Uncertainty and sequential choice*

John A. Bather
Mathematics Division
University of Sussex

§1. Introduction

Suppose that you are offered the choice of one from a set of n possible prizes. Each prize has a value represented by a random variable. We shall compare three different rules for revealing particular values to you. Let us assume that the random variables $X_1, X_2, \ldots, X_n$, representing the values, are non-negative and independent of one another. Their probability distributions are given, so we know the expected values in advance. Let

$$\mu_j = E(X_j)$$

for $j = 1, 2, \ldots, n$. Each $\mu_j \geq 0$ and we suppose that all these expectations are finite.

The first rule to be considered is that you may inspect all the prizes before making your choice. In this case, it is clear that the maximum possible expectation is

$$u_n = E\{\max(X_1, X_2, \ldots, X_n)\}.$$

Now suppose that you must make up your mind without inspecting any of the prizes. Then the best you can do is to find the maximum of $\mu_1, \mu_2, \ldots, \mu_n$ and the expected reward is

$$w_n = \max(\mu_1, \mu_2, \ldots, \mu_n).$$

The third rule is more interesting : suppose that you see the prizes one at a time but, at each stage, you must reject all those examined so far before

---

* Written version of invited lecture delivered to the Singapore Mathematical Society and the Department of Mathematics, National University of Singapore, on 11 April 1990.

you are allowed to inspect another one. For convenience, let us imagine that the values $X_n, X_{n-1}, \ldots, X_1$ are revealed sequentially in that order, until you decide to stop and select the one you have just observed. Let v_n be the maximum expectation obtained by using an optimal stopping procedure. We have $v_1 = \mu$, and $v_2 = E\{\max(X_2, v_1)\}$. This follows by considering the decision whether to accept X_2 or not. In general,

$$v_n = E\{\max(X_n, v_{n-1})\}. \tag{1}$$

It is intuitively obvious and not difficult to prove that

$$u_n \geq v_n \geq w_n \tag{2}$$

for $n \geq 1$, with equality when $n = 1$. As we shall see, u_n and v_n can be much larger than w_n, even when all the underlying random variables have the same distribution.

The aim here is to illustrate the growth of u_n and v_n as n increases by giving some examples. Then we shall investigate how much of an advantage can be obtained from having full information, rather than sequential information. The main result is that

$$u_n \leq 2v_n \tag{3}$$

always. This relation was conjectured by R. Haydon, who discovered it while investigating problems in mathematical analysis not directly related to probability. He mentioned it during a seminar at Sussex in 1989 and one of my colleagues, G. B. Trustrum, produced a proof soon afterwards.

It now seems surprising that the general inequality (3) was not discovered earlier. Many sequential decision procedures and optimal choice problems have been investigated in recent years : for example, see [1] or [2]. The proof given later is based on analysing the recurrence relation (1). The relation can also be used in a different way to find optimal policies for selling an asset and this will be the final topic.

§2. Preliminaries

A random variable X can be described by a probability density function f or by specifying discrete probabilities at certain points. We can include both types by using the distribution function F defined by

$F(x) = P(X \leq x)$ for all real x. For non-negative random variables $F(x) = 0$ when $x < 0$ and, if X has a continuous density f,

$$F(x) = \int_0^x f(y)dy, \quad F'(x) = f(x),$$

for $x > 0$. In such cases, the mean μ is given by

$$E(X) = \int_0^\infty (1 - F(x))dx. \tag{4}$$

This will be useful because it can be extended to show that, for any positive constant v,

$$E\{\max(X, v)\} = v + \int_v^\infty (1 - F(x))dx. \tag{5}$$

The first two examples illustrate the behaviour of the sequences $\{u_n\}$ and $\{v_n\}$ when the independent random variables $X_1, X_2, \ldots, X_n$ all have the same density f and distribution function F. Note that the distribution function of $\max(X_1, X_2, \ldots, X_n)$ is given by the product law:

$$P(\max(X_1, X_2, \ldots, X_n) \leq x) = P(X_1 \leq x) \ldots P(X_n \leq x) = F^n(x).$$

Hence, by using (4),

$$u_n = \int_0^\infty (1 - F^n(x))dx. \tag{6}$$

Example 1: Let $f(x) = 1$ for $0 \leq x \leq 1$ and $f(x) = 0$, otherwise. Thus, $X_1, X_2, \ldots, X_n$ are uniformly distributed on the interval [0,1].

