The Annals of Applied Probability
1995, Vol. 5, No. 1, 310–321

A STOCHASTIC GAME OF OPTIMAL STOPPING AND ORDER SELECTION

BY ALEXANDER V. GNEDIN[1] AND ULRICH KRENGEL

University of Göttingen

We study the following two-person zero-sum game. n random numbers are drawn independently from a continuous distribution known to both players. Player 2 observes all the numbers and selects an order to present them to the opponent. Player 1 learns the numbers sequentially as they are presented and may stop learning whenever he/she pleases. If the stop occurred at the number that is the kth largest among all n numbers, Player 1 pays the amount $q(k)$ to Player 2, where $q(1) \leq \cdots \leq q(n)$ is a given payoff function. Player 1 aims to minimize the expected payoff; Player 2 aims to maximize it. We find an explicit solution of the game for a wide class of payoff functions including those q's typically considered in the context of best choice problems.

1. Introduction. A number of minimax versions of the best choice problem can be viewed as particular forms of the following two-person zero-sum game. Let $\mathscr{F}$ be a family of n-dimensional distributions F of the random vector $(Y_1, \dots, Y_n)$ with zero probability of ties and let T be a class of stopping rules with possible values in $\{1, \dots, n\}$. Let $q(1) \leq \cdots \leq q(n)$ be a nondecreasing sequence. Independently of each other, the participants make their choice: Player 1 selects $\tau \in T$ and Player 2 selects $F \in \mathscr{F}$. Player 1 observes successively the values $Y_1, \dots, Y_n$ and applies τ to stop the observation process. The loss of Player 1, hence the reward of Player 2, is $q(k)$ if Y_τ has the kth largest value among $Y_1, \dots, Y_n$. The antagonistic objectives of the participants are minimization, respectively, maximization, of the expected payoff.

The assumption that the payoff depends on observations only through their ranks, but not otherwise on their actual values, is specific for the best choice problems. Two classical goals, maximizing the probability of stopping at the largest observation and minimizing the expected rank, correspond to the payoff functions $q(k) = 1_{\{k>1\}}$ and $q(k) = k$, respectively. We will speak of the *unrestricted* game if T is the class of all nonanticipating stopping rules and of the *rank* game if T is the class of stopping rules based solely on the observation of the relative ranks.

Received July 1993; revised May 1994.

[1]Supported by Deutsche Forschungsgemeinschaft.

AMS 1991 *subject classification*. Primary 60G40.

Key words and phrases. Optimal stopping, rank, order selection, arrangement, best-choice problem, minimax strategy.

The best known game of this kind is Martin Gardner's Googol [cf. Ferguson (1989) and Gnedin (1994)]. This is an unrestricted game, where $\mathscr{F}$ is the class of all exchangeable distributions and Player 1 aims to maximize the probability of stopping at the largest observation.

Another example relevant to the subject of the present paper is the rank game where Player 2 chooses the order in which random numbers must be demonstrated to Player 1. The two classical payoff functions appear in Gilbert and Mosteller (1966) and Chow, Moriguti, Robbins and Samuels (1964). One of the possibilities to bring this model into agreement with the above scheme is to identify $\mathscr{F}$ with the family of distributions on the set of permutations, represented as integer n-vectors.

Recent results of Hill and Kennedy (1992) are naturally interpreted in the framework of the unrestricted game in which the observations are independent and Player 1 knows the distribution of $\max(Y_1, \dots, Y_n)$. Using a convexity idea, Hill and Kennedy obtained general estimates of the performance of threshold rules prescribing to stop at the first observation that exceeds a fixed level. To construct minimax distributions for certain payoff functions, they introduced a sequence of random variables, called *Bernoulli pyramid*, that generates the process of relative ranks with an independence property.

Unrestricted games with various parametric classes of exchangeable distributions have been studied in the framework of the so-called partial information problems [cf. Berezovskiy and Gnedin (1984) and Section 3 of Samuels (1991)]. Other related work appears in Hill and Krengel (1991, 1992) and Irle and Schmitz (1978).

