IEEE Control Systems Letters paper presented at
2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Deception by Motion: The Eater and the Mover Game

Violetta Rostobaya[1], Yue Guan[2], James Berneburg[1], Michael Dorothy[3], and Daigo Shishika[1]

***Abstract*— This paper studies the idea of "deception by motion" through a two-player dynamic game played between a Mover who must reach a goal to retrieve resources, and an Eater who can consume resources from two candidate goals. The Mover seeks to minimize the resource consumption at the true goal it must reach, while the Eater tries to maximize it without knowing which one the true goal is. Unlike existing works on deceptive motion control that measures the deceptiveness through the quality of inference made by a distant observer (an estimator), we incorporate agents' actions to directly measure the efficacy of deception through the outcome of the game. An equilibrium concept is then proposed without the notion of an estimator. We further identify a pair of equilibrium strategies and demonstrate that if the Eater optimizes for the worst-case scenario, hiding the intention (deception by ambiguity) is still effective, whereas trying to fake the true goal (deception by exaggeration) is not.**

I. Introduction

In competitive games with asymmetric information, players can sometimes leverage deception to alter the decisions made by the opponent and achieve a higher payoff [1]. Common forms of deception include sensor jamming [2], controlling shared information [3], etc. In this work we draw attention to deception via direct perception, where a player does not have a communication channel, but instead, tries to deceive its opponent by moving in a particular way.

A number of existing works have considered deception in the context of motion control. A typical formulation optimizes the path of a moving agent to reach its goal while minimizing the quality of the inference an observer is trying to make about the location of the agent's goal [4]–[7]. Specifically, [5] and [6] formalized the notion of *ambiguity* (hiding information about true goal), and *exaggeration* (moving towards a decoy goal to send a false signal) as two ways to measure deceptiveness.

A common assumption made in these works is that the observing agent uses a *prescribed* estimator/inference policy, and the deceiving agent leverages the knowledge of its structure. The deceived agent is also often so naive that it is not aware of the possibility of being deceived. Furthermore, since no decision is made by the observing agent on its action or inference policy, such formulation normally boils down to a one-sided optimization. The works on goal recognition by a passive observer also fall into this category [8]–[14].

We gratefully acknowledge the support of ARL grant ARL DCIST CRA W911NF-17-2-0181. The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Army, Department of Defense, or the United States Government.

[1]Violetta Rostobaya, James Berneburg and Daigo Shishika are with College of Engineering and Computing at George Mason University, Emails: {vrostoba,jbernebu,dshishik}@gmu.edu

[2]Yue Guan is with the School of Aerospace Engineering, Georgia Institute of Technology, Email: yguan44@gatech.edu

[3]Michael Dorothy is with DEVCOM Army Research Laboratory, Email: michael.r.dorothy.civ@army.mil

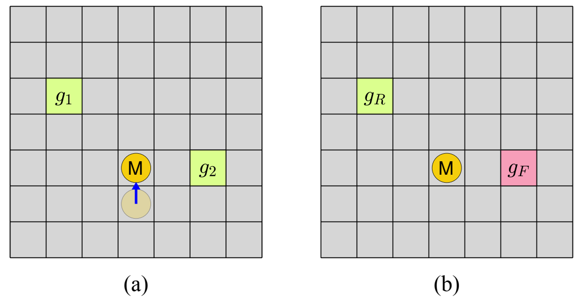

Fig. 1. Illustration of the asymmetric information. (a) Perspective of the Eater, who knows the most recent Mover's action (blue arrow), (b) Perspective of the Mover who knows the true goal g_R.

In this paper, we are interested in studying the possibility of deception without making the assumptions discussed above. We do not prescribe an estimator for the observing player, but instead allow it to select its own policy, which makes the problem a two-player game. Since the observing agent must choose its action based on the observed motion of the deceiver, the success of deception can be directly measured through the influence on decisions of the observer, and ultimately by the outcome of the game. While existing works use observer's belief in the objective function (making deception itself to be the goal [4], [5], [7], [9]), our formulation views deception as a tool to accomplish underlying mission objectives (i.e., improve the game outcome).

The contributions of this work are: (i) the formulation of a novel game that explores the effectiveness of deception against an observing agent that actively takes actions; (ii) the identification of equilibrium strategies, where the observer does not utilize an estimator to predict the opponent's behaviors; and (iii) analytic characterizations of the associated game outcome. Our results indicate that deception by exaggeration does *not* work in our problem when the payoff function for the observer captures the worst-case scenario. Nevertheless, the deceiving agent can still use ambiguity to improve its payoff. This result is the first step towards investigating the existence of *deception by motion* when the observing agent has the ability to select its own policy.

