
Independent Learning in Constrained Markov Potential Games

Philip Jordan
ETH Zürich

Anas Barakat
ETH Zürich

Niao He
ETH Zürich

Abstract

Constrained Markov games offer a formal mathematical framework for modeling multi-agent reinforcement learning problems where the behavior of the agents is subject to constraints. In this work, we focus on the recently introduced class of constrained Markov Potential Games. While centralized algorithms have been proposed for solving such constrained games, the design of converging independent learning algorithms tailored for the constrained setting remains an open question. We propose an independent policy gradient algorithm for learning approximate constrained Nash equilibria: Each agent observes their own actions and rewards, along with a shared state. Inspired by the optimization literature, our algorithm performs proximal-point-like updates augmented with a regularized constraint set. Each proximal step is solved inexactly using a stochastic switching gradient algorithm. Notably, our algorithm can be implemented independently without a centralized coordination mechanism requiring turn-based agent updates. Under some technical constraint qualification conditions, we establish convergence guarantees towards constrained approximate Nash equilibria. We perform simulations to illustrate our results.

1 INTRODUCTION

In multi-agent reinforcement learning (RL), several agents interact within a shared dynamic and uncertain environment evolving over time depending on the individual strategic decisions of all the agents. Each

agent aims to maximize their own individual reward which may however depend on all players¹ decisions. Besides reward maximization, agents may also contend with satisfying constraints that are often dictated by multi-agent RL applications. Prominent such real-world applications include multi-robot control on cooperative tasks (Gu et al., 2023) as well as autonomous driving (Shalev-Shwartz et al., 2016; Liu et al., 2023) where physical system constraints and safety considerations such as collision avoidance are of primary importance. In other applications, agents may be subject to soft constraints such as average users' total latency thresholds in wireless networks or average power constraints in signal transmission. Each agent seeks to maximize their reward while also accounting for constraints which are coupled among agents. Constrained Markov games (Altman and Shwartz, 2000) offer a mathematical framework to model multi-agent RL problems incorporating coupled constraints.

In this work, we focus on a particular class of structured constrained Markov games: constrained Markov Potential Games (CMPGs). Recently introduced in Alatur et al. (2023) to incorporate constraints, CMPGs naturally extend the class of Markov Potential Games (MPGs) that has been actively investigated in the last few years (Macua et al., 2018; Leonardos et al., 2022; Fox et al., 2022; Zhang et al., 2022b; Song et al., 2022; Ding et al., 2022; Zhang et al., 2022a; Maheshwari et al., 2023; Zhou et al., 2023). Interestingly, this class of games is a class of mixed cooperative/competitive Markov games including pure identical interest Markov games (in which all the reward and cost functions of the agents are identical) as a particular case. The ability to cooperate between learning agents is crucial to improve their joint welfare and achieve social welfare for artificial intelligence (see Dafoe et al. (2020, 2021) for an extensive discussion about the need for promoting cooperative AI).

Independent learning has recently attracted increasing attention thanks to its versatility as a learning protocol. We refer the reader to a recent nice survey on the topic (Ozdaglar et al., 2021). In this protocol,

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹We will use *player* and *agent* interchangeably.

agents can only observe the realized state and their own reward and action in each stage to individually optimize their return. In particular, each agent does not observe actions or policies from any other agent. This protocol offers several advantages including the following aspects: (a) Scaling: independent learning dynamics do not scale exponentially with the number of players in the game (also known as the curse of multi-agents); (b) Privacy protection: agents may avoid sharing their local data and information to protect their privacy and autonomy; (c) Communication cost: a central node that can bidirectionally communicate with all agents may not exist or may be too expensive to afford. Therefore, this protocol is particularly appealing in several applications where agents need to make decisions independently, in a decentralized manner. For example, dynamic load balancing, which consists in evenly assigning clients to servers in distributed computing, demands for learning algorithms that minimize communication overhead to enable low-latency response times and scalability across large data centers. This task has been modeled as an MPG (Yao and Ding, 2022). In other applications such as the pollution tax model and distributed energy marketplace detailed in section 5, coordination is inherently ruled out due to the competitive nature of the players’ interactions. Independent learning algorithms have been proposed for unconstrained multi-agent RL problems such as zero-sum Markov games (Daskalakis et al., 2020; Sayin et al., 2021; Chen et al., 2023) as well as for unconstrained MPGs in a recent line of works (Leonardos et al., 2022; Zhang et al., 2022a,b; Ding et al., 2022; Maheshwari et al., 2023).

However, for *constrained* MPGs, existing algorithms with convergence guarantees require coordination between players. Indeed, inspired by Song et al. (2022), Alatur et al. (2023) recently proposed a coordinate ascent algorithm for CMPGs in which each agent updates their policy in turn. At each time step, the policies of other agents are fixed while the updating agent faces a constrained Markov Decision Process (CMDP) to solve. When this coordination is not possible as in the independent learning protocol, the problem becomes more challenging as the environment is no longer stationary from the viewpoint of each agent and the problem does not reduce to solving a CMDP at each time step. This motivates the following question:

*Can we design an **independent** learning algorithm for **constrained** MPGs with non-asymptotic global convergence guarantees?*

In this paper, we answer this question in the affirmative. Our contributions are as follows:

- We design an algorithm for independent learning of constrained ϵ -approximate Nash equilibria (NE) in CMPGs. Inspired by recent works in nonconvex optimization under nonconvex constraints, our algorithm implements an inexact proximal-point update augmented with a regularized constraint set. In particular, the inexact proximal step is computed using a stochastic gradient switching algorithm for solving the resulting subproblem where both the objective and the constraint functions are strongly convex. Notably, the algorithm can be run independently by the different agents without taking turns.
- We analyze the proposed algorithm and establish its sample complexity to converge to an ϵ -approximate NE of the CMPG with polynomial dependence on problem parameters. Our analysis requires new technical developments that do not rely on results from the CMDP literature.
- We illustrate the performance of our algorithm on two simple CMPG applications: a pollution tax model and a marketplace for distributed energy resources.

Table 1: Position of our work in the literature. ‘centralized’ means that the algorithm requires coordination between players who take turns in updating their policy; for ‘independent’ learning, see section 2.

	centralized	independent
MPG	Nash-CA Song et al. (2022)	Independent PGA Leonardos et al. (2022) Zhang et al. (2022b) Ding et al. (2022)
CMPG	CA-CMPG Alatur et al. (2023)	Algorithm 1 <i>This work</i>

Related Works We refer the reader to Table 1 for a schematic positioning of our work in the recent literature. We next discuss some closely related work.

Markov Potential Games MPGs have been introduced as a natural extension of normal form potential games (Monderer and Shapley, 1996) to the dynamic setting starting with state-based potential games (Marden, 2012) and later Markov games (Macua et al., 2018). Leonardos et al. (2022) introduced a variant of MPGs and proposed independent stochastic policy gradient methods with an $\mathcal{O}(\epsilon^{-6})$

sample complexity to reach an ϵ -approximate NE. Similar results were shown in Zhang et al. (2022b) with model-based algorithms. This result was later improved to an $O(\epsilon^{-5})$ sample complexity for large state-action spaces with linear function approximation (Ding et al., 2022) and further to an $O(\epsilon^{-4.5})$ by reducing the variance of the agent-wise stochastic policy gradients (Mao et al., 2022). Zhang et al. (2022a) explored the use of the softmax policy parametrization instead of the direct parametrization. In particular, they established an $O(\epsilon^{-2})$ iteration complexity in the deterministic setting and showed the benefits of using regularization to improve the convergence rate. Maheshwari et al. (2023) proposed a fully independent and decentralized two timescale algorithm for MPGs with asymptotic guarantees where players may not even know the existence of other players. Narasimha et al. (2022) provided verifiable structural assumptions under which a Markov game is an MPG and further provided several algorithms for solving MPGs in the deterministic setting. Song et al. (2022) proposed an $O(\epsilon^{-3})$ sample complexity coordinate ascent algorithm (Nash-CA) which requires coordination between players. Guo et al. (2023) recently introduced the class of α -MPGs which relaxes the definition of MPGs by allowing α -deviations with respect to (w.r.t.) the potential function. More recently, Zhou et al. (2023) introduced a class of networked MPGs for which they proposed a localized actor-critic algorithm with linear function approximation. All the aforementioned works focused on the unconstrained setting.

Constrained Markov Games and CMPGs

There has been a vast array of works in multi-agent RL with safety constraints in practice (see e.g., ElSayed-Aly et al. (2021); Gu et al. (2023) and the references therein). Altman and Shwartz (2000) defined constrained Markov games and provided sufficient conditions for the existence of stationary constrained NE. Nonasymptotic theoretical convergence guarantees to game theoretic solution concepts for constrained multi-agent RL are relatively scarce in the literature. Chen et al. (2022) introduced a notion of correlated equilibria for general constrained Markov games and provided a primal-dual algorithm for learning those equilibria. Ding et al. (2023) established regret guarantees for episodic two-player zero-sum constrained Markov games. Alatur et al. (2023) introduced the class of constrained MPGs. Inspired by Nash-CA (Song et al., 2022), they proposed a constrained variant of the algorithm which enjoys an $O(\epsilon^{-5})$ sample complexity. Crucially, this algorithm requires coordination between agents and cannot be implemented independently by the agents.

Inexact Proximal-Point The idea of using inexact proximal-point methods to solve *nonconvex* problems has been fruitfully exploited in the literature for a couple of decades (see e.g., Hare and Sagastizábal (2009); Davis and Grimmer (2019)). A recent line of works (Boob et al. (2023); Ma et al. (2020); and also Jia and Grimmer (2023)) extended this idea in order to solve nonconvex optimization problems with nonconvex functional constraints. The initial nonconvex problem is transformed into a sequence of convex problems by adding quadratic regularization terms to *both* the objective and constraints. These works also established convergence rates to Karush–Kuhn–Tucker (KKT) points under constraint qualification conditions. Our present work is inspired by this recent line of research. We point out though that we deal with a multi-agent RL problem and we provide convergence guarantees to approximate constrained NE. In these regards, our independent algorithm design and our analysis require several new technical developments.

2 PRELIMINARIES

We consider an m -player constrained Markov Game where the players repeatedly select actions for maximizing their individual value functions while satisfying some constraints defined as cost value function bounds. More formally, the tabular game with random stopping, which we focus on, is described by a tuple $\mathcal{G} = (\mathcal{S}, \mathcal{N}, \{\mathcal{A}_i, r_i, c_i\}_{i \in \mathcal{N}}, \alpha, \mu, P, \kappa)$ with:

- A finite set of states \mathcal{S} of cardinality $S := |\mathcal{S}|$ and a finite set of m agents $\mathcal{N} := \{1, \dots, m\}$.
- A finite set of actions \mathcal{A}_i of cardinality $A_i := |\mathcal{A}_i|$ for all $i \in \mathcal{N}$ with $A_{\max} := \max_{i \in \mathcal{N}} A_i$. The joint action space is denoted by $\mathcal{A} := \prod_{i \in \mathcal{N}} \mathcal{A}_i$.
- A reward function $r_i : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ and a cost function $c_i : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ for each agent $i \in \mathcal{N}$. Throughout this paper, we will suppose that all the cost functions are identical across the agents and equal to a single cost function c .²
- A distribution μ over states from which the initial state of the game is drawn.
- A probability transition kernel P : For any state $s \in \mathcal{S}$ and any joint action $a \in \mathcal{A}$, the game transitions from state s to a state $s' \in \mathcal{S}$

²The case of multiple such common costs can be addressed with our approach with minor modifications. The case where cost functions may differ between players is more challenging and left for future work. See Remark 1 for details.

with probability $P(s'|s, a)$ and the game terminates with probability $\kappa_{s,a} > 0$. We further define $\kappa := \min_{s \in \mathcal{S}, a \in \mathcal{A}} \kappa_{s,a}$ and $\gamma := 1 - \kappa$.

At each time step $t \geq 0$ of a given episode of the game, all the agents observe a shared state $s_t \in \mathcal{S}$ and choose a joint action $a_t \in \mathcal{A}$. Then, each agent $i \in \mathcal{N}$ receives a reward $r_i(s_t, a_t)$ and incurs a cost $c(s_t, a_t)$. The game either stops at time t with probability κ_{s_t, a_t} or proceeds by transitioning to a state s_{t+1} drawn from the distribution $P(\cdot | s_t, a_t)$. We denote by T_e the random stopping time when the episode terminates.³ For a similar setting in the unconstrained case, see [Daskalakis et al. \(2020\)](#); [Giannou et al. \(2022\)](#).

Policies and Value Functions Each agent $i \in \mathcal{N}$ chooses their actions according to a randomized stationary policy denoted by $\pi_i \in \Pi^i := \Delta(\mathcal{A}_i)^{\mathcal{S}}$ where $\Delta(\mathcal{A}_i)$ is the probability simplex over the finite action space \mathcal{A}_i . The set of joint policies $\pi = (\pi_i)_{i \in \mathcal{N}}$ is denoted by $\Pi := \times_{i \in \mathcal{N}} \Pi^i$ and we further use the notation $\pi_{-i} = (\pi_j)_{j \in \mathcal{N} \setminus \{i\}} \in \Pi^{-i} := \times_{j \in \mathcal{N} \setminus \{i\}} \Pi^j$ for joint policies of all agents other than i . For any $u \in \{r_i \mid i \in \mathcal{N}\} \cup \{c\}$ and any joint policy $\pi \in \Pi$, we define the value function $V_u(\pi)$ for every state $s \in \mathcal{S}$ by $V_{u,s}(\pi) := \mathbb{E}[\sum_{t=0}^{T_e} u(s_t, a_t) | s_0 = s]$. The shorthand notation $V_u(\pi)$ will stand for $V_{u,\mu}(\pi) := \mathbb{E}_{s \sim \mu}[V_{u,s}(\pi)]$. For any policy $\pi \in \Pi$ and $s, s' \in \mathcal{S}$, the state visitation distribution is defined by $d_s^\pi(s') := \mathbb{E}[\sum_{t=0}^{T_e} \mathbf{1}_{\{s_t = s'\}} | s_0 = s]$ where $\mathbf{1}$ is the indicator function and we write $d_\mu^\pi(s') = \mathbb{E}_{s \sim \mu}[d_s^\pi(s')]$.

In the rest of this paper, we will aim to minimize both rewards and costs to align with conventions from the constrained optimization literature. The equivalence to the common RL reward maximization formulation follows from considering reward functions $1 - r_i$ instead of r_i for each $i \in \mathcal{N}$.

Constrained MPGs In this paper, we consider an m -player constrained MPG (CMPG) ([Alatur et al., 2023](#)) which is a constrained version of a Markov Potential Game ([Macua et al., 2018](#); [Leonardos et al., 2022](#)). In an MPG, for each state $s \in \mathcal{S}$, there exists a so-called potential function $\Phi_s : \Pi \rightarrow \mathbb{R}$ such that for all $i \in \mathcal{N}$, it holds that $V_{r_i, s}(\pi_i, \pi_{-i}) - V_{r_i, s}(\pi'_i, \pi_{-i}) = \Phi_s(\pi_i, \pi_{-i}) - \Phi_s(\pi'_i, \pi_{-i})$ for any policies $(\pi_i, \pi_{-i}) \in \Pi$, and $\pi'_i \in \Pi^i$. We will also use the notation $\Phi(\pi) := \mathbb{E}_{s \sim \mu}[\Phi_s(\pi)]$. Notice that the fully cooperative setting when all the reward functions of the players are identical is a particular instance of an MPG. Note also that the potential function is typically unknown for the players interacting in the game.

³The discounted infinite horizon setting can also be addressed with minor adaptations.

The joint policies of the agents are constrained to the set $\Pi_c := \{\pi \in \Pi \mid V_c(\pi) \leq \alpha\}$ of feasible policies. The set of feasible policies for agent $i \in \mathcal{N}$ when the policy of the other agents is fixed to $\pi_{-i} \in \Pi^{-i}$ is denoted by $\Pi_c^i(\pi_{-i}) := \{\pi_i \in \Pi^i \mid (\pi_i, \pi_{-i}) \in \Pi_c\}$.

Nash Equilibria For any $\epsilon \geq 0$, a joint policy $\pi^* \in \Pi_c$ is called an ϵ -approximate constrained NE if for every $i \in \mathcal{N}$ and any policy $\pi'_i \in \Pi_c^i(\pi_{-i}^*)$, we have $V_{r_i}(\pi^*) - V_{r_i}(\pi'_i, \pi_{-i}^*) \leq \epsilon$. When $\epsilon = 0$, such a policy π^* is called a constrained NE policy and no agent has an incentive to deviate unilaterally from a NE policy π^* . Observe that unilateral deviations are only allowed within the set of feasible policies in our constrained setting. We refer the reader to [Altman and Shwartz \(2000\)](#) for the existence of stationary constrained NEs.

Independent Learning Protocol All the players interact via executing their policies for a fixed number of episodes in order to find an approximate constrained NE. Importantly, during the learning procedure, each player executes their policy at each episode of the game to sample a trajectory and exclusively observes their own trajectory $(s_t, a_{i,t}, r_i(s_t, a_t), c(s_t, a_t))_{0 \leq t \leq T_e}$. In particular, a player does not have access to the policies of other players or their chosen actions. Such a protocol was considered for instance in two-player zero-sum Markov games in [Daskalakis et al. \(2020\)](#); [Chen et al. \(2023\)](#) as well as for unconstrained MPGs ([Leonardos et al., 2022](#); [Ding et al., 2022](#); [Maheshwari et al., 2023](#)).

3 INDEPENDENT ALGORITHM FOR CONSTRAINED MPGs

In this section, we present our independent iProxCMPG algorithm for learning constrained NE in CMPGs. Before describing our approach, we discuss an alternative, natural but unsuccessful, approach to motivate our algorithm design. This will allow us to highlight the challenges arising from the combination of (a) the presence of coupled constraints, (b) the multi-player setting, and (c) the independent learning protocol.

Our starting point is the known result that any maximizer of the potential function is a NE of the game. This result was initially proved by [Monderer and Shapley \(1996\)](#) for normal form potential games and later generalized to MPGs by [Leonardos et al. \(2022\)](#) and to constrained MPGs more recently ([Alatur et al., 2023](#)). Therefore, in order to find an (approximate) constrained NE for our CMPG⁴, we will consider solv-

⁴Approximate KKT points of this problem will be related to approximate constrained NE of our CMPG.

ing the following constrained optimization problem:

$$\min_{\pi \in \Pi_c} \Phi(\pi), \quad (1)$$

where Φ is the potential function for our CMPG using the notations introduced in section 2. This problem involves a nonconvex objective with a nonconvex constraint since the value function is a nonconvex function of the policy in general (see e.g., Lemma 1 in Agarwal et al. (2021)). However, although nonconvex optimization problems with nonconvex constraints are notoriously hard, it turns out that problem (1) is still tractable in the single agent setting. In this case, the problem boils down to a CMDP problem. Despite its nonconvexity, the problem can be recast as a linear program in the space of occupancy measures which is a convex set (see Chapter 3 in Altman (1999)). Then, strong duality permits to design primal-dual policy gradient algorithms to solve the problem with convergence guarantees (see e.g., Paternain et al. (2019)).

