

Full Length Article

Nash equilibria in the Showcase Showdown game with unlimited spins



L. Bayón^a, P. Fortuy Ayuso^{a, }, J.M. Grau^{a, },* , A.M. Oller-Marcén^{b, },
M.M. Ruiz^{a, }

^a Departamento de Matemáticas, Universidad de Oviedo, Avda. Calvo Sotelo, s/n 33007, Oviedo, Spain

^b Departamento de Matemáticas - IUMA, Universidad de Zaragoza, Pedro Cerbuna, 12, 50006 Zaragoza, Spain

ARTICLE INFO

Keywords:

Nash equilibrium
Optimal stopping
 n -Person game
Threshold strategy

ABSTRACT

The game of *Showcase Showdown* with unlimited spins is investigated as an n -players continuous game, and the Nash Equilibrium strategies for the players are obtained. The sequential game with information on the results of the previous players is studied, as well as three variants: no information, possibility of draw, and different modalities of winner payoff.

1. Introduction

A well known example of applied probability in TV is the game show: “The Price is Right”. One of the games of the show is called “The Showcase Showdown.” In it, each of three players spins a wheel in turns. The wheel has 20 values, between 5c and \$1, in increments of 5c. Each player spins once, and then, after seeing the result, has two options: either spinning again, or stopping. There are a maximum of two spins and if the player spins twice, both values are added. If a player exceeds a dollar, he is immediately eliminated. Otherwise, the turn passes to the next player. The aim of each player is to obtain the highest score less than or equal to 1. Each player knows the results of the previous ones. The natural question is: What is the best strategy for playing Showcase Showdown?

Coe and Butterworth [1] define the optimal stopping time for player i ($i = 1, 2, 3$) as the smallest value such that it is better for player i to stop after the first spin rather than spinning a second time. Tenorio and Cason [11] also analyzed this discrete game with three players, assuming that the result of each spin is a discrete random variable, uniformly distributed on the set $\{0.05, 0.10, \dots, 1.00\}$, and that the spins are independent random events.

Kaynar [2] considers a variant in which each player draws one or two random numbers between 0 and 1, and where each player has no information about the results and actions of the previous ones. The optimal solution with two possible attempts and two and three players is also given there.

In several studies, Sakaguchi [4,6] covers the game for two players, each of whom can play once or twice, and assuming each sampling follows a uniformly distributed random variable in $[0, 1]$. The aim of the version of “Showcase Showdown” he calls (GSS) is to obtain the highest score among all of the players in the game, from one or two chances of sampling. He also studies three

* Corresponding author.

E-mail addresses: bayon@uniovi.es (L. Bayón), fortunypedro@uniovi.es (P. Fortuy Ayuso), grau@uniovi.es (J.M. Grau), oller@unizar.es (A.M. Oller-Marcén), mrui@uniovi.es (M.M. Ruiz).

<https://doi.org/10.1016/j.amc.2025.129368>

Received 6 November 2024; Received in revised form 17 February 2025; Accepted 18 February 2025

different versions of this game, depending on the scoring function: “Keep-or-Exchange” (GKE), “Competing Average” (GCA), and “Risky Exchange” (GRE). In [4], he solves GSS, while in [6] he does so for GKE and GRE. The version GCA has not been solved yet. Also, in [5,8] the same author studies the games with two sampling possibilities but three players.

Elsewhere, also Sakaguchi [7], solves the GKE and GRE versions with two players and three different types of information: no information sharing at all, that each player informs the other one of his results, and that the first one informs the second one of his results but not the other way around.

Another generalization is analyzed by Swenson [10], with the same rules as Showcase Showdown with n players, but where each spin follows a continuous random variable uniformly distributed on $[0, 1]$: the n -player continuous game. He raises the question of the existence of optimal cutoff values for the case in which more than two spins are allowed.

Our aim is to study the generalization of Showcase Showdown to any number of samplings (called *spins*) and of players. This has been scarcely studied. Mazalov and Ivashko [3] analyze the GSS game with no information (each player makes his decisions with no knowledge of the results of the others). Using dynamic programming theory, they find the Nash equilibrium of the n -player GSS with an infinite number of spins: a state when it is unprofitable for all players to individually deviate from their strategies, assuming that if the scores of all players are more than 1, then the winner is the one whose score is closest to 1. Seregina, Ivashko and Mazalov [9] study the same no-information version but with n spins, providing the optimal payoffs. The optimal strategies of the players in the version with complete information are studied and the optimal payoffs for the first player are computed.

This work covers the n -player GSS with an infinite number of spins, where the aim of each player is to obtain the highest total score less than (or equal to) 1, allowing for the possibility of draw among players when all their scores are strictly greater than 1. We study the sequential version of the game, with information on the previous players results, studying their optimal strategies and obtaining formulas which allow us to compute the optimal payoffs not only of the first player but of all of them. We also state and solve three variants of the no-information version with different payoff modes under the condition if all the scores are above 1, then the payoff is 0. Finally, we study the Nash equilibria as the expected gains in each variant.

In summary, our main contributions are:

- (A) We provide formulas allowing the computation of the winning probability of each player in the sequential game, when all of them play maximizing their own winning probability (Proposition 2).
- (B) We describe and provide examples of how, in the sequential case, coalitions between players can reduce other player’s winning probabilities (Section 5.1).
- (C) We compute the formulas giving the Nash equilibria in different versions of the game; these formulas show that, despite the essential aim of the game being the same, the diverse payoffs can give rise to great changes in the equilibrium strategies (Propositions 5, 6, and 7, one for each variant of the game we study).

The paper is organized as follows: In Section 2, we describe the versions of the n -person Showcase Showdown game with unlimited spins we cover. In section 3 we present a version *ad hoc* of the one-stage look-ahead (OLA) stopping rule. In section 4 we compute the cumulative distribution function of the random variable given by the score of each player depending on his greed threshold. The optimal strategies and expected payoffs of all players are given for the sequential game in Section 5. In Section 6, we study three variants of the no-information game, with different payoffs for the winner. Finally, in section 7, some prospects for the future are presented that we consider interesting.

2. The Showcase Showdown game

We shall consider several cases of the following version of the Showcase Showdown game, with set of players $\{A_i\}_{i=1}^n$: each player A_i starts the game with a value $S(A_i) = 0$. Then he successively retrieves a value from a uniformly distributed random variable in $[0, 1]$, and adds this value to $S(A_i)$ (this act will be called a *play* of the player); this retrieval is repeated until either he stops or $S(A_i) > 1$. If he stops with $S(A_i) \leq 1$, then his score is $S(A_i)$, otherwise $S(A_i)$ is irrevocably set to $S(A_i) = 0$, and player A_{i+1} starts his turn. The winner is the player with the greatest score, and there is a draw if all the scores are 0.

The variants we shall consider are the following, in all of which the final payoff for the winner is 1.

1. **GAME I — Sequential game.** Player A_i knows the values $S(A_j)$ for $j = 1, \dots, i - 1$.
2. **GAME II — No-information game.** There is no information available about $S(A_j)$ for any other player. There are three sub-variants: the first two depending on where the payment comes from (either an external agent or the rest of the players), and the third one, in which one pre-selected player, say A_n , is the winner in case of draw.
 - **GAME II.1 — Non-constant-sum.** The winner’s payoff is provided by an external agent. If $S(A_i) = 0$ for all i , then the payoff is 0 for all players. Thus, the sum of all payoffs may be 0.
 - **GAME II.2 — Zero-sum.** The payoff of the winner is collected from the other players, each providing $1/(n - 1)$. As in the previous case, the payoff is 0 if $S(A_i) = 0$ for all i .
 - **GAME II.3 — Non-symmetric and constant-sum.** There is a known player A_j with advantage: if $S(A_i) = 0$ for all i , then A_j wins. We shall set $j = n$ as the index is irrelevant.

We shall show how the reasonable strategies are all based on establishing a *greed threshold* for each player A_i : a value $\kappa_i \in (0, 1)$ such that A_i continues playing (i.e. retrieving a random number and adding it to $S(A_i)$) until $S(A_i) > \kappa_i$. This way, the no-information

variants can be understood as continuous games in which each player A_i computes his greed threshold κ_i , and where the final payoff is given by some functions $\mathbb{P}_i(\kappa_1, \dots, \kappa_n)$. These payoffs depend, obviously, on the specific game. In the sequential version (GAME I), we shall also see how the optimal policy is of this type and depends on the score of the previous players, and the number of players still to play.

3. Optimal strategy. Threshold strategy

In this section we restrict ourselves to the Showcase Showdown game with a single player with payoff $h(x)$ if he stops with score $S = x$. The function $h(x)$ is defined in $[0, 1]$ and assumed non-decreasing. This includes, for instance, the payoff $h(x) = x$ (ordinary single-player Showcase Showdown). This situation can also happen with several players in the sequential game (GAME I) and where $h(x)$ is the probability that no later player gets a score greater than x . In no-information games, $h(x)$ will represent the expected payoff when stopping at x , assuming certain rival strategies.

Roughly speaking, one can say that the optimal policy in this case (single player) consists in: given the value S , decide which of playing again or stopping has better expected payoff, and act accordingly. Let $G(x)$ be the expected payoff resulting from playing again and following the optimal policy (whatever this may be) from that point on. When $S = x$, the expected payoff following that strategy is necessarily $\max(h(x), G(x))$, so that the expected payoff $G(x)$ of continuing playing with $S = x$ and following, later on, the optimal policy, must satisfy:

$$G(x) = h(0)x + \int_0^{1-x} \max(h(x+t), G(x+t))dt,$$

where $h(0)$ is the payoff when $S > 1$ (i.e. final score 0); the probability of this when $S = x$ is x . This gives the following integral equation for $G(x)$:

$$G(x) = h(0)x + \int_x^1 \max(h(t), G(t))dt.$$

The following lemma proves that $G(x)$ exists and can be explicitly defined (so that it is also unique) in terms of $h(x)$. This is an *ad-hoc* version of the one-stage look-ahead (OLA) stopping rule ([3]), which compares the payoff if stopping when $S = x$ with the expected payoff of making a single play more, and stopping.

