

A Scenario Approach to the Robustness of Nonconvex–Nonconcave Minimax Problems[★]

Huan Peng^{*} Guanpu Chen^{*} Karl H. Johansson^{*}

^{*} *School of Electrical Engineering and Computer Science,
KTH Royal Institute of Technology, SE-100 44, Stockholm, Sweden
e-mail: {huanp, guanpu, kallegj}@kth.se*

Abstract: This paper investigates probabilistic robustness of nonconvex–nonconcave minimax problems via the scenario approach. Inspired by recent advances in scenario optimization (Garatti and Campi, 2025), we obtain robustness results for key equilibria with nonconvex–nonconcave payoffs, overcoming the dependence on the non-degeneracy assumption. Specifically, under convex strategy sets for all players, we first establish a probabilistic robustness guarantee for an ε -stationary point by proving the monotonicity of the stationary residual in the number of scenarios. Moreover, under nonconvex strategy sets for all players, we derive a probabilistic robustness guarantee for a global minimax point by invoking the extreme value theorem and Berge’s maximum theorem. A numerical experiment on a unit commitment problem corroborates our theoretical findings.

Keywords: Uncertainty, robustness, scenario approach, nonconvex, minimax problem

1. INTRODUCTION

Minimax problems model a two-player zero-sum game where one player aims to minimize a payoff function while the other aims to maximize it. This structure underpins a vast array of important applications, from optimal control (Başar and Bernhard, 2008) to machine learning (Sutton and Barto, 2018). For the convex case, the minimax problem is well-studied (von Neumann, 1928; Fan, 1953; Sion, 1958). Recent advances in machine learning have spurred growing interest in nonconvex minimax problems. Prominent examples include generative adversarial networks (GANs) (Goodfellow et al., 2014) and adversarial training (Madry et al., 2018), which can often be formulated as nonconvex minimax problems.

Uncertainty affects the decision-making process in many applications. The scenario approach provides a data-driven methodology for addressing uncertainty: it yields probabilistic robustness guarantees without requiring prior knowledge of the probability distribution or geometric structure of the uncertainty set. Building on this approach, a substantial body of work investigated uncertain convex games, e.g., Fele and Margellos (2020); Pantazis et al. (2024); Chen et al. (2025). Most existing analyses, however, critically rely on the assumption of convexity; when this assumption is violated, the resulting bounds may not be directly applied. To address this limitation, Campi et al. (2015, 2018) first studied nonconvex scenario optimization, while Garatti and Campi (2025) subsequently achieved significantly tighter risk bounds across a broad class of

problems, albeit with more technically involved derivations.

The robustness of minimax problems has typically been studied under limited settings. While Carè et al. (2015) addressed the convex case and Assif et al. (2020); Garatti and Campi (2025) extended the analysis to nonconvex cases, research has predominantly focused on probabilistic robustness with uncertainty confined to the max-player’s strategy set. A critical gap remains, namely, when the strategy sets of both players are affected by uncertainty. Addressing this problem, particularly when the payoff is nonconvex–nonconcave with respect to the two players’ strategies, is relevant and mathematically challenging. Moreover, a Nash equilibrium (NE) in this case may fail to exist or may lack standard well-posedness properties. Therefore, the robustness of equilibria under these conditions remains an open question.

In this paper, we consider nonconvex–nonconcave minimax problems in which both players’ strategy sets are subject to uncertainty. The primary challenge lies in establishing tight probabilistic robustness guarantees without the non-degeneracy assumption. We address this challenge by leveraging the latest developments in nonconvex scenario optimization (Garatti and Campi, 2025). In this way, we characterize the robustness of key equilibria under both convex and nonconvex strategy sets of the players.

The main contributions of this paper are twofold. Firstly, we establish the robustness of an ε -stationary point, the first-order condition for a local NE, under nonconvex–nonconcave payoffs and convex strategy sets. By demonstrating that the stationary residual is non-increasing as the number of scenarios increases, we derive a tight probabilistic robustness guarantee for the ε -stationary point (Theorem 1). Secondly, to further study the case with non-

[★] This work was supported by the Swedish Research Council Distinguished Professor Grant 2017-01078, the Knut and Alice Wallenberg Foundation Wallenberg Scholar Grant, and the Swedish Strategic Research Foundation SUCCESS Grant FUS21-0026. The authors are also affiliated to Digital Futures.

convex strategy sets, we derive a probabilistic robustness guarantee for a global minimax point by considering the sequential order of players' decision-making. Leveraging the extreme value theorem and Berge's maximum theorem, we show a tight robustness bound in this fully nonconvex setting (Theorem 2).

The remainder of the paper is organized as follows. Section 2 formulates the minimax problem and revisits a robustness result for NE in the convex-concave case. Section 3 establishes probabilistic robustness guarantees for an ε -stationary point and a global minimax point. Section 4 presents a numerical experiment, while Section 5 concludes the paper.

2. THE SCENARIO MINIMAX PROBLEM

In this section, we introduce the minimax problem and revisit a robustness result from the literature.

2.1 Problem Formulation

Let θ be a random parameter taking values in a support set Θ , equipped with a σ -algebra, and let $\mathbb{P}$ be the associated probability measure on Θ . Let $\mathcal{X}_\theta \subseteq \mathbb{R}^p$ and $\mathcal{Y}_\theta \subseteq \mathbb{R}^q$ be θ -dependent uncertain strategy sets, and let $f : \mathcal{X}_\theta \times \mathcal{Y}_\theta \rightarrow \mathbb{R}$ be a continuously differentiable payoff function. The minimax problem

$$\min_{x \in \mathcal{X}_\theta} \max_{y \in \mathcal{Y}_\theta} f(x, y) \quad (1)$$

represents a two-player zero-sum game under uncertainty. The min-player selects a strategy $x \in \mathcal{X}_\theta$ to minimize the payoff, whereas the max-player selects a strategy $y \in \mathcal{Y}_\theta$ to maximize it.

