



Search for a moving target in a competitive environment

Benoit Duvocelle¹ · János Flesch¹ · Hui Min Shi¹ · Dries Vermeulen¹

Accepted: 24 January 2021 / Published online: 17 February 2021

© The Author(s) 2021

Abstract

We consider a discrete-time dynamic search game in which a number of players compete to find an invisible object that is moving according to a time-varying Markov chain. We examine the subgame perfect equilibria of these games. The main result of the paper is that the set of subgame perfect equilibria is exactly the set of greedy strategy profiles, i.e. those strategy profiles in which the players always choose an action that maximizes their probability of immediately finding the object. We discuss various variations and extensions of the model.

Keywords Game theory · Search game · Optimal search · Greedy strategy · Subgame perfect equilibrium

1 Introduction

In this paper we consider a dynamic search game, in which an object moves according to a time-varying Markov chain across a finite set of positions. The set of competing players can be either finite or infinite. At each period, an active player is drawn according to a fixed distribution, and he chooses one of the possible positions. If the object is in this position, then the active player finds the object and wins the game. Otherwise, the object moves according to the transition matrix, and the game enters

We would like to thank Aditya Aradhya, Luca Margaritella, Niels Mourmans, the Associate Editor and two anonymous Reviewers for their precious comments and suggestions.

✉ Benoit Duvocelle
b.duvocelle@maastrichtuniversity.nl

János Flesch
j.flesch@maastrichtuniversity.nl

Hui Min Shi
hm.shi@student.maastrichtuniversity.nl

Dries Vermeulen
d.vermeulen@maastrichtuniversity.nl

¹ Department of Quantitative Economics, Maastricht University, Maastricht, The Netherlands

the next period. Each player observes the action chosen by his opponents, and the transition probabilities, initial probabilities and probabilities of being the active player are known to all players. The goal of each player is to find the object and win the game. So each player prefers the winning outcome and is indifferent between the outcomes in which one of his opponents wins or in which nobody finds the object.

The main result of this paper is that the set of subgame perfect equilibria is exactly the set of greedy strategy profiles, also known as myopic strategy profiles, i.e. those strategy profiles in which the active player always selects one of the most likely positions containing the object. The key to this result is to show that by playing a greedy strategy, each player can guarantee that he wins with a probability at least as much as his probability of being active. This implies that, in each subgame, under a greedy strategy profile each player wins with exactly the same probability as his probability of being active. Hence, under a greedy strategy profile, all subgames are identical as far as the winning probabilities are concerned, and each greedy strategy profile is robust to one-deviations: the best each player can do at a certain period is to choose an action that maximizes the probability to win immediately. All these arguments hinge on the fact that the active player is chosen according to the same distribution at each period. Intuitively, this assumption ensures that the active player cannot invest into the future by finding the object with low probability at the present period, as this would not only benefit him but also the other players, proportionally to the fixed probabilities with which each player becomes active.

The rest of this paper is divided as follows. In Sect. 1.1 we refer to related literature. In Sect. 1.2 we introduce the model and in Sect. 1.3 we present an illustrative example. In Sect. 2 we study the winning probabilities of players playing a greedy strategy, and we present a characterization of the subgame perfect equilibria. We show that the set of subgame perfect equilibria is exactly the set of greedy strategy profiles. In Sect. 3 we discuss some extensions of the model and see to which extend the main result still holds or not. The conclusion is in Sect. 4.

1.1 Related literature

The field of search problems is one of the original disciplines of Operations Research, with various applications such as military problems, R&D problems or patent races, and many of these models involve multiple decision makers. In the basic settings, the searcher's goal is to find a hidden object, also called the target, with maximal probability or as soon as possible. By now, the field of search problems has evolved 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, 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 are multiple searchers, they 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 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. For an introduction to search games, we refer to Alpern and Gal (2006), Gal (1979, 2010, 2013), Garnaev (2012), and for surveys see Benkoski et al. (1991) and Hohzaki (2016).

