

Non-parametric Search.

Chien-fu Chou and Gabriel Talmain* †

State University of New York at Albany

December 10, 1990

Abstract

We analyze the search problem of a consumer who derives information only from the sequential search process. The paper considers the case of a consumer who uses a non-parametric procedure to estimate the probability distribution. It is shown that a solution to the consumer's problem is a very simple strategy which depends only on the order statistics, on the discounting factor, and on the duration of the search. It leads to a finite search almost surely. This optimal strategy is a myopic rule which is computable and which is characterized by a sequence of strictly increasing reservation prices.

Keywords: Search, Dynamic Optimization, Non-Parametric Estimation.

JEL Classification Number: 022, 026, 213

^{0*}Please address all correspondence to Gabriel Talmain, Department of Economics, State University of New York at Albany, Albany, NY 12222.

^{0†}We would like to thank Carlos Rodriguez, Michael Sattinger, the participants to the 1991 conference of the Society for Economic Dynamics and Control in Capri, Italy, and two anonymous referees for numerous comments and suggestions. All remaining errors are the authors' sole responsibility.

1. Introduction

The purpose of this paper is to study the search strategy of a consumer who has no idea at all of what the distribution of prices is.

The elementary search model, in which the distribution of prices is common knowledge, has been extensively studied [see Lippman and McCall (1981) for a survey]. As Lippman and McCall remarked, “Perhaps the most restrictive and least palatable assumption of the elementary search model is the supposition that the offer distribution F is known.” The problem of search from an unknown distribution has been studied by Rothschild (1974), Rosenfield and Shapiro (1981) and Talmain (1991). Both Rothschild and Talmain considered the case when prices can only take a finite number of values. Rosenfield and Shapiro analyzed the case when the unknown distribution is continuous. They considered the cases when the distribution is exponential or normal. All these papers take a *parametric Bayesian* approach: even though the true distribution is unknown, the consumer knows that the unknown distribution belongs to some family of distributions that can be described by a finite number of parameters, such as a normal distribution with unknown mean. If the consumer really knows that the true distribution belongs to the family mentioned above, then this is a powerful approach to the problem of search. On the other hand, it is conceivable that the consumer cannot say *a priori* that the unknown distribution belongs to any particular parametric family.

One option is for the consumer to apply stopping rules (based on a parametric Bayesian approach) available from the existing literature, even though the underlying assumptions that justify these stopping rules are not satisfied. However the result can be disastrous as illustrated by the following example. Suppose that the consumer uses the stopping rule from DeGroot (1970, p. 345), which assumes a normal distribution with unknown mean and fixed variance, and he assumes a prior mean μ_0 , when in

fact (unknown to him) the true distribution is degenerated at μ_0 , then he will search forever. Intuitively, the searcher always observes μ_0 , which convinces him that the true mean is indeed at μ_0 . Since he erroneously believes the distribution is normal, this means that he believes he has 50% chances of improving on μ_0 , making him willing to continue. Notice that if he knew the true distribution, the consumer would stop at once.

The problem with parametric Bayesian estimation in this example is that it excludes *a priori* some distributions (including the true one) when in fact the consumer has no information to support this exclusion. The potential for this problem is always present when using a parametric Bayesian approach without having prior information to justify the restriction to a certain parametric family. Another option is to recognize the non-parametric nature of the problem: the consumer does not put *a priori* restrictions on his estimator of the true distribution, such as restricting it to be normal. This is the approach adopted in this paper. We consider the case of a first time buyer who only knows, at the beginning of search, that prices can take non-negative values. As the search proceeds, the consumer accumulates observations and uses this information to update his estimation of the true distribution. At each stage of the search process, the consumer uses the maximum entropy principle to estimate this distribution. We show that the search problem admits an optimal solution, which is a strategy that terminates in finite time. We derive an explicit solution to the search problem. The solution is a myopic rule characterized by a sequence of increasing reservation prices. In the context of job search, our model explains why workers have decreasing reservation wages and why they may go back to an offer they initially refused. These results, which are consistent with empirical evidences, could not be explained by the elementary search model.

In the next section, the model is presented. In section 3, we construct the optimal strategy. Section 4 compares our estimation procedure to the standard parametric

Bayesian updating methods. Section 5 concludes the paper.

2. The Model

A consumer wants to buy one unit of a commodity. The consumer is willing to pay at most one dollar for the good, in which case his consumer surplus is zero.¹

He searches sequentially from an infinite number of sellers. Each search takes one period. Let p_t be the price quoted in the store he visits in period t . The consumer believes that quotes are independently drawn from the same probability distribution which has support $[0, 1]$. We assume that prices are recallable.

