



HAL
open science

Leveraging Noisy Observations in Zero-Sum Games

Emmanouil M Athanasakos, Samir M Perlaza

► **To cite this version:**

Emmanouil M Athanasakos, Samir M Perlaza. Leveraging Noisy Observations in Zero-Sum Games. 2024. hal-04432819

HAL Id: hal-04432819

<https://inria.hal.science/hal-04432819>

Preprint submitted on 3 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Leveraging Noisy Observations in Zero-Sum Games

Emmanouil M. Athanasakos* and Samir M. Perlaza*^{†‡}

*INRIA, Centre Inria d’Université Côte d’Azur, Sophia Antipolis, France.

[†]ECE Dept. Princeton University, Princeton, 08544 NJ, USA.

[‡]GAATI, Université de la Polynésie Française, Faaa, French Polynesia.

Abstract—This paper studies an instance of zero-sum games in which one player (the leader) commits to its opponent (the follower) to choose its actions by sampling a given probability measure (strategy). The actions of the leader are observed by the follower as the output of an arbitrary channel. In response to that, the follower chooses its action based on its current information, that is, the leader’s commitment and the corresponding noisy observation of its action. Within this context, the equilibrium of the game with noisy action observability is shown to always exist and the necessary conditions for its uniqueness are identified. Interestingly, the noisy observations have important impact on the cardinality of the follower’s set of best responses. Under particular conditions, such a set of best responses is proved to be a singleton almost surely. The proposed model captures any channel noise with a density with respect to the Lebesgue measure. As an example, the case in which the channel is described by a Gaussian probability measure is investigated.

I. INTRODUCTION

Zero-sum games (ZSGs) have proven to be a valuable tool across an array of disciplines, ranging from machine learning (ML) to economics. ZSGs can be utilized to proficiently model interactions between mutually adversarial decision makers (players) and have taken on a pivotal role in the advancement of techniques that enhance the robustness of machine learning algorithms against adversarial attacks. Game theoretic adversarial learning/training techniques are introduced as a competitive game between the learner and the adversary in [1] and [2]. Additionally, the fundamental concept of generative adversarial networks (GANs) draws inspiration from a two-player zero-sum game played between the discriminator and the generator [3]–[6].

The majority of the research at the intersection of game theory and adversarial ML utilizes the properties of ZSGs to forecast game results based on mixed strategies rather than individual actions. To put it simply, the core premise of the Stackelberg equilibrium (SE) [7] in mixed strategies assumes that the leader’s chosen strategy is completely observable by the follower, while actions remain unobservable. Conversely, in the context of the SE in pure strategies, the assumption is that the actions are perfectly observable, rendering the concept of committing to a strategy irrelevant [8]–[10]. Regardless, these assumptions are rarely fulfilled in the majority of ML applications and can be challenging to substantiate their realism. Indeed, data is acquired through a data acquisition system,

which is prone to quantization errors, additive noise, and distortions arising from data transmission and storage. These supplementary imperfections are not necessarily attributable solely to adversarial intent but rather stem from the inherent characteristics of transmission channels. This work makes a step towards this direction and studies new equilibrium concepts. This work extends the game theoretic framework that was introduced in [11] by considering arbitrary channels. Driven by the growing relevance of game theory in the progress of ML [12]–[14], particularly in the context of GANs and adversarial learning, this work presents the subsequent contributions:

- Under perfect observation of the commitment (strategy of the leader), a new game formulation is introduced in which the follower observes a noisy version of actions played by the leader. In order to capture the impact of imperfect data acquisition, leader’s actions are the inputs of an arbitrary channel, while the observations of the follower are the outputs of the channel.
- The formulation is not restricted to specific types of noise distributions; rather, it is general enough to facilitate any noise distribution that can be described by a probability density function (pdf).
- The set of best responses of the follower is characterized and found that is a singleton almost surely. Moreover, the cases in which the noise is beneficial or irrelevant for the follower are investigated and the optimal commitments of the leader are characterized.
- The equilibrium is shown to always exist and the key conditions that affect its uniqueness are identified.
- An example for the case in which the channel is an additive white Gaussian noise channel (AWGN) is investigated and the equilibrium solution is characterized.

