On an Erdős-Szekeres Game

Lara Pudwell

November 26, 2024

Abstract

We consider a 2-player permutation game inspired by the celebrated Erdős-Szekeres Theorem. The game depends on two positive integer parameters a and b and we determine the winner and give a winning strategy when $a \geq b$ and $b \in \{2, 3, 4, 5\}$.

1 Introduction

Let $\mathcal{S}_n$ be the set of all permutations of $\{1, 2, \ldots, n\}$. We say that permutation $\pi \in \mathcal{S}_n$ *contains* $\rho \in \mathcal{S}_m$ if π contains a subsequence $\pi_{i_1}\pi_{i_2}\cdots\pi_{i_m}$ such that $i_1 < i_2 < \cdots < i_m$ and $\pi_{i_a} < \pi_{i_b}$ if and only if $\rho_a < \rho_b$. Otherwise we say π *avoids* ρ. One of the oldest theorems that can rephrased in terms of permutation patterns is the Erdős-Szekeres Theorem [1], which was first published in 1935, and is phrased in terms of patterns as Theorem 1.

Theorem 1. *Any permutation of length $n \geq (a-1)(b-1)+1$ contains either an increasing subsequence of length a or a decreasing subsequence of length b.*

Since we are primarily interested in the monotone patterns, we refer to $12\cdots a$ as I_a and $b\cdots 1$ as J_b.

While there are a variety of proofs of 1, one of the most concise was given by Seidenberg [7] in 1959, using an application of the pigeonhole principle, which we revisit here:

Proof of Theorem 1. Let $\pi = \pi_1 \cdots \pi_n \in \mathcal{S}_n$. For each π_i associate an ordered pair (a_i, d_i) where a_i is the length of the longest increasing subsequence of π ending in π_i and d_i is the length of the longest decreasing subsequence of π ending in π_i. Clearly, for all i, $a_i \geq 1$ and $d_i \geq 1$ since the digit π_i is itself an increasing (resp. decreasing) subsequence of length 1.

However, we also have that if $i \neq j$, then $(a_i, d_i) \neq (a_j, d_j)$. This follows from the fact that the digits of π are distinct members of $\{1, \ldots, n\}$. Without loss of generality, suppose $i < j$. If $\pi_i < \pi_j$, then $a_i < a_j$ since appending π_j onto the increasing subsequence of length a_i ending at π_i produces an increasing subsequence of length $a_i + 1$ ending at π_j. Similarly, if $\pi_i > \pi_j$, then $d_i < d_j$.

We have n distinct ordered pairs of positive integers associated with π. If $a_i \geq a$ or $d_i \geq b$ for some i then π contains an I_a or a J_b pattern. So, if π

avoids I_a and J_b, then $1 \leq a_i \leq a-1$ and $1 \leq d_i \leq b-1$ for all i, which means $n \leq (a-1)(b-1)$. Taking the contrapositive, if $n \geq (a-1)(b-1)+1$, then π contains either I_a or J_b as a pattern.

□

In 1983, Harary, Sagan, and West [2] studied a game based on Theorem 1 with rules as follows: Consider the set of integers $\{1, 2, \ldots, ab+1\}$. Two players take turns selecting numbers from this set until either an increasing subsequence of length $a+1$ or a decreasing subsequence of length $b+1$ is formed. In the achievement version of the game, the first player to complete such a subsequence wins. In the avoidance version of the game, the first player to complete such a subsequence loses. Their analysis is computer-aided, and is limited by the computer memory available at the time. In particular, each state of the game can be labeled as winning or losing for player 1 based on an analysis of subsequent possible moves. They determined the winning player for games where $ab+1 \leq 15$, and since the tree of game states grows exponentially in a and b, they predicted that it is prohibitive to push computer analysis much further. While they determined the winner of the game for particular small values of a and b, they were unable to find a general winning strategy.

In 2021, the current author proposed a variation on this game as part of a public lecture [4]. Instead of proposing a 2-player competitive game, the new version is collaborative, and instead of picking specific digits from a set, players append a new digit onto a permutation prefix. In particular, after the first n moves, the current game state is a permutation $\pi \in \mathcal{S}_n$, and the next player can play any digit in $\{1, \ldots, n+1\}$ as their move. After the first move, the game is always in state $\pi = 1$. However, on subsequent moves, if a player chooses $m \in \{1, \ldots, n+1\}$, we let $A = \{i | \pi_i \leq m-1\}$ and let $B = \{i | \pi_i \geq m\}$ and the new game state becomes $\pi' = \pi'_1 \pi'_2 \cdots \pi'_n m \in \mathcal{S}_{n+1}$ where $\pi'_i = \pi_i$ if $i \in A$ and $\pi'_i = \pi_i + 1$ if $i \in B$. Notice that $\pi'_1 \pi'_2 \cdots \pi'_n$ forms a π pattern in π', that is, the relative order of the initial digits is unchanged. The game ends when the most recent move completes an I_a pattern or a J_b pattern. The collaborative version of this game was generated as a teaching tool to build intuition about Theorem 1 for audience members who were unfamiliar with the theorem. Optimal collaborative game play achieves a permutation of length $(a-1)(b-1)+1$.

In [5], the current author analyzed a question that arose from continuing to use this collaborative game as a teaching tool: how many ways are there to play the game optimally? For example, if $a \geq b = 2$, optimal game play is to build an increasing permutation of length $(a-1)(2-1)+1 = a$, and there is only one such permutation. However, if $a \geq b > 2$, there are more maximum length permutations. These permutations can be counted by mapping permutations avoiding I_a and J_b to pairs of same-shaped standard Young tableaux with at most $a-1$ columns and $b-1$ rows using the RSK algorithm, and then by using the hook length formula to count the number of such pairs. This process was elaborated on by Stanley [6], who cited work of Schensted [8] in a solution to a problem posed in *American Mathematical Monthly* in 1969. While this strategy answers the general enumeration question it is worth noting that the $a \geq b = 3$ case has a particularly nice enumeration: there are $(C_{a-1})^2$ permutations of length $2(a-1)$ avoiding both I_a and J_3, where C_n denotes the nth Catalan number, i.e., $C_n = \frac{\binom{2n}{n}}{n+1}$, and therefore there are $(C_{a-1})^2 \cdot a$ maximum length permutations in the $a \geq b = 3$ game. The current author proved this known enumeration via simpler techniques, i.e., using a bijection between the $(C_{a-1})^2$ permutations in question and pairs of parentheses arrangements. Of note, the bijection focused on tracking the positions and the values of the left-to-right maxima of the optimal play permutations.

