Stealthy strategies for deception in hypergames with asymmetric information

Bahman Gharesifard Jorge Cortés

***Abstract*— This paper considers games with incomplete asymmetric information, where one player (the deceiver) has privileged information about the other (the mark) and intends to employ it for belief manipulation. We use hypergames to represent the asymmetric information available to players and assume a probabilistic model for the actions of the mark. This framework allows us to formalize various notions of deception in a precise way. We provide a necessary condition and a sufficient condition for deceivability when the deceiver is allowed to reveal information to the mark as the game evolves. For the case when the deceiver acts stealthily, i.e., restricts her actions to those that do not contradict the belief of the mark, we are able to fully characterize when deception is possible. Moreover, we design the `worst-case max-strategy` that, when such a sequence of deceiving actions exists, is guaranteed to find it. An example illustrates our results.**

I. Introduction

Informational asymmetries in strategic scenarios provide opportunities for manipulating beliefs or inducing certain desired perceptions. In this paper, we consider a class of games where one player (the deceiver) wishes to misrepresent certain information in order to gain a strategic advantage over the opponent (the mark). In our framework, the deceiver can anticipate the effect that her actions will have on the mark's belief structure. In this sense, the deception goal can be understood as steering the evolution of a particular dynamical system into a desired set of outcomes. Scenarios of interest includes bargaining, cybersecurity, military operations, and human behavior modeling.

Literature review: In strategic scenarios with informational asymmetries [1], players may decide not to disclose some information (passive deception) or lie about a value of interest to the opponent (active deception). Within the context of games of incomplete information, deception has not been studied in a systematic way with the exception of a few references. [2] demonstrates that the inconsistent structure of beliefs can lead to counterintuitive behaviors. [3] studies deception via strategic communication, in which a 'sophisticated' player sends either truthful or false messages to the opponents. [4] investigates the vulnerability of strategic decision makers to persuasion. The recent work [5] constructs a theory of deception for games with incomplete information where players form expectations about the average behavior of the other players based on past histories. [6], [7] consider scenarios where one player has access to certain information and can distort it before it is passed on to others. In this paper, we make use of hypergames [8], [9], [10], since they provide a natural framework for modeling strategic situations with asymmetric information among players. Early references on deception in dynamic games with imperfect information include [11], [12]. The works [13], [14], [15] provide examples of how informational asymmetries can be used to induce false perceptions in the opponent and lead to strategic deception. The works [16], [17] provide deception-robust schemes for a class of discrete dynamic stochastic games under imperfect observations.

Statement of contributions: We consider games of incomplete information where players have different perceptions about the scenarios they are involved in. Specifically, we study a class of 2-player hypergames where the deceiver has full information about the mark's game and intends to plant a certain belief in her. The mark is a rational player that observes the actions taken by the deceiver and assumes she acts rationally (although she may not), and updates her perception about the opponent's preferences accordingly. From the deceiver's viewpoint, the mark's actions are rational and probabilistic.

This framework sets the stage for the first contribution of the paper, which is the introduction of precise notions of deception to capture different forms of belief manipulation. These notions allow us to identify a necessary condition and a sufficient condition for deceivability on the mark's belief structure. Next, we study scenarios where the deceiver purposefully restricts her set of actions to those that do not contradict the mark's belief structure. We term these actions stealthy and fully characterize when deception via such actions is possible. We show how the problem of finding a stealthy sequence of actions is equivalent to finding a longest path in an appropriate digraph that encodes the mark's belief structure. Our third contribution is then the design of the `worst-case max-strategy` that, given a desired deception objective, determines a stealthy sequence of actions that achieves it. We end the paper with an example to illustrate the notions and the results of the paper.

II. Preliminaries

This section introduces some basic notions regarding graph theory, Markov chains, and hypergames. We denote the set of real and positive real numbers by $\mathbb{R}$ and $\mathbb{R}_{>0}$, respectively. We denote by $\mathbb{Z}_{\geq 0}$ and $\mathbb{Z}_{\geq 1}$ the set of nonnegative and positive integers, respectively. A nonempty set X along with a preorder $\succeq$, i.e., a reflexive and transitive binary relation, is called a *directed set* if for every pair of elements in X there exists an upper bound with respect to the preorder. We use $\sigma = (x_1, x_2, \ldots)$, where $x_1, x_2, \ldots \in X$, to denote

The authors are with the Department of Mechanical and Aerospace Engineering, University of California, San Diego, California, USA, {bgharesifard,cortes}@ucsd.edu

a sequence of elements in X. Note that a finite sequence of $k \in \mathbb{Z}_{\geq 1}$ elements is simply a k-tuple.

A. Graph theory

We recall some basic notions from [18]. A *digraph* G is a pair (V, E), where V is a finite set, called the vertex set, and $E \subseteq V \times V$, called the edge set. Given an edge $(u, v) \in E$, u is an *in-neighbor* of v and v is an *out-neighbor* of u. The set of in-neighbors and out-neighbors of v are denoted, respectively, by $\mathcal{N}^{\text{in}}(v)$ and $\mathcal{N}^{\text{out}}(v)$. The *in-degree* and *out-degree* of v are the number of in-neighbors and out-neighbors of v, respectively. $\mathcal{A}$ is an adjacency matrix for $G = (V, E)$ if the following holds: for each $v_i, v_j \in V$, $a_{ij} > 0$ if and only if $(v_i, v_j) \in E$. A *(directed) path* is an ordered sequence of vertices so that any two consecutive vertices in the sequence are an edge of the digraph. A *cycle* in a digraph is a directed path that starts and ends at the same vertex and has no other repeated vertex. A digraph is called *acyclic* if it does not contain any cycle.

B. Markov chains

We recall here some basic notions from Markov chains following [19]. We denote by $(\Omega, \mathcal{F}, \mathbb{P})$ a *probability space*, where Ω is a countable set, $\mathcal{F}$ is a σ-algebra over Ω, and $\mathbb{P}$ is a probability measure. An *E-valued random* variable is a measurable mapping $X : (\Omega, \mathcal{F}, \mathbb{P}) \to (E, \mathcal{E})$, where $\mathcal{E}$ is a σ-algebra over E and $(E, \mathcal{E})$ is a measurable space. A *Markov chain* is a sequence of random variables $(X_1, X_2, \ldots)$ such that, for all $n \in \mathbb{Z}_{\geq 1}$ and $x \in \Omega$,

$$\mathbb{P}(X_{n+1} = x \,|\, X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n) = \mathbb{P}(X_{n+1} = x \,|\, X_n = x_n).$$

The *probability transition kernel* $T_{\mathbb{P}}$ is

$$T_{\mathbb{P}}(x_i, x_j) = \mathbb{P}(X_{n+1} = x_i \,|\, X_n = x_j),$$

where $x_i, x_j \in \Omega$. Note that for every $x \in \Omega$, $T_{\mathbb{P}}(x, .)$ is also a probability measure on Ω. One can inductively define

$$\begin{aligned} T^k_{\mathbb{P}}(x_i, x_j) &:= \mathbb{P}(X_{n+k} = x_i \,|\, X_n = x_j) \\ &= \sum_{z \in \Omega} T_{\mathbb{P}}(x_i, z) T^{k-1}_{\mathbb{P}}(z, x_j). \end{aligned}$$

If there exists $k \in \mathbb{Z}_{\geq 1}$ such that $T^k_{\mathbb{P}}(x_i, x_j) > 0$, the state x_i is *reachable* from x_j (or, equivalently, that x_j *communicates* with x_i). We denote the set of all states reachable from x_j, with respect to the transition probability $T_{\mathbb{P}}$, by

$$\mathcal{R}_{T_{\mathbb{P}}}(x_j) = \{x_i \in \Omega \mid \exists k_i \in \mathbb{Z}_{\geq 1},\ T^{k_i}_{\mathbb{P}}(x_i, x_j) > 0\}.$$

C. Hypergame theory

We consider games with inconsistent perceptions across the players in the framework of hypergames [20], [8], [9]. A 0-level hypergame is simply a *(finite) game*, i.e., a triplet $\mathbf{G} = (V, \mathbf{S}_{\text{outcome}}, \mathbf{P})$, where V is a set of $n \in \mathbb{Z}_{\geq 1}$ players, $\mathbf{S}_{\text{outcome}} = S_1 \times \ldots \times S_n$ is the outcome set with finite cardinality $N = |\mathbf{S}_{\text{outcome}}| \in \mathbb{Z}_{\geq 1}$ and $\mathbf{P} = (P_1, \ldots, P_n)$, with $P_i = (x_1, \ldots, x_N)^T \in \mathbf{S}_p$ the preference vector of player v_i, $i \in \{1, \ldots, n\}$. Here, S_i is a finite set of actions available to player $v_i \in V$ and $\mathbf{S}_p \subset \mathbf{S}^N_{\text{outcome}}$ is the set of all elements in $\mathbf{S}^N_{\text{outcome}}$ with pairwise different entries. We denote by π_i the projection of $\mathbf{S}_{\text{outcome}}$ onto S_i.