II. PRELIMINARIES AND NOTATION

A. Game Formulation

Consider a two-player ZSG in normal form in which Player 1 and Player 2 choose their actions from the sets \mathcal{A}_1 and \mathcal{A}_2 , respectively. These sets are countable, and for all $k \in \{1, 2\}$, $\mathcal{A}_k \triangleq \{a_{k,1}, a_{k,2}, \dots, a_{k,m_k}\}$ exhibits cardinality $m_k > 1$. Let the function

$$u : \mathcal{A}_1 \times \mathcal{A}_2 \rightarrow \mathbb{R}, \quad (1)$$

This work is supported by the Inria Exploratory Action–Information and Decision Making(AExIDEM).

be such that if Player 1 plays $a_{1,i}$ and Player 2 plays $a_{2,j}$, the payoff achieved by both players is $u(a_{1,i}, a_{2,j})$. Note that the function u in (1) is a simple function as its image is formed by the elements in the set $\{u(a_{1,i}, a_{2,j}) : (i, j) \in \{1, 2, \dots, m_1\} \times \{1, 2, \dots, m_2\}\}$. For the ease of notation, for all $(i, j) \in \{1, 2, \dots, m_1\} \times \{1, 2, \dots, m_2\}$ consider that $u_{i,j} \triangleq u(a_{1,i}, a_{2,j})$. Player 1 and Player 2 choose their actions to maximize and minimize their payoffs, respectively.

When the game is played with commitments and one of the players observes the action played by its counterpart subject to noise, the notion of memoryless channels in the sense of [15], [16] is utilized. Let $\Delta(\tilde{\mathcal{A}}_2|\mathcal{A}_2)$ be the set of all probability measures on $\tilde{\mathcal{A}}_2$ conditioned on the elements of \mathcal{A}_2 . The channel is represented by $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2} \in \Delta(\tilde{\mathcal{A}}_2|\mathcal{A}_2)$, for some given set $\tilde{\mathcal{A}}_2$. More specifically, given a channel input $b \in \mathcal{A}_2$, the channel output is $\tilde{b} \in \tilde{\mathcal{A}}_2$ with probability $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=b}(\tilde{b})$. A strategy for Player 1, denoted by $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2} \in \Delta(\mathcal{A}_1|\tilde{\mathcal{A}}_2)$, is a probability measure on \mathcal{A}_1 conditioned on the elements of $\tilde{\mathcal{A}}_2$. Hence, when the channel output is \tilde{b} , Player 1 chooses its actions by sampling the probability measure $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}} \in \Delta(\mathcal{A}_1)$, where $\Delta(\mathcal{A}_1)$ denotes the set of all probability measures on the set \mathcal{A}_1 . Alternatively, a strategy for Player 2 is a probability measure denoted by $P_{\mathcal{A}_2} \in \Delta(\mathcal{A}_2)$. That is, Player 2 chooses its actions by sampling $P_{\mathcal{A}_2}$.

The extended ZSG that includes the noisy channel $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2}$ is represented by the tuple:

$$\mathcal{G}(u, P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2}) \triangleq (\mathcal{A}_1, \mathcal{A}_2, u, P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2}). \quad (2)$$

The game $\mathcal{G}(u, P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2})$ is played in three stages. In the first stage, Player 2 announces its strategy $P_{\mathcal{A}_2} \in \Delta(\mathcal{A}_2)$ to Player 1 and commits to choose its actions by using such a strategy. In the second stage, Player 2 plays action $b \in \mathcal{A}_2$ with probability $P_{\mathcal{A}_2}(b)$ and Player 1 observes the channel output $\tilde{b} \in \tilde{\mathcal{A}}_2$ with probability $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=b}(\tilde{b})$. In the final stage, Player 1 plays action $a \in \mathcal{A}_1$, with probability $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}}(a)$ and both players obtain equal payoffs $u(a, b)$.

Therefore, the expected payoff obtained by the players is determined by the function $v : \Delta(\mathcal{A}_1|\tilde{\mathcal{A}}_2) \times \Delta(\mathcal{A}_2) \rightarrow \mathbb{R}$, such that given the strategy $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2}$ of Player 1 and the strategy $P_{\mathcal{A}_2}$ of Player 2, the expected payoff is

$$v(P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2}, P_{\mathcal{A}_2}) = \int_{\mathcal{A}_2} \left(\int_{\tilde{\mathcal{A}}_2} \left(\int_{\mathcal{A}_1} u(a, b) dP_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}}(a) \right) dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=b}(\tilde{b}) \right) dP_{\mathcal{A}_2}(b). \quad (3)$$

While the sets \mathcal{A}_1 and \mathcal{A}_2 are assumed to be countable, the set $\tilde{\mathcal{A}}_2$ is assumed to be a subset of \mathbb{R} . For all $j \in \{1, \dots, m_2\}$, the channel $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2}$ satisfies

$$P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,j}} \lll \lambda, \quad (4)$$

where \lll denotes mutual absolute continuity; and λ denotes the Lebesgue measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$.

Under these assumptions, the expected payoff function v in (3) satisfies

$$v(P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2}, P_{\mathcal{A}_2}) = \int_{\tilde{\mathcal{A}}_2} \left(\sum_{i=1}^{m_1} P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}}(a_{1,i}) \left(\sum_{j=1}^{m_2} u_{i,j} P_{\mathcal{A}_2}(a_{2,j}) \frac{dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,j}}(\tilde{b})}{dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,1}}(\tilde{b})} \right) \right) dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,1}}(\tilde{b}), \quad (5)$$

where for all $j \in \{1, \dots, m_2\}$, the function $\frac{dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,j}}}{dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,1}}} : \tilde{\mathcal{A}}_2 \rightarrow (0, +\infty)$ is the Radon-Nikodym derivative of $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,j}}$ with respect to $P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,1}}$.