In this paper, we study the unrestricted version of the game introduced by Gilbert and Mosteller (1966) and Chow, Moriguti, Robbins and Samuels (1964), in which Player 2 compiles $Y_1, \dots, Y_n$ by rearranging the values of an iid sequence with continuous distribution known to both players. The class $\mathscr{F}$ appearing in this model can be characterized by the property that the cumulative distribution of the order statistics of $(Y_1, \dots, Y_n)$ is the same as the distribution of the order statistics of a vector with iid components. Our main construction of the *travellers' process* yields a sequence of observations with the same rank properties as that of the Bernoulli pyramid, leading in some cases to a solution analogous to the solution of Hill and Kennedy (1992).

It is surprising that rearranging iid variables may give the same effect as selecting distributions of independent variables. This is the case for the above-mentioned classical payoff functions. We make a step forward by characterizing explicitly those payoff functions that admit a solution via Bernoulli pyramid or travellers' process.

2. Description of the game. The game is played by two participants. There are n iid continuously distributed random variables $X_1, \dots, X_n$. The integer n and the distribution of the X_i's are known to both players. Player 2 observes all random variables $X_1, \dots, X_n$ and arranges the observed values in some order. The resulting sequence $Y_1, \dots, Y_n$ is shown successively to Player 1, who must stop the process at one of the observations solely using the information collected so far. That is, having observed $Y_1, \dots, Y_i$, Player 1 must decide whether to stop at Y_i or to make the next observation, if there is at least one more observation available; in the case $i = n$, Player 1 must stop at Y_n. If the stop occurs at the observation that is the kth largest among all n observations, Player 1 has to pay the opponent the amount $q(k)$. We assume that the payoff function $q(\cdot)$ is nondecreasing and satisfies $q(1) < q(n)$. Player 1 aims to minimize the expected payoff; Player 2 wants to maximize it. We are interested in the value of the game and in the minimax strategies of both players.

The monotonicity assumption means that Player 1 prefers to select observations with small ranks, trying to recognize them in the course of the observation process. Player 2 tries to make the search as difficult as possible. We can view the game as a worst-case study of the "full-information" best-choice problem [cf. Gilbert and Mosteller (1966)] when the common assumption of arrival of the objects in random order is violated.

Because the distribution of the X_i's is known exactly, and because the ranking of variables is invariant under all monotonic transformations respecting continuity of the distribution, we lose no generality by assuming that the distribution is uniform on $I = [0, 1]$. Hence, we can identify $X = (X_1, \dots, X_n)$ with a random point of the unit n-cube sampled from the Lebesgue measure.

We first give a formal description of the available strategies and of the payoff structure.

We define a *stopping rule* to be a measurable function $\tau\colon I^n \to \{1, \dots, n\}$ such that the inverse image of any $i = 1, \dots, n$ is a subset of $I^n = I^i \times I^{n-i}$ of the form $B_i \times I^{n-i}$, where $B_i \subset I^i$. This condition formalizes the anticipating property: The decision to stop at y_i in the series $y_1, \dots, y_n$ must depend solely on $y_1, \dots, y_i$. Clearly, each stopping rule can be identified with a measurable partition of I^n into n pieces $B_i \times I^{n-i}$.

We define an *arrangement* to be a measurable mapping $\phi\colon I^n \to I^n$ satisfying $\{x_1, \dots, x_n\} = \{y_1, \dots, y_n\}$ for $y = \phi(x)$. This means that there exists a permutation s depending on x such that $y_1 = x_{s(1)}, \dots, y_n = x_{s(n)}$. For almost all x this permutation is unique. Hence, ignoring a null set, any arrangement can be seen to be a partition of I^n into $n!$ pieces.

