

UNIVERSIDADE NOVA DE LISBOA

Faculdade de Economia

ISSN 0871-2573

SEARCH WITH LEARNING

Gabriel Talmain

Working Paper Nº 122

Section 1: Introduction.

The economic literature has demonstrated a long standing interest in the question of search. In the elementary search model, the distribution of prices is common knowledge (see S. Lippman and J. McCall (1981)). The optimal rule is then myopic (i. e. the consumer stops if and only if he cannot improve his expected utility by searching exactly one more time) and consists in a reservation price (i. e. the consumer only buys if he gets a better price than his reservation benchmark), also the rule for a consumer who forfeits the option to buy at the current price if he continues searching (no recall) is the same as the rule with recallable prices. As Diamond (1971) remarked in the context of a job search model, such an assumption may lead to a degenerate distribution in which the firms set the discriminating monopsony wage, as the unique wage rate for the economy¹.

In view of this limitation the consumer may well decide that his decision rule should be independent of it. Criteria for optimality of behavior can be drawn over the set of *outcomes*, for instance the minimax rule (see Savage (1954)). However most of the research has assumed that when the consumer who does not know the distribution of prices will try to gather information about it from the search process itself. Rothschild (1974) considered the problem of a sequential Bayesian search with no recall, and established the weaker reservation-price property (i. e. a strategy which defines for each round a

¹Such a result can be interpreted in the light of the recent bargaining literature: assume as in Rubinstein and Wolinsky (1985) that buyers and sellers meet at random. Also assume, unlike the model presented there, that buyers always quote the price, and sellers cannot make counter-offers (Sellers act as Stackelberg leaders in the bargaining game). Then it is easy to see that the only perfect equilibrium for this economy is for the sellers to extract all consumer surplus from the buyers. This result is not entirely satisfactory because it is true even when the time discount factor of the buyer is very small (he is very patient), at the same time as the time discount factor of the seller is very large (he is very impatient). Therefore the buyer disposes of a very serious (for his opponent) low-cost (to himself) threat, but the requirement for game perfectness prevents him from using it.

benchmark price, the consumer buys if he gets a price better than the benchmark, and continues otherwise) for consumers who hold Dirichlet beliefs. Rosenfield and Shapiro (1981) extended this approach to a larger set of priors, establishing conditions for the optimal strategy to be myopic and exhibit the reservation price property. Putting together both sides of the market, Jovanovic (1979) showed with a job-matching model how price dispersion could hold when both sides of the market were taken into account. Further interest in equilibrium price dispersion was attested by the articles of Burdett and Judd (1983) where price dispersion resulted from randomness in the sample size, and Albrecht and Axell (1984) which used heterogeneous agents to obtain the result. Inquiries about the optimality of the search procedure were conducted by Morgan and Manning (1985) who looked into searches with non-constant sample size, Wolinsky (1987) who studied an economy in which agents bargained while conducting search on the side, and McAfee and McMillan (1988) who considered buying mechanisms mixing search and auction procedures for agents facing small population of traders.

The main emphasis of the existing literature in search has been to prove the existence, characterize and study the properties the optimal strategies, but it is notably short on explicit solutions to the search problem when the distribution of prices is unknown. The purpose of this paper is to find such an explicit optimal strategy when the consumer already had some experience in the market, or even when he comes as a first time buyer and knows nothing about the distribution of prices.

We will only consider searches with recall, that is a consumer buy at any price he observed in the past. We consider two types of search: one in which the consumer knows the set of prices he is facing is finite (the discrete case), and other in which this set is a continuum of prices (the continuous case). For the discrete case, we solve the problem of a Bayesian searcher who holds Dirichlet priors, though we will show that the solution we present can be extended to a more general class of beliefs. Dirichlet priors are interesting on at least two grounds: First because as Rothschild (1974) wrote ‘... as experience

accumulates, all searchers come to behave as if their priors were Dirichlet' (page 697). And second, because we will show that these priors enable us to justify the attitude: 'I will assign a probability to an event according to the frequency I observed it in the past' as a Bayesian approach.

For the continuous case, the Bayesian approach runs into technical difficulties, unless one restricts the admissible set of price distributions in order to have only a finite number of parameters to estimate. A typical example is to assume that it belongs to the class of normally distributed density of probability functions, some parameters of which are unknown. We propose a methodology which allows us to explicitly solve the search problem without having to impose such a constraint on the distribution of prices¹: Instead of relying on some unexplained priors and Bayes rule, the consumer uses the Maximum Entropy Principle to estimate the true distribution of prices. Interestingly, in both the discrete and the continuous cases the solution is defined as the myopic rule.

We will also discuss some of the properties of these solutions. We will show the optimal strategy exhibits a monotonous price reservation property. We will also show that as the consumer becomes more patient, the length of his search will increase, and that the consumer may need to take advantage of the recall option. All these are reasonable properties one may expect from a search strategy. They also tally with some empirically observed facts:

- The wage an applicant is willing to accept decreases with the length of the tenure of his search (mentioned by Vishwanath (1984)), this can be explained by the decreasing reservation wage predicted by our solution;

¹If the consumer knows that the true distribution is normal, then he is better off using a parametric Bayesian search (as in Rosenfield and Shapiro (1981)), he could pretend he ignores this fact and use our methodology, but he would then lose some efficiency. On the other hand, the converse is not true: a consumer who just postulates that the true price distribution is normally distributed without having information to that effect exposes himself to problems of inconsistency (his estimate of the true distribution does not improve with time).

- For long duration stints of unemployment the rate of mortality is far greater than the rate of new job offers (see Blau (1988)), this can be explained by job applicants returning to take jobs they did not select at first.

This paper is organized as follows: Section 2 introduces the model and presents the consumer's maximization problem. Section 3 is concerned with consumers holding Dirichlet prior beliefs. Some properties of Dirichlet priors are presented, and a statement presenting the optimal search strategy is provided (proof in section 6). An example is worked out. Section 4 presents an optimal strategy for a class of beliefs which includes the Dirichlet distribution (proof in section 6). The optimal strategy is shown to be a myopic rule, to satisfy the monotonic reservation price property, the use of the recall option is shown to occur in some circumstances, and an increase in the time discount factor is shown to increase the length of the search. Section 5 discusses some of the limitations inherent to the previous framework. A very short introduction to the Maximum Entropy Principle follows. Next is discussed how this principle can be used in the context of search, the optimal strategy is then presented, and an example follows. Section 6 contains the proof of the optimality of the strategies presented in the three previous sections. Section 7 draws the conclusions to this paper.

Section 2: Framework, and definition of search strategies.

2.1 - Framework.

A consumer wants to buy one unit of a good. This good is available in an infinite set of stores that the consumer can visit sequentially. He believes that the prices he will observe during this sequential search are independent realizations of the same random variable which takes its value on some compact set P . We will examine two possibilities: either prices can take a finite number of values $P = \{p^0, \dots, p^N\}$ with $p^0 = 0 < p^1 < \dots < p^N = 1$ ¹, this is the discrete case, or prices belong to the interval $P = [0, 1]$ ², the continuous case. The consumer derives utility from consuming the good, from income and he discounts time as well. More precisely, if the consumer buys the good after searching t stores and at a price p , his utility is³:

$$\delta^t u(1 - p), \text{ with } 0 < \delta < 1, \text{ and } u(0) = 0$$

where δ is his time discount factor, and u is a strictly increasing function; if the consumer does not buy the good, his utility is zero. We will assume that the consumer can always return to a store he visited earlier in his search, that is price quotations are recallable. As a consequence, the consumer's utility only depends on the lowest price he observed at the time he decides to stop his search.

Actually the material presentation of the rest of the paper is much simplified by using the following notation: let $q = 1 - p$ which we will refer to as a q-price quotation. Since we will exclusively deal with q-prices in the reminder of this paper, we will often

¹The assumption that $p^0 = 0$ and $p^N = 1$ is just a matter of convenience so that both cases, discrete and continuous distributions of probabilities, can be treated simultaneously without additional notations.

²This assumption can be interpreted as follows: If the commodity is not a bad, the price must be greater than zero; on the other hand the price the consumer is willing to pay is bounded either by his consumer's surplus (this is the discriminating monopoly price), or by his total income. Therefore the range of interest for prices is compact, and setting the upper bound to 1 is only a normalization convention.

refer to them simply as prices, using sometimes the term q -price for emphasize. The consumer is to look for the highest q -price, and his utility is $\delta^t u(q)$, which is an increasing function of q , so the only modification required in order to transform our model into a job search model is to interpret q as a wage. The set of q -prices is the set Q , with either $Q = \{q^0, \dots, q^N\}$ with $q^0 = 0 < q^1 < \dots < q^N = 1$, or $Q = [0, 1]$.

2.2 - Definition of the consumer's problem.

First we shall introduce some notations. An infinite search would produce an infinite sequence of price observations or quotations, let $h = \{q_1, q_2, \dots\}$ ¹ be such a sequence; we will call it a realization of the history of price search. At time t , only the finite subsequence of the t first components of the eventual realization h is known, and we will denote by $h_t(h)$ the function which associate to a realization the subsequence of prices revealed at time t ². We will also denote by $h_t = \{q_1, \dots, q_t\}$ the subsequence of price itself and call it the past history. Of course at time t the past history h_t does not indicate unambiguously what the whole realization h would eventually be. Let $\hat{q}_t(h_t) = \text{Max}_{\tau \leq t} \{q_1, \dots, q_t\}$ ³ be the order statistic associated with past history h_t , and $\hat{q}_t$ denotes the order statistics itself, which is a random variable at time $\tau < t$, and the realization of this random variable at times $\tau \geq t$.

Our consumer may have some prior information about the distribution of prices he is receiving quotes from, but he does not know the distribution for sure. It is logical for him to use whatever information he gets from his search for the best price to try to improve

¹Let Ω be the set of underlying events, a quote at time t : q_t is a random variable: $\Omega \rightarrow Q$. An infinite search sequence $\{q_1, q_2, \dots\}$ is a random variable: $\Omega^{(1, 2, \dots)} \rightarrow Q^{(1, 2, \dots)}$. h is a realization of such a sequence, and the set of all h is $Q^{(1, 2, \dots)}$.

² $h_t(\cdot)$ is a function from $Q^{(1, 2, \dots)} \rightarrow Q^t$; it is a projection of h on its first- t components. h_t is an element of Q^t , and $\{q_1, \dots, q_t\}$ is either the random variable: $\Omega^t \rightarrow Q^t$, or its realization h_t . As is customary, a random variable will often be identified with its realization when such an identification does not induce confusion.

³ $\hat{q}_t(\cdot)$ is a function: $Q^t \rightarrow Q$, $\hat{q}_t$ is a random variable: $\Omega^t \rightarrow Q$.

his knowledge of this distribution. The consumer we will be considering uses the following scheme: He starts with an initial belief on the cumulative distribution of prices $\tilde{F}_1(.)$ ¹, at this stage he believes his first price quotation will be drawn from this distribution; he then observes this first price quotation q_1 . He then uses an updating mechanism to update his belief on the distribution of prices and obtains a new distribution $\tilde{F}_2(.|q_1)$, at this stage he believes that his next price quotation q_2 will be drawn from this distribution. In his mind, the consumer can keep on repeating this process, determining for all possible past histories what his belief on the distribution of prices is for the next period. Using the law of conditional probabilities, he can then determine what is the joint probability distribution of any realization². We will refer to these probabilities as the **stochastic process** generated by the initial beliefs and the updating mechanism of the consumer.