Generally, one cannot expect a solution to (1) unless the support set Θ is exactly known or additional structural assumptions on the distribution $\mathbb{P}$ are imposed, e.g., Chen et al. (2021); Fochesato et al. (2023). To overcome this obstacle, we employ a data-driven framework, replacing (1) with its empirical approximation based on an independent and identically distributed (i.i.d.) multi-sample consisting of M samples $\boldsymbol{\theta}^M = (\theta_1, \dots, \theta_M) \in \Theta^M$ drawn from Θ . This leads to the following deterministic *scenario minimax problem*:

$$\min_{x \in \mathcal{X}^M} \max_{y \in \mathcal{Y}^M} f(x, y), \quad (2)$$

where $\mathcal{X}^M := \bigcap_{i=1}^M \mathcal{X}_{\theta_i}$ and $\mathcal{Y}^M := \bigcap_{i=1}^M \mathcal{Y}_{\theta_i}$, provided that $\mathcal{X}^M \times \mathcal{Y}^M$ is nonempty for any $M \in \mathbb{N}$.

We first revisit the convex-concave case, where the payoff f is convex in x and concave in y . In this setting, the scenario minimax problem (2) is well-studied (von Neumann, 1928; Fan, 1953; Sion, 1958), with NE serving as a central concept.

Definition 1. A collective strategy (x^*, y^*) is an NE of problem (2) if, for all $(x, y) \in \mathcal{X}^M \times \mathcal{Y}^M$:

$$f(x^*, y) \leq f(x^*, y^*) \leq f(x, y^*).$$

In the convex-concave case, the NE is equivalent to a global saddle point, satisfying the condition that no player has incentive to unilaterally deviate. The following assumption is widely employed to address uncertainty in strategy sets (Chen et al., 2021; Fochesato et al., 2023; Xu et al., 2023).

Assumption 1. (Strategy convexity). $\mathcal{X}_\theta$ and $\mathcal{Y}_\theta$ are convex, compact, and nonempty for all $\theta \in \Theta$.

Lemma 1. (Fan, 1953). Supposing f is convex-concave, and, under Assumption 1, the scenario minimax problem (2) admits an NE on $\mathcal{X}^M \times \mathcal{Y}^M$. Furthermore, the minimax equality holds:

$$\min_{x \in \mathcal{X}^M} \max_{y \in \mathcal{Y}^M} f(x, y) = \max_{y \in \mathcal{Y}^M} \min_{x \in \mathcal{X}^M} f(x, y).$$

Under Assumption 1, the scenario minimax problem (2) can be framed as a *simultaneous game*, in which the order of play is immaterial. In this convex-concave setting, the NE serves as the natural notion of stationarity.

2.2 Robustness of an NE

Once an NE of the scenario minimax problem (2) is found, a natural concern is its robustness against new and unseen samples outside the chosen multi-sample $\boldsymbol{\theta}^M$. Before presenting its robustness result, we first provide some background on the scenario approach, as adopted from Campi and Garatti (2018). The *probability of violation* serves as a measure of robustness.

Definition 2. The probability of violation $V : \mathbb{R}^{p+q} \rightarrow [0, 1]$ of a point (x, y) is defined as

$$V(x, y) := \mathbb{P}\{\theta \in \Theta : (x, y) \notin \mathcal{X}_\theta \times \mathcal{Y}_\theta\}.$$

Definition 2 quantifies the risk that a point (x, y) in the scenario minimax problem (2) may be infeasible for some $\theta \in \Theta$. Central concepts in the scenario approach also include *support list* and *complexity*, defined as follows.

Definition 3. Given a multi-sample $\boldsymbol{\theta}^M$, a support list is a sequence of elements $\theta_{i_1}, \dots, \theta_{i_k}$ of $\boldsymbol{\theta}^M$ with $i_1 < \dots < i_k$ and $k \leq M$ such that

- (1) removing all elements of $\boldsymbol{\theta}^M$ except those in the sequence will not change the equilibrium;
- (2) no more element can be removed from the sequence while leaving the equilibrium unchanged.

Definition 4. The complexity of a multi-sample is the minimal cardinality among all its support lists.

We make the following non-degeneracy assumption. This is considered a mild assumption in convex settings, a point discussed comprehensively by Campi and Garatti (2018, Sec. 8) and Garatti and Campi (2025, Sec. 1.3).

Assumption 2. (Non-degeneracy). For any multi-sample $\boldsymbol{\theta}^M$ with $M \in \mathbb{N}$, with probability 1 over the random draws $\theta_1, \dots, \theta_M$, there exists a unique support list.

Given Lemma 1 and a tie-break rule ensuring uniqueness, e.g., by selecting the NE with minimal ℓ_2 -norm, a direct application of Campi and Garatti (2018, Th. 3) yields a robustness guarantee for the NE.

