# Local (coarse) correlated equilibria in non-concave games

Mete Seref Ahunbay<sup>\*1</sup>

<sup>1</sup>Technische Universität München, School of Computation, Information and Technology, Department of Computer Science. Boltzmannstraße 3, 85748, Garching bei München, Germany.

March 28, 2024

#### Abstract

We investigate local notions of correlated equilibria, distributions of actions for smooth games such that players do not incur any regret against modifications of their strategies along a set of continuous vector fields. Our analysis shows that such equilibria are intrinsically linked to the projected gradient dynamics of the game. We identify the equivalent of coarse equilibria in this setting when no regret is incurred against any gradient field of a differentiable function. As a result, such equilibria are approximable when all players employ online (projected) gradient ascent with equal step-sizes as learning algorithms, and when their compact and convex action sets either (1) possess a smooth boundary, or (2) are polyhedra over which linear optimisation is "trivial". As a consequence, primal-dual proofs of performance guarantees for local coarse equilibria take the form of a generalised Lyapunov function for the gradient dynamics of the game. Adapting the regret matching framework to our setting, we also show that general local correlated equilibria are approximable when the set of vector fields is finite, given access to a fixed-point oracle for linear or conical combinations. For the class of affine-linear vector fields, which subsumes correlated equilibria of normal form games as a special case, such a fixedpoint turns out to be the solution of a convex quadratic minimisation problem. Our results are independent of concavity assumptions on players' utilities.

# 1 Introduction

The central question we study is on notions of correlated equilibria in a smooth game; given a set of players N, compact & convex action sets  $X_i$  and sufficiently smooth payoffs  $u_i : \times_{i \in N} X_i \to \mathbb{R}$ , a correlated equilibrium is a probability distribution  $\sigma$  on  $\times_{i \in N} X_i$  satisfying  $\phi(\sigma) \leq 0 \forall \phi \in \Phi(u)$ for some family of linear equilibrium constraints  $\Phi$ . Our motivation is two-fold; (1) we seek an answer to a question posed in [28] regarding what an appropriate theory of non-concave games should be, and (2) we aim to refine the notions of (coarse) correlated equilibria of normal form games to strengthen the analysis of their learnable outcomes via appeals to linear programming duality. As a consequence, we demand our concept of equilibrium to be tractably computable given access to payoff gradients independent of concavity assumptions on the payoffs, and contain (coarse) correlated equilibria of normal form games as special cases.

<sup>\*</sup>E-mail: mete.ahunbay@cit.tum.de

Towards this end, we build upon the recent work of [16], who considered a *local* variant of *swap regret* traditionally considered in literature (cf. [39]), and showed two such families of local strategy modifications are tractably computable via online gradient ascent. Our key insight is that the  $\delta$ -strategy modifications contained within are generated by gradient fields of a suitable family of functions. We thus take a differential perspective, and identify two notions of "first-order" correlated equilibria with respect to a family of vector fields F. For local correlated equilibria players do not have any incentive to modify their strategies along any vector field  $f \in F$ , while for stationary correlated equilibria we demand the equilibrium satisfy the properties of a fixed point of the gradient dynamics of the game *in expectation*. The main difficulty in analysis, and the necessity of defining two distinct classes of equilibria, arises from dealing with projections onto the feasible action set.

**Contributions.** We establish settings in which local or stationary equilibria are approximable. We first observe that the approximation results of [16] for online gradient ascent are in fact for families of *gradient fields*, and that such equilibria contain coarse correlated equilibria of normal form games as a special case. We thus identify *coarseness* of our equilibrium notion so.

A high level argument shows that curves in the action space generated via the continuous time projected gradient dynamics of the game universally approximates a local or stationary coarse correlated equilibrium; the regret against  $\nabla h$  for any differentiable function h vanishes linearly in time, and is proportional to a bound on  $\|\nabla h\|$ . The question is then whether discrete time projected gradient ascent possesses similar approximation properties.

We answer in the affirmative, and show that the approximability of local correlated equilibria follows from that of the stationary one. Our methods are based on extending the outcome of online gradient ascent to a piecewise curve, and thus our results depend on the geometry of the action sets  $X_i$ ; (1) whenever each  $X_i$  has a smooth boundary of bounded principal curvature K, and (2) whenever  $X_i$  is a "acute" polyhedron, taking step sizes  $1/\sqrt{\bar{\tau}}$  results in a  $O(\log(\bar{\tau})/\bar{\tau})$  regret bound for each h after  $\bar{\tau}$  iterations. For (1), the approximation bound decreases in K, while the class of polyhedra considered in (2) contain the simplex and the hypercube as special cases.

Next, we inspect the primal-dual framework arising from stationary coarse correlated equilibria. Here, we observe that for any function q on  $\times_{i \in N} X_i$  (say, welfare), a dual lower bound  $\gamma$  on the expectation of q in equilibrium is provided by a function h such that, under the continuous gradient dynamics of the game, h is strictly decreasing whenever  $q(x) < \gamma$ . We denote such a function h as a generalised Lyapunov function; when  $\max_x q(x) = \gamma$  and q has a unique maximiser, the function h is then a Lyapunov function in the traditional sense. Our analysis thus sheds light on the form of primal-dual performance bounds (as considered in the price of anarchy literature) when players all employ online gradient ascent. We also show convergence bounds to performance guarantees for fixed  $\overline{\tau}$ .

We then drop the coarseness assumption. In this case, for finite |F|, we show that the usual regret matching framework (cf. [71, 39, 40]) may be applied to obtain  $O(1/\sqrt{T})$ -stationary correlated equilibria after T iterations, given access to a fixed point oracle for every linear combination of vector fields in F. An identical result holds for local correlated equilibria, if each  $f \in F$  is everywhere tangent to  $\times_{i \in N} X_i$ . We show that when each f is also affine-linear, the associated fixed point computation reduces to convex quadratic minimisation. This setting contains correlated equilibria of normal form games as a special case, but by example we show that further refinements to correlated equilibria are possible.

### 1.1 Related work

Game theoretic analysis of multi-agent systems is often based on first the assertion of a concept of equilibrium as an expected outcome, followed by an investigation of its properties. The classical assumption is that with self-interested agents, the outcome of the game should be a Nash equilibrium, which exists for finite normal-form games [57], or more generally, concave games [63].

Diverging from classical equilibrium theory for convex markets, the outcome of a game need not be socially optimal. The algorithmic game theory perspective has then been to interpret the class of equilibria in question from the lens of approximation algorithms. First proposed by Koutsoupias & Papadimitriou [47], the *price of anarchy* measures the worst-case ratio between the outcome of a Nash equilibrium and the socially optimal one. A related concept measuring the performance of the best-case equilibrium outcome, the *price of stability* was first employed by [68] and named so in [6]. Since then, a wealth of results have followed, proving constant welfare approximation bounds for e.g. selfish-routing [64], facility location [73], a variety of auction settings including single-item [46] and simultaneous [24] first-price auctions, simultaneous second-price auctions [23], hinting at good performance of the associated mechanisms in practice. Meanwhile, deteriorating price of anarchy bounds, such as those for sequential first-price auctions [48], are interpreted as arguments against the use of such mechanisms.

However, the assertion that the outcome of a game should be a Nash equilibrium is problematic for several reasons, despite the guarantee that it exists in concave games. It is often the case that multiple equilibria exists in a game, in which case we are faced with an *equilibrium selection problem* [42]. Moreover, Nash equilibrium computation is computationally hard in general, even in the setting of concave games where its existence is guaranteed. In general determining whether a pure strategy NE exists in a normal-form game is NP-complete [26], and even for two player normal form games finding an exact or approximate mixed Nash equilibrium is PPAD-complete [19, 29, 27]. Similar results extend to auction settings; for instance, (1) finding a Bayes-Nash equilibrium in a simultaneous second-price auction is PP-hard [17], even when buyers' valuations are restricted to be either additive or unit-demand, and (2) with subjective priors computing an approximate equilibrium of a single-item first-price auction is PPAD-hard [36], a result that holds also with common priors when tie-breaking rule is also part of the input [21].

Some answers to this problem come from learning theory. First, for special classes of games, when each agent employs a specific learning algorithm in repeated instances of the game, the outcome converges in *average play* to Nash equilibria. This is true for fictitious play [15] for the empirical distribution of players' actions in a variety of classes of games, including zero-sum games [62],  $2 \times n$  non-degenerate normal form games [10], or potential games [53]. For the case of zero-sum games, the same is true for a more general class of no-(internal [18] or external [74]) regret learning algorithms, while for general normal-form games they respectively converge to correlated or coarse correlated equilibria – convex generalisations of the set of Nash equilibria of a normal form game. The price of anarchy / stability approach can then be extended to such coarse notions of equilibria, with the smoothness framework of [65] for *robust* price of anarchy bounds being a prime example.

Unfortunately, this perspective falls short of a complete picture for several reasons. First, learning dynamics can exhibit arbitrarily high complexity. Commonly considered learning dynamics may cycle about equilibria, as is the case for fictitious play [70] or for multiplicative weights update [7]. Worse, learning dynamics can exhibit formally chaotic behaviour [59, 22], bimatrix games may approximate arbitrary dynamical systems [4]. In fact, replicator dynamics on matrix games is Turing complete [5], and reachability problems for the dynamics is in general undecideable.

In the converse direction, the usual notion of no-regret learning can be too weak to capture *learnable* behaviour. In this case, the associated price of anarchy bounds may misrepresent the efficiency of actual learned outcomes. One motivating example here is that of auctions. For first-price single item auctions, in the complete information setting there may exist coarse correlated equilibria with suboptimal welfare, even though the unique equilibrium of the auction is welfare optimal [33]. Meanwhile, in the incomplete information setting with symmetric priors, whether coarse correlated equilibria coincide with the unique equilibrium depends on informational assumptions on the equilibrium structure itself [9] and on convexity of the priors [1]. This is in apparent contradiction with recent empirical work which suggest the equilibria of an even wider class of auctions are learnable when agents implement deep-learning or (with full feedback) gradient based no-regret learning algorithms [11, 12, 50].

This motivates the necessity of a more general notion of equilibrium analysis, stronger than coarse correlated equilibria for normal-form games and weaker than Nash equilibria, which never-theless captures the guarantees of the above-mentioned settings and is tractable to approximate or reason about. For the case of auctions, one recent proposal has been that of *mean-based learning algorithms* [14], but even in that case convergence results of [31, 34] are conditional.

There are two approaches towards the resolution of this question which, while not totally divorced in methodology, can be considered distinct in their philosophy. One approach has been to consider "game dynamics as the meaning of a game" [61], inspecting the existence of dynamics which converge to Nash equilibria, and extending price of anarchy analysis to include whether an equilibrium is reached with high probability. The work on the former has demonstrated impossibility results; there are games such that any gradient dynamics have starting points which do not converge to a Nash equilibrium, and for a set of games of positive measure no game dynamics may guarantee convergence to  $\epsilon$ -Nash equilibria [8, 52]. Meanwhile, [60, 66] proposed the average price of anarchy as a refined performance metric accounting for the game dynamics. The average price of anarchy is defined as the expectation over the set of initial conditions for the welfare of reached Nash equilibria, for a fixed gradient-based learning dynamics for the game.

Another approach has been to establish the computational complexity of *local* notions of equilibria. This has attracted attention especially in the setting non-concave games, where the existence of Nash equilibria is no longer guaranteed, due to the success of recent practical advances in machine learning via embracing non-convexity [28]. However, approximate minmax optimisation is yet again *PPAD*-complete [30]. As a consequence, unless  $PPAD \subseteq FP$ , a tractably approximable local equilibrium concept for non-concave games with compact & convex action sets must necessarily be *coarser*. Towards this end, [44, 41] define a notion of regret that is based on a sliding average of players' payoff functions. Meanwhile, a recent proposal by [16] has been to define a local correlated equilibrium, a distribution over action profiles and a set of local deviations such that, approximately, no player may significantly improve their payoffs when they deviate locally pointwise in the support of the distribution. They are then able to show for two classes of local deviations, such an approximate local correlated equilibrium is tractably computable.

Our goal in this paper is then to address the question of an equilibrium concept which is (1) also valid for non-concave games, (2) is stronger than coarse correlated equilibria for normal form games, (3) is tractable to approximate, and (4) admits a suitable extension of the usual primal-dual framework for bounding the expectation of quantities over the set of coarse correlated equilibria. The latter necessitates not only a distribution of play, but also incentive constraints which are specified only depending on the resulting distribution and not its time-ordered history. We remark

that framework of [44] falls short in the latter aspect when the cyclic behaviour of projected gradient ascent results in a non-vanishing regret bound. We thus turn our attention to generalising the work of [16].

Strikingly, in doing so, we demonstrate the intrinsic link between such local notions of coarse equilibria and the dynamical systems perspective on learning in games. In particular, the two local correlated notions of equilibria defined in [16] are subclasses of what we dub *local coarse correlated equilibria*, distribution of plays such that agents *in aggregate* do not have any strict incentive for infinitesimal changes of the support of the distribution along *any gradient field* over the set of action profiles. The history of play induced by online (projected) gradient ascent then approximates such an equilibrium, by virtue that it approximates a time-invariant distribution for the game's gradient dynamics. Extending the usual primal-dual scheme for price of anarchy bounds then reveals that any dual proof of performance bounds is necessarily of the form of a "generalised" Lyapunov function for the quantity whose expectation is to be bounded. The usual LP framework for coarse correlated equilibria is in fact contained in this approach, its dual optimal solutions corresponding to a "best-fit" quadratic Lyapunov function.

Our approach in proving our results combines insights previously explored in two previous works. The existence and uniqueness of solutions to projected dynamical systems over polyhedral sets is due to [32], and in our approximation proofs we also define a history over a continuous time interval. However, our analysis differs as we are not interested in approximating the continuous time projected dynamical system itself over the entire time interval; an approach that would doom our endeavour for tractable approximations in the presence of chaotic behaviour, which is not ruled out under our assumptions [22, 4]. Instead, we are interested in showing the approximate stationarity of expectations of quantities. Therefore, for projected gradient ascent, we suitably extrapolate the history of play into a piecewise differentiable curve. We then identify settings in which this curve moves approximately along the payoff gradients at each point in time via consideration of the properties of the boundary of the action sets in question.

Meanwhile, whereas Lyapunov-function based arguments is not new in analysis of convergence to equilibria in economic settings (e.g. [56]), in evolutionary game theory (c.f. [67] for an in-depth discussion), and in learning in games (e.g. [51, 75]), our perspective in bounding expectations of quantities appears to be relatively unexplored. Most Lyapunov-function based arguments in literature are concerned with pointwise (local) convergence to a unique Nash equilibrium, and work under the assumption of monotonicity or variational stability, or the existence of a potential function. The former two conditions imply the existence of a quadratic Lyapunov function for the game's projected gradient dynamics, from which Lyapunov functions for alternate learning processes may be constructed. One exception is [37], which deals explicitly with the problem of bounding expectations of stable distributions of Markov processes. Continuous time gradient dynamics is of course a Markov process, and a rather "boring" one in the sense that it is fully deterministic. Moreover, it is there that the dual solution of the LP bounding the expectation of some function is dubbed a "Lyapunov function", which motivates us to denote any of our dual solutions as a generalised Lyapunov function. However, their results do not include approximations of the Markov process and how Lyapunov-function based dual proofs extend to such approximations. Moreover, it is unclear how their results apply to projected gradient ascent, with polyhedral action sets. We are able to provide positive answers on both fronts.

In turn, our analysis of local correlated equilibria depends on the more established framework of regret matching. Our techniques here are essentially those used in [40, 39]. One key difference is our vector field formulation; usual swap-regret minimisation [71, 39, 40] considers mappings of the action space onto itself, while we measure regret against its differential generators. The consequences are reminiscent of the result of [43] on the equivalence of no regret learning and fixed point computation, in that we require access to a fixed-point oracle for the linear (or conical) combinations of vector fields in question to implement the regret matching algorithms. On the other hand, our vector field formulation allows us to extend the notion of correlated equilibria to a family of vector fields for which fixed-point computation is tractable; we are unaware of any similar observation.

### 1.2 Overview

In Section 2, we introduce our notation as well as the our definitions of local and stationary correlated equilibria, and their coarseness. In Section 3, we study the approximability of coarse correlated equilibria; we start by showing that continuous-time projected gradient dynamics approximates both generalisations of CCE which hints that online projected gradient ascent might provide such an approximation too. Positive results follow; in Section 3.1.1 we show that both local and stationary CCE are approximable whenever all players have compact and convex sets with a smooth boundary of bounded curvature, while in Section 3.1.2 we extend the approximability results to a class of polytopes which include the hypercube and the simplex. Section 3.2 considers the form of price of anarchy bounds for stationary coarse correlated equilibria, and shows that differentiable generalised Lyapunov functions provide such bounds for the outcomes of both approximate projected gradient dynamics and for online projected gradient ascent. Next we study correlated equilibria in Section 4; Section 4.1 contains our approximation bounds and generalisations of usual correlated equilibria, and in Setcion 4.2 we provide a brief remark on duality. In Section 5 we discuss further directions that remain to be investigated.

