinations of Voronoi best responses, which are defined for all (u, v) (thus satisfying non-emptiness), and since they are convex combinations, they are convex subsets of $\Delta^{m-1} \times \Delta^{n-1}$.

Finally for the fourth item, let us consider such a sequence where $\sigma_n = (u_n, v_n)$, $\sigma = (u, v)$, $\sigma'_n = (u'_n, v'_n)$, and $\sigma' = (u', v')$. To show the claim, it suffices to consider the sequences $\{u_n\}$ and $\{v'_n\}$ with respective limits u and v' , and show that $v' \in \text{conv}(V^C(u))$

To show this however, it suffices to use the fact that Voronoi best response sets are closed. Suppose that u has a certain set $S \subset [n]$ of Voronoi best responses. Then there exists a constant $\mu > 0$ such that $B_\mu(u) \cap VC_i \neq \emptyset$ if and only if $i \in S$; namely the μ neighbourhood around u only intersects Voronoi best response sets from u 's Voronoi best responses.

To explicitly construct such a μ , let us consider $D_i = d(u, \widehat{C}_i)$ to be the distance between u and the empirical best response set \widehat{C}_i . This means that $S = \text{argmin}_i D_i$, so let us define $D = \min_i D_i$ and $\mu = \frac{\min_{j \notin S} D_j - D}{3}$ which is positive due to the fact that there are finitely many partial best response sets. Now suppose that $x \in B_\mu(u)$, then for any $j \notin S$ we have $d(u, \widehat{C}_j) \leq d(u, x) + d(x, \widehat{C}_j)$ by the triangle inequality, which rearranging gives us: $d(x, \widehat{C}_j) \geq d(u, \widehat{C}_j) - d(u, x) \geq (D + 3\mu) - \mu = D + 2\mu$. On the other hand, for any $i \in S$, $d(x, \widehat{C}_i) \leq d(x, u) + d(u, \widehat{C}_i) \leq D + \mu$. It thus follows that x can only have Voronoi best responses from S .

Now from the fact that $u_n \rightarrow u$, then for some $N > 0$, if $n > N$ then $u_n \in B_\mu(u)$. This in turn means that $v_n \in \text{conv}(S)$ by assumption, which means that $v \in \text{conv}(S)$ as well, which is what we

wanted to show. To extend this to σ and σ' , it suffices to repeat the previous argument in each component of the correspondence.

Now that the conditions of Kakutani's fixed point Theorem are satisfied, we know of the existence of an (u^*, v^*) such that $(u^*, v^*) \in B^*(u^*, v^*)$. As in the statement of the Theorem, suppose that Voronoi best responses for B^* arise from an $\frac{\varepsilon}{2\sqrt{m-1}}$ -close labelling of Δ^m and an $\frac{\varepsilon}{2\sqrt{n-1}}$ of Δ^n , then we know that all Voronoi best responses are ε best responses for both players. The conditions of the fixed point amount to saying that both players are playing convex combinations of Voronoi best responses, therefore (u^*, v^*) is an ε -WSNE. \square

With Theorem 8 in hand and our algorithms for constructing ε -close labellings, we can put everything together and prove our desired results regarding the query complexity of computing an ε -WSNE in general bimatrix games.

Theorem 9. *Suppose that G is an $m \times n$ bimatrix game and let n be constant. We can compute an ε -WSNE using $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon})) = \text{poly}(m, \log(\frac{1}{\varepsilon}))$ adversarial best response queries.*

Proof. Suppose that \mathcal{C} and \mathcal{R} are the polytope partitions arising from best-response sets in G . This means that \mathcal{C} is a $(m-1, n)$ -polytope partition and \mathcal{R} is a $(n-1, m)$ -polytope partition. Let $\varepsilon_C = \frac{\varepsilon}{2\sqrt{m-1}}$ and $\varepsilon_R = \frac{\varepsilon}{2\sqrt{n-1}}$. From Theorem 8, we know that computing an ε_C -close labelling of \mathcal{C} and a ε_R -close labelling of \mathcal{R} suffice to compute an ε -WSNE of G . We use adversarial CR-GBS on \mathcal{C} and adversarial CD-GBS on \mathcal{R} .

n is the number of polytopes in the partition \mathcal{C} , which is assumed to be constant. Consequently, Theorem 4 states that computing an ε_C -close labelling of \mathcal{C} using CR-GBS uses $O((m-1)^k \log^k(\frac{m-1}{\varepsilon_C}))$ adversarial queries, where $k = \binom{n}{2}$. Since $k \leq n^2$, we can upper bound the number of queries by $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon}))$.