B. Equilibrium

The necessary mathematical objects for characterizing the solution concept of the extended game $\mathcal{G}(u, P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2})$ in (2) are defined in this section.

Definition 2.1: (Best Responses of Player 1) Given a commitment $P \in \Delta(\mathcal{A}_2)$ and a channel output $\tilde{b} \in \tilde{\mathcal{A}}_2$, the set of best responses of Player 1, is determined by the correspondence $\text{BR}_1 : \Delta(\mathcal{A}_2) \times \tilde{\mathcal{A}}_2 \rightarrow \mathcal{P}(\Delta(\mathcal{A}_1))$, where $\mathcal{P}(\Delta(\mathcal{A}_1))$ denotes the power set of \mathcal{A}_1 , such that

$$\text{BR}_1(P, \tilde{b}) = \arg \max_{Q \in \Delta(\mathcal{A}_1)} \sum_{i=1}^{m_1} Q(a_{1,i}) \left(\sum_{j=1}^{m_2} u_{i,j} P(a_{2,j}) \frac{dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,j}}(\tilde{b})}{dP_{\tilde{\mathcal{A}}_2|\mathcal{A}_2=a_{2,1}}(\tilde{b})} \right). \quad (6)$$

A strategy $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2} \in \Delta(\mathcal{A}_1|\tilde{\mathcal{A}}_2)$ of Player 1 is a best response to the commitment $P \in \Delta(\mathcal{A}_2)$, if for all $\tilde{b} \in \tilde{\mathcal{A}}_2$ the probability measure $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}} \in \text{BR}_1(P, \tilde{b})$.

Let the function $\hat{v} : \Delta(\mathcal{A}_2) \rightarrow \mathbb{R}$ be such that for a given commitment $P \in \Delta(\mathcal{A}_2)$,

$$\hat{v}(P) = \max_{Q_{\mathcal{A}_1|\tilde{\mathcal{A}}_2} \in \Delta(\mathcal{A}_1|\tilde{\mathcal{A}}_2)} v(Q_{\mathcal{A}_1|\tilde{\mathcal{A}}_2}, P), \quad (7)$$

s.t. $\forall \tilde{b} \in \tilde{\mathcal{A}}_2 \quad Q_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}} \in \text{BR}_1(P, \tilde{b})$,

where the function v is defined in (5) and the correspondence BR_1 in (6).

Accordingly, the set of optimal commitments is the set of minimizers of the function \hat{v} given in (7). This observation leads to the following notion of equilibrium.

Definition 2.2: (Equilibrium) The strategies $P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2}^\dagger \in \Delta(\mathcal{A}_1|\tilde{\mathcal{A}}_2)$ and $P_{\mathcal{A}_2}^\dagger \in \Delta(\mathcal{A}_2)$ for Player 1 and Player 2, respectively, are said to form an equilibrium of the game $\mathcal{G}(u, P_{\tilde{\mathcal{A}}_2|\mathcal{A}_2})$ in (2) if

$$P_{\mathcal{A}_2}^\dagger \in \arg \min_{P \in \Delta(\mathcal{A}_2)} \hat{v}(P) \quad (8)$$

and for all $\tilde{b} \in \tilde{\mathcal{A}}_2$

$$P_{\mathcal{A}_1|\tilde{\mathcal{A}}_2=\tilde{b}}^\dagger \in \text{BR}_1(P_{\mathcal{A}_2}^\dagger, \tilde{b}), \quad (9)$$

where the function \hat{v} is defined in (7) and the correspondence BR_1 is defined in (6).

C. Existing Special Cases of Game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$

Consider that the channel $P_{\tilde{A}_2|A_2}$ is such that for all $\tilde{b} \in \tilde{\mathcal{A}}_2$, $P_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b}) = P_{\tilde{A}_2|A_2=a_{2,r}}(\tilde{b})$, with $(j, r) = \{1, \dots, m_2\}^2$ and $j \neq r$. This is the case in which the channel input and output are independent. In such case, the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ simplifies to the game denoted by

$$\mathcal{G}(u) \triangleq (\mathcal{A}_1, \mathcal{A}_2, u). \quad (10)$$

When players simultaneously choose their actions in the absence of commitment the game in (10) has been studied in [17]. The solution concept is the Nash Equilibrium (NE), as introduced in [18].