A **strategy** σ of the consumer is a rule which associate to each possible past history h_t an action of the set $\{S, C\}$, either to stop the search and buy at the best price so far $\hat{q}_t(h_t)$, or to continue searching³, let Σ be the set of all feasible strategies. To each realization h , the strategy σ associates a pair $[q_\sigma(h), t_\sigma(h)]$ ⁴ which are:

¹ $\tilde{F}_1$ is an increasing function $Q \rightarrow [0, 1]$, with $\tilde{F}_1(0) \geq 0$, and $\tilde{F}_1(1) = 1$. In a Bayesian framework, $\tilde{F}_1$ is the total cumulative probability distribution derived from the Bayesian priors (see next section).

²This is the Kolgomorov construction of a probability space for the infinite dimensional space of price quotations (see Gihman and Skorohod (1974)). It is a stochastic process, that is: Consider a possible past history h_t and the let $H|h_t$ be the set of all possible realizations which admit h_t as their first- t components; this construction insures that the unconditional probability of h_t is equal to the probability of $H|h_t$, that is that the consumer is always consistent in his assignment of probabilities form one period to the next.

³Consider a function from $Q^{(1,2,\dots)} \rightarrow Q \times Q^2 \times Q^3 \times \dots$ which associates to each realization h the infinite sequence of all its past histories $\{h_1(h), h_2(h), \dots\}$. Let TQ be the image of Q by this transformation, the set of trces generated by Q . σ is a function from $TQ \rightarrow \{S, C\}^{(1,2,\dots)}$, with the property that $\forall h \in Q, \sigma[h_t(h)] = S \Rightarrow \forall \tau \geq t, \sigma[h_\tau(h)] = S$. σ is really defined on the set of realizations, but the previous property implies that $\{\forall h \in H|h_t, \sigma(h)$ leads to the same decision at time $t\}$ which means

$$t_{\sigma}(h) = \text{Min}_{\tau \in I} \{ \tau / \sigma[h_{\tau}(h)] = S \}$$

$$\text{Let } t_{\sigma}(h) = \tau, q_{\sigma}(h) = \hat{q}_{\tau}(h_{\tau})$$

t_{σ} , the **stopping time** associated with the rule σ , and the **buying price** q_{σ} are random variables which distribution of probability is derived from the joint probability distribution of realizations. This allows us to define $W(\sigma)$ ¹ the value of the search associated with the strategy σ as:

$$W(\sigma) = E[\delta^{t_{\sigma}} u(q_{\sigma})]$$

where the expectations are computed according to the probability distribution we just discussed². The consumer's problem is then to find a strategy which would maximize his expected utility (if such a strategy exists), that is to find σ^* such that:

that one can define the **strategy at time t** $\sigma_t(H|h_t)$ as a function, and the t-th component of σ only depends on the past history h_t .

⁴The functions $q_{\sigma}(\cdot): Q^{\{1,2,\dots\}} \rightarrow Q$, and $t_{\sigma}(\cdot): Q^{\{1,2,\dots\}} \rightarrow \{1, 2, \dots\}$, and the random variables: $q_{\sigma}: \Omega^{\{1,2,\dots\}} \rightarrow Q$, and $t_{\sigma}: \Omega^{\{1,2,\dots\}} \rightarrow \{1, 2, \dots\}$.

¹The function $W: \Sigma \rightarrow \mathbb{R}$. Notice that W is not a random variable.

²This remark illustrates the limitation of the probabilistic approach to search. Consider a consumer who is 'sure' that the price distribution is concentrated at $\{1\}$, that is he believes $\Pr\{q \neq 1\} = 0$, and has a linear utility function and a discount factor of $\delta = 0.9$. Suppose he receives as a first quote $q_1 = 0.5$, then he will interpret this as 'I have been infinitely unlucky', and will continue his search. If the beliefs of the consumer are really represented by the previous statement, then after having observed 432,000 times the price 0.5, the consumer will say 'I have been infinitely unlucky to the power 432,000', and will still continue searching. Of course non-psychotic consumers would not behave that way, which is an indication that real people do expect the impossible to occur with a non-zero probability. The point however is that the consumer only maximizes expected discounted utility with respect to the stochastic process defined by his beliefs, whereas what he would really want to do is to maximize with respect to the true distribution. Notice that this is true quite independently of whether the consumer conducts a Bayesian search or not. Now one can always invoke asymptotic properties of some well chosen estimator of the distribution of prices to

$$W(\sigma^*) = \text{Max}_{\sigma \in \Sigma} \{W(\sigma)\}$$

The value of the search to the consumer is then defined as $V = \text{Max}_{\sigma \in \Sigma} \{W(\sigma)\}$; this value, which the consumer computes prior to receiving any price quotation, depends on one hand on his initial beliefs and on the updating mechanism, on the other hand on his utility $u(.)$ and on his time discount factor¹.

Section 3: Interpretation of the Dirichlet distribution, and derivation of its associated optimal strategy.

3.1 - Interpretation of the Dirichlet distribution.

One of the key to our interest in the Dirichlet distribution is that it can provide a theoretical underpinning to a very appealing behavior. Consider a consumer who faces a discrete distribution of prices $Q = \{q^0, \dots, q^N\}$. Prior to initiating the current search, either in the course of a previous search or because he inadvertently saw price quotes even though he was not in the market at the time, he has observed each price $(\alpha_0, \dots, \alpha_N)$ times, all α 's being strictly positive, but these prices are not available to him any longer ('he has forgotten where the store was'). He then decides on the empirical rule:

'I will assign to the event $\{q = q^n\}$ the probability $\frac{N_n}{N}$ where N_n is the number of times I have observed the price q^n either before or during the course of the present search, and $N = \sum_{n=1}^{n=N} N_n$.'

Proposition 1: The empirical rule above can be interpreted as the Bayesian rule for a consumer with Dirichlet priors $(N_1, \dots, N_N)$.

The proof of this statement follows. Let's consider a consumer whose prior beliefs about this distribution, at this stage of the search, are captured by a Dirichlet prior distribution function with parameters $(\alpha_0, \dots, \alpha_N)$ all strictly positive. This means the following. A given distribution of probability on the set Q is characterized by the vector

where q is the next realization of the price quote¹. The Dirichlet priors assigns a density of probability to each possible probability distribution on Q according to²:

$$f[x_0, \dots, x_N; \alpha_0, \dots, \alpha_N] = \frac{\Gamma(\sum_{i=0}^N \alpha_i)}{\Gamma(\alpha_0) \dots \Gamma(\alpha_N)} x_0^{\alpha_0-1} \dots x_N^{\alpha_N-1}$$

where Γ is the gamma function:

$$\Gamma(\alpha) = \int_0^{\infty} u^{\alpha-1} e^{-u} du \quad \text{for } \alpha > 0$$

Using this priors, we can compute the probability the consumer associates at this stage of his search with the event $\{q = q^n\}$:

$$\begin{aligned} \pi_n(\alpha_0, \dots, \alpha_N) &= \int_0^1 dx_0 (\dots \int_0^1 dx_N (x_n f[x_0, \dots, x_N; \alpha_0, \dots, \alpha_N]) \dots) \\ &= E[x_n; \alpha_0, \dots, \alpha_N] \end{aligned}$$

$[\pi_0, \dots, \pi_N]$ is the vector of **total probabilities**³, it characterizes the belief of the consumer about the cumulative distribution of probability of prices $\tilde{F}$ of the previous section. We can also use Bayes' rule to compute how the consumer revises his beliefs as he receives more price quotes. The results to these two computations is given by the next lemma:

¹The set of all possible distributions of probability on Q is the $N+1$ -Simplex
 $\Delta = \{(x_0, \dots, x_N) \in [0, 1]^{N+1}, \sum_{i=0}^N x_i = 1\}.$

² $f(\cdot; \alpha_0, \dots, \alpha_N)$ is a density of probability measure on Δ , that is a function $\Delta \rightarrow \mathbb{R}^+$ such that $\int_{\Delta} f(x) dx = 1.$

³For all $n = 0, \dots, N$, π_n is an element of $[0, 1]$, and $\sum_{n=0}^N \pi_n = 1.$

Lemma 1: $\pi_n(\alpha_0, \dots, \alpha_N) = \frac{\alpha_n}{\sum_{i=0}^N \alpha_i}$, and if the consumer

observes $\{t_0, \dots, t_N\}$ times the events

$\{\Pr\{q = q^0\}, \dots, \Pr\{q = q^N\}\}$, he will revise his beliefs to a

Dirichlet distribution with parameters $\{\alpha_0 + t_0, \dots, \alpha_N + t_N\}$.

Proof: See DeGroot page 51 equation (6) and page 174 Theorem 1. Q. E. D. •

Proof of Proposition 1: The proposition is true initially by construction, and after search is underway by Lemma 1. Q. E. D. •

There is however a limitation on the interpretation we presented: the consumer cannot start with a set of parameters that includes a zero. So, in our interpretation, we must require that any consumer who wants to apply it to itself must have observed every possible price at least once prior to the inception of the search. The reason that no parameter can be a zero is that the integral:

$$\int_0^1 dx_0(\dots \int_0^1 dx_N (f[x_0, \dots, x_N; \alpha_0, \dots, \alpha_N]) \dots)$$

which is integral of a probability density on the whole space, and is required to be equal to one, diverges if one of the α' s is less or equal to zero.

There is some logic to why this restriction appears in our model, other than a problem with the mathematics. The consumer attaches probabilities to a restricted set of prices. He excludes the possibility that prices other than the ones figuring in the set Q may be observed. Indeed if such a price where to be observed, the consumer would have to rethink his search from scratch. Why would then the consumer entertain the possibility that some prices may appear but not others? It must be the case that he has some information on the prices which may show-up. In our interpretation, the information on a price was qualified by the number of past observations of this price, hence the requirement about the parameters α . What would then do a consumer who does not feel he can exclude some prices from the set $[0, 1]$? We refer the reader to the section on search based on the maximum entropy principle for a solution to this problem.

3.2 - Optimal search strategy for a consumer who holds Dirichlet prior beliefs.

Consider a consumer who starts his search with Dirichlet priors with parameters $(\alpha_0, \dots, \alpha_N)$; let's define the sequence of functions¹:

$$\forall i \in \{0, \dots, N\}, f_t(q^i) = u(q^i) \frac{\sum_{j=0}^{j=i} \alpha_j + t - 1}{\sum_{i=0}^i \alpha_i + t - 1} + \sum_{j=i+1}^{j=N} \frac{\alpha_j u(q^j)}{\sum_{i=0}^i \alpha_i + t - 1}$$

This sequence of functions only depends on the utility function of the consumer and on the parameters of its initial beliefs, they do not depend on the strategy adopted by the consumer. In section 6, we will prove that for all $t > 0$, $f_t(\cdot)$ is an increasing function of q^i , and that for all q^i , $f_t(q^i)$ is a decreasing function of t . For each period t , we can find a benchmark $\xi_t \in Q$ such that:

$$\forall q^i \geq \xi_t, u(q^i) \geq \delta f_t(q^i), \text{ and } \forall q^i < \xi_t, u(q^i) < \delta f_t(q^i)$$

Notice that these benchmarks are the solution to an ordinary equation. Because $f_t(\cdot)$ is a decreasing function in t , the benchmarks ξ_t form a decreasing sequence, and their limit is 0 as will be shown in the same section. The next proposition will define an optimal strategy for the consumer:

Proposition 2. An optimal strategy for the consumer who holds Dirichlet priors is, for each period t :

Stop at t if and only if $\hat{q}_t \geq \xi_t$

Continue at t if and only if $\hat{q}_t < \xi_t$

where $\hat{q}_t$ is the best price observed at time t .