II. Problem Formulation

We consider a two-player discrete-time dynamic game with asymmetric information played in a grid world between the Mover who controls its position and the Eater who controls the amount of resources at the goal locations. At the beginning of the game, two goal locations are specified

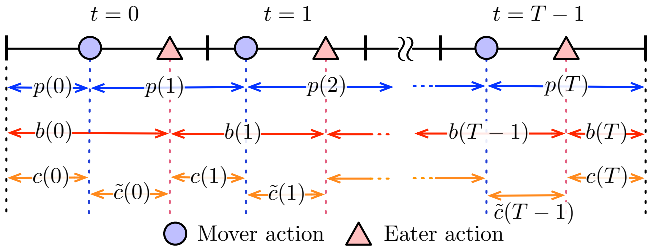

Fig. 2. Illustration of the game timeline.

by nature: $g_1, g_2 \in \mathbb{Z}^2$. One is the true goal, g_R, and the other one is the fake goal, g_F. Each goal is initialized with sufficiently large number of resources (bananas). Although in this formulation the Eater decreases the amount of resources, the same analysis can be used when the resource increases. Various strategic resource allocation scenarios fall under this framework, where the observing player must stock up or evacuate resources at an infrastructure in the face of an incoming adversary.

States: The game evolves with two types of states. One is the Mover's position $p(t) \in P$ where $P = \mathbb{Z}^2$. The positions could be equivalently denoted as nodes on a graph that represents the grid world. Although we restrict our analysis to the grid world for simplicity and intuitive understanding, we believe that the obtained results can be generalized to more generic graphs, which could capture, e.g., Mover with different dynamics and environments with obstacles. The other state is the number of bananas at each goal location, which the Eater consumes. We use $b_i(t) \geq 0$ to denote the consumption of bananas from g_i for $i \in \{1,2\}$, and define $\mathbf{b}(t) = [b_1(t), b_2(t)] \in B$ as the *consumption vector*,[1] where $B = \mathbb{R}^2_{\geq 0}$ and $\mathbf{b}(0) \in \mathbb{Z}^2_{\geq 0}$.

Actions: The Mover selects an action $a^M(t)$ from its action space $A^M = \{$`up`, `down`, `left`, `right`$\}$ at each time step, which updates its position from $p(t)$ to $p(t+1)$. The Eater's action $\mathbf{a}^E(t)$ is drawn from $A^E = \{[1,0],[0,1],[0.5,0.5]\}$, corresponding to eating one banana from g_1, eating one from g_2, or eating half from both goals.[2] The consumption vector has the following dynamics:

$$\mathbf{b}(t+1) = \mathbf{b}(t) + \mathbf{a}^E(t).$$

Terminal condition: The game terminates at time T, when the Mover reaches the true goal g_R:

$$T = \min_{t \in \mathbb{Z}_{\geq 0}} \{t | p(t) = g_R\},$$

which solely depends on the strategy used by the Mover. Note that the true goal is specified by the nature and does not change throughout the game. The Eater makes its last action at $t = T-1$ and the banana consumption is updated for one last time (see Fig. 2). The Mover receives the remaining bananas at the true goal.

[1] The difference between the numbers of bananas at the two goals at $t = 0$ can be properly reflected by setting $b_i(0) > 0$ for one of the goals.

[2] We assume that the bananas at both goals do not run out, but removing this assumption does not significantly affect our analysis and the policies since the Eater's decision when one goal has no banana is trivial.

Information structure: The set of goal locations $G = \{g_1, g_2\}$ is common knowledge. We consider sequential actions as illustrated in Fig. 2. When choosing an action at t, the Mover has access to the states $p(t)$ and $\mathbf{b}(t)$. The Mover also has private information set $I = \{1,2\}$ which indicates the true goal. The Eater has access to the updated states $p(t+1)$ and $\mathbf{b}(t)$, and the most recent Mover action $a^M(t)$ (or equivalently, one-step memory $p(t)$).

Strategy sets: We use $\mathscr{G}_i$, $i \in \{1,2\}$, to denote the game where $g_i = g_R$. The strategy of the Mover for $\mathscr{G}_i$ is denoted as: $\pi_i^M : P \times B \to A_M$, where $i \in \{1,2\}$. The overall Mover's strategy uses I (the knowledge of true goal) as follows:

$$\pi^M = \begin{cases} \pi_1^M & \text{if } g_R = g_1, \\ \pi_2^M & \text{if } g_R = g_2. \end{cases} \tag{1}$$

On the other hand, the Eater's strategy is a mapping $\pi^E : P \times A_M \times B \to A_E$. Note that the Eater deploys the same strategy for both games, because it cannot differentiate the true goal from the fake one.

Objective functions: We define the outcome of game $\mathscr{G}_i$ induced by strategy pair (π^M, π^E) as

$$J_i(p, \mathbf{b}; \pi_i^M, \pi^E) = b_i(T(\pi_i^M)), \tag{2}$$

which denotes the banana consumed at the true goal when the game terminates.The implicit dependency of b_i on π^E is omitted. Since the Mover knows which game it is playing, we let it directly minimize the consumption at the true goal, and hence its objective function is

$$J_i^M(p, \mathbf{b}; \pi_i^M, \pi^E) = J_i(p, \mathbf{b}; \pi_i^M, \pi^E). \tag{3}$$

Although the Eater wants to maximize J_i as well, it cannot tell which goal is real, or which game is being played. Consequently, there are a variety of candidate metrics for the Eater to optimize, e.g., average, best-case, or worst-case performance. In this paper, we set the worst-case performance as the Eater's objective function and let the Eater maximize the following:

$$J^E(p, \mathbf{b}; \pi_i^M, \pi^E) = \min_{i \in \{1,2\}} J_i(p, \mathbf{b}; \pi_i^M, \pi^E). \tag{4}$$

This Eater's objective provides the worst-case guarantee among the two possible games: $\mathscr{G}_1$ and $\mathscr{G}_2$. We will later see that this objective prevents the Eater from being deceived by exaggeration (Lem. 3). In the sequel, we will omit the dependence of the objective functions on the initial conditions p and $\mathbf{b}$.