A n*-person* 1*-level hypergame* is a set $H^1 = \{\mathbf{G}_1, \ldots, \mathbf{G}_n\}$, where $\mathbf{G}_i = (V, (\mathbf{S}_{\text{outcome}})_i, \mathbf{P}_i)$, for $i \in \{1, \ldots, n\}$, is the subjective finite game of player $v_i \in V$, and V is a set of n players; $(\mathbf{S}_{\text{outcome}})_i = S_{1i} \times \ldots \times S_{ni}$, with S_{ji} the finite set of strategies available to v_j, as perceived by v_i; $\mathbf{P}_i = (P_{1i}, \ldots, P_{ni})$, with P_{ji} the preference vector of v_j, as perceived by v_i. In a 1-level hypergame, each player $v_i \in V$ plays the game $\mathbf{G}_i$ with the perception that she is playing a game with complete information. The definition of 1-level hypergame can be extended to higher-order hypergames as follows: a n*-person* k*-level hypergame*, $k \geq 1$, is a set $H^k = \{H^{k_1}_1, \ldots, H^{k_n}_n\}$, where $k_i \leq k - 1$ and at least one k_i is equal to $k - 1$.

1) Stability and equilibria: Here we recall the notion of stability for 2-person 1-level hypergames. This class of hypergames is the simplest one that explicitly models the perception of players about their opponents' preferences (the reader is referred to [8] for the extension to higher-order hypergames). Let $H^1 = \{H^0_A, H^0_B\}$. Here, $H^0_A = (\mathrm{P}_{AA}, \mathrm{P}_{BA})$ is the 0-level hypergame for player A, where P_{AA} and P_{BA} are, respectively, the preferences of player A and player B perceived by player A. The same convention holds for $H^0_B = (\mathrm{P}_{AB}, \mathrm{P}_{BB})$. For simplicity, the 0-level hypergames have the same set of outcomes $\mathbf{S}_{\text{outcome}}$. We denote by $\succeq_{\mathrm{P}_{IJ}}$ the binary relation on $\mathbf{S}_{\text{outcome}}$ induced by P_{IJ}, where $I, J \in \{A, B\}$. For convenience, we let $\mathbf{S}_{\text{outcome}}|_{\pi_I(x)} = \{y \in \mathbf{S}_{\text{outcome}} \mid \pi_I(y) = \pi_I(x)\}$ and refer to it as a restricted outcome set. We also find it useful to use I' to denote the opponent of I in $\{A, B\}$. We assign $\mathrm{rank}(x, \mathrm{P}_{IJ}) \in \mathbb{R}_{>0}$ to each outcome $x \in \mathbf{S}_{\text{outcome}}$ such that $\mathrm{rank}(y, \mathrm{P}_{IJ}) > \mathrm{rank}(x, \mathrm{P}_{IJ})$ if and only if $x \succ_{\mathrm{P}_{IJ}} y$ (players prefer the outcomes with lower ranks). We use the set $\{1, \ldots, |\mathbf{S}_{\text{outcome}}|\}$ to rank the outcomes.

Given two distinct outcomes $x, y \in \mathbf{S}_{\text{outcome}}$, y is an *improvement* from x for $I \in \{A, B\}$, perceived by $J \in \{A, B\}$ in H^0_J, if and only if $\pi_{I'}(y) = \pi_{I'}(x)$ and $y \succ_{\mathrm{P}_{IJ}} x$. An outcome $x \in \mathbf{S}_{\text{outcome}}$ is called *rational* for $I \in \{A, B\}$, as perceived by $J \in \{A, B\}$ in H^0_J, if there exists no improvement from x for I. The common notion of rationality in hypergames is the notion of sequential rationality [21], [9], [22]. An outcome $x \in \mathbf{S}_{\text{outcome}}$ is *sequentially rational* for $I \in \{A, B\}$ with respect to H^0_J, $J \in \{A, B\}$, if and only if for each improvement y for I, perceived by J in H^0_J, there exists an improvement z for I', perceived by J in H^0_J, such that $x \succ_{\mathrm{P}_{IJ}} z$. Whenever this holds, we say that the improvement z from y for I' *sanctions* the improvement y from x for I in H^0_J. By definition, a rational outcome is also sequentially rational. An outcome $x \in \mathbf{S}_{\text{outcome}}$ is *unstable* for player I with respect to H^0_J if it is not sequentially rational for player I, as perceived by player J and is an *equilibrium of* H^0_J if it is sequentially rational for both J and J', perceived by player J. An outcome x is an *equilibrium of* H^1 if it is sequentially rational for player A in H_A^0 and also for player B in H_B^0. Note that x can be an equilibrium for H^1 and not an equilibrium of H_A^0.

2) H-digraphs: The notion of H-digraph encodes the stability information of hypergames. Formally, the *H-digraph* associated to H_A^0 is $\mathcal{G}_{H_A^0} = (\mathbf{S}_{\text{outcome}}, \mathcal{E}_{H_A^0})$, where there exists an edge $(x, y) \in \mathcal{E}_{H_A^0}$ iff either there exists an improvement y from x for A for which there is no sanction of B in H_A^0, or there exists an improvement y from x for B for which there is no sanction of A in H_A^0. One can similarly construct $\mathcal{G}_{H_B^0}$. By definition, an outcome x is sequentially rational for A (respectively for B) if and only if $\mathcal{N}^{\text{out}}(x) \cap \mathbf{S}_{\text{outcome}}|_{\pi_B(x)} = \emptyset$ (respectively $\mathcal{N}^{\text{out}}(x) \cap \mathbf{S}_{\text{outcome}}|_{\pi_A(x)} = \emptyset$). Moreover, an outcome is an equilibrium for the hypergame H_A^0 if and only if its out-degree in the associated H-digraph is zero.

3) Learning in hypergames: Suppose players A and B take actions that change the outcome from x to y. If player A can perfectly observe B's action and believes that the opponent is rational, she concludes that player B prefers $(\pi_A(x), \pi_B(y))$ over x. Therefore, player A can incorporate this information into her hypergame and update her perception about the preferences of player B. Here, we recall a method called *swap learning* to do this, see [20]. These notions can similarly be defined for player B.

We start by an algebraic construction. Let V be a set of cardinality N and let W be the subset of V^N with pairwise different elements. For $x_1, x_2 \in V$, let $\text{swap}_{x_1 \mapsto x_2} : W \to W$ be defined by

$$(\text{swap}_{x_1\mapsto x_2}(v))_k = v_k \quad \text{if } v_k \neq x_1, x_2,$$
$$(\text{swap}_{x_1\mapsto x_2}(v))_i = \begin{cases} v_j & \text{if } v_i = x_1, v_j = x_2 \text{ and } i < j, \\ v_i & \text{if } v_i = x_1, v_j = x_2 \text{ and } i > j, \end{cases}$$
$$(\text{swap}_{x_1\mapsto x_2}(v))_j = \begin{cases} v_i & \text{if } v_i = x_1, v_j = x_2 \text{ and } i < j, \\ v_j & \text{if } v_i = x_1, v_j = x_2 \text{ and } i > j. \end{cases}$$

We refer to $\text{swap}_{x_1\mapsto x_2}$ as the x_1 *to* x_2 *swap map*. The *swap learning maps* $\mathbf{Sw}^A_{x,y} : \mathbf{S}_\text{P} \to \mathbf{S}_\text{P}$ for player A is given by

$$\mathbf{Sw}^A_{x,y}(\text{P}) = \text{swap}_{x\mapsto(\pi_A(x),\pi_B(y))}(\text{P}).$$

One can show [20] that if players are rational, swap learning is guaranteed to decrease the mismatch between a player's perception and the true payoff structure of other players.

