Warps and Squares

Hari Desikan

10/20/2025

1 The problem

Now that GPT can solve this problem, there is no use hiding it from the internet.

Alice and Bob play a game on a board of length $n+2$, labelled $0, \ldots, n+1$. Alice starts at square 0. At any time, if Alice is at square x, Alice must jump to square $y > x$. Alice's goal is to jump to as many squares as possible. Once Alice jumps to square $n+1$, Alice's score is tallied. Alice's score is the number of squares they visited besides the first and last squares.

However, Alice must send Bob their (mixed) strategy to play this game. Then, Bob (who hates Alice) may choose to place an obstruction on any square $z \in [1, n]$. If Alice ever lands on this obstruction, Alice's game immediately ends and Alice does not get points for square z, but gets points for the other squares she visited.

Find (with proof) the expected percentage of the squares Alice visits, if Alice chooses a strategy maximizing this expectation. You may assume n is sufficiently large

1.1 Presentation

This pdf is presented to make sure any reader can learn something. As such, I've purposefully included extraneous information so that many different ideas are discussed and so readers from different backgrounds can glean something before they are lost. For example,

- An "elementary" problem that is closely tied to this problem
- Incorrect approaches that were presented to me when I tried this problem, as well as
- Elementary game theoretical foundations that are crucial to this problem,
- All observations present in a correct proof

In particular, I'm not just solving the problem. If for every problem I spent much time solving I simply wrote up the solution, I would need to comparatively solve many more problems to include a similar volume of content on my webpage. I have a few other hobbies.

Also, sometimes I write

$$\sum_{i=m}^{1} f(i).$$

This is identical to

$$\sum_{i=1}^{m} f(i).$$

m is not negative. This is a stylistic choice in some sections to make the summation fit neater into your head, say, if the $i = m$ case intuitively occurs "before" the $i = 1$ case.

2 An introductory problem

It's clear Alice's strategy needs to be mixed (ie include randomness.) Otherwise, Bob will just place the obstruction where Alice's deterministic first move plays to. Let's investigate this randomness.

Given the same setup, suppose at any instance Alice randomly jumps to one of the valid available squares, and Bob does not place an obstruction. We'll also impose that Alice is not entirely stupid, as is stopped from jumping to $n+1$ unless Alice is at n as any other jump is better.

- In this setup, Alice is not smart and Bob does not exist. Do we expect Alice to hit fewer or more squares with these competing effects?
- Find $\lim_{n\to\infty} S(A)/n$, the expected proportion of the squares Alice hits.
- Find the expected number of squares Alice lands on for arbitrary n.

The first two problems are meant to be intuition checks. The third problem is meant to be computed.

3 An incorrect solution

We've already seen that uniformly picking a point to jump to gives up too much EV (it gets a value asymptoting to 0.) Consider now a strategy that can be considered the next-simplest form of randomization. At any point, Alice randomly jumps one or two steps. Clearly, the best counterstrategy is to place the block on one of the first two squares, because placing the block further out just gives Alice free sqaures for no gain, so placing the block further out is strictly dominated. Alice then gets half the squares (jump one or two) half the time (Alice doesn't jump on the block) for a guarantee of

$$\frac{1}{4}$$

Clearly this strategy is completely random and unlikely to be correct. Try to find all the mistakes/issues before reading on!

1. Alice does not get half the squares. Since Alice moves either 1 or 2 squares, Alice on average reaches $\frac{2}{3}$ of the squares, so Alice actually gets $\frac{1}{3}$ with this strategy.

2. The optimal counterstrategy here is quite simple. A good Alice strategy should ideally force our opponent to not have a trivial counterstrategy. Here, we introduce a quick aside:

3.1 Aside

In perfect information games between actors Alice, Bob, the following algorithm converges to a nash equilibria:

1. Come up with a strategy for Alice (any strategy) S_1

2. Find a maximally exploitative counterstrategy for S_1 "S_2"

3. Find a maximally exploitative counterstrategy for S_2 "S_3.

4. Recurse to infinity

5. The limit of Alice's strategies is her optimal strategy.

Unfortunately, this is not a perfect information game. And guarantees do not hold for imperfect information games. Consider Rock-paper-scissors with $S_1 = R, S_2 = P, S_3 = S, \ldots$ clearly, this doesn't converge.

Also, before you even mention it, even taking the average of the odd S_is to construct a NE doesn't work in general.

However, in trying to find optimal strategies, empirically true that something like this converges to a correct strategy:

1. Pick S_1 intelligently

2. Pick S_2 maximally exploiting S_1

3. At any point $t = i$, while recursing as above, suppose we prove that the optimal strategy belongs to some set of strategies Q. Henceforth enforce that $S_{i+j} \forall j \geq 1$ maximally exploits S_{i+j-1} of the strategies in Q.

In other words, humans can proceed as if the rock-paper-scissors example above isn't an issue.