Lemma 1. *Let $h : [0, 1] \rightarrow \mathbb{R}$ be a non-decreasing monotone function, and let $\tilde{h}(x) = h(0)x + \int_x^1 h(t)dt$, for $x \in [0, 1]$. Define*

$$\kappa := \inf \{x \in [0, 1] : h(x) \geq \tilde{h}(x)\}.$$

Then the function

$$G(x) := \begin{cases} (\tilde{h}(\kappa) - h(0))e^{\kappa-x} + h(0) & \text{if } x < \kappa, \\ \tilde{h}(x) & \text{if } x \geq \kappa, \end{cases} \tag{1}$$

is a non-increasing monotone $C[0, 1]$ function, and the only one satisfying the integral equation

$$y(x) = h(0)x + \int_x^1 \max(h(t), y(t))dt. \tag{2}$$

Proof. Before proceeding, notice that $h(x)$ being non-decreasing implies that it is Riemann integrable in $[0, 1]$, so that the definition of $\tilde{h}(x)$ and Equation (2) make sense. By definition, $\tilde{h}(x)$ is a continuous function, so that $\tilde{h}(\kappa)$ is well-defined.

Let us first verify that $G(x)$ as defined in (1) is a solution of (2). By construction, $G(x)$ is continuous, as both parts are continuous and they coincide at $x = \kappa$. Also by construction, it is non-increasing for $x < \kappa$. Now, if $x \geq \kappa$ and $\epsilon \geq 0$, then, by definition of κ and $G(x)$, and because $h(x)$ is non-decreasing, we have:

$$\begin{aligned} G(x + \epsilon) &= \tilde{h}(x + \epsilon) = h(0)(x + \epsilon) + \int_{x+\epsilon}^1 h(t)dt = h(0)x + h(0)\epsilon + \int_{x+\epsilon}^1 h(t)dt \leq \\ &h(0)x + h(x)\epsilon + \int_{x+\epsilon}^1 h(t)dt \leq h(0)x + \int_x^{x+\epsilon} h(t)dt + \int_{x+\epsilon}^1 h(t)dt = \\ &h(0)x + \int_x^1 h(t)dt = \tilde{h}(x) = G(x), \end{aligned}$$

so that $G(x)$ is non-increasing for $x \geq \kappa$. Thus, $G(x)$ is non-increasing in the whole interval $[0, 1]$.

By construction, and by definition of κ , $G(x) = \tilde{h}(x) \leq h(x)$ for $x > \kappa$. Also, as $h(x)$ is non-decreasing and $G(x)$ is non-increasing, we infer that $G(x) \geq h(x)$ for $x < \kappa$. Thus, κ satisfies also:

$$\kappa = \inf \{x \in [0, 1] : G(x) \leq h(x)\},$$

which implies that for $x \in (\kappa, 1]$, we have $\max(h(x), G(x)) = h(x)$, which gives:

$$G(x) = h(0)x + \int_x^1 \max(h(t), G(t))dt \text{ for } x > \kappa,$$

that is: $G(x)$ satisfies (2) for $x > \kappa$ and obviously too for $x = \kappa$.

On the other hand, as $G(x)$ is non-increasing, we obtain $G(x) \geq h(x)$ for $x < \kappa$. This gives, for $x \in [0, \kappa)$:

$$G(x) = h(0)x + \int_x^\kappa G(t)dt + \int_\kappa^1 h(t)dt,$$

which is, by construction,

$$G(x) = h(0)x + \int_x^1 \max(h(t), G(t))dt,$$

as required. Thus, $y(x) = G(x)$ is a solution of (2).

In order to show its uniqueness, we apply Banach’s fixed point theorem. Let $\alpha \in (0, 1]$ and consider the map:

$$\begin{aligned} C[\alpha, 1] &\xrightarrow{\varphi} C[\alpha, 1] \\ y(x) &\longmapsto h(0)x + \int_x^1 \max(h(t), y(t))dt \end{aligned}$$

and, in $C[\alpha, 1]$ consider the supremum metric. Take $y_0, y_1 \in C[\alpha, 1]$. We have

$$\|y_0 - y_1\| = \max_{x \in [\alpha, 1]} \left| h(0)x + \int_x^1 \max(y_0(t), h(t))dt - h(0)x - \int_x^1 \max(y_1(t), h(t))dt \right|,$$

so that

$$\|y_0 - y_1\| \leq \max_{x \in [\alpha, 1]} \int_x^1 |\max(y_0(t), h(t)) - \max(y_1(t), h(t))|dt.$$

Given three real numbers a, b, c , $|\max(a, b) - \max(a, c)| \leq \max(b - c)$: the only possible values of $|\max(a, b) - \max(a, c)|$ are: 0, $|b - c|$, $|a - c|$ and $|a - b|$, but the last two can only happen if a is between b, c , so that in any case $|\max(a, b) - \max(a, c)| \leq |b - c|$. As a consequence,

$$\|y_0 - y_1\| \leq \max_{x \in [\alpha, 1]} \int_x^1 |y_0(t) - y_1(t)|dt \leq (1 - \alpha) \max_{t \in [x, 1]} |y_0(t) - y_1(t)| \leq (1 - \alpha)\|y_0 - y_1\|.$$

As $\alpha > 0$, we deduce that φ is a contraction map, and has a single fixed point in $C[\alpha, 1]$ for any α . As G is continuous in $C[0, 1]$, its restriction to any $C[\alpha, 1]$ is that unique solution. This gives the uniqueness of G , as $G(0)$ is determined by its continuity. \square

In the following result, “optimal” means “the expected payoff is maximum”. Notice that both “stopping” and “continuing playing” might be optimal at the same time.

Proposition 1. Consider the single-player Showcase Showdown game with payoff function $h(x)$ for score x . Assume $h : [0, 1] \rightarrow \mathbb{R}$ is non-decreasing monotone and let $\tilde{h}(x) = h(0)x + \int_x^1 h(t)dt$, for $x \in [0, 1]$. Then, there exists an optimal policy which is of threshold type. In other words, there is $\kappa \in [0, 1]$ (optimal threshold) such that if at some point the total score is $S < \kappa$, then the optimal decision consists in continuing playing, whereas if $S > \kappa$ then the optimal decision is to stop. Finally, if $S = \kappa$, the optimal decision is to stop if and only if $h(\kappa) \geq \tilde{h}(\kappa)$. Furthermore,

$$\kappa = \inf \{x \in [0, 1] : h(x) \geq \tilde{h}(x)\},$$

and the expected payoff following this policy (that is, the optimal expected payoff) is

$$E = (\tilde{h}(\kappa) - h(0))e^\kappa + h(0). \tag{3}$$

Proof. The expected payoff of continuing playing with score x is $G(x)$ as defined in (1) in Lemma 1. We know that $G(x) > \tilde{h}(x)$ for $x \in [0, \kappa)$, and that $G(x) = \tilde{h}(x)$ for $x \in [\kappa, 1]$.

For any $x \in [0, 1]$, a strategy which consists in stopping for $S = x$ is optimal if and only if the expected payoff when stopping is greater than or equal to the expected payoff when continuing. Thus, stopping is optimal if and only if $h(x) \geq G(x)$. If $x < \kappa$, then $h(x) < \tilde{h}(x)$, which implies that $h(x) < G(x)$ and the optimal strategy is to continue playing and not stopping. If $x \in (\kappa, 1]$, then clearly stopping is optimal because $h(x) \geq \tilde{h}(x) = G(x)$. When $x = \kappa$, stopping is optimal if and only if $h(\kappa) \geq \tilde{h}(\kappa)$ because $G(\kappa) = \tilde{h}(\kappa)$.

The expected payoff following the strategy in the statement is, clearly, the expected payoff at the start of the game: $G(0) = (\tilde{h}(\kappa) - h(0))e^\kappa + h(0)$, which is E . \square

Obviously, if h is continuous and $h(x) \neq h(0)$ for some $x \in (0, 1)$, then the optimal threshold κ is the unique root of the equation $h(x) = h(0)x + \int_x^1 h(t)dt$. Notice that, in this case, if at some stage the score is κ , then both stopping and continuing are optimal strategies.

We provide two simple examples showing how the previous results apply to different payoff functions, one continuous and the other discontinuous.

Example 1. Assume the payoff in a single-player Showcase Showdown game like above is $h(x) = x$. The optimal threshold is $\kappa = \sqrt{2} - 1 \simeq 0.41421$, the only root of the equation $x = \int_x^1 tdt$. The expected payoff is $\kappa e^\kappa \simeq 0.62678$.

Example 2. Assume now the non-continuous payoff function

$$h(x) := \begin{cases} x, & \text{if } x < 1/2, \\ 7x, & \text{if } x \geq 1/2. \end{cases}$$

The optimal threshold is

$$\kappa = \inf \left\{ x \in [0, 1] : h(x) \geq \int_x^1 h(t)dt \right\} = \frac{1}{2}.$$

The expected payoff is (following the notation of Proposition 1, equation (3)), $(\tilde{h}(\kappa) - h(0))e^\kappa + h(0) = 21e^{1/2}/8 \simeq 4.3278$.

4. Cumulative distribution function of the score

We can compute the cumulative distribution function of the random variable ξ_τ representing the score of a player who follows the threshold strategy given by the threshold τ .

