June 1988 Report No. STAN-CS-88-1208

Toetjes

by

Tomás Feder

Department of Computer Science

Stanford University

Stanford, California 94305

Toetjes

Tomás Feder*
Computer Science Department
Stanford University
Stanford, CA 94305

Abstract

A number is secretly chosen from the interval [0, 1], and n players try in turn to guess this number. When the secret number is revealed, the player with the closest guess wins. We describe an optimal strategy for a version of this game.

1. Introduction

The Toetjes Problem was posed at the Programming and Problem Solving Seminar [1] held at Stanford University in the Winter of 1957. It was suggested by Sape Mullender, who described it as follows:

> In Amsterdam, where I grew up, dessert is usually referred to as "toetje" (Dutch for "afters"). The problem of allocating a left-over toetje to one of the children in my family became the Toetjes Problem. The algorithm was the following: First my mother would choose a secret number between one and a hundred. Then the children, in turn, youngest to oldest, could try to guess the number. After the last guess my mother would tell whose guess was closest to her secret number and the winner would get the toetje.
>
> Now that I have a degree in mathematics, the problem still puzzles me: Given that the secret number is chosen randomly from the interval $[0, 1]$, what is the optimal strategy for choosing a number for the ith child in a family of n children? The ith child knows what the first $i - 1$ children chose, and knows that all the children choose optimally (i.e., choose to maximize their own chance without consideration for the chances of any other child in particular).

We will study Mullender's "continuous" version of the problem, in which the secret number is chosen randomly from the interval [0, 1]. Notice that the description we have given so far is in fact incomplete. A first ambiguity arises because the optimal move for a player may

*This research was primarily supported by a Bell Laboratories Scholarship, and in part by the National Science Foundation under grant CCR-86-10181, and by Office of Naval Research contract N00014-87-K-0502.

not be unique. If this is the case, the criterion that the player uses to choose among equally desirable moves can influence the strategy followed by earlier players. We will resolve this ambiguity by assuming that there is a given *tie-breaking rule* that all players must follow when choosing among several equally desirable moves.

The choice of the tie-breaking rule affects the play quite drastically. To illustrate this? let $x_1, \ldots, x_n$ be the (distinct) numbers chosen by the players in the interval [0, 1]. Suppose player n plays between two earlier players i and j such that no other player has already played between them. Then his payoff (i.e., his probability of winning) is the same no matter where in the interval (x_i, x_j) he plays: it equals $(x_j - x_i)/2$, half the interval size. However, his choice will affect the payoff (and hence the strategy) for players i and j.

A second difficulty occurs because, in this continuous version of the problem, each player has an infinite number of choices, so there may not be an optimal move. Instead, it can happen that for each possible choice there is another choice that gives a slightly better payoff. For example, in the two-player game, suppose the first player chooses $1/2$; then the second player will get a better payoff the closer he plays to the first player, so there is no "optimal" move, but instead a family of moves arbitrarily close to optimal. We will later see that there are situations where a player can obtain a better payoff the closer he plays to a certain number ξ, while playing exactly at ξ would give him a very low payoff. We handle this difficulty by defining an optimal *limiting play* which is such that all players can obtain payoffs arbitrarily close to optimal by playing sufficiently close to this limiting play.

In this paper we characterize the optimal strategy under two simple tie-breaking rules: the "closest to the first player" rule and the "rightmost" rule. In particular, we prove that under the first of these tie-breaking rules, the payoff for the first and last players is exactly half the payoff for each of the remaining players. Under the second tie-breaking rule, the last two players get a payoff equal to half the payoff of each of the other players.

The remainder of the paper is organized as follows. In section 2 we give a rigorous definition of the game, which resolves the two ambiguities mentioned abvove. In section 3 we develop a plausible optimal playing strategy by making some assumptions about the tie-breaking rule. -Then in section 4 we prove that the strategy of section 3 is in fact optimal under a simple tie-breaking rule, the "closest to the first player" rule. Section 5 extends the results of the previous section to the "rightmost" tie-breaking rule and explores the "random" tie-breaking rule.

