

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

Strategic disclosure of random variables

János Flesch, Andrés Perea*

Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands

ARTICLE INFO

Article history:

Received 2 June 2009

Revised 7 June 2010

Accepted 24 August 2010

Available online 19 September 2010

Keywords:

Secretary problem

Moser's problem

Incomplete information

Lack of information on one side

Optimal strategies

ABSTRACT

We consider a game G_n played by two players. There are n independent random variables Z_1, \dots, Z_n , each of which is uniformly distributed on $[0,1]$. Both players know n , the independence and the distribution of these random variables, but only player 1 knows the vector of realizations $\mathbf{z} := (z_1, \dots, z_n)$ of them. Player 1 begins by choosing an order z_{k_1}, \dots, z_{k_n} of the realizations. Player 2, who does not know the realizations, faces a stopping problem. At period 1, player 2 learns z_{k_1} . If player 2 accepts, then player 1 pays z_{k_1} euros to player 2 and play ends. Otherwise, if player 2 rejects, play continues similarly at period 2 with player 1 offering z_{k_2} euros to player 2. Play continues until player 2 accepts an offer. If player 2 has rejected $n - 1$ times, player 2 has to accept the last offer at period n . This model extends Moser's (1956) problem, which assumes a non-strategic player 1.

We examine different types of strategies for the players and determine their guarantee-levels. Although we do not find the exact max–min and min–max values of the game G_n in general, we provide an interval $I_n = [a_n, b_n]$ containing these such that the length of I_n is at most 0.07 and converges to 0 as n tends to infinity. We also point out strategies, with a relatively simple structure, which guarantee that player 1 has to pay at most b_n and player 2 receives at least a_n . In addition, we completely solve the special case G_2 where there are only two random variables. We mention a number of intriguing open questions and conjectures, which may initiate further research on this subject.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

For many years, scientists from different disciplines have explored the well-known “secretary problem”. This is a stopping problem in which n secretaries are invited, in a random order, for an interview to fill a secretarial position. The employer knows the number of secretaries, and is aware that the order is random. After every interview, the employer can rank the secretaries interviewed so far from best to worst without ties, and must decide whether or not to hire the last candidate. His task is to find a stopping rule that maximizes the probability of hiring the best secretary. The optimal stopping rule has the following form: Reject the first r_n secretaries, and then hire the first secretary who is better than all the preceding ones. If no such secretary arrives after round r_n , then the best candidate was among the first r_n secretaries, and it therefore does not make a difference whether or not to hire the last secretary. For large n , the optimal choice of r_n is approximately n/e , and the probability of hiring the best secretary is approximately $1/e$. For a historical overview of this classical secretary problem the reader is referred to Ferguson (1989).

The secretary problem has been extended in a number of important directions. In particular, versions have been studied in which the payoff depends on the rank of the selected candidate, even if he or she is not the best. This seems more realistic than the classical scenario, as hiring the second best candidate is obviously better than hiring the third best. We can further extend this situation by assuming that every secretary has a cardinal value distributed according to some probability measure, but where the payoff solely depends on the rank of the selected candidate (see Gnedin and Krengel (1995) and the references therein, and Bearden (2006)). In this case, however, it is perhaps more natural to assume that the payoff is exactly equal to the cardinal value of the selected candidate, instead of its relative rank. For instance, if there are two secretaries with neighbouring ranks, then selecting the best amongst these two is less relevant if their values are close, and more relevant if the difference in values is high. This is exactly the model as studied by Moser (1956), who assumes that the values are independently and uniformly distributed on $[0,1]$. In fact, Moser's model is a variant of a problem considered by Arthur Cayley in the nineteenth century. (See Ferguson (1989) for a description of Cayley's problem.)

In the present paper we take Moser's model, but assume in addition that there is an adversary who knows the values of the secretaries, and chooses the order of the secretaries strategically. The employer, on the other hand, does not know these values,

* Corresponding author. Tel.: +31 43 3883922.

E-mail addresses: J.Flesch@maastrichtuniversity.nl (J. Flesch), A.Perea@maastrichtuniversity.nl (A. Perea).

but only knows the number of secretaries, and the distribution of their respective values. We thus obtain a zero-sum game with incomplete information on the employer's side. We are not the first to take a game theoretic approach to the problem.¹ See, for instance, Gilbert and Mosteller (1966), Gnedin and Krenzel (1995) and de Carvalho et al. (2008).

The adversary's main problem is how to optimally exploit his private information. This is a difficult problem since the adversary, by using his private information, would make choices that would reveal part of his private information to the employer. In our analysis, however, we mainly focus on the employer, in line with the literature on the secretary problem. In particular, we will be interested in his max–min strategies, that is, strategies for which the worst-case expected payoff is as high as possible.

For the case of two secretaries, we prove that the value exists, and show that the employer's unique² optimal strategy is to hire the first secretary precisely when her value is at least 0.5. This strategy guarantees an expected payoff of 7/12 to the employer. An optimal strategy for the adversary is to first send the secretary whose value is closer to 0.5.

If there are more than two secretaries, we are not able to find exact max–min strategies for the employer. However, we provide strategies with a simple structure that approach the max–min value within a distance of at most 0.07. The class of strategies we focus on are *threshold strategies*, and they work as follows: For every period k choose a threshold a_k , which may depend on the values of the rejected secretaries, and hire the current secretary precisely when her value is at least a_k . Such a threshold strategy is called *stationary* if a_k is constant throughout the game, except for the last period where the employer must accept the last candidate. The strategy is a *Markov threshold strategy* if a_k depends on the period k , but not on the values of the rejected secretaries.

We show that the best stationary threshold strategy is to choose the threshold equal to $(1/n)^{1/(n-1)}$, where n is the number of secretaries. This threshold converges (slowly) to 1 if n tends to infinity. Interestingly, this is also the best stationary threshold strategy in Moser's model, where the order of the secretaries is not chosen strategically. We show that this stationary threshold strategy performs relatively well in general, as it approximates the max–min value by at most 0.08.

We then turn to Markov threshold strategies. We show that the best amongst these involves thresholds that are non-increasing over time. For the case of two and three secretaries, this strategy is in fact the best stationary threshold strategy discussed above. So, for these cases choosing different thresholds over time does not yield higher payoffs. We conjecture, supported by numerical simulations, that this remains to be true for more than three secretaries as well.

However, for at least three secretaries, we show that the best threshold strategy must base its thresholds not only on the period, but also on the values of the rejected secretaries. Nevertheless, it remains true that the thresholds should be non-increasing over time. We prove that the employer, by using such general threshold strategies, can approach the max–min value by at most 0.07.

It turns out to be very difficult to provide effective strategies for the adversary. We do, however, provide some suggestions at the end of the paper.

1.1. Some open problems

We end this introduction by providing an overview of some interesting open problems, which may initiate some further research in this area. (1) Find out if the value and optimal strategies exist for more than two secretaries. And if these exist, determine the exact value, and optimal strategies, of the game; (2) Verify whether, in general, pure optimal strategies always exist, and whether this can be a threshold strategy for player 2; (3) Explore if, in general, Markov threshold strategies are not better for player 2 than the best stationary threshold strategy.

1.2. Outline of the paper

In Section 2 we introduce the model. In Section 3 we describe the optimal strategy for the employer if the order of secretaries is not chosen strategically. After this section we will explore the situation where the adversary is strategic, that is, chooses the order of secretaries to his own advantage. Section 4 covers the case of two secretaries. In Section 5 we turn to the case of more than two secretaries, and examine stationary threshold strategies for the employer. General threshold strategies are explored in Section 6. Section 7 contains some concluding remarks, also on effective strategies for the adversary. Some technical proofs have been moved to the appendix.

2. The Model

2.1. The game

Consider the following game G_n , where $n \in \mathbb{N}$, played by two players. There are n independent random variables Z_1, \dots, Z_n , each of which is uniformly distributed on $[0, 1]$. We assume that both players know n , the independence and the distribution of these random variables, but only player 1 knows the vector of realizations $\mathbf{z} := (z_1, \dots, z_n)$ of these. The game is played as follows. Let $N = \{1, \dots, n\}$. At period 1, player 1 chooses one of $\{z_i\}_{i \in N}$, say z_{k_1} , and offers z_{k_1} euros to player 2. If player 2 accepts, then player 1 pays z_{k_1} euros to player 2 and play ends. Otherwise, if player 2 rejects, play continues at period 2, where player 1 chooses one of the remaining amounts $\{z_i\}_{i \in N - \{k_1\}}$, say z_{k_2} . Player 1 subsequently offers z_{k_2} euros to player 2, who either accepts or rejects. If player 2 accepts, then player 1 pays z_{k_2} euros to player 2 and play ends, whereas if player 2 rejects, then player 1 has to offer one of the remaining amounts $\{z_i\}_{i \in N - \{k_1, k_2\}}$. This continues until player 2 accepts an offer. If player 2 has rejected $n - 1$ times, player 2 has to accept the last offer at period n .