Lemma 2. Let $\{Z_n\}_{n \in \mathbb{N}}$ be a sequence of independent uniformly distributed random variables in $[0, 1]$. Define

$$\chi_0 := 0 \text{ and } \chi_n := \chi_{n-1} + Z_n,$$

and, for $\tau \in [0, 1]$,

$$\xi_\tau := \begin{cases} \chi_n, & \text{if } \chi_{n-1} < \tau \leq \chi_n \leq 1, \\ 0, & \text{if } \chi_{n-1} < \tau \text{ and } \chi_n > 1. \end{cases}$$

Then the cumulative distribution function of ξ_τ is

$$F_\tau(x) := P(\xi_\tau \leq x) = \begin{cases} 0, & \text{if } x < 0, \\ 1 + e^\tau (\tau - 1), & \text{if } 0 \leq x \leq \tau, \\ 1 + e^\tau (x - 1), & \text{if } \tau < x \leq 1, \\ 1, & \text{otherwise.} \end{cases} \tag{4}$$

Proof. The sequence of random variables above consists in following a threshold strategy with greed threshold τ . Thus, we shall reason with a player following such a strategy.

Assume the player has played j times, for $j \in \mathbb{Z}_{\geq 0}$ and has score $S = t$. For $t \leq \tau$, let $f(t)$ be the probability that the final score is $S \in [\tau, x]$.

$$f(t) := P(\tau \leq \xi_\tau \leq x | \chi_j = t).$$

In order to compute $f(t)$, we need to consider the two possibilities leading to $S \in [\tau, x]$ assuming that $\chi_j = t$:

- i) Either $Z_{j+1} \in [\tau - t, x - t]$ (i.e. the process ends after playing once more), which happens with probability $x - \tau$.

ii) More than just another play is required. The probability we are trying to compute is the expected value of the probability of ending in $[\tau, x]$ (that is, $f(t + s)$) starting with $\chi_j = t + s$ for $s = [0, \tau - t]$:

$$\int_0^{\tau-t} f(t + s) ds.$$

Thus,

$$f(t) = (x - \tau) + \int_0^{\tau-t} f(t + s) ds,$$

which, as f is necessarily continuous if it satisfies that equation, gives the differential equation

$$f'(t) = -f(t),$$

with the condition $f(\tau) = x - \tau$, whose solution is $f(t) = e^{\tau-t} (x - \tau)$. We obtain:

$$P(\tau \leq \xi_\tau \leq x) = f(0) = e^\tau (x - \tau).$$

Therefore, the cumulative distribution function of ξ_τ is given, for $x \in [0, \tau]$, by:

$$F_\tau(x) = P(0 \leq \xi_\tau \leq \tau) = 1 - P(\tau \leq \xi_\tau \leq 1) = 1 - e^\tau (1 - \tau), \tag{5}$$

while, for $x \in [\tau, 1]$, we have:

$$F_\tau(x) = P(0 \leq \xi_\tau \leq x) = P(0 \leq \xi_\tau \leq \tau) + P(\tau \leq \xi_\tau \leq x) = 1 + e^\tau (x - 1), \tag{6}$$

and the rest of the statement follows. \square

Remark 1. Notice that $F_\tau(x)$ is the distribution of the product of two independent random variables, one of which is Bernoulli with parameter $e^\tau (1 - \tau)$, and the other uniform on $[\tau, 1]$:

$$\xi_\tau \sim \text{Be}(e^\tau (1 - \tau)) \mathbf{U}[\tau, 1].$$

The “success” event in the Bernoulli random variable $\text{Be}(e^\tau (1 - \tau))$ means getting a payoff greater than 0 using τ as greed threshold, which is the same as obtaining $S \in [x, 1]$ with uniform distribution for S in that interval.

Observe that if $\tau \in [0, 1]$, then by (5) and (6), we obtain $P(0 \leq \xi_\tau \leq \tau) = P(\xi_\tau = 0)$. In what follows, we shall use the notation

$$\mathbf{P}(x) := P(\xi_x = 0) = 1 + e^x (x - 1). \tag{7}$$

to represent the probability of a player with greed threshold x to get score $S = 0$.

The following easy corollary will be useful.

Corollary 1. Let h be a Riemann integrable function on $[0, 1]$. Then,

i) The expected value $\mathbb{E}[h(\xi_x)]$ satisfies:

$$\mathbb{E}[h(\xi_x)] = \mathbf{P}(x)h(0) + e^x \int_x^1 h(t) dt.$$

ii) The conditional expected value $\mathbb{E}[h(\xi_x) | \xi_x > 0]$ satisfies:

$$\mathbb{E}[h(\xi_x) | \xi_x > 0] = \int_x^1 \frac{h(t)}{1 - x} dt, x \in [0, 1). \tag{8}$$

Proof. Since $\xi_x \sim \text{Be}(e^x (1 - x)) \cdot \mathbf{U}[x, 1]$, we have that

$$\mathbb{E}[h(\xi_x)] = \mathbf{P}(x)h(0) + (1 - \mathbf{P}(x)) \int_x^1 \frac{h(t)}{1 - x} dt = \mathbf{P}(x)h(0) + e^x \int_x^1 h(t) dt,$$

while

$$\mathbb{E}[h(\xi_x)|\xi_x > 0] = \mathbb{E}[h(U[x, 1])] = \int_x^1 \frac{h(t)}{1-x} dt. \quad \square$$

5. Game I. Sequential game

In this game, there are n players which play just once each, sequentially. Each one knows the score obtained by the previous ones. The winner receives a unitary payoff, and it is irrelevant whether this comes from an external payer or from the other players. Our aim is to find the optimal policy for each player, that is, the threshold strategy which maximizes his expected payoff—or, what is the same, his probability of winning.

The Showcase Showdown game without draw is studied in [9]: *if all the scores exceed 1, then the winner is the player with the lowest one*. In the sequential game, if all previous players had score 0, then the last player (assuming a rational behavior) wins simply stopping after the first play. Thus, the present section contains parts of section 3.2 of that reference, in which the optimal policy for each player is provided, as well as the probability of the first player winning (but just the first one). In this work, using some delicate arguments, we have been able to define a recursive procedure for computing each players' probability of winning, assuming all of them act optimally. We also include, for the case of three players, a study of the possible coalitions that can be made that increase the joint probability of winning for the coalesced.

Lemma 3. For all $n > 0$, the equation $\mathbf{P}(x)^{n-1} = \int_x^1 \mathbf{P}(t)^{n-1} dt$, i.e.,

$$(1 - e^x + xe^x)^{n-1} = \int_x^1 (1 - e^t + te^t)^{n-1} dt, \tag{9}$$

has a single solution in $[0, 1]$, which will be denoted θ_n .

Proof. Put $f_n(x) = (1 - e^x + xe^x)^{n-1}$ and $g_n(x) = \int_x^1 (1 - e^t + te^t)^{n-1} dt$. If $n = 1$, then $f_1(x) = 1$ and $g_1(x) = 1 - x$ so $x = 0$ is the only solution. Now, for $n > 1$, we have that $f_n(0) = 0$, $f_n(1) = 1$, $g_n(0) > 0$ and $g_n(1) = 0$. Since f_n is increasing and g_n is decreasing in $[0, 1]$ for every $n > 0$, the result follows. \square

Lemma 4. The sequence $\{\theta_n\}_{n>0}$ is strictly increasing.

Proof. With the previous notation, it is enough to show that $f_{n+1}(\theta_n) < g_{n+1}(\theta_n)$. Then the result follows just like in Lemma 3. To do so, first note that

$$(1 - e^{\theta_n} + \theta_n e^{\theta_n})(1 - e^t + te^t)^{n-1} < (1 - e^t + te^t)^n,$$

for every $t \in (\theta_n, 1]$. Then,

$$f_{n+1}(\theta_n) = (1 - e^{\theta_n} + \theta_n e^{\theta_n})f_n(\theta_n) = (1 - e^{\theta_n} + \theta_n e^{\theta_n})g_n(\theta_n) < g_{n+1}(\theta_n),$$

and the claim follows. \square

Proposition 2. In Game I, let us assume that there are r players still to play. If M_r is the maximum score of the players who have already played, then the optimum threshold for the next player is $\max\{\theta_r, M_r\}$.

Proof. We argue by induction on r . The case $r = 1$ is trivial because the optimum threshold of the last player is obviously $M_1 = \max\{\theta_1, M_1\}$ (notice that, as θ_n satisfies (9), we have $\theta_1 = 0$). Thus, let $r > 1$ and assume that the result holds for $r - 1$.

We will call $h_r(x)$ the probability of winning stopping with score $S = x$. By induction hypothesis, $h_r(x)$ is the probability that the following players finish with a score of 0 when trying to exceed their respective optimal thresholds. And, obviously, if $x \leq M_r$ the probability of winning is 0, so

$$h_r(x) = \begin{cases} \prod_{s=1}^{r-1} \mathbf{P}(\max(\theta_s, x)) & \text{if } x > M_r, \\ 0 & \text{if } M_r \geq x. \end{cases}$$

Now, to be in the conditions of Proposition 1, we must see that

$$\max(M_r, \theta_r) = \inf \left\{ x \in [0, 1] : h_r(x) \geq \int_x^1 h_r(t) dt \right\}.$$

- If $M_r < \theta_r$, then, $h_r(x) = \mathbf{P}(x)^{r-1}$ for all $x > \theta_r$. Thus, taking into account that $\mathbf{P}(\theta_r)^{r-1} = \int_{\theta_r}^1 \mathbf{P}(t)^{r-1} dt$, it follows that

$$\inf \left\{ x \in [0, 1] : h_r(x) \geq \int_x^1 h_r(t) dt \right\} = \inf \left\{ x \in [0, 1] : \mathbf{P}(x)^{r-1} \geq \int_x^1 \mathbf{P}(t)^{r-1} dt \right\} = \theta_r,$$

and ultimately

$$\inf \left\{ x \in [0, 1] : h_r(x) \geq \int_x^1 h_r(t) dt \right\} = \max(M_r, \theta_r).$$

- If $M_r \geq \theta_r$, then $h_r(x) \geq \int_x^1 h_r(t) dt$ for all $x \geq M_r$, and $0 = h_r(x) < \int_x^1 h_r(t) dt$ for all $x < M_r$, so that

$$\max(M_r, \theta_r) = M_r = \inf \left\{ x \in [0, 1] : h_r(x) \geq \int_x^1 h_r(t) dt \right\}. \quad \square$$

Remark 2. In Proposition 2, if $r = n$ (i.e. it is the first player’s turn) we will obviously consider $M_0 = 0$. Therefore, the optimum threshold for the first player is θ_n .