There are many different types of search games in the literature due to variations in characteristics of the searchers and of the object. Pollock (1970) introduced a simple search game with one searcher, in which an object is moving across two locations referred to as “state 1” and “state 2” respectively, according to a discrete-time Markov chain. The objective of the searcher can be either to minimize the expected number of looks to find the object, or to maximize the probability of finding the object within a given horizon. The model allows for overlooking probabilities, which means that even if the searcher chooses the correct location, he may fail to find the object there. Instead, Nakai (1973) investigates the search problem with three states, while assuming perfect detection, so not accounting for overlooking probabilities. Flesch et al. (2009) investigates a search game similar to Pollock (1970), but in addition to searching for the object in state 1 or state 2, the player has a third option which is to wait. The point of waiting is that it is costless and can induce a favorable probability distribution over the two states at the next period. It is shown that there is a unique optimal strategy, which is characterized by two thresholds. Jordan (1997) studies the structural properties of the optimal strategy, where the goal is to find the object while minimizing the search costs. Thereby he derived some properties of the optimal strategy for the search problem with a finite set of states in the no-overlook case and for the case where each state has the same overlooking probability and cost. Assaf and Sharlin-Bilitzky (1994) investigates a search problem with two states in which the object moves according to a continuous-time Markov process. The objective is to find the object with a minimal expected cost, where the “real time” until the object is found is also taken into account in the cost structure. A competitive environment with more searchers and a static object is considered in Nakai (1986). Finally, in the companion paper (Duvocelle et al. 2020), we consider a two-player competitive search game where players play by turns. We show that an equilibrium does not always exist, but that we can find subgame perfect ε -equilibria for all $\varepsilon > 0$. A classical reference for an overview of search games is the book of Alpern et al. (2013). For a recent paper on search games, we refer to Garrec and Scarsini (2020). They model the search game as a zero-sum two-person stochastic game where one player is looking for the other one. They provide upper and lower bounds on the value of the game.

1.2 The model

1.2.1 The game

An object is moving over a finite set $S = \{1, \dots, n\}$ of states according to a time-varying Markov chain. The initial distribution of the object is given by $\pi = (\pi_s)_{s \in S}$, and the transition probabilities at each period $t \in \mathbb{N}$ are given by the $S \times S$ transition matrix P_t , where entry $P_t(s, s')$ is the probability that the object moves to state s' , given it is in state s at period t .

Let I denote a set of players, who compete to find the object. We assume that $I \subseteq \mathbb{N}$ and $|I| \geq 2$; so the set I can be either finite or countably infinite. The players do not observe the current state of the object, but they know the initial distribution π and the transition matrices P_t , for each $t \in \mathbb{N}$. At each period $t \in \mathbb{N}$ one of the players is active: player i is active with probability $q_i > 0$, where $\sum_{i \in I} q_i = 1$. The active player chooses a state $s_t \in S$. If the object is in state s_t , then the active player finds the object and wins the game. Otherwise the game enters period $t + 1$. We assume that each player observes the actions chosen by his opponents, and knows the probabilities q_i , for each $i \in I$.

The goal of each player is to find the object and win the game. For each player the game has three possible outcomes. The first possible outcome is that the player himself finds the object and wins the game. The second outcome is that one of his opponents finds the object and wins the game. The third outcome is that no one finds the object. In this game each player prefers the first outcome, but is indifferent between the second and third outcomes. This means that players do not have opposite interests, which makes it a non-zero-sum game.

1.2.2 Actions and histories

The action set for each player i is $A_i = S$. Thus, a history at period $t \in \mathbb{N}$ is a sequence $h_t = (i_1, s_1, \dots, i_{t-1}, s_{t-1}) \in S^{2(t-1)}$ of past active players and past actions. By $H_t = S^{2(t-1)}$ we denote the set of all histories at period t . Note that H_1 consists of the empty sequence. Given a history h_t , with the knowledge of the initial probability distribution π and the transition matrices P_1, \dots, P_{t-1} , the players can calculate the probability distribution of the location of the object at period t . This probability distribution does not depend on the past active players, only on the past actions in h_t .