2. Definitions

We now give a precise definition of the "optimal" playing strategy. Given the k distinct numbers $x_1, \ldots, x_k$ chosen by the first k players, we consider *limiting plays,* of the form $(x_{k+1}, \ldots, x_n, \prec)$, for the remaining $n-k$ players. Intuitively, a. limiting play of this form indicates that the last $n-k$ players play arbitrarily close to $x_{k+1}, \ldots, x_n$ and in such a way that the resulting ordering of the n players in the [0, 1] interval is given by the total ordering $\prec$. While the numbers played are all distinct. some of the $n-k$ limiting values can in fact be equal to each other or to the numbers chosen by the first k players. In the case

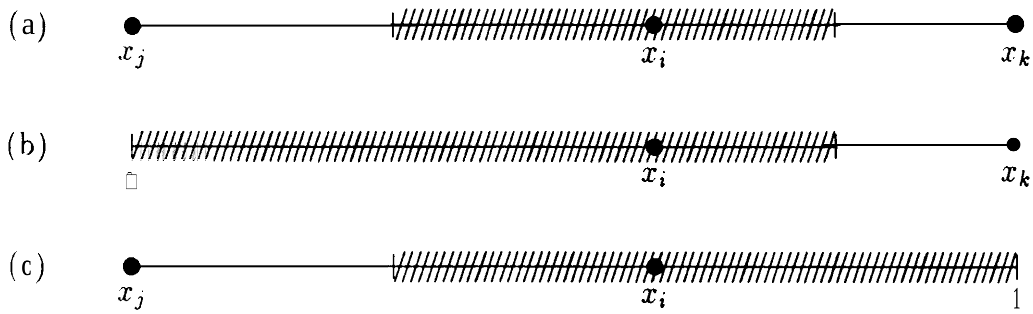

Figure 1: Payoff for player i.

where two or more values are equal, the left-to-right order of the n moves in the [0, 1] interval needs to be specified: only then can the payoff for each player corresponding to this limiting play be evaluated. This is the reason for including in the limiting play $(x_{k+1}, \dots, x_n, \prec)$ an ordering $\prec$ that extends the partial ordering $<$ on the numbers $x_1, \dots, x_n$ to a total ordering.

The winning probability (or payoff) $p(i, x_1, \dots, x_n, \prec)$ for player i under such a limiting play can now be easily computed [see Fig. 1]: (a) If the move x_i is between x_j and x_k in the $\prec$ ordering, then the payoff is $(x_i - x_j)/2 + (x_k - x_i)/2 = (x_k - x_j)/2$. (b) If x_i is the smallest move and is followed by x_k, then the payoff is $(x_i - 0) + (x_k - x_i)/2 = (x_i + x_k)/2$. (c) If x_i is the largest move and is preceded by x_j, then the payoff is $(x_i - x_j)/2 + (1 - x_i) = 1 - (x_i + x_j)/2$.

We use this payoff function to define the optimal limiting play $\mathrm{opt}(x_1, \dots, x_k) = (x_{k+1}, \dots, x_n, \prec)$ for the last $n - k$ players, given the k distinct numbers $x_1, \dots, x_k$ chosen by the first k players. The definition is by induction on the number $n - k$ of players left to play. The base case $k = n$ is simply $\mathrm{opt}(x_1, \dots, x_n) = (\prec)$, where $\prec$ is the $<$ ordering on the distinct values $x_1, \dots, x_n$. Now we assume inductively that the function $\mathrm{opt}(x_1, \dots, x_k)$ has been defined for a given k, and we define $\mathrm{opt}(x_1, \dots, x_{k-1})$.

So fix the numbers $x_1, \dots, x_{k-1}$ chosen by the first $k - 1$ players. We say that $(x_k, \dots, x_n, \prec)$ is *a limiting play available to player k* if for all $\epsilon > 0$ there are numbers $y_k, \dots, y_n$ with $|y_i - x_i| < \epsilon$ such that $\mathrm{opt}(x_1, \dots, x_{k-1}, y_k) = (y_{k+1}, \dots, y_n, \prec)$. In other words, the limiting play is available to player k if player k can play close to x_k so that the remaining players will then play close to $x_{k+1}, \dots, x_n$.

The optimal limiting payoff α for player k is now defined to be the least upper bound of the payoff $p(k, x_1, \dots, x_n, \prec)$ *over all limiting plays* $(x_k, \dots, x_n, \prec)$ available to player k. We can show that there exists a limiting play available to player k which achieves this maximum payoff α by considering a sequence of limiting plays available to player k that achieve payoffs arbitrarily close to α, and then using the Bolzano-Weierstrass theorem to obtain a limit point $(x_k, \ldots, x_n, \prec)$ for this sequence. We let S be the set of limiting plays $(x_k, \ldots, x_n, \prec)$ available to player k which give player k this maximum payoff α; thus S is the set of optimal limiting plays for player k. It is not difficult to show that S is in fact a closed set [1].

