



Mathellaneous

Norman Do

How to Win at Tic-Tac-Toe

1 More Than Child's Play

Tic-Tac-Toe as a Tree

Undoubtedly one of the most popular pencil and paper games in the world is tic-tac-toe, also commonly known as noughts and crosses. The game has a longstanding history in many cultures across the globe. It has been posited that it may even have been played under the name of “terni lapilli” in Ancient Rome, where examples of the tic-tac-toe grid have been found etched in stone throughout the empire. No doubt almost every reader will have played a game of tic-tac-toe, perhaps as a child subjected to a less than exhilarating class at school. And no doubt just as many will be well aware of the fact that if both players adopt their optimal strategy, then neither player can force a win.

Most tic-tac-toe enthusiasts will come to the conclusion that the game ends in a theoretical draw through intuition and experience — but how does one prove such a fact? A naive brute force approach is to treat the game of tic-tac-toe as a tree. More precisely, consider a graph with a root vertex corresponding to the empty tic-tac-toe board, with adjacent edges connecting it to all of the possible states of the game after one move, which are in turn connected to vertices corresponding to all of the possible states of the game after two moves, and so on. Since every game must end after at most nine moves, the graph constructed in this way does not contain any cycles and is therefore a tree. The edges of the tree correspond to all of the possible moves while the vertices correspond to all of the possible states of the game. Furthermore, the vertices of degree one correspond to completed games, so they can be labelled with an O, X or D depending on whether they correspond to a win for O, a win for X, or a draw, respectively. A large number of two player games of no chance can be considered as a tree in this way. These games are the realm of combinatorial game theory, an area of mathematics with an extensive, though growing, theory.

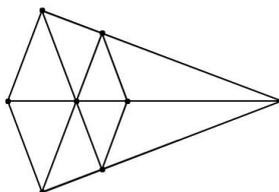
Now determining whether X has a winning strategy or O has a winning strategy or whether the game is a theoretical draw can be found by a method known as tree-pruning. In this process, branches are repeatedly deleted from the tree until it has been pruned all the way back to the root. The first pruning of the tic-tac-toe tree removes all moves in which a player can win but does not. The second pruning removes those branches in which a player could have blocked those wins, but did not. To comprehend just how brutish this brute force approach actually is, consider the fact that the tic-tac-toe tree contains 255,168 vertices of degree one, corresponding to all of the different possible tic-tac-toe games that can be played. Of course, this number can be significantly reduced by taking into account the symmetry of the board. For example, rather than considering all nine edges from the root vertex corresponding to the nine squares in which the first move can be, we need only consider the three edges corresponding to a move in the corner, edge or centre. Despite this reduction, it seems unlikely that anyone would want to perform the full analysis by hand. And this is just for tic-tac-toe — imagine the size of a tree for a game as complicated as Nim, or chess, for that matter!

Show that there exist 255,168 possible tic-tac-toe games, where the symmetry of the board has not been taken into account.

Games to Beat your Friends With

Since the optimal strategy is so well-known, a tic-tac-toe duel between intelligent players will almost certainly involve them playing out an unavoidable draw. Thus, tic-tac-toe is a solved game and, like a solved maths problem, we can sweep it under the proverbial rug and turn our attention to other more interesting pastimes. The competitive reader may like to challenge a formidable opponent at the following three pencil and paper games.

Jam: The diagram below shows a map of towns and roads, represented by points and line segments, respectively. Two players take turns to select a road and a road may not be chosen if it has already been used. The first player to take all of the roads passing through a town wins. If all roads have been selected without one of the players winning, then the game is declared a draw.



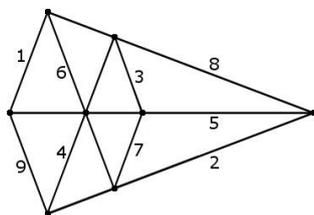
Magic Fifteen: Two players take turns to select an integer from 1 to 9. An integer may not be chosen if it has already been used and a player wins once they have chosen three distinct numbers which add to 15. If all of the nine numbers have been selected without one of the players winning, then the game is declared a draw.