Proof: The proof will be presented in section 4 and section 6. •

This strategy can also be presented more elegantly in terms of stopping times, that is given the best price so far $\hat{q}_t = q^i$, the stopping time $t(q^i)$ tells the consumer to stop now (at time t) if and only if $t \geq t(q^i)$. A straightforward manipulation shows that:

¹For all $t = 1, 2, \dots$, $f_t(\cdot): Q \rightarrow [u(0), u(1)]$.

$$t(q^i) = \frac{1}{(1 - \delta)u(q^i)} \left\{ \left[\delta \sum_{j=0}^{j=i} \alpha_j - \sum_{j=0}^{j=N} \alpha_j \right] u(q^i) + \delta \sum_{j=i+1}^{j=N} \alpha_j u(q^j) \right\}$$

Therefore we have the corollary:

Corollary 1: The previous optimal strategy can be expressed as:

At time t , if $\hat{q}_t = q^i$, Stop if and only if: $t \geq t(q^i)$

Continue if and only if: $t < t(q^i)$

where $t(q^i)$ are the above benchmarks.

Notice that the total probability the consumer, who has observed q^i as his best price in $t - 1$ previous observations, will attribute to the event $\{\hat{q}_t = q^i\}$ is:

$$F_t(q^i | \hat{q}_{t-1} = q^i) = \frac{\sum_{j=0}^{j=i} \alpha_j + t - 1}{\sum_{i=0}^{i=N} \alpha_i + t - 1}$$

and that the probability the same consumer will attribute to the event $\{\hat{q}_t = q^j\}$ for $j > i$ is:

$$\Pr \{ \{\hat{q}_t = q^j\} | \hat{q}_{t-1} = q^i \} = \frac{\alpha_j}{\sum_{i=0}^{i=N} \alpha_i + t - 1}$$

Actually, using these probabilities, this result can be generalized to a larger class of consumers as we will see in the next section.

3.3 - An example: The case of the linear utility function.

Let's consider a consumer whose utility function $u(\cdot)$ is linear:

$$u(q) = aq, a > 0$$

who is facing $N + 1$ prices uniformly distributed on the interval $[0, 1]$, that is:

$$\forall i \in \{0, \dots, N\}, q^i = \frac{i}{N}$$

and who starts with Dirichlet prior beliefs with identical parameters $\alpha_i \equiv \alpha$.

A short calculation shows that the stopping times are:

$$t(q^i) = \frac{\alpha}{q^i(1 - \delta)} \left[\frac{\delta}{2} N [q^i]^2 - (N + 1 - \frac{\delta}{2}) q^i + \frac{\delta(N + 1)}{2} \right]$$

One can check that $t(q^i)$ is a decreasing function of q^i , and that therefore a consumer who stops now at a low price would have stopped as well if he had gotten a better price, just as stated in the previous proposition.

Example: Let's suppose that the set of prices is $Q = \{0, 1/4, 1/2, 3/4, 1\}$. The stopping times for this case are:

$$t(1) = -5\alpha < 0$$

$$t(3/4) = \frac{\alpha}{1 - \delta} \frac{16\delta - 15}{3}$$

$$t(1/2) = \frac{\alpha}{1 - \delta} \frac{13\delta - 10}{2}$$

$$t(1/4) = \frac{\alpha}{1 - \delta} (11\delta - 5)$$

$$t(0) = \infty$$

A negative stopping time means the consumer will buy at this price at the first occasion, or that if he could have chosen between proceeding with the search or buying at this price with no search at all, he would elect to renounce searching. We will not discuss how the time discount factor affects the strategy of the consumer as this will be discussed for the more general case in the next section.

Section 4: Search with a discrete distribution of prices.

4.1 - The discrete case.

In this section we will generalize the proposition established for a consumer with Dirichlet priors, and discuss some properties of the optimal strategy.

Consider the consumer we described in section 2. Let's define $\tilde{F}_t(x| h_{t-1})$ as the cumulative distribution of probabilities the consumer believes in at time $t - 1$ conditional on the observed past history h_{t-1} when forecasting the price quote for time t : q_t . We will only consider distributions such that $\tilde{F}_t(x| h_{t-1})$ only depends on the initial beliefs of the consumer $\tilde{F}_1(\cdot)$, the time t and the number of observations below x in the past history. Actually the consumer is only interested in forecasting the distribution of the highest q -price, the order statistic $\hat{q}_t$. Since $\hat{q}_t \geq \hat{q}_{t-1}$, and since the price quotes are independent random variables, the cumulative distribution of the order statistic $\hat{q}_t$, forecasted at time $t - 1$ conditional on the observed past history h_{t-1} , only depends on the sufficient statistic $\hat{q}_{t-1}$. Let $F_t(x| \hat{q}_{t-1} = q) = \Pr\{\hat{q}_t \leq x | \hat{q}_{t-1} = q\}$ be the cumulative distribution of the order statistic $\hat{q}_t$:

$$F_t(x| \hat{q}_{t-1} = q) = 0 \quad \text{for } x < q$$

$$F_t(x| \hat{q}_{t-1} = q) = \Pr\{q_t \leq x | h_{t-1}\} \quad \text{for } x \geq q$$

$$F_t(1| \hat{q}_{t-1} = q) = \Pr\{q_t \leq 1 | h_{t-1}\} = 1$$

We will impose the following assumptions on the stochastic process defined by the initial beliefs and on the updating mechanism of the consumer.

Assumptions (A):

- (A-1) $F_t(x| h_t) = F_t(x| \hat{q}_{t-1} = q)$, and $\forall x \geq q' > q$,
 $F_t(x| \hat{q}_{t-1} = q') = F_t(x| \hat{q}_{t-1} = q)$
- (A-2) $\forall x \in Q \cap [q, 1]$, $\forall t \geq 1$, $F_{t+1}(x| \hat{q}_t = q) \geq F_t(x| \hat{q}_{t-1} = q)$
- (A-3) $\forall x \in Q \cap [q, 1]$, $\Pr\{\hat{q}_{t+1} = x | \hat{q}_t = q\} \leq \Pr\{\hat{q}_t = x | \hat{q}_{t-1} = q\}$
- (A-4) $\forall \varepsilon > 0$, $\exists T(\varepsilon) \geq 1$, $\forall t \geq T(\varepsilon)$, $\forall q \in [0, 1]$,
 $0 < 1 - F_t(q | \hat{q}_{t-1} = q) < \varepsilon$
- (A-5) $\forall \varepsilon > 0$, $\forall t \geq 1$, $\exists T(t) \geq 1$, $\Pr\{\hat{q}_{T(t)+t} = 0 | \hat{q}_t = 0\} < \varepsilon$
- (A-6) The utility function is a strictly increasing function of q .

Interpretation of the assumptions:

- (A-1) states that $\hat{q}_{t-1}$ is a sufficient statistic to determine the cumulative distribution of $\hat{q}_t$, and that F_t only depends on the number of observations below $\hat{q}_{t-1}$, not its value.
- (A-2) states that the distribution of the order statistic of period $t - 1$ first order stochastically dominates the order statistic of period t . That is as the search progresses the probability to improve over what has already been achieved decreases with time.
- (A-3) Assumption (A-2) implies that the probability that $\hat{q}_{t+1}$ strictly improves over $\hat{q}_t$ is not as great as the probability that $\hat{q}_t$ strictly improves on $\hat{q}_{t-1}$. This assumption adds that the lower probability for $\hat{q}_{t+1}$ holds not only globally for the all set of realizations which would constitute an improvement, but also for each individual realization in that set.
- (A-4) states that after a consumer has spent enough time searching, he believes his chances of improving tend to zero.
- (A-5) states that no matter how long the consumer has spent searching an encountered only zero-prices, he still believes that, given enough time, he will most certainly find a non-zero price. This assumption does not contradict (A-4).

Lemma 2: The distribution generated by Dirichlet beliefs satisfy assumptions (A).

Proof: Since the cumulative distribution and probability distribution of the order statistic generated by Dirichlet priors are:

$$F_t(q^i | \hat{q}_{t-1} = q^i) = \frac{\sum_{j=0}^{j=i} \alpha_j + t - 1}{\sum_{i=0}^N \alpha_i + t - 1}$$

and
$$\Pr \{ \{ \hat{q}_t = q^j \} | \hat{q}_{t-1} = q^i \} = \frac{\alpha_j}{\sum_{i=0}^N \alpha_i + t - 1}$$

the assumptions are satisfied. Q. E. D. •

4.2 - Optimal Strategy and its properties.

Let's define the function $f_t(\cdot)$ as:

$$\forall q^i \in Q, f_t(q^i) = u(q^i)F_t(q^i | \hat{q}_{t-1} = q^i) + \sum_{q^j > q^i} u(q^j) \Pr \{ \hat{q}_t = q^j | \hat{q}_{t-1} = q^i \}$$

Proposition 3: Under assumptions (A), for all $t \geq 1$, there exists a unique sequence of decreasing benchmarks $\{\xi_1, \xi_2, \dots\} \in Q^{(1, 2, \dots)}$ such that:

$$\forall x \in Q, x \geq \xi_t \Leftrightarrow u(q^i) \geq \delta f_{t+1}(q^i),$$

and
$$x < \xi_t \Leftrightarrow u(q^i) < \delta f_{t+1}(q^i)$$

$$1 \geq \xi_1 \geq \dots \geq \xi_t \geq \dots > 0$$

$$\lim_{t \rightarrow \infty} \xi_t = 0$$

and the following strategy σ^* is optimal:

$$\forall t \geq 1, \sigma^* \{ \hat{q}_t < \xi_t \} = C, \sigma^* \{ \hat{q}_t \geq \xi_t \} = S.$$

Proof: Section 6 will prove this proposition. Q. E. D. •

Notice that the benchmarks are the solution to ordinary equations involving the utility function of the consumer and the conditional probability beliefs of the consumer. The next propositions establishes some properties of this optimal solution.

Proposition 4: The optimal strategy imply that the consumer may need to use the recall option.

Proof: The benchmarks are decreasing; suppose that at time t the price quote turns out to be the best price so far $0 < q_t = \hat{q}_t$, but that $q_t < \xi_t$ so that the consumer will continue searching. We know that $\exists T \geq t$, such that $\xi_T < q_t$ since $\lim_{t \rightarrow \infty} \xi_t = 0$. Consider the history up to time T for which:

$$\forall \tau \in \{t+1, t+2, \dots, T\}, q_\tau < \hat{q}_T = q_t$$

Then at T , the optimal strategy of the consumer directs him to stop, and the consumer will return to the store he visited at time t to buy. Q. E. D. •

Proposition 5: This optimal strategy is a myopic rule, and it exhibits a decreasing reservation q -price property.

Proof: Since $f_{t+1}(\cdot)$ is the expected value of the utility of the consumer for the next quote, and since the consumer stops if and only if $\delta f_{t+1}(\hat{q}_t) \leq u(\hat{q}_t)$, then his strategy is myopic.

The reservation q -prices are the benchmarks ξ_t 's, and these benchmarks are decreasing. Q. E. D. •

Now consider a sequence of consumers, each with the same utility function $u(\cdot)$ and the same initial prior beliefs and updating mechanism, but with increasing time discount factors $\delta_1 < \dots < \delta_n < \dots < 1$, and $\lim_{n \rightarrow \infty} \delta_n = 1$. Increasing time discount factors corresponds to increasingly patient consumers, and $\delta = 1$ corresponds to a consumer who does not discount time. Let's assume that we are conducting an experiment in which all these consumers independently visit the same string of stores in the same order. They receive identical realized price quotes $\{q_1, q_2, \dots\}$, and let's assume that:

$$\forall t \geq 1, q_t < 1 \text{ and } \lim_{t \rightarrow \infty} q_t = 1.$$

The next proposition is a statement about the duration of the search as a function of the time discount factor.

Proposition 6: Increasingly patient consumers such as in the above experiment will conduct a search of increasing duration, and in the limit, the consumer who does not discount time will search forever.

