

Strategically revealing capabilities in General Lotto games[★]

Keith Paarporn^{*} Philip N. Brown^{*}

^{*} *University of Colorado, Colorado Springs, Colorado Springs, CO
80918 USA (e-mail: kpaarporn@uccs.edu, pbrown2@uccs.edu).*

Abstract:

Can revealing one's competitive capabilities to an opponent offer strategic benefits? In this paper, we address this question in the context of General Lotto games, a class of two-player competitive resource allocation models. We consider an asymmetric information setting where the opponent is uncertain about the resource budget of the other player, and holds a prior belief on its value. We assume the other player, called the *signaler*, is able to send a noisy signal about its budget to the opponent. With its updated belief, the opponent then must decide to invest in costly resources that it will deploy against the signaler's resource budget in a General Lotto game. We derive the subgame perfect equilibrium to this extensive-form game. In particular, we identify necessary and sufficient conditions for which a signaling policy improves the signaler's resulting performance in comparison to the scenario where it does not send any signal. Moreover, we provide the optimal signaling policy when these conditions are met. Notably we find that for some scenarios, the signaler can effectively double its performance.

Keywords: Information design, Signaling, General Lotto games, Resource allocation

1. INTRODUCTION

The advancement of communication technologies and platforms have fundamentally shifted how information is sent, perceived, and ultimately utilized to make decisions. From a system-level perspective, revealing, concealing, or manipulating information can be a viable and impactful method for the control of multi-agent systems, particularly when those systems include strategic decision-makers. For example, travel information systems may broadcast recommended routes to drivers in a transportation network with the objective of lowering overall congestion [Zhu and Savla, 2022]. Advertisers may influence user behavior (e.g. purchasing decisions, engaging with certain content) on online platforms by making personalized recommendations [Ke et al., 2022; Candogan and Drakopoulos, 2020]. In these systems, a central authority implements a signaling policy that carefully determines what information to broadcast to a collection of uninformed users.

Due to its wide-ranging potential applications, the problem of strategically revealing or concealing information has received a great deal of research attention in recent years, often under the name of *Bayesian persuasion* [Kamenica and Gentzkow, 2011] or *information design* [Bergemann and Morris, 2019]. Under these frameworks, an informed agent crafts the information that is revealed to an uninformed agent (or agents); by selecting the revelation policy carefully, the informed agent can influence the posterior beliefs of the uninformed agent and can often thus control its resulting decisions. These techniques have primarily been applied to study the influence of multi-agent systems such as transportation networks, epidemic management,

and social networks [Massicot and Langbort, 2019; Zhu and Savla, 2022; Liu and Zhu, 2022; Wu and Amin, 2019; Candogan and Drakopoulos, 2020].

In this paper, we focus on an information designer (called the *signaler*) that has an opportunity to reveal information about its own competitive capabilities to an *adversary*, with which it is in direct competition. We study such a scenario in the context of General Lotto games, a popular game-theoretic model of competitive resource allocation. We base our analysis on a General Lotto game where the resource budget of the signaler is randomly drawn from a publicly-known Bernoulli distribution (studied recently in Paarporn et al. [2021a]). The true budget is the private information of the signaler. The signaler adopts a *signaling policy*; once the signaler's resource budget is realized, the signaling policy delivers a noisy signal of its budget to the adversary. The adversary then updates its beliefs about the signaler's budget, chooses an amount of resources to invest in (which is publicly disclosed), and subsequently engages in a General Lotto game with the signaler using its invested resources. We are primarily concerned with deriving the signaler's optimal policy, and identifying conditions on environmental parameters for which it offers performance improvements compared to not signaling at all. We note that our study is a departure from the traditional Bayesian Persuasion setup, for which the signaler does not participate in strategic interactions after signaling to the receiver.

Our main results are given as follows.

- We fully characterize the signaler's optimal signaling policy in all instances of the game.

[★] This work was supported in part by NSF Award #ECCS-2013779.

- We fully characterize necessary and sufficient conditions for which the optimal signaling policy provides performance improvements for the signaler.

We find that the optimal signaling policy derived maximizes the probability that the signal sent deters the adversary from competing at all in the General Lotto game. Interestingly, we find there are parameters under which the signaler can effectively double its performance by using an optimal policy.

Related works: Much recent work has been devoted to strategic signaling, particularly in the area of cyber-physical-human systems such as transportation networks [Massicot and Langbort, 2019; Wu et al., 2021; Gould and Brown, 2022; Ferguson et al., 2022]. In these works, a system planner who is informed about the state of the world (e.g., the presence of traffic accidents on a highway) decides what information to transmit to a user or group of uninformed self-interested system users, with the typical goal of improving system performance. Thus, this line of research is broadly focused on using strategic information provision in a benevolent way to act as a coordination mechanism among a large population of disorganized decision-makers. In contrast, our paper studies the information design problem in a competitive setting; i.e., rather than attempting to coordinate behavior among a group of users, our signaler is attempting to manipulate the uncertainty of an adversary.

Another line of research studies information design in competitive settings, but focuses mainly on competition in markets: e.g., two firms competing over market share may reveal/conceal information about their competitors' product quality [Kamenica and Gentzkow, 2011; Ivanov, 2013; Board and Lu, 2018; Li and Norman, 2018]. These works generally focus on the effects of competitive information provision on metrics such as consumer surplus. In contrast, our model considers providing information directly to an adversary, rather than to market participants.