Proposition 1. Let Assumptions 1 and 2 hold, and let $h(m)$ be any $[0, 1]$ -valued function for $m = 0, 1, \dots, p+q$. Given a multi-sample $\boldsymbol{\theta}^M$, the NE (x^*, y^*) of the scenario minimax problem (2) satisfies

$$\mathbb{P}\{V(x^*, y^*) > h(S_M^*)\} \leq \gamma^*,$$

where S_M^* is the complexity of the multi-sample $\boldsymbol{\theta}^M$, and

$$\gamma^* = \inf_{\xi(\cdot) \in \mathbf{P}_M} \xi(1) \quad (3)$$

subject to

$$\frac{1}{k!} \frac{d^k}{dt^k} \xi(t) \geq \binom{M}{k} t^{M-k} \cdot \mathbb{1}\{t \in [0, 1 - h(k)]\}$$

for all $t \in [0, 1]$ and all $k = 0, 1, \dots, M$, where $\mathbb{1}\{\cdot\}$ denotes the indicator function, which equals 1 if the condition inside is true, and 0 otherwise; $\mathbf{P}_M$ denotes the class of polynomials of degree M .

After revisiting the convex-concave case, we now turn to the central contribution of this paper: establishing robustness results for nonconvex-nonconcave minimax problems. Notably, non-degeneracy in Assumption 2 is generally difficult to verify, even in convex optimization, and this challenge is amplified in the context of games (Fele and Margellos, 2020). Furthermore, the existence of an NE is no longer guaranteed once convexity is violated.

Accordingly, we state our problem as follows.

Problem 1. For general nonconvex-nonconcave minimax problems, characterize the robustness of key equilibria.

3. ROBUSTNESS OF NONCONVEX-NONCONCAVE MINIMAX PROBLEMS

In the nonconvex-nonconcave setting, demonstrating a robustness guarantee like Proposition 1 is challenging because Assumption 2, a crucial non-degeneracy condition, typically fails to hold; indeed, non-degeneracy is almost the norm in nonconvex problems (Garatti and Campi, 2025). This raises a key question: how can we guarantee the robustness of an equilibrium without the assumption of non-degeneracy? We find inspiration from Garatti and Campi (2025), who successfully removed the assumption of convexity in the scenario approach. Their main result can be stated as follows:

Proposition 2. Given a multi-sample $\boldsymbol{\theta}^M$, if there is a map $\mathcal{M}_M : \Theta^M \rightarrow \mathcal{Z}_{\theta_M}$ satisfying the following *consistency property*:

- (1) if $\theta_{i_1}, \dots, \theta_{i_M}$ is a permutation of $\theta_1, \dots, \theta_M$, then $\mathcal{M}_M(\theta_{i_1}, \dots, \theta_{i_M}) = \mathcal{M}_M(\theta_1, \dots, \theta_M)$;
- (2) if $z_M^* \in \mathcal{Z}_{\theta_{M+i}}$ for all $i = 1, \dots, N$, then $z_M^* = \mathcal{M}_M(\theta_1, \dots, \theta_M) = \mathcal{M}_{M+N}(\theta_1, \dots, \theta_{M+N}) = z_{M+N}^*$;
- (3) if $z_M^* \notin \mathcal{Z}_{\theta_{M+i}}$ for at least one $i = 1, \dots, N$, then $z_M^* = \mathcal{M}_M(\theta_1, \dots, \theta_M) \neq \mathcal{M}_{M+N}(\theta_1, \dots, \theta_{M+N}) = z_{M+N}^*$;

then we have

$$\mathbb{P}\{V(z_M^*) > h(S_M^*)\} \leq \gamma^*,$$

where S_M^* is the complexity of the multi-sample $\boldsymbol{\theta}^M$, $h(k), k = 0, 1, \dots, M$ is any $[0, 1]$ -valued function, and γ^* is defined in (3).

In the context of minimax problems, the map $\mathcal{M}_M$ serves as an algorithm to seek an equilibrium $z_M^* \in \mathcal{Z}_{\theta_M}$ based on the multi-sample $\boldsymbol{\theta}^M$. Building upon the elegant result of Proposition 2, in the remainder of this section we study the robustness of two different equilibria, since properly defining stationarity for nonconvex-nonconcave minimax problems remains a fundamental challenge (Li et al., 2025). Specifically, we investigate the robustness of minimax problems with nonconvex-nonconcave payoffs, focusing on (i) an ε -stationary point under convex strategy

sets (Section 3.1) and (ii) a global minimax point under nonconvex strategy sets (Section 3.2).

3.1 Robustness of an ε -Stationary Point

We first examine the robustness of solutions to (2) under Assumption 1, namely, minimax problems with nonconvex-nonconcave payoffs and convex strategy sets. In this setting, a comprehensive description of various stationarity concepts can be found in Zhang et al. (2022). One approach is ε -stationary point, which provides necessary conditions for a local NE; this concept is often analyzed under either the Polyak-Łojasiewicz (PL) condition (Nouiehed et al., 2019; Doan, 2022; Yang et al., 2022) or the Kurdyka-Łojasiewicz (KL) condition (Zheng et al., 2023; Li et al., 2025).

Following Zheng et al. (2023); Lin et al. (2024), we first define the concept of ε -stationary point.

Definition 5. A collective strategy $(\hat{x}, \hat{y}) \in \mathcal{X}^M \times \mathcal{Y}^M$ is an ε -stationary point of (2), if

$$\begin{aligned} \text{dist}(\mathbf{0}_p, \nabla_x f(\hat{x}, \hat{y}) + \mathcal{N}_{\mathcal{X}^M}(\hat{x})) &\leq \varepsilon, \\ \text{dist}(\mathbf{0}_q, -\nabla_y f(\hat{x}, \hat{y}) + \mathcal{N}_{\mathcal{Y}^M}(\hat{y})) &\leq \varepsilon, \end{aligned}$$

where $\mathcal{N}_{\mathcal{S}}$ is the normal cone operator associated with a set $\mathcal{S} \subseteq \mathbb{R}^d$, and $\text{dist}(z, \mathcal{S}) := \inf_{v \in \mathcal{S}} \|z - v\|$ denotes the distance from $z \in \mathbb{R}^d$ to $\mathcal{S}$.

