A competitive search game with a moving target*†

Benoit Duvocelle ‡§ János Flesch¶ Mathias Staudigl‖ Dries Vermeulen **

September 22, 2021

Abstract

We introduce a discrete-time search game, in which two players compete to find an invisible object first. The object moves according to a time-varying Markov chain on finitely many states. The players are active in turns. At each period, the active player chooses a state. If the object is there then he finds the object and wins. Otherwise the object moves and the game enters the next period.

We show that this game admits a value, and for any error-term $\varepsilon > 0$, each player has a pure (subgame-perfect) ε-optimal strategy. Interestingly, a 0-optimal strategy does not always exist. We derive results on the analytic and structural properties of the value and the ε-optimal strategies. We devote special attention to the important time-homogeneous case, where additional results hold.

Keywords: Search game; two-player zero-sum game; optimal strategies; discrete time-varying Markov process.

1 Introduction

The field of search problems is one of the original disciplines of Operations Research. In the basic settings, the searcher's goal is to find a hidden object, also called the target, either with maximal probability or as soon as possible. By now, the field of search problems has developed into a wide range of models. The models in the literature differ from each other by the characteristics of the searchers and of the objects. Concerning objects, there might be one or several objects, mobile or not,

*We would like to thank Steve Alpern, Miquel Oliu-Barton, Dietmar Berwanger, Michael Eichler, Rida Laraki, Jérôme Renault, Bernhard von Stengel, Xavier Venel, Yannick Viossat, Bruno Ziliotto, and the audience of the Paris game theory seminar for their precious comments and referring us to related literature. We also thank the Associate Editor and three anonymous reviewers for reading the paper very carefully and providing many useful comments that helped us greatly improve the paper.

†Declarations of interest: none

‡Corresponding author. Address: School of Business and Economics, Quantitative Department, Tongersestraat 53, 6211 LM, Maastricht, The Netherlands

§Department of Quantitative Economics, Maastricht University. Email: benoit.duvocelle@hotmail.fr

¶Department of Quantitative Economics, Maastricht University. Email: j.flesch@maastrichtuniversity.nl

‖Department of Data Science and Knowledge Engineering, Maastricht University. Email: m.staudigl@maastrichtuniversity.nl

**Department of Quantitative Economics, Maastricht University. Email: d.vermeulen@maastrichtuniversity.nl

and they might have no aim or their aim is to not be found. Concerning the searchers, there might be one or more. When there is only one searcher, the searcher faces an optimization problem. When there is more than one searcher, searchers might be cooperative or not. If the searchers cooperate, their aim is similar to the settings with one player: they might want to minimize the expected time of search, the worst case time, or some search cost function. If the searchers do not cooperate, the problem becomes a search game with at least two strategic non-cooperative players, and hence game theoretic solution concepts and arguments will play an important role.

In almost all existing search games with more than one searcher, the searchers are assumed to cooperate in order to achieve a common goal. In real life it is often the case that the different searchers involved do not cooperate, for various reasons. For instance in nature, when several predators are looking for the same prey. Another relevant situation is when several companies have to dig to find a resource on a given surface (gold, coal, oil, lithium). The different companies do not have incentives to cooperate, and base their search on the choices that the companies have done before. It is clear in those examples that each searcher involved has to take into account the possible change of the object (a moving pray, new technologies to locate resources).

We introduce a competitive search game, played at discrete periods in $\mathbb{N}$. An object is moving according to a time-varying Markov chain on finitely many states. Two players compete to find the object first. They both know the Markov chain and the initial probability distribution of the object, but do not observe the current state of the object. Player 1 is active at odd periods, and player 2 is active at even periods. The active player chooses a state, and this choice is observed by the other player. If the object is in the chosen state, this player wins and the game ends. Otherwise, the object moves according to the Markov chain and the game continues at the next period. If the object is never found, the game lasts indefinitely. In that case, neither player wins.

When the active player chooses a state, he needs to take two opposing effects into account. First, if the object is at the chosen state, then he wins immediately. This aspect makes choosing states favorable where the object is located with a high probability. Second, if the object is not at the chosen state, then knowing this, the opponent gains information: the opponent can calculate the conditional probability distribution of the location of the object at the next period. This aspect makes choosing states favorable where, on condition that the object not being there, the induced conditional distribution at the next period disfavors the opponent. In particular, this conditional distribution should not be too informative, and for example it should not place too high a probability on a state. Clearly, in some cases there is no state that would be optimal for both scenarios at the same time, and hence the active player somehow needs to aggregate the two scenarios in order to make a choice.

Each player's goal is to maximize the probability to win the game, that is, to find the object first. In our model, we do not assume that the players take into account the period when the object is found. Of course, in most cases, maximizing the probability to win will entail at least partially that each player would prefer to find the object at earlier periods, thereby preventing the other player from finding the object. We refer to Duvocelle et al. [2020] for the finite horizon and on the discounted versions of the search game, where the period when the object is found also matters.

The two players have opposite interests, up to the event when the object is never found. More precisely, each player's preferred outcome is that he finds the object, but he is indifferent between the outcome that the other player finds the object and the outcome that the object is never found. As we will see, the possibility that neither player finds the object will only have minor role, and hence the two players have essentially opposite interests in the search game.

Main results. Our main results can be summarized as follows.

[1] We study the existence of ε-equilibria and subgame-perfect ε-equilibirum. A strategy profile is called an ε-equilibrium if neither player can increase his winning probability by more than ε with a unilateral deviation. A subgame-perfect ε-equilibrium is an ε-equilibrium in each subgame. We prove that each competitive search game admits a subgame-perfect ε-equilibrium in pure strategies, for all error-terms $\varepsilon > 0$ (cf. Theorem 3). The proof is based on topological properties of the game (cf. Appendix A). Interestingly, a 0-equilibrium does not always exist, not even in mixed strategies. This is demonstrated with an example (cf. Example 2). In the special case of time-homogeneous processes, if the Markov chain is aperiodic and irreducible, then there exists a subgame-perfect 0-equilibrium in pure strategies (cf. Theorem 4).

[2] We examine the properties of (subgame-perfect) ε-equilibria. We show that in each subgame-perfect ε-equilibrium where $\varepsilon > 0$ is small enough, the object is found with probability 1 (cf. Proposition 15), and that the set of ε-equilibrium payoffs converge to a singleton $(v_1, 1 - v_1)$, with $v_1 \in (0, 1]$ as ε vanishes (cf. Proposition 9 and Proposition 11). This implies that the two players have essentially opposite interests, and we may consider v_1 to be the value of the game (cf. Definition 10) and the strategies of ε-equilibria as ε-optimal strategies (cf. Definition 10 and Proposition 11).

[3] We investigate the properties of the value and the ε-optimal strategies. We show that the payoff functions have linear properties (cf. Proposition 6), which implies that the value is a Lipschitz continuous function with respect to the initial probability distribution of the location of the object (cf. Theorem 8). We also provide inequality properties of the value (cf. Propositions 6 and 12, and

Corollary 13).

[4] We present some geometric structures of optimal actions and optimal regions of competitive search games. An optimal action is a state among the ones necessary to choose in period 1 in order for the active player to guarantee the value. Optimal regions are the subsets of the set of probability distributions associated to each optimal action. We show that optimal regions are star-convex centered at the vertices of the unit simplex, and that the intersection of all the optimal regions is non-empty (cf. Theorem 16).

Related literature

In the literature, search games have been studied under many different assumptions. The models differ in various characteristics. For example, the number of searchers, the number of objects, the aim of the objects, and the search space. von Neumann [1953] studied a discrete version of the model where the search space is a matrix. Several variations of this game were studied by Neuts [1963], Efron [1964], Gittins and Roberts [1979], Roberts and Gittins [1978], Sakaguchi [1977], Subelman [1981], Berry and Mensch [1986], and Baston et al. [1990], among others.

The search game with an immobile hider was introduced by Isaacs [1965]. Beck and Newman [1970] considered a search game with a hider hiding on a line according to some distribution and a searcher, starting from an origin and moving at fixed speed, who tries to find the hider as soon as possible. The continuous model was then generalized by Gal [1972], Gal [1974] and Gal and Chazan [1976] who extended the state space from a line to a surface.

There is a large literature on search games on graphs with an immobile hider. Among them, Cao [1995] and Lidbetter [2020] have studied a search games on trees. Gal [1979], Reijnierse and Potters [1993], Dagan and Gal [2008] and Alpern et al. [2008] examined search games on Eulerian networks. Pavlović [1995], Gal [2010], Kikuta [2004], Alpern et al. [2008], and Alpern et al. [2009] extended the analysis to more general networks. Jotshi and Batta [2008] proposed an algorithm to find a hider hidden uniformly at random on a network. More recently Garrec and Scarsini [2020] proposed a search game in a stochastic network. They proved that the value of such games always exists, and found upper and lower bounds of the value, and optimal strategies for certain types of games. von Stengel and Werchner [1997] proved that a particular search game played on a graph is NP-hard.

More relevantly to our paper, some authors dealt with discrete search problems with a moving object. Pollock [1970], Schweitzer [1971], Dobbie [1974] and Kan [1974] study the two-state problem. Assuming perfect detection, Nakai [1973] investigates the three-state problem. Brown [1980] considers the search for a target with Markov motion in discrete time and space using an exponential detection function. He provides a necessary and sufficient condition for an optimal search plan and an efficient iterative algorithm for generating optimal plans. Washburn [1983] studies a discrete effort analogue of Brown [1980], in which searchers decide the effort they want to invest in order to find the object at each location they visit. General necessary and sufficient conditions which extend Brown's results to an arbitrary stochastic process for any mixture of discrete and continuous time and space are given in Stone [1976]. Hohzaki and Iida [2001] investigates a search problem for a moving target in which a searcher can anticipate the probabilities of routes selected by the targets but does not have any time information about when the target transits the route. Zoroa et al. [2011] study a pray-predator model in which the prey can move. They find optimal strategies for both the prey and the predator and compute the value of the game. Abramovskaya et al. [2016], Angelopoulos and Lidbetter [2020] and Delavernhe et al. [2020] are recent papers which study search games with a mobile object. We should also mention the very recent PhD thesis of Clarkson [2020], which contributes to the study of both search games with a mobile hider and to search games with an immobile object.

Most of the search games focus on the case of one searcher, or several cooperative searchers. Some problems with several cooperative searchers and one or several moving targets are mentioned in the book of Stone et al. [2016].

To the best of our knowledge, only two models consider several non-cooperative searchers. Nakai [1986] investigates a non-zero-sum game in which two searchers compete with each other for quicker detection of an object hidden in one of n boxes, with exponential detection functions. Each player wishes to maximize the probability that he detects the object before the opponent detects it. The author shows the existence of an equilibrium point of the form of a solution of simultaneous differential equations, and gets explicit solution results showing that both players have the same equilibrium strategy even though the detection rates are different. In Nakai [1990], two Searchers compete to find different objects before the other. Flesch et al. [2009] investigate the problem in which an agent has to find an object that moves between two locations according to a discrete Markov process, with the additional costless option to wait instead of searching. They find a unique optimal strategy characterized by two thresholds and show that, in a clear contrast with our model, it can never be optimal to search the location with the lower probability of containing the object. They also analyze the case of multiple agents, where the agents not only compete against time but also against each other in finding the object. They find different kinds of subgame-perfect equilibria.

As in Nakai [1973], we investigate functional and structural properties of the objective function. Nakai proved that the function that allocates to a probability distribution the average number of looks before finding the object is continuous, concave and enjoy some linear properties. They also show that the optimal regions (see Section 6) are star-convex sets. These properties have also been studied in MacPhee and Jordan [1995] and in the PhD thesis of Jordan [1997].

For an introduction to search games, we recommend the books of Alpern and Gal [2006], Gal [1979], Gal [2010] and Garnaev [2012]. We also refer to Gal [2013], and to the recent surveys Benkoski et al. [1991] and Hohzaki [2016].

Structure of the paper. In Section 2, we present the model. In Section 3, we examine the existence of (subgame-perfect) ε-equilibrium, for $\varepsilon \geq 0$. In Section 4, we argue that the two players have essentially opposite interests, we define the value and the notion of ε-optimal strategies, we present some properties of the value of the game and we show the existence of ε-optimal strategies for both players, for all $\varepsilon > 0$. In Section 6, we define optimal actions and optimal regions of the game, and we give geometric properties of these sets. The conclusion is in Section 7.

2 The Model

The Game. We study a competitive search game G played by two players. Let $\mathbb{N} = \{1, 2, 3, \ldots\}$. An object is moving according to a discrete-time Markov chain $(X_t)_{t\in\mathbb{N}}$ on a finite state space S. The initial probability distribution of the object over the set S is given by $p \in \Delta(S)$, and the transition probabilities in period t are given by an $S \times S$ transition matrix $P_t = [P_t(i,j)]_{(i,j)\in S^2}$, where $P_t(i,j)$ is the probability for the object to move from state i to state j in period t.