Given those positive results for single agent CMDPs, a natural approach is to derive a primal-dual algorithm for our multi-agent problem (1) as it was proposed by Diddigi et al. (2020). In the latter work, a primal-dual policy gradient algorithm was proposed using the Lagrangian function $\mathcal{L}(\pi, \lambda) := \Phi(\pi) + \lambda(V_c(\pi) - \alpha)$ where $\lambda \geq 0$ is a Lagrange multiplier. This algorithm can then be run independently by the different agents using existing independent learning algorithms for the unconstrained setting (Leonardos et al., 2022; Zhang et al., 2022b; Ding et al., 2022). Unfortunately, it has been recently shown by Alatur et al. (2023) that strong duality does not hold in general for the CMPG problem. As a consequence, it is not clear how to obtain guarantees for convergence to constrained NE using this duality approach. This is due to the multi-agent nature of our problem. In particular, since the constraint couples the agents’ individual policies, the set of state-action occupancy measures induced by joint policies of the players cannot be obviously split into several convex problems involving the occupancy measures induced by each one of the players’ policies. The well-known challenge of nonstationarity of the environment in multi-agent RL makes the design of independent learning algorithms difficult. As a remedy, Alatur et al. (2023) resort to coordination among players and propose a coordinate ascent algorithm for CMDPs. At each time step and for every player i , by fixing the policy of other players but player i to π_{-i} , player i can learn a “best-response” policy by solving a CMDP since the environment now becomes stationary from agent i ’s viewpoint.

Recall now that our main objective is to design an *independent* learning algorithm in the sense of section 2 in order to learn constrained NE for our CMPG.

We now describe our approach which takes a different route. Our algorithm is inspired by recent work in nonconvex optimization under nonconvex constraints (Boob et al., 2023; Ma et al., 2020; Jia and Grimmer, 2023). Following their ideas, we consider the following proximal update with penalized constraints:

$$\pi^{(t+1)} = \arg \min_{\pi \in \Pi} \left\{ \Phi(\pi) + \frac{1}{2\eta} \left\| \pi - \pi^{(t)} \right\|^2 \mid V_c(\pi) + \frac{1}{2\eta} \left\| \pi - \pi^{(t)} \right\|^2 + \beta \leq \alpha \right\} \quad (2)$$

where $\pi^{(0)}$ is a given initial joint policy, $\eta > 0$ is a step size and $\beta > 0$ an additional slack. Observe that $V_c(\pi^{(t+1)}) + \beta - \alpha \leq -\|\pi^{(t+1)} - \pi^{(t)}\|^2/2\eta$. Hence, the policy $\pi^{(t)}$ is feasible with slack β , i.e., $V_c(\pi^{(t)}) + \beta \leq \alpha$, for every $t \geq 0$. We introduce two additional notations for convenience. Define for any joint policies $\pi, \pi' \in \Pi$, and $\eta > 0$,

$$\begin{aligned} \Phi_{\eta, \pi'}(\pi) &:= \Phi(\pi) + \frac{1}{2\eta} \|\pi - \pi'\|^2, \\ V_{\eta, \pi'}^c(\pi) &:= V_c(\pi) + \frac{1}{2\eta} \|\pi - \pi'\|^2, \\ \Pi_{\eta, \pi'}^c &:= \left\{ \pi \in \Pi \mid V_{\eta, \pi'}^c(\pi) + \beta \leq \alpha \right\}. \end{aligned}$$

Our update rule in (2) can then be rewritten as:

$$\pi^{(t+1)} = \arg \min_{\pi \in \Pi_{\eta, \pi^{(t)}}^c} \Phi_{\eta, \pi^{(t)}}(\pi). \quad (3)$$

We immediately observe that the above update rule is well-defined since $\Phi_{\eta, \pi^{(t)}}$ and $V_{\eta, \pi^{(t)}}^c$ are strongly convex for every $t \geq 0$ for a suitable step size η . This is in contrast with the original problem where both the potential function Φ and the constraint function V_c are smooth but nonconvex. We also remark that if $\pi^{(t)}$ converges, then the regularization term $\|\pi^{(t+1)} - \pi^{(t)}\|$ becomes small and the surrogate feasible region $\Pi_{\eta, \pi^{(t)}}^c$ approaches the original constraint set Π_c up to the additional slack β .

Now, we discuss how to solve the proximal problem in (3) defining our main update rule. To solve this strongly convex problem with strongly convex constraint, we adapt a gradient switching algorithm proposed in Lan and Zhou (2020). At each iteration k , our algorithm performs a projected gradient descent step along either the gradient of the (regularized) objective or the gradient of the constraint function depending on whether an estimate of the constraint function satisfies the relaxed constraint $V_c(\pi^{(t,k)}) + \beta - \alpha \leq \delta_k$ where (δ_k) is a decreasing sequence converging to zero and hence progressively enforcing the constraint. However, it is not immediate from the above procedure how to obtain an independent learning algorithm specifying an update rule for each player without coordination

between the players. Recall for instance that the potential function Φ is unknown to the players in general and full gradients of both the potential and constraint functions w.r.t. the joint policy cannot be available to each agent since we exclude coordination and centralization. To obtain our independent iProxCMPG, see Algorithm 1, we propose to use agent-wise updates where each agent runs the gradient switching algorithm independently using only partial gradients of the potential and constraint functions w.r.t. their individual policy. Notice that our subroutine algorithm deviates from the one proposed in Lan and Zhou (2020) in that we use the estimate of the constraint function V_c instead of the regularized constraint function $V_{\eta, \pi^{(t)}}$. This is because the regularized constraint function involves the *joint policy* in the regularization while the constraint value function can be estimated independently. We further remark that for our analysis, the index \hat{k} sampled in line 7 of Algorithm 1 is supposed to be picked the same by all the players (see also Remark 3 in Appendix A for further details).

Stochastic Setting When exact gradients and value functions are not available, we estimate them using sampled trajectories. For each joint policy $\pi^{(t,k)}$, every player i samples a trajectory $\tau_i := (s_j^{(t,k)}, a_{i,j}^{(t,k)}, r_{i,j}^{(t,k)}, c_j^{(t,k)})_{0 \leq j \leq T_e}$ of length $T_e + 1$ by executing their own policy $\pi_i^{(t,k)}$. Here, $s_0^{(t,k)} \sim \mu$ and $r_{i,j}^{(t,k)}, c_j^{(t,k)}$ respectively refer to the reward and cost incurred by the i -th player at the j -th step. The gradients $\nabla_{\pi_i} V_{r_i}(\pi^{(t,k)})$ and $\nabla_{\pi_i} V_c(\pi^{(t,k)})$ are replaced by their sample estimates

$$\begin{aligned} \hat{\nabla} V_{\pi_i}^{r_i}(\pi^{(t,k)}) &:= R_i^{(T_e, t, k)} \psi_{\pi_i^{(t,k)}}^{T_e}, \\ \hat{\nabla} V_{\pi_i}^c(\pi^{(t,k)}) &:= C^{(T_e, t, k)} \psi_{\pi_i^{(t,k)}}^{T_e}, \end{aligned} \quad (4)$$

where $R_i^{(T_e, t, k)} := \sum_{j=0}^{T_e} r_{i,j}^{(t,k)}$, $C^{(T_e, t, k)} := \sum_{j=0}^{T_e} c_j^{(t,k)}$ and $\psi_{\pi_i^{(t,k)}}^{T_e} := \sum_{j=0}^{T_e} \nabla_{\pi_i} \log \pi_i^{(t,k)}(a_{i,j}^{(t,k)} | s_j^{(t,k)})$.

Each agent estimates $V_c(\pi^{(t,k)})$ by $\hat{V}_c(\pi^{(t,k)}) := C^{(T_e, t, k)}$ independently, using the cost feedback information they receive. Note that details of trajectory sampling are omitted in Algorithm 1 for more compact presentation (for the full version, see Algorithm 2 in Appendix A).

Remark 1. *As potential avenues for future work, we would like to point out two possible generalizations of the considered CMPG setting in which our current iProx-CMPG algorithm and analysis are not directly applicable:*

- **Potential cost constraints.** Suppose we do not require the cost functions c_i to be identical across all players $i \in \mathcal{N}$ but instead as-

sume that for each $s \in \mathcal{S}$, there exists a so-called cost potential function $\Phi_{c,s} : \Pi \rightarrow \mathbb{R}$ such that for all $i \in \mathcal{N}$, it holds that $V_{c_i, s}(\pi_i, \pi_{-i}) - V_{c_i, s}(\pi_i', \pi_{-i}) = \Phi_{c,s}(\pi_i, \pi_{-i}) - \Phi_{c,s}(\pi_i', \pi_{-i})$ for any policies $(\pi_i, \pi_{-i}) \in \Pi$, and $\pi_i' \in \Pi^i$.

Note that in order to use the gradient switching subroutine in our algorithm, it is essential that all agents are able to estimate whether or not the constraint holds for the current joint policy. In the case of a potential cost, it is not clear how to provide such estimates unless agents have knowledge of the potential (e.g. as a known function of the cost). This is an interesting question that merits further investigation.

- **Playerwise cost thresholds.** Suppose each player $i \in \mathcal{N}$ has an individual feasibility threshold α_i . The set of feasible policies is then redefined as $\Phi_c := \{\pi \in \Pi \mid \forall i \in \mathcal{N}, V_{c_i}(\pi) \leq \alpha_i\}$.

In this case, if all the cost functions c_i are identical and if the thresholds α_i can be communicated among players, one can consider the hardest constraint $\alpha := \min_{i \in \mathcal{N}} \alpha_i$ and use our approach to find a policy solving this stricter problem (if such policy exists). Otherwise, if cost functions are not identical or if each agent has their private threshold, our algorithm and analysis need further adjustments.

4 CONVERGENCE ANALYSIS AND SAMPLE COMPLEXITY

In this section, we establish the iteration complexity of Algorithm 1 in the deterministic setting before stating its sample complexity in the stochastic setting. We first introduce our assumptions. The first one guarantees the existence of a strictly feasible policy that is available to the agents for initialization.

Assumption 1. *The initial policy $\pi^{(0)}$ satisfies $V_c(\pi^{(0)}) < \alpha$.*

A few remarks are in order regarding this assumption:

- Similar assumptions have been made in the related constrained optimization literature when dealing with nonconvex constraints (Boob et al., 2023; Ma et al., 2020; Jia and Grimmer, 2023). Otherwise, satisfying a constraint may require finding a global minimizer which is computationally intractable in a general nonconvex setting. In our case, this corresponds to finding the global minimizer of a potential function in a fully cooperative unconstrained MPG. While this can be achieved in a single agent setting thanks to the gradient dominance property (Agarwal et al.,

Algorithm 1 iProxCMPG: independent Proximal-policy algorithm for CMPGs

-
- 1: **initialization:** $\pi^{(0)} \in \Pi^\xi$ s.t. $V_c(\pi^{(0)}) < \alpha$ and suitably chosen $\eta, \beta, T, K, \{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$
 - 2: **for** $t = 0, \dots, T - 1$ **do**
 - 3: $\pi_i^{(t,0)} = \pi_i^{(t)}$ for $i \in \mathcal{N}$
 - 4: **for** $k = 0, \dots, K - 1$ and $i \in \mathcal{N}$ **simultaneously do**
 - 5:
$$\pi_i^{(t,k+1)} = \begin{cases} \mathcal{P}_{\Pi^{i,\xi}} \left[\pi_i^{(t,k)} - \nu_k \hat{\nabla}_{\pi_i} V_{r_i}(\pi^{(t,k)}) - \frac{\nu_k}{\eta} (\pi_i^{(t,k)} - \pi_i^{(t)}) \right] & \text{if } \hat{V}_c(\pi^{(t,k)}) + \beta - \alpha \leq \delta_k \\ \mathcal{P}_{\Pi^{i,\xi}} \left[\pi_i^{(t,k)} - \nu_k \hat{\nabla}_{\pi_i} V_c(\pi^{(t,k)}) - \frac{\nu_k}{\eta} (\pi_i^{(t,k)} - \pi_i^{(t)}) \right] & \text{otherwise} \end{cases}$$
 - 6: $\mathcal{B}^{(t)} = \{\lfloor K/2 \rfloor \leq k \leq K \mid \hat{V}_c(\pi^{(t,k)}) \leq \delta_k\}$
 - 7: $\pi_i^{(t+1)} = \pi_i^{(t,\hat{k})}$ where $\hat{k} = 1$ if $\mathcal{B}^{(t)} = \emptyset$ and else sampled s.t. for $k \in \mathcal{B}^{(t)}$, $\mathbb{P}(\hat{k} = k) = (\sum_{k \in \mathcal{B}^{(t)}} \rho_k)^{-1} \rho_k$
 - 8: **output:** $\pi_i^{(T)}$ for $i \in \mathcal{N}$
-

2021; Xiao, 2022), such a global optimality result is not available in the literature for our multi-agent setting to the best of our knowledge.

- While finding a strictly feasible policy is involved in general, it may be possible to find such a policy in some special cases, such as when the state space can be factored, the probability transitions are independent across agents and the constraint cost functions are separable (see examples 1 and 2 in Alatur et al. (2023) for more details).

In addition to initial feasibility, we require that Slater's condition holds for each subproblem given by a proximal-point update. This is ensured by the following uniform Slater's condition.

Assumption 2. Let $\eta = \frac{1}{2L_\Phi}$ where L_Φ is the smoothness parameter⁵ of Φ . Then, there exists $\zeta > 0$ such that for any strictly feasible $\pi' \in \Pi$, i.e., $V_c(\pi') < \alpha$, there exists $\pi \in \Pi$ with $V_{\eta,\pi'}^c(\pi) \leq \alpha - \zeta$.

We make the following comments:

- First, we point out that a strictly feasible π' satisfies $V_{\eta,\pi'}^c(\pi') = V_c(\pi') < \alpha$, i.e., existence of a strictly feasible policy for the regularized constraint function $V_{\eta,\pi'}^c$ is trivially given. Assumption 2 additionally ensures that strict feasibility holds with slack ζ where ζ is independent of π' .
- Similar constraint qualification conditions have been widely used in the nonconvex constrained optimization literature, see Boob et al. (2022), Table 1 for an overview. In particular, Assumption 2 is similar to the uniform Slater's condition of Ma et al. (2020). Assumption 3 in Boob et al. (2023) is a strong feasibility assumption which implies Assumption 2, and hence could also replace it here. Strong feasibility assumes existence of a

⁵See also Lemma 8, item 3 for an expression of L_Φ in terms of m, γ , and A_{\max} .

policy π such that $V_c(\pi) \leq \alpha - \frac{\text{diam}(\Pi)^2}{\eta}$ where $\text{diam}(\Pi) := \max_{\pi, \pi' \in \Pi} \|\pi - \pi'\|$.

- A uniform strict feasibility assumption similar to Assumption 2 was used for centralized NE-learning, see Alatur et al. (2023), Assumption 2.

Exact Gradients Case In the noiseless setting with access to exact gradients, we achieve the following iteration complexity result.

Theorem 1. Let Assumptions 1 and 2 hold and let the distribution mismatch coefficient $D := \max_{\pi \in \Pi} \|d_\mu^\pi / \mu\|_\infty$ be finite. For any $\epsilon > 0$, after running iProxCMPG, Algorithm 1, for $\xi = 0$, suitably chosen η, β, T, K , and $\{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$, there exists $t \in [T]$, such that $\pi^{(t)}$ is a constrained ϵ -NE. The total iteration complexity is given by $\mathcal{O}(\epsilon^{-4})$ where $\mathcal{O}(\cdot)$ hides polynomial dependencies in $m, S, A_{\max}, D, 1 - \gamma$, and ζ .

The full proof of Theorem 1 is deferred to Appendix B.1. We briefly outline the key steps below.

Proof idea. First, we show that $K = \mathcal{O}(\epsilon^{-2})$ iterations of the inner loop yield a policy that is feasible and achieves potential value sufficiently close to the exact proximal update (3). For $T = \mathcal{O}(\epsilon^{-2})$, standard arguments then imply existence of $t \in [T]$ such that $\|\pi^{(t+1)} - \pi^{(t)}\| = \mathcal{O}(\epsilon)$. It can further be shown that such $\pi^{(t+1)}$ satisfies a particular form of approximate CMPG-specific KKT conditions for the original constrained optimization problem (1). We then leverage the multi-agent structure to argue that for all $i \in \mathcal{N}$, similar KKT conditions also hold w.r.t. the playerwise problem $\min_{\pi_i \in \Pi_i(\pi_{-i}^{(t+1)})} V_{r_i}(\pi_i, \pi_{-i}^{(t+1)})$

where $\pi_{-i}^{(t+1)}$ is fixed. Finally, using playerwise gradient dominance (see e.g., Lemma D.3 in Leonardos et al. (2022) or Lemma 2 in Giannou et al. (2022)), one can bound the duality gap of player i 's constrained problem for all $i \in \mathcal{N}$ which implies that $\pi^{(t+1)}$ is a con-

strained ϵ -NE. The total iteration complexity is given by $T \cdot K = \mathcal{O}(\epsilon^{-4})$. \square

Finite Sample Case In the stochastic setting, when exact gradients are not available, the variance of the stochastic policy gradients in (4) can be unbounded if the policies get closer to the boundaries of the simplex (see e.g., Eq. (13) in Giannou et al. (2022)). Therefore, we consider exploratory ξ -greedy policies to address this issue as in prior work (Daskalakis et al., 2020; Leonardos et al., 2022; Ding et al., 2022; Giannou et al., 2022). Define for any $\xi \geq 0, i \in \mathcal{N}$ the subset of ξ -greedy policies

$$\Pi^{i,\xi} := \{\pi \in \Pi \mid \forall s \in \mathcal{S} : \pi_i(\cdot \mid s) \geq \xi/A_i\},$$

which is used in Algorithm 1. We are now ready to state our sample complexity result.

Theorem 2. *Let Assumptions 1 and 2 hold, and let D (as in Theorem 1) be finite. Then, for any $\epsilon > 0$, after running iProxCMPG based on finite sample estimates (see Algorithm 2) for suitably chosen $\eta, \beta, \xi, T, K, B$, and $\{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$, there exists $t \in [T]$, such that in expectation, $\pi^{(t)}$ is a constrained ϵ -NE. The total sample complexity is given by $\tilde{\mathcal{O}}(\epsilon^{-7})$ where $\tilde{\mathcal{O}}(\cdot)$ hides polynomial dependencies in $m, S, A_{\max}, D, 1-\gamma$, and ζ , as well as logarithmic dependencies in $1/\epsilon$.*

We refer the reader to Appendix B.2 for the proof of Theorem 2. Below, we briefly explain how we obtain our sample complexity result.

Proof idea. As in the exact gradients case, we require $T = \mathcal{O}(\epsilon^{-2})$ iterations of the outer loop. In the stochastic setting, our independent implementation of the CSA algorithm (Lan and Zhou, 2020) still converges at a $\mathcal{O}(1/K)$ -rate due to strong convexity, but requires sampling a batch of size $B = \mathcal{O}(\epsilon^{-2})$ for estimating constraint function values at each iteration. To counteract the variance of ξ -greedy gradient estimates (which in our case grows as $\mathcal{O}(\epsilon^{-1})$), we need to set $K = \mathcal{O}(\epsilon^{-3})$. All in all, we end up with sample complexity $T \cdot K \cdot B = \mathcal{O}(\epsilon^{-7})$ for proving existence of $t \in [T]$ such that $\mathbb{E}[\|\pi^{(t)} - \pi^{(t+1)}\|] = \mathcal{O}(\epsilon)$. Using similar arguments as for Theorem 1, this implies that $\pi^{(t+1)}$ is a constrained ϵ -NE in expectation. \square

Remark 2. *Comparing our result to the state-of-the-art in the unconstrained case ($\mathcal{O}(\epsilon^{-5})$, Ding et al. (2022)), accounting for constraints comes at a cost, increasing the sample complexity by a $\mathcal{O}(\epsilon^{-2})$ -factor. In the centralized setting, a similar gap can be observed between best known results for unconstrained ($\mathcal{O}(\epsilon^{-3})$, Song et al. (2022)) vs. constrained ($\mathcal{O}(\epsilon^{-5})$, Alatur et al. (2023)) NE-learning. Whether this $\mathcal{O}(\epsilon^{-2})$ -gap can be narrowed is an interesting open question for both centralized and independent learning.*

5 SIMULATIONS

We test our stochastic iProxCMPG algorithm in two simple applications that can be modeled as CMPGs and for which coordination among players is unrealistic. Both examples are inspired by unconstrained variants presented in Narasimha et al. (2022) who study MPGs. Our code is publicly available⁶.