Under Assumption 1, often in conjunction with the PL or KL condition, convergence to an ε -stationary point can often be achieved, typically via the gradient descent-ascent (GDA) algorithm and its variants (Li et al., 2025; Laguel et al., 2024; Yang et al., 2022). Despite its importance, the robustness of the resulting ε -stationary point remains largely unexplored, and this work aims to fill this research gap.

Quantifying the robustness of ε -stationary points is challenging in the nonconvex-nonconcave case, where the complex geometry and potential for degeneracy make direct analysis intractable. Our core strategy is to first prove the existence of an ε -stationary point under Assumption 1, and then establish the consistency property introduced in Proposition 2 and in that way characterize the stability of ε -stationary points against sample changes. Note that the result holds even without a non-degeneracy assumption.

Theorem 1. Let Assumption 1 hold. For a given multi-sample $\boldsymbol{\theta}^M$, there exists an ε -stationary point $(\hat{x}, \hat{y})$ of (2) for some $\varepsilon > 0$ such that

$$\mathbb{P}\{V(\hat{x}, \hat{y}) > h(S_M^*)\} \leq \gamma^*,$$

where S_M^* is the complexity of the multi-sample $\boldsymbol{\theta}^M$, and γ^* is defined in (3).

Proof. The ε -stationary point's invariance to permutation of the multi-sample $\boldsymbol{\theta}^M$ means that condition (1) of the consistency property is naturally met.

To make the dependence on the sample set explicit for a fixed $\boldsymbol{\theta}^M$, define the stationary residual

$$r^M(x, y) := \max \left\{ \text{dist}(\mathbf{0}_p, \nabla_x f(x, y) + \mathcal{N}_{\mathcal{X}^M}(x)), \text{dist}(\mathbf{0}_q, -\nabla_y f(x, y) + \mathcal{N}_{\mathcal{Y}^M}(y)) \right\}.$$

The continuously differentiability of f implies the continuity of both $\nabla_x f$ and $\nabla_y f$. By Rockafellar and Wets (1998,

Prop. 6.6), the set-valued mapping $x \mapsto \mathcal{N}_{\mathcal{X}^M}(x)$ is outer semicontinuous with closed values. Therefore, Rockafellar and Wets (1998, Prop. 5.11) guarantees that for fixed y , the map

$$x \mapsto \text{dist}(\mathbf{0}_p, \nabla_x f(x, y) + \mathcal{N}_{\mathcal{X}^M}(x))$$

is lower semicontinuous (lsc) in x , and similarly for the y -component. Consequently, since $r^M(x, y)$ is the maximum of two lsc functions, it follows directly that $r^M(x, y)$ is lsc in (x, y) . Under Assumption 1, Rockafellar and Wets (1998, Th. 1.9) guarantees that there exists

$$(\hat{x}_M, \hat{y}_M) \in \arg \min_{(x, y) \in \mathcal{X}^M \times \mathcal{Y}^M} r^M(x, y)$$

such that $r^M(\hat{x}_M, \hat{y}_M) = \varepsilon$.

Now consider augmenting the $\boldsymbol{\theta}^M$ by an additional multi-sample $\boldsymbol{\theta}^N$. Observe that both $\mathcal{X}^{M+N}$ and $\mathcal{X}^M$ are convex, and the subset relation $\mathcal{X}^{M+N} \subseteq \mathcal{X}^M$ leads to

$$\mathcal{N}_{\mathcal{X}^M}(x) \subseteq \mathcal{N}_{\mathcal{X}^{M+N}}(x), \quad \forall x \in \mathcal{X}^M,$$

and similarly for the y -component. Hence, for every (x, y) in the common domain $\mathcal{X}^M \times \mathcal{Y}^M$,

$$\begin{aligned} & \text{dist}(\mathbf{0}_p, \nabla_x f(x, y) + \mathcal{N}_{\mathcal{X}^{M+N}}(x)) \\ & \leq \text{dist}(\mathbf{0}_p, \nabla_x f(x, y) + \mathcal{N}_{\mathcal{X}^M}(x)), \end{aligned}$$

and an analogous inequality holds for the y -component. It follows that for all (x, y) ,

$$r^{M+N}(x, y) \leq r^M(x, y).$$

Take the pair

$$(\hat{x}_{M+N}, \hat{y}_{M+N}) \in \arg \min_{(x, y) \in \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}} r^{M+N}(x, y)$$

as a minimizer for the augmented multi-sample $\boldsymbol{\theta}^{M+N}$. If for all $j \in \{1, \dots, N\}$ it holds that

$$(\hat{x}_M, \hat{y}_M) \in \mathcal{X}^{M+j} \times \mathcal{Y}^{M+j},$$

then in particular

$$(\hat{x}_M, \hat{y}_M) \in \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}$$

and therefore

$$r^{M+N}(\hat{x}_{M+N}, \hat{y}_{M+N}) \leq r^{M+N}(\hat{x}_M, \hat{y}_M) \leq r^M(\hat{x}_M, \hat{y}_M) = \varepsilon.$$

This shows that the ε -stationary point remains valid after adding new appropriate samples; i.e., condition (2) of the consistency property holds.