Count Foxy Words: Two players take turns to select one of the following words.

COUNT FOX WORDS AND STAY AWAKE USING LIVELY WIT

A word may not be chosen if it has already been used and a player wins once they have chosen all of the words which contain a given letter. If all of the words have been selected without one of the players winning, then the game is declared a draw.

The astute reader may have noticed a remarkable similarity between the three aforementioned games. They all require two players to alternately select an object from a set of size nine with the aim being to obtain one of eight possible three-element sets. Sound familiar? Of course it does, since all of these games are simply tic-tac-toe in disguise — the game's the same by any name! These are presented in *Winning Ways for Your Mathematical Plays* [1] where the authors bet that you can fool your friends, at least for a short while, by playing these three variations on the tic-tac-toe theme. The diagram below reveals the isomorphisms which demonstrate the equivalence.



8	1	6
3	5	7
4	9	2

COUNT	FOXY	WORDS
AND	STAY	AWAKE
USING	LIVELY	WIT

2 How to Get N in a Row

Stealing Strategies and Pairing Strategies

For those who find the 3×3 board too restrictive, let us turn our attention to a game played in the wide open spaces of the infinite square grid. This tic-tac-toe variant is known as N -in-a-row and involves two players alternately marking cells of the board. The aim of the game, as the name suggests, is to mark N cells in a row, either horizontally, vertically or diagonally. The primary objective in analyzing N -in-a-row is to determine for each value of N whether the first player has a winning strategy, the second player has a winning strategy, or both players can prevent the other from winning, thus rendering the game a draw. It seems difficult to imagine that the second player can possibly have the advantage after beginning the game one move behind. For how can having an extra occupied square on the board possibly hurt the first player's chances of winning? Indeed, this intuition is correct and we can prove that a winning strategy for the second player does not exist by a clever argument known as strategy stealing. We will later see that strategy stealing applies to many tic-tac-toe variations.

Theorem: The second player does not have a winning strategy for N -in-a-row.

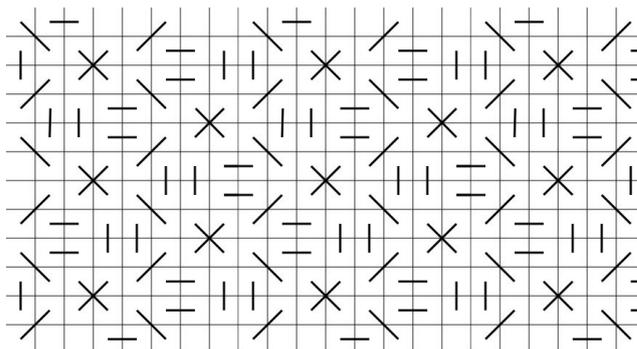
Proof. Let us suppose that the second player has a winning strategy. But now the first player can win by making his or her first move at random and thereafter adopting the second player's winning strategy. If this calls for the first player to play in an already occupied square, he or she just makes another random move. Since having an extra square on the board cannot possibly hurt the first player, this gives the contradiction that both players can force a win. So we must conclude that the second player cannot have a winning strategy, as desired. \square

We have now shown that N -in-a-row is either a first player win or a theoretical draw — but which of these cases arises for which values of N ? It is a simple matter to prove that the game of N -in-a-row is a first player win for small values of N such as 1, 2, 3 or 4. Could it be that the first player always has a winning strategy, no matter how large N is? It surely seems unlikely that a player could conceivably achieve, say, one million squares in a row, but then again, the infinite square grid is large indeed. As it turns out, a simple strategy discovered by Hales and Jewett [5] can be used to prove that 9-in-a-row is a theoretical draw, and hence, so is N -in-a-row for all larger values of N .