Let $q_t \equiv 1 - p_t$ be the consumer surplus and $\hat{q}_t \equiv \max_{1 \leq i \leq t} q_i$. The consumer searches for the highest surplus and discounts time. If the consumer stops at period t , he will achieve a utility of $U(\hat{q}_t, t) \equiv \delta^t u(\hat{q}_t)$, where $0 < \delta < 1$, u is concave and differentiable, and $u(0) \geq 0$.² If the consumer never stops, then his utility is 0. Alternatively, we will also consider the easier case of fixed search cost: $U(\hat{q}_t, t) \equiv u(\hat{q}_t) - ct$, where c is the cost of search.

We consider the problem of a first time buyer. The only information this agent has at the inception of search is that the surplus is between zero and one.³ The consumer has to estimate the probability distribution of surplus, updating this estimate with the information provided by the search process itself. This problem is made easier by the fact that in period t , the consumer only cares about surpluses q_{t+1} : $1 \geq q_{t+1} > \hat{q}_t$. Using an estimation procedure that is described at the end of this section, given \hat{q}_t

¹Any price higher than 1 is treated by the consumer as if it were 1 (no purchase). So without loss of generality, we can assume that price vary between 0 and 1.

²This is the family of all monotonic, continuous and stationary time preferences. The utility is cardinal which means that a nonlinear transformation of the utility function will change the consumer's decision problem.

³In Chou and Talmain (1988), we also consider the case of an *experienced* consumer who has some information about the price distribution before the beginning of the current search. That is, in the past, the consumer observed a finite number of prices from the same distribution, but these prices are not available anymore.

the highest surplus observed up to time t , the consumer has the following estimator π for the conditional distribution of \hat{q}_{t+1} : a probability mass of $t/(t+1)$ at \hat{q}_t , and a constant density of $1/(t+1)(1-\hat{q}_t)$ between \hat{q}_t and 1.

At time t , the searcher has obtained a sequence of surplus. He can either decide to continue searching, or stop and buy at the highest surplus. A strategy is a rule which tells the consumer in every period t and for every possible sequence (q_1, \dots, q_t) whether to stop or to continue. Let σ be the strategy of the consumer, and let us suppose that the consumer has observed (q_1, \dots, q_t) . Let us assume that the strategy σ directs the consumer to continue for all $1 \leq \tau \leq t-1$, and stop at time t . Then $t_\sigma \equiv t$ is the stopping time and $\hat{q}_\sigma \equiv \hat{q}_t$ is the stopping surplus. The consumer evaluates the strategy σ by the expected utility

$$W_\sigma \equiv E(U(\hat{q}_\sigma, t_\sigma)), \quad (1)$$

where the expectation is defined using the above probability distribution. The consumer will choose the strategy which maximizes W_σ , if it exists.

Estimation procedure: The spirit of our procedure is best described by an example. Prior to any observation, a consumer estimates the unknown true distribution by the uniform distribution. Updating of this estimate proceeds as follows. Suppose that the consumer has observed, not necessarily in that order, surpluses: $(q_{(1)}, q_{(2)}) = (0.25, 0.30)$. These two observations divide the interval of possible surpluses $[0, 1]$ into three sub-intervals: $[0, 0.25]$, $(0.25, 0.30]$, and $(0.30, 1]$. Our agent will (i) assign equal probability mass (one third) to each sub-interval, and (ii) distribute this probability mass uniformly within each sub-interval.

The rationale for this procedure, in the case in which the true distribution of surplus is continuous, is as follows:

(i) Before any sampling, the unconditional probability of the event $A \equiv \{q_{(1)} \leq q_3 \leq q_{(2)}\}$

is one third. Without any knowledge of the true distribution, the conditional probability of A , once q_1 and q_2 have been observed, is unknown. We propose to estimate this conditional probability by the previous unconditional probability. This is equivalent to using the quantiles of the empirical distribution to estimate the quantile of the true distribution: with our procedure, the 1/3 and 2/3 quantiles of estimated distribution coincide with those of the empirical distribution (0.25, 0.30). For this reason, we label this procedure the *quantile preserving* estimation.

(ii) The uniform distribution of probability mass within each interval is justified by application of the maximum entropy principle, which will be described in the appendix.