3.2 Not an aside

Returning to the strategy above, note that Bob has a trivial strategy and limits us to just $\frac{1}{3}$ of the points. I will now point out the following slightly unintuitive result: the optimal counterstrategy for Bob is not to play at the beginning, as the first square has 0.5 chance of being chosen. But a square just, say, 20 squares in has $\frac{2}{3} - \epsilon$ chance of being chosen. For n large, giving Alice 20 squares for free to have a higher chance of Nuking Alice is quite the deal. Then, asymptotically we are only guaranteed (C is 20 from before,)

$$\approx (1 - (\frac{2}{3} - \epsilon)\frac{2}{3} + C \approx \frac{2}{9}$$

of the squares, with equality in the limit.

For the rest of this section, we pretend the squares are on a number line. So Bob just needs to play arbitrarily close to 1 to obstruct us $\frac{2}{3}$ of the time. As its clear in the limit that this value is irrelevant, we just say Bob plays at 0. We can however make Bob's life a little bit more difficult by using the intuition in the aside. Instead of always jumping one square or two squares, we instead jump one or two squares up to a point p, after which we jump 1 square always. So if $p = .5$, the first half of the way we always randomize, but if we get to the second half we play one square at a time.

The idea behind this is so if Bob plays at 0, then if we pass the obstruction we realize all of the squares past p. We can make p smaller "for free" as long as Bob is still incentivized to play at 0, and we can keep doing this until Bob is indifferent. Our EV if bob plays at the beginning is

$$\frac{1}{3}(\frac{2}{3}p + 1 - p),$$

and our EV if bob plays at p is

$$\frac{2}{3}p.$$

So if

$$p = \frac{3}{7},$$

Bob is indifferent and we always gain

$$\frac{2}{7} > \frac{2}{9}.$$

3.3 Jump m?

We can extend the idea further, by saying that we should partition the segment $[0, 1]$ into m segments, and jump

$$U(1, m - i + 1)$$

squares in the ith segment, where U is the discrete, integer-valued uniform distribution. We'd want to make Bob indifferent between playing at the beginning of any of these segments.

How much does this get?

For $m = 3$, we get that the segment lengths S_m (where S_m is the segment where Alice jumps $U(1, m)$ are

$$S_3 = \frac{4}{11}, S_2 = \frac{3}{11}, S_1 = \frac{4}{11},$$

and Alice gets

$$\frac{4}{11} > \frac{2}{9},$$

which is a big improvement!

For arbitrary m, we get the following equations:

$$\sum_{i=1}^{n} S_i = 1,$$

$$\frac{m-1}{m+1}\sum_{i=m}^{1} S_i\frac{2}{i+1} = S_m\frac{2}{m+1}+\frac{m-2}{m}\sum_{i=m-1}^{1} S_i\frac{2}{i+1} = S_m\frac{2}{m+1}+S_{m-1}\frac{2}{m}+\frac{m-3}{m-1}\sum_{i=m-2}^{1} s_i\frac{2}{i+1} = \dots.$$

We really just want to know the value of any one of these things that are equal. Note the last term is

$$S_m\frac{2}{m+1} + S_{m-1}\frac{2}{m} + \cdots + S_2\frac{2}{3} + \frac{0}{2}\sum_{i=1}^{1} S_1\frac{2}{2} = \sum_{i=m}^{2} S_i\frac{2}{i+1}.$$

Replacing $A_i = S_i\frac{2}{i+1}$ as is very natural yields

$$\frac{m-1}{m+1}\sum_{i=m}^{1} A_i = A_m + \frac{m-2}{m}\sum_{i=m-1}^{1} A_i = \cdots = (\sum_{i=m}^{m-j} A_i) + \frac{m-j-2}{m-j}\sum_{i=m-j-1}^{1} A_i = \cdots = \sum_{i=m}^{2} A_i.$$

We can notice with our great powers of observation (eg. noticing that the A_is are all forced to be $\frac{\Sigma}{m+1}$ successively due to symmetry, and then A_1 must be double this) that

$$A_m = \cdots = A_2 = \frac{A_1}{2},$$

and so

$$S_m\frac{2}{m+1} = S_{m-1}\frac{2}{m} = \cdots = S_j\frac{2}{j+1} = \cdots = S_2\frac{2}{3}, = S_1\frac{2}{2}/2 \implies \frac{S_m}{m+1} = \frac{S_{m-1}}{m} = \cdots = \frac{S_2}{3}, = \frac{S_1}{4}$$

Let's not forget; we want the EV of this. The EV of this is the EV if our opponent snipes plays at the beginning of the last segment, which is

$$\sum_{i=m}^{2}\frac{2S_i}{i+1}$$

$$= \frac{S_1}{2}(m-1)$$

and I'm wondering where I went so wrong so as to not instantly see this. Finally, note that the S_i are in the ratio $m+1 : m : m-1 : \cdots : 3, : 4$ in order from $S_m : S_1$, so

$$S_1 = \frac{4}{(m+1)(m+2)/2 = 1} \implies \frac{S_1}{2}(m-1) = \frac{2(m-1)}{(m+1)(m+2)/2+1}.$$

Wow! For $m = 5$, we get $\frac{8}{22} = \frac{4}{11}$, the same value as $m = 4$. And going further is actually even worse!

That's not very good.

4 No more guessing

So far we've just been guessing things. And none of these guesses can really constitute a proof. But we have established that Alice can get at least $\frac{4}{11}$ of value, which is something.