When the outcome changes from x to y and player A updates her perception via swap learning, her H-digraph changes from $\mathcal{G}_{H_A^0}$ to $\mathbf{Sw}^A_{x,y}(\mathcal{G}_{H_A^0})$. Similarly, if players A and B repeatedly take actions such that the hypergame outcomes are $\sigma = (x_1, \dots, x_n) \in \mathbf{S}^n_{\text{outcome}}$, then the associated H-digraph of player A is denoted $\mathbf{Sw}^A_{x_1,\sigma}(\mathcal{G}_{H_A^0})$, where

$$\mathbf{Sw}^A_{x_1,\sigma} = \mathbf{Sw}^A_{x_1,x_2} \circ \mathbf{Sw}^A_{x_2,x_3} \circ \cdots \circ \mathbf{Sw}^A_{x_{n-1},x_n}.$$

We denote by $\mathbf{Sw}^A_{x_1,\sigma}(\mathcal{E}_{H_A^0})$ the edge set of $\mathbf{Sw}^A_{x_1,\sigma}(\mathcal{G}_{H_A^0})$.

III. Problem statement

In this paper, we consider 2-person 2-level hypergame. We assume player B has perfect knowledge about the preferences of player A, while A perfectly observes the actions of B and uses the swap learning map to update her perception. We focus on swap learning, although the analysis could also be carried out for other learning mechanisms. Formally, the situation described above corresponds to a 2-person 2-level hypergame $H^2 = \{H_A^0, H_B^1\}$, with $H_B^1 = \{H_{AB}^0, H_{BB}^0\}$ such that $H_{AB}^0 = H_A^0$. Since $H_{BB}^0 = H_B^0$, we actually have

$$H^2 = \{H_A^0, \{H_A^0, H_B^0\}\}.$$

Because of the special form of H^2, it is not difficult to see that the equilibria of H^2, as defined in [8], are exactly the same as the equilibria of the hypergame $H_B^1 = \{H_A^0, H_B^0\}$.

We assume that players take their actions sequentially, one after each other. This assumption matches up with the notion of sequential rationality and guarantees that the repeated play of any 0-level hypergame converges to an equilibrium [23]. Note that scenarios where one players takes multiple actions before the other player acts can also be accommodated. We formalize this concept next.

Definition 3.1 (Admissible sequence): A sequence of outcomes $\sigma = (x_0, x_1, x_2, \dots)$ in $\mathbf{S}_{\text{outcome}}$ is *admissible* if

(i) $\pi_I(x_{2i}) = \pi_I(x_{2i+1})$,
(ii) $\pi_{I'}(x_{2i+1}) = \pi_{I'}(x_{2i+2})$,

for all $i \in \mathbb{Z}_{\geq 0}$, where $I \in \{A, B\}$. The set of all admissible sequences on $\mathbf{S}_{\text{outcome}}$ is denoted by $\mathscr{S}_{\text{adm}}(\mathbf{S}_{\text{outcome}})$.

When convenient, we use the notation σ_B and σ^B to denote admissible sequences where player B is the first and last, respectively, to take an action. The notation σ_B^B then means that B is the first and last to take an action. Similar notations can be defined for A. Given an admissible sequence $\sigma = (x_0, x_1, \dots, x_k)$, $k \in \mathbb{Z}_{\geq 1}$, we say that $z \in \mathbf{S}_{\text{outcome}}$ is aligned with σ at time i if $z = x_i$. In this paper, without loss of generality, we assume that B is the first to take an action. In order to formalize the problem of deception we set out to study, we first introduce some basic notions.

A. Modeling player actions via probability distributions

Although player B has complete information about player A's game, she does not know the strategy that A follows to decide her actions. For instance, if multiple sanction-free improvements from an outcome are available to A, she might not necessarily pick her most preferred sequentially rational outcome (a less favorite improvement now may allow her to achieve a larger payoff in the future). Formally, this scenario can be captured by assigning a probability distribution to the edges of the H-digraph of A. Let $\mathbb{P}_{AB}(X_{n+1} = y \mid X_n = x)$, for $y \in \mathbf{S}_{\text{outcome}}|_{\pi_B(x)}$, denote the probability that the outcome of the hypergame changes from x to y by the action $\pi_A(y)$ of player A, as perceived by player B. Given what B knows about A's game, we have that for all $(x, y) \notin \mathcal{E}_{H_A^0}$,

$$\mathbb{P}_{AB}(X_{n+1} = y \mid X_n = x) = 0.$$

Note that, for all $x \in \mathbf{S}_{\text{outcome}}$,

$$\sum_{y \in \mathbf{S}_{\text{outcome}}|_{\pi_B(x)}} \mathbb{P}_{AB}(X_{n+1} = y \mid X_n = x) = 1.$$

The probability distribution $\mathbb{P}_{AB}$ is selected by player B by applying some rule (e.g., 'assign more probability to the most preferred outcome') to the H-digraph of the opponent. The results of the paper are independent of the specific rule used and so we leave it unspecified.

Player B can choose her own actions based on her preferences in any way she sees fit. We formally describe this via a probability distribution $\mathbb{P}_B$ on any action $\pi_B(y)$ which changes the outcome from x to y. Note that this can, in particular, be a vector with one entry of 1 and the rest 0, and that it can be re-selected at each round of the game. Modeling B's actions in this way will later be helpful in the statement of the results. Since players only use the current state of the game to decide about their next action, the sequence of repeated outcomes of the game is a Markov chain. This Markov chain can possibly be time-varying, since the H-digraph of player A can possibly evolve with time.

B. Notions of deception

Here, we introduce several definitions to capture different forms of deception. The first definition encodes a situation where the deceiver wishes to make sure that the mark will not take a certain action from a given outcome.

Definition 3.2 (Edge-deceivability): Suppose players A and B play sequentially a hypergame $H^2 = \{H_A^0, H_B^1\}$, with $H_{AB}^0 = H_A^0$. An edge $(x, y) \in \mathcal{E}_{H_A^0}$, $\pi_B(x) = \pi_B(y)$, is *deceivable* by B in H_A^0 from $x_0 \in \mathbf{S}_{\text{outcome}}$ if there exists an admissible sequence of outcomes $\sigma_B = (x_0, x_1, x_2, \dots, x_{2k+1})$, $k \in \mathbb{Z}_{\geq 0}$, where

(i) $(x_{2i-1}, x_{2i}) \in \mathbf{Sw}^A_{x_{2i-2},x_{2i-1}} \circ \cdots \circ \mathbf{Sw}^A_{x_0,x_1}(\mathcal{E}_{H_A^0})$ and
(ii) $T_{\mathbb{P}_{AB}}(x_{2i}, x_{2i-1}) > 0$,

for all $i \in \{1, \dots, k\}$, such that $(x, y) \notin \mathbf{Sw}^A_{x,\sigma_B}(\mathcal{E}_{H_A^0})$. We refer to σ_B a *deceiving sequence* and we use the term 'B deceives A' if the hypergame evolves according to σ_B. We denote by $E_{\text{dec}}^{B,x_0}(H_A^0) \subseteq \mathcal{E}_{H_A^0}$ the set of all deceivable edges by B in H_A^0 from x_0. We say that (x, y) is *surely deceivable* by B in H_A^0 from $x_0 \in \mathbf{S}_{\text{outcome}}$ if it is deceivable with probability one and we denote the set of all such edges by $E_{\text{sdec}}^{B,x_0}(H_A^0) \subseteq \mathcal{E}_{H_A^0}$.

Let us elaborate more on the properties of the deceiving sequence σ_B in the above definition. (i) states that A uses her updated H-digraph and takes an action to shift the outcome to a sanction-free improvement. (ii) states that B perceives a positive probability to the actions of A contained in σ_B. There is an abuse of notation due to the fact that $\mathbb{P}_{AB}$ can change with the evolution of the H-digraph. Also, here we have assumed that B takes the last action. This is without loss of generality; if the edge (x, y) is deceived by B, it remains deceived afterwards, unless B reveals new information.

