Journal of Economic Behavior & Organization 123 (2016) 122–137

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

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

A two-agent model of sequential search and choice☆

Eeva Mauring*

Department of Economics, University College London, Gower Street, London WC1E 6BT, United Kingdom

ARTICLE INFO

Article history:
Received 16 April 2015
Received in revised form
18 December 2015
Accepted 28 December 2015
Available online 21 January 2016

JEL classification:
D83
D79
D99

Keywords:
Sequential search
Multi-agent search
Threshold policy
Discouragement effect

ABSTRACT

This paper extends the standard sequential search model by allowing the agent who compiles the choice set via search (the “searcher”) to differ from the agent who chooses from the set (the “chooser”). I show for a general joint distribution of the agents' preferences that the searcher's optimal policy is a threshold rule. In contrast to the standard model, the threshold is weakly decreasing in time (i.e., exhibits the “discouragement effect”), although the search horizon is infinite and the search environment stationary. I characterise the threshold and discuss the testable implications of the discouragement effect. The characteristics of my model differ from two single-agent search models that feature a time-varying threshold (convex search costs or deadline). In particular, my model features a threshold that decreases endogenously over time and never generates return to an item rejected earlier, in contrast to the other models.

1. Introduction

Standard sequential search models with recall build on the assumption that the search and choice stages comprise an undivided whole: the person who searches can stop and choose an item from the accumulated choice set at any time during the search process. This is an innocuous assumption if the preferences of the person are stable over time. In this paper, I extend the standard search model by allowing the preferences according to which the final choice is made to differ from the preferences according to which search is conducted. The set-up has two natural interpretations. First, the preferences belong to different parties: a “searcher” compiles a choice set via sequential search and a “chooser” chooses from the collected choice set. Second, the preferences belong to one individual, but change between the search and choice stages. I show that the searcher's optimal policy is a threshold rule and characterise the threshold.

Examples of such search problems are an HR manager collecting applications for a boss who wants to hire a new worker and a real estate agent collecting offers for a client interested in buying a flat. An example involving a person and a set of individuals is a spouse looking for a job that determines the living place of the couple. A person who is interested in the return while searching for an investment opportunity, but later tempted to invest in an option that involves the least paperwork

☆ I thank my advisors Ran Spiegler and Martin Cripps for support and advice, and the editor, two anonymous referees, Jan Eeckhout, Andrea Pirrone, Nikita Roketskiy, Ludo Visschers, Lukas Wenner, and audiences at UCL, RES Junior Symposium, SaM Annual Conference, and RES Women's Meeting for comments. Financial support from ERC grant no. 230251; ESRC; Kristjan Jaak Scholarship funded by the Ministry of Education and Research, Estonia, and the Archimedes Foundation; and Krediidipank is gratefully acknowledged. Earlier title: “Two-Person Search”.

* Tel.: +44 75838 38329.
E-mail address: eevamauring@gmail.com

http://dx.doi.org/10.1016/j.jebo.2015.12.010

is a "multi-selves" example. More generally, many household decisions, organisational decisions involving different phases and multiple agents, choice processes partially outsourced to external partners, and political decisions involving advisers feature one party compiling a choice set for another party via search.

In this paper I analyse the optimal policy of a searcher who compiles a choice set for a chooser. I describe the model as a two-agent search problem. The searcher (he) and the chooser (she) have preferences over all items in some grand set of alternatives and the preferences are distributed according to a general full-support distribution function. The searcher has access to an arrival process. In each period, one item arrives and the searcher discovers how much utility both he and the chooser receive from the item if it is chosen. The searcher decides in each period whether to stop or continue the search process. If he stops the process, all the items that have arrived are presented to the chooser. The chooser then chooses the best item in the choice set according to her preferences, unless all the items in the set yield her less utility than her exogenous outside option. Utilities are realised when an item or the outside option is chosen. The process ends after the chooser moves. The searcher's problem is to choose an optimal policy, knowing the chooser's choice rule.

