

The Secretary Problem with a Hazard Rate Condition

Mohammad Mahdian, R. Preston McAfee, and David Pennock

Yahoo! Research

{mahdian,mcafee,pennockd}@yahoo-inc.com

Abstract. In the classical secretary problem, the objective is to select the candidate of maximum value among a set of n candidates arriving one by one. The value of the candidates come from an unknown distribution and is revealed at the time the candidate arrives, at which point an irrevocable decision on whether to select the candidate must be made. The well-known solution to this problem, due to Dynkin, waits for n/e steps to set an “aspiration level” equal to the maximum value of the candidates seen, and then accepts the first candidate whose value exceeds this level. This guarantees a probability of at least $1/e$ of selecting the maximum value candidate, and there are distributions for which this is essentially the best possible. One feature of this algorithm that seems at odds with reality is that it prescribes a long waiting period before selecting a candidate. In this paper, we show that if a standard hazard rate condition is imposed on the distribution of values, the waiting period falls from n/e to $n/\log(n)$, meaning that it is enough to observe a diminishingly small sample to set the optimal aspiration level. This result is tight, as both the hazard condition and the optimal sampling period bind exactly for the exponential distribution.

1 Introduction

The classical secretary problem, dating back to late 1950’s and early 1960’s, asks whether one can hire the best secretary among a set of n candidates with unknown values arriving for interviews in a random order. Only one offer can be made, and the decision about hiring a candidate has to be made immediately after the interview and before seeing the next candidate. The surprising result (often attributed to Dynkin [2]) is that it is indeed possible to select the best secretary with a constant probability using the following strategy: do not make any offer to the first n/e candidates that arrive, and after that, make an offer to the first candidate whose value exceeds the value of all the candidates seen so far. In other words, the algorithm starts with an *exploration phase* that sees the first n/e candidates and sets an *aspiration level* equal to the highest value seen in the exploration phase. After that, the algorithm hires the first candidates that exceeds the aspiration level. We call such a strategy an “*aspiration strategy*”. Many variants of the secretary problem and this solution, and their applications to problems such as online mechanism design for auctions are studied in the literature; see, for example, the survey papers [1, 4].

An alternative formulation of the secretary problem (and according to Ferguson [3], the formulation closer to the original “game of googol” puzzle due to Martin Gardner in 1960 [5]) is as follows: An adversary selects the distribution from which the n values will be drawn independently.¹ The algorithm learns about this distribution (but not the actual values), and then sees the values one by one. The algorithm must decide when to stop, at which point it receives the last value seen. The objective is to maximize the probability that this value is the maximum value, or, equivalently, the expectation of the ratio of this value and the ex post maximum value in the sequence. As observed by Ferguson [3], this formulation of the problem is equivalent to the original one, and the same aspiration strategy achieves a factor of $1/e$. Furthermore, this result is tight, that is, there is a distribution of values for which, even knowing the distribution, the optimal strategy is the aspiration strategy with a factor of $1/e$.

One unsatisfactory feature of the aspiration strategy is that it requires a long exploration phase (observing a constant fraction of the candidates) before setting the aspiration level. One might suspect that this is due to the fact that there is no penalty (except the possibility of missing the maximum value) in the secretary problem for waiting, but surprisingly, similar results hold for variants of the secretary problem where waiting is punished by discounting the values or by setting an explicit cost for each period that no one is hired [4, 7]. The main message of this paper is to show that a long exploration phase is indeed an artifact of the assumption that the values are coming from a worst-case distribution. We show that if we impose a hazard rate condition, frequently used in engineering analyses and the analysis of auctions, the picture changes drastically: even the worst value distribution satisfying this condition requires an exploration phase of length $n/\log(n)$, significantly smaller than n/e .

2 The Model

Consider a setting where one needs to hire a secretary among a set of n potential candidates. These potential candidates arrive one per time period. Candidate i has a value v_i . v_i 's are independently and identically distributed according to a distribution F with cdf f . The value of v_i is revealed when the i 'th candidate arrives. The objective is to give a *termination rule* (i.e., an algorithm that after observing the first i values decides whether to stop or to continue) that stops at the candidate with the maximum value among the n candidates. The algorithm does not have any knowledge of the distribution F .

A standard approach to solve the above problem is to collect data by observing v_i 's for k periods, set an *aspiration level* equal to the maximum value

¹ To be precise, Gardner's formulation does not assume the values are selected independently; instead, it assumes that the adversary selects the values from any joint distribution that is permutation invariant (exchangeable). However, the assumption that the values are chosen independently is nearly without loss of generality since an infinite sequence of exchangeable random variables is conditionally independent by Kingman's theorem [6].

observed during those k periods, and then accept the first candidate observed with value exceeding the aspiration level, or the last candidate. We call such a strategy an *aspiration strategy*. A classical result due to Dynkin [2] shows that an aspiration strategy with $k = n/e$ picks the candidate with the maximum value among all n candidates with probability at least $1/e$. Furthermore, this is essentially the best possible, that is, there are distributions for which no algorithm can guarantee a probability more than $1/e$ of picking the candidate with the maximum value.