You might be wondering why we'd even try guessing strategies and go through all that algebra above. To be honest, when at this point I'd already taken into account the content further on in this section. I'd also made some computation errors, and got $m = 2 \to 0.4, m = 3 \to 0.\overline{4}$ and thought this random strategy, as $m \to \infty$, would achieve a value of 0.5. As you will see in this section, if this were actually true then we would be done!

4.1 Memoryless-ish?

You might recognize that all the strategies we've discussed so far only take into account Alice's current position in order to determine how to randomize Alice's next move. It stands to reason that Alice wants to take into account the *history* of moves she's made so far when making her next move. That is, Alice's next-move generator M should be a function $M(H, x)$ rather than just a function of x.

You might think, from the seeming simplicity of this problem, that there's some trivial way to rule it out. I couldn't come up with one, and I have the following reason why you shouldn't either:

Claim 1. *The strategies discussed in the previous section, for $m = 2, m = 3$, can both be improved by taking history into account.*

Proof. For $m = 2$, precisely when we jump 2 for every single roll in the segment S_2 and we are 2 squares away from $\frac{3}{7}n$, start moving 1 at a time pre-emptively. □

Is this really a proof? No, because we haven't shown that this is an improvement. But clearly some idea like this should be an improvement; if we get to $\frac{3}{7}n - 2$ by jumping 2 every time, vs by jumping 1 every time, we are much more likely to have the block behind us. So since we are more likely to have already jumped over the block, we should be more aggressive in jumping 1-at-a-time.

But if we have no way to rule out histories, guessing a strategy $M(x, H)$ is entirely hopeless. Instead, we turn to one of the most powerful ideas in game theory. While I often thought this theorem to be trivial, applying it to this problem made me realize the theorem's power.

4.2 Nash Equilibria

Theorem 1. *Every two player constant-sum imperfect information game has a nash-equilibria*

Is this really helpful? Well, yes. We've already thrown around NE expecting the reader to understand what they are. But the crucial realization is that NE exist for all games. Including the following game:

4.2.1 Cool game

In this new game, Alice does not give Bob her strategy. But Bob still places a block. Alice supplies a possibly- mixed strategy, hidden to bob, on how to traverse the [n+2] board. Now, Bob *also* supplies a strategy, hidden to Alice; a possibly-mixed strategy on where to place his block.

Consider the nash-equilibria of this game. By the mini-max and maxi-min equivalence, the fair value of this game is a value that Alice can achieve in expectation regardless of what Bob does, and that Bob can hold Alice to no matter what Alice does. That is; there's an Alice-strategy that

no matter what Bob does, gets at least some value v. And there's a Bob strategy that no matter what Alice does, makes sure Alice gets at most some value v.

What is v? Well, since there's an Alice-strategy that gets v no matter what Bob does, *even if Bob knows Alice's strategy precisely,* Alice still gets v. That sounds like the answer to our problem!

But v is now guaranteed to also be something else: the minimum amount Bob can hold Alice to if *Bob has to supply Alice his strategy on where to place his block, and Alice can respond with any traversal strategy to maximally exploit Bob.*

4.3 Concrete details

It is guaranteed that whenever we discuss maximally exploitative strategies, it suffices to discuss deterministic maximally exploitative strategies.

If we are playing rock paper scissors, and you throw scissors 35% of the time and rock and paper 32% of the time, my maximally exploitative counterstrategy is trivially to always throw rock. There is no need for me to randomize. This generalizes to the Nash Equilibria guarantees. So in our original problems, if we supply Bob our strategy, it suffices to consider the n Bob counterstrategies

$$z = 1, \ldots, z = n.$$

Even still, we had a staggering infinitude of Alice strategies that were hard to work with.

When we now flip the problem, consider Alice counterstrategies. When Alice sees a Bob mixed strategy, Alice has a maximally exploitative first move. Given that her first move did not end in ruin, she has a maximally exploitative second move. In particular, any deterministic strategy by Alice can be modeled as a set $S \in 2^{[n]}$, which is the subset of the squares Alice intends to travel to.

Bob's strategy, on the other hand, just amounts to a probability distribution $p_1, \ldots, p_n$ on where to drop the block. Of course, this is still infinite strategies, but intuitively this problem should feel much more tractable than the previous problem. We have no issues with histories (that admittedly are also subsets of $2^{[n]}$, but have weird probabilites of being reached,) for example.

5 Observations on the dual

Let Bob propose $p_1, \ldots, p_n$.

Then suppose Alice proposes $A \in 2^{[n]}$. Let's call the indices of A

$$A_1, \ldots, A_m,$$

and let's quickly note for sanity that $A_m = n$ (if you haven't already chosen the last square, you might as well add it to the subset.)

Now consider Alice's expected number of traversed squares. To do this, we create indicators $I_1, \ldots, I_m$ such that $I_i = 1$ if Alice traverses square A_i and 0 otherwise.