Pollution Tax Model Consider a simple environment with m agents representing e.g. factories, two states, *pollution-free* and *polluted*, and two actions, *clean* and *dirty* corresponding to low and high production volume. Starting in the *pollution-free* state, in each round, the environment transitions to the *polluted* state if and only if at least one agent chooses *dirty*. Each agent’s reward is the sum of its profit minus a pollution tax. In either state, the profit is P_c when choosing *clean* and P_d when choosing *dirty*. The pollution tax is zero in the *pollution-free*, and T_p in the *polluted* state. As pointed out by Narasimha et al. (2022), due to rewards being separable in the sense that $r_i(s, a_i, a_{-i}) = r'_i(s) + r''_i(a_i, a_{-i})$ and state transition probabilities being state independent, the pollution tax model satisfies a sufficient condition under which a Markov game is an MPG. For our simulations, we set $P_c = 2, P_d = 4$, and $T_p = 4$. Due to the lack of incentives for agents to cooperate when promoting environmental sustainability, requiring coordination is unrealistic in this example. Moreover, note that the purpose of the pollution tax is to counteract pollution by penalizing *dirty* actions. However, in practice, there may be additional global requirements on the minimum total production volume. To model this as a CMPG, we charge a cost C per agent that chooses *clean* and impose the constraint $V_c(\pi) \leq \alpha_C$ for appropriately chosen α_C .

We run iProxCMPG on the resulting m -agent CMPG for $m \in \{2, 4, 8\}$ and with $C = 1, \alpha_C = 12$. Hyperparameter choices are reported in Appendix E and Table 2. Fig. 1 shows the mean and standard deviation (shaded region) across independent runs of per-iteration potential and constraint values. Note that unlike in the theory part, we use the potential maximization perspective for experiments. We observe convergence to a constrained NE under which the minimum production requirements are approximately satisfied.

Marketplace for Distributed Energy Resources

As more and more small-scale electricity producers enter the electrical grids, a marketplace emerges. Each participant needs to decide how much energy to sell

⁶<https://github.com/philip-jordan/iProx-CMPG>

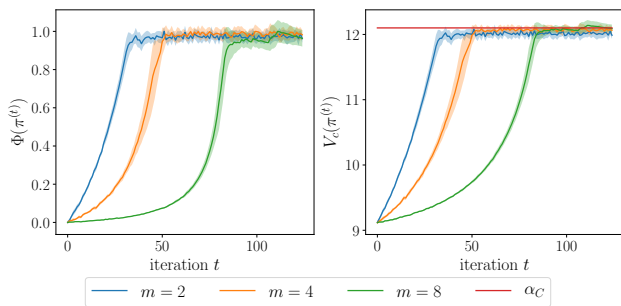


Figure 1: Potential (left, scaled to $[0, 1]$) and constraint (right) values of iProxCMPG for the m -agent pollution tax model.

given the current supply and demand. The competitive nature of such marketplaces motivates studying the convergence of independent algorithms to NEs under the constraints imposed by market rules. The CMPG we consider has states $\mathcal{S} = \{0, \dots, S-1\}$ indicating the grid’s current energy demand from high at 0 to low at $S-1$. Action $a_i \in \mathcal{A}_i = \{0, \dots, A_i-1\}$ represents the units of energy agent i contributes, for which it is rewarded with profit $r_i(s, a_i, a_{-i}) = c_0 a_i^2 - c_1 a_i^2 \sum_{i \in \mathcal{N}} a_i - a_i c_2^s$ where c_0, c_1, c_2 are model parameters. State transitions are modeled by first sampling $w \sim \mathcal{U}(\{0, 1, \dots, W\})$ which models uncertainty due to e.g. weather, and then setting $s' = \max\{0, \min\{S-1, \sum_{i \in \mathcal{N}} a_i - w\}\}$ with probability 0.9 and $s' = w$ otherwise. For our simulations, we set $S = A = W = 5, c_0 = 2, c_1 = 0.25$, and $c_2 = 1.25$. Narasimha et al. (2022) show that the described game is indeed an MPG with $\Phi(\pi) = \mathbb{E}_{\pi, s_0 \sim \mu} [\sum_{t=0}^{T_e} \phi_{s_t}(a_t)]$ and $\phi_s(a_i, a_{-i}) = c_0 \sum_{i \in \mathcal{N}} a_i - c_1 \sum_{i \in \mathcal{N}} a_i^2 - c_2 \sum_{1 \leq i < j \leq m} a_i a_j - m c_2^s$.




Figure 2: Potential (left) and constraint (right) values of iProxCMPG for the m -agent energy marketplace.

We extend this game into a CMPG by having the system incur a cost per unit of energy provided to the grid, i.e., by defining $c(s, a) = \sum_{i \in \mathcal{N}} a_i$ for all $s \in \mathcal{S}$, and requiring $V_c(\pi) \leq \alpha_e$ where we set $\alpha_e = 16$. Fig. 2 shows convergence to a constrained NE where players satisfy the energy provision bound on average.

6 CONCLUSION

In this paper, we proposed an independent learning algorithm for learning constrained NE in CMPGs. Our work opens up a number of avenues for future work. It would be interesting to investigate whether our sample complexity can be improved to match the better sample complexity of centralized algorithms. Our algorithm and theoretical guarantees require the agents to run the same algorithm: This may be seen as implicit coordination between agents. Designing fully independent learning dynamics for our constrained setting, where the players may not even be aware of the existence of other players is an interesting direction. Going beyond the class of CMPGs for learning constrained NE is another research direction that is worth exploring. Using function approximation to scale to large state-action spaces beyond the tabular setting is also a promising prospect for future work.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable comments. The work is supported by ETH research grant, Swiss National Science Foundation (SNSF) Project Funding No. 200021-207343, and SNSF Starting Grant. A.B. acknowledges support from the ETH Foundations of Data Science (ETH-FDS) postdoctoral fellowship.

References

- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. (2021). On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76. 5, 6, 20, 22, 23
- Alatur, P., Ramponi, G., He, N., and Krause, A. (2023). Provably learning nash policies in constrained markov potential games. In *Sixteenth European Workshop on Reinforcement Learning*. 1, 2, 3, 4, 5, 7, 8
- Altman, E. (1999). *Constrained Markov decision processes*, volume 7. CRC press. 5
- Altman, E. and Shwartz, A. (2000). Constrained Markov Games: Nash Equilibria. In Filar, J. A., Gaitsgory, V., and Mizukami, K., editors, *Advances in Dynamic Games and Applications*, Annals of

- the International Society of Dynamic Games, pages 213–221, Boston, MA. Birkhäuser. 1, 3, 4
- Boob, D., Deng, Q., and Lan, G. (2022). Level constrained first order methods for function constrained optimization. *arXiv preprint arXiv:2205.08011*. 7
- Boob, D., Deng, Q., and Lan, G. (2023). Stochastic first-order methods for convex and nonconvex functional constrained optimization. *Mathematical Programming*, 197(1):215–279. 3, 5, 6, 7, 17
- Chen, Z., Ma, S., and Zhou, Y. (2022). Finding correlated equilibrium of constrained markov game: A primal-dual approach. In *Advances in Neural Information Processing Systems*. 3
- Chen, Z., Zhang, K., Mazumdar, E., Ozdaglar, A. E., and Wierman, A. (2023). A finite-sample analysis of payoff-based independent learning in zero-sum stochastic games. In *Thirty-seventh Conference on Neural Information Processing Systems*. 2, 4
- Dafoe, A., Bachrach, Y., Hadfield, G., Horvitz, E., Larson, K., and Graepel, T. (2021). Cooperative ai: machines must learn to find common ground. *Nature*, 593(7857):33–36. 1
- Dafoe, A., Hughes, E., Bachrach, Y., Collins, T., McKee, K. R., Leibo, J. Z., Larson, K., and Graepel, T. (2020). Open problems in cooperative ai. *arXiv preprint arXiv:2012.08630*. 1
- Daskalakis, C., Foster, D. J., and Golowich, N. (2020). Independent policy gradient methods for competitive reinforcement learning. *Advances in neural information processing systems*, 33:5527–5540. 2, 4, 8, 24
- Davis, D. and Grimmer, B. (2019). Proximally guided stochastic subgradient method for nonsmooth, nonconvex problems. *SIAM Journal on Optimization*, 29(3):1908–1930. 3
- Diddigi, R. B., Danda, S. K. R., J., P. K., and Bhatnagar, S. (2020). Actor-Critic Algorithms for Constrained Multi-agent Reinforcement Learning. *arXiv:1905.02907 [cs]*. 5
- Ding, D., Wei, C.-Y., Zhang, K., and Jovanovic, M. (2022). Independent Policy Gradient for Large-Scale Markov Potential Games: Sharper Rates, Function Approximation, and Game-Agnostic Convergence. In *Proceedings of the 39th International Conference on Machine Learning*, pages 5166–5220. PMLR. ISSN: 2640-3498. 1, 2, 3, 4, 5, 8
- Ding, D., Wei, X., Yang, Z., Wang, Z., and Jovanovic, M. (2023). Provably efficient generalized lagrangian policy optimization for safe multi-agent reinforcement learning. In *Learning for Dynamics and Control Conference*, pages 315–332. PMLR. 3
- ElSayed-Aly, I., Bharadwaj, S., Amato, C., Ehlers, R., Topcu, U., and Feng, L. (2021). Safe multi-agent reinforcement learning via shielding. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '21*, page 483–491, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems. 3
- Fox, R., McAleer, S. M., Overman, W., and Panageas, I. (2022). Independent natural policy gradient always converges in markov potential games. In *International Conference on Artificial Intelligence and Statistics*, pages 4414–4425. PMLR. 1
- Giannou, A., Lotidis, K., Mertikopoulos, P., and Vlatakis-Gkaragkounis, E.-V. (2022). On the convergence of policy gradient methods to nash equilibria in general stochastic games. *Advances in Neural Information Processing Systems*, 35:7128–7141. 4, 7, 8
- Gu, S., Grudzien Kuba, J., Chen, Y., Du, Y., Yang, L., Knoll, A., and Yang, Y. (2023). Safe multi-agent reinforcement learning for multi-robot control. *Artificial Intelligence*, 319:103905. 1, 3
- Guo, X., Li, X., Maheshwari, C., Sastry, S., and Wu, M. (2023). Markov α -potential games: Equilibrium approximation and regret analysis. *arXiv preprint arXiv:2305.12553*. 3
- Hare, W. and Sagastizábal, C. (2009). Computing proximal points of nonconvex functions. *Mathematical Programming*, 116(1-2):221–258. 3
- Jia, Z. and Grimmer, B. (2023). First-Order Methods for Nonsmooth Nonconvex Functional Constrained Optimization with or without Slater Points. *arXiv:2212.00927 [math]*. 3, 5, 6, 17
- Lan, G. and Zhou, Z. (2020). Algorithms for stochastic optimization with function or expectation constraints. *Computational Optimization and Applications*, 76(2):461–498. 5, 6, 8, 25, 26, 27
- Leonardos, S., Overman, W., Panageas, I., and Piliouras, G. (2022). Global convergence of multi-agent policy gradient in markov potential games. In *International Conference on Learning Representations*. 1, 2, 4, 5, 7, 8, 15, 23, 24
- Liu, M., Kolmanovsky, I., Tseng, H. E., Huang, S., Filev, D., and Girard, A. (2023). Potential game-based decision-making for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*. 1
- Ma, R., Lin, Q., and Yang, T. (2020). Quadratically regularized subgradient methods for weakly convex optimization with weakly convex constraints. pages 6554–6564. PMLR. 3, 5, 6, 7, 18

- Macua, S. V., Zazo, J., and Zazo, S. (2018). Learning parametric closed-loop policies for markov potential games. In *International Conference on Learning Representations*. 1, 2, 4
- Maheshwari, C., Wu, M., Pai, D., and Sastry, S. (2023). Independent and Decentralized Learning in Markov Potential Games. arXiv:2205.14590 [cs, eess]. 1, 2, 3, 4
- Mao, W., Yang, L., Zhang, K., and Basar, T. (2022). On improving model-free algorithms for decentralized multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 15007–15049. PMLR. 3
- Marden, J. R. (2012). State based potential games. *Automatica*, 48(12):3075–3088. 2
- Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and economic behavior*, 14(1):124–143. 2, 4
- Narasimha, D., Lee, K., Kalathil, D., and Shakkottai, S. (2022). Multi-agent learning via markov potential games in marketplaces for distributed energy resources. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 6350–6357. IEEE. 3, 8, 9
- Ozdaglar, A., Sayin, M. O., and Zhang, K. (2021). Independent learning in stochastic games. *Invited chapter for the International Congress of Mathematicians 2022 (ICM 2022)*, arXiv preprint arXiv:2111.11743. 1
- Paternain, S., Chamon, L., Calvo-Fullana, M., and Ribeiro, A. (2019). Constrained reinforcement learning has zero duality gap. *Advances in Neural Information Processing Systems*, 32. 5
- Polyak, B. T. (1967). A general method for solving extremal problems. In *Doklady Akademii Nauk*, volume 174, pages 33–36. Russian Academy of Sciences. 25
- Sayin, M., Zhang, K., Leslie, D., Basar, T., and Ozdaglar, A. (2021). Decentralized q-learning in zero-sum markov games. *Advances in Neural Information Processing Systems*, 34:18320–18334. 2
- Shalev-Shwartz, S., Shammah, S., and Shashua, A. (2016). Safe, multi-agent, reinforcement learning for autonomous driving. arXiv preprint arXiv:1610.03295. 1
- Song, Z., Mei, S., and Bai, Y. (2022). When can we learn general-sum markov games with a large number of players sample-efficiently? In *International Conference on Learning Representations*. 1, 2, 3, 8
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press. 25
- Xiao, L. (2022). On the convergence rates of policy gradient methods. *Journal of Machine Learning Research*, 23(282):1–36. 7
- Yao, Z. and Ding, Z. (2022). Learning distributed and fair policies for network load balancing as markov potential game. *Advances in Neural Information Processing Systems*, 35:28815–28828. 2
- Zhang, R., Mei, J., Dai, B., Schuurmans, D., and Li, N. (2022a). On the global convergence rates of decentralized softmax gradient play in markov potential games. *Advances in Neural Information Processing Systems*, 35:1923–1935. 1, 2, 3, 23
- Zhang, R. C., Ren, Z., and Li, N. (2022b). Gradient play in stochastic games: Stationary points and local geometry. *IFAC-PapersOnLine*, 55(30):73–78. 25th International Symposium on Mathematical Theory of Networks and Systems MTNS 2022. 1, 2, 3, 5
- Zhou, Z., Chen, Z., Lin, Y., and Wierman, A. (2023). Convergence rates for localized actor-critic in networked Markov potential games. In Evans, R. J. and Shpitser, I., editors, *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pages 2563–2573. PMLR. 1, 3

Checklist

- For all models and algorithms presented, check if you include:
 - A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
- For any theoretical claim, check if you include:
 - Statements of the full set of assumptions of all theoretical results. [Yes]
 - Complete proofs of all theoretical results. [Yes]
 - Clear explanations of any assumptions. [Yes]
- For all figures and tables that present empirical results, check if you include:
 - The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]

- (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Supplementary Materials

CONTENTS

A iProxCMPG: FULL STOCHASTIC ALGORITHM	13
B PROOFS FOR SECTION 4	13
B.1 Proof of Theorem 1 — Exact Gradients Case	13
B.2 Proof of Theorem 2 — Finite Sample Case	20
B.3 Other Technical Lemmas	23
C STRONGLY CONVEX STOCHASTIC OPTIMIZATION WITH STRONGLY CONVEX EXPECTATION CONSTRAINT	25
C.1 A Primal Gradient Switching Algorithm	26
C.2 Convergence and Sample Complexity Guarantee	27
D BACKGROUND IN CONSTRAINED OPTIMIZATION AND A NOVEL TECHNICAL LEMMA	30
D.1 Approximate KKT Points in Constrained Optimization	30
D.2 A Novel Technical Lemma for Approximate Optimality under Gradient Dominance	31
E ADDITIONAL DETAILS ABOUT SIMULATIONS	32

A iProxCMPG: FULL STOCHASTIC ALGORITHM

In this section, for the convenience of the reader, we report the full pseudo-code of Algorithm 1 in the stochastic setting where exact gradients are not available. See Algorithm 2.

Remark 3. *For our analysis, the index \hat{k} sampled in line 11 of Algorithm 2 is supposed to be picked the same by all the players. This sampling step does not require any state-action-reward samples, it is just an index sampling that is useful for our proofs (even in the exact gradients case), see the constraint satisfaction inequality in the proof of Theorem 3, page 28.*

B PROOFS FOR SECTION 4

Notation For any integer $n \geq 1$, we use the notation $[n] := \{1, \dots, n\}$ throughout the proofs.

In this section, we provide complete proofs of our main results. We begin with the exact gradients case before addressing the more involved finite sample case.

B.1 Proof of Theorem 1 — Exact Gradients Case

First, we restate Theorem 1.

Algorithm 2 iProxCMPG: independent Proximal-policy algorithm for CMPGs

- 1: **initialization:** $\pi^{(0)} \in \Pi^\xi$ s.t. $V_c(\pi^{(0)}) < \alpha$ and suitably chosen $\eta, \beta, \xi, T, K, \{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$
- 2: **for** $t = 0, \dots, T - 1$ **do**
- 3: $\pi_i^{(t,0)} = \pi_i^{(t)}$
- 4: **for** $k = 0, \dots, K - 1$ and $i \in \mathcal{N}$ **simultaneously do**
- 5: sample B trajectories $\{\{(a_{i,j}^{(b)}, s_j^{(b)}, r_{i,j}^{(b)}, c_j^{(b)})\}_{j=0}^{\hat{T}_e^{(b)}}\}_{b=1}^B$ by following $\pi_i^{(t,k)}$
- 6: set $\hat{V}_{r_i}(\pi^{(t,k)}) = \frac{1}{B} \sum_{b=1}^B \sum_{j=0}^{\hat{T}_e^{(b)}} r_{i,j}^{(b)}$ and $\hat{V}_c(\pi^{(t,k)}) = \frac{1}{B} \sum_{b=1}^B \sum_{j=0}^{\hat{T}_e^{(b)}} c_j^{(b)}$
- 7: $\hat{\nabla} V_{\pi_i}^{r_i}(\pi^{(t,k)}) = \hat{V}_{r_i}(\pi^{(t,k)}) \cdot \frac{1}{B} \sum_{b=1}^B \sum_{j=1}^{\hat{T}_e^{(b)}} \nabla \log \pi_i(a_{i,j}^{(b)} | s_j^{(b)})$
- 8: $\hat{\nabla} V_{\pi_i}^c(\pi^{(t,k)}) = \hat{V}_c(\pi^{(t,k)}) \cdot \frac{1}{B} \sum_{b=1}^B \sum_{j=1}^{\hat{T}_e^{(b)}} \nabla \log \pi_i(a_{i,j}^{(b)} | s_j^{(b)})$
- 9: $\pi_i^{(t,k+1)} = \begin{cases} \mathcal{P}_{\Pi^{i,\xi}} \left[\pi_i^{(t,k)} - \nu_k \hat{\nabla}_{\pi_i} V_{r_i}(\pi^{(t,k)}) - \frac{\nu_k}{\eta} (\pi_i^{(t,k)} - \pi_i^{(t)}) \right] & \text{if } \hat{V}_c(\pi^{(t,k)}) + \beta - \alpha \leq \delta_k \\ \mathcal{P}_{\Pi^{i,\xi}} \left[\pi_i^{(t,k)} - \nu_k \hat{\nabla}_{\pi_i} V_c(\pi^{(t,k)}) - \frac{\nu_k}{\eta} (\pi_i^{(t,k)} - \pi_i^{(t)}) \right] & \text{otherwise} \end{cases}$
- 10: $\mathcal{B}^{(t)} = \{\lfloor K/2 \rfloor \leq k \leq K \mid \hat{V}_c(\pi^{(t,k)}) \leq \delta_k\}$
- 11: $\pi_i^{(t+1)} = \pi_i^{(t,\hat{k})}$ where $\hat{k} = 1$ if $\mathcal{B}^{(t)} = \emptyset$ and else sampled s.t. for $k \in \mathcal{B}^{(t)}$, $\mathbb{P}(\hat{k} = k) = \left(\sum_{k \in \mathcal{B}^{(t)}} \rho_k\right)^{-1} \rho_k$
- 12: **output:** $\pi_i^{(T)}$ for $i \in \mathcal{N}$

Theorem 1. Let Assumptions 1 and 2 hold and let the distribution mismatch coefficient $D := \max_{\pi \in \Pi} \|d_\mu^\pi / \mu\|_\infty$ be finite. For any $\epsilon > 0$, after running iProxCMPG, Algorithm 1, with $\xi = 0$, suitably chosen η, β, T, K , and $\{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$, there exists $t \in [T]$, such that $\pi^{(t)}$ is a constrained ϵ -NE in expectation⁷. The total iteration complexity is given by $\mathcal{O}(\epsilon^{-4})$ where $\mathcal{O}(\cdot)$ hides polynomial dependencies in $m, S, A_{\max}, D, 1 - \gamma$, and ζ .