Conversely, if for some $j \in \{1, \dots, N\}$ the point $(\hat{x}_M, \hat{y}_M)$ fails to belong to $\mathcal{X}^{M+j} \times \mathcal{Y}^{M+j}$, then in particular

$$(\hat{x}_M, \hat{y}_M) \notin \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}.$$

Recall that

$$(\hat{x}_{M+N}, \hat{y}_{M+N}) \in \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}.$$

Thus, we conclude that

$$(\hat{x}_M, \hat{y}_M) \neq (\hat{x}_{M+N}, \hat{y}_{M+N}),$$

thereby verifying condition (3) of the consistency property. Combining the above observations with Proposition 2 completes the proof. $\square$

Theorem 1 establishes a generalization guarantee of an ε -stationary point to problem (2). Given a solution $(\hat{x}_M, \hat{y}_M)$ obtained via a suitable algorithm, such as the one proposed in Zheng et al. (2023), Theorem 1 provides a probabilistic upper bound on its generalization error.

From a practical perspective, one typically fixes a high confidence level $1 - \beta$ (i.e., very close to 1) and, in

return, desires a deterministic upper bound $h(S_M^*)$ for the probability of violation $V(\hat{x}_M, \hat{y}_M)$. Since the function $h(\cdot)$ as in Proposition 2 and Theorem 1 is indefinite for a fixed γ^* , there may exist infinitely many solutions $h(k)$ to (3). A practical corollary resolves this issue as follows.

Corollary 1. Under Assumption 1, for a given $\beta \in (0, 1)$ and a multi-sample $\boldsymbol{\theta}^M$, the ε -stationary point $(\hat{x}, \hat{y})$ of (2) satisfies

$$\mathbb{P}\{V(\hat{x}, \hat{y}) > g(S_M^*)\} \leq \beta,$$

where S_M^* is the complexity of $\boldsymbol{\theta}^M$ and

$$g(k) = \begin{cases} 1 - t(k), & \text{if } k = 0, 1, \dots, M-1, \\ 1, & \text{if } k = M, \end{cases} \quad (4)$$

with $t(k)$ being the unique solution in the interval $(0, 1)$ of

$$\frac{\beta}{M+1} \sum_{m=k}^M \binom{m}{k} t^{m-k} - \binom{M}{k} t^{M-k} = 0.$$

Equivalently, the result in Corollary 1 can be stated in terms of the complementary probability:

$$\mathbb{P}\{V(\hat{x}, \hat{y}) \leq g(S_M^*)\} \geq 1 - \beta.$$

This probabilistic bound gives a powerful robustness guarantee for an obtained ε -stationary point $(\hat{x}, \hat{y})$ from a multi-sample $\boldsymbol{\theta}^M$. By first selecting an arbitrarily high confidence level $1 - \beta$ and then computing the complexity S_M^* of $\boldsymbol{\theta}^M$, one can uniquely determine the upper bound $g(S_M^*)$ using (4). That is, we have at least $1 - \beta$ confidence to ensure that $V(\hat{x}, \hat{y})$ will not exceed this computed and unique $g(S_M^*)$. Note that the probability of violation $V(\hat{x}, \hat{y})$ serves as a measure of robustness: the smaller its value, the more robust the ε -stationary point $(\hat{x}, \hat{y})$ is. Therefore, Corollary 1 formalizes this probabilistic robustness guarantee.

3.2 Robustness of a Global Minimax Point

Let us consider the minimax problem (2) with nonconvex–nonconcave payoffs and nonconvex strategy sets. In Section 3.1, we investigate the robustness of an ε -stationary point by playing a simultaneous game. However, in this simultaneous setting, it is not straightforward to extend the result of Theorem 1 to minimax problems with nonconvex strategy sets. This is because the monotonicity of the stationary residual with respect to the number of scenarios is difficult to maintain when the convexity assumption on the strategy sets (Assumption 1) is violated.

Interestingly, the GANs framework implements a sequential two-player game: it first optimizes the discriminator D while keeping the generator G fixed, and then updates G using the optimized D . More generally, adversarial training describes the minimax problem (2) in the frame of a *sequential game*. These cases lead to the equilibrium concept known as *global minimax point*, which captures global optimality and is widely employed in machine learning (Jin et al., 2020; Zhang et al., 2022; Chen et al., 2024). Intuitively, the min-player acts as the leader, first determining their optimal strategy, while the max-player then acts as the follower, choosing a best response to the min-player’s choice.

Following Jin et al. (2020), we begin by defining the concept of global minimax point. In this paper, we focus

exclusively on a global minimax point of (2), since any global maximin point corresponds to its global minimax point with the payoff $-f(y, x)$.

Definition 6. A collective strategy $(x^*, y^*) \in \mathcal{X}^M \times \mathcal{Y}^M$ is a global minimax point of (2), if for all (x, y) :

$$f(x^*, y) \leq f(x^*, y^*) \leq \max_{y' \in \mathcal{Y}^M} f(x, y').$$

The next assumption removes all convexity requirements, including the strategy sets originally assumed to be convex in Assumption 1.

Assumption 3. (Strategy nonconvexity). $\mathcal{X}_\theta$ and $\mathcal{Y}_\theta$ are compact and nonempty for all $\theta \in \Theta$.

Lemma 2. (Jin et al., 2020). Under Assumption 3, there exists at least one global minimax point (x^*, y^*) of the scenario minimax problem (2).

Given the existence established, we next investigate the robustness property of such global minimax point.

Theorem 2. Let Assumption 3 hold. For a given a multi-sample θ^M , the global minimax point (x^*, y^*) of the scenario minimax problem (2) satisfies

$$\mathbb{P}\{V(x^*, y^*) > h(S_M^*)\} \leq \gamma^*,$$

where S_M^* is the complexity of the multi-sample θ^M , and γ^* is defined in (3).

Proof. The condition (1) in the consistency property naturally holds, since the global minimax point is invariant under any permutation of the multi-sample θ^M .

