Journal of Behavioral Decision Making
J. Behav. Dec. Making, **19**: 229–250 (2006)
Published online 20 June 2006 in Wiley InterScience
(www.interscience.wiley.com). **DOI**: 10.1002/bdm.521

Experimental Studies of Sequential Selection and Assignment with Relative Ranks

J. NEIL BEARDEN[1]*, AMNON RAPOPORT[1,2] and RYAN O. MURPHY[3]
[1]*University of Arizona, USA*
[2]*Hong Kong University of Science and Technology, Hong Kong*
[3]*Columbia University, USA*

ABSTRACT

We study a class of sequential selection and assignment problems in which a decision maker (DM) must sequentially assign applicants to positions with the objective of minimizing expected cost. In modeling this class of problems, we assume that on each period the DM is only informed of the *rank* of the present applicant relative to the applicants that she previously observed and assigned. We first present the optimal decision policy that we subsequently use as a normative benchmark, and then report results from three experiments designed to study sequential assignment behavior. In comparing the aggregate results from all three experiments to the optimal decision policy, we identify a systematic bias, called the *middleness bias*, to over-assign applicants to intermediate positions. The results also reveal a strong bias for early applicants to be over-assigned to important positions. Copyright © 2006 John Wiley & Sons, Ltd.

KEY WORDS selection and assignment; secretary problem; optional stopping; experiments; normative benchmarks; bias; middleness bias; framing effects

INTRODUCTION

Consider the problem faced by safety officials who must respond to catastrophic events, such as earthquakes, hurricanes, and suicide bombings. Casualties are likely to be massive and resources limited. Further, the extent and nature of the casualties is not known instantaneously; rather, they are revealed sequentially in realtime. Consequently, the assignment of resources to victims encountered early must be made in the absence of knowledge of what resources future, perhaps more badly injured, victims might need. The decision problems faced by these officials are obviously difficult. In the current paper, we examine behavior in a formal assignment problem and ask: How well do people make sequential assignment decisions? To help

* Correspondence to: J. Neil Bearden, Brain Imaging Analysis Center, Duke University, 163 Bell Research Building, Box 3918, Duke University Medical Center, Durham, NC 27710, USA. E-mail: jneilb@gmail.com

Contract grant/sponser: AFOSR/MURI; contract grant/number: F49620-03-1-0377.

determine the answer, we first introduce a class of formal assignment problems for which optimal assignment policies can be derived.

Derman, Lieberman, and Ross (1972) formulated the following sequential stochastic assignment problem: A decision maker (DM) observes a sequence of n jobs, one at a time, for which there are m machines available. Associated with each job j, $j = 1, 2, \ldots, n$, is a *value*, which is a random variable X_j that takes on the values x_j. When a job arrives and its value x_j is revealed, the DM must decide whether to assign it to one of the available machines (selection). If she decides to do so, she must next decide which of the m machines to assign it to (assignment). Assignment decisions are irrevocable; once a job is assigned it may not be reassigned to another machine. Costs or rewards reflect the quality of the match between the type of job (which is revealed upon observation) and the type of machine (which is known *a priori*). The objective criterion is the minimization of expected cost, which is operationally the same as maximizing fit.

This formulation is quite general; it describes the essential characteristics of sequential selection and assignment problems in other domains. For example, the m ''machines'' may correspond to m organizations (e.g., academic departments) that differ from one another in their specialization, expectations, and requirements, and the n ''jobs'' to n applicants (e.g., post-doctoral fellows) who also differ from one another on several attributes. Interviewing and subsequently assigning them one at a time, the DM seeks to achieve the best match between the organizations and applicants. For another example, the DM may correspond to a plant manager, the m machines to m subordinates, and the n jobs to n projects that arrive one at a time. In assigning projects to subordinates, the manager may wish to maximize the fit between the duration, importance, difficulty, and other requirements of the projects and the qualifications, expertise, and experience of the available subordinates. David and Yechiali (1985) have recognized yet another application of the model to the allocation of kidney transplants. In this application, there are m candidates for kidney transplants and n kidneys that become available one at a time. Kidneys and candidates have their own types (e.g., tissue type of the candidate and of the organ donor). The candidate types are known ahead of time, but the organ types are only revealed upon their arrival. Once a kidney is assigned, the reward is a function of the match between the kidney's and the candidate's types. The DM's objective is to maximize the total expected reward (Su & Zenios, 2002). In our experimental set-up, we used yet another example of sequential triage assignment decisions. The m machines correspond to m hospitals, the n jobs to n injured patients who wait for sorting and assignment, and the DM to a triage officer who must sequentially assign patients to hospitals. Each patient j is assigned a score x_j that measures the severity of his injuries, and each hospital k, $k = 1, 2 \ldots, m$, has an associated cost, b_k, that represents the quality of care, distance of the hospital from the DM, or some combination of the above. The DM's objective is to assign patients to hospitals, one at a time, in an attempt to minimize the total expected cost over the n patients, $\sum_{j=1}^{n} x_j b_k$.

The sequential assignment model of Derman, Lieberman, and Ross (1972) has been extended in a number of different directions. Albright and Derman (1972) studied the asymptotic behavior of the optimal assignment policy derived by Derman et al. Albright (1974) considered an analogue of the discrete time model in continuous time, where jobs are assumed to arrive according to a non-homogenous Poisson process. Sakaguchi (1984) considered the case in which the number of jobs is unknown, and Kennedy (1986) studied yet another case of serially correlated job arrivals. See, for example, David and Yechiali (1990), David (1995), and Su and Zenios (2002) for more recent extensions and applications. In deriving the optimal decision policy for their class of sequential assignment problems, Derman et al. (1972) assumed that the n jobs are compared to one another in terms of a *single* dominant attribute whose distribution function is known to the DM before the assignments commence. In particular, they assumed that associated with each job j is a single X_j, which is an i.i.d. variable randomly drawn from a common univariate probability distribution $f(x \mid \theta)$ with parameters θ. Their approach is known as the *full information* approach (Gilbert & Mosteller, 1966).

In the next section we discuss what we consider to be the limiting features of this approach. To do so, we first present and discuss a simpler problem—often referred to as *sequential search* or optional stopping—and characterize the optimal decision policy derived under the full information assumption. Our criticism of the full information approach in this case applies *ipso facto* to the optimal decision policy of the more complicated sequential selection and assignment problem. Following Chun and Sumichrast (in press, hereafter CS), we next describe the *no information* approach to the sequential assignment problem, in which the assignment decision of the *j*th job is solely based on its relative rank among the jobs observed and assigned thus far. (See also, Hornberger & Ahn, 1997, who present the no information approach for deriving optimal decision policies for the assignment of transplant organs.)

Under the no information approach, the optimal policy for the sequential assignment problem is an extension of the celebrated secretary problem (see, e.g., Freeman, 1983; Ferguson, 1989; Gilbert & Mosteller, 1966; Samuels, 1991). Because the problem is not very familiar to researchers in the social sciences, we include a brief presentation of the classical secretary problem (CSP) before presenting a detailed formulation of the sequential assignment problem based on relative ranks. Next, we describe and illustrate the optimal decision policy for the no information case, which is due to CS, that applies dynamic programming to numerically compute the optimal assignment of each job *j*. We use the optimal decision policy as a benchmark for the analysis of the observed assignments.

After formally introducing the assignment problem, we present three experiments designed to study selection and assignment decisions of financially motivated subjects in a series of computer-controlled tasks. These experiments are aimed at uncovering systematic and replicable patterns of sequential assignment behavior. We are interested in both the overall quality of the assignment decisions and the ways in which the assignments may depart from optimality. Observed assignment decisions are compared to the optimal decisions to answer these questions. Implications of the results and future research directions are discussed in the final section.

LIMITING FEATURES OF THE FULL INFORMATION APPROACH

To illustrate the shortcomings of the full information approach, we apply it to the simpler case of searching for the best of *n* objects. A sample of *n* observations is randomly drawn from a population with continuous cumulative distribution $F(\theta)$ with parameters θ. Knowing $F(\theta)$ and *n*, the DM is informed of the value of each observation *x*, whereupon she must choose whether to select this observation. Selecting an observation terminates the search. The DM's objective is to maximize the probability of selecting the *largest* observation in the sample. Because the distribution $F(\theta)$ is known and the largest observation in the sample remains the largest under all monotone transformations of *x*, there is no loss of generality in assuming that $F(\theta)$ is the standard uniform $F(\theta) = x$ and density function $f(x \mid \theta) = 1$ on $0 \leq x \leq 1$ (Gilbert & Mosteller, 1966).

The optimal decision policy (Gilbert & Mosteller, 1966) consists of a vector $s = (s_1, s_2, \ldots, s_n)$ of cutoff values satisfying $s_{j+1} \geq s_j$. These cutoff values, one per stage, are computed recursively. The DM should stop the search and take the *j*th observation if $x_j \geq s_j$, and continue, otherwise. As the cutoff values are non-increasing in the draws, the optimal policy allows the DM to pass over an observation at stage *j* but stop and select an observation with a smaller value at a later stage *j*', *j*' > *j*. Selected values of the probability of winning, *p*(win), are presented in Table 7 of Gilbert and Mosteller. An early experiment designed to test the optimal policy when $f(x \mid \theta)$ is a normal distribution with known parameter values was conducted by Rapoport and Tversky (1970).