When there exist a deterministic one-to-one mapping between the inputs and the outputs of the channel, the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ is reduced to the game studied in [19] and [20], [21]. This is the case in which for all $a_{2,j} \in \mathcal{A}_2$, with $j = \{1, \dots, m_2\}$, it holds that $P_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b}) = 1$. That is, regardless the commitment, Player 1 can always select an optimal action in response to Player 2's action.

When the commitment is the only information that Player 1 has regarding the action played by Player 2, *i.e.*, the channel input and output are independent, the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ is equivalent to the one investigated in [22] and [23], and the solution concept is the SE in mixed strategies.

Other special cases of the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ are studied in [24] and [11]. In the former, the strategies of the players are restricted to pure strategies and the sets \mathcal{A}_2 and $\tilde{\mathcal{A}}_2$ are finite and assumed to be identical. In the latter, the analysis is generalized to include the consideration of mixed strategies.

III. MAIN RESULTS

For the rest of this work, it is assumed that for all $(i, l) \in \{1, \dots, m_1\}^2$,

$$\lambda \left(\left\{ \tilde{b} \in \tilde{\mathcal{A}}_2 : \sum_{j=1}^{m_2} (u_{i,j} - u_{l,j}) P(a_{2,j}) f_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b}) = 0 \right\} \right) = 0, \quad (11)$$

where $f_{\tilde{A}_2|A_2=a_{2,j}}$ is the corresponding pdf of the measure $P_{\tilde{A}_2|A_2=a_{2,j}}$.

The assumption in (11) is always satisfied when for all $(i, l) \in \{1, \dots, m_1\}^2$, it holds that $\prod_{j=1}^{m_2} (u_{i,j} - u_{l,j}) P(a_{2,j}) > 0$. In [11], it has been shown that imposing such condition restricts the game $\mathcal{G}(u)$ to exhibit a unique NE in mixed strategies.

The following lemma describes the set of best responses of Player 1.

Lemma 3.1: Given a commitment $P \in \Delta(\mathcal{A}_2)$ and a channel output $\tilde{b} \in \tilde{\mathcal{A}}_2$, the set of best responses of Player 1 is

$$\text{BR}_1(P, \tilde{b}) \triangleq \{Q \in \Delta(\mathcal{A}_1) : Q(a) > 0 \text{ if } a \in \mathcal{S}(P, \tilde{b})\}, \quad (12)$$

where

$$\mathcal{S}(P, \tilde{b}) \triangleq \text{argmax}_{a \in \mathcal{A}_1} \left\{ \sum_{j=1}^{m_2} u(a, a_{2,j}) P(a_{2,j}) \frac{dP_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b})}{dP_{\tilde{A}_2|A_2=a_{2,1}}} \right\}. \quad (13)$$

Proof: The proof of is presented in Appendix A of [25]. ■

The preceding lemma indicates the fact that given a commitment $P \in \Delta(\mathcal{A}_2)$ and a channel output $\tilde{b} \in \tilde{\mathcal{A}}_2$, Player 1's optimal strategy is to concentrate the probability measure $P_{A_1|\tilde{A}_2=\tilde{b}}$ over the set $\mathcal{S}(P, \tilde{b}) \subseteq \mathcal{A}_1$, with $|\mathcal{S}(P, \tilde{b})| = 1$. If $|\mathcal{S}(P, \tilde{b})| = 1$, there is a unique best response $P_{A_1|\tilde{A}_2=\tilde{b}}^* \in \text{BR}_1(P, \tilde{b})$ that assigns probability one to the singleton $\mathcal{S}(P, \tilde{b})$. The set of best responses of Player 1 might also contain infinitely many strategies. In the specific case in which $\text{BR}_1(P, \tilde{b}) = \Delta(\mathcal{A}_1)$, then as observed in [11], Player 1 opts for indifferently choosing its actions. Interestingly, under assumptions (4) and (11), for each channel output, Player 1 deterministically chooses a specific action almost surely, which will be discussed in the sequel.

For all $i \in \{1, \dots, m_1\}$ and for all $P \in \Delta(\mathcal{A}_2)$, let the set $\mathcal{H}_i(P)$ be

$$\mathcal{H}_i(P) \triangleq \left\{ \tilde{b} \in \tilde{\mathcal{A}}_2 : i \in \text{argmax}_{j=1}^{m_2} \left\{ \sum_{j=1}^{m_2} u_{i,j} P(a_{2,j}) \frac{dP_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b})}{dP_{\tilde{A}_2|A_2=a_{2,1}}} \right\} \right\} \setminus \mathcal{H}_{i-1}(P), \quad (14)$$

with $\mathcal{H}_0(P) \triangleq \emptyset$. Let also for all $(i, l) \in \{1, \dots, m_1\}^2$, with $i \neq l$, the set $\mathcal{H}_{i,l}$ be

$$\mathcal{H}_{i,l}(P) \triangleq \left\{ \tilde{b} \in \tilde{\mathcal{A}}_2 : \sum_{j=1}^{m_2} u_{i,j} P(a_{2,j}) \frac{dP_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b})}{dP_{\tilde{A}_2|A_2=a_{2,k}}} \right. \\ \left. = \sum_{j=1}^{m_2} u_{l,j} P(a_{2,j}) \frac{dP_{\tilde{A}_2|A_2=a_{2,j}}(\tilde{b})}{dP_{\tilde{A}_2|A_2=a_{2,1}}} \right\}. \quad (15)$$