In terms of the secretary problem as described in the introduction, player 1 corresponds to the adversary whereas player 2 plays the role of employer. The realizations of the random variables Z_1, \dots, Z_n are the values of the n secretaries.

2.2. Strategies

Let $\Omega := [0, 1]$ be the space of realizations. For every $k \in N$, let $\Omega^k := [0, 1]^k$ be endowed with the Borel σ -algebra. Let $\phi(N)$ denote the set of all permutations of N . Note that player 1, after a vector of realizations, could as well just choose one permutation in $\phi(N)$ in advance, instead of choosing period by period. A pure strategy for player 1 is a measurable function s^1 which specifies a permutation in $\phi(N)$ for each possible vector of realizations $\mathbf{z} \in \Omega^n$. By S^1 we denote the set of pure strategies for player 1, which we endow with the product topology. Let $\mathcal{B}(S^1)$ be the induced Borel σ -algebra on S^1 . A mixed strategy for player 1 is a probability measure σ^1 on $(S^1, \mathcal{B}(S^1))$. Denote by Σ^1 the set of mixed strategies for player 1.

¹ Most of the models are interested in the relative rank of the chosen secretary, and not in the cardinal value as we are. There is, however, a relationship between the two approaches. Bruss and Ferguson (1993) show, namely, that there is a strong correlation between the cardinal values and their associated ranks. See also Bruss (2005).

² To be precise, unique up to behavior on a set of measure zero.

Note that a pure strategy s^1 can be identified with a mixed strategy σ^1 putting probability 1 on $\{s^1\}$.

At any decision point for player 2, the history observed by player 2 is the sequence consisting of all amounts that player 1 has offered before the current period. A pure strategy for player 2 is a collection s^2 of functions $s_k^2, k \in \{1, \dots, n\}$, where s_k^2 is a measurable function which assigns a member from $\{\text{Accept}, \text{Reject}\}$ to any vector of realizations in Ω^k . The interpretation of s_k^2 is that player 2 makes a decision at period k , based on the history and the currently offered amount. By S^2 we denote the set of pure strategies for player 2, which again we endow with the product topology. Let $\mathcal{B}(S^2)$ be the induced Borel σ -algebra on S^2 . A mixed strategy for player 2 is a probability measure σ^2 on $(S^2, \mathcal{B}(S^2))$. Denote by Σ^2 the set of mixed strategies for player 2. Again, a pure strategy s^2 can be identified with a mixed strategy σ^2 putting probability 1 on $\{s^2\}$.

A pure strategy s^2 for player 2 is called a threshold strategy if, after any history h of player 2, there exists a threshold $a(h) \in [0, 1]$ such that s^2 prescribes to accept the current offer y when $y \geq a(h)$ and prescribes to reject it when $y < a(h)$. If these thresholds only depend on the period, then s^2 is called a Markov threshold strategy, whereas if there is just one threshold then s^2 is called a stationary threshold strategy.³ Markov threshold strategies for player 2 are given and denoted by the sequence of thresholds $\mathbf{a} := (a_1, \dots, a_{n-1})$ for the first $n - 1$ periods, whereas a stationary threshold strategy of player 2 is simply a threshold $a \in [0, 1]$.

2.3. Utility

With respect to a pair of (mixed) strategies (σ^1, σ^2) and a vector of realizations $\mathbf{z} \in \Omega^n$, let $U_{\mathbf{z}}(\sigma^1, \sigma^2)$ denote the expected amount that player 1 has to pay to player 2. Denote by P^n the uniform distribution on Ω^n , and let

$$U(\sigma^1, \sigma^2) := \int_{\mathbf{z} \in \Omega^n} U_{\mathbf{z}}(\sigma^1, \sigma^2) dP^n,$$

be the expected amount that player 1 must pay to player 2, prior to the realization. Note that $U(\sigma^1, \sigma^2)$ is well-defined because $U_{\mathbf{z}}(\sigma^1, \sigma^2)$ is a bounded and measurable function of $\mathbf{z} \in \Omega^n$, and therefore integrable on Ω^n with respect to P^n . We also refer to $U(\sigma^1, \sigma^2)$ as the expected utility. We evaluate every strategy σ^1 of player 1 by

$$\psi^1(\sigma^1) = \sup_{\sigma^2} U(\sigma^1, \sigma^2),$$

which is the worst-case scenario for what player 1 has to pay in expectation. Similarly, for every strategy σ^2 of player 2, let

$$\psi^2(\sigma^2) = \inf_{\sigma^1} U(\sigma^1, \sigma^2).$$

A strategy σ^1 for player 1 is called a best reply to a strategy σ^2 of player 2, if $U(\sigma^1, \sigma^2) = \psi^2(\sigma^2)$. Similarly, a strategy σ^2 for player 2 is called a best reply to a strategy σ^1 of player 1, if $U(\sigma^1, \sigma^2) = \psi^1(\sigma^1)$.

2.4. The min-max and max-min values

The min-max value of the game is given by

$$\psi^1 := \inf_{\sigma^1} \psi^1(\sigma^1).$$

So, player 1 can guarantee not to pay more than ψ^1 by appropriately choosing a strategy σ^1 (up to an arbitrarily small error term). A strategy σ^1 is called an ϵ -min-max strategy, where $\epsilon \geq 0$, if

$$\psi^1(\sigma^1) \leq \psi^1 + \epsilon.$$

Similarly, we define the max-min value by

$$\psi^2 := \sup_{\sigma^2} \psi^2(\sigma^2),$$

which is the highest possible payment that player 2 can guarantee for himself by appropriately choosing a strategy σ^2 (up to an arbitrarily small error term). A strategy σ^2 is called an ϵ -max-min strategy, where $\epsilon \geq 0$, if

$$\psi^2(\sigma^2) \geq \psi^2 - \epsilon.$$

It is known that $\psi^1 \geq \psi^2$. If they are equal, then this amount is called the value of the game, and is denoted by v_n . If v_n exists, then a strategy σ^1 for player 1 is called optimal if $\psi^1(\sigma^1) = v_n$, whereas a strategy σ^2 for player 2 is called optimal if $\psi^2(\sigma^2) = v_n$. It is clear that σ^1 and σ^2 are optimal if and only if they are best replies to each other.

We note that we are able to prove that the value exists for $n = 2$ (see Section 4), but it is still an open problem whether it exists for $n \geq 3$. Our conjecture is that it exists for $n \geq 3$ as well.

3. Playing against a non-strategic player 1

In this section, we examine the situation in which player 1 does not manipulate the order of the realizations z_1, \dots, z_n , and simply chooses z_k for period k . For every n , let τ_n^1 denote this strategy for player 1, and let \tilde{v}_n denote the best utility player 2 can achieve against τ_n^1 , i.e. $\tilde{v}_n = \psi^1(\tau_n^1)$.

The most important properties of this situation are summarized below. Most of these were already proven by Moser (1956).

Theorem 1 (Non-strategic player 1).

- (1) Player 2's best reply to τ_n^1 , unique up to a set of measure zero, is the Markov threshold strategy which, for period $k \in \{1, \dots, n\}$, prescribes threshold $b_k := \tilde{v}_{n-k}$. (Recall that \tilde{v}_{n-k} is player 2's best utility against τ_{n-k}^1).
- (2) Player 2's best utility \tilde{v}_n satisfies the recursion $\tilde{v}_1 = \frac{1}{2}$, and $\tilde{v}_n = \frac{1}{2} + \frac{1}{2}(\tilde{v}_{n-1})^2$ for all $n \geq 2$.
- (3) The sequence \tilde{v}_n is strictly increasing and $\lim_{n \rightarrow \infty} \tilde{v}_n = 1$.
- (4) Player 2's best amongst the stationary threshold replies to τ_n^1 is $a_n^* = (\frac{1}{n})^{\frac{1}{n-1}}$, i.e. for any stationary threshold strategy a for player 2 we have $U(\tau_n^1, a_n^*) \geq U(\tau_n^1, a)$. The strategy a_n^* , while not being a best reply to τ_n^1 for any $n \geq 3$, is asymptotically a best reply, i.e. $\lim_{n \rightarrow \infty} U(\tau_n^1, a_n^*) = \lim_{n \rightarrow \infty} \tilde{v}_n (= 1)$.