At each period $t \in \mathbb{N}$, one of the players is active: At odd periods player 1 is the active player, and at even periods player 2 is the active player. The active player chooses a state $s_t \in S$, which we call the action in period t. If the object is at state $X_t = s_t$, then the active player finds the object and wins the game. Otherwise, the object moves according to the transition matrix P_t and the game enters period $t+1$. We assume that each player observes the actions chosen by his opponent. The transition matrices $(P_t)_{t\in\mathbb{N}}$ and the initial distribution p are known to the players.

The aim of each player is to maximize the probability that he finds the object first.

Histories. A history in period $t \in \mathbb{N}$ is a sequence $h_t = (s_1, \ldots, s_{t-1}) \in S^{t-1}$ of past actions with the property that it has a positive probability that the object is not found before period t if the players choose their actions according to h_t. By $H_t \subseteq S^{t-1}$ we denote the set of all histories in period t. Note that H_1 consists of the empty sequence. Let $\mathbb{N}^{\text{odd}} = \{1, 3, 5, \ldots\}$ and $\mathbb{N}^{\text{even}} = \{2, 4, 6, \ldots\}$. We denote by $H^{\text{odd}} = \cup_{t\in\mathbb{N}^{\text{odd}}} H_t$ the set of histories at odd periods, and by $H^{\text{even}} = \cup_{t\in\mathbb{N}^{\text{even}}} H_t$ the set of histories at even periods.

For a distribution $q \in \Delta(S)$ for the location of the object and a state $s \in S$ for which $q(s) < 1$, let $q^{\neg s} \in \Delta(S)$ denote the distribution q conditioned on the object not being in state s. That is, $q^{\neg s}(s) = 0$ and $q^{\neg s}(s') = \frac{q(s')}{1-q(s)}$ for each state $s' \neq s$. With the help of these conditional distributions, the players can update the distribution for the current location of the object. Indeed, the initial distribution for the location of the object is p. If player 1 chooses state s_1 but he does not find the object there, then the players can update the distribution of the object in period 1 to $p^{\neg s_1}$. This implies that the distribution of the object in period 2 is $p_2 = p^{\neg s_1} P_1$, as the object moves once according to the transition matrix P_1. The update procedure continues in a similar fashion.

Strategies. The action sets for both players are $A_1 = A_2 = S$. A strategy $\sigma = (\sigma_t)_{t \in \mathbb{N}^{\text{odd}}}$ for player 1 is a sequence of functions $\sigma_t \colon H_t \to \Delta(S)$. The interpretation is that, at each period $t \in \mathbb{N}^{\text{odd}}$, given the history h_t, the strategy σ_t chooses to search state $s \in S$ with probability $\sigma_t(h_t)(s)$. Similarly, a strategy $\tau = (\tau_t)_{t \in \mathbb{N}^{\text{even}}}$ for player 2 is a sequence of functions $\tau_t \colon H_t \to \Delta(S)$.[1] We denote by Σ and $\mathcal{T}$ the set of strategies for players 1 and 2, respectively. Note that $\Sigma = \prod_{h \in H^{\text{odd}}} \Delta(S)$ and $\mathcal{T} = \prod_{h \in H^{\text{even}}} \Delta(S)$. We say that a strategy σ for player 1 is pure if it uses no randomization: for each history $h \in H_t$ with t odd, $\sigma(h)$ places probability 1 on some action $a_h \in S$. Pure strategies are defined similarly for player 2.

We endow the strategy spaces Σ and $\mathcal{T}$ with the topology of pointwise convergence. This is identical with the product topology on Σ and the product topology on $\mathcal{T}$. Under this topology, the spaces Σ and $\mathcal{T}$ are compact, and as H^{odd} and H^{even} are countable, Σ and $\mathcal{T}$ are also metrizable.

Winning probabilities. We define the stopping time[2] of the game by $\Theta = \min\{t \in \mathbb{N} \mid s_t = X_t\}$. Consider a strategy profile (σ, τ). The probability under (σ, τ) that player 1 wins is denoted by $u_1(\sigma, \tau) = \mathbb{P}_{\sigma,\tau}\left(\Theta \in \mathbb{N}^{\text{odd}}\right)$, and that player 2 wins is denoted by $u_2(\sigma, \tau) = \mathbb{P}_{\sigma,\tau}\left(\Theta \in \mathbb{N}^{\text{even}}\right)$. Note that $u_1(\sigma, \tau) + u_2(\sigma, \tau) = 1 - \mathbb{P}_{\sigma,\tau}(\Theta = \infty)$.

If the play is in period t and the current distribution of the object is $q \in \Delta(S)$, we will write $u_1^t(\sigma, \tau)(q)$ and $u_2^t(\sigma, \tau)(q)$ for the continuation winning probabilities of the players from period t onwards.

Similarly, if the play is in period t and the history is h, we will use the notations $u_1(\sigma, \tau)(h) = u_1^t(\sigma, \tau)(p_h)$ and $u_2(\sigma, \tau)(h) = u_2^t(\sigma, \tau)(p_h)$, where p_h is the conditional probability distribution for the location of the object at history h.

Subgame-perfect ε-equilibrium. Let $\varepsilon \geq 0$ be an error-term. A strategy σ for player 1 is an ε-best response against strategy τ for player 2 if $u_1(\sigma, \tau) \geq u_1(\sigma', \tau) - \varepsilon$ for every strategy σ' of

[1] A strategy (for either player) could also be defined as a function $\rho : \Delta(S) \to \Delta(S)$, with the interpretation that for each distribution $p \in \Delta(S)$ for the location of the object, the player should choose a state according to $\rho(p)$.

[2] With the convention that $\min\{\emptyset\} = +\infty$.

player 1. Similarly, a strategy τ for player 2 is an ε-best response against strategy σ for player 1 if $u_2(\sigma,\tau) \geq u_2(\sigma,\tau') - \varepsilon$ for every strategy τ' of player 2. A strategy profile (σ,τ) is called an ε-equilibrium if σ is an ε-best response against τ and τ is an ε-best response against σ. A subgame-perfect ε-equilibrium is a strategy profile (σ,τ) which is an ε-equilibrium in each subgame. That is, for each history h, for each $\sigma' \in \Sigma$, for each $\tau' \in \mathcal{T}$, $u_1(\sigma,\tau)(h) \geq u_1(\sigma',\tau)(h) - \varepsilon$ and $u_2(\sigma,\tau)(h) \geq u_2(\sigma,\tau')(h) - \varepsilon$.

An alternative interpretation of the game. We call the previous game Model [1], which is with imperfect information as the players do not observe the location of the object, unless a player searches the location where the object is staying. Now we present an alternative model of this game in perfect information, which we call Model [2]. We use Model [2] only as an auxiliary model in our paper, so that we can apply existence results in the literature that are formulated for this type of models (cf. Theorem 3).

Model [2]: Another way to describe our game is as follows. One could imagine that the game consists of two phases. In the first phase the players choose actions sequentially. More precisely, in the first phase, player 1 chooses an action at odd periods and player 2 chooses an action at even periods sequentially, just as before. This results in an infinite sequence of states $(s_1, s_2, \ldots)$. The set of infinite histories is S^∞. Every pure strategy profile (σ,τ) induces a unique infinite history $h^\infty_{\sigma,\tau} \in S^\infty$. In a second phase, players receive a payoff. Now, for $i = 1,2$, consider the payoff function $f_i : S^\infty \to [0,1]$ defined as follows. Consider an infinite history $(s_1, s_2, \ldots)$. Take any pure strategy profile (σ,τ) such that $h^\infty_{\sigma,\tau} = (s_1, s_2, \ldots)$ and define $f_i(s_1, s_2, \ldots) = u_i(\sigma,\tau)$; note that this definition only depends on the realized history. The goal of each player is to maximize his payoff. Note that this is a game without an object. This way we obtain a two-player perfect information game.

Discussion. We briefly argue that the above descriptions are equivalent. For each pure strategy profile (σ,τ), for each player $i = 1,2$, we have $u_i(\sigma,\tau) = f_i(h^\infty_{\sigma,\tau})$. Then, a strategy profile in one of the models leads to the same payoff in the other game. The difference is that Model [1] is in imperfect information, as players only know the probability distribution of the object, while Model [2] is in perfect information.

Model [1] gives a very clear, intuitive and concrete description of the game. This is the reason why we usually work with this model in the paper. Model [2] is used as a tool to prove existence of subgame-perfect ε-equilibrium for all $\varepsilon > 0$, as in Theorem 3.

3 Existence of equilibrium

In this section, we examine equilibria in competitive search games. In the first subsection, we show that there are search games for which there exist no 0-equilibrium, not even in mixed strategies. From a technical point of view, this is caused by discontinuity in the payoff functions of the players. In the second subsection, we focus on the notion of subgame-perfect ε-equilibrium, where $\varepsilon > 0$ is an error-term, and prove that each search game admits a subgame-perfect ε-equilibrium in pure strategies, for all $\varepsilon > 0$. In the third subsection, we present sufficient conditions for the existence of a subgame-perfect 0-equilibrium in pure strategies.

3.1 Search games with no 0-equilibrium

Theorem 1. *There exist time-homogeneous competitive search games which admit no 0-equilibrium, not even in mixed strategies.*

We provide an example below of a time-homogeneous competitive search game which admits no 0-equilibrium, not even in mixed strategies (for another example, we refer to Duvocelle et al. [2020]). The main idea of this example is that during the game the active player is incentivised to choose a transient state of the Markov chain, as these are the only states where moving first gives an advantage. However, if the players always choose a transient state, then they do not find the object with probability 1, and hence this cannot constitute a 0-equilibrium.

Example 2. Consider the game in Figure 1. In this game, $\eta \in (0, \frac{1}{4})$ and the initial probability distribution is $p = (q, q, \frac{1}{2} - q, \frac{1}{2} - q)$, where $q \in (0, \frac{1}{4})$. Notice that states 1 and 2 have the same transition probabilities, and so do states 3 and 4. States 1 and 2 are transient, whereas states 3 and 4 are absorbing.

We show that this game admits no 0-equilibrium. The intuition for this claim is as follows. Consider period 1. Player 1 has a choice between a transient state (i.e., states 1 and 2) and an absorbing state (i.e., states 3 and 4).

Suppose first that player 1 chooses an absorbing state, say state 3. We will show (see Claim 1 below) that this, with optimal follow-up play, induces a winning probability of exactly $\frac{1}{2}$. If player 1 does not find the object in period 1, then it is very likely that the object is in state 4, and player 2 should respond by choosing state 4 in period 2. If player 2 does not find the object, then due to the transition probabilities, the object is very likely to be in state 3 in period 3. Indeed, we know that the object was not in state 3 in period 1 and not in state 4 in period 2, so the object has one more period to move to state 3 compared to state 4. Continuing this argument shows that the optimal continuation

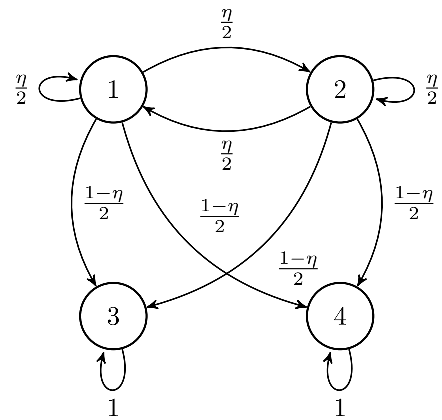

Figure 1: A game without 0-equilibrium.

play consists of player 1 always choosing state 3 and player 2 always choosing state 4. This gives a winning chance of exactly $\frac{1}{2}$ to each player.

Now suppose that player 1 chooses a transient state, say state 1. In addition, suppose that if period 3 is reached, then player 1 will choose state 3, unless player 2 chose state 3 in period 2, in which case player 1 chooses state 4. We show that this strategy guarantees a winning probability of strictly more than $\frac{1}{2}$ for player 1 (see Claim 2 below). Indeed, player 1 wins immediately with probability $q > 0$, and as we show, the probability of winning in period 3 is greater than $\frac{1}{2} - q$, so that the total probability that player 1 wins is strictly more than 1/2.

Based on the above discussion, player 1 should choose a transient state in period 1. If player 1 does not find the object, by an inductive argument, player 2 should also choose a transient state in period 2. Continuing this way, the players should choose transient states in all periods. Then, however, they do not find the object with probability 1, and hence this cannot constitute a 0-equilibrium (see Claim 3 below).