Note that $\mathcal{H}_1(P), \mathcal{H}_2(P), \dots, \mathcal{H}_{m_1}(P)$ form a partition of $\tilde{\mathcal{A}}_2$. The following theorem presents a property exhibited by the set $\mathcal{H}_{i,l}(P)$ in (15).

Theorem 3.1: Given a commitment $P \in \Delta(\mathcal{A}_2)$, under assumptions (4) and (11), if for all $(i, l) \in \{1, \dots, m_1\}^2$, with $i \neq l$, and for all $a_{2,j} \in \mathcal{A}_2$

$$(u_{i,j} - u_{l,j}) P(a_{2,j}) \neq 0, \quad (16)$$

then it holds that

$$P_{\tilde{A}_2}(\mathcal{H}_{i,l}(P)) = 0, \quad (17)$$

where $\mathcal{H}_{i,l}(P)$ is defined in (15); and the probability measure $P_{\tilde{A}_2}$ is such that for all measurable subsets \mathcal{A} of $\tilde{\mathcal{A}}_2$, it holds that

$$P_{\tilde{A}_2}(\mathcal{A}) = \int P_{\tilde{A}_2|A_2=a}(\mathcal{A}) dP(a). \quad (18)$$

Proof: The proof of is presented in Appendix B of [25]. ■

A first observation from Theorem 3.1, is that if Player 2's commitment is in strict mixed strategies, then the elements of the set $\mathcal{H}_{i,l}(P)$ are observed with probability zero with respect to $P_{\tilde{A}_2}$ in (18). Note that an element \tilde{b} of the set $\mathcal{H}_{i,l}(P)$ is the output of the channel for which actions $a_{1,i}$ and $a_{1,l}$ might be both best responses for Player 1 to the commitment P and the channel output \tilde{b} . Hence, under the assumptions of Theorem 3.1, the set $\mathcal{S}(P, \tilde{b})$ is always a singleton and thus $|\text{BR}_1(P, \tilde{b})| = 1$. On the other hand, if (16) does not hold, then the set of best responses of Player 1 includes infinitely many strategies. This indicates that Player 1 indifferently chooses its actions from the set $\mathcal{S}(P, \tilde{b})$. This is formalized by the following lemma.

Lemma 3.2: Under assumptions (4) and (11), given a commitment $P \in \Delta(\mathcal{A}_2)$ that satisfies (16), for all $\tilde{b} \in \tilde{\mathcal{A}}_2$ and for all $P_{A_1|\tilde{A}_2} \in \text{BR}_1(P, \tilde{b})$, with $\text{BR}_1(P, \tilde{b})$ in (6), it holds that for all $i \in \{1, \dots, m_1\}$

$$P_{A_1|\tilde{A}_2=\tilde{b}}(a_{1,i}) = \begin{cases} 1, & \text{if } \tilde{b} \in \mathcal{H}_i(P) \\ 0, & \text{otherwise,} \end{cases} \quad (19)$$

where $\mathcal{H}_i(P)$ is defined in (14) and $a_{1,i} \in \mathcal{A}_1$. Moreover, the set of best responses satisfies

$$P_{\tilde{A}_2} \left(\left\{ \tilde{b} \in \tilde{\mathcal{A}}_2 : |\text{BR}_1(P, \tilde{b})| = 1 \right\} \right) = 1. \quad (20)$$

Proof: The proof of is presented in Appendix C of [25]. ■

The following theorem characterizes the equilibrium of the extended game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ in (2).

Theorem 3.2: Under assumptions (4) and (11), and if (16) holds for all $j \in \{1, \dots, m_2\}$ and for all $(i, l) \in \{1, \dots, m_1\}^2$, with $i \neq l$, then the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ possesses a unique equilibrium almost surely with respect to $P_{\tilde{A}_2}$.