Perhaps most closely aligned with our work is the recent studies on strategic information provision in *contests*. Several recent works have shown that the strategic revelation of information to opponents can serve as a viable competitive strategy, ranging from pre-commitments to full revelation of information [Paarporn et al., 2021b; Epstein and Mealem, 2013]. Zhang and Zhou [2016] consider a scenario in which a *contest organizer* can force contestants to reveal private information in an attempt to maximize the effort expended by contestants; in that work, if contestants have binary valuations, the optimal policy is either no or full disclosure. Other papers study similar problems; Fu et al. [2011] study disclosure policy when contest entry is stochastic and Denter et al. [2011] examine the strategic effects of time-delayed information revelation. Similar to our work (but using a different contest model), Epstein and Mealem [2013] consider a contest in which the informed player can choose to either conceal or fully reveal its abilities to the uninformed player.

In this paper, we consider the role that information signaling has in adversarial interactions. The driving question is: can signaling one's capabilities to an uncertain opponent offer strategic benefits?

2. PRELIMINARIES ON GENERAL LOTTO GAMES

To build up to our signaling General Lotto game, we provide some background on two-player simultaneous-move General Lotto games. First, we present the classic complete information formulation in Section 2.1. Then we present an incomplete information setting with asymmetric budget uncertainty from the recent literature in Section 2.2.

2.1 Complete information Lotto games

A (complete information) General Lotto game consists of two players, $\mathcal{A}$ and $\mathcal{B}$. Each player is tasked with allocating their endowed resource budgets $A, B > 0$ across a set of n battlefields. Each battlefield has an associated value $v_j > 0$, $j \in [n]$. An allocation for $\mathcal{A}$ is any vector $\mathbf{x}_{\mathcal{A}} \in \mathbb{R}_{\geq 0}^n$, and similarly for $\mathcal{B}$. An admissible strategy for $\mathcal{A}$ is a randomization $F_{\mathcal{A}}$ over allocations such that the expended resources do not exceed the budget A in expectation. Specifically, $F_{\mathcal{A}}$ is an n -variate (cumulative) distribution that belongs to the family

$$\mathbb{F}(A) \triangleq \left\{ F : \mathbb{E}_{\mathbf{x}_{\mathcal{A}} \sim F} \left[\sum_{j=1}^n x_{\mathcal{A},j} \right] \leq A \right\}. \quad (1)$$

and similarly, $F_{\mathcal{B}} \in \mathbb{F}(B)$. Given a strategy profile $(F_{\mathcal{A}}, F_{\mathcal{B}})$, the utility of player $\mathcal{A}$ is

$$u_{\mathcal{A}}(F_{\mathcal{A}}, F_{\mathcal{B}}) \triangleq \mathbb{E}_{\substack{\mathbf{x}_{\mathcal{A}} \sim F_{\mathcal{A}} \\ \mathbf{x}_{\mathcal{B}} \sim F_{\mathcal{B}}}} \left[\sum_{j=1}^n v_j \cdot \mathbf{1}_{\{x_{\mathcal{A},j} \geq x_{\mathcal{B},j}\}} \right] \quad (2)$$

where $\mathbf{1}_{\{\cdot\}}$ is 1 if the statement in the bracket is true, and 0 otherwise¹. It follows that the utility of player $\mathcal{B}$ is

$$u_{\mathcal{B}}(F_{\mathcal{A}}, F_{\mathcal{B}}) \triangleq \phi - u_{\mathcal{A}}(F_{\mathcal{A}}, F_{\mathcal{B}}) \quad (3)$$

where $\phi \triangleq \sum_{j=1}^n v_j$ is the total sum of battlefield values. An instance of the complete information General Lotto game is denoted by $\text{GL}(A, B)$. An equilibrium is a strategy profile $(F_{\mathcal{A}}^*, F_{\mathcal{B}}^*)$ such that

$$\begin{aligned} u_{\mathcal{A}}(F_{\mathcal{A}}^*, F_{\mathcal{B}}^*) &\geq u_{\mathcal{A}}(F_{\mathcal{A}}, F_{\mathcal{B}}^*), \quad \forall F_{\mathcal{A}} \in \mathbb{F}(A) \\ u_{\mathcal{B}}(F_{\mathcal{A}}^*, F_{\mathcal{B}}^*) &\geq u_{\mathcal{B}}(F_{\mathcal{A}}^*, F_{\mathcal{B}}), \quad \forall F_{\mathcal{B}} \in \mathbb{F}(B). \end{aligned} \quad (4)$$

The unique equilibrium payoffs in General Lotto games is well-established in the literature.

Theorem 2.1. (Kovenock and Roberson [2021]). Consider any General Lotto game $\text{GL}(A, B)$. The payoff to player $\mathcal{A}$ in any equilibrium is given by

$$\pi_{\mathcal{A}}^{\text{ci}}(A, B) \triangleq \phi \cdot \begin{cases} \frac{A}{2B}, & \text{if } A \leq B \\ 1 - \frac{B}{2A}, & \text{if } A > B \end{cases}, \quad (5)$$

and the payoff to player $\mathcal{B}$ is $\pi_{\mathcal{B}}^{\text{ci}}(A, B) = \phi - \pi_{\mathcal{A}}^{\text{ci}}(A, B)$.