CLAIM 1: If player 1 starts with choosing state 3 or state 4, then with optimal follow-up play, the winning probability of player 1 is exactly $\frac{1}{2}$. More precisely, let $\Sigma^{3,4}$ denote the set of strategies for player 1 that looks at state 3 or state 4 in period 1. Then,

$$\sup_{\sigma\in\Sigma^{3,4}} \inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau) = \frac{1}{2}.$$

PROOF OF CLAIM 1: Let $\tau = (\tau_t)_{t\in\mathbb{N}^{\text{even}}}$ be the strategy of player 2 defined as follows. For all

$t \in \mathbb{N}^{\text{even}}$, for all $h_t \in H_t$,

$$\tau_t(h_t) = \begin{cases} \text{state 1} & & \text{if } h_t \in \{1,2\}^{t-1}, \\ \text{state 3} & & \text{if } h_t(t-1) = 4, \\ & \text{or} & \text{if } h_t(t-1) \in \{1,2\} \text{ and } h_t(t-2) = 4, \\ \text{state 4} & & \text{if } h_t(t-1) = 3, \\ & \text{or} & \text{if } h_t(t-1) \in \{1,2\} \text{ and } h_t(t-2) = 3, \\ \text{arbitrary} & & \text{otherwise,} \end{cases}$$

where $h_t(t-2)$ and $h_t(t-1)$ are the second-to-last and the last actions chosen under history h_t, respectively.

The idea is that τ looks at state 1 if player 1 has never played state 3 or state 4, and plays the most likely state otherwise. Assume for simplicity that player 1 looks at state 3 in period 1. Assume that player 1 does not find the object in period 1. The conditional probability of the object being in state 4 in period 2 is then equal to

$$\frac{\frac{1}{2} - q}{\frac{1}{2} + q} + 2 \cdot \frac{1-\eta}{2} \cdot \frac{q}{\frac{1}{2} + q} = \frac{\frac{1}{2} - q \cdot \eta}{\frac{1}{2} + q},$$

which is strictly higher than $\frac{1}{2}$ by our assumption that $q < \frac{1}{4}$ and $\eta < \frac{1}{4}$. Then, in the continuation of the game, player 2 guarantees strictly more than $\frac{1}{2}$ if he looks at state 4 in period 2. If he does not, player 2 will get strictly less than $\frac{1}{2}$ if player 1 looks at state 4 in period 3. For similar reasons, if period 3 is reached, it is better for player 1 to look at state 3. By repeating this argument, it is better for player 1 to always look at state 3 against τ.

With probability $\frac{1}{2} - q$, player 1 finds the object in period 1. With probability $\frac{1}{2} - q + q \cdot (1-\eta)$, player 2 finds the object in period 2. With probability $q \cdot (1-\eta) + q \cdot (1-\eta) \cdot \eta$, player 1 finds the object in period 3. With probability $q \cdot (1-\eta) \cdot \eta + q \cdot (1-\eta) \cdot \eta^2$, player 2 finds the object in period 4. And so on. Then, player 1 finds the object with probability

$$\frac{1}{2} - q + q \cdot (1-\eta) + q \cdot (1-\eta) \cdot \eta + q \cdot (1-\eta) \cdot \eta^2 + q \cdot (1-\eta) \cdot \eta^3 + \ldots = \frac{1}{2} - q + q \cdot (1-\eta) \cdot \frac{1}{1-\eta} = \frac{1}{2}.$$

So, by playing state 3 or state 4 in period 1, player 1 gets exactly $\frac{1}{2}$ against τ. $\square$

Let $\sigma = (\sigma_t)_{t \in \mathbb{N}^{\text{odd}}}$ be the strategy of player 1 defined as follows. For all $t \in \mathbb{N}^{\text{odd}}$, for all $h_t \in H_t$,

$$\sigma_t(h_t) = \begin{cases} \text{state 1} & & \text{if } t = 1, \\ \text{state 3} & & \text{if } t \geq 3 \text{ and } h_t(t-1) = 4, \\ & \text{or} & \text{if } t \geq 3 \text{ and } h_t(t-1) \in \{1,2\} \text{ and } h_t(t-2) = 4, \\ \text{state 4} & & \text{if } t \geq 3 \text{ and } h_t(t-1) = 3, \\ & \text{or} & \text{if } t \geq 3 \text{ and } h_t(t-1) \in \{1,2\} \text{ and } h_t(t-2) = 3, \\ \text{arbitrary} & & \text{otherwise.} \end{cases}$$

The idea is that from period 3 onward, σ chooses the most likely location of the object.

CLAIM 2: When player 1 uses σ he guarantees himself strictly more than $\frac{1}{2}$: $u_1(\sigma, \tau) > \frac{1}{2}$ for every τ.

PROOF OF CLAIM 2: Under σ, player 1 looks at state 1 in period 1 and finds the object with probability q in period 1. The object is not found in period 1 with probability $1-q$. In that case, with probability $\frac{q}{1-q}$ the object is in state 2 in period 1, and it moves with probability $\eta/2$ to state 1 and with probability $1-\eta$ to state 3 or state 4 and in period 2. So the distribution of the location of the object at the beginning of period 2 is $p_2 = (\frac{q\cdot\eta}{2(1-q)}, \frac{q\cdot\eta}{2(1-q)}, \frac{1}{2} - \frac{q\cdot\eta}{2(1-q)}, \frac{1}{2} - \frac{q\cdot\eta}{2(1-q)})$.

Suppose player 2 chooses state 1 or state 2 in period 2. With probability $\frac{q\cdot\eta}{2}$, player 2 finds the object in period 2, and with probability $\frac{1}{2} - q + \frac{q\cdot(1-\eta)}{2} + \frac{q\cdot\eta\cdot(1-\eta)}{4}$, player 1 finds the object in period 3. So in total, player 1 finds the object over periods 1 and 3 with probability $q + \frac{1}{2} - q + \frac{q\cdot(1-\eta)}{2} + \frac{q\cdot\eta\cdot(1-\eta)}{4} > \frac{1}{2}$.

Suppose player 2 chooses state 3 (resp. state 4) in period 2. With probability $\frac{1}{2} - q + \frac{q\cdot(1-\eta)}{2}$, player 2 finds the object in period 2, and with probability $\frac{1}{2} - q + \frac{q\cdot(1-\eta)}{2} + \frac{q\cdot\eta\cdot(1-\eta)}{2}$, player 1 finds the object in period 3. So in total, player 1 finds the object over periods 1 and 3 with probability $q + \frac{1}{2} - q + \frac{q\cdot(1-\eta)}{2} + \frac{q\cdot\eta\cdot(1-\eta)}{2} > \frac{1}{2}$.

Thus, in every case, player 1 guarantees strictly more than $\frac{1}{2}$ by using strategy σ. $\square$

CLAIM 3: There is no 0-equilibrium.

PROOF OF CLAIM 3: Assume by way of contradiction that there is a 0-equilibrium (σ', τ'). From CLAIM 1 and CLAIM 2, player 1 chooses state 1 or state 2 with probability 1 in period 1. In both cases, in period 2 the current probability distribution is $p_2 = (\frac{q\cdot\eta}{2(1-q)}, \frac{q\cdot\eta}{2(1-q)}, \frac{1}{2} - \frac{q\cdot\eta}{2(1-q)}, \frac{1}{2} - \frac{q\cdot\eta}{2(1-q)})$. Then, in period 2, the game is similar to the original one, with a parameter $q' = \frac{q\cdot\eta}{2(1-q)}$ instead of q, which still satisfies $q' \in (0, \frac{1}{4})$, and where the roles of the players are exchanged. Then, as τ is a 0-best response, it follows from the previous reasoning that player 2 plays state 1 or state 2 with probability 1. By following this process recursively, players will choose states 1 and 2 with probability 1 forever. This leads to the payoff $\frac{4\cdot q}{4-\eta^2}$ for player 1. Then, player 1 has an incentive to deviate from σ' and to choose state 3 in period 1 to get a payoff of at least $\frac{1}{2} - q > \frac{4\cdot q}{4-\eta^2}$, a contradiction. $\square$

3.2 Existence of pure subgame-perfect ε-equilibrium

In this subsection we are interested in the existence of subgame-perfect ε-equilibrium, where $\varepsilon > 0$. In the next theorem, we show that all competitive search games admit a subgame-perfect ε-equilibrium in pure strategies, for each $\varepsilon > 0$. The proof relies on existence results for subgame-perfect ε-equilibria in games with bounded and lower semi-continuous payoff functions (see Flesch et al. [2010] orFlesch and Predtetchinski [2016]).

Theorem 3. *Every competitive search game admits a pure subgame-perfect ε-equilibrium, for each $\varepsilon > 0$.*

Proof. Consider Model [2] of a competitive search game, as described in Section 2. Note that: (i) this is a multiplayer perfect information game, and (ii) the payoffs are bounded and lower semi-continuous, in view of Proposition 18. Thus, by applying Theorem 2.3 of Flesch et al. [2010] (or by Theorem 4.1 of Flesch and Predtetchinski [2016]), the game admits a pure subgame-perfect ε-equilibrium for each $\varepsilon > 0$. □

Revisiting Example 2. In view of Theorem 3, the game in Example 2 has a subgame-perfect ε-equilibrium in pure strategies, for each $\varepsilon > 0$. We now construct such a strategy profile, for each $\varepsilon > 0$.

Let $\varepsilon > 0$. The idea of the subgame-perfect ε-equilibrium described here is as follows. Take a sufficiently large period n. Let the players choose state 1 until period n, and then let them choose the most likely between state 3 or state 4 in the remaining game. If a player deviates to an absorbing state before period n, then also continue with states 3 and 4 as described before.

More formally, for each $n \in \mathbb{N}$, let (σ^n, τ^n) be the pure strategy profile defined as follows. For all $t \in \mathbb{N}$, for all history h_t in period t, for all $n \in \mathbb{N}$, let $f_t^n : H_t \to S$ be defined by

$$f_t^n(h_t) = \begin{cases} \text{state } 1 & & \text{if } h_t \in \{1,2\}^{t-1} \text{ and } t < n, \\ \text{state } 3 & & \text{if } h_t(t-1) = 4, \\ & \text{or} & \text{if } h_t(t-1) \in \{1,2\} \text{ and } h_t(t-2) = 4, \\ & \text{or} & \text{if } h_t \in \{1,2\}^{t-1} \text{ and } t \geq n, \\ \text{state } 4 & & \text{if } h_t(t-1) = 3, \\ & \text{or} & \text{if } h_t(t-1) \in \{1,2\} \text{ and } h_t(t-2) = 3, \\ \text{arbitrary} & & \text{otherwise.} \end{cases}$$

Then, we define $\sigma_t^n(h_t) = f_t^n(h_t)$ for all $t \in \mathbb{N}^{\text{odd}}$, and $\tau_t^n(h_t) = f_t^n(h_t)$ for all $t \in \mathbb{N}^{\text{even}}$ and all history h_t at time t. The idea of σ^n and τ^n is to look at state 1 until period n (if the other player does the same) and from period n onward (or before if the other player deviates) to look at the most likely state. We argue that if $n \geq \frac{\ln\left(\frac{q}{\eta\varepsilon}\right)}{\ln\left(\frac{2}{\eta}\right)}$, the difference between those two expressions is smaller than ε so (σ^n, τ^n) is a subgame-perfect ε-equilibrium. For simplicity, we assume that n is odd.

It follows from Claim 1 and Claim 2 of the proof of Theorem 1 that τ^n is a 0-best-response against σ^n. It is then sufficient to show that σ^n is an ε-best response against τ^n when n is large enough. From Claim 2 of the proof of Theorem 1 it follows that a 0-best response against τ^n is to follow the strategy σ^{n+1}, which only differs from σ^n in period n. Under (σ^n, τ^n), with probability q player 1 finds the object in period 1, with probability $q.\left(\frac{\eta}{2}\right)$ player 2 finds the object in period 2, with probability $q.\left(\frac{\eta}{2}\right)^2$ player 1 finds the object in period 3, and so on until period $n-1$. Then in the continuation game that starts in period n it follows from the proof of Claim 2 in Theorem 1 that both players find the object with probability $\frac{1}{2}$. So, with probability $q + q \cdot \left(\frac{\eta}{2}\right)^2 + q \cdot \left(\frac{\eta}{2}\right)^4 + \ldots + q \cdot \left(\frac{\eta}{2}\right)^{n-3}$ player 1 finds the object before period n, with probability $q \cdot \left(\frac{\eta}{2}\right) + \ldots + q \cdot \left(\frac{\eta}{2}\right)^{n-2}$ player 2 finds the object before period n, and with probability $\frac{1}{2}\cdot\left[1-\left(q+q\cdot\left(\frac{\eta}{2}\right)+q\cdot\left(\frac{\eta}{2}\right)^2+\ldots+q\cdot\left(\frac{\eta}{2}\right)^{n-2}\right)\right]$ each player finds the object from period n onward. This implies that under (σ^n,τ^n) the expected payoff of player 1 is

$$q+q\cdot\left(\frac{\eta}{2}\right)^2+q\cdot\left(\frac{\eta}{2}\right)^4+\ldots+q\cdot\left(\frac{\eta}{2}\right)^{n-3}+\left[1-\left(q+q\cdot\left(\frac{\eta}{2}\right)+q\cdot\left(\frac{\eta}{2}\right)^2+\ldots+q\cdot\left(\frac{\eta}{2}\right)^{n-2}\right)\right]\cdot\frac{1}{2}$$

and under $(\sigma^{n+1},\tau^{n+1})$, the expected payoff of player 1 is

$$q+q\cdot\left(\frac{\eta}{2}\right)^2+q\cdot\left(\frac{\eta}{2}\right)^4+\ldots+q\cdot\left(\frac{\eta}{2}\right)^{n-1}+\left[1-\left(q+q\cdot\left(\frac{\eta}{2}\right)+q\cdot\left(\frac{\eta}{2}\right)^2+\ldots+q\cdot\left(\frac{\eta}{2}\right)^{n-1}\right)\right]\cdot\frac{1}{2}$$

Those two terms converge to the same limit $\frac{q}{1-\left(\frac{\eta}{2}\right)^2}+\left[1-\frac{q}{1-\left(\frac{\eta}{2}\right)}\right]\cdot\frac{1}{2}$ which is the value of the game. Moreover, the difference between these two expressions is

$$\left|q\cdot\left(\frac{\eta}{2}\right)^{n-1}-\frac{1}{2}\cdot q\cdot\left(\frac{\eta}{2}\right)^{n-1}\right|=\frac{q}{2}\cdot\left(\frac{\eta}{2}\right)^{n-1}.$$

Hence, when $n\geq\frac{\ln\left(\frac{q}{\eta\varepsilon}\right)}{\ln\left(\frac{2}{\eta}\right)}$, the difference between those two expressions is smaller than ε so (σ^n,τ^n) is an ε-equilibrium. □

3.3 Sufficient conditions for existence of pure subgame-perfect 0-equilibrium

In this subsection, we present sufficient conditions for the existence of a pure subgame-perfect 0-equilibrium.