### 2 Preliminaries

In what follows,  $\mathbb{N}$  denotes the set of natural numbers<sup>1</sup>, and we identify also with each  $N \in \mathbb{N}$  the set  $\{n \in \mathbb{N} | 1 \leq n \leq N\}$ . Following standard notation, for a given  $N \in \mathbb{N}$ , and any tuple  $(X_j)_{j \in N}$ indexed by  $N^2$  and any  $i \in N$ ,  $X_{-i} \equiv (X_j)_{j \in N \setminus \{i\}}$  denotes the tuple with the *i*'th coordinate dropped. Meanwhile, for a tuple  $x \equiv (x_j)_{j \in N}$ , some  $i \in N$ , and  $y_i, (y_i, x_{-i})$  is the tuple where  $x_i$ is replaced by  $y_i$  in x. In addition, given some  $D \in \mathbb{N}$ , for each  $i \in D$  we will denote by  $e_i$  the standard basis vector in  $\mathbb{R}^D$  whose *i*'th component equals one and all others zero. Given a compact and convex set  $X \in \mathbb{R}^D$ , and some  $x \in X$ , we let  $\mathcal{TC}_X(x)$  and  $\mathcal{NC}_X(x)$  respectively denote the *tangent* and *normal* cones to X at x, i.e.

$$\mathcal{TC}_X(x) = \overline{\operatorname{conv}}\{t \cdot (y - x) \mid t \ge 0 \land y \in X\},\\ \mathcal{NC}_X(x) = \overline{\operatorname{conv}}\{z \in \mathbb{R}^D \mid \forall y \in \mathcal{TC}_X(x), \langle y, z \rangle \le 0\}.$$

Here,  $\overline{\text{conv}}$  denotes the *convex closure* of a set, and  $\langle x, y \rangle$  is the standard inner product of x and y in  $\mathbb{R}^D$ . In turn, for any  $D \in \mathbb{N}$ , and any  $y \in \mathbb{R}^D$ , we write  $||y|| = \sqrt{\langle y, y \rangle}$  for the standard

<sup>&</sup>lt;sup>1</sup>Including 0.

<sup>&</sup>lt;sup>2</sup>For instance, vectors or families of sets indexed by N.

Euclidean norm on  $\mathbb{R}^D$ , and  $\mathbb{P}_X[y] \equiv \arg\min_{x \in X} ||x - y||_2^2$  for the projection of y onto  $X \subseteq \mathbb{R}^D$ . Finally, given  $X \equiv \times_{i \in N} X_i$  such that each  $X_i \subseteq \mathbb{R}^{D_i}$  for some natural number  $D_i$ , some  $x \in X$ , and a differentiable function  $f : X \to \mathbb{R}$ ,  $\nabla f(x)$  denotes the usual gradient of f, while  $\nabla_i f(x)$  is the vector  $(\partial f/\partial x_{ij})_{i \in D_i} \in \mathbb{R}^{D_i}$ .

### **Definition 1.** The data of a smooth game is specified by a tuple $(N, (X_i)_{i \in N}, (u_i)_{i \in N})$ , where

- 1.  $N \in \mathbb{N}$  is the number (and by choice of notation, the set) of players.
- 2. For each  $i \in N$ ,  $X_i$  is the action set of player i. We assume that  $X_i$  is a compact and convex subset of  $\mathbb{R}^{D_i}$  for some  $D_i \in \mathbb{N}$ , and denote by  $X \equiv \times_{i \in N} X_i$  the set of outcomes of the game. We shall also write as the total dimension of the set of outcomes,  $D = \sum_{i \in N} D_i$ .
- 3. For each  $i \in N$ ,  $u_i : X \to \mathbb{R}$  is the utility function of player *i*. Each  $u_i$  is assumed to be continuously differentiable and have Lipschitz gradients, that is to say, there exists  $G_i, L_i \in \mathbb{R}_+$  such that for any  $x, x' \in X$ ,

$$\begin{aligned} \|\nabla u_i(x)\| &\leq G_i, \\ \|\nabla u_i(x) - \nabla u_i(x')\| &\leq L_i \|x - x'\|. \end{aligned}$$

We will denote  $\vec{G} \equiv (G_i)_{i \in N}$  and  $\vec{L} \equiv (L_i)_{i \in N}$ , the full vector of the bounds on players' gradients and Lipschitz coefficients.

Theoretic analysis of a game is often done by endowing it with an *equilibrium concept*, which specifies the "expected, stable outcome" of a given game. The standard equilibrium concept is that of a **Nash equilibrium** (NE), an outcome  $x \in X$  of a game such that for any player  $i \in N$ , and any action  $x'_i \in X_i$ ,  $u_i(x) \ge u_i(x'_i, x_{-i})$ . Whenever all  $u_i$  are concave over  $X_i$  given any fixed  $x_{-i} \in X_{-i}$ , such an equilibrium necessarily exists. However, for a generic smooth game, an NE need not exist, which motivates the notion of a **local Nash equilibrium** (local NE); an outcome  $x \in X$  is called a local NE if for every player  $i, \nabla_i u_i(x) \in \mathcal{NC}_{X_i}(x_i)$ .

By fixed point arguments, a local NE necessarily exists for a smooth game. However, the computation of such an equilibrium is not known to be tractable as of present date. This has lead to questions on whether there exists a *correlated* variant of a local NE; keeping in mind that correlated and coarse correlated equilibria (respectively, CE and CCE) of a finite game are computable in time polynomial in the size of its normal-form representation, one would ask whether such a local CCE is also tractable to (at least approximately) compute. Towards this end, Cai et al. [16] propose one definition of local CE, analogous to the definition of swap regret for finite games; for some  $\delta > 0$ , for each player i,  $\Phi^{X_i}(\delta)$  denotes a family of  $\delta$ -local strategy modifications  $\phi_i: X_i \to X_i$ , satisfying  $\|\phi(x_i) - x_i\| \leq \delta$  for any  $x_i \in X_i$ . Two proposals for such families are

$$\Phi_{\text{Int}}^{X_i}(\delta) = \{x_i \mapsto \delta x_i^* + (1-\delta)x_i \mid x_i^* \in X_i\},\\ \Phi_{\text{Proj}}^{X_i}(\delta) = \{x_i \mapsto \mathbb{P}_{X_i}[x_i + \delta v] \mid v \in \mathbb{R}^{D_i}, \|v\| \le 1\}.$$

The notion of a local CE is then defined by requiring players to not increase their payoff via such  $\delta$ -local strategy modifications with respect to the support of the distribution of play:

**Definition 2** ([16], Definition 2). For  $\epsilon, \delta > 0$ , a distribution  $\sigma$  over X is said to be an  $(\epsilon, \Phi(\delta))$ local CE if for any  $i \in N$  and any  $\phi_i \in \Phi^{X_i}(\delta)$ ,  $\mathbb{E}_{x \sim \sigma}[u_i(\phi_i(x_i), x_{-i}) - u_i(x)] \leq \epsilon$ . An immediate observation is that both  $\phi_{\text{Int}}^{X_i}$  and  $\phi_{\text{Proj}}^{X_i}$  as  $\delta$ -strategy modifications are provided by families of *gradient fields* of functions over  $X_i$ . Specifically, we may write

$$\Phi_{\text{Int}}^{X_i}(\delta) = \{ x_i \mapsto x_i - \delta \nabla_i \| x_i - x_i^* \|^2 / 2 \mid x_i^* \in X_i \},$$
(1)

$$\Phi_{\text{Proj}}^{X_i}(\delta) = \{ x_i \mapsto \mathbb{P}_{X_i}[x_i + \delta \nabla_i \langle x_i, v \rangle] \mid v \in \mathbb{R}^{D_i}, \|v\| \le 1 \}.$$
(2)

In what follows, we will exploit this observation and define alternate, *differential* definitions of local correlated equilibria, and its coarseness. In particular, we shall identify the appropriate notion of coarseness when the equilibrium constraints are all provided by gradient fields – as in this case, online gradient ascent will provide a *universal* approximation algorithm. To define our differential notion of equilibria, we first modify Definition 2 to only apply to sufficiently smooth vector fields.

**Definition 3.** For  $\epsilon, \Delta > 0$ , a distribution  $\sigma$  over X is said to be an  $(\epsilon, \Delta)$ -local CE with respect to a family F of Lipschitz continuous vector fields  $X \to \mathbb{R}^D$ , if for every  $f \in F$ ,

$$\sum_{i \in N} \mathbb{E}_{x \sim \sigma} \left[ u_i(\mathbb{P}_{X_i}[x_i + \delta f_i(x)], x_{-i}) - u_i(x) \right] \le \epsilon \delta \cdot \operatorname{poly}(\vec{G}, \vec{L}, G_f, L_f) + o(\delta) \ \forall \ \delta \in [0, \Delta].$$
(3)

Here,  $G_f$  and  $L_f$  are respectively bounds on the magnitude of ||f(x)|| and the Lipshitz coefficient of f, analogous to  $G_i$  and  $L_i$ . If  $\epsilon = 0$ ,  $\sigma$  is simply called a local CE.

**Definition 4** (Coarseness). For Definition 3 and in all subsequent definitions in this section, if the family F is given as a family of gradient fields  $\{\nabla h \mid h \in H\}$  for a subset  $H \subseteq C^1(X)$  of continuously differentiable functions,  $\sigma$  is said to be a **coarse** equilibrium.

We remark on the difference between the roles of  $\epsilon$  in Definition 2 and Definition 3; in the former,  $\epsilon$  is an absolute bound on the violation of the equilibrium constraints, while in the latter  $\epsilon$  is a multiplicative bound on the expectation of sum over the set of players of the derivatives of  $u_i$  in direction  $(\mathbb{P}_{X_i}[x_i + \delta f(x_i)] - x_i) / \delta$ . Another difference between the two definitions is that Definition 3 generalises to arbitrary suitably smooth vector fields  $f : X \to \mathbb{R}^D$  and not those defined solely on  $X_i$  for some  $i \in N$ .

Dividing both sides of (3) and taking the limit  $\delta \downarrow 0$  for an  $(\epsilon, \Delta)$ -local CE gives the following differential definition of an  $\epsilon$ -local CCE.

**Definition 5.** For  $\epsilon > 0$ , a distribution  $\sigma$  over X is said to be an  $\epsilon$ -local CE with respect to a family F of Lipschitz continuous vector fields  $X \to \mathbb{R}^D$ , if for every  $f \in F$ ,

$$\sum_{i \in N} \mathbb{E}_{x \sim \sigma} \left[ \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[f_i(x)], \nabla_i u_i(x) \right\rangle \right] \le \epsilon \cdot \operatorname{poly}(\vec{G}, \vec{L}, G_f, L_f).$$
(4)

Via an appeal to the projected gradient dynamics of the smooth game, we will show that such notions of local *coarse* correlated equilibria are in fact weakenings of the equilibrium concept we obtain by demanding our distribution  $\sigma$  to be invariant under time translation for the *projected* gradient dynamics of the smooth game. Such dynamics are provided by the system of differential equations,

$$\frac{dx_i(t)}{dt} = \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \ \forall \ i \in N,$$
(5)

with initial conditions  $x_i(0) \in X_i$  for each player *i*. By [25], for any initial conditions  $x(0) \in X$ , there is a unique absolutely continuous solution x(t) for  $t \in [0, \infty)$  such that (5) holds almost everywhere.

Now, suppose that a distribution  $\sigma$  is invariant under time translation, that is to say, for any measurable set  $A \subseteq X$  and any t > 0, the set  $A^{-1}(t) = \{x(0) \in X \mid x(t) \in A\}$  is measurable, and moreover,  $\sigma(A^{-1}(t)) = \sigma(A)$ . In this case, for any continuously differentiable function  $h: X \to \mathbb{R}$ ,  $\frac{d}{dt}\mathbb{E}_{x(0)\sim\sigma}[h(x(t))] = 0$  at t = 0. In particular, whenever the expectation and the time derivative commute,

$$\sum_{i \in N} \mathbb{E}_{x \sim \sigma} \left[ \left\langle \nabla_i h(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)} [\nabla_i u_i(x)] \right\rangle \right] = \sum_{i \in N} \mathbb{E}_{x(0) \sim \sigma} \left[ \left\langle \nabla_i h(x), dx_i(t)/dt \right\rangle \Big|_{t=0} \right] = 0.$$

In particular, distributions which are stationary under gradient dynamics satisfy a different regret property, in which the expected value of a function may not be modified via time translation.

In turn, if  $x \in X$  is a fixed-point of the gradient dynamics of the game, then  $\mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] = 0$  for every player *i*. Therefore, for any vector  $v \in \mathbb{R}^D$ ,  $\langle v, \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \rangle = 0$ . Together, these observations motivate our notion of *stationary* correlated equilibria.

**Definition 6.** For  $\epsilon > 0$ , a distribution  $\sigma$  is said to be an  $\epsilon$ -stationary CE with respect to a family F of Lipschitz continuous vector fields  $X \to \mathbb{R}^D$ , if for every  $f \in F$ ,

$$\left|\sum_{i\in N} \mathbb{E}_{x\sim\sigma}\left[\left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)], f_i(x)\right\rangle\right]\right| \leq \epsilon \cdot \operatorname{poly}(\vec{G}, \vec{L}, G_f, L_f).$$
(6)

## 3 On Local Coarse Correlated Equilibria

A primary objective of our work is to show that the above notions of  $\epsilon$ -local or  $\epsilon$ -stationary CCE are *universally* approximable in some "game theoretically relevant" settings. Here, by universal we mean that in Definition 4 we may take H be the set of all differentiable functions  $h: X \to \mathbb{R}$  with Lipschitz gradients and nevertheless tractably compute approximate local or stationary equilibria. We shall establish this by showing that projected gradient dynamics obtains an approximate stationary CCE, and such stationary CCE are necessarily also local CCE. Whereas for the purposes of practical algorithms we will need to investigate when all players take discrete steps in time, the form of analysis is motivated by how projected gradient dynamics yields our desired approximation.

To wit, let  $h \in H$  be given, and consider sampling uniformly from the trajectory of the projected gradient dynamics given initial conditions x(0). That is, we consider the distribution  $\sigma$  on X by drawing  $t \in [0, T]$  and sampling x(t). In this case,

$$\begin{split} \mathbb{E}_{x \sim \sigma} \left[ \sum_{i \in N} \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)], \nabla_i h(x) \right\rangle \right] &= \mathbb{E}_{t \sim U[0,T]} \left[ \sum_{i \in N} \left\langle dx_i(t)/dt, \nabla_i h(x(t)) \right\rangle \right] \\ &= \int_0^T \frac{1}{T} \frac{dh(x(t))}{dt} \cdot dt \\ &= \frac{h(x(T)) - h(x(0))}{T}. \end{split}$$

Now, let  $C \equiv \max_{x,x'} ||x - x'||$ . Then  $|h(x(T)) - h(x(0))| \leq C \cdot G_h$ , which immediately shows that we obtain an  $\epsilon$ -stationary CCE with  $\epsilon = C/T$  and poly $(\vec{G}, \vec{L}, G_h, L_h) = G_h$ . Interestingly, we do not yet need to invoke any bounds on the magnitude of the utility gradients nor the Lipschitz moduli.

Now, suppose for a simple example that each  $X_i$  is the  $D_i$ -dimensional [0, 1]-hypercube. In this case, we will argue that

$$\mathbb{E}_{x \sim \sigma} \left[ \sum_{i \in N} \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)} [\nabla_i h(x)], \nabla_i u_i(x) \right\rangle \right] \leq \mathbb{E}_{x \sim \sigma} \left[ \sum_{i \in N} \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)} [\nabla_i u_i(x)], \nabla_i h(x) \right\rangle \right],$$

which implies that the given distribution is an  $\epsilon$ -local CCE with respect to H with the very same  $\epsilon$ . We shall do so by showing that, for each  $i \in N$  and each  $t \in [0, T]$ ,

$$\left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i h(x)], \nabla_i u_i(x) \right\rangle \le \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)], \nabla_i h(x) \right\rangle$$
(7)

except on a subset of measure zero of [0, T]. Here, we supress the time-dependence for the purpose of parenthesis economy. In this case, first note that

$$\nabla_i h(x) = \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i h(x)] + \mathbb{P}_{\mathcal{NC}_{X_i}(x_i)}[\nabla_i h(x)], \text{ and}$$
$$\nabla_i u_i(x) = \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] + \mathbb{P}_{\mathcal{NC}_{X_i}(x_i)}[\nabla_i u_i(x)].$$

As a consequence,

$$\left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\nabla_{i}h(x)], \nabla_{i}u_{i}(x)\right\rangle \leq \left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\nabla_{i}u_{i}(x)], \nabla_{i}h(x)\right\rangle$$
$$\Leftrightarrow \left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\nabla_{i}h(x)], \mathbb{P}_{\mathcal{NC}_{X_{i}}(x_{i})}[\nabla_{i}u_{i}(x)]\right\rangle \leq \left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\nabla_{i}u_{i}(x)], \mathbb{P}_{\mathcal{NC}_{X_{i}}(x_{i})}[\nabla_{i}h(x)]\right\rangle.$$

Notice that, by the definition of the tangent and normal cones, both the left-hand side and the right-hand side specify quantities which are non-positive. Therefore, it is sufficient to show that the right-hand side equals 0 for almost every  $t \in [0, T]$ .

