

EUROPEAN UNIVERSITY INSTITUTE
DEPARTMENT OF ECONOMICS

EUI Working Paper ECO No. 99/4

Learning while Searching for the Best Alternative

Klaus Adam

BADIA FIESOLANA, SAN DOMENICO (FI)

All rights reserved.

No part of this paper may be reproduced in any form
without permission of the author.

©1999 Klaus Adam

Printed in Italy in February 1999

European University Institute

Badia Fiesolana

I-50016 San Domenico (FI)

Italy

Learning While Searching for the Best Alternative

Klaus Adam*

European University Institute

Florence, Italy

e-mail: kadam@datacomm.iue.it

January 4, 1999

Abstract

This paper delivers the solution to an optimal search problem with learning where the searcher has distinguishable search opportunities. The optimal sampling strategy is characterized by simple reservation prices that determine which of the search alternatives to sample and when to stop search. The reservation price criterion is optimal for a large class of learning rules having the so-called falling reservation price property, including Bayesian, non-parametric and ad-hoc learning rules. The considered search problem contains as special cases many earlier contributions to the search literature and thereby unifies and generalizes two directions of research: search with learning from identical search alternatives and search without learning from distinguishable search alternatives.

JEL-Class.No.: D81, D83

Keywords : Optimal Search, Systematic Search, Learning, Reservation Prices, Uncertainty

*I would like to thank my supervisor Ramon Marimon, Pierpaolo Battigalli, Matthias Brückner, my colleagues at the EUI, and seminar participants at the EEA-Meeting 1998 in Berlin. All errors are mine.

1 Introduction

Economic problems involving search due to uncertainty about the location of objects are copious and hence have received a considerable amount of attention. After the igniting article of Stigler [18] economists themselves have been searching, namely for sampling strategies that are optimal in different situations involving such uncertainty (Lippman and McCall [10] or McKenna [11]). This paper stands in this tradition and determines the optimal search strategy for a class of search problems that is characterized by two main features: Learning during the search process and distinguishable search alternatives.

To be explicit, consider the following job search example falling into the class of problems I consider. A job searching unemployed worker faces a number of job offering firms where each firm might either be willing to hire this worker and offer some wage or reject the worker's application. The fundamental uncertainty in the worker's search process consists of the fact that the worker does not know which firms are willing to hire at which wage and which ones would reject the application. Thus, the worker has to search for a good offer by filing applications to firms, observing the outcomes and deciding whether to accept an offer or whether to continue searching.

Learning is introduced by allowing for the natural possibility that the searcher is not only uncertain about which firm offers which wage but also uncertain about the prevailing wage offer distribution. The searcher, possessing priors about the offer distribution, can use a search outcome, i.e. a job offer of a particular firm, to learn about the wage offer distribution by updating these priors.

It is equally natural to suppose that the searcher can distinguish firms along some dimension and has different priors about the type of vacancies offered by different firms. The distinction could be based upon firms belonging to different sectors or local markets or upon any other observable characteristic of firms. As a result, the searcher faces distin-

guishable search alternatives and has to choose to which sector or which local market to apply.

In abstract terms, a search problem involving learning adds to the uncertainty about the location of objects the uncertainty about the objects' values while the presence of distinguishable search alternatives captures the fact that search opportunities typically differ from each other and that search involves a thorough choice among the available alternatives.

If search is sequential with full recall of previous offers, then I find that the optimal search strategy for the class of search problems involving learning and distinguishable search alternatives is characterized by a simple reservation price for each search alternative. The reservation price of an alternative is simply a real number that is assigned to the alternative and the higher this number, the more attractive it is to search the corresponding alternative. The reservation prices for all alternatives together determine both, which of the search alternatives to sample, and when to stop search. The optimal strategy is very simple and prescribes to search always the search alternative with the highest reservation price and to stop search as soon as the best offer exceeds the reservation prices of all available alternatives.

The reservation prices keep changing during the search process as new information arrives through new search outcomes and learning takes place. In this way, it is optimal for the searcher to stay reactive to the search outcomes and, for example, direct search towards another search alternative, if the outcomes of the previously searched alternative have been disappointing.

The optimality of the search strategy holds for a large class of learning rules for which, roughly speaking, the reservation prices keep decreasing as additional search outcomes are observed. Learning rules with this property include Bayesian learning as well as non-parametric and ad-hoc learning.

In addition to answering the question on how to search optimally in a situation involving learning and distinguishable search alternatives, the result of this paper should be of twofold interest to economists.

First, the answer to the normative question allows for positive modeling of economic behavior within the neoclassical maximization paradigm. There are many situations of economic interest that involve both of the above features and where the findings of this paper are applicable. Besides job search these are consumers' search for the best price or firms' research for new products or technologies. Examples of the latter include oil companies searching for new oil fields to exploit or pharmaceutical companies' research for medical drugs. The results are equally applicable to any kind of investment decision if investment is interpreted as the search for good investments projects.

Second, the result contains several earlier contributions to the search literature as special cases and thereby contributes to the unification and generalization of the search theoretical framework. Although learning and distinguishable search alternatives have already been considered in the literature only one of these features was present at a time (Rothschild [15], Rosenfield and Shapiro [14], Morgan [12], Talmain [19], Chou and Talmain [3], Bikchandani and Sharma [1] considered learning but assumed indistinguishable search alternatives; Salop [16], Weitzman [21], Vishwanath [20] studied distinguishable search alternatives but abstracted from learning) and many of the search problems studied in earlier contributions are contained in the class of problems considered in this paper.¹

It is worth noticing that removing learning or distinguishable search alternatives both reduce the complexity and realism of search problems considerably. On one hand, assuming indistinguishable search alternatives removes the choice decision from the search problem. All search

¹Exceptions from the listed articles are Vishwanath [20] dealing with non-sequential search, Morgan [12] dealing mainly with the existence of reservation price functions and Rothschild [15] not allowing for recall of previous offers.

alternatives are (at least believed to be) the same and the search problem then reduces to the question on when to stop search optimally. On the other hand, abstracting from learning implies that the value of a search outcome (e.g. of a job offer) consists solely in its payoff (i.e. the wage), since search outcomes do not convey any valuable information (e.g. about the wage offer distribution). As a result the optimal search strategy has to condition only on the best of all observed offers (i.e. the best wage offered so far) and not on the whole sequence of observed offers.

Finally, notice that the problem considered in this paper differs from simple armed bandit problems but that it is related to bandit superprocesses.

First, consider the difference to the simple bandit problem. In such a decision problem the player receives a reward every time the arm of a bandit is pulled and nothing otherwise. In contrast to this, in the considered search problem a number of arms are pulled without actually receiving a reward. Only when the searcher decides not to pull any further arms (i.e. to stop search) the best of all previously observed rewards is obtained.

Next, consider bandit superprocesses which are a generalization of armed bandit processes allowing for multiple arms per bandit. Adding a second 'stopping arm' to a standard bandit (as in Glazebrook [8]) allows for the possibility that the payoff is obtained at the end of search when the stopping arm is pulled. Glazebrook shows that if the value of the stopping option is non-decreasing in the number of searches, then the optimal policy is characterized by some simple selection rule between the arms and the indices given by Gittins and Jones [5] for simple bandits. However, while I allow for a finite or an infinite number of search opportunities, Glazebrook's result fails to hold, if there is not an infinite number of search opportunities of each search alternative.² Even if there are infinite numbers, the indices in Gittins and Jones are not particularly

²It is easy to see that already in the simple example given in section 3 the solution given by Glazebrook [8] does not hold anymore. A finite number of search opportunities is like an additional constraint on the action space of the superbandid.

explicit and the monotonicity conditions that allow for a straightforward explicit calculation (e.g. as the ones in propositions 4.2 and 4.5 in Gittins [4]) fail to hold in our case.³ Thus, the contribution of this paper could also be considered in delivering an explicit expression for these indices in the absence of such monotonicity.⁴

The next section sets up the search problem I consider and explains how other search problems with identical search alternatives or without learning are special cases of the one considered here. Section 3 describes as a benchmark the optimal search strategy when the searcher knows the payoff distributions and is not learning. Section 4 contains the main part of the paper. I delineate the class of admitted learning rules and present the optimal search strategy for the case with learning. I also explain why the sampling rule of the benchmark problem generalizes to the case with learning. In Section 5 I ask whether one can also hope for optimality of the search rule with more general learning rules than the ones I considered. Unless for a very special case the answer is found to be negative. A conclusion summarizes the findings. The appendix contains the proofs.

2 The Model

A search problem is characterized by a searcher facing a (possibly infinite) number of search opportunities. Each search opportunity can be thought of as a box that contains an uncertain reward. The searcher has the possibility to open any box at a cost and find out what reward is contained in the box. I want to allow the boxes to differ from each other, not only with respect to the actual reward they contain but also with respect to the probability with which they contain (or are believed to

³Note that although we have decreasing reservation prices with our learning rules there is always a positive probability that the search outcome is above the reservation price.