We have $F(x) = x$ for $0 \leq x \leq 1$ and $F(x) = 1$ for $x \geq 1$, so (6) reduces to

$$u_n = \int_0^1 (1 - x^n)dx = \frac{n}{n+1}.$$

In this case, the common mean of all the random variables is $\mu = \frac{1}{2}$ so $w_n = \frac{1}{2}$ for all $n \geq 1$. The sequence $\{v_n\}$ starts with $v_1 = \frac{1}{2}$ and then for $n \geq 2$,

$$v_n = v_{n-1} + \int_{v_{n-1}}^1 (1 - x)dx = \frac{1}{2}(1 + v_{n-1}^2),$$

by using (1) and (5). This formula can be used to calculate v_2, v_3 etc., and we obtain the following table:

n	1	2	3	4	5	10	20
u_n	.500	.667	.750	.800	.833	.909	.952
v_n	.500	.625	.695	.742	.775	.861	.920

It can be shown that $v_n \geq (n+1)/(n+3)$ and this bound is close, for large n. In other words,

$$u_n = 1 - \frac{1}{n+1} \quad \text{and} \quad v_n \approx 1 - \frac{2}{n+3}.$$

Example 2: Take $f(x) = e^{-x}$ and $F(x) = 1 - e^{-x}$ for $x \geq 0$. Here, the mean $\mu = 1$, so $w_n = 1$ for all $n \geq 1$. The formula (6) shows that

$$u_n = \int_0^\infty \{1 - (1 - e^{-x})^n\} dx$$

and this can be evaluated by using the substitution $y = 1 - e^{-x}$.

We find that

$$u_n = 1 + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{n}.$$

The recurrence relation for the sequential gain v_n is easily obtained from (5):

$$v_n = v_{n-1} + e^{-v_{n-1}}$$

for $n \geq 2$ and $v_1 = u_1 = 1$. In this example, both u_n and v_n are of order $\log n$ for larger n. It can be shown that

$$\log(a+n) \leq v_n \leq u_n \leq 1 + \log n$$

always holds, where $a = e - 1 \approx 1.718$. The sequence $\{u_n\}$ is well known: in particular, the difference $u_n - \log n$ converges to a limit $\gamma \approx 0.577$, known as Euler's constant. It can also be shown, with some difficulty, that $u_n - v_n$ increases very slowly with n and $u_n - v_n \to \gamma$ as $n \to \infty$.

We now consider a simple discrete distribution which shows that u_n and v_n can be of order n. Note that, since

$$\max(X_1, X_2, \ldots, X_n) \leq X_1 + X_2 + \ldots + X_n,$$

the corresponding expectations must satisfy $u_n \leq \mu_1 + \mu_2 + \ldots + \mu_n$ and so $u_n \leq n\mu$ when each $\mu_j = \mu$.

Example 3: Let K be a large constant and suppose that each of the random variables $X_1, X_2, \ldots, X_n$ takes the values 0 and K with probabilities $1 - \frac{1}{K}$ and $\frac{1}{K}$, respectively. Then $\mu_j = 1$ always. In general, $w_n = 1$ and we also know that $u_n \leq n$. Now $\max(X_1, X_2, \ldots, X_n)$ has only two possible values here, 0 and K. Hence,

$$u_n = K\{1 - P(X_1 = X_2 = \ldots = X_n = 0)\},$$
$$u_n = K\{1 - (1 - \frac{1}{K})^n\}.$$

It is also easy to check that, for this example $v_n = u_n$ for each $n \geq 1$. This can be established by using (1). Since v_{n-1} always lies between 0 and K, we have

$$v_n = 1 + (1 - \frac{1}{K})v_{n-1}.$$

This holds for $n = 2, 3, \ldots$ with $v_1 = 1$ and it can be shown by induction that

$$v_n = u_n = K\{1 - (1 - \frac{1}{K})^n\}.$$

For any fixed $n \geq 1$, this expression can be made as close as we please to n by choosing a sufficiently large value for the constant K.

The next two examples suggest lower bounds for the ratio v_n/u_n. Both are concerned with the case $n = 2$. In Example 4, the random variables X_1 and X_2 have the same distribution and the distribution is chosen to indicate the lowest possible value of v_2/u_2. Example 5 shows that when X_1 and X_2 are allowed to have different distributions, v_2/u_2 can be reduced to levels arbitrarily close to $\frac{1}{2}$ and, according to (3), $v_n/u_n \geq \frac{1}{2}$ always holds.

