Discrete Mathematics and Theoretical Computer Science DMTCS vol. **12:2**, 2010, 363–380

On a Class of Optimal Stopping Problems with Mixed Constraints

F. Thomas Bruss

Université Libre de Bruxelles, Département de Mathématique, Bruxelles, Belgium. tbruss@ulb.ac.be

received Sep. 24, 2009, revised March 17, 2010, accepted May 10, 2010.

Let $X_1, X_2, \cdots, X_n$ be independent, identically distributed uniform random variables on $[0, 1]$. We can observe the outcomes sequentially and must select online at least r of them, and, moreover, in expectation at least $\mu \geq r$. Here μ need not be integer. We see X_k as the *cost* of selecting item k and want to minimize the expected total cost under the described combined (r, μ)-constraint. We will see that an optimal selection strategy exists on the set $\mathcal{S}_n$ of all selection strategies for which the decision at instant k may depend on the value X_k, on the number N_k of selections up to time k and of the number $n-k$ of forthcoming observations. Let $\sigma_{r,\mu}(n)$ be the corresponding $\mathcal{S}_n$-optimal selection strategy and $v_{r,\mu}(n)$ its value. The main goal of this paper is to determine these and to understand the limiting behavior of $v_{r,\mu}(n)$. After discussion of the specific character of this combination of two types of constraints we conclude that the S_n-problem has a recursive structure and solve it in terms of a double recursion. Our interest will then focus on the limiting behavior of $nv_{r,\mu}(n)$ as $n \to \infty$. This sequence converges and its limit allows for the interpretation of a normalized limiting cost $\mathcal{L}(r, \mu)$ of the (r, μ)-constraint. Our main result is that $\mathcal{L}(r, \mu) = g_r((\mu - r)^2/2)$ where g_r is the r^{th} iterate of the function $g(x) = 1 + x + \sqrt{1 + 2x}$.

Our motivation to study mixed-constraints problems is indicated by several examples of possible applications. We also shortly discuss the intricacy of the expectational part of the constraint if we try to extend the class of strategies $\mathcal{S}_n$ to the set of full-history-dependent and/or randomized strategies.

Keywords: Selection Problem, multiple selection, full information, definite constraint, expectational constraint, double-recursive functions, Riccati differential equation, knapsack problem, minimal spanning tree.

This study is dedicated to
Philippe Flajolet, INRIA,
Dr. h.c. of the Université Libre de Bruxelles,
at the occasion of his 60th birthday
for his outstanding contributions to the
Theory and Applications of Algorithms.

1 Introduction

Let $X_1, X_2, \cdots, X_n$ be i.i.d. $U[0, 1]$-variables. We can observe the X_k's sequentially and must decide at each time (instant) k whether to select X_k, or else to refuse it and to continue with further observations. The number of selections is subject to two different types of constraints. On the one hand there is a prescribed minimal number r of selections, and, on the other hand, the selection strategy should accept, in expectation, a minimum number μ of the X_k's. We interpret X_i as the cost of item i. Our objective is to minimize the expected total cost of selected items under the two constraints over the set of all non-anticipative online selection strategies. *Non-anticipative* means that we have no prophetical skills whatsoever, and *online*, that each decision to select or refuse must be taken on the spot, without recall on preceding observations.

Mathematical formulation

For $n \in \mathbb{N}$ fixed let $X_1, X_2, \cdots, X_n$ be i.i.d. uniform random variables on $[0,1]$ defined on the probability space $(\Omega, \mathcal{F}, P)$. Further let $I_1, I_2, \cdots, I_n$ be indicator functions fully adapted to the X_k in the sense that the events $\{I_k = 0\}, \{I_k = 1\}$ are in the σ-field $\mathcal{F}_k$ generated by $(X_1, X_2, \cdots, X_k)$. We say that X_k is *selected* if $I_k = 1$, and refused if $I_k = 0$. Let $\mathcal{T}_n$ be the set of all such stopping vectors $\tau := \tau_n = (I_1, I_2, \cdots, I_n)$. Our objective is to find, provided it exists, the value

$$u_{r,\mu}(n) = \min_{(I_1,\cdots,I_n)\in\mathcal{T}_n} \mathrm{E}\left(\sum_{k=1}^{n} I_k X_k\right),\ n \geq \mu \geq r,\ \ r, n \in \mathbb{N} \tag{1.1}$$

and the corresponding selection rule

$$\tau_{r,\mu}(n) = \mathrm{argmin}_{\tau\in\mathcal{T}_n} \mathrm{E}\left(\sum_{k=1}^{n} I_k X_k\right) \tag{1.2}$$

achieving this value, subject to the constraints

$$\sum_{k=1}^{n} I_k \geq r,\ r \in \mathbb{N},\ 1 \leq r \leq n \tag{1.3}$$

and

$$\mathrm{E}\left(\sum_{k=1}^{n} I_k\right) \geq \mu,\ \mu \in \mathbb{R},\ \mu \geq r. \tag{1.4}$$

The constraints (1.3) and (1.4) will be referred to as D-constraint and E-constraint respectively, and any non-anticipative stopping rule satisfying them both as *admissible*. The corresponding strategy $\tau_{r,\mu}(n)$ minimizing the total cost (loss) in the class of admissible rules, provided it exists, is called *optimal*.

1.1 Classification of problem-type

In the terminology of optimal stopping problems, this problem is a finite-horizon i.i.d. full-information multiple stopping problem. The meaning of *multiple stopping* is obvious. *Full-information* refers to the hypothesis that the distribution of the X_k's is known to us, and that we are allowed to observe the actual *values* of these and not only some statistic of these as for instance in rank-based stopping problems. We should and do add the description *with mixed constraints*, because we have here two different *types* of constraints on the number of required selections: a *definite* constraint (D-constraint (1.3)) as well as an *expectational* constraint (E-constraint (1.4)).

1.2 Motivation

As far as the author is aware, this combination of constraints is new. The author's interest in this problem was instigated by David Aldous (private communication, 2006) who asked what is in general known about such problems. Aldous formulated a problem of improving a lower bound of interest for spanning trees in terms of such a stopping problem for $r = 1$ and $\mu = 2$ and we will return to it later on. (See also Aldous et al. (2008)). Our motivation is now also going further because, once one thinks about it, one sees that the proposed mixed-constraints setting offers a variety of applications: In certain buying/selling problems, for instance, some contract specifications require not only a minimum number of purchases, but also a long-run average "chiffres d'affaires". In the case of buying, our proposed model fits the problem directly; in the case of selling we re-interpret the costs as benefits resulting from sales and would like to solve the equivalent maximization problem. Another example of selection problems where this setting is relevant is the online-knapsack problem: Items, each having utility one, arrive in random order. Time constraints force us to select online from a sequence of incoming different items and to pack without recall. Costs are measured in terms of weight or space, and the total combined utility is supposed to be at least r, with an average of at least μ. There are many other problems along these lines which come to our mind.

In *practical* problems the verifiability of the E-constraint would have clear limitations, of course. Moreover, as we shall see in Section 2, this constraint needs also special attention from the theoretical point of view. Having said this, we may all agree that the mixed-constrained objective has applicational appeal and should merit our interest.

1.3 Related work

In our problem the cost of each selection is simply the value of the variable, that is *we pay for what we select*. There is no cost of inspection, and the problem has its interest through the new combination of constraints. Models with a fixed inspection cost are different in their structure. Examples for related problems with cost for inspection are the classical house-selling problem studied by MacQueen and Miller (1960) and generalized by Stadje (1990) as well as the multiple buying-selling problem of Bruss and Ferguson (1997).

