ARTICLE IN PRESS

Organizational Behavior and Human Decision Processes xxx (2014) xxx–xxx

Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

journal homepage: www.elsevier.com/locate/obhdp

Sequential search by groups with rank-dependent payoffs: An experimental study

Vincent Mak [a,*], Amnon Rapoport [b], Darryl A. Seale [c]

[a] Cambridge Judge Business School, University of Cambridge, United Kingdom
[b] A. Gary Anderson Graduate School of Management, University of California, Riverside, United States
[c] Lee Business School, University of Nevada, Las Vegas, United States

ARTICLE INFO

Article history:
Received 11 April 2013
Accepted 25 March 2014
Available online xxxx

Keywords:
Secretary problem
Sequential search by groups
Strategic behavior
Contrast effects

ABSTRACT

In many sequential search situations, decisions are reached by groups. We examine behavior in such situations experimentally using an extension of the "secretary problem". In our setup, group members (players) with non-aligned preferences inspect alternatives or "applicants" one at a time with no backward solicitation. A minimal information structure is assumed where players are only informed of the *relative ranks* of the alternatives as they inspect them sequentially. We present the equilibrium solution, and then use it as a benchmark for our analysis. We report the results from a controlled experiment showing that subjects over-searched relative to equilibrium. Decisions were affected by theoretically irrelevant observations including the relative rank of the previous alternative and the other player's relative rank of the current alternative. For managers engaged in committee sequential search tasks, our findings highlight the importance of being aware to reach compromises early on, among other implications.

Introduction

In this study, we propose an experimental paradigm for investigating search behavior that combines two streams of research. The first stream concerns sequential search of alternatives or "applicants" that arrive in a random order and are inspected or "interviewed" one at a time; the second concerns decision making by groups. Group members in our setup have non-aligned preferences over the applicants, and the decision at each stage of the search process – whether to accept or reject the current applicant – is reached collectively by some pre-determined mechanism.

To exemplify our sequential search process, consider first the simpler case of a single decision maker (DM) who has to choose one of *n* alternatives. A common example is of an employee who has newly moved to a city and is looking for an apartment to rent. He is provided with a finite list of rental apartments that he may visit and inspect one at a time; when visiting an apartment on the list, he may only rank its overall quality with respect to the apartments he already has visited. For a second example, consider an employer who faces the problem of determining the time to stop interviewing additional job applicants in order to maximize the probability of hiring the best applicant or, alternatively, maximize the expected rank of the hired applicant among all applicants (rank 1 being the best). Both examples can, and have been, categorized under a class of sequential search problems, often called *secretary problems* (SP; see e.g., DeGroot, 1970), that are common in human resource management, labor economics, operations management, marketing, and decision making in general. An important, critical, and realistic characteristic of the SP framework is that the DM has virtually no knowledge about the distribution of the quality of the searchable options. The DM only knows that she can rank all the options, and the options could present themselves in any order of rankings as she inspects them one by one. She could only base her (binary) decision to either stop or continue search on such paltry information.

There is a vast theoretical literature in applied probability and operations research on SP dating back to the 1960s (see Ferguson, 1989; Freeman, 1983; Samuels, 1991 for reviews); on the other hand, experimental work on behavioral decision making in SP contexts is still in its infancy (cf. Bearden & Rapoport, 2005). Moreover, most theoretical and experimental work in this area is concerned with a single DM. By contrast, in this study we report an experimental study on an SP-type sequential search by a *group*. The extension from individual to group search has obvious realistic relevance and potential applications. Apartments are often

* Corresponding author. Address: Cambridge Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, United Kingdom. Tel: +44 (0)1223 764295; Fax: +44 (0)1223 339701.
E-mail address: v.mak@jbs.cam.ac.uk (V. Mak).

http://dx.doi.org/10.1016/j.obhdp.2014.03.004
0749-5978/

collectively rented or purchased by families and not by an individual DM; members of the family evaluate different attributes of the apartment under inspection (e.g., price, location, number of bathrooms) and express their preferences that are combined in some way into a single decision. Job applicants are frequently interviewed by committees of managers who evaluate the applicants on several different attributes (e.g., appearance, proficiency, social skills). Decision by groups has been a substantial topic of research in social psychology, decision making, and organizational behavior (Kerr & Tindale, 2004). But rarely, if at all, have researchers in that literature examined sequential search decisions. Sequential search and group decision making have hitherto been discussed separately in organizational behavior, decision making, experimental economics, and other disciplines. Here we attempt to bridge the gap between them by applying an SP framework to investigate group decision making in a sequential search setting.

Our research objectives are primarily to find out whether binary decisions to accept or reject the current applicant deviate from benchmark theoretical predictions; and if so, what might be the drivers of these deviations, especially regarding the transparency of preference information among the DMs in the group. We also aim to find out if behavioral regularities such as the contrast effects (cf. Tversky & Griffin, 1991) of theoretically irrelevant ranking observations might have impacted their decisions in our setting; this is especially valuable in that contrast effects have hardly been explored in the domain of sequential search. Through examining these issues in our study, we have obtained managerial implications into group decision making in a sequential search context. These notably include the insight – as will be discussed later – that group DMs may not be sufficiently alert to reaching compromises early on, and that this problem can be mitigated by sharing preference information among group members.