Note that the payoffs depend on the total value ϕ , and not on individual values of the valuation vector $\mathbf{v}$. Hence, we omit specifying $\mathbf{v}$ in the notation $\text{GL}(A, B)$.




Fig. 1. The signaling General Lotto game. In Stage 1, the signaler (player $\mathcal{A}$) selects a signaling policy σ , which becomes common knowledge. Before Stage 2, the signaler's budget type is realized according to the prior distribution p , and a signal s is sent to player $\mathcal{B}$ according to the distribution specified by σ . The receiver (player $\mathcal{B}$) updates its belief. In Stage 2, player $\mathcal{B}$ decides how many resources $B \geq 0$ to invest in, paying the cost $c \cdot B$. In Stage 3, both players simultaneously engage against each other in the Bernoulli Lotto game. The sequence of choices and events is standard in the information design literature [Kamenica, 2019].

2.2 Lotto games with asymmetric budget uncertainty

We present an asymmetric information General Lotto game from the recent literature [Paarporn et al., 2021a], which will serve as the basis of our signaling General Lotto game.

Player $\mathcal{A}$ has two possible budget types, $t \in T = \{h, \ell\}$. The type is drawn according to a Bernoulli distribution with parameter p . Specifically, with probability $p \in [0, 1]$, $\mathcal{A}$'s type is h and it is endowed with a high budget A_h . With probability $1 - p$, $\mathcal{A}$'s type is ℓ and it is endowed with a low budget A_ℓ , with $A_h \geq A_\ell \geq 0$. The realized type is known only to player $\mathcal{A}$, but the distribution and values A_h, A_ℓ are common knowledge. Player $\mathcal{B}$ thus holds the prior belief p on the high type A_h . The budget endowment $B \geq 0$ of player $\mathcal{B}$ is common knowledge.

An admissible strategy for player $\mathcal{A}$ is a pair $\mathbf{F}_\mathcal{A} = \{F_\mathcal{A}^h, F_\mathcal{A}^\ell\} \in \mathbb{F}(A_h) \times \mathbb{F}(A_\ell)$, where $F_\mathcal{A}^h$ or $F_\mathcal{A}^\ell$ is implemented depending on which budget type is realized. An admissible strategy for player $\mathcal{B}$ is a single strategy $F_\mathcal{B} \in \mathbb{F}(B)$ that is implemented regardless of which type is realized. If $\mathcal{A}$'s private type is $t \in \{h, \ell\}$, the ex-interim expected utilities given the strategy profile $(\mathbf{F}_\mathcal{A}, F_\mathcal{B})$ are defined as

$$\begin{aligned}
 U_\mathcal{A}(\mathbf{F}_\mathcal{A}, F_\mathcal{B}; t) &\triangleq u_\mathcal{A}(F_\mathcal{A}^t, F_\mathcal{B}) \\
 U_\mathcal{B}(F_\mathcal{B}, \mathbf{F}_\mathcal{A}) &\triangleq p \cdot u_\mathcal{B}(F_\mathcal{B}, F_\mathcal{A}^h) + (1 - p) \cdot u_\mathcal{B}(F_\mathcal{B}, F_\mathcal{A}^\ell)
 \end{aligned} \tag{6}$$

where $u_\mathcal{A}$ and $u_\mathcal{B}$ are defined from (2). A Bayes-Nash equilibrium is a strategy profile $(\mathbf{F}_\mathcal{A}^*, F_\mathcal{B}^*)$ such that for every type $t \in \{h, \ell\}$,

$$\begin{aligned}
 U_\mathcal{A}(\mathbf{F}_\mathcal{A}^*, F_\mathcal{B}^*; t) &\geq U_\mathcal{A}(\mathbf{F}_\mathcal{A}, F_\mathcal{B}^*; t), \quad \forall \mathbf{F}_\mathcal{A} \in \mathbb{F}(A_h) \times \mathbb{F}(A_\ell) \\
 U_\mathcal{B}(F_\mathcal{B}^*, \mathbf{F}_\mathcal{A}^*) &\geq U_\mathcal{B}(F_\mathcal{B}, \mathbf{F}_\mathcal{A}^*), \quad \forall F_\mathcal{B} \in \mathbb{F}(B)
 \end{aligned} \tag{7}$$

We refer to this setup as a Bernoulli Lotto game, where a particular instance is characterized by the parameters $\mathbb{P} \triangleq (A_h, A_\ell, p)$ and B . We will denote a Bernoulli Lotto game as $\text{BL}(\mathbb{P}, B)$.