$n-1$ is the dimension of the ambient simplex in the partition \mathcal{R} , which is assumed to be finite. Consequently, Theorem 3 states that computing an ε_R -close labelling of \mathcal{R} using CD-GBS uses $O(m^{(n-1)^2} \log^{(n-1)^2}(\frac{1}{\varepsilon}))$ queries. We trivially upper bound this quantity by $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon}))$.

Putting everything together, the total query usage is thus $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon})) = \text{poly}(m, \log(\frac{1}{\varepsilon}))$ as desired. \square

9 A Brief Foray into Multiplayer Games

In this section we partially extend our results from sections 7 and 8. In particular, we show that in multiplayer games utility functions are Lipschitz continuous as well, which allows us to uncover approximate best-response information when observing different best responses at “nearby” mixed strategy profiles. In addition we generalise our definitions of best response sets and of ε -close labellings to obtain a result similar to Theorem 8 where we showed that obtaining a precise enough empirical labelling provides enough information to compute a well-supported approximate Nash equilibrium.

The main difference in the multiplayer setting however is that best response sets are no longer polytopes nor convex, which means that our algorithms for computing empirical labellings via generalisations of binary search no longer apply. This does not preclude us however from simply querying an ε -net, which as we will show will suffice for computing an ε -WSNE.

9.1 Notation for Multiplayer Games

For simplicity we will focus on games with n players where each player has a strategy set A_i consisting of $|A_i| = k$ pure strategies. It is straightforward to extend our results to more general games where different players have action sets of different cardinalities.

In general, we let $A = \prod_{i=1}^n A_i$ be the space of all pure strategy profiles of all players. For the i -th player, we also denote $A_{-i} = \prod_{j \neq i} A_j$ as the space of all pure strategy profiles of players other than i . Since every player has k actions, it is straightforward to see that all A_{-i} are isomorphic, hence without loss of generality we can assume that for all i , A_{-i} is canonical representation. For a given pure strategy profile $a \in A$, we may wish to distinguish the pure strategy taken by the i -th player, and this is done by writing $a = (a_i, a_{-i})$ with $a_i \in A_i$ and $a_{-i} \in A_{-i}$.

We denote the i -th player's mixed strategy space by $\Delta(A)_i$, and we note it is equivalent to Δ^{k-1} . This means that the space of mixed strategy profiles of all players is $\Delta(A) = \prod_{i=1}^n (\Delta^{k-1}) = (\Delta^{k-1})^n$. In addition, for the i -th player, we also denote $\Delta(A)_{-i} = \prod_{i=1}^n (\Delta^{k-1}) = (\Delta^{k-1})^{n-1}$ as the space of all mixed strategy profiles of players other than i . Once again, since all players have k actions, we can assume without loss of generality that all $\Delta(A)_{-i}$ have the same canonical representation. For a mixed strategy profile $x \in \Delta(A)$, we may wish to distinguish the mixed strategy of the i -th player by writing $x = (x_i, x_{-i})$ with $x_i \in \Delta(A)_i$ and $x_{-i} \in \Delta(A)_{-i}$.

Definition 22 (Multiplayer Utility Functions). *For any player i and action $r \in A$, we denote $U_i^r : A_{-i} \rightarrow [0, 1]$ as the i -th player's utility for playing r . If $a = (a_i, a_{-i}) \in A$ is a pure strategy profile, the utility player i receives is $U_i^{a_i}(a_{-i})$ which we denote by $U_i(a)$.*

Utility functions are defined for pure strategy profiles, but with a slight abuse of notation we extend the domain to include mixed strategy profiles. In particular, if $x = (x_i, x_{-i})$ is a mixed strategy profile, we let $U_i^r(x_{-i}) = \mathbb{E}_{a_{-i} \sim x_{-i}}(U_i^r(a_{-i}))$ and by extension $U_i(x) = \mathbb{E}_{a \sim x}(U_i(a))$.