The probability Alice traverses square A_i is the probability squares $A_1, \ldots, A_i$ don't have the block, which occurs with probability

$$E(I_i) = 1 - \sum_{j=1}^{i} p_{A_j},$$

and so Alice's expected score is

$$\sum_{i=1}^{n} E(I_i) = \sum_{i=1}^{m}(1 - \sum_{j=1}^{i} p_{A_j}) = m - \sum_{i=1}^{m}\sum_{j=1}^{i} p_{A_j},$$

$$= m - \sum_{i=1}^{m}(m - i + 1)p_{A_i}.$$

This is progress! Given a probability distribution $p = p_1, \ldots, p_n$ and an Alice strategy $A = A_1, \ldots, A_m =$, we have computed Alice's expected value.

Theorem 2. *There exists a NE where Bob's supplied strategy is nonincreasing.*

Proof. Why on earth would it increase somewhere? This should be intuitive to the reader at this point.

Suppose Bob's strategy increases somewhere, so $p_i < p_{i+1}$ somewhere. There are three cases: For all the cases, let p' be the probability distribution $p_1, \ldots, p_{i-1}, p_{i+1}, p_{i+2}, \ldots, p_n$ (e.g. swap p_i, p_{i+1}). We construct some strategy (not necessarily the optimal one) α on p' that obtains at least as much on p' as A does on p.

1. Neither $i, i+1$ appears in the optimal counterstrategy A. In this case, taking $\alpha = A$ suffices.

2. Exactly i or $i+1$ appears in A. Then taking

   $$\alpha = (A \setminus \{i, i+1\}) \cup (\{i, i+1\} \setminus (\{i, i+1\} \cap A))$$

   Achieves the same value on p' as A on p. (This means: just take the other one.)

3. Both $i, i+1$ appear in A. Then taking $\alpha = A$ yields a strictly higher value on p' than A on p. The difference in the values we obtain is

   $$(-2p_{i+1} - p_i) - (-2p_i - p_{i+1}) = p_i - p_{i+1} > 0,$$

   which exceeds 0 by assumption.

So we can construct a strategy that does at least as well on p' as our optimal counterstrategy on p, so it suffices to look at decreasing distributions p. □

Theorem 3. *Any optimal counterstrategy $A = A_1, \ldots, A_m$ satisfies the requirement that $A_m = n$, and that between any two consecutive values A_i, A_{i+1} such that $A_i + 1 \neq A_{i+1}$,*

$$\min_{j \in [A_i+1, A_{i+1}-1]} p_j \geq \max(p_{A_i}, p_{A_{i+1}})$$

Proof. This theorem basically says that if Alice picks two points A_i, A_{i+1} and doesn't pick any points in between, then none of the interior probabilities are smaller than p_{A_i} or $p_{A_{i+1}}$. This is true as if one of the interior points j had a smaller probability than (WLOG) A_i, we could just pick

$$\alpha : \alpha_i = j, \alpha_i = A_i, \alpha_{j \neq i} = A_j$$

as a better strategy on p, as the squares in α just shares all the probabilities of being blocked with A except one of the squares has strictly lower probability of being blocked than before. To be clear, the difference in the obtained values between this constructed α and A is, with m defined as the length of A,

$$(m - i + 1)(p_{A_i} - p_j) > 0.$$

Also, it's trivial that $A_m = n$. □

We now combine our previous two theorems to obtain the next result:

Theorem 4. *If p is nonincreasing and A satisfies the requirement from the above theorem, then there exists a maximally exploitative counterstrategy A to p that is a tail of $[n]$. In other words,*

$$A = \{n - (m-1), n - (m-2), \ldots, n\}.$$

Proof. We have that A includes n. Now, suppose $|A| = m$, and A is not a tail. Let $n - j : j \in [m-1]$ be the largest index that isn't included in A. Let u be the largest index in A less than j. Then A is of the form

$$\ldots, u, n - (j-1), n - (j-2), \ldots, n.$$

Note that removing u from A and adding $n - j$ to A increases Alice's score by

$$j(p_u - p_{n-j}) > 0,$$

contradicting optimality of A.

So there is no largest index $n - j : j \in [m-1]$ that isn't included in A, so in particular the elements $n - (m-1), \ldots, n$ are all in A. Since by assumption A has m elements, these are all the elements of A. □

We have a lot of progress in the dual. Alice's optimal counterstrategy is a tail, Bob's distribution is nonincreasing. So if Bob supplies Alice a nonincreasing distribution, all Alice has to do if find the best tail to use. So fixing p, the set of possible counterstrategies is

$$\{\{n\}, \{n, n-1\}, \ldots, \{n, n-1, \ldots, 1\}\}.$$

Let's call these strategies $A_1, A_2, \ldots, A_n$ by the sizes of the sets they describe. So A_i is the i-length tail of $[n]$. These $A_1, \ldots, A_n$ give Alice expectations

$$1 - p_n, 2 - 2p_{n-1} - p_n, \ldots, m - \sum_{i=m}^{1} ip_{n-(i-1)}, \ldots, n - \sum_{i=n}^{1} ip_{n-(i-1)}$$

The difference between the Alice-EV of consecutive strategies A_{m-1}, A_m is

$$(m - \sum_{i=m}^{1} ip_{n-(i-1)}) - ((m-1) - \sum_{i=m-1}^{1} ip_{n-(i-1)})$$

$$= 1 - mp_{n-(m-1)}.$$

Also, note that the above difference is a decreasing function, as

$$(1-mp_{n-(m-1)})-(1-(m+1)p_{n-m}) = m(p_{n-m}-p_{n-(m-1)})+p_{n-m} \geq m(p_{n-m}-p_{n-(m-1)}) = p_{n-m}-p_{n-m+1},$$

$$\geq 0$$

as p is nonincreasing.