The full information approach is rather appealing as it yields point predictions that are experimentally testable. It also yields a qualitative prediction $s_{j+1} \geq s_j$) that is testable and intuitively compelling. However, it suffers from three major shortcomings. The first limiting feature is the assumption that $f(x \mid \theta)$ is known to the DM *before* the search commences. Under this *no learning* assumption, the DM is supposed to numerically determine the cutoff values s_j before taking the first observation and extract no information from the observations. But in reality decision behavior is adaptive, and people adjust their criteria for terminating the search as they gather more information.

A second and related shortcoming is that in the full information case, the parameters θ of $f(x \mid \theta)$ are assumed to be known by the DM with precision. It is a rare case in the real world that a DM knows with certainty the distribution from which her options are drawn. We are hard pressed to think of examples, even in the case where the sequentially observed objects vary on a single dimension (e.g., sequential search for a good wage). Nor can robustness and "flat maxima" arguments be invoked to justify the full information assumption. The values of s_j are located at the extreme tail of the density function and are consequently very sensitive to even minor changes in the parameter values θ.

The third, and in our opinion the most limiting, feature of the full information approach, which is also shared by the model of Derman et al., is the assumption of single dimensionality. As pointed out by Seale and Rapoport (2000), Zwick, Rapoport, Lo, and Muthokrishnan (2003), and others, there exist sequential search situations where "jobs" can be evaluated and compared to one another in terms of a single attribute (price is the best example) whose distribution function may be known approximately or can be learned with experience. But in most other search situations it may not be possible to specify the value function because the objects have *multiple*, possibly conflicting, attributes. In a sequential search for an academic position, offers are usually compared to one another in terms of several dimensions like starting salary, fringe benefits, availability of research facilities, geographical location, education facilities for the children, etc. Applicants for a secretarial position vary from one another in their experience, knowledge of computers, social facility, and so on. The tissue type of kidneys for transplantation consists of six proteins, and perfect matching occurs when the proteins of the donor and recipient are identical (Su & Zenios, 2002). One may possibly propose to extend the full information approach to the multi-attribute case by postulating some multi-dimensional distribution function. However, it is difficult, if not impossible, to identify the overall value function $f(x \mid \theta)$ of multiple attributes x and estimate its parameters θ. A major attractive feature of the no information approach to modeling sequential search and sequential assignment is that no assumption is made about the value function.

SEQUENTIAL OBSERVATION AND SELECTION PROBLEMS

The "Secretary Problem" (SP) was first discussed in the late 1950's. Since then, it has been extended and generalized in many different directions and given rise to a "field" of study in applied probability (see reviews by Ferguson, 1989; Freeman, 1983; Samuels, 1991). The SP has also stimulated experimental, rather than theoretical, research on sequential search by Bearden, Murphy, and Rapoport (2005), Corbin, Olson, and Abbondanza (1975), Seale and Rapoport (1997, 2000), Stein, Seale, and Rapoport (2003), and Zwick et al. (2003). Although there are many variations and extensions of the problem, in its simple form the Classical Secretary Problem (CSP) satisfies the following assumptions:

1. There is only a single position to be filled.
2. The number of applicants for this position, n, is finite and known to the DM before the search commences.
3. The n applicants are interviewed (observed, inspected) by the DM, one at a time, in a random order. Consequently, each of the $n!$ orderings is equally likely.
4. The DM can rank order all the n applicants in terms of their quality (desirability, attractiveness, suitability) from best to worst with no ties. Her decision to either accept or reject an applicant in any given period only depends on the relative rank of the applicant with respect to all the applicants that have already been interviewed.
5. Once rejected, an applicant cannot later be recalled.
6. The DM's payoff is 1 if she chooses the best of the n applicants (in terms of their absolute ranks), and is 0 otherwise. (This objective function is known as "nothing but the best.")

Managerial and marketing sequential decision and selection problems modeled by the CSP include, among others, selling an asset in the open market, hiring an employee for a job, purchasing some product on the market, assigning a job to a single machine, choosing an investment alternative, or searching to rent an apartment (CS, in press; Zwick et al., 2003). Often, these search situations are characterized by multi-attribute decision alternatives (''applicants'') that are presented and inspected sequentially over time. Following an evaluation of the alternative, the DM may either choose it and thereby terminate the search or reject it and continue searching for at least one more period. The decision is based on the overall relative rank of the alternative under consideration with respect to its various attributes and the possibility of finding a better alternative if the search is continued.

Previous empirical research on the sequential selection problems has most often shown that DMs stop searching too early (e.g., Seale & Rapoport, 1997, 2000; Rapoport & Tversky, 1970). This research is suggestive as to the biases that might be observed in DMs in contexts, such as asset markets, where the selection decision involves when to sell or when to buy an asset; however, it does not speak to the much broader class of sequential decision problems in which the DM must make more consequential decisions for each encountered alternative. Not purchasing an asset at some time t, when the price is relatively high, seems considerably less troubling than assigning someone who is moderately wounded to a facility whose beds are valuable to those who have grave wounds. Our objective in the current paper is to evaluate decision behavior in this broader context, where each decision really counts.

For most applications, the CSP is too restrictive. Consequently, every one of the assumptions above has been relaxed. Since the generalizations of the CSP are of no importance to the present research, we do not discuss them here. For reviews of these extensions and their optimal decision policies, see the references above. The sequential assignment problem that we describe next encompasses the CSP under a specific formulation—and is, therefore, of greater applicability.

In designing the experiments, we opted to couch the experimental task in terms of ''patients'' and ''hospitals'' rather than in the more abstract terms of ''applicants'' and ''positions.'' The sequential triage model that we describe below is clearly stylized and fails to capture some essential aspects of the current guidelines for triage assignments. We use the term ''triage'' in the common sense of classifying patients according to medical needs and matching of these patients to available care resources. Our sequential model recognizes that for optimal allocation of resources in the treatment of man-made or natural disasters it is useful to decide as early as possible which patient will benefit most from transport to a dedicated trauma center (Bond et al., 1997). However, the common guidelines for triaging patients dictate a four-color classification system, where red stands for serious but salvageable life threatening injury/illness, yellow stands for moderate to serious injury/illness, green stands for ''walking-wounded'' patients who can stand a delay in transportation, and black indicates victims who are found to be clearly deceased at the scene or having fatal injuries. These four ''values'' clearly differ from the information structure that we assume. But, then, our purpose is not to mimic current guidelines directed at well-trained professionals but, rather, to simplify an already complicated sequential decision task presented to untrained subjects.

In the formulation of the *sequential selection and assignment problem* (SSAP) as a triage assignment, a set of h hospitals, their capacities, and their costs are assumed to be known by the DM. A set of n patients arrive at an emergency station in a random order, one at a time, and their ''values'' (degree of injury), denoted by the random variable X_j, are treated as independent. A major assumption of our formulation sharply differentiates it from the parametric selection and assignment problem. Our formulation is based on the assumption that the DM, who has to assign patients to hospitals, can only rank patient j in terms of the severity of her injuries *relative* to the patients who have already been assigned to the hospitals; she cannot determine the severity of the patient's injuries relative to those patients who are yet to arrive.

Each time the DM inspects a patient j, she must decide to which of the h hospitals not yet filled to capacity the patient should be assigned. (The script j corresponds to the number of not-yet-assigned patients.) Each hospital position has an associated ''cost'' denoted by b_k. The cost b_k may represent the quality of care and consequently the chances of full recovery, the time of evacuation which increases in the geographical distance of the hospital from the scene of the disaster, or some combination of the above. Each patient j has a ''value'' score x_j—a realization of the random variable X_j—that reflects the severity of her injury. The cost of assigning patient j to hospital position k is given by $f(x_j)b_k$. (In the classical selection and assignment problem, $f(\bullet)$ is an identity function; $f(\bullet)$ for the version we study will be discussed below.) The DM's objective is to assign each patient j to one of the h hospitals in order to minimize the *total expected cost* $(\sum_j f(x_j)b_k)$. This objective would easily be obtained if the patients were to be assigned to hospitals simultaneously: one merely assigns the smallest $f(x_j)$ to the largest b_k, the second smallest $f(x_j)$ to the second largest b_k, and so on. When the patients must be assigned sequentially, and the assignment is based only on their overall value relative to those patients who have already been assigned, the problem becomes non-trivial.

Terminology and notation

The SSAP assumes that there are n objects (e.g., jobs, patients) that have to be assigned to m positions (e.g., machines, hospitals). In the triage context that we study, a ''position'' is interpreted as a bed in a hospital, and an ''object'' as a patient. Thus, $m \geq h$ with equality holding only in the case that each hospital has only a single bed. The number of patients n could be larger or smaller than the number of positions m. If $n > m$, then some patients will necessarily be rejected and receive no treatment (or, alternatively, will receive treatment on the spot and then be released). We assume that there is no cost associated with treating these patients. On the other hand, if $n < m$, then some of the beds will remain unoccupied. Hereafter, we shall assume with no loss of generality (see CS, in press) that $n = m$. Each hospital s, $s = 1, \ldots, h$, has its own capacity m_s; thus, $m = \sum_s m_s$. In addition, each hospital position k has an associated cost b_k. Suppose that $h = 3$, $m_A = 3$, $m_B = 2$, $m_C = 1$, and that hospitals A, B, and C have associated costs of 150, 50, and 25, respectively. That is, within a hospital, all positions have the same cost. An example of three hospitals with their corresponding capacities and costs is shown in Table 1. We adopt the convention of ordering hospital positions in ascending order from the first position of the highest cost hospital to the last position of the lowest cost hospital. This will be useful below when we describe the optimal solution formally.