We consider the equilibrium concept defined as follows.

Definition 1. *A pair of strategies (π^{M*}, π^{E*}) constitutes an equilibrium, if for all $\pi^E \in \Pi^E$ and $\pi^M \in \Pi^M$, it satisfies:*

$$J_1^M(\pi_1^{M*}, \pi^{E*}) \leq J_1^M(\pi_1^M, \pi^{E*}) \tag{5a}$$
$$J_2^M(\pi_2^{M*}, \pi^{E*}) \leq J_2^M(\pi_2^M, \pi^{E*}) \tag{5b}$$
$$J^E(\pi^{M*}, \pi^{E*}) \geq J^E(\pi^{M*}, \pi^E) \tag{5c}$$

where Π^E and Π^M are the sets of admissible strategies for the Eater and the Mover respectively.[3]

[3] In general, the choice of J^E affects the equilibrium. For example, best-case performance $J^E = \max_i J_i$ leads to a trivial Eater behavior that simply takes from the goal which gives higher J_i^M in the full information game.

In the following we propose a pair of strategies and prove that they constitute an equilibrium.

III. Main Results

A. Preliminary Analysis

We classify the Mover's action based on the change in the distances it induces. Let $d(\cdot,\cdot) : P \times P \to \mathbb{Z}_{\geq 0}$ denote the 1-norm (Manhattan distance) between two points on the grid. For the case of a general graph, this metric would be the length of the shortest path between two nodes. We use $d_i(t) \triangleq d(p(t), g_i)$, $i \in \{1,2\}$ to denote the distance from $p(t)$ to g_i. We also denote the change in the distance as

$$\delta d_i(t) \triangleq d_i(t+1) - d_i(t). \tag{6}$$

Based on $\delta d_1(t)$ and $\delta d_2(t)$, we categorize the Mover's action at time t into two classes: *ambiguous* and *explicit*.

Definition 2 (Ambiguous Move)**.** *An action of the Mover $a^M(t)$ is an ambiguous move if $\delta d_1(t) = \delta d_2(t)$.*

Definition 3 (Explicit Move)**.** *An action of the Mover $a^M(t)$ is an explicit move if $\delta d_1(t) \neq \delta d_2(t)$.*

Note that for the dynamics studied in this paper, we only have explicit moves with $\delta d_1(t) = -\delta d_2(t)$. However, the above definition accommodates a more general case (e.g., with diagonal moves) where it is possible to have, e.g., $\delta d_1(t) = 0$ and $\delta d_2 = -1$.

Definition 4 (Number of Steps)**.** *For any given $p(t)$, the minimum numbers of ambiguous and explicit moves required to reach g_i are denoted as $n^a(t)$ and $n_i^e(t)$.*

Notice that the number of ambiguous moves for both games is $n^a(t)$, and hence $n^a(t)$ does not have a subscript. Also note that $T \geq n^a(0) + n_R^e(0)$, where the equality holds when the Mover follows a shortest path to g_R. Although shortest paths are non-unique, they all contain the same number of ambiguous and explicit moves. Finally, the numbers $n^a(t)$ and $n_i^e(t)$ are independent in the sense that a single Mover action changes only one of them.

We define a quantity that describes the maximum possible banana consumption when the Mover uses a shortest path.

Definition 5. *Let $c_i(t)$ be defined as*

$$c_i(t) \triangleq b_i(t) + d_i(t), \tag{7}$$

and $\tilde{c}_i(t)$ be the one measured between the Mover's and Eater's action at time step t (also see Fig. 2):

$$\tilde{c}_i(t) \triangleq b_i(t) + d_i(t+1). \tag{8}$$

If the Eater had the knowledge of the true goal, it would always eat from that location. Hence one can interpret $c_i(t)$ as the value of the *complete information* game, in which the Eater knows g_R and always eats from the true goal.

Definition 6. *Let $\Delta c_i(t)$ and $\Delta \tilde{c}_i(t)$ denote the difference functions for $c_i(t)$ and $\tilde{c}_i(t)$ respectively:*

$$\Delta c_i(t) \triangleq c_i(t) - c_{-i}(t) \tag{9}$$

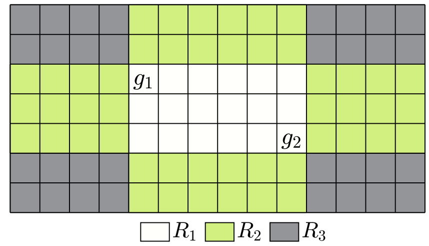

Fig. 3. Partition of the game environment.

$$\Delta \tilde{c}_i(t) \triangleq \tilde{c}_i(t) - \tilde{c}_{-i}(t), \tag{10}$$

where $-i = \{1,2\} \setminus \{i\}$.