For $y \in I^n$, $i \in \{1, \dots, n\}$, we define the *absolute rank* of y_i as

$$R_{i,n}(y) = \#\{j\colon j \le n, y_j \ge y_i\}$$

and the *relative rank* of y_i as

$$R_i(y) = \#\{j\colon j \le i, y_j \ge y_i\}.$$

Note that the ith relative rank is completely determined through comparisons of y_i with the first i coordinates.

Assume that all components of y are different. Then the sequence of absolute ranks $(R_{1,n}(y), \dots, R_{n,n}(y))$, which in this case is a permutation of n integers, uniquely determines $(R_1(y), \dots, R_n(y))$. The converse of this connection is less obvious: Absolute ranks are unambiguously reconstructable from the sequence of relative ranks.

With these definitions, the game looks as follows: Player 2 selects an arrangement ϕ and Player 1 selects a stopping rule τ. The arrangement ϕ is applied to transform X into $\phi(X) = Y = (Y_1, \dots, Y_n)$. Then τ is applied to Y. The payoff is $q(R_{\tau,n}(Y))$, where $R_{\tau,n}(Y) = R_{k,n}(Y)$ on $\{\tau = k\}$. Thus the expected payoff is

$$v(\phi, \tau) \stackrel{\text{def}}{=} Eq(R_{\tau,n}(Y)) = \sum_{i=1}^{n} \int_{\{x \in I^n:\, \tau(\phi(x)) = i\}} q(R_{i,n}(\phi(x)))\, dx.$$

Player 2 determines the distribution of Y by his choice of the arrangement. For example, if ϕ is the identity, then Y is uniformly distributed in the cube; if ϕ puts the components in increasing order, then Y is uniformly distributed in the unit simplex. By backward induction, for any ϕ, there exists an optimal counterstrategy that minimizes $v(\phi, \cdot)$. On the other hand, also for any τ, there exists ϕ maximizing $v(\cdot, \tau)$. Indeed, for each x just pick a permutation maximizing $R_{\tau,n}(y)$, taking special care of measurability.

The game will be solved if we find two strategies ϕ^*, τ^* such that ϕ^* maximizes $v(\cdot, \tau^*)$ and τ^* minimizes $v(\phi^*, \cdot)$. If such strategies exist, they are minimax and the value of the game is

$$v(\phi^*, \tau^*) = \inf_{\tau} \sup_{\phi} v(\phi, \tau) = \sup_{\phi} \inf_{\tau} v(\phi, \tau).$$

REMARK. Standard minimax theorems of the game theory, which state the existence of the value of a game, rely on convexity of the strategy spaces. Because the set of partitions of I^n has no obvious convex structure, it does not seem possible to apply the general results to the game under consideration. One would expect in this situation that the strategy spaces must be convexified properly to guarantee the existence of the value or, in other words, that *randomized* strategies should be allowed. In what follows we isolate a class of payoff functions such that the randomization is not needed, that is, both players have minimax pure strategies and the value of the game exists.

3. Threshold rules. Define the *threshold rule* with threshold $\theta \in I$ by

$$\tau_\theta(y) = \begin{cases} \min\{i\colon y_i > \theta\}, & \\ n, & \text{if there is no such } i. \end{cases}$$

That is, stop at the first exceedance over θ or proceed to the last observation.

It is easy to understand which arrangements are the best counterstrategies against τ_θ. Assume first that some $k > 1$ elements of $x_1, \dots, x_n$ exceed θ. Then Player 2 cannot prevent stopping at one of the exceedances, whose absolute rank is not greater than k. On the other hand, Player 2 forces the opponent to stop at the observation with the absolute rank equal to k by showing the smallest exceedance prior to the other exceedances. If there are no exceedances at all, then τ_θ prescribes stopping at the last observation; hence it is optimal for Player 2 to put the minimal value (= rank n) at the end of the sequence. Because the number of exceedances is binomially distributed, this argument yields

$$\max_\phi v(\phi, \tau_\theta) = \sum_{k=1}^{n} \binom{n}{k}(1-\theta)^k \theta^{n-k} q(k) + \theta^n q(n) \overset{\text{def}}{=} f(\theta). \tag{1}$$