The procedure above can be applied even in the case in which the consumer does not know whether the true distribution is continuous, or discrete, or even a mixture of both. Suppose that next surplus observed by the consumer is $q_3 = 0.30$. With three observations, our procedure is to pour one fourth of probability mass between each observed surplus, except that the interval $[q_{(2)}, q_{(3)}]$ is degenerate, i.e. it is the single point $q = 0.30$. The resulting estimator of the probability distribution becomes a *mixture*: uniform distribution with one fourth probability mass in the intervals: $[0, 0.25]$, $[0.25, 0.30]$, and $(0.30, 1]$, and one fourth probability mass on 0.30.⁴ The rationale for this assignment of probability mass is still the quantile preserving argument.

One justification for our procedure is that our estimator of the distribution of probability is *consistent*: it converges uniformly to the true distribution (with probability one) by the Glivenko-Cantelli version of the uniform strong law of large numbers [see Pollard (1984, p. 13)] and it is asymptotically *efficient*.⁵

⁴In the event the true distribution may have probability mass concentrated at some points, the consumer must decide whether the true distribution is in fact continuous or not. An *atom* is a point with a probability mass. If a distribution has atom(s) [resp. is atomless] the probability of a repeat observation is strictly greater than [resp. equal to] zero. So the likelihood ratio test (at the 100% level) indicates that, if a repeat observation occurs, the true distribution is not atomless.

⁵The estimator of the distribution at every point reaches asymptotically the Cramer-Rao bound. This is because the estimator converges at every point to the Bernoulli sample mean, which is an efficient estimator, see DeGroot (1975, p. 366, exercise 7).

3. Derivation of the optimal strategy

In this section we want to construct the optimal strategy. The decision either to continue or to stop in period t depends only on \hat{q}_t , because the estimation of future \hat{q}_{t+k} , $k > 0$, depends only on \hat{q}_t . Let $V_t(q)$ be the expected value of search discounted back to period t conditional on $\hat{q}_t = q$:⁶

$$V_t(q) \equiv \max_{\sigma \in \Sigma_t} E[\delta^{t\sigma-t} u(q_\sigma) \mid \hat{q}_t = q], \quad (2)$$

where Σ_t is the set of all strategies which tell the consumer to continue for $\tau < t$. If in period t the consumer receives an offer of a zero surplus, $q_1 = 0$, then obviously he should continue. On the other hand, if $q_1 > 0$, then the following proposition shows that he will complete search in a finite time.

Proposition 1 *If $q_1 > 0$, then there exists an upper bound $T(q_1) < \infty$ to the tenure of search.*

Proof: Let $\pi_t(x) = Pr[\hat{q}_{t+1} \leq x \mid \hat{q}_t = q]$. Since $V_{t+1} \leq u(1)$ and $V_t(q) \geq V_{t+1}(q)$,⁷ we have

$$\begin{aligned} E[V_{t+1}(\hat{q}_{t+1}) \mid \hat{q}_t = q] &= \int_0^1 V_{t+1}(x) d\pi_t(x) \\ &= \pi_t(q) V_{t+1}(q) + \int_q^1 V_{t+1}(x) d\pi_t(x) \\ &\leq \pi_t(q) V_t(q) + [1 - \pi_t(q)] u(1). \end{aligned}$$

$$\text{Therefore } V_t(q) - \delta E[V_{t+1}(\hat{q}_{t+1}) \mid \hat{q}_t = q] \geq [1 - \delta \pi_t(q)] u(q) + [1 - \pi_t(q)] \delta u(1).$$

Since $\lim_{t \rightarrow \infty} \pi_t(q) = 1$,

$$\lim_{t \rightarrow \infty} V_t(q) - \delta E[V_{t+1}(\hat{q}_{t+1}) \mid \hat{q}_t = q] \geq (1 - \delta) u(q) \geq (1 - \delta) u(q_1) > 0.$$

Hence, $\exists T(q_1), \forall t \geq T(q_1), [V_t(q) - \delta E[V_{t+1}(\hat{q}_{t+1}) \mid \hat{q}_t = q] > (1 - \delta) u(q)/2 > 0$. Since $V_t(q) = \max\{u(q), \delta E[V_{t+1}(\hat{q}_{t+1}) \mid \hat{q}_t = q]\}$, $V_t(q) = u(q)$ for all $t > T(q_1)$. *Q. E. D.*

⁶For a proof of the existence of $V_t(q)$, see Chou and Talmain (1988).

⁷ $V_t(q) \geq V_{t+1}(q)$ because $\pi(\hat{q}_{t+2} \mid \hat{q}_{t+1} = q)$ is first order stochastically dominated by $\pi(\hat{q}_{t+1} \mid \hat{q}_t = q)$, see Chou and Talmain (1988).