Suppose that among the n patients, the DM has already observed and assigned to hospitals $n - j$ patients, and therefore j more patients are to be assigned. Upon observing patient j, the SSAP assumes that the DM can only determine the relative rank of this patient among the $n - j + 1$ patients. Thus, the *relative rank* of patient j, which we denote by R_j, is simply her rank among the $n - j + 1$ patients that have already been observed by the DM. In contrast, the *absolute rank* of patient j, which we denote by A_j, is her rank among *all* the n patients (including the ones who have not yet been observed by the DM). In accordance with Assumption 2 of the CSP, all $n!$ possible sequences of patient arrivals are assumed to be equally likely. Both the absolute and relative ranks are random variables whose realization depend on which of the $n!$ sequences is presented. We shall denote their realizations by r_j and a_j, respectively. Under the assumption that all $n!$ sequences of patient arrivals are equally probable, the random variable R_j follows the discrete uniform distribution

$$\Pr[R_j = r_j] = 1/(n - j + 1), \quad \text{for} \quad r_j = 1, 2, \ldots, n - j + 1$$

Table 1. Three hospitals with a total capacity of six beds

Hospital A (Cost = 150)	Hospital B (Cost = 50)	Hospital C (Cost = 25)
Bed 1 (1)	Bed 1 (4)	Bed 1 (6)
Bed 2 (2)	Bed 2 (5)	
Bed 3 (3)		

Numbers in parentheses represent position number.

Chun (1996) has shown that the conditional probability of A_j, given that $R_j = r_j$, is given by

$$\Pr[A_j = a_j \mid R_j = r_j] = \frac{\binom{a_j - 1}{r_j - 1}\binom{n - a_j}{n - j + 1 - r_j}}{\binom{n}{n - j + 1}}, \quad \text{for} \quad a_j = r_j, r_j + 1, \ldots, r_j + j - 1$$

Thus, given the observed value of r_j, the DM can compute the distribution of A_j as well as its expected value.

Example 1. To illustrate the difference between relative and absolute ranks, and the way relative ranks are derived from absolute ranks, consider the following example.

	Patient Position								
Rank	1st	2nd	3rd	4th	5th	6th	7th	8th	9th
a_j	6	2	5	3	8	1	9	7	4
r_j	1	1	2	2	5	1	7	6	4

There are nine patients in this example with absolute ranks (a_j) from 1 to 9. The absolute ranks reflect the severity of injury, with rank 1 assigned to the most severely injured patient and rank 9 to the least severely injured patient. The patient with absolute rank 6 is the first to be brought to the triage station, the patient with absolute rank 2 is the second, and so on. The DM can only assess the relative ranks (bottom row). Thus, the first patient is clearly the most severely injured of the ones she has seen so far, and is assigned a relative rank 1. The second patient whose injuries are more severe than the first (as absolute rank 2 is higher than rank 6) is consequently assigned the relative rank 1. The third patient (with absolute rank 5) is more severely injured than the first (with absolute rank 6) but less severely injured than the second (with absolute rank 2). Therefore, she is assigned the relative rank 2. This process of assessing relative ranks continues until the ninth and last patient is assigned the relative rank 4, which necessarily corresponds to her absolute rank.

After the DM observes a patient with relative rank r_j, she must decide to which hospital to assign the patient. Associated with the kth position (bed in the hospital) is the cost b_k, $k = 1, 2, \ldots, n$. With no loss of generality, we assume that the positions are arranged such that $b_1 \geq b_2 \geq \ldots \geq b_n$. For example, suppose that there are three major hospitals A, B, and C, where the maximum capacities are 2, 3, and 4, positions, respectively. Then, the costs of these hospitals would be $b_1 = b_2 \geq b_3 = b_4 = b_5 \geq b_6 = b_7 = b_8 = b_9$. If a patient with absolute rank a_j is assigned to the kth position, then the SSAP assumes that the cost $a_j b_k$ is incurred. (In the example in Table 1, $b_1 = b_2 = b_3 = 150, b_4 = b_5 = 50$, and $b_6 = 25$.)

THE OPTIMAL ASSIGNMENT POLICY

The DM's objective is to ensure that the available medical resources are used to treat the greatest number of patients in need of these resources. In our formulation, this is tantamount to minimizing the total expected cost. Denoting a particular assignment of the n patients to the m positions by Γ, where Γ is a permutation of the integers from 1 to n, the DM wishes to find $\Gamma^* = \underset{\Gamma}{\operatorname{argmin}}\left(\sum_{j=1}^{n} a_j b_k\right)$. However, since the patients arrive sequentially, Γ^* cannot easily be found. Instead, given the constraints of the problem, the optimal decision policy minimizes the expected value of the total cost: $E\left(\sum_{j=1}^{n} a_j b_k\right)$. A decision policy dictates to which open position to assign the jth patient when her relative rank is r_j.

Using dynamic programming, CS constructed an optimal assignment scheme that takes into account the severity of injury and the location, capacity, and medical facilities available at each of h hospitals. To describe it, define the *critical relative rank* $r^*_{i,j}$ at stage j by

$$r^*_{i,j} = \begin{cases} 0 & \text{if } \ i = 0 \\ \dfrac{1}{n-j+3} \displaystyle\sum_{r_{j-1}=1}^{n-j+2} \min[\max(r_{j-1}, r^*_{i-1,j-1}) r^*_{i,j-1}] & \text{if } \ 1 \le i \le j \\ \infty & \text{if } \ i = j \end{cases} \tag{1}$$

Equation (1) shows that the critical relative ranks do not depend on the costs b_k except through their order. They satisfy the condition

$$0 < r^*_{0,j} < r^*_{1,j} <, \ldots, < r^*_{j,j} < \infty$$

Based on the critical relative ranks, the optimal assignment policy at each stage j is to assign the patient of relative rank r_j to the ith position if $r^*_{i-1,j} < r_j < r^*_{i,j}$

Example 2. Consider the same parameter values as in Example 1 with $n = 9$ patients who arrive in the order

$$\{a_j\} = (6, 2, 5, 3, 8, 1, 9, 7, 4)$$

As shown in Example 1, the corresponding relative ranks observed one at a time are

$$\{r_j\} = (1, 1, 2, 2, 5, 1, 7, 6, 4)$$

The critical relative ranks for Example 2, computed from Eq. (1), are presented in Table 2.

Continuing with this example, assume that there are three hospitals, A, B, and C, with a corresponding number of positions (beds) 2, 3, and 4. Recall that j indicates the number of patients remaining to be assigned and i indicates the position. Consider the first patient ($j = 9$). This patient has absolute rank 6 and the corresponding relative rank 1 is between $r^*_{4,9} = 0.9153$ and $r^*_{5,9} = 1.0847$ at stage 9. Therefore, the first patient is assigned to position 5 (in hospital B). Following this assignment, hospital B is left with only two beds. Consider next the second patient ($j = 8$) of absolute rank 2 and the corresponding relative rank 1, which is

Table 2. Critical relative ranks for the sequential assignment problem in Example 2. Under the optimal policy, the DM should assign patient j to position i whenever the patient's rank r_j is such that $\boldsymbol{r^*_{i-1,j} < r_j < r^*_{i,j}}$

State	$j=9$	$j=8$	$j=7$	$j=6$	$j=5$	$j=4$	$j=3$	$j=2$	$j=1$
$i=0$	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
$i=1$	0.4992	0.7488	0.9984	1.3307	1.7460	2.3056	3.1111	4.5000	∞
$i=2$	0.6390	0.9585	1.4169	1.9180	2.5873	3.5000	4.8889	∞	
$i=3$	0.7487	1.2460	1.7836	2.5000	3.4127	4.6944	∞		
$i=4$	0.9153	1.5000	2.2164	3.0820	4.2540	∞			
$i=5$	1.0847	1.7540	2.5831	3.6693	∞				
$i=6$	1.2513	2.0415	3.0016	∞					
$i=7$	1.3610	2.2512	∞						
$i=8$	1.5008	∞							
$i=9$	∞								

Table 3. Optimal assignment to hospitals in Example 2

Hospital A	Hospital B	Hospital C
1	2	5
4	3	7
	6	9
		8

between $r^*_{2,8} = 0.9585$ and $r^*_{3,8} = 1.2460$ at stage $j = 8$. This patient is assigned to position 3 (also in hospital B). This leaves hospital B with only a single bed. Consider next the third patient ($j = 7$) of absolute rank 5 and corresponding relative rank 2, which is between $r^*_{3,7} = 1.17836$ and $r^*_{4,7} = 2.2164$. Because Hospital B has only one position left, the fourth available position is now the first in Hospital C. Therefore, patient 3 is assigned to position 6 (which is in Hospital C). Continuing in this way, the optimal assignment of the nine patients to the nine positions is

$$\mathrm{P} = (5, 3, 6, 4, 9, 1, 8, 7, 2)$$

or in terms of the three hospitals:

$$\mathrm{A} = \{1, 4\}, \quad \mathrm{B} = \{2, 3, 6\}, \quad \mathrm{C} = \{5, 7, 9, 8\}$$

This optimal assignment is depicted in Table 3 that shows the absolute ranks of the nine patients. Note that this is not the best possible assignment. If all nine patients were to be assigned *simultaneously* rather than sequentially, then the optimal assignment procedure would be to assign patients with absolute ranks 1 and 2 to hospital A, patients 3, 4, and 5 to hospital B, and the remaining four patients with absolute ranks 6, 7, 8, and 9 to hospital C.