1.2.3 Strategies

A strategy for player i is a sequence of functions $\sigma_i = (\sigma_{i,t})_{t \in \mathbb{N}}$ where $\sigma_{i,t}: H_t \rightarrow \Delta(S)$ for each period $t \in \mathbb{N}$. The interpretation is that, at each period $t \in \mathbb{N}$, if player i becomes the active player, then given the history h_t , the strategy $\sigma_{i,t}$ recommends to search state $s \in S$ with probability $\sigma_{i,t}(h_t)(s)$. We denote by Σ_i the set of strategies of player i . We say that a strategy is pure if, for any history, it places probability 1 on one action.

A strategy is called greedy if, for any history, it places probability 1 on the most likely states. In the literature the greedy strategy is sometimes also called the myopic strategy.

1.2.4 Winning probabilities

Consider a strategy profile $\sigma = (\sigma_i)_{i \in I}$. The probability under σ that player i wins is denoted by $u_i(\sigma)$. Note that

$$\sum_{i \in I} u_i(\sigma) = 1 - \text{Prob}_\sigma[\text{object never found}].$$




Fig. 1 An illustrative example

The aim of player i is to maximize $u_i(\sigma)$. If the object has not been found before period t , and the history is h_t , the continuation winning probability from period t onward is denoted by $u_i(\sigma)(h_t)$ for player i .

1.2.5 Subgame perfect equilibrium

A strategy σ_i for player i is a best response to a profile of strategies σ_{-i} for all other players if $u_i(\sigma) \geq u_i(\sigma'_i, \sigma_{-i})$ for every strategy $\sigma'_i \in \Sigma_i$. A strategy profile $\sigma = (\sigma_i)_{i \in I}$ is called an equilibrium of the game if σ_i is a best response to σ_{-i} for each player $i \in I$.

A strategy σ_i for player i is a best response in the subgame at history h to a profile of strategies σ_{-i} for all other players, if $u_i(\sigma)(h) \geq u_i(\sigma'_i, \sigma_{-i})(h)$ for every strategy $\sigma'_i \in \Sigma_i$. A strategy profile $\sigma = (\sigma_i)_{i \in I}$ is called a subgame perfect equilibrium if, at each history h , in the subgame at history h the strategy σ_i is a best response to σ_{-i} for each player $i \in I$. In other words, σ is a subgame perfect equilibrium if it induces an equilibrium in each subgame.

1.3 An illustrative example

In the model we introduced the greedy strategies. However, it is not clear a priori if those strategies are relevant, nor if one greedy strategy can be better or worst than another one. Now we examine an example in order to try to answer those questions. Consider the following game with a parameter $c \in (0, 1)$, with two states and two players. At each period player 1 is active with probability $q_1 = 0.99 = 1 - q_2$. The initial probability of the location of the object is $\pi = (c, 1 - c)$ and the transition matrices are defined as follows: for all $t \in \mathbb{N}$,

$$P_t = P = \begin{pmatrix} c & 1 - c \\ 1 & 0 \end{pmatrix}.$$

The induced Markov chain of this game is depicted in Fig. 1.

We discuss two cases.

Case 1 Consider the case in which $0 < c < \frac{1}{3}$. At period $t = 1$ it is more likely for the object to be in state 2 as $c < 1 - c$.

Intuitively, if player 2 gets the chance to be active at period 1, he should look in state 2 as his chance to play later is quite low. However, it is not immediately clear what player 1 should do at period 1. On the one hand, if he looks in state 1 and he does not find the object, then he finds it in period 2 with probability 1 if he can play,

which is likely to happen as $q_1 = 0.99$. On the other hand he might also simply want to maximise his chance to find the object at period 1 by looking at state 2.

More precisely, assume first that player 2 is active at period 1. If he looks at state 1, with probability c he finds the object and with probability $1 - c$ the active player at period 2 can find the object by looking at state 1. In this case, player 2 finds the object with probability $c + q_2 \cdot (1 - c)$. However, if he looks at state 2, with probability $1 - c$ he finds the object and with probability c the active player at period 2 can find the object with probability $1 - c$ by looking at state 2. In this case, player 2 finds the object with probability at least $1 - c + q_2 \cdot c \cdot (1 - c) > c + q_2 \cdot (1 - c)$ as $q_2 = 0.01$ and $c < \frac{1}{3}$. Thus, it is strictly better for player 2 to look at state 2 at period 1. The same holds for each subgame in which the probability distribution of the object is $(c, 1 - c)$; in particular this is the case after a player choosing state 2 and not finding the object there.