Proof: Let $\{\xi_1(\delta_n), \xi_2(\delta_n), \dots\}$ be the sequence of benchmarks for consumer n . Because all consumers share the same utility function $u(\cdot)$ and the same stochastic process, they also share the same functions $f_t(\cdot)$. Consider the benchmark $\xi_t(\delta_n)$, it is such that:

$$\forall x \in Q, x < \xi_t(\delta_n) \Leftrightarrow u(x) < \delta_n f_{t+1}(x)$$

Suppose that $m > n$, then $\delta_m > \delta_n$, therefore $u(x) < \delta_n f_{t+1}(x) < \delta_m f_{t+1}(x)$, so that $x < \xi_t(\delta_m)$. Henceforth $\xi_t(\delta_n) \leq \xi_t(\delta_m)$, so any q -price quote that would have stopped consumer m would also have stopped consumer n . Since these two consumers face the same sequence of price realizations, this means that consumer n stops before consumer m does.

We have also:

$$\begin{aligned} \forall q \in Q, \delta_n f_t(q) - u(q) &= \sum_{q^i > q} [\delta_n u(q^i) - u(q)] \Pr\{\hat{q}_t = q^i | \hat{q}_{t-1} = q\} \\ &\quad - (1 - \delta_n) F_t(q | \hat{q}_{t-1} = q) u(q) \end{aligned}$$

Therefore $\forall t \geq 1$,

$$\lim_{n \rightarrow \infty} [\delta_n f_t(q) - u(q)] = \sum_{q^i > q} \lim_{n \rightarrow \infty} [\delta_n u(q^i) - u(q)] \Pr\{\hat{q}_t = q^i | \hat{q}_{t-1} = q\}.$$

If $\delta_n \rightarrow 1$, then $\forall q \in Q, q \neq 1, \forall q^i > q, \exists N \geq 1, \forall n \geq N, \delta_n u(q^i) - u(q) > 0$, and therefore for all $t \geq 1$, and all $q \in Q, q \neq 1, \lim_{n \rightarrow \infty} [\delta_n f_t(q) - u(q)] > 0$. Therefore $\forall t \geq 1, \lim_{n \rightarrow \infty} [\xi_t(\delta_n)] = 1$, and the consumer who does not discount time searches forever. Q. E. D. •

Section 5: Search with a continuous distribution of prices.

5.1 - Bayesian analysis and the Maximum Entropy Principle.

Since prices are always expressed in terms of a smallest indivisible unit (say in terms of cents), it would appear that the previous section gives a reasonable answer to the problem of search, save for the special case when the consumer has not been able to observe all prices beforehand. We are going to argue that this special case is however the most relevant one: Consider the opposite case, in which the consumer was able to observe a 'large' number of prices¹. Then the histogram of the empirical distribution, which produces a consistent estimator of the true cumulative distribution of probabilities by the strong law of large numbers, is a good estimator of this true cumulative distribution. Then the consumer operates as if the distribution of prices were known, and Diamond's argument should apply for economies with duplicate agents.

On the other hand, if the consumer has not observed a 'large' number of prices, he cannot be sure of having observed all prices. He can then either ignore unobserved prices, but at the risk that his search will breakdown if one of these shows up², or he can assign them arbitrary priors. In the latter case, the solution to the search problem depends on an element which is not explained by the model.

The strength of the methodology we will use, the Maximum Entropy Principle, is that it offers a guide on how to form initial beliefs. In addition, consider the case of consumer who has no experience with the market, and who is unwilling to make assumptions about the shape of the distribution of probability. Then at any time in his search, the consumer divides the set of prices between the ones which are below the maximum he already observed, and whose distribution of probability is of no concern to

¹By large we mean that the strong law of large numbers applies to the sample.

²Bayes rule does not allow to update a distribution with an observation which had zero total prior probability.

him, and the one above this maximum, which by definition he has not observed. The problem of the consumer is clearly to try and forecast the probability of an event he never observed. In that sense the consumer never learns anything about prices he did not observe, but only on the range these prices may belong to. The estimator of the price distribution we propose is particularly well suited to take advantage of the latter information.

On an other hand, the restriction of a discrete set of prices is not entirely welcome either¹. In most of the rest of economic literature the range of prices is continuous, therefore for the sake of comparison, one would not want to discard this assumption when studying search. As we will see, the Maximum Entropy Principle can be harnessed to solve this problem.

¹There are considerable technical difficulties in using Bayesian analysis with a continuous random variable. The set Δ mentioned in the footnote of section 2 becomes a set whose dimension has the power of the continuous. Defining a probability distribution on such a set is a problem: for instance there are no uniform priors on such a set (just as there are no uniform priors for the real line). A standard way around this problem is to assume the true distribution belongs to some restricted set which can be finitely parameterized, such as the class of normally distributed densities of probability. However, the consumer should not assume such a restriction, unless he really knows it holds, first because the estimator he would get would not be consistent, second because having a continuous set of prices does not mean that no price should be assigned a probability mass, hence the possibility of repeated observation that the Bayesian consumer cannot handle.

5.2 - The Maximum Entropy Principle.

Since some readers may not be familiar with the Maximum Entropy Principle, we shall give a brief intuitive introduction to the matter, based on an example¹. Our presentation is based on the notion of information as developed by Shannon (1948).

Let's suppose that a statistician has placed a bet on a soccer match between team A and team B, and is waiting for the reliable message informing him who won. If the two team are, in the statistician opinion, of equal strength he can not say that the message "A has won" surprises him more than the message "B has won", both of them carry the same 'amount of information'. On the other hand if team A is the Bayern of Munich (a leading European soccer team) and team B is the local ball club of Knocklezout-les-Oies, if he receives the message "team A won" he may say 'I already knew that', whereas if he receives the message "team B won", he may feel he is living through Homeric times. Indeed, if an event has a probability one to occur, the message that this event actually occurred provides no information to the statistician. This remarks leads to the notion that if a random variable can have N possible realizations say $\{x_1, \dots, x_N\}$ to which the statistician assigns subjective probabilities $\{p_1, \dots, p_N\}$ with $\sum_{n=1}^{n=N} p_n = 1$, the information content of the message "event x_n occurred" is assigned an information content $f(1/p_n)$, where f is an increasing function. In order to obtain some desirable properties, for instance that the information provided by two independent messages is the sum of the information provided by each event, f is taken to be the logarithmic function $f(1/p_n) = \log 1/p_n$.

Next is defined the expected value of the information that the statistician is waiting to receive a.k.a. the entropy of the message. It depends on the subjective assignment of probability the statistician makes prior to receiving the message:

$$I(p_1, \dots, p_N) = - \sum_{n=1}^{n=N} p_n \log p_n$$

¹Interested readers are referred to Theil (67) for a more thorough introduction, to Jaynes (83) for a discussion of the place of the Maximum Entropy in the development of probability and statistics, and to Shore (1980) and Skilling (1986) about an axiomatic derivation of the Maximum Entropy Principle.

The statistician may have some information about the true distribution of probabilities. This information is particularly easy to process with the Maximum Entropy Principle when it is provided in the form of a constraint on the moments of the distribution, like the expected value, and so on. These constraints plus the equality $\sum_{n=1}^{n=N} p_n = 1$ define a set of admissible distributions of probability for the problem.

The Maximum Entropy Principle then directs the statistician to choose among these the distribution which maximizes the entropy, that is the statistician ought to maximize the information content of the message he is about to receive . Since the entropy is a concave function of the probability distribution, providing the constraints mentioned above define a convex set, the solution to this problem is unique.

A useful remark is that, when the statistician does not have any information about the distribution of probability of the random variable under consideration, the Maximum Entropy Principle tells him to assign equal probability to all events.

This idea can be extended to distributions of probability on continuous sets. Suppose the set of realizations is the interval $[0, 1]$, and that the subjective density of probability of observing event x is $f(x)$, then the entropy of this distribution is:

$$I[f] = - \int_0^1 f(x) \log f(x) dx$$

and the statistician proceeds as before. If the the statistician does not have any information about the distribution of probability of the random variable under consideration, then the Maximum Entropy Principle directs him to select the uniform distribution on that interval. We will make use of this remark next.

5.3 - Search for a consumer who uses the Maximum Entropy Principle.

Let' s consider a consumer who knows the q -prices belong to the interval $[0, 1]$, but who has no ground to restrict his beliefs to a finite set of prices. The simplest case is when the consumer has no further information about the price distribution, we will refer to him as the 'naive' consumer. We may also want to consider the case of a consumer who has in the past observed some prices which are not recallable now, he would then be an 'experienced' consumer.

Distribution function of the beliefs of a naive consumer:

The consumer we are considering has no previous experience with, or prior information about the market he is in, except that price quotes belong to $[0, 1]$, and that quotes are independently identically distributed random variables. He uses the Maximum Entropy Principle to attach probabilities to observing price quotes. Prior to receiving any price quote the consumer will opt for the distribution of probability F_1 whose density function maximizes:

$$\text{Max} \int_0^1 f(x) \log f(x) dx$$

subject to: $\int_0^1 f(x) dx = 1$

It is clear that the solution to this problem is:

$$f(x) \equiv 1, \text{ or } F_1(x) = \int_0^x f(x) dx = x$$

Let's now assume that the consumer has received a price quote q_1 . He needs now to forecast what the probability distribution of the order statistic will be if he draws an additional quote. It is clear that this order statistic $\hat{q}_2$ will be greater or equal to q_1 . Consider two prices $q_1 < q < q' \leq 1$, under which circumstances does the information that q_1 has been observed in the first drawing incite the consumer to think q or q' are more likely to be drawn next? If the consumer has some information or beliefs about how the probability attached to q_1 is related to the ones attached to q and q' . Since we are interested in a consumer who does not draw this kind of inference, we must conclude that the observation of q_1 should not lead the consumer to attach different probabilities to q and q' , that is the consumer must opt for the uniform distribution of probability on $(q_1, 1]$. The question then is how much weight should the consumer attach to the event $\{q_2 \leq q_1\} = \{\hat{q}_2 = q_1\}$, as opposed to the event

$\{q_2 > q_1\} = \{\hat{q}_2 > q_1\}$. Suppose the true distribution of prices is F^* , and that the price quotes really are i.i.d. random variable from $[0, 1]$. Then the true expected value of the probability of the event $\{q_2 > q_1\}$ is:

$$E[\Pr\{q_2 > q_1\}] = \int_0^1 \Pr\{q_2 > q_1\} dF^*(q_1) = \int_0^1 [1 - F^*(q_1)] dF^*(q_1) = \frac{1}{2}$$

We propose that this be the probability the consumer will attach to the event $\{\hat{q}_2 > q_1\}$.