So if we get to some m, and

$$EV(A_m - A_{m-1}) \geq 0, EV(A_{m+1}) - EV(A_m) \leq 0,$$

then A_m is an optimal counterstrategy. This is because the above reasoning shows us why A_m is at least as good as all other A_i.

6 Guessing with constraints

To recap, we now have that:

1. p is nonincreasing
2. Maximally exploitative A is some tail
3. if A_i is the $i-$ length tail, we can find the maximally exploitative tail A_m as the minimal m for which $1-(m+1)p_{n-m} \leq 0$, or

$$p_{n-m} \geq \frac{1}{m+1}.$$

Let's guess! Maybe Alice is indifferent to any A_i. This means $P_{n-m} = \frac{1}{m+1} \forall m$. This won't work because the harmonic series diverges to more than 1. This makes sense, too, because if this were the solution then Alice would get 0 squares. Put another way, Alice may as well choose A_1, which gets her at most 1 square for a score of

$$\frac{1}{n} \to 0 < \frac{4}{11}$$

(recall $\frac{4}{11}$, which we got from some strategy in the primal, is a lower bound on Alice's value.)

6.1 Adding constraints

We will now impose another restriction on p. Note that once it is $-EV$ for Alive to go from $A_m \to A_{m+1}$, Alice must stop. So Alice's optimal counterstrategy only sees $A_n, \ldots, A_{n-m}$ (A_{n-m} seen to check A_{m+1}, So the probabilities $p_1, \ldots, p_{n-(m+1)}$ are never viewed! So Bob should minimize them to put more probability into the squares that Alice does see. However, Bob's p must still be nonincreasing, so we require that

$$\min(p_1, \ldots, p_{n-(m+1)}) \geq p_{n-m} \geq \frac{1}{m+1},$$

and so minimizing these values sets them all to $\frac{1}{m+1}$.

To be precise, here is what we are doing:

1. We suppose one of Alice's maximally exploitative counterstrategy has length m, which requires exactly that
   (a)
   $$p_{n-(i-1)} \leq \frac{1}{i} \forall i \leq m,$$
   (b)
   $$p_n - m \geq \frac{1}{m+1}$$
   (c) p is nonincreasing (from before)
2. We construct a Bob strategy satisfying the above three bullet points that minimizes Alice's expectation. So Alice isn't incentivized to switch to another strategy, but we minimize Alice's expectation given that constraint.

So Bob's strategy p looks like

$$p_1 = \frac{1}{m+1}, \ldots, p_{n-m} = \frac{1}{m+1}, p_{n-(m-1)} \leq \frac{1}{m+1}, \ldots, p_n \leq \frac{1}{m+1}$$

for some m which will be the size of Alice's maximally exploitative tail strategy.

Now, remember that Alice's expectation given this strategy is penalized most by probability being concentrated in the earlier squares. So Bob should try to place as much probability as possible in the earliest squares, which means Bob's strategy p must place $\frac{1}{m+1}$ on squaress after square p_{n-m} until the amount of probability "left over" is less than $\frac{1}{m+1}$, at which point all the residual probability goes on the first square with no probability assigned to it yet.

But note that

$$\frac{1}{m+1} | m,$$

so there is actually no residual probability. Instead, the first

$$m + 1 \geq n - m$$

squares have probability $\frac{1}{m+1}$, and the last

$$n - (m + 1)$$

squares have 0 probability. Also note that from simply rearranging an inequality above, we get that

$$m \geq \frac{n-1}{2}.$$

We also have that $m < n (m \neq n)$ as assigning $\frac{1}{n+1}$ to every square doesn't even fully assign all of our probability. And $m = n - 1$, where Bob's distribution is uniform, does strictly better vs Alice as it assigns higher probability to every square. So

$$n > m \geq \frac{n-1}{2}.$$

7 The final computation

Then, let's consider for each m the Bob-strategy that has $m+1$ $\frac{1}{m+1}s$. Alice is incentivized to play A_m (take the last m squares) to maximally counterexploit Bob. Then Alice's EV is (from section 5: "Observations on the Dual", and plugging in A equal to the m-tail, and setting p as the distribution with the first $m+1$ elements equal to $\frac{1}{m+1}$) is

$$m - (\frac{m}{m+1} + \frac{m-1}{m+1} + \cdots + \frac{m-(2m-n)}{m+1})$$

As the first $m+1$ squares and the last m squares share $2m+1-n$ squares, so exactly $2m+1-n$ of the last n squares have probability $\frac{1}{m+1}$ assigned to them.