Alternative Assignment Procedures

Suppose that the costs per position in hospitals A, B, and C are 250, 100, and 40 units, respectively. Then, the total cost associated with the optimal assignment policy in Table 3 is

$$250 \times 5 + 100 \times 11 + 40 \times 29 = 3510$$

If, for example, patients are assigned to the three hospitals in the order of their arrival, the resulting assignment is

$$\mathrm{A} = \{6, 2\}, \quad \mathrm{B} = \{5, 3, 8\}, \quad \mathrm{C} = \{1, 7, 9, 4\}$$

and the total cost associated with this assignment procedure is

$$250 \times 8 + 100 \times 16 + 40 \times 21 = 4440$$

The total cost of this assignment procedure is 26.5% higher than the one associated with the optimal policy. The worst possible assignment procedure is when the least injured patients are assigned to hospital A and the most seriously injured to hospital C:

$$\mathrm{A} = \{8, 9\}, \quad \mathrm{B} = \{5, 6, 7\}, \quad \mathrm{C} = \{1, 2, 3, 4\}$$

The total cost associated with this worst possible assignment procedure is

$$250 \times 17 + 100 \times 18 + 40 \times 10 = 6450$$

This last assignment almost doubles the total cost associated with the optimal assignment.

EXPERIMENTS 1 AND 2: INITIAL EXPERIMENTAL STUDIES OF THE SSAP

Overview

We had financially motivated subjects perform the SSAP in a computerized decision task that simulates sequential triage assignment following a fictitious disaster. Repeating the task for multiple trials, our design allows us to assess the subjects' assignments relative to the optimal decision policy and the changes in behavior, if any, over time. We first conducted an experiment in which the number of patients was relatively small ($n = 12$); then, to test the generality of our findings, we conducted a second experiment in which this number was doubled ($n = 24$).

Method

Subjects

Thirty-six financially motivated subjects participated in Experiment 1 and 39 in Experiment 2. The experiments were conducted in a large computerized laboratory. All the subjects were students recruited by advertisements asking for volunteers to participate in a decision-making experiment with payoffs contingent on performance. The mean payoff per session, which typically lasted 1 to 1.5 hours, was about \$25 (minimum \$15, maximum \$40). In addition to the monetary payoff, the subjects could request class credit for their participation.

Procedure

Each subject was provided with a written (hard copy) of the instructions for the SSAP task. The general task was described as follows:

> The experiment is concerned with the decision making process following a natural or man-made disaster (e.g., earthquake, fire, terrorist attack). When the number of people injured overwhelms the medical resources of area hospitals, there is typically a central emergency allocation system that handles the flow of patients. We have designed an experiment that simulates emergency situations whose purpose is to study how people make sequential assignment decisions.

The instructions placed special emphasis on the computation of relative rank information for the patients and on the way in which assignment costs were to be determined. Several examples of the SSAP were presented in detail. The first involved a simple assignment problem in which three patients had to be assigned to three hospitals, each of which had one position. The costs of each of the six possible assignments were presented to the subject, along with a detailed explanation of how they were computed. Two more examples were then presented to the subjects, each involving six patients. Finally, the instructions explained the payoff scheme. Specifically, the subjects were told that they would randomly select three trials for payment at the conclusion of the task by drawing numbers from a hat. Their payments for each trial would be based on their net earnings for these trials (Net Point Earnings = Endowment Points − Cost Incurred on Trial). The total points earned on these three trials was converted to a cash payment at a rate of \$1 per 1000 points.

After reading the instructions, two practice problems were presented to verify the subject's understanding of the task. The experimental problems were presented once the subjects successfully completed these practice problems. Each subject completed 60 trials (replications) of the SSAP.

Manipulations

We conducted two separate experiments that varied the number of patients: $n = 12$ (Experiment 1) and $n = 24$ (Experiment 2). The number of hospitals was fixed at $h = 4$. The number of openings for the hospitals in Experiment 1 was $m_1 = 3$, $m_2 = 4$, $m_3 = 3$, and $m_4 = 2$; for Experiment 2 it was $m_1 = 6$, $m_2 = 7$, $m_3 = 6$, and $m_4 = 5$. The costs for all positions within a particular hospital were the same in both experiments. Hospital 1, the most costly one, had an associated cost of 100, and Hospitals 2, 3, and 4 had associated costs of 50, 25, and 10, respectively. Participants in Experiment 1 received an endowment of 10 000 units ($10) on each trial; those in Experiment 2 received 20 000 units ($20).

Each trial consisted of a SSAP with n (12 or 24) patients, each presented separately on a different *period*. Each trial was structured in the same way: The subject was sequentially shown the relative rank for each patient and required to assign the patient to an open hospital position. Subjects assigned a patient to a hospital by clicking on the button for that hospital. The patient's ID—simply her number in the sequence of n patients—was then placed in the open position for that hospital and remained there until the termination of the trial. Each hospital button gave the hospital name and the cost associated with that hospital (e.g., ''Hospital 1 (cost = 100)''). Hospital buttons were ''dimmed'' and could not be clicked if the hospital was full. Once placed in a particular hospital, a patient could not be moved to another hospital.

To allow comparison between subjects, in both experiments we generated a *sample* of 60 different random sequences of patients (out of a *population* of 60! sequences). In each experiment, each subject was presented with the same 60 sequences in the sample, but the presentation order was randomized across subjects.

Results

Overview

Using Equation 1, for each trial and each subject separately we determined the hospitals to which the n patients *should be* assigned. Recall that the optimal assignment algorithm instructs the DM into which *open position* a patient should be put. In our task, however, subjects could only assign patients to *open hospitals* (i.e., open hospitals with open positions); they *could not* select the bed in a particular hospital a patient should be placed. Therefore, we simply mapped the optimal assignment positions (integers from 1 to n) given by Equation 1 to hospitals 1 through 4. For example, if the third open position is in Hospital 2, then an assignment to (any position in) Hospital 2 is optimal; assignments to any other hospital are not.

We define *unconditionally optimal assignment* of the n patients as the assignment dictated by Equation 1. Given the arrival ordering of the patients, these can be computed *a priori* without consideration of the subject's previous assignments. Unconditionally optimal assignments have no place for subject errors. Once a subject assigns a single patient erroneously, all the remaining patients are out of line. To allow for the presence of errors in assignment, we computed for each subject separately *conditionally optimal assignments* that explicitly recognize the subject's previous assignments. Specifically, the assignment of patient j is *conditionally optimal* if and only if she is assigned to the optimal open hospital (as determined by Equation 1), given the availability of open hospital positions, which are determined by the vacancies left after the patients previously assigned by the subject. Unconditional and conditional assignments coincide if and only if the subject adheres to the optimal policy on each period. All the results reported below are based on comparisons to the conditionally optimal assignments. Necessarily, they are conducted on the individual level.

Quality of patient assignments

The subject's performance in the SSAP was measured by the root-mean-squared-deviation (RMSQ) of the actual assignments of the n patients from the (conditionally) optimal assignments. Denoting the hospital to

