Competition and Recall in Selection Problems

Gensbittel Fabien[1], Pizarro Dana[2*] and Jérôme Renault[1]

[1]Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France.

[2]Instituto de Ciencias de la Ingeniería, Universidad de O'Higgins, Rancagua, Chile.

*Corresponding author(s). E-mail(s): dana.pizarro@uoh.cl; Contributing authors: fabien.gensbittel@tse-fr.eu; jerome.renault@tse-fr.eu;

Abstract

We extend the prophet inequality problem to a competitive setting. At every period, a new realization of a random variable with a known distribution arrives, which is publicly observed. Then two players simultaneously decide whether to pick an available value or to pass and wait until the next period (ties are broken uniformly at random). As soon as a player gets a value, he leaves the market and his payoff is the value of this realization. In the first variant, namely the "no recall" case, the agents can only bid at each period for the current value. In a second variant, the "full recall" case, the agents can also bid for any of the previous realizations which has not been already selected. For each variant, we study the subgame-perfect Nash equilibrium payoffs of the corresponding game. More specifically, we give a full characterization in the full recall case and show in particular that the expected payoffs of the players at any equilibrium are always equal, whereas in the no recall case the set of equilibrium payoffs typically has full dimension. Regarding the welfare at equilibrium, surprisingly the best equilibrium payoff a player can have may be strictly higher in the no recall case. However, the sum of equilibrium payoffs is weakly larger when the players have full recall. Finally, we show that in the case of 2 arrivals and arbitrary distributions, the prices of Anarchy and Stability in the no recall case are at most 4/3, and this bound is tight.

Keywords: Optimal stopping, Competing agents, Recall, Prophet inequalities, Subgame-perfect equilibria, Game theory, Job search problem, House selling problem.

1 Introduction

1.1 Context

The theory of optimal stopping has a vast history and is concerned with the problem of a decision-maker who observes a sequence of random variables arriving over time and has to decide when to stop optimizing a particular objective. Probably, the two best-known problems in optimal stopping are the secretary problem (also known as marriage problem or best choice problem) and the prophet inequality problem (also known as house selling problem or job search problem). In the classical model of the former, introduced in the '60s, a decision-maker observes a sequence of values arriving in a random order over time and has to pick one in a take-it-or-leave-it fashion maximizing the probability of picking the highest one. In other words, after observing an arrival, he has to decide whether to pick this value or to pass and continue observing the sequence. Once a value is picked the game ends. The goal of the decision-maker is to maximize the probability of getting the highest value. Lindley (1961) proves that an optimal stopping rule for this problem consists in rejecting a particular amount of values first and then accepting the first value higher than the maximum observed so far. When the number of arrivals goes to infinity, the probability of picking the best value approaches $1/e$. Since then, several variants have been studied in the literature (see, e.g., Ezra et al 2021a; Freeman 1983; Gilbert and Mosteller 1966; Immorlica et al 2006). For a survey of the secretary problem with one decision-maker and their variants see, e.g., Freeman (1983) and Ferguson (1989) .

Related to the secretary problem is the optimal selection problem called prophet inequality problem. Here, the decision-maker knows not only the total number of arrivals but also the distribution behind them. The goal of the decision-maker is to maximize the expected value of what he gets. Instead of looking at this problem as an optimal stopping problem, in the '70s researchers started to answer the question of how good can a decision-maker play compared to what a prophet can do, where a prophet is someone who knows all the realizations of the random variables in advance and simply picks the maximum. It was in the '70s when Krengel and Sucheston (1978) proved that the decision-maker can get at least $1/2$ of what a prophet gets and that this bound is tight. Later, in 1984 Samuel-Cahn (1984) proved that instead of looking at all feasible stopping rules, it is enough to look at a single threshold strategy to get the $1/2$ bound. These results are for a general setting where the random variables are independent but not necessarily identically distributed, and then one natural question that arose was if this bound could be improved assuming i.i.d. random variables. The works by Hill and Kertz (1982), Kertz (1986), and Correa et al (2017) answered this question positively, obtaining a tight bound of 0.7451. A lot of work has appeared considering different model variants (infinitely many arrivals, feasibility constraints, multi-selection, etc) but it was in the last decades when this problem gained particular attention due to its surprising connection with online mechanisms (see, e.g., Chawla et al 2010; Correa et al 2019; Hajiaghayi et al 2007).

1.2 Our paper

In the i.i.d. setting of the prophet inequality problem, there is a decision-maker who sequentially observes realizations $X_1, \ldots, X_n$ from a known distribution F. Only one value can be selected, and the benchmark is $\mathbb{E}(\max_{i \in \{1,\ldots,n\}} X_i)$. After each arrival, the decision-maker must make an immediate and irrevocable decision whether to pick this value or not. If he selects it, he leaves the market and the game ends. Otherwise, he continues observing the next arrival. This problem is commonly motivated by applications to online auction theory, where at each time period a seller, who wants to sell an item, receives an offer from a potential buyer and has to decide whether to accept it and sell the item or reject it and wait for the next offer. Most of the work on the prophet inequality problem relies on the fact that there is only one decision-maker and the decision must be taken immediately after observing the offer. However, in some situations of interest– a person who wants to sell a house, or a company hiring an employee, among others– it seems reasonable to allow more than one decision-maker, as well as to be able to make the decision later in time. Driven by this fact, in this paper we study two variants of the classic setting. More specifically, on one hand, we consider the setting where two decision-makers compete to get the best possible value, and we call it *competitive selection problem with no recall*. On the other hand, we consider the problem where the two decision-makers are allowed to select any available value that appeared in the past and not only the one just arrived, and we call it *competitive selection problem with full recall*.

For each variant, we study the two-player game induced by the optimal stopping problem, focusing on subgame-perfect Nash equilibria (SPE for short). The main contributions of this paper can be divided into three lines, which are summarized in what follows.

Description of the subgame-perfect equilibrium payoffs. The first stream of contributions refers to the study of the set of SPE payoffs (SPEP) for both settings. Regarding the full recall case, we fully characterize the set of SPEP (Theorems 1 and 2) and we obtain that each such payoff is symmetric, meaning that every SPE gives the same payoff to both players. In the no-recall case, the set of SPEP is clearly symmetric with respect to the identity line (which we call diagonal throughout the paper) but contains points outside it, inducing different payoffs for the players. In this case, we give in Theorem 3 recursive formulas to compute the best and the worst SPEP (in terms of the sum of payoffs of the players), as well as the worst payoff a single player can get at a subgame-perfect equilibrium. We highlight here that the analysis of SPE in both variants is based on auxiliary games which exhibit different structures, namely coordination games in the model with recall and anti-coordination games in the variant without recall (see section 3.3). To illustrate these results, in Section 3.5 we provide a detailed study of the particular case where F is the uniform distribution on $[0, 1]$.