By taking the derivative, it follows that the minimum of $f(\theta)$ in $[0,1]$ is attained by the unique value θ^* satisfying the equation

$$q(n) - q(1) = \sum_{k=1}^{n-1} \binom{n-1}{k} \left(\frac{1-\theta}{\theta}\right)^k (q(k+1) - q(k)). \tag{2}$$

Therefore, the upper value of the game satisfies

$$\inf_\tau \sup_\phi v(\phi, \tau) \le \sup_\phi v(\phi, \tau_{\theta^*}) = f(\theta^*). \tag{3}$$

4. Travellers' process. Think of the unit interval as a road connecting 0 and 1 and passing through n cities with unknown locations $x_1, \dots, x_n$. Two travellers start walking at time $t = 0$ from some point θ strictly between zero and one. They move in opposite directions with constant speeds and reach the endpoints at the same moment $t = 1$. The travellers report the locations of the cities as soon as they reach them on the road, and the locations become known as soon as they are reported.

The story above determines an arrangement, defined for almost all $x \in I^n$. Indeed, consider the function

$$t(x) = \begin{cases} \dfrac{x-\theta}{1-\theta}, & \text{for } x \in [\theta, 1], \\ \dfrac{\theta - x}{\theta}, & \text{for } x \in [0, \theta]. \end{cases} \tag{4}$$

[The function $t(x)$ is the moment when one of the travellers reaches x.] Assuming that $x = (x_1, \dots, x_n)$ has different components and

$$\frac{x_i - \theta}{1-\theta} \ne \frac{\theta - x_j}{\theta} \qquad \forall\, i, j \colon x_i \le \theta < x_j,$$

there exists a unique permutation s such that $t(x_{s(1)}) < \cdots < t(x_{s(n)})$. Define

$$\phi_\theta(x) = (x_{s(1)}, \dots, x_{s(n)}).$$

This arrangement is invariant under permutations of the arguments $x_1, \dots, x_n$. Applying ϕ_θ to the random vector X we obtain a random sequence $Y = (Y_1, \dots, Y_n) = \phi_\theta(X)$ called *the travellers' process with starting point θ.*

For $i = 1, \dots, n$, set $J_i = 1_{\{Y_i \ge \theta\}}$. An important distributional property of the travellers' process is given next.

LEMMA 1. *The random variables* $J_1, \dots, J_n$ *are iid, with* $P\{J_i = 1\} = 1 - \theta$. *For any* $i = 1, \dots, n-1$, *the variables* $J_{i+1}, \dots, J_n$ *are independent of* $Y_1, \dots, Y_i$.

PROOF. Introduce random variables $H_i = 1_{\{X_i \geq \theta\}}$ and $U_i = t(X_i)$, $i = 1, \dots, n$. Clearly, the H_i's are iid Bernoulli random variables. We prove now that the variables $H_1, \dots, H_n; U_1, \dots, U_n$ are independent. Indeed, independence of the pairs $(H_1, U_1), \dots, (H_n, U_n)$ follows from the independence of the X_i's. The variable X_i, conditioned on $\{X_i \geq \theta\}$, is uniformly distributed on the interval $[1 - \theta, 1]$. Therefore, the distribution of U_i conditioned on $\{H_i = 1\}$ is standard uniform. The same holds for conditioning on $\{H_i = 0\}$, whence H_i and U_i are independent as well.

Set $U'_i = U_j$ and $H'_i = H_j$ if U_j has the ith smallest value among $U_1, \dots, U_n$. The above independence implies that the vector $(H'_1, \dots, H'_n)$ is independent of $U'_1, \dots, U'_n$ and has the same distribution as $(H_1, \dots, H_n)$. It remains to note that $H'_i = J_i$ and Y_i is a function of H'_i and U'_i. □

With a little additional effort, we can establish an interesting regenerative property of the travellers' process. Let $\gamma_i\colon I \to I$ be the (random) nondecreasing mapping that shrinks the interval $(\theta - \theta t(Y_i), \theta + (1 - \theta)t(Y_i))$ to θ and expands linearly $[0, \theta - \theta t(Y_i)]$ to $[0, \theta]$ and $[\theta + (1 - \theta)t(Y_i), 1]$ to $[\theta, 1]$.