Consider a time-homogeneous competitive search game. In this case the transition matrix at each period is the same matrix P. For all $r\in\mathbb{N}$ we denote by P^r the matrix P to the power r. We call P *irreducible* if for each entry (i,j), there exists $r\in\mathbb{N}$ such that the entry (i,j) of P^r is positive. We call P *periodic* of period $r\geq 2$ if $r=\gcd\{n\geq 2\,|\,P^n(x,x)>0\}$. When $r=1$, we say that P is *aperiodic*. A subset $S'\subseteq S$ is *ergodic* if for $(i,j)\in S'\times(S\backslash S')$, $P(i,j)=0$ and the transition matrix P restricted to the set S' is irreducible.

A probability distribution $\pi\in\Delta(S)$ over the set S is called a *stationary distribution* for the transition matrix P if $\pi P=\pi$.

It is known (see Levin and Peres [2017], Corollary 1.17 page 13 and Theorem 4.9 page 52) that if the transition matrix P is irreducible, then there exists a unique stationary distribution $\pi\in\Delta(S)$ and $\pi(s)>0$ for all $s\in S$. If P is also aperiodic, then there exist constants $\beta\in(0,1)$ and $c>0$ such that for all distributions p for the initial location of the object and all periods $t\in\mathbb{N}$,

$$||pP^t-\pi||_{TV}\leq c\cdot\beta^t,$$

where $||p-q||_{TV} = \max\limits_{A \subset S} \sum_{s \in A}(p(s)-q(s))$ is the total variation distance over $\Delta(S)$. Note that for each $p, q \in \Delta(S)$, it holds that $||p-q||_{TV} = \frac{1}{2} \cdot \sum_{s \in S} |p(s)-q(s)|$ (see Levin and Peres [2017], pages 47-48).

Theorem 4. *Consider a time-homogeneous competitive search game. Assume that the transition matrix P is irreducible and aperiodic. Then, no matter the initial probability distribution p, the object is found with probability 1 under every strategy profile (σ, τ), i.e., $\mathbb{P}_{\sigma,\tau}(\Theta < \infty) = 1$. Hence, the payoff functions are continuous in this game, and there exists a subgame-perfect 0-equilibrium in pure strategies.*

Proof. As mentioned, the transition matrix P has a unique stationary distribution $\pi \in \Delta(S)$ and $\pi(s) > 0$ for all $s \in S$. Moreover, there exist constants $c > 0$ and $\beta \in (0,1)$ such that $|pP^t(s) - \pi(s)| \leq c \cdot \beta^t$ for all $t \in \mathbb{N}$, for all $s \in S$ and for all $p \in \Delta(S)$. Hence, there exists $t^* \in \mathbb{N}$ with the following property: for all $p \in \Delta(S)$, for all $s \in S$, for all $t \geq t^*$, we have $(pP^{t-1})(s) > \frac{\delta}{2}$, where $\delta = \min_{s \in S} \pi(s)$. Without loss of generality we can assume that $t^* \geq 2$.

Let $\alpha = \frac{\delta}{4(t^*-1)}$. The proof is divided into four steps.

STEP 1: Let (σ, τ) be a pure strategy profile, and let $(s_t)_{t \in \mathbb{N}}$ denote the induced sequence of actions. We show that the object is found during the first t^* periods with probability at least α.

PROOF: For each $t \in \mathbb{N}$, let $p_t = (p_t(s))_{s \in S} \in \Delta(S)$ denote the probability distribution of the location of the object in period t, conditional on not being found through the history $(s_1, \ldots, s_{t-1})$.

If there is a period $t \leq t^*$ such that $p_t(s_t) \geq \alpha$, then under (σ, τ), the object is found in period t with probability at least α, if it has not been found before. Hence, the claim of step 1 is true.

Therefore, it suffices to show that if at each period $t \leq t^* - 1$ we have $p_t(s_t) < \alpha$, then $p_{t^*}(s_{t^*}) \geq \alpha$. So assume that at each period $t \leq t^* - 1$ we have $p_t(s_t) < \alpha$. The idea of the calculation below is that, since the object is found with low probabilities at the first $t^* - 1$ periods, the probability distribution for the object in period t^* conditioned on not being found during the first $t^* - 1$ periods is almost the same as the unconditioned probability distribution. That is, p_{t^*} is close to pP^{t^*-1}, which is in turn close to the stationary distribution π.

Note that, if the players do not condition on the past, the probability distribution of the location of the object in period t^* is simply pP^{t^*-1}. Recall that, if $p(s) < 1$, then $p^{\neg s}$ denotes the probability distribution p conditioned on the object not being in state s. We have

$$\begin{aligned}
||p_{t^*} - pP^{t^*-1}||_{TV} &\leq ||p_{t^*} - p_{t^*-1}P||_{TV} + ||p_{t^*-1}P - pP^{t^*-1}||_{TV} \\
&= ||p_{t^*-1}^{\neg s_{t^*-1}}P - p_{t^*-1}P||_{TV} + ||p_{t^*-1}P - pP^{t^*-1}||_{TV} \\
&\leq ||p_{t^*-1}^{\neg s_{t^*-1}} - p_{t^*-1}||_{TV} + ||p_{t^*-1} - pP^{t^*-2}||_{TV} \\
&= p_{t^*-1}(s_{t^*-1}) + ||p_{t^*-1} - pP^{t^*-2}||_{TV} \\
&< \alpha + ||p_{t^*-1} - pP^{t^*-2}||_{TV} \\
&< \alpha \cdot (t^* - 1) + ||p_1 - pP^0||_{TV} \\
&= \alpha \cdot (t^* - 1) \\
&= \frac{\delta}{4}.
\end{aligned}$$

Here, in the first inequality we used the triangle inequality. In the first equality we used $p_{t^*} = p_{t^*-1}^{\neg s_{t^*-1}}P$, as p_{t^*} is the location of the object in period t^* conditioned on not being found up to period $t^* - 1$ and $p_{t^*-1}^{\neg s_{t^*-1}}P$ expresses the distribution that arises when object moves once according to P after not being found up to period $t^* - 1$. The second inequality is true as $||qP - q'P||_{TV} \leq ||q - q'||_{TV}$ for all $q, q' \in \Delta(S)$. The second equality follows from the above interpretation of $p_{t^*-1}^{\neg s_{t^*-1}}$ and the definition of the total variation norm. The third inequality is due to the assumption that at each period $t \leq t^* - 1$ we have $p_t(s_t) < \alpha$. The fourth inequality then follows by induction. The last two equalities are due to $p_1 = p$ and the choice of α.

Therefore,

$$p_{t^*}(s_{t^*}) \;\geq\; (pP^{t^*-1})(s_{t^*}) - ||p_{t^*} - pP^{t^*-1}||_{TV} \;\geq\; \frac{\delta}{2} - \frac{\delta}{4} \;=\; \frac{\delta}{4} \;\geq\; \alpha.$$

This completes the proof of Step 1.

STEP 2: Consider any strategy profile (σ, τ). We show that the object is found during the first t^* periods with probability at least α.

PROOF: On the finite horizon t^*, each strategy can be equivalently represented as a mixed strategy, i.e. a probability distribution on the finite set of pure strategies on horizon t^* (see for example Maschler et al. [2013]). Hence, Step 2 follows from Step 1.

STEP 3: Consider any strategy profile (σ, τ). We show that the object is found with probability 1 under (σ, τ). By Proposition 18, this will imply that the payoff functions are continuous in this game.

PROOF: By Step 2, the object is found during the first t^* periods with probability at least α. Since t^* and therefore α do not depend on the initial distribution of the object, if the object is not found in the first t^* periods, then it will be found between periods $t^* + 1$ and $2t^*$ with probability at least α. By repeating this argument, the object is found with probability 1 under (σ, τ).

STEP 4: We show that there exists a subgame-perfect 0-equilibrium in pure strategies.

PROOF:[3] In view of Theorem 3, for each $n \in \mathbb{N}$, there exists a subgame-perfect $\frac{1}{n}$-equilibrium (σ^n, τ^n) in pure strategies. Since the spaces of strategies Σ and $\mathcal{T}$ are compact and metrizable, by taking a subsequence if necessary, we can assume that the sequence $(\sigma^n, \tau^n)_{n\in\mathbb{N}}$ converges to a strategy profile (σ, τ) in pure strategies as $n \to \infty$.

For each $n \in \mathbb{N}$, and for every history h we have $u_1(\sigma^n, \tau^n)(h) \geq u_1(\sigma', \tau^n)(h) - \frac{1}{n}$ and $u_2(\sigma^n, \tau^n)(h) \geq u_2(\sigma^n, \tau')(h) - \frac{1}{n}$ for all $\sigma' \in \Sigma$ and $\tau' \in \mathcal{T}$. Since by Step 3 the payoff functions u_1 and u_2 are continuous, by taking the limits as $n \to \infty$, we obtain $u_1(\sigma, \tau)(h) \geq u_1(\sigma', \tau)(h)$ and $u_2(\sigma, \tau)(h) \geq u_2(\sigma, \tau')(h)$ for every history h, for all $\sigma' \in \Sigma$ and $\tau' \in \mathcal{T}$. Hence, (σ, τ) is a subgame-perfect 0-equilibrium in pure strategies. $\square$

4 Existence and properties of the value

In competitive search games, the winning probabilities do not always add up to 1, as under certain strategy profiles it may have a positive probability that the object is never found. For instance, this is the case in the game in Example 2, if the players always choose state 1. However, neither player is interested in the outcome when the object is never found. Therefore, as we argue in this section, in essence the players have opposite interest and the value is a natural solution concept for competitive search games.

We denote $v_1 = \sup\limits_{\sigma\in\Sigma} \inf\limits_{\tau\in\mathcal{T}} u_1(\sigma, \tau)$ and $v_2 = \sup\limits_{\tau\in\mathcal{T}} \inf\limits_{\sigma\in\Sigma} u_2(\sigma, \tau)$. Intuitively, a strategy σ guarantees for player 1 a winning probability of $\inf_{\tau\in\mathcal{T}} u_1(\sigma, \tau)$, and therefore v_1 can be interpreted as the largest winning probability that player 1 can guarantee in the game. The interpretation of v_2 is similar.

More generally, when we wish to emphasize the initial distribution p of the object as a parameter, we will write $v_1(p) = \sup\limits_{\sigma\in\Sigma} \inf\limits_{\tau\in\mathcal{T}} u_1(\sigma, \tau)(p)$ and $v_2(p) = \sup\limits_{\tau\in\mathcal{T}} \inf\limits_{\sigma\in\Sigma} u_2(\sigma, \tau)(p)$.

Let Σ^s denote the set of those strategies for player 1 that choose state s in period 1. We define $v_1(p, s) = \sup\limits_{\sigma\in\Sigma^s} \inf\limits_{\tau\in\mathcal{T}} u_1(\sigma, \tau)(p)$. Note that $v_1(p) = \max_{s\in S} v_1(p, s)$.

As discussed in Section 2, each history h induces a conditional probability distribution for the location of the object, say p_h. We will use the notation $v_1(h) = v_1(p_h)$ and $v_2(h) = v_2(p_h)$.

Proposition 5. *We have*

$$v_1 = \inf_{\tau\in\mathcal{T}} \sup_{\sigma\in\Sigma} u_1(\sigma, \tau) \quad \textit{and} \quad v_2 = \inf_{\sigma\in\Sigma} \sup_{\tau\in\mathcal{T}} u_2(\sigma, \tau).$$

[3] The claim of Step 4 follows from Fudenberg and Levine [1983], but for completeness, we give a proof based on Theorem 3.