This equals

$$m - (\frac{(n)(\frac{2m-n+1}{2})}{m+1}).$$

We want to minimize this over all m. Note from before that $n > m \geq \frac{n-1}{2}$, but we can also note that plugging in

$$m = 0 \to \frac{n(n-1)}{2}$$

, and plugging in

$$m = n \to \frac{n}{2}.$$

The point of plugging in these two invalid values is to show that the values at the edges are large, so the expression is minimized when the derivative is zero. Plugging in $m = n-1$ directly yields

$$m = n-1 \to \frac{n-1}{2}$$

which is smaller than the value at n, and is intuitively the payoff when Alice plays A_{m-1} vs Bob's uniform. (It's also trivially the payoff when Alice plays A_m, which might be unintuitive. Why?) , Let's compute:

$$\frac{d}{dm}(m - (\frac{(n)(\frac{2m-n+1}{2})}{m+1}))$$

$$= 1 - \frac{n(m+1) - \frac{2mn-n^2+n}{2}}{(m+1)^2},$$

and setting this to 0 yields

$$1 = \frac{n(m+1) - \frac{2mn-n^2+n}{2}}{(m+1)^2}$$

$$\implies (m+1)^2 = n(m+1) - \frac{2mn-n^2+n}{2}$$

$$\implies m^2 + 2m + 1 = nm + n - mn + \frac{n^2}{2} - \frac{n}{2}$$

$$\implies m^2 + 2m + 1 = \frac{n^2+n}{2} \implies m^2 + 2m + 1 - \frac{n^2+n}{2} = 0,$$

and solving the quadratic yields

$$m = -1 + \sqrt{\frac{n^2+n}{2}}.$$

Let us suspend for a moment the fact that this is not necessarily an integer. We should really be taking the floor of this value, but for n large we can get close enough.

Plugging this into the other equation yields

$$(-1+\sqrt{\frac{n^2+n}{2}})-(\frac{(n)(\frac{2(-1+\sqrt{\frac{n^2+n}{2}})-n+1}{2})}{(-1+\sqrt{\frac{n^2+n}{2}})+1})$$

Recall we want the *proportion* of squares Alice gets to, so we want

$$\begin{aligned}
&\lim_{n\to\infty}\frac{(-1+\sqrt{\frac{n^2+n}{2}})-(\frac{(n)(\frac{2(-1+\sqrt{\frac{n^2+n}{2}})-n+1}{2})}{(-1+\sqrt{\frac{n^2+n}{2}})+1})}{n}\\
&=\lim_{n\to\infty}\frac{1}{\sqrt{2}}-\frac{\frac{-2+2\frac{n}{\sqrt{2}}-n+1}{2}}{-1+\frac{n}{\sqrt{2}}+1}\\
&=\lim_{n\to\infty}\frac{1}{\sqrt{2}}-\frac{\frac{n\sqrt{2}-n-1}{2}}{\frac{n}{\sqrt{2}}}\\
&=\lim_{n\to\infty}\frac{1}{\sqrt{2}}-\frac{n\sqrt{2}-n-1}{n\sqrt{2}}\\
&=\lim_{n\to\infty}\frac{1}{\sqrt{2}}-1+\frac{1}{\sqrt{2}}\\
&=\boxed{\sqrt{2}-1}.
\end{aligned}$$

7.1 An easier approach

One slightly easier way to approach this last computation, which is what I actually did when I solved this problem, is to treat the problem as though the space were continuous on [0,1] rather than discrete on $[1,2,\ldots,n]$ near the end. Then if Alice takes an $m=pn$ tail, Bob must play a uniform that extends $m+1=pn$ far. Then Alice's payout for some value of p is

$$E(X)=E(X|B\in(1-p,p))P(B\in(1-p,p))+E(x|B\in(0,1-p))P(B\in(0,1-p))$$

$$=\frac{2p-1}{2}\cdot\frac{2p-1}{p}+p\cdot\frac{1-p}{p}=p-1+\frac{1}{2p},$$

which is easily minimized via AM-GM to the value of

$$\sqrt{2}-1$$

as well.

8 Concise Proof

Consider the game where Alice does not give Bob her strategy. So Alice picks her strategy without informing Bob, and Bob places his block without informing Alice where.

Consider the NE of this game. Clearly, both strategies are mixed.

By the mini-max and maxi-min equivalence, the fair value of this game v is both equal to the max value Alice can achieve if she tells Bob her strategy, and the min value Bob can hold Alice to if Bob tells Alice his strategy. The former is precisely the answer to our question, and the latter is how we approach the problem.

Let Bob propose $p_1, \ldots, p_n$ and let Alice proposes $A \in 2^{[n]}$ maximally exploiting Bob. Let's call the indices of A

$$A_1, \ldots, A_m$$

where A_i is the ith largest element of A, and let's quickly note for sanity that $A_m = n$.

Now consider Alice's expected number of traversed squares. To do this, we create indicators $I_1, \ldots, I_m$ such that $I_i = 1$ if Alice traverses square A_i and 0 otherwise.

The probability Alice traverses square A_i is

$$E(I_i) = 1 - \sum_{j=1}^{i} p_{A_j},$$

and so Alice's expected score is

$$\sum_{i=1}^{n} E(I_i) = \sum_{i=1}^{m}(1 - \sum_{j=1}^{i} p_{A_j}) = m - \sum_{i=1}^{m}\sum_{j=1}^{i} p_{A_j},$$

$$= m - \sum_{i=1}^{m}(m - i + 1)p_{A_i}.$$

Theorem 5. *There exists a NE where Bob's supplied strategy is nonincreasing.*

Proof. Suppose Bob's strategy increases somewhere, so $p_i < p_{i+1}$ somewhere. There are three cases: For all the cases, let p' be the probability distribution $p_1, \ldots, p_{i-1}, p_{i+1}, p_{i+2}, \ldots, p_n$ (e.g. swap p_i, p_{i+1}). We construct some strategy (not necessarily the optimal one) α on p' that obtains at least as much on p' as A does on p.