PROPOSITION 2. *For any* $i = 1, \dots, n - 1$, *the distribution of the vector* $(\gamma_i(Y_{i+1}), \dots, \gamma_i(Y_n))$ *conditioned on the values of* $Y_1, \dots, Y_i$ *does not depend on these values and coincides with the distribution of the* $(n - i)$ *point travellers' process, with the same starting point* θ.

PROOF. Use the above representation through the H'_i's and U'_i's and the following fact: The distribution of $n - i$ top order statistics of $(X_1, \dots, X_n)$ conditioned on the values of i bottom order statistics coincides with the cumulative distribution of the order statistics corresponding to a uniform iid sample. □

We can say that the travellers' process starts anew at each discovery, with a reduced number of points. In particular, the future of the process depends on its prehistory only through $n - i$ and $t(Y_i)$.

Note that each Y_i is either the greatest or the smallest point among $Y_1, \dots, Y_i$, depending on whether $Y_i \geq \theta$ or $Y_i < \theta$. Equivalently, the ith relative rank assumes the value 1 or i. Indeed, the traveller moving in the positive direction reports an increasing sequence located over θ, whereas the other traveller reports the points below θ in decreasing order. We have

$$R_i = J_i + i(1 - J_i) \quad \text{and} \quad R_{i,n} = R_i + \sum_{j=i+1}^{n} J_j \tag{5}$$

(for shorthand notation, the dependence on Y is omitted). By Lemma 1, the distribution of the forthcoming relative ranks (hence also of absolute ranks) conditioned on the history of the process coincides with their unconditional distribution. Observing the travellers' process provides no information about the ranking of future items.

5. A solution. The argument of Section 3 shows that ϕ_θ is always an optimal counterstrategy against τ_θ. Our principal result is essentially an answer to the reciprocal question regarding the optimality of the threshold rule as a counterstrategy against ϕ_θ.

THEOREM 3. *The pair of strategies ϕ_θ, τ_θ provides a solution of the game iff $\theta = \theta^*$ and (in addition to the required monotonicity) the payoff function satisfies the inequalities*

$$q(n) - q(i) \le \sum_{k=1}^{n-i} \binom{n-i}{k} \left(\frac{1-\theta^*}{\theta^*}\right)^k \times (q(k+i) - q(k)), \quad i = 2, \dots, n-1, \tag{6}$$

where θ^* is found from (2). *In this case, the value of the game equals $f(\theta^*)$.*

PROOF. Consider the problem of optimal stopping of the travellers' process $Y_1, \dots, Y_n$, with fixed starting point θ. By Lemma 1 and by (5), the expected payoff for stopping at the ith observation is

$$E\big(q(R_{i,n}) \mid Y_1, \dots, Y_i\big) = a_i J_i + b_i(1 - J_i),$$

where

$$a_i = \sum_{k=0}^{n-i} \binom{n-i}{k} (1-\theta)^k \theta^{n-i-k} q(k+1),$$

$$b_i = \sum_{k=0}^{n-i} \binom{n-i}{k} (1-\theta)^k \theta^{n-i-k} q(k+i).$$

Define stopping rules

$$\tau(i) = \begin{cases} \min\{k : k \ge i, Y_k > \theta\}, \\ n, \quad \text{if there is no such } k. \end{cases}$$

All $\tau(i)$ are relative-rank-based rules, except $\tau(1) = \tau_\theta$. Arguing as above we show that the conditional payoff for exploiting $\tau(i)$ is

$$E\big(q(R_{\tau(i),n}) \mid Y_1, \dots, Y_{i-1}\big) = w_i,$$

where

$$w_i = \sum_{k=1}^{n-i+1} \binom{n-i+1}{k} (1-\theta)^k \theta^{n-i+1-k} q(k) + \theta^{n-i+1} q(n).$$