We can now apply any chosen *tie-breaking rule* to pick an optimal limiting play $(x_k, \ldots, x_n, \prec)$ from S. We could, for example, always pick a limiting play from S that has a largest x_k value (such a play exists because S is a closed set); this is in fact a particular case of the family of "rightmost" tie-breaking rules that we will consider in section 5. The chosen $(x_k, \ldots, x_n, \prec)$ is then the value of $\mathrm{opt}(x_1, \ldots, x_{k-1})$, completing the inductive definition.

Now, if all players play optimally and follow the given tie-breaking rule to choose arnong equally desirable moves, their numbers will be given in the limit by $\mathrm{opt}() = (x_1, \ldots, x_n, \prec)$, with corresponding winning probabilities $p(i, x_1, \ldots, x_n, \prec)$.

3. A candidate strategy

Finding a constructive characterization of the optimal strategy $\mathrm{opt}(x_1, \ldots, x_k)$ is, under many tie-breaking rules, a difficult task. It turns out that for the "closest to the first player" tie-breaking'rule considered in the next section the optimal strategy is, surprisingly, fairly simple. However, instead of just giving the strategy and proving that it is optimal (we will do this in the next section), we would like to describe here the properties of the tie-breaking rule which can lead us to guess such a strategy. These properties might in turn be useful for finding optimal strategies under other possible tie-breaking rules.

After examining the problem for small values of n, one observes that many tie-breaking rules seem to obey some form of the following *optimality principle:*

A player who does not play optimally for himself only increases the chances of winning for the players that follow him.

We use restricted versions of this principle to obtain a plausible playing strategy. When applied to the last player, this principle says that the chances for the last player are worst when the players before him play optimally (since his chances only increase when those players play suboptimally). We thus make the assumption that all players play so as to make the payoff of the last player as small as possible.

The optimal payoff for the last player given the numbers chosen by the players before him is easy to compute. Suppose the moves of the first $n-1$ players have partitioned $[0, 1]$ into n intervals of respective lengths $a_1, \ldots, a_n$ from left to right. Then the optimal limiting move for the last player gives him a payoff of $\max(a_1, a_2/2, \ldots, a_{n-1}/2, a_n)$, since the payoff when playing anywhere between two other players is half the distance between the two players, while the payoff when playing in the leftmost (rightmost) interval can be made arbitrarily close to the interval size by playing sufficiently close to the leftmost (resp. rightmost) player.

[1] That is, if for all $\epsilon > 0$ there are y_i with $|y_i - x_i| < \epsilon$ such that $(y_k, \ldots, y_n, \prec) \in S$, then $(x_k, \ldots, x_n, \prec) \in S$.

Now suppose that the first $k-1$ players have already played, for some $k \geq 2$, partitioning $[0,1]$ into k intervals $J_1, \ldots, J_k$ of respective sizes $c_1, \ldots, c_k$, from left to right. Given any real $B > 0$, the payoff of the last player can be forced to be B/2 or less by the $n-k$ players before him that are left to play, provided that at least $\lceil \rho_i \rceil - 1$ players move into the interval J_i for each $1 \leq i \leq k$. Here ρ_i is the ratio

$$\rho_i = \begin{cases} c_i/B & \text{if } 1 < i < k, \\ c_i/B + 1/2 & \text{if } i = 1 \text{ or } i = k. \end{cases} \tag{1}$$

This is so because in the case $1 < i < k$ the interval J_i must be partitioned into subintervals of size B or less, while in the cases i = 1 and i = k the subintervals at both ends of [0, 1] must be of size $B/2$ or less, to ensure that the payoff for the last player is at most $B/2$. This can be achieved if and only if

$$\sum_{1 \leq i \leq k} (\lceil \rho_i \rceil - 1) < r, \tag{2}$$

where r = n-k+ 1 is the number of players left to play. If, following the optimality principle, the players before the last play so as to make the payoff $\hat{B}/2$ of the last player as small as possible, then $\hat{B}$ must be the least value of B satisfying (2). An equivalent definition for this key parameter is that $\hat{B}$ is the largest B satisfying

$$\sum_{1 \leq i \leq k} \lfloor \rho_i \rfloor \geq r; \tag{3}$$

in fact $\hat{B}$ is the only value of B which satisfies (2) and (3) simultaneously. Notice that the summands in (2) and (3) are equal unless ρ_i is an integer, in which case they differ by 1. Since only r players are left to play, it is natural in view of (3) to allocate at most $\lfloor \rho_i \rfloor$ players to the interval J_i, including player n. Then if r_i out of the last r players move into the interval J_i, we will have

$$\lceil \rho_i \rceil - 1 \leq r_i \leq \lfloor \rho_i \rfloor \text{ for all } 1 \leq i \leq k, \text{ with } \textstyle\sum_{1 \leq i \leq k} r_i = r. \tag{4}$$