We classify the environment into three regions (See Fig. 3). R_1 is the *no-ambiguity* region. Any move from R_1 to R_1 is an explicit move. R_2 is the *partial-ambiguity* region. Among the two subsets {`right, left`} and {`up, down`}, only one will be ambiguous in R_2. R_3 is the *full-ambiguity* region, where any action the Mover takes is ambiguous. Note that an action that transfers the Mover from one region to another is *always* ambiguous.

Remark 1. *For $p(0) \notin R_1$, there always exists a shortest path to g_i that uses all the necessary ambiguous moves first, which takes the Mover to the boundary of R_1.*

B. Equilibrium Strategies and Outcome

We propose the following strategy pair.

The Mover's strategy π_i^{M*}***:*** Move on a path with the following two properties.

- The path is one of the shortest paths to g_i.
- All $n^a(0)$ ambiguous moves are made before $n_i^e(0)$ explicit moves.

The Eater's strategy π^{E*}***:*** Observe $a^M(t)$ and identify if it was ambiguous or explicit based on Def. 2 and 3.

- If $a^M(t)$ was ambiguous, then use the *conservative* action: take from the goal with higher risk (smaller $\tilde{c}$)

$$\mathbf{a}^E(t) = \begin{cases} [1,0], & \text{if } \Delta\tilde{c}_1(t) < 0 \\ [0,1], & \text{if } \Delta\tilde{c}_1(t) > 0 \\ [0.5, 0.5], & \text{if } \Delta\tilde{c}_1(t) = 0. \end{cases} \tag{11}$$

- If $a^M(t)$ was explicit then use an *exploiting* action: take from the goal that the Mover approached

$$\mathbf{a}^E(t) = \begin{cases} [1,0], & \text{if } \delta d_1(t) < \delta d_2(t), \\ [0,1], & \text{if } \delta d_1(t) > \delta d_2(t). \end{cases} \tag{12}$$

Notice that π_i^{M*} does not explicitly use $\mathbf{b}(t)$, but π^{E*} uses it to compute $\Delta\tilde{c}_1(t)$.

Theorem 1 (Equilibrium)**.** *The strategies π^{M*} and π^{E*} form an equilibrium, i.e., they satisfy equations* (5a)*,* (5b) *and* (5c)*.*

Proof. We prove this main result in three steps in Sec. IV. Step I: We prove equilibrium for $p(0) \in R_1$ (Lem. 1). Step II: We prove the optimality of π^{M*} outside of R_1. Specifically, we show that the Mover should use a shortest path (Lem. 2 and 3), and it has no incentive to use explicit moves before ambiguous ones (Lem. 4). Step III: Optimality of π^{E*} outside of R_1 is proved (Lem. 5). $\square$

Before providing the equilibrium outcome, we present observations that facilitate the analysis throughout the paper.

Remark 2. *If $a^M(t)$ is ambiguous, then*

$$\Delta\tilde{c}_i(t) = \Delta c_i(t). \tag{13}$$

Furthermore, if the Eater uses π^{E*}, then $|\Delta c_i(t)|$ approaches zero by one for every ambiguous move: i.e.,

$$\Delta c_i(t+1) - \Delta c_i(t) = \begin{cases} 1 & \Delta c_i(t) < 0 \\ -1 & \Delta c_i(t) > 0 \\ 0 & \Delta c_i(t) = 0. \end{cases} \tag{14}$$

Remark 3. *If $a^M(t)$ is an explicit move toward g_i, then*

$$\Delta\tilde{c}_i(t) = \Delta c_i(t) - 2. \tag{15}$$

Furthermore, if the Eater uses π^{E*}, then $\Delta c_i(t)$ decreases by one: i.e.,

$$\Delta c_i(t+1) = \Delta c_i(t) - 1 \tag{16}$$

Theorem 2. *The equilibrium outcome under (π^{M*}, π^{E*}) for game $\mathscr{G}_i$ is:* [4]

$$V_i^{M*}(p, \mathbf{b}) = \begin{cases} c_i - 0.5n^a(1+\mathrm{sgn}(\Delta c_i)) & n^a \le |\Delta c_i| \quad (17) \\ c_i - 0.5(n^a + \Delta c_i) & n^a > |\Delta c_i| \quad (18) \end{cases}$$

The Eater's performance at the equilibrium is then given by $V^{E*}(p, \mathbf{b}) = \min_{i\in\{1,2\}} V_i^{M*}(p, \mathbf{b})$.

It is worth noting that the terms $0.5n^a(1+\mathrm{sgn}(\Delta c_i))$ and $0.5(n^a + \Delta c_i)$ in (17) and (18) are non-negative, and they represent the reduction in consumption when compared to the complete-information scenario, $c_i(0)$. In this sense, we can interpret these quantities as the *value of information*. Note that this value of information is zero if $n^a \le |\Delta c_i|$ and $\Delta c_i < 0$, which is the case when the real goal has higher risk and there is not enough ambiguous moves to achieve $\Delta c_i(t) = 0$.