Proof. First, we show part 1. Consider period 1. If player 2 decides to reject, then $n - 1$ amounts will remain, yielding \tilde{v}_{n-1} in expectation at a best continuation. Hence, if player 2 is offered at least \tilde{v}_{n-1} at period 1, then he accepts, otherwise he rejects. This argument holds for any later period, which proves part 1.

Next, we prove part 2. At period 1, with regard to the strategy prescribed in part 1, player 2 accepts with probability $1 - \tilde{v}_{n-1}$, with conditional expected amount $\frac{1}{2}(\tilde{v}_{n-1} + 1)$, and rejects with probability \tilde{v}_{n-1} . Therefore,

$$\tilde{v}_n = (1 - \tilde{v}_{n-1}) \cdot \frac{1}{2}(\tilde{v}_{n-1} + 1) + \tilde{v}_{n-1} \cdot \tilde{v}_{n-1} = \frac{1}{2} + \frac{1}{2}(\tilde{v}_{n-1})^2. \quad (1)$$

It is obvious that $\tilde{v}_1 = \frac{1}{2}$, so part 2 has been verified.

Part 3 is simple and intuitive. Take some $n \geq 2$. Then, $\tilde{v}_n > \tilde{v}_{n-1}$ because by (1)

$$\begin{aligned} \tilde{v}_n &= (1 - \tilde{v}_{n-1}) \cdot \frac{1}{2}(1 + \tilde{v}_{n-1}) + \tilde{v}_{n-1} \cdot \tilde{v}_{n-1} \\ &> (1 - \tilde{v}_{n-1}) \cdot \tilde{v}_{n-1} + \tilde{v}_{n-1} \cdot \tilde{v}_{n-1} = \tilde{v}_{n-1}. \end{aligned}$$

³ In accordance with the literature on dynamic games, we use the term Markov to emphasize that the thresholds only depend on the current period but not on the specific history. Similarly, we use stationary to emphasize the time and history independence of the thresholds.

Here we used that $\tilde{v}_{n-1} \in (0, 1)$. Since the sequence \tilde{v}_n is strictly increasing and $\tilde{v}_n \leq 1$ for all n , we may conclude that $\lim_{n \rightarrow \infty} \tilde{v}_n$ exists. By part 2,

$$\lim_{n \rightarrow \infty} \tilde{v}_n = \frac{1}{2} + \frac{1}{2} \lim_{n \rightarrow \infty} (\tilde{v}_n)^2,$$

yielding $\lim_{n \rightarrow \infty} \tilde{v}_n = 1$.

Finally, we show part 4. Take a stationary threshold strategy a for player 2. With probability a^{n-1} , we have $z_i < a$ for all $i \in \{1, \dots, n-1\}$, in which case player 2 rejects all z_1, \dots, z_{n-1} and must accept z_n , yielding a conditional expectation of $\frac{1}{2}$. On the other hand, with probability $1 - a^{n-1}$, we have $z_i \geq a$ for at least one $i \in \{1, \dots, n-1\}$, hence player 2 will accept the first amount above a , yielding a conditional expectation of $\frac{1}{2}(a+1)$. Thus, strategy a gives

$$U(\tau_n^1, a) = a^{n-1} \cdot \frac{1}{2} + (1 - a^{n-1}) \cdot \frac{1}{2}(a+1) = \frac{1}{2}(1 + a - a^n).$$

By taking derivatives, it easily follows that $U(\tau_n^1, a)$ has a unique maximum at $a_n^* = (\frac{1}{n})^{\frac{1}{n-1}}$, which is in $[0, 1]$.

Note that

$$U(\tau_n^1, a_n^*) = \frac{1}{2} \left(1 + a_n^* - \frac{1}{n} a_n^{*n} \right),$$

which, in view of Lemma 9, implies

$$\lim_{n \rightarrow \infty} U(\tau_n^1, a_n^*) = \frac{1}{2} \left(1 + \lim_{n \rightarrow \infty} a_n^* \right) = 1.$$

Hence, a_n^* is an asymptotically best reply to τ_n^1 .

Finally, it is clear in view of parts 1 and 3 that a_n^* is not optimal when $n \geq 3$, since in this case different thresholds must be used at periods 1 and 2. \square

Remark. The following table shows an approximation of \tilde{v}_n for some values of n :

n	2	3	4	5	10	20	50	100
\tilde{v}_n	≈ 0.63	≈ 0.70	≈ 0.74	≈ 0.78	≈ 0.86	≈ 0.92	≈ 0.96	≈ 0.98

In fact, Moser (1956) and Gilbert and Mosteller (1966) showed that

$$\tilde{v}_n \approx 1 - \frac{2}{n + \ln(n) + b},$$

if n is large. Here, b is a constant approximately equal to 1.7680.

4. The special case of two random variables (The Game G_2)

From now on, we will focus on the situation in which there is a strategic adversary. In this section, we examine in detail the case when we have two random variables, that is, $n = 2$. So, player 1 must choose an order for the realizations z_1 and z_2 . In this game both players only have to make a choice at period 1. Therefore, whenever we speak about a player's choice we always mean his choice at period 1. We show the following results.

Theorem 2 (Properties of the game G_2).

- (1) An optimal strategy for player 1 is to choose (in period 1) the amount closer to $\frac{1}{2}$, i.e. to choose z_1 if $|z_1 - \frac{1}{2}| \leq |z_2 - \frac{1}{2}|$, and to choose z_2 otherwise.
- (2) Player 2's optimal strategy, unique up to a set of measure zero, is the stationary threshold strategy $a = \frac{1}{2}$.
- (3) The value is $v_2 = \frac{7}{12}$.

Proof. Let σ^1 and a denote the strategies described in part 1 and part 2, respectively. It is sufficient to show that (1) σ^1 is a best reply to a , (2) a is a best reply to σ^1 , unique up to a set of measure zero, and (3) the induced expected utility is $U(\sigma^1, a) = \frac{7}{12}$. In the following, let $x_1 := \min\{z_1, z_2\}$ and $x_2 := \max\{z_1, z_2\}$.

Step 1: We show that σ^1 is a best reply to a . We distinguish the following cases (we assume that $x_1 \neq x_2$, otherwise player 1's strategy is surely a best reply):

Case 1: $x_1 < x_2 < \frac{1}{2}$. In this case, σ^1 offers x_2 , which a rejects, yielding x_1 as the outcome, which is the best possible amount for player 1.

Case 2: $\frac{1}{2} \leq x_1 < x_2$. In this case, σ^1 offers x_1 , which a accepts, yielding x_1 as the outcome, which is the best possible amount for player 1.

Case 3: $x_1 < \frac{1}{2} \leq x_2$. In this case, σ^1 offers either x_1 or x_2 , depending on the exact values of x_1 and x_2 , but a is going to reject x_1 and accept x_2 . Thus, the outcome is x_2 . Observe that player 1 cannot achieve x_1 given player 2's threshold strategy $a = \frac{1}{2}$.

In conclusion, σ^1 is a best reply to a in all cases.

Step 2: We show that a is a best reply to σ^1 . It will be clear from the proof that the best reply is unique up to a set of measure zero. Suppose player 1 offers some amount x (which is either x_1 or x_2). Let y denote the other amount.

Assume first that $x \geq \frac{1}{2}$. Given player 1's strategy, either $y \leq 1 - x$ or $y \geq x$, otherwise player 1 would offer y instead of x . Thus, given x is offered, y is uniformly distributed over $[0, 1 - x] \cup [x, 1]$, and therefore has conditional expectation $\frac{1}{2}$.

Assume now that $x < \frac{1}{2}$. Given x is offered, we obtain similarly that y is uniformly distributed over $[0, x] \cup [1 - x, 1]$, and therefore has conditional expectation $\frac{1}{2}$ again.

Hence, we see that the conditional expectation of y is always $\frac{1}{2}$, and therefore a best reply for player 2 is to accept x if and only if $x \geq \frac{1}{2}$, in accordance with a .

Step 3: We prove that $U(\sigma^1, a) = \frac{7}{12}$. We distinguish the following cases:

Case 1: $z_1 < \frac{1}{2}$ and $z_2 < \frac{1}{2}$ (cf. case 1 in step 1): This occurs with probability $\frac{1}{4}$, and the outcome is $x_1 = \min\{z_1, z_2\}$ with conditional expectation $\frac{1}{6}$ (cf. Lemma 8 in appendix).