The literature in the domain of so-called secretary problems shows many papers on selection problems with full-information and varying single constraints as e.g. Chen et al. (1984), Kurushima and Ano (2009), Mazalov and Tamaki (2006), and Tamaki (1986). Sometimes selection means investment (as in Lebek and Szajowski (2007) and others) and then the cost is measured in terms of limited resources. (For a review of earlier work see e.g. Samuels (1990).) The single constraint makes each of these problems quite different from ours. The same holds for the selection problem of Baryshnikov and Gnedin (2000) of vector-valued observations and the competitive selection problem with two decision makers studied by Enns and Ferenstein (1985) and Bruss et al. (1998). And then there are multiple selection problems where the number of selections need not be specified such as e.g. the online monotone subsequence problem studied by Samuels and Steele (1981), Gnedin (2000), and Bruss and Delbaen (2001). We note that, here again, the constraint, although being a "moving constraint" for each selection, is still a single-type constraint, namely monotonicity. Our current problem combining two types of constraints seems very different in nature from all of these.

One special case of the described problem has been studied in the literature, and, interestingly, it appeared before secretary problems started to attract attention. This is the so-called Caley/Moser problem (Moser (1956)). In this problem the decision maker observes sequentially n i.i.d. $U[0, 1]$-random variables, and exactly one selection must be made, that is $r = 1$. The Caley/Moser objective was to maximize the expected value of the selected observation. With X_k being $U[0,1]$, the transformation $Z_k = 1 - X_k$ stays a $U[0,1]$ variable, which makes it equivalent to our minimization problem. No E-constraint on the number of selections was mentioned in the Caley/Moser problem but it concurs with the implicit constraint $\mu = 1$, of course.

1.4 Organization of this paper

In Section 2 we first analyze the problem as stated in (1.1)-(1.4) and see that the problem to find $u_{r,\mu}(n)$ cannot be solved by separating the two constraints. We then argue that we should be somewhat restrictive by confining the optimization problem to a subclass $\mathcal{S}_n$ of selection strategies which, at each instant k, may depend on the current observation X_k, on $n-k$, and on the total number N_k of observations selected so far. Hence we replace the problem of finding the overall optimum $u_{r,\mu}(n)$ (see (1.1)-(1.4)) by the problem of finding the corresponding optimal value $v_{r,\mu}(n)$ over all strategies $s \in \mathcal{S}_n$. In Section 3. we use a threshold-value argument to find a double recursion for the optimal value $v_{r,\mu}(n)$ and prove the existence of an $\mathcal{S}_n$-optimal strategy $\sigma_{r,\mu}(n)$ which achieves $v_{r,\mu}(n)$. In Section 4. we collect basic properties of $v_{r,\mu}(n)$ which will be used in the following sections. Section 5. proposes an algorithm to compute $v_{r,\mu}(n)$ and $\sigma_{r,\mu}(n)$ simultaneously. We also give preliminary estimates for the asymptotic behavior of $v_{r,\mu}(n)$ with the purpose of preparing the Section 6. This section presents the main asymptotic result: we prove that $nv_{r,\mu}(n)$ converges for all $r \in \mathbb{N}$ and $\mu \geq r$. We then show via a well defined approximation by a Riccati differential equation that each $V_{r,\mu} := \lim_{n\to\infty} nv_{r,\mu}(n)$ can be explicitly computed by the rth iterate of a function g evaluated in $(\mu - r)^2/2$.

2 Restricting the set of possible strategies

Let $X_1, X_2, \cdots, X_n$ be i.i.d. uniform random variables on $[0,1]$ defined on $(\Omega, \mathcal{F}, P)$, and let, as before, I_k be the indicator of the event that X_k is *selected*. A closer analysis of the problem as stated in (1.1)-(1.4) shows that the E-constraint raises some questions which may be difficult to answer. (The author is indebted to the referee for drawing his attention to this point.) To indicate possible problems and to add something possibly new to the referee's discussion we consider the case $r = 0$. Then the problem becomes that of finding the minimal expected cost under the pure E-constraint

$$\mathrm{E}\left(\sum_{k=1}^{n} I_k\right) = \sum_{k=1}^{n} P(I_k = 1) \geq \mu.$$

A frequentist may justifiably argue that the only way to verify that a given strategy fulfills this constraint is to apply it to a sequence of identical games, that is, to apply it consecutively to m blocks of n i.i.d. $U[0,1]$-random variables

$$\left\{(X_1^j, X_2^j, \cdots, X_n^j)\right\}_{j=1,2,\cdots,m}$$

and to check whether for $m \to \infty$ the $\liminf$ of the arithmetic mean of the number of selections per block is at least μ. This raises the question whether a good strategy may allow to select in certain blocks many more observations than necessary - for instance if there is an opportunity to collect them cheaply - in order to compensate for other blocks where one may prefer a particular small value of the loss and select nothing.

Note that the problem stays the same for general $r \in \mathbb{N}, r > 0$, because, inherently, any admissible strategy must end up in a random index k, at which the D-constraint becomes void. The author does not claim that the problem is, for such reasons, ill-posed, but he does not see how to make things precise to affirm, that an optimal strategy for (1.1)-(1.4) must always exist. We therefore confine the study to the set of strategies in which the "memory" of the indicators $\{I_k\}_{k=1,2,\ldots,n}$ is just sufficient to guarantee admissibility of the applied strategy at each instant. In this vein, let

$$\mathcal{S}_n = \big\{\sigma = (I_1, I_2, \cdots, I_n) : I_k \text{ may depend on } X_k, n-k \text{ and } N_{k-1}\big\}, \tag{2.1}$$

where

$$N_k = I_1 + I_2 + \cdots + I_k. \tag{2.2}$$

Our objective is to find the value

$$v_{r,\mu}(n) := \min_{\sigma \in \mathcal{S}_n} \mathrm{E}\left(\sum_{k=1}^{n} I_k X_k\right),\ n \geq \mu \geq r,\ \ r, n \in \mathbb{N}, \tag{2.3}$$

and the corresponding selection rule

$$\sigma_{r,\mu}(n) = \mathrm{argmin}_{\sigma \in \mathcal{S}_n} \mathrm{E}\left(\sum_{k=1}^{n} I_k X_k\right) \tag{2.4}$$

achieving this value subject to the D-constraint $N_n \geq r$ and the E-constraint $\mathrm{E}(N_n) \geq \mu$.

2.1 Difference between D-constraint and E-constraint

We have already pointed out that the E-constraint needs special attention. It is quite different in nature from the D-constraint for which no ambiguity is possible. But the difference should also give something new, of course. We therefore allow that, once the D-constraint is satisfied under a given selection rule the reduced E-constraint is no longer sequentially updated in order to just achieve a minimum of $\mu - r$ more selections. Otherwise the mixed constraints setting would be equivalent to a D-constraint of $\lceil \mu \rceil$, where $\lceil x \rceil$ denotes the ceiling of x. This, however, would be the same case as the case where $r = \lceil \mu \rceil$ which is already covered through our model with the E-constraint being the same as the D-constraint. Therefore it is more general to allow that once the D-constraint is satisfied, the Problem becomes the problem of minimizing additional cost imposed by a pure E-constraint of $\mu - r$. Hence we can continue from that moment onwards with any (memoryless) strategy which selects, in expectation, $\mu - r$ items.