As long as the consumer observes only zero surpluses ($\hat{q}_t = 0$), it is obvious that search should continue. Once a strictly positive surplus is observed, then the previous proposition applies and the search horizon becomes finite, and we can use backward induction to construct the optimal search strategy.

Proposition 2 *The optimal search strategy is characterized by a sequence of decreasing reservation surpluses $\xi_1 > \xi_2 > \dots > \xi_t > \dots > 0$, where ξ_t is the solution to the equation*

$$u(\xi_t) = \delta \left[u(\xi_t) \frac{t}{t+1} + \int_{\xi_t}^1 \frac{u(x)}{(1-\xi_t)(t+1)} dx \right]. \quad (3)$$

Proof: Given $\hat{q}_t = q$, let $f_t(q)$ be the undiscounted expected utility if the consumer searches one more time, and let $g_t(q)$ be the expected loss from one more search:

$$f_t(q) \equiv u(q) \frac{t}{t+1} + \int_q^1 \frac{u(x)}{(1-q)(t+1)} dx \quad \text{and} \quad g_t(q) \equiv u(q) - \delta f_t(q).$$

First, since $f_t(q)$ is a decreasing function of t , $g_t(q)$ is increasing in t . Next, we show that $g_t(q)$ is also an increasing function of q . First:

$$(t+1)f'_t(q) = tu'(q) + \frac{1}{1-q} \left[\int_q^1 \frac{u(x)}{1-q} dx - u(q) \right].$$

By the intermediate value theorem, there exists a $c_q \in (q, 1)$ such that $\int_q^1 u(x) dx = (1-q)u(c_q)$. And again by the same theorem there exists a $z_q \in [q, c_q]$ such that $u(c_q) - u(q) = u'(z_q)(c_q - q)$. Notice that $u'(z_q) \leq u'(q)$ by concavity of u . And since $c_q < 1$, we have

$$(t+1)f'_t(q) = tu'(q) + \frac{u(c_q) - u(q)}{1-q} = tu'(q) + u'(z_q) \frac{c_q - q}{1-q} < (t+1)u'(q).$$

For all t , $g_t(0) < 0$ and $g_t(1) > 0$. Therefore, for each t , the equation $g_t(q) = 0$ has a unique solution $\xi_t \in (0, 1)$. Also because the function g_t is increasing in both t and q , the sequence of solutions is decreasing: $1 > \xi_1 > \xi_2 > \dots > 0$.

Suppose that the consumer first observes $q_1 > 0$. From proposition 1, we know that there is an upper bound $T = T(q_1)$ to the tenure of search. At T the search stops. At

$T - 1$, the search continues if and only if $g_{T-1}(\hat{q}_{T-1}) < 0$, that is, if $\hat{q}_{T-1} < \xi_{T-1}$. Let us suppose that for $\tau = T - 1, \dots, t + 1$, the optimal decision rule is characterized by the sequence of decreasing reservation surpluses $\xi_{T-1} < \xi_{T-2} < \dots < \xi_{t+1}$. We want to show that the optimal decision rule in period t is characterized by the reservation surplus ξ_t . If $\hat{q}_t \geq \xi_t > \xi_{t+1}$, then the search will stop for sure at $t + 1$. Therefore, the consumer should stop if $g_t(\hat{q}_t) \geq 0$. So if $\hat{q}_t \geq \xi_t$ the consumer should stop at t . If on the other hand $\hat{q}_t < \xi_t$, then the expected value of one more search is higher than the utility of the present surplus. So, indeed ξ_t is the reservation surplus at t . Notice that $\xi_{T-1}, \xi_{T-2}, \dots, \xi_t, \dots, \xi_1$ are independent of q_1 . Therefore the complete sequence $\{\xi_t\}_{t=1}^\infty$ indeed characterizes the optimal strategy. Q. E. D.