The following consequence is obvious:

Corollary 2. For every $1 \leq r \leq n$, if there are r players still to play and all the previous ones got a score less than θ_r , then the optimum threshold for the next player (the one whose turn it is) is θ_r .

In what follows, we will assume that all the players follow their *optimal strategy* described in Proposition 2. Given $1 \leq r < n$, $1 \leq m \leq r$, and $x \geq \theta_r$, we shall denote by $F_r^m(x)$ the winning probability of the $(n - r + m)$ -th player when the maximum score of the first $n - r$ players is x .

Proposition 3. The following equalities hold:

$$F_r^m(x) = \begin{cases} e^x \int_x^1 (1 - e^t + te^t)^{r-1} dt & \text{if } m = 1, \\ (1 - e^x + xe^x)F_{r-1}^{m-1}(x) + e^x \int_x^1 F_{r-1}^{m-1}(t) dt & \text{if } m \geq 2. \end{cases}$$

Proof. For $m = 1$, in order for the $(n - r + 1)$ -th player to win, all the remaining $r - 1$ players must obtain a score of 0. We know that, if $\xi_x \neq 0$, then $\xi_x > x \geq \theta_r$. Consequently, due to Lemma 4 all the subsequent players will use ξ_x as their greed threshold because of Proposition 2. Thus, using Corollary 1 and (7), we get that

$$F_r^1(x) = \mathbb{E}[\mathbf{P}(\xi_x)^{r-1}] = e^x \int_x^1 (1 - e^t + te^t)^{r-1} dt.$$

Now, let us assume that $m > 1$ and denote by W_m the event “the $(n - r + m)$ -th player wins”. Since the $(n - r + 1)$ -th player used x as optimum threshold, his score is ξ_x . Then, the Law of total probability gives:

$$F_r^m(x) = P(W_m) = P(W_m | \xi_x = 0)P(\xi_x = 0) + P(W_m | \xi_x > 0)P(\xi_x > 0).$$

Furthermore, on one hand, it is straightforward that

$$P(W_m | \xi_x = 0) = F_{r-1}^{m-1}(x),$$

while on the other,

$$P(W_m | \xi_x > 0) = \mathbb{E}[(F_{r-1}^{m-1}(\xi_x) | \xi_x > 0)],$$

because, in this case, the maximum score obtained by the first $n - r + 1$ players is $\xi_x \geq \theta_{r-1}$. Combining both equalities and using Corollary 1 we obtain the desired formula:

$$F_r^m(x) = \mathbf{P}(x)F_{r-1}^{m-1}(x) + (1 - \mathbf{P}(x)) \int_x^1 \frac{F_{r-1}^{m-1}(t)}{1 - x} dt = (1 - e^x + xe^x)F_{r-1}^{m-1}(x) + e^x \int_x^1 F_{r-1}^{m-1}(t) dt. \quad \square$$

Table 1
Approximate values of θ_n and winning probabilities ($n \leq 10$) for each player.

n	2	3	4	5	6	7	8	9	10
θ_n	0.5706	0.6879	0.7487	0.7871	0.8141	0.8342	0.8499	0.8626	0.8730
P_n^1	0.4250	0.2859	0.2176	0.1764	0.1486	0.1285	0.1133	0.1013	0.0917
P_n^2	0.5750	0.3248	0.2357	0.1866	0.1551	0.1329	0.1165	0.1037	0.0936
P_n^3		0.3893	0.2570	0.1978	0.1619	0.1375	0.1197	0.1061	0.0954
P_n^4			0.2897	0.2104	0.1691	0.1422	0.1230	0.1085	0.0972
P_n^5				0.2289	0.1770	0.1470	0.1263	0.1109	0.0991
P_n^6					0.1883	0.1523	0.1297	0.1133	0.1008
P_n^7						0.1596	0.1333	0.1158	0.1026
P_n^8							0.1382	0.1184	0.1044
P_n^9								0.1218	0.1063
P_n^{10}									0.1088

With this proposition, we are in the condition to prove the main result of this section. Let us denote by P_n^m the winning probability of the m -player in Game I if there are n players. Then, we have the following.

Proposition 4. *The winning probability of player m in Game I, $1 \leq m \leq n$, using the optimal policy described in Proposition 2 is:*

$$P_n^m = \begin{cases} e^{\theta_n} (1 - e^{\theta_n} + \theta_n e^{\theta_n})^{n-1} & \text{if } m = 1, \\ (1 - e^{\theta_n} + \theta_n e^{\theta_n}) P_{n-1}^{m-1} + e^{\theta_n} \int_{\theta_n}^1 F_{n-1}^{m-1}(t) dt & \text{if } m \geq 2. \end{cases}$$

Proof. The proof is essentially the same as for Proposition 3, using ξ_{θ_n} as greed thresholds for each player, and expected values instead of distribution functions. \square

Remark 3. Assume none of the first $n - r$ players got a score higher than θ_r . Then the winning probability of player $n - r + 1$ is the same as the winning probability of the first player in the r -player game.

Table 1 shows the optimal greed threshold θ_n of the first player in an n -player game (or when there are still n players and none of the previous ones got a positive score), and the winning probabilities P_n^m of the m -th player in Game I, for $n = 1, \dots, 10$.

Remark 4. The threshold strategies described in Proposition 2 constitute a Nash equilibrium: no unilateral deviation will increase the winning probability of the defector. However, though optimal for each single player, each value P_n^r is the winning probability for player r only if every player plays optimally according to his interests without mistakes or collaboration (collusion) within some group. The deviation of a group of players from their individual optimum policy can increase the winning probability of some players and decrease that of others. To illustrate this, we now carry out a detailed study of the 3-player game, in which collusion between two players can modify the winning probabilities of the other one.

Notice, however, that the winning probability of each player using his optimal threshold θ_i (shown in Table 1) does not represent the value of the game for each, as their winning probabilities depend on the behavior of the other players. There exist coalitions, already in the 3-player game which decrease the winning probability of the remaining player, thus increasing the probability of one of the allies win. This is what we study in the next two subsections.

5.1. Coalition between first and second player

The first and second players can decrease the third one’s winning probability using some strategy. Let $\theta_2(x)$ be the second player’s optimal threshold conditional to the first one having obtained score x , assuming their common aim is to decrease the third player’s winning chance. Let $\mathbf{p}_3(x)$ be the third player’s losing probability assuming the first one has obtained score x and the second one uses threshold $\theta_2(x)$, which is:

$$\mathbf{p}_3(x) = P(\xi_{\theta_2(x)} = 0) \cdot P(\xi_x = 0) + P(\xi_{\theta_2(x)} > 0) \cdot \mathbb{E}(\mathbf{P}(\xi_{\theta_2(x)}) | \xi_{\theta_2(x)} > 0), \tag{10}$$

and by Corollary 1:

$$\mathbf{p}_3(x) = \mathbf{P}(\theta_2(x))\mathbf{P}(x) + e^{\theta_2(x)} \int_{\theta_2(x)}^1 \mathbf{P}(t) dt.$$

Assume the first player has score x and plays once more. The probability of the third player not winning under this assumption is given by:

$$x \cdot \mathbf{p}_3(0) + \int_x^1 \mathbf{p}_3(t)dt.$$

Thus, by Proposition 1, the first player’s optimal threshold Θ_1 , satisfies:

$$\mathbf{p}_3(\Theta_1) = \Theta_1 \mathbf{p}_3(0) + \int_{\Theta_1}^1 \mathbf{p}_3(t)dt,$$

and the probability of the third player not winning is by the same Proposition 1:

$$(\mathbf{p}_3(\Theta_1) - \mathbf{p}_3(0))e^{\Theta_1} + \mathbf{p}_3(0).$$

Computing the value of Θ_1 is very laborious, as one needs to calculate $\theta_2(x)$ and $\mathbf{p}_3(x)$ beforehand.

Calculation of $\theta_2(x)$: If the first player has stopped with score x and the second with $y > x$, the probability that the third one does not win is $\mathbf{P}(y)$. Applying Proposition 1, $\theta_2(x)$ is the solution of the equation:

$$\mathbf{P}(y) = y\mathbf{P}(x) + \int_y^1 \mathbf{P}(t)dt,$$

which together with (7), gives $\theta_2(x)$ implicitly:

$$-e^{\theta_2(x)}(2\theta_2(x) - 3) + \theta_2(x)e^x(x - 1) = e.$$

Foregoing the uninteresting details, we have obtained:

$$\Theta_1 = 0.63386\dots$$

so that the third player’s winning probability is:

$$1 - (\mathbf{p}_3(\Theta_1) - \mathbf{p}_3(0))e^{\Theta_1} - \mathbf{p}_3(0) = 0.3867\dots$$

slightly less than 0.3893 (P_3^3 in Table 1), which is the one under the Nash equilibrium.

