The secretary problem

Uriel Feige

July 5, 2021

Chapter 11 in the BWCA book considers random order models. From that chapter, we sketch here the well known secretary problem.

Suppose that we need to hire one secretary among n candidates. The candidates arrive online one by one for interviews. Only at the interview itself the quality of the candidate is revealed to us. After every interview we need to decide if to hire the candidate (and cancel all future interviews), or to move on to the next interview (in which case the candidate leaves and accepts a job elsewhere). Our goal is to hire the candidate of highest quality among the n candidates.

In the worst case, no strategy has probability higher than $1/n$ of achieving this goal, against the following adversary. The adversary picks a random index $t \in [n]$, and then arranges the candidates so that quality increases up to candidate t, and then decreases.

We consider a "semi-random" model where qualities of candidates are arbitrary (set by an adversary), but the order of arrival is random.

Consider the following threshold strategy, with parameter t. The first t candidates are ignored, and the first candidate that is better than (for simplicity, assume that there are no ties) all the first t candidates is hired. If there is no such candidate, hire the last candidate.

The strategy succeeds if at a certain step $s > t$, the best candidate arrives at step s, and the second best candidate among those that arrive earlier arrives in a step up to t. Hence the probability of success is:

$$p(t) = \sum_{s=t+1}^{n} \frac{1}{n} \frac{t}{s-1} = \frac{t}{n} \sum_{s=t+1}^{n} \frac{1}{s-1} \simeq \frac{t}{n}(\ln n - \ln t)$$

To find the optimal value of t, compute $np'(t) = \ln n - \ln t - t\frac{1}{t}$. Equating $np'(t) = 0$ we get $\ln \frac{n}{t} = 1$, implying that $t = \frac{n}{e}$. This gives $p(t) \simeq \frac{1}{e}$.

Threshold strategies are optimal among all strategy if interviews only give ordinal information(the relative standing of a candidate compared to previous candidates). This follows from symmetry (if the candidate in time s is better than every candidate in the prefix, then the permutation before s has no effect on the probability that s is the global maximum, and hence either we always pick s or never), and monotonicity (the probability of such an s being the global maximum only increases with s).

If the qualities of candidates are selected by the adversary from an unbounded range, then one can show that for every strategy, even one that sees the actual values of candidates, the expected value of the selected candidate is at most roughly a $\frac{1}{e}$ fraction of the value of the best candidate.