Proof. The proof is omitted due to page limit. One can easily verify the claim by checking the two cases: $n^a(0) \le |\Delta c_i(0)|$ and $n^a(0) > |\Delta c_i(0)|$. $\square$

IV. Proof of Equilibrium

A. Equilibrium in Region R_1

Lemma 1. *The strategy pair (π^{M*}, π^{E*}) forms an equilibrium in R_1, i.e., equations* (5a)*,* (5b) *and* (5c) *hold in R_1.*

Proof. Assuming that $p(t) \in R_1$, we start by analyzing Eater's deviation in R_1 (condition (5c)). From Thm. 2, Eater's outcome in R_1 is given by $J^E(\pi^{M*}, \pi^{E*}) = \min_{i\in\{1,2\}} c_i(0)$. Consider a different Eater's strategy $\pi^{E'}$, which takes $x \in \{0.5, 1\}$ banana from g_{-i} at least once, even if the Mover is approaching g_i. The Mover's outcome under such strategy is $J_i^M(\pi_i^{M*}, \pi^{E'}) \le c_i(t) - x$, and therefore, we have $J^E(\pi^{M*}, \pi^{E'}) \le \min_{i\in\{1,2\}}(c_i(t) - x) = V^{E*} - x$, which implies that the Eater has no incentive to deviate.

Next, we show that the Mover has no incentive to deviate from π_i^{M*} in R_1. From Thm. 2, the Mover's outcome under (π_i^{M*}, π^{E*}) is given by $J_i^M(\pi_i^{M*}, \pi^{E*}) = c_i(t) = b_i(t) + n_i^e(t)$. Consider a different strategy $\pi^{M'}$, which increases the path length by adding additional explicit or ambiguous moves.[5] In either case, the total number of explicit moves, $n_i^e(0)$, can only increase but never decrease. Therefore, the banana consumption increases, i.e., $J_i^M(\pi_i^{M'}, \pi^{E*}) \ge b_i(0) + n_i^e(0) = V_i^{M*}$, and thus the Mover has no incentive to deviate. $\square$

B. Mover's Strategy

This section proves that the Mover has no incentive to deviate from π^{M*} when $p(t) \notin R_1$, if the Eater uses π^{E*}.

Lemma 2. *The Mover has no incentive to deviate from shortest path by moving away from both goals if the Eater uses π^{E*}.*

Proof. Consider $\mathscr{G}_i$ and $p(0) \notin R_1$. Under π^{M*}, the Mover makes $n^a(0) > 0$ ambiguous moves first and $n_i^e(0) \ge 0$ later. Consider $\pi^{M'} \ne \pi^{M*}$, which makes at least one ambiguous move away from both goals. Clearly, this strategy cannot reduce the minimum number of explicit moves, $n_i^e(0)$. The question then becomes: by making additional moves away from both goals, can the Mover decrease the resource consumption due to ambiguous moves?

We compare the Mover's trajectory with $n^a(0)$ ambiguous moves and the one with $n^a(0) + 2m$ steps, where $m \in \mathbb{N}$ is the number of ambiguous steps *away* from g_R. For simplicity, consider the case where $\pi^{M'}$ makes the $n^a(0)$ ambiguous moves first. Rem. 2 implies that in the first $n^a(0)$ steps, $\Delta c_i(t)$ evolves in the same way for the two trajectories. Therefore, the consumption in this phase is the same for both trajectories. The trajectory with $2m$ additional steps has either: (i) no additional consumption if $\Delta c_i(t) > 0$ for $t \le n^a(0) + 2m$, (ii) $2m$ additional consumption if $\Delta c_i(t) < 0$ for $t \le n^a(0) + 2m$, or (iii) somewhere between the above two if $\Delta c_i(t) = 0$ is achieved at some point. In all these three cases, the consumption at the true goal g_i does not decrease. This analysis extends to the case where $\pi^{M'}$ does not apply all $n^a(0)$ ambiguous moves first. The detailed proof is omitted due to page constraint. $\square$

One might expect the existence of a deceptive Mover strategy that approaches g_F at the beginning and misleads the Eater to consume more from the fake goal. Such behavior is known as an exaggeration [6]. The next lemma shows that this is not effective under our problem setting.

Lemma 3 (No Exaggeration)**.** *The Mover has no incentive to make any explicit moves towards g_F if the Eater uses π^{E*}.*

Proof. Without loss of generality consider $\mathscr{G}_i$ where real goal is g_i and fake goal is g_{-i}. We assume $p(0) \notin R_1$ because we already discussed R_1 in Lemma 1. The Mover cannot make any explicit move in R_3, so we can restrict our attention to the situation where $p(0) \in R_2$.

[4] The time arguments for p, $\mathbf{b}$, c_i, n^a and Δc_i are omitted for conciseness.

[5] Additional ambiguous moves are possible when the Mover is at the boundary of R_1.

Suppose that the Mover uses a different strategy $\pi^{M'} \neq \pi^{M*}$ which makes at least one explicit move towards g_{-i} in R_2. We can see that this "exaggeration" move will result in $\tilde{c}_i(t) = c_i(t)+1$ and $\tilde{c}_{-i}(t) = c_{-i}(t)-1$. From Def. 6 we have $\Delta\tilde{c}_i(t) = \Delta c_i(t)+2$. Since π^{E*} will respond by taking from g_{-i}, we will have $\Delta c_i(t+1) = \Delta c_i(t)+1$.