More generally speaking, we can compute the expected value of $\Pr\{q_{t+1} \leq \hat{q}_t\}$ as:

$$\int_0^1 \Pr\{q_{t+1} \leq q\} d\Pr\{\hat{q}_t = q\} = \int_0^1 F^*(q) t [F^*(q)]^{t-1} dF^*(q) = \frac{t}{t+1}$$

since
$$d\Pr\{\hat{q}_t = q\} = \sum_{s=1}^{s=t} \Pr\{\forall s' \neq s, q_{s'} \leq q\} d\Pr\{q_s = q\} = t [F^*(q)]^{t-1} dF^*(q)$$

and we propose that after having observed a history h_t such that $\hat{q}_t = q$, the consumer processes this information as the constraint $F_{t+1}(q | \hat{q}_t = q) = \frac{t}{t+1}$ on the distribution of prices quotes for time $t+1$. The consumer then derives the density of distribution f_{t+1} as the solution to:

$$\begin{aligned} & \text{Max } \int_0^1 f_{t+1}(x) \log f_{t+1}(x) dx \\ \text{subject to: } & \int_0^1 f_{t+1}(x) dx = 1, \text{ and } \int_q^1 f_{t+1}(x) dx = \frac{1}{t+1}, \quad f_{t+1}(q) = \text{Dirac}\left(\frac{t}{t+1}\right) \end{aligned}$$

The distribution function of $\hat{q}_{t+1}$ is then:

$$\begin{aligned} F_{t+1}(x | \hat{q}_t = q) &= 0 & \text{if } 0 \leq x < q \\ F_{t+1}(x | \hat{q}_t = q) &= \frac{t}{t+1} + \frac{x - q}{(t+1)(1-q)} & \text{if } 1 \geq x \geq q \end{aligned}$$


Distribution function of the beliefs of an experienced consumer:

Let's suppose the experienced consumer has observed prices $\{q^0, \dots, q^i, \dots, q^N\}$ each a number $\{\alpha_0, \dots, \alpha_i, \dots, \alpha_N\}$ of times, with the convention $q^0 = 0$ and $q^N = 1$ (and therefore α_0 and α_N may be equal to 0). We will assume the consumer treats these pre-search price quotes as he does for current ones. This leads him to attribute the following probability distribution to $\hat{q}_{t+1} | \hat{q}_t = q$, when $q^i \leq q < q^{i+1}$:

$$\begin{aligned} F_{t+1}(x | \hat{q}_t = q) &= 0 & \text{if } 0 \leq x < q \\ F_{t+1}(x | \hat{q}_t = q) &= \frac{t + \sum_{j=0}^{j=i} \alpha_j}{t + \sum_{j=0}^{j=N} \alpha_j + 1} + \frac{x - q}{(t + \sum_{j=0}^{j=N} \alpha_j + 1)(q^{i+1} - q)} & \text{if } q^{i+1} > x \geq q \\ F_{t+1}(x | \hat{q}_t = q) &= \frac{t + \sum_{j=0}^{j=i'} \alpha_j}{t + \sum_{j=0}^{j=N} \alpha_j + 1} + \frac{x - q^{i'}}{(t + \sum_{j=0}^{j=N} \alpha_j + 1)(q^{i'+1} - q^{i'})} & \text{if } q^{i'+1} > x \geq q^{i'} \text{ for } i' > i \\ F_{t+1}(1 | \hat{q}_t = q) &= 1 \end{aligned}$$

Notice that this distribution is stepwise continuous¹, the set of jump points being $J \cap [\hat{q}_t, 1]$, where $J = \{q^0, \dots, q^i, \dots, q^N\}$ if $\alpha_0 \neq 0$ and $\alpha_N \neq 0$, and the obvious alteration if either of these parameters is zero.

The graph below represents this distribution when there are two such discontinuity points at q^1 and q^2 :



Assumptions (A'):

The stochastic process generated by the believes of the consumer is characterized by the above distribution of probability, and the utility function of the consumer is a strictly increasing concave differentiable function.

These assumptions imply (but are not equivalent to) assumptions below:

¹In the consumer's opinion the true distribution of probability may well be a mixture of a continuous and of a discrete distribution of probability. For repeated observations, the consumer believes they correspond to a price to which the true distribution of probability assigned a probability mass.

$$(A'-1) F_t(x| h_t) = F_t(x| \hat{q}_{t-1} = q)$$

$$(A'-3) \forall x \in (Q \setminus J) \cap [q, 1], dF_t(x| \hat{q}_{t-1} = q) \geq dF_{t+1}(x| \hat{q}_t = q) \text{ and}$$

$$\forall x \in J \cap [q, 1], F_t(x| \hat{q}_{t-1} = q) - F_t(x-1| \hat{q}_{t-1} = q) \geq F_{t+1}(x| \hat{q}_t = q) - F_{t+1}(x-1| \hat{q}_{t-1} = q)$$

(A'-6) The utility function is a strictly increasing, differentiable, strictly concave function of q .

Assumptions (A-2), (A-4) and (A-5) without modifications.

$F_t(x-)$ is the left hand limit of the function F_t at the point x , dF_t is the differential of F_t at points at which it exists and $Q \setminus J = \{x \in Q, x \notin J\}$.

5.4 - Optimal Strategy and its properties.

Let's define the function $f_t(\cdot)$ as:

$$f_t(q) = u(q)F_t(q| \hat{q}_{t-1} = q) + \int_q^1 u(x)dF_t(x| \hat{q}_{t-1} = q)$$

Proposition 7: Under assumptions (A'), for all $t \geq 1$, there exists a unique sequence of decreasing benchmarks $\{\xi_1, \xi_2, \dots\} \in Q^{(1, 2, \dots)}$ such that:

$$\forall x \in Q, x \geq \xi_t \Leftrightarrow u(q^i) \geq \delta f_{t+1}(q^i),$$

$$\text{and } x < \xi_t \Leftrightarrow u(q^i) < \delta f_{t+1}(q^i)$$

$$1 \geq \xi_1 \geq \dots \geq \xi_t \geq \dots > 0$$

$$\lim_{t \rightarrow \infty} \xi_t = 0$$

and the following strategy σ^* is optimal:

$$\forall t \geq 1, \sigma^*\{\hat{q}_t < \xi_t\} = C, \sigma^*\{\hat{q}_t \geq \xi_t\} = S.$$

Proof: Section 6 will prove this proposition. Q. E. D.

•

$$t(q) = \frac{1}{(1 - \delta)u(q)} \left\{ \left[\delta \sum_{j=0}^{i-1} \alpha_j - \sum_{j=0}^{i-1} \alpha_j u(q) + \delta \sum_{j=i+1}^N \alpha_j u(q^j) + [\delta g(q) - u(q)] \right] \right\}$$

with

$$g(q) = \int_q^{q^{i+1}} \frac{u(x)}{q^{i+1} - q} dx + \sum_{j=i+1}^{j=N-1} \int_{q^j}^{q^{j+1}} \frac{u(x)}{q^{j+1} - q^j} dx$$

This leads to the two corollaries:

Corollary 2: The previous optimal strategy can be expressed as:

At time t , if $\hat{q}_t = q$,

Stop if and only if: $t \geq t(q)$

Continue if and only if: $t < t(q)$

where $t(q^i)$ are the above benchmarks.

Corollary 3: The solution to the search problem of a consumer using a continuous distribution model verifies the properties mentioned for the solution to the discrete case, namely propositions 4, 5 and 6.

5.5 - An example: The case of the linear utility function.

Let's consider a consumer whose utility function $u(\cdot)$ is linear:

$$u(q) = aq, a > 0$$

and who is a naive consumer. Then a short calculation shows that the reservation prices for this consumer are:

$$\forall t \geq 1, \xi_t = \frac{\delta}{2(1 - \delta)t + 2 - \delta}$$

Section 6: Existence and characterization of the optimal strategy.

6.1 - Intuition for the proof of the optimality of the myopic rule.

The construction of the optimal strategy is based on the following remark. Consider the myopic rule, it is defined by a sequence of decreasing reservation q -prices (lemma 6 and 7), that is the consumer is less and less finicky. If the consumer were to search with a finite horizon, then by backward induction the myopic rule would be the optimal rule (see theorem 4). Such a finite horizon is provided by theorem 3 for an optimal strategy, the existence of which is established by theorem 2. Lemma 5 is instrumental in establishing the finiteness of search: it is clear that the most arduous problem in finding the optimal strategy is to compare strategies which do not stop under the same circumstances. Lemma 5 is the place where this is done in this proof for one particular case. It relies on the assumption that in the sequence of estimators the consumer uses to estimate the true distribution of probability, each estimator first order stochastically dominates the next.

6.2 - Existence of a Regular Optimal Strategy.

An optimal strategy is said to be **regular** if it stops in finite time almost surely and if, at any time, it directs the consumer to continue the search only when the expected value of the continuing the search under the said strategy is strictly greater than the value of stopping immediately.

Given a past history $h_t = (q_1, q_2, \dots, q_t)$ in general a strategy would associate a decision at time t to the whole vector h_t . However, since the consumer is only interested in the best price, and since the stochastic process defined by his beliefs only depends on the sufficient statistic $\hat{q}_t$ because of assumption (A-1) or (A'-1), his decision to stop or continue at time t should only depend on this sufficient statistic. A strategy which at time t only depends on the sufficient order statistic $\hat{q}_t$ is said to be **history independent**. It is clear that an optimal strategy must be history independent.

Given a set of realizations A and a past history h_t , let $A|h_t$ be the set (possibly empty):

$$\{A|h_t\} = \{h \in A, h_t(h) = h_t\}$$

Given a set A of realizations, the expected value of some random variable $z(h)$ which is a function of the realization of price quotes is defined as:

$$E[z|A] = \int_A z(h) d\Pr(h|h \in A)$$

Lemma 3. Let's suppose assumptions (A) or (A') hold. Suppose that $\Pr\{A|h_t\} = \alpha > 0$, then there exists a time T and $\delta_0 > 0$ such that $\forall \tau \geq T, E[u(\hat{q}_\tau)|\{A|h_t\}] > \delta_0$.¹

Proof: Assumption (A-5) insures that there exists a time $T \geq 1$ such that

$\Pr\{\hat{q}_T > 0|h_t\} > 1 - \frac{\alpha}{2}$. We have:

$$\begin{aligned} \Pr[\{\hat{q}_T > 0\} \cap \{A|h_t\}] &= \Pr\{\hat{q}_T > 0|h_t\} + \Pr\{A|h_t\} - \Pr[\{\hat{q}_T > 0\} \cup \{A|h_t\}] \\ &\geq 1 - \frac{\alpha}{2} + \alpha - 1 = \frac{\alpha}{2} > 0. \end{aligned}$$

Since $\lim_{n \rightarrow \infty} \{\hat{q}_T > \frac{1}{n}\} = \{\hat{q}_T > 0\}$, there exists an n_0 such that

$\Pr[\{\hat{q}_T > \frac{1}{n_0}\} \cap \{A|h_t\}] > \frac{\alpha}{4}$. Let $\delta_0 = \frac{\alpha u(1/n_0)}{4}$. Then

$$\forall \tau \geq T, E[u(\hat{q}_\tau)|\{A|h_t\}] \geq E[u(\hat{q}_T)|\{A|h_t\}] > \delta_0 \text{ Q. E. D.}$$

We are now going to extend the notion of stopping times to finite past histories.

Given a past history h_t , for $\sigma(h_t) = S$, let's define $t_\sigma[h_t]$ to be the stopping time as before, and let's denote by $t_\sigma(h) > t$ the case $\sigma(h_t) = C$. Given a past history h_t and a strategy σ , such that $t_\sigma(h_t) > t$, let's define the set of realizations:

$$A_\tau(\sigma, h_t) = \{h \in Q^{(1, 2, \dots)} / h_t(h) = h_t, t_\sigma(h) \geq \tau\}$$

and $A_\infty(\sigma, h_t) = \bigcap_{\tau \in \{1, 2, \dots\}} A_\tau(\sigma, h_t)$

¹Notice that the expectation which is defined is not the value of some strategy, since no discount factor is involved. In particular this lemma should not be interpreted as urging the consumer to continue past time T .