Now assume that player 1 is active at period 1. If he looks at state 1, then similarly to the analysis of player 2, he finds the object with probability $c + q_1 \cdot (1 - c)$. However, if player 1 decides to choose state 2 at each period he is active, by our analysis of player 2, it follows that player 2 will also do the same. Hence player 1 finds the object with probability $(1 - c) + q_1 \cdot c \cdot (1 - c) + q_1 \cdot c^2 \cdot (1 - c) + \dots = 1 - c + q_1 \cdot c > c + q_1 \cdot (1 - c)$ as $q_1 < 1$ and $0 < c < \frac{1}{3}$. So it is better for player 1 to choose state 2 at period 1.

From our discussion, it follows that every Nash equilibrium induces the play in which at each period the active player chooses state 2. If a deviation occurs to state 1 and the object is not found, the next active player chooses state 1 and wins the game. This is a greedy strategy profile as at each period the most likely state is chosen.

Case 2 Consider the case in which $c = 1/2$.

In this case, both states are equally likely, so a greedy strategy can choose any state at period 1. Note that at any period, if state 1 is chosen and the object is not found, a greedy strategy finds the object at the next period in state 1. However, if state 2 is chosen and the object is not found, a greedy strategy can choose any of the two states at the next period. Therefore, there are many greedy strategy profiles, and all can induce different plays. It is then natural to ask whether every greedy strategy profile induces the same winning probabilities. As we will show later, every greedy strategy profile leads to the winning probability q_1 for player 1 and q_2 for player 2.

2 Greedy strategies and subgame perfect equilibria

In this section we prove our main result, which is that the set of subgame perfect equilibria is exactly the set of greedy strategy profiles. We start by introducing an intermediate result related to the winning probability guarantees of the players who play a greedy strategy.

Proposition 1 *Consider a strategy profile $\sigma = (\sigma_j)_{j \in I}$ and a player $i \in I$. If σ_i is a greedy strategy, then under σ , the object is found with probability 1 and player i wins with probability at least q_i .*

Proof Let $\sigma = (\sigma_j)_{j \in I}$ be a strategy profile in which player i plays a greedy strategy (the strategies of the other players are arbitrary).

STEP 1. Under σ , the object is found with probability 1.

PROOF OF STEP 1. Consider an arbitrary period $t \in \mathbb{N}$ and suppose that the object has not been found yet. Then, player i plays with probability q_i , and when he plays he wins with probability at least $1/n$. This implies that the object is found at period t with probability at least q_i/n . Since this holds for each period t , under σ the object is found with probability 1.

STEP 2. Under σ , player i wins with probability at least q_i .

PROOF OF STEP 2. Consider a period $t \in \mathbb{N}$ and a history $h \in H_t$. For each state $s \in S$, let $z_h(s)$ denote the probability that the object is in state s at period t given the history h . Let $z_h^* = \max_{s \in S} z_h(s)$.

For each player $j \in I$, let $w_j(h)$ denote the probability that after history h player j becomes the active player and by using the mixed action $\sigma_j(h)$ he wins immediately. Since σ_i is a greedy strategy, we have $w_i(h) = q_i \cdot z_h^*$. For each other player $j \neq i$ we have $w_j(h) \leq q_j \cdot z_h^*$. Hence, given that the object is found at period t after the history h , the conditional probability that player i finds it is

$$\frac{w_i(h)}{\sum_{j \in I} w_j(h)} \geq \frac{q_i \cdot z_h^*}{\sum_{j \in I} q_j \cdot z_h^*} = q_i.$$

Since this holds for every period t and every history $h \in H_t$, and since by Step 1 the object is found under σ with probability 1, player i wins with probability at least q_i . □

Proposition 2 *If σ is a greedy strategy profile or σ is an equilibrium, then under σ , each player i wins with probability q_i .*

Proof By Proposition 1, each player i can guarantee that he wins with probability at least q_i . Since $\sum_{i \in I} q_i = 1$, the statement follows. □

Now we state and prove the main result of the paper.

