Non-Bayesian Optimal Search and Dynamic Implementation

Alex Gershkov and Benny Moldovanu *

January 21, 2010

Abstract

We show that two non-Bayesian learning procedures lead to very permissive implementation results concerning the efficient allocation of resources in a dynamic environment where impatient, privately informed agents arrive over time, and where the designer gradually learns about the distribution of agents' values. This contrasts the rather restrictive results that have been obtained for Bayesian learning in the same environment.

1 Introduction

We analyze the implementation of the efficient dynamic policy in a model where impatient, privately informed agents arrive over time, and where the designer gradually learns about the distribution of agents' values using non-Bayesian updating procedures. We show that some simple but not trivial non-Bayesian updating procedures that were used in the classical literature lead to very permissive implementation results, contrasting the rather restrictive results that have been obtained for Bayesian learning in the same mechanism design environment. This highlights the role of the learning procedure in dynamic mechanism design problems.

*We wish to thank Philippe Jehiel for helpful remarks. We are grateful to the German Science Foundation for financial support. Moldovanu: Department of Economics, University of Bonn, mold@uni-bonn.de; Gershkov: Department of Economics, Hebrew University of Jerusalem, alexg@huji.ac.il

The allocation model studied here is based on a classical model due to Derman, Lieberman and Ross [5] (DLR hereafter). In the DLR model, a finite set of heterogenous, commonly ranked objects needs to be assigned to a set of agents who arrive one at a time. After each arrival, the designer decides which object (if any) to assign to the present agent[1]. Both the attribute of the present agent (that determines his value for the various available objects) and the future distribution of attributes are known to the designer in the DLR analysis.

Learning about future values in the complete-information DLR model has been analyzed by Albright [1]. Gershkov and Moldovanu [6] added incomplete information to Albright's model, and derived an implicit condition on the structure of the allocation policy ensuring that efficient implementation is possible. Roughly speaking, implementation is possible if the impact of currently revealed information on today's values is higher than the impact on option values. This insight replaces in the dynamic framework with learning the *single-crossing* idea appearing in the theory of static efficient implementation with interdependent values[2].

Gershkov and Moldovanu [7] focused on Bayesian learning and offered conditions on the exogenous parameters of the model - the prior beliefs about the agents' values- that allow efficient implementation. Since these results are relatively restrictive, they also characterized the incentive-efficient, second-best mechanism.

In the present paper we study two adaptive, non-Bayesian learning processes that have been used for the classical, complete-information one-object search framework (see Rothschild [11]) by Bickchandani and Sharma [2], and by Chou and Talmain [3], respectively[3]. The first learning process constructs a posterior that is a convex combination of a prior and the empirical distribution, with more and more weight given to the empirical distribution. The second process starts with a maximum entropy prior and constructs a quantile preserving posterior based on the observations made so far.

For both processes, we prove that the efficient allocation is always im-

[1] In a framework with several homogenous objects the decision is simply whether to assign an object or not.

[2] See for example Dasgupta and Maskin [4] and Jehiel and Moldovanu [9] who analyzed static models with direct informational externalities.

[3] Both processes are consistent in the sense that they uniformly converge to the true distribution as the number of observations goes to infinity. In both cases, this is a consequence of the well known Glivenko-Cantelli Theorem.

plementable since new information is incorporated in option values at a slow rate, so that the impact of new information on present values is always higher. As in the case of standard Bayesian learning, the efficient allocation maximizes at each decision period the sum of the expected utilities of all agents, given all the available information. The only difference to the Bayesian approach is in the inference made from new information.

A word of caution is needed here: Our results do not imply that the considered non-Bayesian procedures are "better" than Bayesian updating for the purposes of efficient implementation! They just say that the complete information efficient allocation - whose calculation proceeds **given** the assumed learning procedure - can always be implemented for the particular adaptive processes studied here. An example below will illustrate this issue.

In a one-object search model with complete information, Rothschild focused on *reservation price search policies*, i.e., policies where for each information state s there exists a price $R(s)$ such that prices above are rejected and prices below are accepted. Our implementation results hinge on a monotonicity property that generalize such reservation prices for settings with several heterogenous objects. Rothschild showed that the optimal Bayesian search rule need not generally have this property[4]. He also computed an example where the reservation property holds: the searcher obtains price quotations from a multinomial distribution with a parameter that follows a Dirichlet distribution[5]. Interestingly, it turns out that in this special case, the Bayesian learning process coincides in fact with one of the non-Bayesian procedures analyzed here.

The paper is organized as follows: In Section 2 we present the dynamic allocation and learning model. In Section 3 we first recall two results: 1. The characterization of the efficient allocation policy under complete information due to Albright [1]; 2. An implicit condition on the structure of the efficient policy ensuring that this policy can be implemented also under incomplete information, due to Gershkov & Moldovanu [6]. In Section 4 we focus on the two non-Bayesian learning models. Theorems 2 and 3 show that, given these learning models, the implicit condition is always satisfied, and hence the corresponding efficient allocation policy is always implementable. Section

[4]See Rosenfield and Shapiro [10], Seierstad [12] for conditions where optimal search in the Rothschild model displays the reservation price property.

[5]The Dirichlet is the conjugate prior of the multinomial distribution, so the posterior is also Dirichlet in this case.