Definition 3.3: (Strong edge-deceivability): The edge (x, y) is *strong deceivable* by B in H_A^0 if it is deceivable from any outcome $x_0 \in \mathbf{S}_{\text{outcome}}$ and is *surely strong deceivable* if it is strong deceivable with probability one. The set of strong deceivable and surely strong deceivable edges are denoted, respectively, by $E_{\text{stdec}}^B(H_A^0) \subseteq \mathcal{E}_{H_A^0}$ and $E_{\text{sstdec}}^B(H_A^0) \subseteq \mathcal{E}_{H_A^0}$.

Note that Definitions 3.2 and 3.3 are a stepping stone towards the deceiver being able to make sequentially rational an (in principle) unstable outcome for the mark. One can indeed similarly define a notion of outcome-deceivability: an outcome is deceivable if all the out-edges corresponding to the opponent's sanction-free improvements can be deceived. In this paper, we restrict our attention to the problem of edge-deceivability.

Lemma 3.4 (Deceivability inclusions): For all $x_0 \in \mathbf{S}_{\text{outcome}}$, the following inclusions hold

$$E_{\text{sstdec}}^B(H_A^0) \subseteq E_{\text{sdec}}^{B,x_0}(H_A^0),\ E_{\text{stdec}}^B(H_A^0) \subseteq E_{\text{dec}}^{B,x_0}(H_A^0).$$

We are now ready to formally state the problem we set out to study. Consider a 2-person 2-level hypergame $H^2 = \{H_A^0, H_B^1\}$, with $H_{AB}^0 = H_A^0$. We wish to provide answers to the following two problems:

(i) given $(x, y) \in \mathcal{E}_{H_A^0}$, with $\pi_B(x) = \pi_B(y)$, what are the set of outcomes $x_0 \in \mathbf{S}_{\text{outcome}}$ from which the edge is (surely) deceivable by B? When is the edge (surely) strong deceivable?
(ii) given an answer to the previous question, design an strategy that B can implement in order to deceive A.

IV. When is it possible to perform deception?

In this section, we identify a necessary condition and a sufficient condition for the notions of deceivability introduced in Section III-B. We also define a class of admissible sequences of outcomes, termed stealthy, and characterize conditions for deceivability that are both necessary and sufficient when the allowable sequences are restricted to this family.

A. Necessary conditions for deceivability and sufficient conditions for surely deceivability

We first identify a necessary condition for deceivability.

Lemma 4.1: (Necessary condition for edge-deceivability): Let $x_0 \in \mathbf{S}_{\text{outcome}}$ and assume $(x, y) \in E_{\text{dec}}^{B,x_0}(H_A^0)$. Then

$$\mathsf{H}_{\text{dec}}^A(x, y) = \{u \in \mathbf{S}_{\text{outcome}}|_{\pi_A(y)} \mid u \prec_{\mathrm{P}_{AA}} x\} \neq \emptyset.$$

Proof: We reason by contradiction. Suppose $u \succeq_{\mathrm{P}_{AA}} x$ for all $u \in \mathbf{S}_{\text{outcome}}|_{\pi_A(y)}$. Therefore, player B has no sanction against the improvement from x to y for player A, and thus the edge $(x, y) \in \mathbf{Sw}^A_{x,\sigma}(\mathcal{E}_{H_A^0})$, for any sequence of outcomes σ and any initial outcome $x_0 \in \mathbf{S}_{\text{outcome}}$. ∎

Note that Lemma 4.1 also gives a necessary condition for strong deceivability, c.f. Lemma 3.4. Next, we give a sufficient condition for surely deceivability.

Lemma 4.2: (Sufficient conditions for surely deceivability): Let $(x, y) \in \mathcal{E}_{H_A^0}$, $\pi_B(x) = \pi_B(y)$, and suppose $\mathsf{H}_{\text{dec}}^A(x, y) \neq \emptyset$. Then $(x, y) \in E_{\text{sdec}}^{B,\tilde{y}}(H_A^0)$, for all $\tilde{y} \in \mathsf{T}_{\text{dec}}^A(y) = \{w \in \mathbf{S}_{\text{outcome}}|_{\pi_A(y)} \mid w \succeq_{\mathrm{P}_{BA}} y\}$.

Proof: Note that $z \prec_{\mathrm{P}_{BA}} y$ for $z \in \mathsf{H}_{\text{dec}}^A(x, y)$, since otherwise, the improvement y from x of A would be sanctioned by the perceived improvement z from y of B and this would imply $(x, y) \notin \mathcal{E}_{H_A^0}$. Suppose B takes an action from $\tilde{y} \in \mathsf{T}_{\text{dec}}^y(y)$ that changes the outcome to $z \in \mathsf{H}_{\text{dec}}^A(x, y)$. Since $(\tilde{y}, z) \notin \mathcal{E}_{H_A^0}$, A uses the swap learning map to update her perceptions about B. But then $(x, y) \notin \mathbf{Sw}^A_{\tilde{y},z}(\mathcal{E}_{H^0_A})$, since the outcome z with $z \succ_{\mathbf{Sw}^A_{\tilde{y},z}(\mathrm{P}_{BA})} y$ is now perceived by A as a sanction of B against the improvement y from x by A. As a result, $\sigma_B = (\tilde{y}, z)$ is a deceiving sequence for B and thus the result follows. ■

B. Stealthy sequences of actions

If B takes an action not aligned with the perception of A, and A updates her perception (using for instance swap learning), then the structure of the H-digraph of A will change. Therefore, for B, the complexity of selecting a sequence of actions to deceive the opponent greatly grows with the length of the sequence. Here, instead, we focus on a particular family of sequences, which we term *stealthy*, that B can employ to achieve her goal without revealing any information to A, up to the moment that the 'deceiving action' takes place. Let us formally define this notion.

Definition 4.3: (Stealthy sequence): An admissible sequence of outcomes $\sigma_B = (x_0, x_1, \ldots, x_k)$, $k \in \mathbb{Z}_{\geq 1}$, is *stealthy* if the following holds:

- $(x_i, x_{i+1}) \in \mathcal{E}_{H^0_A}$, for all $i < k-1$;
- $(x_{k-1}, x_k) \notin \mathcal{E}_{H^0_A}$.

A consequence of the definition is that, if $\sigma_B = (x_0, x_1, \ldots, x_k)$, $k \in \mathbb{Z}_{\geq 1}$, is a stealthy sequence, then

$$\mathbf{Sw}^A_{x_{i-1},x_i}(\mathcal{E}_{H^0_A}) = \mathcal{E}_{H^0_A},$$

for all $i \in \{1, \ldots, k-1\}$, i.e., player A does not see her perception contradicted when the outcomes of the game correspond to σ_B. Moreover, at the last outcome,

$$\mathbf{Sw}^A_{x_{k-1},x_k}(\mathcal{E}_{H^0_A}) = \mathbf{Sw}^A_{x_0,\sigma_B}(\mathcal{E}_{H^0_A}) \neq \mathcal{E}_{H^0_A}. \quad (1)$$

Note that with this definition, the probability distribution $\mathbb{P}_{AB}$ does not change when the games is played according to a stealthy sequence. This definition motivates us to define the set $\mathscr{S}^{\mathbb{P}_{AB}}_{\mathrm{adm}}(\mathbf{S}_{\text{outcome}}) \subseteq \mathscr{S}_{\mathrm{adm}}(\mathbf{S}_{\text{outcome}})$ with

$$\mathscr{S}^{\mathbb{P}_{AB}}_{\mathrm{adm}}(\mathbf{S}_{\text{outcome}}) = \{(x_0, x_1, x_2, \ldots) \in \mathscr{S}_{\mathrm{adm}}(\mathbf{S}_{\text{outcome}}) \mid T_{\mathbb{P}_{AB}}(x_{i+1}, x_i) > 0, \forall i \in \mathbb{Z}_{\geq 0}, \pi_B(x_i) = \pi_B(x_{i+1})\}.$$

If $\sigma \in \mathscr{S}^{\mathbb{P}_{AB}}_{\mathrm{adm}}(\mathbf{S}_{\text{outcome}})$, we call σ a $\mathbb{P}_{AB}$*-admissible sequence*. With this definition, B perceives a positive probability to the actions of A contained in σ. From now on, whenever we use the term 'stealthy sequence' we mean to say '$\mathbb{P}_{AB}$-admissible stealthy sequence'.