Theorem 3 *The set of subgame perfect equilibria is exactly the set of greedy strategy profiles.*

Proof The proof is divided into two steps.

STEP 1. Every greedy strategy profile is a subgame perfect equilibrium.

PROOF OF STEP 1. Let σ be a greedy strategy profile and consider the subgame at a history $h \in H$. We show that σ induces an equilibrium in this subgame. Consider a player i and a deviation σ'_i . By Proposition 2, we have $u_i(\sigma)(h) = q_i$. By Proposition 1, the strategy σ_j guarantees to each player $j \in I$ that he wins with probability at least q_j , even if another player deviates. Since $\sum_{j \in I} q_j = 1$, we find $u_i(\sigma'_i, \sigma_{-i}) \leq q_i$. Thus, the deviation σ'_i is not profitable.

STEP 2. Every subgame perfect equilibrium is a greedy strategy profile.

PROOF OF STEP 2. Assume by way of contradiction that there exists a subgame perfect equilibrium $\sigma = (\sigma_i)_{i \in I}$, which is not a greedy strategy profile. Suppose that σ_i is not greedy at history h . Let strategy σ'_i be a one-deviation from σ_i at history h , under which player i plays greedy at history h . So $\sigma'_i(h) \neq \sigma_i(h)$ and $\sigma'_i(h') = \sigma_i(h')$ for every $h' \in H \setminus \{h\}$.

Let $z_h(\sigma_i(h))$ denote the probability that player i finds the object immediately with σ_i at history h given that he is the active player. Similarly, let $z_h(\sigma_{-i}(h))$ denote the probability that one of the opponents of player i finds the object immediately with σ_{-i} at history h given that one of the opponents is active.

Then we have

$$\begin{aligned} u_i(\sigma_i, \sigma_{-i})(h) &= q_i \cdot [z_h(\sigma_i(h)) + (1 - z_h(\sigma_i(h))) \cdot q_i] + (1 - q_i) \cdot (1 - z_h(\sigma_{-i}(h))) \cdot q_i \\ &= q_i \cdot z_h(\sigma_i(h)) \cdot (1 - q_i) + q_i \cdot [q_i + (1 - q_i) \cdot (1 - z_h(\sigma_{-i}(h)))]. \end{aligned}$$

The first equality follows from the following argument. Player i has a probability of q_i of becoming the active player. Then, he wins immediately with probability $z_h(\sigma_i(h))$ or he does not find the object immediately with probability $1 - z_h(\sigma_i(h))$ and still wins in the future with probability q_i by Proposition 2. However, with probability $1 - q_i$ one of his opponents is active. Then, player i still has a chance to win. His opponents fail to find the object immediately with probability $1 - z_h(\sigma_{-i}(h))$ and then player i can again win in the future with probability q_i .

Similarly, for the deviation σ'_i we have

$$\begin{aligned} u_i(\sigma'_i, \sigma_{-i})(h) &= q_i \cdot [z_h(\sigma'_i(h)) + (1 - z_h(\sigma'_i(h))) \cdot q_i] + (1 - q_i) \cdot (1 - z_h(\sigma_{-i}(h))) \cdot q_i \\ &= q_i \cdot z_h(\sigma'_i(h)) \cdot (1 - q_i) + q_i \cdot [q_i + (1 - q_i) \cdot (1 - z_h(\sigma_{-i}(h)))]. \end{aligned}$$

Since $z_h(\sigma'_i(h)) > z_h(\sigma_i(h))$ holds due to the fact that σ'_i is greedy at h but σ_i is not, we find $u_i(\sigma'_i, \sigma_{-i})(h) > u_i(\sigma_i, \sigma_{-i})(h)$. This however contradicts the assumption that σ is a subgame perfect equilibrium. \square

3 Extensions and variations

In this section we consider several extensions of the model and discuss if the main theorem still holds or not.