Theorem: The second player can force a draw in 9-in-a-row.

Proof. In the diagram below, the squares of the infinite square grid are paired by line segments joining their centres and the pattern repeats periodically over the entire board.

Note that any horizontal, vertical or diagonal row of nine squares must contain both squares of a pair — this suggests the following strategy for the second player. Wherever the first player moves, play in the corresponding square of the pair. In this way, they can never secure both squares of a pair, and hence, can never occupy nine squares in a row. Thus, the second player can force a draw in 9-in-a-row by using this pairing strategy.



□

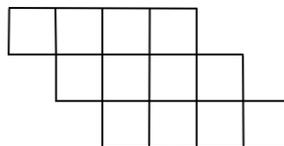
Prove that the first player has a winning strategy for N -in-a-row when $N = 1, 2, 3$ or 4 .

8-in-a-Row is a Draw

In the April 1979 issue of *The American Mathematical Monthly*, Guy and Selfridge proposed the problem of proving that 9-in-a-row is a draw, presumably expecting the pairing strategy shown above. At that time, no significant results had been obtained for N -in-a-row for $5 \leq N \leq 8$. Thus, they must have been both pleased and surprised to receive the following proof from a T. G. L. Zettlers of Amsterdam that 8-in-a-row is also a draw [4].

Theorem: The second player can force a draw in 8-in-a-row.

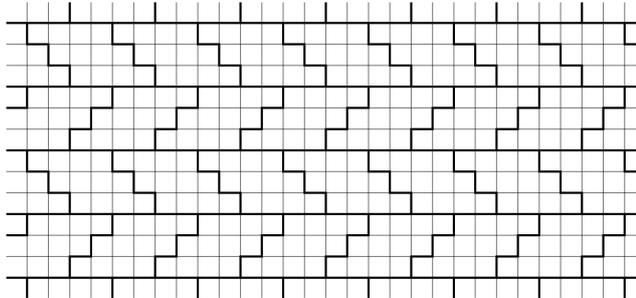
Proof. First, consider the game of Zettlers played on the following board consisting of twelve squares. As in tic-tac-toe, two players take turns to mark squares. However, the aim of the game is to occupy all of the squares from one of the three rows of four squares, one of the four diagonals of three squares, or one of the two columns of two squares.



The second player can always force a draw in this game with the general strategy being to occupy squares in each of the two columns containing two squares and then playing defensively. The details are left to the interested reader.

Returning to the game of 8-in-a-row, let us tessellate the plane with copies of the Zettlers boards as shown in the diagram below. The key aspect of this tiling is that any eight squares in a row must occupy a winning path on one of the Zettlers boards. Thus, the second player

can adopt the drawing strategy on each of these smaller boards and prevent the first player from achieving eight squares in a row.



□

Prove that the second player can always force a draw in the game of Zetters described above.

Go Moku Solved

From a player's perspective, the most interesting version of N -in-a-row is the $N = 5$ case. This game has been played since the 7th century BC in Japan, where the game is known as Go Moku and was originally played on the 19×19 grid used for the board game Go. Since the turn of the twentieth century, it became apparent that there was a distinct advantage for the first player in Go Moku. Thus, extra rules and handicaps for the first player were introduced to remove this discrepancy.

- The board was reduced from the traditional 19×19 to a smaller 15×15 grid.
- The first player was forbidden to make certain configurations, such as the "double three attack".
- An overline of six or more in a row was not counted as a win for either player.

Even with all of these restrictions and handicaps, many of the world's leading Go Moku experts continued to believe that the game was a theoretical win for the first player, to the extent that this became a "folklore theorem". However, it was not until 1993 when Allis, van den Herik and Huntjens used a new technique known as threat space search along with proof number search and hundreds of hours worth of CPU time to show that Go Moku on a 15×15 board without restrictions is a win for the first player. This solution to Go Moku then yields the weaker fact that 5-in-a-row is a first player win. So there remain only two unsolved cases of N -in-a-row.