In this paper, we consider the 2-player competitive version of this game. In particular, as in the collaborative version, after the first n moves, the current game state is a permutation of length n, and the next player can play any digit in $\{1, \dots, n+1\}$ as their move, appending it to the current permutation. We will primarily analyze the game where the first player who completes an I_a pattern or a J_b pattern loses. In contrast to the analysis of Harary, Sagan, and West's game, where no general strategy was found, we give strategies that work for specific choices of b but for any $a \geq b$.

The organization of the rest of this paper is as follows. In Section 2 we develop a visual way of representing game moves, inspired by Seidenberg's proof of Theorem 1, which was given above. In Sections 3, 4, 5, and 6 we give a general winning strategy for the cases where $a \geq b$ and $2 \leq b \leq 5$. In Section 7 we consider the version of the game where the first player who completes an I_a pattern or a J_b pattern wins the game. Finally, we conclude with topics for future investigation.

2 Representation of Moves

As described in Section 1, we consider a 2-player game where players take turns appending a new digit onto a permutation. On the nth turn of the game, a player may play any number in $\{1, \dots, n\}$. After the nth turn of the game, the current game state is a permutation π of length n. For all $i \geq 1$, the pattern formed by the first i digits remains unchanged as the game progresses, but we need to track increasing and decreasing subsequences within the permutation being constructed. We call a game that ends when a player completes an I_a or a J_b pattern an (a, b)-*game*. The Seidenberg proof of Theorem 1 tracks increasing and decreasing subsequences using ordered pairs of positive integers. Motivated by this representation we define the *board* of an (a, b)-game to be a grid of cell cells with $b-1$ rows and $a-1$ columns. Each cell is indexed by the ordered pair (c, r) where c denotes the column number and r denotes the row number of the cell.

Now, for each move of an (a, b)-game, we shade a cell (c, r) if the longest increasing subsequence of π ending in π_n has length c and the longest decreasing subsequence of π ending in π_n has length r. Figure 1 shows the board in the $(6, 5)$-game corresponding to $\pi = 163425$. To check, $\pi_1 = 1$ corresponds to $(1, 1)$,

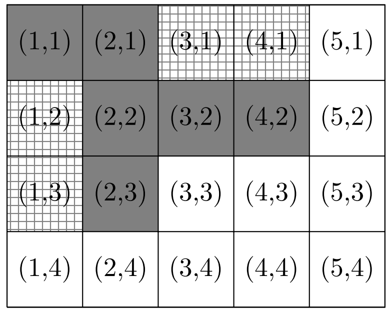

Figure 1: The board corresponding to $\pi = 163425$ in a $(6,5)$-game

$\pi_2 = 6$ corresponds to $(2,1)$, $\pi_3 = 3$ corresponds to $(2,2)$, $\pi_4 = 4$ corresponds to $(3,2)$, $\pi_5 = 2$ corresponds to $(2,3)$, and $\pi_6 = 5$ corresponds to $(4,2)$. Note that the increasing subsequence length is given first for consistency of notation with the name of (a,b)-game, while these values correspond to column numbers (rather than row numbers) in the board for vertically efficient use of the page throughout this manuscript.

At the nth move, we know that for each $j < n$ either $\pi_j < \pi_n$ (in which case the column number of the nth shaded cell is larger than the column number of the jth shaded cell) or $\pi_j > \pi_n$ (in which case the row number of the nth shaded cell is larger than the row number of the jth shaded cell). So, each newly shaded cell must be pairwise further to the right or further below each previously shaded cell. To this end, in addition to the shaded cells we call a cell (c^*, r^*) *eliminated* if (c,r) is shaded and both $c^* \leq c$ and $r^* \leq r$. The cells $(3,1)$, $(4,1)$, $(1,2)$, and $(1,3)$ are eliminated in Figure 1 and thus marked with hatching. These are cells that are ineligible to become shaded in future turns, based on the permutation formed so far in the game. The shaded cells are a subset of the eliminated cells at any point of the game.

We now make some observations about the shaded region of a board at any point in an (a,b)-game. First, by definition of eliminated cells, at any point in the game, the set of shaded and eliminated board cells forms one contiguous region. Moreover, this region tells us clearly which cells are available to be shaded by the next move of the game. We refer to cells that are not shaded and not eliminated as *open* cells.

Proposition 1. Suppose S is the set of shaded and eliminated cells after the nth turn of the (a,b)-game. Then move $n+1$ corresponds to shading an open cell that is edge-adjacent to a cell is S. Moreover, every open cell that is edge-adjacent to S corresponds to a possible next move.

Proof. We consider the possible values of (c,r) corresponding to π_{n+1}. Let i be the maximum column number of a cell in S and let d be the maximum row number of a cell in S.

First, consider the extreme cases where $\pi_{n+1} = 1$ or $\pi_{n+1} = n + 1$. If $\pi_{n+1} = 1$, then playing π_{n+1} corresponds to cell $(1, d+1)$. This is edge adjacent to S since some cell in S has row number d, and therefore cell $(1, d)$ is either shaded or eliminated. Similarly, if $\pi_{n+1} = n + 1$, playing π_{n+1} corresponds to cell $(i+1, 1)$. This is also edge adjacent to S since some cell has column number i, and therefore cell $(i, 1)$ is either shaded or eliminated.