3.1 Extensions where our results still hold

3.1.1 History dependent transitions

In the model description we assumed that the object moves according to a time-varying Markov chain. A more general situation is when the transition probabilities at each period t can also depend on (i) the sequence of past active players, (ii) the sequence of past choices of the active players, and (iii) the sequence of states visited by the object in the past. Without any modification in the proofs, our main result, Theorem 3, remains valid.

3.1.2 Overlooking

Our main result, Theorem 3, can be extended to a model with overlooking probabilities. Overlooking means that the object is in the chosen state, but the active player “overlooks” it and therefore does not find it. Note that players cannot distinguish between overlooking and searching in a vacant state. For this extension, let $\delta_s < 1$ denote the probability of overlooking the object in state s , for each $s \in S$.

The overlooking probabilities need to be taken into account when defining greedy strategies. A greedy strategy should choose a state s that maximizes the probability of containing the object times the probability of not overlooking the object, i.e. it should maximize the immediate probability of finding the object.

3.1.3 No active player

The model could also be adjusted for the probability that no player is active, i.e. $\sum_{i \in I} q_i < 1$. Then, $r := 1 - \sum_{i \in I} q_i$ is the probability that no player is active at a certain period. We still assume that $q_i > 0$ for each player i , so that $r < 1$. Our main result, Theorem 3, would still hold, as the key properties of Proposition 1 remain valid.

3.2 Variations where results break down or need adjustment

3.2.1 Robustness to finite horizon and discounting

In view of Theorem 3, the greedy strategy profiles are natural solutions in our search games. We will now examine two closely related variations of our model, search games on finite horizon and discounted search games, and investigate how the greedy strategy profiles perform in them. Consider a search game as described in Sect. 1.2, and let σ be an arbitrary greedy strategy profile.

The search game on horizon T , where $T \in \mathbb{N}$, is played in the same way at periods $1, \dots, T$, but if the object has not been found by the end of period T , then the game ends. Let $\varepsilon > 0$ be an arbitrary error-term. It follows from standard arguments that if the horizon T is sufficiently large, then the greedy strategy profile σ is an ε -equilibrium in the search game on horizon T . Here, an ε -equilibrium is a strategy profile such that no player can increase his probability of finding the object by more than ε with a unilateral deviation.

The discounted search game with discount factor $\delta \in (0, 1)$ is played in the same way as the original search game, but now each player i maximizes $\sum_{t=1}^{\infty} \delta^{t-1} \cdot z_{i,t}$, where $z_{i,t}$ is the probability that player i finds the object at period t . Let $\varepsilon > 0$ be an arbitrary error-term. If the discount factor δ is close to 1, then the greedy strategy profile σ is an ε -equilibrium in the search game with discount factor δ . It is quite intuitive, and the main reason is that, as long as at least one player plays a greedy strategy, the object will be found at an exponential rate, so essentially in finite time.

Note however that a greedy strategy profile is not necessarily a 0-equilibrium on finite horizon. Indeed, consider the game represented in Fig. 2 over $T = 2$ periods. There are three states. The initial probability distribution of the object is given by




Fig. 2 A 2-period game

$\pi = (1/3 + \varepsilon, 1/3 + 2\varepsilon, 1/3 - 3\varepsilon)$, where $\varepsilon > 0$ is small enough so that $0.99 \cdot (2/3 + 3\varepsilon) \leq 2/3$, and the transitions at period 1 are given by the arrows in Fig. 2. There are two players, and player 1 plays at each period with probability $q_1 = 0.99$.

The greedy strategy profile is to choose state 2 at period 1 and state 1 at period 2. Under this strategy profile, player 1 finds the object with probability $q_1 \cdot (2/3 + 3\varepsilon) \leq 2/3$. However, it would be a profitable deviation for player 1 to choose state 3 at period 1 and state 1 at period 2. Indeed, player 1 is the active player at both periods with probability $(q_1)^2$, and in that case this strategy finds the object with probability 1.