For the case of the hypercube, we may compute projections coordinate-wise. So suppose that at time t,  $\mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)]_j \cdot \mathbb{P}_{\mathcal{NC}_{X_i}(x_i)}[\nabla_i h(x)]_j < 0$ . This can only be when  $|\mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)]_j| > 0$ , hence for some  $\gamma(t) > 0$ , on the interval  $t' \in (t, t + \gamma(t))$ , no hypercube constraint for coordinate j binds. As a consequence, for  $t' \in (t, t + \gamma(t))$ ,  $\mathbb{P}_{\mathcal{NC}_{X_i}(x_i)}[\nabla_i h(x)]_j = 0$ . Now, denote by  $\overline{T}_{ij}$  the set of  $t \in [0, T]$  such that  $\mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)]_j \cdot \mathbb{P}_{\mathcal{NC}_{X_i}(x_i)}[\nabla_i h(x)]_j < 0$ . Then  $\sum_{t \in \overline{T}_{ij}} \gamma(t) < T$ , as each interval  $(t, t + \gamma(t))$  is disjoint. This in turn implies that  $\overline{T}_{ij}$  is countable, as the sum of a set of strictly positive numbers of uncountable cardinality is unbounded. Therefore,  $\bigcup_{i \in N, j \in D_i} \overline{T}_{ij}$  is countable. This is a superset (with potential equality) of  $t \in [0, T]$  for which (7) fails to hold for some player i, which implies our desired result.

These arguments form the basis of the proof methods by which we shall show that projected online gradient ascent (with discrete steps) can be used to find an  $\epsilon$ -local or stationary CCE. In particular, we will first show the approximability of  $\epsilon$ -stationary CCE by working with a *piecewiselinear* approximation of the underlying projected dynamical system; such an approximation will require invocation of the Lipschitz coefficients of utility gradients. That  $\epsilon$ -local CCE itself is approximable will then follow, by arguing that the zero-measure property above is maintained for the analogue of (7).

### 3.1 Tractable approximations via projected gradient ascent

#### 3.1.1 Compact & convex sets of smooth boundary

The first question we answer is that concerning approximability. We will approach the problem for classes of compact and convex action sets for two classes of action sets for which the problem is straightforward. The first case concerns the setting in which each  $X_i$  possesses a smooth boundary, of curvature bounded by some K > 0. Dealing with this setting will necessitate the definition of our online learning algorithms, as well as the use of the local quadratic approximation for surfaces of bounded curvature in Euclidean space.

We define our approximate projected gradient dynamics first. Given a sequence of decreasing step sizes  $\eta_{\tau} > 0$  such that  $\sum_{\tau=0}^{\infty} \eta_{\tau} = \infty$  and  $\sum_{\tau=0}^{\overline{\tau}} \eta_{\tau}^2 = o\left(\sum_{\tau=0}^{\overline{\tau}} \eta_{\tau}\right)$ , projected gradient ascent with equal learning rates has each player  $i \in N$  play  $x_i^0$  at time  $\tau = 0$ , and update their strategies

$$x_i^{\tau+1} = \mathbb{P}_{X_i}[x_i^{\tau} + \eta_\tau \nabla_i u_i(x^{\tau})].$$

After  $\overline{\tau} > 0$  time steps, the learning dynamics outputs a history  $(x^{\tau})_{\tau=0}^{\overline{\tau}}$  of play. We proceed by defining via this history an **approximate projected gradient dynamics** for the game, via extending this distribution in a piecewise fashion to a curve  $x(t) : [0,T] \to X$  where  $T = \sum_{\tau=0}^{\overline{\tau}-1} \eta_{\tau}$ . Towards this end, if there exists  $\tau' \in \mathbb{N}$  such that  $t = \sum_{\tau=0}^{\tau'-1} \eta_{\tau}$ , then  $x(t) = x^{\tau'}$ . Otherwise, for  $t \in [0,T]$ , let  $\underline{t} = \max\left\{\sum_{\tau=0}^{\tau'-1} \eta_{\tau} | \tau' \in \mathbb{N}, \sum_{\tau=0}^{\tau'-1} \eta_{\tau} \leq t\right\}$ , and fix

$$x(t) = \mathbb{P}_{X_i}[x_i + (t - \underline{t}) \cdot \nabla_i u_i(x(\underline{t}))].$$

We will then denote  $\eta_t \equiv \eta_{\tau'}$  for the  $\tau'$  which determines <u>t</u>.

The first step of the analysis comprises of bounding the difference between  $\frac{dx_i(t)}{dt}$  and the tangent cone projection  $\mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(t))]$ . This turns out to be a straightforward matter in this setting, thanks to the following observation.

**Proposition 3.1.** Let  $X_i$  be a compact and convex set of smooth boundary, and its boundary  $\delta X_i$  has non-negative principal curvature at most K. Then for any  $t \in [0, T]$ ,

$$\left\|\frac{dx_i(t)}{dt} - \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(\underline{t}))]\right\| \le \frac{KG_i \cdot (t - \underline{t})}{1 + KG_i \cdot (t - \underline{t})} \cdot \left\|\mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(\underline{t}))]\right\|$$

*Proof.* We shall work in coordinates where  $x_i(t)$  is at the origin, and  $\Delta x_i(t) = x_i(\underline{t}) + (t - \underline{t}) \cdot \nabla_i u_i(x(\underline{t})) = h \cdot e_1$  for some  $h \ge 0$ , where each  $e_2, e_3, \dots, e_{d_i}$  are principal directions of curvature for  $\delta X_i$  at  $x_i(t)$ . In this case, note that the smoothness of the boundary implies that, for a small neighbourhood about  $x_i(t)$ ,  $X_i$  is well-approximated (up to second order in  $(w_\ell)_{\ell=2}^{d_i}$ ) by the convex body,

$$\tilde{X}_i = \left\{ w \in \mathbb{R}^{d_i} \mid w_1 \le \sum_{\ell=2}^{d_i} -k_\ell w_\ell^2 \right\},\$$

where  $k_{\ell}$  is the (principal) curvature in direction  $e_{\ell}$  for the surface  $\delta X_i$  at point  $x_i(t)$ . Note that by assumption,  $k_{\ell} \leq K$  for every  $2 \leq \ell \leq d_i$ . As a consequence, at time  $t + \epsilon$  for small  $\epsilon > 0$ ,  $x_i(t + \epsilon)$  and the solution to the projection problem

$$\min_{w} (w_1 - h - \epsilon g_1)^2 + \sum_{\ell=2}^{d_i} (w_\ell - \epsilon g_\ell)^2 \text{ subject to } w_1 \le \sum_{\ell=2}^{d_i} -k_\ell w_\ell^2$$

agree up to first-order terms in  $\epsilon$ . At any solution, the constraint will bind, leading us to the unconstrained optimisation problem

$$\min_{w} \left( -h - \epsilon g_1 - \sum_{\ell=2}^{d_i} k_{\ell} w_{\ell}^2 \right)^2 + \sum_{\ell=2}^{d_i} (w_{\ell} - \epsilon g_{\ell})^2.$$

Now, the first-order optimality conditions are

$$\forall \ 2 \le \ell \le d_i, k_\ell w_\ell \left( -h - \epsilon g_1 - \sum_{\ell=2}^{d_i} k_\ell w_\ell^2 \right) + w_\ell - \epsilon g_\ell = 0.$$

For sufficiently small  $\epsilon > 0$ , at an optimal solution  $w_{\ell} = O(\epsilon)$  for each  $\ell \ge 2$  and  $\sum_{\ell=2}^{d_i} k_{\ell} w_{\ell}^2 = O(\epsilon^2)$ . Therefore, up to first-order in  $\epsilon$ , for each  $\ell \ge 2$ ,

$$w_{\ell} = \frac{\epsilon g_{\ell}}{1 + k_{\ell} h}.$$

This in turn implies that

$$\frac{dx_{i\ell}(t)}{dt} = \begin{cases} 0 & \ell = 1, \\ \frac{g_\ell}{1+k_\ell h} & \ell \neq 1. \end{cases}$$

Therefore, the difference between the velocity of motion and the projection to the tangent cone of the utility gradient is, for any principal direction  $\ell$ ,

$$\left(\frac{dx_i(t)}{dt} - \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(\underline{t}))]\right)_{\ell} = -g_\ell \cdot \left(\frac{k_\ell h}{1 + k_\ell h}\right).$$

By the bound on the magnitude of  $\nabla_i u_i(x(\underline{t}))$  and the feasibility of  $x_i(\underline{t})$ ,  $h \leq G_i \cdot (t - \underline{t})$ . In turn,  $\frac{k_\ell h}{1+k_\ell h} \leq \frac{Kh}{1+Kh}$ , by which we have the desired result.

The Lipschitz continuity of  $u_i$  along with Proposition 3.1 then yields proof of approximation for an  $\epsilon$ -stationary CCE.

**Theorem 3.2.** Whenever all  $X_i$  are compact and convex sets with smooth boundary of bounded principal curvature K, the outcome of approximate projected gradient dynamics satisfies for every player i,

$$\frac{1}{T} \cdot \int_0^T dt \cdot \left| \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{T}\mathcal{C}_{X_i}(x_i(t))}[\nabla_i u_i(x(t))] - \frac{dx_i(t)}{dt} \right\rangle \right| \le \frac{1}{T} \cdot \frac{1}{2} \left( \sum_{\tau=0}^{\overline{\tau}-1} \eta_\tau^2 \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2 + L_i \sum_{i \in N} G_i \right) \cdot G_h \left( KG_i^2$$

Then by choice of step-size  $\eta_{\tau} = \frac{1}{\sqrt{\tau+1}}$ , after  $\overline{\tau}$  time steps approximate projected gradient dynamics yields a  $\left(\frac{2d(X)+\log(\overline{\tau})}{2\sqrt{\overline{\tau}}}\right)$ -stationary CCE with respect to the set of all functions  $h: X \to \mathbb{R}$  with Lipschitz gradients, where  $d(X) = \max_{x,x' \in X} ||x - x'||$  is the diameter of X.

*Proof.* To see the result, we shall write

$$\left\langle \nabla_{i}h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}u_{i}(x(t))] - \frac{dx_{i}(t)}{dt} \right\rangle$$
$$= \left\langle \nabla_{i}h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}u_{i}(x(t))] - \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}u_{i}(x(\underline{t}))] \right\rangle$$
(8)

$$+\left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(\underline{t}))] - \frac{dx_i(t)}{dt} \right\rangle.$$
(9)

We then proceed by taking the absolute value of both sides and applying the triangle inequality. The magnitude of (8) is bounded, by the Cauchy-Schwarz inequality,

$$\begin{aligned} \left| \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{T}\mathcal{C}_{X_i}(x_i(t))} [\nabla_i u_i(x(t))] - \mathbb{P}_{\mathcal{T}\mathcal{C}_{X_i}(x_i(t))} [\nabla_i u_i(x(\underline{t}))] \right\rangle \right| \\ &\leq \|\nabla_i h(x(t))\| \|\mathbb{P}_{\mathcal{T}\mathcal{C}_{X_i}(x_i(t))} [\nabla_i u_i(x(t))] - \mathbb{P}_{\mathcal{T}\mathcal{C}_{X_i}(x_i(t))} [\nabla_i u_i(x(\underline{t}))] \| \\ &\leq \|\nabla_i h(x(t))\| \|\nabla_i u_i(x(t)) - \nabla_i u_i(x(\underline{t}))\| \leq (t - \underline{t}) G_h L_i \sum_{j \in N} G_j. \end{aligned}$$

Here, the very last inequality uses the bounds on the magnitude of the gradients and the Lipschitz coefficient of  $u_i$ , that  $||x(t) - x(\underline{t})|| \leq (t - \underline{t}) \cdot \sum_{j \in N} G_j$ , and that the projection operators onto closed convex sets in Euclidean space are non-expansive. Meanwhile, the second term (9) satisfies

$$\left| \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))} [\nabla_i u_i(x(\underline{t}))] - \frac{dx_i(t)}{dt} \right\rangle \right|$$
  

$$\leq \|\nabla_i h(x(t))\| \|\nabla_i u_i(x(\underline{t}))\| \cdot \frac{KG_i(t-\underline{t})}{1+KG_i(t-\underline{t})}$$
  

$$\leq G_h KG_i^2(t-\underline{t}).$$

Combining the above bounds, we have

$$\frac{1}{T} \cdot \int_{0}^{T} dt \cdot \left| \left\langle \nabla_{i}h(x(t)), \mathbb{P}_{\mathcal{T}\mathcal{C}_{X_{i}}(x_{i}(t))}[\nabla_{i}u_{i}(x(t))] - \frac{dx_{i}(t)}{dt} \right\rangle \right| \\
\leq \frac{1}{T} \cdot \int_{0}^{T} dt \cdot (t - \underline{t}) \cdot G_{h} \left( KG_{i}^{2} + L_{i} \sum_{j \in N} G_{j} \right) \\
= \frac{1}{T} \cdot \sum_{\tau=0}^{\overline{\tau}-1} \int_{0}^{\eta_{\tau}} dt \cdot t \cdot G_{h} \left( KG_{i}^{2} + L_{i} \sum_{j \in N} G_{j} \right) \\
= \frac{1}{T} \cdot \frac{1}{2} \left( \sum_{\tau=0}^{\overline{\tau}-1} \eta_{\tau}^{2} \right) \cdot G_{h} \left( KG_{i}^{2} + L_{i} \sum_{j \in N} G_{j} \right).$$

Then the final approximation bound is obtained,

$$\begin{aligned} \left| \frac{1}{T} \int_0^T dt \cdot \sum_{i \in N} \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))} [\nabla_i u_i(x(t))] \right\rangle \right| \\ &\leq \left| \frac{1}{T} \sum_{i \in N} \int_0^T dt \cdot \left\langle \nabla_i h(x(t)), \frac{dx_i(t)}{dt} \right\rangle \right| \\ &+ \frac{1}{T} \sum_{i \in N} \int_0^T dt \cdot \left| \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))} [\nabla_i u_i(x(t))] - \frac{dx_i(t)}{dt} \right\rangle \right| \\ &\leq \frac{1}{T} |h(x(T)) - h(x(0))| + \frac{1}{T} \cdot \frac{1}{2} \left( \sum_{\tau=0}^{\overline{\tau}-1} \eta_\tau^2 \right) \sum_{i \in N} G_h \left( KG_i^2 + L_i \sum_{j \in N} G_j \right). \end{aligned}$$

Of course,  $|h(T) - h(0)| \leq G_h d(X)$ , so setting poly $(\vec{G}, \vec{L}, G_h, L_h) = G_h \sum_{i \in N} \left( 1 + KG_i^2 + L_i \sum_{j \in N} G_j \right)$  for Definition 6 yields the desired result.

It remains to show the approximability of an  $\epsilon$ -local CCE with respect to the set of all differentiable functions  $h : X \to \mathbb{R}$  with Lipschitz gradients. This, again, follows from a tangency argument.

**Theorem 3.3.** Whenever all  $X_i$  are compact and convex sets with smooth boundary of bounded principal curvature K, the outcome of approximate projected gradient dynamics satisfies for every player i,

$$\frac{1}{T} \cdot \int_0^T dt \cdot \left\langle \nabla_i u_i(x(t)) - \frac{dx_i(t)}{dt}, \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i h(x(t))] \right\rangle \le \frac{1}{T} \cdot \frac{1}{2} \left( \sum_{\tau=0}^{\overline{\tau}-1} \eta_\tau^2 \right) \cdot G_h \left( KG_i^2 + L_i \sum_{j \in N} G_j \right)$$

Then by choice of step-size  $\eta_{\tau} = \frac{1}{\sqrt{\tau+1}}$ , after  $\overline{\tau}$  time steps approximate projected gradient dynamics yields a  $\left(\frac{2d(X)+\log(\overline{\tau})}{2\sqrt{\overline{\tau}}}\right)$ -local CCE with respect to the set of all functions  $h: X \to \mathbb{R}$  with Lipschitz gradients, where  $d(X) = \max_{x,x' \in X} ||x - x'||$  is the diameter of X.