5 concludes. All proofs are relegated to an Appendix

2 The Model

There are m items and n agents. Each item i is characterized by a "quality" q_i, and each agent j is characterized by a "type" x_j. If an item with quality $q_i \geq 0$ is assigned to an agent with type x_j and this agent is asked to pay p, then this agent enjoys a utility given by $q_i x_j - p$. Getting no item generates utility of zero**.** The goal is to find an assignment that maximizes total welfare.

Agents arrive sequentially, one agent per period of time, and each agent can transact (in both physical and monetary terms) only upon arrival.

Note that in a static problem, total welfare is maximized by assigning the item with the highest quality to the agent with the highest type, the item with the second highest quality to the agent with the second highest type, and so on (*assortative matching).*

Let period n denote the first period, period $n-1$ denote the second period, ..., period 1 denote the last period. If $m > n$ we can obviously discard the $m - n$ worst items without welfare loss. If $m < n$ we can add "dummy" objects with $q_i = 0$. Thus, we can assume without loss of generality that $m = n$.

While the items' properties $0 \leq q_1 \leq q_2 ... \leq q_m$ are assumed to be known, the agents' types are assumed to be independent and identically distributed random variables X_i on $[0, +\infty)$ with common cumulative distribution function F.

The function F is not known to the designer nor to the agents. At the beginning of the allocation process the designer has a prior Φ_n over possible distribution functions, and he updates his beliefs after each additional observation. Denote by $\Phi_k\left(x_n, ..., x_{k+1}\right)$ the designer's beliefs about the distribution function F after observing types $x_n, ..., x_{k+1}$. Given such beliefs, let $\widetilde{F}_k(x|x_n, ..., x_{k+1})$ denote the distribution of the next type x_k, conditional on observing $x_n, ..., x_{k+1}$. Finally, we assume that each agent, upon arrival observes the whole history of the previous play.

3 The Dynamic Efficient Allocation

Albright [1] derived the efficient dynamic policy under complete information, i.e., when the agent's type is revealed to the designer upon the agent's arrival. The efficient allocation maximizes, at each decision period, the sum of the expected utilities of all agents, given all the information available at that period, and is defined in terms of cutoffs. Gershkov and Moldovanu [6] displayed an implicit sufficient condition on these cutoffs ensuring that the efficient allocation is implementable also under incomplete information. These observations are gathered in the next Theorem.

Let the history at period k, H_k, be the ordered set of all signals reported by the agents that arrived at periods $n, ..., k+1$, and of allocations to those agents. Let $\mathcal{H}_k$ be the set of all histories at period k. Denote by χ_k the ordered set of signals reported by the agents that arrived at periods $n, ..., k+1$.

Theorem 1 *1. (Albright, 1977) Assume that types $x_n, .., x_{k+1}$ have been observed, and consider the arrival of an agent with type x_k in period $k \geq 1$. There exist functions $0 = a_{0,k}(\chi_k, x_k) \leq a_{1,k}(\chi_k, x_k) \leq a_{2,k}(\chi_k, x_k) ... \leq a_{k,k}(\chi_k, x_k) = \infty$ such that the efficient dynamic policy - which maximizes the expected value of the total reward - assigns the item with the $i-th$ smallest type if $x_k \in (a_{i-1,k}(\chi_k, x_k),\ a_{i,k}(\chi_k, x_k)]$. The functions $a_{i,k}(\chi_k, x_k)$ do not depend on the q's.*

2. These functions are related to each other by the following recursive formulae:

$$\begin{aligned} a_{i,k+1}(\chi_{k+1}, x_{k+1}) &= \int_{A_{i,k}} x_k d\widetilde{F}_k(x_k|\chi_{k+1}, x_{k+1}) \\ &+ \int_{\underline{A}_{i,k}} a_{i-1,k}(\chi_k, x_k) d\widetilde{F}_k(x_k|\chi_{k+1}, x_{k+1}) \\ &+ \int_{\overline{A}_{i,k}} a_{i,k}(\chi_k, x_k) d\widetilde{F}_k(x_k|\chi_{k+1}, x_{k+1}) \end{aligned} \tag{1}$$

where[6]

$$\begin{aligned} \underline{A}_{i,k} &= \{x_k : x_k \leq a_{i-1,k}(\chi_k, x_k)\} \\ A_{i,k} &= \{x_k : a_{i-1,k}(\chi_k, x_k) < x_k \leq a_{i,k}(\chi_k, x_k)\} \\ \overline{A}_{i,k} &= \{x_k : x_k > a_{i,k}(\chi_k, x_k)\} \end{aligned}$$

[6] We set $+\infty \cdot 0 = -\infty \cdot 0 = 0$.