Before analyzing the outer loop of Algorithm 1, we begin by focusing on the proximal-point update step. We first introduce some useful notation. Then, we explain how we can use the switching gradient algorithm in Appendix C for approximately solving the proximal-point update step independently. We proceed by establishing guarantees that will be important in the analysis of the outer loop of Algorithm 1.

Notation Recall that for any policies $\pi, \pi' \in \Pi$ and any $\eta > 0$, $\Phi_{\eta, \pi'}(\pi) = \Phi(\pi) + \frac{1}{2\eta} \|\pi - \pi'\|^2$, $V_{\eta, \pi'}^c(\pi) = V_c(\pi) + \frac{1}{2\eta} \|\pi - \pi'\|^2$ and $\Pi_{\eta, \pi'}^c = \{\pi \in \Pi \mid V_{\eta, \pi'}^c(\pi) \leq \alpha - \beta\}$. Moreover, recall the following constrained optimization problem:

$$\min_{\pi \in \Pi_{\eta, \pi'}^c} \Phi_{\eta, \pi'}(\pi). \quad (\text{ProxPb}(\eta, \pi'))$$

In the following “ \lesssim ” denotes inequality up to numerical constants. Moreover, let L_Φ be the smoothness constant of the functions Φ and V_c (see Lemma 8) and let Φ_{\max} be an upper bound⁸ on Φ . Recall that under Assumption 1, the initial policy $\pi^{(0)}$ is strictly feasible. We denote the respective slack by $\bar{\zeta}_0 > 0$, i.e., $\bar{\zeta}_0 := \alpha - V_c(\pi^{(0)})$.

Next, we state and prove the guarantees provided by our proximal-point update subroutine.

Lemma 1. Let Assumption 1 hold and let $0 < \bar{\epsilon} \leq \bar{\zeta}_0$. Set $\beta = \bar{\epsilon}$, $\eta = \frac{1}{2L_\Phi}$, and $\xi = 0$. Denote by $\tilde{\pi}^{(t+1)}$ the unique optimal solution to $(\text{ProxPb}(\eta, \pi^{(t)}))$. There exist $K = \mathcal{O}(\bar{\epsilon}^{-2})$ and suitable choices of $\{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$, such that lines 4-6 of Algorithm 1 guarantee that for any $t \in [T - 1]$,

$$\begin{aligned} \mathbb{E} \left[\Phi_{\eta, \pi^{(t)}}(\pi^{(t+1)}) - \Phi_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}) \right] &\leq \bar{\epsilon}^2, \\ \mathbb{E} \left[V_c(\pi^{(t+1)}) \right] &\leq \alpha, \end{aligned} \quad (5)$$

where the expectation is with respect to the randomness induced by the sampling of \hat{k} in line 11 of Algorithm 2.

Proof. We divide the proof into two steps.

⁷Notice that here we take the expectation w.r.t. the randomness which is induced by the sampling of \hat{k} in line 11 of Algorithm 2.

⁸Such a bound is always trivially available.

- **Step 1: Equivalent centralized update rule for our algorithm.** First, we argue that independently running the subroutine given by the inner loop of Algorithm 1, i.e., lines 4-6, is equivalent to a centralized execution of the stochastic switching subgradient algorithm (see Algorithm 3) applied to our proximal-point update problem. Crucially, as observed by Leonardos et al. (2022), Proposition B.1, for any $i \in \mathcal{N}$ and $\pi \in \Pi$, it holds that $\nabla_{\pi_i} \Phi(\pi) = \nabla_{\pi_i} V_{r_i}(\pi)$. We can extend this observation to our regularized potential and value functions, namely for any $\pi' \in \Pi$,

$$\begin{aligned} \nabla_{\pi_i} \Phi_{\eta, \pi'}(\pi) &= \nabla_{\pi_i} \Phi(\pi) + \frac{1}{\eta} (\pi_i - \pi'_i) \\ &= \nabla_{\pi_i} V_{r_i}(\pi) + \frac{1}{\eta} (\pi_i - \pi'_i), \end{aligned}$$

which is an expression that can be evaluated independently by player i , since access to the joint policy π is not required. Together with separability of the projection operator $\mathcal{P}_{\Pi^\varepsilon}$, see e.g. Leonardos et al. (2022), Lemma D.1, we have

$$\left(\mathcal{P}_{\Pi^\varepsilon} \left[\pi_i^{(t,k)} - \nu_k \nabla_{\pi_i} V_{\eta, \pi^{(t,k)}}^{r_i}(\pi^{(t,k)}) \right] \right)_{i \in \mathcal{N}} = \mathcal{P}_{\Pi^\varepsilon} \left[\pi^{(t,k)} - \nu_k \nabla_{\pi} \Phi_{\eta, \pi^{(t,k)}}(\pi^{(t,k)}) \right],$$

and similarly, for the constraint value function,

$$\left(\mathcal{P}_{\Pi^\varepsilon} \left[\pi_i^{(t,k)} - \nu_k \nabla_{\pi_i} V_{\eta, \pi^{(t,k)}}^c(\pi^{(t,k)}) \right] \right)_{i \in \mathcal{N}} = \mathcal{P}_{\Pi^\varepsilon} \left[\pi^{(t,k)} - \nu_k \nabla_{\pi} V_{\eta, \pi^{(t,k)}}^c(\pi^{(t,k)}) \right].$$

Moreover, since $V_c(\pi^{(t,k)})$ can be estimated equally by each player due to the cooperative nature of our constraint, we can conclude that Algorithm 1 is equivalent to a centralized version where the independent, simultaneous update in line 5 is replaced by the following centralized version:

$$\pi^{(t,k+1)} = \begin{cases} \mathcal{P}_{\Pi^\varepsilon} \left[\pi^{(t,k)} - \nu_k \hat{\nabla}_{\pi} \Phi_{\eta, \pi^{(t,k)}}(\pi^{(t,k)}) \right] & \text{if } \hat{V}_c(\pi^{(t,k)}) + \beta - \alpha \leq \delta_k, \\ \mathcal{P}_{\Pi^\varepsilon} \left[\pi^{(t,k)} - \nu_k \hat{\nabla}_{\pi} V_{\eta, \pi^{(t,k)}}^c(\pi^{(t,k)}) \right] & \text{otherwise.} \end{cases}$$

- **Step 2: Induction on t .** Next, to prove the claimed guarantee for all $t \in [T-1]$, we proceed by induction on t . We will invoke results on the stochastic switching gradient algorithm (see CSA, Algorithm 3) that are separately presented in Appendix C in the context of constrained optimization. By Assumption 1, since $\bar{\varepsilon} \leq \bar{\zeta}_0$ and $\beta = \bar{\varepsilon}$, we have $V_c(\pi^{(0)}) \leq \alpha - \beta$. That is, for $t = 0$, the initial feasibility condition of our CSA result, Theorem 3 in Appendix C, holds for $\pi^{(t)}$. Note further that in our deterministic case, Assumption 3 (which is required for Theorem 3) holds, since by Lemma 8 we have a bound on objective and constraint gradient norms.

Hence, we can apply Theorem 3 in the deterministic setting, i.e., with batch size $J = 1$ and access to exact gradients and constraint function values, to $\Phi_{\eta, \pi^{(t)}}$ and $V_{\eta, \pi^{(t)}}^c$ with $\mu = L_\Phi$ and $M^2 \lesssim \max \{M_G^2 + \mu_G^2 \Delta^4, M_F^2 + \mu_F^2 \Delta^4\} \lesssim M_c^2 + L_\Phi^2 \text{diam}(\Pi)^4$, in the notation of Theorem 3. After plugging in the bounds on M_c , L_Φ , and $\text{diam}(\Pi)$ from Lemma 8, and choosing K as in the statement of this lemma, Theorem 3 implies the desired bounds on constraint violation and optimality gap w.r.t. $\tilde{\pi}^{(t+1)}$ in (5). This concludes the base case of the induction.

As induction hypothesis, suppose now that (5) holds for some $t \in [T-1]$. Then, due to $\beta \geq \bar{\varepsilon}$, $V_c(\pi^{(t+1)}) + \beta \leq \alpha + \bar{\varepsilon}$ implies that the initial feasibility condition of Theorem 3 is satisfied and hence with the same argument as above regarding Assumption 3, we can apply Theorem 3 to conclude that at the end of iteration $t+1$ of Algorithm 1, the inner loop guarantees that

$$\begin{aligned} \mathbb{E} \left[\Phi_{\eta, \pi^{(t+2)}}(\pi^{(t+2)}) - \Phi_{\eta, \pi^{(t+2)}}(\tilde{\pi}^{(t+2)}) \right] &\leq \bar{\varepsilon}^2, \\ \mathbb{E} \left[V_c(\pi^{(t+2)}) \right] &\leq \alpha, \end{aligned}$$

i.e., the inductive hypothesis also holds for $t+1$.

□

We next determine the number of iterations of the outer loop of Algorithm 1 required for convergence in the following sense.

Lemma 2. *Let $\epsilon > 0$ and set $\eta = \frac{1}{2L_\Phi}$. Suppose K is chosen such that the guarantee from Lemma 1 holds for $\bar{\epsilon}^2 = \frac{\epsilon^2}{4\eta}$. Then, after $T = \frac{4\eta\Phi_{\max}}{\epsilon^2}$ iterations of the outer loop of Algorithm 1 where Φ_{\max} is an upper bound of the potential function (i.e., $\forall \pi \in \Pi, \Phi(\pi) \leq \Phi_{\max}$), there exists $0 \leq t \leq T - 1$ such that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \epsilon$.*

Remark 4. *Notice that we need an expectation in Lemma 2 in the exact gradients case because of the random index sampling in Algorithm 1. See also Remark 3.*

Proof. Let \mathcal{F}_t denote the σ -field generated by the random variables given by the iterates $\pi^{(t)}$ up to iteration t . Notice that this randomness is induced by the sampling of \hat{k} in line 11 of Algorithm 2. By Lemma 1, the inner loop of Algorithm 1 guarantees that for any $0 \leq t \leq T - 1$,

$$\begin{aligned} \mathbb{E} \left[\Phi(\pi^{(t+1)}) + \frac{1}{2\eta} \|\pi^{(t+1)} - \pi^{(t)}\|^2 \mid \mathcal{F}_t \right] &= \mathbb{E} \left[\Phi_{\eta, \pi^{(t)}}(\pi^{(t+1)}) \mid \mathcal{F}_t \right] \\ &\leq \mathbb{E} \left[\Phi_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}) \mid \mathcal{F}_t \right] + \bar{\epsilon}^2 \\ &\leq \Phi_{\eta, \pi^{(t)}}(\pi^{(t)}) + \bar{\epsilon}^2 \\ &= \Phi(\pi^{(t)}) + \bar{\epsilon}^2 \end{aligned}$$

where the second inequality is due to $\mathbb{E} \left[V_{\eta, \pi^{(t)}}^c(\pi^{(t)}) \mid \mathcal{F}_t \right] = \mathbb{E} \left[V_c(\pi^{(t)}) \right] \leq \alpha$. Taking total expectation in the above inequality, we obtain

$$\mathbb{E} \left[\|\pi^{(t+1)} - \pi^{(t)}\|^2 \right] \leq 2\eta \left(\mathbb{E} \left[\Phi(\pi^{(t)}) \right] - \mathbb{E} \left[\Phi(\pi^{(t+1)}) \right] \right).$$

Summing the above inequality over $0 \leq t \leq T - 1$, using the upper bound Φ_{\max} on the potential function and plugging in our choices of η , T , and $\bar{\epsilon}$, we obtain

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\|\pi^{(t+1)} - \pi^{(t)}\|^2 \right] &\leq 2\eta \left(\bar{\epsilon}^2 + \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\Phi(\pi^{(t)}) \right] - \mathbb{E} \left[\Phi(\pi^{(t+1)}) \right] \right) \\ &\leq \frac{2\eta\Phi_{\max}}{T} + 2\eta\bar{\epsilon}^2 \\ &\leq \epsilon^2. \end{aligned}$$

Using Jensen's inequality, we conclude that there exists $t \in [T - 1]$ such that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \epsilon$. \square

Next, we aim to prove that the event $\|\pi^{(t+1)} - \pi^{(t)}\| \leq \epsilon$ implies $\text{Nash-gap}(\pi^{(t+1)}) = \mathcal{O}(\epsilon)$ where the constrained Nash-gap is defined as

$$\text{Nash-gap}(\pi^*) := \max_{i \in \mathcal{N}} \max_{\pi'_i \in \Pi_c^i(\pi_{-i}^*)} V_{r_i}(\pi^*) - V_{r_i}(\pi'_i, \pi_{-i}^*). \quad (6)$$

As a result, we will be able to argue that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \epsilon$ implies $\mathbb{E}[\text{Nash-gap}(\pi^{(t+1)})] = \mathcal{O}(\epsilon)$, i.e., that the policy $\pi^{(t+1)}$ is a constrained $\mathcal{O}(\epsilon)$ -NE in expectation.

Towards this goal, we first show that a policy $\pi^{(t+1)}$ satisfying $\|\pi^{(t+1)} - \pi^{(t)}\| = \mathcal{O}(\epsilon)$ (as in the previous lemma) is a $\mathcal{O}(\epsilon)$ - $\widetilde{\text{KKT}}$ policy for our initial constrained minimization problem. The ϵ - $\widetilde{\text{KKT}}$ conditions are a slight modification of the standard ϵ -KKT conditions adapted to our specific requirements (see Definition 1 and Definition 2 in Appendix D). In the following lemma, we will be referring to (in-)exact solutions as well as KKT and $\widetilde{\text{KKT}}$ conditions for different problems. Therefore, we first introduce additional useful notation for clarity.

Notation We refer to the following constrained optimization problem as (**InitPb**):

$$\min_{\pi \in \Pi} \Phi(\pi). \quad (\text{InitPb})$$

For the previously introduced ($\text{ProxPb}(\eta, \pi^{(t)})$), we distinguish between the inexact solution resulting from the update which we denote by $\pi^{(t+1)}$, and the exact solution which will be denoted by $\tilde{\pi}^{(t+1)}$ in the proof below. Furthermore, we define the Lagrangians for the two problems as

$$\begin{aligned}\mathcal{L}(\pi, \lambda) &= \Phi(\pi) + \lambda (V_c(\pi) - \alpha), & (\text{InitPb}-\mathcal{L}) \\ \mathcal{L}_{\eta, \pi'}(\pi, \lambda) &= \Phi_{\eta, \pi'}(\pi) + \lambda (V_{\eta, \pi'}^c(\pi) - \alpha + \beta). & (\text{ProxPb}(\eta, \pi')-\mathcal{L})\end{aligned}$$

Using Lemma 1 and Lemma 2, the following lemma shows that Algorithm 1 is guaranteed to generate an $\mathcal{O}(\epsilon)$ - $\widetilde{\text{KKT}}$ policy. Parts of the proof have appeared in a similar form in the optimization literature (see Lemma 3.5 and Theorem 3.2 in Jia and Grimmer (2023), and Theorem 5 in Boob et al. (2023)). The lemma below differs from these results, since we are in a smooth setting and prove convergence w.r.t. our notion of $\widetilde{\text{KKT}}$ conditions rather than towards a point that is near an ϵ - KKT point. Moreover, our guarantee for the proximal update subroutine is somewhat weaker due to the relaxed constraint satisfaction condition that we use to switch between update types in the inner loop, see Lemma 1. Additionally, in order to achieve exact primal feasibility (instead of ϵ -approximate), we employ a feasibility margin β .

Lemma 3. *Let Assumptions 1 and 2 hold. Let $\epsilon > 0$ and choose $\bar{\epsilon}, K, \beta$ as in Lemma 1 and 2. If $\pi^{(t+1)}$ is a policy such that $\|\pi^{(t+1)} - \pi^{(t)}\| \leq \epsilon$ for some $t \in [T - 1]$, then $\pi^{(t+1)}$ is a $(C_{\text{KKT}}\epsilon)$ - $\widetilde{\text{KKT}}$ policy of (InitPb) where C_{KKT} is a positive constant such that $C_{\text{KKT}} \lesssim \frac{m^{2.5} A_{\max}^{1.5} S}{(1-\gamma)^{4.5} \sqrt{\zeta}}$.*

Proof. First, note that ($\text{ProxPb}(\eta, \pi^{(t)})$) is a strongly convex optimization problem with strongly convex constraints, which is sufficient for the existence of a unique optimum $\tilde{\pi}^{(t+1)}$. Since by Assumption 2, Slater's condition holds for ($\text{ProxPb}(\eta, \pi')$) for any $\pi' \in \Pi$, strong duality is given for ($\text{ProxPb}(\eta, \pi^{(t)})$) and hence there exists a finite dual variable $\tilde{\lambda}^{(t+1)} \geq 0$ forming a KKT pair with $\tilde{\pi}^{(t+1)}$. We first claim that $\|\tilde{\pi}^{(t+1)} - \pi^{(t+1)}\| \leq \epsilon$. This can be seen as follows: By optimality of $(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)})$ for ($\text{ProxPb}(\eta, \pi^{(t)})$), we have

$$\tilde{\lambda}^{(t+1)} \left(V_{\eta, \pi^{(t)}}^c(\tilde{\pi}^{(t+1)}) - \alpha + \beta \right) = 0, \quad (7)$$

$$\left\langle \nabla_{\pi} \mathcal{L}_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}), \pi^{(t+1)} - \tilde{\pi}^{(t+1)} \right\rangle \geq 0. \quad (8)$$

From Lemma 8, we know that $\mathcal{L}_{\eta, \pi^{(t)}}(\cdot, \tilde{\lambda}^{(t+1)})$ is $L_{\Phi}(1 + \tilde{\lambda}^{(t+1)})$ -strongly convex. Therefore, after rearranging the standard strong convexity lower bound, we get

$$\begin{aligned}& \frac{L_{\Phi}}{2} \left(1 + \tilde{\lambda}^{(t+1)} \right) \|\tilde{\pi}^{(t+1)} - \pi^{(t+1)}\|^2 \\ & \leq \mathcal{L}_{\eta, \pi^{(t)}}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) - \mathcal{L}_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) - \left\langle \nabla_{\pi} \mathcal{L}_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}), \pi^{(t+1)} - \tilde{\pi}^{(t+1)} \right\rangle \\ & \stackrel{(a)}{\leq} \underbrace{\Phi_{\eta, \pi^{(t)}}(\pi^{(t+1)}) - \Phi_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)})}_{\leq \bar{\epsilon}^2} + \tilde{\lambda}^{(t+1)} \underbrace{\left(V_c(\pi^{(t+1)}) - \alpha - \beta \right)}_{\leq 0} \\ & \stackrel{(b)}{\leq} \bar{\epsilon}^2,\end{aligned}$$

where step (a) follows by applying (7) and (8), and step (b) by Lemma 1, i.e. the guarantee for $\pi^{(t+1)}$ provided by the algorithm's inner loop. Then, it follows from the previous inequality that

$$\|\tilde{\pi}^{(t+1)} - \pi^{(t+1)}\| \leq \sqrt{\frac{2\bar{\epsilon}^2}{L_{\Phi}}} \leq \frac{\epsilon}{\sqrt{2}} \leq \epsilon. \quad (9)$$

Using the fact that $(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)})$ is a KKT pair for ($\text{ProxPb}(\eta, \pi^{(t)})$), we now argue that $(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)})$ is a $(C_{\text{KKT}}\epsilon)$ - $\widetilde{\text{KKT}}$ pair for (InitPb) (see Definition 2 in Appendix D). We check each one of the requirements of the definition in what follows.