Now we will try to determine more precisely where these r_i players should play inside the interval J_i, for $1 < i < k$. The first $r_i - 1$ players to play in J_i partition this interval into r_i subintervals. Furthermore, there must be a subinterval J_i' among these of size at least $\hat{B}$, since $r_i \leq \lfloor c_i/\hat{B} \rfloor$ by (4). The last player p to play in J_i should then be able to obtain a payoff of at least $\hat{B}/2$ by playing inside J_i'. Now the optimality principle applied to player p says that the $r_i - 1$ players that play in the interval J_i before player p will play so as to make the size of the subinterval J_i' as small as possible. However, the other subintervals of J_i should not be made larger than $\hat{B}$ because otherwise player n could move into J_i and obtain a payoff greater than $\hat{B}/2$. Thus the first $r_i - 1$ players in J_i should partition this interval into $r_i - 1$ subintervals of size $\hat{B}$, and one subinterval J_i' of size at least $\hat{B}$ for player p. Now if we assume that player p plays in J_i' so as to create a subinterval of size $\hat{B}$ as well, we obtain a simple playing strategy:

The r_i players that play in J_i, $1 < i < k$, partition this interval into r_i subintervals of size $\hat{B}$ and a leftover subinterval K_i of size at most $\hat{B}$.

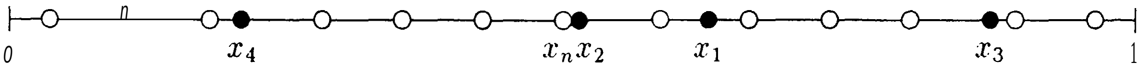

Figure 2: Optimal play $(x_k, \ldots, x_n, \prec)$ given $(x_1, \ldots, x_{k-1})$ for $n = 17$ and $k = 5$, under the "closest to the first player" tie-breaking rule. Here $x_n = x_2$ with $x_n \prec x_2$.

This is illustrated in Fig. 2. A subinterval of size $\hat{B}/2$ (instead of $\hat{B}$) is used at both ends of [0, 1] to keep the potential payoff for the last player bounded by $\hat{B}/2$. Also notice that we have placed the leftover subinterval K_i in each J_i as close as possible to the number x_1 played by player 1; this will be justified by the tie-breaking rule of the next section.

4. Optimal strategy under the "closest to the first player" tie- breaking rule

In this section we characterize the optimal playing strategy under the "closest to the first player" tie- breaking rule. This tie-breaking rule assumes that the optimal limiting play $(x_k, \ldots,$ x,, $\prec)$ chosen by a player $k \geq 2$ is such that no other optimal limiting play that has the first k players in the same relative order in the interval [0, 1] as the chosen play can have player k closer to the first player. Notice that we make no assumptions about how the tie-breaking rule chooses among equivalent moves for the first player, or about how it chooses the relative posit ion of a player $k \geq 2$ with respect to earlier players if several optimal alternatives are available: It turns out that the optimal strategy is essentially the same regardless of how these choices are made. Hence our solution applies not just to a single tie-breaking rule, but rather to a family of possible tie-breaking rules.

We recall some definitions from the previous section. Suppose for now that $k \geq 2$. The numbers $x_1, \ldots, x_{k-1}$ chosen by the first k-l players partition [0, 1] into k intervals $J_1, \ldots, J_k$, from left to right? of respective lengths $c_1, \ldots, c_k$. Let $\hat{B}$ be the unique value of B satisfying both (2) and (3),where the ρ_i are given as functions of B by (1), and r = $n - k + 1$; then $\hat{B}$ is also the smallest B satisfying (2) and the largest B satisfying (3).