Similar equalities hold for $v_1(p)$ and $v_2(p)$ for all $p \in \Delta(S)$.

Proof. We refer to the Appendix B. □

The next theorem discusses properties of the payoff functions and the functions $v_1(p)$ and $v_2(p)$ along line segments in $\Delta(S)$. For each state $s \in S$, let $e^s \in \Delta(S)$ denote the probability distribution that allocates probability 1 to state s and probability 0 to all other states.

Proposition 6.

[1] *Let (σ,τ) be a strategy profile. Then, the payoff functions are linear in the initial probability distribution of the object: for every $p, q \in \Delta(S)$, for every $\lambda \in [0,1]$, and for every player $i = 1,2$:*

$$u_i(\sigma,\tau)(\lambda \cdot p + (1-\lambda)\cdot q) \;=\; \lambda \cdot u_i(\sigma,\tau)(p) + (1-\lambda)\cdot u_i(\sigma,\tau)(q). \tag{1}$$

[2] *For every $s \in S$, the map $p \mapsto v_1(p,s)$ is linear over every line passing through e^s: for every $p \in \Delta(S)$ and for every $\lambda \in [0,1]$:*

$$v_1(\lambda \cdot e^s + (1-\lambda)\cdot p, s) \;=\; \lambda + (1-\lambda)\cdot v_1(p,s).$$

[3] *For every $p, q \in \Delta(S)$, for every $\lambda \in [0,1]$, and for every player $i = 1,2$:*

$$v_i(\lambda \cdot p + (1-\lambda)\cdot q) \;\geq\; \lambda \cdot v_i(p).$$

Proof. We refer to the Appendix B. □

Remark 7. In part [2] of Proposition 6, the linearity of $p \mapsto v_1(p,s)$ relies on the following fact: Let $p \in \Delta(S) \setminus \{e^s\}$, and let ℓ denote the line going through p and e^s. If by choosing state s player 1 does not find the object, then the conditional distribution of the location of the object, $p^{\neg s}$, stays on the line ℓ. Note, however, that the function $p \mapsto v_1(p,s)$ is generally non-linear. Indeed, consider a game that has 4 states, each of which is absorbing. For $p = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ and $q = (\frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3})$, we have $v_1(p,1) = \frac{2}{3}$ and $v_1(q,1) = \frac{2}{3}$, but $v_1(\frac{1}{2}\cdot p + \frac{1}{2}\cdot q, 1) = \frac{1}{2}$. $\Diamond$

We recall the definition of the total variation distance: for $p, q \in \Delta(S)$, the total variation distance between p and q is the non-negative number $||p-q||_{TV} = \max_{A \subset S} \sum_{s \in A}(p(s) - q(s))$.

Theorem 8. *For each player $i = 1,2$, each strategy profile (σ,τ), each state $s \in S$, the functions $p \mapsto u_i(\sigma,\tau)(p)$, $p \mapsto v_i(p)$, and $p \mapsto v_1(p,s)$ are 1-Lipschitz continuous with respect to the total variation distance.*

Proof. We only prove it for player 1. By part [1] of Proposition 6, we have

$$u_1(\sigma,\tau)(p) = \sum_{s\in S} p(s)\cdot u_1(\sigma,\tau)(e^s),$$

$$u_1(\sigma,\tau)(q) = \sum_{s\in S} q(s)\cdot u_1(\sigma,\tau)(e^s).$$

Hence,

$$u_1(\sigma,\tau)(p) - u_1(\sigma,\tau)(q) = \sum_{s\in S}[p(s)-q(s)]\cdot u_1(\sigma,\tau)(e^s) \le \sum_{\substack{s\in S,\\ p(s)>q(s)}} [p(s)-q(s)] = ||p-q||_{TV},$$

and similarly

$$u_1(\sigma,\tau)(q) - u_1(\sigma,\tau)(p) \le ||p-q||_{TV}.$$

Hence, $p \mapsto u_1(\sigma,\tau)(p)$ is 1-Lipschitz-continuous.

Taking the infimum over τ and the supremum over σ on both sides of the inequality $u_1(\sigma,\tau)(p) \le u_1(\sigma,\tau)(q)+||p-q||_{TV}$ gives $v_1(p) \le v_1(q)+||p-q||_{TV}$, which can be written $v_1(p)-v_1(q) \le ||p-q||_{TV}$. Similarly, $v_1(q)-v_1(p) \le ||p-q||_{TV}$. Hence, $p \mapsto v_1(p)$ is 1-Lipschitz-continuous too.

The proof for $p \mapsto v_1(p,s)$ is similar, but the supremum has to be taken over $\sigma \in \Sigma^s$. □

Proposition 9. *We have $v_1 + v_2 = 1$, and in general, $v_1(p) + v_2(p) = 1$ for all $p \in \Delta(S)$.*

Proof. We refer to the Appendix B. □

Based on the Proposition 9, in essence the players have opposite interests in competitive search games. This leads us to defining the value of competitive search games.

Definition 10. Consider a competitive search game. We call v_1 the value of the game. For each $\varepsilon \ge 0$, a strategy σ for player 1 is called ε-optimal, if $u_1(\sigma,\tau) \ge v_1 - \varepsilon$ for all strategies $\tau \in \mathcal{T}$. Similarly, a strategy τ for player 2 is called ε-optimal, if $u_2(\sigma,\tau) \ge v_2 - \varepsilon$ for all strategies $\sigma \in \Sigma$.

The following proposition relates the notions of profiles of ε-optimal strategies and ε-equilibrium, and shows that all ε-equilibria give almost the same payoffs for small ε.

Proposition 11. *Consider a competitive search game. Let $\varepsilon \ge 0$.*

[1] If (σ,τ) is an ε-equilibrium, then σ and τ are ε-optimal strategies. Consequently, when $\varepsilon > 0$, each player has a pure ε-optimal strategy.

[2] If σ and τ are ε-optimal strategies, then (σ,τ) is a 2ε-equilibrium.

[3] A strategy profile (σ,τ) is a 0-equilibrium if and only if σ and τ are 0-optimal strategies.

[4] If (σ, τ) is an ε-equilibrium then, under (σ, τ), the object is found with probability at least $1-\varepsilon \cdot |S|$, and $|u_1(\sigma, \tau) - v_1| \leq \varepsilon$ and $|u_2(\sigma, \tau) - v_2| \leq \varepsilon$.

Proof. We refere to the Appendix B. □

Proposition 12. *Looking at a state in which the object is with zero probability is never better than looking anywhere else. That is, for all states $s, s' \in S$, for all $p \in \Delta(S)$, if $p(s') = 0$ then $v_1(p, s') \leq v_1(p, s)$.*

Proof. If player 1 chooses a state $s \in S$ then with probability $p(s)$ he finds the object immediately, and with probability $1 - p(s)$ he does not find it and then the distribution of the location of the object becomes $p^{\neg s} P_1$ in period 2, where player 2 is the active player. Hence

$$v_1(p, s) \;=\; p(s) + (1 - p(s)) \cdot v_1^2(p^{\neg s} P_1) \;=\; 1 - (1 - p(s)) \cdot v_1^2(p^{\neg s} P_1).$$

In particular, if player 1 chooses state s' then by $p(s') = 0$ we find

$$v_1(p, s') \;=\; v_1^2(pP_1).$$

Note that $pP_1 = p(s) \cdot e^s P_1 + (1 - p(s)) \cdot p^{\neg s} P_1$. By part [3] of Proposition 6,

$$v_1^2(pP_1) \;\geq\; (1 - p(s)) \cdot v_1^2(p^{\neg s} P_1).$$

This implies that $v_1(p, s) \;\geq\; v_1(p, s')$, as claimed. □

Corollary 13. *Consider a time-homogenenous competitive search game. If p is an invariant distribution of P, then $v_1(p) \geq \frac{1}{2}$.*

Proof. We refer to the Appendix B. □

Remark. We conjecture that if P is irreducible and aperiodic, then $v_1(p) > \frac{1}{2}$. The value $v_1(p)$ can be smaller than $\frac{1}{2}$ if p is not the invariant distribution. Indeed, for example with three states, initial probability distribution $p = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and a transition matrix P such that at the second period the object is in state 1 with probability 1.

5 Subgame optimal strategies

An ε-optimal strategy is a relevant solution concept, but it has the drawback that if the opponent makes a mistake, the continuation strategy does not have to be ε-optimal. Hence, in this section we examine *subgame ε-optimal* strategies.

A strategy σ for player 1 is called subgame ε-optimal if, in each subgame, the continuation strategy of σ is ε-optimal. More precisely, for each history $h \in H^{\text{odd}}$ and strategy $\tau \in \mathcal{T}$ for player 2

$$u_1(\sigma, \tau)(h) \geq v_1(h) - \varepsilon.$$

The definition of a subgame ε-optimal strategy for payer 2 is similar. Note that each subgame ε-optimal strategy is ε-optimal.

Example 14. In this example, we show that there are ε-optimal strategies that are not subgame ε-optimal. The set of states is $S = \{1, 2\}$, with each state being absorbing, and the initial probability distribution is $p = (1, 0)$. The value is $v_1 = 1$ and any optimal strategy of player 1 looks at state 1 in period 1. Hence, $v_2 = 0$ and all the strategies of player 2 are 0-optimal. In particular, it is optimal for player 2 to always choose state 2. Let τ denote this strategy.

Now suppose that player 1 makes a mistake and chooses state 2 in period 1. Then, the continuation strategy of τ from period 2 is not optimal. In fact, it would be the best for player 2 to choose state 1 in period 2 and win the game. $\square$

Proposition 15. *Consider a competitive search game.*

1. *For every $\varepsilon > 0$, each player has a pure strategy which is subgame ε-optimal.*
2. *Let $\varepsilon \in (0, \frac{1}{|S|})$. If σ is a subgame ε-optimal strategy for player 1, then for every strategy τ of player 2, the object is found with probability 1 under the strategy profile (σ, τ). A similar statement holds for player 2.*

Proof. [1] Let $\varepsilon > 0$. By Theorem 3, there exists a subgame perfect ε-equilibrium (σ, τ) in pure strategies. Now consider a subgame at a history h. Since the continuation strategies of σ and τ at h form an ε-equilibrium, it follows similarly to part [1] of Proposition 11 that the continuation strategy of σ at h is ε-optimal in the subgame, and similarly the continuation strategy of τ at h is ε-optimal in the subgame. Hence, σ and τ are subgame ε-optimal.

[2] Let $\varepsilon \in (0, \frac{1}{|S|})$ and let σ be a subgame ε-optimal strategy. Consider a history h at an odd period. The strategy for player 1 which looks at a state with the highest probability guarantees $1/|S|$ in the subgame at h. So, $v_1(h) \geq 1/|S|$.

Now consider a strategy τ for player 2. Then, we have $u_1(\sigma, \tau)(h) \geq 1/|S| - \varepsilon > 0$. In particular, in the subgame at h, the object is found with probability at least $1/|S| - \varepsilon > 0$ under (σ, τ). Since this holds for every history h at an odd period, by Lévy's zero-one law, the object is found with probability 1 under (σ, τ). $\square$

6 Optimal actions

For the initial distribution $p \in \Delta(S)$, we call an action $s \in S$ optimal if it is optimal for player 1 to look at state s in period 1: $v_1(p,s) = v_1(p)$. For a given action $s \in S$, we denote by A_s the set of the initial distributions for which s is optimal: $A_s = \{p \in \Delta(S) \mid v_1(p,s) = v_1(p)\}$. We call A_s the optimality region of action s. Note that $\cup_{s\in S} A_s = \Delta(S)$.

Theorem 16. *The optimality regions A_s have the following properties.*

[1] If the initial probability distribution p is sufficiently close to e^s, for some state s, then choosing state s is the only optimal action. That is, the set $A_s \setminus \cup_{j\neq s} A_j$ is a neighborhood of e^s in $\Delta(S)$.

[2] For each subset $N \subseteq S$, the convex hull of the vertices e^s with $s \in N$ is included in the set $\cup_{s\in N} A_s$.

[3] There is an initial distribution at which choosing any state is optimal. That is, $\cap_{s\in S} A_s \neq \emptyset$.

[4] For all $s \in S$, the region A_s is star convex sets centered in e^s. That is, if $p \in A_s$ then the whole line segment between p and e^s is included in A_s.

Proof.

[1] The statement follows from the facts that each $v_1(p,s)$ is continuous (cf. Theorem 8) in p and that $v_1(e^s, s) = 1$ and $v_1(e^s, j) < 1$ for all $j \neq s$.

[2] Let $p \in conv(\{e^s | s \in N\})$. Then $p(s) = 0$ for all $s \notin N$. By Proposition 12, there is an optimal action $j \in N$, and hence $p \in \cup_{s\in N} A_s$.