*Proof.* Note that for any  $t \in [0, T]$  and for every player i,

$$\left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}h(x)], \nabla_{i}u_{i}(x(t))\right\rangle$$

$$\leq \left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}h(x)], \nabla_{i}u_{i}(x(t)) - \nabla_{i}u_{i}(x(\underline{t}))\right\rangle$$
(10)

$$+ \left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}h(x)], \mathbb{P}_{\mathcal{NC}_{X_{i}}(x_{i}(t))}[\nabla_{i}u_{i}(x(\underline{t}))] \right\rangle$$
(11)

$$+ \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i h(x)], \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(\underline{t}))] - \frac{dx_i(t)}{dt} \right\rangle$$
(12)

$$+\left\langle \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i}(t))}[\nabla_{i}h(x)], \frac{dx_{i}(t)}{dt} \right\rangle.$$
(13)

The terms (10) and (12) yield the very same bounds as in the proof of Theorem 3.2. Meanwhile, (11) is  $\leq 0$  by definition of the tangent and the normal cones. As for (13), the normal vector component of  $\frac{dx_i(t)}{dt}$  equals 0 except potentially on a set of measure zero by an argument analogous to the one made for hypercubes in Section 3. As a consequence, for almost every  $t \in [0, T]$ ,

$$\left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i h(x)], \frac{dx_i(t)}{dt} \right\rangle = \left\langle \nabla_i h(x), \frac{dx_i(t)}{dt} \right\rangle.$$

The discussion in this section demonstrates common themes regarding the connection between local CCE, stationary CCE, and the projected gradient dynamics of the underlying smooth game. Approximate projected gradient dynamics provides an approximate stationary CCE, and simultaneously an approximate local CCE, as the tangency of the motion through the strategy space implies that as far as the gradient dynamics are concerned, a local CCE is a relaxation of a stationary CCE. Meanwhile, the complexity of approximation is determined via the properties of the boundary of the strategy space, with approximation guarantees deteriorating linearly with the maximum curvature. All these features will be reobserved in the discussion that follows, where we turn our attention to polyhedral action sets.

### 3.1.2 On polyhedral sets

In this section, we will discuss the approximability of stationary and local CCE for smooth games, when the action set  $X_i$  of each player i is polyhedral; that is to say, for each player i, there exists  $A_i \in \mathbb{R}^{m_i \times d_i}$  and  $b_i \in \mathbb{R}^{m_i}$  such that  $X_i = \{x \in \mathbb{R}^{d_i} \mid A_i x \leq b_i\}$ . We shall assume, without loss of generality, that all involved polyhedra have volume in  $\mathbb{R}^d_i$ , and each  $A_i x \leq b_i$  has no redundant inequalities. Moreover, the rows  $a_{ij}$  of each  $A_i$  will be assumed to be normalised,  $\|a_{ij}\| = 1 \forall i \in N, j \in m_i$ . With that in mind, we begin by defining the class of polyhedra of interest, over which approximation of stationary or local CCE is "easy".

**Definition 7.** A polyhedron  $\{x \in \mathbb{R}^d \mid \langle a_j, x \rangle \leq b_j \forall j \in m\}$  is said to be **acute** if for every distinct  $j, j' \in m, \langle a_j, a_{j'} \rangle \leq 0$ .

Examples of acute polyhedra include potentially the most "game-theoretically relevant" polyhedra, the hypercube and the simplex. That the hypercube is acute in this sense is straightforward: for each  $\ell \in d_i$ , the constraints  $-x_{i\ell} \leq 0$  and  $x_{i\ell} \leq 1$  have negative inner product for the corresponding rows of A, while any other distinct pairs of rows of A are orthogonal. In turn for the simplex, factoring out the constraint  $\sum_{\ell=1}^{d_i} x_i = 1$ , we are left with the set of inequalities,

$$\forall \ \ell \in d_i, x_{i\ell} - \sum_{\ell'=1}^{d_i} \frac{1}{n} x_{i\ell'} \ge -\frac{1}{n}.$$

For any distinct  $\ell, \ell''$ , the inner product of the vectors associated with the left-hand side of these constraints is then

$$\left\langle e_{i\ell} - \sum_{\ell'=1}^{d_i} \frac{1}{n} e_{i\ell'}, e_{i\ell''} - \sum_{\ell'=1}^{d_i} \frac{1}{n} e_{i\ell'} \right\rangle = -\frac{1}{n} - \frac{1}{n} + \frac{1}{n} < 0.$$

As mentioned, the importance of acute polyhedra is that linear optimisation over them is trivial; a greedy algorithm suffices. Moreover, over such polyhedra  $x_i(t)$  always follows the tangent cone projection of  $\nabla_i u_i(x(\underline{t}))$ , rendering the approximate projected gradient dynamics *faithful*. The latter statement is the one we need for the desired approximation bounds, which is proven in the following lemma.

**Lemma 3.4.** Suppose that  $X = \{x \in \mathbb{R}^d \mid \langle a_j, x \rangle \leq b_j \forall j \in m\}$  is an acute polyhedron in  $\mathbb{R}^d$ ,  $x^* \in X$ , and  $g \in \mathbb{R}^d$ . Then for  $x(t) = \mathbb{P}_X[x^* + t \cdot g]$ ,

$$\frac{dx(t)}{dt} = \mathbb{P}_{\mathcal{TC}_X(x(t))}[g].$$

*Proof.* Without loss of generality, we will work in coordinates such that  $x^* = 0$ . Consider the projection problem for any  $t' \ge 0$ ,

$$\min_{x} \frac{1}{2} \|x - t' \cdot g\|^2 \text{ subject to } Ax \le b.$$

Its dual problem is given,

$$\max_{\mu \ge 0} -\frac{1}{2} \|A^T \mu\|^2 + \langle \mu, t' \cdot Ag - b \rangle.$$

Now, given a sequence t' > t decreasing to  $t' \downarrow t$ , there is an infinitely repeated set of active rows  $J = \{j \in m \mid \mu_j(t') > 0\}$  for the optimal solutions  $\mu(t')$  to the dual projection problem. This is necessarily the set of active and relevant constraints at time t, and for the corresponding  $|J| \times d$  submatrix B of A,

$$\mu_j(t') = ((BB^T)^{-1}[t' \cdot Bg - b])_j,$$

for  $t' \ge t$  sufficiently close to t. In particular,  $d\mu(t)_j/dt = [(BB^T)^{-1}Bg]_j$  whenever  $j \in J$  and zero otherwise, while  $dx(t)/dt = g - B^T (BB^T)^{-1}Bg$ .

Now, the acuteness condition implies that the off-diagonal elements of  $BB^T$  are non-positive, whereas it is both positive semi-definite and invertible; implying it is positive definite. Therefore, by [35] (Theorem 4.3, 9° and 11°),  $(BB^T)^{-1}$  has all of its entries non-negative. Meanwhile, by feasibility of  $x^*$ , b has all of its entries non-negative, which implies that  $(BB^T)^{-1}b$  also has only non-negative entries. Therefore,  $d\mu_j(t)/dt = (\mu_j(t) + [(BB^T)^{-1}b]_j)/t > 0$  for every row j which is active, and zero otherwise.

Thus all that remains to show is that  $\mathbb{P}_{\mathcal{TC}_X(x)}[g] = g - B^T (BB^T)^{-1} Bg$ . Let  $I \subseteq m$  be the set of constraints which bind at x(t) (but potentially, for  $j \in I, d\mu(t)_j/dt = 0$ ), and C the associated  $|I| \times m$  submatrix of A. Then consider the tangent cone projection problem

$$\min_{x} \frac{1}{2} \|x - g\|^2 \text{ subject to } Cx \le 0.$$

The dual projection problem is then,

$$\max_{\nu \ge 0} -\frac{1}{2} \|C^T \nu\|^2 + \langle \nu, Cg \rangle.$$

We need to show that  $\nu = (BB^T)^{-1}Bg$  is an optimal solution. But this is immediate now, as the feasibility of  $x(t) + \Delta t \cdot (g - B^T (BB^T)^{-1}Bg)$  for small  $\Delta t$  implies that  $x = g - B^T (BB^T)^{-1}Bg$  is a feasible solution to the tangent cone projection problem, with solution value  $g^T B^T (BB^T)^{-1}Bg/2$ . Meanwhile, the given  $\nu = d\mu(t)/dt$  is dual feasible, with the same solution value. By weak duality, both solutions are necessarily optimal.

**Remark.** The implications of Lemma 3.4 were already proven for the hypercube and the simplex in [54], in which the authors study approximations of projected gradient dynamics. Indeed, the analysis here is in a similar vein in that we track the time evolution of the projection curve, though we identify and exploit features of the polyhedron which makes linear optimisation over it trivial.

It turns out that for acute polyhedra, it is possible to prove stronger approximation guarantees than the general case. Intuitively, that the approximate motion never fails to move along the tangent of utility gradients results in well-approximability of the projected gradient dynamics of the smooth game.

**Theorem 3.5.** Whenever all  $X_i$  are acute polyhedra, the outcome of the approximate gradient dynamics satisfies

$$\frac{1}{T} \cdot \int_0^T dt \cdot \left| \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(t))] - \frac{dx_i(t)}{dt} \right\rangle \right| \le \frac{1}{T} \cdot \frac{1}{2} \left( \sum_{\tau=0}^{\overline{\tau}-1} \eta_\tau^2 \right) \cdot G_h L_i \sum_{i \in N} G_i$$

As a consequence, by choice of step-size  $\eta_{\tau} = 1/\sqrt{1+\tau}$ , approximate projected gradient dynamics yields a  $\left(\frac{2d(X)+\log(\tau)}{2\sqrt{\tau}}\right)$ -stationary CCE with respect to the set of all functions  $h: X \to \mathbb{R}$  with Lipschitz gradients. The same bound also holds for the approximability of local CCE. *Proof.* Near identical to the proof of Theorem 3.2 and Theorem 3.3, except by Lemma 3.4

$$\mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(\underline{t}))] - \frac{dx_i(t)}{dt} = 0,$$

implying that there are no projection losses. However, we nevertheless need an alternative argument that for almost every  $t \in [0, T]$ ,

$$\left\langle \mathbb{P}_{\mathcal{NC}_{X_i}(x_i(t))}[\nabla_i h(x(t))], \frac{dx_i(t)}{dt} \right\rangle = 0.$$

Now, by [54] (Lemma 2), the curve  $x_i(t)$  is piecewise linear whenever  $X_i$  is polyhedral. Then, a constraint  $\langle a_{ij}, x_j \rangle \leq b_{ij}$  binds over an interval  $(t, t+\epsilon)$  only if  $\langle a_{ij}, \frac{dx_i(t)}{dt} \rangle = 0$  for such t. However,  $\mathbb{P}_{\mathcal{NC}_{X_i}(x_i(t))}[\nabla_i h(x(t))] = \sum_j \mu_j(t) a_{ij}$  for such binding constraints  $a_{ij}$  and some  $\mu_j(t) \geq 0$ , which implies the desired result.

The interesting case is thus when the polyhedra we consider are not acute. This is counterintuitive; acuteness in a sense implies that smooth boundary approximations of the convex body would have very high curvature, a condition under which the differential regret guarantees of approximate projected gradient dynamics *deteriorates*. However, acute polyhedra in turn enjoy even better convergence guarantees. Meanwhile, the case for general polyhedra remains open.

### 3.2 Duality and Lyapunov arguments

After our proofs of approximability, we proceed by discussing the appropriate primal-dual framework for  $\epsilon$ -stationary CCE. Our approach is motivated by the possibility of strengthening primaldual efficiency analysis for the outcomes of learning algorithms. Often, in the analysis of games, the object of interest is performance guarantees attached to an equilibrium concept and not necessarily the exact form of equilibrium. That is to say, given an equilibrium concept  $E \subseteq 2^{\Delta(X)}$ , and a continuous "welfare" function  $q: X \to \mathbb{R}$ , one may consider a bound on the worst case performance

$$\frac{\inf_{\sigma \in E} \mathbb{E}_{x \sim \sigma}[q(x)]}{\max_{x \in X} q(x)}$$

This quantity is often referred to as the *price of anarchy* of the game, while the related notion of *price of stability* bounds the best case performance in equilibria,

$$\frac{\sup_{\sigma \in E} \mathbb{E}_{x \sim \sigma}[q(x)]}{\max_{x \in X} q(x)}$$

Most methods of bounding such expectations fall under two classes of a so-called primal-dual method. One class of such methods argues the bounds via the lens of primal-dual approximation algorithms, writing a primal LP ("configuration LP") corresponding to performance (e.g. welfare) maximisation, and obtaining via the equilibrium conditions a solution for its dual minimisation LP. If the equilibrium is good, the constructed dual solution has value a constant factor of the performance of the equilibrium, which provides the desired result. An in-depth discussion on this approach might be found in [58]; we remark that this is not the primal-dual framework of interest here.

Another class of primal-dual methods argue directly via a primal problem over the set of equilibria. Here, the primal LP has as its variables probability distributions over X, and is subject to equilibrium constraints. The objective is to minimise or maximise q(x), corresponding respectively to computing the price of anarchy or stability up to a factor of  $\max_{x \in X} q(x)$ . For instance, (by the arguments in [55]) the smoothness framework of [65] as well as the price of anarchy bounds for congestion games by [13] both fall under this umbrella. The question is whether such arguments may be extended for the outcomes of gradient ascent.

The answer, of course, is *yes*; and for the concept of equilibrium of choice it is more enlightening to invoke  $\epsilon$ -stationary CCE constraints rather than  $\epsilon$ -local CCE constraints. To wit, given the function q, consider the following measure valued (infinite dimensional) LP,

$$\inf_{\sigma \ge 0} \int_X d\sigma(x) \cdot q(x) \text{ subject to}$$
(14)

$$\int_X d\sigma(x) = 1 \tag{(\gamma)}$$

$$\forall h \in H, \sum_{i \in N} \int_X d\sigma(x) \cdot \left\langle \nabla_i h(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle = 0. \tag{(}\mu(h))$$

The Lagrangian dual of (14) may then be naively written,

$$\sup_{\gamma,\mu} \gamma \text{ subject to} \tag{15}$$

$$\forall x \in X, \gamma + \int_{H} d\mu(h) \cdot \sum_{i \in N} \left\langle \nabla_{i} h(x), \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\nabla_{i} u_{i}(x)] \right\rangle \leq q(x). \tag{\sigma(x)}$$

Here,  $\gamma$  is some real number, while  $\mu$  is assumed to be "some measure" on H. Of course, we may simply pick a dual solution  $\mu$  which places probability 1 on some element  $h \in H$ . Under such a restriction, the dual problem is then

$$\sup_{\gamma \in \mathbb{R}, h \in H} \gamma \text{ subject to} \tag{16}$$

$$\forall \ x \in X, \gamma + \sum_{i \in N} \left\langle \nabla_i h(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle \le q(x). \tag{\sigma(x)}$$

Note that the dual is always feasible. Whether weak duality between the primal LP (14) and either one of the dual problems (15) or (16) depends on the specification of the vector spaces the primal and dual variables live in, and as a consequence whether they form a dual pair in the sense of [3, 69].