We define $\lfloor \rho_i \rfloor$ *marked positions* $J_i(1), \ldots, J_i(\lfloor \rho_i \rfloor)$ inside each J_i. Let $J_i(0)$ and $J_i(\lfloor \rho_i \rfloor + 1)$ be the endpoints of J_i which are farthest from and closest to x_1, respectively. Then the marked positions inside J_i are given, for j = $1, 2, \ldots, \lfloor \rho_i \rfloor$, by

$$|J_i(j) - J_i(0)| = \begin{cases} j\hat{B} & \text{if } 1 < i < k, \\ (j - 1/2)\hat{B} & \text{if } i = 1 \text{ or } i = k. \end{cases} \tag{5}$$

We associate with each marked position J;(j) the payoff that a player at j;(j) would get, if each marked position was played by one player. Therefore the payoff of every marked position equals $\hat{B}$, except for the last marked position $J_i(\lfloor \rho_i \rfloor)$ of each interval J_i; this last position has a payoff of $(\hat{B}+l(i))/2$, where Z(i) is the size of the leftover interval $K_i = (J_i(\lfloor\rho_i\rfloor), J_i(\lfloor\rho_i\rfloor \$1))$ and is given by

$$l(i) = \begin{cases} c_i \bmod \hat{B} & \text{if } 1 < i < k, \\ (c_i + \hat{B}/2) \bmod \hat{B} & \text{if i} = 1 \text{ or } i = k. \end{cases} \quad (6)$$

We can now state our main result:

Theorem 1 *The optimal limiting play under the "closest to the first player" tie-breaking rule is given for $k \geq 2$ by* $\text{opt}(x_1, \ldots, x_{k-1}) = (x_k, \ldots, x_n, \prec)$, *where $x_k, \ldots, x_n$ are $n - k + 1$ marked positions of highest associated payoff, in order of decreasing payoff.*

Proof. Let us first determine which marked positions will be played according to this theorem. Notice that the total number of marked positions is at least r, the number of players left to play, by equation (3). Among these marked positions, those of least payoff are the positions $J_i(\lfloor\rho_i\rfloor)$ with ρ_i integer, which have I(i) = 0 and a payoff of $\hat{B}/2$. The other marked positions will thus be used first; however, there are only $\lceil\rho_i\rceil - 1$ of these other marked positions in each J_i, and hence less than r in total by (2), so the theorem implies that one or more of the marked positions of low payoff $\hat{B}/2$ will have to be used as well. From these remarks, it follows that the number of players that will end up in each J_i (according to the theorem) satisfies (4).

We will prove the theorem by induction on the number $n - k + 1$ of players left to play. For the base case $k = n$, the marked positions are endpoints $J_i(\lfloor\rho_i\rfloor + 1)$ of highest payoff, and these are also the optimal moves for the last player under the "closest to the first player" tie- breaking rule. For the inductive step, suppose without loss of generality that player $k < n$ plays in an interval J_i with the endpoint $J_i(0)$ to the left and the endpoint $J_i(\lfloor\rho_i\rfloor + 1)$ to the right. Then J_i is split into a left interval $L = (J_i(0), x_k)$ and a right interval $R = (x_k, J_i(\lfloor\rho_i\rfloor+1))$. The new set of $k + 1$ intervals will then have a new parameter $\hat{B}'$ satisfying (2) and (3), and we can assume inductively that the last n $- k$ players will play as indicated by the theorem for this new parameter $\hat{B}'$. Notice that $\hat{B}' \geq \hat{B}$, since $\hat{B}$ is still a solution to (3) for this new set of intervals, while $\hat{B}'$ is by definition the largest solution to (3).

Lemma 1 (Low payoff for unmarked positions) ***If*** *player k plays between $J_i(j - 1)$ and J;(j) (or at $J_i(j)$), his payoff will be no greater than the* ***payoff*** *associated with marked position $J_i(j)$, for $1 \leq j \leq \lfloor\rho_i\rfloor$, and smaller than $\hat{B}/2$ for $j = \lfloor\rho_i\rfloor + 1$.*

Proof. We claim that if player k plays at x_k, between $J_i(\text{ j} - 1)$ and *J;(j)* (or at *J;(j)),* then for j > 1 some player will play in L at $J_i(j - 1)$ or higher, and for j < $\lfloor\rho_i\rfloor$ some player will play in R at $J_i(j + 1)$ or lower. This claim implies the lemma: For $1 \leq \text{j} \leq \lfloor\rho_i\rfloor$, the payoff associated with marked position J;(j) is at least as high, since this payoff corresponds to a situation where x_k = J;(j) and no player plays between $J_i(j - 1)$ and $J_i(j + 1)$. For j = $\lfloor\rho_i\rfloor$ + 1, either j > 1, in which case the payoff is less than ***B/2*** because I(i) < $\hat{B}$, *or* j = 1, in which case the payoff is less than $\hat{B}/2$ because $\lfloor\rho_i\rfloor$ = 0 so that J_i is a small interval.