Comparison of payoffs for the two variants. Another stream of results is focused on comparing the highest SPEP in both settings. Surprisingly, an example (see Section 3.1) shows that the best SPEP a player can obtain may be higher in the no recall case than in the full recall case. This can be explained by the possible existence of an asymmetric SPE in the no recall case which significantly favors one player at the expense of the other, who picks a value early in the game. However, if we restrict our attention to symmetric SPEP, or to the sum of the payoffs of the players at SPE, Theorem 4 shows that in the case of a continuous distribution, players are always better off in the variant with recall. Furthermore, this advantage can be significant: If F is the uniform distribution on $[0, 1]$, and $n = 5$, the payoff of players in the full recall case is at least 7% higher than in the no recall case.

Efficiency of equilibria. To analyze efficiency, we use in Section 3.4 the standard notions of Price of Anarchy and Price of Stability. We also introduce the notion of Prophet Ratio, defined as the sum of payoffs a team of two prophets would obtain, divided by the best sum of payoffs at equilibrium. When F is uniform on $[0, 1]$, we find numerically that equilibria are quite efficient. Moreover, we observe that having full recall gives an advantage to the players, in the sense that, in the full recall case, the expected payoff is closer to the payoff obtained by playing the best possible feasible strategy, than in the no recall case. However, this advantage is small (see Section 3.5.3). Finally, we consider the case $n = 2$ without recall, and let F be any distribution with support contained in [0,1]. We prove in Theorem 5 that both the prices of anarchy and stability are not greater than 4/3 and that this bound (reminiscent of the bound for routing problems with linear latencies, see Roughgarden and Tardos 2002) is tight in both cases.

It is worth noting that, although competition and recalling variants have already been considered in the literature for some optimal stopping problems (as we will discuss in the next section), to the best of our knowledge, our paper is the first considering them from a Game Theory perspective with a focus on the study of the set of SPEP, and thus it constitutes a good starting point for future research work of interest across academic communities in Operations Research, Computer Science, and Economics (see Section 5 for details in possible future directions).

1.3 Related literature

As it was aforementioned, the literature in optimal stopping theory is extensive and mainly focused on finding optimal or near-optimal policies for the different model variants, as well as on studying the guarantees of some simple strategies, such as single threshold strategies, even when they are not optimal.

However, this paper introduces a game-theoretic approach for a model with competition and where recall is allowed. In what follows, we revisit some of the existing literature regarding optimal stopping problems with some of these two particular features.

Optimal Stopping with competition. Abdelaziz and Krichen (2007) survey the literature on optimal stopping problems with more than one decision-maker until the 2000s. More recently, Immorlica et al (2006), Karlin and Lei (2015) and Ezra et al (2021a) study the secretary problem with competition. Immorlica et al. consider a classical setting where decisions are made in a take-it-or-leave-it fashion and ties are broken uniformly at random, and they show that as the number of competitors grows, the moment at which "accept" is played for the first time in an equilibrium decreases. Karlin and Lei assume that the ties are broken according to a global ranking and they derived a recursive formula to compute the optimal strategy for the players. Ezra et al., incorporate the recalling option and studies the structure and performance of equilibria in this game when the ties are broken uniformly at random or according to a global ranking.

Our paper considers a different model, since the problem is more related to the prophet inequality setting. In this sense, the work closest to ours is the recent paper by Ezra et al (2021b), who introduce, independently of our paper, the no recall case, again with the two variants for tie-breaking. However, the novelty of our paper is the incorporation of recalling in addition to the competition between players. Moreover, instead of studying the reward guarantees under single-threshold strategies, we focus on the study of equilibria of the game.

There exists another stream of related literature considering selection problems with competition from a game theory perspective. The more relevant for our work are Baston and Garnaev (2005); Sakaguchi (2005); and Ramsey and Szajowski (2006). In the first two papers, the authors consider the game with competition, no recall, and uniform distribution on $[0,1]$. In the latter, arrivals are considered to follow a Markov chain, and the authors allow for a correlation between players. The main difference between these papers and ours is that they restrict their study to particular types of Nash equilibria (or correlated equilibria for Ramsey and Szajowski (2006)). For instance, this set does not include the best symmetric equilibrium payoff of the game, which is characterized in our paper. To the best of our knowledge, we are the first to study the set of all SPEP, and not only a particular finite subset of them. We refer the reader to the works by Mazalov and Banin (2003); Radzik and Szajowski (1990); and Sakaguchi and Mazalov (2004) for other related papers.

Optimal Stopping with recall. Allowing decision-makers to choose between any of the values arrived so far is a variant of the classic problem that may have interesting applications. If this extension is considered without competition, it is easy to see that the optimum is just to wait until the end and pick the best value. However, adding competition to the model makes the problem interesting, and characterizing the set of equilibrium payoffs and studying their efficiency are challenging questions we address in this paper. Notice that this notion of recalling is not new for some optimal stopping problems. For example, Yang (1974) considers a variant of the secretary problem, where the interviewer is allowed to make an offer to any applicant already interviewed. In his model, the applicant rejects the offer with some probability that depends on when the offer is made and he studies the optimal stopping rules in this context. Thereafter, different authors have been studying other variants of the secretary problem with recall (see, e.g., Ezra et al 2021a; Kesselheim et al 2023; Sweet 1994). Our work differs from most of them not only in the model we consider but also in terms of questions since we are more interested in a game-theoretic approach to the problem.

1.4 Roadmap

The remainder of this paper is organized as follows: We start presenting the model in Section 2, including the description of the games for the full recall and no recall cases in Section 2.1. The results of our paper are presented formally in Section 3. This section is divided into four parts: In Section 3.2, we study the set of SPEP, starting with its computation for a very simple example (see Section 3.1), and making then a complete analysis for the full recall case (see Section 3.2.1) and for the no recall case (see Section 3.2.2). In Section 3.3 we present the results concerning the comparison of the two variants, whereas in Section 3.4 we study the efficiency of equilibria. In Section 3.5 we make a detailed analysis of the case where F is uniform on $[0, 1]$. In Section 4 we present the proofs of the results. Finally, we close the paper with some possible future directions, which will be given in Section 5.