In the game 6-in-a-row, does the first player have a winning strategy or can the second player force a draw? How about in the game 7-in-a-row?

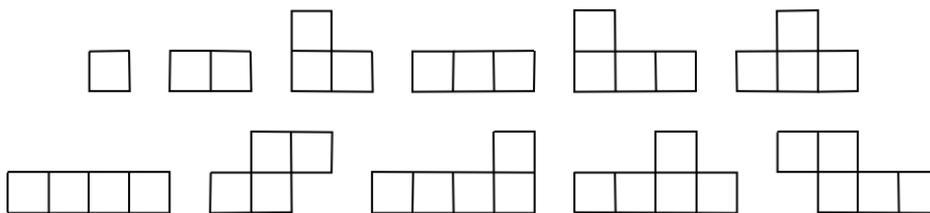
3 Games with Animals

Animal Tic-Tac-Toe

Almost thirty years ago, the famous graph theorist Frank Harary introduced another variation of tic-tac-toe to be played on an infinite square grid. Players take turns to mark cells of the board with the aim of creating a predetermined animal, or polyomino, as it is often

referred to in the modern literature. We will allow any translations and rotations of the animal and, if asymmetrical, also allow reflections. Once again, a strategy stealing argument shows that the second player cannot possibly have a winning strategy so to even up the game, let us decree that the first player wins if they can create the animal, and the second player wins if they can prevent the first player from doing so. This game was popularized by Martin Gardner in his column entitled *Harary's generalized ticktacktoe* [2], where the following result appears.

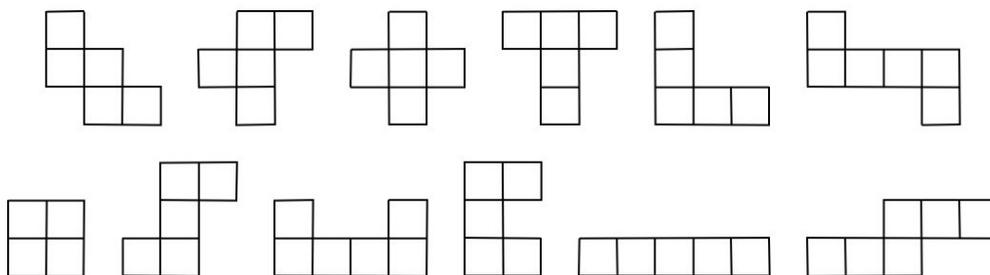
Prove that the first player can win animal tic-tac-toe when the chosen animal is one of the twelve pictured below.



Winners and Losers

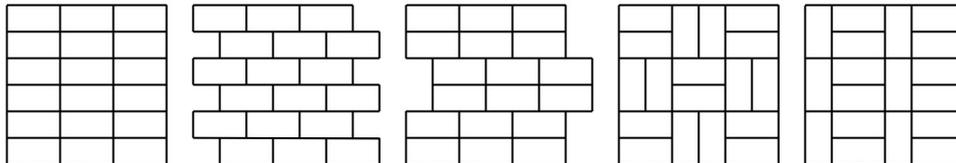
The strategy stealing proof used earlier for N -in-a-row can be translated verbatim to show that the second player cannot have a winning strategy for animal tic-tac-toe either. This observation prompted Harary to divide all animals into winners and losers depending on whether the first player had a winning strategy or not. For example, the problem above gives a list of twelve winners, including all of those animals of size 1, 2 or 3. The list also includes all of the animals of size 4, apart from the 2×2 square, which is affectionately known as “fatty”. As we shall see shortly, fatty is a loser which implies that every larger animal which contains fatty is also a loser. Let us call fatty a “basic loser”, since it is a loser which does not contain a smaller one.

Theorem: The twelve animals in the figure below are all basic losers.



