

Warp Tile Problem


Neal Sodhi

November 2025

1 The Game

A board consists of tiles $0, 1, \dots, N$. Player 1 begins on tile 0 and on each turn jumps forward any positive number of tiles; the game ends when Player 1 reaches tile N . The *score* is the number of jumps, which Player 1 seeks to maximize.

Player 2, who seeks to minimize the expected score, places a *warp* on one of the tiles $1, \dots, N - 1$. Should Player 1 land on the warp, Player 1 is sent directly to tile N and the game ends. A warp therefore does the most damage where Player 1 is likely to land and still has much of the board left to lose.



$N = 10$, warp on tile 6. Player 1 jumps $0 \rightarrow 2 \rightarrow 3 \rightarrow 6$, lands on the warp, and is sent to the end: score 3.

Player 2 chooses the warp with full knowledge of Player 1's strategy. Every deterministic strategy therefore scores 1, since the warp is simply placed on the first tile Player 1 would land on. Player 1 must randomize, and we assume that Player 2 knows the distribution Player 1 plays but not the outcome of Player 1's source of randomness.

Formulation. A pure strategy for Player 1 is the set $S \subseteq \{1, \dots, N - 1\}$ of tiles on which it lands, and a mixed strategy is a random such set. For a warp on tile w ,

$$\text{score}(S, w) = \begin{cases} |S \cap [1, w]| & \text{if } w \in S, \\ |S| + 1 & \text{if } w \notin S. \end{cases}$$

The value of the game is $v(N) = \max_S \min_w \mathbb{E}[\text{score}(S, w)]$, the maximum being taken over all distributions of S .

Question. What is $\lim_{N \rightarrow \infty} \frac{v(N)}{N}$?

Example. Suppose Player 1 jumps 1 or 2 tiles, each with probability $\frac{1}{2}$. The probability u_i of landing on tile i satisfies $u_i = \frac{1}{2}u_{i-1} + \frac{1}{2}u_{i-2}$, so $u_1 = \frac{1}{2}$, $u_2 = \frac{3}{4}$, and $u_i \rightarrow \frac{2}{3}$. Player 2 places the warp on tile 2. With probability $\frac{3}{4}$ the game ends at once; otherwise Player 1 makes $\frac{2}{3}N + O(1)$ jumps in expectation. The expected score is therefore $\frac{1}{4} \cdot \frac{2}{3}N + O(1) = \frac{1}{6}N + O(1)$.

2 The Continuous Version

Rescale the board to $[0, 1]$, so that tile i sits at $x = i/N$, and measure scores as a fraction of N . Write $v = \lim_{N \rightarrow \infty} v(N)/N$ for the value.

A strategy for Player 1 is now a random landing set $S \subseteq [0, 1]$, described to first order by its *landing profile*

$$p(x) = \mathbb{P}(x \in S),$$

the probability of landing at (a tile near) x . For a warp at y , the normalized score is $|S \cap [0, y]|$ if $y \in S$ and $|S|$ if $y \notin S$, where $|A|$ denotes the normalized number of landings in A . Taking expectations,

$$\mathbb{E}[\text{score}] = \int_0^y p(x) dx + \int_y^1 \mathbb{P}(x \in S, y \notin S) dx.$$

Independent Landing. Every profile $p : [0, 1] \rightarrow [0, 1]$ is realized by independent landing: Player 1 lands on each tile independently, with probability $p(x)$. Landing at y then carries no information about later tiles, so $\mathbb{P}(x \in S, y \notin S) = p(x)(1 - p(y))$, and a warp at y yields

$$\mathcal{S}(p, y) = \int_0^y p dx + (1 - p(y)) \int_y^1 p dx = \int_0^1 p dx - p(y) Q(y), \quad Q(y) := \int_y^1 p dx.$$

Hence Player 1 can guarantee the value of the variational problem

$$v \geq \sup_{p: [0,1] \rightarrow [0,1]} \left[\int_0^1 p dx - \sup_{y \in [0,1]} p(y) Q(y) \right]. \quad (1)$$

The penalty $p(y)Q(y)$ is exactly the tension noted in Section 1: the chance of landing at y , times the board that remains to be lost. We shall see in Section 3 that (1) holds with equality, so that correlated strategies do no better.

Example (Continued). “Jumping 1 or 2 tiles” has no literal meaning on the continuous board, but the walk’s landing probabilities settle at $u_i \rightarrow \frac{2}{3}$, and its natural analogue is the constant profile $p \equiv \frac{2}{3}$. This profile gives $\mathcal{S} = \frac{2}{3} - \frac{2}{3} \cdot \frac{2}{3} = \frac{2}{9}$ for the worst warp, near $y = 0$. The walk itself achieves only $\frac{1}{6}$, because its landing probabilities have an early spike ($u_2 = \frac{3}{4}$) which the warp exploits. Independent landing has no such spikes, and for it the continuous formula is exact.

3 The Case of One Warp

Theorem 1. $v = \sqrt{2} - 1 \approx 0.4142$. More precisely, $v(N) = (\sqrt{2} - 1)N + O(1)$.