To prove the claim, we consider two cases. If $\hat{B}' = \hat{B}$, then some player after player k will play at $J_i(\ j - 1\ )$, for j > 1. Also, some player will play at distance $\hat{B}$ from x_k in R, and hence at $J_i(j+1)$ or lower, for $j < \lfloor \rho_i \rfloor$. The remaining case is $\hat{B}' > \hat{B}$. Since the new parameter $\hat{B}'$ is greater than $\hat{B}$, the number of players that can play in each interval is now at most $j-1$ for L, at most $\lceil \rho_i \rceil - j$ for R, and at most $\lceil \rho_{i'} \rceil - 1$ for each of the other intervals $J_{i'}$ with $i' \neq i$. By equation (2), these numbers can only add up to $r-1$, the number of players left to play after player k, if exactly $j-1$ players play in L, exactly $\lceil \rho_i \rceil - j$ in R, and exactly $\lceil \rho_{i'} \rceil - 1$ in the remaining $J_{i'}$. Then, for $j > 1$, the $(j-1)$th player in L (from left to right) is above $J_i(j-1)$. Also, for $j < \lfloor \rho_i \rfloor$, if the leftmost player in R were at $J_i(j+1)$ or above, then there could be no more than $\lceil \rho_i \rceil - j - 1$ players in R, since $\hat{B}' > \hat{B}$, so the leftmost player in R must in fact be below $J_i(j+1)$. This proves the claim and completes the proof of the lemma. □

Lemma 2 (Preservation of marked positions) *Let $J_i(j)$ be a marked position of highest associated payoff. Then among the limiting moves $(x_k, \ldots, x_n, \prec)$ available to player k there is one in which $x_k = J_i(j)$. and $x_k, \ldots, x_n$ are $n-k+1$ marked positions of highest associated payoff, in order of decreasing payoff.*

Proof. If player k plays at a marked position $J_i(j)$ (this is feasible unless $l(i) = 0$), then $\hat{B}$ also satisfies (2) for the new set of $k+1$ intervals, so $\hat{B}' = \hat{B}$. The new set of $k+1$ intervals therefore has a new marked position $L(j)$ at $J_i(j)$, with $L(j) \prec J_i(j)$. If none of the remaining players plays at $L(j)$, then they will all play at the old marked positions in order of decreasing payoff, as desired. On the other hand, if one of them plays at $L(j)$, then the condition of the lemma is not satisfied and in fact player k obtains a low payoff of $\hat{B}/2$ or less.

However, it is still possible for player k to obtain the limiting play described in the lemma by playing just below $J_i(j)$, at $J_i(j) - \epsilon$, where $\epsilon < \hat{B} - l(i)$. Now if $\hat{B}' = \hat{B}$, then the marked positions remain the same after player k has played, except for the marked positions ● $J_i(j')$ with $j' > j$, which change to $R(j'-j) = J_i(j') - \epsilon$. Letting ϵ go to zero gives a limiting play as stated in the lemma. The remaining case is $\hat{B}' > \hat{B}$. Then, as in the proof of Lemma 1, we must have $j-1$ players in L, $\lceil \rho_i \rceil - j$ players in R, and $\lceil \rho_{i'} \rceil - 1$ players in the remaining $J_{i'}$. As before, the leftmost player in R must be below $J_i(j+1)$ (provided that $\lceil \rho_i \rceil - j > 0$ so that R has at least one player), so $\hat{B}' < \hat{B} + \epsilon$; if we let ϵ go to zero then $\hat{B}'$ approaches $\hat{B}$ and the positions of the players approach the marked positions of highest payoff. If $\lceil \rho_i \rceil - j = 0$, then $J_i(j)$ is a marked position with payoff $B/2$ and can only be a position of highest payoff if $\lceil \rho_{i'} \rceil - 1$ equals zero for all intervals $J_{i'}$. But $\lceil \rho_{i'} \rceil - 1$ is the number of players in each ● $J_{i'}$, so player k would be the last player, cont rary to the assumption that $k < n$. □

We use these two lemmas to prove the theorem. If α is the payoff associated with a marked position $J_i(j)$ of highest payoff, then player k can obtain a payoff of α in the limit by playing at $J_i(j)$ or just below $J_i(j)$ as in Lemma 2. On the other hand, by Lemma 1, the only plays that could possibly give the same payoff α to player k involve playing between $J_i(j-1)$ and $J_i(j)$, where $J_i(j)$ is a marked position of payoff α, but then position $J_i(j)$ is closer to x_1 and hence should be preferred by player k under the "closest to the first player" tie-breaking rule. Thus the optimal limiting play is the one given in Lemma 2, completing the induction and proving the theorem. □

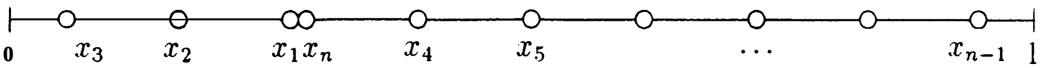

Figure 3: Optimal play $(x_1, \ldots, x_n, \prec)$ for n = 10 under the "closest to the first player" tie-breaking rule. Here $x_n = x_1$ with $x_1 \prec x_n$.