2 Model

Consider a sequence of realizations X_1,...,X_n of a random variable X distributed according to a c.d.f. F with support contained in [0,1].[1] There are two players, or decision-makers, competing to select the highest value among X_1,...,X_n. More explicitly, at each time period $t = 1, ..., n$, the decision-makers – who know the distribution F– observe X_t and simultaneously decide whether or not to select one element in the currently *feasible set* $\mathcal{F}_t$. Once a decision-maker chooses a value X_j from stage j, he leaves the market obtaining a payoff of X_j and the element of stage j no longer belongs to the feasible set. Decision-makers must decide when to stop, maximizing the expected value of their payoff. If at time t both agents want to select the same value, we break

[1]We assume that the support of F is contained in $[0, 1]$ to not overload the notation but the results can be easily extended to the general case.

the tie uniformly at random. That is, each of them gets the value with probability 1/2, and the decision-maker who gets it leaves the market, whereas the other passes to the next period.

It remains to specify the feasible sets, distinguishing the two cases we consider. In the first one, namely *no recall case*, the decision-makers play in a take-it-or-leave-it fashion. That is, after observing the value just arrived they should decide whether they take it or not. If both players discard it, it cannot be chosen later. The second variant is the *full recall case*. In this case, the discarded values can be chosen in future periods if they are still available. That is, at each time period, the decision-makers can choose the element that has just arrived or any of the elements that have already arrived and have not yet been chosen. We assume that at stage n all players who are present select the element corresponding to the highest value. If both players are present, there is a tie, and the player who loses the toss gets the second best value.

Notice that in the 1-player case (decision problem), the optimal strategy with full recall is simply to wait until the end and to pick the element which corresponds to the maximum of $\{X_1, ..., X_n\}$. And in the 1-player case with no recall we have a standard prophet problem, with value smaller than the expectation of the maximum of $\{X_1, ..., X_n\}$. Obviously, having the possibility of getting a value observed in the past is beneficial to the player. In our multi-player setting, we cannot say *a priori* that the full recall case is beneficial to the players, as there are examples where having more information or more actions decreases the sum of the payoffs of the players at equilibrium. This motivated us to ask how important is to have the power of being able to choose a value observed in the past. To answer this question we study the games behind the two model variants described above. In particular, we study the set of SPEP as well as the Price of Anarchy (PoA) and the Price of Stability (PoS).

We remark here that throughout the paper we will use "the decision-maker selects value X_i" to refer to the decision-maker selecting element i from $\mathcal{F}_t$.

2.1 Description of the games

We now formally describe the games induced by the full recall case and the no recall case, denoted by Γ_n^{FR} and Γ_n^{NR}, respectively. Then, we specify the notions of equilibrium that will be used throughout the paper.

Full recall case. Game Γ_n^{FR}

For each $t \in \{0, 1, ..., n+1\}$ we denote by $\mathcal{H}_t$ the set of possible histories up to stage t. $\mathcal{H}_0$ only contains the empty history. $\mathcal{H}_1$ contains what happens at stage 1, i.e., the realization X_1, who tried to get X_1, and who got X_1 (possibly nobody). $\mathcal{H}_2$ contains everything that happened at stages 1 and 2, and so on.

As usual a strategy for player $j \in \{1, 2\}$ is an element $\sigma_j = (\sigma_{j,t})_{t=0,\ldots,n}$ where $\sigma_{j,t}$ is a measurable map that associates to every history in $\mathcal{H}_t$ an available action, that is a probability distribution over $\mathcal{F}_t \cup \{\emptyset\}$. We will assume that in a strategic pair, (σ_1, σ_2), the first element σ_1 is for the row player, Player 1, and the second element σ_2 is for the column player, Player

2. A strategy profile (σ_1, σ_2) induces a probability distribution over the set of possible full histories $\mathcal{H}_{n+1}$, and the payoff (or utility) of each player is defined as the expectation of the value he gets, with the convention that getting none of the values yields a payoff of 0.

No recall case. Game Γ_n^{NR}

Here, the set of available elements at stage t is $\mathcal{F}_t = \{t\}$, and we only need to consider histories $\mathcal{H}_t$ for $t \in \{0, ..., n\}$ where $\mathcal{H}_t$ contains everything that happened up to stage t under the no recall assumption. A strategy for player $j \in \{1, 2\}$ is an element $\sigma_j = (\sigma_{j,t})_{t=0,...,n-1}$ where $\sigma_{j,t}$ is a measurable map which associates to every history in $\mathcal{H}_t$ an available action, that is a probability distribution over $\{t, \emptyset\}$. We will assume that in a strategic pair, (σ_1, σ_2), the first element σ_1 is for the row player, Player 1, and the second element σ_2 is for the column player, Player 2. A strategy profile (σ_1, σ_2) induces a probability distribution over the set of possible plays $\mathcal{H}_n$, and payoffs are defined as in the full recall case.

Equilibrium notions. We recall here the usual notions of Nash equilibrium (NE) and subgame perfect equilibrium (see, e.g., Fudenberg and Tirole 1991). The following definitions apply to both games Γ_n^{FR} and Γ_n^{NR}.

Definition 1 *A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a Nash equilibrium (NE) of the game if for every agent i and every strategy σ_i', player's i utility when (σ_i', σ_{-i}) is played is not greater than the one obtained if σ is played.*

Given a stage number t and a finite history h_t mentioning everything that happened up to stage t, we can define the continuation game after h_t.

Definition 2 *A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a subgame perfect equilibrium (SPE) if it induces a NE for every proper subgame of the game (i.e. for any continuation game after a finite history).*

When studying SPE, we will assume w.l.o.g. that as soon as a player is alone in the game, he plays optimally in the remaining decision problem. By best (resp. worse) equilibrium, we mean an equilibrium maximizing (resp. minimizing) the sum of the payoffs of the players.

2.2 Notation

In this section, we introduce some notation that we will use throughout the paper.

Given 3 real numbers x_1, x_2 and x_3, we denote the median of the numbers, i.e. the middle-most number among x_1, x_2, x_3, by $med[x_1, x_2, x_3]$.

For any event A, the indicator function of A, denoted by $\mathbb{1}_A$ is the random variable taking value 1 if the event A occurs, and 0 otherwise.

Given two real numbers x and y, we will denote by $x \vee y$ the maximum between x and y, and by $x \wedge y$ the minimum. That is, $x \vee y = \max\{x, y\}$, and $x \wedge y = \min\{x, y\}$.