Unique equilibrium payoffs to every Bernoulli Lotto game were completely characterized in recent work [Paarporn et al., 2021a]. For an equilibrium $(\mathbf{F}_\mathcal{A}^*, F_\mathcal{B}^*)$, we will write

¹ An arbitrary tie-breaking rule may be selected, without changing our results. This is generally true in General Lotto games [Kovenock and Roberson, 2021]. For simplicity, we will assume ties are awarded to player $\mathcal{A}$.

player $\mathcal{A}$'s equilibrium ex-interim payoff in $\text{BL}(\mathbb{P}, B)$ given the budget type $t \in \{h, \ell\}$ as

$$\pi_\mathcal{A}(\mathbb{P}, B; t) \triangleq U_\mathcal{A}(\mathbf{F}_\mathcal{A}^*, F_\mathcal{B}^*; t) \tag{8}$$

Player $\mathcal{B}$ has a single (public) budget type and must reason (based on its belief) about the budget type of player $\mathcal{A}$. Its equilibrium payoff is written as

$$\pi_\mathcal{B}(\mathbb{P}, B) \triangleq \phi - (p\pi_\mathcal{A}(\mathbb{P}, B; h) + (1 - p)\pi_\mathcal{A}(\mathbb{P}, B; \ell)). \tag{9}$$

3. THE SIGNALING GENERAL LOTTO GAME

Consider a Bernoulli Lotto game. Before engaging in competition, player $\mathcal{A}$ has an opportunity to re-shape the belief of player $\mathcal{B}$ by sending a (noisy) signal $s \in \{h, \ell\}$ directly to player $\mathcal{B}$. Specifically, player $\mathcal{A}$ adopts a *signaling policy* $\sigma = \{\sigma(s|t)\}_{s \in S, t \in T}$, where $S = \{h, \ell\}$ is the set of possible signals that can be sent to player $\mathcal{B}$. Here, $\sigma(s|t)$ is the probability that the signal $s \in S$ is sent to player $\mathcal{B}$, given that budget type $t \in T$ is realized.

In this paper, we will focus on the sub-class of signaling policies that truthfully signal 'high' when the budget type is actually high. Such an assumption is warranted – intuitively, a competitor would not want to signal that it is weaker than it actually is. Future work will analyze the scenario where this assumption is eliminated².

Assumption 1. Player $\mathcal{A}$'s set of admissible signaling policies is restricted to the sub-class of signaling policies $\sigma \in \Sigma_h$, where Σ_h is the set of policies that satisfy $\sigma(h|h) = 1$.

Any signaling policy in this class $\sigma \in \Sigma_h$ satisfies

$$\begin{aligned}
 \sigma(h|h) &= 1, \quad \sigma(\ell|h) = 0 \\
 \sigma(h|\ell) + \sigma(\ell|\ell) &= 1
 \end{aligned} \tag{10}$$

Note that any signaling policy $\sigma \in \Sigma_h$ is completely characterized by a single number: $q = \sigma(h|\ell)$.

The interaction unfolds in the following three-stage extensive form game.

Stage 1: Player $\mathcal{A}$ selects a feasible signaling policy $\sigma \in \Sigma_h$. The budget type $t \in T$ is drawn, and a signal $s \sim \sigma(\cdot|t)$ is sent to player $\mathcal{B}$. Player $\mathcal{B}$ performs a Bayesian update on its prior belief. Specifically, if $s = h$, σ induces

² Computational results (not shown) indeed suggest that even when Assumption 1 is eliminated, the same SPE signaling policy for player $\mathcal{A}$ detailed in the main result, Theorem 4.1, still holds.

the posterior belief $\mu_h \triangleq \frac{p}{p+q(1-p)}$ on the high budget type. If $s = \ell$, σ induces the posterior belief $\mu_\ell \triangleq 0$ on the high budget type.

Stage 2: Player $\mathcal{B}$ selects an amount of resources $B \geq 0$ to invest in. It pays the cost $c \cdot B$, where $c > 0$ is its *per-unit cost*.

Stage 3: Player $\mathcal{A}$ (with budget A_t) and player $\mathcal{B}$ (with budget B) engage in the ex-interim stage of the Bernoulli Lotto game $\text{BL}(\mathbb{P}_s, B)$, where $\mathbb{P}_s \triangleq (A_h, A_\ell, \mu_s)$ ($s \in \{h, \ell\}$) is $\mathcal{B}$'s posterior distribution on budget types. The final payoff to $\mathcal{A}$ is its ex-interim equilibrium payoff in $\text{BL}(\mathbb{P}_s, B)$, denoted $\pi_{\mathcal{A}}(\mathbb{P}_s, B; t)$. The final payoff to $\mathcal{B}$ is its (Bayes-Nash) equilibrium payoff under the belief $\mathbb{P}_s$ (denoted $\pi_{\mathcal{B}}(\mathbb{P}_s, B)$) minus the investment cost $c \cdot B$ from Stage 2.

The prior distribution on budget types $\mathbb{P} = (A_h, A_\ell, p)$ and player $\mathcal{B}$'s per-unit cost c are defining parameters in the extensive-form signaling game. We denote an instance of the game as $\text{SG}(\mathbb{P}, c)$. We consider the following standard solution concept for extensive-form games.