As in bimatrix games, for a given player i and $x_{-i} \in \Delta(A)_{-i}$, we let $E_i(x_{-i}) = \max_{r \in A_i} U_i^r(x_{-i})$ be the maximal expected utility player i can obtain against the mixed strategy profile x_{-i} of all other players.

Definition 23 (Best Response Query Oracles). *Let G be an n -player game where each player has k pure strategies:*

- *There are n strong best response oracles denoted $BR_s^i : \Delta(A)_{-i} \rightarrow \mathcal{P}(A_i)$ for $i = 1, \dots, n$. Each of these is defined by $BR_s^i(x) = \{r \in A_i \mid U_i^r(x) = E_i(x)\}$.*
- *There are n lexicographic best response oracles denoted $BR_\ell^i : \Delta(A)_{-i} \rightarrow A_i$ for $i = 1, \dots, n$. Each of these is defined by $BR_\ell^i(x) = \operatorname{argmin}_{r \in A_i} r \in BR_s^i(x)$ for some consistent order on each A_i .*
- *An adversarial best response oracle collection is a collection of n functions denoted $BR_A^i : \Delta(A)_{-i} \rightarrow A_i$ for $i = 1, \dots, n$. Each of these satisfies $BR_A^i(x) \in BR_s^i(x)$.*

For completeness we have defined all best response oracles, but we focus on adversarial best response oracles. As mentioned before, they are the weakest from the fact that a lexicographic oracle is a valid adversarial oracle and the fact that adversarial oracles can be simulated with strong best response oracles. This implies that our results for adversarial oracles carry over to other oracle models.

For a given mixed strategy, $x \in \Delta(A)_i$, we say the support of x is the set of pure strategies that are played in x with non-zero probability. It will be useful to formulate this as a function in order to define Nash equilibria combinatorially again.

Definition 24 (Support Functions). *Let $S^i : \Delta(A)_i \rightarrow \mathcal{P}(A_i)$ be the function which returns the support of the i -th player's mixed strategy.*

We are now in a position to define what a Nash equilibrium is in multiplayer games.

Definition 25 (Nash Equilibrium). *We say $x \in \Delta(A)$ is a Nash Equilibrium (NE) if for any player i , and $x'_i \in \Delta(A)_i$ it holds that $U_i(x) \geq U_i(x'_i, x_{-i})$*

As before, this definition involves utility values of players at their mixed strategy profiles. Once again, there is an equivalent combinatorial formulation.

Proposition 10. *$x \in \Delta(A)$ is a NE if and only if for all i , when $x = (x_i, x_{-i})$, $S^i(x_i) \subseteq BR_s^i(x_{-i})$.*

Finally, we define what it means for a mixed strategy profile to be an ε -WSNE in the multiplayer setting. Before proceeding, we say that for a given $x_{-i} \in \Delta(A)_{-i}$, strategy $x \in \Delta(A)_i$ is an ε best response if $U_i(x, x_{-i}) \geq U_i(x', x_{-i}) - \varepsilon$ for all $x' \in \Delta(A)_i$. Intuitively, an ε best response is a mixed strategy where a player has only an ε incentive to deviate.

Definition 26 (ε -Well-Supported Nash Equilibrium). *Suppose that $x \in \Delta(A)$ is a mixed strategy profile of all players. We say that x is an ε -well-supported Nash equilibrium (ε -WSNE) if and only if for every player i , all pure strategies in $S^i(x_i)$ are ε best responses to x_{-i} .*

9.2 Lipschitz Continuity of Utility Functions

As in Section 7, we will show that for each player i and each $r \in A_i$, U_i^r is a Lipschitz continuous function. In order to do so, we must regard the domain of U_i^r , which is $\Delta(A)_{-i} = (\Delta^{k-1})^{n-1}$, as a subset of Euclidean space endowed with the ℓ_1 norm.