Now suppose that $\pi_{n+1} = j$ $(1 < j < n + 1)$ corresponds to shading the cell (c, r). We consider what cell would be shaded if $\pi_{n+1} = j + 1$ instead in four cases. In all cases, the digits $\pi_1 \cdots \pi_n$ have the same relative order. The only difference is that if $\pi_{n+1} = j$, there exists $1 \leq k \leq n$ such that $\pi_k = j + 1$, and now when $\pi_{n+1} = j + 1$, instead $\pi_k = j$.

- π_k played no role in the increasing subsequence of length c or the decreasing subsequence of length r ending in $\pi_{n+1} = j$, and so $\pi_{n+1} = j$ and $\pi_{n+1} = j + 1$ result in shading the same (c, r) cell.

- π_k corresponds to shading cell $(c, r - 1)$ and π_k was part of the decreasing subsequence of length r when $\pi_{n+1} = j$. In the situation where $\pi_{n+1} = j+1$, $\pi_k\pi_{n+1}$ are the last two digits in an increasing subsequence of length $c + 1$, while π_{n+1} ends a decreasing subsequence of length $r - 1$. In other words, $\pi_{n+1} = j + 1$ corresponds to shading $(c + 1, r - 1)$.

- π_k corresponds to shading cell $(c^*, r-1)$ where $c^* < c$ and $\pi_k = j+1$ played a role in forming the decreasing subsequence of length r when $\pi_{n+1} = j$. Now that $\pi_{n+1} = j + 1$, π_{n+1} completes a decreasing subsequence of length $r-1$, while c remains unchanged. In other words, $\pi_{n+1} = j+1$ corresponds to shading $(c, r - 1)$.

- π_k corresponds to shading cell (c, r^*) where $r^* < r - 1$. Since $r^* < r - 1$, there must be a different subsequence that contributed to the decreasing pattern of length r when $\pi_{n+1} = j$, and so when $\pi_{n+1} = j + 1$, the decreasing subsequence ending in π_{n+1} remains the same. However, the increasing subsequence length goes up by 1. In other words, $\pi_{n+1} = j + 1$ corresponds to shading cell $(c + 1, r)$.

These four cases decribe the set of possible cells that can be shaded by various choices of π_{n+1}. We gave specific examples that showed $(1, d + 1)$ and $(i + 1, 1)$ are possible. We also saw that each open cell that could be shaded by a choice of π_{n+1} differs in row and/or column number by at most 1 from another open cell. In fact, the only time when both numbers change is when they are immediately below and immediately right of a shaded cell that forms a southeast corner of S. This uniquely describes the open cells that are edge adjacent to S.

□

As an example, consider Figure 2, which shows the possible ordered pairs corresponding to next moves in a $(6, 5)$-game whose current permutation is $\pi = 163425$. The left side of the figure shows the plots of the points (i, π_i) with

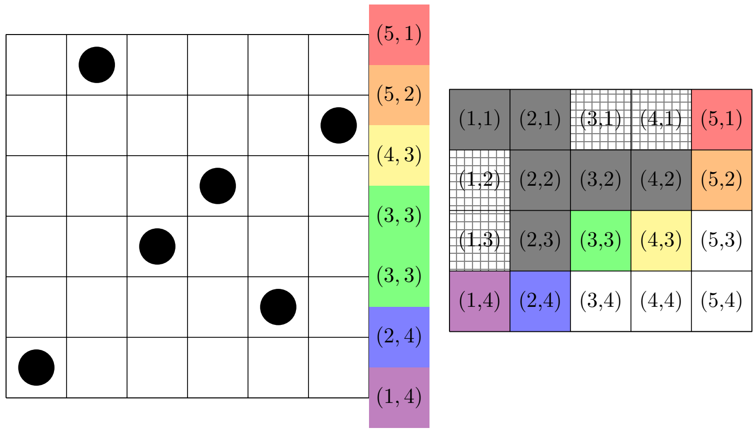

Figure 2: Ordered pairs representing possible next moves in a $(6, 5)$-game where $\pi = 163425$.

ordered pairs given for various choices of π_7, while the right side of the figure shows the shaded and eliminated cells after the sixth turn, and highlights the cell corresponding to choices of π_7.

Finally, the player who shades the $(a-1, b-1)$ cell is the winner since this cell being shaded means that the entire board is either shaded or eliminated, and the next player must add a digit to the permutation that either completes an I_a pattern or a J_b pattern. Now that we have made this translation from digits to boards, we may play the Erdős-Szekeres game as a game of shading cells on a board, rather than thinking merely in terms of one-line permutation notation.

In summary, we can rephrase the Erdős-Szekeres game as follows:

Game Rules. Consider a $(b-1) \times (a-1)$ array of cells. Players 1 and 2 take turns as follows:

- Player 1 begins by shading the cell in the top left corner.
- For each subsequent move, a player shades a cell that is edge-adjacent to the shaded region. All cells that are above and/or right of their chosen cell should also be shaded (i.e. eliminated).
- Players alternate taking turns until the board is full. The player who claims the bottom right corner wins.

Notice that there are many permutations that may correspond to the same shaded board. As a small example, both $\pi = 132$ and $\pi = 312$ correspond to a board where a 2×2 region has been shaded. However, these are both permutations where the most recent digit corresponds to the label $(2, 2)$, and so the same amount of progress has been made towards forming an I_a or J_b pattern.

In terms of tracking a win or loss in the permutation game, no information has been lost.

In the following sections, we articulate a strategy for winning this permutation game in terms of board shading.

3 Strategy when b=2

Suppose $a \geq b = 2$. Using the representation described in Section 2, the game board can be visualized as a one-row, $(a-1)$-column grid, and each player colors the left-most unclaimed cell on their turn. Once the cell in column $(a-1)$ is claimed (by player 2 if a is odd, or by player 1 if a is even), the other player has no more legal moves and loses.

In terms of permutations, each player will play a new largest digit on their turn since playing any smaller digit creates a J_2 pattern and automatically loses the game. The resulting game permutation is $\pi = I_a$, and the game ends on the ath turn, leading to a loss for player 1 if a is odd and a loss for player 2 if a is even.

4 Strategy when b=3

We now consider the $(a, 3)$-game ($a \geq 3$), and we give a winning strategy both in terms of board shading and in terms of permutation digits.