The intuition for the nature of the optimal search strategy is as follows. Consider a *myopic* searcher who thinks only one period ahead, i.e. who stops if, and only if, the expected utility of searching exactly one more time is less than or equal to the current utility. If in period t such a searcher has received \hat{q}_t as his best quote, this myopic searcher will stop if, and only if, $\hat{q}_t \geq \xi_t$, where ξ_t is as defined in (3). The expected gain from one more search is a decreasing function of \hat{q}_t , the higher the initial surplus the less scope for improvement, and equation (3) defines the values for the break-even \hat{q}_t . Also, the cut-off values ξ_t are decreasing with t : the longer the tenure of search the less hope for improvement. What is less obvious is that the *optimal* strategy should coincide with such a *myopic* rule. Consider side-by-side a myopic searcher and an optimal searcher, i.e. both observe the same quotes at the same time. If in period t the myopic searcher wants to continue, then it is obvious that the optimal searcher also wants to continue. More difficult is to see why the optimal searcher would want to stop when the myopic searcher stops. For simplicity, let us consider two searchers with current highest quote $\hat{q}_t \geq \xi_t$: searcher A can search at most one more time; searcher B can search at most two more times. Searcher A , like a myopic searcher, will not continue. Would searcher B want to continue? Suppose B draws one more

time. Then, in period $t + 1$, he is like a myopic searcher. Also his highest quote \hat{q}_{t+1} is such that $\hat{q}_{t+1} \geq \hat{q}_t \geq \xi_t > \xi_{t+1}$, so no matter what he observes in period $t + 1$, he stops. Therefore the option of searcher B to search twice instead of once is worthless in this case. This explains why searcher B stops when A stops. Assuming that an optimal strategy exists that stops in finite time, one can extend this argument by backward induction to show that the optimal searcher always stops when the myopic searcher stops.

It is clear that a searcher who faces no *a priori* restrictions on the tenure of his search can achieve a higher value than a searcher who is limited to search at most once more. The previous argument shows that the gain is not due to the possibility for the optimal searcher to search when the restricted searcher wants to stop. Instead, it is due to the possibility for the optimal searcher to search more when the restricted searcher is forced to stop. With our procedure, “news” are interpreted in the same way by the myopic searcher and by the optimal searcher. This is not true in general. Consider the following alternative search framework: there are *a priori* only two possible distributions, distribution F_a is concentrated at $q = 0.5$, distribution F_b is uniform over the unit interval $[0, 1]$. At $t = 1$ we have observed $q_1 = 0.5$. Now consider searchers A and B as before and suppose they both decide to continue. If they next observe $q_2 = 0.1$, searcher A will feel this is bad news, whereas searcher B may well feel this is good news: B has now 50% chances of receiving $\hat{q}_3 > 0.5$ in the next period, whereas A is stuck at $\hat{q}_2 = 0.5$. Notice that in this case, upon observing $q_1 = 0.5$ the optimal response for A may well be to stop, whereas the optimal strategy for B is to continue. This example is useful in understanding why Bayesian parametric search with normal prior does not exhibit in general the decreasing reservation wage property, a point which is elaborated upon in the next section.

Remark: More patient consumers will search longer. This is the result of the comparative statics of the reservation surplus with respect to δ as can be seen from (3).

Furthermore $\forall t > 0, \lim_{\delta \rightarrow 1} \xi_t(\delta) = 1$. Given a sequence of surpluses quotes (q_1, q_2, \dots) such that $\sup_{t \in N} q_t = 1$, and $\forall t > 0, q_t < 1$, increasingly patient consumers facing such a sequence of surpluses will conduct a search of increasing duration, and in the limit, the consumer who does not discount time will search forever.

Example: Let's consider a consumer with a linear utility function: $u(x) = ax$. Then the optimal strategy at the stage t of the search is:

1. accept the best offer so far if its value \hat{q}_t is higher than or equal to the reservation surplus ξ_t , where $\xi_t = \delta / ((1 - \delta)(2t + 1) + 1)$,
2. continue searching otherwise.

For instance, suppose $\delta = 0.75$, then $(\xi_1, \xi_2, \xi_3) = (0.429, 0.333, 0.273)$. With the numerical example of section 2, $(q_1, q_2, q_3) = (0.25, 0.30, 0.30)$, the consumer will continue after the first and second observations, and stop at the third one.

It is clear that (as was claimed) $\xi_{t+1} < \xi_t$, that $\xi_t \rightarrow 0$ when $t \rightarrow \infty$, so that eventually the consumer gives up all of his consumer's surplus.

The case of fixed cost: We now consider the case when the utility is given by $U(q, t) = u(q) - ct$. Define $\bar{u} \equiv \int_0^1 u(x) dx$ and $T \equiv [(\bar{u}/c)] - 1$, where $[x]$ represents the integer part of x . Search ends in finite time because no searcher will go past T : even if $\hat{q}_T = 0$, the expected gain from one more search is $\bar{u}/(T + 1)$ which is less than the cost c .