⁴For similar exercises see Glazebrook [6] and [7].

contain) certain rewards. One can think of this as different boxes having different colors on the outside, while equal boxes are of equal color. Each color then represents a search alternative and the searcher, being able to observe these colors, has to choose among them in every search step.

More formally, let boxes be indexed by the natural numbers and let the set $J = \{1, 2, \dots\}$ contain all available boxes. Each box $j \in J$ has some color $i \in \{1, 2, \dots, I\}$, i.e. there are I different colors or search alternatives. The color of a box is observable for the searcher at no cost. To simplify language a box of color i will sometimes be referred to as an i -box.

There are M^i boxes of color i where M^i can be finite or infinite. Boxes of the same color are identical and are characterized by the triple $\{c^i, t^i, d^i(\cdot)\}$ where c^i are the costs for opening an i -box, t^i is the time span that passes from opening the box until its reward is observed and the function $d^i : R \mapsto [0, 1]$ describes the probability distribution of rewards from opening the box. The parameters c^i and t^i are known to the searcher while $d^i(\cdot)$ is unknown. The functions $d^i(\cdot)$ can have support on \mathbf{R} and the random variables described by them are assumed to have finite mean if $M^j < \infty$ for all $j = 1, 2, \dots, I$ and to have finite variance in all other cases.

For a given point in time I denote by r^i the number of already opened i -boxes. x_n^i is the outcome from opening the n -th box of color i . The vector $X_{r^i}^i = (x_1^i, x_2^i, \dots, x_{r^i}^i)$ contains the r^i so far observed outcomes from opening i -boxes.

The searcher samples sequentially for boxes and can open a closed box of color i by paying the amount c^i . He has to wait a time span t^i and then receives an offer drawn from $d^i(\cdot)$.⁵ Recall of previously drawn offers is allowed. If search stops, the searcher gets y which is the maximum of the so far drawn offers and some outside opportunity x^o the searcher

⁵That search costs c^i have to be paid some time t^i before the search result is observed is not restrictive. Problems where c^i is payed at the time when search results are observed fit into the problem by appropriately discounting search costs.

possesses independently from the search outcomes:

$$y = \max \{x^o, x_1^1, x_2^1, \dots, x_{r^1}^1, \dots, x_1^I, x_2^I, \dots, x_{r^I}^I\}$$

The searcher maximizes discounted expected payoffs minus costs with a discount rate $0 \leq r < \infty$.

Uncertainty has two sources. First, offers from boxes are drawn from some probability distribution. Second, there is uncertainty about the prevailing distribution from which offers are drawn. Uncertainty about $d^i(\cdot)$ may be represented by beliefs in form of a probability distribution $p^i(\theta)$ over some parameter θ that indexes the set of possible true probability distributions $d^i(\cdot | \theta)$ for boxes of color i , where the true distribution function $d^i(\cdot)$ is equal to $d^i(\cdot | \theta^i)$ for some specific value θ^i of the parameter. Beliefs $p^i(\theta)$ about boxes of color i are updated using the observed search outcomes $X_{r^i}^i$ from i -boxes. Updated beliefs are denoted by $p^i(\theta | X_{r^i}^i)$. Given these beliefs one can calculate an expected true probability distribution $f^i(x | X_{r^i}^i)$ for the boxes of each color by integrating out for the uncertainty about the parameter θ :

$$f^i(x | X_{r^i}^i) = E [d^i(x | \theta) | X_{r^i}^i] = \int_{\Theta} d^i(x | \theta) p^i(\theta | X_{r^i}^i) d\theta$$

For expository reasons, $f^i(\cdot | \cdot)$ has been derived from a Bayesian learning mechanism above. Since I do not want to confine myself to rational learning, I equally allow $f^i(\bullet | \bullet)$ to be directly specified by some non-rational ad-hoc learning rule.⁶ In both cases, rational and non-rational learning, the functions $f^i(\cdot | \cdot)$ determine a joint prior probability for any sequence $(x_1^i, x_2^i, \dots, x_n^i)$ of search outcomes with

$$\Pr(x_1^i, x_2^i, \dots, x_n^i) = f^i(x_1^i) \cdot f^i(x_2^i | x_1^i) \cdot \dots \cdot f^i(x_n^i | x_1^i, x_2^i, \dots, x_{n-1}^i) \quad (1)$$

Given the probability distribution (1) defined by the learning rule, the searcher maximizes the discounted expected payoff minus costs

$$\max_S E [e^{-r\tau_s} y_{\tau_s} - C_s] \quad (2)$$

⁶The random variables described by $f^i(\cdot | \cdot)$ are assumed to have finite mean if $M^j < \infty$ for all $j = 1, 2, \dots, I$ and to have finite variance otherwise.

with τ_s being the stopping time under sampling rule S , y_{τ_s} being the offer that got accepted in τ_s and $E[C_s]$ being the expected discounted sampling costs under S . Clearly, if learning is non-rational, then (2) differs from expected utility maximization because the searcher is only optimal for the given learning rule he uses. If learning is Bayesian, then (2) is identical to expected utility maximization.

I want to make two comments with regard to the above setup. First, it is a quite restrictive but crucial assumption that the functions F^i depend only on observations of i -boxes, i.e. outcomes of boxes of color $j \neq i$ do not reveal information about the parameter θ^i of i -boxes. For a Bayesian learner this is an implicit assumption on having prior one on the parameters $(\theta^1, \dots, \theta^I)$ being chosen independently.

Second, the setup comprises as special cases models without learning and several search alternatives and models with learning but identical boxes. In case that there is only one box of each color, no learning will take place and the model reduces to the one studied by Weitzman [21].⁷ In case that all boxes have the same color, the model reduces to the search problems considered (amongst other problems) in Rosenfield and Shapiro [14], Talmain [19], Bikchandani and Sharma [1], Chou and Talmain [3].

3 Benchmark: Optimal Strategy Without Learning

This section presents the optimal sampling rule when there is only one box of each color and hence no learning taking place.⁸ Such a problem is equivalent to a search problem with full information when the searcher's expected payoff distributions equal the true payoff distributions. The results presented here will serve as a helpful reference point for our later considerations and the main result is due to Weitzman [21].

⁷The searcher might still learn about the box of a particular color by opening it, yet at the time learning takes place there are no other boxes of that color left.

⁸Remember that we ruled out learning across boxes of different color.

For expository reasons consider the following simple but instructing example.

Example 1 *Suppose that there are only 2 boxes, a red one and a green one. Table 1 describes the payoff distributions $d^i(\cdot)$ of each box. For simplicity I will refer to the zero outcome as a "failure" and to the strictly positive outcome as a "success". With search costs for opening a box*

Table 1:

<i>Red</i>	<i>payoff</i>	<i>0</i>	<i>70</i>
	<i>with probability</i>	<i>0.1</i>	<i>0.9</i>
<i>Green</i>	<i>payoff</i>	<i>0</i>	<i>200</i>
	<i>with probability</i>	<i>0.85</i>	<i>0.15</i>

equal to 20, no discounting and the value of the outside option equal to zero, the expected payoffs from opening a single box are shown in table 2. Since the red box has a higher expected value than the green one, it

Table 2:

	<i>Expected Payoff</i>
<i>Red</i>	<i>43</i>
<i>Green</i>	<i>10</i>

might seem better to sample the red box first. If the result of doing so is a failure, it is clear that it pays to sample the green box as well because it has positive expected payoff. If the result of sampling the red box was a success, then sampling the green box yields a negative expected gain. The expected payoff of this sampling order is therefore readily calculated to be

$$-20 + 0.9 \cdot 70 + 0.1(-20 + 0.15 \cdot 200) = 44$$

Yet, sampling the green box first and then in case of a failure the red box is the optimal sampling order. Its expected value is

$$-20 + 0.15 \cdot 200 + 0.85(-20 + 0.9 \cdot 70) = 46.55$$

A simple intuition exists as to why the expectation criterion does not work in deciding upon which box to open first: It ignores the option value of the possibility to continue search in case of a low search outcome. This option value is relatively small in the case of a failure of the red box, namely $0.1 \cdot (-20 + 0.15 \cdot 200) = 1$ (the probability of a failure of the red box times the expected value of opening the green box), but relatively high in case of a failure of the green box, namely $0.85 \cdot (-20 + 0.9 \cdot 70) = 36.55$. Adding the first option value to the expected value of the red box gives 44, which is the value of the non-optimal sampling order. Adding the second option value to the expected value of the green box gives 46.55, the value of the optimal sampling order. Thus, although the immediate payoff from sampling the green box is lower than the immediate payoff from sampling the red box, the higher option value of continued search more than compensates for this.

It turns out that it is not necessary to calculate the option values of continued search to determine the right sampling order. There is a simple way of calculating an index for every search alternative that is based on the payoff distribution of the respective alternative alone. This is important to know because the option value of continued search can be a fairly complicated object, especially if one has many boxes of many different colors and, as in the next section, learning going on during the search process. The index has already been suggested by Lippman and McCall [10]. In the following I will describe how it is calculated and give some intuition on why it works.