[3] We will use the Knaster-Kuratowski-Mazurkiewicz (KKM) theorem[4], see Knaster et al. [1929]. Note that by Theorem 8, the function $p \mapsto v_1(p,s)$ is continuous for all $s \in S$. Thus, each region A_s is closed. From this fact and from [2], we can apply the KKM Theorem. We conclude from the KKM Theorem that $\cap_{s\in S} A_s \neq \emptyset$.

[4] Let $s \in S$, let $p \in A_s$ and let $\lambda \in [0,1]$. We want to show that $\lambda \cdot e^s + (1-\lambda)\cdot p \in A_s$. Let (σ,τ) be a strategy profile. By equation (1)

$$\begin{aligned}\sup_{\sigma\in\Sigma} u_1(\sigma,\tau)(\lambda e^s + (1-\lambda)p) &= \sup_{\sigma\in\Sigma}\ [\lambda \cdot u_1(\sigma,\tau)(e^s) + (1-\lambda)\cdot u_1(\sigma,\tau)(p)] \\ &\leq \lambda \cdot \left[\sup_{\sigma\in\Sigma}\ u_1(\sigma,\tau)(e^s)\right] + (1-\lambda)\cdot\left[\sup_{\sigma\in\Sigma}\ u_1(\sigma,\tau)(p)\right] \\ &= \lambda + (1-\lambda)\cdot\left[\sup_{\sigma\in\Sigma}\ u_1(\sigma,\tau)(p)\right],\end{aligned}$$

[4] The KKM theorem states: Let $n \in \mathbb{N}$ be the cardinality of the set of states S, in other words $|S| = n$. Let Δ^n be the unit simplex in $\mathbb{R}^n$. A KKM covering is defined as a collection $C_1, \ldots, C_n$ of closed sets such that for any $N \subseteq \{1,\ldots,n\}$, the convex hull of the vertices corresponding to N is covered by $\cup_{s\in N} C_s$. Then any KKM covering has a non-empty intersection, i.e.: $\cap_{s\in S} C_s \neq \emptyset$.

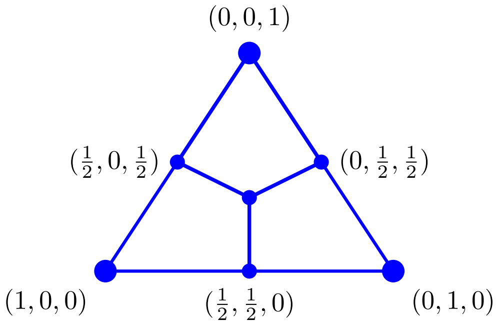

(a) The three optimal regions

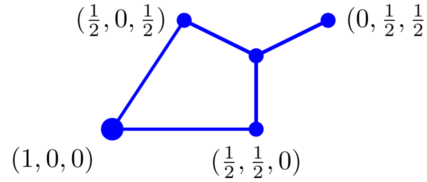

(b) Optimal region A_1.

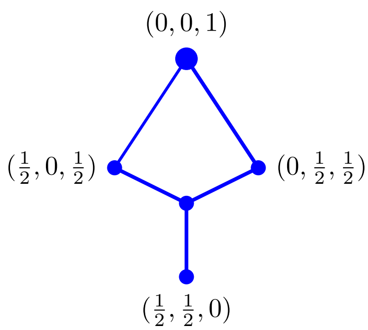

(c) Optimal region A_3

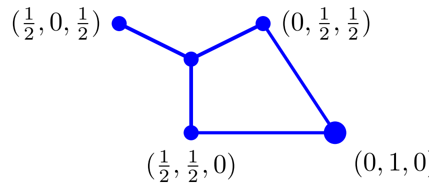

(d) Optimal region A_2

Figure 2: Optimal regions when $P = I_3$.

where we used that $u_1(\sigma, \tau)(e^s) = 1$ for any strategy σ that looks at state s in period 1. Hence

$$\begin{aligned} v_1(\lambda \cdot e^s + (1-\lambda) \cdot p) &= \inf_{\tau \in \mathcal{T}} \sup_{\sigma \in \Sigma} u_1(\sigma, \tau)(\lambda \cdot e^s + (1-\lambda) \cdot p) \\ &\leq \lambda + (1-\lambda) \cdot \left[\inf_{\tau \in \mathcal{T}} \sup_{\sigma \in \Sigma} u_1(\sigma, \tau)(p) \right] \\ &= \lambda + (1-\lambda) \cdot v_1(\sigma, \tau). \end{aligned}$$

On the other hand, by Proposition 6, $v_1(\lambda \cdot e^s + (1-\lambda) \cdot p, s) = \lambda + (1-\lambda) \cdot v_1(p, s)$. So, choosing s when the initial probability distribution is $\lambda \cdot e^s + (1-\lambda) \cdot p$ is optimal. □

Example 17. Consider the case in which the set of states is $S = \{1, 2, 3\}$. Let $Q = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 \end{bmatrix}$. The sets A_1, A_2 and A_3 are represented in the time-homogeneous case where the transition matrix is the identity matrix in Figure 2, and the matrix Q in Figure 3.

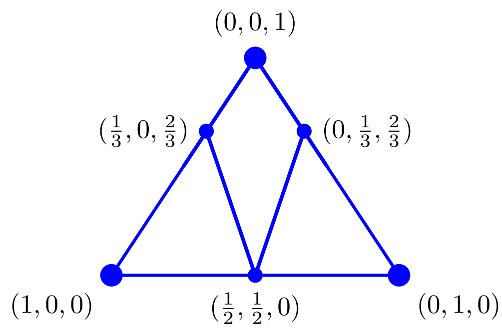

Figure 3: Optimal regions when $P = Q$.

Example 17 illustrates the statements of Theorem 16. In particular here are some remarks.

- It makes intuitive sense that if the object is in a certain state with probability close to 1, then it is optimal to look at this state. Geometrically, this means that for all states $s \in S$, the set A_s contains a neighborhood of e^s in $\Delta(S)$.
- Looking at a state s' such that $p(s') = 0$ can still be (weakly) optimal. For example, in Figure 2 and Figure 3 with initial probability distribution $p = (\frac{1}{2}, \frac{1}{2}, 0)$, looking at state 3 is just as good as looking at either state 1 or state 2. Figure 3 shows that in general, if $N \subset S$, then $\cup_{s \in N} A_s$ does not necessarily include an open neighborhood of $conv\{e^s | s \in N\}$ in $\Delta(S)$.
- Figure 2 illustrates that the intersection of the regions A_i can be more than a single point.
- Figure 2 illustrates the fact that the sets A_s are not always convex. We conjecture that their relative interior is convex. It would imply that the closure of the relative interior of the sets A_s are polytopes.

7 Concluding remarks and future work

We introduced an infinite horizon search game, in which two players compete to find an object that moves according to a time-varying Markov chain. We proved that these games always admit a subgame-perfect ε-equilibrium in pure strategies, for all error-terms $\varepsilon > 0$, but not necessarily a 0-equilibrium. We showed that the ε-equilibrium payoffs converge to a singleton $(v_1, 1 - v_1)$ as ε vanishes, and therefore the game is essentially a one-sum game with value v_1. We examined the analytical and structural properties of the solutions, and devoted attention to the important special case when the Markov chain is time-homogeneous, where stronger results hold.

We remark that, in these search games, the ε-optimal strategies are robust in the following sense: they are 2ε-optimal if the horizon of the game is finite but sufficiently long, and they are also 2ε-optimal in the discounted version of the game, provided that the discount factor is close to 1. For the precise statements and their proofs we refer to Duvocelle et al. [2020].

It would be interesting to generalize the results when the active player is chosen according to an arbitrary stochastic process. In the companion paper Duvocelle et al. [2021], we examine the variation in which the active player is chosen according to a fixed probability distribution at each period.

Also, one could introduce overlook probabilities to the model. In that case, even if the active player chooses the state that currently contains the object, there is a positive probability that the player fails to find it.

Appendices

A Topological properties of search games

Let X be a topological space. A function $f : X \to \mathbb{R}$ is called *lower semi-continuous at* $x \in X$ if, for every sequence $x^n \to x$, we have $\liminf_{n\to\infty} f(x^n) \geq f(x)$. A function $f : X \to \mathbb{R}$ is called *upper semi-continuous* at $x \in X$ if, for every sequence $x^n \to x$, we have $\limsup_{n\to\infty} f(x^n) \leq f(x)$. A function $f : X \to \mathbb{R}$ is called *continuous* at $x \in X$ if it is lower semi-continuous at x and upper semi-continuous at x.

A function $f : X \to \mathbb{R}$ is called *lower semi-continuous* (resp. *upper semi-continuous*, resp. *continuous*) if f is lower semi-continuous at all $x \in X$ (resp. upper semi-continuous at all $x \in X$, resp. continuous at all $x \in X$).

Proposition 18. *Take a player $i \in \{1, 2\}$.*

[1] The payoff function $u_i : \Sigma \times \mathcal{T} \to \mathbb{R}$ is lower semi-continuous.

[2] Assume that (σ, τ) is a strategy profile under which the object is found with probability 1. Then, u_i is continuous at (σ, τ).

Proof.

[1] For each strategy profile $(\sigma, \tau) \in \Sigma \times \mathcal{T}$, for each period $n \in \mathbb{N}$, we denote by $u_i^n(\sigma, \tau)$ the probability

that player i finds the object during the first n periods under the strategy profile (σ, τ). Note that $u_i^n(\sigma, \tau)$ is non-decreasing in n and converges to $u_i(\sigma, \tau)$ as $n \to \infty$.

Let $(\sigma^k, \tau^k)_{k \in \mathbb{N}}$ be a sequence in $\Sigma \times \mathcal{T}$ converging to a strategy profile (σ, τ). We have for each $n \in \mathbb{N}$

$$u_i^n(\sigma, \tau) = \lim_{k \to \infty} u_i^n(\sigma^k, \tau^k) = \liminf_{k \to \infty} u_i^n(\sigma^k, \tau^k) \leq \liminf_{k \to \infty} u_i(\sigma^k, \tau^k).$$

Since $u_i^n(\sigma, \tau)$ converges to $u_i(\sigma, \tau)$ as $n \to \infty$, we obtain

$$u_i(\sigma, \tau) \leq \liminf_{k \to \infty} u_i(\sigma^k, \tau^k),$$

which proves that u_i is lower semi-continuous.

[2] Assume that under the strategy profile (σ, τ) the object is found with probability 1. Thus, $u_1(\sigma, \tau) + u_2(\sigma, \tau) = 1$. Due to part 1, we only need to show that u_1 and u_2 are upper semi-continuous at (σ, τ). We will prove it for u_1; the proof for u_2 is similar.

Let $(\sigma^k, \tau^k)_{k \in \mathbb{N}}$ be a sequence in $\Sigma \times \mathcal{T}$ converging to (σ, τ). Then

$$\limsup_{k \to \infty} u_1(\sigma^k, \tau^k) = 1 - \liminf_{k \to \infty} (1 - u_1(\sigma^k, \tau^k)) \leq 1 - \liminf_{k \to \infty} u_2(\sigma^k, \tau^k) \leq 1 - u_2(\sigma, \tau) = u_1(\sigma, \tau),$$

where the first equality is a classic supinf equality applied to a limit, the first inequality comes from $u_1 + u_2 \leq 1$, the second inequality follows from part 1, and the second equality comes from the assumption we made on (σ, τ). Hence, u_1 is upper semi-continuous at (σ, τ), as desired. □

B Technical proofs

Proof of Proposition 5. In the expression $\inf_{\tau \in \mathcal{T}} \sup_{\sigma \in \Sigma} u_1(\sigma, \tau)$, player 1 is maximizing $u_1(\sigma, \tau)$ and player 2 is minimizing the same expression. Note that $(\sigma, \tau) \mapsto u_1(\sigma, \tau)$ is bounded, and by Proposition 18, it is lower semi-continuous, and hence Borel measurable. Thus, the equality $v_1 = \inf_{\tau \in \mathcal{T}} \sup_{\sigma \in \Sigma} u_1(\sigma, \tau)$ follows from Martin [1975]. The equality $v_2 = \inf_{\sigma \in \Sigma} \sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau)$ follows similarly. □

Proof of Proposition 6. First we prove part [1]. The probability distribution $\lambda \cdot p + (1 - \lambda) \cdot q$ can be interpreted as follows: with probability λ the initial probability distribution is p and induces the expected payoff $u_i(\sigma, \tau)(p)$ for player i, and with probability $1 - \lambda$ the probability distribution is q and induces the expected payoff $u_i(\sigma, \tau)(q)$ for player i. Hence, the equality (1) holds.

