When should you stop waiting?
A mathematical exploration of optimal decision making

How many times have you wondered whether to stop waiting and just make a decision, whether it was looking for a new apartment, hunting for the best parking spot, or even choosing an ice cream flavour? We have all struggled with the paralysis of choice. Now, would it be crazy if I told you that maths has found an optimal strategy for all situations just like this? Sounds impossible, right? This is the secretary problem, popularised by Martin Gardner in 1960, and the answer is as surprising as it is elegant, yet simultaneously raises a multitude of deeper questions.

The problem

The problem is set out like this. Imagine you were the CEO of Apple looking for a new head of marketing. You must interview n candidates in sequence; after each interview, you either offer them the job immediately, ending the search, or reject them and move on. You cannot go back, after rejecting a candidate. The only information available at each stage is how the current candidate compares to those already seen. So what do we do? Do we simply keep interviewing until someone feels right? Or do we follow a structured strategy?

Strategy and proof

A natural strategy is to reject the first k candidates outright, using them purely as a baseline, then offer the job to the first candidate who outperforms everyone in that rejection window. We can define the event Bj as:

B_j = "j'th person is the best candidate"

Now we can write the probability of picking the best candidate using the law of total probability:

best candidate

$$p(\text{best candidate}) = \sum_{i=1}^{n} \underbrace{P(BC \mid B_j)}_{2\ \text{cases}} \cdot \underbrace{P(B_j)}_{\frac{1}{n}}$$

Since each candidate is equally likely to be the best, P(Bj) = 1/n for all j. The conditional probability P(BC|Bj) is split into two cases:

<u>case 1:</u>

if $j \leq k$, $p(BC \mid B_j) = 0$

<u>case 2:</u>

$j > k$ assume B_j occurs $p(BC / B_j) = \frac{k}{i-1}$

Case 1 states that if jth person is between the interval 1 and k, then there is 0 chance of picking the best candidate as j is in our rejection interval, therefore no matter what position j is between 1 and k, the sum will always be 0 (0 x 1/n + 0 x 1/n +... all the way to k sums to 0).

Case 2 states that if jth person is greater than k, we now have a chance of picking the best candidate only if we don't stop early at anything before. You stop early if anything between k plus 1 and j-1 was the best you've seen so far. So the only time when we don't stop early before j, is if the best candidate so far was seen in our “must reject” window. Therefore the probability that the best of those j-1 items happens to be in your rejection window is simply, k(favourable slots) over j-1 (total slots).

Now the probability of selecting the best candidate becomes:

$$P(Bc) = \sum_{\boxed{i=k+1}}^{n} \frac{k}{j-1} \times \frac{1}{n}$$

our beginning of our sum has changed due to the probability sum from 1 to k being 0, ∴ we can start from k+1.

As n becomes large, the sum

k/n × Σ 1/(j-1) for i = k+1 to n

can be approximated by an integral. This is because for large n the terms 1/(j-1) become increasingly fine-grained, resembling a Riemann sum, as shown below.

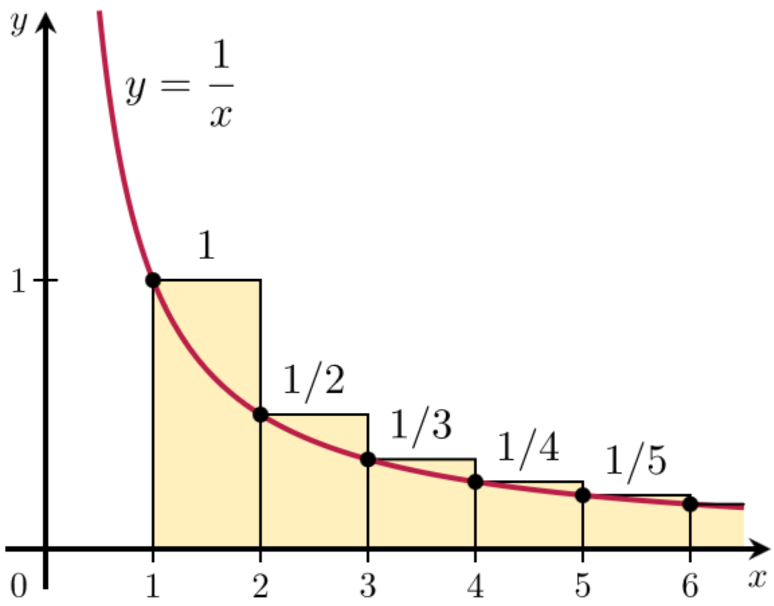

Therefore we can do the approximation:

Σ 1/(j-1) ≈ ∫ from k to n of 1/x dx

Integration working out:

$$P(Be) = \sum_{\boxed{i=k+1}}^{n} \frac{k}{j-1} \times \frac{1}{n}$$

our beginning of our sum has changed due to the probability sum from 1 to k being 0, ∴ we can start from k+1.