Suppose the best offer from previous searches is y , then the expected gain over y from opening an i -box and stopping search with what is best then can be calculated to be

$$\begin{aligned} Q^i(y) &= \left(\beta^i \int_{-\infty}^y y dF^i(x) + \beta^i \int_y^{\infty} x dF^i(x) - c_i \right) - y \\ &= \beta^i \int_y^{\infty} (x - y) dF^i(x) - (1 - \beta^i)y - c_i \end{aligned}$$

where $\beta^i = e^{-rt^i} \leq 1$ is the discount factor.

Define as the reservation price R^i of an i -box that value of the best offer y at which the searcher would be indifferent between the following two actions: 1. Stopping search with y , and 2. Sampling an i -box and stopping thereafter with what is the best offer then, i.e.

$$Q^i(R^i) \equiv 0$$

Notice that R^i can be calculated using the payoff distribution of i -boxes only, ignoring any value from continued search.

The values R^i are the indices characterizing the optimal search strategy. The optimal sampling rule for the search problem without learning (later on also referred to as the benchmark rule) based on these indices is as follows:

- Step 1** Calculate the reservation prices for each box.
- Step 2** If there is no closed box with a reservation price higher than the current best offer y , then stop search and accept y , otherwise continue with step 3.
- Step 3** Open the box with the highest reservation price and go back to step 2.

A simple check of the reservation prices of the two boxes in our previous example reveals that $R^{red} = 47.8 < 66.7 = R^{green}$.⁹ The rule therefore confirms the optimality of sampling the green box first.

A simple intuition exists on why the above sampling rule should be the optimal one. Consider the following alternative interpretation of the reservation prices. It is well known that the optimal strategy for a search problem with an infinite number of i -boxes (and no other alternatives) is a reservation price strategy. The optimal reservation price for such a problem is the same as the one calculated above. Moreover, the reservation price is the value of a secure payoff that makes the searcher

⁹In the case of example 1 the reservation price formula boils down to $R^i = x_h^i - \frac{c^i}{p_h^i}$ where x_h^i is the value of the positive payoff and p_h^i is the probability of obtaining it.

indifferent between accepting a secure payoff and having the opportunity to sample i -boxes. $R^i > R^j$ can then be understood as the return from sampling i -boxes being higher than the return from sampling j -boxes. Search opportunities with higher reservation prices should therefore be sampled first.

4 Optimal Strategy with Learning

This section contains the main results of this paper. I begin by presenting the reservation prices and discussing their properties. Then I delineate the class of admitted learning rules and present the optimal sampling strategy for a learning searcher. Since the optimal strategy is a generalization of the benchmark strategy I explain in the last subsection why this is the case. The section is rather technical and can be skipped by readers mostly interested in the results.

4.1 The Reservation Prices

As in the case of known distributions, one can define the expected gain Q^i of opening one more i -box and stopping search thereafter over stopping immediately. With learning the expected distribution of search outcomes of i -boxes, $F^i(\cdot | X_{r^i}^i)$, now depends on the information contained in the previously observed search outcomes $X_{r^i}^i$. Therefore, the expected gain Q^i is now a function of the available information:

$$\begin{aligned}
 Q^i(X_{r^i}^i, y) &= \beta^i \int_{-\infty}^y y dF^i(x_{r^i+1}^i | X_{r^i}^i) \\
 &\quad + \beta^i \int_{-\infty}^y x_{r^i+1}^i dF^i(x_{r^i+1}^i | X_{r^i}^i) - c_i - y \\
 &= \beta^i \int_y^y (x_{r^i+1}^i - y) dF^i(x_{r^i+1}^i | X_{r^i}^i) - (1 - \beta^i)y - c_i
 \end{aligned}$$

where $\beta^i = e^{-rt^i} \leq 1$ is the discount factor.

Analogously to the full information case, one can define the reservation price of boxes from alternative i .¹⁰

Definition The reservation price $R^i(X_{r^i}^i)$ for boxes from alternative i is the value of y that solves $Q^i(X_{r^i}^i, y) = 0$

Again, the reservation price of i -boxes is that value of the best offer y which makes the searcher indifferent between stopping and doing one more search step.

Notice that reservation prices R^i are now also a function of the current information $X_{r^i}^i$. Reservation prices may therefore change over time as new information becomes available. Yet, how they might change in the future does not enter into the calculation of the reservation prices. Therefore, for given beliefs and hence given expected distribution function $F(\cdot | X_{r^i}^i)$, the reservation prices are independent from the searcher's learning rule.

The $R^i(X_{r^i}^i)$ have again an alternative interpretation as the reservation price of an optimally behaving (non-learning) searcher facing an infinite number of boxes with payoff distribution $F^i(\cdot | X_{r^i}^i)$.

4.2 Learning Rules

We saw in the previous section that the reservation prices depend only on current beliefs and are independent from the potential future evolution of these beliefs, i.e. from the learning rule. If we want to characterize the optimal search strategy based on this momentary picture of beliefs, we have to restrict the admitted learning rules in a way that this picture is sufficiently informative about the future.

We can express the necessary requirements on the learning rules in terms of an assumption on the evolution of reservation prices as learning

¹⁰Existence and uniqueness is guaranteed by the conditions of lemma 2 in the appendix.

proceeds. All learning rules with falling reservation prices are admitted. Formally,

Assumption A1 Let $X_{r^i+1}^i = (X_{r^i}^i, x_{r^i+1}^i)$, then

$$R^i(X_{r^i+1}^i) \leq R^i(X_{r^i}^i) \text{ or } R^i(X_{r^i+1}^i) \leq x_{r^i+1}^i \quad \forall i, X_{r^i}^i, x_{r^i+1}^i$$

Assumption A1 requires that after observing an additional search outcome of an i -box, the new reservation price $R^i(X_{r^i+1}^i)$ is either smaller than the old reservation price $R^i(X_{r^i}^i)$ or smaller than new offer $x_{r^i+1}^i$.¹¹

This can be interpreted as follows: Either the searcher gets a low search outcome and lowers in response to that the beliefs about the attractiveness of the sampled search alternative, which in turn leads to a lower reservation price, or the searcher gets a high outcome indicating that the search alternative is more attractive than thought before and increases the reservation price. In the latter case, it is important that the increase in the reservation price is moderate enough to ensure that the second of the above inequalities holds.

What is ruled out are so-called strong positive learning effects. These are search outcomes revealing a lot of good news about the attractiveness of a search alternative. In fact, so much that if the searcher were given the value of such a search outcome as the outside option, he would terminate search, but as one told him that this outside option is a draw from the search alternative, he would want to continue search.

In the following I give examples of learning rules that fulfill A1 and that have been used in the search literature dealing with identical boxes.¹² The optimal sampling strategy I derive holds for any of the following learning rules. The searcher might even apply different learning schemes to different search alternatives.

¹¹Since $R^i(X_{r^i}^i) < y$ implies $Q^i(X_{r^i}^i, y) < 0$, A1 insures that the one period gains $Q^i(X_{r^i}^i, y)$ stay negative, once they have become negative at some point of time. A1 therefore implies the sufficient condition used in Rosenfield and Shapiro ([14], Theorem1) to establish the optimality of a myopic stopping rule.

¹²In many of the following references increasing reservation prices can be found because the search problem is posed in terms of search for the lowest price of some good.

- i. Let the offer distribution be multinomial with N possible outcomes x_1, x_2, \dots, x_N and the probability of observing outcome x_i be equal to θ_i . If learning is Bayesian and the searcher has Dirichlet priors about the vector θ , i.e.

$$p(\theta \mid \alpha_1, \alpha_2, \dots, \alpha_N) \propto \theta_1^{\alpha_1-1} \theta_2^{\alpha_2-1} \dots \theta_N^{\alpha_N-1} \text{ with } \alpha_i > 0$$

then reservation prices are decreasing (e.g. Talmain [19]). The generalization of the multinomial Dirichlet case to an infinite number of possible outcomes by a Dirichlet process also implies declining reservation prices (see Bikchandani and Sharma [1]).

- ii. A class of ad-hoc learning rules (generalizing the learning rule of the previous point) where the posterior distribution is a convex combination of the prior and the empirical distribution with the weight on the empirical distribution non-decreasing with additional observations:

$$F^i(x \mid X_{r^i}^i) = (1 - a_{r^i})F(x) + a_{r^i}H(x \mid X_{r^i}^i)$$

with $a_{r^{i+1}} \geq a_{r^i}$, $F(x)$ being the prior distribution before search started and $H(\cdot \mid X_{r^i}^i)$ being the empirical distribution based on the observations $X_{r^i}^i$ (Bikchandani and Sharma [1]).