Case 2: $z_1 \geq \frac{1}{2}$ and $z_2 \geq \frac{1}{2}$ (cf. case 2 in step 1): This occurs with probability $\frac{1}{4}$, and the outcome is $x_1 = \min\{z_1, z_2\}$ with conditional expectation $\frac{2}{3}$ (cf. Lemma 8 in appendix).

Case 3: $z_1 < \frac{1}{2} \leq z_2$ or $z_2 < \frac{1}{2} \leq z_1$ (cf. case 3 in step 1): This occurs with probability $\frac{1}{2}$, and the outcome is $x_2 = \max\{z_1, z_2\}$ with conditional expectation $\frac{3}{4}$.

Hence,

$$U(\sigma^1, a) = \frac{1}{4} \cdot \frac{1}{6} + \frac{1}{4} \cdot \frac{2}{3} + \frac{1}{2} \cdot \frac{3}{4} = \frac{7}{12},$$

completing the proof. \square

Remark. In the game G_2 , player 1 in fact has many different optimal strategies, one of which is described in the theorem above. All optimal strategies coincide, up to a set of measure zero, with σ^1 in cases 1 and 2 of step 1, but may show different behavior in case 3. The reason is that in case 3 any behavior for player 1 is a best reply against $a = \frac{1}{2}$.

5. Stationary threshold strategies for player 2

In this section, we identify the best strategy that player 2 has amongst all stationary threshold strategies. It turns out that this

strategy is exactly the same as the one we found in Section 3. This means that player 2, if he is restricted to stationary threshold strategies, will behave identically irrespective of whether player 1 chooses strategically or not.

Theorem 3 (Best stationary threshold strategy). Consider the game G_n . Player 2's best stationary threshold strategy is $a_n^* = (\frac{1}{n})^{\frac{1}{n-1}}$, i.e. for any stationary threshold strategy a of player 2 we have $\psi^2(a_n^*) \geq \psi^2(a)$. Moreover, a_n^* guarantees

$$\psi^2(a_n^*) = \frac{1}{n+1} + a_n^* - (a_n^*)^n = \frac{1}{n+1} + \frac{n-1}{n} a_n^*.$$

Proof. It is clear that $a = 0$ or $a = 1$ cannot be player 2's best stationary threshold strategy. Therefore, take an arbitrary $a \in (0, 1)$. In order to find $\psi^2(a)$, we identify a best reply for player 1. In the following, let x_1, \dots, x_n be a permutation of z_1, \dots, z_n such that $x_1 \leq x_2 \leq \dots \leq x_n$. We distinguish two cases:

Case 1: $x_1 \leq \dots \leq x_n < a$. This case occurs with probability a^n , and the outcome with best play by player 1 is x_1 (player 1 should keep x_1 for the last period when it has to be accepted). The conditional expectation of x_1 is

$$\frac{n \cdot 0 + a}{n+1} = \frac{a}{n+1},$$

(cf. Lemma 8 in appendix).

Case 2: $x_1 \leq \dots \leq x_k < a \leq x_{k+1} \leq \dots \leq x_n$ with some $k \in \{0, \dots, n-1\}$. This case occurs with probability

$$\binom{n}{k} \cdot a^k \cdot (1-a)^{n-k},$$

and the outcome with best play by player 1 is x_{k+1} (all x_1, \dots, x_k are rejected by player 2, and the lowest amount player 2 accepts is x_{k+1}). The conditional expectation of x_{k+1} is

$$\frac{(n-k) \cdot a + 1}{n-k+1},$$

(cf. Lemma 8 in appendix).

So the expected utility when player 1 uses a best reply σ^1 to a is

$$U(\sigma^1, a) = a^n \cdot \frac{a}{n+1} + \sum_{k=0}^{n-1} \binom{n}{k} \cdot a^k \cdot (1-a)^{n-k} \cdot \frac{(n-k) \cdot a + 1}{n-k+1}.$$

Somewhat surprisingly, Lemma 10 guarantees that the above expression reduces to

$$U(\sigma^1, a) = \frac{1}{n+1} + a - a^n.$$

By taking derivatives, it easily follows that $U(\sigma^1, a)$ has a unique maximum at $a_n^* = (\frac{1}{n})^{\frac{1}{n-1}}$, which is in $[0, 1]$. \square

Remark. Lemma 9 shows that $a_n^* = (\frac{1}{n})^{\frac{1}{n-1}}$ converges to 1 when n tends to infinity. Note also that the probability of the maximal amount $\max\{z_1, \dots, z_n\}$ being at least a_n^* is exactly

$$\alpha_n := 1 - (a_n^*)^n = 1 - \frac{1}{n} a_n^*.$$

Hence, if player 2 uses the stationary threshold strategy a_n^* , then the probability that player 2 eventually accepts an amount above the threshold a_n^* is also exactly α_n . Since α_n converges to 1 when n tends to infinity, the strategy a_n^* will accept an amount above a_n^* with probability close to 1, for large n . This, of course, also means that $\psi^2(a_n^*)$ converges to 1, although this also follows directly from the

expression for $\psi^2(a_n^*)$ in Theorem 3. The following table shows an approximation of a_n^* and $\psi^2(a_n^*)$ for some values of n :

n	2	3	4	5	10	20	50	100
a_n^*	=0.5	≈0.58	≈0.63	≈0.67	≈0.77	≈0.85	≈0.92	≈0.95
$\psi^2(a_n^*)$	≈0.58	≈0.63	≈0.67	≈0.70	≈0.79	≈0.86	≈0.92	≈0.95

Interestingly, $a_{n+1}^* = (\frac{1}{n+1})^{\frac{1}{n}}$ is a very good approximation of $\psi^2(a_n^*)$ for all n , i.e. $a_{n+1}^* \geq \psi^2(a_n^*)$ for all $n \geq 2$ and

$$\max_{n \geq 2} (a_{n+1}^* - \psi^2(a_n^*)) = a_3^* - \psi^2(a_2^*) \approx 0.006.$$

Regarding the comparison between $\psi^2(a_n^*)$ and \bar{v}_n (cf. Section 3), we remark that $\bar{v}_n \geq \psi^2(a_n^*)$ for all $n \geq 2$ and that

$$\max_{n \geq 2} (\bar{v}_n - \psi^2(a_n^*)) \leq 0.08. \tag{2}$$

Since the max–min value of the game G_n belongs to the interval $[\psi^2(a_n^*), \bar{v}_n]$, the stationary threshold strategy a_n^* is quite effective for all $n \geq 2$.

6. General threshold strategies for player 2

In this section, we examine general threshold strategies for player 2. First we show that we may restrict our investigation to threshold strategies σ^2 with the following property: if σ^2 prescribes threshold b_k at period k , and the offered amount is below b_k and gets rejected by σ^2 , then the new threshold b_{k+1} at period $k+1$ satisfies $b_k \geq b_{k+1}$. Thus, the thresholds are non-increasing during any play.

Theorem 4 (Thresholds are non-increasing). Consider an arbitrary threshold strategy σ^2 for player 2. Define another threshold strategy $\bar{\sigma}^2$ for player 2 as follows: At period 1, the strategy $\bar{\sigma}^2$ prescribes the same threshold as σ^2 , i.e. $\bar{\sigma}^2(\emptyset) := \sigma^2(\emptyset)$, where \emptyset denotes the empty history at period 1. At any period $k \geq 2$, if (y_1, \dots, y_{k-1}) denotes the sequence of past rejected amounts, then let

$$\bar{\sigma}^2(y_1, \dots, y_{k-1}) := \min_{\ell=1, \dots, k} \sigma^2(y_1, \dots, y_{\ell-1}).$$

Then, $\bar{\sigma}^2$ has the following properties:

1. The threshold strategy $\bar{\sigma}^2$ is at least as good as σ^2 , i.e. $\psi^2(\bar{\sigma}^2) \geq \psi^2(\sigma^2)$.
2. With respect to $\bar{\sigma}^2$, the thresholds are non-increasing during any play: at any period $k \geq 2$ and for any sequence (y_1, \dots, y_{k-1}) of past amounts that $\bar{\sigma}^2$ has rejected, we have

$$\bar{\sigma}^2(y_1, \dots, y_{k-2}) \geq \bar{\sigma}^2(y_1, \dots, y_{k-1}).$$

Proof. Property (2) is obvious, so we only have to show the property (1). Let τ^1 be a pure best reply of player 1 to $\bar{\sigma}^2$. Take an arbitrary realization vector $\mathbf{z} = (z_1, \dots, z_n)$, and suppose that τ^1 prescribes the realizations in the order y_1, \dots, y_n . Let m denote the period at which $\bar{\sigma}^2$ accepts y_m . Thus, $y_m = U_z(\tau^1, \bar{\sigma}^2)$ and $y_m \geq \bar{\sigma}^2(y_1, \dots, y_{m-1})$. If $m > 1$, we may assume without loss of generality that $y_m < \bar{\sigma}^2(y_1, \dots, y_{m-2})$, because otherwise τ^1 could just as well offer y_m already at period $m-1$ since $\bar{\sigma}^2$ would accept it. This means

$$\bar{\sigma}^2(y_1, \dots, y_{m-1}) \leq y_m < \bar{\sigma}^2(y_1, \dots, y_{m-2}),$$

which implies

$$\bar{\sigma}^2(y_1, \dots, y_{m-1}) = \sigma^2(y_1, \dots, y_{m-1}). \tag{3}$$

We now show that $U_z(\tau^1, \sigma^2) = y_m$. Since the threshold prescribed by $\tilde{\sigma}^2$ is never higher than the threshold prescribed by σ^2 , it is clear that σ^2 also rejects y_1, \dots, y_{m-1} up to period $m - 1$. But then, at period m , the strategy σ^2 also accepts y_m in view of (3). Thus, $U_z(\tau^1, \sigma^2) = y_m$ indeed.