Definition 1. A pair (σ^*, B^*) , where $\sigma^* \in \Sigma_h$ and $B^* : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$, is a *subgame perfect equilibrium (SPE)* of $\text{SG}(\mathbb{P}, c)$ if:

(1) For any signaling policy $\sigma \in \Sigma_h$ and $s \sim \sigma(\cdot|t)$ with $t \in T$,

$$B^*(\mu_s) \in \arg \max_{B \geq 0} \{\pi_{\mathcal{B}}(\mathbb{P}_s, B) - c \cdot B\} \quad (11)$$

(2) The signaling policy σ^* solves

$$\max_{\sigma \in \Sigma_h} \Pi_{\mathcal{A}}(\sigma) \triangleq \mathbb{E}_t [\mathbb{E}_{s \sim \sigma(\cdot|t)} [\pi_{\mathcal{A}}(\mathbb{P}_s, B^*(\mu_s); t)]] \quad (12)$$

Note that player $\mathcal{B}$'s investment decision in Stage 2 is contingent on a realization of the signal s (11). Player $\mathcal{A}$'s choice of signaling policy σ in Stage 1 is taken before its budget type is realized – it thus considers final payoffs in expectation with respect to t and s (12). These are standard formulations in the Bayesian persuasion literature [Kamenica, 2019].

Definition 2. The *trivial signaling policy*, denoted $\emptyset$, is the signaling policy for which $q = 1$. In particular, this policy always sends $s = h$ regardless of the budget type, leaving player $\mathcal{B}$'s posterior belief unchanged from the prior, i.e. $\mu_h = \mu_\ell = p$. We denote the payoff player $\mathcal{A}$ obtains by implementing a trivial policy as $\Pi_{\mathcal{A}}^{\text{ns}} \triangleq \Pi_{\mathcal{A}}(\emptyset)$.

The trivial signaling policy is thus equivalent to not signaling at all. The primary goal of this paper is to identify conditions on $(\mathbb{P}, c)$ for which player $\mathcal{A}$'s payoff in an SPE of $\text{SG}(\mathbb{P}, c)$ exceeds $\Pi_{\mathcal{A}}(\emptyset)$.

4. MAIN RESULTS

The following result identifies necessary and sufficient conditions for which the signaling policy from the SPE outperforms the trivial policy.

Theorem 4.1. The SPE signaling policy $\sigma^* \in \Sigma_h$ of $\text{SG}(\mathbb{P}, c)$ outperforms the trivial policy, i.e. $\Pi_{\mathcal{A}}(\sigma^*) > \Pi_{\mathcal{A}}^{\text{ns}}$, if and only if (c, p) satisfies either of the following:




Fig. 2. Illustration of the range of parameters where signaling is beneficial for player $\mathcal{A}$. The labels indicate the regions specified in Theorem 4.1. The intensity of the plot indicates the value $100 \times (\frac{\Pi_{\mathcal{A}}^*}{\Pi_{\mathcal{A}}^{\text{ns}}} - 1)$, which is the percent improvement attainable from the SPE signaling policy. In this plot, we have set $A_h = 1.2$, $A_\ell = 0.5$, and $\phi = 1$.

$$(1) \quad \frac{\phi}{2A_h} \leq c < \min \left\{ \frac{\phi}{2A(p)}, \lambda(p) \right\} \text{ and } p \in [0, 1], \text{ where} \quad (13)$$

$$\lambda(p) \triangleq \frac{\phi}{2A_h^2} \left(\sqrt{(1-p)A_\ell} + \sqrt{pA_h + (1-p)A_\ell} \right)^2.$$

The SPE signaling policy is $\sigma^*(h|\ell) = \frac{p}{1-p} \frac{2cA_h - \phi}{\phi - 2cA_\ell}$.

$$(2) \quad \max \left\{ \lambda(p), \frac{\phi}{2A_h} \right\} \leq c < \frac{\phi}{2(A_h - A_\ell)}, p < \frac{A_h - 2A_\ell}{A_h - A_\ell}, \text{ and} \quad (14)$$

$$f_p > \frac{\phi - 2cA_\ell}{2cA_h - \phi} \cdot \frac{\sqrt{c\phi A_\ell/2}}{\phi - \sqrt{c\phi A_\ell/2}}$$

where $f_p \triangleq \frac{p}{\sqrt{1-p-(1-p)}} \cdot \frac{p}{1-p} \frac{2cA_h - \phi}{\phi - 2cA_\ell}$. The SPE signaling policy is $\sigma^*(h|\ell) = \frac{p}{1-p} \frac{2cA_h - \phi}{\phi - 2cA_\ell}$.

$$(3) \quad \max \left\{ \frac{\phi}{2(A_h - A_\ell)}, \frac{\phi}{2A_\ell} \left(\frac{f_p - \sqrt{f_p^2 - f_p + 1}}{f_p - 1} \right)^2 \right\} \leq c < \frac{(1-p)\phi}{2A_\ell}$$

and $p < \frac{A_h - 2A_\ell}{A_h - A_\ell}$. The SPE signaling policy is $\sigma^*(h|\ell) = \frac{p}{1-p} \frac{2cA_h - \phi}{\phi - 2cA_\ell}$.

Figure 2 depicts an example of the regions given by the above conditions, as well as the percent improvement in $\mathcal{A}$'s performance that it obtains by implementing the SPE signaling policy.

Discussion of main results: There are several interesting observations of Theorem 4.1. For parameters near the top border of the region in Figure 2 (i.e. setting $c = \frac{\phi}{2A_h}$ and $p \rightarrow 1$), player $\mathcal{A}$'s performance approaches a two-fold improvement compared to the performance of the trivial policy. Additionally, the sharp discontinuities in the plot indicate that player $\mathcal{A}$'s performance, and in turn, player $\mathcal{B}$'s investment decision, is highly sensitive to changes in the underlying parameters.