Lemma 17. *For any player i and action $r \in A_i$, U_i^r is 1-Lipschitz when the domain is endowed with the ℓ_1 norm.*

Proof. Consider an arbitrary player j with $j \neq i$. If we endow $\Delta(A)_j$ with the ℓ_1 norm, from Lemma 13, we know that U_i^r as a function of $x_j \in \Delta(A)_j$ (which is a component of $\Delta(A)_{-i}$) is 1-Lipschitz. This is because if all mixed strategies other than those of player i and j are fixed, we obtain a $k \times k$ bimatrix game between player i and j . Now let us consider $x, y \in \Delta(A)_{-i}$.

$$\begin{aligned}
|U_i^r(x) - U_i^r(y)| &= \left| \sum_{j=1}^{n-1} U_i^r(y_1, \dots, y_j, x_{j+1}, \dots, x_n) - U_i^r(y_1, \dots, y_{j+1}, x_{j+2}, \dots, x_n) \right| \\
&\leq \sum_{j=1}^{n-1} |U_i^r(y_1, \dots, y_j, x_{j+1}, \dots, x_n) - U_i^r(y_1, \dots, y_{j+1}, x_{j+2}, \dots, x_n)| \\
&\leq \sum_{j=1}^n \|x_i - y_i\|_1 \\
&= \|x - y\|_1
\end{aligned} \tag{1}$$

□

Since Lipschitz continuity is maintained over maxima, with the same Lipschitz constant, we get the following result that says that best responses to nearby mixed strategy profiles are in fact approximate best responses.

Lemma 18. $E_i : \Delta(A)_{-i} \rightarrow [0, 1]$ is 1-Lipschitz when its domain is endowed with the ℓ_1 norm. In particular, if $x, y \in \Delta(A)_{-i}$ and $\|x - y\|_1 \leq \frac{\varepsilon}{2}$, then any $r \in BR_s^i(y)$ is an ε best response to x for player i .

9.3 Nash's Theorem with Discrete Approximations in Multiplayer Games

Just as in bimatrix games, we have a notion of best response sets.

Definition 27 (Multiplayer Best Response Sets). Let G be a game with n players, each with k strategies. For a player i and pure strategy $r \in A_i$, we define $P_i^r = \{x \in \Delta(A)_{-i} \mid BR_s^i(x) = r\}$. We say that P_i^r is a best response set corresponding to strategy r for player i , and note that $\{P_i^r\}_{r \in A_i}$ cover $\Delta(A)_{-i}$.

In bimatrix games our goal was to learn ε -close labellings of the polytope partitions induced by best response sets. In the multiplayer setting that goal can be generalised.

Definition 28 (Multiplayer ε -close labellings). Let G be a game with n players, each with k actions, and let i be a specific player in the game. Suppose that for each action $r \in A_i$, $\hat{P}_i^r \subseteq P_i^r$ is a closed set. We say the collection $\{\hat{P}_i^r\}_{r \in A_i}$ is an empirical labelling of $\Delta(A)_{-i}$. If in addition $\bigcup_{r \in A_i} \hat{P}_i^r$ is an ε -net of $\Delta(A)_{-i}$ in the ℓ_1 norm, we say that the collection $\{\hat{P}_i^r\}_{r \in A_i}$ is an ε -close labelling of $\Delta(A)_{-i}$.

As we will see shortly, if we manage to compute an $\frac{\varepsilon}{2}$ -close labelling for all $\Delta(A)_{-i}$, we have enough information to compute an approximate equilibrium.

In bimatrix games, each P_i^r is a polytope, but in multiplayer games, expected utilities are no longer linear in $\Delta(A)_{-i}$. Consequently, best response sets are semi-algebraic sets instead. This means that in general best response set are not connected and thus not convex. Without convexity and polytope structure we can no longer use binary search methods to learn ε -close labellings. As mentioned before, this does not preclude us from computing an ε -close labelling via a brute force method of querying an entire ε -net of $\Delta(A)_{-i}$. Before we show this suffices however, we show the key result of this section: computing $\frac{\varepsilon}{2}$ -close labellings for all $\Delta(A)_{-i}$ suffices to compute ε -WSNE. To do so we revisit Voronoi best response sets. In what follows we let $d(x, S)$ denote the infimum distance of a point, x to a set S in the ℓ_1 norm.

Definition 29 (Multiplayer Voronoi Best Response Functions and Best Response Sets). Let G be a game with n players, each with k actions, and let i be a specific player in the game. Suppose that $\{\hat{P}_i^r\}_{r \in A_i}$ is an empirical labelling of $\Delta(A)_{-i}$. Player i 's Voronoi Best Response function is denoted by $V^i : \Delta(A)_{-i} \rightarrow A_i$. The function is defined as $V^i(x) = \operatorname{argmin}_{r \in A_i} d(x, \hat{P}_i^r)$. We also define player i 's Voronoi Best Response Sets as $V_i^r = \{x \in \Delta(A)_{-i} \mid V^i(x) = r\}$.