3. *(Gershkov & Moldovanu, 2009) If for any k, χ_k and for any $i \in \{0, .., k\}$, the cutoff $a_{i,k}(\chi_k, x_k)$ is differentiable with respect to the signal of the agent arriving at k, x_k, and if $\frac{\partial}{\partial x_k} a_{i,k}(\chi_k, x_k) < 1$ for any x_k and χ_k, then, the first-best policy can be implemented also under incomplete information.*

The above policy is the dynamic analogue of the assortative matching policy that is optimal in the static case where all agents arrive simultaneously. The associated cutoffs have a natural interpretation: for each object i and period k the cutoff $a_{i,k}(\chi_k, x_k)$ equals the expected value of the agent's type to which the item with $i - th$ smallest type is assigned in a problem with $k-1$ periods before the period $k-1$ signal is observed.[7] The last point requires the effect of the current information on the current value to be stronger than the effect on the option value. If this is the case, an agent with a higher value obtains a higher quality at each period and for any remaining inventory - a monotonicity property that allows implementation under incomplete information.

4 Non-Bayesian Optimal Search

In this Section we study two adaptive, non-Bayesian learning process that have been analyzed in the classical one-object search framework by Bickchandani and Sharma [2], and by Chou and Talmain [3], respectively. For both processes, we prove that the efficient allocation is always implementable.

4.1 Learning Based on the Empirical Distribution

Assume that before stage n (the first stage), the designer's prior belief about the distribution of the first type x_n is given by a distribution H. Then, conditional on sequentially observing $x_n, x_{n-1}..., x_{k+1}$ at stages $n, n-1, ...k+1$, the designer's belief about the distribution of the next type $x = x_k$ is given by:

$$\widetilde{F}_k(x|x_n, ..., x_{k+1}) = (1-\beta_k^n)H(x) + \beta_k^n \frac{1}{n-k} \sum_{i=k+1}^{n} \mathbf{1}_{[x_i,\infty)}(x), \; k = 1, 2, ...n-1$$

[7] Note also that if $\widetilde{F}_k(x_k|\chi_{k+1}, x_{k+1})$ is symmetric with respect to the observed signals, then $a_{i,k+1}(\chi_{k+1}, x_{k+1})$ is symmetric as well.

where $0 < \beta_k^n < 1$, and where $\mathbf{1}_{[z,\infty)}(x)$ denotes the indicator function of the set $[z,\infty)$. Thus at each stage, the posterior distribution is a convex combination of the prior distribution and of the empirical distribution. Since, by the Glivenko-Cantelli theorem, the empirical distribution uniformly converges to the true underlying distribution, the posterior distribution also converges to the true distribution if the weight on the empirical distribution satisfies: $\forall k$, $\lim_{n\to\infty} \beta_k^n = 1$.

Theorem 2 *Assume that the designer learns based on the empirical distribution. Then, the efficient dynamic policy can always be implemented under incomplete information.*

Proof. See Appendix. ■

For special prior distributions, the process studied above does in fact coincide with the standard Bayesian learning. This is the case, for example, for a multinomial Dirichlet prior or for a Dirichlet process prior. Thus, for such priors, Theorem 2 also asserts the implementability of the efficient dynamic allocation under Bayesian learning. Bickchandani and Sharma [2] showed that the above learning model induces optimal search with the reservation price property in Rothschild's one object model with complete information (where implementation issues do not play any role). Our result can also be interpreted as saying that their insight continues to hold for the case with several heterogenous objects.

4.2 Maximum Entropy/Quantile Preserving Learning

We now assume that designer believes that types distribute continuously on a finite interval, which we normalize here to be the interval $[0,1]$. It is well known that the *maximum entropy* distribution among all continuous distributions with support on an interval $[a,b]$ is the uniform distribution on this interval[8]. More generally, consider a sub-division $a = a_0 < a_1 < ..a_m = b$ and probabilities $p_1, ..p_m$ which add up to one, and consider the class of all continuous distributions supported on $[a,b]$ such that

$$\Pr\{a_{i-1} \le X \le a_i\} = p_i,\ i = 1, .., m$$

[8]Similar exercises can be performed starting with other initial beliefs. For example, the maximum entropy distribution given that a continuous random variable is known to have (normalized) zero mean and unit variance is the standard normal distribution. See Jaynes [8].

Then, the density of the maximum entropy distribution for this class is constant on each of the intervals $[a_{j-1}, a_j)$. Guided by this principle, Chou and Talmain [3] looked at the following *quantile preserving* updating procedure[9]: Prior to any observation, the designer estimates the unknown distribution by the uniform distribution. Suppose that m observations were observed, and order them in increasing order $\{x_{(1)}, ..., x_{(m)}\}$. Let $x_{(0)} = 0$ and $x_{(m+1)} = 1$. Then, the type of the next arrival is estimated according to the density

$$f_k(x|x_n, ..., x_{n-m+1}) = \sum_{i=1}^{m+1} \frac{\mathbf{1}_{[x_{(i-1)}, x_{(i)})}(x)}{(m+1)\left(x_{(i)} - x_{(i-1)}\right)}.$$

In other words, each interval of the form $[x_{(i-1)}, x_{(i)})$ gets assigned a probability $p_i = \frac{1}{m+1}$, and the density within the interval is constant[10]. The rationale behind the equal weights of $\frac{1}{m+1}$ for each interval becomes apparent by recalling that, for m large,

$$E[X_{i,m}] \approx F^{-1}(\frac{i}{m+1}) \text{ and } F(E[X_{i,m}]) - F(E[X_{i-1,m}] \approx \frac{1}{m+1}$$

where the $X_{i,m}$ is the $i - th$ highest order statistic, $i = 1, ..m$, of a random variable X distributed according to distribution F.