3 Results

The goal of this section is to present the main results of the paper. First, we study the structure of the sets of SPEP. We start by introducing a simple example, where the computation for both the recall and no recall cases is easy. Then, we present our main result in the full recall case which fully characterizes the set of SPEP. Regarding the no recall case, we provide recursive formulas to compute the worst payoff a player can get in equilibrium, as well as the sum of payoffs of players for both the best and worst SPE, when the distribution function F is continuous. After that, we move on to understand what is the relationship between the payoff a player gets in both problems. In other words, we answer the question: Is the recall case always beneficial for the players? Finally, we study the efficiency of equilibria for the full recall and no recall cases. We leave the proof of the results in Section 4.

3.1 Motivating example.

Let us consider the particular instance of the problem where n realizations of X arrive sequentially over time, where the law F of X is defined by:

$$X = \begin{cases} 1/3 & \text{with probability } 1/2, \\ 2/3 & \text{with probability } 1/2. \end{cases}$$

We compute the set of SPEP in the full recall and no recall case, denoted by E_n^{FR} and E_n^{NR}, respectively.

Full recall case. Note that in this case, the unique SPE is to bid at any time if $X = 2/3$ and to pass otherwise until the last period, when the best available value is chosen. Then, if all observations take the value $1/3$, the players obtain a payoff of $1/3$; if only one value is equal to $2/3$, both players bid for this value and then they obtain $2/3$ with probability $1/2$ and $1/3$ with probability $1/2$; and if more than one realization of X has value $2/3$, both players obtain a payoff of $2/3$. Thus, the expected payoff of a player in equilibrium, namely h_n, can be computed as follows

$$h_n = \frac{1}{3}\mathbb{P}(X_{(n)} = 1/3) + \frac{1}{2}\mathbb{P}(\exists!\ i : X_i = 2/3) + \frac{2}{3}\mathbb{P}(\exists i \neq j : X_i = X_j = 2/3),$$

where $X_{(n)}$ denotes the maximum of n i.i.d. realizations of $X \sim F$. The probabilities above are easy to compute, having that $\mathbb{P}(X_{(n)} = 1/3) = 1/2^n$, $\mathbb{P}(\exists!\ i : X_i = 2/3) = n/2^n$ and $\mathbb{P}(\exists i \neq j : X_i = X_j = 2/3) = 1 - (n + 1)/2^n$.

Putting all together we conclude that $E_n^{FR} = \{(h_n, h_n)\}$, with

$$h_n = \frac{2}{3} - \frac{n+2}{3 \cdot 2^{n+1}}.$$

No recall case. Because we are considering an instance with n time periods and the random variable only takes two possible values, for the first $n-2$ arrivals, at equilibrium, players bid if and only if a value $2/3$ is observed. Therefore, at equilibrium: if at least two observations have value $2/3$ up to time $n-2$, players obtain $2/3$; if only one has value $2/3$ up to time $n-2$, each player obtains $2/3$ with probability $1/2$ and the expected value of $X_{(2)}$ with probability $1/2$; and if all values up to $n-2$ are $1/3$, the players obtain an expected payoff $(e_n^1, e_n^2) \in E_2^{NR}$. Note that $\mathbb{E}(X_{(2)})$ is the value of the decision problem in a standard prophet setting with two arrivals from X, and thus $\mathbb{E}\left(X_{(2)}\right) = 7/12$. Therefore, an expected payoff of player 1 can be computed as

$$\begin{aligned}
&\mathbb{P}\left(X_{(n-2)} = 1/3\right) e_2^1 + \mathbb{P}(\exists!\ i \leq n-2 : X_i = 2/3)\frac{15}{24} \\
&+\mathbb{P}(\exists i \neq j : X_i = X_j = 2/3, i, j \leq n-2)\frac{2}{3},
\end{aligned}$$

whereas for player 2 we have

$$\begin{aligned}
&\mathbb{P}\left(X_{(n-2)} = 1/3\right) e_2^2 + \mathbb{P}(\exists!\ i \leq n-2 : X_i = 2/3)\frac{15}{24} \\
&+\mathbb{P}(\exists i \neq j : X_i = X_j = 2/3, i, j \leq n-2)\frac{2}{3}.
\end{aligned}$$

Again, the probabilities above are easy to compute, having $\mathbb{P}\left(X_{(n-2)} = 1/3\right) = 1/2^{n-2}$, $\mathbb{P}(\exists!\ i \leq n-2 : X_i = 2/3) = (n-2)/2^{n-2}$ and $\mathbb{P}(\exists i \neq j : X_i = X_j = 2/3, i, j \leq n-2) = 1 - (n-1)/2^{n-2}$. Thus, we obtain that an expected payoff of player i is given by:

$$\frac{e_2^i}{2^{n-2}} + \frac{15}{24} \cdot \frac{n-2}{2^{n-2}} + \frac{2}{3} \cdot \left(1 - \frac{n-1}{2^{n-2}}\right) = \frac{2}{3} + \frac{e_2^i}{2^{n-2}} - \frac{n+14}{3 \cdot 2^{n+1}}. \tag{1}$$

It remains to compute the set E_2^{NR}. To this end, for $a \in \{1/3, 2/3\}$ we consider the game, called $\Gamma_1^{NR}(a)$, defined as Γ_1^{NR} but with an initial available value a. That is, there is a time period "zero" where players choose between selecting the value a or pass, before observing the value of the single realization in the game. We denote by $E_1^{NR}(a)$ the set of SPEP of the game $\Gamma_1^{NR}(a)$. Notice that

$$E_2^{NR} = \left\{(\gamma_1, \gamma_2) : \gamma_i = \sum_{a \in \{1/3, 2/3\}} \mathbb{P}(X = a) e_1^i(a), \text{s.t. } \left(e_1^1(a), e_1^2(a)\right) \in E_1^{NR}(a)\right\}, \tag{2}$$

and then E_2^{NR} can be easily computed from $E_1^{NR}(a)$. To study $E_1^{NR}(a)$, we consider the payoff matrix for the game $\Gamma_1^{NR}(a)$, represented in Fig. 1.

$$\begin{array}{c c c} & a & \emptyset \\ \begin{array}{c} a \\ \emptyset \end{array} & \multicolumn{2}{c}{\begin{pmatrix} \left(\frac{1}{2}a+\frac{1}{4},\frac{1}{2}a+\frac{1}{4}\right) & \left(a,\frac{1}{2}\right) \\ \left(\frac{1}{2},a\right) & \left(\frac{1}{4},\frac{1}{4}\right) \end{pmatrix}} \end{array}$$

Fig. 1: Payoff matrix for the game $\Gamma_1^{NR}(a)$.