Therefore, $U_z(\tau^1, \tilde{\sigma}^2) = y_m = U_z(\tau^1, \sigma^2)$, which implies

$$\psi^2(\tilde{\sigma}^2) = U(\tau^1, \tilde{\sigma}^2) = U(\tau^1, \sigma^2) \geq \psi^2(\sigma^2),$$

completing the proof of property (1). \square

Corollary 5. Let $\mathbf{a} = (a_1, \dots, a_{n-1})$ be a Markov threshold strategy for player 2. Then, there exists a Markov threshold strategy $\mathbf{b} = (b_1, \dots, b_{n-1})$ for player 2 such that $b_k \geq b_{k+1}$ for all $k \in \{1, \dots, n - 2\}$ and for which $\psi^2(\mathbf{b}) \geq \psi^2(\mathbf{a})$. Consequently, a best strategy amongst the Markov threshold strategies for player 2 consists of a non-increasing sequence of thresholds.⁴

The above corollary can also be shown in a more direct way. First, one can prove the following statement about transpositions of neighboring thresholds. Let $\mathbf{a} = (a_1, \dots, a_{n-1})$ be a Markov threshold strategy for which $a_k < a_{k+1}$ holds for some $k \in \{1, \dots, n - 2\}$. Let $\mathbf{b} = (b_1, \dots, b_{n-1})$ denote the Markov threshold strategy obtained by $b_k = a_{k+1}$ and $b_{k+1} = a_k$, while $b_m = a_m$ for all $m \in \{1, \dots, k - 1, k + 2, \dots, n - 1\}$. Then, it can be shown that $\psi^2(\mathbf{b}) \geq \psi^2(\mathbf{a})$. Given this result, the corollary above follows by the well known theorem in algebra that any permutation can be written as a product of transpositions of two neighboring elements.

Theorem 6 (Stationary threshold strategies are not optimal). When $n \geq 3$, player 2 has a threshold strategy which is strictly better than all stationary threshold strategies.

Proof. Let $n \geq 3$. In view of Theorem 3, it suffices to construct a threshold strategy for player 2 which is strictly better than the stationary threshold strategy $a = (\frac{1}{n})^{\frac{1}{n-1}}$. Let $b \in (0, a)$ be arbitrary. Let σ_b^2 be the threshold strategy for player 2 which prescribes threshold a at every period, except in the following case: at period $n - 1$ (which is the last period when player 2 has a choice), if all $n - 2$ previously rejected amounts are in the interval $[b, a)$, then use threshold b at period $n - 1$.

Now we show that σ_b^2 is strictly better than strategy a for player 2, if b is sufficiently close to threshold a . Let (z_1, \dots, z_n) denote the realizations of the amounts. We distinguish the following cases:

Case 1: Less than $n - 2$ amounts in z_1, \dots, z_n are in the interval $[b, a)$. In this case, σ_b^2 is the same as a .

Case 2: Precisely $n - 2$ amounts in z_1, \dots, z_n are in the interval $[b, a)$. In this case, σ_b^2 can only prescribe threshold b at period $n - 1$ if player 1 offers precisely these $n - 2$ amounts in $[b, a)$ at periods up to $n - 2$. But then, for periods $n - 1$ and n , there is no amount left in $[b, a)$. Hence, in this case, σ_b^2 is equally good as a .

Case 3: Precisely $n - 1$ amounts in z_1, \dots, z_n are in the interval $[b, a)$. Let \underline{z} denote the minimum of these $n - 1$ amounts in $[b, a)$ and let w denote the amount outside $[b, a)$. We show that, with best play by player 1, strategy σ_b^2 yields outcome \underline{z} and a yields outcome w .

Case 3i: $w < b$. In this case, with best play by player 1, strategy σ_b^2 yields outcome \underline{z} . Indeed, if player 1 offers all amounts in $[b, a)$ except \underline{z} at periods up to $n - 2$ and offers \underline{z} at period $n - 1$, strategy σ_b^2 will accept \underline{z} as $\underline{z} \geq b$. On the other hand, w cannot be the outcome for the following reason. Player 2 would only accept

at the last period, as $w < b$. Thus, to achieve w , player 1 would have to offer the $n - 1$ amounts in $[b, a)$ at periods up to $n - 1$. But then, σ_b^2 accepts the amount at period $n - 1$. The strategy a , on the other hand, yields w as outcome, because player 1 can reserve w for the last period, when it has to be accepted.

Case 3ii: $w \geq a$. In this case, with best play by player 1, strategy σ_b^2 yields outcome \underline{z} . Indeed, if player 1 offers all amounts in $[b, a)$ except \underline{z} at periods up to $n - 2$ and offers \underline{z} at period $n - 1$, strategy σ_b^2 will accept \underline{z} as $\underline{z} \geq b$. The strategy a , on the other hand, yields w , as a rejects all other amounts.

Case 4: All n amounts in z_1, \dots, z_n are in the interval $[b, a)$. In this case, with best play by player 1, strategy σ_b^2 yields $\min\{z_1, \dots, z_n\}$ as outcome, since player 1 can offer the minimal amount at period $n - 1$. The strategy a also yields $\min\{z_1, \dots, z_n\}$ as outcome, because player 1 can reserve the minimal amount for the last period, when it has to be accepted. Hence, in this case, σ_b^2 is equally good as a .

In conclusion, σ_b^2 is equally good as a in all cases except for case 3. On condition that case 3 occurs, we obtain the following. Recall that, with best play by player 1, strategy σ_b^2 yields outcome \underline{z} and strategy a yields outcome w . Let $E_{\underline{z}}(b)$ denote the conditional expected value of \underline{z} and let $E_w(b)$ denote the conditional expected value of w in case 3. In order to show that strategy σ_b^2 is strictly better than strategy a , for b sufficiently close to a , we need to show that

$$\lim_{b \uparrow a} E_{\underline{z}}(b) > \lim_{b \uparrow a} E_w(b).$$

Since $\underline{z} \in [b, a)$ in case 3, we have

$$\lim_{b \uparrow a} E_{\underline{z}}(b) = a.$$

We calculate $\lim_{b \uparrow a} E_w(b)$ in the following way. Subcase 3i appears with conditional probability

$$p_b = \frac{b}{b + (1 - a)},$$

and the conditional expectation of w is $\frac{b}{2}$. Subcase 3ii appears with conditional probability $1 - p_b$, and the conditional expectation of w is

$$\frac{a + 1}{2}.$$

Hence,

$$E_w(b) = p_b \cdot \frac{b}{2} + (1 - p_b) \cdot \frac{a + 1}{2}.$$

By taking the limit, we obtain

$$\lim_{b \uparrow a} E_w(b) = \frac{a^2}{2} + \frac{1 - a^2}{2} = \frac{1}{2}.$$

In conclusion,

$$\lim_{b \uparrow a} E_{\underline{z}}(b) = a > \frac{1}{2} = \lim_{b \uparrow a} E_w(b),$$

which completes the proof. \square

Theorem 7 (Markov threshold strategies are not optimal). When $n \geq 3$, player 2 has a threshold strategy which is strictly better than all Markov threshold strategies.