Theorem 3 *Assume that the designer uses the maximum entropy/quantile preserving learning procedure. Then the efficient dynamic policy can always be implemented under incomplete information.*

Proof. See Appendix. ■

It is illustrative to compare Bayesian and non-Bayesian learning in a simple example where the dynamically efficient allocation is **not** implementable under Bayesian learning.

Example 4 *1. There are two periods and one indivisible object. Before starting the allocation process, the designer believes that the distribution of values is uniform on the interval* $[0, \frac{1}{2}]$ *with probability* $\frac{1}{2}$, *while with probability* $\frac{1}{2}$ *he believes that it is uniform on* $[\frac{1}{2}, 1]$. *Under Bayesian*

[9] They studied search with recall and did not look at the reservation price property for search without recall.

[10] As above, the Glivenko-Cantelli theorem implies that the above estimated distribution uniformly converges to the true distribution.

learning, the posterior after observing $x_2 < (>)\frac{1}{2}$*, is that* x_1 *is uniformly distributed on* $[0, \frac{1}{2}]$ $([\frac{1}{2}, 1])$*. This yields*

$$a_{12}^{B}(x_2) = \begin{cases} \frac{1}{4} & \textit{if} \quad x_2 < \frac{1}{2} \\ \frac{1}{2} & \textit{if} \quad x_2 = \frac{1}{2} \\ \frac{3}{4} & \textit{if} \quad x_2 > \frac{1}{2} \end{cases}$$

Thus, the first arriving agent should efficiently get the object if $x_2 \in [\frac{1}{4}, \frac{1}{2}] \cup [\frac{3}{4}, 1]$*. This policy is not monotone and* ***cannot*** *be implemented (see GM [6])*

2. *Consider now the learning process based on the empirical distribution with weight* $0 < \beta < 1$ *on the empirical distribution, and with the same prior as above. Then, after having observed* x_2*, the beliefs of the designer are given by* $F(x_1|x_2) = (1-\beta)U([0,1]) + \beta \mathbf{1}_{[x_2,1]}$*, which yields*

$$a_{12}^{ED}(x_2) = \frac{1}{2}(1-\beta) + \beta x_2.$$

   Thus, the first arriving agent should get the object if and only if $x_2 \geq a_{12}^{ED}(x_2) \Leftrightarrow x_2 \geq \frac{1}{2}$*, which can be implemented by a take-it-or-leave-it offer at a price of* $\frac{1}{2}$*. Note however that the implemented allocation differs here from the one that needs to be implemented under Bayesian learning.*

3. *Finally consider the maximum entropy/quantile preserving procedure. Then, after having observed* x_2*, the beliefs of the designer are given by the density*

$$f_k\left(x_1|x_2\right) = \frac{1}{2}\left(\frac{\mathbf{1}_{[0,x_2]}(x)}{x_2} + \frac{\mathbf{1}_{[x_2,1]}(x)}{1-x_2}\right).$$

   which yields

$$a_{12}^{ME}(x_2) = \frac{1}{4} + \frac{1}{2}x_2.$$

   Thus, the first arriving agent should get the object if and only if $x_2 \geq a_{12}^{ME}(x_2) \Leftrightarrow x_2 \geq \frac{1}{2}$*, which can again be implemented by a take-it-or-leave-it offer at a price of* $\frac{1}{2}$[11]*.*

[11]The fact that the efficient policies given the non-Bayesian procedures coincide is a mere coincidence, due to the uniform prior assumed in this example for the learning based on the empirical distribution.

5 Conclusion

We have displayed several non-Bayesian learning models that always allow the dynamic implementation of the corresponding efficient allocation. This is in sharp contrast to the rather restrictive conditions under which efficient dynamic implementation is possible under Bayesian learning. Our results highlight the importance of the learning/updating procedure for dynamic mechanism design models, and point to other fruitful combinations of non-Bayesian statistical methods and strategic dynamic models, e.g., in the area of revenue/yield management.

6 Appendix

For the proof of Theorem 2 we need the following well-known Lemma:

Lemma 5 *Let $u(x)$ be a function on the interval $[a, b]$ such that there exist a division of the interval $a = z_0 < z_1.. < z_n = b$ and values $c_1, ...c_n$ with $u(x) = c_i$ for $z_i < x < z_{i+1}$, $i = 0, 1, ..n - 1$. Then, for any continuous function $v(x)$ on $[a, b]$, it holds that*

$$\int_a^b v(x)du(x) = \sum_{i=0}^{n} v(z_i)(c_{i+1} - c_i)$$

where $\int$ denotes here the Stieltjes integral.