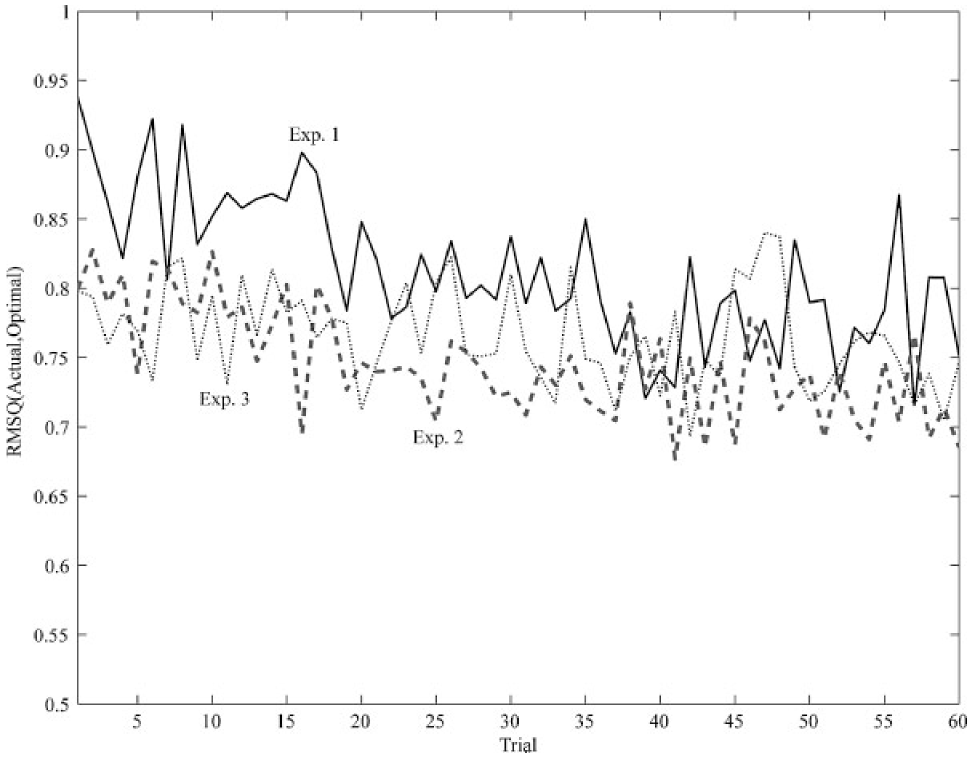

Figure 1. Mean RMSQ error of assignments across 60 experimental trials for Experiments 1, 2, and 3.

which the jth patient would be optimally assigned s_j^*, and the hospital to which the patient was actually assigned s_j, the subject's performance can be assessed by the measure

$$RMSQ(Actual, Optimal) = \left[n^{-1} \sum_{j=1}^{n} (s_j - s_j^*)^2 \right]^{.5} \tag{2}$$

which is 0 for a perfectly optimal assignment solution (conditional or unconditional) and grows as the quality of the assignment decreases.

Individual RMSQ scores were computed for each subject and each trial. Figure 1 displays the mean RMSQ scores across the 60 experimental trials for Experiments 1, 2, and 3. (The results for Experiment 3 will be presented later.) Three characteristics of the performance scores are noteworthy. First, the mean RMSQ score is non-zero across all 60 trials, indicating suboptimal solutions to the SSAP. Secondly, however, we observe that the mean RMSQ scores in Experiments 1 and 2 slowly decrease across trials, providing evidence for learning. To establish an alternative baseline for performance (i.e., one other than optimal), we computed the expected RMSQ under fully random assignment to open positions. For the configurations used in Experiments 1 and 2 the expectations are 1.27 and 1.42, respectively, both of which significantly exceed the observed RMSQ in each experiment. Despite the difficulty of the assignment task—the computations required for optimal assignment are lengthy—the subjects in both experiments seem to have developed decision rules that allowed them to make reasonably good assignments. Additionally, the subjects in all three experiments learned to improve their assignments with experience. Finally, there is considerable

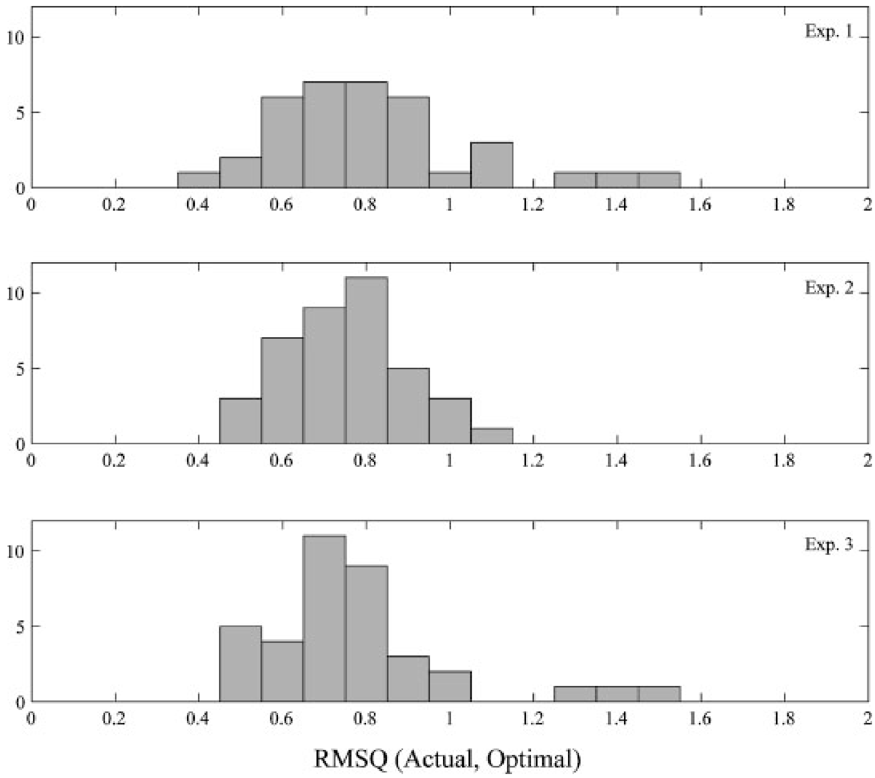

Figure 2. RMSQ error of assignments of individual subjects for Experiments 1, 2, and 3.

subject heterogeneity in the RMSQ measure. Frequency distributions of individual mean RMSQ scores (i.e., RMSQ scores averaged over all 60 assignment problems for each subject) in Experiments 1, 2, and 3 are exhibited in Figure 2.

Figure 3 displays the mean difference between optimal and actual assignment positions across the *n* periods for each subject in Experiments 1 (top panel), 2 (middle panel), and 3 (bottom panel). Although there is some variability in the plots, the general pattern across the *n* positions is quite similar for most subjects. Specifically, $s_j - s_j^*$ is quite often positive for the first patient in Experiments 1 and 2 (but not in Experiment 3). This indicates a slight bias to place the first patient in a worse hospital than she ought to be placed in. The differences then tend to be negative for the next few patients before converging to 0 for roughly the last half of patients. For the configuration used in Experiment 1, the expected value of $s_j - s_j^*$ under random assignment to open positions is positive for the first patient and negative for patients 2 through 11. The expectation is zero for the final patient under all possible assignment strategies, using the conditionally optimal assignment. The expectation under random assignment is also positive for the first patient in Experiment 2, but is then (very close to) zero for patients 2 through 23. Hence, the positive $s_j - s_j^*$ difference for patient 1 in these experiments may be an artifact of the particular configurations used. (This is supported by the results of Experiment 3 presented later.) Further, we should be cautious in drawing strong inferences from the negative $s_j - s_j^*$ for early patients in Experiment 1: this result could obtain for a variety of reasons. However, we can say more about the $s_j - s_j^*$ results in Experiment 2; specifically, there is a pronounced bias to over-assign early patients to good hospitals, i.e., to overtriage early patients. The generality of this bias will be addressed in Experiment 3, where we used multiple hospital configurations.

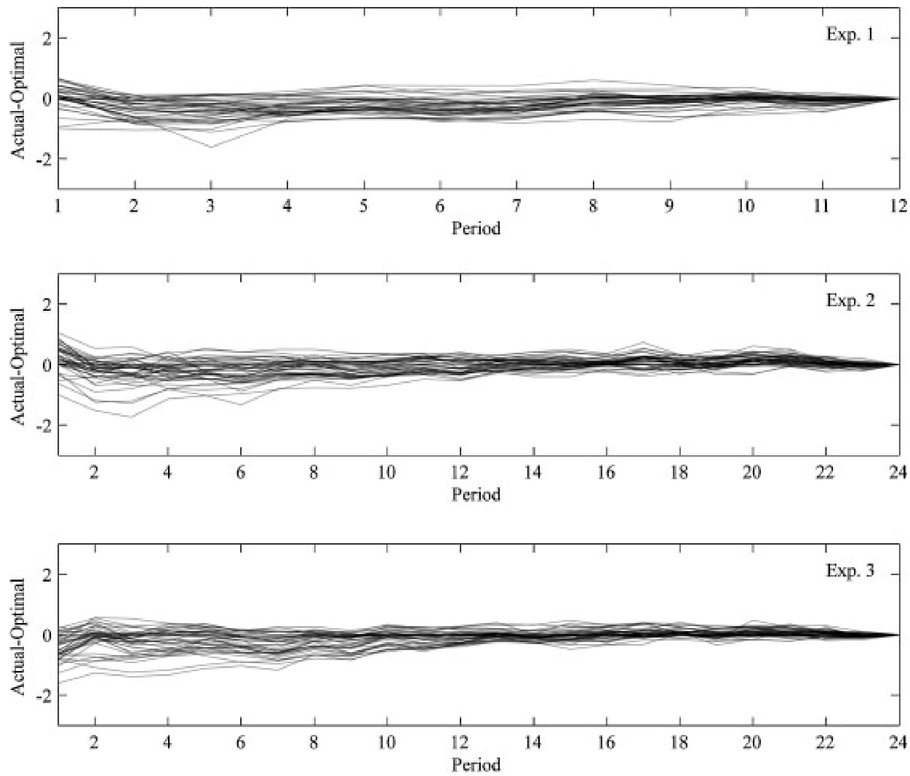

Figure 3. Mean difference in actual and optimal hospital assignments by patient position and subject across periods for Experiments 1, 2, and 3. [Each curve corresponds to an individual subject. Each point on each curve represents the mean difference between actual and optimal assignment positions taken over the 60 experimental trials. Negative values correspond to assigning patients to positions that are "too good."]