In this paper, we focus on a special class of value distributions F . First, we assume that values are all positive, that is, F has support on $[0, \infty)$ and $F(0) = 0$. This is essentially without loss of generality, as any candidate with negative or zero value can be discarded. More importantly, we assume that the distribution F satisfies the *monotone hazard rate condition*, which means that the hazard function $\frac{f(y)}{1-F(y)}$ is a monotone non-decreasing function of y . This assumption is standard in engineering applications and auction theory [8], and intuitively means that given that the value of a candidate is not less than y , the likelihood that it is equal to y increases as y increases. As an example, Gaussian, uniform, and exponential distributions satisfy this property.

In addition to the above, we make the simplifying assumption that the values are bounded. Therefore, the inverse function $F^{-1}(z) := \inf\{y \in [0, \infty) : F(y) \geq z\}$ is well-defined on its domain $[0, 1]$ and satisfies $F^{-1}(0) = 0$ and $F^{-1}(1) < \infty$.²

3 Main Result

Our main result is to analyze the optimal sample size in an aspiration strategy, assuming that the value distribution satisfies the hazard rate condition. This is stated in the following theorem.

Theorem 1. *Suppose the value distribution F satisfies the above conditions. Then the optimal sample size in an aspiration strategy does not exceed*

$$\frac{n-1}{\sum_{i=1}^{n-1} \frac{1}{i}} - 1.$$

Proof. Consider the optimal aspiration strategy, and let y denote the realized maximum value of the first k observations. The probability density function associated with this event can be written as $kF(y)^{k-1}f(y)$. The strategy then performs a search from candidates $k+1$ to n . Period j ($j > k$) is reached if none of the earlier values are acceptable. This happens with probability $F(y)^{j-k-1}$, and in this case if the value x of the j 'th candidate exceeds y (or if $j = n$), it is accepted. Thus the payoff in an aspiration strategy associated with setting the aspiration level with k observations is

² It is not hard to see that our proof works for a more general class of distributions, satisfying the condition that $\lim_{z \rightarrow 1} (1-z)F^{-1}(z) = 0$. This includes distributions such as the exponential distribution. Details of the proof is left to the final version of the paper.

$$\begin{aligned}
\pi_k &= \int_0^\infty kF(y)^{k-1}f(y) \left(\sum_{j=k+1}^{n-1} F(y)^{j-k-1} \int_y^\infty xf(x)dx + F(y)^{n-k-1} \int_0^\infty xf(x)dx \right) dy \\
&= \int_0^\infty kf(y) \left(\sum_{j=k+1}^{n-1} F(y)^{j-2} \int_y^\infty xf(x)dx + F(y)^{n-2} \int_0^\infty xf(x)dx \right) dy \\
&= k \sum_{j=k+1}^{n-1} \frac{F(y)^{j-1}}{j-1} \int_y^\infty xf(x)dx \Big|_{y=0}^\infty + k \int_0^\infty \left(\sum_{j=k+1}^{n-1} \frac{F(y)^{j-1}}{j-1} yf(y) \right) dy \\
&\quad + \frac{k}{n-1} F(y)^{n-1} \Big|_{y=0}^\infty \int_0^\infty xf(x)dx \\
&= 0 + k \sum_{j=k+1}^{n-1} \int_0^\infty y \frac{F(y)^{j-1}}{j-1} f(y) dy + \frac{k}{n-1} \int_0^\infty xf(x)dx \\
&= k \sum_{j=k}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz + \frac{k}{n-1} \int_0^1 F^{-1}(z) dz,
\end{aligned}$$

where the third equation follows from integration by parts, and the last equation is using the change of variables $z = F(y)$. Thus,

$$\begin{aligned}
\pi_{k+1} - \pi_k &= (k+1) \sum_{j=k+1}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz - k \sum_{j=k}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz + \frac{1}{n-1} \int_0^1 F^{-1}(z) dz \\
&= \sum_{j=k+1}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz - \int_0^1 F^{-1}(z) z^k dz + \frac{1}{n-1} \int_0^1 F^{-1}(z) dz \\
&= \int_0^1 F^{-1}(z) \left(\sum_{j=k+1}^{n-2} \frac{z^j}{j} - z^k + \frac{1}{n-1} \right) dz \\
&= F^{-1}(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1} \right) \Big|_{z=0}^1 \\
&\quad - \int_0^1 F^{-1}'(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1} \right) dz \\
&= - \int_0^1 F^{-1}'(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1} \right) dz, \tag{1}
\end{aligned}$$

where the last equation follows from the fact that $F^{-1}(1) < \infty$ and the identity

$$\frac{1}{n-1} + \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} - \frac{1}{k+1} = 0. \quad (2)$$