By the same argument as in the proof of proposition 2, the optimal strategy is a sequence of decreasing reservation surpluses $\{\xi_t\}_{t=1, \dots, T}$ that are the solutions to the equations:

$$h_t(\xi_t) \equiv u(q) - (f_t(q) - c) = 0.$$

4. Comparison with parametric Bayesian updating

We are going to compare the Bayesian estimation method to the procedure adopted here. In the economic literature, a Bayesian statistician who wants to estimate a probability distribution first commits to a family of distributions that can be indexed by a parameter $\theta \in \Theta \subset \mathcal{R}^N$. An example is the normal family with unknown mean and fixed variance, in which case $\Theta = \{\mu \mid -\infty < \mu < \infty\}$ and $N = 1$, or the normal family with both mean and variance unknown, in which case $\Theta = \{\mu, \sigma \mid -\infty < \mu < \infty, 0 < \sigma\}$ and $N = 2$. He then must assign a probability distribution over the set of parameters (the Bayesian priors), and revises this assignment of probability with each observation using Bayes rule.

Our searcher does not want to commit to any particular family of distributions. Instead our searcher will do non-parametric estimation, and will commit to the maximum entropy principle, which directs him to select the uniform distribution in the absence of any information.⁸ The searcher uses the quantile preserving step to interpret the additional information provided by a new observation in terms of restriction of the distribution, so that this information can be processed by the maximum entropy principle.

This difference in estimation procedure has implications for the properties of the optimal search strategy. Because of the assumption of search with recall, at each stage t of the search the consumer only cares about the distribution of probability over $[\hat{q}_t, 1]$. The Bayesian search literature tends to concentrate on unimodal symmetric

⁸Alternatively, the consumer could do *non-parametric* Bayesian estimation. Notice that the dimension N of the set Θ is infinite; defining a Bayesian prior (a probability measure) on such a complicated set raises technical difficulties. One such prior is the Dirichlet process which is an extension of the Dirichlet priors to the case with infinite number of parameters, see Ferguson (1973). However the infinite dimensional nature raises some technical problems. For instance Diaconis and Freedman (1986) show that, in the context of estimating a location parameter, there exists a prior which makes the Bayes estimator inconsistent (whereas a Bayes estimator is always consistent in the finite parametric context).

distributions, especially the normal distributions. In this case the reservation surplus property may not hold, and even in the case it holds, it may not be characterized by decreasing reservation surpluses. Suppose such a Bayesian searcher has received several quotes close to $q = 0.5$. An additional quote of say $q = 0.55$ will be interpreted as an increase in the precision of the subjective mean and is likely to induce the searcher to stop. On the other hand, an additional quote of say $q = 0.1$ will be interpreted as a substantial increase in the dispersion of the subjective distribution, and is likely to encourage the searcher to continue. This observation $q = 0.1$ provides *global* information about the whole distribution of probability: it changes the *shape* of the distribution at say $q = 0.5$ as well as at $q = 0.1$. Note that with our procedure the information is *local*, concept that is best illustrated on an example. Suppose we observe $q_t = 0.1$ and that the closest previous observations on each side are 0.05 and 0.15. The ratio of the densities at any two points outside the interval $[0.05, 0.15]$ is the same before and after observing $q_t = 0.1$.

Suppose that the consumer observes a low q_{t+1} . If the Bayesian searcher knows the mean of the distribution, he will interpret this observation as an increased variance of his estimation and consequently will become more optimistic about improving over \hat{q}_t . If he knows the variance, the information is interpreted as a downward shift of the estimated distribution, which will make him more pessimistic about improving over \hat{q}_t . With our procedure, an observation $q_{t+1} < \hat{q}_t$ always makes the searcher more pessimistic about improving over \hat{q}_t . Furthermore, it does not matter whether q_{t+1} is very low or barely lower than \hat{q}_t . To our searcher all probability distributions are possible. For example the true distribution could have a peak beyond \hat{q}_t , observing q_{t+1} much lower than \hat{q}_t does not convey any information on where this peak could be. This property, that observations only provide local information on the conditional distribution is responsible for the decreasing reservation surplus property of our procedure.⁹

⁹This is a consequence of the subset independence axiom (see appendix).

In the absence of a well-established non-parametric Bayesian procedure, we have developed a procedure that is well-adapted to the context of search. From the point of view of estimation, our procedure is natural, it is consistent and efficient. On the other hand, even if an estimation procedure is consistent and efficient, it may not be suitable for the purpose of modelling search. Consider a searcher who uses the empirical distribution to estimate the true unknown distribution. As an estimator it is consistent and efficient. However with such an estimator the consumer will stop at the first observation: if he observes $q_1 = 0.3$, then his estimator is degenerate at $q = 0.3$, so there is no point to search any further. Notice that this is also the minimax search strategy.