Outside of the indicated regions, the SPE signaling policy is the trivial policy. That is, no signaling policy can strictly improve upon $\Pi_{\mathcal{A}}^{\text{ns}}$. For parameters to the right of the indicated regions, player $\mathcal{B}$'s cost to invest in resources is sufficiently expensive such that its SPE investment is zero when $\mathcal{A}$ uses the trivial signaling policy. In particular, player $\mathcal{A}$ is able to win the entire contest without signaling

at all. Now, consider parameters to the left of the indicated regions. Player $\mathcal{B}$'s cost to invest in resources is cheap enough such that its SPE investment is high, and no signaling policy is able to induce $\mathcal{B}$ to invest in a lower amount of resources.

Under Assumption 1, a signaling policy $\sigma \in \Sigma_h$ is determined by a single number $\sigma(h|\ell) = q \in [0, 1]$, i.e. the probability the signal $s = h$ is sent when the budget type is low. The SPE policy is given by q^* (values specified in the statement). Interestingly, when the parameters satisfy any of the conditions in Theorem 4.1, the policies in the range $q \in [0, q^*]$ induce player $\mathcal{B}$ to invest zero resources in Stage 2, given that it received the signal $s = h$. In a sense, q^* is intuitively the highest fraction of time $\mathcal{A}$ can “lie” about its low budget type, such that player $\mathcal{B}$ will refrain from competing in the General Lotto game in Stage 3. Moreover, we find that for condition (1) in Theorem 4.1, the signaling policy $q = 0$ still provides an improvement over a trivial policy. In other words, full revelation of one's budget type performs better than not revealing any additional information for certain parameters.

5. ANALYSIS

In this section, we provide the derivation of the SPE of $\text{SG}(\mathbb{P}, c)$, and the proof of Theorem 4.1. The analysis hinges on using the equilibrium characterizations from previous work [Paarporn et al., 2021a]. We first provide the benchmark payoff that player $\mathcal{A}$ obtains with a trivial policy.

Lemma 5.1. (Lemma C.2 in [Paarporn et al., 2021a]).

Consider the extensive form game $\text{SG}(\mathbb{P}, c)$. If $p \geq \frac{A_h - 2A_\ell}{A_h - A_\ell}$, then

$$\Pi_{\mathcal{A}}^{\text{ns}} = \begin{cases} \sqrt{\frac{c\phi\bar{A}(p)}{2}}, & \text{if } c < \frac{\phi}{2\bar{A}(p)} \\ \phi, & \text{if } c \geq \frac{\phi}{2\bar{A}(p)} \end{cases} \quad (15)$$

If $p < \frac{A_h - 2A_\ell}{A_h - A_\ell}$, then

$$\Pi_{\mathcal{A}}^{\text{ns}} = \begin{cases} \sqrt{\frac{c\phi\bar{A}(p)}{2}}, & \text{if } c \in [0, \lambda(p)) \\ p\phi + \sqrt{\frac{c\phi(1-p)A_\ell}{2}}, & \text{if } c \in [\lambda(p), \frac{(1-p)\phi}{2A_\ell}) \\ \phi, & \text{if } c \geq \frac{(1-p)\phi}{2A_\ell} \end{cases} \quad (16)$$

where $\bar{A}(p) \triangleq pA_h + (1-p)A_\ell$ is the expected budget of $\mathcal{A}$ under belief p , and $\lambda(p)$ was defined in (13).

5.1 SPE investment level

The SPE of $\text{SG}(\mathbb{P}, c)$ can be derived by backwards induction. Hence, we first derive the optimal investment $B^*(\mu)$ (11) in Stage 2, given any belief $\mu \in [0, 1]$ on the high budget. The characterization of $B^*(\mu)$ is derived from prior work:

Lemma 5.2. (Lemma C.1 in [Paarporn et al., 2021a]).

Consider any posterior belief $\mu \in [0, 1]$. If $\mu \geq \frac{A_h - 2A_\ell}{A_h - A_\ell}$, then the SPE investment for player $\mathcal{B}$ is

$$B^*(\mu) = \begin{cases} \sqrt{\frac{\bar{A}(\mu)\phi}{2c}}, & \text{if } c < \frac{\phi}{2\bar{A}(\mu)} \\ 0, & \text{if } c \geq \frac{\phi}{2\bar{A}(\mu)} \end{cases} \quad (17)$$

If $\mu < \frac{A_h - 2A_\ell}{A_h - A_\ell}$, then the SPE investment for player $\mathcal{B}$ is

$$B^*(\mu) = \begin{cases} \sqrt{\frac{\bar{A}(\mu)\phi}{2c}}, & \text{if } c \in [0, \lambda(\mu)) \\ \sqrt{\frac{(1-\mu)A_\ell\phi}{2c}}, & \text{if } c \in [\lambda(\mu), \frac{(1-p)\phi}{2A_\ell}) \\ 0, & \text{if } c \geq \frac{(1-\mu)\phi}{2A_\ell} \end{cases} \quad (18)$$

where $\lambda(\mu)$ is defined in (13).