Since real function $f(x, \cdot)$ is continuous on the nonempty compact set $\mathcal{Y}^M$ for any fixed x , the extreme value theorem (Rudin, 1976, Th. 4.16) guarantees the existence of a point $y_M^*(x) \in \mathcal{Y}^M$ (note the dependency on x) such that, for all $(x, y) \in \mathcal{X}^M \times \mathcal{Y}^M$:

$$f(x, y) \leq f(x, y_M^*(x)). \quad (5)$$

Indeed, provided the set $\mathcal{Y}^M$ (or $\mathcal{X}^M$) is nonempty, the intersection of an arbitrary collection of compact sets $\mathcal{Y}_\theta$ (or $\mathcal{X}_\theta$) remains compact. Therefore, after adding a new multi-sample θ^N , there exists

$$y_{M+N}^*(x) \in \mathcal{Y}^{M+N} \subseteq \mathcal{Y}^M \ni y_M^*(x)$$

such that for all $(x, y) \in \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}$:

$$f(x, y) \leq f(x, y_{M+N}^*(x)). \quad (6)$$

If $y_M^*(x) \in \mathcal{Y}^{M+j}$ for all $j = 1, \dots, N$, (6) is held for all

$$(x, y) \in \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N} \subseteq \mathcal{X}^M \times \mathcal{Y}^M. \quad (7)$$

Note that under Assumption 3, the extreme value theorem ensures the existence, but not the uniqueness of the maximizer $y_M^*(x)$. However, since $y_M^*(x)$ was already ranked first in the larger feasibility set $\mathcal{Y}^M$ before the new constraints were introduced, we choose the tie-break rule to rank it first according to (7). Then, we can conclude that $y_M^*(x) = y_{M+N}^*(x)$, for all x .

Define the envelope function as

$$\phi(x) := \max_{y' \in \mathcal{Y}^M} f(x, y').$$

Since f is continuous and $\mathcal{Y}_\theta$ is compact and nonempty, the Berge's maximum theorem (Aliprantis and Border, 2006, Th. 17.31) directly induces the continuity of $\phi(x)$.

The extreme value theorem guarantees that there exists a $x_M^* \in \mathcal{X}^M$ such that for any x ,

$$\phi(x_M^*) \leq \phi(x),$$

that is,

$$\max_{y' \in \mathcal{Y}^M} f(x_M^*, y') \leq \max_{y' \in \mathcal{Y}^M} f(x, y'), \quad \forall x.$$

It follows from

$$f(x_M^*, y_M^*(x_M^*)) \leq \max_{y' \in \mathcal{Y}^M} f(x_M^*, y')$$

and (5) with $x = x_M^* \in \mathcal{X}^M$ that the existence of the global minimax point $(x_M^*, y_M^*(x_M^*))$ is guaranteed. A similar argument gives $x_M^* = x_{M+N}^*$. Therefore, the condition (2) in the consistency property holds.

Conversely, if for some $j \in \{1, \dots, N\}$ the point (x_M^*, y_M^*) does not belong to $\mathcal{X}^{M+j} \times \mathcal{Y}^{M+j}$, then in particular it is not contained in $\mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}$, whereas

$$(x_{M+N}^*, y_{M+N}^*) \in \mathcal{X}^{M+N} \times \mathcal{Y}^{M+N}.$$

We thereby obtain

$$(x_M^*, y_M^*) \neq (x_{M+N}^*, y_{M+N}^*),$$

which establishes condition (3) of the consistency property. By invoking Proposition 2, we conclude that the proof is complete. $\square$

While Theorem 1 focuses on a local ε -stationary point, Theorem 2 shows that the same high-probability generalization bound also holds for a global minimax point, even in fully nonconvex settings. By choosing an arbitrarily high confidence level $1 - \beta$ and computing the corresponding upper bound $h(S_M^*)$, one can assert, with probability at least $1 - \beta$, that the probability of violation $V(x^*, y^*)$ does not exceed $h(S_M^*)$. Analogous to Corollary 1, we present a practical and specific bound below.

Corollary 2. Let Assumption 3 hold. Given $\beta \in (0, 1)$ and a multi-sample θ^M , the global minimax point (x^*, y^*) of the scenario minimax problem (2) satisfies

$$\mathbb{P}\{V(x^*, y^*) > g(S_M^*)\} \leq \beta,$$

where S_M^* is the complexity of θ^M and the function $g(k)$, for $k = 0, 1, \dots, M$, is defined as

$$g(k) = \begin{cases} 1 - t(k), & \text{if } k = 0, 1, \dots, M-1, \\ 1, & \text{if } k = M, \end{cases}$$

with $t(k)$ being the unique solution in the interval $(0, 1)$ of

$$\frac{\beta}{M+1} \sum_{m=k}^M \binom{m}{k} t^{m-k} - \binom{M}{k} t^{M-k} = 0.$$

Compared with Theorem 2, which provides a general theoretical characterization, Corollary 2 (derived from Theorem 2) is more relevant in practice. It extends the high-probability generalization guarantee from a local ε -stationary point to a global minimax point while preserving practical applicability.