The following result gives necessary and sufficient conditions for deceivability using stealthy sequence.

Theorem 4.4: (Necessary and sufficient conditions for deceivability via stealthy sequences): Let $x_0 \in \mathbf{S}_{\text{outcome}}$ and $(x, y) \in \mathcal{E}_{H^0_A}$, $\pi_B(x) = \pi_B(y)$. The following are equivalent:

(i) (x, y) is deceivable from x_0 via a stealthy sequence;

(ii) $\mathsf{H}^A_{\text{dec}}(x, y) \neq \emptyset$ and

$$\mathcal{T}^A_{\text{dec}}(y, x_0) = \mathsf{T}^A_{\text{dec}}(y) \cap \left(\{x_0\} \cup \mathcal{R}_{T_{\mathbb{P}_{AB}} T_{\mathbb{P}_B}}(x_0)\right) \neq \emptyset,$$

for a probability distribution $\mathbb{P}_B$ such that $\mathbb{P}_B(X_{n+1} = z \mid X_n = r) > 0$ for any $(r, z) \in \mathcal{E}_{H^0_A}$.

Proof: We first show that (i) implies (ii). Suppose $(x, y) \in \mathcal{E}_{H^0_A}$. First of all, note that since, by assumption, (x, y) is deceivable from x_0, the necessary conditions of Lemma 4.1 hold, i.e., $\mathsf{H}^A_{\text{dec}}(x, y) \neq \emptyset$. If $x_0 \in \mathsf{T}^A_{\text{dec}}(y)$, then $\mathcal{T}^A_{\text{dec}}(y, x_0) \neq \emptyset$ and the result follows. Suppose $x_0 \notin \mathsf{T}^A_{\text{dec}}(y)$. By Definitions 4.3 and 3.2, there exists a stealthy sequence $\sigma_B = (x_0, \ldots, x_{k-1}, x_k)$ such that $(x, y) \notin \mathbf{Sw}^A_{x_0,\sigma_B}(\mathcal{E}_{H^0_A})$ and $\mathfrak{S} = (x_0, \ldots, x_{k-1})$ is a path in $\mathcal{G}_{H^0_A}$. Note that $x_{k-1} \in \mathsf{T}^A_{\text{dec}}(y)$ and $x_k \in \mathsf{H}^A_{\text{dec}}(x, y)$, since otherwise, by definition of a swap learning map and Equation (1), $(x, y) \in \mathbf{Sw}^A_{x_0,\sigma_B}(\mathcal{E}_{H^0_A})$, a contradiction with the assumption. Next, let $\tilde{y} = x_{k-1}$ (observe that $k \geq 3$). We show that $\tilde{y} \in \mathcal{R}_{T_{\mathbb{P}_{AB}} T_{\mathbb{P}_B}}(x_0)$. Suppose $\mathbb{P}_B$ is a probability distribution such that $\mathbb{P}_B(X_{n+1} = z \mid X_n = r) > 0$, for all r and all $z \in \mathbf{S}_{\text{outcome}}|_{\pi_A(r)}$ with $(r, z) \in \mathcal{E}_{H^0_A}$. By definition of a $\mathbb{P}_{AB}$-admissible sequence, $T_{\mathbb{P}_{AB}}(x_i, x_{i+1}) > 0$, $\pi_B(x_i) = \pi_B(x_{i+1})$, for all $i \in \{1, \ldots, k\}$. Thus there is a strictly positive probability that $\tilde{y}$ is reachable from x_0 via the path $\mathfrak{S} = (x_0, \ldots, x_{k-2}, \tilde{y})$, $\pi_B(x_{k-2}) = \pi_B(\tilde{y})$. Thus there exists some $K \in \mathbb{Z}_{\geq 1}$ such that

$$(T_{\mathbb{P}_{AB}} T_{\mathbb{P}_B})^K (\tilde{y}, x_0) > 0,$$

i.e., $\tilde{y} \in \mathcal{R}_{T_{\mathbb{P}_{AB}} T_{\mathbb{P}_B}}(x_0)$. As a result, $\mathcal{T}^A_{\text{dec}}(y, x_0) \neq \emptyset$.

Conversely, let us show that (ii) implies (i). The results hold by Lemma 4.2 if $x_0 \in \mathsf{T}^A_{\text{dec}}(y)$. Suppose $x_0 \notin \mathsf{T}^A_{\text{dec}}(y)$. We need to show that there exists a sequence of outcomes σ_B that satisfies the conditions of Definition 4.3 and $(x, y) \notin \mathbf{Sw}^A_{x_0,\sigma_B}(\mathcal{E}_{H^0_A})$. First, note that, by Lemma 4.1, there exists an outcome $z \in \mathbf{S}_{\text{outcome}}|_{\pi_A(y)}$ such that $z \prec_{\mathrm{P}_{AA}} x$. By assumption, there exists an outcome $\tilde{y} \in \mathsf{T}^A_{\text{dec}}(y)$ that can be reached from x_0, for a probability distribution $\mathbb{P}_B$ described above, i.e., there exists a path $\mathfrak{S} = (x_0, \ldots, x_{k-1}, \tilde{y})$ in $\mathcal{G}_{H^0_A}$ such that $T_{\mathbb{P}_{AB}}(x_{i+1}, x_i) > 0$, for all $i \in \{0, \ldots, k-2\}$ with $\pi_B(x_i) = \pi_B(x_{i+1})$. If player B takes an action that changes the outcome from $\tilde{y}$ to z, then, by definition, z is a sanction against the perceived improvement y from x for player A; thus $(x, y) \notin \mathbf{Sw}^A_{\tilde{y},z}(\mathcal{E}_{H^0_A})$. Next, we define $\sigma_B = (x_0, \ldots, x_{k-1}, \tilde{y}, z)$. By Definition 4.3, σ_B is a stealthy sequence starting from x_0 and since $(x, y) \notin \mathbf{Sw}^A_{\tilde{y},z}(\mathcal{E}_{H^0_A})$, (x, y) is also a deceiving sequence, as claimed. ■

The choice of $\mathbb{P}_B$ in Theorem 4.4(ii) ensures that all actions of player B are considered when determining if a stealthy sequence exists to deceive A. Once such sequence is found, B will assign probability one to each of the actions for her prescribed in the sequence (cf. Section V).

Theorem 4.4 shows that, given $x_0 \in \mathbf{S}_{\text{outcome}}$, any action of B from $\mathcal{T}^A_{\text{dec}}(y, x_0)$ to $\mathsf{H}^A_{\text{dec}}(x, y)$ removes the edge (x, y) from the H-digraph $\mathcal{G}_{H^0_A}$. Thus, if these two sets are nonempty, finding a stealthy sequence is equivalent, by definition of $\mathcal{T}^A_{\text{dec}}(y, x_0)$, to finding a path in $\mathcal{G}_{H^0_A}$ that reaches $\mathcal{T}^A_{\text{dec}}(y, x_0)$ from x_0. One can characterize the set of all initial outcomes from which the edge (x, y) is deceivable as

$$\mathcal{I}^A_{\text{dec}}(x, y) = \{x_0 \in \mathbf{S}_{\text{outcome}} \mid \mathsf{H}^A_{\text{dec}}(x, y) \neq \emptyset, \mathcal{T}^A_{\text{dec}}(y, x_0) \neq \emptyset\}. \quad (2)$$

V. THE `WORST-CASE MAX-STRATEGY`

Here, we provide an algorithmic approach that can be used by player B to determine a stealthy sequence to deceive A.