Proof. Let $n \geq 3$. Take a best Markov threshold strategy $\mathbf{a} = (a_1, \dots, a_{n-1})$ for player 2. We construct a threshold strategy σ^2 for player 2 which is strictly better than \mathbf{a} . In view of Corollary 5 and Theorem 6, we may assume that $a_1 \geq \dots \geq a_k > a_{k+1} \geq \dots \geq a_{n-1}$. It can easily be verified that $a_1 = 1$ or $a_{n-1} = 0$ can never yield a best strategy amongst all Markov threshold strategies, and hence we assume that $a_1 < 1$ and $a_{n-1} > 0$.

⁴ It can be shown, also by the argument after the corollary, that all best Markov threshold strategies have non-increasing thresholds.

Consider the threshold strategy σ^2 for player 2 which prescribes the same thresholds as \mathbf{a} except in the following case: at period $k + 1$, if the rejected amount at period k was in the interval $[0, a_{k+1})$, then use threshold a_k at period $k + 1$. We show that, with best play by player 1, the strategy σ^2 is strictly better than \mathbf{a} for player 2.

Step 1: We show that, for every vector of realizations $\mathbf{z} = (z_1, \dots, z_n)$, the strategy σ^2 is at least as good as \mathbf{a} for player 2. Let $\mathbf{z} = (z_1, \dots, z_n)$ be a vector of realizations. Let σ^1 denote a pure best reply for player 1 to σ^2 , and suppose that σ^1 offers these amounts in the order (y_1, \dots, y_n) . Let τ^1 be the strategy for player 1 which also uses the order (y_1, \dots, y_n) , except in the following case: if $y_k < a_{k+1}$ and $y_{k+1} \in [a_{k+1}, a_k)$, then use the order

$$(y_1, \dots, y_{k-1}, y_{k+1}, y_k, y_{k+2}, \dots, y_n).$$

We show that $U_z(\sigma^1, \sigma^2) = U_z(\tau^1, \mathbf{a})$, which will imply that $\psi^2(\sigma^2) \geq \psi^2(\mathbf{a})$. We may assume that all $y_1 < a_1, \dots, y_{k-1} < a_{k-1}$ (i.e. they are all rejected) and that $y_k < a_{k+1}$, otherwise (σ^1, σ^2) and (τ^1, \mathbf{a}) lead to the same outcome. We distinguish the following cases:

- Case 1: If $y_{k+1} \geq a_k$. In this case, $U_z(\sigma^1, \sigma^2) = U_z(\tau^1, \mathbf{a}) = y_{k+1}$.
- Case 2: If $y_{k+1} \in [a_{k+1}, a_k)$. In this case, σ^2 will use threshold a_k at period $k + 1$. Thus, both (σ^1, σ^2) and (τ^1, \mathbf{a}) lead to rejections at periods k and $k + 1$, and hence $U_z(\sigma^1, \sigma^2) = U_z(\tau^1, \mathbf{a})$.
- Case 3. If $y_{k+1} < a_{k+1}$. Also in this case, both (σ^1, σ^2) and (τ^1, \mathbf{a}) lead to rejections at periods k and $k + 1$, and hence $U_z(\sigma^1, \sigma^2) = U_z(\tau^1, \mathbf{a})$.

Step 2: We show that there exists a set W of realization vectors such that W has a positive probability and that, for every realization vector in W , the strategy σ^2 is strictly better than \mathbf{a} for player 2.

Let W denote the set of realization vectors in which exactly k amounts are in the interval $[0, a_{n-1})$, exactly 1 amount is in the interval $[a_{k+1}, a_k)$, and exactly $n - k - 1$ amounts are in the interval $[a_1, 1]$. Of course, W has a positive probability, since $a_1 < 1$ and $a_{n-1} > 0$. Take an arbitrary realization vector in W . Notice that, against \mathbf{a} , it is a best reply for player 1 to offer the k amounts in $[0, a_{n-1})$ at periods up to k , which all get rejected, and then at period $k + 1$ to offer the amount in $[a_{k+1}, a_k)$, which is accepted. This does not work against σ^2 , since in this case, σ^2 uses threshold a_k at period $k + 1$. It is easy to see that σ^2 leads to an outcome in $[a_1, 1]$, regardless player 1's strategy. \square

Two important questions arise:

Question (1): Which are the best threshold strategies for player 2?

Question (2): Does player 2 have better strategies than threshold strategies?

Question (1) is already challenging for $n = 3$. The best threshold strategy for player 2 that we could find is the following. Let $a \in (0, 1)$ and $b \in (0, a)$ be arbitrary. Let σ_{ab}^2 be the threshold strategy for player 2 which prescribes threshold a at period 1, and prescribes threshold a at period 2 if the rejected amount was in interval $[0, b)$ and threshold b at period 2 if the rejected amount was in interval $[b, a)$. (This is very similar to strategy σ_b^2 in the proof of Theorem 6, with the only difference being that a here is also a variable.) Let a^* and b^* denote optimal values for a and b . According to a numerical approximation by the program package Mathematica, $a^* \approx 0.5838$ and $b^* \approx 0.4975$. The strategy $\sigma_{a^*b^*}^2$ guarantees $\psi^2(\sigma_{a^*b^*}^2) \approx 0.6354$ (the best stationary threshold strategy found in Theorem 3 yields 0.6349, which is just “slightly” less). We did not manage to find an improvement on $\sigma_{a^*b^*}^2$ for player 2. It seems natural to try to find an improvement by splitting $[0, a)$ into more

than 2 subintervals. Thus, instead of a and b only, now player 2 can choose $a, b, c, d \in (0, 1)$, with $d \leq c \leq b \leq a$, and a map $\alpha : \{a, b, c, d\} \rightarrow \{a, b, c, d\}$. Then, player 2 will use the following threshold strategy: At period 1, use threshold a . If the offered amount y^1 is rejected, then the new threshold at period 2 will depend on which of the subintervals $[0, d)$, $[d, c)$, $[c, b)$, $[b, a)$ contains y^1 . More precisely, if y^1 is contained in subinterval $[u, w)$, then the new threshold is $\alpha(w)$. With the help of the program package Mathematica, with a numerical precision of 10^{-10} , we found the surprising conclusion that $\sigma_{a^*b^*}^2$ is still a best amongst these strategies, i.e. one of the optimal choices is $a = a^*, b = c = d = b^*$ and $\alpha(a) = b, \alpha(b) = a$. We do not see now how one could improve upon $\sigma_{a^*b^*}^2$. It is not even really clear to us why $\sigma_{a^*b^*}^2$ is so effective, although the proof of Theorem 6 provides some ideas.

We do not know the answer to Question (2). We only know that an optimal threshold strategy exists for player 2 in G_2 , cf. Section 4.

7. Concluding remarks

7.1. Are Markov threshold strategies really better for player 2 than stationary threshold strategies?

Notice that corollary 5 does not exclude the possibility that the best Markov threshold strategy is the stationary threshold strategy which we derived in Section 5. In fact, we conjecture that this holds true for all n . For $n = 1$ and $n = 2$, this is trivial, since player 2 uses at most one threshold. For $n = 3$, we verified this conjecture in the following way. For an arbitrary Markov threshold strategy $\sigma^2 = (a_1, a_2)$ for player 2 with non-increasing thresholds $a_1 \geq a_2$, we determined a best response for player 1, and calculated the corresponding expected outcome, as a function of a_1 and a_2 . Then, we checked that this expected outcome is indeed maximal when $a_1 = a_2 = (\frac{1}{3})^{\frac{1}{2}}$. For $n = 4$ and $n = 5$, the program package Mathematica confirms this conjecture (with a numerical precision of 10^{-5}). For a general n , it is difficult to prove this conjecture. First of all, the proof we used for $n = 3$ produces huge polynomials when n is large. It seems more natural to try a proof based on induction. Perhaps, the best candidate is to try to show the following statement: if $\mathbf{a} = (a_1, a_2, \dots, a_{n-1})$ is a Markov threshold strategy for player 2 with thresholds

$$a_1 \geq \dots \geq a_k > a_{k+1} = \dots = a_{n-1},$$

then there is a strictly better Markov threshold strategy of the form $\tilde{\mathbf{a}} = (a_1, \dots, a_k, b, \dots, b)$ with an appropriate threshold $b \in [a_{k+1}, a_k]$.

7.2. Effective strategies for player 1

In this section, we focus on player 1. The main difference is that here we have to deal with more complex strategies, as player 1 can base his decisions on the realization vector. In general for $n \geq 3$, we have not been able to find a 0-min-max strategy for player 1. Nevertheless, we provide some insight, and present effective strategies for player 1 with a simple structure (although just offering amount z_k at period k for every k , as in Section 3, is already quite effective, especially for large n).