For any strategy σ , let's define $\sigma|h_t$ as its **conditional strategy** that is the strategy which 'starts' at time t , that is:

$$\forall \tau \geq 1, \tau < t \Rightarrow (\sigma|h_t)[h_\tau(h_t)] = C, \text{ and } \tau \geq t \Rightarrow (\sigma|h_t)[h_\tau] = \sigma[h_\tau].$$

Let $\Sigma|h_t$ be the set of all these conditional strategies. Let's define values of the search after some history has been observed:

$$W_t(\sigma|h_t) = E[\delta^{(t_{\sigma|h_t}-t)} u(q_{\sigma|h_t}) | h_t]$$

Notice that W_t reflects value discounted back to time t , and not to the inception of the search. For history independent strategies, h_t is summarized by $\hat{q}_t$; since optimal strategies are history independent, we can define the value of the search at time t conditional on $\hat{q}_t = q$ as:

$$V_t(q) = \text{Max}_{\sigma \in \Sigma|h_t} \{W_t(\sigma|h_t)\}$$

Theorem 1. Let's suppose assumptions (A) or (A') hold. Given a past history h_t and a strategy σ , such that $t_\sigma(h_t) > t$ and $\Pr\{A_\infty(\sigma, h_t)\} = \alpha > 0$, there exists a strategy σ' such that $t_{\sigma'}(h_t) > t$ and $\Pr\{A_\infty(\sigma', h_t)\} = 0$ and $W_t(\sigma'|h_t) > W_t(\sigma|h_t)$.

Proof: By lemma 3, there exists T_0 and δ_0 such that $\forall \tau \geq T_0 E[u(\hat{q}_\tau) | A_\infty(\sigma, h_t)] > \delta_0$.

Because $A_\tau(\sigma, h_t)$ is a decreasing sequence of sets converging to $A_\infty(\sigma, h_t)$,

$$\Pr\{A_\tau(\sigma, h_t)\} \rightarrow \Pr\{A_\infty(\sigma, h_t)\}, \text{ and } \exists T_1, \forall \tau \geq T_1, \Pr\{A_\tau(\sigma, h_t) \setminus A_\infty(\sigma, h_t)\} < \frac{\delta_0 \alpha}{u(1)},$$

where the set $A \setminus B = \{x \in A, x \notin B\}$.

Let $T = \text{Max}\{T_0, T_1\}$, and let's define the strategy σ' as follows: First we impose $t_{\sigma'}(h_t) > t$. Then on the set of realizations $H(h_t) = \{h \in Q^{(1, 2, \dots)} / h_t(h) = h_t\}$ compatible with past history h_t , we define σ' as:

$$\begin{aligned}
W_t(\sigma' | h_t) &= E[W_t(\sigma' | h_t) | H(h_t) \setminus A_T(\sigma, h_t)] \Pr\{H(h_t) \setminus A_T(\sigma, h_t)\} + \\
&\quad E[W_t(\sigma' | h_t) | A_T(\sigma, h_t) \setminus A_\infty(\sigma, h_t)] \Pr\{A_T(\sigma, h_t) \setminus A_\infty(\sigma, h_t)\} + \\
&\quad E[W_t(\sigma' | h_t) | A_\infty(\sigma, h_t)] \Pr\{A_\infty(\sigma, h_t)\} \\
&= W_t(\sigma | h_t) + \\
&\quad E[W_t(\sigma' | h_t) - W_t(\sigma | h_t) | A_T(\sigma, h_t) \setminus A_\infty(\sigma, h_t)] \Pr\{A_T(\sigma, h_t) \setminus A_\infty(\sigma, h_t)\} + \\
&\quad E[W_t(\sigma' | h_t) - W_t(\sigma | h_t) | A_\infty(\sigma, h_t)] \Pr\{A_\infty(\sigma, h_t)\} \\
&\geq W_t(\sigma | h_t) - \delta^T u(1) \Pr\{A_T(\sigma, h_t) \setminus A_\infty(\sigma, h_t)\} + E[\delta^T u(q_T) | A_\infty(\sigma, h_t)] \alpha \\
&\geq W_t(\sigma | h_t) - \frac{\delta^T u(1) \delta_0 \alpha}{u(1)} + \delta^T \delta_0 \alpha \\
&= W_t(\sigma | h_t) \quad \text{Q. E. D.} \quad \bullet
\end{aligned}$$

Lemma 4: Given any $t \geq 0$, and any past history h_t , let $W_t^* = \sup_{\sigma | h_t} \{W_t(\sigma | h_t)\}$. Then there exists a strategy $\sigma^* | h_t$ such that $W_t(\sigma^* | h_t) = W_t^*$.

Proof: See DeGroot (1970), Theorem 1, p 347. Q. E. D. •

Theorem 2. Let's suppose assumptions (A) or (A') hold. Given any $t \geq 0$, and any past history h_t , there exists a regular strategy σ^* such that $W_t(\sigma^* | h_t) = \sup_{\sigma} \{W_t(\sigma | h_t)\}$.

Proof: Lemma 4 establishes the existence of a strategy whose expected value reaches the supremum of expected utility, whereas Theorem 1 proves that this strategy stops in finite time with probability one. It is clear that an optimal strategy which stops in finite time cannot direct a consumer to continue searching if the value of the subsequent search is less than the value of stopping immediately, hence this strategy is regular. Q. E. D. •

6.3 - Existence of a Finite Time Horizon When a Strictly Positive Price Has Been Observed.

Lemma 5: Both assumptions (A) and (A') imply that $V_{t+1}(q) \leq V_t(q)$.

Proof: Since we are looking at optimal strategies, we can restrict ourselves to history independent histories. Consider the conditional strategies associated with the history

summarized by $\hat{q}_{t+1} = q$, we will note them $\sigma^t q$. Let $\sigma^{t+1} q$ be the optimal strategy¹ which yields $V_{t+1}(q)$ given that $\hat{q}_{t+1} = q$. We want to construct a strategy $\sigma^t q$ which will yield a higher value to the search given that $\hat{q}_t = q$. Let's define τ^{t+1} as the stopping time associated with the strategy $\sigma^{t+1} q$, and $q_{\sigma^{t+1}}$ the associated buying price. Then:

$$V_{t+1}(q) = E[\delta^{\tau^{t+1} - (t+1)} u(q_{\sigma^{t+1}}) | \hat{q}_{t+1} = q]$$

We are going to show that we can define a strategy denoted by $\sigma^t q$ which defines a stopping time τ^t which is a random variable with the same distribution as τ^{t+1} and whose associated buying price q_{σ^t} is always greater than or equal to $q_{\sigma^{t+1}}$, this would establish the claim.

If $V_{t+1}(q) = u(q)$, then define $(\sigma^t q)[\hat{q}_t = q] = S$, and the claim is true. For the rest of the proof, we will assume that $(\sigma^{t+1} q)[\hat{q}_t = q] = C$.

To simplify notations let $F_{t+2}^{t+1}(x) = F_{t+2}(x | \hat{q}_{t+1} = q)$, and let $F_{t+1}^t(x) = F_{t+1}(x | \hat{q}_t = q)$. Let's establish a relationship z^t between the set of realizations of $\hat{q}_{t+1} | \hat{q}_t = q : [q, 1] \cap Q$ and its image by the probability measure: $(0, 1]$. Let's define the correspondence:

If x is not a jump point of F_{t+1}^t , $z^t(x) = F_{t+1}^t(x)$.

If x is a jump point of F_{t+1}^t , $z^t(x) = \vartheta(F_{t+1}^t(x-), F_{t+1}^t(x))$

where ϑ is the uniformly distributed random variable, $f(x-)$ is the left hand limit of the function f at the point x . z^t is a correspondence from $[q, 1] \cap Q$ onto $(0, 1]$. Let's now establish another correspondence r^t from $(0, 1]$ onto the set of realizations of $\hat{q}_{t+2} | \hat{q}_{t+1} = q : [q, 1] \cap Q$. For $z \in (0, 1]$:

¹Actually $\sigma^{t+1} q$ is the restriction of the optimal strategy to periods $t+1, t+2, \dots$ for an history which specifies $\hat{q}_{t+1} = q$. Because optimal strategies are history independent this strategy so restricted does not depend on prices $q_1, \dots, q_t$, and we can disregard entirely these prices. $\sigma^t q$ must be interpreted in the same way.

If $\exists y \in [q, 1] \cap Q$, such that $F_{t+2}^{t+1}(y) = z$,

$$r^t(z) = \text{Min}\{y \in [q, 1] \cap Q / F_{t+2}^{t+1}(y) = z\}$$

If $\nexists y' \in [q, 1] \cap Q$, $F_{t+2}^{t+1}(y') = z$, then $\exists y \in [q, 1] \cap Q$, $F_{t+2}^{t+1}(y) > z \geq F_{t+2}^{t+1}(y-)$, and:

$$r^t(z) = y$$

r^t is onto the support of $\hat{q}_{t+2}$. Now define the correspondence $\sigma^t = r^t \circ z^t$; it is a

correspondence from the set of realizations of $\hat{q}_{t+1}$ | $\hat{q}_t = q$ onto the the set of realizations of $\hat{q}_{t+2}$ | $\hat{q}_{t+1} = q$.

By construction, $\Pr\{\sigma^t(\hat{q}_{t+1}) | \hat{q}_t = q \leq z\} = F_{t+2}^{t+1}(z)$, therefore $\sigma^t(\hat{q}_{t+1}) | \hat{q}_t = q$

has the same distribution as $\hat{q}_{t+2}$ | $\hat{q}_{t+1} = q$. Furthermore $\sigma^t(\hat{q}_{t+1}) \leq \hat{q}_{t+1}$: Let x be the

realization of $\hat{q}_{t+1}$, and y be the realization of $\sigma^t(x)$. Suppose $y > x$. If y is not a jump point for F_{t+2}^{t+1} , then $F_{t+2}^{t+1}(x) < F_{t+2}^{t+1}(y) = z^t(x) \leq F_{t+1}^t(x)$ which contradicts assumption (A-2). If y

is a jump point for F_{t+2}^{t+1} , then $F_{t+2}^{t+1}(x) \leq F_{t+2}^{t+1}(y-) \leq z^t(x) \leq F_{t+1}^t(x)$. So the only solution is

that $F_{t+2}^{t+1}(x) = F_{t+2}^{t+1}(y-) = z^t(x) = F_{t+1}^t(x) < F_{t+2}^{t+1}(y)$, and $(x, y) \notin Q$. But then

$z^t(x) = F_{t+2}^{t+1}(x) \Rightarrow \sigma^t(x) = x$ which contradicts $\sigma^t(x) = y$.

Let's define the strategy $\sigma^t q$ for period $t + 1$ as: $(\sigma^t q)(\hat{q}_{t+1}) = S$ if and only if

$(\sigma^{t+1} q)[\sigma^t(\hat{q}_{t+1})] = S$. From what precedes it is clear that $\sigma^t q$ will stop in period $t + 1$ with the same probability that $\sigma^{t+1} q$ stops in period $t + 2$, and that it only stops at a higher price.

Hence when $\sigma^t q$ stops in period $t + 1$ it contributes a higher value to the search than its counterpart for strategy $\sigma^{t+1} q$.

For histories h^t such that $(\sigma^t q)(h^t)$ at time $t + 1$ is continue, we shall define

similarly a correspondence σ^{t+1} between the set of realizations of

$\hat{q}_{t+2}$ | $\{\sigma^t(\hat{q}_{t+1}) = q_1, \hat{q}_t = q\}$, and $\hat{q}_{t+3}$ | $\{\hat{q}_{t+2} = q_1, \hat{q}_{t+1} = q\}$. Since $\sigma^t(\hat{q}_{t+1}) \leq \hat{q}_{t+1}$,



$\Pr\{\hat{q}_{t+2} \leq a | \sigma^t(\hat{q}_{t+1}) = q_1, \hat{q}_t = q\} \leq \Pr\{\hat{q}_{t+2} \leq a | \hat{q}_{t+1} = q_1, \hat{q}_t = q\}$ which implies as

before that $\sigma^{t+1}(\hat{q}_{t+2}) \leq \hat{q}_{t+2}$. The result is then established by forward induction. Q. E. D. •


Representation of the construction of $\sigma^t q$ for the case of a continuous distribution.