Consider the scenario described in Section III. Suppose at time $t \geq 0$ the outcome of the 2-person 2-level hypergame is $\mathbf{x}(t)$. Without loss of generality, assume that player B takes actions when $t \in 2\mathbb{Z}_{\geq 0}$ and player A takes actions when $t \in 2\mathbb{Z}_{\geq 0}+1$. In this situation, Theorem 4.4 characterizes the edges of the H-digraph of A that are deceivable by B via a stealthy sequence. To model the fact that the outcome of the hypergame is influenced by the actions of player A, let us introduce the map $\Phi_{\mathbb{P}_{AB}} : \mathscr{S}_{\text{adm}}^{\mathbb{P}_{AB}}(\mathbf{S}_{\text{outcome}}) \to \mathbb{R}$,

$$\Phi_{\mathbb{P}_{AB}}(x_0,\dots,x_k) = \sum_{\substack{i=0\\ \pi_B(x_i)=\pi_B(x_{i+1})}}^{k-1} \ln\left(T_{\mathbb{P}_{AB}}(x_{i+1},x_i)\right). \tag{3}$$

This map captures the probability of reaching an outcome via a $\mathbb{P}_{AB}$-admissible sequence.

In this scenario, after making sure that the necessary condition for deception is satisfied, a reasonable strategy for B at each round is to take an action that maximizes the minimum probability of achieving the deception goal. Informally, this strategy can be described as follows:

> *[Informal description]:* Initially, B has a stealthy sequence (possibly empty) stored in its memory. At each round,
>
> (i) if there is a deceiving action that takes the current outcome to $\mathsf{H}_{\text{dec}}^A(x,y)$, player B takes it to deceive A, c.f. Lemma 4.2;
>
> (ii) otherwise, B checks if the A's last action is aligned with the stored sequence. If it is, B takes the next action prescribed by the sequence. If it is not, B considers the outcomes $w \in \mathbf{S}_{\text{outcome}}$ where she can take the game to by an action aligned with A's H-digraph, and computes the stealthy sequence with minimum probability of reaching an outcome in $\mathcal{T}_{\text{dec}}^A(y,x_0)$ from each w. B stores the sequence that maximizes these probabilities and takes the action prescribed by it.

We call this strategy the `worst-case max-strategy` and formally describe it in Algorithm 1. The rationale behind its name is made explicit in the next result.

Lemma 5.1: (The `worst-case max-strategy` maximizes the minimum probability of deception): The following are equivalent:

(i) $\sigma_B = (x_0,x_1,x_2,\dots,x_{2k}) \in \mathscr{S}_{\text{adm}}^{\mathbb{P}_{AB}}(\mathbf{S}_{\text{outcome}})$, where $k \in \mathbb{Z}_{\geq 1}$, $x_{2k} \in \mathsf{T}_{\text{dec}}^A(y)$, and $(x_i,x_{i+1}) \in \mathcal{E}_{H_A^0}$ for $i \in \{0,\dots,2k-1\}$, is a minimizer of $\Phi_{\mathbb{P}_{AB}}$;

(ii) σ_B corresponds to the longest path from x_0 to $x_{2k} \in \mathsf{T}_{\text{dec}}^A(y)$, in the digraph $(\mathbf{S}_{\text{outcome}}, \mathcal{E}_{H_A^0}, \mathcal{A}_{H_A^0})$, where, for $i,j \in \{1,\dots|\mathbf{S}_{\text{outcome}}|\}$,

$$(\mathcal{A}_{H_A^0})_{ij} = \begin{cases} |\ln\left(T_{\mathbb{P}_{AB}}(z_j,z_i)\right)|, & \pi_B(z_i)=\pi_B(z_j),\\ 0, & \text{otherwise.}\end{cases}$$

Algorithm 1: `worst-case max-strategy`

Input: $\mathcal{G}_{H_A^0}$, $\mathbb{P}_{AB}$, $(x,y) \in \mathcal{E}_{H_A^0}$, $x_0 \in \mathbf{S}_{\text{outcome}}$, $\mathcal{N}^{\text{out}}(x_0) \cap \mathbf{S}_{\text{outcome}}|_{\pi_A(x_0)} \neq \emptyset$
Initialization: $\alpha^{\text{maxmin}} = -\infty$, $\sigma_B = \emptyset$, $\mathbf{x}(0) = x_0$

check that $\mathsf{H}_{\text{dec}}^A(x,y) \neq \emptyset$; otherwise, announce that (x,y) is not deceivable
at time: $t \in 2\mathbb{Z}_{\geq 0}$
if $\mathbf{x}(t) \in \mathsf{T}_{\text{dec}}^A(y)$ **then**
    take action that makes $\mathbf{x}(t+1) \in \mathsf{H}_{\text{dec}}^A(x,y)$
else
    if $\sigma_B \neq \emptyset$ and $\mathbf{x}(t)$ is aligned with σ_B **then**
        take action prescribed by σ_B
    else
        foreach $w \in \mathbf{S}_{\text{outcome}}|_{\pi_A(\mathbf{x}(t))}$, $(\mathbf{x}(t),w) \in \mathcal{E}_{H_A^0}$ **do**
            $\alpha^{\text{min}} = +\infty$
            foreach $\tilde{y} \in \mathsf{T}_{\text{dec}}^A(y)$ **do**
                if there is path in $\mathcal{G}_{H_A^0}$ from w to $\tilde{y}$ **then**
                    find sequence σ_A^A from w to $\tilde{y}$ minimizing $\Phi_{\mathbb{P}_{AB}}$
                    if $\Phi_{\mathbb{P}_{AB}}(\sigma_A^A) \leq \alpha^{\text{min}}$ **then**
                        $\alpha^{\text{min}} = \Phi_{\mathbb{P}_{AB}}(\sigma_A^A)$
                    end
                end
            end
            if $\alpha^{\text{min}} \neq +\infty$ and $\alpha^{\text{min}} \geq \alpha^{\text{maxmin}}$ **then**
                $\alpha^{\text{maxmin}} = \alpha^{\text{min}}, \quad \eta = \sigma$
            end
        end
        if $\alpha^{\text{maxmin}} \neq -\infty$ **then**
            $\sigma_B = (\mathbf{x}(t), \eta)$
            take action prescribed by σ_B
        else
            (x,y) is not deceivable from $\mathbf{x}(t)$
        end
    end
end

Note that, in Lemma 5.1, (i) is equivalent to stating that σ_B is a minimizer of $\Pi_{i=1}^k T_{\mathbb{P}_{AB}}(x_{2i},x_{2i-1})$, and (ii) implies that finding solutions to the `worst-case max-strategy` is equivalent to finding a longest path on a digraph.

Remark 5.2: (Complexity of Algorithm 1): The digraph $(\mathbf{S}_{\text{outcome}}, \mathcal{E}_{H_A^0}, \mathcal{A}_{H_A^0})$ was recently shown in [23] to be acyclic, and therefore, the problem of finding a longest path is well-posed and can be solved efficiently. For example, since $\left|\mathbf{S}_{\text{outcome}}|_{\pi_A(\mathbf{x}(t))}\right|$ and $|\mathsf{T}_{\text{dec}}^A(y)|$ are bounded by $|\mathbf{S}_{\text{outcome}}|$, if one uses a version of Dijkstra's algorithm [24] to find the longest path, the computational complexity of the `worst-case max-strategy` in each round is in $O(|\mathbf{S}_{\text{outcome}}|^4)$. Furthermore, since, by acyclicity, the length of the longest stealthy path in the digraph is bounded by $|\mathbf{S}_{\text{outcome}}|$, the time complexity is in $O(|\mathbf{S}_{\text{outcome}}|^5)$. •

Next, we show that the `worst-case max-strategy` is complete, in the sense that it always finds a stealthy sequence that deceives a surely deceivable edge.