- iii. A non-parametric learning procedure used in Chou and Talmain ([3]) that makes no assumptions on the underlying class of probability distributions and is constructing $F(\cdot \mid X_{r^i}^i)$ according to the

The searcher implicitly obtains some utility U from consuming the good and minimizes over all search strategies σ the expectation of the price paid plus search costs, i.e. $\min_{\sigma} E[p_{\sigma} + c_{\sigma}]$. Rephrasing the search problem as one of looking for rewards with $r_{\sigma} = U - p_{\sigma}$, the above minimization problem is equivalent to $\max_{\sigma} E[r_{\sigma} - c_{\sigma}]$ which is the problem considered in this paper. Furthermore, if the optimal search strategy σ^* of the minimization problem is a sequence of increasing reservation prices $\{p_i^*\}_{i=1}^N$ such that search is continued if the best offer $p > p_i^*$ and search is terminated if $p \leq p_i^*$, this implies a sequence of reservation rewards $\{r_i^*\}_{i=1}^N$ with $r_i^* = U - p_i^*$ that is decreasing and where search is continued if the best offer $r < r_i^*$ and search is terminated if $r \geq r_i^*$.

maximum entropy principle. Suppose the searcher knows that outcomes are distributed between some interval $[a, b]$. The conditional probability of some outcome x , having observed $x^1 \leq x^2 \leq \dots \leq x^{r^i}$ (not necessarily in this order) is obtained by assigning to each interval $[a, x^1], [x^1, x^2], \dots, [x^{r^i}, b]$ probability mass $\frac{1}{r^i+1}$ uniformly distributed (and a point mass if $x^i = x^{i+1}$).

iv. Let the offer distribution be exponential with unknown origin θ :

$$f(x | \theta) = ae^{a(x-\theta)} \text{ for } x \leq \theta$$

Learning is Bayesian and priors are such that the logarithm of the prior distribution $\log(p(\theta))$ is concave (see Rosenfield and Shapiro [14]).

4.3 Results

The following theorem states the optimal sampling strategy for the search problem with learning and contains the main result of this paper. Its proof is deferred to the appendix. The optimal rule is just the benchmark rule applied to repeatedly updated reservation prices.

Theorem 1 *Given A1 holds, the following sampling strategy is optimal:*

- Step 1** *With the available observations calculate the reservation prices for each alternative and go to step 2.*
- Step 2** *If there is no closed box with a reservation price higher than the current best offer y , then stop search and accept y , otherwise continue with step 3.*
- Step 3** *Search the alternative (or one of them, if there are several) with the highest reservation price and go back to step 1.*

The theorem tells us that the reservation prices which are based solely on current beliefs are sufficient to determine the optimal sampling strategy. The optimality of such a focus on current beliefs might be

surprising. In fact, using the rate of return interpretation of reservation prices from section 3, the rule tells us to sample the alternatives with the *currently* highest returns.

In a learning context information is valuable as well, since it enables the searcher to make better search decisions in the future. In general, it might therefore be worth to give up payoffs in the short term to obtain information that allows to make decisions with a higher payoff in the long run.

In the considered search problem there is no such trade-off between the information gain and the payoff gain and focusing on the payoff gain alone is sufficient to obtain optimality. The reason for this is to be found in the restrictions on the learning rules I imposed. They exhibit enough monotonicity to prevent the searcher from optimally going through a 'payoff-valley' to potentially reach a higher 'payoff-mountain' later on.

Obviously, the possibility of strong learning could give an incentive to go through the 'payoff-valley', and therefore I had to rule it out. However, it is not immediately clear why the remaining learning processes do not give such an incentive. To get some intuition on this point consider the following example.

Imagine to have two search alternatives, a blue one and an orange one. Suppose that at current information both have identical expected distribution functions and thereby equal reservation prices. In terms of payoffs the boxes are therefore identical. There is, however, only one blue box left, while there are still many orange boxes. Sampling the blue box therefore reveals no information on any other search opportunity, while sampling an orange box reveals information about all the remaining orange boxes. Thus, in addition to the payoff, opening an orange box provides information. It therefore seems better to open an orange box than to open the blue box.

Surprisingly, the optimal search rule in theorem 1 states that it does not matter whether a blue or an orange box is opened first. The intuition behind this result can be obtained by considering the rate of

return interpretation of reservation prices more carefully: Sampling the blue box or an orange box has equal rates of return but after sampling an orange box, the remaining orange boxes will have a lower rate of return (due to $A1$).¹³

Reinterpret this search problem as a search problem without learning (where the benchmark rule is optimal): There are in fact two boxes with a high and equal rate of return, the blue box and the first sampled orange box, and many boxes with lower rates of return, the remaining orange boxes. It is irrelevant for determining the sampling order of the first two boxes to know how much lower the rate of return for the remaining orange boxes is: We know that it is optimal to sample the boxes with the highest reservation rate of return and one can do this two times without this knowledge. Or, equivalently, it is sufficient to know the rates of return for the remaining orange boxes after *both* high rate of return boxes have been sampled and not important to know it already after the first of them has been sampled. Therefore, it does not matter whether the blue box or an orange box is sampled first. I will come back to the reinterpretation of the search problem with learning as one without learning in much more detail in section 4.4.

The optimal sampling procedure above has changed only slightly when compared to the benchmark sampling rule. An informed searcher had to calculate reservation prices only once, while a learning searcher has to permanently adapt them in the light of new information. Step 3 of the rule therefore points back to step 1. For the rest, the rule remains unchanged. This slight change, however, alters optimal search behavior substantially, as illustrates the following example.

Example 2 *Suppose that there are only two search alternatives, a red one and a green one, but many boxes of each alternative. Boxes have only two kinds of outcomes: "success", identified with a payoff equal to 1, or "failure", identified with a payoff equal to zero. The true probabilities for*

¹³We abstract here from the possibility that search stops to make the argument as simple as possible.

success and failure for the respective alternatives are indicated in table 3 below.

Table 3:

	payoff	0	1
Red	with probability	0.5	0.5
Green	with probability	0.7	0.3

In addition, assume a discount factor equal to 1, sampling costs for both boxes equal to 0.1 and the value of the outside option equal to 0.

a.) Optimal sampling strategy under full information

Consider first the sampling strategy under full information. Knowing the true probabilities of outcomes, the reservation prices are $R^{red} = 0.8$ and $R^{green} = 0.\bar{6}$. Hence, an informed searcher prefers to open red boxes and stops with the first success. Suppose that the searcher encounters a sequence of failures. Optimally, his strategy is to continue opening red boxes until they have all been opened and to switch then to the opening of green boxes. Green boxes are opened until a success is encountered or all of them have been searched. Notice the following feature of the optimal strategy: Since the ranking of alternatives is constant during the search process, the searcher does not switch sampling from one alternative to another, unless there are no boxes of that alternative left.

b.) Optimal Sampling Strategy with Learning

Now consider a searcher that is uncertain about the true underlying probability distribution and is learning by taking a convex combination between his prior distribution and the empirical distribution function (This is the second learning rule in section 4.2):

$$F^i(x | X_{r^i}^i) = (1 - a_{r^i})F(x) + a_{r^i}H(x | X_{r^i}^i)$$

Let the weight on the empirical distribution be $a_{r^i} = \frac{r^i}{1+r^i}$ and the searcher's priors $F(x)$ be unbiased in the sense that they are equal to the true underlying probability function as shown in table 3.

At the beginning of search, reservation prices are therefore equal to the ones of an informed searcher, but as the searcher makes additional observations, they are adjusted downwards. The ranking of alternatives is therefore changing during the search process. The searcher might well search green boxes before all red boxes have been opened. Negative results from searching red boxes 'bid' down their reservation price and make the searcher believe that green boxes are more interesting. The same reasoning applied to green boxes might cause a switch back to sampling red boxes again. In further contrast to the full information case, sampling might even stop with a failure and not all boxes been searched because of beliefs having worsened so much that the outside option looks more profitable than continued search.

The previous effects can be seen in table 2 for the above learning rule and a sequence of failures. The table reads as follows. The first column indicates the search stage, the second the number of so far made observations of red and green boxes (i.e. the number of observed failures of each), the following two columns show the current reservation prices. The last column gives the optimal search strategy according to theorem 1. The searcher bids down reservation prices and switches between sampling red and green boxes in response to failures until finally the reservation prices of both boxes are so low that the outside option appears more attractive than continued search.