Now assume that τ_θ minimizes $v(\cdot, \phi_\theta)$. Then

$$E\big(q(R_{i,n}) \mid Y_1, \dots, Y_i\big) \le E\Big(q\big(R_{\tau(i+1),n}\big) \mid Y_1, \dots, Y_i\Big) \quad \text{on } \{\tau_\theta = i\},$$

because stopping with τ_θ should be better than proceeding with $\tau(i+1)$, and also

$$E\Big(q\big(R_{\tau(i+1),n}\big) \mid Y_1, \dots, Y_i\Big) \le E\big(q(R_{i,n}) \mid Y_1, \dots, Y_i\big) \quad \text{on } \{\tau_\theta > i\},$$

because proceeding with τ_θ should be better than stopping and $\tau_\theta = \tau(i+1)$ on $\{\tau_\theta > i\}$. Because τ_θ accepts any of the values from 1 to n with positive probability, we have

$$(7) \qquad a_i \le w_{i+1} \le b_i, \qquad i = 1, \dots, n-1.$$

In particular,

$$(8) \qquad a_1 = b_1 = w_2$$

whence the rule $\tau(2)$ and the constant rule $\tau \equiv 1$ are both optimal as well.

Conversely, assume (7). Letting i run backward from $n-1$ to 1 we prove the optimality of each $\tau(i)$ among the rules always passing the first $i-1$ observations. This implies that all three rules $\tau \equiv 1$, $\tau(2)$ and τ_θ are optimal, with $Eq(R_{\tau_\theta}) = Eq(R_1) = a_1$.

We have proved that (7) holds iff τ_θ is optimal. Substituting the formulas for a_1 and w_2 into (8) we arrive at the familiar equation (2). Hence optimality can take place only if $\theta = \theta^*$. Further, for $i = 2, \dots, n-1$, the inequalities $a_i \le w_{i+1}$ can be readily written as

$$q(n) - q(1) \ge \sum_{k=1}^{n-i} \binom{n-i}{k} \left(\frac{1-\theta}{\theta}\right)^k (q(k+1) - q(k))$$

and are always satisfied for $\theta = \theta^*$, by (2). Thus (7) is equivalent to the system

$$a_1 = w_2, \qquad b_i \ge w_{i+1} \quad \text{for } i = 2, \dots, n-1.$$

Transforming the last inequalities, we obtain (6).

To complete the proof recall the upper bound (3). □

REMARK. We have seen that if τ_θ is optimal against ϕ_θ, then there are also other optimal rules (in fact, infinitely many of them). Similarly, there are infinitely many optimal counterstrategies against τ_θ. This, however, in no way answers the question whether the solution of the game is unique.

EXAMPLES. For the two classical payoffs the problem is solved.

For the payoff function $q(k) = 1_{\{k>1\}}$, the optimal threshold value is $\theta^* = 1 - 1/n$ and (6) is obviously satisfied. Consequently, the minimax probability of stopping at the largest observation is

$$1 - f(\theta^*) = (1 - 1/n)^{n-1},$$

tending to the famous e^{-1} as $n \to \infty$. Thus asymptotically the probability of the best choice is the same as in the standard problem, though for finite n these values are different.

For the expected rank problem, where $q(k) = k$, we have $\theta^* = n^{-1/(n-1)}$. It is easy to reduce (6) to the inequality $i^{1/(i-1)} \geq n^{1/(n-1)}$, which holds by the monotonicity of the function $z^{1/(z-1)}$ for $z > 1$. The minimax expected rank is

$$1 + (n-1)(1 - n^{-1/(n-1)}) \sim \log n.$$

The analogy with the standard problem of minimizing the expected rank is no longer the case [the classical asymptotic value is 3.8695..., as found in Chow, Moriguti, Robbins and Samuels (1964)].

It is easy to check that our solution works for all payoff functions for $n = 2, 3$ and, with more effort, also for $n = 4$.