Now we prove part [2]. Let $p \in \Delta(S)$, $p \neq e^s$, and let $\lambda \in (0, 1)$. Then $p^{\neg s}$ is the linear projection of p from e^s to the face $\{q \in \Delta(S) | q(s) = 0\}$. Thus

$$(\lambda \cdot e^s + (1 - \lambda) \cdot p)^{\neg s} = p^{\neg s}.$$

Indeed, $(\lambda \cdot e^s + (1-\lambda)\cdot p)^{\neg s}(s) = 0 = [p^{\neg s}](s)$ and for all $j \neq s$:

$$(\lambda \cdot e^s + (1-\lambda)\cdot p)^{\neg s}(j) = \frac{(\lambda \cdot e^s + (1-\lambda)\cdot p)(j)}{1-(\lambda \cdot e^s + (1-\lambda)\cdot p)(s)} = \frac{(1-\lambda)\cdot p(j)}{1-(\lambda + (1-\lambda)\cdot p(s))}$$
$$= \frac{(1-\lambda)\cdot p(j)}{(1-\lambda)\cdot(1-p(s))} = \frac{p(j)}{(1-p(s))} = \frac{p(j)}{1-p(s)} = [p^{\neg s}](j).$$

Hence, by using $(\lambda \cdot e^s + (1-\lambda)\cdot p)(s) = \lambda + (1-\lambda)\cdot p(s)$ we have

$$\begin{aligned}
&v_1(\lambda \cdot e^s + (1-\lambda)\cdot p, s)\\
&= (\lambda \cdot e^s + (1-\lambda)\cdot p)(s) + (1-(\lambda \cdot e^s + (1-\lambda)\cdot p)(s))\cdot(1-v_1((\lambda \cdot e^s + (1-\lambda)\cdot p)^{\neg s}P))\\
&= (\lambda \cdot e^s + (1-\lambda)\cdot p)(s) + (1-(\lambda \cdot e^s + (1-\lambda)\cdot p)(s))\cdot(1-v_1(p^{\neg s}P))\\
&= \lambda + (1-\lambda)\cdot(p(s) + (1-p(s))\cdot(1-v_1(p^{\neg s}P)))\\
&= \lambda + (1-\lambda)\cdot v_1(p,s),
\end{aligned}$$

which completes the second part of the proof

Finally, we prove part [3]. Let $\varepsilon > 0$, and let σ denote an ε-optimal strategy for player 1 for the initial distribution p. Take any strategy τ for player 2. Then

$$\begin{aligned}
u_1(\sigma,\tau)(\lambda\cdot p + (1-\lambda)\cdot q) \;&=\; \lambda\cdot u_1(\sigma,\tau)(p) + (1-\lambda)\cdot u_1(\sigma,\tau)(q)\\
&\geq\; \lambda\cdot u_1(\sigma,\tau)(p)\\
&\geq\; \lambda\cdot(v_1(p)-\varepsilon).
\end{aligned}$$

Since $\varepsilon > 0$ and τ were arbitrary, the proof is complete. □

Proof of Proposition 9. We prove $v_1 + v_2 = 1$.

Step 1. Let $\widetilde{\Sigma}$ denote the set of strategies σ for player 1 such that for every strategy τ for player 2, the object is found with probability 1 under the strategy profile (σ,τ), i.e. $\mathbb{P}_{\sigma,\tau}(\Theta < \infty) = 1$. Note that $\widetilde{\Sigma}$ is nonempty; for example, $\widetilde{\Sigma}$ contains the strategy for player 1 that always chooses a state according to the uniform distribution on S. We claim that

$$v_1 \;=\; \sup_{\sigma\in\widetilde{\Sigma}}\inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau).$$

Proof of Step 1. Since $\widetilde{\Sigma} \subseteq \Sigma$, we have $v_1 \;=\; \sup_{\sigma\in\Sigma}\inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau) \;\geq\; \sup_{\sigma\in\widetilde{\Sigma}}\inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau)$.

Now we prove the opposite inequality:

$$\sup_{\sigma\in\widetilde{\Sigma}}\inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau) \;\geq\; v_1. \tag{2}$$

Let $\varepsilon \in (0, \frac{1}{|S|})$. By Mashiah-Yaakovi [2015] (or alternatively, from Mertens [1990], or from Flesch et al. [2021]), there is a strategy σ for player 1 such that for every history h

$$\inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau)(h) \;\geq\; v_1(h) - \varepsilon. \tag{3}$$

We argue that $\sigma \in \widetilde{\Sigma}$. So, take a strategy τ for player 2. Assume that h is a history at an odd period. In the subgame starting at h, player 1 can immediately win with probability at least $\frac{1}{|S|}$ if he chooses a state with the highest probability. So, $v_1(h) \geq \frac{1}{|S|}$, and hence $u_1(\sigma,\tau)(h) \geq \frac{1}{|S|} - \varepsilon > 0$. In particular, in the subgame at h, the object is found with probability at least $\frac{1}{|S|} - \varepsilon > 0$ under (σ,τ). Since this holds for every history h at an odd period, by Lévy's zero-one law, the object is found with probability 1 under (σ,τ). This proves that $\sigma \in \widetilde{\Sigma}$, as desired.

Since $\sigma \in \widetilde{\Sigma}$, by applying (3) to the empty history in period 1, we find $\inf_{\tau\in\mathcal{T}} u_1(\sigma,\tau) \;\geq\; v_1 - \varepsilon$. Since $\varepsilon > 0$ is arbitrary, (2) follows.

Step 2. Recall that Σ^s denotes the set of those strategies for player 1 that choose state s in period 1. Define

$$v_2(p,s) \;=\; \inf_{\sigma\in\Sigma^s}\sup_{\tau\in\mathcal{T}} u_2(\sigma,\tau)(p).$$

We claim that for every $\delta > 0$ there is $w(\delta) > 0$ such that if $p \in \Delta(S)$ satisfies $p(s') \leq w(\delta)$ for some state $s' \in S$, then for all states $s \in S$ we have

$$v_2(p,s') + \delta \;\geq\; v_2(p,s). \tag{4}$$

Intuitively, if player 1 wishes to minimize player 2's winning chances, then choosing a state s' with a very low probability of containing the object cannot be much better than choosing any state s.

Proof of Step 2. For each state s we have $pP_1 = p(s)\cdot e^s P_1 + (1-p(s))\cdot p^{\neg s}P_1$. Hence, By part [3] of Proposition 6

$$v_1(pP_1) \;\geq\; (1-p(s))\cdot v_1(p^{\neg s}P_1). \tag{5}$$

Note that if player 1 chooses a state $s \in S$ then with probability $p(s)$ he finds the object immediately, and with probability $1-p(s)$ he does not find it and then the distribution of the location of the object becomes $p^{\neg s}P_1$ in period 2, where player 2 is the active player. Hence, $v_2(p,s) \;=\; (1-p(s))\cdot v_1(p^{\neg s}P_1)$.

Thus, for all states $s, s' \in S$ we find

$$\begin{aligned}
v_2(p,s) - v_2(p,s') &= (1-p(s)) \cdot v_1(p^{\neg s}P_1) - (1-p(s')) \cdot v_1(p^{\neg s'}P_1)) \\
&= (1-p(s)) \cdot v_1(p^{\neg s}P_1) - v_1(p^{\neg s'}P_1) + p(s') \cdot v_1(p^{\neg s'}P_1) \\
&\leq v_1(pP_1) - v_1(p^{\neg s'}P_1) + p(s') \\
&\leq ||(p - p^{\neg s'}) \cdot P_1||_{TV} + p(s') \\
&= \frac{1}{2} \cdot \sum_{s \in S} |[(p - p^{\neg s'}) \cdot P_1](s)| + p(s') \\
&= \frac{1}{2} \cdot \sum_{s \in S} \Big| \sum_{r \in S} (p(r) - p^{\neg s'}(r)) \cdot P_1(r,s) \Big| + p(s') \\
&\leq \frac{1}{2} \cdot \sum_{s \in S} \sum_{r \in S} \big| (p(r) - p^{\neg s'}(r)) \big| \cdot P_1(r,s) + p(s') \\
&= \frac{1}{2} \cdot \sum_{r \in S} \big| (p(r) - p^{\neg s'}(r)) \big| \cdot \sum_{s \in S} P_1(r,s) + p(s') \\
&= \frac{1}{2} \cdot \sum_{r \in S} \big| (p(r) - p^{\neg s'}(r)) \big| + p(s') \\
&= ||p - p^{\neg s'}||_{TV} + p(s') \\
&= 2 \cdot p(s'),
\end{aligned}$$

where in the first inequality we used (5) and $v_1(p^{\neg s'}P_1) \leq 1$, and in the second inequality we used Theorem 8. Now the claim of Step 2 follows immediately.

Step 3. We claim that

$$v_2 = \inf_{\sigma \in \widetilde{\Sigma}} \sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau).$$

Proof of Step 3. Since $\widetilde{\Sigma} \subseteq \Sigma$, we have $v_2 = \inf_{\sigma \in \Sigma} \sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau) \leq \inf_{\sigma \in \widetilde{\Sigma}} \sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau)$.

Now we prove the opposite inequality: $v_2 \geq \inf_{\sigma \in \widetilde{\Sigma}} \sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau)$. Let $\varepsilon > 0$. By Mashiah-Yaakovi [2015] (or alternatively, from Mertens [1990], or from Flesch et al. [2021]), there is a pure strategy σ for player 1 such that for every history h

$$\sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau)(h) \leq v_2(h) + \varepsilon.$$

Since σ is pure, $\sigma(h)$ places probability 1 on a state $s_h \in S$ for each history h.

For each $n \in \mathbb{N}$, let $w_n = w\big(\frac{\varepsilon}{2^n}\big)$ where $w\big(\frac{\varepsilon}{2^n}\big)$ is as in Step 3. Note that $w_n > 0$ for each n.

For each $n \in \mathbb{N}$, we define a strategy σ_n^* for player 1 as follows: start using σ, until a history h_1 occurs (if it occurs at all) such that s_{h_1} contains the object with probability at most w_1. Then, at h_1, choose a state which contains the object with the highest probability (which is at least $1/|S|$). Then, follow σ again, until a history h_2 occurs (if it occurs at all) such that s_{h_2} contains the object with probability at most w_2. Then, at h_2, choose a state which contains the object with the highest probability (which is at least $1/|S|$). Then, follow σ again, and so on, until a history h_n occurs (if it occurs at all) such that s_{h_n} contains the object with probability at most w_n. Then, at h_n, choose a state which contains the object with the highest probability (which is at least $1/|S|$), and then follow σ in the remaining game. Note that, along the play of the game, σ_n^* deviates from σ at most n times.

We define σ^* as the limit of the strategies σ_n^* when $n \to \infty$. This means that, along the play of the game, σ^* may deviate from σ at infinitely many histories. We argue that $\sigma^* \in \widetilde{\Sigma}$ and

$$\sup_{\tau \in \mathcal{T}} u_2(\sigma^*, \tau) \;\leq\; v_2 + 2\varepsilon,$$

which will then complete the proof of Step 3.

We show first that $\sigma^* \in \widetilde{\Sigma}$. Take any strategy τ for player 2, and consider the play according to (σ^*, τ). By construction of σ^*: (i) whenever σ^* deviates from σ, the object is found immediately with probability at least $1/|S|$, (ii) if after a period t, σ^* never deviates from σ, then at every period $t' \geq t$, σ finds the object immediately with probability at least w_n for some fixed value of n. This implies that, under (σ^*, τ), the object is found with probability 1. That is, $\sigma^* \in \widetilde{\Sigma}$.

Now we show that

$$\sup_{\tau \in \mathcal{T}} u_2(\sigma^*, \tau) \;\leq\; v_2 + 2\varepsilon.$$

Take any strategy τ for player 2. By the previous argument, there is a period n such that the object is found with probability close to 1 within n periods. This implies that σ^* deviates from σ at most n times, and hence, $u_2(\sigma^*, \tau)$ is close to $u_2(\sigma_n^*, \tau)$. By iteratively applying Step 3, we find that

$$u_2(\sigma_n^*, \tau) \;\leq\; u_2(\sigma, \tau) + \varepsilon + \frac{\varepsilon}{2} + \cdots + \frac{\varepsilon}{2^n} \;<\; u_2(\sigma, \tau) + 2\varepsilon.$$

Hence, the proof is complete.

Step 4. We prove that $v_1 + v_2 = 1$. We have

$$v_1 \;=\; \sup_{\sigma \in \widetilde{\Sigma}} \inf_{\tau \in \mathcal{T}} u_1(\sigma, \tau) \;=\; \sup_{\sigma \in \widetilde{\Sigma}} \inf_{\tau \in \mathcal{T}} \big[1 - u_2(\sigma, \tau)\big] \;=\; 1 - \inf_{\sigma \in \widetilde{\Sigma}} \sup_{\tau \in \mathcal{T}} u_2(\sigma, \tau) \;=\; 1 - v_2.$$

Here, the first equality follows from Step 2, the second equality follows from the fact that if σ belongs to $\widetilde{\Sigma}$ then the object is found with probability 1, and the last equality follows from Step 3. Hence, $v_1 + v_2 = 1$ as desired. $\square$

Proof of Proposition 11. Let $\varepsilon \geq 0$.