Using the same identity, we can write

$$\begin{aligned} \pi_{k+1} - \pi_k &= - \int_0^1 F^{-1'}(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1} - z}{j(j+1)} - \frac{z^{k+1} - z}{k+1} \right) dz \\ &= \int_0^1 F^{-1'}(z)(1-z) \left(\sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} \sum_{i=1}^j z^i - \frac{1}{k+1} \sum_{i=1}^k z^i \right) dz. \end{aligned}$$

This can be written as

$$\pi_{k+1} - \pi_k = \int_0^1 F^{-1'}(z)(1-z)\beta(z)dz, \quad (3)$$

where

$$\beta(z) = \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} \sum_{i=1}^j z^i - \frac{1}{k+1} \sum_{i=1}^k z^i.$$

Since $F^{-1'}(z) = \frac{1}{f(F^{-1}(z))}$, the function $(1-z)F^{-1'}(z) = \frac{1-z}{f(F^{-1}(z))}$ is the composition of the function $\frac{1-F(y)}{f(y)}$, which is non-increasing by the hazard rate condition, and the function F^{-1} , which is non-decreasing. Therefore, $(1-z)F^{-1'}(z)$ is a non-increasing function of z . The function $\beta(z)$ can be simplified as follows:

$$\begin{aligned} \beta(z) &= \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} \sum_{i=k+1}^j z^i + \left(\sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} - \frac{1}{k+1} \right) \sum_{i=1}^k z^i \\ &= \sum_{i=k+1}^{n-2} \sum_{j=i}^{n-2} \frac{1}{j(j+1)} z^i - \frac{1}{n-1} \sum_{i=1}^k z^i \\ &= \sum_{i=k+1}^{n-2} \left(\frac{1}{i} - \frac{1}{n-1} \right) z^i - \frac{1}{n-1} \sum_{i=1}^k z^i \\ &= \sum_{i=k+1}^{n-2} \frac{z^i}{i} - \frac{1}{n-1} \sum_{i=1}^{n-2} z^i, \end{aligned} \quad (4)$$

where the second and the third equations follow from the identity (2). This function has the following properties:

1. $\beta(0) = 0$.
2. For $z > 0$ sufficiently small, $\beta(z) < 0$. This follows from $\beta'(0) = \frac{-1}{n-1} < 0$.

3. There is a $z^* \in (0, 1]$ with $\beta(z) > 0$ if and only if $z > z^*$.

The last property follows from expressing β as

$$\beta(z) = z^k \left(\sum_{i=k+1}^{n-2} z^{i-k} \left(\frac{1}{i} - \frac{1}{n-1} \right) - \frac{1}{n-1} \sum_{i=1}^k z^{i-k} \right)$$

and noting that the term in parentheses is increasing in z , so that once the term is positive it remains positive.

Claim. Assume β is a function satisfying the above properties. If $\int_0^1 \beta(z) dz \leq 0$, then for every non-negative non-increasing function $x(\cdot)$ we have $\int_0^1 x(z)\beta(z) dz \leq 0$.

Proof. Consider the value z^* guaranteed to exist by property (3). We have

$$\begin{aligned} \int_0^1 x(z)\beta(z) dz &= \int_0^{z^*} x(z)\beta(z) dz + \int_{z^*}^1 x(z)\beta(z) dz \\ &\leq \int_0^{z^*} x(z^*)\beta(z) dz + \int_{z^*}^1 x(z^*)\beta(z) dz \\ &= x(z^*) \int_0^1 \beta(z) dz \\ &\leq 0. \end{aligned}$$

The first inequality replaced x in the first integral with the lowest possible value given the other values, increasing the objective function since β was negative, and x in the second integral with the highest possible, again increasing the objective function since β was positive.