Table 2

t	(r^{red}, r^{green})	R^{red}	R^{green}	Optimal strategy
0	(0,0)	0.8	0.66	search a red box
1	(1,0)	0.6	0.66	search a green box
2	(1,1)	0.6	0.33	search a red box
3	(2,1)	0.4	0.33	search a red box
4	(3,1)	0.2	0.33	search a green box
5	(3,2)	0.2	0.0	search a red box
6	(4,2)	0.0	0.0	stop search and take outside option

4.4 An Equivalent Search Problem

In this section I explain why the benchmark rule generalizes to the case with learning. The argument is quite abstract but the general idea is as follows: To the original search problem with learning P one can construct an equivalent search problem P^e without learning in the following sense: To any search rule S of the original problem exists a corresponding search rule S^e , such that S^e yields in P^e the same payoff as S in P . It follows that the optimal rule S^* for the original problem P is then just the rule corresponding to the optimal rule S^{e*} in P^e (given the corresponding rule S^* to S^{e*} exists). It is easy to show that the optimal rule S^{e*} in P^e is the benchmark rule. One can show that the rule S^* corresponding to S^{e*} exists and is just the generalization of the benchmark rule found to be optimal in theorem 1.¹⁴

Consider a search problem with learning P . P is described by the number I of alternatives, the numbers M^i of boxes of each alternative $i = 1, \dots, I$, prior beliefs and the learning rule. A sampling rule S for P is a mapping from the set of available information $(X_{r1}^1, X_{r2}^2, \dots, X_{rn}^n)$ to the set of integers $\{0, 1, \dots, I\}$, where $S = 0$ indicates to stop search and $S = i$ for $i \geq 1$ indicates to continue search with an i -box.

At the beginning of search the M^i boxes of alternative i have an expected distribution function $F^i(\cdot)$ and an associated reservation price R_0^i . As soon as one i -box has been opened, the expected distribution of outcomes for the remaining $M^i - 1$ boxes from i becomes $F^i(\cdot|x_1^i)$ and the associated reservation price $R^i(x_1^i)$. If still another box of this alternative is opened, then the remaining $M^i - 2$ boxes from alternative i have expected distribution $F^i(\cdot|x_1^i, x_2^i)$ and associated reservation price $R^i(x_1^i, x_2^i)$, etc.

Alternatively, one could interpret the previous observations as fol-

¹⁴What follows is basically a sketch of an alternative proof for the optimality of the sampling strategy of theorem 1.

lows: In fact, the searcher has only one box with reservation price R_0^i (after sampling one i -box the reservation price of the remaining i -boxes changes), one box with reservation price $R^i(x_1^i)$, another one with $R^i(x_1^i, x_2^i)$, and so on.

Based on this alternative interpretation I will construct the equivalent search problem P^e : There is an equal number of boxes as in P , namely $\sum_i M^i$. The first M^1 boxes in P^e have reservation prices

$$R_0^1, R^1(w_1^1), R^1(w_1^1, w_2^1), \dots, R^1(w_1^1, w_2^1, \dots, w_{M^1-1}^1)$$

respectively, where the $R^1(\cdot)$ are the reservation price functions of 1-boxes in P . For the moment take the values w_j^1 as given. From assumption A1 we know that (for any values of the w_j^1) reservation prices of the boxes can be ordered as¹⁵

$$R_0^1 \geq R^1(w_1^1) \geq R^1(w_1^1, w_2^1) \geq \dots \geq R^1(w_1^1, w_2^1, \dots, w_{M^1-1}^1)$$

Similarly, let the next M^2 boxes in P^e have reservation prices

$$R_0^2 \geq R^2(w_1^2) \geq R^2(w_1^2, w_2^2) \geq \dots \geq R^2(w_1^2, w_2^2, \dots, w_{M^2-1}^2)$$

Continue to assign reservation prices to boxes in the above manner until each box in P^e got one reservation price.

For simplicity, I will refer to the first M^1 boxes in P^e also as "1-boxes" (always in quotation marks), to the next M^2 boxes as "2-boxes", etc., since their reservation price functions correspond to the respective alternatives in P . Notice that for given values w_j^i , P^e is the benchmark search problem and hence the benchmark strategy is the optimal sampling rule.

Consider the M^i "alternative i boxes" in P^e and suppose the following informational structure: At the beginning of search the values w_j^i are unknown to the searcher. R_0^i is hence the only known reservation

¹⁵We ignore the potential increase in the reservation price admitted by A1 because it leads to termination of search.

price of "i-boxes". However, it is known to the searcher that, whatever the value of the w_j^i , the remaining "i-boxes" have some lower reservation value ordered as listed above.

As search proceeds the searcher observes gradually the variables w_j^i . The first value w_1^i is observed after the first "i-box" with reservation price R_0^i has been opened and w_2^i after the second with $R^i(w_1^i)$ has been opened, and so on. In short, the searcher knows only the reservation price of the best unopened box of each "alternative". The reservation price of the next best box of some "alternative" is revealed only after the previously best box of the same "alternative" has been opened. The information available to the searcher is sufficient to execute the benchmark rule in P^e , since the highest reservation price is just the reservation price of the best of all best "i-boxes".

As already mentioned for given sequences the benchmark rule is clearly optimal. However, I am interested in stochastic sequences, since the x_j^i in P are stochastic as well. For stochastic $w_1^i, w_2^i, \dots, w_{M^i-1}^i$ the optimality of the benchmark rule in P^e is in general not guaranteed. If the distribution of the w_j^i could be influenced by sampling decisions, then the searcher could change the expected reservation prices of closed boxes and thereby the value of search. However, as long as the stochastic nature of the sequences can not be influenced by the searcher's sampling decisions, it is optimal to sample according to the benchmark rule because it is optimal to do so for all given sequences.¹⁶

To make P^e equivalent to P it remains to specify a particular distribution for the w_j^i . This can be done by choosing the $w_1^i, w_2^i, \dots, w_{M^i-1}^i$ to have a true distribution equal to the expected distribution of the $x_1^i, x_2^i, \dots, x_{M^i-1}^i$ in P , i.e. w_1^i to be drawn from $F^i(\cdot)$, w_2^i from $F^i(\cdot|w_1^i)$,

¹⁶The assumption that the distribution of the w_j^i can not be influenced by the searcher will be important when considering in the next section the potential optimality of the search rule with more general learning rules.

w_3^i from $F(\cdot|w_1^i, w_2^i)$, and so on.¹⁷ P^e is now equivalent to P in the following sense:

- i. There is an 'informational' equivalence: Having sampled r^i times "i-boxes" in P^e the searcher has observed $W_{r^i}^i = w_1^i, w_2^i, \dots, w_{r^i}^i$ while when having sampled r^i times i-boxes in P the searcher has observed $X_{r^i}^i = x_1^i, x_2^i, \dots, x_{r^i}^i$.
- ii. Suppose that $W_{r^i}^i$ has been observed in P^e and the same sequence $X_{r^i}^i = W_{r^i}^i$ in P . Then the reservation price of the best "i-box" in P^e is equal to the reservation price of all i-boxes in P .
- iii. Suppose that $W_{r^i}^i$ has been observed in P^e and the same sequence $X_{r^i}^i = W_{r^i}^i$ in P . Then opening the best "i-box" in P^e has (objective) expected utility equal to the (subjective) expected utility of opening an i-box in P .
- iv. The (objective) probability to observe some particular sequence $W_{r^i}^i$ in P^e equals the (subjective) probability to observe the same sequence $X_{r^i}^i = W_{r^i}^i$ in P .

It is now easy to define a sampling rule S^e corresponding S . S^e must be the same as S but evaluated at $W_{r^1}^1, W_{r^2}^2, \dots, W_{r^n}^n$ and specifying to sample the best "i-box", whenever S would specify to sample some i-box, i.e. $S^e = S(W_{r^1}^1, W_{r^2}^2, \dots, W_{r^n}^n)$ with $S^e = 0$ indicating to stop search and $S^e = i$ with $i \geq 1$ indicating to sample the *best* "i-box". From observations (3) and (4) above should be clear that S^e in P^e achieves an (objective) expected payoff equal to the (subjective) expected payoff of S in P .

The optimal search rule S^{e*} in P^e is the benchmark rule which states to sample the best of all best "i-boxes". The optimal rule S^* in

¹⁷Clearly, with such a specification the distribution of any w_j^i is independent from sampling decisions.

P must be the rule corresponding to S^{e*} .¹⁸ From the above definition is easily seen that the strategy S^* corresponding to S^{e*} is to sample the search alternative with the currently highest reservation price which is precisely the optimal sampling rule from theorem 1.

5 Limitations and Extensions

In this section I discuss the possibility to generalize the class of admitted learning rules such that the search rule of theorem 1 preserves its optimality. I mainly consider a relaxation of the assumption on the independence of boxes from different alternatives, since it is the most restrictive and unrealistic one. Unfortunately, the optimality of the proposed search rule turns out to be relatively sensitive to it.

The following simple example shows that in general the proposed sampling strategy is not optimal when the reservation price of some search alternative is affected by the search outcomes of another alternative.

Example 3 *Suppose that there are only two alternatives, red and green, and only one box of each alternative.¹⁹ Currently, the searcher's expected payoff distributions of the respective alternatives are as given in table 4 below. Assume the current best offer to be $y = 0.5$. Without discounting and search costs equal to 0.1 for both alternatives, the reservation prices are $R^1 = 0.7$ and $R^2 = 0.65$ for alternative 1 and 2, respectively.*

Suppose that the searcher first samples a red box, as theorem 1 suggests. Furthermore, suppose that learning is such that the new reservation

¹⁸The optimality of R^* in P follows from the following considerations: The expected value of search in P^e can take on at least all expected values of P , since to every R in P there exists a corresponding R^e in P^e taking on the same value. Therefore, if a R^* corresponding to R^{e*} exists it must be optimal in P .