Here, we observe that the SPE investment $B^*(\mu)$ is increasing in the belief μ for low investment costs (first entry of (17), (18)), is decreasing for intermediate investment costs (second entry of (18)), and is zero for high investment costs (second entry of (17), third entry of (18)).

5.2 SPE signaling policy

Under Assumption 1, any signaling policy $\sigma \in \Sigma_h$ is uniquely determined by a single variable $q \triangleq \sigma(h|\ell)$. For any $q \in [0, 1]$, we can write the objective in (12) as

$$\begin{aligned} \Pi_{\mathcal{A}}(\sigma) &= p \cdot \pi_{\mathcal{A}}(\mathbb{P}_h, B^*(\mu_h); h) \\ &\quad + (1-p) \cdot q \cdot \pi_{\mathcal{A}}(\mathbb{P}_h, B^*(\mu_h); \ell) \\ &\quad + (1-p) \cdot (1-q) \cdot \pi_{\mathcal{A}}(\mathbb{P}_\ell, B^*(\mu_\ell); \ell) \end{aligned} \quad (19)$$

where under Assumption 1, $\mu_h = \frac{p}{p+q(1-p)}$ and $\mu_\ell = 0$. Using Lemma 5.2 and the characterizations $\pi_{\mathcal{A}}$ from [Paarporn et al., 2021a], we obtain the following expressions for player $\mathcal{A}$'s ex-interim equilibrium payoffs $\pi_{\mathcal{A}}(\mathbb{P}_s, B^*(\mu_s); t)$.

Lemma 5.3. ([Paarporn et al., 2021a]). Consider any posterior belief $\mu \in [0, 1]$, and denote $\mathbb{P}_\mu = (A_h, A_\ell, \mu)$. If $\mu \geq \frac{A_h - 2A_\ell}{A_h - A_\ell}$, then for $t \in \{h, \ell\}$:

$$\pi_{\mathcal{A}}(\mathbb{P}_\mu, B^*(\mu); t) = \begin{cases} \sqrt{\frac{c\phi A_t^2}{2\bar{A}(\mu)}}, & \text{if } c < \frac{\phi}{2\bar{A}(\mu)} \\ \phi, & \text{if } c \geq \frac{\phi}{2\bar{A}(\mu)} \end{cases} \quad (20)$$

If $\mu < \frac{A_h - 2A_\ell}{A_h - A_\ell}$, then for $t \in \{h, \ell\}$:

$$\begin{aligned} \pi_{\mathcal{A}}(\mathbb{P}_\mu, B^*(\mu); t) &= \\ &\begin{cases} \sqrt{\frac{c\phi A_t^2}{2\bar{A}(\mu)}}, & \text{if } c \in [0, \lambda(\mu)) \\ \mathbb{1}_{t=h} \cdot \phi + \mathbb{1}_{t=\ell} \cdot \sqrt{\frac{c\phi A_\ell}{2(1-\mu)}}, & \text{if } c \in [\lambda(\mu), \frac{(1-\mu)\phi}{2A_\ell}) \\ \phi, & \text{if } c \geq \frac{(1-\mu)\phi}{2A_\ell} \end{cases} \end{aligned} \quad (21)$$

where $\lambda(\mu)$ is defined in (13).

Remark 1. We point out that in the second entry of (21), player $\mathcal{A}$ secures the entire prize ϕ when endowed with the high budget $t = h$, even though player $\mathcal{B}$ invests non-zero resources to the competition (second entry of (18)). This is due to the players' equilibrium allocation profile $(F_{\mathcal{A}}^*, F_{\mathcal{B}}^*)$:

in this regime, the support of $F_{\mathcal{A},j}^{h*}$ on the allocation to any battlefield j is the interval $[2B, 2(A_h - B)] \cdot \frac{v_j}{\phi}$, whereas the support of $F_{\mathcal{B},j}^*$ is the interval $[0, 2B] \cdot \frac{v_j}{\phi}$ (cf. Section A.5, the “ $\mathcal{R}_4$ ” region [Paarporn et al., 2021a]). Thus, player $\mathcal{B}$ only competes with the low budget type $t = \ell$ in this case.

In regimes where $\mathcal{B}$ ’s investment cost is low, player $\mathcal{A}$ ’s payoff is decreasing in the belief μ (first entry of (20),(21)) regardless of the budget type. A final technical Lemma is stated below.

Lemma 5.4. The following properties hold.

- $\min \left\{ \frac{\phi}{2\bar{A}(p)}, \lambda(p) \right\} = \begin{cases} \lambda(p), & \text{if } p < \frac{A_h - 2A_\ell}{A_h - A_\ell} \\ \frac{\phi}{2\bar{A}(p)}, & \text{if } p \geq \frac{A_h - 2A_\ell}{A_h - A_\ell} \end{cases}$
- For $p < \frac{A_h - 2A_\ell}{A_h - A_\ell}$, $\max \left\{ \frac{\phi}{2\bar{A}(p)}, \lambda(p), \frac{\phi(1-p)}{2A_\ell} \right\} = \frac{\phi(1-p)}{2A_\ell}$.
- It holds that $\lambda(p)$ is strictly increasing on $p \in [0, \frac{A_h - 2A_\ell}{A_h - A_\ell})$ and strictly decreasing on $p \in (\frac{A_h - 2A_\ell}{A_h - A_\ell}, 1]$.