Proof: The proof is presented in Appendix D of [25]. ■

The game $\mathcal{G}(u)$ in (10) holds significant importance due to the fact that its the expected payoff can be used as reference point when analyzing the game in (2). Let the function $\omega : \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2) \rightarrow \mathbb{R}$ represent the expected payoff of game $\mathcal{G}(u)$ such that

$$\omega(P_{A_1}, P_{A_2}) = \sum_{i=1}^{m_1} \sum_{j=1}^{m_2} P_{A_1}(a_{1,i}) u_{i,j} P_{A_2}(a_{2,j}). \quad (21)$$

Let the function $\hat{u} : \Delta(\mathcal{A}_2) \rightarrow \mathbb{R}$ be such that for all commitments $P \in \Delta(\mathcal{A}_2)$,

$$\hat{u}(P) = \max_{Q \in \Delta(\mathcal{A}_1)} \omega(Q, P), \quad (22)$$

where the function ω is defined in (21). The term $\hat{u}(P)$ in (22) represents the expected payoff of the game $\mathcal{G}(u)$ in (10), when Player 2 commits to play strategy P . Note that, the minimum value over all commitments P of the function \hat{u} is the payoff

at the NE. The following lemma compares the payoffs at the equilibria of the games $\mathcal{G}(u)$ in (10) and $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ in (2).

Lemma 3.3: Let the probability measures $(P_{A_1|\tilde{A}_2}^\dagger, P_{A_2}^\dagger)$ form an equilibrium of the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ in (2) and let the pair of strategies $(P_{A_1}^*, P_{A_2}^*) \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2)$ be an NE of the game $\mathcal{G}(u)$ in (10). Then for all commitments $P \in \Delta(\mathcal{A}_2)$, it holds that

$$\hat{u}(P_{A_2}^*) \leq v \left(P_{A_1|\tilde{A}_2}^\dagger, P_{A_2}^\dagger \right) \leq \hat{v}(P) \leq \min_j \max_i u_{i,j}, \quad (23)$$

where the functions \hat{u} , \hat{v} and v are given in (22), (7) and (5), respectively.

Proof: The proof is presented in Appendix E of [25]. ■ Lemma 3.3 unveils the fact that the payoff of the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ at the equilibrium is lower bounded by the payoff of the game $\mathcal{G}(u)$ at the NE and upper bounded by the payoff of such a game at the SE in pure strategies. The former is the case in which Player 1 does not observe the actions of Player 2 (e.g. very noisy channel, $I(P_{\tilde{A}_2|A_2}; P_{A_2}) = 0$); and the latter corresponds to the case in which Player 1 perfectly observes the actions of its opponent (e.g. channel with zero noise, $I(P_{\tilde{A}_2|A_2}; P_{A_2}) = H(P_{A_2})$).

IV. EXAMPLE: GAUSSIAN CHANNEL

Thus far, the probability measure that describes the channel $P_{\tilde{A}_2|A_2}$ of the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ in (2) is not yet specified. This generality signifies a key feature of the theoretical model, which captures several data processing impairments. Moreover, while the specific game solution concept relies on the characteristics of the channel, the generality of the game formulation presented in Section II highlights its capability to effectively address any type of noise which has a density with respect the Lebesgue measure. In this section, the channel is assumed to be an AWGN channel and the action sets of Player 1 and Player 2 have the same cardinality and contain two actions. In this scenario, the channel is such that for all $a_{2,j} \in \mathcal{A}_2$ and all measurable subsets \mathcal{A} of $\tilde{\mathcal{A}}_2$

$$P_{\tilde{A}_2|A_2=a_{2,j}}(\mathcal{A}) = \frac{1}{\sqrt{2\pi\sigma_j^2}} \int_{\mathcal{A}} \exp\left(-\frac{1}{2\sigma_j^2}(\tilde{b} - \mu_j - a_{2,j})^2\right) d\lambda(\tilde{b}), \quad (24)$$

where λ is the Lebesgue measure in \mathbb{R} ; $\mu_j - a_{2,j}$ and σ_j^2 are the mean and the variance, respectively. From (24), it holds that the Radon-Nikodym derivative of $P_{\tilde{A}_2|A_2=a_{2,j}}$ with respect to $P_{\tilde{A}_2|A_2=a_{2,1}}$, and for all $\tilde{b} \in \tilde{\mathcal{A}}_2$ is

$$\frac{dP_{\tilde{A}_2|A_2=a_{2,2}}}{dP_{\tilde{A}_2|A_2=a_{2,1}}}(\tilde{b}) = \exp\left(-\left(\frac{(a_{2,2}^2 - a_{2,1}^2)}{2\sigma^2} + \frac{\tilde{b}(a_{2,1} - a_{2,2})}{\sigma^2}\right)\right), \quad (25)$$

where it is assumed that $\mu_j = 0$ and $\sigma_j = \sigma > 0$, with $j \in \{1, 2\}$. From (15), for all $P_{A_2} \in \Delta(\mathcal{A}_2)$ the set $\mathcal{H}_{i,l}(P_{A_2})$,