If we invoke Lemma 18, we obtain the same result as in bimatrix games whereby Voronoi best responses are actually approximate best responses.

Lemma 19. Suppose that $\{\hat{P}_i^r\}_{r \in A_i}$ is an $\frac{\varepsilon}{2}$ -close labelling of $\Delta(A)_{-i}$, then Voronoi Best Response for player i are ε Best Responses in G .

In addition, our definition of empirical labellings stipulated that \widehat{P}_i^r are all closed sets. This is a property which is inherited by Voronoi Best Response Sets.

Lemma 20. *Voronoi best response sets are closed.*

Proof. The proof is identical to Lemma 15 since ℓ_1 distance is still a continuous function of the relevant domain $\Delta(A)_{-i}$. \square

We now define the generalisation to the Voronoi Best Response Correspondence from before.

Definition 30 (Multiplayer Voronoi Best Response Correspondence). *Suppose that $x \in \Delta(A)$ is a mixed strategy profile. We define the Voronoi Best Response Correspondence $B^* : \Delta(A) \rightarrow \mathcal{P}(\Delta(A))$ as follows:*

$$B^*(x) = \prod_{i=1}^n \text{conv}(V^i(x_{-i})) \subseteq \Delta(A)$$

Theorem 11. *B^* satisfies all the conditions of Kakutani's fixed point Theorem, and hence there exists a mixed strategy profile $x^* \in \Delta(A)$ such that $x^* \in B^*(x^*)$. In particular, if the Voronoi best response for B^* arise from $\frac{\varepsilon}{2}$ -close labellings for all $\Delta(A)_{-i}$, then this in turn implies that x^* is an ε -WSNE.*

Proof. As in the proof of Theorem 8, the first three conditions of Kakutani's fixed point theorem are trivial. We focus on proving that B^* has a closed graph. It suffices to show the following for any player i : if $\{x_n\}$ is a sequence in $\Delta(A)_{-i}$ that converges to x , and $\{y_n\}$ is a sequence in $\Delta(A)_i$ converging to y with the property that $y_n \in \text{conv}(V^i(x_n))$ for all n , then $y \in \text{conv}(V^i(x))$.

To show this, we prove that there exists a constant $\delta > 0$ with the property that $B_\delta^1(x) \cap V_i^r$ if and only if $r \in V_i(x)$. Here $B_\delta^1(x)$ denotes the ℓ_1 ball of radius δ around x . We prove this claim by contradiction.

Suppose instead that for every $\delta > 0$, $B_\delta^1(x)$ contains some point from a V_i^r , where $r \notin V^i(x)$. Let $\delta_1 > 0$ be arbitrary, and let z_1 be the guaranteed point in $B_{\delta_1}^1(x)$ that is contained in a collection of V_i^r where none belong to $V^i(x)$. Let $\delta_2 = \|x - z_1\|_1$. We continue in this fashion where for a given $\delta_k > 0$, we let $z_k \in B_{\delta_k}^1(x)$ be a point contained in a collection of V_i^r where none belong to $V^i(x)$. Accordingly we define $\delta_{k+1} = \|x - z_k\|_1$. Since we can always continue this process, we recover a sequence $\{z_n\}$ of elements in $\Delta(A)_{-i}$ that converges to x . There are only a finite number of Voronoi Best Response sets, hence this sequence must contain an infinite subsequence of points belonging to the same voronoi best response set, say $V_i^{r'}$, where r' does not belong to $V^i(x)$. This however implies that x is a limit point of $V_i^{r'}$, and since this set is closed, this implies that $x \in V_i^{r'}$, which contradicts the fact that $r' \notin V^i(x)$, thus proving our desired claim.

Returning to the sequences $\{x_n\}$ and $\{y_n\}$, the existence of a fixed $\delta > 0$ with the property that $B_\delta^1(x) \cap V_i^r$ if and only if $r \in V_i(x)$ means that for some $N > 0$, if $n > N$, $x_n \in B_\delta^1(x)$, which in turn means that $y_n \in \text{Conv}(V^i(x))$. This in turn means that $y \in \text{Conv}(V^i(x))$, as desired. Applying this result to each i yields the graph-closedness of B^* , thus establishing that B^* satisfies all the properties of Kakutani's fixed point theorem.