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

riemann integral

$$\approx \frac{k}{n} \int_{k}^{n} \frac{1}{x}\, dx$$

$$= \frac{k}{n} \left[\ln x\right]_{k}^{n}$$

$$= \frac{k}{n} \left[\ln(n) - \ln(k)\right]$$

$$= \frac{k}{n} \ln\left(\frac{n}{k}\right)$$

Now we have the function p(r) is equal to rln(1/r), which can be sketched, as shown in the below figure.

$$r = \frac{k}{n}$$

$$p(r) = r\ln\left(\frac{1}{r}\right)$$

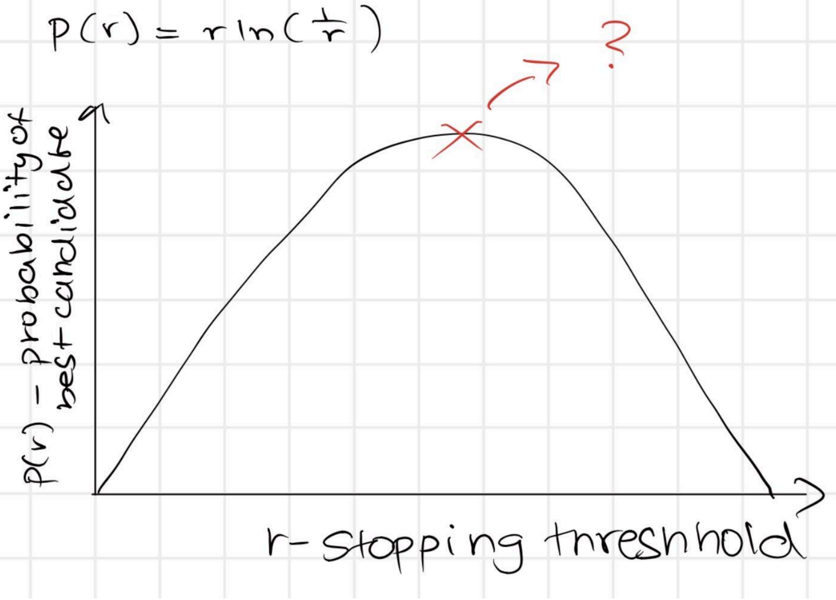

This sketch shows p(r) as a function of the stopping threshold r. The curve rises as r increases from 0, peaks at a single maximum then falls which suggests there exists an optimal stopping point, which can be worked out using a method we all know. Differentiation.

$$p(r) = r\ln\left(\frac{1}{r}\right)$$

$$= r\ln(r^{-1})$$

$$= -r\ln(r)$$

$$p'(r) = \ln r + 1$$

$$\ln r + 1 = 0$$

Now this proof has shown us that our solution to the secretary problem is to reject the first n/e candidates (i.e. approximately the first 36.8 percent of all candidates to be interviewed), then accept the first candidate better than those in the n/e interval. This yields the highest possible probability of selecting the best candidate, a probability of 1/e, approximately 36.8%.

But what is the point of this proof if we don't test it?

Right now you can do an experiment to test this "optimal solution", using a deck of uno cards, (or even a normal deck of cards). To verify this result experimentally, you can replicate the secretary problem using a deck of Uno cards. First you must make sure to remove all special cards such as plus 2s, plus 4s and reverses etc. and keep only the numbered cards from 1 to 9, removing the 0s too, leaving 72 cards in total. The rules are exactly the same as the secretary problem. Flip cards one at a time, note the number, and do not stop until the cutoff point, after which you stop at the first card that is better than everything seen in the rejection window. A 9 counts as a win; anything else is a loss. Running 30 rounds each at a 37% cutoff (k = 27), a 25% cutoff (k = 18), and a 50% cutoff (k = 36) allows direct comparison between thresholds - if it gets a bit tedious, just do at least 10 rounds - you can do even more different cutoffs if your feeling more adventurous!

These were my results after going at it for an arduous 45 mins.

37% - 27 card cutoff

round 1: L	round 10: L	round 19: L	round 28: W
round 2: W	round 11: W	round 20: W	round 29: L
round 3: L	round 12: L	round 21: W	round 30: L
round 4: W	round 13: L	round 22: L	
round 5: L	round 14: W	round 23: L	
round 6: W	round 15: W	round 24: L	
round 7: L	round 16: L	round 25: W	
round 8: L	round 17: L	round 26: L	
round 9: L	round 18: L	round 27: L	

10 wins $\frac{10}{30} \times 100 = 33.\dot{3}\%$ win rate

50% - 36 cards cutoff

round 1: W	round 10: L	round 19: L	round 28: L
round 2: L	round 11: L	round 20: L	round 29: L
round 3: L	round 12: L	round 21: L	round 30: L
round 4: L	round 13: L	round 22: L	
round 5: L	round 14: W	round 23: L	
round 6: L	round 15: W	round 24: L	
round 7: L	round 16: L	round 25: L	
round 8: L	round 17: L	round 26: W	
round 9: L	round 18: L	round 27: W	

$\frac{5}{30}$ wins = 16.7% win rate.