with $(i, l) \in \{1, 2\}^2$, satisfies

$$\mathcal{H}_{1,2}(P_{A_2}) \triangleq \left\{ \tilde{b} \in \tilde{\mathcal{A}}_2: \frac{dP_{\tilde{A}_2|A_2=a_{2,2}}(\tilde{b})}{dP_{\tilde{A}_2|A_2=a_{2,1}}(\tilde{b})} = \frac{(u_{1,1}-u_{2,1})P_{A_2}(a_{2,1})}{(u_{2,2}-u_{1,2})P_{A_2}(a_{2,2})} \right\} \quad (26)$$

$$= \left\{ \frac{\sigma^2}{a_{2,1} - a_{2,2}} \left(\log \left(\frac{(u_{1,1} - u_{2,1}) P_{A_2}(a_{2,1})}{(u_{2,2} - u_{1,2}) P_{A_2}(a_{2,2})} \right) + D \left(P_{\tilde{A}_2|A_2=a_{2,2}} \| P_{\tilde{A}_2|A_2=a_{2,1}} \right) \right) \right\}, \quad (27)$$

where D denotes the Kullback-Leibler divergence. Note that the set $\mathcal{H}_{1,2}(P_{A_2})$ is a singleton. This said, it holds that $\lambda(\mathcal{H}_{1,2}(P_{A_2})) = 0$ and thus, $P_{\tilde{A}_2}(\mathcal{H}_{1,2}(P_{A_2})) = 0$, which demonstrates the validity of Theorem 3.1.

Let the complementary cumulative distribution function (ccdf) of the Gaussian pdf be

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{+\infty} \exp\left(-\frac{z^2}{2}\right) dz. \quad (28)$$

Using this notation, the following lemma provides an explicit expression for the function \hat{v} in (7)

Lemma 4.1: Given a commitment $P_{A_2} \in \Delta(\mathcal{A}_2)$, if for all $a_{2,j} \in \mathcal{A}_2$, the probability measure $P_{\tilde{A}_2|A_2=a_{2,j}}$ satisfies (24), then it holds

$$\begin{aligned} \hat{v}(P_{A_2}) &= \max_{Q_{A_1|\tilde{A}_2} \in \text{BR}_1(P_{A_2}, \tilde{b})} v(Q_{A_1|\tilde{A}_2}, P_{A_2}) \\ &= \sum_{j=1}^{m_2} P_{A_2}(a_{2,j}) \left(u_{2,j} + (u_{1,j} - u_{2,j}) Q \left(\frac{b_{P_{A_2}}^* - a_{2,j}}{\sigma} \right) \right), \end{aligned} \quad (29)$$


where function v is defined in (5); $b_{P_{A_2}}^* \in \mathcal{H}_{1,2}(P_{A_2})$; and function Q is in (28).

Proof: The proof is presented in Appendix F of [25]. ■


In Fig. 1a, the strategies $(P_{A_1}^*, P_{A_2}^*)$ with $P_{A_1}^*(a_{1,1}) = 1 - P_{A_1}^*(a_{1,2}) = 0.23$ and $P_{A_2}^*(a_{2,1}) = 1 - P_{A_2}^*(a_{2,2}) = 0.45$ form the unique NE of the game $\mathcal{G}(u)$ whose payoff is $\hat{u}(P_{A_2}^*)$ (red star). Examining the behavior of the function \hat{v} in (7) for a range of noise variance values σ^2 , it is observed that; (i) the equilibrium of the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ is always lower bounded by the equilibrium of \hat{u} (dashed blue line in the zoomed subfigure) as Lemma 3.3 revealed; (ii) the function \hat{v} is monotonically decreasing with respect to σ^2 ; (iii) letting σ^2 grow to infinity, results in the equivalence of the equilibrium of $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ and the NE of game $\mathcal{G}(u)$ (green curve approaching the red star); and (iv) letting σ^2 to be close to zero, the equilibrium payoff of the game $\mathcal{G}(u, P_{\tilde{A}_2|A_2})$ in (2) becomes identical to the SE in pure strategies of the game $\mathcal{G}(u)$ in (10) which is the min-max solution in pure strategies (red line).

In Fig. 1b, the game $\mathcal{G}(u)$ exhibits a unique equilibrium in pure strategies and the NE payoff of the game $\mathcal{G}(u)$ is $\hat{u}(P_{A_2}^*)$ (red star). Here, the equilibrium payoffs of both games are identical. In both examples it holds that $\hat{u}(P_{A_2}^*) \leq$

$v(P_{A_1}^\dagger|_{\tilde{A}_2}, P_{A_2}^\dagger)$, which signifies that even subject to noise, the equilibrium payoff increases when Player 1 observes a noisy version of the action played by Player 2.



(a) Payoff matrix $u = (8, -6; -2, 2)$.



(b) Payoff matrix $u = (-5, 1; -6, 3)$.