To complete the analysis, we must examine the optimal strategy for the first player. Consider the $n - 1$ positions $p_i = (2i - 1)/2(n - 1)$ for $1 \leq i \leq n - 1$. The above theorem implies that if player 1 plays at one of these positions, then the remaining n − 1 players will play in the limit at these $n - 1$ positions (one of them coinciding with player 1), giving player 1 a payoff of $1/2(n - 1)$ [see Fig. 3]. If on the other hand player 1 does not play at one of these positions, the remaining $n - 1$ players will play near these $n - 1$ positions but closer to player 1, giving him a payoff smaller than $1/2(n - 1)$. This gives:

Corollary 1 *In the optimal limiting play* $\text{opt}() = (x_1, \ldots, x_n, \prec)$ *under the "closest to the first player" tie-breaking rule, the limiting moves* $x_1, \ldots, x_{n-1}$ *of the first* $n - 1$ *players are the numbers* $1/2(n-1), 3/2(n-1), \ldots, (2n-3)/2(n-1)$*, in some order, and the limiting move* x_n *of the last player equals* x_1*. This gives a limiting payoff of* $1/2(n - 1)$ *for the first and last players, and* $1/(n - 1)$ *for the remaining players.*

Thus the youngest and the oldest children will be less likely to get the left-over toetje.

The preceding analysis can in fact be carried through even if the choice among equally desirable intervals is made at random or nondeterministically (i.e., chosen arbitrarily by each player when it is his turn to play), and even if only the last players — those that are left to play once each interval has at most one marked position — use the "closest to the first player" tie-breaking rule to select a position inside the chosen interval.

5. The "rightmost" and "random" tie-breaking rules

We first consider the "rightmost" tie-breaking rule. This tie-breaking rule assumes that the optimal limiting play $(x_k, \ldots, x_n, \prec)$ chosen by a player k is such that no other optimal limiting play that has the first k players in the same relative order in the interval [0, 1] as the chosen play can have player k further to the right. Again, we make no assumptions about how the tie-breaking rule chooses the relative position of a player with respect to earlier players if several optimal alternatives are available, so our solution actually applies to a family of tie-breaking rules.

It turns out that the solution for the "rightmost" tie-breaking rule is very similar to the solution for the "closest to the first player" rule, provided that we make a few changes to the marked positions. We again consider the intervals $J_1, \ldots, J_k$, and define marked positions as follows. Let $J_i(0)$ and $J_i(\lfloor \rho_i \rfloor + 1)$ be the left and the right endpoints of J_i respectively.

If $i < k$, the $\lfloor \rho_i \rfloor$ marked positions for J_i and their associated payoffs are defined as before. For J_k, on the other hand, we define $\lfloor \rho_k \rfloor$ marked positions by letting $J_k(j) = J_k(0) + j\hat{B}$ if $1 \leq j < \lfloor \rho_k \rfloor$ and $J_k(\lfloor \rho_i \rfloor) = J_k(\lfloor \rho_i \rfloor + 1) - \hat{B}/2$. Therefore the associated payoff of every marked position in J_k equals $\hat{B}$, except for the last two marked positions $J_k(\lfloor \rho_k \rfloor - 1)$ and ● lk($\lfloor \rho_k \rfloor$); these two positions have a payoff of $(B+l(k))/2$, where $l(k)$ is the size of the leftover interval $K_i = (J_k(\lfloor \rho_k \rfloor - 1), J_k(\lfloor \rho_k \rfloor))$ and is given by $l(k) = (c_k + \hat{B}/2) \bmod \hat{B}$. In particular, if ρ_k is an integer with $\rho_k \geq 2$, then *we* have a *double marked position* $J_k(\lfloor \rho_k \rfloor - 1) = J_k(\lfloor \rho_k \rfloor)$ with $J_k(\lfloor \rho_k \rfloor - 1) \prec J_k(\lfloor \rho_k \rfloor)$, and each of these two positions has an associated payoff of $\hat{B}/2$. For the particular case $\rho_k = 2$, we will also need to consider a *special position* $J_k()$ just to the right of $J_k(0)$, that is, $J_k() = J_k(0)$ with $J_k(0) \prec J_k()$; this special position is not considered a marked position.