5.2. Coalition between first and third player

There is also a possible coalition between the first and third players harmful to the second one. In this case, however, the third player’s only option is to try and improve on the second one’s score, which would become his playing threshold in this strategy. The second player’s optimal strategy is the one described in Proposition 2, and consists in using the threshold $\max(\theta_2, X_1)$, where X_1 is the first player’s score.

The second player’s winning probability assuming the first one’s score is $x \geq \theta_2$, is given by:

$$\hat{h}(x) = (1 - \mathbf{P}(x)) \int_x^1 \frac{\mathbf{P}(t)}{1 - x} dt = e^x \int_x^1 \mathbf{P}(t)dt = -e^x (e^x(x - 2) + x + e - 1).$$

The strategy under discussion aims to lower this value, which requires finding a new greed threshold for the first player.

If the first player has score x and plays once more, then the probability of the second player not winning is $x(1 - \vartheta) + \int_x^1 (1 - \hat{h}(t))dt$, where $\vartheta = 0.4250\dots$. In this formula, ϑ is the winning probability of the second player if he played just against the third one (what in Table 1 is P_2^1). Setting $h(t) := 1 - \hat{h}(t)$, the first player’s required greed threshold will be given by ϱ satisfying, by Proposition 1:

$$h(\varrho) = \varrho(1 - \vartheta) + \int_{\varrho}^1 h(t)dt,$$

whose value is, approximately,

$$\varrho = 0.75017\dots$$

and the second player’s winning probability assuming the first one uses this threshold is:

$$\mathbf{P}(\varrho)\vartheta + (1 - \mathbf{P}(\varrho)) \int_{\varrho}^1 \frac{\hat{h}(t)}{1 - \varrho} dt = 0.32262\dots$$

slightly less than $P_3^2 = 0.3248$ (Table 1), which is the second player’s winning probability in the Nash equilibrium.

6. Game II — no-information game

In this version, the n players act simultaneously with no information on the results of the others. Unlike the studies [3,9], we consider the no-winner possibility, that is, there is a global tie when all the players get a score of 0. We study in this section the varieties introduced in Section 2: For Games II.1 and II.2, if all the players get a score of 0, there is no payoff, while Game II.3 is an asymmetric version in which a single player has advantage: he wins the payoff in case of global tie at 0. Throughout the section the number of players is denoted, as above, by n .

6.1. Game II.1: non-constant sum

In this version, if there is a winner, he receives a payoff of 1 from an external agent. Hence, each player’s expected payout is his probability of winning.

Lemma 5. For all $n > 0$, the equation

$$(1 + e^x(x - 1))^{n-1} = \frac{1 - (1 + e^x(x - 1))^n}{ne^x}, \tag{11}$$

has a single solution $\alpha_n \in (0, 1)$.

Proof. Put $f_n(x) = (1 + e^x(x - 1))^{n-1}$ and $g_n(x) = \frac{1 - (1 + e^x(x - 1))^n}{ne^x}$. If $n = 1$, then $f_1(x) = 1$ and $g_1(x) = 1 - x$ so $x = 0$ is the only solution. Now, for $n > 1$, we have that $f_n(0) = 0$, $f_n(1) = 1$, $g_n(0) = 1/n > 0$ and $g_n(1) = 0$. Since f_n is increasing and g_n is decreasing in $[0, 1]$ for every $n > 1$, the result follows. \square

Lemma 6. The sequence $\{\alpha_n\}_{n>0}$ from Lemma 5 is strictly increasing.

Proof. With the previous notation, it is enough to show that $f_{n+1}(\alpha_n) < g_{n+1}(\alpha_n)$. First note that $f_{n+1}(\alpha_n) = (1 + e^{\alpha_n}(\alpha_n - 1)) f_n(\alpha_n)$ and by definition of α_n (Lemma 5),

$$f_{n+1}(\alpha_n) = (1 + e^{\alpha_n}(\alpha_n - 1)) g_n(\alpha_n).$$

Then, after some straightforward computations, we see that $f_{n+1}(\alpha_n) < g_{n+1}(\alpha_n)$ if and only if

$$(n + 1) (1 + e^{\alpha_n}(\alpha_n - 1)) - (1 + e^{\alpha_n}(\alpha_n - 1))^{n+1} < n.$$

Now, for every $t \in (0, 1)$, we have that $(n + 1)t - t^{n+1} < n$ because the left-hand side is an increasing function in t , whose value for $t = 1$ is n , and we obtain the result. \square

Proposition 5. Game II.1 admits a Nash equilibrium with equal thresholds α_n for all players. Moreover, the winning probability of each player in that Nash equilibrium is

$$P_n := \frac{1 - (1 - e^{\alpha_n} + \alpha_n e^{\alpha_n})^n}{n}. \tag{12}$$

For this game, we shall call Nash threshold the value α_n .

Proof. Assuming that all rivals use thresholds α_n , we will call $\xi_{\alpha_n}^{(i)}$ the random variable that represents the score obtained by the i -th player using the threshold α_n and $h(x)$ the probability of winning for a player when stopping with score x . We will assume, without loss of generality, that the reference player is the first and that the rest use a greed threshold α_n . It is about proving that, under these conditions, the optimal threshold for the first player is precisely α_n .

- If $x \geq \alpha_n$, then, by Lemma 2:

$$h(x) = P(x > \xi_{\alpha_n}^{(i)} : i = 2, \dots, n) = \prod_{i=2}^n P(x > \xi_{\alpha_n}^{(i)}) = (1 + e^{\alpha_n} (x - 1))^{n-1}.$$

- If $x < \alpha_n$, then

$$h(x) = P(x > \xi_{\alpha_n}^{(i)} : i = 2, \dots, n) = \prod_{i=2}^n P(\xi_{\alpha_n}^{(i)} = 0) = \mathbf{P}(\alpha_n)^{n-1}.$$

In any case, $h(x)$ is continuous and non-decreasing in $[0, 1]$ so that by Proposition 1 it is enough to verify that

Table 2
Approximate Nash threshold and winning probability for Game II.1.

n	2	3	4	5	6	7	8	9	10
α_n	0.5887	0.6989	0.7562	0.7927	0.8184	0.8377	0.8528	0.8650	0.8751
P_n	0.4665	0.3129	0.2366	0.1907	0.1598	0.1375	0.1208	0.1077	0.0972

$$h(\alpha_n) = \int_{\alpha_n}^1 h(t)dt.$$

We have, by definition of α_n , that

$$(1 + e^{\alpha_n} (\alpha_n - 1))^{n-1} = \frac{1 - (e^{\alpha_n}(\alpha_n - 1) + 1)^n}{ne^{\alpha_n}}.$$

And finally, one just needs to keep in mind that

$$(1 + e^{\alpha_n} (\alpha_n - 1))^{n-1} = h(\alpha_n),$$

$$\frac{1 - (e^{\alpha_n}(\alpha_n - 1) + 1)^n}{ne^{\alpha_n}} = \int_{\alpha_n}^1 h(t)dt.$$

In this equilibrium, since all players have the same probability of winning, the winning probability of each player is:

$$P_n = \frac{1 - P(\xi_{\alpha_n} = 0)^n}{n} = \frac{1 - (1 - e^{\alpha_n} + \alpha_n e^{\alpha_n})^n}{n}. \quad \square$$

Table 2 contains the Nash thresholds and the winning probabilities for each player in this version of the n -player game (no-information and external payer). A simple example follows for $n = 2$.

Example 3. For $n = 2$, the Nash threshold is 0.5887+, and $P_2 = 0.4665+$. Notice, however, that the use of this threshold does not guarantee a winning probability of 0.4665: the other player abandoning the equilibrium strategy can be harmful for *both of them*, as we will now show.

When the players use respective thresholds x and y , then the winning probability of the first player, $P_1(x, y) = P(\xi_x > \xi_y)$, is:

$$P_1(x, y) = P(\xi_y = 0)P(\xi_x > 0) + P(\xi_y > 0)P(\xi_x > 0)P(\xi_x > \xi_y | \xi_x > 0 \wedge \xi_y > 0),$$

that is,

$$P_1(x, y) = \mathbf{P}(y)(1 - \mathbf{P}(x)) + (1 - \mathbf{P}(y))(1 - \mathbf{P}(x))P(\xi_x > \xi_y | \xi_x > 0 \wedge \xi_y > 0)$$

which, recalling (7), gives:

If $x \leq y$, then

$$P(\xi_x > \xi_y | \xi_x > 0 \wedge \xi_y > 0) = P(\mathbf{U}[x, 1] > \mathbf{U}[y, 1]) = \frac{1 - y}{2(1 - x)}$$

so that:

$$P_1(x, y) = \frac{1}{2}e^x (e^y(y - 1)(-2x + y + 1) - 2x + 2)$$

Otherwise, if $x > y$, then

$$P(\xi_x > \xi_y | \xi_x > 0 \wedge \xi_y > 0) = P(\mathbf{U}[x, 1] > \mathbf{U}[y, 1]) = \frac{x - y}{1 - y} + \frac{1 - x}{2(1 - y)}$$

so that:

$$P_1(x, y) = -\frac{1}{2}e^x(x - 1)((x - 1)e^y + 2)$$

In short:

$$P_1(x, y) = \begin{cases} \frac{1}{2}e^x (e^y(y - 1)(-2x + y + 1) - 2x + 2), & \text{if } x \leq y \\ -\frac{1}{2}e^x(x - 1)((x - 1)e^y + 2), & \text{otherwise.} \end{cases} \tag{13}$$

If the first player chooses threshold $\alpha_2 = 0.5887+$, then it is easy to verify that, for $y \in (\alpha_2, 0.745+)$,

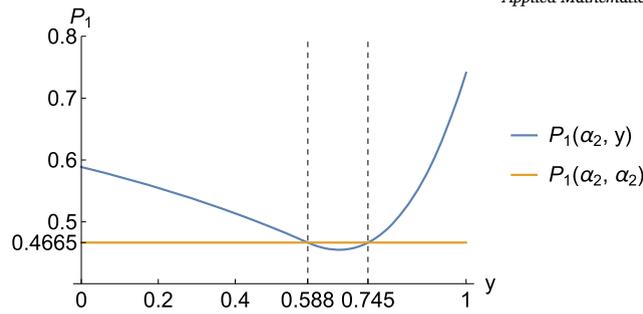


Fig. 1. Winning probability of Player 1 (blue) depending on the threshold of Player 2 (y). In orange: winning probability of each player in the Nash equilibrium.