¹⁹With only one box of each type learning can take place only across different types of boxes.

Table 4:

<i>Red</i>	<i>payoff</i>	0	0.5	1
	<i>with probability</i>	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
<i>Green</i>	<i>payoff</i>	0	0.6	0.95
	<i>with probability</i>	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$

price of the green box drops below 0.5 when the outcome of the red box is 0 or 0.5 and that it is anything smaller than 1 if the outcome of the red box is 1. Interpret this as low outcomes of the red box revealing that low outcomes of the green box are more likely to occur.

With these assumptions search optimally stops after sampling the red box, independently from the search result. The new reservation price of the green box is always below the new best offer and a sampling of the green box would result in an expected loss.²⁰ The expected payoff of sampling the red box first is therefore $\frac{2}{3}0.5 + \frac{1}{3}1 = \frac{2}{3}$.

Consider the alternative strategy of opening the green box first and stopping search thereafter. The expected payoff is $\frac{1}{3}0.5 + \frac{1}{3}0.6 + \frac{1}{3}0.95 = 0.68\bar{3} > \frac{2}{3}$. Clearly, opening the red box first is not optimal although its reservation price is higher than that of the green box.

To see why the sampling rule might be sub-optimal in this more general setting consider the equivalent search problem P^e to the original search problem P I constructed in section 4.4. For the benchmark rule to be optimal in the P^e (and its corresponding rule in P), it was crucial that the searcher could not influence the distribution of the sequences $w_1^i, w_2^i, \dots, w_{M^i-1}^i$ by his sampling decisions. In the above example this assumption is not fulfilled. By sampling the red box the distribution of the green box changes. The benchmark sampling rule is therefore not necessarily optimal in P^e . The same holds in turn for its corresponding rule in P .

²⁰This is easily verified looking at the definition of the reservation price.

In one special case the optimal sampling rule generalizes to dependent alternatives. Recognize that theorem 1 requires only the reservation prices R^i of i -boxes not to depend on observations x^j of j -boxes ($j \neq i$).

This is not equivalent with the expected payoff distributions F^i being independent from x^j . It is possible that observations of j -boxes affect the distribution F^i without changing R^i . Clearly, as long as x^j leaves the values of F^i above the current best offer y unchanged, then the reservation price R^i of i -boxes will be unaffected, given that $R^i > y$. If $R^i < y$, then R^i might change but $R^i < y$ will hold also after the change.²¹ Since boxes with reservation prices below the current best offer are irrelevant for the search problem and all reservation prices above the best current offer are unaffected, this special case of dependent boxes is covered by theorem 1.

6 Conclusions

This paper constructed the optimal sampling strategy for a search problem where the searcher faces different search alternatives and is learning about these alternatives during the search process. I thereby unified and generalized two kinds of earlier contributions: search problems with learning but identical search opportunities and search problems with distinguishable search alternatives but without learning.

The optimal sampling rule is characterized by a simple reservation price criterion. The rule implies that search opportunities with higher reservation prices should be sampled before ones with lower reservation prices. In contrast to the full information case, the ordering of different search alternatives in terms of reservation prices keeps changing during the search process. Learning therefore makes a substantial difference for the optimal sampling order. At the same time the sampling rule retains its simple structure and learning can be accounted for without complicating the analysis.

²¹To verify these claims simply check the definition of the reservation price.

The independence of different search alternatives has been found to be crucial for the optimality of the sampling rule and finding conditions on the learning process that allow for an extension of the results to the case of dependent search alternatives is left for future research.

7 Appendix

Lemma 2 *If either $\beta_i < 1$ or $c_i > 0$, then a unique reservation price exists.*

Proof of Lemma 2: The function $Q^i(X_{r^i}^i, \cdot)$ is continuous, differentiable and decreasing.

$$\frac{d}{dy}Q^i(X_{r^i}^i, y) = \frac{d}{dy} \left[\beta_i \int_y^\infty (x^i - y) dF^i(x^i | X_{r^i}^i) \right] - (1 - \beta_i) \quad (3)$$

$$= \beta_i \int_y^\infty -1 dF^i(x^i | X_{r^i}^i) \quad (4)$$

$$- [(x^i - y) dF^i(x^i | X_{r^i}^i)]_{x^i=y} - (1 - \beta_i) \quad (5)$$

$$= -\beta_i(1 - F^i(y | X_{r^i}^i)) - (1 - \beta_i) \quad (6)$$

$$\leq 0 \quad (7)$$

Since

$$Q^i(X_{r^i}^i, -\infty) = \infty \quad (8)$$

$$Q^i(X_{r^i}^i, +\infty) = \left\{ \begin{array}{l} -\infty \text{ if } \beta_i < 1 \\ -c^i \text{ if } \beta_i = 1 \end{array} \right\} \quad (9)$$

a solution exists. If $\beta_i < 1$, then $\frac{d}{dy}Q^i < 0$ and the solution is also unique. If $\beta_i = 1$, then $\frac{d}{dy}Q^i < 0$ only if $F^i(y | X_{r^i}^i) < 1$. With $c^i > 0$ this is guaranteed at the reservation price: $F^i(R^i(X_{r^i}^i) | X_{r^i}^i) = 1$ implies $Q^i(X_{r^i}^i, R^i(X_{r^i}^i)) < 0$ which contradicts the definition of the reservation price. ■

Proof of Theorem 1:²² I begin by proving the optimality of the stopping rule (i.e. step 2 of the theorem). If there is some i -box with

²²The proof owes the construction of the strategy S' to Weitzman [21]. An alternative proof along the lines of section 4.4 could be given reducing the problem to one where Weitzman's results apply. However, since he did not consider the case of infinitely many search boxes, we would then not cover this case.

$R^i > y$ ($R^i < y$), then the one period gain $Q^i > 0$ ($Q^i < 0$). Therefore, as long as there is some closed i -box with $R^i > y$ stopping cannot be optimal, since opening an i -box and stopping then gives already a higher payoff. If all closed boxes have a reservation value below y , then $A1$ insures that reservation prices will also be below the best offer in all future search steps. Gains from continued search will always be negative and stopping is therefore optimal.

Suppose that S is a sampling rule where stopping is optimal as derived above. In addition, suppose that S specifies at some search stage to sample a k -box with reservation price R^k and in case that the stopping rule prescribes continuation in the next search step an l -box with $R^l > R^k$. I will show that S cannot be optimal. To do so I will construct an alternative sampling rule S' and show that S' has higher expected valued than S . S' is like S but interchanges the sampling order such that the box with the higher reservation price R^l is sampled first and the one with the lower reservation price R^k thereafter.²³

Before constructing S' and proving the claim I have to introduce some notation. At the search stage where S specifies to sample a k -box, let the previous observations of search outcomes be $\{X_{r^i}^i\}_{i=1}^I$ and the current best offer $y = \max\{X_{r^1}^1, X_{r^2}^2 \dots X_{r^I}^I\}$. Define

$$R^j = \max_{i|r^i < M^i} R^i(X_{r^i}^i)$$

$$R^{h(x_{r^{j+1}}^j)} = \max \left\{ R^j(X_{r^{j+1}}^j) | r^j + 1 < M^j, \max_{i|i \neq j \wedge r^i < M^i} \{R^i(X_{r^i}^i)\} \right\}$$

j -boxes have currently the highest reservation price of all closed boxes and h -boxes are the ones that have the highest reservation price after one j -box has been sampled and the search outcome $x_{r^{j+1}}^j$ been observed.²⁴ h may depend on $x_{r^{j+1}}^j$ because the decrease of the reservation price of j -boxes depends on $x_{r^{j+1}}^j$.

²³Notice that the sampling order of S' is feasible. $k = l$ is not possible, since reservation prices k -boxes are decreasing with additional information. Therefore, the box with the reservation price R^l is already available before having sampled the k -box.

²⁴ $r^i < M^i$ is a condition insuring that there is still an unopened i -box.

By assumption we know that

$$R^k(X_{r^k}^k) < R^l(X_{r^l}^l) \leq R^j$$

We should distinguish two cases: $l \neq j$ and $l = j$. The first case is the easier one: The highest reservation price R^j remains unaffected by the sampling of a k - and an l -box. The optimal stopping criterion is therefore the same in both search stages: Stop if the current best offer is larger than R^j and continue otherwise. In the second case the best reservation price drops to $R^{h(x_{r^j}^j)}$ after sampling the l -box ($l = j$). The stopping criterion therefore changes when sampling the l -box. I will only consider this more complicated case.²⁵

Recall that the rule S specifies to sample first a k -box and in case of continuation a j -box with the stopping decision being optimal as derived above. Figure 1a gives a graphical representation of the strategy for the first two search steps. Depending on the search outcome several cases can be distinguished that are represented by branches. The values written at the end of these branches represent the payoffs for the respective cases. If search outcomes fall into the case represented by the lowest branch, then search continues. Φ represents the value of continued search with rule S for this case.