As for the second part of the theorem, suppose that x^* is the guaranteed fixed point of B^* as guaranteed by Kakutani's fixed point theorem. From Lemma 19, we know that Voronoi Best Responses are ε -Best responses. By the definition of B^* , the support of each $x_i^* \in \Delta(A)_i$ consists of ε best responses to x^* , thus establishing the fact that x^* is an ε -WSNE. \square

With the previous theorem in hand, we have established that computing a $\frac{\varepsilon}{2}$ -close labellings of all $\Delta(A)_{-i}$ suffices to compute an ε -WSNE. Although the lack of convexity in best response sets prevents us from using binary search techniques, we can always query an ε -net of $\Delta(A)_{-i}$ in the ℓ_1 norm. In this vein, we construct an explicit ε -net for $\Delta(A)_{-i}$.

Lemma 21. $M_\varepsilon^n = \left(\frac{2\varepsilon}{n}\mathbb{Z}\right)^n \cap \Delta^n$ is an ε -net in the ℓ_1 norm for Δ^n . Furthermore, $|M_\varepsilon^n| = O\left(\left(\frac{n}{2\varepsilon} + n\right)^n\right)$.

Proof. The first claim follows from noting that lattice points lie on vertices of axis-aligned hypercubes of side-length $\frac{2\varepsilon}{n}$. These cubes have a diagonal of length 2ε in the ℓ_1 norm, hence their centres are at most ε -away from a given queried vertex.

As for the cardinality, using a stars and bars argument, one can see that if $S = \frac{1}{\kappa}\mathbb{Z}^n \cap \Delta^n$ for some $\kappa \in \mathbb{N}$, then $|S| = \binom{\kappa+n}{n} = O((\kappa+n)^n)$. \square

Since $\Delta(A)_{-i}$ is a product of simplices, we can take products of the above ε -net constructions to in turn obtain an ε -net for $\Delta(A)_{-i}$.

Lemma 22. Suppose that $\varepsilon > 0$ and let $\varepsilon' = \frac{\varepsilon}{n-1}$. Furthermore, let us define the set $H_\varepsilon^{n,k} = (M_{\varepsilon'}^{k-1})^{n-1} \subset \Delta(A)_{-i} \cong (\Delta^{k-1})^{n-1}$. Then $H_\varepsilon^{n,k}$ is an ε net for $\Delta(A)_{-i}$ in the ℓ_1 norm. Furthermore $|H_\varepsilon^{n,k}| = O\left(\left(\frac{nk}{2\varepsilon}\right)^{nk}\right)$.

Querying all points in an ε -net trivially gives rise to an ε -close labelling. Consequently, with Lemma 22 and Theorem 11 in hand, we have proven our main result regarding the query complexity of computing an ε -WSNE using adversarial Best Response Queries.

Theorem 12. Suppose that G is a game with n players with k pure strategies each. One can compute an ε -WSNE of G using $O\left(n\left(\frac{nk}{\varepsilon}\right)^{nk}\right)$ adversarial Best Response Queries.

10 Conclusion and Future Directions

In this paper we introduced the concept of learning ε -close labellings of (m, n) -polytope partitions with membership queries, and derived query efficient algorithms for when either the dimension of the ambient simplex in the polytope partition, m , is held constant, or when the number of polytopes in the partition, n , is held constant.

Most importantly, we introduced a novel reduction from computing ε -WSNE with best response queries to this geometric problem, thus allowing us to show that in the best response query model, computing ε -WSNE of a bimatrix game has a finite query complexity. More specifically, for $m \times n$ games with $\min(m, n)$ constant, the query complexity is polynomial in $\max(m, n)$ and $\log\left(\frac{1}{\varepsilon}\right)$. Furthermore, we partially extended our results from bimatrix games to multi-player games. Although the underlying geometry in multi-player games prevents us from using our results from learning polytope partitions, we were still able to show that querying a fine enough ε -net of the mixed strategy space of all players suffices to compute an ε -WSNE.