Although the change in $\Delta c_i(t)$ has no effect on the Eater's behavior after an explicit move, notice that a larger $\Delta c_i(t)$ is favorable for the Mover based on (11): i.e., the consumption after the ambiguous move. In this context, the exaggeration move improved $\Delta c_i(t)$ by 1. This increase either: (i) delays a positive $\Delta c_i(t)$ reaching 0 by one time step, (ii) makes $\Delta c_i(t) = 0$ increase to 1, or (iii) expedites a negative $\Delta c_i(t)$ reaching zero by 1 time step. Again, by recalling (11), we can see that this translates to a reduced consumption from g_i after an ambiguous move by at most 1.

However, notice that the above "benefit" causes an increase in the number of explicit moves remaining. Since this always penalizes the Mover by 1 consumption from g_i, we can conclude that deception by exaggeration will at best cancel the penalty, but never result into lower $b_i(T)$. □

Corollary 1 (Shortest Path)**.** *The Mover has no incentive to use a non-shortest path if the Eater uses π^{E*}.*

Proof. This result directly follows from Lem. 2 and 3. □

Lemma 4 (Ambiguous Moves First)**.** *The Mover has no incentive to deviate from making all the ambiguous moves first, given that the Eater uses π^{E*}.*

Proof. Without loss of generality suppose the Mover plays $\mathscr{G}_i$. From Cor. 1 we know that the Mover must stay on the shortest path towards g_i, which gives $T = n^a(0)+n_i^e(0)$.

We compare $b_i(T)$ under the equilibrium sequence generated from π^{M*}, with the one under $\pi^{M'} \neq \pi^{M*}$, which takes the same number of steps T, but in a different order: i.e., at least one explicit move before the final ambiguous move.

Since π^{E*} always takes from g_i after explicit move towards g_i regardless of its timing, consumption due to explicit moves under π^{M*} and $\pi^{M'}$ will be the same. Therefore, we focus on the Eater's behavior after the ambiguous moves. We will prove that making explicit moves earlier will never reduce the subsequent banana consumption from the real goal associated to the ambiguous moves. Let t_k^* and t_k' for $k \in \{1,...,n^a(0)\}$ denote the times when the k-th ambiguous move is used by π^{M*} and $\pi^{M'}$, respectively. Trivially, $t_k^* = k-1$ for all k. The timing for t_k' is delayed by the number of explicit moves used before the k-th ambiguous move, which we denote by $z_k \geq 0$.

Now, we will analyze the difference in the banana consumption under the two strategies by looking at $\Delta c_i(t)$ defined in (10). We will use $\Delta\tilde{c}_i^*$ and $\Delta\tilde{c}_i'$ to denote the ones for π^{M*} and $\pi^{M'}$, respectively. Recalling the discussion in Rem. 2, we know that $\Delta\tilde{c}_i(t) \leq 0$ will result in consumption from the real goal, g_i. Therefore, all we need to show is that

$$\Delta\tilde{c}_i^*(t_k^*) \geq \Delta\tilde{c}_i'(t_k'),\ \forall k \in \{1,...,n^a(0)\}. \tag{19}$$

The above inequality implies that for the k-th ambiguous move, the one from $\pi^{M'}$ will lead to the same or more consumption on the real goal compared to the one from π^{M*}.

To see (19), recall Rem. 3 and see that an explicit move towards g_i at time t will always reduce $\Delta c_i(t)$ by one. Also recall that an ambiguous move will make $\Delta c_i(t)$ approach zero (either from positive or negative side). Therefore, if $\Delta c_i(0) \geq 0$, then we have

$$\Delta c_i^*(t_k^*) = \max\{0,\, \Delta c_i(0)-k+1\},\ \text{whereas} \tag{20}$$

$$\Delta c_i'(t_k') \leq \max\{0,\, \Delta c_i(0)-k+1-z_k\}. \tag{21}$$

From Rem. 2, we know that $\Delta\tilde{c}_i^*(t_k^*) = \Delta c_i^*(t_k^*)$, and from Rem. 3, we have $\Delta\tilde{c}_i'(t_k') \leq \Delta c_i'(t_k')$, where the equality holds when $a^M(t_k')$ is ambiguous.

Thus (19) holds for $\Delta c_i(0) \geq 0$. With a similar argument, we can also show (19) for the case when $\Delta c_i(0) \leq 0$. □

With Cor. 1 and Lem. 4, we have shown that the Mover should stick to π^{M*} if the Eater uses π^{E*}.

C. Eater's Strategy

Now we study the optimality of π^{E*} when $p(t) \notin R_1$. Observe that π^{M*} uses explicit moves only inside R_1, which implies that π^{E*} will use the *exploiting* action (12) only in R_1. We therefore focus our attention to the *conservative* action.

Lemma 5. *For $p(t) \notin R_1$, the Eater has no incentive to deviate from its conservative action* (11)*, given that the Mover makes only ambiguous moves according to π^{M*}.*

Proof. Consider the effect of Eater's actions outside of R_1. The Eater must consider both $\mathscr{G}_i$ and $\mathscr{G}_{-i}$, but notice that the Mover using π^{M*} will make $n^a(0)$ ambiguous moves in both games. Without loss of generality, we assume $\Delta c_i(0) \geq 0$ and examine the two cases presented in Thm. 2.