5. Conclusion

This paper shows that the search problem can be solved without assuming that the consumer knows anything about the price distribution. In fact, the optimal solution is much simpler here than when one assumes that the consumer believes the true distribution belongs to, say, the family of normal distributions.

Our result is consistent with empirical results from the labor market that cannot be explained by the simple search model. Kasper (1967) reported evidence that the reservation wage of workers falls over time. Blau (1988) finds that, for long-duration stints of unemployment, the rate of mortality is far greater than the rate of new jobs offered. This may mean that some searchers go back to offers they previously rejected.

This raises the question of what would happen in a non recallable search set up. The major difference with the present paper is that the searcher would be interested in the whole distribution of surpluses, not just in the possibility of improving over his present best price. This means that the optimal strategy would depend on all the past observations. Preliminary investigation, in the context of a three period non

recallable fixed cost search model, indicates that the optimal strategy is still a sequence of decreasing reservation surpluses. The reservation surplus for period 2 depends on the observation in period 1. The first period reservation surplus is always lower than the corresponding one in the recallable case: this is because the value of search is always higher in the recallable case. However it is possible for the reservation surplus in the second period to be higher than in the recallable case. A searcher under recall who settles, the second time around, for a 'low' surplus has necessarily observed low surpluses in the first and second periods. On the other hand a searcher under no recall who observes the second time around a 'low' surplus may have observed in the first period a surplus barely lower than his reservation. This can make him more 'choosy' about accepting second period 'low' surplus.

References

- Blau, D., 1988, Search for Nonwage Job Characteristics: A Test of the Reservation Wage Hypothesis, Mimeo, University of North Carolina.
- Chou, C. and G. Talmain, 1988, Search without Priors and the Quantile Preserving Maximum Entropy Estimation, SUNY Albany # 239.
- DeGroot, M.H., 1970, Optimal Statistical Decision (McGraw-Hill, New York).
- DeGroot, M.H., 1975, Probability and Statistics (Addison-Wesley, Reading, Mass.).
- Diaconis, P. and D. Freedman, 1986, On the consistency of Bayes estimates, The Annals of Statistics 14, 1-26.
- Ferguson, T.S., 1973, A Bayesian analysis of some nonparametric problems, The Annals of Statistics 1, 209-230.
- Gnedenko, B., 1978, The Theory of Probability, 4th ed. (MIR Publishers, Moscow).
- Jaynes, E.T., 1983, Where Do We Stand on Maximum Entropy, in: Rosenkrantz, R.D., ed., Papers on Probability, Statistics and Statistical Physics (Reidel, Dordrecht, Holland) 210-314.
- Kasper, H., 1967, The asking price of labor and the duration of unemployment, Review of Economics and Statistics 49, 165-172.
- Lippman, S.A. and J.J. McCall, 1981, The Economics of Uncertainty, in: K.J. Arrow and M.D. Intriligator, eds., Handbook of mathematical economics Vol. I (North Holland, New York) 211-227.
- McAfee, R.P. and J. McMillan, 1988, Search Mechanisms, Journal of Economic Theory 44, 99-123.
- Mortensen, D.T., 1984, Job Search and Labor Market Analysis, in: O. Ashenfelter and R. Layard, eds, Handbook of Labor Economics vol. II (North Holland, New York) 849-919.
- Rothschild, M., 1974, Searching for the lowest price when the distribution of prices is unknown, Journal of Political Economy 8, 689-711.
- Pollard, D., 1984, Convergence of Stochastic Processes, (Springer-Verlag, New York).

- Rao, C.R., 1973, *Linear Statistical Inference and its Applications* 2nd ed. (John Wiley and Sons, New York).
- Rosenfield, D.B. and R.D. Shapiro, 1981, Optimal Adaptive Price Search, *Journal of Economic Theory* 25, 1-20.
- Shore, J.E. and R.W. Johnson, 1980, Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy, *IEEE Transactions on Information Theory* 26, 26-37.
- Skilling, J., 1986, The axioms of maximum entropy, mimeo Cambridge, England.
- Talmain, G., 1991, Search from an unknown distribution: an explicit solution, *Journal of Economic Theory*, forthcoming.
- Theil, H., 1967, *Economics and Information Theory* (North-Holland, New York).
- Theil, H. and Fiebig, D.G., 1984, *Exploiting Continuity* (Balinger, Cambridge, Mass.).

