Journal of Economic Behavior & Organization 74 (2010) 104–122

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

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

When is it important to know you've been rejected? A search problem with probabilistic appearance of offers

Sanmay Das[a,*], John N. Tsitsiklis[b]

[a] Dept. of Computer Science, Rensselaer Polytechnic Institute, 110 8th St, Troy, NY 12180, United States
[b] Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Received 7 July 2006
Received in revised form 22 January 2010
Accepted 26 January 2010
Available online 4 February 2010

JEL classification:
D83
J41

Keywords:
Search problems
Secretary problem
Value of information
Sequential decision-making

ABSTRACT

A problem that often arises in the process of searching for a job or for a candidate to fill a position is that applicants do not know if they will receive an offer from any given firm with which they interview, and, conversely, firms do not know whether applicants will definitely take positions they are offered. In this paper, we model the search process as an optimal stopping problem with probabilistic appearance of offers from the perspective of a single decision-maker who wants to maximize the realized value of the offer she accepts. Our main results quantify the *value of information* in the following sense: how much better off is the decision-maker if she knows each time whether an offer appeared or not, compared to the case where she is only informed when offers actually appear? We show that for some common distributions of offer values, she can expect to receive very close to her optimal value even in the lower information case, as long as she knows the probability that any given offer will appear. However, her expected value in the low information case (as compared to the high information case) can fall dramatically when she does not know the appearance probability *ex ante* but must infer it from data. This suggests that hiring and job-search mechanisms may not suffer from serious losses in efficiency or stability from participants hiding information about their decisions, unless agents are uncertain of their own attractiveness as employees or employers.

1. Introduction

Many job markets are structured in a manner where potential employees submit their applications to a number of employing firms simultaneously, and then wait to hear back from these firms. Firms themselves often make exploding offers that employees have to decide on in a short time-frame. Sometimes the firms will tell potential employees as soon as they are no longer under consideration, and in other cases they wait until the end of the search process to provide this information to applicants. The central question that we address in this paper is this: *How much better off is an applicant if she is told every time she has been rejected by a firm, as opposed to only knowing when she receives offers?*

In order to study this problem, we construct a stylized model in which the decision problem faced by agents is a version of the problem variously referred to in the literature as the Cayley–Moser problem, the (job) search problem, the house hunting problem and the problem of selling an asset (Ferguson, 1989). In the original problem, a job applicant knows that there will be exactly n job opportunities, which will be presented to her sequentially. At the time each job is presented, she observes the utility she would receive from taking that job offer (one can think of it purely in terms of wages), and must

* Corresponding author. Tel.: +1 518 276 2782.
E-mail addresses: sanmay@cs.rpi.edu (S. Das), jnt@mit.edu (J.N. Tsitsiklis).

0167-2681/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jebo.2010.01.005

decide immediately whether to accept the job offer or not. If she declines the offer, she may not go back to it. If she accepts it, she may not pick any of the subsequent offers. What is the strategy that maximizes her expected utility? This problem has been addressed for various distributions of offer values, and much of that work is summarized by Gilbert and Mosteller (1966).

The problem we consider is a variant of the above problem in which the total number of possible offers is known, but each offer appears only with a certain probability. This problem is motivated in part by models of two-sided matching markets like labor markets or dating markets. In particular, a problem considered by Das and Kamenica (2005) is one in which men are asked out on dates by women and must respond immediately, but while they have priors on the values of going out with particular women, they do not know the order in which women are going to appear, so they are not aware of whether or not a better option might come along in the future. This is because a better woman than the one currently asking a man out might either have already appeared in the ordering and not asked him out, or might appear later and not ask him out, or might appear later and ask him out. A similar problem can arise in faculty hiring processes for universities and colleges. Universities may not know whether applicants will definitely take positions that are offered, and, conversely, applicants do not know if they will receive an offer from any given university with which they interview. This paper only looks at one side of this process without considering the dynamics involved when multiple agents interact, potentially strategically. Another motivation comes from thinking of the offers as investment opportunities (Gilbert and Mosteller, 1966). In particular, the continuous-time variant we discuss can be interpreted in terms of investment opportunities that arrive as a Poisson process where the decision-maker wants to choose the best one. To simplify the analysis, we assume that the probability that a particular offer appears, p, is the same across all offers and is independent of the actual value of the offer. The value of p may or may not be known to the applicant and can be thought of as a measure of the “attractiveness” of the applicant or decision-maker.

Most of the previous research on search models focuses on solving an agent’s infinite horizon optimal stopping problem when there is either a cost to generating the next offer, or a discount factor associated with future utility (the book by DeGroot (1970) provides an account of much of this line of research). The problem we study here is a finite-horizon search problem with no cost to seeing more offers and no search frictions. The basic questions we pose and attempt to answer relate to how much the expected utility of the decision-maker changes between different information sets and different mechanisms. The question with regard to information sets can be thought of as follows. Suppose you interview with n firms that might want to hire you. Then the companies get ordered randomly and come along in that order and decide whether or not to make you an offer. How much would you pay to go from a situation in which you saw only which companies made you an offer (the *low information* variant) to a situation in which you saw, for each company, whether or not they chose to make you an offer (the *high information* variant)? Generalizing the two informational cases to continuous time provides good approximations for large n and insight into the value of information in these cases. It also allows us to make an interesting connection to a closely related problem called the *secretary problem*. We will also discuss the difference in expected utility between two different mechanisms. The exploding offer mechanism can lead to a substantial decline in the expected utility of a job-seeker compared to a mechanism in which she sees all the offers she will receive simultaneously and can choose from among them. What if you could pay to see the entire set of offers you would get simultaneously so that you could pick among them? How much should you be willing to pay? We will explicitly compare the expected loss in value in going from this *simultaneous choice* mechanism to the *sequential choice* mechanism that generates the stopping problem.

1.1. Related work

In the classical secretary problem (CSP), a decision-maker has to hire one applicant out of a pool of n applicants who will appear sequentially. Again, the decision-maker must decide immediately upon seeing an applicant whether to hire her or not. The key difference between secretary problems and search problems, as Ferguson (1989) notes, is that in secretary problems “the payoff depends on the observations only through their relative ranks and not otherwise on their actual values.” The most studied types of secretary problems are games with 0–1 payoffs, with the payoff of 1 being received if and only if the decision-maker hires the best applicant. The decision-maker’s optimal policy is thus one that maximizes the probability of selecting the best applicant.

A historical review of the early literature on secretary problems, including important references, can be found in the paper by Gilbert and Mosteller (1966), as can solutions to many extensions of the basic problem, including the search problem (with finite and known n and no search costs) for various different distributions over the values of applicants. Many interesting variants of the original problem, mostly focusing on maximizing the probability of hiring the best applicant, have appeared in intervening decades. For instance, Cowan and Zabczyk (1978) introduce a continuous-time version of the problem with applicants arriving according to a Poisson process, which is closely related to the continuous-time problem we describe in Section 4. Their work has been extended by Bruss (1987) and by Kurushima and Ano (2003). Stewart (1981) studies a secretary problem with an unknown number of applicants which is also related to the problem we consider, but differs in the sense that he assumes n to be a random variable and the arrival times of offers to be i.i.d. exponential random variables, so that the decision-maker must maintain a belief distribution on n in order to optimize.

There has been considerable interest in explicitly modeling two-sided search and matching problems in the economics community. In particular, Burdett and Wright (1998) study two-sided search with nontransferable utility, which is relevant to our model because we assume exogenous offer values, implying that an employer cannot make her offer more attractive