3 Existence and structure of the optimal strategy

Our way of attacking the existence of a solution to problem (2.1)-(2.4) benefits from two properties which simplify the problem essentially. The first flows out of the i.i.d. assumption for the X_k''s, namely:

Lemma 3.1 *Let $V_0(n) = v_{r,\mu}(n)$ and let $V(n|\mathcal{H}_k)$ denote the conditional minimal total cost expectation over all strategies in $\mathcal{S}_n$ given the history $\mathcal{H}_k = \{(X_1, I_1), (X_2, I_2), \cdots (X_k, I_k)\}$. If $v_{r-N_k,\mu-N_k}(n-k)$ exists for $0 \le N_k \le r$, then we must have*

$$V(n|\mathcal{H}_k) = v_{r-N_k,\mu-N_k}(n-k) + \sum_{j=1}^{k} I_j X_j \ a.s. \tag{3.1}$$

Proof: : Since the objective function (total cost of selection) is linear in each term (see (2.3)), the optimal strategy from time $k+1$ onwards must be optimal continuation in $\mathcal{S}_n$ after time k, given $\mathcal{H}_k$. The latter is subject to the D-constraint $r - N_k$ and the E-constraint $\mu - N_k$. Note that $N_k = I_1 + \cdots + I_k$ is determined by H_k. Since, by definition, the forthcoming random variables $X_{k+1}, \cdots, X_n$ are independent of $(X_1, I_1), \cdots, (X_k, I_k)$, and since $(X_{k+1}, \cdots, X_n)$ is distributed like $(X_1, \cdots, X_{n-k})$, the corresponding decision process is Markovian, that is, the value of the optimal continuation cannot depend but on $X_k, X_{k+1}, \cdots, X_n$, on $n-k$ and on N_k. The optimal loss increment is therefore $v_{r-N_k,\mu-N_k}(n-k)$ whereas the currently accumulated cost is $\sum_{j=1}^{k} I_j X_j$ a.s. Hence Lemma 3.1 is proved. □

The next Lemma confirms our intuition that (3.1) holds also if we replace the fixed index k by the stopping time defined by the first time instant in which the D-constraint is completely satisfied.

Lemma 3.2 *For given n, let δ be the smallest time at which the D-constraint is completely satisfied under application of the S_n-optimal strategy. Then,*

$$V_\delta(n) = v_{0,\mu-r}(n-\delta) + \sum_{j=1}^{\delta} I_j X_j \ a.s. \tag{3.2}$$

where

$$v_{0,\mu-r}(k) = \frac{(\mu-r)^2}{2k}. \tag{3.3}$$

Proof: Conditioned on $\delta = d$, statement (3.2) follows from Lemma 3.1 because the optimal rule in $\mathcal{S}_n$, σ say, must be admissible. Admissibility also implies that δ satisfies the condition $r \le \delta \le n - \mu + r$. At time δ, the future variables $X_{\delta+1}, \cdots, X_n$ are independent of $X_1, \cdots, X_\delta$. Hence, the first statement is bound to hold unconditionally. It remains to be shown that $v_{0,\mu-r}(k) = (\mu-r)^2/2k$. At time $\delta+$, with $K = n - \delta$ observations to come, the D-constraint vanishes. We are thus confronted with the task of designing a rule which selects in expectation $\mu - r$ more items from K i.i.d $U[0,1]$-random variables.

Now, the pure E-constraint concerns solely the expected number of selections and not their values. These values just intervene in the optimal cost. Hence, if it is optimal to select, for $j > \delta$, a value $X_j = x$, say, then it is optimal to accept any $X'_j < x$, because it is cheaper paying X'_j than paying x. Similarly, if it is optimal to refuse $X_j = x$, it must be optimal to refuse any X'_j with $X_j > x$. It follows that each optimal decision at time j is based on a unique threshold value. This is just the indifference-value with respect to an $\mathcal{S}_n-$optimal strategy on the remaining sequence $X_j, X_{j+1}, \cdots, X_n$.

Let now $t_1, t_2, \cdots, t_K$ be the selection thresholds for $X_{\delta+1}, X_{\delta+2}, \cdots, X_n$, respectively, that is, we select $X_{\delta+j}$ if and only if $X_{\delta+j} \le t_j$. We can clearly confine our interest to threshold values between 0 and 1. Then

$$\mathrm{E}(I_{\delta+j}X_{\delta+j}) = \int_0^1 \mathbf{1}\{x \le t_j\}x dx = \frac{t_j^2}{2}. \tag{3.4}$$

Equations (3.2) and (3.4) imply that we have to minimize $\sum_{j=1}^K t_j^2$ subject to the constraint

$$\sum_{j=1}^K \mathrm{E}(I_{\delta+j}) = \sum_{j=1}^K t_j \ge \mu - r.$$

Note that for all K the E-constraint can here be replaced by

$$\sum_{j=1}^K t_j = \mu - r, \tag{3.5}$$

since all t_j are non-negative. It is now straightforward to show by the Lagrange multiplier method that the optimal thresholds are constant with $t_j = (\mu - r)/K$. Hence

$$v_{0,\mu-r}(K) = K\frac{\mu - r}{K}\,\frac{\mu - r}{2K} = \frac{(\mu - r)^2}{2K},$$

implying (3.3). This completes the proof of Lemma 3.2. □

Existence of an Optimal Strategy

We now are ready to prove the existence of an optimal strategy in $\mathcal{S}_n$. We proceed in two steps.

The first step is based on the following facts: Given the decision process enters the constraints-state $(0, \mu - r)$ at time δ, the optimal future expected cost is, as we have just seen, the *unique* value $v_{0,\mu}(n-\delta)$, which is independent of $X_1, X_2, \cdots, X_\delta$. Then the idea is to look backwards to the first time the decision process enters the constraints-state $(1, \mu - r + 1)$. As we shall show below, the threshold argument used before will go through similarly. Using the independence of the X_k's it follows that the S_n-optimal future loss from that time onwards is just a function of unique threshold values for the next selection and the optimal expected cost increment $v_{0,\mu}(n-\delta)$. Hence $v_{1,\mu-r+1}(k)$ must exist, and the strategy to achieve this value is again unique. By induction, this argument will yield the recursive structure of the optimization problem in $\mathcal{S}_n$.

We now give the details:

Theorem 3.1 *Let, as before, $v_{r,\mu}(n)$ denote the value for the problem with D-constraint r and E-constraint $\mu \ge r$ for n observations. Then, for $\mu \ge r \ge 1$ the $v_{r,\mu}(n)$ must satisfy the (double) recursion*

$$v_{r,\mu}(n) = v_{r,\mu}(n-1) - \frac{1}{2}\left[v_{r,\mu}(n-1) - v_{r-1,\mu-1}(n-1)\right]^2, \tag{3.6}$$

for $n = \lceil\mu\rceil +, \lceil\mu\rceil + 1, \cdots$, with initial conditions

$$v_{r,\mu}(\lceil\mu\rceil) = \frac{\mu}{2}; \quad v_{0,\mu-r}(n) = \frac{(\mu - r)^2}{2n},\ n = 1, 2, \cdots. \tag{3.7}$$

Proof: Suppose we follow the strategy to select X_1 if and only if $X_1 \le t$ and to continue optimally thereafter until the end. Then our value becomes $\tilde{v}_{r,\mu}(n,t)$, say, and we see that this value must satisfy the optimality equation

$$\tilde{v}_{r,\mu}(n,t) = t\Big[\mathrm{E}\,(X|X \le t) + v_{r-1,\mu-1}(n-1)\Big] + (1-t)v_{r,\mu}(n-1). \tag{3.8}$$