However, given a differentiable  $h: X \to \mathbb{R}$  with Lipschitz gradients, and the outcome of approximate gradient dynamics  $x: [0, T] \to X$ , such dual arguments are always valid, as the integral along the curve x exists. In particular, if

$$\gamma = \min_{x \in X} q(x) - \sum_{i \in N} \left\langle \nabla_i h(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle,$$

then the time expectation of q(x(t)) satisfies

$$\begin{aligned} \frac{1}{T} \int_0^T dt \cdot q(x(t)) &\geq \frac{1}{T} \int_0^T dt \cdot \left( \gamma + \sum_{i \in N} \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(t))] \right\rangle \right) \\ &\geq \gamma - \frac{1}{T} \int_0^T \left| \sum_{i \in N} \left\langle \nabla_i h(x(t)), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[\nabla_i u_i(x(t))] \right\rangle \right| \\ &\geq \gamma - \operatorname{poly}(\vec{G}, \vec{L}, G_h, L_h) \cdot \left( \frac{2d(X) + \log(\overline{\tau})}{2\sqrt{\overline{\tau}}} \right) \end{aligned}$$

for step-sizes  $\eta_{\tau} = 1\sqrt{\tau + 1}$ . Moreover, whenever q(x) is also differentiable with Lipschitz gradients, the order of the dependence of the convergence rate on  $\overline{\tau}$  is maintained also for projected gradient ascent. This can be shown via the trivial bound on the difference between the expectations of q(x)for the two resulting distributions; note that for each  $t \in [0, T]$ ,

$$\begin{aligned} |q(x(t)) - q(x(\underline{t}))| &\leq \left| \langle x(t) - x(\underline{t}), \nabla q(x) \rangle + \frac{1}{2} L_q ||x(t) - x(\underline{t})||^2 \right| \\ &\leq (t - \underline{t}) \cdot \left( G_q \sum_{i \in N} G_i + \frac{1}{2} L_q \sum_{i,j \in N} G_i G_j \right), \end{aligned}$$

where we use the fact that by our choice of step-sizes,  $t - \underline{t} \leq 1$  for every  $t \in [0, T]$ . Therefore,

$$\frac{1}{T} \left| \sum_{\tau=0}^{\overline{\tau}-1} \eta_{\tau} q(x^{\tau}) - \int_{0}^{T} dt \cdot q(x(t)) \right| = \frac{1}{T} \cdot \sum_{\tau=0}^{\overline{\tau}-1} \int_{0}^{\eta_{\tau}} dt \cdot t \cdot \left( G_{q} \sum_{i \in N} G_{i} + \frac{1}{2} L_{q} \sum_{i,j \in N} G_{i} G_{j} \right) \\
\leq \frac{1}{T} \left( \sum_{\tau=0}^{\overline{\tau}-1} \eta_{\tau}^{2} \right) \cdot (G_{q} + L_{q}) \operatorname{poly}(\vec{G}, \vec{L}, G_{h}, L_{h}).$$

**Theorem 3.6.** Suppose that q(x) is a continuous function, and  $\gamma$ , h are solutions to (16). Then with step-sizes  $\eta_{\tau} = 1/(\sqrt{\tau} + 1)$ , after  $\overline{\tau}$  steps, the outcome of approximate projected gradient dynamics  $x(t) : [0,T] \to X$  satisfies,

$$\frac{1}{T} \int_0^T dt \cdot q(x(t)) \ge \gamma - \operatorname{poly}(\vec{G}, \vec{L}, G_h, L_h) \cdot \left(\frac{2d(X) + \log(\overline{\tau})}{2\sqrt{\overline{\tau}}}\right).$$

If also q(x) is differentiable with Lipschitz gradients, then the outcome of projected gradient ascent  $(x^{\tau})_{\tau=0}^{\overline{\tau}}$  satisfies

$$\frac{1}{T}\sum_{\tau=0}^{\overline{\tau}-1}\eta_{\tau}q(x^{\tau}) \ge \gamma - \operatorname{poly}(\vec{G}, \vec{L}, G_h, L_h) \cdot \left(\frac{2d(X) + \log(\overline{\tau}) + G_q + L_q}{2\sqrt{\overline{\tau}}}\right).$$

We shall call a dual solution h a (generalised) Lyapunov functions, following the nomenclature in [37]. The insight is that not only such functions can describe worst- and best- case behaviour of approximate or exact gradient dynamics of the game as far as bounding the performance of  $\epsilon$ -stationary CCE is concerned, they are necessarily the form of dual solutions in the primal-dual framework one may consider constructing from such CCE. Moreover, whenever a dual solution certifies uniqueness of equilibrium it is necessarily a Lyapunov function in the traditional sense, as the following example demonstrates. **Example 1.** Here, we shall show that the CCE of finite normal-form games and an  $(0, \Phi_{\text{Proj}}(\delta))$ local CE (in the sense of Definition 2) are equivalent<sup>3</sup>. Inspection of a optimal dual proof of uniqueness of equilibrium will then take the form of a quadratic Lyapunov function.

Consider a finite normal-form game, with a set of players N, players' pure action sets  $(A_i)_{i \in N}$ , and utilities  $(U_i : \times_{i \in N} A_i \to \mathbb{R})_{i \in N}$ . The continuus extension of the game has action sets  $X_i = \Delta(A_i)$ , the probability simplex over  $A_i$ , and utilities  $u_i(x) = \sum_{a \in A} (\prod_{i \in N} x_i(a_i)) u_i(a)$  are given via expectations. Then, for an arbitrary  $\delta \in (0, 1)$ , a  $(0, \Phi_{\text{Proj}}(\delta))$ -local CE is a distribution  $\sigma$  which satisfies

$$\forall i \in N, \forall x_i^* \in X_i, \mathbb{E}_{x \sim \sigma}[u_i((1-\delta)x_i + \delta x_i^*, x_{-i}) - u_i(x)] \le 0.$$

Expanding the left-hand side, we have

$$\begin{split} \mathbb{E}_{x\sim\sigma}[u_i((1-\delta)x_i+\delta x_i^*,x_{-i})-u_i(x)] \\ &= \int_X d\sigma(x)\cdot\sum_{a\in A} \left(\prod_{j\in N\setminus\{i\}} x_j(a_j)\right)\cdot((1-\delta)x_i(a_i)+\delta x_i^*(a_i)-x_i(a_i))\cdot u_i(a) \\ &= \delta\cdot\sum_{a\in A} \left(\int_X d\sigma(x)\cdot(x_i^*(a_i)-x_i(a_i))\cdot\left(\prod_{j\in N\setminus\{i\}} x_j(a_j)\right)\right)\cdot u_i(a) \\ &= \delta\cdot\sum_{a\in A} \left(\int_X d\sigma(x)\cdot\left[\sum_{a_i'\in A_i} x^*(a_i')\cdot(\delta(a_i',a_i)-x_i(a_i))\right]\cdot\left(\prod_{j\in N\setminus\{i\}} x_j(a_j)\right)\right)\cdot u_i(a) \\ &= \delta\cdot\sum_{a_i'\in A_i} x^*(a_i')\cdot\sum_{a\in A} \left(\int_X d\sigma(x)\cdot(\delta(a_i',a_i)-x_i(a_i))\cdot\left(\prod_{j\in N\setminus\{i\}} x_j(a_j)\right)\right)\cdot u_i(a) \\ &= \delta\cdot\sum_{a_i'\in A_i} x^*(a_i')\cdot\sum_{a\in A} \sigma'(a)\cdot(u_i(a_i',a_{-i})-u_i(a)), \end{split}$$

where we write  $\delta(a'_i, a_i)$  for the Kronecker delta and define  $\sigma'(a) = \int_X d\sigma(x) \cdot \left(\prod_{j \in N} x_j(a_j)\right)$ . Also, to improve readability we denote  $\int_X d\mu(x) \cdot f(x)$  as  $\int_X f(x) d\mu(x)$ . Finally, the last equality follows since

$$\begin{split} &\sum_{a \in A} \int_X d\sigma(x) \cdot \delta(a'_i, a_i) \cdot u_i(a) \cdot \left(\prod_{j \in N \setminus \{i\}} x_j(a_j)\right) \\ &= \sum_{a_{-i} \in A_{-i}} \int_X d\sigma(x) \cdot u_i(a'_i, a_{-i}) \cdot \left(\prod_{j \in N \setminus \{i\}} x_j(a_j)\right) \\ &= \sum_{a_{-i} \in A_{-i}} \int_X d\sigma(x) \cdot u_i(a'_i, a_{-i}) \cdot \left(\prod_{j \in N \setminus \{i\}} x_j(a_j)\right) \cdot \underbrace{\left(\sum_{a_i \in A_i} x_i(a_i)\right)}_{=1} \\ &= \sum_{a \in A} \sigma'(a) \cdot u_i(a'_i, a_i). \end{split}$$

<sup>3</sup>This was remarked by Constatinos Daskalakis in a private correspondence.

Then as a consequence, for each player i and each  $x_i^* \in \Delta(A_i)$ , the associated local CCE constraint is simply a convex combination of constraints

$$\forall a_i' \in A_i, \sum_{a \in A} \sigma'(a) \cdot (u_i(a_i', a_{-i}) - u_i(a)) \le 0.$$

This is precisely the usual coarse correlated equilibrium constraint for the finite normal-form game.

Now we consider what a primal-dual proof of uniqueness of such a local CCE looks like. Suppose that the probability distribution which places probability 1 on  $x^* \in X$  is the unique local CCE of the game, then the measure valued primal problem

$$\max_{\sigma \ge 0} \int_X d\sigma(x) \cdot \|x - x^*\|^2 \text{ subject to}$$
(17)

$$\forall i \in N, \forall a'_i \in A_i, \sum_{a \in A} \sigma'(a) \cdot (u_i(a'_i, a_{-i}) - u_i(a)) \le 0 \qquad (d(i, a'_i))$$

$$\int_X d\sigma(x) = 1 \tag{(\omega)}$$

has value 0. First, note that by the convexity of the objective in x and the form of constraints  $d(i, a'_i)$ , we may assume that the optimal solution  $\sigma$  concentrates all probability on pure action profiles. But that implies that we simply have the standard LP over the set of CCE (in the usual sense) of the normal-form game,

$$\max_{\sigma' \ge 0} \sum_{a \in A} \sigma'(a) \cdot \sum_{i \in N} \|x_i(a_i) - x_i^*(a_i)\|^2 \text{ subject to}$$
(18)

$$\forall i \in N, \forall a'_i \in A_i, \sum_{a \in A} \sigma'(a) \cdot (u_i(a'_i, a_{-i}) - u_i(a)) \le 0 \qquad (d(i, a'_i))$$

$$\sum_{a \in A} \sigma'(a) = 1 \tag{(\omega)}$$

For this LP to have value 0,  $x_i^*(a_i)$  must necessarily be  $\{0, 1\}$ -valued; otherwise, the objective is strictly positive for the LP. So we have that  $x_i^*(a_i) = \delta(a_i^*, a_i)$  for each player *i*, and some action profile  $a^*$ . The associated dual LP is given, after scaling the primal objective by 1/2,

$$\min_{\omega \in \mathbb{R}, d \ge 0} \omega \text{ subject to}$$
(19)

$$\forall \ a \in A, \omega + \sum_{i \in N, a_i' \in A_i} d(i, a_i') \cdot (u_i(a_i', a_{-i}) - u_i(a)) \ge \sum_{i \in N} 1 - \delta(a_i, a_i^*). \tag{$\sigma'(a)$}$$

By [1] (Proposition 3.1), any tight dual solution  $d^*$  necessarily has  $d^*(i, a'_i) > 0 \Leftrightarrow a'_i = a^*_i$ , and this condition is in fact equivalent to a unique CCE of the finite normal-form game in pure strategies. Letting  $h_i(x_i|a'_i) = -\frac{1}{2}||x_i - e_{ia'_i}||^2$  and noting that

$$\sum_{a \in A} \sigma'(a) \cdot \left( u_i(a'_i, a_{-i}) - u_i(a) \right) = \int_X d\sigma(x) \cdot \left\langle \nabla_i h_i(x_i | a'_i), \nabla_i u_i(x) \right\rangle,$$

a tight solution to the dual problem of  $(17)^4$ ,

$$\min_{\omega \in \mathbb{R}, \mu \ge 0} \omega \text{ subject to}$$
(20)

$$\forall \ x \in X, \omega + \sum_{i \in N, a_i' \in A_i} \mu(i, a_i') \cdot \left\langle \nabla_i [-\|x_i - e_{ia_i'}\|^2 / 2], \nabla_i u_i(x) \right\rangle \ge \frac{1}{2} \|x - x^*\|^2, \qquad (\sigma'(a))$$

is given by  $\mu(i, a'_i) = d^*(i, a'_i)$  and  $\omega = 0$ . Finally, note that

$$\left\langle \nabla_i h_i(x_i|a_i'), \nabla_i u_i(x) \right\rangle \le \left\langle \nabla_i h_i(x_i|a_i'), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle,$$

since  $\nabla_i h_i(x_i|a'_i) \in \mathcal{TC}_{X_i}(x_i)$  and hence  $\left\langle \nabla_i h_i(x_i|a'_i), \mathbb{P}_{\mathcal{NC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle \leq 0$ . As a consequence, we have

$$\sum_{i \in N} d^*(i, a_i^*) \left\langle \nabla_i h_i(x_i | a_i^*), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle \ge \frac{1}{2} \|x - x^*\|^2 > 0 \ \forall \ x \in X, x \neq x^*, \text{ and}$$
$$\sum_{i \in N} d^*(i, a_i^*) \left\langle \nabla_i h_i(x_i | a_i^*), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i^*)}[\nabla_i u_i(x^*)] \right\rangle = 0,$$

where the first inequality is by dual optimality and the second inequality is because  $\nabla_i h_i(x_i^*|a_i^*) = 0$ for any  $i \in N$ .

Example 1 shows that, as  $\mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] = \frac{dx_i(t)}{dt}$  for the projected gradient dynamics of the smooth game,  $\sum_{i \in N} d^*(i, a_i^*) \cdot h_i(x_i|a_i^*)$  is necessarily a *Lyapunov function* in the usual sense.

**Proposition 3.7.** A finite normal-form game has a unique local CCE which assigns probability 1 to a mixed-strategy profile  $x^*$  if and only if  $x^*$  is a Nash equilibrium in pure strategies and the convergence of the game's projected gradient dynamics to  $x^*$  may be proven via a Lyapunov function of the form  $h(x) = \sum_{i \in N} -C_i \cdot ||x_i - x_i^*||^2$  for some constants  $C_i > 0$ .

This motivates the use of the notions of stationary and local CCEs as *refinements* of the classical notion of a CCE, and then employ the resulting incentive constraints in primal-dual arguments. We conclude the section with an example on how the primal-dual framework may be used to argue about performance bounds or qualitative exact or approximate dynamics for standard normal form games. The main takeaway is, to put it lightly, that constructing and checking the validity of dual solutions can be rather laborious.

**Example 2** (Matching Pennies). Consider the canonical example for a game in which the gradient dynamics necessarily cycle, matching pennies. Up to reparametrisation of the strategy space, there are two players, whose utility functions  $[-1, 1]^2 \to \mathbb{R}$  are given

$$u_1(x_1, x_2) = -x_1 x_2,$$
  
 $u_2(x_1, x_2) = x_1 x_2.$ 

The unconstrained projected gradient dynamics then satisfy, if  $r(0)^2 = x_1^2(0) + x_2^2(0) \le 1$ ,  $x_1(t) = r \cdot \cos(t + \phi)$ ,  $x_2(t) = r \cdot \sin(t + \phi)$  for some "phase angle"  $\phi$ . If r(0) > 1, however, projections

<sup>&</sup>lt;sup>4</sup>Recall, after scaling the objective by a factor 1/2.

at the edge of a box suppresses the radius of motion down to 1 by some time t' > 0, and for each  $t \ge t'$ ,  $(x_1(t), x_2(t)) = (\cos(t + \phi), \sin(t + \phi))$ . Therefore, any stationary probability distribution  $\sigma \in \Delta([-1, 1]^2)$  must be rotationally symmetric, with support on the disc  $D = \{(x_1, x_2) | x_1^2 + x_2^2 \le 1\}$ . The goal here is to present the form of generalised Lyapunov functions h which certify these bounds, at least approximately, such that they provide convergence guarantees also for approximate projected gradient dynamics.

For the bound on the radius, by the LP based arguments in this section, we seek a suitable function  $h: [-1,1]^2 \to \mathbb{R}$  and a small  $\delta > 0$  such that

$$\min_{x_1, x_2 \in [-1,1]} x_1^2 + x_2^2 - \underbrace{\left(\nabla_1 h(x) \mathbb{P}_{\mathcal{TC}_{[-1,1]}(x_1)}[-x_2] + \nabla_2 h(x) \mathbb{P}_{\mathcal{TC}_{[-1,1]}(x_2)}[x_1]\right)}_{\text{``dh}(x)/dt''} \le 1 + \delta.$$

Here, we abuse notation by considering time dependent quantities given suitable initial conditions for the projected gradient dynamics. Now, consider setting,

$$h(x) = \begin{cases} 0 & x_1^2 + x_2^2 \le 1\\ -\frac{M_1(r-1)^2}{1+M_1(r-1)} \cdot M_2 \cdot \arctan(x_1/x_2) & x_1^2 + x_2^2 > 1, x_1x_2 > 0\\ -\frac{M_1(r-1)^2}{1+M_1(r-1)} \cdot M_2 \cdot \arctan(-x_2/x_1) & x_1^2 + x_2^2 > 1, x_1x_2 \le 0 \end{cases}$$

for some constants  $M_1, M_2 > 0$  to be determined later. By construction, h is continuously differentiable with Lipschitz gradients. On the disc  $D, \nabla h(x) = (0,0)$ , whereas for any  $x_1^2 + x_2^2 > 1$ , by symmetry of the dynamics under rotations by  $\pi/2$ , it is sufficient to consider the case when  $x_1, x_2 > 0$ . In this case, if  $x_2 < 1$ , then no projections occur, and

$$\frac{dh(x)}{dt} = \frac{d}{dt} \left[ -\frac{M_1(r-1)^2}{1+M_1(r-1)} \cdot M_2 \cdot \arctan(x_1/x_2) \right] = \frac{M_1 M_2(r-1)^2}{1+M_1(r-1)},$$

in which case we have

$$r^{2} - dh(x)/dt = r^{2} - \frac{M_{2}M_{1}(r-1)^{2}}{1 + M_{1}(r-1)}$$

We would like to bound this quantity. Denote x = r - 1, and suppose that  $x \leq 1/M_1$ . Then

$$(1+x)^2 - \frac{M_1 M_2 x^2}{1+M_1 x} \le (1+x)^2 - \frac{M_1 M_2 x^2}{2} = \left(1 - \frac{M_1 M_2}{2}\right) x^2 + 2x + 1.$$

This bound is maximised, when  $M_1M_2/2 > 1$ , whenever  $x = 1/(-1 + M_1M_2/2)$ , with value  $1 + 2/(M_1M_2-2)$ . Meanwhile, if  $x > 1/M_1$ , then  $1 < M_1x$ , and

$$(1+x)^2 - \frac{M_1 M_2 x^2}{1+M_1 x} \le (1+x)^2 - \frac{M_1 M_2 x^2}{2M_1 x} = x^2 + \left(2 - \frac{M_2}{2}\right) x + 1$$

This function is convex, so it attains its maximum over the feasible interval for  $x \in [0, \sqrt{2} - 1]$ . When x = 0 the bound equals 1, whereas if  $x = \sqrt{2} - 1$ , whenever

$$(\sqrt{2}-1)^2 + (\sqrt{2}-1)\left(2 - \frac{M_2}{2}\right) + 1 \le 1 \Leftrightarrow M_2 \ge 2(\sqrt{2}+1)$$

the bound is  $\leq 1$ . As the bounds will generally improve in  $M_1$ , we will also fix  $M_2 = 10$  at this point for convenience.