Example 4: As before, let K be a large positive constant and let p be a probability, $0 < p < 1$. Suppose that X_1 and X_2 are independent, each having possible values 0, 1 or K. The corresponding probabilities are

$$P(X_j = 0) = p, \quad P(X_j = 1) = 1 - \frac{pK}{K-1}, \quad P(X_j = K) = \frac{p}{K-1}.$$

It follows that $E(X_j) = 1$, for $j = 1, 2$, and it is easy to check that

$$u_2 = E\{\max(X_1, X_2)\} = 1 + 2p - p^2 - \frac{p^2}{K-1},$$
$$v_2 = 1 + p.$$

Then we have

$$\frac{u_2}{v_2} \leq \frac{1+2p-p^2}{1+p} = 1 + \frac{p(1-p)}{1+p}$$

and it is clear that we can choose K to get u_2/v_2 arbitrarily close to this. Now it is a straightforward matter to maximise the expression $p(1-p)/(1+p)$ with respect to p. The maximum occurs when $p = \sqrt{2}-1$ and we then obtain $u_2/v_2 \leq 4 - 2\sqrt{2} \approx 1.172$. This means that

$$\frac{v_2}{u_2} \geq \frac{2+\sqrt{2}}{4} \approx 0.854.$$

This example shows that the ratio v_2/u_2 can be very close to $(2+\sqrt{2})/4$. In fact, it is impossible to get below this level if the random variables X_1 and X_2 have the same distribution, but the proof will not be attempted here.

Example 5: Let X_1 take the values 0 or K with probabilities $1 - 1/K$ and $1/K$, respectively. Thus, $\mu_1 = 1$ and we also have $u_1 = v_1 = 1$. Now define the distribution of X_2 so that

$$P(X_2 = 1) = 1 - \frac{\delta}{K-1}, \quad P(X_2 = K) = \frac{\delta}{K-1},$$

where δ is small. Since $v_2 = E\{\max(X_2, v_1)\}$ and $X_2 \geq v_1 = 1$ always, we have

$$v_2 = \mu_2 = 1 + \delta.$$

Notice that $\max(X_1, X_2) = K$, except when $X_1 = 0$ and $X_2 = 1$. This makes it easy to evaluate

$$u_2 = 2 + \delta - \frac{1+\delta}{K}.$$

The expressions for v_2 and u_2 make clear that the ratio v_2/u_2 can be reduced towards $\frac{1}{2}$ by choosing δ very small and K sufficiently large.

§3. Haydon's Conjecture

We can extend Example 5 to show that v_n/u_n can be made arbitrarily close to $\frac{1}{2}$, for any given $n > 2$. Let $X_j = 0$ with probability 1 for each $j = 3, 4, \ldots, n$. Then it is obvious that $u_n = u_2$ and $v_n = v_2$, so $v_n/u_n = v_2/u_2 \approx \frac{1}{2}$, as before. We now establish that the inequality $v_n/u_n \geq \frac{1}{2}$ always holds. This is an immediate consequence of the following proposition. The proof is due to G. B. Trustrum.

Proposition: *Let $X_1, X_2, \ldots, X_n$ be any non-negative random variables with finite means $\mu_1, \mu_2, \ldots, \mu_n$. Let $v_1 = \mu_1$ and define $v_2 = E\{\max(X_2, v_1)\}, \ldots, v_n = E\{\max(X_n, v_{n-1})\}$. Then*

$$u_n = E\{\max(X_1, X_2, \ldots, X_n)\} \leq v_{n-1} + v_n \leq 2v_n.$$

Proof: Let $Y_j = \max(X_j, v_{j-1})$ for $j \leq n$, where $v_0 = 0$ and $Y_1 = X_1$. We also define $Z_j = \max(X_j - v_{j-1}, 0)$, so that $Y_j = v_{j-1} + Z_j$. Now use the fact that $v_0 \leq v_1 \leq \ldots \leq v_{n-1}$ to obtain

$$Y_j \leq v_{n-1} + Z_j$$

for each $j \leq n$. It is clear that $Y_j \geq X_j$ always. Hence,

$$\max(X_1, X_2, \ldots, X_n) \leq \max(Y_1, Y_2, \ldots, Y_n)$$
$$\leq v_{n-1} + \max(Z_1, Z_2, \ldots, Z_n),$$
$$u_n = E\{\max(X_1, X_2, \ldots, X_n)\} \leq v_{n-1} + E\{\max Z_1, Z_2, \ldots, Z_n)\}.$$

