Problem Setup

- n secretaries to interview; our preference is described by a total order, but we can compare only the ones we have interviewed.

Problem Setup

- n secretaries to interview; our preference is described by a total order, but we can compare only the ones we have interviewed.
- After each interview, we have to make an irrevocable decision whether to hire this secretary or not

Problem Setup

- n secretaries to interview; our preference is described by a total order, but we can compare only the ones we have interviewed.
- After each interview, we have to make an irrevocable decision whether to hire this secretary or not
- The secretaries arrive in a uniformly random order.

Problem Setup

- n secretaries to interview; our preference is described by a total order, but we can compare only the ones we have interviewed.
- After each interview, we have to make an irrevocable decision whether to hire this secretary or not
- The secretaries arrive in a uniformly random order.
- What strategy maximizes the probability of hiring the best secretary?

A first attempt

- Interview the first $n/2$ secretaries, and don't hire any of them; then among the remaining, hire the first one who we prefer to the best we saw in the first half.

A first attempt

- Interview the first $n/2$ secretaries, and don't hire any of them; then among the remaining, hire the first one who we prefer to the best we saw in the first half.

Claim

The above strategy picks the best secretary with probability at least $1/4$.

A first attempt

- Interview the first $n/2$ secretaries, and don't hire any of them; then among the remaining, hire the first one who we prefer to the best we saw in the first half.

Claim

The above strategy picks the best secretary with probability at least $1/4$.

Proof.

With probability $1/2$, the second best is in the first half; with probability $1/2$, the best is in the second half.

A first attempt

- Interview the first $n/2$ secretaries, and don't hire any of them; then among the remaining, hire the first one who we prefer to the best we saw in the first half.

Claim

The above strategy picks the best secretary with probability at least $1/4$.

Proof.

With probability $1/2$, the second best is in the first half; with probability $1/2$, the best is in the second half.
The two events are positively correlated, and with probability at least $1/4$ both happen. Whenever this happens, the strategy picks the best secretary. $\square$

Refining the idea

- Instead of picking $\frac{1}{2}$ of the secretaries for "observation", let's reserve the first α fraction for that purpose and optimize α.

Refining the idea

- Instead of picking $\frac{1}{2}$ of the secretaries for "observation", let's reserve the first α fraction for that purpose and optimize α.
- If the second best appears in the observation part, and the best appear afterwards, our strategy picks the best. This happens with probability at least $\alpha(1-\alpha)$.

Refining the idea

- Instead of picking $\frac{1}{2}$ of the secretaries for "observation", let's reserve the first α fraction for that purpose and optimize α.
- If the second best appears in the observation part, and the best appear afterwards, our strategy picks the best. This happens with probability at least $\alpha(1-\alpha)$.
- If the third best appears in the observation part, and the best two appear afterwards, then our strategy picks the best if the best comes before the second best. This happens with probability at least $\frac{1}{2}\alpha(1-\alpha)^2$.

Refining the idea

- Instead of picking $\frac{1}{2}$ of the secretaries for "observation", let's reserve the first α fraction for that purpose and optimize α.
- If the second best appears in the observation part, and the best appear afterwards, our strategy picks the best. This happens with probability at least $\alpha(1-\alpha)$.
- If the third best appears in the observation part, and the best two appear afterwards, then our strategy picks the best if the best comes before the second best. This happens with probability at least $\frac{1}{2}\alpha(1-\alpha)^2$.
- Similarly, if the fourth best appears in the observation part, and the best three appear afterwards, then our strategy picks the best if the best comes before the other two. Altogether this happens with probability at least $\frac{1}{3}\alpha(1-\alpha)^3$.

Getting an optimal α

- Reconizing that these events are disjoint, we see that our strategy succeeds with probability at least

$$\alpha(1-\alpha) + \frac{1}{2}\alpha(1-\alpha)^2 + \frac{1}{3}\alpha(1-\alpha)^3 + \dots$$

Getting an optimal α

- Reconizing that these events are disjoint, we see that our strategy succeeds with probability at least

$$\alpha(1-\alpha) + \frac{1}{2}\alpha(1-\alpha)^2 + \frac{1}{3}\alpha(1-\alpha)^3 + \dots$$

As n goes large, this goes to

$$\begin{aligned}
\alpha \sum_{k=0}^{\infty} \frac{1}{k}(1-\alpha)^k &= \alpha \sum_{k=0}^{\infty} \int_0^{1-\alpha} x^k \, \mathrm{d}x \\
&= \alpha \int_0^{1-\alpha} \sum_{k=0}^{\infty} x^k \, \mathrm{d}x = \alpha \int_0^{1-\alpha} \frac{1}{1-x} \, \mathrm{d}x \\
&= -\alpha \log \alpha.
\end{aligned}$$

Getting an optimal α

- Reconizing that these events are disjoint, we see that our strategy succeeds with probability at least

$$\alpha(1-\alpha) + \frac{1}{2}\alpha(1-\alpha)^2 + \frac{1}{3}\alpha(1-\alpha)^3 + \dots$$

As n goes large, this goes to

$$\begin{aligned}\alpha \sum_{k=0}^{\infty} \frac{1}{k}(1-\alpha)^k &= \alpha \sum_{k=0}^{\infty} \int_0^{1-\alpha} x^k \, \mathrm{d}x \\ &= \alpha \int_0^{1-\alpha} \sum_{k=0}^{\infty} x^k \, \mathrm{d}x = \alpha \int_0^{1-\alpha} \frac{1}{1-x} \, \mathrm{d}x \\ &= -\alpha \log \alpha.\end{aligned}$$

Maximizing this, we get $\alpha = 1/e$, and *alpha* $\log \alpha = 1/e$ as well.