Theorem 2. *Player 1 has a winning strategy in the $(a, 3)$-game where $a \geq 3$.*

Proof. The game board in this situation is a $2 \times (a-1)$ grid. After each of their first $a-2$ moves, player 1 can produce a board where the first i cells of row 1 are shaded and the first $i-1$ cells of row 2 are shaded, as illustrated in Figure 3. To start, player 1's first move shades 1 cell in row 1 and 0 cells in row 2. After that, there are only two possible moves:

- If player 2 shades the leftmost open cell in row 2, player 1 shades the leftmost open cell in row 1.
- If player 2 shades the leftmost open cell in row 1, player 1 shades the leftmost open cell in row 2.

For the end game, after $(a-2)$ moves for player 1 and $(a-3)$ moves for player 2, all cells except for $(a-1, 1)$, $(a-2, 2)$, and $(a-1, 2)$ are shaded.

Because of Proposition 1, player 2 must shade either $(a-1, 1)$ or $(a-2, 2)$. In either case, player 1 can shade $(a-1, 2)$ on their next move and force a player 2 loss.

□

In this straightforward situation, shading a cell in row 1 corresponds to playing a new left-to-right maximum, while shading a cell in row 2 corresponds to playing a non-left-to-right-maximum, in keeping with observations of maximum

Figure 3: A next-player-loss state for the $(a, 3)$-game.

length I_a and 321-avoiders made in [5]. In terms of permutations, the winning strategy is as follows: Let $m_1 < m_2 < \cdots < m_\ell = n$ be the values of the left-to-right maxima of π at the start of player 1's turn.

- If player 2's move was not m_ℓ, play $\pi_{n+1} = n + 1$.
- If player 2's move was m_ℓ, play $m_{\ell-1}$.

Finally, when π has length $2a - 4$, regardless of what player 2's most recent move was, player 1 plays $\pi_{2a-3} = 2a - 4$, corresponding to cell $(a - 1, 2)$. On their next move, player 2 will complete either an I_a pattern (by playing $\pi_{2a-2} \geq 2a - 3$) or J_3 pattern (by playing $\pi_{2a-2} \leq 2a - 4$) to lose the game.

5 Strategy when b=4

We now consider the $(a, 4)$-game $(a \geq 4)$. This is the first case where the game board has sufficiently many rows that the same shading can be obtained by more than one permutation, and so we only describe it in terms of shaded boards. However, an interested player – thinking in terms of permutations – can consider the ordered pairs corresponding to each digit in the game so far to determine a next digit that follows the strategy given here.

Theorem 3. *Player 1 has a winning strategy in the $(a, 4)$-game where $a \geq 4$.*

Proof. The game board in this situation is a $3 \times (a - 1)$ grid.

To begin the game, player 1 shades cell $(1, 1)$. Regardless of whether player 2 shades cell $(1, 2)$ or cell $(2, 1)$, player 1 responds by shading cell $(2, 2)$.

After this opening sequence, player 1 can always end their turn with a shaded board of one of the following two forms, illustrated in Figure 4:

1. Rows 1 and 2 have k shaded cells and row 3 has $k-2$ shaded or eliminated cells for some $k \geq 2$.

2. Row 1 has k shaded cells and rows 2 and 3 have $k-1$ shaded or eliminated cells for some $k \geq 2$.

Note that the opening sequence produces a board that fits the first case. Now we consider each available move to player 2 and how player 1 may respond.

In the first case, player 2 has four available moves: $(k + 1, 1)$, $(k + 1, 2)$, $(k - 1, 3)$, or $(k, 3)$.

- If player 2 plays $(k+1,1)$, player 1 responds by playing $(k,3)$. This results in a board with $k+1$ shaded or eliminated cells in the first row, and k shaded or eliminated cells in rows 2 and 3, which fits the second case.
- If player 2 plays $(k+1,2)$, then $(k+1,1)$ is eliminated, and player 1 responds by playing $(k-1,3)$. This results in a board with $k+1$ shaded or eliminated cells in the first two rows and $k-1$ shaded or eliminated cells in row 3, which fits the first case.
- If player 2 plays $(k,3)$, then $(k-1,3)$ is eliminated, and player 1 responds by playing $(k+1,1)$. This results in a board with $k+1$ shaded or eliminated cells in row 1 and k shaded or eliminated cells in rows 2 and 3, which fits the second case.
- If player 2 plays $(k-1,3)$, then player 1 responds by playing $(k+1,2)$ which eliminates $(k+1,1)$. This results in a board with $k+1$ shaded or eliminated cells in the first two rows, and $k-1$ shaded or eliminated cells in the last row, which fits the first case.

Similarly, in the second case, player 2 has three available moves: $(k+1,1)$, $(k,2)$, or $(k,3)$.

- If player 2 plays $(k+1,1)$, player 1 responds by playing $(k,3)$ which eliminates $(k,2)$. This results in a board with $k+1$ shaded or eliminated cells in row 1 and k shaded or eliminated cells in rows 2 and 3, which fits the second case.
- If player 2 plays $(k,2)$, player 1 responds by playing $(k+1,2)$, which eliminates $(k+1,1)$. This results in a board with $k+1$ shaded or eliminated cells in rows 1 and 2 and $k-1$ shaded or eliminated cells in row 3, which fits the first case.
- if player 3 plays $(k,3)$, then $(k,2)$ is eliminated, and player 1 responds by playing $(k+1,1)$. This results in a board with $k+1$ shaded or eliminated cells in row 1 and k shaded or eliminated cells in rows 2 and 3, which fits the second case.

The endgame begins when row 1 has $a-2$ shaded or eliminated cells.

In the first case, row 2 also has $a-2$ shaded or eliminated cells and row 3 has $a-4$ shaded or eliminated cells, so the remaining unshaded region is as in Figure 4 (a). Player 2's options are $(a-1,1)$, $(a-1,2)$, $(a-3,3)$, or $(a-2,3)$. If player 2 plays $(a-1,2)$ or $(a-2,3)$, player 1 can play $(a-1,3)$ and win the game. So player 2 will play $(a-1,1)$ or $(a-3,3)$. Whichever of these two options player 2 takes, player 1 takes the other move. This forces player 2 to play $(a-1,2)$ or $(a-2,3)$ on their next move, and player 1 takes the corner cell of $(a-1,3)$, forcing a player 2 loss.