- **Exact primal feasibility:** By Lemma 1, we know that $V_c(\pi^{(t+1)}) \leq \alpha$ for any $0 \leq t \leq T - 1$.

- **Dual feasibility:** This immediately holds by dual feasibility of $(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)})$ for $(\text{ProxPb}(\eta, \pi^{(t)}))$.
- **Complementary slackness:** When $\tilde{\lambda}^{(t+1)} = 0$, we clearly have $|\tilde{\lambda}^{(t+1)} (V_c(\pi^{(t+1)}) - \alpha)| = 0 \leq \epsilon$. Otherwise, we have

$$\begin{aligned}
 V_c(\pi^{(t+1)}) &\stackrel{(a)}{\geq} V_c(\tilde{\pi}^{(t+1)}) - M_c \epsilon \\
 &\stackrel{(b)}{=} \alpha - \beta - \frac{1}{2\eta} \|\tilde{\pi}^{(t+1)} - \pi^{(t)}\|^2 - M_c \epsilon \\
 &\stackrel{(c)}{\geq} \alpha - \frac{\epsilon^2}{\eta} - \left(M_c + \frac{1}{2\sqrt{\eta}} \right) \epsilon,
 \end{aligned} \tag{10}$$

where (a) follows from M_c Lipschitz continuity of V_c (see Lemma 8-item (1)) and Eq. (9), (b) stems from complementary slackness of $(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)})$ for $(\text{ProxPb}(\eta, \pi^{(t)}))$ which states that $V_{\eta, \pi^{(t)}}^c(\tilde{\pi}^{(t+1)}) - \alpha + \beta = 0$. To obtain inequality (c), observe that using the bound from (9), our assumption on $\|\pi^{(t+1)} - \pi^{(t)}\|$, and the triangle inequality, we have $\|\tilde{\pi}^{(t+1)} - \pi^{(t)}\| \leq \|\tilde{\pi}^{(t+1)} - \pi^{(t+1)}\| + \|\pi^{(t+1)} - \pi^{(t)}\| \leq 2\epsilon$.

Combining (10) with the upper bound $V_c(\pi^{(t+1)}) \leq \alpha$ from primal feasibility, we get

$$\left| \tilde{\lambda}^{(t+1)} (V_c(\pi^{(t+1)}) - \alpha) \right| \leq \tilde{\lambda}^{(t+1)} \left(\frac{\epsilon^2}{\eta} + M_c \epsilon + \frac{\epsilon}{2\sqrt{\eta}} \right). \tag{11}$$

We now show that the dual variable $\tilde{\lambda}^{(t+1)}$ is bounded by a constant depending on ζ using Assumption 2 and strong duality. Indeed, we have

$$\tilde{\lambda}^{(t+1)} \leq \frac{\|\nabla_{\pi} \Phi(\tilde{\pi}^{(t+1)})\| + \eta^{-1} \|\tilde{\pi}^{(t+1)} - \pi^{(t)}\|}{\sqrt{\zeta \eta^{-1}}} \leq \frac{M_c + 4\epsilon \eta^{-1}}{\sqrt{\zeta \eta^{-1}}}, \tag{12}$$

where the first inequality follows from the proof of Lemma 1 in Ma et al. (2020), whereas the second inequality uses Lipschitzness of the potential function (see Lemma 8) and the fact that $\|\tilde{\pi}^{(t+1)} - \pi^{(t)}\| \leq 2\epsilon$. Combining (11) and (12), and using the bounds on M_c and L_{Φ} from Lemma 8, we obtain the desired $C_{\text{KKT}\epsilon}$ -complementary slackness.

- **Variational Lagrangian stationarity:** Suppose by contradiction that the Lagrangian stationarity condition that comes with the $\frac{2\epsilon(1+\tilde{\lambda}^{(t+1)})}{\eta}$ -KKT conditions does not hold for $\tilde{\pi}^{(t+1)}$ and (InitPb) . Then there exists $\nu \in N_{\Pi}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)})$ (normal cone to the convex set of policies Π) such that

$$\nabla_{\pi} \mathcal{L}_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) + \nu = 0 \quad \text{and} \quad \left\| \nabla_{\pi} \mathcal{L}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) + \nu \right\| > \frac{2\epsilon(1 + \tilde{\lambda}^{(t+1)})}{\eta}$$

where the equality is by Lagrangian stationarity of $\tilde{\pi}^{(t+1)}$ for $(\text{ProxPb}(\eta, \pi^{(t)}))$ and the inequality is due to the above assumed lack of Lagrangian stationarity of $\tilde{\pi}^{(t+1)}$ for (InitPb) . Plugging in the definition of $\mathcal{L}_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)})$ and combining the equality and inequality above, one can conclude that

$$\frac{2\epsilon(1 + \tilde{\lambda}^{(t+1)})}{\eta} < \left\| \nabla_{\pi} \mathcal{L}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) + \nu \right\| = \frac{1 + \tilde{\lambda}^{(t+1)}}{\eta} \left\| \tilde{\pi}^{(t+1)} - \pi^{(t)} \right\|,$$

which contradicts the inequality $\|\tilde{\pi}^{(t+1)} - \pi^{(t)}\| \leq 2\epsilon$. Hence using the bound on $\tilde{\lambda}^{(t+1)}$ from (12), the policy $\tilde{\pi}^{(t+1)}$ is a $\tilde{C}\epsilon$ -KKT policy for (InitPb) with $\tilde{C} = \frac{2}{\eta} \left(1 + \frac{M_c + 2\epsilon \eta^{-1}}{\sqrt{\zeta \eta^{-1}}} \right)$.

By Lemma 13, this implies that

$$\max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle \leq \text{diam}(\Pi) \tilde{C} \epsilon. \tag{13}$$

Then, in view of showing the variational Lagrangian stationarity for the pair $(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)})$ for (InitPb) ,

we can write

$$\begin{aligned}
 & \max_{\pi' \in \Pi} \left\langle \pi^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle \\
 &= \max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle + \left\langle \pi^{(t+1)} - \tilde{\pi}^{(t+1)}, \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle \\
 &\leq \max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle + \left\| \pi^{(t+1)} - \tilde{\pi}^{(t+1)} \right\| \cdot \left\| \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\| \\
 &\leq \max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle + \epsilon(1 + \tilde{\lambda}^{(t+1)})M_c. \tag{14}
 \end{aligned}$$

We now bound the first term in the above inequality by using $2L_{\Phi}(1 + \tilde{\lambda}^{(t+1)})$ -smoothness of $\nabla_{\pi} \mathcal{L}(\cdot, \tilde{\lambda}^{(t+1)})$ and (13). Using the fact that $\max_{\pi \in \Pi} (A(\pi) + B(\pi)) \leq \max_{\pi \in \Pi} A(\pi) + \max_{\pi \in \Pi} B(\pi)$ for any functions $A(\pi), B(\pi)$, we have

$$\begin{aligned}
 & \max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle \\
 &\leq \max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) - \nabla_{\pi} \mathcal{L}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle + \max_{\pi' \in \Pi} \left\langle \tilde{\pi}^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\tilde{\pi}^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle \\
 &\leq \max_{\pi' \in \Pi} \left\| \tilde{\pi}^{(t+1)} - \pi' \right\| \cdot 2L_{\Phi}(1 + \tilde{\lambda}^{(t+1)}) \left\| \tilde{\pi}^{(t+1)} - \pi^{(t+1)} \right\| + \text{diam}(\Pi)\tilde{C}\epsilon \\
 &\leq \left(2\text{diam}(\Pi)L_{\Phi}(1 + \tilde{\lambda}^{(t+1)}) + \text{diam}(\Pi)\tilde{C} \right) \epsilon.
 \end{aligned}$$

Combining the above inequality with (14), we obtain

$$\max_{\pi' \in \Pi} \left\langle \pi^{(t+1)} - \pi', \nabla_{\pi} \mathcal{L}(\pi^{(t+1)}, \tilde{\lambda}^{(t+1)}) \right\rangle \leq \left((1 + \tilde{\lambda}^{(t+1)})M_c + 2\text{diam}(\Pi)L_{\Phi}(1 + \tilde{\lambda}^{(t+1)}) + \text{diam}(\Pi)\tilde{C} \right) \epsilon.$$

Finally, we use the bound on $\tilde{\lambda}^{(t+1)}$ from (12), as well as bounds on $\text{diam}(\Pi)$, L_{Φ} , and M_c from Lemma 8, to conclude that $\pi^{(t+1)}$ is a $C_{\text{KKT}}\epsilon$ - $\widetilde{\text{KKT}}$ policy for (InitPb). \square

To complete the analysis, it now remains to show that an $\mathcal{O}(\epsilon)$ - $\widetilde{\text{KKT}}$ policy of (InitPb) is a constrained $\mathcal{O}(\epsilon)$ -NE. For this, we leverage the playerwise gradient domination property satisfied by the potential function and the constraint value function. We first introduce some notations.

Notation For each player $i \in \mathcal{N}$ and each policy $\pi_{-i} \in \Pi^{-i}$, consider the playerwise constrained optimization problem given by

$$\min_{\pi_i \in \Pi_c^i(\pi_{-i})} V_{r_i}(\pi_i, \pi_{-i}). \tag{PlayerPb}(\pi_{-i})$$

The respective Lagrangian $\mathcal{L}_{\pi_{-i}} : \Pi^i \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is defined for every $\pi_i \in \Pi^i$ and every $\lambda \geq 0$ by

$$\mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) = \Phi(\pi_i, \pi_{-i}) + \lambda(V_c(\pi_i, \pi_{-i}) - \alpha). \tag{PlayerPb}(\pi_{-i})-\mathcal{L}$$

Lemma 4. *Let $\pi \in \Pi$ be an ϵ - $\widetilde{\text{KKT}}$ policy of (InitPb). Then π is a constrained $C_{\text{NE}}\epsilon$ -NE where $C_{\text{NE}} \lesssim \frac{D}{1-\gamma} + 1$.*

Proof. The proof of the lemma proceeds in two steps:

- **Step 1.** We show that if π is an $\mathcal{O}(\epsilon)$ - $\widetilde{\text{KKT}}$ policy of (InitPb), then for all $i \in \mathcal{N}$, π_i is an $\mathcal{O}(\epsilon)$ - $\widetilde{\text{KKT}}$ policy of (PlayerPb)(π_{-i}).
- **Step 2.** We conclude that each player cannot significantly improve its policy π_i while staying within $\Pi_c^i(\pi_{-i})$ which means π is a constrained $\mathcal{O}(\epsilon)$ -NE.

We provide a proof of each one of the steps successively.

- **Step 1:** Let $\lambda \geq 0$ be a dual variable such that (π, λ) is an ϵ - $\widetilde{\text{KKT}}$ pair of **(InitPb)**, and let $i \in \mathcal{N}$ be arbitrary. We show that (π_i, λ) is an ϵ - $\widetilde{\text{KKT}}$ pair of **(PlayerPb)(π_{-i})** by checking that the respective $\widetilde{\text{KKT}}$ conditions hold. For dual and exact primal feasibility, as well as complementary slackness, this is immediate since the conditions are equivalent for **(InitPb)** and **(PlayerPb)(π_{-i})**. For variational Lagrangian stationarity, observe that

$$\begin{aligned} \max_{\pi'_i \in \Pi^i} \langle \pi_i - \pi'_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle &= \max_{\pi'_i \in \Pi^i} \langle \pi - (\pi'_i, \pi_{-i}), \nabla_{\pi} \mathcal{L}(\pi, \lambda) \rangle \\ &\leq \max_{\pi' \in \Pi} \langle \pi - \pi', \nabla_{\pi} \mathcal{L}(\pi, \lambda) \rangle \\ &\leq \epsilon, \end{aligned}$$

where the first equality is due to $\nabla_{\pi_i} \mathcal{L}(\pi, \lambda) = \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)$ and the fact that all terms except for $\pi_i - \pi'_i$ vanish in the first argument of the scalar product. The second inequality is because $(\pi'_i, \pi_{-i}) \in \Pi$, and the final step is by Lagrangian stationarity of π for **(InitPb)**.

- **Step 2:** Let $i \in \mathcal{N}$ and consider the MDP M_λ , $\lambda \geq 0$, with state space \mathcal{S} , action space \mathcal{A}_i , probability transition kernel P_λ , reward r_λ , and initial distribution μ where

$$\begin{aligned} P_\lambda(s' | s, a_i) &:= \mathbb{E}_{a_{-i} \sim \pi_{-i}((a_i, \cdot)|s)} [P(s' | s, (a_i, a_{-i}))] \\ r_\lambda(s, a_i) &:= \mathbb{E}_{a_{-i} \sim \pi_{-i}((a_i, \cdot)|s)} [r_i(s, (a_i, a_{-i})) + \lambda c(s, (a_i, a_{-i}))]. \end{aligned}$$

Observe that $\mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)$ is the value function associated to the policy π_i in the MDP M_λ , and hence gradient domination holds (Agarwal et al., 2021), i.e., we have

$$\mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) - \min_{\tilde{\pi}_i \in \Pi^i} \mathcal{L}_{\pi_{-i}}(\tilde{\pi}_i, \lambda) \leq \frac{D}{1-\gamma} \max_{\pi'_i \in \Pi^i} \langle \pi_i - \pi'_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle,$$

where D is the distribution mismatch coefficient, supposed to be finite. Using Proposition 2 in Appendix D for $C_1 = 0$, and using the definition of the playerwise primal optimum, see **(PlayerPb)(π_{-i})**, we then get

$$V_{r_i}(\pi) - \min_{\pi_i^* \in \Pi_{c_i}(\pi_{-i})} V_i(\pi_i^*, \pi_{-i}) \leq \left(\frac{D}{1-\gamma} + 1 \right) \epsilon.$$

Since additionally we have exact primal feasibility of π for **(InitPb)**, the result follows by definition of the constrained ϵ -NE. □

Finally, we put together above lemmas to prove the main theorem.

Proof of Theorem 1. Let $\tilde{\epsilon}^2 = \frac{L_\Phi \epsilon^2}{C_{\text{KKT}}^2 C_{\text{NE}}^2}$. Then with $K = \mathcal{O}(\epsilon^{-2})$, and $T = \mathcal{O}(\epsilon^{-2})$, Lemma 1 and Lemma 2 imply that there exists $0 \leq t \leq T - 1$ such that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \frac{\epsilon}{C_{\text{KKT}} C_{\text{NE}}}$. We use Lemma 3 to conclude that $\pi^{(t+1)}$ is a ϵ/C_{NE} - $\widetilde{\text{KKT}}$ policy of **(InitPb)**. Then, by Lemma 4, $\pi^{(t+1)}$ is a constrained ϵ -NE. The total iteration complexity is bounded by $T \cdot K = \mathcal{O}(\epsilon^{-4})$. □

B.2 Proof of Theorem 2 — Finite Sample Case

Moving on to the stochastic setting, we first restate Theorem 2.

Theorem 2. *Let Assumptions 1 and 2 hold, and let D (as in Theorem 1) be finite. Then, for any $\epsilon > 0$, after running *iProxCMPG* based on finite sample estimates (see Algorithm 2) for suitably chosen $\eta, \beta, \xi, T, K, B$, and $\{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$, there exists $t \in [T]$, such that in expectation, $\pi^{(t)}$ is a constrained ϵ -NE. The total sample complexity is given by $\tilde{\mathcal{O}}(\epsilon^{-7})$ where $\tilde{\mathcal{O}}(\cdot)$ hides polynomial dependencies in $m, S, A_{\max}, D, 1 - \gamma$, and ζ , as well as logarithmic dependencies in $1/\epsilon$.*

Similar to the deterministic case, we begin by proving the guarantees provided by the inner loop of Algorithm 2.

Lemma 5. *Let Assumption 1 hold, let $\bar{\epsilon} > 0$ and set $\beta = \bar{\epsilon}, \eta = \frac{1}{2L_\Phi}, \xi = \bar{\epsilon}\sqrt{2\eta}$. Then, there exist $K = \tilde{\mathcal{O}}(\bar{\epsilon}^{-3})$, $B = \tilde{\mathcal{O}}(\bar{\epsilon}^{-2})$, and suitable choices of $\{(\nu_k, \delta_k, \rho_k)\}_{0 \leq k \leq K}$, such that lines 4-11 of Algorithm 2 guarantee that for any $t \in [T - 1]$,*

$$\begin{aligned} \mathbb{E} \left[\Phi_{\eta, \pi^{(t)}}(\pi^{(t+1)}) - \Phi_{\eta, \pi^{(t)}}(\tilde{\pi}^{(t+1)}) \right] &\leq \bar{\epsilon}^2, \\ \mathbb{E} \left[V_c(\pi^{(t+1)}) \right] &\leq \alpha, \end{aligned} \quad (15)$$

where $\tilde{\pi}^{(t+1)}$ denotes the unique optimal solution to $(\text{ProxPb}(\eta, \pi^{(t)}))$.

Proof. The result follows similarly as for Lemma 1 in the deterministic case. We hence only point out differences. In order to ensure bounded norms of gradient estimates, we use ξ -greedy policies. Then, according to Lemma 10, the second moment of value and constraint gradient estimates is bounded by $\mathcal{O}(1/\xi)$. The concentration result shown in Lemma 11 ensures that constraint value estimates follow a sub-exponential distribution. Therefore, Assumption 3, see Appendix C on guarantees for our subroutine, is satisfied. We can thus apply the respective Theorem 3 for optimizing over Π^ξ , and with $\mu = L_\Phi$ and $M^2 \lesssim \max\{M_G^2 + \mu_G^2 \Delta^4, M_F^2 + \mu_F^2 \Delta^4\} \lesssim \frac{24A_{\max}^2}{\xi(1-\gamma)^4} + L_\Phi^2 \text{diam}(\Pi)^4$. After plugging bounds on L_Φ , and $\text{diam}(\Pi)$ from Lemma 8, and choosing K and B as stated, Theorem 3 implies the desired bounds via the same arguments as in the proof of Lemma 1. \square

Next, we analyze the convergence of our main proximal-point method, Algorithm 2. More concretely, we bound the required sample complexity for ensuring that for some $\epsilon > 0$, there exist iterates $\pi^{(t)}, \pi^{(t+1)}$ such that $\|\pi^{(t)} - \pi^{(t+1)}\| = \mathcal{O}(\epsilon)$. In the following, we will then prove that this implies convergence to a constrained $\mathcal{O}(\epsilon)$ -NE.