Theoretical context of the present study

We start by evaluating the information structure that underlies our proposed paradigm, and then state the assumptions of the game in detail. Technically, the SP is a *no-information* search problem, as opposed to *full-information* search problems.[1] In full-information search problems (e.g., Rapoport & Tversky, 1966), the DM sequentially observes (with a cost incurred for every additional observation) up to n choice alternatives X_i ($i = 1,2,\ldots,n$), that are drawn independently from a given density function $f(x)$ that is assumed to be known by the player. It is also typically assumed in applications that the DM's objective is to maximize the expected value of the chosen alternative minus the cumulative cost of search. The optimal policy has the form: stop the search and choose alternative i whenever $x_i > x_i^*$, where x_i^* is an optimal cutoff value determined recursively, and continue otherwise. Bearden and Rapoport (2005) and Bearden, Rapoport, and Murphy (2006) have pointed out two shortcomings of this model. The first limiting feature is the assumption that $f(x)$ is known by the DM *before the search commences*. Under the full-information model, the DM is supposed to numerically determine the cutoff values x_i^* before taking the first observation and then extract no further information from subsequent observations (no learning). However, in reality the search is adaptive, learning takes place, and people adjust their cutoff values as they gather more information. The second limiting feature is the assumption that $f(x)$ is known by the DM *with precision*. Yet, DMs are liable to misperceive likelihoods, and values of x_i^* are typically highly sensitive to such misperceptions.

By contrast, the *no-information* search problems have no distributional assumptions. The DM sequentially observes up to n alternatives, and in each stage i of the search it is only assumed that the DM determines the relative rank of alternative i with respect to all the alternatives $1,2,\ldots,i-1$ that he has inspected (i.e., the best, the second best, etc.). The DM's payoff is generally related to the absolute rank of the chosen alternative, where an alternative's absolute rank is defined as its rank among *all* the n alternatives *before* the search commences. It is this class of problems that we focus on in the present study and introduce in the next section.

The standard or classic SP (CSP) may be stated as follows (Bearden et al., 2006):

1. There is only a single position to be filled.
2. There are n applicants for the position. The value of n is known before the search commences.
3. Applicants are ranked from best (rank 1) to worst (rank n) with no ties.
4. The applicants are interviewed sequentially, one at a time (each interview is called a *stage* in the search), in a random order i.e., all $n!$ orderings are equally probable. Once an applicant is interviewed, her relative rank is determined in comparison to the previously observed applicants that have been interviewed and rejected.
5. For each applicant that he interviews, the DM must either accept it, and thereby terminate the search, or reject it and interview the next one, if any.
6. Once rejected, an applicant may not be recalled. If stage n is reached, the nth applicant must be accepted.
7. The DM's objective is to maximize the probability of choosing the best applicant among all the n applicants (i.e., the applicant with an absolute rank of 1). This implies that the DM wins 1, if he selects the best applicant, and 0, otherwise.

The optimal search policy for the CSP is to interview and reject the first $t^* - 1$ applicants, no matter their relative ranks, and then accept the first applicant thereafter with a relative rank 1 (Lindley, 1961). The cutoff point t^* converges from above to ne^{-1}, and the optimal policy selects the best applicant with probability $e^{-1} \approx 0.368$, as $n \to \infty$.

It has long been recognized that the CSP is too restrictive for most applications (Bearden et al., 2006), and just about any of the assumptions above has been relaxed to give rise to new models (see, e.g., Samuels, 1991). As stated earlier, perhaps the most limiting restriction of the CSP is Assumption 7 stating that the DM is satisfied with "nothing but the best". Chow, Moriguti, Robbins, and Samuels (1964) proposed a more general case replacing Assumption 7 by the following:

7'. The DM's objective is to minimize the expected absolute rank of the chosen applicant.

Secretary problems satisfying Assumptions 1–6 and 7' are termed *expected rank secretary problems* (ERSP). The optimal policy is quite different from the single-cutoff policy for the CSP. It takes the form of a multi-threshold rule with cutoff values $x_1^*, x_2^*, \ldots, x_n^* = n$, entailing interviewing and rejecting the first $x_1^* - 1$ applicants with no regard to their relative ranks, from period x_1^* to period $x_2^* - 1$ only selecting an applicant with relative rank 1, from period x_2^* to $x_3^* - 1$ only selecting an applicant with relative rank of either 1 or 2, and so on (see Chow et al., 1964). The values of x_i^* are computed recursively. It is the ERSP version of secretary problems that will form the basis of our experimental setup.

[1] "Information" as used in this section (in the sense of full- vs. no-information) refers to information about the distribution of the utilities/payoffs of the alternatives in an individual search problem. It should not be confused with "information" as used later on in labeling our experimental manipulations (in the sense of common vs. private information), which refers to whether a player is informed about the preferences of the other player in a two-person group search problem.