In the second case, rows 2 and 3 have $a-3$ shaded or eliminated cells, so the remaining unshaded region is as in Figure 4 (b). Player 2's options are $(a-1,1)$,

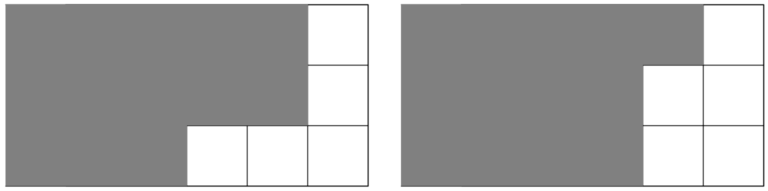

(a) (b)

Figure 4: Two next-player-loss states of the $(a, 4)$-game.

$(a-2, 2)$, or $(a-2, 3)$. If player 2 plays $(a-2, 3)$, player 1 can play $(a-1, 3)$ and win the game, or player 2 will play $(a-1, 1)$ or $(a-2, 2)$. In either case, player 1 plays the other move. This forces player 2 to play $(a-1, 2)$ or $(a-2, 3)$ on their next move, and player 1 takes the corner cell of $(a-1, 3)$, forcing a player 2 loss.

□

Similar to a $(a, 3)$-game, in a $(a, 4)$-game, player 1 has a clear deterministic response no matter what player 2 does. Each of these responses gets closer and closer to player 2 taking either $(a-1, 2)$ or $(a-2, 3)$, so that player 1 claims the bottom right corner of the board and forces a player 2 loss.

6 Strategy when b=5

Finally, we consider the $(a, 5)$-game for $a \geq 5$. Although we give a winning strategy for this situation, there are more options for how player 1 selects a move leading to the final end game. The strategy given in this section was determined in an experimental matter. Since this is a 2-player combinatorial game with perfect information, when a is known, every possible board shading can be labeled as a next-player win or a next-player loss via computer search. Although we are limited to small values of a for a complete computer analysis, once the computer makes this labeling of all states for several specific values of a, an interested human can use the computer data to conjecture a subset of positions that are next-player-loss positions regardless of a and form a strategy that guarantees that regardless of player 2's move, player 1 can respond in such a way to achieve a position in the subset. As in the previous section, there are three clear phases of the game: opening moves, midgame, and endgame.

Theorem 4. *Player 1 has a winning strategy in the $(a, 5)$-game where $a \geq 5$.*

Before we formally prove Theorem 4, we outline player 1's strategy.

For the midgame, player 1 acts to leave the board in one of the following seven states after their turn:

1. Row 1 has an odd number of open cells and k shaded or eliminated cells, while rows 2, 3, and 4 have $k-1$ shaded or eliminated cells where $k \geq 2$.

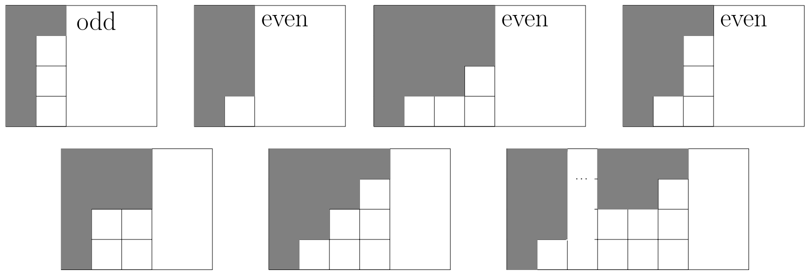

Figure 5: The seven mid-game next-player-loss states for the $(a, 5)$-game

2. Rows 1, 2, and 3 have an even number of open cells and k shaded or eliminated cells, while row 4 has $k-1$ shaded or eliminated cells where $k \geq 1$.

3. Rows 1 and 2 have an even number of open cells and k shaded or eliminated cells, while row 3 has $k-1$ shaded or eliminated cells, and row 4 has $k-3$ shaded/eliminated cells where $k \geq 3$.

4. Row 1 has an even number of open cells and k shaded or eliminated cells, rows 2 and 3 have $k-1$ shaded or eliminated cells, and row 4 has $k-2$ shaded/eliminated cells where $k \geq 2$.

5. Rows 1 and 2 have k shaded or eliminated cells and non-zero open cells, while rows 3 and 4 have $k-2$ shaded or eliminated cells where $k \geq 3$.

6. Row 1 has k shaded or eliminated cells and non-zero open cells, row 2 has $k-1$ such cells, row 3 has $k-2$ such cells, and row 4 has $k-3$ such cells where $k \geq 3$.

7. Row 1 has one more shaded cell than row 2 and non-zero open cells. Row 2 has at least two more shaded cells than row 3, and row 3 has one more shaded cell than row 4.

These seven states are illustrated in Figure 5 and we refer to them as the set $\widehat{S}$. A tedious computer-assisted analysis shows that if the current board is in one of the states from $\widehat{S}$ at the start of player 2's turn, then no matter where player 2 moves, player 1 has a response that returns to a state in $\widehat{S}$.

Further, it is possible to navigate from any of these states to having one of the following three states to prepare for an endgame:

- Both columns $a-1$ and $a-2$ have open cells. Column $a-1$ has one more open cell than column $a-2$. All other columns are completely shaded or eliminated.

- Column $a-1$ has at least two open cells; row 4 has the same number of cells. All other cells are are shaded or eliminated.

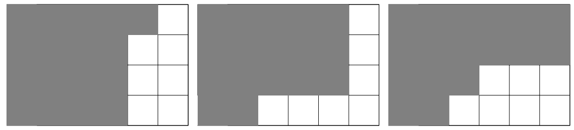

Figure 6: The endgame next-player-loss states for the $(a, 5)$-game

- Rows 1 and 2 have no open cells. Both rows 3 and 4 have open cells. Row 4 has one more open cell than row 3.