Indeed, according to (3.1) we obtain, if we select X_1, the loss $\mathrm{E}(X|X \le t)$ as instantaneous loss plus the optimal loss increment thereafter. The latter is simply $v_{r-1,\mu-1}(n-1)$ because $(I_2, I_3, \cdots, I_n)$ may depend on X_1 only through N_1, that is only through I_1. If we refuse X_1 the optimal loss increment for the following $n-1$ observations is similarly $v_{r,\mu}(n-1)$. Using $\mathrm{E}(X|X \le t) = t/2$, the rhs of (3.8) is clearly differentiable in t for all $t \in]0,1[$, so that the solution of equation $\partial\tilde{v}_{r,\mu}(n,t)/\partial t = 0$ with $\partial^2\tilde{v}_{r,\mu}(n,t)/\partial t^2 > 0$ minimizes $\tilde{v}_{r,\mu}(n,t)$. It is straightforward to see that this *unique* solution t^* is given by

$$t^* = v_{r,\mu}(n-1) - v_{r-1,\mu-1}(n-1). \tag{3.9}$$

There is no other candidate strategy for minimizing the expected future cost in $\mathcal{S}_n$ so that we must have $\tilde{v}_{r,\mu}(n,t^*) = v_{r,\mu}(n)$. Using this fact together with the solution (3.9) in equation (3.6) yields, after some elementary transformations, the recursion equation (3.6) of Theorem 3.1.

To see the initial conditions (3.7), suppose first that $\mu \in \mathbb{N}$ and $n = \mu$. Then, after r compulsory selections, the remaining $\mu - r$ E-constrained selections are also compulsory. Hence the optimal policy must select all observations, and has therefore value $\mu/2$. If $\mu - r \notin \mathbb{N}$ the result turns out the same, because with $\{x\}$ denoting the fractional part of x and $\lfloor x \rfloor$ the floor of x, the E-constraint imposes $(\lfloor \mu - r \rfloor) + \{\mu/2\}$ selections on its own, and hence the value is $r/2 + \lfloor \mu - r \rfloor/2 + \{\mu/2\} = \mu/2$. Finally, the second initial condition follows from equation (3.5), and thus the Theorem is proved. □

Remark 3.1 *It is important to note that in this (Bellman-type) optimality equation the restriction to strategies in the set $\mathcal{S}_n$ is, a priori, not redundant. The optimal future loss depends on $(X_2, I_2), (X_3, I_3), \dots, (X_n, I_n)$ and on the knowledge of N_1. Dependence on X_1 is therefore limited to the dependence on $I_1 = \mathbf{1}\{X_1 \le t\}$. This is why the future optimal expected losses figuring in (3.8) are $v_{r-1,\mu-1}(n-1)$ and $v_{r,\mu}(n-1)$, respectively, with an expected reduction of the E-constraint from μ to $\mu - t$ with $0 \le t \le 1 \le \mu$. There are several reasons why the author believes that $\mathcal{S}_n$ must contain the overall optimal strategy. He has no rigorous proof, however, and must leave it here as a conjecture.*

To describe the optimal rule we introduce the notion of a *state* of the selection process:

Definition 3.1 *For $\rho \in \{0, 1, \cdots, r\}$ and $k \in \{0, 1, \cdots, n\}$ we say that the selection process is in* state *(ρ, k), if ρ selections have been made until time $n-k$ included, that is, with k further unobserved variables to come.*

Note that the current E-constraint is implicit for $0 \le \rho \le r$. Since the continuation thereafter is, by hypothesis, a fixed selection rule, it becomes irrelevant once the D-constraint is satisfied. Hence we need not list it as a separate state-coordinate. From the uniqueness of the optimal thresholds (see (3.9)) and Theorem 6.1 we have then the following corollary.

Corollary 3.1 (Optimal selection strategy) *The optimal selection strategy $\sigma_{r,\mu}(n)$ is a state-dependent threshold selection rule, that is, for each state*

$$(\rho, k) \in \{0, 1, \cdots, r\} \times \{0, 1, \cdots, n\}$$

there exists a unique value $t_{\rho,k}$ such that it is optimal to select X_{n-k+1} if $X_{n-k+1} \leq t_{\rho,k}$, and to refuse X_{n-k+1} otherwise. These optimal thresholds are given by

$$t(\rho, k) := t_{r,\mu}(\rho, k) = v_{r-\rho,\mu-\rho}(k-1) - v_{r-\rho-1,\mu-\rho-1}(k-1). \tag{3.10}$$

4 Elementary monotonicity properties of $v_{r,\mu}(n)$

For all $r \in \mathbb{N}$ and $\mu \geq r$, the values $v_{r,\mu}(n)$ are positive by definition. Before giving in Section 5. the algorithm to implement the optimal rule $\sigma_{r,\mu}(n)$ and the optimal loss $v_{r,\mu}(n)$ we collect elementary properties of $v_{r,\mu}(n)$ which will be used in Sections 5. and Section 6.

We see from (3.6) in Theorem 3.1 that the sequence $(v_{r,\mu}(n))$ decreases in n whenever the sequence $(v_{r-1,\mu-1}(n))$ decreases in n. According to (3.7) we know that $(v_{0,\mu-r}(n))$ decreases in n and thus, as it is bounded below by 0, it must converge to the only possible limit 0. Hence we have by induction:

Corollary 4.1 *For all $\mu \geq 1$ and all integers r, n with $n \geq \mu \geq r$ the sequence $(v_{r,\mu}(n))_{n \geq \mu}$ is monotone decreasing with limit 0.*

This is the monotonicity of $v_{r,\mu}(n)$ with respect to the number n of observations. As we shall see now, v is, for fixed n, also monotone with respect to the both constraints. Note that a reduction of the D-constraint from r to $r-1$ goes, by definition, together with a reduction of the E-constraint from μ to $\mu-1$ (whereas the converse is not true.) The following lemma compares the values in different constraints for the same number of future observations:

Lemma 4.1 *For n fixed with $n \geq \ [\mu]^+$ and $\mu \geq r$ we have*

$$v_{r,\mu}(n) \geq v_{r-1,\mu-1}(n) \tag{4.1}$$

and

$$v_{r,\mu}(n) \geq v_{r,\mu-1}(n). \tag{4.2}$$

Proof: To see inequality (4.1), let $v^h_{\mu,r}(n)$ be the minimal expected total cost of the optimal strategy for the (r, μ)-constraint under the additional hypothesis that we will get, as a bonus, the rth selection for free. Then

$$v^h_{\mu,r}(n) \leq v_{\mu,r}(n) \tag{4.3}$$

because the total cost cannot but decrease by the bonus. However if we play right away optimally under the weaker $(r-1, \mu-1)$-constraint, then we take the minimum expected value over an extended set of strategies: Indeed, the latter covers both the set of strategies for r compulsory selections with the bonus of a free rth selection and those which may forego one selection. Therefore $v_{r-1,\mu-1}(n) \leq v^h_{r,\mu}(n)$, and hence, taking both inequalities together, $v_{r,\mu}(n) \geq v_{r-1,\mu-1}$. This proves (4.1).

Inequality (4.2) follows from the second inequality in (3.7) because $v_{0,\mu-r}(.) > v_{0,\mu-1-r}(.)$ uniformly. $\square$

5 Computing optimal thresholds and values

The optimal thresholds for each state can be computed recursively via the state dependent future values for k observations to come.