Proof. Each of these twelve animals can be shown to be losers by providing a pairing strategy with which the second player can prevent the first from creating the animal. These pairing strategies are indicated by the five tilings of the plane by dominos below. For example, fatty is a loser because whenever the first player moves in a square of the second “brickwork” tiling, the second can retort with a play in the adjacent square belonging to the same

domino. Since every position of a fatty covers one whole domino, it is impossible for the first player to win. Furthermore, every animal properly contained within fatty appears in the list of winners above, so fatty is a basic loser. The remaining animals can be shown to be basic losers analogously.



□

Of the twelve animals of size five, one of them is a loser since it contains fatty, three of them appear in the list of winners, and the remaining eight appear in the list of basic losers. Moving up to the 35 animals of size six, we find that all but four of them contain basic losers of lower order. Of these four, three appear in the list of basic losers and the remaining one we will discuss a little later on. Of the 108 animals of size seven, every single one contains a basic loser and hence, are losers themselves. It follows that every animal of size greater than seven is also a loser since they all contain an animal of size seven.

Snaky

And what about the one animal of size six which has been left unaccounted for? Let us now meet this exotic animal which, in the literature, goes by the name of Snaky.



It turns out that it is unknown whether the elusive animal known as Snaky is a winner or a loser. However, the experts on the problem believe it to be a winner. Furthermore, Harborth and Seemann [6] have shown that there is no pairing strategy available to the second player when trying to stop the first player from creating Snaky. So it seems that considerable insight and further consideration will be required before the following problem is finally resolved.

Is Snaky a winner or a loser?

4 Hypercube Tic-Tac-Toe

Preliminary results

An interesting generalization of tic-tac-toe emerges when we consider the game with boards of arbitrary size and number of dimensions. Hypercube tic-tac-toe is played on a k -dimensional hypercube of side length n divided into n^k unit hypercubes with players taking turns to mark

one of the cells. Of course, the aim of the game is to mark a winning path of n cells whose centres are collinear. We can see that the original game of tic-tac-toe is simply the 3^2 version of hypercube tic-tac-toe. Almost needless to say, the strategy stealing arguments used earlier are equally successful in proving that the second player cannot have a winning strategy. The main problem then is to determine for each n and k whether n^k hypercube tic-tac-toe is a first player win or a theoretical draw. The following are three elementary results about hypercube tic-tac-toe which are indicative of progress in the area.

Hypercube tic-tac-toe on the 4^3 board is a win for the first player.

This version of hypercube tic-tac-toe has been marketed under several names, the most popular being Qubic. In 1980, Oren Patashnik [7] used over 1500 hours of CPU time to prove that Qubic is a win for the first player. His program was a refinement of the tree-pruning algorithm described earlier — for example, rather than the $64 \times 63 \times 62 = 249,984$ possible positions of the board after three moves, Patashnik's reduced tree only considered seven. Some of this reduction comes from the non-trivial observation that there are 192 automorphisms of the 4^3 tic-tac-toe board. In other words, there are 192 permutations of the 64 cells of the board which preserve the 76 winning paths on the board. It would be interesting to determine the number of automorphisms for other hypercube tic-tac-toe boards.

Hypercube tic-tac-toe on the 5^2 board is a theoretical draw.

This particular version of hypercube tic-tac-toe can be proved to be a theoretical draw via the pairing strategy indicated by the diagram below.

V	I	A	A	F
J	B	H	U	B
C	I		G	C
D	U	H	D	F
J	E	E	G	V

Hypercube tic-tac-toe on the 3^3 board is a win for the first player.

It is a simple matter to give a constructive proof that the first player has a winning strategy on the 3^3 board by first moving in the central cube. In fact, a win can be guaranteed in as few as four moves. A more surprising fact is that this particular game of hypercube tic-tac-toe cannot end in a draw. In other words, it is impossible to partition the 27 unit cubes into two parts without one of the parts containing one of the winning paths, of which there are 49. Since a draw is impossible and strategy stealing rules out a second player winning strategy, this observation yields an existence proof that the first player has a winning strategy. This hints at the following result from Ramsey Theory, first proved by Hales and Jewett in [5].