Table 3
Approximate optimal Nash threshold in the Drawless version of Game II.1.

n	2	3	4	5	6	7	8	9	10
u^*	0.633	0.718	0.767	0.800	0.823	0.841	0.856	0.867	0.877

$$P_1(\alpha_2, y) < P_1(\alpha_2, \alpha_2). \tag{14}$$

Which means that the first player’s expected payoff can decrease depending on the greed threshold chosen by the second one (see Fig. 1).

Remark 5. This equilibrium is not a strong equilibrium: the trivial all-player coalition using the same greed threshold 0 provides an expected payoff of $1/n$, which is greater than the one of Proposition 5. However, this is an unstable coalition as each player would be incentivized to abandon it unilaterally in order to obtain a payoff greater than $1/n$. We believe (but have no proof) that the Nash equilibrium of Proposition 5 is coalition-proof, and only binding collaboration agreements can ensure a greater payoff. This Nash equilibrium is not “foolproof” either, in the sense that if some players abandon that strategy (knowingly or by mistake), they may decrease not only their expected payoff but also that of other players, as Example 3 above shows.

In Table 3 we copy the row corresponding to $u^{(inf,n,N)}$ from [9, Table 1] for the no-information game with n players and unlimited spins, but without ties (in that version, if all the scores are above 1, then the least one is the winner), that is a Drawless Game II.1. If we compare it with Table 2, we can observe how the thresholds in this variant are noticeable greater than those of the standard Game II.1 for 2 and 3 players, while they decrease towards the same values as n increases.

6.2. Game II.2: zero-sum game

We still consider a no-information variant with 0 payoff if all the players get score 0. However, when there is a winner, he is paid $1/(n - 1)$ by each player (so that his total payoff is 1). In this variant we will assume that $n > 1$.

Lemma 7. For $n \in (1, \infty)$, the equation

$$(1 + e^x(x - 1))^{n-1} = \frac{1}{1 + e^x(n - 1)}, \tag{15}$$

has a single solution $\gamma_n \in (0, 1)$.

Proof. Put $f_n(x) = (1 + e^x(x - 1))^{n-1}$ and $g_n(x) = \frac{1}{1 + e^x(n - 1)}$. Then, we have that $f_n(0) = 0$, $f_n(1) = 1$, $g_n(0) = 1/n > 0$ and $g_n(1) = \frac{1}{1 + e(n-1)} < 1$. Since $f_n(x)$ is increasing and $g_n(x)$ is decreasing in $[0, 1]$, the result follows. \square

Lemma 8. The sequence $\{\gamma_n\}_{n>1}$, for $n \in \mathbb{N}$, is strictly increasing.

Proof. Write

$$F(x, n) = f_n(x) - g_n(x) = (e^x(x - 1) + 1)^{n-1} - \frac{1}{(n - 1)e^x + 1}.$$

Now, we have

$$\frac{\partial F}{\partial x} = (n - 1)e^x \left(x(e^x(x - 1) + 1)^{n-2} + \frac{1}{((n - 1)e^x + 1)^2} \right)$$

which is positive for all real numbers $x \in (0, 1)$ and $n \in (1, \infty)$. Thus, the curve $F(x, n) = 0$ can be parametrized as $(x(n), n)$ for $x \in (0, 1)$ and $n \in (1, \infty)$. Note that, in particular, $\gamma_n = x(n)$ for every integer $n > 1$.

On the other hand, we also have

$$\frac{\partial F}{\partial n} = \frac{e^x}{((n-1)e^x + 1)^2} + (e^x(x-1) + 1)^{n-1} \log(e^x(x-1) + 1).$$

and, since $F(x, n) = 0$ if and only if $(e^x(x-1) + 1)^{n-1} = \frac{1}{(n-1)e^x + 1}$, it follows that

$$\left. \frac{\partial F}{\partial n} \right|_{F(x,n)=0} = \frac{e^x - \frac{((n-1)e^x + 1) \log((n-1)e^x + 1)}{n-1}}{((n-1)e^x + 1)^2},$$

which is negative for all real numbers $x \in (0, 1)$ and $n \in (1, \infty)$. Thus, the implicit function theorem yields that

$$x'(n) = \frac{-F_n(x(n), n)}{F_x(x(n), n)} > 0$$

which implies that the sequence $(\gamma_n)_{n>1}$ for $n \in \mathbb{N}$ is strictly increasing, as claimed. \square

Proposition 6. *Game II.2 admits a Nash equilibrium with identical strategies γ_n for all players. Moreover, the tying probability is*

$$\widehat{P}_n = (1 - e^{\gamma_n} + \gamma_n e^{\gamma_n})^n,$$

and the winning probability of each player is

$$P_n = \frac{1 - (1 - e^{\gamma_n} + \gamma_n e^{\gamma_n})^n}{n}.$$

The expected payoff of each player is obviously 0.

Proof. Assuming that all rivals use thresholds γ_n , we will call $\xi_{\gamma_n}^{(i)}$ the random variable that represents the score obtained by the i -th player using the threshold γ_n and $h(x)$ the probability of winning for a player when stopping with score x . We will assume without loss of generality that the reference player is the first and that the rest use a greed threshold γ_n . It is about proving that, under these conditions, the optimal threshold for the first player is precisely γ_n .

- If $x \geq \gamma_n$ then the first player wins with probability $\prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} < x)$ and the probability of a tie is zero. Therefore

$$h(x) = \prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} < x) - \frac{1 - \prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} < x)}{n-1} = \frac{n}{n-1} (1 + e^{\gamma_n} (x-1))^{n-1} - \frac{1}{n-1}. \tag{16}$$

- If $x < \gamma_n$ then the first player wins when the remaining players finish with a score of 0 (all exceed 1). Therefore, he wins with probability $\prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} = 0)$ and the probability of a tie is zero, so

$$h(x) = \prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} = 0) - \frac{1 - \prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} = 0)}{n-1} = \frac{n}{n-1} \prod_{i=2}^n P(\xi_{\gamma_n}^{(i)} = 0) - \frac{1}{n-1} = \frac{n}{n-1} \mathbf{P}(\gamma_n)^{n-1} - \frac{1}{n-1},$$

and

$$h(0) = -\frac{1 - \mathbf{P}(\gamma_n)^{n-1}}{n-1}.$$

Furthermore, $h(x)$ is non-decreasing in $[0, 1]$ so that, by Proposition 1, it is enough to verify that

$$h(\gamma_n) = -\frac{\gamma_n}{n-1} (1 - \mathbf{P}(\gamma_n)^{n-1}) + \int_{\gamma_n}^1 h(t) dt. \tag{17}$$

From (16), we have

$$h(\gamma_n) = \frac{n}{n-1} (1 + e^{\gamma_n} (\gamma_n - 1))^{n-1} - \frac{1}{n-1},$$

and

$$\int_{\gamma_n}^1 h(t) dt = \frac{-e^{-\gamma_n} ((e^{\gamma_n} (\gamma_n - 1) + 1)^n - 1) + \gamma_n - 1}{n-1}.$$

By the properties of γ_n (Lemma 7), we have

Table 4
Approximate optimal threshold, tie \hat{P}_n and winning P_n probabilities for Game II.2.

n	2	3	4	5	6	7	8	9	10
γ_n	0.6591	0.7305	0.7744	0.8046	0.8268	0.8440	0.8576	0.8689	0.8783
\hat{P}_n	0.1163	0.0855	0.0680	0.0566	0.0486	0.0426	0.0379	0.0342	0.0312
P_n	0.4419	0.3048	0.2330	0.1887	0.1586	0.1368	0.1203	0.1073	0.0969

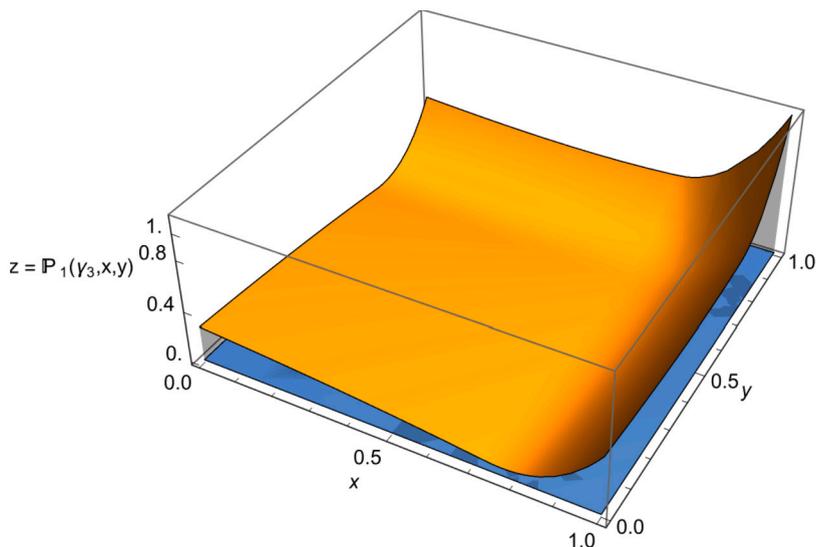


Fig. 2. Payoff $z = P_1(\gamma_3, x, y)$. In blue, $z = 0$: notice how $P_1(\gamma_3, x, y) \geq 0$ everywhere.

$$(1 + e^{\gamma_n}(\gamma_n - 1))^{n-1} = \frac{1}{1 + e^{\gamma_n}(n - 1)},$$

and equality (17) holds.