Assignment probabilities

For each experiment across the n periods we computed the proportion of times (or probabilities) that a patient was assigned to the sth hospital by the subjects (actual) and by the (conditionally) optimal policy. The results for Experiments 1 and 2 are depicted in Figures 4 and 5, respectively. These plots are exceptionally revealing and consistent across experiments. First, note that the optimal assignment probability to Hospital 1 for period 1 is 0 for both conditions. However, the subjects assigned the first patient to Hospital 1 with probability about 0.20 for the $n = 12$ case (Experiment 1) and about 0.15 for the $n = 24$ case (Experiment 2). In both experiments, the first patient should *always* be assigned to Hospital 2. Subsequent actual assignment probabilities for Hospital 1, however, are quite close to the optimal assignment probabilities, though in both experiments they are generally slightly smaller. With the exception of the first patient, in both experiments patients are generally assigned to Hospital 4 less often than they ought to be. Instead, there is a tendency to place patients in Hospitals 2 and 3 with higher probability than is dictated by the optimal policy. We tested the difference between actual and optimal assignment probabilities for positions 1 to $n - 1$ for each hospital in Experiments 1 and 2 with sign tests, using sign(optimal probability – actual probability). The counts used for the sign tests are shown in Table 4. In Experiment 1, the actual assignment probability to Hospital 4 was less than the optimal for 9 of the 11 positions, $p < 0.05$; for Hospital 1, the difference was positive for 6 of 11 positions, $p > 0.05$. The probabilities of negative differences for both Hospitals 2 and 3 were statistically significant, $p < 0.05$. All tests for the data in Experiment 2 were significant and in the direction predicted by what we term the *middleness bias*: that the subjects tend to avoid the best and particularly the worst hospitals more often than they ought to, preferring instead to assign patients to intermediate (or middle) hospitals.

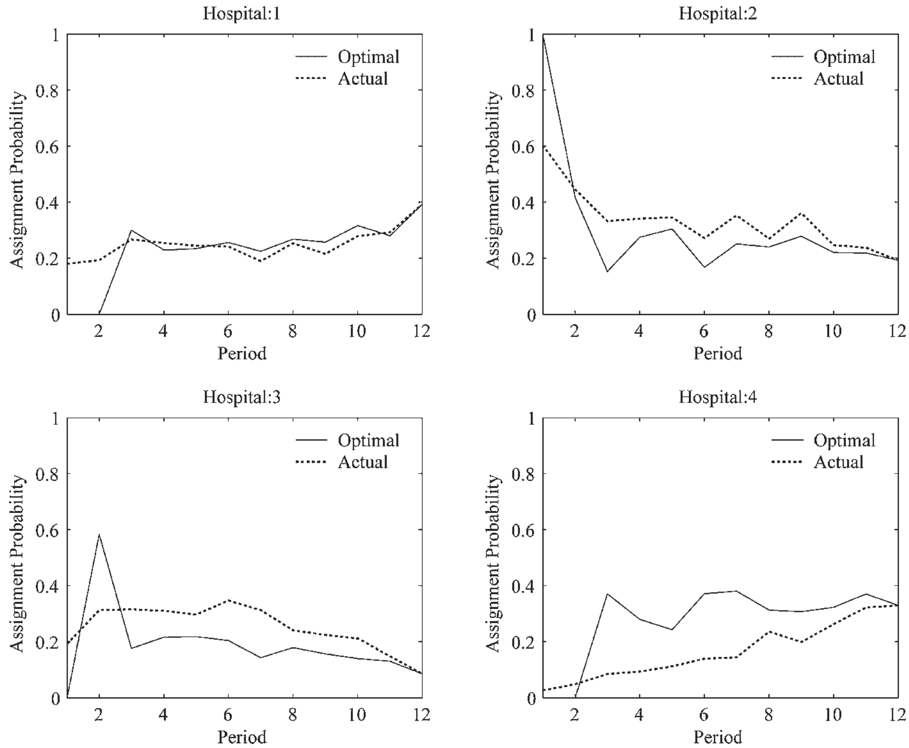

Figure 4. Actual and optimal assignment probabilities for the 12 periods Experiment 1 by hospital.

Discussion

In deciding how to assign the sequentially observed patients to hospitals, or, more generally, how to assign applicants to positions, the DM is pulled by two opposing forces. Because the applicants that appear early in the sequence necessarily have lower relative ranks ($r_1 = 1$, $r_2 =$ either 1 or 2, etc.), DMs who may partly confuse relative with absolute ranks may assign these applicants to the more costly positions. They may realize with experience that this policy results in having applicants with relatively low absolute ranks appearing later in the sequence being assigned to inferior positions. Hence, DMs may be pulled in the opposite direction of reserving the most costly positions to applicants appearing later in the sequence.

Assessing the absolute quality of the subject's performance is not possible, as no obvious standard exists other than the optimal and random assignment. We observe that our subjects' assignments are considerably better than random assignment; however, this is a weak criterion. Better bases for determining mean performance are the assignment probabilities displayed in Figures 4 and 5, which show that the difference between actual and optimal assignment probabilities tends to be relatively small. Although it is difficult to assess the absolute quality of assignments, we can clearly see that the assignments show very systematic departures from optimality (using the assignment probability metric). A major finding, what we refer to above as a *middleness bias*, shows a propensity to place patients in the intermediate quality hospitals. This bias is consistent with subjects saving the extreme hospitals (the best and worst) for patients in whose severity of injuries they have more confidence, *viz.* those who arrive later in the sequence. Further, we see that the DMs learn and make better assignments with experience.

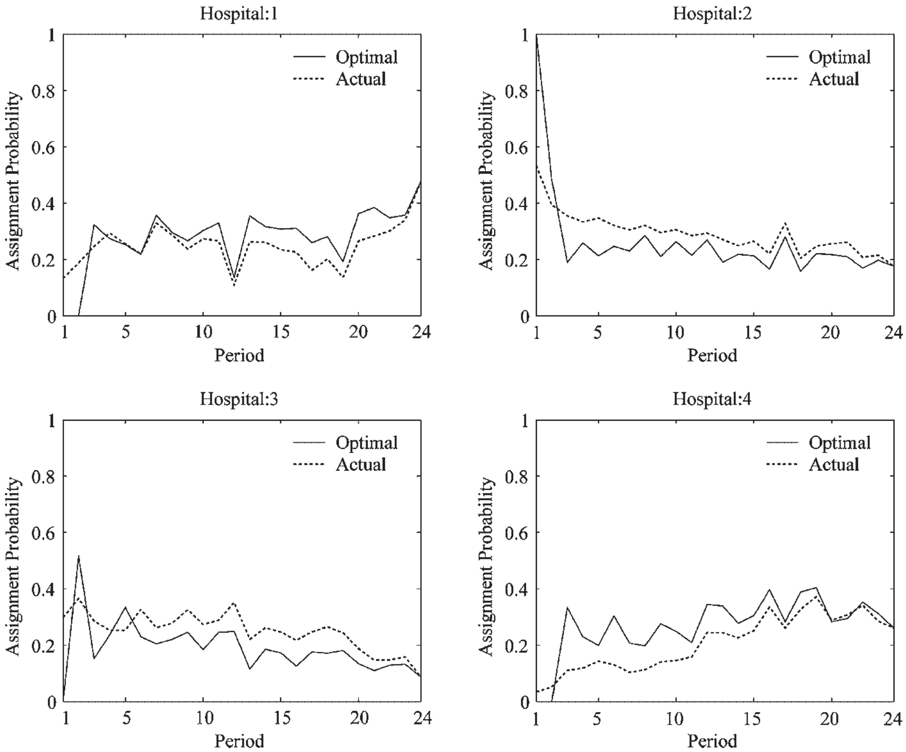

Figure 5. Actual and optimal assignment probabilities for the 24 periods in Experiment 2 by hospital.

Table 4. Counts of optimal assignment probabilities greater than actual assignment probabilities for periods 1 to $n-1$ for Experiments 1, 2, and 3

	Hospital 1 (cost = 100)	Hospital 2 (cost = 50)	Hospital 3 (cost = 25)	Hospital 4 (cost = 10)
Experiment 1	6	1*	1*	9*
Experiment 2	18*	2*	2*	19*
Experiment 3				
Configuration 1	19*	2*	4*	20*
Configuration 2	13	5*	11	13*
Configuration 3	17*	0*	4*	19*
Configuration 4	17*	0*	4*	19*
Configuration 5	17*	1*	3*	21*

Results are broken down by configuration for Experiment 3. The middleness bias implies that the counts for the extreme hospitals (1 and 4) should exceed those for the middle hospitals (2 and 3).

Note: For each configuration and each hospital we used a sign test to compare optimal and actual assignment probabilities across periods $1, \ldots, n-1$. (The difference is always 0 for period n.) For Hospitals 1 and 4, we hypothesized that the difference is more likely to be positive and therefore used a one-tailed test. For Hospitals 2 and 3, we hypothesized that the difference is more likely to be negative and used a one-tailed test. Asterisks (*) denote significance at the 0.05 level.