If instead  $x_2 = 1$  then " $dx_2/dt'' = \mathbb{P}_{\mathcal{TC}_{[-1,1]}(x_2)}[x_1] = 0$ , and we have

$$\begin{split} h(x) &= -\frac{M_1 \left(\sqrt{1+x_1^2}-1\right)^2}{1+M_1(\sqrt{1+x_1^2}-1)} \cdot 10 \arctan(x_1),\\ \frac{dh(x)}{dt} &= -10M_1 \cdot \underbrace{\left(\frac{dx_1}{dt}\right)}_{=-1} \cdot \frac{d}{dx_1} \left[\frac{\arctan(x_1)(\sqrt{1+x_1^2}-1)^2}{1+M_1(\sqrt{1+x_1^2}-1)}\right]\\ &= \frac{10M_1(\sqrt{1+x_1^2}-1)}{\sqrt{1+x_1^2}(1+M_1(\sqrt{1+x_1^2}-1))}\\ &\quad \cdot \left[\frac{\sqrt{1+x_1^2}-1}{\sqrt{1+x_1^2}} + \frac{2\arctan(x_1)x_1+M_1x_1\arctan(x_1)(\sqrt{1+x_1^2}-1)}{1+M_1(\sqrt{1+x_1^2}-1)}\right]\\ &\geq \frac{10M_1(r-1)}{r(1+M_1(r-1))} \left[1-\frac{1}{r}\right]\\ &\geq \frac{5M_1(r-1)^2}{(1+M_1(r-1))}, \end{split}$$

where as  $r \leq \sqrt{2}$ ,  $1 - 1/r \geq (r - 1)/\sqrt{2}$ . To bound  $r^2 - dh(x)/dt$ , we again consider the cases when r - 1 < 1/M and  $r - 1 \geq 1/M$ . In the former case,

$$r^{2} - \frac{5M_{1}(r-1)^{2}}{(1+M_{1}(r-1))} \le r^{2} - \frac{5M_{1}(r-1)^{2}}{2},$$

which is maximised for  $r = 1 + \frac{2}{5M_1-2}$  with the same value. Meanwhile, if  $r - 1 \ge 1/M$ , then

$$r^{2} - \frac{5M_{1}(r-1)^{2}}{1+M_{1}(r-1)} \le r^{2} - \frac{5(r-1)}{2} = (1+x)^{2} - \frac{5x}{2},$$

defining  $x = r - 1 \in [0, \sqrt{2} - 1]$ . Again by the convexity of the expression, it is sufficient to check its values at its endpoints; setting x = 0 results in a bound of 1, while setting  $x = \sqrt{2} - 1$  results in a bound  $2 - 5(\sqrt{2} - 1)/2 < 1$ . As a consequence, h proves a bound of

$$\delta = \min\left\{\frac{2}{10M_1 - 2}, \frac{2}{5M_1 - 2}\right\} = \frac{2}{5M_1 - 2}$$

Now, note that the convergence bound for approximate projected gradient dynamics proven in Theorem 3.5 depend linearly in  $G_h$ . However,  $G_h$  admits a constant bound independent of choice of  $M_1$  in this setting. As a consequence, h can be chosen as an arbitrarily tight dual solution for the gradient dynamics of the matching pennies game by letting  $M_1 \to \infty$ , and approximate projected gradient dynamics after  $\overline{\tau}$  time steps and step sizes  $\eta_{\tau} \sim \frac{1}{\sqrt{\tau+1}}$  results in a distribution with expected square radius at most  $1 + O(\log(\overline{\tau})/\sqrt{\overline{\tau}})$ .

Finally, we would like to provide an approximate dual proof that any stable distribution is approximately rotationally invariant. Towards this end, consider a function  $q: [-1,1] \rightarrow \mathbb{R}$  which

is rotationally symmetric, i.e.  $q(x_1, x_2) = p(r) \cdot \sin(k \cdot \theta + \phi)$  for some function  $p : [0, \sqrt{2}] \to \mathbb{R}$ differentiable with p(0) = 0,  $\theta$  is the angle in polar coordinates corresponding to  $x_1$  and  $x_2$ ,  $k \in \mathbb{Z}$ is the associated frequency, and  $\phi \in \mathbb{R}$  is again some phase angle. For any rotationally symmetric stationary distribution  $\sigma \in \Delta([-1,1]^2)$ , via Fourier decomposition-based arguments, it must be the case that  $\mathbb{E}_{x \sim \sigma}[q(x)] = 0$ .

Towards this end, let  $\ell(x_1, x_2) = p(r) \cdot \cos(k\theta + \phi)/k$ . In the region without projections,  $d\ell(x)/dt = q(x)$ , and therefore  $\ell$  forms an exact Lyapunov function for q whenever projections are not needed. To handle the region where projections are needed, consider  $\ell + A \cdot h$  for the hpreviously defined for some choice of  $M_1 \leftarrow M > 0$  and some A > 0, and by the symmetry of the problem without loss of generality restrict attention to the case when  $x_2 = 1, x_1 > 0$ . In this case,

$$\frac{dr}{dt} = \frac{d}{dt}\sqrt{1+x_1^2} = \frac{-x_1}{r}, \frac{d\theta}{dt} = \frac{d}{dt}\left[\pi/2 - \arctan(x_1)\right] = \frac{1}{1+x_1^2}.$$

Therefore,

$$\frac{d(\ell + Ah)(x)}{dt} = \underbrace{-\frac{x_1 p'(r) \sin(k\theta + \phi)}{k\sqrt{1 + x_1^2}} + \frac{q(x)}{1 + x_1^2}}_{(*)} + A\frac{dh(x)}{dt}.$$

Note that by a similar treatment as the previous setting, via considering  $M \to \infty$  and picking A sufficiently large, we may acquire a family of dual solutions which prove convergence. In the interest of brevity, we shall argue about the asymptotic case. In the large M limit,  $dh(x)/dt \ge 5$  whenever  $1 \ge x_1 > 0, x_2 = 1$ , while (\*) is bounded on that set by  $Cx_1$  for some constant  $x_1$ . Therefore, a fixed choice of constant A is sufficient to provide a family of increasingly tighter dual solutions for bounds on the expected value of |q(x)|. In particular, we again conclude that approximate projected gradient dynamics, after  $\overline{\tau}$  time steps and step sizes  $\eta_{\tau} \sim \frac{1}{\sqrt{\tau+1}}$ , provides a probability distribution for which  $|q(x)| \le O(\log(\overline{\tau})/\sqrt{\overline{\tau}})$  in expectation.

### 4 On Local Correlated Equilibria

We now turn our attention to when local or stationary correlated equilibria are approximable. Unlike the case for local coarse correlated equilibria, here we do not expect a tractable "universal approximation scheme". Indeed, if F contains all vector fields which have Lipschitz modulus  $\leq L$ , for any smooth game F contains the vector field  $(L/L_i) \cdot (\nabla_i u_i)$  for each player i. As a consequence, any such  $\epsilon$ -local (or stationary) CE  $\sigma$  satisfies the inequalities

$$\mathbb{E}_{x \sim \sigma} \left\langle \mathbb{P}_{\mathcal{T}\mathcal{C}_{X_i}(x_i)}[\nabla_i u_i(x)], \nabla_i u_i(x) \right\rangle \leq \epsilon \cdot \operatorname{poly}(\vec{G}, \vec{L}, G, L) \ \forall \ i \in N.$$

For small enough  $\epsilon$ , the support of the approximate local CE contains an approximate local Nash equilibrium, for which no polynomial time algorithm is known.

However, by an analogue of the argument in Example 1,  $\epsilon$ -correlated equilibria for normal form games is equivalent to an  $\epsilon$ -local CE with respect to the set of gradient fields,

$$F = \{x_i(a_i) \cdot (e_{ia'_i} - e_{ia_i}) \mid i \in N, a_i, a'_i \in A_i\}.$$

Indeed, for any player *i* and any  $x \in \times_{i \in N} \Delta(A_i) \equiv X$ ,

$$\begin{aligned} \nabla_i u_i(x)_{a_i} &= \sum_{a_{-i} \in A_{-i}} \left( \prod_{j \neq i} x_j(a_j) \right) \cdot u_i(a_i, a_{-i}) \\ \Rightarrow x_i(a_i) \cdot \left\langle e_{ia'_i} - e_{ia_i}, \nabla_i u_i(x) \right\rangle &= x_i(a_i) \cdot \sum_{a_{-i} \in A_{-i}} \left( \prod_{j \neq i} x_j(a_j) \right) \cdot (u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i})) \\ &= x_i(a_i) \cdot \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i})]. \end{aligned}$$

Therefore, for  $\sigma \in \Delta(X)$ , defining  $\sigma'$  to be the probability distribution on A induced by  $\sigma$ ,

$$\begin{split} &\int_X d\sigma(x) \cdot x_i(a_i) \cdot \left\langle e_{ia'_i} - e_{ia_i}, \nabla_i u_i(x) \right\rangle \\ &= \int_X d\sigma(x) \cdot x_i(a_i) \cdot \sum_{a_{-i} \in A_{-i}} \left( \prod_{j \neq i} x_j(a_j) \right) \cdot \left( u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i}) \right) \\ &= \sum_{a_{-i} \in A_{-i}} \left( \int_X d\sigma(x) \cdot x_i(a_i) \cdot \prod_{j \neq i} x_j(a_j) \right) \cdot \left( u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i}) \right) \\ &= \sum_{a_{-i} \in A_{-i}} \sigma'(a_i, a_{-i}) \cdot \left( u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i}) \right). \end{split}$$

The final term is, of course, the left hand side of the usual (linear) correlated equilibrium constraints for a normal form game. As efficient algorithms for computing approximate correlated equilibria exist, we conclude that there might exist tractable algorithms to compute local or stationary CE for appropriately chosen families of vector fields F. This is indeed the case, independent of the (non)-convexity of the continuous game, as we will show in the following section.

### 4.1 Regret matching and approximability of $\epsilon$ -local and -stationary equilibria

Our promise of the approximability of local correlated equilibria in fact follows from standard regret matching algorithms, e.g. [71, 43, 39, 40]. Whenever the family of vector fields F has finitely many elements, we show that such algorithms are applicable to obtain  $O(1/\sqrt{T})$ -stationary correlated equilibria after T iterations, provided we have access to a fixed-point oracle for every linear combination over F. Such guarantees are also apply to  $O(1/\sqrt{T})$ -local correlated equilibria, provided each  $f \in F$  satisfies a tangency condition. Whenever we are also guaranteed that  $\|\sum_{f \in F} \mu_f f\|^2$  is a convex function for every  $\mu : F \to \mathbb{R}_+$ , a fixed-point may be computed as the solution of a convex problem. This setting subsumes usual correlated equilibria of normal form games as a special case, and we demonstrate by example that such algorithms may incorporate "rotational corrections". Altogether, this implies that there exists a non-trivial strengthening of correlated equilibria which is still tractably approximable.

For the proof of our bounds, we apply the framework of [40] which in turn follows from the analysis of [38]. In our approximation results, for simplicity we restrict attention to quadratic potential functions. Proof of convergence are then an immediate consequence of [40] (Theorem 11); we nevertheless provide the algorithms and the associated proofs for the sake of a complete

exposition. To wit, for a fixed smooth game and a family of vector fields F, recall that an  $\epsilon$ -stationary CE is a distribution  $\sigma$  on X such that

$$\forall f \in F, \left| \sum_{i \in N} \int_X d\sigma(x) \cdot \left\langle \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)], f_i(x) \right\rangle \right| \le \epsilon \cdot \operatorname{poly}(\vec{G}, \vec{L}, G_f, L_f),$$

where the poly( $\vec{G}, \vec{L}, G_f, L_f$ ) factor is fixed in advance. In our further analysis of  $\epsilon$ -stationary (and also  $\epsilon$ -local) CE, we shall fix attention to finite |F|, and fix the poly-factor to 1 – absorbing the bounds to  $\epsilon$ .

We consider the history of our algorithm to be a sequence of probability distributions  $(\sigma_t)_{t \in \mathbb{N}}$ on X. We then denote the **differential stationarity regret** with respect to  $f \in F$  at iteration t as

$$\mu_{tf} = \int_X d\sigma_t(x) \cdot \sum_{i \in N} \left\langle f_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle.$$

Our first regret matching algorithm then is given:

Algorithm 1: Regret matching for  $\epsilon$ -stationary CE Input: Smooth game  $(N, (X_i)_{i \in N}, (u_i)_{i \in N})$ , set of vector fields  $F, \sigma_1 \in \Delta(X), \epsilon$ Output:  $\sigma_t$ , an approximate stationary CE 1  $t \leftarrow 1$ ; 2  $\mu_{1f} \leftarrow \sum_{i \in N} \int_X d\sigma_t(x) \cdot \left\langle f_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle \forall f \in F$ ; 3 while  $\exists f \in F, |\mu_{tf}| > \epsilon \cdot \operatorname{poly}(\vec{G}, \vec{L}, G_f, L_f)$  do 4  $x_{t+1} \leftarrow a$  fixed point of  $\sum_{f \in F} \mu_{tf} f$  in X; 5 Find suitable  $\alpha_t \in (0, 1)$ ; 6  $\sigma_{t+1} \leftarrow (1 - \alpha_t) \cdot \sigma_t + \alpha_t \cdot \delta(x_{t+1});$ 7  $t \leftarrow t+1;$ 8  $|\mu_{tf} \leftarrow \sum_{i \in N} \int_X d\sigma_t(x) \cdot \left\langle f_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle \forall f \in F$ ; 9 end while

Here,  $\delta(x_{t+1})$  in line 6 denotes the point mass distribution on  $x_{t+1}$ . The convergence of Algorithm 1 then follows from usual arguments in the regret matching literature, adjusting for the fact that we are dealing with equality constraints for stationarity and that our action transformations are defined via infinitesimal translations about a vector field in a constrained set.

**Theorem 4.1.** Suppose that in Algorithm 1,  $\alpha_t = 1/(t+1)$  or is chosen optimally to minimize  $\sum_{f \in F} \mu_{tf}^2$  at each time step. Then after t iterations,

$$\max_{f \in F} |\mu_{tf}| \le \sqrt{\frac{|F|}{t+1}} \cdot \left(\sum_{i \in N} G_i \cdot \max_{f \in F} G_f\right).$$

As a consequence, an  $\epsilon$ -stationary CE may be computed after  $O(|F|/\epsilon^2)$  iterations, given a fixedpoint oracle for  $\sum_{f \in F} \mu_f f$  for any real vector  $\mu \in \mathbb{R}^F$ .