Recall that, for any $n \geq 3$, the max-min value of the game belongs to the interval $[\psi^2(a_n^*), \tilde{v}_n]$, where $\psi^2(a_n^*)$ and \tilde{v}_n are determined in Sections 5 and 3, respectively. We reiterate the approximations:

n	2	3	4	5	10	20	50	100
\tilde{v}_n	≈ 0.63	≈ 0.70	≈ 0.74	≈ 0.78	≈ 0.86	≈ 0.92	≈ 0.96	≈ 0.98
$\psi^2(a_n^*)$	≈ 0.58	≈ 0.63	≈ 0.67	≈ 0.70	≈ 0.79	≈ 0.86	≈ 0.92	≈ 0.95

As pointed out in (2),

$$\max_{n \geq 2} (\tilde{v}_n - \psi^2(a_n^*)) \leq 0.08,$$

which means that the strategy for player 1 in which he offers every realization z_k at period k , is already 0.08-min-max. Nevertheless, player 1 can do better than \tilde{v}_n . Below we provide some possible improvements.

7.2.1. Improvement 1

Based on the results for $n = 2$ in Section 4, there is one simple improvement for player 1 for all $n \geq 3$. Let ξ_n^1 be the strategy for player 1 which, for any realization vector $\mathbf{z} = (z_1, \dots, z_n)$, prescribes the following: At any period $k \leq n - 2$, offer z_k to player 2. At the last two periods, i.e. at periods $n - 1$ and n , play the strategy found in Section 4 with the two remaining amounts z_{n-1} and z_n . Now we determine $u_n := \psi^1(\xi_n^1)$. We have $u_2 = v_2 = \frac{7}{12}$ according to Theorem 2. We proceed by calculating u_3 . Just as in the proof of Theorem 1, player 2 should accept the offered amount y at period 1 if $y \geq v_2 = \frac{7}{12}$, and reject it otherwise. Thus, player 2 accepts y with probability $1 - v_2$ and with conditional expected amount $\frac{1}{2}(v_2 + 1)$, and rejects with probability v_2 . Therefore,

$$u_3 = (1 - v_2) \cdot \frac{1}{2}(v_2 + 1) + v_2 \cdot v_2 = \frac{1}{2} + \frac{1}{2}(v_2)^2.$$

Using this argument inductively, we obtain for all $n \geq 3$ that

$$u_n = \frac{1}{2} + \frac{1}{2}(u_{n-1})^2.$$

We obtained a similar recursion for \tilde{v}_n in Theorem 1. As $\tilde{v}_2 = \frac{5}{8} > \frac{7}{12} = u_2$, we may conclude $\tilde{v}_n > u_n$ for all $n \geq 2$. The following table provides an approximation of u_n for some values of n :

n	2	3	4	5	10	20	50	100
$\psi^1(\xi_n^1) = u_n$	≈ 0.58	≈ 0.67	≈ 0.72	≈ 0.76	≈ 0.86	≈ 0.92	≈ 0.96	≈ 0.98

Note that

$$\max_{n \geq 2} (\psi^1(\xi_n^1) - \psi^2(a_n^*)) \leq 0.07, \tag{4}$$

which means that ξ_n^1 is 0.07-min-max for player 1.

7.2.2. Improvement 2

Recall from Section 4 that, for $n = 2$, it was optimal for player 1 to choose the amount closer to $\frac{1}{2}$. We now try to generalize this strategy for $n = 3$. Take some $w \in [\frac{1}{2}, 1]$. Let σ_w^1 be the strategy for player 1 which, for realizations z_1, z_2, z_3 , prescribes the following: At period 1, player 1 should offer the amount closest to w . If this amount is rejected, then at period 2, player 1 should offer the amount amongst the two remaining amounts which is closer to w .

Suppose that, against σ_w^1 , player 2 uses a pure threshold strategy σ_w^2 . (One can show, based on the discussion below, that player 2 has a best reply in threshold strategies.) Let a_w denote the threshold prescribed by σ_w^2 at period 1. If σ_w^2 rejects amount $y_1 \in [0, a_w)$ at period 1, it is relatively easy to determine the best threshold for period 2. Let y_2 and y_3 denote the amounts chosen by σ_w^1 for periods 2 and 3. We distinguish the following cases.

Case 1: If $y_2 \geq w$. In this case, since y_3 is not closer to w than y_2 , we have either $y_3 \in [y_2, 1]$ or $y_3 \in [0, 2w - y_2]$. This gives a conditional expectation of y_3 equal to

$$d_3 = \frac{(1 - y_2) \cdot \frac{y_2 + 1}{2} + (2w - y_2) \cdot \frac{2w - y_2}{2}}{(1 - y_2) + (2w - y_2)}.$$

As $w \geq \frac{1}{2}$, we have $1 - y_2 \leq 2w - y_2$. Moreover, $y_2 \geq w$ implies that

$$\frac{y_2 + 1}{2} \geq \frac{2w - y_2}{2}.$$

Hence,

$$d_3 \leq \frac{1}{2} \left(\frac{y_2 + 1}{2} + \frac{2w - y_2}{2} \right) = \frac{1}{4}(2w + 1) \leq w \leq y_2.$$

This means that player 2 should accept y_2 .

Case 2: If $y_2 < w$. In this case, since y_3 is not closer to w than y_2 , we have either $y_3 \in [0, y_2]$ or $y_3 \in [2w - y_2, 1]$ (with the latter interval being empty when $2w - y_2 > 1$).

Case 2.1: If $2w - y_2 > 1$. In this case, $y_3 \in [0, y_2]$ and player 2 should accept y_2 .

Case 2.2: If $2w - y_2 \leq 1$. In this case, the conditional expectation of y_3 is equal to

$$d_3 = \frac{y_2 \cdot \frac{y_2}{2} + (1 - 2w + y_2) \cdot \frac{2w - y_2 + 1}{2}}{y_2 + (1 - 2w + y_2)}. \tag{5}$$

Here, player 2 should accept y_2 exactly when $y_2 \geq d_3$, which is a quadratic inequality.

By backwards induction, it would be possible to calculate the best threshold a_w for period 1. We have found the following, with the help of a simulation with the program package Mathematica:

- Regarding σ_w^1 : The best value of w for player 1 is about $w = 0.59$.
- Regarding σ_w^2 : Against $\sigma_{0.59}^1$, the best value of threshold $a_{0.59}$ is about $a_{0.59} = 0.59$. (This implies for case 2.2, in view of (5), that player 2 should accept y_2 exactly when $y_2 \leq 0.21$ or $y_2 \geq 0.47$, approximately.)
- Regarding the expected outcome: we have $\psi^1(\sigma_{0.59}^1) \approx U(\sigma_{0.59}^1, \sigma_{0.59}^2) \approx 0.639$.

This numerical simulation indicates that the strategy $\sigma_{0.59}^1$ is very effective, as $v_3 \geq \psi^2(a_3^*) \approx 0.6349$ (or even $v_3 \geq \psi^2(\sigma_{a^*b^*}^2) \approx 0.6354$ where $\sigma_{a^*b^*}^2$ is the strategy given at the end of Section 6).

This also yields a possible improvement for player 1 for all $n \geq 4$, just as above. Let τ_n^1 be the strategy for player 1 which, for any realization vector $\mathbf{z} = (z_1, \dots, z_n)$, prescribes the following: At any period $k \leq n - 3$, offer amount z_k to player 2. At the last three periods, play the strategy $\sigma_{0.59}^1$ with amounts z_{n-2}, z_{n-1} and z_n . With a similar calculation as before we find the following approximate values of $\psi^1(\tau_n^1)$ for different values of n :

n	3	4	5	10	20	50	100
$\psi^1(\tau_n^1)$	≈ 0.64	≈ 0.70	≈ 0.75	≈ 0.85	≈ 0.92	≈ 0.96	≈ 0.98

with

$$\max_{n \geq 2} (\psi^1(\tau_n^1) - \psi^2(a_n^*)) \leq 0.065.$$

We do not know whether, for $n > 3$, a strategy similar to σ_w^1 would be effective for player 1.

7.3. The value v_n and optimal strategies

Except for $n = 2$ in Section 4, we are not sure whether the value and optimal strategies exist. Nevertheless, we know that $v_2 = \frac{7}{12}$ and that the min-max and max-min values are strictly increasing in n and converge to 1 as n tends to infinity. For an estimation of these values, we may use the interval $I_n = [\psi^2(a_n^*), \psi^1(\xi_n^1)]$. The length of I_n is at most 0.07, in view of (4), and converges to 0 as n tends to infinity. Consequently, the stationary threshold strategy a_n^* is 0.07-max-min for player 2 and the strategy ξ_n^1 is also 0.07-min-max for player 1. Note that, at the end of Sections 6

and 7.2, we provided possible improvements on these strategies (i.e. $\sigma_{a,b}^2$ for $n = 3$ and τ_n^1 for $n \geq 3$), which also yield in turn a better estimate for the min–max and max–min values. We remark that, when $n = \infty$ (i.e. the case of countably infinite random variables), the value equals $v_\infty = 1$ and any strategy of player 1 is optimal, whereas player 2 has only near-optimal strategies.