4. NUMERICAL EXPERIMENT

In this section, we validate the theoretical results by considering a unit commitment problem under uncertainty. Our formulation is abstracted from Bertsimas and Koulouras (2024), which also considers additional capacity uncertainty, non-adaptive costs, and much more complex

constraints. For clarity, we simplify or omit these aspects and obtain the following minimax problem:

$$\min_{x \in \mathcal{X}_\theta} \max_{y \in \mathcal{Y}_\theta} \sum_{i=1}^I \left(C_i u_i + \sum_{j=1}^J V_{ij} y_j \right),$$

which aims to minimize the total cost of meeting a fixed level of demand under load uncertainty. In this paper, we consider $I = 5$ generators and $J = 5$ demand nodes. The min-player's decision variable $x = \text{col}(u, \text{vec}(V))$, where $u = \text{col}(u_1, \dots, u_I)$ represents the commitment status of five generating units. The binary variable u_i equals 0 if generator i is turned off and 1 if it is turned on. The matrix $V \in \mathbb{R}^{I \times J}$ represents the electricity contribution of each generator to each demand node under uncertain load $y = \text{col}(y_1, \dots, y_J)$. Specifically, each element V_{ij} is a unit response coefficient, indicating the portion of power from generator i allocated to demand node j when node j experiences an uncertain load level y_j . Thus, u represents the non-adaptive (commitment) part of the decision, while V captures the adaptive (dispatch) response to the uncertainty in y .

We consider the generator turn-on cost vector $C = (C_i)_{i=1}^I = (1.0, 1.1, 0.9, 1.05, 1.2)$. The strategy sets of both players depend on an uncertainty parameter θ . Following the idea of adjustable uncertainty sets in (Wang et al., 2016), we link the aggregate commitment of generating units to θ and impose the following feasible set for the min-player:

$$\mathcal{X}_\theta = \left\{ (u, \text{vec}(V)) : u \in \{0, 1\}^I, \sum_{i=1}^I R_i u_i \geq 3\theta, \right. \\ \left. 0 \leq V_{ij} \leq u_i, \forall i \in I, j \in J, \sum_{i=1}^I \sum_{j=1}^J V_{ij} = 5 \right\},$$

where $R = (R_i)_{i=1}^I = (1.0, 1.2, 0.9, 1.1, 1.3)$. Following Bertsimas and Koulouras (2024), the max-player's strategy is constrained by a box uncertainty set $\mathcal{Y}_\theta = \{y \in \mathbb{R}^J : \|y\|_\infty \leq \theta\}$. The parameter θ represents an uncertain factor affecting the feasible sets $\mathcal{X}_\theta$ and $\mathcal{Y}_\theta$. In this numerical experiment, θ is assumed to be uniformly distributed over $\Theta = [0.5, 1.5]$.

Given the scenario minimax problem associated with a multi-sample θ^M , we compute $\phi(x) = \max_y f(x, y)$ and select the min-player's strategy as $x_M^* = \arg \min_x \phi(x) = \arg \min_x \max_y f(x, y)$, while the max-player chooses $y_M^* = \arg \max_y f(x_M^*, y)$. With this minimax point (x_M^*, y_M^*) , we aim to ensure that $\mathbb{P}\{\theta \in \Theta : (x_M^*, y_M^*) \notin \mathcal{X}_\theta \times \mathcal{Y}_\theta\} > g(S_M^*)$ holds with probability at most β , where the complexity S_M^* is evaluated using the greedy algorithm described in Campi et al. (2018, Sec. 2).

To empirically validate this bound $g(S_M^*)$, a Monte Carlo simulation with $R = 200$ repetitions is performed for the number of scenarios, M , ranging from 1 to 100, and for $\beta = 0.01$. In each repetition, a set of M scenarios is sampled, the global minimax point (x_M^*, y_M^*) is computed, and its complexity S_M^* is recorded to determine the theoretical bound $g(S_M^*)$. Following this, the solution is subject to an out-of-sample test against a large independent set of $N_{\text{test}} = 10^4$ test scenarios. This test yields the empirical probability of violation $\hat{V}_M(x_M^*, y_M^*) =$




Fig. 1. The probability of violation depends on the number of scenarios for $\beta = 0.01$. The figure compares the theoretical bound $g(S_M^*)$ (red solid line) with the range of all possible empirical probabilities of violation $\hat{V}_M(x_M^*, y_M^*)$ (light purple shaded area) and the average empirical probability of violation $\bar{V}_M(x_M^*, y_M^*)$ (blue solid line) across different numbers of scenarios.

$\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \mathbb{1}\{(x_M^*, y_M^*) \notin \mathcal{X}_{\theta_i} \times \mathcal{Y}_{\theta_i}\}$, for that repetition. The average empirical probability of violation over R repetitions in the Monte Carlo simulation is denoted by $\bar{V}_M(x_M^*, y_M^*)$.

Unlike the probability of violation, the complexity S_M^* is identical across all Monte Carlo repetitions for any M . Therefore, $g(S_M^*)$ in Fig. 1 remains constant for a given M . By fixing $\beta = 0.01$ (a 99% confidence level), we expect the empirical violation $\hat{V}_M(x_M^*, y_M^*)$ to exceed the theoretical bound $g(S_M^*)$ in no more than 1% of the repetitions. In Fig. 1, our numerical results are consistent with this guarantee, as the bound was rarely violated (observed only once at $M = 50$ and once at $M = 69$). That is, the true confidence $2/(100 \times 200) = 0.01\% \leq 1\%$. This provides a numerical verification of Theorem 2 and Corollary 2.

5. CONCLUSIONS

This paper developed a data-driven and distribution-free framework to establish tight probabilistic robustness guarantees for an ε -stationary point and a global minimax point in general nonconvex–nonconcave minimax problems. In particular, we obtained a robustness guarantee for a local equilibrium through an ε -stationary point under convex strategy sets, and we demonstrated that a global minimax point satisfied an analogous robustness guarantee with nonconvex strategy sets. We removed the non-degeneracy assumption by leveraging recent advances in nonconvex scenario optimization. We also identified the extension of this framework to nonconvex multiplayer games as the most challenging direction for future research.