Proof. For  $\alpha_t = 1/(t+1)$ , we shall track the time evolution of  $t\mu_{tf}$ , i.e. the cumulative regret. In this case, note that for a potential  $G(\mu) = \|\mu\|^2$ , link function  $g(\mu) = 2\mu$  and error terms  $\gamma(\mu) = \|\mu\|^2$ ,  $(G, g, \gamma)$  is trivially a Gordon triple (cf. [40], Definition 5) as the condition

$$G(\mu + \Delta \mu) \le G(\mu) + \langle g(\mu), \Delta \mu \rangle + \gamma(\Delta \mu)$$

holds with equality for any  $\mu, \Delta \mu : F \to \mathbb{R}$ . Now, for our choice of  $\alpha_t$ ,

$$(t+1) \cdot \mu_{(t+1)f} = t \cdot \mu_{tf} + \sum_{i \in N} \left\langle f_i(x_{t+1}), \mathbb{P}_{\mathcal{TC}_{X_i}(x_{(t+1)i})}[\nabla_i u_i(x_{t+1})] \right\rangle.$$

Therefore,

$$\left\langle g(t \cdot \mu_{tf}), \sum_{i \in N} \langle f_i(x_{t+1}), \nabla_i u_i(x_{t+1}) \rangle \right\rangle$$

$$= 2t \cdot \sum_{i \in N} \left\langle \sum_{f \in F} \mu_{tf} f_i(x_{t+1}), \mathbb{P}_{\mathcal{TC}_{X_i}(x_{(t+1)i})} [\nabla_i u_i(x_{t+1})] \right\rangle$$

$$+ \sum_{f \in F} \left( \sum_{i \in N} \left\langle f_i(x_{t+1}), \mathbb{P}_{\mathcal{TC}_{X_i}(x_{(t+1)i})} [\nabla_i u_i(x_{t+1})] \right\rangle \right)^2$$

$$\leq |F| \cdot (\max_{f \in F} G_f \cdot \sum_{i \in N} G_i)^2.$$
(21)

Here, the bound follows from the bounds on the magnitudes of  $f, \nabla_i u_i$ , and since (21) is nonnegative. The latter is true since by the fixed point assumption,  $\sum_{f \in F} \mu_{tf} f(x_{t+1}) \in \mathcal{NC}_X(x_{t+1})$ , which is true if and only if  $\sum_{f \in F} \mu_{tf} f(x_{(t+1)_i}) \in \mathcal{NC}_{X_i}(x_{(t+1)_i})$  for every player  $i^5$ . As a consequence, by [40] (Theorem 6),

$$(t+1)^2 \sum_{f \in F} \mu_{(t+1)f}^2 \le \sum_{f \in F} \mu_{1f}^2 + t \cdot |F| \cdot (\max_{f \in F} G_f \cdot \sum_{i \in N} G_i)^2 \le (t+1) \cdot |F| (\max_{f \in F} G_f \cdot \sum_{i \in N} G_i)^2.$$

The result follows immediately.

**Remark.** We note that simply substituting the computed fixed points  $x_{t+1}$  with (approximately) stationary distributions  $\sigma'_{t+1} \in \Delta(X)$  with respect to the vector field  $\sum_{f \in F} \mu_{tf} f$  in Algorithm 1 is insufficient for our purposes in general. This is because the "reward system" in our setting formed by the utility gradients, which do not necessarily form a conservative vector field. In particular, even if  $\sigma'_{t+1}$  is induced as a measure via some closed parametrised loop  $c : [0, \ell] \to X$  with  $dc(u)/dt = \mathbb{P}_{\mathcal{TC}_X(c(u))}[\sum_{f \in F} \mu_{tf} f(c(u))]$  almost everywhere, it can be the case that

$$\sum_{f \in F} \mu_{tf} \cdot \sum_{i \in N} \int_X d\sigma'_{t+1}(x) \cdot \left\langle f_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle > 0.$$

For instance, this is true when  $\sum_{f \in F} \mu_{tf} f_i = \nabla_i u_i$  for each player *i*. It is in this sense that access to a fixed point oracle may be considered necessary.

One exception, of course, is if we are dealing with a potential game, wherein  $\nabla_i u_i = \nabla_i V$  for some smooth potential function  $V: X \to \mathbb{R}$ . In this case if X is one of the suitable convex bodies studied in Section 3, by the discussion therein for any stationary distribution  $\sigma$  computed via the

<sup>&</sup>lt;sup>5</sup>Recall that for any convex set X and any  $x \in X$ , if  $\mu$  is an element of the normal cone to X at x and  $\nu$  an element of the tangent cone, then  $\langle \mu, \nu \rangle \leq 0$ .

gradient dynamics for  $\sum_{f \in F} \mu_f F$  given some  $\mu : F \to \mathbb{R}$ ,

$$\sum_{f \in F} \mu_{tf} \cdot \sum_{i \in N} \int_{X} d\sigma(x) \cdot \left\langle f_{i}(x), \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\nabla_{i}u_{i}(x)] \right\rangle$$
$$\leq \sum_{i \in N} \int_{X} d\sigma(x) \cdot \left\langle \nabla_{i}V(x), \mathbb{P}_{\mathcal{TC}_{X_{i}}(x_{i})}[\sum_{f \in F} \mu_{tf}f_{i}(x)] \right\rangle \approx 0.$$

However, in this case a local correlated equilibrium may be computed by simply performing gradient ascent on V and finding an approximate local maximum. Therefore, it is unclear whether a regret matching algorithm would provide any advantages regarding tractable computation.

The approximation of an  $\epsilon$ -local CE given access to a fixed-point oracle follows near identically. One important point of note is that local CE involve tangent cone projections of the associated vector fields f. In general, it is not true that  $\mathbb{P}_{\mathcal{TC}_X(x)}[\sum_{f \in F} \mu_f f] = \sum_{f \in F} \mu_f \mathbb{P}_{\mathcal{TC}_X(x)}[f]$ , except when for every  $x \in X$  and  $f \in F$ ,  $f(x) \in \mathcal{TC}_X(x)$  in which case the equality is assured by the cone property. We shall call F **tangential** whenever this is the case, and denote the **differential local regret** with respect to  $f \in F$  at iteration t as

$$\mu_{tf} = \max\left\{\int_X d\sigma_t(x) \cdot \sum_{i \in N} \left\langle f_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle, 0\right\}.$$

Adapting Algorithm 1 to differential local regret then outputs an approximate local CE:

Algorithm 2: Regret matching for  $\epsilon$ -local CE

**Input:** Smooth game  $(N, (X_i)_{i \in N}, (u_i)_{i \in N})$ , tangential set of vector fields  $F, \sigma_1 \in \Delta(X), \epsilon$ **Output:**  $\sigma_t$ , an approximate local CE 1  $t \leftarrow 1$ ;

 $\begin{array}{l} \mathbf{1} \ t \leftarrow 1, \\ \mathbf{2} \ \mu_{1f} \leftarrow \sum_{i \in N} \int_{X} d\sigma_{t}(x) \cdot \langle f_{i}(x), \nabla_{i}u_{i}(x) \rangle \ \forall \ f \in F; \\ \mathbf{3} \ \mathbf{while} \ \exists \ f \in F, \mu_{tf} > \epsilon \cdot \operatorname{poly}(\vec{G}, \vec{L}, G_{f}, L_{f}) \ \mathbf{do} \\ \mathbf{4} \ | \ x_{t+1} \leftarrow a \ \text{fixed point of } \sum_{f \in F} \max\{\mu_{tf}, 0\}f \ \text{in } X; \\ \mathbf{5} \ | \ \text{Find suitable } \alpha_{t} \in (0, 1); \\ \mathbf{6} \ | \ \sigma_{t+1} \leftarrow (1 - \alpha_{t}) \cdot \sigma_{t} + \alpha_{t} \cdot \delta(x_{t+1}); \\ \mathbf{7} \ | \ t \leftarrow t + 1; \\ \mathbf{8} \ | \ \mu_{tf} \leftarrow \sum_{i \in N} \int_{X} d\sigma_{t}(x) \cdot \langle f_{i}(x), \nabla_{i}u_{i}(x) \rangle \ \forall \ f \in F; \\ \mathbf{9} \ \mathbf{end while} \end{array}$ 

**Theorem 4.2.** Suppose that in Algorithm 2,  $\alpha_t = 1/(t+1)$  or is chosen optimally to minimize  $\sum_{f \in F} \max\{\mu_{tf}, 0\}^2$  at each time step. Then after t iterations,

$$\max_{f \in F} \mu_{tf} \le \sqrt{\frac{|F|}{t+1}} \cdot \left(\sum_{i \in N} G_i \cdot \max_{f \in F} G_f\right).$$

As a consequence, an  $\epsilon$ -local CE may be computed after  $O(|F|/\epsilon^2)$  iterations, given a fixed-point oracle for  $\sum_{f \in F} \mu_f f$  for any non-negative vector  $\mu \in \mathbb{R}^F_+$ .

*Proof.* Identical to the proof of Theorem 4.1, by noting  $G(\mu) = \sum_{f \in F} \max\{\mu_f, 0\}^2$ ,  $g(\mu)_f = 2 \max\{\mu_f, 0\}$  and  $\gamma(\Delta \mu) = \|\Delta \mu\|^2$  form a Gordon triple ([40], Lemma 12), and that the tangentiality of F implies that at any fixed-point  $x_{t+1}, \sum_{f \in F} \max\{\mu_{tf}, 0\}f(x_{t+1}) = 0$ .

The question that remains is, then, whether there exists a family of vector fields F such that an  $\epsilon$ -stationary or -local CE are tractably approximable. The answer is easily shown to be affirmative for the case of  $\epsilon$ -local CE when each f is affine linear in each component.

**Proposition 4.3.** Suppose that F is a family of tangential vector fields, such that  $f(x) = P_f x + q_f$ for some matrix  $P_f \in \mathbb{R}^{d \times d}$  and some vector  $q_f \in \mathbb{R}^d$ . Then for any  $\mu : F \to \mathbb{R}_+$ ,  $\|\sum_{f \in F} \mu_f f(x)\|^2$ is a convex quadratic function. Moreover, any  $\arg \min_{x \in X} \|\sum_{f \in F} \mu_f f(x)\|^2$  is a fixed point of  $\sum_{f \in F} \mu_f f(x)$ .

Proof. As  $\sum_{f \in F} \mu_f f(x)$  is linear in x,  $\|\sum_{f \in F} \mu_f f(x)\|^2$  is a sum of squares. Since F is tangential, at any fixed point  $x^* \in X$  of  $\sum_{f \in F} \mu_f f$ ,  $\sum_{f \in F} \mu_f f(x^*) \in \mathcal{TC}_X(x^*) \cap \mathcal{NC}_X(x^*) = \{0\}$ .

**Corollary 4.4.** For a family of tangential affine linear vector fields F, a  $\epsilon$ -local CE with respect to F can be approximated via solving  $O(|F|/\epsilon^2)$  convex quadratic minimisation problems.

By the discussion at the beginning of Section 4, this subsumes the correlated equilibria of normal form games; but such equilibria remain tractably approximable even for non-concave games. As the following example demonstrates, for the hypercube equilibrium refinements are possible.

**Example 3.** Suppose in a two player game that  $X_i = [-1,1]$  for each player *i*, and hence  $X = [-1,1]^2$ . Then the usual correlated equilibrium constraints are provided by vector fields

$$f^{1+} = \begin{pmatrix} 1-x_1\\0 \end{pmatrix}, f^{1-} = \begin{pmatrix} -1-x_1\\0 \end{pmatrix}, f^{2+} = \begin{pmatrix} 0\\1-x_2 \end{pmatrix}, \text{ and } f^{2-} = \begin{pmatrix} 0\\-1-x_2 \end{pmatrix}$$

Note that these vector fields are all conservative, and as a result they are gradient fields. We therefore consider extending our set of vector fields by considering

$$g^{1-} = \begin{pmatrix} -x_1 - x_2 \\ 0 \end{pmatrix}, g^{1+} = \begin{pmatrix} x_2 - x_1 \\ 0 \end{pmatrix}, g^{2-} = \begin{pmatrix} 0 \\ -x_1 - x_2 \end{pmatrix}, \text{ and } g^{2+} = \begin{pmatrix} 0 \\ x_1 - x_2 \end{pmatrix}.$$

The vector fields g have non-zero curl, and thus none of them arise as the gradient field of a quadratic function; although, they are coordinate projections of suitable quadratic functions. As a consequence, no vector field  $g^{i\pm}$  may be expressed as a conical combination of the vector fields  $f^{i\pm}$ . Setting  $F = \{f^{ij}, g^{ij} \mid i \in \{1, 2\}, j \in \{+, -\}\}$  thus provides us a refinement of the usual correlated equilibria of  $2 \times 2$  normal form games. The refinement can be strict; consider the matching pennies game from Example 2 where  $\nabla_i u_i(x) = (-1)^i x_{-i}$ . Then

$$\sum_{i \in N} \left\langle g^{1-}(x) + g^{2+}(x), \nabla_i u_i(x) \right\rangle = x_1^2 + x_2^2,$$

which implies that the only local CE with respect to F is the unique Nash equilibrium at x = (0, 0).

We note that when we consider linear (instead of conical) combinations of vector fields  $f \in F$  in Example 3, we may obtain any arbitrary affine-linear vector field on [-1, 1];  $\{f^{i\pm}, g^{i\pm}\}$  is a linearly dependent set for each player *i*. As a consequence, for a 2 × 2 normal form game, any stationary CE with respect to F is necessarily the convex hull of its Nash equilibria.

### 4.2 Remarks on duality

We end this section with a short note on the associated primal-dual framework to local correlated equilibria, akin to Section 3.2. Given local correlated equilibria against a finite set of vector fields F, the analogue of (22) for stationary CE is given

$$\inf_{\sigma \ge 0} \int_X d\sigma(x) \cdot q(x) \text{ subject to}$$
(22)

$$\int_X d\sigma(x) = 1 \tag{(\gamma)}$$

$$\forall f \in F, \sum_{i \in N} \int_X d\sigma(x) \cdot \left\langle f_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle = 0, \qquad (\mu_f)$$

and its Lagrangian dual is simply

$$\sup_{\gamma,\mu} \gamma \text{ subject to} \tag{23}$$

$$\forall x \in X, \gamma + \sum_{f \in F, i \in N} \left\langle \mu_f f(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle \le q(x). \tag{\sigma(x)}$$

Here, an analogue of Theorem 3.6 is immediate. However, at present, we do not see a clear intuitive interpretation of a dual solution  $\mu: F \to \mathbb{R}$ . When we take  $q(x) = -\left\langle \nabla_i u_i(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle$ , the optimal value of (22) equals 0 if and only if the only stationary CE of the game are probability distributions over the set of its (first-order) Nash equilibria. Then if  $(0, \mu)$  is a solution of (23),

$$\sum_{f \in F, i \in N} \left\langle \mu_f f(x), \mathbb{P}_{\mathcal{TC}_{X_i}(x_i)}[\nabla_i u_i(x)] \right\rangle < 0 \ \forall \ x \in X, \exists \ i \in N, \nabla_i u_i(x) \notin \mathcal{NC}_{X_i}(x).$$

In words, the vector field  $-\sum_{f \in F} \mu_f f$  is suitably aligned with the utility gradient at every nonequilibrium point x, whether the gradient dynamics cycle or not. Moreover, fixed points of  $-\sum_{f \in F} \mu_f f$  and  $(\nabla_i u_i(x))_{i \in N}$  necessarily coincide. This does not necessarily provide an obvious computational advantage though, since whenever F is rich enough we may set  $-\sum_{f \in F} \mu_f f_i = \nabla_i u_i$ (cf. remark after Example 3).

## 5 Further Directions & Extensions

In this paper, we have established the question of computing an approximate local or stationary CCE, by identifying such outcomes as the natural property of the gradient dynamics of the underlying game and considering the weighted history of play when all players employ projected gradient ascent as their learning algorithm with equal learning rates. Appealing to the properties of these dynamics and extending the usual convex optimisation based primal-dual framework for price of anarchy / stability bounds, we were then able to argue that performance guarantees for the resulting distribution may be proven via constructing Lyapunov functions for the quantity in question. For the setting of finitely many vector fields, we have shown how regret matching may be performed for stationary or local CE, and discussed tractable settings. However, our results raise many questions yet unanswered.

### Approximability of local & stationary (C)CE

Some questions pertain directly to the computation of approximate local or stationary (C)CE. First, we have not established how  $(\epsilon, \Delta)$ -approximations are approximable. This is because our approach has been to consider the tangent and normal cones pointwise at a given  $x \in X$ , but it can be the case that for  $\delta > 0$ , the respective projections involving a step in direction  $f_i(x)$  or  $\nabla_i u_i(x)$  involve projections on constraints distinct each other, or the ones which bind at x. In fact, this holds true even for gradient ascent when  $x_i(t)$  is in the interior of  $X_i$ . How our appeals to the gradient dynamics of the game may be extended to cover such cases thus remains open. We remark that, unlike the approximability of  $\epsilon$ -local or stationary CCE, the associated approximation bounds of online projected gradient ascent would depend on the Lipschitz moduli of f in general.

Similarly, we do not yet know approximability of CCE in settings with general compact and convex action sets  $X_i$ , or even in the setting where all  $X_i$  are polyhedral. We have proven our results in settings where approximate projected gradient dynamics are faithful to the tangent cone projection of the direction of motion for the unconstrained gradient dynamics at each point in time, but [54] demonstrates that this need not be the case for polyhedral  $X_i$ . A deeper question, perhaps, is "What is the correct parameter of complexity for compact and convex  $X_i$  and approximability of local / stationary CCE via approximate projected gradient dynamics?". As previously remarked, when  $X_i$  has a smooth boundary, the approximation guarantee deteriorates linearly in K, the bound on the principal curvature of the boundary  $\delta X_i$ . However, when  $X_i$  are acute polyhedra – a condition intuitively in contradiction to  $\delta X_i$  having bounded curvature – projected gradient ascent models approximate projected gradient dynamics perfectly. It is thus yet unknown how to bound the approximation guarantees in the setting with general compact and convex action sets.