The proposed alternative strategy S' differs from S for the first two search steps but is identical to S for later search steps: S' specifies to first sample a j -box (instead of a k -box). If the new best offer $\max\{y, x_{r_{j+1}}^j\} \geq R^{h(x_{r_{j+1}}^j)}$, then S' specifies to stop search. Otherwise it prescribes to sample a k -box and to continue as prescribed by the rule S . Figure 2a represents the sampling rule S' graphically. Again, Φ represents the value of continued search with rule S' when search is optimally continued. This value is the same as the value of search with rule S because by definition S' equals S for all steps after the second.

The following notation will prove useful to calculate the expected payoffs of S and S' :

²⁵The results for the first case can be obtained by replacing j by l and $R^{h(x_{r^j}^j)}$ by R^j in the following.

$$\begin{aligned}
\Pi^k &= \Pr(x_{r^k+1}^k \geq R^j) \\
\Pi^j &= \Pr(x_{r^j+1}^j \geq R^j) \\
\lambda_k &= \Pr(R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)}) \\
\lambda_j &= \Pr(R^j > x_{r^j+1}^j \geq R^{h(x_{r^j+1}^j)}) \\
\mu^k &= \Pr(R^{h(x_{r^j+1}^j)} > x_{r^k+1}^k \geq R^k) \\
w^k &= E \left[x_{r^k+1}^k | x_{r^k+1}^k \geq R^j \right] \\
w^j &= E \left[x_{r^j+1}^j | x_{r^j+1}^j \geq R^j \right] \\
\tilde{v}^k &= E \left[\max \{ x_{r^k+1}^k, y \} | R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)} \right] \\
\tilde{v}^j &= E \left[\max \{ x_{r^j+1}^j, y \} | R^j > x_{r^j+1}^j \geq R^{h(x_{r^j+1}^j)} \right] \\
v^k &= E \left[x_{r^k+1}^k | R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)} \right] \\
u^k &= E \left[x_{r^k+1}^k | R^{h(x_{r^j+1}^j)} > x_{r^k+1}^k \geq R^k \right] \\
d &= E \left[\max \left\{ x_{r^j+1}^j, x_{r^k+1}^k, y \right\} | R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)}, \right. \\
&\quad \left. R^j > x_{r^j+1}^j \geq R^{h(x_{r^j+1}^j)} \right] \\
\Phi &= E \left[\Psi(\bar{S} \setminus \{j, k\}, \max \{ x_{r^j+1}^j, x_{r^k+1}^k, y \}, S | R^{h(x_{r^j+1}^j)} > x_{r^k+1}^k, \right. \\
&\quad \left. R^{h(x_{r^j+1}^j)} > x_{r^j+1}^j \right]
\end{aligned}$$

All probabilities and expectation operators are conditional on the information $\{X_{r^i}^i\}_{i=1}^I$. The function $\Psi(\bar{S} \setminus \{j, k\}, \max \{x_{r^j+1}^j, x_{r^k+1}^k, y\}, S)$ represents the value of continued search when the set of closed boxes is \bar{S} less one j - and one k -box, the best offer is $\max \{x_{r^j+1}^j, x_{r^k+1}^k, y\}$ and the sampling rule S or S' . Figure 1b describes the probabilities and the expected payoffs of strategy S for the cases distinguished in Figure 1a using the above notation. Similar does Figure 2b for strategy S' and the cases of Figure 2a. Looking at these figures reveals that the expected payoffs of the strategies S and S' are

$$S = -c^k + \beta^k \Pi^k w^k + \lambda^k \beta^k (-c^j + \Pi^j \beta^j w^j + \lambda^j \beta^j d) \quad (10)$$

$$+ (1 - \Pi^j - \lambda^j) \beta^j \tilde{v}^k + (1 - \Pi^k - \lambda^k) \beta^k (-c^j + \Pi^j \beta^j w^j) \quad (11)$$

$$+ \lambda^j \beta^j \tilde{v}^j + (1 - \Pi^j - \lambda^j) \beta^j \Phi \quad (12)$$

$$S' = -c^j + \Pi^j \beta^j w^j + \lambda^j \beta^j w^j + (1 - \Pi^j - \lambda^j) \beta^j (-c^k + \Pi^k \beta^k w^k) \quad (13)$$

$$+ \lambda^k \beta^k w^k + (1 - \Pi^k - \lambda^k) \beta^k \Phi \quad (14)$$

The payoff difference between S' and S is

$$\begin{aligned}
S' - S &= (c^k - \Pi^k \beta^k w^k) (1 - (1 - \Pi^j - \lambda^j) \beta^j) \\
&\quad - (c^j - \Pi^j \beta^j w^j) (1 - (1 - \Pi^k) \beta^k) - d \\
&\quad + (1 - (1 - \Pi^k - \lambda^k) \beta^k) \lambda^j \beta^j
\end{aligned} \tag{15}$$

From the definition of the reservation prices we have

$$c^j = \Pi^j \beta^j (w^j - R^j) - (1 - \beta^j) R^j \tag{16}$$

$$\begin{aligned}
c^k &= \beta^k (\Pi^k (w^k - R^k) + \lambda^k (v^k - R^k)) \\
&\quad + \mu^k (u^k - R^k) - (1 - \beta^k) R^k
\end{aligned} \tag{17}$$

Substituting (16) and (17) into (15) gives:

$$\begin{aligned}
S' - S &= (R^j - R^k) (1 - \beta^j (1 - \Pi^j)) (1 - \beta^k (1 - \Pi^k)) \\
&\quad + (v^k - R^k) (\lambda^k \beta^k (1 - \beta^j (1 - \Pi^j))) \\
&\quad + (\tilde{v}^j - R^k) (\lambda^j \beta^j (1 - \beta^k (1 - \Pi^k))) \\
&\quad + (u^k - R^k) (\mu^k \beta^k (1 - \beta^j (1 - \Pi^j - \lambda^j))) \\
&\quad + (\tilde{v}^j + v^k - R^k - d) \lambda^k \beta^k \lambda^j \beta^j
\end{aligned} \tag{18}$$

Furthermore,

$$\begin{aligned}
d &= E \left[\max \left\{ \max \{x_{r^j+1}^j, y\} - R^{h(x_{r^j+1}^j)}, x_{r^k+1}^k - R^{h(x_{r^j+1}^j)} \right\} \right. \\
&\quad \left. + R^{h(x_{r^j+1}^j)} \mid R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)}, R^j > x_{r^j+1}^j \geq R^{h(x_{r^j+1}^j)} \right] \\
&\leq E \left[\max \left\{ x_{r^j+1}^j, y \right\} - R^{h(x_{r^j+1}^j)} + x_{r^k+1}^k - R^{h(x_{r^j+1}^j)} \right. \\
&\quad \left. + R^{h(x_{r^j+1}^j)} \mid R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)}, R^j > x_{r^j+1}^j \geq R^{h(x_{r^j+1}^j)} \right] \\
&= \tilde{v}^j + v^k \\
&\quad - E \left[R^{h(x_{r^j+1}^j)} \mid R^j > x_{r^k+1}^k \geq R^{h(x_{r^j+1}^j)}, R^j > x_{r^j+1}^j \geq R^{h(x_{r^j+1}^j)} \right] \\
&\leq \tilde{v}^j + v^k - R^k
\end{aligned}$$

The last inequality is due to the fact that by definition $R^{h(x_{r^j+1}^j)} \geq R^k$ for any realization of $x_{r^j+1}^j$. Therefore, any term in (18) is greater or equal zero with the first term being strictly greater than zero. This proves the sub-optimality of any strategy of the form S . Optimal strategies must

sample boxes in the order of decreasing reservation prices. However, this does not mean that at each search stage the box with the highest reservation price has to be sampled as the theorem prescribes. I will turn attention to this point in the following.

Suppose T is a sampling rule that stops according to the optimal stopping rule and samples boxes in order of decreasing reservation prices. However, suppose that T specifies at some search stage not to sample the box with the currently highest reservation price. I will show that T cannot be optimal by proving that there exists a strategy T' that has a higher expected value.