Similarly to the deterministic case, we next determine the number of updates needed until convergence in the following sense. The next lemma is analogous to its deterministic counterpart Lemma 2.

Lemma 6. *Let $\epsilon > 0$ and set $\eta = \frac{1}{2L_\Phi}$. Suppose K is chosen such that the guarantee from Lemma 5 holds for $\bar{\epsilon}^2 = \frac{\epsilon^2}{4\eta}$. Then after $T = \frac{4\eta\Phi_{\max}}{\epsilon^2}$ iterations of the outer loop of Algorithm 2 where Φ_{\max} is an upper bound of the potential function (i.e., $\forall \pi \in \Pi, \Phi(\pi) \leq \Phi_{\max}$), there exists $0 \leq t \leq T - 1$ such that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \epsilon$.*

Proof. The proof follows the same lines as the proof of Lemma 2 upon replacing Lemma 1 by Lemma 5. We do not reproduce it here for conciseness. \square

Next, as in the exact gradients case, we aim to prove that the event $\|\pi^{(t+1)} - \pi^{(t)}\| \leq \epsilon$ implies $\text{Nash-gap}(\pi^{(t+1)}) = \mathcal{O}(\epsilon)$, in order to argue that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \epsilon$ implies $\mathbb{E}[\text{Nash-gap}(\pi^{(t+1)})] = \mathcal{O}(\epsilon)$.

Recall that in Lemma 3 we have already shown $\|\pi^{(t+1)} - \pi^{(t)}\| \leq \epsilon$ to imply that $\pi^{(t+1)}$ is a $\widetilde{C_{\text{KKT}}\epsilon\text{-KKT}}$ policy of (InitPb) which equivalently holds in the stochastic ξ -greedy setting. Arguing that a $\epsilon\text{-KKT}$ policy is a constrained $\mathcal{O}(\epsilon)$ -NE however requires an adapted proof, since here in each iteration we solve the subproblem over Π^ξ instead of Π , i.e., the $\widetilde{\text{KKT}}$ conditions hold w.r.t. Π^ξ . The following lemma is an adjustment of Lemma 4 for this fact.

Lemma 7. *Let $\pi \in \Pi^\xi$ be an $\epsilon\text{-KKT}$ policy of (InitPb) (where $\widetilde{\text{KKT}}$ are w.r.t. Π^ξ) and $\xi = \epsilon$. Then π is a constrained $\hat{C}_{NE}\epsilon\text{-NE}$ where $\hat{C}_{NE} \lesssim \frac{D}{1-\gamma} + \frac{m\sqrt{S}A_{\max}D}{(1-\gamma)^{4.5}} + 1$.*

Proof. We divide the proof into two steps:

- **Step 1:** Analogously to step 1 of Lemma 4, one can show that (π_i, λ) is an $\epsilon\text{-KKT}$ pair of $(\text{PlayerPb}(\pi_{-i}))$.
- **Step 2:** Let $i \in \mathcal{N}$ and consider the MDP \tilde{M}_λ (for $\lambda \geq 0$) with state space \mathcal{S} , action space \mathcal{A}_i , transition probability kernel \tilde{P} , reward \tilde{r}_λ , discount factor γ , and initial distribution μ where

$$\begin{aligned} \tilde{P}(s' \mid s, a_i) &:= \mathbb{E}_{a_{-i} \sim \pi_{-i}(\cdot \mid s)} [P(s' \mid s, (a_i, a_{-i}))], \\ \tilde{r}_\lambda(s, a_i) &:= \mathbb{E}_{a_{-i} \sim \pi_{-i}(\cdot \mid s)} [r_i(s, (a_i, a_{-i})) + \lambda c(s, (a_i, a_{-i}))]. \end{aligned}$$

Observe that $\mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)$ is the value function associated to the policy π_i for \tilde{M}_λ , and hence gradient domination holds (Agarwal et al., 2021). In our particular case of π being a ξ -greedy policy, we can also show a similar inequality, even w.r.t. the non- ξ -greedy optimum. Let

$$\begin{aligned}\hat{\pi}_i &:= \arg \max_{\pi'_i \in \Pi^i} \langle \pi_i - \pi'_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle, \\ \hat{\pi}_i^\xi &:= \arg \max_{\pi'_i \in \Pi^{i, \xi}} \langle \pi_i - \pi'_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle.\end{aligned}$$

Then, we have

$$\begin{aligned}\mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) - \min_{\pi_i^* \in \Pi^i} \mathcal{L}_{\pi_{-i}}(\pi_i^*, \lambda) &\leq \frac{1}{1-\gamma} \left\| \frac{d_{\mu}^{\pi_i^*, \pi_{-i}}}{\mu} \right\|_\infty \langle \pi_i - \hat{\pi}_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle \\ &= \frac{1}{1-\gamma} \left\| \frac{d_{\mu}^{\pi_i^*, \pi_{-i}}}{\mu} \right\|_\infty \left[\langle \pi_i - \hat{\pi}_i^\xi, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle + \langle \hat{\pi}_i^\xi - \hat{\pi}_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle \right].\end{aligned}$$

We further bound the last term above as follows

$$\begin{aligned}\langle \hat{\pi}_i^\xi - \hat{\pi}_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle &= \max_{\pi_i^\xi \in \Pi^{i, \xi}} \langle \pi_i^\xi, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle - \max_{\pi_i \in \Pi^i} \langle \pi_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle \\ &\stackrel{(a)}{\leq} \xi \sqrt{S} \|\nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)\| \\ &\stackrel{(b)}{\leq} \xi \sqrt{S} (1 + \lambda) M_c.\end{aligned}$$

In the above inequalities, (a) follows from using Lemma 9 to obtain

$$\max_{\pi_i^\xi \in \Pi^{i, \xi}} \langle \pi_i^\xi, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle \leq \max_{\pi_i \in \Pi^i} \langle \pi_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle + \sum_{a_i, s} \frac{\xi}{A_i} [\nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)](a_i | s).$$

Using the fact that for any $x \in \mathbb{R}^d$, $\|x\|_1 \leq \sqrt{d} \|x\|_2$, we further get

$$\sum_{a_i, s} \frac{\xi}{A_i} [\nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)](a_i | s) = \frac{\xi}{A_i} \|\nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)\|_1 \leq \frac{\xi}{A_i} \sqrt{A_i S} \|\nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)\| \leq \zeta \sqrt{S} \|\nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda)\|.$$

The bound used in (b) follows from Lipschitz continuity, see Lemma 8. We conclude that

$$\mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) - \min_{\pi_i^* \in \Pi^i} \mathcal{L}_{\pi_{-i}}(\pi_i^*, \lambda) \leq \frac{1}{1-\gamma} \left\| \frac{d_{\mu}^{\pi_i^*, \pi_{-i}}}{\mu} \right\|_\infty \left[\langle \pi_i - \hat{\pi}_i, \nabla_{\pi_i} \mathcal{L}_{\pi_{-i}}(\pi_i, \lambda) \rangle + \xi \sqrt{S} (1 + \lambda) M_c \right]. \quad (16)$$

Applying Proposition 2, see Appendix D.2, with $\xi = \epsilon$, bounding the distribution mismatch coefficient by D , and using the definition of the playerwise primal optimum, see (PlayerPb(π_{-i})), we then get

$$\begin{aligned}V_i(\pi) - \min_{\pi_i^* \in \Pi_{c_i}(\pi_{-i})} V_i(\pi_i^*, \pi_{-i}) &\leq \left(\frac{D}{1-\gamma} + \frac{(1+\lambda)\sqrt{S}M_c D}{1-\gamma} + 1 \right) \epsilon \\ &\leq \left(\frac{D}{1-\gamma} + \frac{m\sqrt{S}A_{\max}D}{(1-\gamma)^{4.5}} + 1 \right) \epsilon.\end{aligned}$$

where for the second inequality we use (12) and our bounds on M_c and L_Φ from Lemma 8. Since additionally, we have exact primal feasibility of π for (InitPb) by the KKT conditions, the result follows by definition of a constrained ϵ -NE. \square

Finally, we put together above lemmas to prove our main theorem in the stochastic setting.

Proof of Theorem 2. Let $\tilde{\epsilon}^2 = \frac{L_\Phi \epsilon^2}{C_{\text{KKT}}^2 \hat{C}_{\text{NE}}^2}$. Then with $K = \tilde{\mathcal{O}}(\epsilon^{-3})$, $B = \tilde{\mathcal{O}}(\epsilon^{-2})$, and $T = \mathcal{O}(\epsilon^{-2})$, Lemma 1 and Lemma 2 imply that there exists $t \in [T - 1]$ such that $\mathbb{E}[\|\pi^{(t+1)} - \pi^{(t)}\|] \leq \frac{\epsilon}{C_{\text{KKT}} \hat{C}_{\text{NE}}}$. We use Lemma 3 to conclude that $\pi^{(t+1)}$ is a $\epsilon/\hat{C}_{\text{NE}}\widetilde{\text{KKT}}$ policy of (InitPb) in expectation. Then, by Lemma 4, $\pi^{(t+1)}$ is a constrained ϵ -NE in expectation. The total sample complexity is bounded by $T \cdot K \cdot B = \tilde{\mathcal{O}}(\epsilon^{-7})$. \square

B.3 Other Technical Lemmas

The next lemma collects standard regularity properties of the value and potential functions.

Lemma 8. *The following statements hold true.*

1. *The functions Φ and V_c are M_c -Lipschitz continuous over Π with $M_c = \frac{\sqrt{m A_{\max}}}{(1-\gamma)^2}$. This immediately implies that $\|\nabla\Phi(\pi)\| \leq M_c$ and $\|\nabla V_c(\pi)\| \leq M_c$, for all $\pi \in \Pi$.*
2. *For any $i \in \mathcal{N}$ and any $\pi_{-i} \in \Pi_{-i}$, the function $V_{r_i}(\cdot, \pi_{-i})$ is L_i -smooth with $L_i = \frac{2\gamma A_i}{(1-\gamma)^3}$ and hence L_i -weakly convex.*
3. *The functions Φ and V_c are L_Φ -smooth with $L_\Phi = m \cdot \max_i L_i = \frac{2m\gamma A_{\max}}{(1-\gamma)^3}$ and hence L_Φ -weakly convex.*
4. *For $\eta = \frac{1}{2L_\Phi}$ and any $\pi' \in \Pi$, the regularized function $\Phi_{\eta, \pi'}(\pi) = \Phi(\pi) + L_\Phi \|\pi - \pi'\|^2$ is L_Φ -strongly convex and the functions $\Phi_{\eta, \pi'}$, $V_{\eta, \pi'}^c$ are both $2L_\Phi$ -smooth.*
5. *For any $\lambda \in \mathbb{R}$, $\pi' \in \Pi$ and $\eta = \frac{1}{2L_\Phi}$, $\mathcal{L}(\cdot, \lambda) = \Phi(\cdot) + \lambda V_c(\cdot)$ is $L_\Phi(1 + \lambda)$ -smooth, and $\mathcal{L}_{\eta, \pi'}(\cdot, \lambda) = \Phi_{\eta, \pi'}(\cdot) + \lambda V_{\eta, \pi'}^c(\cdot)$ is $2L_\Phi(1 + \lambda)$ -smooth. Hence $\mathcal{L}_{\eta, \pi'}(\cdot, \lambda)$ is also $L_\Phi(1 + \lambda)$ -strongly convex.*

Proof. Item 2 has been proved in Agarwal et al. (2021), Lemma D.3. Item 3 has been reported in Leonardos et al. (2022), Lemma D.4. Item 4 immediately follows from item 3. We now prove item 5 for \mathcal{L} , the result for $\mathcal{L}_{\eta, \pi'}$ follows similarly. Using the definition of the Lagrangian and the triangle inequality, for any $\lambda \in \mathbb{R}$ and $\pi, \pi' \in \Pi$,

$$\begin{aligned} \|\nabla_\pi \mathcal{L}(\pi, \lambda) - \nabla_\pi \mathcal{L}(\pi', \lambda)\| &\leq \|\nabla\Phi(\pi) - \nabla\Phi(\pi')\| + \lambda \|\nabla V_c(\pi) - \nabla V_c(\pi')\| \\ &\leq L_\Phi \|\pi - \pi'\| + \lambda L_\Phi \|\pi - \pi'\| \\ &\leq L_\Phi(1 + \lambda) \|\pi - \pi'\|. \end{aligned}$$

To show item 1, Lipschitz continuity of Φ and V_c , observe that for any $i \in \mathcal{N}$, $\pi \in \Pi$ and $\pi'_i \in \Pi^i$, by using Lemma 32 of Zhang et al. (2022a) in the second step, we have

$$\begin{aligned} |\Phi(\pi_i, \pi_{-i}) - \Phi(\pi'_i, \pi_{-i})| &= |V_{r_i}(\pi_i, \pi_{-i}) - V_{r_i}(\pi'_i, \pi_{-i})| \\ &\leq \frac{1}{(1-\gamma)^2} \max_{s \in \mathcal{S}} \|\pi_i(\cdot | s) - \pi'_i(\cdot | s)\|_1 \\ &\leq \frac{\sqrt{A_i}}{(1-\gamma)^2} \max_{s \in \mathcal{S}} \|\pi_i(\cdot | s) - \pi'_i(\cdot | s)\|_2 \\ &\leq \frac{\sqrt{A_i}}{(1-\gamma)^2} \|\pi_i - \pi'_i\|_2 \end{aligned}$$

where in the third step we use the fact that for any $x \in \mathbb{R}^d$, $\|x\|_1 \leq \sqrt{d} \|x\|_2$. Then, the following decomposition

yields the result: For any $\pi, \pi' \in \Pi$,

$$\begin{aligned}
 |\Phi(\pi) - \Phi(\pi')| &= \left| \sum_{i \in \mathcal{N}} \Phi(\pi'_1, \dots, \pi'_{i-1}, \pi_i, \pi_{i+1}, \dots, \pi_m) - \Phi(\pi'_1, \dots, \pi'_{i-1}, \pi'_i, \pi_{i+1}, \dots, \pi_m) \right| \\
 &\leq \sum_{i \in \mathcal{N}} \left| \Phi(\pi'_1, \dots, \pi'_{i-1}, \pi_i, \pi_{i+1}, \dots, \pi_m) - \Phi(\pi'_1, \dots, \pi'_{i-1}, \pi'_i, \pi_{i+1}, \dots, \pi_m) \right| \\
 &\leq \frac{1}{(1-\gamma)^2} \sum_{i \in \mathcal{N}} \sqrt{A_i} \|\pi_i - \pi'_i\| \\
 &\leq \frac{\sqrt{mA_{\max}}}{(1-\gamma)^2} \|\pi - \pi'\|,
 \end{aligned}$$

where in the second inequality we apply the above result for playerwise deviations, and the last step is again due to the fact that for any $x \in \mathbb{R}^d$, $\|x\|_1 \leq \sqrt{d} \|x\|_2$. The result follows similarly for V_c . \square

The next lemma is an immediate result showing that any ξ -greedy playerwise policy $\pi_i \in \Pi^{i,\xi}$ (see definition in the main part p. 7 which defines this set as a set of lower bounded policies away from zero) can be represented as a convex combination of a uniform distribution over the action space \mathcal{A}_i and a policy $\pi_i \in \Pi^i$.

Lemma 9. For any $\xi > 0$, $i \in \mathcal{N}$,

$$\Pi^{i,\xi} \subseteq \left\{ \pi_i \in \Pi^i \mid \exists \pi'_i \in \Pi^i, \forall a_i \in \mathcal{A}_i, \forall s \in \mathcal{S} : \pi_i(a_i | s) = \xi/A_i + (1-\xi)\pi'_i(a_i | s) \right\}.$$

Proof. Let $\xi > 0$, $i \in \mathcal{N}$ and let $\pi_i^\xi \in \Pi^{i,\xi}$. Then for all $a_i \in \mathcal{A}_i$ and $s \in \mathcal{S}$, set

$$\pi_i(a_i | s) := \frac{\pi_i^\xi(a_i | s) - \xi/A_i}{1-\xi}.$$

Indeed $\pi_i \in \Pi^i$, since for all $a_i \in \mathcal{A}_i$ and $s \in \mathcal{S}$, we have $\pi_i(a_i | s) \geq 0$ due to $\pi_i^\xi(a_i | s) \geq \xi$ and

$$\sum_{a_i \in \mathcal{A}_i} \pi_i(a_i | s) = \frac{1}{1-\xi} \left(\underbrace{\sum_{a_i \in \mathcal{A}_i} \pi_i^\xi(a_i | s)}_{=1} - \underbrace{\sum_{a_i \in \mathcal{A}_i} \xi/A_i}_{=\xi} \right) = 1.$$

\square

The following lemma shows that the estimators we use for the playerwise policy gradients are unbiased and enjoy a bounded variance.

Lemma 10 (Daskalakis et al. (2020); Leonardos et al. (2022)). For any $\xi > 0$ and $\pi \in \Pi^\xi$, we have

$$\begin{aligned}
 \mathbb{E}_\pi \left[\hat{\nabla} V_{\pi_i}^{r_i}(\pi) \right] &= \nabla_{\pi_i} V_{r_i}(\pi) = \nabla_{\pi_i} \Phi(\pi), \\
 \mathbb{E}_\pi \left[\left\| \hat{\nabla} V_{\pi_i}^{r_i}(\pi) \right\|^2 \right] &\leq \frac{24A_{\max}^2}{\xi(1-\gamma)^4}, \\
 \mathbb{E}_\pi \left[\hat{V}_c(\pi)^2 \right] &\leq \frac{1}{(1-\gamma)^2}.
 \end{aligned}$$

The same holds for $\hat{\nabla} V_{\pi_i}^c(\pi)$ w.r.t. $\nabla_{\pi_i} V_c(\pi)$.

Finally, the following lemma shows that our constraint function estimates concentrate around their mean.