Proof of Theorem 2. By Theorem 1-2, we can write

$$a_{i,k+1}(\chi_{k+1}, x_{k+1}) = E_{x_k|x_{k+1}} G_{i,k}(x_k, x_{k+1}, \chi_{k+1}) \tag{2}$$

where the function $G_{i,k}(x_k, x_{k+1}, \chi_{k+1})$ is given by:

$$\begin{cases} a_{i-1,k}(\chi_{k+1}, x_{k+1}, x_k) & \text{if} \quad x_k \leq a_{i-1,k}(\chi_{k+1}, x_{k+1}, x_k) \\ x_k & \text{if} \quad a_{i-1,k}(\chi_{k+1}, x_{k+1}, x_k) < x_k \leq a_{i,k}(\chi_{k+1}, x_{k+1}, x_k) \\ a_{i,k}(\chi_{k+1}, x_{k+1}, x_k) & \text{if} \quad x_k > a_{i,k}(\chi_{k+1}, x_{k+1}, x_k) \end{cases}. \tag{3}$$

In words, $G_{i,k}(x_k, x_{k+1}, \chi_{k+1})$ is the second-highest order statistic out of the set $\{a_{i-1,k}(\chi_{k+1}, x_{k+1}, x_k),\ x_k\ ,\ a_{i,k}(\chi_{k+1}, x_{k+1}, x_k)\}$. Let $\mathbf{mx} = (x, x, ...x)$ denote an $m-$vector of x. We show by induction that $\forall m,\ m \leq n-k+1$, the function $a_{i,k}\left(x_n,..x_{k+m},\mathbf{mx}\right)$ is continuously differentiable in the observed signals with

$$\forall i,k,\quad \frac{\partial a_{i,k}\left(x_n,..x_{k+m},\mathbf{mx}\right)}{\partial x}<1.$$

Since the conditional distribution $\widetilde{F}_k(x|x_n,...,x_{k+1})$ does not have here a well-defined density, we use below the notion of *Stieltjes integral.* In the last but one period $k=2$, the only relevant, non-trivial cutoff is:

$$\begin{aligned} a_{1,2}(x_n,...,x_2) &= \int_0^\infty x_1 d\widetilde{F}_1\left(x_1|x_n,..,x_2\right) \\ &= (1-\beta_2^n)\int_0^\infty x_1 dH(x_1)+\beta_2^n\int_0^\infty x_1 d\left(\sum_{i=2}^n 1_{[x_i,\infty)}(x_1)\right) \\ &= (1-\beta_2^n)E(H)+\beta_2^n\frac{1}{n-1}\sum_{i=2}^n x_i \end{aligned}$$

The second equality follows by the additivity property of the Stieltjes integral. The third equality follows by Lemma 5 since $\sum_{i=2}^n 1_{[x_i,\infty)}(x)$ is a step function. Thus, as required, we obtain that $a_{1,2}(x_n,...,x_2)$ is continuously differentiable in the observed signals and that

$$\frac{\partial a_{1,2}\left(x_n,..,x_{2+m},\mathbf{mx}\right)}{\partial x}\leq\frac{m\beta_2^n}{n-1}<1,\ m=1,2,...n-1$$

Assume now that the statement holds for all periods up to k (recall that period 1 is the last period, and so on...), and let us look at period $k+1$, and at $m\leq n-k$. Since $a_{i,k}\left(x_n,..x_k\right)$ is continuous, the induction hypothesis implies that for any $i\in\{1,...,k-1\}$ there exists at most one solution to the equation $a_{i,k}\left(x_n,..,x_{k+1},x\right)=x$. Denote this solution by $a_{i,k}^*(x_n,..x_{k+1})$. By the induction hypothesis, and by the Implicit Function Theorem, we obtain that $a_{i,k}^*(x_n,..x_{k+1})$ is continuously differentiable in the observed signals. If $a_{i,k}\left(x_n,..x_{k+1},x\right)>x$ for any x, define $a_{i,k}^*(x_n,..x_{k+1},x)=\infty$, and if $a_{i,k}\left(x_n,..x_{k+1},x\right)<x$ for any x define $a_{i,k}^*(x_n,..x_{k+1},x)=0$. Then we can

write

$$
\begin{aligned}
a_{i,k+1}(x_n,..x_{k+1}) &= \int_0^{a^*_{i-1,k}(\chi_k)} a_{i-1,k}(x_n,..x_{k+1},x_k)d\widetilde{F}_k\left(x_k|x_n,..,x_{k+1}\right) \\
&+ \int_{a^*_{i-1,k}(\chi_k)}^{a^*_{i,k}(\chi_k)} x_k d\widetilde{F}_k\left(x_k|x_n,..,x_{k+1}\right) \\
&+ \int_{a^*_{i,k}(\chi_k)}^{\infty} a_{i,k}(x_n,..x_{k+1},x_k)d\widetilde{F}_k\left(x_k|x_n,..,x_{k+1}\right).
\end{aligned}
$$

By the induction hypothesis we obtain that $a_{i,k+1}(x_n,..x_{k+1})$ is continuously differentiable. We obtain moreover that:

$$
\begin{aligned}
& a_{i,k+1}\left(x_n,..x_{k+m+1},\mathbf{mx}\right) \\
= & \int_0^{\infty} G_{i,k}(x_n,..x_{k+1+m},\mathbf{mx},x_k)d\widetilde{F}_k(x_k|(x_n,..x_{k+1+m},\mathbf{mx})) \\
= & (1-\beta^n_{k+1})\int_0^{\infty} G_{i,k}(x_n,..x_{k+1+m},\mathbf{mx},x_k)dH(x_k) \\
& +\frac{m\beta^n_{k+1}}{n-k}\left[G_{i,k}(x_n,..x_{k+m},(\mathbf{m+1})\mathbf{x})\right] \\
& +\frac{\beta^n_{k+1}}{n-k}\sum_{j=k+m}^{n} G_{i,k}(x_n,..x_{k+m},\mathbf{mx},x_j)
\end{aligned}
$$