These three states are illustrated in Figure 6 and we refer to this set of states as set $\widehat{E}$.

We will ultimately prove Theorem 4 by a sequence of lemmas.

Lemma 1. By the end of their fourth move in an $(a, 5)$-game with $a \geq 5$, player 1 can leave the board in a state from $\widehat{S}$.

Proof. Player 1's first move is $(1, 1)$. Regardless of whether player 2 chooses $(1, 2)$ or $(2, 1)$, player 1's second move is $(2, 2)$ which results in a 2×2 shaded region.

If player 2 plays cell $(3, 1)$ or $(1, 3)$, then player 1 responds by playing the other option which leaves the board in state 6 from $\widehat{S}$ after player 1's third move.

On the other hand, suppose player 2 plays $(3, 2)$ or $(2, 3)$ as their second move. We proceed in cases.

If $a - 1$ is odd, then player 1 plays whichever of $(3, 2)$ and $(2, 3)$ was not chosen by player 2. This results in state 3 from $\widehat{S}$.

If $a - 1$ is even, then player 1 plays $(3, 3)$ which results in a 3×3 shaded region. Player 2 has 6 choices for what they can play in response: $(4, 1)$, $(4, 2)$, $(4, 3)$, $(1, 4)$, $(2, 4)$, or $(3, 4)$. If $a = 5$, the board is already in a state from $\widehat{E}$ and player 1 responds in a tit-for-tat way described in Lemma 3. Otherwise, player 1 responds as described below.

If player 2 plays $(4, 1)$ or $(2, 4)$, player 1 should play the other of these two cells, which results in state 4 from $\widehat{S}$ on player 1's fourth move.

If player 2 plays $(4, 2)$ or $(1, 4)$, player 1 should play the other of these two cells, which results in state 3 from $\widehat{S}$ on player 1's fourth move.

If player 2 plays $(4, 3)$ or $(3, 4)$, player 1 should play the other of these two cells, which results in state 2 from $\widehat{S}$ on player 1's fourth move. □

Next, we proceed to the midgame, by showing player 1 has a strategy to return the game to a state from $\widehat{S}$ no matter how player 2 acts. The proof we give documents the result of computer search in a manner that is, admittedly tedious, but able to be verified by an interested reader.

Lemma 2. If the board is in a state from $\widehat{S}$ at the end of player 1's turn, regardless of player 2's next move, player 1 can return the game to a state from $\widehat{S}$.

Proof. We proceed in cases by considering each of the seven states in $\widehat{S}$, each of player 2's options, and an appropriate response from player 1.

In state 1, if player 2 plays in row 1, player 1 responds with the cell below to turn the board to state 5. If player 2 plays in row 2, player 1 responds in the cell to the right to to turn the board to state 5. If player 2 plays in row 3, player 1 plays in row 1 to turn the board to state 4. If player 2 plays in row 4, player 1 plays in row 3 to turn the board to state 2.

In state 2, if player 2 plays in row 1, player 1 plays in row 4 and vice versa to get to state 1. If player 2 plays in row 2, player 1 plays in row 1 to get to state 6. If player 2 plays in row 3, player 1 responds by playing in row 2 to get to state 3.

In state 3, if player 2 plays the first open cell in row 1 player 1 plays the first open cell in row 4, and vice versa to get to state 6. If player 2 plays in row 2, player 1 plays the second open cell in row 4, and vice versa to get to state 5. If player 2 plays in row 3, player 2 responds by playing the second open cell in row 4 to get to state 2.

In state 4, if player 2 plays the first open cell in row 1, player 1 responds by playing the first open cell in row 2 and vice versa to get to state 6. If player 2 plays the first open cell in row 3, player 1 responds by playing the first open cell in row 4 and vice versa to get to state 2.

In state 5, if player 2 plays in row 1, player 1 responds by playing the first open cell in row 3 to get state 6. If player 2 plays in row 2, player 1 responds by playing one cell in row 4 to return to state 5. If player 2 plays the second open cell in row 3 and there are an odd number of open cells in row 1, player 1 responds by playing in row 2 to get state 3, while if there are an even number of open cells in row 1, player 1 responds by playing in row 4 to get to state 2. If player 2 plays in the first open cell in row 3, player 1 responds by playing in row 1 to get to state 6. Finally, if player 2 plays in row 4, player 1 responds by playing in row 2 to get back to state 5.

In state 6, if player 2 plays in row 1, player 1 responds by playing the first open cell in row 2 to get to state 7. If player 2 plays in row 2, player 1 responds in row 4 to get state 5. If player 2 plays row 3 and there are an even number of open cells in row 1, player 1 plays in row 2 to get state 3. If player 2 plays row 3 and there are an odd numbers of open cells in row 1, player 1 plays the second open cell in row 4 to get state 1. If player 2 plays in row 4, player 1 plays in row 2 to get state 5.

In state 7, if player 2 plays the first open cell in row 1, player 2 plays the first open cell in row 2 and vice versa, which returns to state 7. If player 2 plays the final available open cell in row 3 and there are an odd number of open cells in row 1, player 1 responds by playing the cell below player 2's move, resulting in state 1. On the other hand, if player 2 plays the final available open cell in row 3 and there are an even number of open cells in row 1, player 1 responds by playing the cell diagonally below player 2's move, resulting in state 4. If player 2 plays anywhere else in row 3, player 1 responds by playing in row 4 to get to state 6 or 7. If player 2 plays one cell in row 4, player 1 responds by playing one cell in row 3 to get to state 6 or 7. □

Lemma 3. If the board is in a state from $\widehat{S}$ with non-zero open cells in row 1, no matter what move player 2 makes, player 1 has a response to turn the board to a state in $\widehat{E}$.