We have to start with the optimal values for which we have two independent lines of initial conditions, namely for $v_{0,\mu-r}(k)$ with $k \geq \mu - r$ and for $v_{s,k}(k)$ with $k \geq s$. The first one stems from the corresponding optimal memoryless threshold rule (see (3.3) in Lemma 3.2). The second one is even easier, because if $k \geq s$ and if we have to satisfy the E-constraint and D-constraint at the same time, we must accept all observations to come. These figure as (A1) and (A2) in part A of the algorithm below and are used in the recursion (A3). Part B below computes the optimal thresholds using A.

5.1 The algorithm

Algorithm 1: Compute the optimal strategy and the optimal value

A Optimal values:

(A1) **for** $k = \lceil\mu\rceil - r, \lceil\mu\rceil - r + 1, \cdots, n - r$ **do**
  $v_{0,\mu-r}(k) = (\mu - r)^2/(2k)$

(A2) **for** $k = \lceil\mu\rceil - r, \lceil\mu\rceil - r + 1, \cdots, n - r; s = 1, 2, \cdots r$ **do**
  $v_{s,k}(k) = k/2,$

(A3) **for** $s = 1, 2, \cdots, r$ *and initial conditions (A1) and (A2)* **do**
  $v_{s,\mu-r+s}(k) = v_{s,\mu-r+s}(k-1) - \frac{1}{2}\left[v_{s,\mu-r+s}(k-1) - v_{s-1,\mu-r+(s-1)}(k-1)\right]^2,$
  where $k = \mu - r + s, \cdots, n - r - 1, n - r$

B Optimal thresholds:

(B1) **for** $k = \lceil\mu\rceil - r, \lceil\mu\rceil - r + 1, \cdots n - r$ **do**
  $t(r,k) = v_{0,\mu-r}(k) = (\mu - r)^2/2,$

(B2) **for** $\rho = 0, 1, \cdots, r - 1$ **do**
  $t(\rho, k) = v_{r-s,\mu-s}(k-1) - v_{r-s-1,\mu-s-1}(k-1)$

Example 5.1
As an example we solve the problem for $r = 1$ *and* $\mu = 2$ *(Aldous' problem). We have* $\mu - r = 1$ *and hence from (3.3) the equation* $v_{0,1}(k) = 1/(2k)$. *Moreover the initial condition for* $v_{1,2}(2) = 1$ *is evident, since any admissible rule must accept both observations. Hence, using (3.6),*

$$v_{1,2}(k) = v_{1,2}(k-1) - \frac{1}{2}\left(v_{1,2}(k-1) - \frac{1}{2(k-1)}\right)^2, \; k = 2, 3, \cdots n.$$

Putting $v(k) := v_{2,1}(k)$ *we see that the latter is a simple recursion, that is*

$$v(k) = v(k-1) - \frac{1}{2}\left(v(k-1) - \frac{1}{2(k-1)}\right)^2, k = 3, 4 \cdots n, \tag{5.1}$$

with initial condition $v(2) = 1$. *The sequence* $(v(n))$ *is clearly decreasing and bounded below by 0. Hence* $v = \lim v(n)$ *exists and taking limits shows* $v = 0$. *Aldous' question concerned the limit* y *of* (y_n) *where* $y(n) = nv(n))$. *As the next section will show, this limit exists and equals* $3/2 + \sqrt{2}$. *Hence*

$$v_{2,1}(n) \sim \frac{1}{n}\left(\frac{3}{2} + \sqrt{2}\right). \tag{5.2}$$

5.2 Bounds of $v_{r,\mu}(n)$ for general r and μ

The motivation of the following simple lemma is to obtain bounds for the values $v_{r,\mu}(n)$ in terms of values stemming from *simple* recursions. The following result is intuitive and easy to prove:

Lemma 5.1 *For all* $0 \leq \rho \leq r, \rho \leq m \leq \mu$ *and* $\max\{\rho, m\} \leq k \leq n$ *we have*

$$v_{r,\mu}(n) \leq v_{\rho,m}(k) + v_{r-\rho,\mu-m}(n-k).$$

Proof: Fix indices ρ, m and k such that the conditions for the Lemma are fulfilled. This is always possible since the set of such triples (ρ, m, k) contains at least the trivial triples (r, μ, n) and $(0, 0, 0)$ for which the statement holds by definition. More generally, consider a two-legged strategy. Leg 1 minimizes the expected total cost of accepting items until time k under the D-constraint ρ and the E-constraint m. The continuation leg 2 remembers the occurred cost at time k and then minimizes (independently) the additional cost of accepting further items under the D-constraint $r - \rho$ and the E-constraint $\mu - m$. This composed strategy is admissible since it fulfills the original constraints, and since $X_{k+1}, X_{k+2} \cdots X_n$ are independent of the past, its value is $v_{\rho,m}(k) + v_{r-\rho,\mu-m}(n-k)$. The inequality follows then by sub-optimality. □

For the special case $\rho = r = m$ we obtain then from equality (3.6) and Lemma 5.1:

Corollary 5.1 *For* $1 \leq r \leq \mu \leq n$*:*

$$v_{r,\mu}(2n) \leq v_{r,r}(n) + \frac{1}{2n}(\mu - r)^2.$$

Lemma 5.1 and Corollary 5.1 have an important consequence (Lemma 5.2) which will make the asymptotic approach in Section 6 rigorous:

Lemma 5.2 *For all* $1 \leq r \leq \mu$ *there exist constants* $\alpha = \alpha(r, \mu)$ *and* $\beta = \beta(r, \mu)$ *such that*

$$\frac{\alpha}{n} \leq v_{r,\mu}(n) \leq \frac{\beta}{n} \tag{5.3}$$

for all $n \geq \mu$, *with* n *sufficiently large.*

Proof: We first prove the lower bound $\alpha(r, \mu)/n$.
By definition of the D-constraint and E-constraint we have $\mu \geq r$. Since, moreover, $v_{r,\mu}(\cdot)$ is non-decreasing in μ, for fixed r and n (see inequality (4.2)), it suffices to show $v_{r,r}(n) \geq \alpha/n$ for some

constant α. Now, the optimal strategy to fulfill the (r,r)-constraints cannot do better than selecting the r items of minimum total cost, that is all r smallest order statistics. The expectation of the sum of these is

$$\frac{1}{n+1}+\frac{2}{n+1}+\cdots+\frac{r}{n+1}=\frac{r(r+1)}{2(n+1)},$$

and is thus greater than $r^2/(2n)$. Hence $v_{r,\mu}(n) \geq v_{r,r}(n) \geq \alpha/n$ for $\alpha = r^2/2$.
Concerning the upper bound β we see (now from from Corollary 5.1) that the statement is true, if it is true for $v_{r,r}(n)$. To see the latter we note that $v_{r,r}(rn) \leq v_{r-1,r-1}((r-1)n) + v_{1,1}(n)$, and hence by induction

$$v_{r,r}(rn)) \leq r v_{1,1}(n). \tag{5.4}$$

The sequence $(v_{1,1}(n))$ coincides with Moser's sequence, which is known to satisfy $v_{1,1}(n) \leq 2/n$ for all n. Therefore $v_{r,r}(rn)) \leq (2r)/n$. But then, for general n we have $v_{r,r}(n) \leq v_{r,r}(\lfloor n/r \rfloor)$, where $\lfloor x \rfloor$ denotes the floor of x. Hence

$$v_{r,r}(n) \leq \frac{2r}{\lfloor n/r \rfloor} \leq \frac{2r^2+\epsilon}{n}$$

for all $\epsilon > 0$ and n sufficiently large, and thus the proof is complete. □