where the second equality follows from Lemma 5. Hence, for any $m \leq n-k$, we obtain that:

$$
\begin{aligned}
& \frac{\partial a_{i,k+1}\left((x_n,..x_{k+m+1},\mathbf{mx}\right)}{\partial x} \\
= & (1-\beta^n_{k+1})\int_0^{\infty}\frac{\partial G_{i,k}(x_n,..x_{n-k-m},\mathbf{mx},x_k)}{\partial x}dH(x_k) \\
& +\frac{m\beta^n_{k+1}}{n-k}\frac{\partial G_{i,k}(x_n,..x_{k+1+m},(\mathbf{m+1})\mathbf{x})}{\partial x} \\
& +\frac{\beta^n_{k+1}}{n-k}\sum_{j=k+m}^{n}\frac{\partial G_{i,k}(x_n,..x_{k+m},\mathbf{mx},x_j)}{\partial x} \\
< & (1-\beta^n_{k+1})+\beta^n_{k+1}\left(\frac{m}{n-k}+\frac{n-k-m}{n-k}\right)=1
\end{aligned}
$$

where the inequality follows by the induction hypothesis. By setting $m = 1$, we obtain from the above that:

$$\forall i,k \quad \frac{\partial a_{i,k}(x_n, ..., x_{k+1})}{\partial x_{k+1}} < 1$$

Together with Theorem1-3, this proves the result. ■

For the proof of Theorem 3 we first need the following Lemma:

Lemma 6 *Assume that for any k, and for any pair of ordered lists of reports $\chi_k \geq \chi'_k$ that differ only in one coordinate $\widetilde{F}_k(x|\chi_k) \succsim_{FOSD} \widetilde{F}_k(x|\chi'_k)$. Then the cutoff $a_{i,k}(\chi_k, x_k)$ is non-decreasing in x_k.*

Proof. The proof is by induction on the number of remaining periods. For $k = 2$ we have

$$\begin{aligned} a_{2,2}(\chi_2, x_2) &= \infty \\ a_{1,2}(\chi_2, x_2) &= \int_0^\infty x_1 d\widetilde{F}_1(x_1|\chi_2, x_2) \\ a_{0,2}(\chi_2, x_2) &= 0 \end{aligned}$$

[12]First order stochastic dominance immediately implies that the cutoffs are non-decreasing in x_2. We now apply the induction argument, and assume that, for any χ_k and for any i, $a_{i,k}(\chi_k, x_k)$ is non-decreasing in x_k. This implies that the function $G_{i,k}(x_k, x_{k+1}, \chi_{k+1})$ (see the proof of Theorem 2 for its definition) is non-decreasing in x_k , and that for any i we have:

$$\begin{aligned} a_{i,k}(\chi_{k+1}, x_{k+1}, x_k) &= a_{i,k}(\chi_{k+1}, x_k, x_{k+1}) \geq \\ a_{i,k}(\chi_{k+1}, x_k, x'_{k+1}) &= a_{i,k}(\chi_{k+1}, x'_{k+1}, x_k) \end{aligned}$$

where both equalities follow from the symmetry property whereby switching the order of the observations does not affect the final beliefs[13]. Therefore we obtain $G_{i,k}(x_k, x_{k+1}, \chi_{k+1}) \geq G_{i,k}(x_k, x'_{k+1}, \chi_{k+1})$ for any x_k. Moreover we have that

$$\begin{aligned} a_{i,k+1}(\chi_{k+1}, x_{k+1}) &= E_{x_k|x_{k+1}} G_{i,k}(x_k, x_{k+1}, \chi_{k+1}) \\ &\geq E_{x_k|x_{k+1}} G_{i,k}(x_k, x'_{k+1}, \chi_{k+1}) \\ &\geq E_{x_k|x'_{k+1}} G_{i,k}(x_k, x'_{k+1}, \chi_{k+1}) = a_{i,k+1}(\chi_{k+1}, x'_{k+1}) \end{aligned}$$

[12]This property also holds for Bayesian updating and for the non-Bayesian updating procedure analyzed in the previous Section.

[13]This property holds for both procedures analyzed in this paper. It also holds for Bayesian updating.

where the second inequality follows from first order stochastic dominance, and from the fact that, by the induction argument, $G_k(x_k, x'_{k+1}, \chi_{k+1})$ is non-decreasing in x_k. ■

Proof of Theorem 3. We shall show that for any i, k, χ_k and x_k, the cutoff $a_{i,k}(\chi_k, x_k)$ is continuously differentiable in x_k with $\frac{\partial}{\partial x_k} a_{i,k}(\chi_k, x_k) \leq \frac{1}{n-k+2}$. We prove this result by induction on k, the number of remaining periods. Note first that Lemma 6 above yields the monotonicity of $a_{i,k+1}(\chi_{k+1}, x_{k+1})$ in x_{k+1}.