Proof. First notice that if we consider a move by player 2 followed by the prescribed response of player 1 in Lemma 2, at most two cells from row 1 are shaded or eliminated in that pair of turns. In fact, the only time two cells from row 1 are shaded or eliminated in the same pair of turns is starting from state 2. So, it is sufficient to follow the response prescribed in Lemma 2 until there is one open cell in row 1 for states 1, 5, 6, and 7, or until there are two open cells in row 1 for states 2, 3, and 4, and then consider how one may adapt strategies to obtain a state in $\widehat{E}$ at that point of game play.

From state 1, if there is only one open cell in row 1, this is already a state in $\widehat{E}$, so respond as in Lemma 3.

From state 2, if there are two open cells in row 1, and player 2 plays in row 1 or row 4, player 1 responds as in Lemma 2 to get to state 1 with one open cell in row 1, which is a state in $\widehat{E}$. If player 2 plays in row 2, player 1 plays the last square in row 2 to get to a state in $\widehat{E}$. If player 2 plays in row 3, player 1 responds in row 1 to get to a state in $\widehat{E}$.

From state 3, if there are two open cells in row 1, and player 2 plays in row 1, row 2, or row 4, player 1 responds as in Lemma 2 to get to state 5 or state 6 with one open cell in row 1, both of which we consider below. If player 2 plays in row 3, player 1 responds as in Lemma 2 to get to state 2, still with two open cells in row 1, which was considered above.

From state 4, if there are two open cells in row 1, following game play as in Lemma 2 leads to state 6 with one open cell in row 1, considered below, or state 2 with two open cells in row 1, considered above.

From state 5, if there is one open cell in row 1, and player 2 takes the last cell in row 1, then player 1 takes penultimate cell in row 3 and vice versa to get to a state in $\widehat{E}$. If player 2 takes the last cell in row 2, then player 1 takes the first square in row 3 and vice versa to get to a state in $\widehat{E}$. If player 2 moves in row 4, player 1 takes the last square in row 1 to get to a state in $\widehat{E}$.

From state 6 when there is one open cell in row 1, if player 2 moves in row 1 or row 2, player 1 takes the final cell of row 2 to get to a state in $\widehat{E}$. If player 2 moves in row 3 or row 4, player 1 takes cell $(a-3,4)$ to get to a state in $\widehat{E}$.

From state 7 when there is one open cell in row 1, if player 2 moves in row 1 or row 2, player 1 takes the final cell of row 2 to get to a state in $\widehat{E}$. If player 2 takes the cell $(a-3,3)$, then player 1 takes cell $(a-3,4)$ to get to a state in $\widehat{E}$. Otherwise if player 2 takes a different cell $(c,3)$ with $c < a-3$, then player 1 responds by taking cell $(c-1,4)$, and if player 2 plays in row 4, player 1 takes one cell in row 3, returning to state 7. □

Proof of Theorem 4. By Lemma 1, Player 1 has a sequence of moves leading to a state in $\widehat{S}$.

By Lemma 2, if the board is in a state in $\widehat{S}$ at the end of Player 1's turn, no matter what move player 2 makes, player 1 has a response to return the board to a state in $\widehat{S}$.

By Lemma 3, if the board is in a state in $\widehat{S}$ with at most two open cells remaining in row 1, no matter what move player 2 makes, player 1 has a response to turn the board to a state in $\widehat{E}$.

Once the board is in a state from $\widehat{E}$, player 1 has a tit-for-tat response to any move player 2. In particular, there are either two rows (or two columns, or one row and one column) with open cells. Without loss of generality consider the case with two rows of open cells. When player 2 takes cell $(a-1,3)$ or cell $(a-2,4)$, player 1 takes cell $(a-1,4)$ and wins the game. Until then, when player 2 takes cells from one row, player 1 takes the same number of cells from the other row. □

While we provided a strategy for Player 1 to win, we admit that the proof of this strategy is a long list of case work. This strategy was determined by having a computer search through all $\binom{(a-1)+4}{4}$ possible shadings of an $(a-1)\times 4$ board and recursively label each as a winning or losing position for player 1, and then combing through winning states by hand to describe patterns within the of winning states that could be used as $\widehat{S}$. This is certainly not the only winning strategy for player 1 or the only set that could be used for $\widehat{S}$ in a similar strategy. This strategy is also notably more complex than the arguments for $b<5$, and while we conjecture a strategy for a player 1 win exists for larger b, the current methodology becomes increasingly cumbersome.

The set of states $\widehat{S}$ contains states where moves are available in any of the four rows. These are convenient descriptions of families of board shadings, where the lattice path dividing the shaded and eliminated cells from the open cells is translated horizontally, depending on k. However, these kinds of states are only convenient descriptions of families of states for player 1 to aim for once cells have been shaded or eliminated in all but the final row, so that moves are available in all four rows, rather than restricted to the top 2 or 3 rows of the board. This means that to come up with a similar strategy for larger b, the starting strategy of the game will take longer, to give enough moves for the top $b-2$ rows of the board to have non-zero shading.

7 Strategy for the misère game

In most of this paper, we considered a permutation game where the first player to complete either an I_a or a J_b pattern loses. In this section we consider the same rules, except that the first player to complete either an I_a pattern or a J_b pattern *wins*.

We describe the strategy in terms of the same boards as before. In both games, the first player who plays off of the $(b-1)\times(a-1)$ board ends the game. If this results in a loss, the goal is to play on the board as long as possible. The player who claims the lower right corner, indexed as $(a-1, b-1)$ is therefore the winner, since the entire board is eliminated, and their opponent must play off the board.

However, if playing off the board results in a win, the penultimate move will be in column $a-1$ or row $b-1$, giving the player an opportunity to finish the appropriate monotone subsequence. Since playing in column $a-1$ or row $b-1$ gives the opponent a win, players seek to *not* play in this row and column. In other words, players wanting to win, try to restrict themselves to the $(b-2) \times (a-2)$ subboard. A player who claims cell $(a-2, b-2)$ then forces their opponent to play in the last row or last column, which allows the player who claimed $(a-2, b-2)$ to win. In other words, the misère game requires the same strategy as the original game, but on a board with one fewer row and one fewer column. This means that player 1 has a misère winning strategy when $4 \leq b \leq 6$. The parity of a determines the winner when $b = 3$. And player 2 is guaranteed a misère win when $b = 2$ merely by completing a decreasing J_2 pattern on their first turn.