Observe that if $a = 2/3$, (a, a) is the unique NE with payoff $(7/12, 7/12)$, and thus $E_1^{NR}(2/3) = \{(7/12, 7/12)\}$. Otherwise, if $a = 1/3$, there are three NE: $(a, \emptyset), (\emptyset, a)$ and a symmetric mixed equilibrium in which both agents play a with probability $1/2$ and pass with probability $1/2$. The equilibrium payoffs are $(1/3, 1/2), (1/2, 1/3)$, and $(3/8, 3/8)$, respectively, and then $E_1^{NR}(1/3) = \{(1/3, 1/2), (1/2, 1/3), (3/8, 3/8)\}$. We now can compute the set E_2^{NR} by using (2), obtaining

$$E_2^{NR} = \{(11/24, 13/24), (13/24, 11/24), (23/48, 23/48)\}.$$

Finally, we go back to equations (1) and we conclude that E_n^{NR} is the three-element set

$$E_n^{NR} = \{(p_n, q_n), (r_n, r_n), (q_n, p_n)\}, \tag{3}$$

with

$$p_n := \frac{2}{3} - \frac{n+1}{3 \cdot 2^{n+1}}, \quad q_n := \frac{2}{3} - \frac{n+3}{3 \cdot 2^{n+1}} \quad \text{and} \quad r_n := \frac{2}{3} - \frac{n+5/2}{3 \cdot 2^{n+1}}.$$

Note that p_n and q_n represent the best and worst possible payoff of a player in equilibrium, respectively, whereas r_n represents the payoff in between.

Recall that $h_n = 2/3 - (n+2)/(3 \cdot 2^{n+1})$ and therefore $p_n > h_n > r_n > q_n$. Moreover $p_n + q_n = 2h_n > 2r_n$. The latter means that the sum of equilibrium payoffs of the players in the full recall case is equal to the sum of the payoffs of players in the no recall case if an asymmetric equilibrium is played. Otherwise, i.e., if in the no recall case the symmetric equilibrium is played, then the sum of payoffs of the players is strictly lower than in the full recall case, which is somehow not surprising.

However, for this example we observe that, surprisingly, the best payoff a single player can obtain at equilibrium is strictly higher in the no recall case than in the full recall case. Intuitively, what happens here is that the lack of recalling power implies that with a positive probability one player will take the "bad" value in period $n-1$ and therefore it gives an advantage to the other player, who is now alone in the game.

A natural question is to ask how efficient are the equilibria of the games compared with the best feasible strategy. We use the well-known notions of Price of Stability (PoS) and Price of Anarchy (PoA) which are defined as the ratio between the maximal sum of payoffs obtained by players under any feasible strategy and the best and the worst sum of payoffs of SPEs respectively. For the example, this analysis is easy to do, and we start by noting that in both model variants, there exists a feasible strategy (not the same for both problems) where players obtain the first and the second best values. This implies that the numerator of PoA and PoS will be the same and equal to $\mathbb{E}(X_{(n)} + X_{(2:n)})$, where $X_{(2:n)}$ denotes the second best realization, which is given by $4/3 - (n+2)/(3 \cdot 2^n)$. Observe that this value is exactly $2h_n$ and $p_n + q_n$ and then PoS is 1 for both settings and PoA is 1 for the full recall case, meaning that in the full recall case the resulting equilibrium is efficient. On the other hand, the PoA for the no recall case is given by $\mathbb{E}(X_{(n)} + X_{(2:n)})/(2r_n)$ which is strictly greater than 1 but it converges fast to one when n grows.

This example highlights the relevance of our work in three ways. First, even in this case where the random variable only takes two possible values, we observe that the computation of the SPEP requires a detailed analysis and then computing these sets for a general probability distribution, even under some mild assumptions, seems difficult and challenging. On the other hand, this example shows us that the recalling power is not always favorable to all players, and then it is an interesting question to try to understand under which assumptions one can ensure that the payoff of a player in the full recall case in equilibria will be at least as good as the one in the no recall case. Finally, we see in this example that the SPEP are efficient and it motivates us to study if that still holds for general distributions.

3.2 Description of the subgame-perfect equilibrium payoffs

Here we will fully characterize the set of SPEP for the full recall case. For the no recall case, we provide recursive formulas to compute the best and the worst sum of SPEP. We prove these results in Section 4.1.

3.2.1 Full recall case

We go back to the general case and study here the game with full recall, that is, at time t, any of the values so far arrived that have not been selected before can be selected.

We introduce a two-player game $\Gamma_n^{FR}(a,b)$, where for each natural number n and $1 \geq a \geq b \geq 0$, $\Gamma_n^{FR}(a,b)$ is defined as Γ_n^{FR} with n values to arrive and a and b two available values present at the beginning of the game. That is, we have a time period "zero" where players choose between getting a, b or pass, before the sequential arrival of the realizations. We denote by $E_n^{FR}(a,b) \subset \mathbb{R}_+^2$ the set of the SPEP of the game $\Gamma_n^{FR}(a,b)$. Note that the set of the SPEP of the game Γ_n^{FR} is just the set $E_n^{FR}(0,0)$.

The next theorem states that the set of SPEP for the defined auxiliary game $\Gamma_n^{FR}(a,b)$ is contained in the diagonal, i.e., at each SPE both players get the same payoff. We see this as a surprising result.

Theorem 1. *Consider an instance of the game $\Gamma_n^{FR}(a,b)$, for a,b real numbers such that $0 \leq b \leq a \leq 1$ and $n \in \mathbb{N}$. The set of SPE payoffs is contained in the diagonal, that is $E_n^{FR}(a,b) \subset \{(u,u), u \in [0,1]\}$.*

Furthermore, if we define $PE_n^{FR}(a,b)$ as the projection of $E_n^{FR}(a,b)$ to $\mathbb{R}$, we have $\min\ PE_n^{FR}(a,b) = l_n(a,b)$ and $\max PE_n^{FR}(a,b) = h_n(a,b)$, where l_n and h_n are defined recursively as follows:

- $l_0(a,b) = h_0(a,b) = \frac{a+b}{2}$;
- *for* $n \geq 1$*:*
  - $l_n(a,b) = L\left(a, \mathbb{E}\left(X_{(n)} \vee b\right), d_n^-(a,b)\right)$, *with* $d_n^-(a,b) = \mathbb{E}_X\left(l_{n-1}(a \vee X, med[a,b,X])\right)$,
  - $h_n(a,b) = H\left(a, \mathbb{E}\left(X_{(n)} \vee b\right), d_n^+(a,b)\right)$ *with* $d_n^+(a,b) = \mathbb{E}_X\left(h_{n-1}(a \vee X, med[a,b,X])\right)$,

  where $X_{(n)} = \max\{X_1, \ldots, X_n\}$, med denotes the median, and $L : \mathbb{R}^3 \to \mathbb{R}$ and $H : \mathbb{R}^3 \to \mathbb{R}$ are defined by:

$$L(x,y,z) = \begin{cases} z & \text{if } x \leq y, \\ \frac{1}{2}(x+y) & \text{if } x > y, \end{cases} \quad \text{and} \quad H(x,y,z) = \begin{cases} \frac{1}{2}(x+y) & \text{if } x > y \vee z, \\ z & \text{if } x \leq y \vee z. \end{cases}$$

Using the theorem above, we can prove the following result which fully characterizes the set of SPEP of the games Γ_n^{FR} when F is continuous (i.e. when the corresponding distribution is atomless).