Case 1: $n^a(0) \leq |\Delta c_i(0)|$. Note that $\Delta c_i(0)$ must be positive in this case. Under the equilibrium strategies, the Eater will take only from g_{-i} in the first $n^a(0)$ steps, and based on Thm. 2, we know the outcome is $V^{E*} = c_i(0)-\Delta c_i(0)$. Now, we consider a deviation $\pi^{E'} \neq \pi^{E*}$. Recalling Rem. 2, we can state $\Delta\tilde{c}_i(t) > 0, \forall t \leq n^a(0)$. The only way for $\pi^{E'}$ to deviate from π^{E*} is by eating $x \in \{0.5, 1\}$ from the goal with higher $\tilde{c}(t)$ (g_i in our case) at least once. Based on (17) the Mover's outcome under this strategy will be $J_i^M(\pi_i^{M*}, \pi^{E'}) \geq c_i(0)-n^a(0)+x$ and $J_{-i}^M(\pi_{-i}^{M*}, \pi^{E'}) \leq c_{-i}(0)-x = c_i(0)-\Delta c_i(0)-x$. Since $n^a(0) \leq |\Delta c_i(0)|$ the Eater's outcome results in $J^E(\pi^{M*}, \pi^{E'}) \leq c_i(0)-\Delta c_i(0)-x < V^{E*}$.

Case 2: $n^a(0) > |\Delta c_i(0)|$. Notice that $V_i^{M*} = V_{-i}^{M*}$ in Case 2, which is easy to see from (18) and the fact that $c_i(t) = c_{-i}(t)+\Delta c_i(t)$ and $\Delta c_i(t) = -\Delta c_{-i}(t)$. Now any deviation $\pi^{E'} \neq \pi^{E*}$ in the first $n^a(0)$ steps will only cause the above equality to break: i.e., $J_i^M(\pi_i^{M*}, \pi^{E'}) = V_i^{M*}+0.5x$ and $J_{-i}^M(\pi_{-i}^{M*}, \pi^{E'}) = V_{-i}^{M*}-0.5x$ for $x \in \mathbb{Z}$. Such deviation results in $J^E(\pi^{M*}, \pi^{E'}) = V^{M*}-0.5|x|$, which is suboptimal. □

To summarize, the Eater has no incentive to deviate from its conservative action after Mover's ambiguous moves. This concludes discussion of Eater's and Mover's equilibrium strategies, and provides justification of Theorem 1.

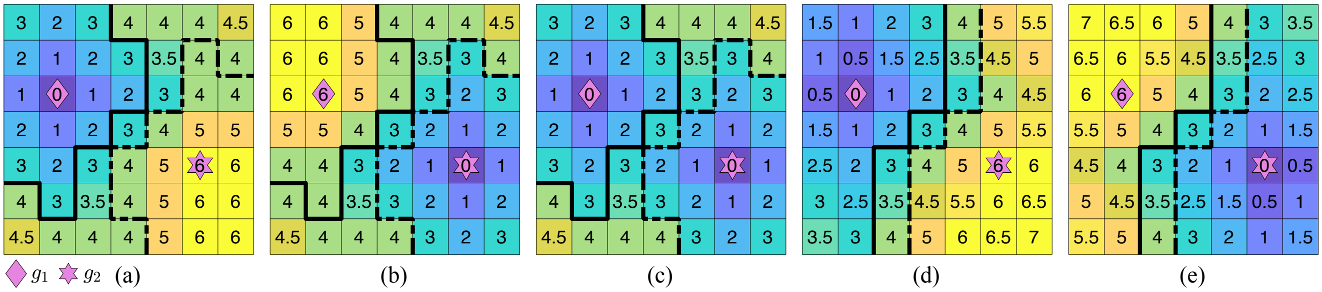

Fig. 4. The game outcome under different strategies: (a) V_1^{M*}, (b) V_2^{M*}, (c) V^{E*}; (d) $J_1^M(\pi_1^{M*}, \pi^{E'})$, (e) $J_2^M(\pi_1^{M*}, \pi^{E'})$.

V. Numerical Illustration

This section presents the numerical solution to the game environment shown earlier in Fig. 1. Figures 4a and 4b show the Mover's performance under π^{M*} and π^{E*} calculated based on Thm. 2. We highlight the boundary where the minimum in (2) switches between $\mathscr{G}_1$ and $\mathscr{G}_2$. The area between the two surfaces is where $V_1^{M*} = V_2^{M*}$. Figure 4c shows the Eater's equilibrium performance.

Suboptimal Eater: Figures 4e and 4d show the outcome under an Eater strategy $\pi^{E'}$ that uses exploiting action after explicit moves and $\mathbf{a}^E = [0.5, 0.5]$ after ambiguous moves. This strategy may lead to higher banana consumption than π^{E*} in certain scenarios. For example, in $\mathscr{G}_2$, if the Mover starts at the top left cell then $J_2^M(\pi_2^{M*}, \pi^{E'}) = 7$, while $J_2^{M*}(\pi_2^{M*}, \pi^{E*}) = 6$. However, if $\mathscr{G}_1$ is actually played, $J_1^M(\pi_1^{M*}, \pi^{E'}) = 1.5$, which worse than $J_1^{M*}(\pi_1^{M*}, \pi^{E*}) = 3$. Since the Eater does not know which game is played, we let the Eater optimize its worst-case performance as in (4), and in this sense π^{E*} indeed achieves a better worst-case guarantee.