REFERENCES

- Aliprantis, C.D. and Border, K.C. (2006). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3rd edition.
- Assif, M., Chatterjee, D., and Banavar, R. (2020). Scenario approach for minmax optimization with emphasis on

- the nonconvex case: Positive results and caveats. *SIAM Journal on Optimization*, 30(2), 1119–1143.
- Başar, T. and Bernhard, P. (2008). *H-infinity Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Springer Science & Business Media, 2nd edition.
- Bertsimas, D. and Koulouras, A.G. (2024). Marginal pricing in adaptive robust unit commitment under load and capacity uncertainty. *IEEE Transactions on Power Systems*, 40(1), 341–354.
- Campi, M.C. and Garatti, S. (2018). Wait-and-judge scenario optimization. *Mathematical Programming*, 167(1), 155–189.
- Campi, M.C., Garatti, S., and Ramponi, F.A. (2015). Non-convex scenario optimization with application to system identification. In *Proceedings of the 54th IEEE Conference on Decision and Control (CDC)*, 4023–4028. IEEE.
- Campi, M.C., Garatti, S., and Ramponi, F.A. (2018). A general scenario theory for nonconvex optimization and decision making. *IEEE Transactions on Automatic Control*, 63(12), 4067–4078.
- Carè, A., Garatti, S., and Campi, M.C. (2015). Scenario min-max optimization and the risk of empirical costs. *SIAM Journal on Optimization*, 25(4), 2061–2080.
- Chen, G., Ming, Y., Hong, Y., and Yi, P. (2021). Distributed algorithm for ε -generalized Nash equilibria with uncertain coupled constraints. *Automatica*, 123, 109313.
- Chen, G., Xu, G., He, F., Hong, Y., Rutkowski, L., and Tao, D. (2024). Approaching the global Nash equilibrium of non-convex multi-player games. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12), 10797–10813.
- Chen, G., Xu, G., He, F., Tao, D., Parisini, T., and Johansson, K.H. (2025). Inverse learning of black-box aggregator for robust Nash equilibrium. *IEEE Transactions on Automatic Control*. doi:10.1109/TAC.2025.3634219.
- Doan, T. (2022). Convergence rates of two-time-scale gradient descent-ascent dynamics for solving nonconvex min-max problems. In *Learning for Dynamics and Control Conference*, 192–206.
- Fan, K. (1953). Minimax theorems. *Proceedings of the National Academy of Sciences*, 39(1), 42–47.
- Fele, F. and Margellos, K. (2020). Probably approximately correct Nash equilibrium learning. *IEEE Transactions on Automatic Control*, 66(9), 4238–4245.
- Fochesato, M., Fabiani, F., and Lygeros, J. (2023). Generalized uncertain Nash games: Reformulation and robust equilibrium seeking. In *Proceedings of the 21st European Control Conference (ECC)*, 1–6. IEEE.
- Garatti, S. and Campi, M.C. (2025). Non-convex scenario optimization. *Mathematical Programming*, 209(1), 557–608.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- Jin, C., Netrapalli, P., and Jordan, M. (2020). What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning*, 4880–4889.
- Laguel, Y., Syed, Y., Aybat, N.S., and Gürbüzbalaban, M. (2024). High-probability complexity guarantees for nonconvex minimax problems. In *Advances in Neural Information Processing Systems*, 140937–140969.
- Li, J., Zhu, L., and So, A.M.C. (2025). Nonsmooth nonconvex-nonconcave minimax optimization: Primal-dual balancing and iteration complexity analysis. *Mathematical Programming*, 1–51.
- Lin, Z., Xue, C., Deng, Q., and Ye, Y. (2024). A single-loop robust policy gradient method for robust markov decision processes. In *International Conference on Machine Learning*, 30392–30426.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 4138–4160.
- Nouiehed, M., Sanjabi, M., Huang, T., Lee, J.D., and Razaviyayn, M. (2019). Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*, 14934–14942.
- Pantazis, G., Fele, F., and Margellos, K. (2024). A priori data-driven robustness guarantees on strategic deviations from generalised Nash equilibria. *Automatica*, 167, 111746.
- Rockafellar, R.T. and Wets, R.J. (1998). *Variational Analysis*. Springer.
- Rudin, W. (1976). *Principles of Mathematical Analysis*. McGraw-Hill, 3rd edition.
- Sion, M. (1958). On general minimax theorems. *Pacific Journal of Mathematics*, 8(1), 171–176.
- Sutton, R.S. and Barto, A.G. (2018). *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition.
- von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1), 295–320.
- Wang, C., Liu, F., Wang, J., Qiu, F., Wei, W., Mei, S., and Lei, S. (2016). Robust risk-constrained unit commitment with large-scale wind generation: An adjustable uncertainty set approach. *IEEE Transactions on Power Systems*, 32(1), 723–733.
- Xu, G., Chen, G., and Qi, H. (2023). Algorithm design and approximation analysis on distributed robust game. *Journal of Systems Science and Complexity*, 36(2), 480–499.
- Yang, J., Orvieto, A., Lucchi, A., and He, N. (2022). Faster single-loop algorithms for minimax optimization without strong concavity. In *International Conference on Artificial Intelligence and Statistics*, 5485–5517.
- Zhang, G., Poupart, P., and Yu, Y. (2022). Optimality and stability in non-convex smooth games. *Journal of Machine Learning Research*, 23(35), 1–71.
- Zheng, T., Zhu, L., So, A.M.C., Blanchet, J., and Li, J. (2023). Universal gradient descent ascent method for nonconvex-nonconcave minimax optimization. In *Advances in Neural Information Processing Systems*, 54075–54110.