3.2.2 Period-dependent probabilities q_i

It is crucial for our proofs that the probabilities with which the players become active do not depend on the period. Indeed, if two players alternate in searching for the target, the (unique) greedy strategy profile in the leading example of Sect. 1.3 with $c = 0.51$ fails to be an equilibrium. Suppose that player 1 is active at period 1. The greedy strategy would dictate to choose state 1. Then either player 1 wins immediately or player 2 wins at period 2 by choosing state 1. Thus, choosing state 1 would make player 1 win with probability $c = 0.51$. However, by choosing state 2 at period 1, player 1 would win with probability much more than 0.51, as this would make player 2 uncertain about the location of the object at period 2. Intuitively, the current decision of the active player should not only depend on the immediate probability to win, but also on the conditional probability distribution of the object at the next period.

3.2.3 Infinitely many states

In our model, we assumed that the set of states is finite. If there are infinitely many states, the greedy strategy profiles may fail to be an equilibrium. In fact, this can happen even when there is only one player. As an example, consider the following game. The set of states is $S = \{0\} \cup \mathbb{N}$. The object starts in state 0 with probability 0.51 and in state 1 with probability 0.49. State 0 is absorbing. If the object is in state 1, it then moves to states 2, 3, 4 with equal probability $\frac{1}{3}$, and from these states it moves to states 5, ..., 13, with equal probability $\frac{1}{9}$, and in general, at each period t it moves to 3^{t-1} new states with equal probability.

If the player first chooses state 2 and then state 1, he finds the object with probability 1. However, any greedy strategy would first choose state 1, and in that case the object cannot be found with probability 1 any more, as the transition law of the object when starting in state 2 is diffuse.

4 Conclusion

In this paper we examined a discrete-time search game with multiple competitive searchers who look for one object moving over finitely many locations. We showed that if the probability to play for each player may vary among players but it is the same at each period, the set of subgame perfect equilibria is exactly the set of greedy strategy profiles (cf. Theorem 3). We discussed several variations, such as the finite truncation of the game, the discounted version of the game, cases with infinitely many states, overlooking probabilities and we examined the possibility that no player is active at a period. A challenging task would be to investigate stochastic search games when the probability of a player to be active depends on the history.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alpern S, Gal S (2006) The theory of search games and rendezvous. Springer Science & Business Media, New York
- Alpern S, Fokkink R, Gasieniec L, Lindelauf R, Subrahmanian VS (2013) Search theory. Springer, Berlin
- Assaf D, Sharlin-Bilitzky A (1994) Dynamic search for a moving target. *J Appl Probab* 31(2):438–457
- Benkoski SJ, Monticino MG, Weisinger JR (1991) A survey of the search theory literature. *Naval Res Logist (NRL)* 38(4):469–494
- Duvocelle B, Flesch J, Staudigl M, Vermeulen D (2020) A competitive search game with a moving target. arXiv preprint, [arXiv:2008.12032](https://arxiv.org/abs/2008.12032)
- Flesch J, Karagözoğlu E, Perea A (2009) Optimal search for a moving target with the option to wait. *Naval Res Logist (NRL)* 56(6):526–539
- Gal S (1979) Search games with mobile and immobile hider. *SIAM J Control Optim* 17(1):99–122
- Gal S (2010) Search games. *Wiley Encyclopedia of Operations Research and Management Science*. Wiley, New York
- Gal S (2013) Search games: a review. In: Search theory. Springer, Berlin, pp 3–15
- Garnaev A (2012) Search games and other applications of game theory. Springer Science and Business Media, New York
- Garrec T, Scarsini M (2020) Search for an immobile hider on a stochastic network. *Eur J Oper Res* 283(2):783–794
- Hohzaki R (2016) Search games: literature and survey. *J Oper Res Soc Jpn* 59(1):1–34
- Jordan BP (1997) On optimal search for a moving target. PhD thesis, Durham University
- Nakai T (1973) Model of search for a target moving among three boxes: some special cases. *J Oper Res Soc Jpn* 16:151–162
- Nakai T (1986) A search game with one object and two searchers. *J Appl Probab* 23(3):696–707
- Pollock SM (1970) A simple model of search for a moving target. *Oper Res* 18(5):883–903

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.