EXPERIMENT 3: TESTING THE GENERALITY OF THE RESULTS OF EXPERIMENTS 1 AND 2

Overview

A potential shortcoming of Experiments 1 and 2 is that our conclusions may be restricted to the particular hospital configurations that we used. In both experiments, the hospital configuration was fixed for all 60 trials. Further, the configuration in Experiment 2 was similar to the one in Experiment 1; the only difference was that each hospital had three more positions in Experiment 2. To overcome this shortcoming, in Experiment 3 we varied the configurations of hospitals from trial to trial. Doing so allows us to determine whether the behavioral patterns observed in the previous experiments are general biases that subjects exhibit in the SSAP. In addition, we can test whether the learning that was reported in Experiments 1 and 2 is due to subjects learning to assign patients to a particular hospital configuration. Thus, Experiment 3 provides a critical test of the generality of our previous conclusions that is not limited to a fixed hospital configuration.

Method

Subjects

Thirty-seven financially motivated subjects participated in Experiment 3. The experiment was conducted at the same computerized laboratory, and the subjects were recruited in the same manner as in the previous two experiments. The mean payoff per session was about \$31 (minimum \$15, maximum \$40).

Procedure

Except for the changing of hospital configurations across trials, the procedure was identical to the one used in the previous two experiments.

Manipulation

The subjects were presented with five different hospital configurations each repeated 12 times. Each subject observed the same random orderings of patients for the 12 repetitions of each configuration, but the ordering of trials was randomized separately for each subject. We set $n = 24$ for all configurations. Each subject received an endowment of 20 000 payoff units on each trial. The conversion rate was \$1 for every 1000 points earned on three randomly selected trials.

Table 5 presents the five configurations used in Experiment 3. These cover much of the space of possible configurations: positions in Configuration 1 are uniform over the hospitals; positions in Configuration 2 are peaked in the center, whereas those in Configuration 3 are peaked on the extremes; Configuration 4 (5) has a negatively (positively) skewed distribution of openings.

Table 5. Hospital configurations used in Experiment 3

Configuration	Hospital 1 (cost = 100)	Hospital 2 (cost = 50)	Hospital 3 (cost = 25)	Hospital 4 (cost = 10)
1	6	6	6	6
2	4	8	8	4
3	8	4	4	8
4	4	6	6	8
5	8	6	6	4

Cell entries represent the number of beds available in each hospital for each configuration.

Results

Overview

As for Experiments 1 and 2, the results are based on the conditionally optimal assignments.

Quality of patient assignments

Figure 1 shows that the quality of patient assignments—as measured by the RMSQ scores—is of the same order of magnitude as in Experiments 1 and 2. (The expected value of RMSQ under random assignment for the configurations used in Experiment 3 is 1.46.) In all three experiments, the RMSQ scores decrease across trials at roughly the same rate. The distribution of RMSQ over subjects in Experiment 3 is shown in the bottom panel of Figure 2. It is similar to those reported for Experiment 1 (upper panel) and Experiment 2 (middle panel).

The mean difference between actual and optimal assignments (averaged over configurations) tends to be negative for early arriving patients and then converges to zero (see lower panel of Figure 3). The expected value of $s_j - s_j^*$ under random assignment is -0.25 for patient 1 and near 0 for patients 2–23. We can, therefore, be confident that the consistently negative $s_j - s_j^*$ scores for early patients represent a genuine assignment bias and are not an artifact of the problem structure. Further, we can now confidently conclude that the positive difference for the first patient in Experiments 1 and 2 is, indeed, an artifact of the configurations used in those two experiments. Thus, early patients tend to be *overtriaged*: On average they are placed into better hospitals that should be reserved for other potentially more severely injured patients.

Assignment probabilities

Mean actual and optimal assignment probabilities (averaged over the five configurations) are exhibited in Figure 6. Comparison of Figures 4, 5, and 6 shows the same middleness bias: on average, patients are assigned to the intermediate hospitals (2 and 3) with higher probabilities than they ought to be. We used sign tests to test the assignment probabilities for each hospital in each configuration. From the middleness bias observed in Experiment 1, we predict that the actual assignment probabilities to Hospitals 1 and 4 will more often be less than those of the optimal assignment policy and that the actual probabilities will more often be greater for Hospitals 2 and 3. These predictions were statistically upheld for all hospitals in configurations 1, 3, 4, and 5 and also for two of the four positions in configuration 2. Again, the counts used for the sign tests are shown in the lower part of Table 4. Configuration 2, which provides the least support for the middleness bias prediction, has the greatest number of positions in the middle hospitals. It seems that the structure of the problem with Configuration 2 accommodates or corrects for the middleness bias. In sum, the middleness bias does seem to be a general property of subjects' assignments in the SSAP, except in cases in which the greatest number of open positions is in the intermediate hospitals. In these cases, the structure of the problem compensates for the bias.

An additional observation is in order. For all configurations, the subjects assigned the first patient to the best hospital (Hospital 1) with non-zero probability, though under optimal assignment for all five configurations the first patient should never be assigned to Hospital 1. Although we observe a general learning trend across trials, the first patient error tends to persist. The probability of assigning the first patient to Hospital 1 in the first 30 trials is 0.08 and in the last 30 trials is 0.10. The difference between the two is not significant.

Discussion

The consistency of the results across Experiments 1 through 3 is striking. The major findings can easily be summarized. First, there is a general bias to assign patients to the two middle hospitals with greater probability than is dictated by the (conditionally) optimal policy. Subjects seem to be trying to reserve the best

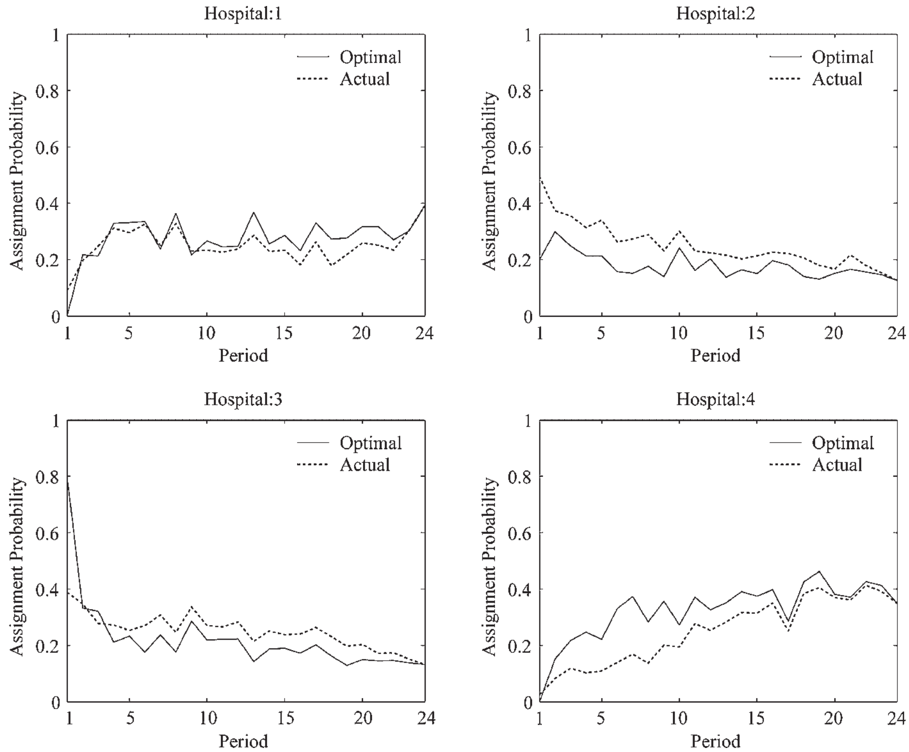

Figure 6. Actual and optimal assignment probabilities for the 24 periods in Experiment 3 by hospital.

and worst hospitals for later patients, about whom they will have more information (as relative ranks are more informative about absolute rank for patients later in the sequence). The second finding, which complements the first, is that there is a bias to place early patients into hospitals that are better than the ones to which they ought to be assigned. Put differently, early patients tend to be overtriaged. Taken together, the first and second findings reveal that there is a bias to over-assign patients to middle quality hospitals, and further that the middleness bias is in the direction of the better middle hospital. Finally, and most importantly, we observe learning across repeated play of the SSAP: The subjects' assignments move closer to the optimal assignments with experience. But, again, the observed biases do not fully disappear.