Strategy of Proof: Duality. For fixed N the game is finite and zero-sum, so by von Neumann's minimax theorem

$$v(N) = \max_S \min_w \mathbb{E}[\text{score}(S, w)] = \min_w \max_S \mathbb{E}[\text{score}(S, w)],$$

the outer optimization in each case being over probability distributions. On the left Player 1 commits first to a random landing set; on the right, in the dual problem, Player 2 commits first to a random warp. It is therefore enough to exhibit one strategy for each player: a random landing set that scores at least V against every warp proves $v \geq V$, and a random warp that holds every landing set to at most V proves $v \leq V$. When the two bounds agree, V is the value, and the two strategies form a saddle point, each a best response to the other.

Player 1: Choosing the Landing Set in Advance. Player 1 decides before the game begins which tiles to land on, and then simply walks the chosen path: each tile i is placed in S independently with probability $p(i/N)$, and Player 1 jumps from each chosen tile to the next, finishing at tile N . The strategy is thus a single function, the landing profile p , and by Section 2 a warp at y yields $\int_0^1 p - p(y)Q(y)$. To keep the penalty pQ below a bound c while making $\int p$ as large as possible, take $p = 1$ on $[1 - c, 1]$, where $pQ = 1 - y \leq c$, and $pQ = c$ before that. Then $-QQ' = c$, so $Q(0)^2 = 2c - c^2$, and the payoff $\sqrt{2c - c^2} - c$ is maximized at $c = 1 - 1/\sqrt{2}$, giving

$$p^*(x) = \frac{\sqrt{2} - 1}{\sqrt{1 - (4 - 2\sqrt{2})x}} \quad (0 \leq x \leq \tfrac{1}{\sqrt{2}}), \quad p^*(x) = 1 \quad (x \geq \tfrac{1}{\sqrt{2}}).$$

Then $p^*Q \equiv 1 - \frac{1}{\sqrt{2}}$ on $[0, \frac{1}{\sqrt{2}}]$, p^*Q is smaller thereafter, and $\int_0^1 p^* = \frac{1}{\sqrt{2}}$. Every warp therefore leaves Player 1 a score of at least $\sqrt{2} - 1$, and $v \geq \sqrt{2} - 1$.

Player 2: Placing the Warp in the Dual. In the dual problem Player 2 commits first, to a random position for the warp, and Player 1 replies with the best landing set against it. Let the warp be uniform on the first $a = 1/\sqrt{2}$ of the board (on tiles $1, \dots, \lfloor N/\sqrt{2} \rfloor$). Against this warp only the amount of $[0, a]$ that a landing set covers matters, not which tiles it contains or how they are correlated. Indeed, let S be any landing set covering a fraction u of $[0, a]$. The warp meets S with probability u , and when it does, it falls on a uniformly random landing of S . Hence


$$\mathbb{E}[\text{score}] \leq a \left(u - \frac{u^2}{2} \right) + (1 - u)(1 - a), \quad (2)$$

which is maximized at $1 - u = (1 - a)/a$, where it equals $\frac{a}{2} + \frac{(1-a)^2}{2a} = \sqrt{2} - 1$. Thus $v \leq \sqrt{2} - 1$.

The Value. The two bounds coincide, so $v = \sqrt{2} - 1$ and $(p^*, \text{uniform warp})$ is a saddle point. Against p^* , every warp in $[0, \frac{1}{\sqrt{2}}]$ does the same damage and later warps do less; against the uniform warp, a landing set is optimal precisely when it covers a fraction $2 - \sqrt{2}$ of $[0, \frac{1}{\sqrt{2}}]$ and all of the board beyond, as p^* does. Since the upper bound holds for every strategy, correlation cannot help Player 1, and (1) is an equality.

N	10	20	40	80	∞
$v(N)$	4.4286	8.5714	16.8571	33.4286	
$v(N)/N$	0.4429	0.4286	0.4214	0.4179	0.4142

Exact values on finite boards, computed by linear programming, approach the limit from above.



Player 1's optimal landing profile p^* ; Player 2's warp is uniform on the shaded region.

4 The Case of k Warps


Suppose now that Player 2 places k warps, on a set $W \subseteq \{1, \dots, N-1\}$ of k distinct tiles, and that landing on any of them ends the game. With $w^* = \min(S \cap W)$ the first warp Player 1 lands on,

$$\text{score}(S, W) = \begin{cases} |S \cap [1, w^*]| & \text{if } S \cap W \neq \emptyset, \\ |S| + 1 & \text{if } S \cap W = \emptyset, \end{cases} \quad v_k(N) = \max_S \min_{|W|=k} \mathbb{E}[\text{score}(S, W)].$$

In the continuous version with independent landing, warps at $y_1 < \dots < y_k$ yield

$$\mathcal{S}(p; y_1, \dots, y_k) = \int_0^1 p(x) \prod_{i: y_i < x} (1 - p(y_i)) dx.$$

Challenge. For every $k \geq 1$, determine $v_k = \lim_{N \rightarrow \infty} \frac{v_k(N)}{N}$, together with optimal strategies for both players.



$N = 14, k = 3$. Player 1 jumps over the warp on tile 3 but lands on the one on tile 8: score 4.

Remark. With several warps the landing profile no longer determines the payoff. What matters is the probability of meeting at least one warp, which depends on the correlations among Player 1's landings, and it is not evident that independent landing remains optimal.