Finally, the tying probability \hat{P}_n is obviously the probability of all the players getting a score of 0 when using γ_n as their greed threshold; i.e., $\hat{P}_n = P(\gamma_n)^n$. Consequently, the winning probability of each player is given by

$$P_n = \frac{1 - P(\gamma_n)^n}{n},$$

and Lemma 2 gives then

$$P_n = \frac{1 - (1 - e^{\gamma_n} + \gamma_n e^{\gamma_n})^n}{n},$$

as desired. \square

Table 4 shows γ_n and the tying and winning probabilities in Game II.2 for $2 \leq n \leq 10$.

Remark 6. The Nash equilibrium found in Proposition 6 is the only possible one with identical strategies for all players. It is easy to see that for $n = 2$ it is the only possible equilibrium (symmetric or not), and that the strategies are maximin (because the game is zero-sum). For $n > 2$ we are not sure but it seems reasonable to think that there are no other (non-symmetric) Nash equilibria, and neither do there exist any collusive coalitions yielding negative payoffs to the players out of it. Thus, we conjecture that it is both a strong and a coalition-proof Nash equilibrium. In Fig. 2 we plot the expected payoff of the first player $P_1(u_1, u_2, u_3)$ in the 3-player game when they use respective greed thresholds u_1, u_2, u_3 : if the first one sets $u_1 = \gamma_3$, no combined strategy of the other two yields him a negative payoff.

Remark 7. Unlike in Game II.1, the equilibrium strategy is not to maximize the winning probability but to have a score at least the same as the rest of the players. As a matter of fact, if all players used the Nash thresholds of Game II.1, they would have a greater winning probability than with the γ_n of Proposition 6. However, that would not provide a Nash equilibrium in Game II.2 because all players would be better off abandoning that strategy in order to have a greater probability of success than the others, despite their own probability of winning being less.

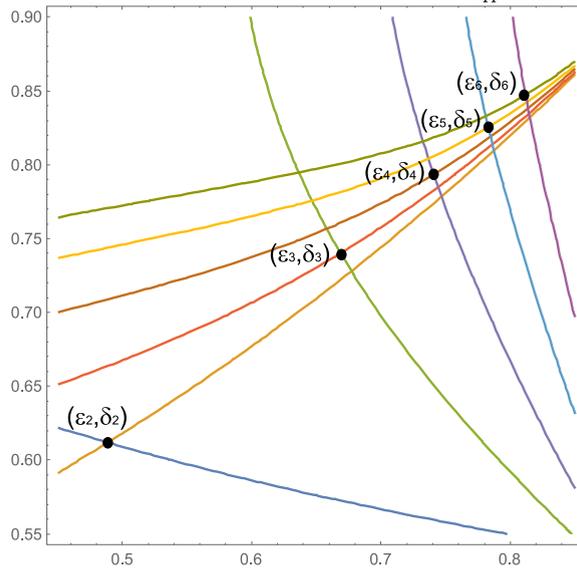


Fig. 3. Solutions of Equation (18) for $n = 2, 3, \dots, 6$. The decreasing plots correspond to a) in (18), while the increasing ones to b).

Remark 8. If instead of having the losing players pay the winner, we consider an external payer, and that if a tie happens, the game is to be repeated until it is broken, then the Nash equilibrium would be the same and the expected payoff would be $1/n$. This is because this variant is a constant-sum game. It is remarkable that the equilibrium strategies are different when there is no payoff in case of tie, and when, in this case, the game is repeated until the tie is broken.

6.3. Game II.3: non-symmetric and constant-sum

We now turn our attention to the no-information, non-symmetric version of Showcase Showdown with $n > 1$ players in which one has advantage in the sense that if there is a tie in scores (i.e., if all players get a score of 0) then he is the winner. In this case the payer being external or the players is irrelevant, as this is a constant-sum game and the only relevant point is the winning probability.

Lemma 9. For all $n > 1$, the system of equations

$$\begin{cases} \text{a) } (1 + e^x(x - 1))^{n-2} = \frac{e^y (-1 + (1 + e^x (y - 1))^n) + n e^x}{n e^x (1 + e^y (y - 1)) (1 + e^x (n + x - 2))}, \\ \text{b) } (1 + e^x(x - 1))^{n-1} = \frac{e^{-x} (n e^x (e^x(y - 1) + 1)^{n-1} + (e^x(y - 1) + 1)^n - 1)}{n y}. \end{cases} \tag{18}$$

has a single solution in $[0, 1] \times [0, 1]$ which we will denote from now on (ϵ_n, δ_n) with $\delta_n > \epsilon_n$.

Proof. It can be seen, using implicit differentiation and some computational effort, that the first equation defines y as a decreasing function of x , while the second equation defines y as an increasing function of x . Then, it is enough to compare the values at $x = 0$ and $x = 1$ to conclude the proof. \square

Fig. 3 shows the situation of Lemma 9 for $2 \leq n \leq 6$. The curves represent the points (x, y) satisfying either a) or b) in Lemma 9, for $n = 2, \dots, 6$, so that their respective intersections are the points (ϵ_n, δ_n) . The figure clearly suggests that both sequences $\{\epsilon_n\}_{n>1}$ and $\{\delta_n\}_{n>1}$ are strictly increasing. Since this fact is not required in the sequel, and we have not been able to find a direct proof for every $n > 1$, this remains a conjecture.

Conjecture 1. The sequences $\{\epsilon_n\}_{n>1}$ and $\{\delta_n\}_{n>1}$ of solutions of the system of equations (18) introduced in Lemma 9 are strictly increasing.

Proposition 7. Game II.3 admits a Nash equilibrium with greed threshold δ_n for the player with advantage and identical greed thresholds ϵ_n for the remaining players. Moreover, the winning probability for the player with advantage is

$$P_n^A = (1 - e^{\epsilon_n} + \epsilon_n e^{\epsilon_n})^{n-1} (1 - e^{\delta_n} + \delta_n e^{\delta_n}) + \frac{e^{\delta_n} (1 - (1 + e^{\epsilon_n} (\delta_n - 1))^n)}{e^{\epsilon_n} n}. \tag{19}$$

Proof. Without loss of generality we assume that the player with advantage is the n -th one. We will show that the following two statements **A)** and **B)** are verified:

A) If the n -th player uses the threshold δ_n and the rest except one of them (which we will assume without loss of generality is the first player) use ϵ_n as a threshold, then the optimal threshold of the first is also ϵ_n .

Thus, assume that player n uses threshold δ_n and that players $2, \dots, n-1$ use the identical threshold ϵ_n . Assuming that the rivals use the thresholds already mentioned, we will call $\xi_{\epsilon_n}^{(i)}$, $2 \leq i \leq n-1$, the random variable that represents the score obtained by player i using the threshold ϵ_n , $\xi_{\delta_n}^{(n)}$ the random variable that represents the score obtained by player n using the threshold δ_n and let $h(x)$ be the probability of winning for the first player when stopping with score x . It is about proving that under these conditions the optimal threshold for the first player is precisely ϵ_n .

- For all $x \in [\epsilon_n, \delta_n]$,

$$h(x) = P(\xi_{\delta_n}^{(n)} = 0) \prod_{i=2}^{n-1} P(\xi_{\epsilon_n}^{(i)} < x) = (1 + e^{\delta_n} (\delta_n - 1)) (1 + e^{\epsilon_n} (x - 1))^{n-2}.$$

- For $x < \epsilon_n$,

$$h(x) = P(\xi_{\delta_n}^{(n)} = 0) \prod_{i=2}^{n-1} P(\xi_{\epsilon_n}^{(i)} = 0) = (1 + e^{\delta_n} (\delta_n - 1)) (1 + e^{\epsilon_n} (\epsilon_n - 1))^{n-2}.$$

- For $x > \delta_n$,

$$h(x) = P(\xi_{\delta_n}^{(n)} < x) \prod_{i=2}^{n-1} P(\xi_{\epsilon_n}^{(i)} < x) = (1 + e^{\delta_n} (x - 1)) (1 + e^{\epsilon_n} (x - 1))^{n-2}.$$

We have that $h(x)$ is continuous and non-decreasing in $[0, 1]$ so to be in the conditions of Proposition 1 we need to see that

$$h(\epsilon_n) = \int_{\epsilon_n}^1 h(t) dt.$$

To see this, we take into account that

$$h(\epsilon_n) = (1 + e^{\delta_n} (\delta_n - 1)) (1 + e^{\epsilon_n} (\epsilon_n - 1))^{n-2}.$$

On the other hand,

$$\begin{aligned} \int_{\epsilon_n}^1 h(t) dt &= \int_{\epsilon_n}^{\delta_n} h(t) dt + \int_{\delta_n}^1 h(t) dt = \int_{\epsilon_n}^{\delta_n} (1 + e^{\delta_n} (\delta_n - 1)) (1 + e^{\epsilon_n} (t - 1))^{n-2} dt + \\ &\int_{\delta_n}^1 (1 + e^{\delta_n} (t - 1)) (1 + e^{\epsilon_n} (t - 1))^{n-2} dt = \\ &\frac{e^{-\epsilon_n} (1 + e^{\delta_n} (\delta_n - 1)) \left[(1 + e^{\epsilon_n} (\delta_n - 1))^{n-1} - (1 + e^{\epsilon_n} (\epsilon_n - 1))^{n-1} \right]}{n - 1} + \\ &\frac{e^{-2\epsilon_n} \left[-e^{\delta_n} + e^{\epsilon_n} n + (1 + e^{\epsilon_n} (\delta_n - 1))^{n-1} (e^{\delta_n} - e^{\epsilon_n} (e^{\delta_n} (\delta_n - 1) (n - 1) + n)) \right]}{n(n - 1)}. \end{aligned}$$

Taking into account now the definition of ϵ_n and δ_n in Lemma 9, the equality $h(\epsilon_n) = \int_{\epsilon_n}^1 h(t) dt$ holds.