6. Bernoulli pyramid. In what follows we establish a link with the results of Hill and Kennedy (1992). This will be done by exploiting the travellers' process to solve a more general optimal stopping problem.

Let $Y_1, \dots, Y_n$ be an arbitrary sequence of continuous random variables that has the same rank properties as the travellers' process with starting point θ, that is, the associated relative ranks satisfy

$$(9) \qquad R_{i+1}, \dots, R_n \text{ are independent of } Y_1, \dots, Y_i, \qquad i = 1, \dots, n-1,$$

and

$$(10) \qquad P\{R_i = i\} = 1 - P\{R_i = 1\} = \theta, \qquad i = 2, \dots, n.$$

These properties guarantee that there exists a stopping rule based solely on the relative ranks, which minimizes $Eq(R_{\tau,n})$ among all stopping rules based on the Y_i's, as it follows from a well-known fact of the stopping theory about randomized rules. Thus the upper bound

$$\inf_{\tau} Eq(R_{\tau,n}) \leq f(\theta^*)$$

holds in general *because it is valid for the travellers' process* [recall (3)], and the infinum depends only on $q(\cdot)$ and θ. The same idea works in the extreme case.

THEOREM 4. *Assume the variables* $Y_1, \dots, Y_n$ *satisfy* (9) *and* (10) *for some* θ. *The following are equivalent*:

(i) $\inf_\tau Eq(R_{\tau,n}) = f(\theta^*)$.
(ii) $\theta = \theta^*$ *and the payoff function satisfies* (6).

PROOF. Assume (i). Then τ_{θ^*} is the optimal stopping rule for the travellers' process with starting point θ, and (ii) follows by Theorem 3. The converse is similar. □

Denote

$$G(y) = P\{\max(Y_1, \dots, Y_n) \le y\}$$

the distribution function of the maximal observation. Extending our old definition of the threshold rule set,

$$\tau_\theta(y) = \begin{cases} \min\{i: y_i > G^{\leftarrow}(\theta^n)\}, \\ n, \qquad \text{if there is no such } i, \end{cases}$$

where $G^{\leftarrow}(\cdot)$ denotes the generalized inverse.

Consider the following game. Player 2 selects the distributions $F_1, \dots, F_n$ of independent random variables $Y_1, \dots, Y_n$ and tells the distribution of the maximum, $G = \prod F_i$, to Player 1. Player 1 selects a stopping rule and applies it to $Y = (Y_1, \dots, Y_n)$. The expected loss of Player 1 is

$$v(F_1, \dots, F_n; \tau) \stackrel{\text{def}}{=} Eq(R_{\tau, n}(Y)).$$

(Using the transformation technique we can reduce the game to the case where G is some fixed distribution, known to both players, and Player 2 picks the F_i's subject to the constraint $G = \prod F_i$. A particular form of G does not matter.)

Theorem 1.1 by Hill and Kennedy (1992) asserts that

$$\sup_{F_1, \dots, F_n} v(F_1, \dots, F_n; \tau_{\theta^*}) \le f(\theta^*), \tag{11}$$

providing an estimate of the upper value of the game.

Extending the original construction of Hill and Kennedy (1992), we define the *Bernoulli pyramid* with parameter $\theta \in [0,1]$ to be a sequence of independent random variables satisfying (9) and (10). For example, take independent variables such that Y_1 is uniformly distributed on $[0,1]$ and Y_i is uniformly distributed on

$$[1 + (i-2)\theta, 1 + (i-1)\theta] \cup [-(i-2)(1-\theta), -(i-1)(1-\theta)]$$

for $i = 2, \dots, n$.

Combining (11) with Theorem 4 yields the following result: The pair (Bernoulli pyramid, threshold rule) provides a solution of the game (when Player 2 selects the distributions of independent observations) iff the parameters of both strategies equal to θ^* and (6) holds.