We denote by $x_{(i)}$, the $i-$th lowest observation among the $n-k+1$ observations made up to and including period k, with $x_{(0)} = 0$ and $x_{(n-k+2)} = 1$. For $k = 2$, we have:

$$\begin{aligned} a_{1,2}(x_n, ..., x_2) &= \sum_{i=1}^{n} \int_{x_{(i-1)}}^{x_{(i)}} \frac{x}{n\left(x_{(i)} - x_{(i-1)}\right)} dx \\ &= \frac{1 + 2\sum\limits_{i=1}^{n-1} x_{(i)}}{2n} = \frac{1 + 2\sum\limits_{i=2}^{n} x_i}{2n} \Rightarrow \\ \frac{\partial a_{1,2}(x_n, ..., x_2)}{\partial x_2} &= \frac{1}{n}. \end{aligned}$$

Assume now that the statement holds for all periods up to k. This implies that there exists at most one solution to the equation $a_{i,k}(\chi_k, x_k) = x_k$, denoted by $a^*_{i,k}(\chi_k)$. Let $l = \max\left\{j : x_{(j)} \leq a^*_{i-1,k}(\chi_k)\right\}$ and $m = \max\left\{j : x_{(j)} \leq a^*_{i,k}(\chi_k)\right\}$, and assume, for simplicity, that $m > l$ (the case $m = l$ is analogous). Using the definition of $a_{i,k+1}(\chi_{k+1}, x_{k+1})$ we obtain

$$
\begin{aligned}
& a_{i,k+1}(\chi_{k+1}, x_{k+1}) \\
= \; & \sum_{j=1}^{l} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} a_{i-1,k}(\chi_k, x_k) dx_k}{(n-k+1)\left(x_{(j)} - x_{(j-1)}\right)} + \frac{\int\limits_{x_{(l)}}^{a^*_{i-1,k}(\chi_k)} a_{i-1,k}(\chi_k, x_k) dx_k}{(n-k+1)\left(x_{(l+1)} - x_{(l)}\right)} \\
& + \frac{\int\limits_{a^*_{i-1,k}(\chi_k)}^{x_{(l+1)}} x_k dx_k}{(n-k+1)\left(x_{(l+1)} - x_{(l)}\right)} + \sum_{j=l+2}^{m} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} x_k dx_k}{(n-k+1)\left(x_{(j)} - x_{(j-1)}\right)} \\
& + \frac{\int\limits_{x_{(m)}}^{a^*_{i,k}(\chi_k)} x_k dx_k}{(n-k+1)\left(x_{(m+1)} - x_{(m)}\right)} + \frac{\int\limits_{a^*_{i,k}(\chi_k)}^{x_{(m+1)}} a_{i,k}(\chi_k, x_k) dx_k}{(n-k+1)\left(x_{(m+1)} - x_{(m)}\right)} \\
& + \sum_{j=m+1}^{n-k+1} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} a_{i,k}(\chi_k, x_k) dx_k}{(n-k+1)\left(x_{(j)} - x_{(j-1)}\right)}.
\end{aligned}
$$

Let j **be** the index satisfying $x_{(j)} = x_{k+1}$.There are three different cases: 1. $x_{k+1} \leq x_{(l)}$; 2. $x_{(m)} \geq x_{k+1} > x_{(l)}$; 3. $x_{k+1} > x_{(m)}$. We prove the result for the first case; the proofs of the other two cases are very similar, and we omit them here. We obtain:

$$
\begin{aligned}
& \frac{\partial a_{i,k+1}(\chi_{k+1}, x_{k+1})}{\partial x_{k+1}} \\
= \; & \frac{a_{i-1,k}(\chi_k, x_{(j)})}{(n-k+1)\left(x_{(j)} - x_{(j-1)}\right)} - \frac{a_{i-1,k}(\chi_k, x_{(j)})}{(n-k+1)\left(x_{(j+1)} - x_{(j)}\right)} \\
& + \sum_{j=1}^{l} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} \frac{\partial a_{i-1,k}(\chi_k, x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(j)} - x_{(j-1)}\right)} + \frac{\int\limits_{x_{(l)}}^{a^*_{i-1,k}(\chi_k)} \frac{\partial a_{i-1,k}(\chi_k, x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(l+1)} - x_{(l)}\right)}
\end{aligned}
$$

$$+\frac{\int\limits_{a^*_{i,k}(\chi_k)}^{x_{(m+1)}} \frac{\partial a_{i,k}(\chi_k,x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(m+1)}-x_{(m)}\right)} + \sum_{j=m+1}^{n-k+1} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} \frac{\partial a_{i,k}(\chi_k,x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)}$$

$$-\frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} a_{i-1,k}(\chi_k,x_k)dx_k}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)^2} + \frac{\int\limits_{x_{(j)}}^{x_{(j+1)}} a_{i-1,k}(\chi_k,x_k)dx_k}{(n-k+1)\left(x_{(j+1)}-x_{(j)}\right)^2}$$