Theorem 2. *Assume F is continuous. Then, for $n \geq 1$, the set E_n^{FR} of SPE payoffs of the game Γ_n^{FR} is the segment:*

$$E_n^{FR} = \{\lambda(l_n, l_n) + (1-\lambda)(h_n, h_n), \lambda \in [0,1]\},$$

where $l_n = l_n(0,0)$ and $h_n = h_n(0,0)$.

That is, E_n^{FR} is convex, contained in the diagonal, and its extreme points are $(l_n(0,0), l_n(0,0))$ and $(h_n(0,0), h_n(0,0))$, where $l_n(0,0)$ and $h_n(0,0)$ are defined in the statement of Theorem 1.

3.2.2 No recall case

We now consider the no recall variant, where players can only play in a take-it-or-leave-it fashion, without being able to select a value that arrived in the past. In this section, we assume that F is continuous, which ensures that the set of SPEP is convex, and allows us to derive recursive formulas for some points in the boundary of this set.

We introduce here the two-player game $\Gamma_n^{NR}(a)$, where for each natural number n and $a \in [0,1]$, $\Gamma_n^{NR}(a)$ is defined as Γ_n^{NR} with n realizations to arrive, but with a an available value present at the beginning. That is, we have a time period "zero" where players choose between getting a or pass, before the sequential arrival of the values. Calling $E_n^{NR}(a) \subset \mathbb{R}_+^2$ the set of SPEP of the game $\Gamma_n^{NR}(a)$, we have that the set of the SPEP of Γ_n^{NR} is just $E_n^{NR} := E_n^{NR}(0) = \mathbb{E}_{a\sim F}(E_{n-1}^{NR}(a))$, where the latter is defined by

$$\left\{ \int_0^1 f(a)\mathrm{d}F(a), f : [0,1] \to \mathbb{R}^2 \text{ measurable with } f(a) \in E_{n-1}^{NR}(a) \text{ for each } a \right\}.$$

Below we present a technical result, stating that the set of SPE payoffs for the no recall case is symmetric with respect to the diagonal, and when F is continuous, it is also convex and compact.

Proposition 1 *For each natural number n, the set of the SPE payoffs E_n^{NR} is symmetric with respect to the diagonal. If F is continuous, E_n^{NR} is convex compact.*

The main result of this section is Theorem 3 below, where we give explicit recursive formulas to compute the sum of the SPE payoffs for the best and worst equilibria under the no recall setting. Recall that by best (resp. worst) equilibrium we mean a SPE that maximizes (minimizes) the sum of the payoffs of the 2 players. We will see that the recursive structure of the SPEP in the no recall case is more complex than in the full recall case. For instance, although the set of SPEP for the no recall case is convex and symmetric with respect to the diagonal, it is not a subset of the diagonal. Moreover, the proof of the theorem shows that to achieve the smallest sum of equilibrium payoffs, the players need to consider continuation equilibrium payoffs which do *not* minimize the sum of the payoffs of the players.

Before presenting the result, we introduce some notation. We first define by induction (with $X \sim F$):

$$c_1 = \mathbb{E}(X), \text{ and } \forall n > 1, \ c_n = \mathbb{E}(X \vee c_{n-1}).$$

Note that c_n is the value of the decision problem in a standard prophet setting.

On the other hand, we denote by α_n (β_n) the smallest (highest) coordinate value of a point on E_n^{NR} belonging to the diagonal, and by α'_n the smallest coordinate of a point belonging to E_n^{NR}. That is, $\alpha_n := \min\{x : (x,x) \in E_n^{NR}\}$, $\beta_n := \max\{x : (x,x) \in E_n^{NR}\}$ and $\alpha'_n = \min\{\min\{x,y\} : (x,y) \in E_n^{NR}\}$ (see Fig. 2). It is easy to see that: $\alpha'_1 = \alpha_1 = \beta_1 = 1/2 \cdot \mathbb{E}(X)$.

Theorem 3. *Assume F is continuous. In the game Γ_n^{NR}, the following hold:*

a) the **worst payoff** *a player can get at equilibrium is α'_n, where for $n > 1$:*

$$\alpha'_n = \frac{c_{n-1}+1}{2} - \int_{\alpha'_{n-1}}^{c_{n-1}} F(a)\mathrm{d}a - \frac{1}{2}\int_{c_{n-1}}^1 F(a)\mathrm{d}a.$$

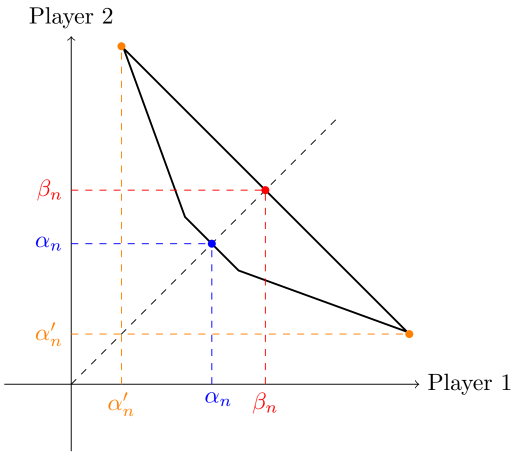

Fig. 2: Representation of α_n, β_n and α'_n if the set E_n^{NR} is given by a trapezoid.