We can now complete Example 5.1. Recall $y_k = kv(k)$ where $v(k) = v_{1,2}(k)$ for $k = 2, 3, \cdots$. It follows from Lemma 5.2 that (y_k) is bounded by some positive constants α and β. Hence, to show that $y = \lim y_k$ exists it suffices to show that y_k is monotone for all k sufficiently large. Replacing k by $k+1$ in equation (5.1) and multiplying this equation then by k yields after some elementary transformations

$$y_{k+1} - y_k = v(k+1) - \frac{k}{2}\left(v(k) - \frac{1}{2k}\right)^2, \tag{5.5}$$

so that it suffices to show that the rhs does not change signs for k sufficiently large. Indeed, it is always positive for all $k \geq 2$, as A. Dutrifoy (private communication) has shown by a skillful direct estimation. However, we do not give his proof here because we give below an independent and purely probabilistic proof of the more general result that $nv_{r,\mu}(n)$ is increasing in n. Indeed, this monotonicity is true for all $r \in \mathbb{N}, \mu \geq r$ and all $n \geq \lceil \mu \rceil$.

Hence we know that $y = \lim_n y(n)$ exists, so that we can eliminate y in the limiting equation and pass on to take limits again. This yields in terms of y the new limit equation $y = (1/2)(y - 1/2)^2$ with solutions $3/2 \pm \sqrt{2}$. Since $y_n \geq 1$ for all $n \geq 2$, and hence $y \geq 1$, we can exclude the smaller root, and therefore $y = 3/2 + \sqrt{2}$.

6 Approximation of general solution

We now prove the important result, that $nv_{r,\mu}(n)$ converges for all $r \in \mathbb{N}$ and all $\mu \geq r$ to a limit $c(r,\mu)$ which we can compute explicitly. Hence $v_{r,\mu}(n) \sim c(r,\mu)/n$. This will be done in several steps.

6.1 Monotonicity

Lemma 6.1 *For all $r \in \mathbb{N}$ and $\mu \geq r$ the sequence $\left(n\, v_{r,\mu}(n)\right)_{n \geq \mu, n \in \mathbb{N}}$ is increasing in n.*

Proof: In order to prove the statement of the Lemma we have to show that

$$v_{r,\mu}(n) \geq \frac{n-1}{n}\, v_{r,\mu}(n-1), \text{ for } n \geq \mu+1, n \in \mathbb{N}. \tag{6.1}$$

To see this we use a refinement of the "half-prophet"-trick of Bruss and Ferguson. See (Bruss and Ferguson, 1993, Theorem 2, pp 623-624). Imagine a decision maker, H say, which has some "half-prophetical" capacity in the sense that he or she can predict beforehand the largest of the n outcomes, that is $M_n = \max\{X_1, X_2, \cdots, X_n\}$. Suppose now that H wants to minimize, exactly as we do, the expected total selection cost under the (r,μ)-constraint over the set of strategies $\mathcal{S}_n$. Then H can use at least all the admissible strategies we can use, so that his or her minimal expected total cost, $h_{r,\mu}(n)$ say, must satisfy

$$h_{r,\mu}(n) \leq v_{r,\mu}(n), \text{ for } n \geq \mu. \tag{6.2}$$

It follows from the i.i.d. assumption for the uniform $[0,1]$- random variables X_k that conditioned on $M_n = x$ all X_k's other than M_n are i.i.d. uniform random variables on $[0,x]$.

Now suppose that H has no further prophetical capacities apart from the gift of being able to foresee M_n. Then, as we shall see, H has exactly the same type of optimization problem as we would have for $n-1$ i.i.d. uniform random variables on $[0,1]$. Indeed, let $\{U_1, U_2, \cdots, U_{n-1}\}$ be the (unordered) set of the $n-1$ smallest order statistics of the sample$\{X_1, X_2, \cdots, X_n\}$, and let $X'_k = X_k/M_n$ for these $1 \leq k \leq n-1$. Since the $X'_k s$ are again i.i.d and uniform on $[0,1]$, and since there is a one-to-one correspondence between the U_k's and the X'_k's, an optimal strategy forces H to select (respectively, refuse) U_k whenever optimal behavior would force us to accept (respectively, refuse) X'_k. Our optimal value (that is our optimal total expected selection cost) for this selection problem is by definition $v_{r,\mu}(n-1)$ because it suffices to replace in equation (2.3) the number n by $n-1$.

Now, using the tower property of conditional expectations in (2.3) we have

$$v_{r,\mu}(n) = \min_{(I_1,\cdots,I_n)\in\mathcal{S}_n} \mathrm{E}\left(\mathrm{E}\left(\sum_{k=1}^{n} I_k X_k\right) \Big| M_n\right), \; n \geq \mu \geq r, \;\; r, n \in \mathbb{N}. \tag{6.3}$$

This is our optimal value whereas H, knowing M_n, can replace the conditional expectation (6.3) by the simple expectation taken on the sum of the $I_k U_k$. Hence, given $M_n = x$, the corresponding optimal value for H reduces by definition of the U_k's to $x v_{r,\mu}(n-1)$. If we denote the conditional value for H given M_n by $h_{r,\mu}(n \,|\, M_n)$ and the density of M_n by f_{M_n} then we get accordingly for $x \in [0,1]$

$$f_{M_n}(x) h_{r,\mu}(n \,|\, M_n = x) = x f_{M_n}(x) v_{r,\mu}(n-1). \tag{6.4}$$

Since the density of M_n is given by $f_{M_n}(x) = n x^{n-1}$, integrating out this equation yields the identity

$$h_{r,\mu}(n) = \int_0^1 n x^n v_{r,\mu}(n-1) dx = \frac{n}{n+1} v_{r,\mu}(n-1). \tag{6.5}$$

Using (6.2) we get from (6.5)

$$v_{r,\mu}(n) \geq h_{r,\mu}(n) = \frac{n}{n+1}\, v_{r,\mu}(n-1) \;\geq\; \frac{n-1}{n}\, v_{r,\mu}(n-1)$$

which is (6.1), and the proof is complete. □

Corollary 6.1 $(nv_{r,\mu}(n))_{n\geq\mu,\, n\in\mathbb{N}}$ *converges for all* $r\in\mathbb{N}, \mu\geq r$.

Proof: We know from Lemma 5.2 that, for all $r\in\mathbb{N}, \mu\geq r$ there exists a constant $\beta(r,\mu)$ such that $v_{r,\mu}(n)\leq\beta(r,\mu)/n$. Hence $nv_{r,\mu}(n)$ is bounded. Convergence follows thus immediately from Lemma 6.1. □

6.2 Asymptotic behavior of $nv_{r,\mu}(n)$

We are now ready to tackle the problem of how to compute $\lim_{n\to\infty} n\, v_{r,\mu}(n)$. Corollaries 3.1 and 6.1 justify trying to replace the difference equations determining $v_{r,\mu}(n)$ by a differential equation, and this approach will indeed be successful.