#### Generalisations of local / stationary CCE

Our local coarse notions of equilibria are also, in some sense, "bespoke" for projected gradient dynamics. We note that, while we have proven our results for projected gradient dynamics of the smooth game, we have never actually used the fact that the direction of motion is the gradient of some utility function for each player. Therefore, our results apply for any projected gradient dynamics of the form  $dx_i(t)/dt = \mathbb{P}_{\mathcal{TC}_{X_i}(x_i(t))}[F_i(x(t))]$  for each player *i*, with the resulting "equilibrium constraints" obtained by swapping any expression  $\nabla_i u_i(x(t))$  with  $F_i(x(t))$ . This implies a different notion of local or stationary CCE for each distinct time-independent gradient based learning algorithm; time-independent in the sense that no  $F_i$  has explicit time or history dependence, but are determined solely via x(t) at time *t*. How to simultaneously capture outcomes reached via such dynamics within a class of equilibria is thus open. Such an equilibrium concept would of course be a relaxation of equilibrium notions discussed in this paper.

We observe that mean-based learning algorithms [14] obtain a generalisation of various dual averaging based learning algorithms, and comparing the results of [33] and [31] we know that the outcomes of mean-based learning algorithms can be a strict subset of the (non-local) CCE of a game. However, to our knowledge projected gradient ascent is not known to be such an algorithm. Moreover, our Lyapunov function based primal-dual approach is ill-suited for dual averaging based methods in the first place. The associated continuous dynamics for dual averaging may be obtained [51] by setting action sets  $Y_i = \mathbb{R}^{d_i}$  for each player *i*, and letting

$$x_i(t) = \max_{x_i \in X_i} \langle x_i, y_i \rangle - r_i(x_i), \text{ and}$$
(24)

$$dy_i(t)/dt = \nabla_i u_i(x(t)), \tag{25}$$

where  $r_i : X \to \mathbb{R} \cup \{\infty\}$  is a lower semi-continuous and convex regulariser. In settings where  $r_i(x_i) = \infty$  if and only if  $x_i \in \delta X_i$  and  $r_i$  is differentiable in the interior of  $X_i$ , we are effectively faced with a smooth game where players' actions sets are  $d_i$  dimensional Euclidean spaces, and their equations of motion satisfy  $dx_i(t)/dt \to 0$  as  $||y_i(t)|| \to \infty$ . As a consequence, our approximability results break down completely, and how to extend our specific framework for price of anarchy / stability bounds to such settings at present unknown. The average price of anarchy approach of [60, 66] offers one solution, yet at present the performance metric does not come with convergence rate guarantees, even in expectation.

Similar presently unexplored directions for notions of equilibria relate to relaxing other assumptions. If we drop the previously mentioned time-independence assumption, we obtain settings where agents may learn at asymptotically different rates as a subclass; this question was investigated in [72] for monotone games. However, such games admit a Lyapunov function  $||x - x^*||^2$ which proves convergence to their unique equilibrium – a strong assumption that does not hold in many settings of interest. The time-independence assumption is also violated for any algorithm that uses history information, such that (a weighted) average history of play. Another assumption that may be dropped is deterministic motion, say for stochastic gradient ascent, in which case the continuous time dynamical system would be a Markov process. While Glynn and Zeevi [37] have already investigated Brownian motion, which can be taken as a model of SGD (cf. [45, 49]), it is unclear how the results within apply with respect to approximation guarantees in our constrained setting.

### On tractable local CE

Whereas we have demonstrated a generalisation of correlated equilibrium is tractable, it remains unknown how to procedurally generate a basis for such a family. Moreover, the affine-linearity of a vector field is a rather strong condition, and such vector fields in general fail to be tangential over a more general polytope. Finding wider settings in which fixed point computation for Algorithm 2 is tractable and the associated choice of F is meaningful is thus an open problem. Finally, we remark that our results are probably not necessarily optimal. Recent work [20, 2] shows that swap regret may be minimised in O(1/T) (up to polylog factors) iterations via more specialised algorithms, and we suspect similar methods might apply in our setting.

### Finding dual solutions

The final question pertains to how useful our primal-dual framework could be. Whereas we "have a primal-dual framework" to prove bounds on the performance of gradient ascent, Example 2 demonstrates that finding dual solutions and verifying their validity can be a daunting task. It is therefore of interest whether there exist systematic ways of constructing generalised Lyapunov functions to prove dual bounds, tight or approximate. One approach would be to consider relaxations to subclasses of differentiable functions with Lipschitz gradients for which the problem is tractable. As usual CCE of normal form games are tractable in the sense that performance bounds take the form of an LP, we infer by the discussion in Section 3.2 that there exists at least one such relaxation.

### Acknowledgements

Special thanks to Constantinos Daskalakis for the introduction to the question and discussions, and to Martin Bichler, Alexandros Hollender and Matthias Oberlecher for their feedback. This project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Grant No. BI 1057/9.

### References

- [1] Mete Şeref Ahunbay and Martin Bichler. On the uniqueness of Bayesian coarse correlated equilibria in standard first-price and all-pay auctions, 2024.
- [2] Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *Proceedings of the 54th Annual ACM SIGACT Symposium* on Theory of Computing, pages 736–749, 2022.
- [3] Edward J. Anderson and Peter Nash. *Linear Programming in Infinite-dimensional Spaces: Theory and Applications.* A Wiley-Interscience publication. Wiley, 1987.
- [4] Gabriel P. Andrade, Rafael Frongillo, and Georgios Piliouras. Learning in matrix games can be arbitrarily complex. In Mikhail Belkin and Samory Kpotufe, editors, Proceedings of Thirty Fourth Conference on Learning Theory, volume 134 of Proceedings of Machine Learning Research, pages 159–185. PMLR, 15-19 Aug 2021.
- [5] Gabriel P. Andrade, Rafael Frongillo, and Georgios Piliouras. No-regret learning in games is turing complete. arXiv preprint arXiv:2202.11871, 2022.
- [6] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. SIAM Journal on Computing, 38(4):1602–1623, 2008.
- [7] James P. Bailey and Georgios Piliouras. Multiplicative weights update in zero-sum games. In Proceedings of the 2018 ACM Conference on Economics and Computation, pages 321–338. ACM, 2018.
- [8] Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Perturbations of set-valued dynamical systems, with applications to game theory. *Dynamic Games and Applications*, 2:195–205, 2012.
- [9] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017.
- [10] Ulrich Berger. Fictitious play in  $2 \times n$  games. Journal of Economic Theory, 120(2):139–154, 2005.
- [11] Martin Bichler, Maximilian Fichtl, Stefan Heidekrüger, Nils Kohring, and Paul Sutterer. Learning equilibria in symmetric auction games using artificial neural networks. *Nature Ma-chine Intelligence*, 3:687–695, 2021.

- [12] Martin Bichler, Maximilian Fichtl, and Matthias Oberlechner. Computing Bayes-Nash equilibrium in auction games via gradient dynamics. Operations Research, to appear, 2023.
- [13] Vittorio Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory of Computing Systems*, 62:1288–1317, 2018.
- [14] Mark Braverman, Jieming Mao, Jon Schneider, and Matt Weinberg. Selling to a no-regret buyer. In Proceedings of the 2018 ACM Conference on Economics and Computation, pages 523–538, 2018.
- [15] George W Brown. Iterative solution of games by fictitious play. Activity analysis of production and allocation, 13(1):374–376, 1951.
- [16] Yang Cai, Constantinos Daskalakis, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. On local equilibrium in non-concave games, 2024.
- [17] Yang Cai and Christos Papadimitriou. Simultaneous Bayesian auctions and computational complexity. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, page 895–910, New York, NY, USA, 2014. Association for Computing Machinery.
- [18] Nicolo Cesa-Bianchi and Gábor Lugosi. Prediction, learning, and games. Cambridge university press, 2006.
- [19] Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, FOCS '06, page 261–272, USA, 2006. IEEE Computer Society.
- [20] Xi Chen and Binghui Peng. Hedging in games: Faster convergence of external and swap regrets. Advances in Neural Information Processing Systems, 33:18990–18999, 2020.
- [21] Xi Chen and Binghui Peng. Complexity of equilibria in first-price auctions under general tie-breaking rules. arXiv preprint arXiv:2303.16388, 2023.
- [22] Yun Kuen Cheung and Georgios Piliouras. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In Proceedings of Thirty Second Conference on Learning Theory, volume 99 of Proceedings of Machine Learning Research, pages 807–834. PMLR, 2021.
- [23] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. Journal of the ACM (JACM), 63(2):1–19, 2016.
- [24] George Christodoulou, Annamária Kovács, Alkmini Sgouritsa, and Bo Tang. Tight bounds for the price of anarchy of simultaneous first-price auctions. ACM Transactions on Economics and Computation (TEAC), 4(2):1–33, 2016.
- [25] Monica-Gabriela Cojocaru and Leo Jonker. Existence of solutions to projected differential equations in hilbert spaces. Proceedings of the American Mathematical Society, 132(1):183– 193, 2004.

- [26] Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. Games and Economic Behavior, 63(2):621–641, 2008. Second World Congress of the Game Theory Society.
- [27] Constantinos Daskalakis. On the complexity of approximating a Nash equilibrium. ACM Transactions on Algorithms (TALG), 9(3):1–35, 2013.
- [28] Constantinos Daskalakis. Non-concave games: A challenge for game theory's next 100 years. In *Cowles Preprints*, 2022.
- [29] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. SIAM Journal on Computing, 39(1):195–259, 2009.
- [30] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, pages 1466–1478, 2021.
- [31] Xiaotie Deng, Xinyan Hu, Tao Lin, and Weiqiang Zheng. Nash convergence of mean-based learning algorithms in first price auctions. In *Proceedings of the ACM Web Conference 2022*, pages 141–150, 2022.
- [32] Paul Dupuis and Anna Nagurney. Dynamical systems and variational inequalities. Annals of Operations Research, 44:7–42, 1993.
- [33] Michal Feldman, Brendan Lucier, and Noam Nisan. Correlated and coarse equilibria of singleitem auctions. In Yang Cai and Adrian Vetta, editors, Web and Internet Economics, pages 131–144, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [34] Zhe Feng, Guru Guruganesh, Christopher Liaw, Aranyak Mehta, and Abhishek Sethi. Convergence analysis of no-regret bidding algorithms in repeated auctions. volume 35, pages 5399–5406, 2021.
- [35] Miroslav Fiedler and Vlastimil Pták. On matrices with non-positive off-diagonal elements and positive principal minors. *Czechoslovak Mathematical Journal*, 12(3):382–400, 1962.
- [36] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. On the complexity of equilibrium computation in first-price auctions. In *Proceedings* of the 22nd ACM Conference on Economics and Computation, pages 454–476, 2021.
- [37] Peter W. Glynn and Assaf Zeevi. Bounding stationary expectations of markov processes. In Markov processes and related topics: a Festschrift for Thomas G. Kurtz, volume 4, pages 195–215. Institute of Mathematical Statistics, 2008.
- [38] Geoffrey J Gordon. No-regret algorithms for structured prediction problems. Carnegie Mellon University. Center for Automated Learning and Discovery, 2005.
- [39] Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of the 25th International Conference on Machine Learning*, pages 360–367, 2008.
- [40] Amy Greenwald, Zheng Li, and Casey Marks. Bounds for regret-matching algorithms. In AI&M, 2006.

- [41] Nadav Hallak, Panayotis Mertikopoulos, and Volkan Cevher. Regret minimization in stochastic non-convex learning via a proximal-gradient approach. In *International Conference on Machine Learning*, pages 4008–4017. PMLR, 2021.
- [42] John C. Harsanyi and Reinhard Selten, editors. A general theory of equilibrium selection in games. MIT Press, 1st edition, 1988.
- [43] Elad Hazan and Satyen Kale. Computational equivalence of fixed points and no regret algorithms, and convergence to equilibria. Advances in Neural Information Processing Systems, 20, 2007.
- [44] Elad Hazan, Karan Singh, and Cyril Zhang. Efficient regret minimization in non-convex games. In International Conference on Machine Learning, pages 1433–1441. PMLR, 2017.
- [45] Wenqing Hu, Chris Junchi Li, Lei Li, and Jian-Guo Liu. On the diffusion approximation of nonconvex stochastic gradient descent. Annals of Mathematical Sciences and Applications, 4(1), 2019.
- [46] Yaonan Jin and Pinyan Lu. Settling the efficiency of the first-price auction. SIGecom Exch, 20:69–74, 2022.
- [47] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. Computer science review, 3(2):65–69, 2009.
- [48] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, pages 869–886. SIAM, 2012.
- [49] Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *The 22nd international conference on artificial intelligence and statistics*, pages 983–992. PMLR, 2019.
- [50] Carlos Martin and Tuomas Sandholm. Finding mixed-strategy equilibria of continuous-action games without gradients using randomized policy networks. arXiv preprint arXiv:2211.15936, 2022.
- [51] Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Math. Program.*, 173(1–2):465–507, jan 2019.
- [52] Jason Milionis, Christos Papadimitriou, Georgios Piliouras, and Kelly Spendlove. Nash, Conley, and computation: Impossibility and incompleteness in game dynamics, 2022.
- [53] Dov Monderer and Lloyd S Shapley. Potential games. Games and economic behavior, 14(1):124–143, 1996.
- [54] Hassan Mortagy, Swati Gupta, and Sebastian Pokutta. Walking in the shadow: A new perspective on descent directions for constrained minimization. Advances in neural information processing systems, 33:12873–12883, 2020.

- [55] Uri Nadav and Tim Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *International Workshop on Internet and Network Economics*, pages 319–326. Springer, 2010.
- [56] Anna Nagurney and Ding Zhang. Projected dynamical systems and variational inequalities with applications, volume 2. Springer Science & Business Media, 2012.
- [57] John F. Nash. Equilibrium points in n-person games. Proceedings of the National Academy of Sciences, 36(1):48–49, 1950.
- [58] Kim Thang Nguyen. Primal-Dual Approaches in Online Algorithms, Algorithmic Game Theory and Online Learning. Habilitation à diriger des recherches, Université Paris Sorbonne, June 2019.
- [59] Gerasimos Palaiopanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. Advances in Neural Information Processing Systems, 30, 2017.
- [60] Ioannis Panageas and Georgios Piliouras. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 703–720, 2016.
- [61] Christos Papadimitriou and Georgios Piliouras. Game dynamics as the meaning of a game. SIGecom Exch., 16(2):53–63, May 2019.
- [62] Julia Robinson. An iterative method of solving a game. Annals of mathematics, pages 296–301, 1951.
- [63] J. Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. Econometrica: Journal of the Econometric Society, pages 520–534, 1965.
- [64] Tim Roughgarden. Selfish routing and the price of anarchy. MIT press, 2005.
- [65] Tim Roughgarden. Intrinsic robustness of the price of anarchy. Journal of the ACM (JACM), 62(5):1–42, 2015.
- [66] Iosif Sakos, Stefanos Leonardos, Stelios Andrew Stavroulakis, William Overman, Ioannis Panageas, and Georgios Piliouras. Beating price of anarchy and gradient descent without regret in potential games. In *The Twelfth International Conference on Learning Representations*, 2024.
- [67] William H Sandholm. Evolutionary game theory. Complex Social and Behavioral Systems: Game Theory and Agent-Based Models, pages 573–608, 2020.
- [68] Andreas S. Schulz and Nicolás E. Stier Moses. On the performance of user equilibria in traffic networks. In SODA, volume 3, pages 86–87, 2003.
- [69] Alexander Shapiro. On duality theory of conic linear problems. Nonconvex Optimization and its Applications, 57:135–155, 2001.
- [70] Lloyd Shapley. Some topics in two-person games. Advances in game theory, 52:1–29, 1964.

- [71] Gilles Stoltz and Gábor Lugosi. Learning correlated equilibria in games with compact sets of strategies. Games and Economic Behavior, 59(1):187–208, 2007.
- [72] Shaolin Tan, Ye Tao, Maopeng Ran, and Hao Liu. On the convergence of distributed projected gradient play with heterogeneous learning rates in monotone games. *Systems & Control Letters*, 182:105654, 2023.
- [73] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 416–425. IEEE, 2002.
- [74] H. Peyton Young. Strategic learning and its limits. Oxford University Press, 2004.
- [75] Zhengyuan Zhou, Panayotis Mertikopoulos, Susan Athey, Nicholas Bambos, Peter W Glynn, and Yinyu Ye. Learning in games with lossy feedback. Advances in Neural Information Processing Systems, 31, 2018.