The first diagram shows how to associate to a realization x of $\hat{q}_{t+1}$ a value y for $\sigma^t(\hat{q}_{t+1})$; the consumer does for $\hat{q}_{t+1}$ what he would have done under σ^{t+1} when $\hat{q}_{t+2} = y$. The second

diagram shows what happens when the consumer decides to continue under σ^{t+1} . The distribution for $\hat{q}_{t+2} | \hat{q}_{t+1} = x, \hat{q}_t = q$ is denoted by F_{t+2} , whereas the distribution for $\hat{q}_{t+3} | \hat{q}_{t+2} = y, \hat{q}_{t+1} = q$ is denoted by F_{t+3} .



The next diagram presents the idea for a discrete distribution; it illustrates how the strategy σ^q can become a random rule for jump point of F_{t+1}^t : For $\hat{q}_{t+1} = q_2$, $\sigma^t(\hat{q}_{t+1}) = q_1$ with 20% probability, $\sigma^t(\hat{q}_{t+1}) = q_2$ with 80% probability.



Theorem 3: Under assumptions (A) or (A'), if a price $q > 0$ is observed, then there exists a time horizon $T(q)$ past which an optimal strategy will not proceed.

Proof: At time t , if the price $\hat{q}_t = q$ has been observed an optimal strategy will stop if:

$$V_t(q) \geq \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q]$$

$$\begin{aligned} \text{We have}^1: \quad E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q] &= \int_0^1 V_{t+1}(x) dF_{t+1}(x | \hat{q}_t = q) \\ &= V_{t+1}(q) F_{t+1}(q | \hat{q}_t = q) + \int_q^1 V_{t+1}(x) dF_{t+1}(x | \hat{q}_t = q) \\ &\leq V_t(q) F_{t+1}(q | \hat{q}_t = q) + u(1) [1 - F_{t+1}(q | \hat{q}_t = q)] \end{aligned}$$

the last inequality being true because of lemma 5, and because $u(1)$ is an upper bound for the value of the search. Hence:

$$V_t(q) - \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q] \geq [1 - \delta F_{t+1}(q | \hat{q}_t = q)] V_t(q) - [1 - F_{t+1}(q | \hat{q}_t = q)] \delta u(1)$$

¹This proof is presented for the continuous case. For the discrete case, the integral is replaced by a finite sum.

$$\geq (1 - \delta)u(q) - [1 - F_{t+1}(q | \hat{q}_t = q)]\delta u(1)$$

because an optimal strategy will always verify $V_t(q) \geq u(q)$ and $F_{t+1}(q | \hat{q}_t = q) \leq 1$. Take $\varepsilon = \frac{(1 - \delta)u(q)}{2\delta u(1)} > 0$, and take $T(q)$ as the time horizon specified in assumption (A-4) for this ε . Then:

$$\forall t \geq T(q), V_t(q) - \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q] \geq \frac{(1 - \delta)u(q)}{2} > 0$$

and $\forall q' \geq q, \forall t \geq T(q)$,

$$V_t(q') - \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q'] \geq (1 - \delta)[u(q') - \frac{u(q)}{2}] > \frac{(1 - \delta)u(q')}{2} > 0$$

So for all $t \geq T(q)$ and all $q' \in [q, 1]$, $V_t(q') > \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q']$, so a consumer who uses an optimal strategy, and has observed $\hat{q}_t = q'$ will stop. Q. E. D. •

6.4 - Construction of the Optimal Regular Strategy:

For the case of a distribution function with a continuous support, consider the function:

$$f_t(q) = u(q)F_t(q | \hat{q}_{t-1} = q) + \int_q^1 u(x)dF_t(x | \hat{q}_{t-1} = q)$$

For the case of a discrete support, define the function $f_t(q)$ for values of q which belong to the support of the distribution function as:

$$f_t(q^i) = u(q^i)F_t(q^i | \hat{q}_{t-1} = q^i) + \sum_{i' > i} u(q^{i'})Pr(\hat{q}_t = q^{i'} | \hat{q}_{t-1} = q^i)$$

Lemma 6: Under assumptions (A) or (A'), the benchmarks ξ_t form a decreasing sequence of t , and tend to 0 when t tends to infinity.

Proof: The benchmark ξ_t is the solution the point such that:

$$\forall q \geq \xi_t, u(q) \geq \delta f_t(q), \text{ and } \forall q < \xi_t, u(q) < \delta f_t(q)$$

Suppose that $f_{t+1}(q) \leq f_t(q)$, then $\delta f_{t+1}(\xi_t) \leq \delta f_t(\xi_t) \leq u(\xi_t)$. Hence

$\xi_t \in \{q \in Q / \delta f_{t+1}(q) \leq u(q)\}$, so $\xi_{t+1} \leq \xi_t$. Furthermore the intermediate value theorem implies that:

$$\exists c_{q,t} \in [q, 1], f_t(q) = u(q)F_t(q | \hat{q}_{t-1} = q) + u(c_{q,t})[1 - F_t(q | \hat{q}_{t-1} = q)]$$

Since $F_t(q|\hat{q}_{t-1} = q) \rightarrow 1-$ when $t \rightarrow \infty$, therefore $f_t(q) \rightarrow u(q)-$ when $t \rightarrow \infty$. Therefore for $\delta > 0$, $\forall q > 0$, $\lim_{t \rightarrow \infty} \delta f_t(q) < u(q)$, and $\lim_{t \rightarrow \infty} \delta f_t(0) = u(0) = 0$. Hence the lemma would be proved.

For the continuous case:

$$\begin{aligned} f_{t+1}(q) - f_t(q) &= u(q)[F_{t+1}(q|\hat{q}_t = q) - F_t(q|\hat{q}_{t-1} = q)] \\ &+ \int_q^1 u(x)[dF_{t+1}(x|\hat{q}_t = q) - dF_t(x|\hat{q}_{t+1} = q)] \\ &= \int_q^1 [u(x) - u(q)][dF_{t+1}(x|\hat{q}_t = q) - dF_t(x|\hat{q}_{t+1} = q)] \\ &\leq 0 \text{ by assumption (A'-3).} \end{aligned}$$

For the discrete case, the integral is replaced by a finite sum, and assumption (A-3) gives the result. Q. E. D. •

Lemma 7: Under assumptions (A) or (A'), the functions $u(q)$ and $\delta f_t(q)$ are increasing functions of q , and $\exists \xi_t \in [0, 1]!$, $\forall x \geq \xi_t$, $x \in Q$, $u(x) \geq \delta f_t(x)$, $\forall x' < \xi_t$, $x' \in Q$, $u(x') < \delta f_t(x')$.

Proof: For the continuous case. The function $f_t(q)$ is continuous on $[0, 1]$. It is differentiable only on $[0, 1] \setminus J$, where J is the finite set of jump points of Q . Let q be such a point, $q^i < q < q^{i+1}$:

$$f'_t(q) = u'(q)F_t(q|\hat{q}_{t-1} = q) + \int_q^1 u(x)dF_{t;q}(x|\hat{q}_{t-1} = q) - u(q)F_{t;x}(q+|\hat{q}_{t-1} = q)$$

where $F_{t;q}$ is the derivative of $F_t(x|\hat{q}_{t-1} = q)$ with respect to q and $F_{t;x}(x+|.)$ is the right-hand derivative of $F_t(x|\hat{q}_{t-1} = q)$ with respect to x . We have:

$$\begin{aligned} F_{t;q}(x|\hat{q}_{t-1} = q) &= -\frac{q^{i+1} - x}{(t + \sum_{i=0}^{i=N} \alpha_i)(q^{i+1} - q)^2} && \text{for } q < x < q^{i+1} \\ &= 0 && \text{for } q^{i+1} \leq x \end{aligned}$$

$$F_{t;x}(q+|\hat{q}_{t-1} = q) = \frac{1}{(q^{i+1} - q)(t + \sum_{i=0}^{i=N} \alpha_i)}$$

$$\begin{aligned}
\text{Hence } f'_t(q) &= \frac{1}{t + \sum_{i=0}^N \alpha_i} \left\{ (t-1 + \alpha_i)u'(q) + \frac{1}{(q^{i+1} - q)^2} \int_q^{q^{i+1}} u(x)dx - \frac{1}{q^{i+1} - q} u(q) \right\} \\
&= \frac{1}{t + \sum_{i=0}^N \alpha_i} \left\{ (t-1 + \alpha_i)u'(q) + \frac{1}{(q^{i+1} - q)^2} \int_q^{q^{i+1}} [u(x) - u(q)]dx \right\}
\end{aligned}$$

Therefore $f'_t(q)$ is always strictly positive, and since $f_t(q)$ is continuous, it is a strictly increasing function of q . Also because $u(q)$ is concave:

$$\begin{aligned}
f'_t(q) &\leq \frac{1}{t + \sum_{i=0}^N \alpha_i} \left\{ (t-1 + \alpha_i)u'(q) + \frac{u'(q)}{(q^{i+1} - q)^2} \int_q^{q^{i+1}} [x - q]dx \right\} \\
&= \frac{1}{t + \sum_{i=0}^N \alpha_i} \left\{ (t-1 + \alpha_i)u'(q) + \frac{u'(q)}{2} \right\} \\
&< u'(q)
\end{aligned}$$

Hence since both $f_t(q)$ and $u(q)$ are continuous and strictly increasing functions, and $f_t(\cdot)$ always increases less than $u(\cdot)$, therefore the functions $\delta f_t(\cdot)$ and $u(\cdot)$ intersect at most once. In addition, we have:

$$\forall t > 0, \forall \delta \in [0, 1], \quad \delta f_t(0) = \int_0^1 \delta u(x) dF_t(x | \hat{q}_{t-1} = 0) dx > 0 = u(0)^1,$$

$$\text{and } u(1) \geq \delta f_t(1) = \delta u(1)$$

Therefore by the intermediate value theorem, there exists a ξ_t such that $u(\xi_t) = \delta f_t(\xi_t)$, and since $\delta f'_t(q) < u'(q)$, the result is established.

For the discrete case. The function $f_t(q^i)$ can be written as:

$$f_t(q^i) = u(q^i)F_t(q^i | \hat{q}_{t-1} = q^i) + u(q^{i+1})\Pr(q^{i+1} | \hat{q}_{t-1} = q^i) + \sum_{i' > i+1} u(q^{i'})\Pr(\hat{q}_t = q^{i'} | \hat{q}_{t-1} = q^i)$$

$$\text{also } f_t(q^{i+1}) = u(q^{i+1})F_t(q^{i+1} | \hat{q}_{t-1} = q^{i+1}) + \sum_{i' > i+1} u(q^{i'})\Pr(\hat{q}_t = q^{i'} | \hat{q}_{t-1} = q^{i+1})$$

$$\text{But for } i' > i+1, \Pr(\hat{q}_t = q^{i'} | \hat{q}_{t-1} = q^{i+1}) = \Pr(\hat{q}_t = q^{i'} | \hat{q}_{t-1} = q^i)$$

$$\text{and } F_t(q^{i+1} | \hat{q}_{t-1} = q^i) = F_t(q^{i+1} | \hat{q}_{t-1} = q^{i+1})$$

because of assumption (A-1); also:

$$F_t(q^{i+1} | \hat{q}_{t-1} = q^i) = F_t(q^i | \hat{q}_{t-1} = q^i) + \Pr(q^{i+1} | \hat{q}_{t-1} = q^i)$$

¹This is a consequence of assumption (A-4).