[1] Let (σ, τ) be an ε-equilibrium. Since $u_1(\sigma, \tau) \geq u_1(\sigma', \tau) - \varepsilon$ for all $\sigma' \in \Sigma$, we have $u_1(\sigma, \tau) \geq v_1 - \varepsilon$, which means that σ is ε-optimal. It follows similarly that τ is ε-optimal as well.

When $\varepsilon > 0$, by Theorem 3, there is an ε-equilibrium (σ, τ) in pure strategies. Hence, σ and τ are ε-optimal strategies for the players.

[2] Assume now that σ and τ are ε-optimal strategies for player 1 and player 2. Let $\sigma' \in \Sigma$. Then, $u_2(\sigma', \tau) \geq v_2 - \varepsilon$. By Proposition 9, we get that

$$u_1(\sigma', \tau) \;\leq\; 1 - u_2(\sigma', \tau) \;\leq\; 1 - (v_2 - \varepsilon) \;=\; v_1 + \varepsilon.$$

This implies that $u_1(\sigma, \tau) \;\geq\; v_1 - \varepsilon \;\geq\; u_1(\sigma', \tau) - 2\varepsilon$. Similarly, we obtain $u_2(\sigma, \tau) \geq u_2(\sigma, \tau') - 2\varepsilon$ for every $\tau' \in \mathcal{T}$. So, (σ, τ) is a 2ε-equilibrium.

[3] This is a direct consequence of [1] and [2].

[4] Let (σ, τ) be an ε-equilibrium.

First we prove that under (σ, τ), the object is found with probability at least $1 - \varepsilon \cdot |S|$. Suppose the opposite: it is found with probability $z := \mathbb{P}_{\sigma,\tau}(\theta < \infty) < 1 - \varepsilon \cdot |S|$. Let $\widetilde{H}$ denote those histories after which the object is found with probability at most $\gamma := \frac{1}{2}(\frac{1}{|S|} - \frac{\varepsilon}{1-z})$:

$$\widetilde{H} \;=\; \{h \in H \colon \mathbb{P}_{\sigma,\tau}(\theta < \infty | h) \leq \gamma.\}$$

By Lévy's zero-one law, under (σ, τ), it has probability at least $1 - z$ that a history in $\widetilde{H}$ arises during the play. Now consider the strategy σ' for player 1 that follows σ as long as no history arises in $\widetilde{H}$, but as soon as this happens, σ' chooses each state with equal probability and plays arbitrarily afterwards. Then,

$$u_1(\sigma', \tau) - u_1(\sigma, \tau) \;\geq\; (1 - z) \cdot (\frac{1}{|S|} - \gamma) \;>\; \varepsilon.$$

This contradicts the fact that (σ, τ) be an ε-equilibrium.

Now prove the inequalities in [4]. As we argued in the proof of part [1], $u_1(\sigma, \tau) \geq v_1 - \varepsilon$ and similarly, $u_2(\sigma, \tau) \geq v_2 - \varepsilon$. Thus, by Proposition 9

$$u_1(\sigma, \tau) \leq 1 - u_2(\sigma, \tau) \leq 1 - (v_2 - \varepsilon) = v_1 + \varepsilon.$$

Similarly, $u_2(\sigma, \tau) \leq v_2 + \varepsilon$. These inequalities give $|u_1(\sigma, \tau) - v_1| \leq \varepsilon$ and $|u_2(\sigma, \tau) - v_2| \leq \varepsilon$. □

Proof of Corollary 13. Assume first that there is a state $s \in S$ for which $p(s) = 0$. Then $p^{\neg s}P = \pi P = \pi$. Since $\pi(s) = 0$ we have $v_1(p, s) = 1 - v_1(p)$. As $v_1(p) \geq v_1(p, s)$, we obtain $v_1(p) \geq 1 - v_1(p)$. Hence, $v_1(p) \geq \frac{1}{2}$.

Assume there is no state $s \in S$ for which $p(s) = 0$. Consider the game G' that arises by adding a state w to G. More precisely, G' is the game with set of states $S' = S \cup \{w\}$, initial probability distribution p' such that $p'(s) = p(s)$ for each state $s \in S$ and $p'(w) = 0$, and transition matrix P' that has the same transition probabilities between states in S and makes w absorbing. Then, with probability 1, the object will never be in w. By Proposition 12, the players may ignore state w during the game. Then, p' is an invariant distribution of P', and hence by the first part we find $v_1(p) = v'(p') \geq \frac{1}{2}$. $\square$

References

Tatjana V Abramovskaya, Fedor V Fomin, Petr A Golovach, and Michał Pilipczuk. How to hunt an invisible rabbit on a graph. *European Journal of Combinatorics*, 52:12–26, 2016.

Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*, volume 55. Springer Science & Business Media, 2006.

Steve Alpern, Vic Baston, and Shmuel Gal. Network search games with immobile hider, without a designated searcher starting point. *International Journal of Game Theory*, 37(2):281–302, 2008.

Steve Alpern, Vic Baston, and Shmuel Gal. Searching symmetric networks with utilitarian-postman paths. *Networks: An International Journal*, 53(4):392–402, 2009.

Spyros Angelopoulos and Thomas Lidbetter. Competitive search in a network. *European Journal of Operational Research*, 2020.

Vic Baston, FA Bostock, and William H Ruckle. The gold-mine game. *Journal of optimization theory and applications*, 64(3):641–650, 1990.

Anatole Beck and Dava J Newman. Yet more on the linear search problem. *Israel journal of mathematics*, 8(4):419–429, 1970.

Stanley J Benkoski, Michael G Monticino, and James R Weisinger. A survey of the search theory literature. *Naval Research Logistics (NRL)*, 38(4):469–494, 1991.

Donald A Berry and Roy F Mensch. Discrete search with directional information. *Operations research*, 34(3):470–477, 1986.

Scott Shorey Brown. Optimal search for a moving target in discrete time and space. *Operations research*, 28(6):1275–1289, 1980.

Buyang Cao. Search-hide games on trees. *European journal of operational research*, 80(1):175–183, 1995.

Jake Clarkson. *Optimal search in discrete locations: extensions and new findings.* PhD thesis, Lancaster University, 2020.

Arnon Dagan and Shmuel Gal. Network search games, with arbitrary searcher starting point. *Networks: An International Journal*, 52(3):156–161, 2008.

Florian Delavernhe, Patrick Jaillet, André Rossi, and Marc Sevaux. Planning a multi-sensors search for a moving target considering traveling costs. *European Journal of Operational Research*, 2020.

James M Dobbie. A two-cell model of search for a moving target. *Operations Research*, 22(1):79–92, 1974.

Benoit Duvocelle, János Flesch, Mathias Staudigl, and Dries Vermeulen. A competitive search game with a moving target. *arXiv preprint arXiv:2008.12032*, 2020.

Benoit Duvocelle, János Flesch, Hui Min Shi, and Dries Vermeulen. Search for a moving target in a competitive environment. *International Journal of Game Theory*, pages 1–11, 2021.

Bradley Efron. Optimum evasion versus systematic search. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):450–457, 1964.

János Flesch and Arkadi Predtetchinski. Subgame-perfect epsilon-equilibria in perfect information games with common preferences at the limit. *Mathematics of Operations Research*, 41(4):1208–1221, 2016.

János Flesch, Emin Karagözoğlu, and Andrés Perea. Optimal search for a moving target with the option to wait. *Naval Research Logistics (NRL)*, 56(6):526–539, 2009.

János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operations Research*, 35(4):742–755, 2010.

János Flesch, P Jean-Jacques Herings, Jasmine Maes, and Arkadi Predtetchinski. Subgame maxmin strategies in zero-sum stochastic games with tolerance levels. *Dynamic Games and Applications*, pages 1–34, 2021.

Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite-and infinite-horizon games. *Journal of Economic Theory*, 31(2):251–268, 1983.

Shmuel Gal. A general search game. *Israel Journal of Mathematics*, 12(1):32–45, 1972.

Shmuel Gal. Minimax solutions for linear search problems. *SIAM Journal on Applied Mathematics*, 27(1):17–30, 1974.

Shmuel Gal. Search games with mobile and immobile hider. *SIAM Journal on Control and Optimization*, 17(1):99–122, 1979.

Shmuel Gal. Search games. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.

Shmuel Gal. Search games: a review. In *Search Theory*, pages 3–15. Springer, 2013.

Shmuel Gal and Dan Chazan. On the optimality of the exponential functions for some minimax problems. *SIAM Journal on Applied Mathematics*, 30(2):324–348, 1976.

Andrey Garnaev. *Search games and other applications of game theory*, volume 485. Springer Science & Business Media, 2012.

Tristan Garrec and Marco Scarsini. Search for an immobile hider on a stochastic network. *European Journal of Operational Research*, 283(2):783–794, 2020.

JC Gittins and DM Roberts. The search for an intelligent evader concealed in one of an arbitrary number of regions. *Naval Research Logistics Quarterly*, 26(4):651–666, 1979.

Ryusuke Hohzaki. Search games: Literature and survey. *Journal of the Operations Research Society of Japan*, 59(1):1–34, 2016.

Ryusuke Hohzaki and Koji Iida. Optimal ambushing search for a moving target. *European Journal of Operational Research*, 133(1):120–129, 2001.

Rufus Isaacs. *Differential games: a mathematical theory with applications to warfare and pursuit, control and optimization.* Control and Warfare Wiley, New York, 1965.

Benjamin P Jordan. *On optimal search for a moving target.* PhD thesis, Durham University, 1997.

Arun Jotshi and Rajan Batta. Search for an immobile entity on a network. *European Journal of Operational Research*, 191(2):347–359, 2008.

YC Kan. A counterexample for an optimal search-and-stop model. *Operations Research*, 22(4):889–892, 1974.

Kensaku Kikuta. A search game on a cyclic graph. *Naval Research Logistics (NRL)*, 51(7):977–993, 2004.

Bronisław Knaster, Casimir Kuratowski, and Stefan Mazurkiewicz. Ein beweis des fixpunktsatzes für n-dimensionale simplexe. *Fundamenta Mathematicae*, 14(1):132–137, 1929.

David A Levin and Yuval Peres. *Markov chains and mixing times.* American Mathematical Society, 2017.

Thomas Lidbetter. Search and rescue in the face of uncertain threats. *European Journal of Operational Research*, 285(3):1153–1160, 2020.

Iain M MacPhee and Benjamin P Jordan. Optimal search for a moving target. *Probability in the Engineering and Informational Sciences*, 9(2):159–182, 1995.

Donald A Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

Michael Maschler, Eilon Solan, and Shmuel Zamir. Game theory. *Cambridge University Press, Cambridge*, 2013.

Ayala Mashiah-Yaakovi. Correlated equilibria in stochastic games with borel measurable payoffs. *Dynamic Games and Applications*, 5(1):120–135, 2015.

Jean-François Mertens. Repeated games. In *Game Theory and Applications*, pages 77–130. Elsevier, 1990.

Teruhisa Nakai. Model of search for a target moving among three boxes: Some special cases. *Journal of Operations Research Society of Japan*, 16:151–162, 1973.

Teruhisa Nakai. A search game with one object and two searchers. *Journal of applied probability*, 23 (3):696–707, 1986.

Teruhisa Nakai. A preemptive detection game. *Journal of Information and Optimization Sciences*, 11 (1):1–15, 1990.

Marcel F Neuts. A multistage search game. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):502–507, 1963.

Ljiljana Pavlović. A search game on the union of graphs with immobile hider. *Naval Research Logistics (NRL)*, 42(8):1177–1189, 1995.

Stephen M Pollock. A simple model of search for a moving target. *Operations Research*, 18(5):883–903, 1970.

JH Reijnierse and Jos AM Potters. Search games with immobile hider. *International Journal of Game Theory*, 21(4):385–394, 1993.

Darby M Roberts and John Charles Gittins. The search for an intelligent evader: Strategies for searcher and evader in the two-region problem. *Naval Research Logistics Quarterly*, 25(1):95–106, 1978.

Minoru Sakaguchi. A sequential allocation game for targets with varying values. *Journal of the Operations Research Society of Japan*, 20(3):182–193, 1977.

Paul J Schweitzer. Threshold probabilities when searching for a moving target. *Operations Research*, 19(3):707–709, 1971.

Lawrence D Stone. *Theory of optimal search*, volume 118. Elsevier, 1976.

Lawrence D Stone, Johannes O Royset, and Alan R Washburn. Optimal search for moving targets (international series in operations research & management science 237). *Cham, Switzerland: Springer*, 2016.

Eduardo J Subelman. A hide-search game. *Journal of Applied Probability*, pages 628–640, 1981.

John von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games*, 2(0):5–12, 1953.

Bernhard von Stengel and Ralph Werchner. Complexity of searching an immobile hider in a graph. *Discrete Applied Mathematics*, 78(1-3):235–249, 1997.

Alan R Washburn. Search for a moving target: The FAB algorithm. *Operations research*, 31(4): 739–751, 1983.

Noemí Zoroa, M José Fernández-Sáez, and Procopio Zoroa. A foraging problem: sit-and-wait versus active predation. *European Journal of Operational Research*, 208(2):131–141, 2011.