Appendix: Description of the maximum entropy principle

First we will introduce the intuition behind the concept of entropy using the information theory presentation, then we will present the axiomatic approach to justify the maximum entropy principle.

Suppose that statistician A has wrecked his car. Let us suppose that A knows that the repair will cost \$2,000 with probability $\pi_1 = 99\%$ and \$3,000 with probability $\pi_2 = 1\%$. If A receives a bill for \$2,000 he is not surprised. So the information contained in the message: “the bill is \$2,000” is small. On the other hand a bill for \$3,000 would be a big surprise: it conveys a lot of information to A . More generally speaking the information received by observing the occurrence of a particular outcome decreases with the probability of this outcome. We will define the *information content* of the message “the bill is \$2,000 [resp. \$3,000]” as $-\log(\pi_1) = 0.01$ [resp. $-\log(\pi_2) = 4.60$].

Next consider statistician B who has also wrecked his car, and who knows that the repair will cost: \$2,000 [resp. \$3,000] with probability 0.50. We can say that B is more ‘uncertain’ about the cost of repair than A . We define the entropy of the distribution facing A as the expected value of the information content: $H_A \equiv -[\pi_1 \log(\pi_1) + \pi_2 \log(\pi_2)] = 0.056$, whereas $H_B = 0.693$; entropy is a measure of the ‘uncertainty’ embodied in a distribution of probability.

Suppose now that statistician C has wrecked his car and wonders how to estimate the probability distribution of the cost of repair. In the context of the maximum entropy principle, the information that C may have is treated as functional restrictions on the probability distribution: for instance C may know the support of the distribution, its mean, its variance and so on. All the knowledge of the statistician must be put in the form of these *informational* constraints. The maximum entropy principle directs C to choose the distribution of probability with the highest entropy subject to the informational constraints. For instance the maximum entropy distribution on the

real line subject to known mean μ and known variance σ^2 is the normal distribution $\mathcal{N}(\mu, \sigma)$ [see Rao (1973, p. 162)].

The maximum entropy principle can be justified from an axiomatic basis. Shore and Johnson (1980) proved that any method for estimating a distribution function that satisfied the following four axioms was equivalent to the maximum entropy principle.

- 1- *Uniqueness*: the method should provide a unique solution,
- 2- *Invariance*: the choice of coordinate system should not matter,
- 3- *System Independence*: if the method is applied to estimate the joint distribution of two random variables, and if there is no information *a priori* to the contrary, then the axiom postulates that the two variables are independent,
- 4- *Subset Independence*: Consider an unknown distribution F and two disjoint subsets A_1 and A_2 of the support of F , and assume that there is no constraint directly linking A_1 and A_2 . The axiom postulates that the estimator of the conditional distribution $F(\cdot | A_1)$ should not be affected by restrictions put on the conditional distribution $F(\cdot | A_2)$ (i.e. in case the estimation problem decomposes naturally into estimating F onto disjoint intervals, the estimation of $F(\cdot | A_1)$ only depends on the information pertaining to A_1 , not on the information pertaining to A_2).

Remark: As a referee judiciously pointed out, there is a question on whether the ignorance of the searcher would lead him to a uniform distribution as the estimator of \hat{q}_{t+1} or as the estimator of $u(\hat{q}_{t+1})$. After all, if the consumer is totally ignorant, there is no reason why he should be less ignorant of the distribution of his utility level than that of his surplus. (Notice that a non-uniform distribution contains more information than a uniform distribution from the point of view of entropy.) This is a relevant question because if $u(q)$ is not linear, then choosing one estimator or the other would lead to different optimal search strategies. The apparent symmetry between utility and surplus does not hold in our framework because we have explicitly assumed that searchers observe surpluses, not a utility level. This apparently innocuous assumption

can be exploited with the help of the next axiom to justify why the consumer must estimate the distribution over surpluses rather than over utility levels.

Consider two agents who observe the drawing of balls marked with real numbers from a statistical urn. The two agents must estimate the probability distribution of the numbers.

Axiom (A): If at time t the agents have the same estimator of the probability distribution, and at $t + 1$ they observe the same drawing, then they will have the same estimator of the probability distribution at $t + 1$.

It is easy to see that if consumers apply the maximum entropy principle to the utility levels instead of to the surpluses, then two consumers who have the same estimator of the distribution of q at time t and observe the same q_{t+1} will not agree on the distribution of q at time $t+1$ if their utility functions are not linear transformations of each other. The issue here is closely related to Bertrand's Paradox [see Gnedenko (1978, pp. 34-36)].