We rewrite (3.6) for $t\in\mathbb{N}$ and $\delta t=1$ as

$$\frac{1}{\delta t}\Big(v_{r,\mu}(t)-v_{r,\mu}(t-\delta t)\Big)=-\frac{1}{2}\Big(v_{r,\mu}(t-\delta t)-v_{r-1,\mu-1}(t-\delta t)\Big)^2 \tag{6.6}$$

with initial condition (3.7). We fix r and μ and can then simplify the notation by writing $v_{r-1,\mu-1}(t)=: v(t)$ and $v_{r,\mu}(t)=: w(t)$, say. Let $\tilde{v}(t)$ and $\tilde{w}(t)$ be differentiable functions which coincide with $v(t)$ and $w(t)$ for $t\in\mathbb{N}$ with $t\geq\mu$. It follows from Corollary 3.1 and Corollary 6.1 that the differential equation

$$\tilde{w}'(t)=-\frac{1}{2}\left(\tilde{w}(t)-\tilde{v}(t)\right)^2 \tag{6.7}$$

defined for $t\in[\mu,\infty]$ must (provided that we can solve it) catch the asymptotic behavior of both $w(t)$ and $tw(t)$. Note that equation (6.7) is a general Riccati differential equation. The idea is now to use our results on $nv_{r,\mu}(n)$ in order to show that only one class of solutions of equation (6.6) is compatible with the established properties of $v_{r,\mu}(n)$, and then that, within this class, all solutions are asymptotically equivalent.

Theorem 6.1 *If* $\tilde{v}(t)=c/t$ *for some constant* $c\geq 2$ *then the only solutions* $\tilde{w}(t)$ *of equation (15) satisfying* $\lim_{t\to\infty} v_{r,\mu}(t)/\tilde{w}(t)=1$ *are those functions which are asymptotically equivalent to*

$$\tilde{w}(t)=\frac{1}{t}\left(1+c+\sqrt{1+2c}\right).$$

Proof: We first prove that the function

$$\tilde{w}_1(t)=\left(1+c+\sqrt{1+2c}\right)/t \tag{6.8}$$

is a particular solution of equation (6.7). Indeed, there must exist a constant, c_1 say, such that c_1/t is a particular solution, because plugging this into (6.7) yields a simple quadratic equation in c_1. We get the equation

$$-\frac{c_1}{t^2}=-\frac{1}{2t^2}(c_1^2-2cc_1+c^2)$$

with the two solutions $\{1+c\pm\sqrt{1+2c}\}$. We now show that only the solution

$$c_1 = 1 + c + \sqrt{1 + 2c}$$

is meaningful for our problem. Indeed, with c being strictly positive, we would obtain $1+c-\sqrt{1+2c} < c$ otherwise. This, however, would contradict $\tilde{w}(t) \geq \tilde{v}(t)$ which we must require in order to stay compatible with the general inequality $v_{r,\mu}(n) \geq v_{r-1,\mu-1}(n)$ (see (4.1)). This shows that $\tilde{w}_1(t)$ is a particular solution of (6.7).

From the general theory of Riccati differential equations (see e.g. (H. Grauert und W. Fischer , 1968, pp.109-112)) we know that we can now generate a general solution $\{\tilde{w}_2\}$ from a particular solution by solving (substitution $u(t) = 1/(w_2(t) - w_1(t))$ the first order linear equation

$$u'(t) = -(Q(t) + 2R(t)\tilde{w}_1(t))u(t) - R(t)$$

where, in our case, $R(t) = -1/2$ and $Q(t) = c/t$. The set $\{\tilde{w}_2\}$ of solutions is then the set $\{\tilde{w}_2(t) = \tilde{w}_1(t) + u(t)^{-1}\}$ with a single undetermined constant. Plugging our particular solution $\tilde{w}_1(t)$ into the first order equation yields, after straightforward simplification, the equation $u'(t) = u(t)\left(1+\sqrt{1+2c}\right)/t + \frac{1}{2}$. We solve its associated homogeneous equation and then apply the method of the variation of constants. Using this yields after re-substitution (including an independent check by Mathematica) that all solutions which are compatible with the monotonicity of $v_{r,\mu}(n)$ (decreasing) and the monotonicity of $nv_{r,\mu}(n)$ (increasing) lie in the class

$$\mathcal{S} = \left\{ \frac{c}{t} - \frac{1}{t}\left(-1 + \sqrt{1+2c}\left(-1 + \frac{2\kappa}{t^{\sqrt{1+2c}} + \kappa}\right)\right) : \kappa \in \mathbb{R} \right\}.$$

Note that for all $s \in \mathcal{S}$ we have the asymptotic relationship

$$s(t) \sim \frac{1}{t}\left(1 + c + \sqrt{1+2c}\right) = \tilde{w}(t) \tag{6.9}$$

as $t \to \infty$, and hence the proof. □

Since we have indeed identified the only possible limiting behavior of $nv_{r,\mu}(n)$, and thus its correct limiting behavior, our main result Theorem 6.2 will now easily follow by induction.

Theorem 6.2 *Let $g : \mathbb{R}^+ \to \mathbb{R}^+$ and g_r be defined by*

(i) $g(x) = 1 + x + \sqrt{1+2x}$

(ii) $g_0(x) = g(x);\ g_r(x) = g\left(g_{r-1}(x)\right)$ *for* $r \in \mathbb{N}$.

Then for all $r \in \mathbb{N}$ and $\mu \in \mathbb{R}, \mu \geq r$,

$$\mathcal{L}(r,\mu) := \lim_{n\to\infty} n\, v_{r,\mu}(n) = g_r\left(\frac{(\mu-r)^2}{2}\right).$$

Proof: Let $\tilde{w}(t) = v_{1,\mu-r+1}(t)$ and $\tilde{v}(t) = v_{0,\mu-r}(t)$. Recall from Lemma 3.2 that $v_{0,\mu-r}(t) = (\mu - r)^2/(2t)$ so that from equation (6.7)

$$\frac{dv_{1,\mu-r+1}(t)}{dt} = -\frac{1}{2}\left(v_{1,\mu-r+1} - \frac{(\mu-r)^2}{2t}\right)^2 \tag{6.10}$$

and thus from (6.8)

$$v_{1,\mu-r+1}(t) \sim \frac{1+c_0+\sqrt{1+2c_0}}{t} =: \frac{c_1}{t}, \tag{6.11}$$

where $c_0 = (\mu - r)^2/2$. Renaming successively $\tilde{w}(t) = v_{1,\mu-r+1}(t),\ v_{2,\mu-r+2}(t),\cdots$ and $\tilde{v}(t) = v_{0,\mu-r}(t),\ v_{1,\mu-r+1}(t),\cdots$ we get by induction according to (6.8) and (6.9)

$$v_{r,\mu}(t) \ \sim\ \frac{c_r}{t} \ \text{ and } \ c_r := g(c_{r-1}) \tag{6.12}$$

where $g(x) = 1 + x + \sqrt{1+2x}$. From Corollary 6.1 and Theorem 6.1 the function g in (6.12) is unique. Hence

$$\mathcal{L}(r,\mu) = \lim_{t\to\infty} t v_{r,\mu}(t) = c_r = g_r\left(\frac{(\mu-r)^2}{2}\right) \tag{6.13}$$

and Theorem 6.2 is proved. □

6.3 Cost comparison for D-constraint and E-constraint

It seems interesting to compare the normalized limiting costs $\mathcal{L}(r,\mu) = \lim_{t\to\infty} t v_{r,\mu}(t)$ for different r and μ. It is somewhat intuitive that the D-constraint is "more costly" than the E-constraint. One feels that one has always an advantage if one is not forced to select an observation. This is true and best understood if we fix μ (integer, say) and then compare the values $\mathcal{L}(1,\mu), \mathcal{L}(2,\mu), \cdots, \mathcal{L}(\mu,\mu)$. They are increasing in r, as we know already, and indeed by increasing increments, as one can easily show (see also Table 1 below). Hence this part of our intuition is correct.