b) the sum of payoffs for the **best SPE** *is* $2\beta_n$*, where for* $n > 1$*:*

$$\begin{aligned}2\beta_n = &\int_0^{\alpha'_{n-1}} 2\beta_{n-1}\mathrm{d}F(a) + \int_{\alpha'_{n-1}}^{\beta_{n-1}} \max\{a + c_{n-1}, 2\beta_{n-1}\}\mathrm{d}F(a)\\ &+ \int_{\beta_{n-1}}^{1} (a + c_{n-1})\mathrm{d}F(a).\end{aligned}$$

c) the sum of payoffs for the **worst SPE** *is* $2\alpha_n$*, where for* $n > 1$*:*

$$\begin{aligned}2\alpha_n = &\int_{c_{n-1}}^{1} (a + c_{n-1})\mathrm{d}F(a) + \int_{\beta_{n-1}}^{c_{n-1}} \frac{4ac_{n-1} - 2\beta_{n-1}(a + c_{n-1})}{c_{n-1} + a - 2\beta_{n-1}}\mathrm{d}F(a)\\ &+ \int_{\alpha_{n-1}}^{\beta_{n-1}} 2a\mathrm{d}F(a) + \int_{\alpha'_{n-1}}^{\alpha_{n-1}} \min\{2\alpha_{n-1}, a + c_{n-1}\}\mathrm{d}F(a)\\ &+ \int_0^{\alpha'_{n-1}} 2\alpha_{n-1}\mathrm{d}F(a).\end{aligned}$$

3.3 Comparison of the two model variants

We now come back to the general case and no longer assume that F is continuous, and want to compare the SPEP obtained by the players with and without recalling power. The proof of the results are in Section 4.2

We have that the best SPE payoff of a player under full recall is $h_n = h_n(0, 0)$ defined recursively in Theorem 1. We now denote by

$$\beta'_n = \max\{\max\{x, y\} : (x, y) \in E_n^{NR}\}$$

the best possible SPEP of a player under no-recall.

We know from the simple example presented in Section 3.1 that we may have: $\beta'_n > h_n$, that is the best possible SPE payoff for a player may be strictly higher under no recall than under full recall. Unsurprisingly this is not a general result, and we will give here one example when $\beta_2' < h_2$ and one example when $\beta_2' = h_2$.

Then, we will consider the maximal sum of the payoffs of the players at equilibrium. In the full recall case, we know that the best SPE payoff is (h_n, h_n) and the worst SPE payoff is (l_n, l_n). In the no-recall case when F is continuous, the SPE payoff set E_n^{NR} is convex, compact, and symmetric with respect to the diagonal, hence the best sum of payoffs is obtained at the symmetric equilibrium (β_n, β_n). Theorem 4 will show that $l_n \geq \beta_n$, implying that when F is continuous, at a symmetric equilibrium payoff, players are always better off under full recall than under no recall.

Example 1 An example with $h_2 > \beta_2'$. Consider the following discrete random variable:

$$X = \begin{cases} 1/10 & \text{with probability } 1/2, \\ 1/2 & \text{with probability } 1/2. \end{cases}$$

We will show that the expected payoff of a player at equilibrium in the full recall case is always higher than the expected payoff of a player at equilibrium in the no recall case, when there are 2 realizations of X arriving sequentially. In other words, we will prove that $h_2 > \beta_2'$.

Let us start with the full recall case. In this setting, the only SPE is to bid in the first stage if and only if $X = 1/2$. Then, the expected payoff of a player playing SPE is given by:

$$\begin{aligned} h_2 &= \frac{1}{10}\mathbb{P}\left(X_{(2)} = \frac{1}{10}\right) + \left(\frac{1}{2}\cdot\frac{1}{2} + \frac{1}{2}\cdot\frac{1}{10}\right)\mathbb{P}\left(\exists! i : X_i = \frac{1}{10}\right) + \frac{1}{2}\mathbb{P}\left(X_{(2:2)} = \frac{1}{2}\right) \\ &= \frac{3}{10}. \end{aligned}$$

In the no recall case, Fig. 3 represents the matrix of expected payoffs for the game if the first arrival is x.

$$\begin{array}{c c} & \begin{array}{cc} \quad\quad x \quad\quad & \quad\quad \emptyset \quad\quad \end{array} \\ \begin{array}{c} x \\ \emptyset \end{array} & \begin{pmatrix} \left(\frac{1}{2}x + \frac{3}{20}, \frac{1}{2}x + \frac{3}{20}\right) & \left(x, \frac{3}{10}\right) \\ \left(\frac{3}{10}, x\right) & \left(\frac{3}{20}, \frac{3}{20}\right) \end{pmatrix} \end{array}$$

Fig. 3: Payoff matrix for the game with $n = 2$ if $X_1 = x$.

Note that if $x = 1/10$, $(\emptyset, \emptyset)$ is the unique NE with payoff $(3/20, 3/20)$, and if $x = 1/2$, (x, x) is the unique NE with payoff $(2/5, 2/5)$. Then, the SPE payoff of each player is $1/2 \cdot 3/20 + 1/2 \cdot 2/5 = 11/40$. This means that $\beta_2' = 11/40$, and we then conclude that $h_2 > \beta_2'$.

Example 2 An example with $h_2 = \beta_2'$. Let us now consider the game with two observations of X arriving sequentially over time, with $X \sim \text{Unif}[0, 1]$.

Using Theorem 1 we can easily obtain that $h_2 = \mathbb{E}(X)$, and then in this case we have $h_2 = 1/2$.

Regarding the no recall case, Fig. 1 represents the expected payoff matrix if the first arrival is x. We now compute the set of NE depending on the value of x:

*C*ase 1. If $x < 1/4$, $(\emptyset, \emptyset)$ is the unique NE and the payoff is given by $(1/4, 1/4)$.

*C*ase 2. If $x > 1/2$, (x, x) is the unique NE with payoff $(x/2 + 1/4, x/2 + 1/4)$.

*C*ase 3. If $x \in [1/4, 1/2]$, $(x, \emptyset)$ and $(\emptyset, x)$ are the pure NE with payoff $(x, 1/2)$ and $(1/2, x)$, respectively, and there is also a mixed NE with payoff $(3/4 - 1/(8x), 3/4 - 1/(8x))$.

Therefore we have that,

$$E_1^{NR}(x) = \begin{cases} \{(1/4, 1/4)\} & \text{if } x < 1/4 \\ \{(x, 1/2), (1/2, x), (3/4 - 1/(8x), 3/4 - 1/(8x))\} & \text{if } x \in [1/4, 1/2] \\ \{(x/2 + 1/4, x/2 + 1/4)\} & \text{if } x > 1/2. \end{cases}$$

Note that the maximum possible expected payoff for one player is obtained if he passes for every $x \in [1/4, 1/2]$, and the payoff obtained is given by

$$\beta_2' = \int_0^{1/4} \frac{1}{4} \mathrm{d}x + \int_{1/4}^{1/2} \frac{1}{2} \mathrm{d}x + \int_{1/2}^{1} \frac{x}{2} + \frac{1}{4} \mathrm{d}x = \frac{1}{2},$$

concluding that $h_2 = \beta_2'$.