Therefore $f_t(q^{i+1}) - f_t(q^i) = [u(q^{i+1}) - u(q^i)]\Pr(q^{i+1} | \hat{q}_{t-1} = q^i)$

and $0 < f_t(q^{i+1}) - f_t(q^i) < u(q^{i+1}) - u(q^i)$

So $f_t(q)$ is an increasing function with a smaller rate of increase on its support set than $u(q)$; $f_t(q)$ can then be extended to the whole interval $[0, 1]$ so as to apply the same argument as for the continuous case. Then ξ_t is defined as the element of Q which is the minimum of the set $\{q \in Q, \delta f_t(q) \leq u(q)\}$. Q. E. D. •

Theorem 4: Under the assumptions (A) or (A'), the following strategy σ^* is an optimal strategy which is history independent and regular:

For $t > 0$, let $h_t = \{q_1, \dots, q_t\}$, and let $\hat{q}_t = \sup_{\tau \in \{1, \dots, t\}} \{q_\tau\}$.

Then: $\sigma^*(h_t) = S \Leftrightarrow \hat{q}_t \geq \xi_t$,

and $\sigma^*(h_t) = C \Leftrightarrow \hat{q}_t < \xi_t$.

Proof: From theorem 2, we know that the search problem admits an optimal strategy which is regular. This strategy σ must satisfy the principal of optimality: Given a history h_t with an order statistic $\hat{q}_t = q$:

$$V_t(\hat{q}_t) = \text{Max}\{u(q), \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q]\}$$

The expected value (computed at time t) of continuing search in period $t + 1$ and beyond is¹:

$$E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q] = V_{t+1}(q)F_{t+1}(q | \hat{q}_t = q) + \int_q^1 V_{t+1}(x) dF_{t+1}(x | \hat{q}_t = q)$$

Let's suppose first that the consumer was searching with a finite horizon T . Let σ^T be optimal regular strategy given the limited horizon T . It is clear that in period $T - 1$, the value of continuing search is given by the equation above in which the function $V_T(\cdot)$ is replaced by the function $u(\cdot)$ since the consumer stops anyway at T . Hence in period $T - 1$, $\sigma^T(\hat{q}_{T-1} = q)$ must select the $\max\{u(q), \delta f_{T-1}(q)\}$, i. e. $\sigma^T(\hat{q}_{T-1}) = S$ if and only if $\hat{q}_{T-1} \geq \xi_{T-1}$, $\sigma^T(\hat{q}_{T-1}) = C$ otherwise. We are going to use backward induction to establish the result (still under finite horizon T) for all periods.

¹This proof is presented for the continuous case. For the discrete case, the integral is replaced by a finite sum.

Let' s suppose that for $\tau = t + 1, \dots, T$, $\sigma^T(\hat{q}_\tau) = S$ if and only if $\hat{q}_\tau \geq \xi_\tau$, $\sigma^T(\hat{q}_\tau) = C$ otherwise. At time t , for any history h_t with an order statistic $\hat{q}_t = q$, we have:

$$V_t(\hat{q}_t) = \text{Max}\{u(q), \delta E[V_{t+1}(\hat{q}_{t+1}) | \hat{q}_t = q]\}$$

Suppose $q < \xi_t$. Then any strategy σ_S^T which includes the statement: $\sigma_S^T(\hat{q}_t = q) = S$ is, by definition of ξ_t , dominated by the strategy σ_C^T which is defined as:

$$\forall h_t \text{ such that } \hat{q}_t = q, \sigma_C^T(\hat{q}_t = q) = C \text{ exactly one more time,}$$

$$\forall h_t \text{ such that } \hat{q}_t \neq q, \sigma_C^T(h_t) = \sigma_S^T(h_t).$$

Hence an optimal strategy at time t must include the statement: if h_t is such that $\hat{q}_t < \xi_t$ then $\sigma^T(h_t) = C$.

Suppose $q \geq \xi_t$, then lemma 6 implies that $q \geq \xi_{t+1}$, therefore by the induction hypothesis: $\forall x \geq q, V_{t+1}(x) = u(x)$, and $V_t(\hat{q}_t) = \text{Max}\{u(q), \delta f_t(q)\}$. Hence if $\hat{q}_t \geq \xi_t$ then $\sigma^T(h_t) = S$, and the backward induction has been carried one step back. Notice that the benchmarks ξ_t ' s used to construct the strategies under finite search horizon do not depend on the horizon chosen. We are left to remove the time horizon constraint. We will proceed to do this with forward induction.

Let' s suppose that in period $t = 1$ the consumer observes a price $q_1 > 0$. Then by theorem 3 we know that the search will not proceed past an horizon $T(q_1)$. It is also clear that $u(q_1) \geq \delta f_{T(q_1)}(q_1)$, otherwise the consumer would have opted to continue at least one more time at time $T(q_1)$ if $\hat{q}_{T(q_1)} = q_1$, which would have contradicted the time horizon $T(q_1)$. Hence the consumer could use the strategy $\sigma^{T(q_1)}$. Notice that the strategy σ^* would have lead to the same outcome, given the initial observation of q_1 . Let' s suppose that the consumer observes the price $q_1 = 0$ in period $t = 0$. If the consumer stops, his pay-off is $u(0) = 0$. This pay-off is less than the discounted expected pay-off of a strategy which directs him to continue exactly one more time, hence a strategy which includes the statement $\sigma(q_1 = 0) = S$ cannot be optimal. So for an optimal strategy $\sigma(q_1 = 0) = C$; notice that σ^* verifies this condition.

Let' s suppose that we have shown that the optimal strategy of the consumer is defined for all history up to period $\tau, \tau \leq t$ as:

- if $\forall \tau' < \tau, \hat{q}_{\tau'} = 0$, and $\hat{q}_{\tau} > 0$, then adopt the strategy σ^*
- if $\hat{q}_t = 0$, then $\sigma(\hat{q}_t = 0) = C$

Then if $\hat{q}_{t+1} > 0$, there exists a finite horizon for the search i. e. $\min\{t + 1, T(\hat{q}_{t+1})\}$ and using the strategy σ^* yields the same outcome as the use of such a finite horizon strategy; if $\hat{q}_{t+1} = 0$, then it is clear that we must have for an optimal strategy $\sigma(\hat{q}_{t+1} = 0) = C$, hence in both cases the use of the strategy σ^* is equivalent to the use of an optimal strategy, hence the forward induction is established, and the strategy σ^* is proved to be optimal. Q. E. D. •

Section 7: Conclusion.

This paper has presented explicit solutions to a range of approaches to search with recall. An interesting feature of this solution is that it can be described as ‘choose the myopic rule’ (by the way, this was also a way to characterize the solution of search with a known distribution). There is something very satisfying about such a result. The search problem is complex, and one could have thought that the optimal solution would require the consumer to take into account the consequences of action far removed in time. If this had been the case, one would hardly have expected to observe such a behavior in practice, and one would have concluded that if, in theory, economic agents optimize, in reality they do not, and that economic theory is a hopeless representation of actual economic activity. As it turns out, the optimal rule sounds like a rule of thumb, a rule that economic agents could have learned through experience, and apply without realizing its optimality properties. This would have been a weak point if we had found a myopic rule hopelessly difficult to compute. But at least for the case of Dirichlet priors, or for the case of the Maximum Entropy estimator we described, this rule does not require much pondering.

An other appealing aspect of these models is related with their management of information. As search is sequential, the past history of search is described by an order sequence of bits of information. Because the price quotes are assumed to be i. i. d. the ordering can be discarded. This still requires the consumer to possibly store quite a few numbers. However for the two cases we mentioned the consumer only needs to keep a record of the highest price and the number of stores searched saving him memory storage space, which is again a consideration that could hardly be overlooked in practice.

The simplicity of the behavior of the consumer, and the fact that they are impervious to the type of blackmail outlined by Diamond could prove of considerable help in designing new models of search and matching, and of equilibrium price dispersion.

References

- [1] Albrecht, J. and Axell, B. (1984), "An Equilibrium Model of Search Unemployment," *Journal of Political Economy*, **92** (5), 825-839.
- [2] ——— and Jovanovic, B. (1986), "The Efficiency of Search under Competition and Monopsony," *Journal of Political Economy*, **94** (6), 1246-1257.
- [3] Blau, D. (1988), "Search for Nonwage Job Characteristics: A Test of the Reservation Wage Hypothesis," Mimeo, University of North Carolina.
- [4] Burdett, K. and Judd, K. L. (1983), "Equilibrium Price Dispersion," *Econometrica*, **51** (4), 955-969.
- [5] DeGroot, M. (1970), *Optimal Statistical Decisions*. New York: McGraw-Hill.
- [6] Gihman, I.I. and Skorohod (1974), *The Theory of Stochastic Processes I*. Berlin: Springer-Verlag.
- [7] Jaynes, E. T. (1983), "Where Do We Stand on Maximum Entropy," in *Papers on Probability, Statistics and Statistical Physics*, Rosenkrantz R.D. Editor, 210-314. Dordrecht: Reidel.
- [8] Jovanovic, B. (1979), "Job Matching and the Theory of Turnover," *Journal of Political Economy*, **87** (5), 972-990.
- [9] Kagan, A. M., Linnik, Y. V. and Rao, C. R. (1973), *Characterization of Problems in Mathematics Statistics*, New York: John Wiley & Sons.
- [10] Lippman, S. A. and McCall, J. J. (1981), "The Economics of Uncertainty," in *Handbook of Mathematical Economics, Vol I*, Arrow, K. J. and Intriligator, M. D. Editors. Chapter 6, Section 2, 217-227. Amsterdam: North Holland.
- [11] McAfee, R. P. and McMillan, J. (1988), "Search Mechanisms," *Journal of Economic Theory*, **44**, 99-123.
- [12] Morgan, P. and Manning, R. (1985), "Optimal Search," *Econometrica*, **53** (4), 923-944.
- [13] Mortensen, D. T. (1984), "Job Search and Labor Market Analysis," in *Handbook of Labor Economics*, Amsterdam: North Holland.

- [14] Pollard, D. (1984), *Convergence of Stochastic Processes*, Berlin: Springer-Verlag.
- [15] Rosenfield, D. B. and Shapiro, R. D. (1981), "Optimal Price Search with Bayesian Extensions," *Journal of Economic Theory*, **25**, 1-20.
- [16] Rothschild, M. (1974), "Searching for the Lowest Price when the Distribution of Prices is Unknown," *Journal of Political Economy*, **82**, 689-711.
- [17] Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, **50** (1), 97-109.
- [18] ——— (1985), "A Bargaining Model with Incomplete Information about Time Preferences," *Econometrica*, **53**, 1151-1172.
- [19] ——— and Wolinsky, A. (1985), "Equilibrium in a market with Sequential Bargaining," *Econometrica*, **53** (5), 1133-1150.
- [20] Sargent, T. J. (1987), *Dynamic Economic Theory*, Cambridge MA: Harvard University Press.
- [21] Savage, L. J. (1954), *The Foundations of Statistics*, New York: John Wiley & Sons.
- [22] Selten, R. (1975), "Re-examination of the Perfectness Concept for equilibrium Points in Extensive Games," *International Journal of Game Theory*, **4**, 25-55.
- [23] Shannon, C. E. (1948), "A Mathematical Theory of Communication," *Bell System Technical Journal*, **27**, 631-635.
- [24] Shore J. E. and Johnson R. W. (1980), "Axiomatic Derivation of the Principle of Maximum Entropy and the Principle of Minimum Cross-Entropy," *IEEE Transactions on Information Theory*, **26** (1), 26-37.
- [25] Skilling, J. (1986), "The Axioms of Maximum Entropy," Mimeo, Cambridge University.
- [26] Theil, H. (1967), *Economics and Information Theory*. Amsterdam: North-Holland.
- [27] ——— and Fiebig, D.G. (1984), *Exploiting Continuity*. Cambridge MA: Balinger.
- [28] Vishwanath, T. (1984), "Reservation Wages and Quit Rates from Unemployment with Unknown Offer Distributions," Mimeo, Cornell University.
- [29] Wolinsky, A. (1987), "Matching, Search and Bargaining," *Journal of Economic Theory*, **42** (2), 311-333.