B) If all but the n -th player use the threshold ϵ_n then the optimal threshold for the last one is δ_n .

Assuming that all the players (except the last one) use thresholds ϵ_n , we will consider the random variable $\xi_{\epsilon_n}^{(i)}$, $1 \leq i < n$, which represents the score obtained by the i -th player using the threshold ϵ_n and $h(y)$ to be the probability of winning of the n -th player by stopping with score y . It is about proving that under these conditions the optimal threshold for the n -th player is precisely δ_n .

- If $y \geq \epsilon_n$, then we have that the n -th player wins with probability

$$h(y) = \prod_{i=1}^{n-1} P(\xi_{\epsilon_n}^{(i)} < y) = (1 + e^{\epsilon_n} (y - 1))^{n-1}. \tag{20}$$

- If $y < \epsilon_n$, then the n -th player wins if the remaining players finish with score 0 (all exceed 1)

$$h(y) = \prod_{i=1}^{n-1} P(\xi_{\epsilon_n}^{(i)} = 0) = (1 + e^{\epsilon_n} (\epsilon_n - 1))^{n-1}.$$

Table 5
Nash thresholds (ϵ_n, δ_n) , and winning probability P_n^A of the player with advantage, and winning probability P_n^N of the other players in Game II.3.

n	2	3	4	5	6	7	8	9	10
ϵ_n	0.4887	0.6687	0.7408	0.7832	0.8119	0.8329	0.8491	0.8621	0.8728
δ_n	0.6118	0.7401	0.7936	0.8256	0.8475	0.8637	0.8763	0.8865	0.8948
P_n^A	0.5366	0.3720	0.2879	0.2357	0.1998	0.1736	0.1535	0.1377	0.1249
P_n^N	0.4634	0.3140	0.2374	0.1911	0.1600	0.1378	0.1209	0.1078	0.0972

We have that $h(y)$ is continuous and non-decreasing in $[0, 1]$ so that by Proposition 1 it is enough to verify that

$$h(\delta_n) = \delta_n \mathbf{P}(\epsilon_n)^{n-1} + \int_{\delta_n}^1 h(t) dt.$$

To do this, using (20) for $y \geq \epsilon_n$, we get:

$$\int_y^1 h(t) dt + y \mathbf{P}(\epsilon_n)^{n-1} = -\frac{e^{-\epsilon_n} (-1 + (1 + e^{\epsilon_n} (y - 1))^n)}{n} + y (1 + e^{\epsilon_n} (\epsilon_n - 1))^{n-1}.$$

Performing simple calculations, taking into account the second equation in (18) in Lemma 9 with $x = \epsilon_n$, it follows that

$$h(\delta_n) = \delta_n \mathbf{P}(\epsilon_n)^{n-1} + \int_{\delta_n}^1 h(t) dt.$$

Now, the winning probability of the advantageous player is the sum of the probabilities of the following complementary events:

- a) All the players have a score of 0. This happens with probability

$$\mathbf{P}(\epsilon_n)^{n-1} \mathbf{P}(\delta_n) = (1 + e^{\epsilon_n} (\epsilon_n - 1))^{n-1} (1 + e^{\delta_n} (\delta_n - 1)).$$

- b) The n -th player has a score of $t \in [\delta_n, 1]$ and the rest of them (using threshold ϵ_n) do not exceed this score. This happens with probability

$$e^{\delta_n} (1 - \delta_n) \int_{\delta_n}^1 \frac{(P(\xi_{\epsilon_n} < t))^{n-1}}{1 - \delta_n} dt = e^{\delta_n} \frac{1 - (1 + e^{\epsilon_n} (\delta_n - 1))^n}{n e^{\epsilon_n}}.$$

Then, P_n^A is just the sum of the previous probabilities and the result follows. \square

Table 5 shows the Nash thresholds α_n for the “normal” players and β_n for the one with advantage, and the winning probability P_n^A of the latter in Game II.3. Note that the winning probability of the remaining players is just $P_n^N = \frac{1 - P_n^A}{n - 1}$.

Remark 9. The above Nash equilibrium is the only possible one with identical thresholds for the disadvantageous players. And, as in the previous versions, it seems reasonable that there are no more Nash equilibria. This equilibrium seems also (as in the case of Game II.2) a strong and coalition-proof Nash equilibrium. There is no possibility of collusion in this game, either.

7. Conclusions and future perspectives

Several versions of the Showcase Showdown game with unlimited number of spins have been studied, including cases with possibility of draw. For the sequential n -player game, optimal thresholds for each player, and their winning probabilities have been computed, thus improving on the results in [9], which apply only to the first player. We have also discovered that there is a possibility of coalitions that decrease the probability of winning for certain players, so that the Nash equilibrium does not guarantee the winning probabilities computed.

We have studied three versions of the game in which the players have no information on the score of the others. Despite the underlying game being the same, the Nash equilibria vary considerably, especially for games with few players (see, for instance, cases $n = 2$ and $n = 3$ in Table 6). The greedest strategies (those with greater threshold) seem to be those of the player with advantage in the asymmetric game (except when $n = 2$, where the greedest threshold happens in the zero-sum game). The equilibrium strategies with lesser threshold seem to happen for all n for the players without advantage, in the asymmetric game. See Table 6 for a summary

Table 6
Nash thresholds for the Non-zero-sum game (α_n), the Zero-sum game (γ_n), and the Asymmetric game (ϵ_n, δ_n).

n	2	3	4	5	6	7	8	9	10
α_n	0.5887	0.6989	0.7562	0.7927	0.8184	0.8377	0.8528	0.8650	0.8751
γ_n	0.6591	0.7305	0.7744	0.8046	0.8268	0.8440	0.8576	0.8689	0.8783
ϵ_n	0.4887	0.6687	0.7408	0.7832	0.8119	0.8329	0.8491	0.8621	0.8728
δ_n	0.6118	0.7401	0.7936	0.8256	0.8475	0.8637	0.8763	0.8865	0.8948

of these results. We have also established that the equilibrium thresholds are increasing in n in the Non-zero-sum and in the Zero-sum games (we conjecture that this also happens in the Asymmetric one).

Some open questions remain for the case $n > 2$: it seems likely that the symmetric Nash equilibria found in the no-information games are unique, due to the inherent symmetry in the game; furthermore, in the constant-sum cases, it seems reasonable to think that the computed strategies are optimal in the sense that they guarantee at least the expected payoff when all the players use their Nash thresholds. But we have not been able to tackle these problems yet.

We suggest the following future research:

1. To consider a maximum value for the scores (accumulated sums) $M > 1$. It may be very interesting to study the limits of the optimal thresholds and winning probabilities in the sequential game of each player when M tends to infinity.
2. Let the payoffs be a function of the score. This would intertwine the struggle to win and the aim of getting as great a payoff as possible.
3. To choose other underlying random variables different from $U[0, 1]$. For instance, the exponential distribution looks promising.
4. A version of the sequential game with more information might be as follows: instead of each player making all his spins in his (single) turn, he spins once every turn (or stops), and then the next player gets the turn iteratively, in the same order (unless he has stopped, in which case he does not play any more). The game will end when all the players have stopped. Players who get a score greater than 1 are eliminated. The game stops when all have been eliminated or stopped, and the winner is the one with the highest score.

Data availability

No data was used for the research described in the article.

References

- [1] P.R. Coe, W. Butterworth, Optimal stopping in “the showcase showdown”, *Am. Stat.* 49 (3) (1995) 271–275.
- [2] B. Kaynar, Optimal stopping in a stochastic game, *Probab. Eng. Inf. Sci.* 23 (1) (2009) 51–60.
- [3] V.V. Mazalov, A.A. Ivashko, Equilibrium in n -person game of showcase showdown, *Probab. Eng. Inf. Sci.* 24 (3) (2010) 397–403.
- [4] M. Sakaguchi, Two-player games of “score showdown”, *Sci. Math. Jpn.* 60 (2) (2004) 411–421.
- [5] M. Sakaguchi, Equilibrium in the three-player game of “risky exchange”, *Sci. Math. Jpn.* 61 (2) (2005) 335–344.
- [6] M. Sakaguchi, Equilibrium in two-player games of showcase showdown, *Sci. Math. Jpn.* 61 (1) (2005) 145–151.
- [7] M. Sakaguchi, Players’ information in two-player games of “score showdown”, *Int. J. Math. Game Theory Algebr.* 15 (6) (2006) 657–669, with errata on floppy disk.
- [8] M. Sakaguchi, Three-player game of ‘keep-or-exchange’, *Int. J. Math. Game Theory Algebr.* 15 (3) (2006) 275–287.
- [9] T.V. Seregina, A.A. Ivashko, V.V. Mazalov, Optimal stopping strategies in the game “the price is right”, *Proc. Steklov Inst. Math.* 307 (Suppl 1) (2019) 127–141.
- [10] D. Swenson, Optimal strategy in “the price is right” showcase showdown: a module for students of calculus and probability, *PRIMUS* 25 (7) (2015) 578–595.
- [11] R. Tenorio, T.N. Cason, To spin or not to spin? Natural and laboratory experiments from the price is right, *Econ. J.* 112 (476) (2002) 170–195.