First, I derive and characterise the searcher's optimal policy. His optimal policy is a threshold rule and the threshold depends on the item that would be chosen by the chooser were the searcher to stop immediately, x_m. The searcher's threshold is the lower the higher the chooser values this item because it acts as a restriction for the searcher: if he is unsatisfied with the utility he would receive from x_m, a new item is chosen only if the chooser's utility from it exceeds her utility from x_m. This has two implications. First, if the searcher finds an item that has a very high value for the chooser, he optimally stops searching regardless of the value that the item yields him. Second, the observed threshold that the searcher uses is weakly decreasing in time, although the search horizon is infinite and the search environment stationary. I call this the "discouragement effect". This is in contrast with the standard single-agent search model where a stationary environment translates into a stationary threshold. The searcher's threshold in my model is defined implicitly by a differential equation. I use a specific joint distribution where the utilities' correlation is captured by a single parameter to numerically show that an increase in the correlation parameter unambiguously increases the searcher's threshold, in line with intuition.

Second, I compare the optimal threshold of the searcher in the main model with imperfectly correlated preferences to the benchmark where the agents' preferences are perfectly aligned. I first show that in the main model the searcher's threshold is always lower than in the benchmark: the searcher is "less picky". The reason is that the chooser chooses according to her preferences not the searcher's, which lowers the latter's continuation value, thus, his threshold. I then provide an example where, as a result of a mean-preserving spread, the searcher's threshold decreases in the main model, while it always increases in the benchmark: the searcher is "more conservative" in the main model.[1] A mean-preserving spread increases the probability that an item arrives that yields very high utility to the chooser, which restricts the searcher and lowers his continuation value.

Third, I explain how the model's characteristics differ from those of two single-agent search models that feature a threshold that varies in time. The first model has convex search costs and the second, a deadline. Both models result in an optimal threshold that decreases over time (for a fixed searcher-preferred item in the choice set) because they assume non-stationarity of the environment. My model features a time-decreasing threshold in a stationary environment. Also, in those models returning to an item found earlier is possible, while search always stops with the item found last in my model. The searcher returns to an item found earlier in the models with convex search costs or a deadline because his threshold decreases exogenously over time. In my model, the decrease is endogenous: it happens only if there is a change in the item that would be chosen if the searcher stopped. I suggest three tests on data that allow us to reject one or more of the three models.

Finally, I extend the model to allow the searcher to hide items. The searcher can hide an item only upon its arrival and succeeds with some given probability that is strictly less than one. He observes whether he succeeded before making the decision whether to stop or continue. As in the main model, I find that the searcher's threshold is unambiguously decreasing in the value that the chooser receives from the chooser-preferred element in the choice set. For independent uniform utilities, I show that the threshold is strictly increasing in the probability that the searcher succeeds in hiding. The constraint of having to account for the chooser's preferences becomes the less restrictive the likelier that the searcher can ignore those preferences.

1.1. Related literature

This paper is closely related in spirit to other papers on multi-agent search. In "committee" search problems the committee has a common arrival process and must agree on when to stop.[2] In "couple" search problems each person has his own arrival process, but they pool income.[3] In these papers some part of the entire search process is joint, while distinct parties are engaged in distinct stages of the process in my paper. I borrow the terms "less picky" and "more conservative" in their specific meaning from Albrecht et al. (2010) (AAV henceforth). AAV analyse a committee search problem, where M members

[1] The terms "less picky" and "more conservative" are borrowed from Albrecht et al. (2010).

[2] For example, see Albrecht et al. (2010), Compte and Jehiel (2010), Bergemann and Välimäki (2011), Kamada and Muto (2011), and Moldovanu and Shi (2013).

[3] For example, see Dey and Flinn (2008), Ek and Holmlund (2010), Flabbi and Mabli (2012), and Guler et al. (2012).