Suppose again that available observations are $\{X_{r^i}^i\}_{i=1}^I$ and that j -boxes have the highest reservation price equal to R^j . T specifies to sample a k -box with $R^k < R^j$. Thereafter (in case of continued search), T prescribes to sample l, m, n, \dots -boxes with $R^k \geq R^l \geq R^m \geq R^n \geq \dots$.²⁶ Since a sampling of a j -box is incompatible with the assumption of sampling in order of decreasing reservation prices, j -boxes will never be sampled. The optimal stopping rule then implies that search stops only if $y \geq R^j$. To calculate the expected value of search rule T define for $\alpha = j, k, l, m, \dots$:

$$\begin{aligned}\Pi^\alpha &= \Pr(x_{r^\alpha + \#\alpha}^\alpha \geq R^j) \\ w^\alpha &= E[x_{r^\alpha + \#\alpha}^\alpha | x_{r^\alpha + 1}^\alpha \geq R^j]\end{aligned}$$

where $\#\alpha$ is the number of alternative α -boxes in the sequence $k, l, m, \dots \alpha$.²⁷ Π^α is the probability that search stops when sampling box α . w^α is the expected value of x^α given that search stops.²⁸

The expected value of T is easily calculated to be

$$T = [-c^k + \beta^k \Pi^k w^k]$$

²⁶ m might depend on the outcome $x_{r^{k+1}}^k$, similarly the types n, l, \dots might depend on previous observations. For notational simplicity, we will ignore this dependence.

²⁷Remember that each alternative k, l, m, \dots is a number from the set $\{1, 2, \dots, I\}$. For example, if $\alpha = n$ and $k, l, m, n = 1, 4, 3, 4$ then $\#\alpha = 2$, i.e. it is the second box of alternative 4.

²⁸Probabilities and expectations are again conditional on the available information $\{X_{r^i}^i\}_{i=1}^I$.

$$\begin{aligned}
& +\beta^k(1 - \Pi^k) [-c^l + \beta^l \Pi^l w^l] \\
& +\beta^k(1 - \Pi^k)\beta^l(1 - \Pi^l) [-c^m + \beta^m \Pi^m w^m] \\
& +\beta^k(1 - \Pi^k)\beta^l(1 - \Pi^l)\beta^m(1 - \Pi^m) [-c^n + \beta^n \Pi^n w^n] \\
& + \dots
\end{aligned} \tag{19}$$

Now consider the following alternative strategy T' . T' uses the same stopping rule as T : Stop if $y \geq R^j$ and continue otherwise. However, T' samples first a j -box and then (in case of continuation) k, l, m, n, \dots boxes. The expected value of T' is

$$T' = [-c^j + \beta^j \Pi^j w^j] + \beta^j(1 - \Pi^j)T$$

Remembering from the definition of the reservation price that

$$c^\alpha = \beta^\alpha \Pi^\alpha (w^\alpha - R^j) + \beta^\alpha \lambda^\alpha (v^\alpha - R^\alpha) - (1 - \beta^\alpha) R^\alpha \tag{20}$$

where

$$\begin{aligned}
\lambda^\alpha &= \Pr(R^j > x_{r^\alpha+1+\#\alpha}^\alpha \geq R^\alpha) \\
v^\alpha &= E[x_{r^\alpha+1+\#\alpha}^\alpha | R^j > x_{r^\alpha+1+\#\alpha}^\alpha \geq R^\alpha]
\end{aligned}$$

I can calculate the payoff difference between T' and T to be

$$\begin{aligned}
T' - T &= [-c^j + \beta^j \Pi^j w^j] + [\beta^j(1 - \Pi^j) - 1] T \\
&= \beta^j \Pi^j R^j + (1 - \beta^j) R^j + [\beta^j(1 - \Pi^j) - 1] T \\
&= [\beta^j(\Pi^j - 1) + 1] (R^j - T)
\end{aligned} \tag{21}$$

The first bracket in the last line of (21) is strictly positive.²⁹ It remains to show that $R^j > T$. Substituting (20) into (19) and recognizing that

²⁹From the definition of the reservation price we have $\Pi^j > 0$.

$\beta^\alpha \lambda^\alpha (v^\alpha - R^\alpha) \geq 0$ we obtain

$$\begin{aligned}
T &= [\beta^k \Pi^k R^k + (1 - \beta^k) R^k - \beta^k \lambda^k (v^k - R^k)] \\
&\quad + \beta^k (1 - \Pi^k) [\beta^l \Pi^l R^l + (1 - \beta^l) R^l - \beta^l \lambda^l (v^l - R^l)] \\
&\quad + \beta^k (1 - \Pi^k) \beta^l (1 - \Pi^l) [\beta^m \Pi^m R^m + (1 - \beta^m) R^m \\
&\quad \quad - \beta^m \lambda^m (v^m - R^m)] \\
&\quad + \dots \\
&\leq [\beta^k \Pi^k R^k + (1 - \beta^k) R^k] \\
&\quad + \beta^k (1 - \Pi^k) [\beta^l \Pi^l R^l + (1 - \beta^l) R^l] \\
&\quad + \beta^k (1 - \Pi^k) \beta^l (1 - \Pi^l) [\beta^m \Pi^m R^m + (1 - \beta^m) R^m] \\
&\quad + \dots \\
&= [1 - \beta^k (1 - \Pi^k)] R^k + \beta^k (1 - \Pi^k) [1 - \beta^l (1 - \Pi^l)] R^l \\
&\quad + \beta^k (1 - \Pi^k) \beta^l (1 - \Pi^l) [1 - \beta^m (1 - \Pi^m)] R^m
\end{aligned}$$

Defining

$$s^\alpha = \beta^\alpha (1 - \Pi^\alpha) \geq 0$$

we can write

$$\begin{aligned}
T &= [1 - s^k] R^k + s^k [1 - s^l] R^l + s^k s^l [1 - s^m] R^m + \dots \\
&= R^k + s^k \underbrace{[(R^l - R^k)]}_{\leq 0} + s^l \underbrace{[(R^m - R^l)]}_{\leq 0} + s^m [\dots] \\
&\leq R^k \\
&< R^j
\end{aligned}$$

Thus (21) is strictly positive and strategies of the form T cannot be optimal.

The only strategy that is not of the form S or T and that has not been proven to be suboptimal is the sampling strategy of theorem 1. It uses the optimal stopping rule, samples boxes in the order of decreasing reservation prices and always chooses the box with the highest reservation price. Since an optimal strategy exists (either due to the finiteness of expectations in the case of a finite number of search opportunities or due to the assumption of finite variance in the case of infinitely many search opportunities, see DeGroot [2] chap. 12 and 13), this establishes the optimality of the proposed rule. ■

References

- [1] Sushil Bikchandani and Sunil Sharma. Optimal search with learning. *Journal of Economic Dynamics and Control*, 20:333–359, 1996.
- [2] M. DeGroot. *Optimal Statistical Decisions*. McGraw-Hill, New York, 1970.
- [3] Chien fu Chou and Gabriel Talmain. Nonparametric search. *Journal of Economic Dynamics and Control*, 17:771–784, 1993.
- [4] J.C. Gittins. *Multi-armed Bandit Allocation Indices*. John Wiley and Sons, 1989.
- [5] J.C. Gittins and D.M. Jones. A dynamic allocation index for the sequential design of experiments. in *"Progress in Statistics"*, (J.Gani, Ed.), North-Holland, Amsterdam, 1972.
- [6] K.D. Glazebrook. A profitability index for alternative research projects. *OMEGA, International Journal of Management Science*, 4:79–83, 1976.
- [7] K.D. Glazebrook. Some ranking formulae for alternative research projects. *OMEGA, International Journal of Management Science*, 6:193–194, 1978.
- [8] K.D. Glazebrook. Stoppable families of alternative bandit processes. *Journal of Applied Probability*, 16:843–854, 1979.
- [9] Meir G. Kohn and Steven Shavell. The theory of search. *J. Econ. Theory*, 9:93–123, 1974.
- [10] Steven A. Lippman and John J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189 and 347–367, 1976.
- [11] C.J. McKenna. Theories of individual search behavior. in *"Surveys in the Economics of Uncertainty"*, (J.D.Hey and P.J. Lambert, Eds.), Basil Blackwell, New York, 1987.

- [12] Peter B. Morgan. Distributions of the duration and value of job search with learning. *Econometrica*, 53:1199–1232, 1985.
- [13] Jennifer F. Reinganum. Nash equilibrium search for the best alternative. *J. Econ. Theory*, 30:139–152, 1983.
- [14] Donald B. Rosenfield and Roy D. Shapiro. Optimal adaptive price search. *J. Econ. Theory*, 25:1–20, 1981.
- [15] Michael Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82:689–711, 1974.
- [16] S. C. Salop. Systematic job search and unemployment. *Review of Economic Studies*, 40:191–201, 1973.
- [17] Robert Slonim. Learning in a search-for-the-best-alternative experiment. *Journal of Economic Behaviour and Organization*, 25:141–165, 1994.
- [18] J.G. Stigler. The economics of information. *Journal of Political Economy*, 65:213–225, 1961.
- [19] Gabriel Talmain. Search from an unknown distribution: An explicit solution. *J. Econ. Theory*, 57:141–157, 1992.
- [20] Tara Vishwanath. Parallel search for the best alternative. *Economic Theory*, 2:1992, 1992.
- [21] Martin L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47:641–654, 1979.




Figure 1a




Figure 1b




Figure 2a

Figure 2b