8 Future Directions

In this paper we considered a particular two-player Erdős-Szekeres game. Unlike the game of Harary, Sagan, and West, where players chose fixed integers on each turn and where determining a strategy quickly became computationally untractable beyond a specific finite bound on $a + b$, in this game, players add any new nth digit to a permutation while preserving the pattern formed by the first $n-1$ digits. Playing this pattern-informed game allows for more generalized strategies where a is arbitrarily large. Although the problem gets complex to analyze for sufficiently large b, we conjecture that a first player winning strategy exists for $a \geq b \geq 3$. We did, however, find a first player winning strategy for $a \geq b$ when $3 \leq b \leq 5$.

In the remainder of this section we consider several directions for possible future work.

8.1 Permutations and boards

For sufficiently large b it is convenient to phrase the strategies of this paper in terms of board shading rather than in terms of the actual underlying permutations in the original problem statement. However, at any point in the game, we may ask "how many permutations would have produced this particular board shading?" Clearly when $b = 2$ and $b = 3$, the permutations are unique. But for larger boards, multiple permutations would result in the same shaded board.

b	minimum moves	maximum moves	actual moves using strategy
2	2	a	a
3	3	$2a-1$	$2a-2$
4	4	$3a-2$	$2a-3$ or $2a-1$
5	5	$4a-3$	between $2a-C$ and $4a-6$

Table 1: Range of moves used in playing an (a, b)-permutation game.

8.2 Number of moves

Table 1 shows the maximum and minimum number of moves possible in playing the Erdős-Szekeres game for $a \geq b$ and $2 \leq b \leq 5$. The minimum is given by two players forming a decreasing permutation of length b. The maximum is given by the $(a-1)(b-1)+1$ bound given by Theorem 1. However the number of moves actually used by the strategies in this paper are in general somewhere between the two. When $b = 2$, the two players create an I_a permutation. When $b = 3$, player 1 makes a strategic move to take cell $(a-1, 2)$ which results in a final permutation one shorter than the maximum length in Theorem 1. When $b = 4$, in general when player 2 makes a move that shades/eliminates 2 cells, player 1 responds by shading/eliminating 1 cell and vice versa. The difference between $2a-1$ or $2a-3$ total moves is decided by how player 2 handles the endgame. When $b = 5$, although player 1 has a clear response to any move made by player 2, player 2 has far more options for what to do that have an impact on how quickly the game passes. At its slowest, players take turns moving/eliminating one cell at a time, and any turn that eliminates multiple cells come from the starting moves or the end game. At its fastest, player 2's turn and player 1's response shade or eliminate 4 cells, resulting in a permutation half as long as the maximum, minus a small finite number C of extra cells eliminated by choices in the start and end game.

A number of interesting follow up questions remain. For $b = 5$, is there a more efficient winning strategy than the one presented here? For larger b, what strategies exist, and how do they compare, proportionally to the maximum length game guaranteed by Theorem 1?

8.3 Number of winning positions

Considering the board shading interpretation of the game results in additional interesting questions. We can shade a legal region of an $(a-1) \times (b-1)$ board in $\binom{(a-1)+(b-1)}{a-1}$ ways, by choosing the lattice path that divides the shaded and eliminated cells from the open cells. Table 2 gives data for various choices of a and b. The first number is the number of shaded regions that result in a next-player loss, while the number in parentheses is this number divided by $\binom{(a-1)+(b-1)}{a-1}$; in other words, the second number is the percent of board shadings that result in a next-player loss (and thus are candidates to help form a set

$b\backslash a$	2	3	4	5	6	7	8	9
2	1 (0.5)	2 (0.67)	2 (0.5)	3 (0.6)	3 (0.5)	4 (0.57)	4 (0.5)	5 (0.56)
3		2 (0.33)	3 (0.3)	4 (0.27)	5 (0.24)	6 (0.21)	7 (0.19)	8 (0.18)
4			6 (0.3)	10 (0.29)	15 (0.27)	21 (0.25)	28 (0.23)	36 (0.22)
5				18 (0.26)	31 (0.25)	46 (0.22)	67 (0.2)	91 (0.18)
6					58 (0.23)	103 (0.22)	164 (0.21)	253 (0.2)

Table 2: Total number (and percentage) of shadings of an $(a-1)\times(b-1)$ board that result in a next-player loss

analogous to $\widehat{S}$ in the strategy for $a \geq b = 5$. Of note, when b is constant but a increases, these percentages overall decrease. However, they show differences in parity, which is reflected in how the states of $\widehat{S}$ relied on how many open cells were in row 1. Do these values converge on a non-zero value as a increases? If so, what is it?

References

[1] P. Erdős and G. Szekeres, (1935). A combinatorial problem in geometry. *Compos. Math.* **2**: 463–470.

[2] F. Harary, B. Sagan, and D. West, Computer-aided analysis of monotonic sequence games, *Atti Accad. Perolitana Pericolanti Cl. Sci. Fis. Mat. Natur.* **61** (1983), 67–78.

[3] OEIS Foundation Inc. (2024), The On-Line Encyclopedia of Integer Sequences. published electronically at `oeis.org`.

[4] L. Pudwell, Patterns in Permutations: the hidden and surprising structures that emerge from ordered lists, Math Encounters presentation, National Museum of Mathematics (virtual); June 2, 2021. published electronically at `https://www.youtube.com/watch?v=B7DPo9YQTgw`.

[5] L. Pudwell, Catalan Numbers and Permutations, Mathematics Magazine 97.3 (2024), 279–283.

[6] S. Rabinowitz (proposer), R. Stanley (solver). Advanced Problem 5641. Amer. Math. Monthly **76.10** (1969): 1153

[7] A. Seidenberg, A Simple Proof of a Theorem of Erdős and Szekeres. *J. Lond. Math. Soc.* **34.3** (1959): 352.

[8] C. Schensted, Longest increasing and decreasing subsequences. *Canad. J. Math.* **13** (1961): 179—191.