Theorem 5.3: (Surely deceivable edges via `worst-case max-strategy`*):* The edge $(x,y) \in \mathcal{E}_{H_A^0}$, $\pi_B(x) = \pi_B(y)$, is surely deceivable from $x_0 \in \mathbf{S}_{\text{outcome}}$ via a stealthy sequence of B if and only if $\mathsf{H}_{\text{dec}}^A(x,y) \neq \emptyset$ and either $x_0 \in \mathsf{T}_{\text{dec}}^A(y)$ or

$$\max_{x_1 \in \mathbf{S}_{\text{outcome}}|_{\pi_A(x_0)}} \min_{\sigma_B} \Phi_{\mathbb{P}_{AB}}(\sigma_B) = 0,$$

where $\sigma_B = (x_0, x_1, x_2, \ldots, x_{2k}) \in \mathscr{S}_{\text{adm}}^{\mathbb{P}_{AB}}(\mathbf{S}_{\text{outcome}})$, $k \in \mathbb{Z}_{\geq 1}$, $x_{2k} \in \mathsf{T}_{\text{dec}}^A(y)$, $(x_i, x_{i+1}) \in \mathcal{E}_{H_A^0}$ for all $i \in \{0, \ldots, 2k-1\}$.

Proof: Suppose $(x,y) \in \mathcal{E}_{H_A^0}$, $\pi_B(x) = \pi_B(y)$, is surely deceivable from x_0 via a stealthy sequence. By Lemma 4.1, $\mathsf{H}_{\text{dec}}^A(x,y) \neq \emptyset$. Suppose $x_0 \notin \mathsf{T}_{\text{dec}}^A(y)$. By Theorem 4.4 and by the definition of surely deceivability, there exists a sequence of outcomes $\sigma_B = (x_0, x_1, x_2, \ldots, x_{2k}) \in \mathscr{S}_{\text{adm}}^{\mathbb{P}_{AB}}(\mathbf{S}_{\text{outcome}})$, $k \in \mathbb{Z}_{\geq 1}$, $x_{2k} \in \mathsf{T}_{\text{dec}}^A(y)$, where

(i) $(x_i, x_{i+1}) \in \mathcal{E}_{H_A^0}$, for all $i \in \{0, \ldots, 2k-1\}$;

(ii) $T_{\mathbb{P}_{AB}}(x_{2i}, x_{2i-1}) = 1$, for all $i \in \{1, \ldots, k\}$.

By (3), $\Phi_{\mathbb{P}_{AB}}(\sigma_B) = 0$. Since $T_{\mathbb{P}_{AB}}(x_{2i}, x_{2i-1}) = 1$, for all $i \in \{1, \ldots, k\}$, if B chooses her sequential actions aligned with σ_B at each time, then the sequence will reach x_{2k} with probability one, and thus it is the unique sequence starting at x_0 reaching x_{2k} which includes x_1. This, along with the fact that $\Phi_{\mathbb{P}_{AB}}(\sigma_B) \leq 0$ for any σ_B, proves the result.

Conversely, if $x_0 \in \mathsf{T}_{\text{dec}}^A(y)$, since $\mathsf{H}_{\text{dec}}^A(x,y) \neq \emptyset$, the result follows from Lemma 4.2. Suppose $x_0 \notin \mathsf{T}_{\text{dec}}^A(y)$. Then, by assumptions, σ_B is a stealthy sequence from x_0 which reaches $x_{2k} \in \mathcal{T}_{\text{dec}}^A(y, x_0)$, with probability one. Since $\mathsf{H}_{\text{dec}}^A(x,y) \neq \emptyset$, the result follows by Theorem 4.4 and the definition of surely deceivability. ∎

The following results demonstrates that the `worst-case max-strategy` can also characterize the surely deceivable edges when the opponent is using a best-response strategy. The proof is similar to Theorem 5.3 and is omitted here.

Proposition 5.4: (Best-response strategies and the `worst-case max-strategy`*):* If A takes the sanction-free action associated to her most preferred outcome at all times and B knows about this, then $(x,y) \in E_{\text{sdec}}^{B,x_0}(H_A^0)$, $x_0 \in \mathbf{S}_{\text{outcome}}$, via a stealthy sequence of B if and only if $\mathsf{H}_{\text{dec}}^A(x,y) \neq \emptyset$ and either $x_0 \in \mathsf{T}_{\text{dec}}^A(y)$ or

$$\max_{x_1 \in \mathbf{S}_{\text{outcome}}|_{\pi_A(x_0)}} \min_{\sigma_B} \Phi_{\mathbb{P}_{AB}^*}(\sigma_B) = 0,$$

where $\mathbb{P}_{AB}^*$ assigns one to the edges of $\mathcal{G}_{H_A^0}$ associated to the most preferred sanction-free actions of A and $\sigma_B = (x_0, x_1, x_2, \ldots, x_{2k}) \in \mathscr{S}_{\text{adm}}^{\mathbb{P}_{AB}}(\mathbf{S}_{\text{outcome}})$, $k \in \mathbb{Z}_{\geq 1}$, $x_{2k} \in \mathsf{T}_{\text{dec}}^A(y)$, $(x_i, x_{i+1}) \in \mathcal{E}_{H_A^0}$ for all $i \in \{0, \ldots, 2k-1\}$.

Remark 5.5 (Strong deceivability): The execution of the `worst-case max-strategy` from all the outcomes in $\mathbf{S}_{\text{outcome}}$ fully characterizes the set $\mathcal{I}_{\text{dec}}^A(x,y)$. Note that, by definition, $\mathcal{I}_{\text{dec}}^A(x,y) = \mathbf{S}_{\text{outcome}}$ if and only if (x,y) is strongly deceivable via a stealthy sequence. •

VI. An Example

This section illustrates the results presented in the paper. Consider a 2-level $H^2 = \{H_A^0, H_B^1\}$ between A and B, with $H_{AB}^0 = H_A^0$ and outcome set $\mathbf{S}_{\text{outcome}} = S_A \times S_B = \{1, \ldots, 50\}$, where S_A and S_B are the action sets of A and B, respectively, and $|S_A| = 5$ and $|S_B| = 10$. The preference vectors P_{AA} and P_{BA} are shown in Figure 1. The H-digraph

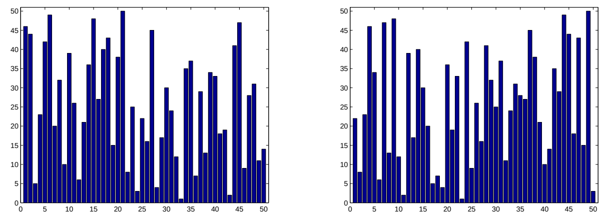

(a) P_{AA} (b) P_{BA}

Fig. 1. Preference vectors P_{AA} and P_{BA}. The horizontal axis shows the outcomes and the vertical axis shows the rank of outcomes.

$\mathcal{G}_{H_A^0}$ is shown in Figure 2(left). Regarding the actions of A, player B perceives that outcomes with lower rank in P_{AA} have higher probability of occurring. Formally, B assigns

$$T_{\mathbb{P}_{AB}}(j,i) = \frac{50 - \operatorname{rank}(j, \mathrm{P}_{AA})}{\sum_{l \in \mathcal{N}^{\text{out}}(i) \cap \mathbf{S}_{\text{outcome}}|_{\pi_B(i)}} (50 - \operatorname{rank}(l, \mathrm{P}_{AA}))},$$

to the event that the outcome changes from i to j by the action $\pi_A(j)$ of A, where $j \in \mathbf{S}_{\text{outcome}}|_{\pi_B(i)}$.