Lemma 11. For $\pi \in \Pi^\xi$, let $\hat{V}_c^{(1)}, \dots, \hat{V}_c^{(B)}$ be independent copies of $\hat{V}_c(\pi)$, and let $\hat{V}_c := \frac{1}{B} \sum_{i=1}^B \hat{V}_c^{(i)}$. Then there exists $C > 0$ such that for any $\lambda \geq 0$,

$$\mathbb{P} \left(\left| \hat{V}_c - V_c(\pi) \right| > \frac{\lambda}{\sqrt{B}} \right) \leq 4 \exp(-C(1-\gamma)\lambda) + 2 \exp(-C^2(1-\gamma)^2 \lambda^2).$$

Proof. For $i \in [B]$, we decompose $\hat{V}_c^{(i)} = \hat{c}_0^{(i)} + \hat{V}_{c,\geq 1}^{(i)}$ where $\hat{c}_0^{(i)}$ is the cost incurred at step 0, and let $\hat{c}_0 := \frac{1}{B} \sum_{i=1}^B \hat{c}_0^{(i)}$, $\hat{V}_{c,\geq 1} := \frac{1}{B} \sum_{i=1}^B \hat{V}_{c,\geq 1}^{(i)}$. Since $0 \leq \hat{c}_0 \leq 1$, by Hoeffding's inequality, there exists $C_0 > 0$ such that for any $\lambda \geq 0$,

$$\mathbb{P} \left(|\hat{c}_0 - \mathbb{E}[\hat{c}_0]| > \frac{\lambda}{2\sqrt{B}} \right) \leq 2 \exp(-C_0 \lambda^2). \quad (17)$$

Moreover, note that since for all $s \in \mathcal{S}, a \in \mathcal{A}$, $0 \leq c(s, a) \leq 1$, we have that for all $i \in [B]$, $V_{c,\geq 1}^{(i)} \leq T_e^{(i)}$ where $T_e^{(i)}$ is the stopping time of the respective episode and an independent copy of T_e . Assuming $\kappa_{s,a} = \min_{s \in \mathcal{S}, a \in \mathcal{A}} \kappa_{s,a} = 1 - \gamma$ for all $s \in \mathcal{S}, a \in \mathcal{A}$, T_e follows a geometric distribution with parameter $1 - \gamma$. By definition of the geometric distribution and elementary computations, we get for any $\lambda \geq 0$,

$$\mathbb{P}(T_e \geq \lambda) \leq \gamma^{\lceil \lambda \rceil} \leq \exp(\lceil \lambda \rceil \log \gamma) \leq \exp\left(-\lceil \lambda \rceil \frac{1 - \gamma}{3}\right) \leq \exp(-(1 - \gamma)\lambda/3)$$

which by a standard characterization of sub-exponential random variables, see Proposition 2.7.1 in Vershynin (2018), implies that T_e , and therefore also $V_{c,\geq 1}^{(i)}$ for all $i \in [B]$ are sub-exponential. Moreover, by the so-called centering lemma for sub-exponential distributions, see section 2.7 in Vershynin (2018), for any random variable X that is sub-exponential with parameter σ , there exists an absolute constant c such that $X - \mathbb{E}[X]$ is sub-exponential with parameter $c\sigma$. Thus for all $i \in [B]$, $V_{c,\geq 1}^{(i)} - \mathbb{E}[\hat{V}_{c,\geq 1}^{(i)}]$ is sub-exponential with parameter in $\mathcal{O}(1/(1 - \gamma))$. Then, we apply Bernstein's inequality, see Theorem 2.8.1 in Vershynin (2018), to show that there exist $C_1, C_2 > 0$ such that for any $\lambda \geq 0$,

$$\mathbb{P} \left(\left| \hat{V}_{c,\geq 1} - \mathbb{E}[\hat{V}_{c,\geq 1}] \right| > \frac{\lambda}{2\sqrt{B}} \right) \leq 2 \exp(-C_1(1 - \gamma)\lambda) + 2 \exp(-C_2^2(1 - \gamma)^2 \lambda^2). \quad (18)$$

Finally, using a union bound we combine (17) and (18) to get the desired bound. \square

C STRONGLY CONVEX STOCHASTIC OPTIMIZATION WITH STRONGLY CONVEX EXPECTATION CONSTRAINT

In this section, we describe a stochastic gradient switching algorithm for stochastic constrained optimization under expectation constraints. Up to the modification of using a relaxed constraint (which is crucial for enabling its independent implementation), our algorithm and analysis follow the Cooperative Stochastic Approximation (CSA) algorithm presented in Lan and Zhou (2020) which is inspired by Polyak's subgradient method (Polyak, 1967). Lan and Zhou (2020) hint at the fact that a $1/K$ convergence rate of the CSA algorithm can be shown in the case of strongly convex objective and under expectation constraints. Here we explicitly carry out this analysis by deriving a result in expectation and under a somewhat weaker assumption on the distribution of the constraint function estimates.

Let $X \subseteq \mathbb{R}^d$ be a convex and compact set with diameter $\Delta := \max_{x, x' \in X} \|x - x'\|$. Suppose θ are random vectors supported on $\Theta \subset \mathbb{R}^p$, and let $F : X \times \Theta \rightarrow \mathbb{R}, G : X \times \Theta \rightarrow \mathbb{R}$ be functions such that $F(\cdot, \theta)$ and $G(\cdot, \theta)$ are μ_F and μ_G -weakly convex, respectively. For any $x' \in X$, we define $F_{\mu, x'}(x, \theta) := F(x, \theta) + \mu_F \|x - x'\|^2$ and $G_{\mu, x'}(x, \theta) := G(x, \theta) + \mu_G \|x - x'\|^2$. Let $f(x) := \mathbb{E}_\theta[F(x, \theta)], g(x) := \mathbb{E}_\theta[G(x, \theta)]$ (where expectations are supposed to be well-defined and finite) and $f_{\mu, x'}(x) := f(x) + \mu_F \|x - x'\|^2, g_{\mu, x'}(x) := g(x) + \mu_G \|x - x'\|^2$ for every $x \in X$.

The problem we aim to solve⁹ is given by

$$\begin{aligned} \min_{x \in X} f_{\mu, x'}(x) &:= \mathbb{E}_\theta[F_{\mu, x'}(x, \theta)] \\ \text{s.t. } g_{\mu, x'}(x) &:= \mathbb{E}_\theta[G_{\mu, x'}(x, \theta)] \leq 0. \end{aligned} \quad (19)$$

Recall that such a problem needs to be solved at each time step in our iProxCMPG algorithm. The point x' is arbitrarily fixed throughout the rest of this section.

Suppose we are only given access to first-order information of $f_{\mu, x'}, g_{\mu, x'}$ and zeroth-order information of g via a stochastic oracle that outputs unbiased and bounded-variance estimates.

⁹Note that the final guarantees we will obtain are actually in terms of a relaxed constraint satisfaction bound. This is due to our modification of the original CSA algorithm.

Assumption 3. For every $x \in X$, the estimators $F'_{\mu,x'}(x,\theta), G'_{\mu,x'}(x,\theta)$ and $G(x,\theta)$ are unbiased, i.e., $\mathbb{E}_\theta [F'_{\mu,x'}(x,\theta)] = \nabla f_{\mu,x'}(x), \mathbb{E}_\theta [G'_{\mu,x'}(x,\theta)] = \nabla g_{\mu,x'}(x)$ and $\mathbb{E}_\theta [G(x,\theta)] = g(x)$. Moreover, there exist $M_F, M_G > 0$ such that

$$\mathbb{E}_\theta \left[\|F'(x,\theta)\|^2 \right] \leq M_F^2; \quad \mathbb{E}_\theta \left[\|G'(x,\theta)\|^2 \right] \leq M_G^2.$$

Furthermore, we suppose to have access to independent unbiased estimators $\hat{G}^{(1)}, \dots, \hat{G}^{(J)}$ of $G(x, \cdot)$ for which there exists $\sigma > 0$ such that for any $\lambda \geq 0$, it holds that

$$\mathbb{P}_\theta \left(|\hat{G} - g(x)| > \lambda/\sqrt{J} \right) \leq 4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2), \quad (20)$$

where $\hat{G} := \frac{1}{J} \sum_{j=1}^J \hat{G}^{(j)}$.

It can be easily seen that Assumption 3 also implies existence of \tilde{M}_F, \tilde{M}_G such that

$$\begin{aligned} \mathbb{E}_\theta \left[\|F'_{\mu,x'}(x,\theta)\|^2 \right] &\leq 2\mathbb{E}_\theta \left[\|F'(x,\theta)\|^2 \right] + 2\mu_F^2 \|x - x'\|^4 \leq 2M_F^2 + 2\mu_F^2 \Delta^4 =: \tilde{M}_F^2, \\ \mathbb{E}_\theta \left[\|G'_{\mu,x'}(x,\theta)\|^2 \right] &\leq 2\mathbb{E}_\theta \left[\|G'(x,\theta)\|^2 \right] + 2\mu_G^2 \|x - x'\|^4 \leq 2M_G^2 + 2\mu_G^2 \Delta^4 =: \tilde{M}_G^2. \end{aligned}$$

Remark 5. Notice that the concentration requirement of (20) is relaxed compared to the sub-Gaussian assumption made in Lan and Zhou (2020) which is too strong to hold in our case. We refer the reader to Lemma 11 where we prove that this weaker tail bound assumption holds for our constraint function estimates.

C.1 A Primal Gradient Switching Algorithm

Algorithm 3 is designed as a primal algorithm that switches between taking a step along the objective or constraint gradient, depending on whether the constraint is currently (estimated to be) satisfied or not.

Algorithm 3 CSA (adapted from Lan and Zhou (2020))

- 1: **initialization:** $x_1 \in X$ s.t. $g(x_1) \leq \epsilon$ and $\{\delta_k\}_{k \in [N]}, \{\nu_k\}_{k \in [N]}, \{\rho_k\}_{k \in [N]}, s \in [N]$
 - 2: **for** $k = 1, \dots, N - 1$ **do**
 - 3: sample $\hat{G}_k^{(1)}, \dots, \hat{G}_k^{(J)}$ from $G(x_k, \cdot)$ and set $\hat{G}_k = \frac{1}{J} \sum_{j=1}^J \hat{G}_k^{(j)}$
 - 4: $x_{k+1} = \begin{cases} \mathcal{P}_X [x_k - \nu_k F'_{\mu,x'}(x_k, \theta_k)] & \text{if } \hat{G}_k \leq \delta_k \\ \mathcal{P}_X [x_k - \nu_k G'_{\mu,x'}(x_k, \theta_k)] & \text{else} \end{cases}$
 - 5: let $\mathcal{B}_s := \{s \leq k \leq N \mid \hat{G}_k \leq \delta_k\}$
 - 6: **output:** $x_{\hat{k}}$ where $\hat{k} = 1$ if $\mathcal{B}_s = \emptyset$ and otherwise sampled s.t. for $k \in \mathcal{B}_s, \mathbb{P}(\hat{k} = k) = (\sum_{k \in \mathcal{B}_s} \rho_k)^{-1} \rho_k$
-

In the analysis, we will denote $\mathcal{M}_s := \{s \leq k \leq N \mid k \notin \mathcal{B}_s\}$ and $\mathcal{B} := \mathcal{B}_1, \mathcal{M} := \mathcal{M}_1$.

We point out the following differences between Algorithm 3 and the original CSA algorithm, see Lan and Zhou (2020), Algorithm 1.

- (a) We relax the switching condition in line 4 by using an estimate of $g(x_k)$ instead of $g_{\mu,x'}(x_k)$ if we were to exactly use the algorithm proposed in Lan and Zhou (2020). This modification is crucial for our application as subroutine of an *independent* learning algorithm, as described in the proof of Lemma 1, see section B.1. As a result, compared to Lan and Zhou (2020), we get a weaker guarantee in terms of constraint violation which however is still sufficient for our purposes.
- (b) Instead of constructing the output as a ρ_k -weighted average over iterates x_k , we sample an iterate from a ρ_k -weighted distribution, see line 6. This is because our relaxed constraint function g is not necessarily convex (unlike $g_{\mu,x'}$) and hence we cannot easily bound the constraint value at an average over iterates.

C.2 Convergence and Sample Complexity Guarantee

The following analysis uses the techniques presented in [Lan and Zhou \(2020\)](#) applied to the strongly convex case with expectation constraint, under our modified Assumption 3 and Algorithm 3. The proofs follow along the same lines, we highlight differences when appropriate.

First, we establish a basic recursion about CSA iterates that will be used repeatedly throughout the rest of the analysis.

Proposition 1. *For any $s \in [N]$, $x \in X$, and a_s as defined by (21), it holds that*

$$\begin{aligned} & \sum_{k \in \mathcal{M}_s} \rho_k (G_{\mu, x'}(x_k, \theta_k) - G_{\mu, x'}(x, \theta_k)) + \sum_{k \in \mathcal{B}_s} \rho_k (F_{\mu, x'}(x_k, \theta_k) - F_{\mu, x'}(x, \theta_k)) \\ & \leq (1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \|F'_{\mu, x'}(x_k, \theta_k)\|^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \|G'_{\mu, x'}(x_k, \theta_k)\|^2. \end{aligned}$$

Proof. Let $s \in [N]$ and $k \in \mathcal{B}_s$. Then, by non-expansiveness of the projection \mathcal{P}_X and strong convexity,

$$\begin{aligned} \|x_{k+1} - x\|^2 & \leq \|x_k - x\|^2 - \nu_k \langle F'_{\mu, x'}(x_k, \theta_k), x_k - x \rangle + \frac{1}{2} \nu_k^2 \|F'_{\mu, x'}(x_k, \theta_k)\|^2 \\ & \leq \|x_k - x\|^2 - \nu_k \left[F_{\mu, x'}(x_k, \theta_k) - F_{\mu, x'}(x, \theta_k) + \frac{\mu_F}{2} \|x_k - x\|^2 \right] + \frac{1}{2} \nu_k^2 \|F'_{\mu, x'}(x_k, \theta_k)\|^2 \\ & \leq \left(1 - \frac{\nu_k \mu_F}{2} \right) \|x_k - x\|^2 - \nu_k [F_{\mu, x'}(x_k, \theta_k) - F_{\mu, x'}(x, \theta_k)] + \frac{1}{2} \nu_k^2 \|F'_{\mu, x'}(x_k, \theta_k)\|^2. \end{aligned}$$

Similarly, if $k \in \mathcal{M}_s$,

$$\|x_{k+1} - x\|^2 \leq \left(1 - \frac{\nu_k \mu_G}{2} \right) \|x_k - x\|^2 - \nu_k [G_{\mu, x'}(x_k, \theta_k) - G_{\mu, x'}(x, \theta_k)] + \frac{1}{2} \nu_k^2 \|G'_{\mu, x'}(x_k, \theta_k)\|^2.$$

After defining

$$a_k = \begin{cases} \mu_F \nu_k & \text{if } k \in \mathcal{B} \\ \mu_G \nu_k & \text{if } k \in \mathcal{M} \end{cases}; \quad A_k = \begin{cases} 1 & \text{if } k = 1 \\ (1 - a_k) A_{k-1} & \text{if } k \geq 2 \end{cases}; \quad \rho_k = \frac{\nu_k}{A_k}; \quad (21)$$

the result follows by application of Lemma 21, [Lan and Zhou \(2020\)](#). \square

The next lemma provides a condition on $\{\nu_k, \delta_k, \rho_k\}_{s \leq k \leq N}$ that guarantees either low regret in terms of objective value or that a large number of iterates satisfy the constraint with high probability.

Lemma 12. *Let x^* be an optimal solution of (19). If for some $s \in [N]$ and $\lambda \geq 0$,*

$$\frac{N - s + 1}{2} \min_{k \in \mathcal{M}_s} \rho_k \delta_k > (1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2 + \frac{\lambda}{\sqrt{J}} \sum_{k \in \mathcal{M}_s} \rho_k, \quad (22)$$

then one of the following statements hold,

- (a) $\mathbb{P}_\theta(|\mathcal{B}_s| \geq (N - s + 1)/2) \geq 1 - |\mathcal{M}_s| (4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2))$, or,
- (b) $\sum_{k \in \mathcal{B}_s} \rho_k (f_{\mu, x'}(x_k) - f_{\mu, x'}(x^*)) \leq 0$.

Note that unlike in [Lan and Zhou \(2020\)](#), due to our modified choice of Algorithm 3's output, well-definedness of x_k does not require $\mathcal{B}_s \neq \emptyset$.

Proof. In Proposition 1, set $x = x^*$, take expectation w.r.t. θ on both sides, and apply $\mathbb{E}_\theta \|F'_{\mu, x'}(x, \theta)\|^2 \leq \tilde{M}_F^2$, $\mathbb{E}_\theta \|G'_{\mu, x'}(x, \theta)\|^2 \leq \tilde{M}_G^2$. Then,

$$\begin{aligned} & \sum_{k \in \mathcal{M}_s} \rho_k (g_{\mu, x'}(x_k) - g_{\mu, x'}(x^*)) + \sum_{k \in \mathcal{B}_s} \rho_k (f_{\mu, x'}(x_k) - f_{\mu, x'}(x^*)) \\ & \leq (1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2. \end{aligned} \quad (23)$$

If $\sum_{k \in \mathcal{B}_s} \rho_k (f_{\mu, x'}(x_k) - f_{\mu, x'}(x^*)) \leq 0$, then (b) holds. Otherwise, we make three observations. First, we have that $g_{\mu, x'}(x^*) \leq 0$. Second, it holds that $g(x_k) \leq g_{\mu, x'}(x_k)$. Third, for $k \in \mathcal{M}_s$, by Assumption 3 and due to $\hat{G}_k > \delta_k$, we get

$$\mathbb{P}_\theta \left(g(x_k) < \delta_k - \frac{\lambda}{\sqrt{J}} \right) \leq 4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2). \quad (24)$$

By a union bound this inequality holds for all $k \in \mathcal{M}_s$ with probability at most $|\mathcal{M}_s| (4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2))$. Combining these three observations with (23) yields that with probability at least $1 - |\mathcal{M}_s| (4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2))$, it holds that

$$\sum_{k \in \mathcal{M}_s} \rho_k \delta_k \leq (1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2 + \frac{\lambda}{\sqrt{J}}.$$

Above inequality then implies (a) because if $|\mathcal{B}_s| < (N - s + 1)/2$, i.e., $|\mathcal{M}_s| \geq (N - s + 1)/2$, then condition (22) implies that

$$\sum_{k \in \mathcal{M}_s} \rho_k \delta_k \geq \frac{N - s + 1}{2} \min_{k \in \mathcal{M}_s} \rho_k \delta_k > (1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2 + \frac{\lambda}{\sqrt{J}} \sum_{k \in \mathcal{M}_s} \rho_k,$$

which is a contradiction. \square

Next, we state and prove the main guarantees provided by Algorithm 3.