Now apply the claim with $x = F^{-1'}(z)(1-z)$. By (4), we have

$$\begin{aligned} \int_0^1 \beta(z) dz &= \sum_{i=k+1}^{n-2} \int_0^1 \frac{z^i}{i} dz - \frac{1}{n-1} \sum_{i=1}^{n-2} \int_0^1 z^i dz \\ &= \sum_{i=k+1}^{n-2} \frac{1}{i(i+1)} - \frac{1}{n-1} \sum_{i=1}^{n-2} \frac{1}{i+1} \\ &= \frac{1}{k+1} - \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{1}{i} \end{aligned}$$

Therefore, when $k \geq \frac{n-1}{\sum_{i=1}^{n-1} \frac{1}{i}} - 1$, by the above equation $\int_0^1 \beta(z) dz \leq 0$. Therefore, since $x = F^{-1'}(z)(1-z)$ is a non-negative and non-increasing function, by Claim 3 and Equation (3) we obtain

$$\pi_{k+1} - \pi_k = \int_0^1 x(z)\beta(z) dz \leq 0.$$

Thus, the sample size k that maximizes π_k cannot be larger than $\frac{n-1}{\sum_{i=1}^{n-1} \frac{1}{i}} - 1$.

Note that $\sum_{i=1}^{n-1} \frac{1}{i} \geq \ln n$. Thus, an upper bound for the optimal k is

$$\frac{n-1}{\ln n} - 1 < \frac{n}{\ln n}.$$

The upper bound is strict in the sense that it is taken on exactly when the distribution is exponential and satisfies the hazard rate constraint with equality.

Our result shows that, for example, when n is a million, it is enough to explore 7.2% of the candidates before setting an aspiration level, rather than 36.8% as in the standard result. Moreover the share of observations devoted to learning falls as the sample size rises, which is a natural conclusion, as the standard mechanism does a lot of search.

How well does the improved calculation do? Table 1 reports estimated values using aspiration strategies with exploration phases of length n/e and $n/\ln n$, for various values of n and an exponential distribution. All values are expressed as a percentage of the expected maximum realized value

$$\int_0^\infty nx F(x)^{n-1} f(x) dx.$$

The column ‘‘Higher’’ gives the proportion of times a strategy beats the other strategy. Thus the n/e strategy wins over $n/\ln n$ frequently, but the $n/\ln n$ produces a higher value.

n	n/e		$n/\ln n$		Optimal Search
	Value	Higher	Value	Higher	
50	64%	22%	69%	12%	88%
100	64%	27%	71%	15%	89%
500	64%	36%	76%	20%	92%
1000	63%	38%	78%	22%	92%
5,000	63%	43%	81%	25%	94%
10,000	63%	44%	82%	26%	94%

Table 1. Comparison between three aspiration strategies with exponential distribution

The last column of the table sets out the average values associated with optimal search, conditional on the exponential distribution. Optimal search involves setting a critical value c_t in each time period t and accepting an offer if it exceeds that critical value. The critical values c_t can be computed using a dynamic programming approach as follows: clearly $c_n = 0$ as any offer better than nothing is acceptable in the last step. This produces a last period expected value v_n of

$$\int_0^\infty x f(x) dx.$$

Let v_t be the value of searching starting from period t . This value is

$$v_t = \int_{c_t}^\infty x f(x) dx + F(c_t) v_{t+1}.$$

Clearly v_t is maximized by $c_t = v_{t+1}$, which gives

$$v_t = \int_{v_{t+1}}^{\infty} xf(x)dx + F(c_t)v_{t+1} = v_{t+1} + \int_{v_{t+1}}^{\infty} (1 - F(x))dx.$$

It is not hard to see that the upper bound on the length of the exploration phase is tight for the exponential distribution, but can be too large for other distributions. For example, for the uniform distribution, by Equation (1) we have

$$\begin{aligned} \pi_{k+1} - \pi_k &= - \int_0^1 \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1} \right) dz \\ &= - \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)(j+2)} + \frac{1}{(k+1)(k+2)} - \frac{1}{2(n-1)} \\ &= \frac{1}{2} \left(\frac{1}{(k+1)(k+2)} - \frac{1}{n} \right). \end{aligned}$$

Thus, the optimal k for the uniform distribution is less than \sqrt{n} . These results give insight into optimal experimentation. The standard approach utilizes 37% of the data for experimental purposes. In contrast, when tails are not too fat, one should use at most a $1/\ln n$ fraction for learning about the distribution. With the uniform distribution, a $1/\sqrt{n}$ fraction is the appropriate length of the experimentation phase.

References

1. M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. Online auctions and generalized secretary problems. to appear in *ACM SIGecom Exchanges*.
2. E.B. Dynkin. The optimal choice of the stopping moment for a markov process. *Dokl. Akad. Nauk. SSSR*, 150:238–240, 1963.
3. T.S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
4. P.R. Freeman. The secretary problem and its extensions: A review. *International Statistical Review*, 51:189–206, 1983.
5. M. Gardner. *Scientific American*, Feb., Mar. 1960.
6. J.F.C. Kingman. Uses of exchangeability. *Annals of Probability*, 6:183–197, 1978.
7. M. Mahdian, R. P. McAfee, and D. Pennock. Secretary problem with durable employment. working paper, 2008.
8. R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738, June 1987.