Suppose the game initially starts at outcome $x_0 = 14$ and B wishes to deceive A by removing the edge $(29, 26) \in \mathcal{E}_{H_A^0}$ via a stealthy sequence. Since

$$\begin{aligned}\mathcal{H}_{\text{dec}}^A(29,26) &= \{z \in \mathbf{S}_{\text{outcome}}|_{\pi_A(26)} \mid z \prec_{\mathrm{P}_{AA}} 29\} \\ &= \{11, 31, 41\} \neq \emptyset,\end{aligned}$$

the necessary condition of Lemma 4.1 is satisfied. According to Theorem 4.4, we compute

$$\begin{aligned}\mathsf{T}_{\text{dec}}^A(26) &= \{w \in \mathbf{S}_{\text{outcome}}|_{\pi_A(26)} \mid w \succeq_{\mathrm{P}_{BA}} 26\} \\ &= \{1, 6, 26, 36\}.\end{aligned}$$

The actions of B from 14 aligned with A's H-digraph are $\mathcal{N}^{\text{out}}(14) \cap \mathbf{S}_{\text{outcome}}|_{\pi_A(14)} = \{9, 24, 39\}$. By executing the `worst-case max-strategy`, B finds that the action that maximizes the minimum probability of reaching any of the outcomes in $\mathsf{T}_{\text{dec}}^A(26)$ is $\pi_B(24)$, where she perceives that the repeated play of the game will reach outcome 36 via the path $\mathfrak{S} = (14, 24, 25, 40, 36)$, with probability 0.52. Note that, by definition, $36 \in \mathcal{T}_{\text{dec}}^A(26, 14)$. If the repeated play of the hypergame goes according to B's perception, after reaching the outcome 36, B takes an action that changes the outcome from 36 to any of the outcomes in $\mathcal{H}_{\text{dec}}^A(29, 26)$. For example, if B chooses to take the action $\pi_B(11)$ (note that $(36, 11) \notin \mathcal{E}_{H_A^0}$), then A's H-digraph after updating her perception via swap learning is shown in Figure 2(right). If A takes an action which is not aligned with the sequence $\mathfrak{S}$ at any round of the hypergame, according to the `worst-case max-strategy`, B will recompute the stealthy sequence and take the ensuing action accordingly.

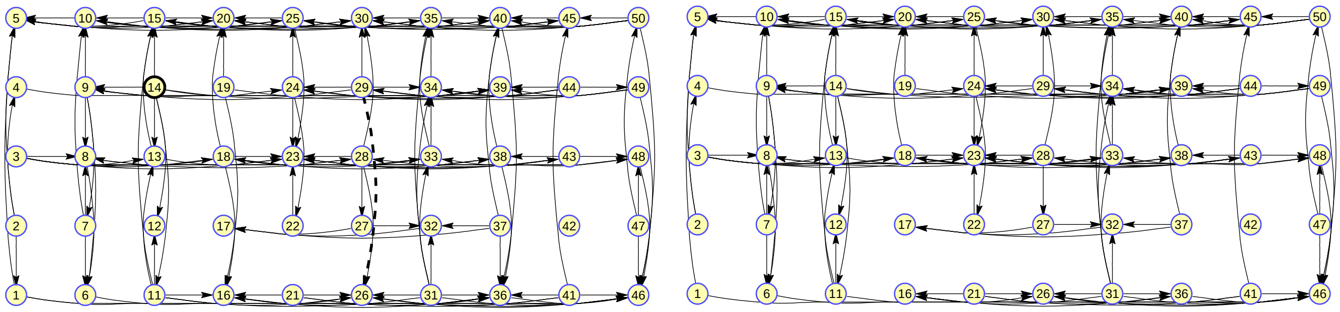

Fig. 2. H-digraphs $\mathcal{G}_{H_A^0}$ (left) and $\mathbf{Sw}^A_{36,11}(\mathcal{G}_{H_A^0})$ (right). Player A plays rows, player B plays columns, and P_{AA} and P_{BA} are given in Figure 1. Player B intends to remove the edge $(29, 26) \in \mathcal{E}_{H_A^0}$ (left plot, dashed) via a stealthy sequence, starting from outcome 14. After reaching the outcome 36, the edge $(29, 26) \in \mathcal{E}_{H_A^0}$ is removed (right plot) by the action $\pi_B(11)$ of player B.

VII. Conclusions

We have studied scenarios of active deception in 2-person 2-level hypergames with asymmetric information. Using the properties of hypergames encoded in the notion of H-digraph, we have introduced formal notions that capture different forms of deception. We have provided a necessary condition and a sufficient condition for deceivability for the case when the deceiver might take actions that contradict the perception of her opponent about the game. When this is not the case, i.e., if the deceiver acts in a stealthy way and only takes actions aligned with her opponent's perception, we have fully characterized when deception is possible. Finally, we have introduced the `worst-case max-strategy` which maximizes the minimum probability that the deceiver achieves the deception goal. We have shown this algorithm to be complete, in the sense that it always finds a stealthy sequence that deceives a surely deceivable edge. Future work will study efficient ways of performing outcome deceivability, the impact of signaling cost on the deceiver's available strategies, and the challenging scenario of deception via non-stealthy strategies, where the H-digraph of the opponent might change as a result of the actions taken by the deceiver.

References

[1] R. J. Lewicki, D. M. Saunders, and B. Barry, *Negotiation*. New York: McGraw-Hill, 6nd ed., 2010.

[2] D. Teneketzis and D. A. Castañón, "Informational aspects of class of subjective games of incomplete information: static case," *Journal of Optimization Theory & Applications*, vol. 54, no. 2, pp. 413–422, 1987.

[3] V. P. Crawford, "Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions," *American Economic Review*, vol. 93, pp. 133–149, 2003.

[4] S. Mullainathan, J. Schwartzstein, and A. Shleifer, "Coarse thinking and persuasion," *Quarterly Journal of Economics*, vol. 123, no. 5, pp. 577–619, 2008.

[5] D. Ettinger and P. Jehiel, "A theory of deception," *American Economic Journal: Microeconomics*, vol. 2, pp. 1–20, 2010.

[6] V. J. Baston and F. A. Bostock, "Deception games," *International Journal of Game Theory*, vol. 17, pp. 129–134, 1988.

[7] K. T. Lee and K. L. Teo, "A game with distorted information," *Naval Research Logistics*, vol. 40, pp. 993–1001, 1993.

[8] N. M. Fraser and K. W. Hipel, *Conflict Analysis: Models and Resolutions*. Systems Science and Engineering, Oxford, UK: Elsevier, 1984.

[9] M. Wang, K. W. Hipel, and N. M. Fraser, "Solution concepts in hypergames," *Applied Mathematics and Computation*, vol. 34, pp. 147–171, 1989.

[10] R. Vane and P. Lehner, "Using hypergames to increase planned payoff and reduce risk," *Autonomous Agents and Multi-Agent Systems*, vol. 5, pp. 365–380, 2002.

[11] P. Bernhard, A. L. Colomb, and G. Papavassilopoulos, "Rabbit and hunter game: Two discrete stochastic formulations," *International Journal of Applied Mathematics and Computer Science*, vol. 13, pp. 205–225, 1987.

[12] G. J. Olsder and G. P. Papavassilopoulos, "About when to use a searchlight," *Journal of Mathematical Analysis and Applications*, vol. 136, pp. 466–487, 1988.

[13] D. A. Castañón, P. R. Chandler, and M. Pachter, "A game of deception," in *IEEE Conf. on Decision and Control*, (Paradise Island, Bahamas), pp. 3364–3369, 2004.

[14] J. P. Hespanha, "Application and value of deception," in *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind* (A. Kott and W. M. McEneaney, eds.), (Boca Raton, FL), pp. 145–165, Chapman and Hall/CRC, 2007.

[15] D. Li and J. B. C. Jr., "Information, decision-making and deception in games," *Decision Support Systems*, vol. 47, pp. 518–527, 2009.

[16] W. M. McEneaney, "Some classes of imperfect information finite state space stochastic games with finite-dimensional solutions," *Applied Mathematics and Optimization*, vol. 50, no. 2, pp. 87–118, 2004.

[17] W. M. McEneaney and R. Singh, "Robustness against deception," in *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind* (A. Kott and W. McEneaney, eds.), (Boca Raton, FL), pp. 167–207, Chapman and Hall/CRC, 2007.

[18] N. Biggs, *Algebraic Graph Theory*. Cambridge University Press, 2 ed., 1994.

[19] S. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*. Cambridge, UK: Cambridge University Press, 2nd ed., 2009.

[20] B. Gharesifard and J. Cortés, "Evolution of players' misperceptions in hypergames under perfect observations," *IEEE Transactions on Automatic Control*, 2010. Submitted.

[21] D. M. Kreps and R. Wilson, "Sequential equilibria," *Econometrica*, vol. 50, no. 4, pp. 863–894, 2004.

[22] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. SIAM, 2 ed., 1999.

[23] B. Gharesifard and J. Cortés, "Learning of equilibria and misperceptions in hypergames with perfect observations," in *American Control Conference*, (San Francisco, California), 2010. Accepted.

[24] G. Tel, *Introduction to Distributed Algorithms*. Cambridge University Press, 2 ed., 2001.