Theorem: For any positive integer values of n and c , there exists an integer k such that whenever the squares of the n^k tic-tac-toe board are coloured in c colours, there exists a monochromatic winning path.

The Number of Winning Paths

As we have seen, pairing strategies provide a successful method for proving that the second player can force a draw. However, note that no pairing strategy exists to show that the 3^2 version of hypercube tic-tac-toe is a draw. This is since a necessary condition for a pairing

strategy to exist is that the number of cells is at least twice the number of winning paths. And the number of winning paths on the 3^2 board is eight, while the number of cells is nine. Thus, it is both useful as well as interesting to know how many winning paths there are on the n^k board.

Theorem: The number of winning paths on the n^k hypercube is

$$\frac{(n+2)^k - n^k}{2}.$$

Proof. Let us represent each cell of the n^k hypercube by a vector $\mathbf{x} = (x_1, x_2, \dots, x_k)$, where $1 \leq x_i \leq n$ for each i . A winning path consists of an ordered sequence of n such vectors $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ where each component either runs from 1 up to n , from n down to 1, or remains constant at one of the values $1, 2, \dots, n$. Since there are $n+2$ possibilities for each of the k components, the number of winning paths is no more than $(n+2)^k$. However, note that n^k of these paths are constant paths, so our number is now reduced to $(n+2)^k - n^k$. It remains only to observe that we have overcounted by a factor of two, since each winning path can be traversed in one of two directions.

The fact that the number of winning paths on the n^k hypercube is $\frac{(n+2)^k - n^k}{2}$ suggests the following “geometric proof”. Embed the n^k hypercube within an $(n+2)^k$ hypercube and notice that each winning path can be extended in each direction to give a pair of cells lying in the outer hypercube but not the inner one. Furthermore, every cell lying in the outer shell corresponds to a unique winning path. Therefore, the number of winning paths is simply half the number of cells in the outer shell — namely, $\frac{(n+2)^k - n^k}{2}$. \square

As discussed earlier, a necessary condition for a pairing strategy to exist is that the number of cells is at least twice the number of winning paths. This occurs when n and k satisfy

$$\begin{aligned} n^k &\geq (n+2)^k - n^k \\ 2n^k &\geq (n+2)^k \\ 2 &\geq \left(1 + \frac{2}{n}\right)^k \\ \sqrt[k]{2} &\geq 1 + \frac{2}{n} \\ n &\geq \frac{2}{\sqrt[k]{2} - 1}. \end{aligned}$$

Results and Conjectures

The first significant results in hypercube tic-tac-toe appeared in the 1963 paper *Regularity and positional games* by Hales and Jewett [5]. One of their main results was to use Hall’s Marriage Theorem to prove that the second player can force a draw by a pairing strategy when $n \geq 2M$ where M denotes the maximum number of winning paths which pass through a square. Erdős and Selfridge managed to improve on this result by showing that the second player can force a draw when $2^k > M + L$, where L is the number of winning paths, although their strategy cannot be described by a simple pairing. More recently, Jozsef Beck has extended the argument to show that the second player can force a draw if $k = O(n^2 / \log n)$.

The following table gives results and conjectures for hypercube tic-tac-toe for small values of n and k . An entry labelled “W” denotes a win for the first player while “D” denotes a

theoretical draw and “?” indicates that the result is merely conjectured but not actually proven.