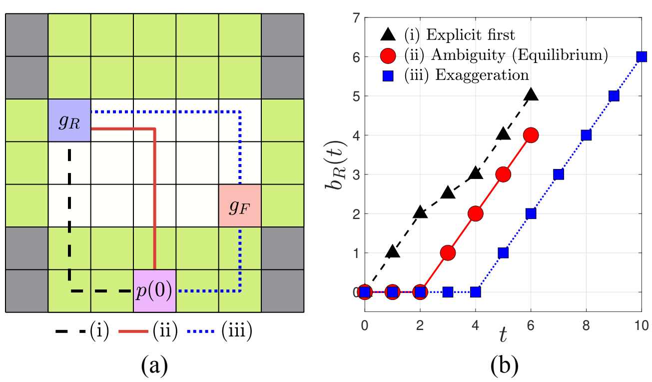

Fig. 5. Mover's paths starting with $b_R(0) = b_F(0) = 0$. (a) Trajectories in the grid world. (b) Corresponding banana consumption from g_R at each time step under π^{E*}.

Suboptimal Mover: Figure 5 compares the performance of three different Mover strategies against π^{E*}. Path (ii) is the shortest path induced by π^{M*}. At the end of Path (ii), the Eater consumed 4 bananas (see Fig. 5b), which is the least amount among the three trajectories, demonstrating the effectiveness of *ambiguity*. Path (i) is also a shortest path, but the Mover makes explicit moves before the ambiguous ones, which leads to 5 bananas consumed. Path (iii) corresponds to an exaggeration strategy, which tries to deceive the Eater by moving towards g_F first but actually results in 6 bananas consumed. This is an example of deviating from the shortest path and is clearly suboptimal as discussed in Lem. 3.

VI. Conclusion

We introduce and solve the Eater and the Mover game, which we use to explore the possibility of deception by motion when the observer must take an action, thereby allowing us to measure the effectiveness of the deception through the game outcome and not prescribing an estimator for the observer. The results demonstrate that the Mover cannot deceive the Eater by exaggeration if the Eater optimizes its worst-case performance. However, the ambiguity is still useful for the Mover to improve its performance. An interesting and useful avenue for future work is to study how the ideas of ambiguity and exaggeration generalize to a broader class of problems: e.g., multiple goals, general graph environments (which accommodates obstacles and different dynamics), different payoff functions, and continuous space.

References

[1] J. P. Hespanha, Y. S. Ateskan, and H. H. Kizilocak, "Deception in non-cooperative games with partial information," in *Proceedings of DARPA-JFACC Symp. on Advances in Enterprise Control*, 2000.

[2] Y. Yavin, "Pursuit-evasion differential games with deception or interrupted observation," *Computers & Mathematics With Applications*, vol. 13, pp. 191–203, 1987.

[3] M. O. Sayin and T. Başar, "Dynamic information disclosure for deception," in *IEEE Conf. Decis. Control (CDC)*, 2018, pp. 1110–1117.

[4] M. Ornik and U. Topcu, "Deception in optimal control," *Annual Allerton Conf. on Comm., Control, and Comp.*, pp. 821–828, 2018.

[5] A. D. Dragan, R. Holladay, and S. S. Srinivasa, "Deceptive robot motion: synthesis, analysis and experiments," *Autonomous Robots*, vol. 39, pp. 331–345, 2015.

[6] Y. Savas, M. O. Karabag, B. M. Sadler, and U. Topcu, "Deceptive planning for resource allocation," *ArXiv*, vol. abs/2206.01306, 2022.

[7] M. O. Karabag, M. Ornik, and U. Topcu, "Deception in supervisory control," *IEEE Trans. Aut. Control*, vol. 67, no. 2, pp. 738–753, 2022.

[8] M. Ramírez and H. Geffner, "Probabilistic plan recognition using off-the-shelf classical planners," *Proceedings of the AAAI Conf. Art. Int.*, pp. 1121–1126, 2010.

[9] P. Masters and S. Sardiña, "Deceptive path-planning," in *Int. Joint Conf. on Art. Int.*, 2017, p. 4368–4375.

[10] P. Masters and S. Sardina, "Cost-based goal recognition in navigational domains," *J. Artif. Intell. Res.*, vol. 64, pp. 197–242, 2019.

[11] C. Wayllace, P. Hou, W. G. S. Yeoh, and T. C. Son, "Goal recognition design with stochastic agent action outcomes," in *Int. Joint Conf. on Art. Int.*, 2016, p. 3279–3285.

[12] A. Kulkarni, M. E. Klenk, S. D. Rane, and H. Soroush, "Resource bounded secure goal obfuscation," in *AAAI Fall Symp. Integr. Plan., Diag. and Causal Reasoning*, 2018.

[13] P. Masters and S. Sardina, "Cost-based goal recognition for the path-planning domain," in *Int. Joint Conf. Art. Int.*, 2018.

[14] R. Pereira, N. Oren, and F. Meneguzzi, "Landmark-based heuristics for goal recognition," in *AAAI Conf. Art. Int.*, 2017.