Since every $Z_j \geq 0$, we have

$$\max(Z_1, Z_2, \ldots, Z_n) \leq Z_1 + Z_2 + \ldots, Z_n.$$

But $Z_j = Y_j - v_{j-1}$ and $E(Z_j) = v_j - v_{j-1}$ from the definitions of Y_j and v_j. It follows that

$$E\{\max(Z_1, Z_2, \ldots, Z_n)\} \leq v_1 - v_0 + v_2 - v_1 + \ldots + v_n - v_{n-1} = v_n.$$

This, together with the earlier inequality for u_n, shows that $u_n \leq v_{n-1} + v_n$ and the proof is complete. □

Previously, we have assumed that the random variables $X_1, X_2, \ldots, X_n$ are independent of one another, but the above proposition holds more generally. The independence assumption is needed for the interpretation of v_n as the maximum expected gain that can be obtained by choosing one of the values $X_n, X_{n-1}, \ldots, X_1$, when they are observed sequentially.

§4. Selling an Asset

Our final topic is based on a slightly different interpretation of the sequence $X_1, X_2, \ldots$. Imagine that you have a valuable object for sale and suppose that offers for it arrive independently at times $1, 2, \ldots$. These are represented by the random variables $X_1, X_2, \ldots$, and you may continue to receive them for as long as you wish before deciding to accept one of them. There is a cost $c > 0$ per unit time for waiting. At each stage, you must decide whether to accept the latest offer or to reject it and wait for the next one at cost c.

Suppose that the independent offers have a common probability density f and distribution function F ; let X represent the latest one. The optimal decision rule for dealing with it does not depend explicitly on the time. We define $g(x)$, for $x \geq 0$, to be the maximum expected net gain that you can obtain, given that you have just observed $X = x$. If you accept it, the gain is x and no further offers are received. Otherwise, you must wait for the next offer Y and then maximise your expectation, given $Y = y$, say. It follows that

$$g(x) = \max\{x, \int_0^\infty g(y)f(y)dy - c\}. \tag{7}$$

Notice that the second term on the right does not depend on x: the constant

$$\lambda = \int_0^\infty g(y)f(y)dy - c \tag{8}$$

represents the best you can do after rejecting x. We now have $g(x) = \max(x, \lambda)$ and the integral can be simplified by using (5):

$$\int_0^\infty g(y)f(y)dy = E\{\max(Y, \lambda)\} = \lambda + \int_\lambda^\infty (1 - F(y))dy$$

Hence, (8) reduces to

$$\int_\lambda^\infty (1 - F(y))dy = c. \tag{9}$$

The equation (9), together with the special form of the function $g(x) = \max(x, \lambda)$, determine the optimal decision rule for this problem. It is easy to check that, for any positive constant c with $c < \mu = \int_0^\infty (1 - F(y))dy$, there is a unique $\lambda > 0$ that satisfies (9). Then the rule is : accept the first offer that exceeds λ. In other words, offers $X_1 < \lambda$, $X_2 < \lambda$, ... should be rejected until the first time n at which $X_n \geq \lambda$. The expected net gain for this policy, starting at time zero and allowing for the cost of waiting, is just $g(0) = \lambda$.

For the uniform distribution of Example 1, the mean offer is $\mu = \frac{1}{2}$ and the solution of (9) is

$$\lambda = 1 - \sqrt{2c}$$

for any $c < \frac{1}{2}$. Similarly, for the exponential distribution of Example 2, we obtain $\mu = 1$ and if $c < 1$, the critical level is

$$\lambda = -\log c.$$

Note that λ represents the effective value of your asset which depends on c and also on the distribution of the offers.

References

[1] T. S. Ferguson, Who solved the secretary problem? *Statistical Science* 1989, Vol 4, no 3, 282-296.

[2] P. Whittle, *Optimization over time*, Vol 1, 1982, Vol 2, 1983, Wiley.

Biographical Note: Professor Bather was born in 1936 and obtained his Ph.D. from Cambridge in 1963. He has been a professor at the University of Sussex since 1969 and has also held visiting positions in USA and Canada. This was his second visit to Singapore as External Examiner in Mathematics at the National University of Singapore. Professor Bather's research interests are in probability, statistics and operational research, especially in making sequential decisions.