Getting an optimal α

- Reconizing that these events are disjoint, we see that our strategy succeeds with probability at least

$$\alpha(1-\alpha) + \frac{1}{2}\alpha(1-\alpha)^2 + \frac{1}{3}\alpha(1-\alpha)^3 + \dots$$

As n goes large, this goes to

$$\begin{aligned}\alpha \sum_{k=0}^{\infty} \frac{1}{k}(1-\alpha)^k &= \alpha \sum_{k=0}^{\infty} \int_0^{1-\alpha} x^k \,\mathrm{d}x \\ &= \alpha \int_0^{1-\alpha} \sum_{k=0}^{\infty} x^k \,\mathrm{d}x = \alpha \int_0^{1-\alpha} \frac{1}{1-x} \,\mathrm{d}x \\ &= -\alpha \log \alpha.\end{aligned}$$

Maximizing this, we get $\alpha = 1/e$, and *alpha* $\log \alpha = 1/e$ as well.

Optimal Algorithm

Theorem

The two-stage strategy of first interviewing n/e secretaries, then hiring the first better than all in the first stage, hires the best secretary with probability at least $1/e$. This is the best guarantee by any algorithm.

Optimal Algorithm

Theorem

The two-stage strategy of first interviewing n/e secretaries, then hiring the first better than all in the first stage, hires the best secretary with probability at least $1/e$. This is the best guarantee by any algorithm.

Proof.

We have shown the first part. The optimality of the algorithm we argue by a linear program.

Optimal Algorithm

Theorem

The two-stage strategy of first interviewing n/e secretaries, then hiring the first better than all in the first stage, hires the best secretary with probability at least $1/e$. This is the best guarantee by any algorithm.

Proof.

We have shown the first part. The optimality of the algorithm we argue by a linear program. For $i = 1, 2, \ldots, n$, let x_i be the probability of hiring the i-th secretary.

Optimal Algorithm

Theorem

The two-stage strategy of first interviewing n/e secretaries, then hiring the first better than all in the first stage, hires the best secretary with probability at least $1/e$. This is the best guarantee by any algorithm.

Proof.

We have shown the first part. The optimality of the algorithm we argue by a linear program. For $i = 1, 2, \ldots, n$, let x_i be the probability of hiring the i-th secretary.
The success probability is $\sum_i \frac{i}{n} x_i$.

Optimal Algorithm

Theorem

The two-stage strategy of first interviewing n/e secretaries, then hiring the first better than all in the first stage, hires the best secretary with probability at least $1/e$. This is the best guarantee by any algorithm.

Proof.

We have shown the first part. The optimality of the algorithm we argue by a linear program. For $i = 1, 2, \ldots, n$, let x_i be the probability of hiring the i-th secretary.

The success probability is $\sum_i \frac{i}{n} x_i$. The algorithm interviews the i-th secretary with probability no more than $1 - \sum_{j<i} x_j$ and no less than $i x_i$.

Optimal Algorithm

Theorem

The two-stage strategy of first interviewing n/e secretaries, then hiring the first better than all in the first stage, hires the best secretary with probability at least $1/e$. This is the best guarantee by any algorithm.

Proof.

We have shown the first part. The optimality of the algorithm we argue by a linear program. For $i = 1, 2, \ldots, n$, let x_i be the probability of hiring the i-th secretary.
The success probability is $\sum_i \frac{i}{n} x_i$. The algorithm interviews the i-th secretary with probability no more than $1 - \sum_{j<i} x_j$ and no less than ix_i. (As the i-th secretary is interviewed, with probability $1/i$ will this he the best so far, and only when this happens can the algorithm hire him.) □

Proof continued

Proof.

Therefore the performance of any algorithm is upper bounded by the following linear program:

$$
\begin{aligned}
\max_{\mathbf{x}} \sum_i \frac{i}{n} x_i \\
s.t. \quad i x_i \le 1 - \sum_{j<i} x_j, \quad i = 1, 2, \ldots, n; \\
\sum_i x_i \le 1; \\
x_i \ge 0, \quad i = 1, 2, \ldots, n.
\end{aligned}
$$

Proof continued

Proof.

Therefore the performance of any algorithm is upper bounded by the following linear program:

$$
\begin{aligned}
&\max_{\mathbf{x}} \sum_i \frac{i}{n} x_i \\
s.t. \quad & i x_i \leq 1 - \sum_{j<i} x_j, \quad i = 1, 2, \ldots, n; \\
& \sum_i x_i \leq 1; \\
& x_i \geq 0, \quad i = 1, 2, \ldots, n.
\end{aligned}
$$

Looking at the dual, one sees an upper bound of $1/e$ on the value of the LP. □