Note that

$$\begin{aligned}
& \sum_{j=1}^{l} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} \frac{\partial a_{i-1,k}(\chi_k,x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)} + \frac{\int\limits_{x_{(l)}}^{a^*_{i-1,k}(\chi_k)} \frac{\partial a_{i-1,k}(\chi_k,x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(l+1)}-x_{(l)}\right)} \qquad (4) \\
& +\frac{\int\limits_{a^*_{i,k}(\chi_k)}^{x_{(m+1)}} \frac{\partial a_{i,k}(\chi_k,x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(m+1)}-x_{(m)}\right)} + \sum_{j=m+1}^{n-k+1} \frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} \frac{\partial a_{i,k}(\chi_k,x_k)}{\partial x_{k+1}} dx_k}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)} \\
\leq\ & \frac{1}{(n-k+1)} \frac{n-k-m+l+1}{n-k+2} \leq \frac{1}{(n-k+1)} \frac{n-k}{n-k+2}
\end{aligned}$$

where the first inequality follows from the inductive assumption ($\frac{\partial a_{i-1,k}(\chi_k,x_k)}{\partial x_{k+1}} \leq \frac{1}{n-k+2}$) while the second inequality follows because $m > l$. In addition,

$$\begin{aligned}
& \frac{a_{i-1,k}(\chi_k,x_{(j)})}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)} - \frac{a_{i-1,k}(\chi_k,x_{(j)})}{(n-k+1)\left(x_{(j+1)}-x_{(j)}\right)} \qquad (5) \\
& -\frac{\int\limits_{x_{(j-1)}}^{x_{(j)}} a_{i-1,k}(\chi_k,x_k)dx_k}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)^2} + \frac{\int\limits_{x_{(j)}}^{x_{(j+1)}} a_{i-1,k}(\chi_k,x_k)dx_k}{(n-k+1)\left(x_{(j+1)}-x_{(j)}\right)^2} \\
\leq\ & \frac{a_{i-1,k}(\chi_k,x_{(j)})}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)} - \frac{a_{i-1,k}(\chi_k,x_{(j)})}{(n-k+1)\left(x_{(j+1)}-x_{(j)}\right)} \\
& -\frac{a_{i-1,k}(\chi_k,x_{(j-1)})}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)} + \frac{a_{i-1,k}(\chi_k,x_{(j+1)})}{(n-k+1)\left(x_{(j+1)}-x_{(j)}\right)} \\
=\ & \frac{a_{i-1,k}(\chi_k,x_{(j)})-a_{i-1,k}(\chi_k,x_{(j-1)})}{(n-k+1)\left(x_{(j)}-x_{(j-1)}\right)} + \frac{a_{i-1,k}(\chi_k,x_{(j+1)})-a_{i-1,k}(\chi_k,x_{(j)})}{(n-k+1)\left(x_{(j+1)}-x_{(j)}\right)} \\
=\ & \frac{1}{n-k+1}\left[\frac{\partial}{\partial x_{k+1}} a_{i-1,k}(\chi_k,x'_k) + \frac{\partial}{\partial x_{k+1}} a_{i-1,k}(\chi_k,x''_k)\right] \\
\leq\ & \frac{1}{n-k+1} \frac{2}{n-k+2}
\end{aligned}$$

where $x_k' \in \left[x_{(j-1)}, x_{(j)}\right]$ and $x_k'' \in \left[x_{(j)}, x_{(j+1)}\right]$. The first inequality follows from the monotonicity of $a_{i-1,k}(\chi_k, x_k)$, and the last inequality follows from the induction argument. Combining (4) and (5) we obtain

$$\frac{\partial a_{i,k+1}(\chi_{k+1}, x_{k+1})}{\partial x_{k+1}} \leq \frac{1}{n-k+1}.$$

as desired. ■

References

[1] Albright, S.C. (1977): "A Bayesian Approach to a Generalized House Selling Problem." *Management Science*, **24**(4), 432-440.

[2] Bickhchandani, S. and S. Sharma, (1996): "Optimal Search with Learning." *Journal of Economic Dynamics and Control* **20**, 333-359.

[3] Chou, C.F, and Talmain, G. (1993): "Nonparametric Search." *Journal of Economic Dynamics and Control*, **17**: 771-784.

[4] Dasgupta, P. and Maskin, E. (2000): "Efficient Auctions." *Quarterly Journal of Economics,* **115**(2): 341-388.

[5] Derman, C., Lieberman, G, and Ross, S. (1972): "A Sequential Stochastic Assignment Problem." *Management Science*, **18**(7): 349-355.

[6] Gershkov, A. and Moldovanu, B. (2009): "Learning About the Future and Dynamic Efficiency." *American Economic Review* **99**(4): 1576-1588.

[7] Gershkov, A. and Moldovanu, B. (2010): "Optimal Search, Learning and Implementation", discussion paper, University of Bonn

[8] Jaynes, E. T. (2003): "*Probability Theory: The Logic of Science.",* Cambridge: Cambridge University Press.

[9] Jehiel, P. and Moldovanu, B. (2001): "Efficient Design with Interdependent Valuations." *Econometrica,* **69**(5): 1237-1259.

[10] Rosenfield, D. and Shapiro, R. (1981): "Optimal Adaptive Price Search." *Journal of Economic Theory,* **25**, 1-20.

[11] Rothschild, M. (1974): "Searching for the Lowest Price When the Distribution of Prices Is Unknown." *Journal of Political Economy*, **82**(4), 689-711.

[12] Seierstad, A. (1992): "Reservation Prices in Optimal Stopping." *Operations Research*, **40**, 409-414