1. Neither $i, i+1$ appears in the optimal counterstrategy A. In this case, taking $\alpha = A$ suffices.

2. Exactly i or $i+1$ appears in A. Then taking

   $$\alpha = (A \setminus \{i, i+1\}) \cup (\{i, i+1\} \setminus (\{i, i+1\} \cap A))$$

   Achieves the same value on p' as A on p. (This means: just take the other one.)

3. Both $i, i+1$ appear in A. Then taking $\alpha = A$ yields a strictly higher value on p' than A on p. The difference in the values we obtain is

   $$(-2p_{i+1} - p_i) - (-2p_i - p_{i+1}) = p_i - p_{i+1} > 0,$$

   which exceeds 0 by assumption.

So we can construct a strategy that does at least as well on p' as our optimal counterstrategy on p, so it suffices to look at decreasing distributions p. □

Theorem 6. *Any optimal counterstrategy $A = A_1, \ldots, A_m$ satisfies the requirement that $A_m = n$, and that between any two consecutive values A_i, A_{i+1} such that $A_i + 1 \neq A_{i+1}$,*

$$\min_{j \in [A_i+1, A_{i+1}-1]} p_j \geq \max(p_{A_i}, p_{A_{i+1}})$$

Proof. This theorem basically says that if Alice picks two points A_i, A_{i+1} and doesn't pick any points in between, then none of the interior probabilities are smaller than p_{A_i} or $p_{A_{i+1}}$. This is true as if one of the interior points j had a smaller probability than (WLOG) A_i, we could just pick

$$\alpha : \alpha_i = j, \alpha_i = A_i, \alpha_{j \neq i} = A_j$$

as a better strategy on p, as the squares in α just shares all the probabilities of being blocked with A except one of the squares has strictly lower probability of being blocked than before. To be clear, the difference in the obtained values between this constructed α and A is, with m defined as the length of A,

$$(m - i + 1)(p_{A_i} - p_j) > 0.$$

Also, it's trivial that $A_m = n$. □

We now combine our previous two theorems to obtain the next result:

Theorem 7. *If p is nonincreasing and A satisfies the requirement from the above theorem, then there exists a maximally exploitative counterstrategy A to p that is a tail of $[n]$. In other words,*

$$A = \{n - (m-1), n - (m-2), \ldots, n\}.$$

Proof. We have that A includes n. Now, suppose $|A| = m$, and A is not a tail. Let $n - j : j \in [m-1]$ be the largest index that isn't included in A. Let u be the largest index in A less than j. Then A is of the form

$$\ldots, u, n - (j-1), n - (j-2), \ldots, n.$$

Note that removing u from A and adding $n - j$ to A increases Alice's score by

$$j(p_u - p_{n-j}) > 0,$$

contradicting optimality of A.

So there is no largest index $n - j : j \in [m-1]$ that isn't included in A, so in particular the elements $n - (m-1), \ldots, n$ are all in A. Since by assumption A has m elements, these are all the elements of A. □

So fixing p, the optimal counterstrategy from Alice is in

$$\{\{n\}, \{n, n-1\}, \ldots, \{n, n-1, \ldots, 1\}\}.$$

Let's call these strategies $A_1, A_2, \ldots, A_n$ by the sizes of the sets they describe, so that A_i is the i-length tail of $[n]$. Then $A_1, \ldots, A_n$ give Alice expectations

$$1 - p_n, 2 - 2p_{n-1} - p_n, \ldots, m - \sum_{i=m}^{1} i p_{n-(i-1)}, \ldots, n - \sum_{i=n}^{1} i p_{n-(i-1)}$$

The difference between the Alice-EV of consecutive strategies A_{m-1}, A_m is

$$(m - \sum_{i=m}^{1} ip_{n-(i-1)}) - ((m-1) - \sum_{i=m-1}^{1} ip_{n-(i-1)})$$
$$= 1 - mp_{n-(m-1)}.$$

Also, note that the above difference is a decreasing function, as

$$(1-mp_{n-(m-1)})-(1-(m+1)p_{n-m}) = m(p_{n-m}-p_{n-(m-1)})+p_{n-m} \geq m(p_{n-m}-p_{n-(m-1)}) = p_{n-m}-p_{n-m+1},$$
$$\geq 0$$

as p is nonincreasing.

This means if we take $A_m : m$ is the minimal integer satisfying

$$EV(A_m - A_{m-1}) \geq 0, EV(A_{m+1}) - EV(A_m) \leq 0,$$

then A_m is an optimal counterstrategy.