	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$	$n = 6$	$n = 7$	$n = 8$	$n = 9$	$n = 10$
$k = 1$	W	D	D	D	D	D	D	D	D	D
$k = 2$	W	W	D	D	D	D	D	D	D	D
$k = 3$	W	W	W	W	D?	D?	D?	D	D	D
$k = 4$	W	W	W	W	W?	W?	W?	D?	D?	D?
$k = 5$	W	W	W	W?	W?	W?	W?	W?	W?	D?
$k = 6$	W	W	W	W?						

These results appear in Golomb and Hales’ article on hypercube tic-tac-toe [3] as well as the following two conjectures which seem intuitively obvious but have so far resisted proof.

- (1) If the n^k game is a draw, then the n^{k-1} game is a draw.
- (2) If the n^k game is a draw, then the $(n + 1)^k$ game is a draw.

Remember that a necessary condition for a pairing strategy to exist is that the number of cells is at least twice the number of winning paths. As we have shown, this is equivalent to $n \geq \left\lceil \frac{2}{\sqrt[k]{2}-1} \right\rceil$. Hales and Jewett conjectured that drawing strategies exist for the second player whenever this inequality holds. Golomb and Hales subsequently noted the interesting fact that there is a fairly accurate linear approximation to the expression $\frac{2}{\sqrt[k]{2}-1}$.

$$\begin{aligned} \frac{2}{\sqrt[k]{2}-1} &= 2 \frac{a^k - 1}{a - 1} = 2(1 + a + a^2 + \cdots + a^{k-1}) \\ &\approx 2 \int_0^k a^t dt = 2 \frac{a^k - 1}{\log_e a} \\ &= \frac{2k}{\log_e 2}, \end{aligned}$$

where $a = a_k = \sqrt[k]{2}$.

From this observation, it is tempting to make the conjecture that

$$\left\lceil \frac{2}{\sqrt[k]{2}-1} \right\rceil = \left\lfloor \frac{2k}{\log_e 2} \right\rfloor.$$

In fact, a computer program can easily verify the conjecture to be true for thousands, even millions of terms, before reaching the incredibly large value of $k = 6,847,196,937$, when the conjecture first fails. Furthermore, this is the only failure until we reach the second counterexample at $k = 27,637,329,632$. In fact, the theory of diophantine approximation can be used to prove that these anomalies can only occur when k is the denominator of a continued fraction convergent for $\frac{2}{\log_e 2}$. Let us now finish with a tic-tac-toe-inspired number theory problem for which, I believe, the answer is unknown.

For how many positive integer values of k does the equation

$$\left\lceil \frac{2}{\sqrt[k]{2}-1} \right\rceil = \left\lfloor \frac{2k}{\log_e 2} \right\rfloor$$

fail to be true?

References

- [1] E.R. Berlekamp, J.H. Conway, and R.K. Guy, *Winning Ways for your Mathematical Plays*, Volume 3, 2nd Edition (A. K. Peters Ltd Wellesley MA 2003).
- [2] M. Gardner, *Harary's generalized ticktacktoe*, in: *The Colossal Book of Mathematics* (W.W. Norton and Company New York 2002), 493–503.
- [3] S. W. Golomb and A. W. Hales, *Hypercube tic-tac-toe*, in: *More Games of No Chance* (MSRI Publications 2002), 167–182.
- [4] R.K. Guy, J.L. Selfridge, T.G.L. Zetters, *S10 — 8 (or more) in a row*, *The American Mathematical Monthly* **87** (1980), 575–576.
- [5] A. W. Hales and R. I. Jewett, *Regularity and positional games*, *Transactions of the American Mathematical Society* **106** (1963), 222–229.
- [6] H. Harborth and M. Seemann, *Snaky is a paving winner*, *Bull. Inst. Combin. Appl.* **19** (1997), 71–78.
- [7] O. Patashnik, *Qubic: $4 \times 4 \times 4$ tic-tac-toe*, *Mathematics Magazine* **53** (1980), 202–216.

Department of Mathematics and Statistics, The University of Melbourne, VIC 3010

E-mail: N.Do@ms.unimelb.edu.au