25% - 18cards cutoff

round 1: L	round 10: L	round 19: W	round 28: W
round 2: L	round 11: L	round 20: L	round 29: L
round 3: L	round 12: L	round 21: L	round 30: L
round 4: L	round 13: W	round 22: L	
round 5: W	round 14: L	round 23: L	
round 6: L	round 15: L	round 24: L	
round 7: L	round 16: L	round 25: W	
round 8: W	round 17: L	round 26: L	
round 9: L	round 18: W	round 27: L	

$\frac{7}{30}$ wins = 23.3% win rate

Across 30 trials, the 37% threshold achieved a success rate of 33.3%, closely matching the theoretical prediction of 36.8%. By contrast, a 25% threshold yielded only 23.3% and a 50% threshold just 16.7%, confirming that the optimal stopping point is neither too early nor too late.

Limitations and extensions

Despite going through the rigorous proof and a test of this 37 percent rule's solution, there will most likely be cases where this strategy may seem to fall apart. The strategy requires us to know n in advance so that we can calculate our k cutoff point. However in real life there will be a plethora of cases where we don't know the number of total options we will encounter. How many apartment viewings will you go to? How many people will apply to the secretary role? How many people do you plan to date in your lifetime? Without knowing our n, the strategy seems to break down.

This limitation was actually resolved by T.J. Stewart (1981) where he used a time threshold approach. So candidates arrive randomly over a time interval [0,T], which can be normalised to [0,1], with t being current time. We want to find the optimal time threshold t* after which we can now start choosing a candidate. At time t, if the current candidate is the best choice so far, you face a choice. The probability of successfully picking the overall best if you stop now is approximately:

$$p(\text{success} \mid \text{stop at time } t) = t$$

This is because for the current candidate to be the best overall candidate, all candidates from the interval [t,1] must be worse than the current one. Stewart (1981) showed that under a continuous arrival model, the optimal success probability V(t) satisfies a differential equation:

$$V(t) = -t\ln(t)$$

This is exactly identical to our original P(r) = r ln(1/r). Setting t ≥ V(t) and solving recovers the threshold t = 1/e, meaning you should observe until 1/e of your time has elapsed before selecting. Stewart had preferred to treat t as remaining time, as the continuous arrival model naturally tracks what opportunity remains rather than what has passed, which gives 1 - 1/e ≈ 0.632. Remarkably, the same constant that governed the classical problem governs this one too, even without any knowledge of n. What I find particularly amazing about this is that two completely different methods, one through combinatorics and one through differential equations, lead to the exact same mathematical solution. It is the kind of result that makes you feel mathematics is uncovering something genuinely deep about the nature of optimal decision making, rather than just solving a made-up puzzle.

Another common question is what if you can go back? In the original problem, the rules are set out that you cannot go back and change your decision of rejecting a candidate. But this doesn't seem to hold in most cases. The apartment that you viewed last week is most likely still up for sale, an employer will most likely be able to reconsider a rejected applicant. This is known as the secretary problem with recall. Yang M.C.K (1974) in his paper "Recognising the maximum of a sequence with backward solicitation", showed that when full recall is permitted, the optimal strategy changes. Early candidates now require a higher bar to be selected, but as you approach the end of your search you become progressively less selective. Now that the cost of rejection is no longer permanent, you can afford to have higher standards and be more strict earlier on. The success probability is now much higher than 1/e, as recall reduces the risk of missing the best candidate entirely in the rejection interval, like the [1,k] interval which was our must reject window, in the original problem. What I find quite remarkable is how Yang was able to produce an entirely different mathematical structure after a single change to one assumption.

These solutions further illustrate that in mathematics limitations aren't dead ends, they are rather doors to richer mathematics.

Furthermore, the 1/e barrier is not an impenetrable frontier. As pointed out by Braun and Sarkar (2024), the inclusion of only one extra bit of information about the hidden process, specifically, the additive gap between the optimal solution and the kth best alternative, is sufficient to break the traditional barrier. This new approach has a competitive ratio of 0.4 irrespective of the size of the gap. It is astonishing how much can be accomplished with such minimal information that is required to breach the barrier, which stood invincible for over six decades. It is unnecessary to have any specific knowledge of the optimal value or even its index number. It is enough to know that there exists a gap of size c.

<u>**Conclusion**</u>

To conclude, the secretary problem transforms an everyday dilemma into a striking mathematical insight; sometimes the best strategy is to be patient rather than deliberately wait before acting. Through proof, experimentation and real-life application, this essay has shown that this problem is more than a theoretical curiosity, but rather a mathematical framework for understanding decision-making. The hidden principle remains powerful; success is not about choosing well but rather when to choose.

References:

The Secretary Problem with an Unknown Number of Options by T.J. Stewart

https://www.kurims.kyoto-u.ac.jp/~kyodo/kokyuroku/contents/pdf/1559-11.pdf

https://arxiv.org/pdf/2306.08340

https://arxiv.org/pdf/2409.20460

https://youtu.be/ZWib5olGbQ0?si=ChFmIZntqfp1ajMY

https://youtu.be/EahsbSMvXrM?si=AgosGV7wdhh03q6e

Recognising the maximum of a sequence with backward solicitation by Yang M.C.K