Note that once it is $-EV$ for Alive to go from $A_m \to A_{m+1}$, Alice must stop. So Alice's optimal counterstrategy only sees $A_n, \ldots, A_{n-m}$ (A_{n-m} seen to check A_{m+1}, So the probabilities $p_1, \ldots, p_{n-(m+1)}$ are never viewed, and Bob should minimize them constrained to our solution class

$$\min(p_1, \ldots, p_{n-(m+1)}) \geq p_{n-m} \geq \frac{1}{m+1}.$$

Minimizing these values sets them all to $\frac{1}{m+1}$.

To be precise:

1. We suppose one of Alice's maximally exploitative counterstrategy has length m, which requires exactly that

   (a)
   $$p_{n-(i-1)} \leq \frac{1}{i} \forall i \leq m,$$

   (b)
   $$p_n - m \geq \frac{1}{m+1}$$

   (c) p is nonincreasing (from before)

2. We consider the class of solutions, indexed by m, with minimal Alice-EV. Bob will only play these strategies.

So Bob's strategy p looks like

$$p_1 = \frac{1}{m+1}, \ldots, p_{n-m} = \frac{1}{m+1}, p_{n-(m-1)} \leq \frac{1}{m+1}, \ldots, p_n \leq \frac{1}{m+1}$$

for some m which will be the size of Alice's maximally exploitative tail strategy.

We can quickly note that Alice's expectation given this strategy is penalized most by probability being concentrated in the earlier squares, so Bob should try to place as much probability as possible in the earliest squares, but

$$\frac{1}{m+1} | m.$$

So the unique way to maximize this is to set the first $m+1$ probabilities to $\frac{1}{m+1}$.

Then, let's consider for each m the Bob-strategy that has $m+1$ $\frac{1}{m+1}s$. Alice is incentivized to play A_m (take the last m squares) to maximally counterexploit Bob. Then Alice's EV is (from plugging in A equal to the m-tail, and setting p as the distribution with the first $m+1$ elements equal to $\frac{1}{m+1}$) is

$$m - (\frac{m}{m+1} + \frac{m-1}{m+1} + \cdots + \frac{m-(2m-n)}{m+1})$$

As the first $m+1$ squares and the last m squares share $2m+1-n$ squares, exactly $2m+1-n$ of the last n squares have probability $\frac{1}{m+1}$ assigned to them.

This equals

$$m - (\frac{(n)(\frac{2m-n+1}{2})}{m+1}).$$

We want to minimize this over all m. Note from before that $n > m \geq \frac{n-1}{2}$, but we can also note that plugging in

$$m = 0 \to \frac{n(n-1)}{2}$$

, and plugging in

$$m = n \to \frac{n}{2}.$$

The values at the edges are large, so the expression is minimized when the derivative is zero. Let's compute:

$$\frac{d}{dm}(m - (\frac{(n)(\frac{2m-n+1}{2})}{m+1}))$$

$$= 1 - \frac{n(m+1) - \frac{2mn-n^2+n}{2}}{(m+1)^2},$$

and setting this to 0 yields

$$1 = \frac{n(m+1) - \frac{2mn-n^2+n}{2}}{(m+1)^2}$$

$$\implies (m+1)^2 = n(m+1) - \frac{2mn-n^2+n}{2}$$

$$\implies m^2 + 2m + 1 = nm + n - mn + \frac{n^2}{2} - \frac{n}{2}$$

$$\implies m^2 + 2m + 1 = \frac{n^2+n}{2} \implies m^2 + 2m + 1 - \frac{n^2+n}{2} = 0,$$

and solving the quadratic yields

$$m = -1 + \sqrt{\frac{n^2+n}{2}}.$$

Let us suspend for a moment the fact that this is not necessarily an integer. We should really be taking the floor of this value, but for n large we can get close enough.

Plugging this into the other equation yields

$$(-1 + \sqrt{\frac{n^2+n}{2}}) - (\frac{(n)(\frac{2(-1+\sqrt{\frac{n^2+n}{2}})-n+1}{2})}{(-1+\sqrt{\frac{n^2+n}{2}})+1})$$

Recall we want the *proportion* of squares Alice gets to, so we want

$$\lim_{n\to\infty} \frac{(-1+\sqrt{\frac{n^2+n}{2}}) - (\frac{(n)(\frac{2(-1+\sqrt{\frac{n^2+n}{2}})-n+1}{2})}{(-1+\sqrt{\frac{n^2+n}{2}})+1})}{n}$$

$$= \lim_{n\to\infty} \frac{1}{\sqrt{2}} - \frac{\frac{-2+2\frac{n}{\sqrt{2}}-n+1}{2}}{-1+\frac{n}{\sqrt{2}}+1}$$

$$= \lim_{n\to\infty} \frac{1}{\sqrt{2}} - \frac{\frac{n\sqrt{2}-n-1}{2}}{\frac{n}{\sqrt{2}}}$$

$$= \lim_{n\to\infty} \frac{1}{\sqrt{2}} - \frac{n\sqrt{2}-n-1}{n\sqrt{2}}$$

$$= \lim_{n\to\infty} \frac{1}{\sqrt{2}} - 1 + \frac{1}{\sqrt{2}}$$

$$= \boxed{\sqrt{2}-1}.$$

9 Alice's Strategy in the primal?

Let me know if you compute Alice's Strategy in the primal! This is a really cool problem in that I don't have the faintest idea of what the strategy looks like.