Theorem 3. *Under Assumption 3, let $\epsilon > 0$, suppose x_1 is such that $g(x_1) \leq \epsilon$, and let $f_{\max} > 0$ such that for all $x \in X$, $0 \leq f(x) \leq f_{\max}$. Choose $s = N/2, \lambda = \sigma^2 \log(N^2/(4f_{\max}))$, set $M = \max\{\tilde{M}_G, \tilde{M}_F\}$, $\mu = \min\{\mu_G, \mu_F\}$, and*

$$\begin{aligned} \nu_k &= \begin{cases} \frac{2}{\mu_F^{(k+1)}} & \text{if } k \in \mathcal{B} \\ \frac{2}{\mu_G^{(k+1)}} & \text{if } k \in \mathcal{M} \end{cases}; & \delta_k &= \frac{\lambda}{\sqrt{J}} + \frac{1}{2k} \left(\frac{4\Delta^2}{k} + \frac{16M^2}{\mu^2} \right) \cdot \begin{cases} \mu_F & \text{if } k \in \mathcal{B} \\ \mu_G & \text{if } k \in \mathcal{M} \end{cases}; \\ a_k &= \begin{cases} \mu_F \nu_k & \text{if } k \in \mathcal{B} \\ \mu_G \nu_k & \text{if } k \in \mathcal{M} \end{cases}; & A_k &= \begin{cases} 1 & \text{if } k = 1 \\ (1 - a_k) A_{k-1} & \text{if } k \geq 2 \end{cases}; & \rho_k &= \frac{\nu_k}{A_k} \\ N &= \max \left\{ \frac{64\mu_F M^2}{\mu^2 \epsilon^2}, \frac{\sqrt{32\Delta^2 \mu_F}}{\epsilon}, \frac{32\sigma \mu_F}{\mu \epsilon^2} \right\}; & J &= \max \left\{ \frac{9\lambda^2}{\epsilon^2}, \frac{32\sigma \mu_F}{\mu \epsilon^2} \right\}. \end{aligned}$$

Then Algorithm 3 guarantees that

$$\mathbb{E} [f_{\mu, x'}(x_{\hat{k}}) - f_{\mu, x'}(x^*)] \leq \epsilon^2, \quad (25)$$

$$\mathbb{E} [g(x_{\hat{k}})] \leq \epsilon. \quad (26)$$

Proof. First, we observe that for any $k \in \mathcal{M}_s$,

$$\begin{aligned} \mathbb{E} \left[\sqrt{J} (g(x_k) - \delta_k) \right] &= \int_0^\infty \left(1 - \mathbb{P} \left(\sqrt{J} (g(x_k) - \delta_k) \leq z \right) \right) dz \\ &\quad - \int_{-\infty}^0 \mathbb{P} \left(\sqrt{J} (g(x_k) - \delta_k) \leq z \right) dz \\ &\geq - \int_{-\infty}^0 4 \exp(z/\sigma) + 2 \exp(z^2/\sigma^2) dz \\ &\geq -6\sigma \end{aligned} \quad (27)$$

where the first inequality is by (24). Therefore, we have $\mathbb{E}[g(x_k)] \geq \delta_k - \frac{6\sigma}{\sqrt{J}}$. Moreover, by an argument similar to our derivation in Lemma 11 but with Bernstein's inequality applied to the sum $(\sum_{k \in \mathcal{M}_s} \rho_k)^{-1} \sum_{k \in \mathcal{M}_s} \rho_k \hat{G}_k$,

$$\mathbb{P} \left(\sum_{k \in \mathcal{M}_s} \rho_k (\hat{G}_k - g(x_k)) > \frac{\lambda}{\sqrt{J} |\mathcal{M}_s|} \sum_{k \in \mathcal{M}_s} \rho_k \right) \leq 4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2).$$

Therefore, following (27), we get

$$\mathbb{E} \left[\sum_{k \in \mathcal{M}_s} \rho_k g(x_k) \right] \geq \sum_{k \in \mathcal{M}_s} \rho_k \delta_k - \frac{6\sigma}{\sqrt{J|\mathcal{M}_s|}} \sum_{k \in \mathcal{M}_s} \rho_k. \quad (28)$$

Next, we derive (25). Note that (22) holds for our choices of $s, \nu_k, \delta_k, \rho_k$. Then, if part (b) of Lemma 12 holds, we have

$$\begin{aligned} \mathbb{E} [f(x_{\hat{k}}) - f(x^*)] &= \mathbb{E}_{\hat{k}} \left[\mathbb{E} [f(x_k) - f(x^*) \mid \hat{k} = k] \right] \\ &\leq \left(\sum_{k \in \mathcal{B}_s} \rho_k \right)^{-1} \sum_{k \in \mathcal{B}_s} \rho_k \mathbb{E} [f(x_k) - f(x^*)] \\ &\leq 0. \end{aligned}$$

Otherwise, if part (a) holds, then using above bound on $\mathbb{E}[g(x_k)]$ together with convexity of $f_{\mu, x'}$, (23) and (28), it follows that

$$\begin{aligned} &\sum_{k \in \mathcal{M}_s} \rho_k \delta_k - \frac{6\sigma}{\sqrt{J|\mathcal{M}_s|}} \sum_{k \in \mathcal{M}_s} \rho_k + \sum_{k \in \mathcal{B}_s} \rho_k \mathbb{E} [f_{\mu, x'}(x_{\hat{k}}) - f_{\mu, x'}(x^*)] \\ &\leq \sum_{k \in \mathcal{M}_s} \rho_k \mathbb{E} [g(x_k)] + \sum_{k \in \mathcal{B}_s} \rho_k \mathbb{E} [f_{\mu, x'}(x_{\hat{k}}) - f_{\mu, x'}(x^*)] \\ &\leq \sum_{k \in \mathcal{M}_s} \rho_k \mathbb{E} [g(x_k)] + \sum_{k \in \mathcal{B}_s} \rho_k \mathbb{E} [f_{\mu, x'}(x_k) - f_{\mu, x'}(x^*)] \\ &\leq (1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2. \end{aligned}$$

Denote by $E_{\mathcal{B}_s}$ the event that $|\mathcal{B}_s| \geq (N - s + 1)/2$. Then, using the law of total expectation, our choice of $\lambda = \sigma^2 \log(N^2/(4f_{\max}))$, $\rho_k \delta_k \geq 0$, and above inequality, we have

$$\begin{aligned} &\mathbb{E} [f(x_{\hat{k}}) - f(x^*)] \\ &\leq \mathbb{E} [f(x_{\hat{k}}) - f(x^*) \mid E_{\mathcal{B}_s}] \cdot \underbrace{\mathbb{P}(E_{\mathcal{B}_s})}_{\leq 1} + \mathbb{E} [f(x_{\hat{k}}) - f(x^*) \mid \bar{E}_{\mathcal{B}_s}] \cdot \underbrace{\mathbb{P}(\bar{E}_{\mathcal{B}_s})}_{\leq |\mathcal{M}_s| (4 \exp(-\lambda/\sigma) + 2 \exp(-\lambda^2/\sigma^2))} \\ &\leq \left(\sum_{k \in \mathcal{B}_s} \rho_k \right)^{-1} \left((1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2 + \frac{6\sigma}{\sqrt{J|\mathcal{M}_s|}} \sum_{k \in \mathcal{M}_s} \rho_k \right) + \frac{1}{N} \\ &\leq \left(\frac{N - s + 1}{2} \min_{k \in \mathcal{B}_s} \rho_k \right)^{-1} \left((1 - a_s) \Delta^2 + \frac{1}{2} \sum_{k \in \mathcal{M}_s} \rho_k \nu_k \tilde{M}_G^2 + \frac{1}{2} \sum_{k \in \mathcal{B}_s} \rho_k \nu_k \tilde{M}_F^2 + \frac{6\sigma}{\sqrt{J|\mathcal{M}_s|}} \sum_{k \in \mathcal{M}_s} \rho_k \right) + \frac{1}{N}. \end{aligned}$$

In order to show the constraint violation bound, note that by a similar argument as (27), for any $k \in \mathcal{B}_s$, $\mathbb{E}[g(x_k)] \leq \delta_k + \frac{6\sigma}{\sqrt{J}}$, and therefore

$$\mathbb{E} [g(x_{\hat{k}})] = \mathbb{E}_{\hat{k}} \left[\mathbb{E} [g(x_k) \mid k = \hat{k}] \right] \leq \mathbb{E}_{\hat{k}} [\delta_{\hat{k}}] + \frac{6\sigma}{\sqrt{J}} = \frac{\sum_{k \in \mathcal{B}_s} \rho_k \delta_k}{\sum_{k \in \mathcal{B}_s} \rho_k} + \frac{6\sigma}{\sqrt{J}}.$$

In order to derive the guarantees (25) and (26), we plug in the choices of $\nu_k, \delta_k, a_k, \rho_k, N$, and J stated in Theorem 3. Observing that for any $s \leq k \leq N$, we have $A_k = \frac{2}{k(k+1)}$, that for any $k \in \mathcal{B}$, we have $\rho_k = \frac{2k}{\mu_F}$ as

well as $\rho_k \nu_k = \frac{4}{\mu_F^2}$, and for any $k \in \mathcal{M}$, $\rho_k = \frac{2k}{\mu_G}$, $\rho_k \nu_k = \frac{4}{\mu_G^2}$.

$$\begin{aligned} \mathbb{E} [f(x_{\hat{k}}) - f(x^*)] &\leq \frac{\Delta^2 + 2N\mu_F^{-2}\tilde{M}_F^2 + 2N\mu_G^{-2}\tilde{M}_G^2 + \frac{2\sigma}{\sqrt{J}}N^{3/2}\mu_G^{-1}}{N^2/4 \cdot \mu_F^{-1}} \\ &\leq \frac{\Delta^2 + 4N\mu^{-2}M^2 + \frac{2\sigma}{\sqrt{J}}N^{3/2}\mu^{-1}}{N^2/4 \cdot \mu_F^{-1}} + \frac{1}{N} \\ &\leq \frac{4\mu_F\Delta^2}{N^2} + \frac{16\mu_F\mu^{-2}M^2}{N} + \frac{8\sigma\mu^{-1}\mu_F}{\sqrt{JN}} + \frac{1}{N} \\ &\leq \epsilon^2/4 + \epsilon^2/4 + \epsilon^2/4 + \epsilon^2/4. \end{aligned}$$

Moreover, for the constraint bound, it holds that

$$\begin{aligned} \mathbb{E} [g(x_{\hat{k}})] &\leq \frac{\sum_{k \in \mathcal{B}_s} (4\Delta^2/k + 16M^2/\mu^2)}{\sum_{k \in \mathcal{B}_s} 2k/\mu_F} + \frac{6\sigma}{\sqrt{J}} \\ &\leq \frac{8\Delta^2\mu_F}{N^2} + \frac{16M^2\mu_F}{\mu^2 N} + \frac{6\sigma}{\sqrt{J}} \\ &\leq \epsilon. \end{aligned}$$

□

D BACKGROUND IN CONSTRAINED OPTIMIZATION AND A NOVEL TECHNICAL LEMMA

Notation For any non-empty subset $Y \subset \mathbb{R}^d$ and any vector $x \in \mathbb{R}^d$, the distance from x to the set Y is defined as $\text{dist}(x, Y) := \inf_{y \in Y} \|x - y\|$ where $\|\cdot\|$ is the standard 2-norm of the Euclidean space \mathbb{R}^d .

In this section, we recall some useful definitions for constrained optimization. In particular, we recall the definition of an approximate Karush-Kuhn-Tucker (KKT) point and a variation thereof. Then we prove a new technical result that will be useful in our analysis.

D.1 Approximate KKT Points in Constrained Optimization

Let $X \subset \mathbb{R}^d$ be a closed convex set. Consider the following constrained optimization problem:

$$\begin{aligned} P^* &= \min_{x \in X} f(x) \\ &\text{s.t. } f_c(x) \leq 0, \end{aligned} \tag{ConstrOpt}$$

where $f, f_c : X \rightarrow \mathbb{R}$ are differentiable (possibly nonconvex) functions.

The associated Lagrangian function $\mathcal{L} : X \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is defined for any $x \in X, \lambda \geq 0$ by $\mathcal{L}(x, \lambda) = f(x) + \lambda f_c(x)$. The primal and dual problems can be respectively written as

$$\begin{aligned} P^* &= \inf_{x \in X} \sup_{\lambda \geq 0} \mathcal{L}(x, \lambda), \\ D^* &= \sup_{\lambda \geq 0} \underbrace{\inf_{x \in X} \mathcal{L}(x, \lambda)}_{=: d(\lambda)}. \end{aligned}$$

By weak duality, we know that $P^* \geq D^*$.

For any $x \in \mathbb{R}^d$, the normal cone to the set X at x is defined by:

$$N_X(x) := \{g \in \mathbb{R}^d \mid \forall y \in X, \langle g, y - x \rangle \leq 0\}.$$

Definition 1. Let $\epsilon \geq 0$. A point $x \in X$ is an ϵ -KKT point of (ConstrOpt) if there exists a real λ such that the following conditions hold:

$$\begin{aligned} f_c(x) &\leq \epsilon, && \text{(primal feasibility)} \\ \lambda &\geq 0, && \text{(dual feasibility)} \\ |\lambda f_c(x)| &\leq \epsilon, && \text{(complementary slackness)} \\ \text{dist}(\nabla_x \mathcal{L}(x, \lambda), -N_X(x)) &\leq \epsilon. && \text{(Lagrangian stationarity)} \end{aligned}$$

We also call (x, λ) an ϵ -KKT pair. The point x is simply a KKT point of (ConstrOpt) if moreover $\epsilon = 0$.

We additionally define a slight modification of the above standard KKT conditions which turns out to be useful in our analysis. More precisely, the definition replaces approximate Lagrangian stationarity by a variational form thereof. Moreover, primal feasibility is now supposed to be exact. Other conditions remain unchanged.

Definition 2. Let $\epsilon \geq 0$. A point $\tilde{x} \in X$ is an ϵ - $\widetilde{\text{KKT}}$ point of (ConstrOpt) if there exists a real $\tilde{\lambda}$ such that the following conditions hold:

$$\begin{aligned} f_c(\tilde{x}) &\leq 0, && \text{(exact primal feasibility)} \\ \tilde{\lambda} &\geq 0, && \text{(dual feasibility)} \\ |\tilde{\lambda} f_c(\tilde{x})| &\leq \epsilon, && \text{(complementary slackness)} \\ \max_{x' \in X} \langle \tilde{x} - x', \nabla_x \mathcal{L}(\tilde{x}, \tilde{\lambda}) \rangle &\leq \epsilon. && \text{(variational Lagrangian stationarity)} \end{aligned}$$

In particular, the point \tilde{x} is said to be a $\widetilde{\text{KKT}}$ point of (ConstrOpt) when $\epsilon = 0$.

The next lemma connects the first stationarity condition with a variational form thereof. In particular, this result allows to connect the two definitions of approximate KKT points above.

Lemma 13. Let $X \subseteq \mathbb{R}^d$ be a convex and compact set. Let $\epsilon > 0$ and let $x, g \in \mathbb{R}^d$. If $\text{dist}(g, -N_X(x)) \leq \epsilon$, then $\max_{x' \in X} \langle x - x', g \rangle \leq \Delta \epsilon$, where $\Delta := \max_{x, x' \in X} \|x - x'\|$ is the diameter of the set X .

Proof. Let $y_0 \in -N_X(x)$. For any $x' \in X$, we have

$$\begin{aligned} \langle x - x', g \rangle &= \langle x - x', g - y_0 \rangle + \langle -y_0, x' - x \rangle, \\ &\leq \langle x' - x, g - y_0 \rangle, \\ &\leq \|x' - x\| \cdot \|g - y_0\|, \end{aligned}$$

where the first inequality follows from the fact that $y_0 \in -N_X(x)$, the second inequality stems from the Cauchy-Schwarz inequality. Taking the infimum with respect to y_0 in the last inequality gives the desired inequality since $\text{dist}(g, -N_X(x)) = \inf_{y \in -N_X(x)} \|g - y\| \leq \epsilon$. \square

D.2 A Novel Technical Lemma for Approximate Optimality under Gradient Dominance

We now state our technical lemma. This results shows that an approximate KKT point of (ConstrOpt) at which a gradient domination inequality holds for the Lagrangian function is approximately optimal for the objective function to be minimized.

Proposition 2. Let $\epsilon > 0$ and let $\tilde{x} \in X$ be an ϵ - $\widetilde{\text{KKT}}$ point of (ConstrOpt) . Suppose there exist constants $C_0, C_1 \geq 0$ such that the Lagrangian function associated to (ConstrOpt) satisfies for all $\lambda \geq 0$ and for all $x \in X$,

$$\mathcal{L}(x, \lambda) - \mathcal{L}(x_\lambda^*, \lambda) \leq C_0 \max_{x' \in X} \langle x - x', \nabla_x \mathcal{L}(\tilde{x}, \tilde{\lambda}) \rangle + C_1 \epsilon, \quad (29)$$

where x_λ^* is a minimizer of $\mathcal{L}(\cdot, \lambda)$. Then, we have

$$f(\tilde{x}) - P^* \leq (C_0 + C_1 + 1)\epsilon.$$

Proof. Let $(\tilde{x}, \tilde{\lambda})$ be an ϵ -KKT pair. Then, we have

$$\begin{aligned}
 D^* &\stackrel{(a)}{=} \max_{\lambda \geq 0} d(\lambda) \geq d(\tilde{\lambda}) \\
 &\stackrel{(b)}{=} \min_{x \in X} \mathcal{L}(x, \tilde{\lambda}) \\
 &\stackrel{(c)}{\geq} \mathcal{L}(\tilde{x}, \tilde{\lambda}) - (C_0 + C_1)\epsilon \\
 &= f(\tilde{x}) + \tilde{\lambda} f_c(\tilde{x}) - (C_0 + C_1)\epsilon \\
 &\stackrel{(d)}{\geq} f(\tilde{x}) - (C_0 + C_1 + 1)\epsilon
 \end{aligned}$$

where (a) and (b) are by definition, and (d) is due to complementary slackness. To see (c), observe that by Lagrangian stationarity and (29),

$$\epsilon \geq \max_{x' \in X} \left\langle \tilde{x} - x', \nabla_x \mathcal{L}(\tilde{x}, \tilde{\lambda}) \right\rangle \geq \frac{1}{C_0} \left(\mathcal{L}(\tilde{x}, \tilde{\lambda}) - \mathcal{L}(x_\lambda^*, \tilde{\lambda}) - C_1 \epsilon \right),$$

which implies that¹⁰

$$\mathcal{L}(\tilde{x}, \tilde{\lambda}) - \mathcal{L}(x_\lambda^*, \tilde{\lambda}) \leq (C_0 + C_1)\epsilon.$$

Finally, we use weak duality, i.e. $P^* \geq D^*$, to conclude that

$$f(\tilde{x}) - P^* = \underbrace{f(\tilde{x}) - D^*}_{\leq (C_0 + C_1 + 1)\epsilon} + \underbrace{D^* - P^*}_{\leq 0} \leq (C_0 + C_1 + 1)\epsilon.$$

□

E ADDITIONAL DETAILS ABOUT SIMULATIONS

We provide additional details regarding the implementation of our iProxCMPG (Algorithm 2) in practice:

- (a) In our experiments, each episode terminates after a fixed number of steps $T_e = 10$ corresponding to a discount factor $\gamma = 0.9$.
- (b) In order to reduce the variance and enable the usage of larger step sizes, all constraint and value (gradient) estimates are obtained by sampling a batch of B trajectories.
- (c) For the subroutine, i.e. as solution to the proximal-point update, we do not consider a ρ_k -weighted average over iterates but simply use the last iterate $\pi^{(t, K)}$.
- (d) We choose $\delta_k = 0$ for all $k \in \mathbb{N}$.

Hyperparameters We report hyperparameter choices for our simulations in Table 2. Note that to ensure convergence, as indicated by our theoretical results, a larger number of players m requires smaller step sizes and larger sample batches. Step sizes η and ν_k were chosen by tuning over the range $[0, 1]$.

Error Bars and Reproducibility The plots in Figs. 1 and 2 show the means of estimated potential values across 10 independent runs, and the corresponding shaded region displays the respective standard deviation. Obtaining results for all presented experiments thus requires simulating 60 runs in total. All experiments are fully reproducible using the provided code and specified seeds.

Computing Infrastructure In order to reduce computation time by executing all runs in parallel, we conducted the simulations within less than 4 hours on a cluster of 15 4-core *Intel(R) Xeon(R) CPU E3-1284L v4* clocked at 2.90GHz and equipped with 8Gbs of memory.

¹⁰If $C_0 = 0$, the same inequality immediately holds from (29).

Table 2: Overview of hyperparameters used in our simulations.

Hyperparameters	Number of players m	Pollution tax	Energy marketplace
Step size η (outer loop)	-	0.1	0.1
	2	0.005	0.002
Step size ν_k (inner loop)	4	0.002	0.001
	8	0.0007	0.0003
Sample batch size B	2	1000	100
	4	1000	150
	8	2500	200
K (#iterations inner loop)	-	20	20
T (#iterations outer loop)	-	20	60
Discount factor γ	-	0.9	0.9
Episode length T_e	-	10	10

Notation As used in the main part, $\mathcal{U}(\{1, \dots, W\})$ refers to the uniform distribution over the finite set $\{1, \dots, W\}$ where $W \geq 2$ is an integer.