Now we may wonder whether it is cheaper to move the D-constraint up by one and keep μ fixed, or is it cheaper vice versa? (Here, as always, it is understood that the D-constraint may not strictly exceed μ.) Our feeling is now that the answer should depend on how r relates to μ: If μ is large and also much larger than r, then it is intuitive that a change from (r,μ) to $(r+1,\mu)$ should increase the normalized limiting cost relatively little. We may reason by a law of large numbers argument that even the pure E-constraint problem (that is, the problem for $(0,\mu)$) would lead, with a probability close to one to a sufficiently large number of selections covering r or $r+1$ all the like. This intuition (also visible in Table 1) is again true, as one can generally show using the properties of the function g and its iterates g_k defined in Theorem 6.2. The remaining question is, whether there are some surprises if r gets close to μ. More precisely, is it is possible that the normalized limiting cost does not change if we move from a mixed constraint (r,μ) to $(r-k,\mu+k)$ for some $k \in \mathbb{N}$. The answer is that it may coincide. However, there is only exactly one non-trivial case (i.e. $k > 0$) where the values coincide, which is the case $r = \mu$ and $k = 1$ as we shall now prove.

Propostion 6.1

$$\mathcal{L}(r-1,r+1) = \mathcal{L}(r,r). \tag{6.14}$$

Furthermore, this is the only equality of the form $\mathcal{L}(r,\mu) = \mathcal{L}(r-k,\mu+k)$ which can hold for $1 \le k \le \mu$.

Proof: Recall that $\mathcal{L}(r,\mu) = g_r((\mu-r)^2/2)$ where $g(x) = 1+x+\sqrt{1+2x}$. Hence $\mathcal{L}(r-1,r+1) = g_{r-1}(2^2/2) = g_{r-1}(2)$. Similarly we obtain $\mathcal{L}(r,r) = g_r(0)$. Now note that $g(0) = 2$ so that

$$\mathcal{L}(r,r) = g_r(0) = g_{r-1}(g(0)) = g_{r-1}(2) = \mathcal{L}(r-1,r+1). \tag{6.15}$$

				$\mathcal{L}(r,\mu)$				
2.000								
2.914	5.236							
5.236	6.527	9.623						
8.662	9.623	11.276	15.123					
13.123	13.943	15.123	17.129	21.713				
18.599	19.342	20.317	21.713	24.067	29.378			
25.083	25.779	26.642	27.770	29.378	32.077	38.108		
32.571	33.236	34.029	35.010	36.289	38.108	41.149	47.895	
41.062	41.704	42.450	43.339	44.438	45.867	47.895	51.276	58.733

Tab. 1: The values $\mathcal{L}(r,\mu)$ rounded to the 3rd decimal for $1 \le r \le \mu,\ 1 \le \mu \le 9$. Here $\mathcal{L}(r,\mu)$ figures in the rth position of the μth row of the array

Hence equality holds for $k = 1$. To see that (6.15) is the only non-trivial value, suppose that $\mathcal{L}(r,r) = \mathcal{L}(r-k, r+k)$. Then $\mathcal{L}(r,r) = g_r(0) = g_{r-k}(g_k(0))$ and $\mathcal{L}(r-k, r+k) = g_{r-k}(2k^2)$. From the definition of g we see that $g(0) = 2$, and $g(2) = 3 + \sqrt{5} < 8 = 2 \times 2^2$. Since g is increasing, all the iterated g_k are increasing. It follows by induction on k that $g_k(0) < 2k^2$ for $k = 2, 3, \cdots, r$. Therefore the equality $\mathcal{L}(r,r) = \mathcal{L}(r-k, r+k)$ cannot hold for k with $2 \le k \le r$, and hence the proof. $\square$

Acknowledgements

The author's instigation to do research on this problem is due to a question of David Aldous whose insight in our initial discussions was very helpful. Further, the author thanks Freddy Delbaen and Alexandre Dutrifoy for their interest in this problem and their comments. Special thanks go to the referee whose remarks lead to a better presentation of this paper.

References

D. Aldous, O. Angel, and N. Berestycki. Online random weight minimal spanning trees and a stochastic coalescent. Draft version June 16, 2008; private communication, 2008.

Y. M. Baryshnikov and A. V. Gnedin. Sequential selection of an increasing sequence from a mutidimensional random sample. *Ann. of Appl. Prob.*, 10(1):258–267, 2000.

F. T. Bruss and F. Delbaen. Optimal rules for the sequential selection of monotone subsequences of maximum expected length. *Stoch. Processes and Th. Applic.*, (96):313–342, 2001.

F. T. Bruss and T. S. Ferguson. Maximizing the expected rank with full information. *J. Appl. Prob.*, (30): 616–626, 1993.

F. T. Bruss and T. S. Ferguson. Multiple buying and selling with vector offers. *J. Appl. Prob.*, (34): 959–973, 1997.

F. T. Bruss, M. Drmota, and G. Louchard. The complete solution of the competitive rank selection problem. *Algorithmica*, (22):413–447, 1998.

R. W. Chen, V. N. Nair, A. M. Odlyzko, L. A. Shepp, and Y. Vardi. Optimal sequential selection of N random variables under a constraint. *J. Appl. Prob.*, 21:537–547, 1984.

E. G. Enns and E. Ferenstein. The horse game. *J. Oper. Res. Soc. Japan*, 28:51–62, 1985.

A. V. Gnedin. Sequential selection of an increasing subsequence from a random sample with geometically distributed sample size. *Lect. Notes IMS - Monograph Series*, 35:101–109, 2000. In F. Thomas Bruss and Lucien LeCam (eds).

H. Grauert und W. Fischer . *Differential- und Integralrechnung II*. Springer Verlag Berlin, Heidelberg, New York, 1968.

A. Kurushima and K. Ano. Maximizing the expected duration of owning a relatively best object in a poisson process with rankable observations. *J. Appl. Probab.*, 46(2):402–414, 2009.

D. Lebek and K. Szajowski. Optimal strategies in high risk investments. *Bull. Belg. Math. Soc.*, 14(1): 143–155, 2007.

J. MacQueen and R. G. Miller. Optimal persistence policies. *Operat. Res.*, (8):363–380, 1960.

V. V. Mazalov and M. Tamaki. An explicit formula for the optimal gain in the full-information problem of owning a relatively best object. *J. Appl. Prob.*, 43:87–101, 2006.

L. Moser. On a problem of cayley. *Scripta Math.*, (22):289–292, 1956.

S. M. Samuels. Secretary problems. *Handbook of Sequential Analysis*, pages 381–405, 1990. edited by B.K. Ghosh and P.K. Sen, Marcel Dekker, Inc., New-York, Basel, Hong Kong.

S. M. Samuels and J. M. Steele. Optimal sequential selection of a monotone sequence from a random sample. *Ann. of Probab.*, 9:937–947, 1981.

W. Stadje. A full-information pricing problem for the sale of several identical commodities. *Z. Operations Res.*, (34):161–181, 1990.

M. Tamaki. A full-information best-choice problem with finite memory. *J. Appl. Probab.*, 23(3):718–735, 1986.