IMPLICATIONS AND FUTURE DIRECTIONS

Using inexperienced subjects, rather than individuals trained in triage assignment, we observed and documented systematic biases in triage assignments. After careful consideration, and following the practice used in most previous experimental studies of dynamic decision making (e.g., fire-fighting by Brehmer & Allard, 1991; vehicle navigation by Jagacinski & Miller, 1978; supervisory control by Kirlik, Miller, & Jagacinski, 1993; health management by Kleinmuntz & Thomas, 1987), we opted to use a scenario to which inexperienced subjects might easily relate rather than an abstract scenario of assigning "applicants" to "positions." One wonders, though, whether the findings observed in the experiments reported here might be contingent on the framing of the problem. Subjects were asked to make (simulated) triage decisions that could possibly result in the (simulated, conceptual) *loss* of life. Conceivably, this emphasis on saving each encountered patient could affect the nature of the assignments in a way that an alternative framing—say, in more positive terms—would not. For instance, suppose that the DM is given the task of allocating rewards to others based on their performance. Since the gravity of the assignment is less extreme than in the life saving scenario, it is possible that the assignment decisions would show a different trend. Similarly, the early stopping results observed in previous work on sequential selection problems (e.g., Bearden et al., 2005; Rapoport & Tversky, 1970; Seale & Rapoport, 1997) might be, at least in part, driven by the framing of the problems: in all cases DMs earned money on the basis of their selection decisions. Would one observe early stopping if DMs were instead trying to minimize their losses? Before general conclusions are drawn about the biases inherent in human selection and assignment decisions, these possibilities should be investigated.[1]

The middleness bias is consistent with choice (preference) data showing that when choosing among multi-attribute alternatives DMs tend to avoid extreme alternatives, opting instead for those with more intermediate-valued attributes (e.g., Simonson, 1989; Simonson & Tversky, 1992). For instance, consider three wines A, B, and C that can be ranked in terms of quality from best to worst ($A < B < C$) and also in terms of price ($A < B < C$). A significant number of studies—mostly in the marketing literature—have shown that the "market share" for C is likely to be increased by adding a new wine D to the choice set that is of both higher quality and price than C. Once C is no longer an extreme choice, it is chosen more often. Perhaps our observed middleness bias is nothing more than the assignment corollary of the "compromise effect" found in studies of choice. Simonson and Tversky explain the preference finding using loss aversion. They suggest that each alternative's attribute values are evaluated relative to a reference point that is dependent on the attribute values of the options in the choice set. And since losses (e.g., paying more than the reference point price) loom larger than gains (paying less), options that offer a compromise tend to be chosen.

The observed middleness bias seems to be quite sensible in many naturalistic settings; it might be a safe rule of thumb to fall back on. Consider some alternatives. If DMs consistently assumed the worst-case scenario, they would quickly inundate valuable resources that should be reserved for high-priority assignments. On the other hand, a bias to allocate assignments to the lowest priority category seems potentially equally disastrous. When the decisions are tough, it seems quite sensible to err in a direction that is conservative both with respect to over- and under-assigning resources. In short, given the complexity of determining optimal assignments, which involve solving complicated dynamic programming problems, our subjects' behavior seems quite sensible.

There are two important features of our decision problem that may not be characteristic of many sequential assignment problems. First, the DMs in our experiments were informed of the number of to-be-assigned patients. A more realistic assumption is that the DM can only estimate the distribution of n (see Seale & Rapoport, 2000, who relaxed the assumption of known number of applicants in an experimental selection task). Currently, the optimal assignment policy for this problem is not known. Another possible extension is to cases in which the DM must assign batches or groups of applicants simultaneously to particular positions; for example, this is relevant when patients must be transported to hospitals in ambulances or helicopters. The optimal policy for this class of problems is yet to be derived.

ACKNOWLEDGEMENTS

We gratefully acknowledge financial support by a contract F49620-03-1-0377 from the AFOSR/MURI to the University of Arizona. We also thank Frank Yates and three anonymous reviewers for their valuable suggestions.

[1]We thank an anonymous reviewer for providing us with this suggestion.

REFERENCES

Albright, S. C. (1974). Optimal sequential assignments with random arrival times. *Management Science, 21*, 60–67.

Albright, S. C., & Derman, C. (1972). Asymptotic optimal policies for the stochastic sequential assignment problem. *Management Science, 19*, 46–51.

Bearden, J. N., Murphy, R. O., & Rapoport, A. (2005). A multi-attribute extension of the secretary problem: theory and experiments. *Journal of Mathematical Psychology, 49*, 410–425.

Bond, R. J., Kortbeck, J. B., & Preshaw, R. M. (1997). Field trauma triage: combining mechanism of injury with pre-hospital index for an improved trauma triage tool. *Journal of Trauma, 43*, 283–287.

Brehmer, B., & Allard, R. (1991). Real-time dynamic decision making: effects of task complexity and feedback delays. In J. Rasmussen, B. Brehmer, & J. Leplat (Eds.), *Distributed decision making: Cognitive models for cooperative work*. Chichester: Wiley.

Chun, Y. H. (1996). On the rank-based selection strategy for the group interview problem. *Decision Sciences, 27*, 801–816.

Chun, Y. H., & Sumichrast, R. T. (in press). A rank-based approach to the sequential selection and assignment problem. *European Journal of Operational Research*.

Corbin, R. M., Olson, C. L., & Abbondanza, M. (1975). Context effects in optimal stopping decisions. *Organizational Behavior and Human Performance, 14*, 207–216.

David, I. (1995). Sequential assignment match process with general renewal arrival times. *Probability of Engineering Information Science, 9*, 475–492.

David, I., & Yechiali, U. (1985). A time-dependent stopping problem with application to live organ transplant. *Operations Research, 33*, 491–504.

David, I., & Yechiali, U. (1990). Sequential assignment match processes with arrival of candidates and offers. *Probability of Engineering Information Science, 4*, 413–430.

Derman, C., Lieberman, G. J., & Ross, S. M. (1972). A sequential stochastic assignment problem. *Management Science, 18*, 349–355.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science, 4*, 282–296.

Freeman, P. R. (1983). The secretary problem and its extensions: a review. *International Statistical Review, 51*, 189–206.

Gilbert, J., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association, 61*, 35–73.

Hornberger, J., & Ahn, J.-H. (1997). Deciding eligibility for transplantation when a donor kidney becomes available. *Medical Decision Making, 17*, 160–170.

Jagacinski, R. J., & Miller, R. A. (1978). Describing the human operator's internal model of a dynamic system. *Human Factors, 20*, 425–433.

Kennedy, D. P. (1986). Optimal sequential assignment. *Mathematics of Operations Research, 11*, 619–626.

Kirlik, A., Miller, R. A., & Jagacinski, R. J. (1993). Supervisory control in a dynamic and uncertain environment: a process model of skilled human-environment interaction. *IEEE Transactions of Systems, Man, and Cybernetics, 23*, 929–952.

Klienmuntz, D., & Thomas, J. (1987). The value of action and inference in dynamic decision making. *Organizational Behavior and Human Decision Processes, 39*, 3341–3364.

Rapoport, A., & Tversky, A. (1970). Effects of search cost on sequential search behavior. *Perception & Psychophysics, 6*, 234–240.

Sakaguchi, M. (1984). A sequential stochastic assignment problem with an unknown number of kidneys. *Mathematical Journal of Japan, 29*, 141–152.

Samuels, S. M. (1991). Secretary problems. In B. K. Ghosh, & P. K. Sen (Eds.), *Handbook of sequential analysis* (pp. 381–405). New York: Marcel Dekker.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: an experimental investigation of the "secretary problem." *Organizational Behavior and Human Decision Processes, 69*, 221–236.

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: the secretary problem with unknown population size. *Journal of Behavioral Decision Making, 13*, 391–411.

Simonson, I. (1989). Choice based on reasons: the case for attraction and compromise effects. *Journal of Consumer Research, 16*, 158–174.

Simonson, I., & Tversky, A. (1992). Choice in context: tradeoff contrast and extremeness aversion. *Journal of Marketing Research, 29*, 281–295.

Stein, W. E., Seale, D. A., & Rapoport, A. (2003). Analysis of heuristic solutions to the best choice problem. *European Journal of Operational Research, 51*, 140–152.

Su, X., & Zenios, S. (2002). Allocation of kidneys to autonomous transplant candidates: a sequential stochastic assignment model. Graduate School of Business, Stanford University, unpublished manuscript.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthokrishnan, A. V. (2003). Consumer sequential search: not enough or too much? *Marketing Science, 22*, 503–519.

Authors' biographies:

J. Neil Bearden is a Postdoctoral Fellow in the Brain Imaging and Analysis Center at Duke University. He received his Ph.D. from the University of North Carolina. His research is in behavioral decision theory and optimization.

Amnon Rapoport is a Karl Eller Professor of Management and Policy in the Eller College of Management at the University of Arizona, and an Adjunct Professor of Economics at the Hong Kong University of Science and Technology. He received his Ph.D. from the University of North Carolina. He studies individual and interactive decision making.

Ryan O. Murphy recently earned his Ph.D. from the Eller Business School at the University of Arizona. His research interests are in individual decision making and behavioral game theory. He currently is a research scientist and the associate director of the Center for Decision Sciences at Columbia University.

Authors' addresses:

J. Neil Bearden, Brain Imaging Analysis Center, Duke University, 163 Bell Research Building, Box 3918, Duke University Medical Center, Durham, NC 27710, USA.

Amnon Rapoport, University of Arizona, Department of Management and Policy, 405 McClelland Hall, Tucson, AZ 85721, USA.

Ryan O. Murphy, Columbia University, Center for the Decision Sciences, 420 West 118th Street, Room 805A, Mail Code 335, New York, NY 10027, USA.