As was shown by Hill and Kennedy (1992), a solution is of this form for the exponential payoff $q(k) = -z^k$, $z \in (0,1)$, and for $q(k) = 1_{\{k > m\}}$ (maximizing the probability of picking one of the m largest items). It follows that these payoff functions satisfy (6).

It follows from Hill and Kennedy (1992) that there is no solution of this form for $(q(1), \dots, q(5)) = (0, 1, 1, 1, 2)$; hence, (6) does not hold in this case (this can be checked directly).

A posteriori, these examples apply, via the travellers' process, to the problem of optimal stopping and order selection.

7. Final remarks. Hill and Kennedy (1992) allow for the possibility to pass all the observations without stopping, in which case Player 1 pays an amount $q(0)$ [the interesting case is when $q(1) \leq q(0) \leq q(1)$]. This generalization can be treated by our methods after appropriate modification of (2) and (6).

The referee suggested the following version of the game. Player 1 selects a stopping rule. Player 2 selects distributions of independent variables, with a fixed distribution of the maximal value *and* the order to present the values to the opponent. Under (6) the threshold rule is again minimax, whereas Player 2 has two surprisingly different minimax strategies: the travellers' process and the Bernoulli pyramid.

It is interesting to compare these results with the rank game, in which Player 2 can arrange the sequence. It is known that the value is always $(q(1) + \cdots + q(n))/n$, the (randomized) minimax strategy of Player 1 is stopping at random with equal probabilities at any of the observations and there is an optimal arrangement inducing independent relative ranks [cf. Samuels (1991), Section 4.2.1].

If Player 2 is not limited at all in the choice of distributions of the observations, then stopping at random happens to be minimax also in the analogous unrestricted game. A worst-case sequence can be taken Markovian, with the same rank properties as in Samuels (1991).

A number of intriguing questions require further attention. Is the above solution unqiue, and how does it look like if (6) violated? Is the threshold rule optimal under the assumption that the induced rank process satisfies (9) and (10), and (6) holds?

Acknowledgment. The initial version of the paper was devoted solely to the classical payoff "either the best or nothing." We are thankful to the referee whose comments stimulated the general solution and helped to improve the presentation. Our special thanks go to Douglas Kennedy for his comments on the case $n = 4$.

REFERENCES

BEREZOVSKIY, B. A. and GNEDIN, A. V. (1984). *The Best Choice Problem.* Nauka, Moscow.

CHOW, Y. S., MORIGUTI, S., ROBBINS, H. and SAMUELS, S. M. (1964). Optimal selection based on relative rank (the "secretary problem"). *Israel J. Math.* **2** 81–90.

FERGUSON, T. (1989). Who solved the secretary problem? *Statist. Sci.* **4** 282–296.

GILBERT, J. and MOSTELLER, F. (1966). Recognizing the maximum of a sequence. *J. Amer. Statist. Assoc.* **61** 35–73.

GNEDIN, A. V. (1994). Pick the largest number: A solution to the game of googol. *Ann. Probab.* **22** 1588–1595.

HILL, T. P. and KENNEDY, D. P. (1992). Sharp inequalities for optimal stopping with rewards based on ranks. *Ann. Appl. Probab.* **2** 503–517.

HILL, T. P. and KRENGEL, U. (1991). Minimax-optimal stop rules and distributions in secretary problems. *Ann. Probab.* **19** 342–353.

HILL, T. P. and KRENGEL, U. (1992). On the game of Googol. *Int. J. Game Theory* **21** 151–160.

IRLE, A. and SCHMITZ, N. (1978). Minimax strategies for discounted "secretary problems." *Meth. Oper. Res.* **30** 77–86.

SAMUELS, S. M. (1991). Secretary problems. In *Handbook of Sequential Analysis* (B. K. Ghosh and P. K. Sen, eds.) 381–402. Dekker, New York.

INSTITUTE FÜR MATHEMATISCHE STOCHASTIK
UNIVERSITY OF GÖTTINGEN
LOTZESTRASSE 13
37083 GÖTTINGEN
GERMANY