5.3 Proof of Theorem 4.1

For space considerations, we will omit the technical details for the proof of Theorem 4.1. However, these may be found in the online version [Paarporn and Brown, 2022]. The general approach is to analyze the expected final payoff for player $\mathcal{A}$, $\Pi_{\mathcal{A}}(q)$ (12), as a function of the signaling policy $q \in [0, 1]$. We characterize the maximizer q^* , and identify conditions on the parameters (c, p) for which $\Pi_{\mathcal{A}}(q^*) > \Pi_{\mathcal{A}}^{\text{ns}}$. Such an analysis is enabled through the four technical Lemmas in this section. We note that $\Pi_{\mathcal{A}}(q)$ has points of discontinuity, due to Lemma 5.3. The analysis may be divided into three distinct cases, which correspond to the items in the statement of Theorem 4.1.

6. CONCLUSION

This paper studies a competitive interaction between a signaler and an adversary, where the signaler has the opportunity to provide additional information about its capabilities to the adversary. We formulated this interaction as an extensive-form game, and used a General Lotto game model as the basis of the competition model. Leveraging recent results of incomplete information General Lotto games, we derived the optimal signaling policies within a sub-class of policies. Moreover, we derived necessary and sufficient conditions under which the optimal policy offers performance improvements to the signaler over not signaling at all. Future work will focus on deriving optimal policies over the entire space of signaling policies.

REFERENCES

- Bergemann, D. and Morris, S. (2019). Information Design: A Unified Perspective. *Journal of Economic Literature*, 57(1), 44–95.
- Board, S. and Lu, J. (2018). Competitive Information Disclosure in Search Markets. *Journal of Political Economy*, 126(5), 1965–2010.
- Candogan, O. and Drakopoulos, K. (2020). Optimal signaling of content accuracy: Engagement vs. misinformation. *Operations Research*, 68(2), 497–515.
- Denter, P., Morgan, J., and Sisak, D. (2011). ‘Where Ignorance is Bliss, ‘Tis Folly to Be Wise’: Transparency in Contests. *SSRN Electronic Journal*.
- Epstein, G.S. and Meale, Y. (2013). Who gains from information asymmetry? *Theory and Decision*, 75(3), 305–337.
- Ferguson, B.L., Brown, P.N., and Marden, J.R. (2022). Avoiding Unintended Consequences: How Incentives Aid Information Provisioning in Bayesian Congestion Games. In *61st IEEE Conference on Decision and Control (to appear)*.
- Fu, Q., Jiao, Q., and Lu, J. (2011). On disclosure policy in contests with stochastic entry. *Public Choice*, 148(3-4), 419–434.
- Gould, B.T. and Brown, P.N. (2022). Information Design for Vehicle-to-Vehicle Communication. *Transportation Research Part C: Emerging Technologies (under review)*.
- Ivanov, M. (2013). Information revelation in competitive markets. *Economic Theory*, 52(1), 337–365.
- Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics*, 11, 249–272.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.
- Ke, T., Lin, S., and Lu, M.Y. (2022). Information Design of Online Platforms. *SSRN Electronic Journal*.
- Kovenock, D. and Roberson, B. (2021). Generalizations of the general Lotto and Colonel Blotto games. *Economic Theory*, 1–36.
- Li, F. and Norman, P. (2018). On Bayesian persuasion with multiple senders. *Economics Letters*, 170, 66–70.
- Liu, S. and Zhu, Q. (2022). Eproach: A population vaccination game for strategic information design to enable responsible covid reopening. In *2022 American Control Conference (ACC)*, 568–573. IEEE.
- Massicot, O. and Langbort, C. (2019). Public Signals and Persuasion for Road Network Congestion Games under Vagaries. *IFAC-PapersOnLine*, 51(34), 124–130.
- Paarporn, K., Chandan, R., Alizadeh, M., and Marden, J.R. (2021a). A general lotto game with asymmetric budget uncertainty. *arXiv preprint arXiv:2106.12133*.
- Paarporn, K. and Brown, P.N. (2022). Strategically revealing capabilities in general lotto games. *arXiv preprint arXiv:2211.14907*.
- Paarporn, K., Chandan, R., Kovenock, D., Alizadeh, M., and Marden, J.R. (2021b). Strategically revealing intentions in general lotto games. *arXiv preprint arXiv:2110.12099*.
- Wu, M. and Amin, S. (2019). Information design for regulating traffic flows under uncertain network state. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 671–678. IEEE.
- Wu, M., Amin, S., and Ozdaglar, A.E. (2021). Value of Information in Bayesian Routing Games. *Operations Research*, 69(1), 148–163.
- Zhang, J. and Zhou, J. (2016). Information Disclosure in Contests: A Bayesian Persuasion Approach. *The Economic Journal*, 126(597), 2197–2217.
- Zhu, Y. and Savla, K. (2022). Information Design in Nonatomic Routing Games With Partial Participation: Computation and Properties. *IEEE Transactions on Control of Network Systems*, 9(2), 613–624.