Acknowledgements

We thank Péter Csóka and Jérôme Renault for discussions on this subject, as well as the editor and two reviewers for their useful comments. We are also grateful to Wolfram Research, Inc., for program package Mathematica, which assisted us with some numerical approximations.

Appendix A

Lemma 8. Let W_1, \dots, W_m denote independent random variables, each of which having a uniform distribution on some interval $[c, d]$. Then

$$\mathbb{E}(\min\{W_1, \dots, W_m\}) = \frac{m \cdot c + d}{m + 1}$$

and

$$\mathbb{E}(\max\{W_1, \dots, W_m\}) = \frac{c + m \cdot d}{m + 1}.$$

Proof. Let $W^* = \min\{W_1, \dots, W_m\}$. We will show for the interval $[c, d] = [0, 1]$ that

$$\mathbb{E}(W^*) = \frac{1}{m + 1}.$$

Since $\max\{W_1, \dots, W_m\} = 1 - \min\{1 - W_1, \dots, 1 - W_m\}$, we obtain part 2 for interval $[0, 1]$. It is easy to check, by using the linearity of the expectation, that the results hold for a general interval $[c, d]$ as well.

So, take $[c, d] = [0, 1]$. Let W denote a random variable having a uniform distribution on interval $[0, 1]$. Then, its density $f_W(t)$ is given by $f_W(t) = 1$ for $t \in [0, 1]$ and $f_W(t) = 0$ otherwise. Also, its cumulative distribution $F_W(t) = \mathbb{P}(W \leq t)$ is given by $F_W(t) = 0$ if $t < 0$, and $F_W(t) = t$ if $t \in [0, 1]$ and $F_W(t) = 1$ if $t > 1$.

As for W^* , we clearly have $F_{W^*}(t) = 0$ if $t < 0$ and $F_{W^*}(t) = 1$ if $t > 1$, whereas for $t \in [0, 1]$

$$\begin{aligned} F_{W^*}(t) &= \mathbb{P}(W^* \leq t) = 1 - \mathbb{P}(W^* > t) \\ &= 1 - \mathbb{P}(\{W_i > t \forall i \in \{1, \dots, m\}\}) = 1 - (\mathbb{P}(W > t))^m \\ &= 1 - (1 - F_W(t))^m = 1 - (1 - t)^m. \end{aligned}$$

Hence, $f_{W^*}(t) = 0$ if $t < 0$ or $t > 1$, while for $t \in [0, 1]$

$$f_{W^*}(t) = \frac{d}{dt} F_{W^*}(t) = m \cdot (1 - t)^{m-1}.$$

So,

$$\begin{aligned} \mathbb{E}(W^*) &= \int_{-\infty}^{\infty} t \cdot f_{W^*}(t) dt = \int_0^1 t \cdot m \cdot (1 - t)^{m-1} dt \\ &= m \cdot \int_0^1 (1 - u) \cdot u^{m-1} du = m \cdot \left[\frac{1}{m} u^m - \frac{1}{m+1} u^{m+1} \right]_{u=0}^{u=1} \\ &= m \cdot \left[\frac{1}{m} - \frac{1}{m+1} \right] = \frac{1}{m+1}, \end{aligned}$$

where we used the substitution $u = 1 - t$. \square

Lemma 9. $\lim_{n \rightarrow \infty} \left(\frac{1}{n}\right)^{\frac{1}{n-1}} = 1$.

Proof. By using the rule of L'Hospital, we have

$$\lim_{n \rightarrow \infty} \frac{\ln(n)}{n-1} = \lim_{n \rightarrow \infty} \frac{1}{n} = 0.$$

Hence,

$$\begin{aligned} \lim_{n \rightarrow \infty} \left(\frac{1}{n}\right)^{\frac{1}{n-1}} &= \lim_{n \rightarrow \infty} n^{-\frac{1}{n-1}} = \lim_{n \rightarrow \infty} e^{\ln\left(n^{-\frac{1}{n-1}}\right)} \\ &= \lim_{n \rightarrow \infty} e^{-\frac{1}{n-1} \ln(n)} = e^0 = 1. \quad \square \end{aligned}$$

Lemma 10. For every $a \in \mathbb{R} - \{0, 1\}$ and every $n \in \mathbb{N}$ it holds that

$$\frac{a^{n+1}}{n+1} + \sum_{k=0}^{n-1} \binom{n}{k} \cdot a^k \cdot (1-a)^{n-k} \cdot \frac{(n-k) \cdot a + 1}{n-k+1} = \frac{1}{n+1} + a - a^n. \tag{6}$$

Proof. Let D_n be equal to the lefthandside of (6). Then,

$$\begin{aligned} D_n &= \frac{a^{n+1}}{n+1} + \sum_{k=0}^{n-1} \frac{n!}{k! \cdot (n-k+1)!} \cdot a^k \cdot (1-a)^{n-k} \cdot ((n-k+1) \cdot a \\ &\quad + (1-a)) \\ &= \frac{a^{n+1}}{n+1} + \sum_{k=0}^{n-1} \frac{n!}{k! \cdot (n-k+1)!} \cdot a^k \cdot (1-a)^{n-k} \cdot ((n-k+1) \cdot a \\ &\quad + (1-a)) - a^n - \frac{a^{n+1}}{n+1} \\ &= \sum_{k=0}^{n-1} \frac{n!}{k! \cdot (n-k+1)!} \cdot a^k \cdot (1-a)^{n-k} \cdot ((n-k+1) \cdot a + (1-a)) \\ &\quad - a^n \\ &= \sum_{k=0}^{n-1} \frac{n! \cdot (n+1-k)}{k! \cdot (n+1-k)!} \cdot a^{k+1} \cdot (1-a)^{n-k} + \sum_{k=0}^{n-1} \frac{n!}{k! \cdot (n-k+1)!} \cdot a^k \\ &\quad \cdot (1-a)^{n+1-k} - a^n. \end{aligned}$$

Notice that

$$\begin{aligned} \sum_{k=0}^{n-1} \frac{n! \cdot (n+1-k)}{k! \cdot (n+1-k)!} \cdot a^{k+1} \cdot (1-a)^{n-k} \\ = a \cdot \sum_{k=0}^n \frac{n!}{k! \cdot (n-k)!} \cdot a^k \cdot (1-a)^{n-k} = a \cdot (a + (1-a))^n = a \end{aligned}$$

and

$$\begin{aligned} \sum_{k=0}^{n-1} \frac{n!}{k! \cdot (n-k+1)!} \cdot a^k \cdot (1-a)^{n+1-k} \\ = \frac{1}{n+1} \sum_{k=0}^{n+1} \frac{(n+1)!}{k! \cdot (n+1-k)!} \cdot a^k \cdot (1-a)^{n+1-k} \\ = \frac{1}{n+1} (a + (1-a))^{n+1} = \frac{1}{n+1}. \end{aligned}$$

Thus,

$$D_n = a + \frac{1}{n+1} - a^n,$$

completing the proof. \square

References

- Bearden, J.N., 2006. A new secretary problem with rank-based selection and cardinal payoffs. *Journal of Mathematical Psychology* 50, 58–59.
- Bruss, F.T., 2005. What is known about Robbins' problem? *Journal of Applied Probability* 42, 108–120.
- Bruss, F.T., Ferguson, T.S., 1993. Minimizing the expected rank with full information. *Journal of Applied Probability* 30, 616–626.
- de Carvalho, M., Chaves, L.M., de Abreu Silva, R.M., 2008. Variations of the secretary problem via game theory and linear programming. *INFOCOMP Journal of Computer Science* 7, 78–82.
- Ferguson, T.S., 1989. Who solved the secretary problem? *Statistical Science* 4, 282–296.
- Gilbert, J., Mosteller, F., 1966. Recognizing the maximum of a sequence. *Journal of the American Statistical Association* 61, 35–73.
- Gnedin, A.V., Krengel, U., 1995. A stochastic game of optimal stopping and order selection. *The Annals of Applied Probability* 5, 310–321.
- Moser, L., 1956. On a problem of Cayley. *Scripta Mathematica* 22, 289–292.