Fig. 1: Plots of \hat{v} in (7) and \hat{u} in (22) as a function of the commitment P_{A_2} . The channel is the AWGN defined in (24) with variance $\sigma_1^2 = \sigma_2^2 = \sigma^2$; and $-a_{2,1} = a_{2,2} = 10^2$.

REFERENCES

- [1] N. Dalvi, P. Domingos, Mausam, S. Sanghai, and D. Verma, "Adversarial classification," in *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (New York, NY, USA), pp. 99–108, 2004.
- [2] D. Lowd and C. Meek, "Adversarial learning," in *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, (New York, NY, USA), pp. 641–647, 2005.
- [3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, A. Ozair, S. and Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Advances in Neural Information Processing Systems*, (Montréal, Canada), pp. 1–9, Dec. 2014.

- [4] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *Proc. of the International Conference on Learning Representations (ICLR)*, (Banff, Canada), Apr. 2014.
- [5] Y. Hsieh, C. Liu, and V. Cevher, "Finding mixed Nash equilibria of generative adversarial networks," in *Proc. Int. Conf. on Machine Learning*, (Long Beach, California, USA), pp. 2810–2819, Jun. 2019.
- [6] F. A. Oliehoek, R. Savani, J. Gallego, E. Pol, and R. Groß, "Beyond local Nash equilibria for adversarial networks," in *Proc. Benelux Conf. on Artificial Intelligence*, ('s-Hertogenbosch, Netherlands), pp. 73–89, Nov. 2018.
- [7] H. von Stackelberg, *Theory of the Market Economy*. Oxford University Press, Mar. 1952.
- [8] P. Huang, M. Xu, F. Fang, and D. Zhao, "Robust reinforcement learning as a Stackelberg game via adaptively-regularized adversarial training," in *Proc. Int. Joint Conf. on Artificial Intelligence*, (Vienna, Austria), pp. 3099–3106, Jul. 2022.
- [9] M. Brückner and T. Scheffer, "Stackelberg games for adversarial prediction problems," in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, (San Diego, California, USA), pp. 547–555, Aug. 2011.
- [10] S. Zuo, C. Liang, H. Jiang, X. Liu, P. He, J. Gao, W. Chen, and T. Zhao, "Adversarial training as Stackelberg game: An unrolled optimization approach," *arXiv preprint arXiv:2104.04886*, 2021.
- [11] K. Sun, S. M. Perlaza, and A. Jean-Marie, "2 x 2 zero-sum games with commitments and noisy observations," in *Proc. of the IEEE International Symposium on Information Theory (ISIT)*, 2023.
- [12] X. Gao, S. Liu, and L. Yu, "Achieving optimal adversarial accuracy for adversarial deep learning using Stackelberg games," *Acta Mathematica Scientia*, vol. 42B, pp. 2399–2418, Aug. 2022.
- [13] C. Jin, P. Netrapalli, and M. Jordan, "What is local optimality in nonconvex-nonconcave minimax optimization?," in *Proc. Int. Conf. on Machine Learning*, (Virtual), pp. 4880–4889, Jul. 2020.
- [14] Y. Bai, C. Jin, H. Wang, and C. Xiong, "Sample-efficient learning of Stackelberg equilibria in general-sum games," in *Proc. Advances in Neural Information Processing Systems*, vol. 34, (Virtual), pp. 25799–25811, Dec. 2021.
- [15] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, pp. 379–423, Jul. 1948.
- [16] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, pp. 623–656, Oct. 1948.
- [17] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior: 60th Anniversary Commemorative Edition*. Princeton University Press, 2007.
- [18] J. F. Nash, "Equilibrium points in n-person games," *Proceedings of the National Academy of Sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [19] H. von Stackelberg, *Marktform und Gleichgewicht*. J. Springer, 1934.
- [20] M. Simaan and J. B. Cruz, "Additional aspects of the Stackelberg strategy in nonzero-sum games," *Journal of Optimization Theory and Applications*, vol. 11, no. 6, pp. 613–626, 1973.
- [21] M. Simaan and J. B. Cruz, "On the Stackelberg strategy in nonzero-sum games," *Journal of Optimization Theory and Applications*, vol. 11, pp. 533–555, May 1973.
- [22] V. Conitzer, "On Stackelberg mixed strategies," *Synthese*, vol. 193, pp. 689–703, Mar. 2016.
- [23] S. Leonardos and C. Melolidakis, "On the commitment value and commitment optimal strategies in bimatrix games," *International Game Theory Review*, vol. 20, p. 1840001, Sep. 2018.
- [24] K. Bagwell, "Commitment and observability in games," *Games and Economic Behavior*, vol. 8, no. 2, pp. 271–280, 1995.
- [25] E. M. Athanaskos and S. M. Perlaza, "Leveraging noisy observations in zero-sum games," Tech. Rep. RR-9538, INRIA, Centre Inria d'Université Côte d'Azur, Sophia Antipolis, France, Jan. 2024.