As mentioned in the introduction, this geometric framework could be of use in other areas where Lipschitz continuous structures appear over domains with convex partitions. Upon further inspection, it is not difficult to see that polytope partitions do not need to be contained in Δ^m , and in fact our algorithms extend to arbitrary ambient polytopes. Furthermore, it would be of great interest to create algorithms with a better query cost, prove lower bounds with regards to

computing ε -close labellings, or simply explore weaker query paradigms, such as noisy membership oracles. Finally, we have mentioned that in the multiplayer setting, best response sets are no longer polytopes, but rather semi-algebraic sets. It would be of interest to create learning algorithms for ε -close labellings of these more complicated geometric objects, since doing so suffices to compute ε -WSNE.

References

- [1] Yakov Babichenko. Query complexity of approximate Nash equilibria. *J. ACM*, 63(4):36:1–36:24, 2016.
- [2] Yakov Babichenko and Aviad Rubinfeld. Communication complexity of approximate Nash equilibria. In *Proceedings of the 49th STOC*, pages 878–889. ACM, 2017.
- [3] Yakov Babichenko, Siddharth Barman, and Ron Peretz. Empirical distribution of equilibrium play and its testing application. *Math. Oper. Res.*, 42(1):15–29, 2017.
- [4] Elizabeth Baldwin and Paul Klemperer. Understanding Preferences: “Demand Types”, and the Existence of Equilibrium with Indivisibilities. Technical report, LSE, October 2016.
- [5] Keith Ball. An elementary introduction to modern convex geometry. In Silvio Levi, editor, *Flavors of Geometry*, chapter 1, pages 1–58. Cambridge University Press, Cambridge, 1997.
- [6] George W. Brown. Some notes on computation of game solutions. *RAND corporation report*, page 78, April 1949.
- [7] Nader H. Bshouty, Paul W. Goldberg, Sally A. Goldman, and H. David Mathias. Exact learning of discretized geometric concepts. *SIAM J. Comput.*, 28(2):674–699, 1998.
- [8] X. Chen, X. Deng, and S-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14:1–14:57, 2009.
- [9] C. Daskalakis and Q. Pan. A counter-example to Karlin’s strong conjecture for fictitious play. In *Procs. of 55th FOCS*, pages 11–20, 2014.
- [10] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- [11] J. Fearnley, M. Gairing, P.W. Goldberg, and R. Savani. Learning equilibria of games via payoff queries. *J. Mach. Learn. Res.*, 16:1305–1344, 2015.
- [12] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [13] Paul W. Goldberg and Stephen Kwek. The precision of query points as a resource for learning convex polytopes with membership queries. In *Procs. of the 13th COLT*, pages 225–235, 2000.
- [14] Paul W. Goldberg and Aaron Roth. Bounds for the query complexity of approximate equilibria. *ACM Trans. Economics and Comput.*, 4(4):24:1–24:25, 2016.

- [15] Paul W. Goldberg, Rahul Savani, Troels Bjerre Sørensen, and Carmine Ventre. On the approximation performance of fictitious play in finite games. *Int. J. Game Theory*, 42(4):1059–1083, 2013.
- [16] P.W. Goldberg and S. Turchetta. Query complexity of approximate equilibria in anonymous games. *Journal of Computer and System Sciences*, 90:80–98, 2017.
- [17] S. Hart and Y. Mansour. How long to equilibrium? the communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1):107–126, 2010.
- [18] Sergiu Hart and Noam Nisan. The query complexity of correlated equilibria. *Games and Economic Behavior*, 2016. ISSN 0899-8256.
- [19] Fritz John. *Extremum Problems with Inequalities as Subsidiary Conditions*, pages 197–215. Springer Basel, Basel, 2014. ISBN 978-3-0348-0439-4. doi: 10.1007/978-3-0348-0439-4_9. URL https://doi.org/10.1007/978-3-0348-0439-4_9.
- [20] Shizuo Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Math. J.*, 8(3):457–459, 09 1941.
- [21] Paul Klemperer. The product-mix auction: a new auction design for differentiated goods. *Journal of the European Economic Association*, 8(2/3):526–536, 2010.
- [22] R.J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Procs. of the 4th ACM-EC, EC ’03*, pages 36–41, New York, NY, USA, 2003. ACM.
- [23] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [24] Julia Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54(2):296–301, 1951.
- [25] B. Von Stengel. Leadership with commitment to mixed strategies. *CDAM Research Report*, 2004.