Using this new set of marked positions, the results of the previous section for the "closest to the first player" tie-breaking rule can be extended to the "rightmost" tie-breaking rule as well, provided that the somewhat subtle situations that arise when a double marked position is present are properly treated.

Theorem 2 *The optimal limiting play under the "rightmost" tie-breaking rule is given for* $k \geq 2$ *by* $\operatorname{opt}(x_1, \ldots, x_{k-1}) = (x_k, \ldots, x_n, \prec)$*, where one of the following conditions holds:*

1. *There is no double marked position, and* $x_k, \ldots, x_n$ *are marked positions of highest payoff, in order of decreasing payoff.*

2. *There is a double marked position, two* x_i *among* $x_k, \ldots, x_n$ *are at the double marked position, and* $x_k, \ldots, x_n$ *are marked positions of highest payoff, in order of decreasing payoff.*

3. *There is a double marked position, only one* x_i *among* $x_k, \ldots, x_{n-1}$ *is at the double marked position, and* $x_k, \ldots,$ Xi-1 $, x_{i+1}, \ldots, x_n$ *are marked positions of highest payoff, in order of decreasing payoff.*

4. *There is a special position,* x_n *is at the special position, no* x_i *is at the double marked position, and* $x_k, \ldots, x_{n-1}$ *are marked positions of highest payoff, in order of decreasing payoff.*

If we now consider, for each possible move of the first player? the optimal play for the remaining players described by the four cases in the theorem, we obtain:

Corollary 2 *In the optimal limiting play* $\operatorname{opt}() = (x_1, \ldots, x_n, \prec)$ *under the "rightmost" tic-breaking rule, the limiting moves* $x_1, \ldots, x_{n-1}$ *of the first* $n - 1$ *players are the numbers* $1/2(n-1), 3/2(n-1), \ldots, (2n-3)/2(n-1)$*, in some order, and the limiting move* x_n *of the last player equals* x_{n-1}*. This gives a limiting payoff of* $1/2(n-1)$ *for the last two players, and* $1/(n-1)$ *for the remaining players.*

For other tie-breaking rules, the situation is more complicated. To simplify things and avoid dealing with the special intervals J_1 and and J_k, let us assume that for all the intervals

J_i, both endpoints of J_i have been taken by earlier players. This will be the case if the numbers $\mathbf{0}$ and 1 are taken at the beginning of the game (or, alternatively, if the game is played on a circle).

Under this simplifying assumption, there is no difference between the "closest to the first player" and the "rightmost" tie-breaking rules. We now examine the situation under the more equitable tie-breaking rule in which the choice among equivalent moves is made uniformly at random.

If we follow the heuristic argument of section 3, we would expect that the number of players r_i in each J_i would still satisfy (4). Under the "random" tie-breaking rule, however. the subintervals of size $\hat{B}$ would be taken by the first $r_i - 1$ players in J_i at random from either the left end or the right end of J_i, leaving a subinterval J_i' of size greater than $\hat{B}$ somewhere in between. Then the last player in J_i would play at random inside J_i' but in such a way as to partition J_i' into two subintervals of size at most $\hat{B}$.

This is indeed the optimal strategy when the total number of players in each interval is small. The inductive step, however, fails when the number of players is large. Thus, in the seven-player game with both endpoints taken, if all players follow the strategy we described, the expected payoff for the first player will be $(1/7)(1 - (1/2)^7) < 1/7$. but if the first player chooses the midpoint of the [O. 1] interval and the remaining six players follow the strategy we described, playing three to each side of the first player, the expected payoff for the first player will be $7/48 > 1/7$ (see [1] for the details of these calculations).

Nevertheless, since this strategy is optimal when the number of players is small, it may provide a good base case from which to extend the solution for the "random" tie-breaking rule to larger values of n.

6. Conclusion

We have shown that a simple zero-sum game in which all players play independently so as to maximize their own payoff can be viewed, under certain conditions, as a game in which the players cooperate to minimize the payoff of the last player. Furthermore, this approach is helpful in developing an optimal strategy for the game. An interesting question is whether this principle can be applied to a broader family of zero-sum games.

Acknowledgments

The author would like to thank Donald Knuth, Tomas Rokicki, and the other students in the problem seminar for their ideas and suggestions on this problem.

References

[1] Tomas G. Rokicki and Donald E. Knuth, "A programming and problem-solving seminar," Stanford University Technical Report STAN-CS-G-11.54 (1987).