Sum of equilibrium payoffs. From the examples above, we conclude that it is not true that at equilibrium players always take advantage of the recalling power. However, the sum of the expected revenues of the players in the full recall case is always at least the one obtained in the no recall case:

Theorem 4. *Assume that F is continuous. Let $(u, u) \in E_n^{FR}$ be a SPE payoff under full recall, and $(v_1, v_2) \in E_n^{NR}$ be a SPE payoff under no recall. Then $2u \geq v_1 + v_2$.*

The proof in Section 4.2 will use the following technical lemma, which gives a lower bound for a SPE of a player in the full recall case.

Lemma 1. *Let a be a positive real number such that $a \leq \mathbb{E}\left(X_{(n+1)}\right)$. If γ_n^{FR} denotes the expected payoff of one player in some SPE in $\Gamma_n^{FR}(a \vee X, a \wedge X)$, then*

$$\gamma_n^{FR} \geq \frac{a + c_{n+1}}{2},$$

where c_n is the value of the decision problem in the no recall case with one decision-maker and n arrivals.

We highlight here that an immediate consequence of Theorem 4 is that in the continuous case if (u, u) and (v, v) are symmetric SPE payoffs in, respectively, the full recall case and no recall case, then $u \geq v$.

Structure of equilibria. The recursive structure of SPE analyzed in the proofs of Theorems 1 and 3 reveals an interesting structural difference between the equilibria in the two model variants. In the model with full recall, the game being played at the first stage after observing the first realization X_1, when both players anticipate a given common continuation equilibrium payoff d if none of them chooses to take X_1, is a coordination game when $c < X_1 < d$ where c is the maximal payoff a player can get if X_1 is taken by the other player.[2] It means that there are two pure Nash equilibria in this game in which the players coordinate by either both choosing to try to take X_1, or both choosing not to take X_1, and the game also has a completely mixed equilibrium. A property of symmetric coordination games is that both players obtain the same expected payoff at equilibrium, and this property is the reason why the SPEP belong to the diagonal in the model with recall. Interestingly, the corresponding game in the model without recall, when both players anticipate continuation equilibrium payoffs (d, e) if none of them chooses to take X_1, is an anti-coordination game when $c > X_1 > d \wedge e$, where c is the maximal payoff a player can get if X_1 is taken by the other player.[3] It means that there are two pure Nash equilibria in this game in which one of the players chooses to take X_1 and gets a smaller payoff X_1, while the other chooses not to take X_1 and gets a larger payoff c, and the game also has a completely mixed equilibrium. Note that pure equilibria give payoffs outside the diagonal even when $d = e$, in which case the anti-coordination game is symmetric. Let us emphasize that these auxiliary games are never anti-coordination games in the variant with recall, and never coordination games in the variant without recall, whatever the values of X_1 and d, e.

3.4 Efficiency of equilibria

The goal of this section is to study how efficient the SPE payoffs are. To this end, we define as usual the Price of Anarchy and Price of Stability, and we introduce what we call the Prophet Ratio of the game. Given an instance of a game, the first two notions refer to the ratio between the maximal sum of payoffs obtained by players under any feasible strategy and the sum of payoffs for the worst and the best SPEs, respectively. On the other hand, we define the Prophet Ratio of an instance of the problem as the ratio between the optimal Prophet value of the problem (that is, the expected sum of the two best values) and the sum of payoffs for the best SPE. We call this quantity the Prophet Ratio because we are comparing the best sum of payoffs obtained by playing a SPE strategy with what a prophet would do if he knows all the values of the realizations in advance.

Next, we formally introduce these definitions.

[2] Precisely, it corresponds to case 2 in the analysis of the game described in Figure 9, see the proof of Theorem 1

[3] Precisely, it corresponds to case 3 in the analysis of the game described in Figure 10, see the proof of Theorem 3

Definition 3 *Consider an instance Γ_n^{FR} or Γ_n^{NR}, where n observations from a distribution F arrive sequentially over time. Denote by Σ the set of all feasible strategy pairs and by SPE the set of subgame-perfect equilibria. We define:*

1. *the Price of Anarchy of this game instance– denoted by $PoA_n(F)$– to be the following ratio*

$$PoA_n(F) := \frac{\max_{\sigma\in\Sigma}\gamma^1(\sigma)+\gamma^2(\sigma)}{\min_{\sigma\in SPE}\gamma^1(\sigma)+\gamma^2(\sigma)},$$

2. *the Price of Stability of this game instance– denoted by $PoS_n(F)$– to be the following ratio*

$$PoS_n(F) := \frac{\max_{\sigma\in\Sigma}\gamma^1(\sigma)+\gamma^2(\sigma)}{\max_{\sigma\in SPE}\gamma^1(\sigma)+\gamma^2(\sigma)},$$

3. *the Prophet Ratio of this game instance– denoted by $PR_n(F)$– to be the following ratio*

$$PR_n(F) := \frac{\mathbb{E}(X_{(1:n)}+X_{(2:n)})}{\max_{\sigma\in SPE}\gamma^1(\sigma)+\gamma^2(\sigma)},$$

   where $X_{(1:n)}$ and $X_{(2:n)}$ represent the first and second-order statistics from the sample $\{X_i\}_{i\in[n]}$.

Clearly, by definition, it holds that for each n and F

$$\min\{\mathrm{PoA}_n(F), \mathrm{PR}_n(F)\} \geq \mathrm{PoS}_n(F) \geq 1.$$

Notice that

$$\mathrm{PoS}_n(F)/\mathrm{PR}_n(F) = \frac{\max_{\sigma\in\Sigma}\gamma^1(\sigma)+\gamma^2(\sigma)}{\mathbb{E}(X_{(1:n)}+X_{(2:n)})}$$

is usually called competitive ratio for the two-sample optimal selection problem.

For each n, we define the Price of Anarchy, Price of Stability, and Prophet Ratio of our competitive selection problems as the worst-case ratio over all possible value distributions F. That is:

$$\mathrm{PoA}_n := \sup_F \mathrm{PoA}_n(F), \quad \mathrm{PoS}_n := \sup_F \mathrm{PoS}_n(F), \quad \mathrm{PR}_n := \sup_F \mathrm{PR}_n(F).$$

In what follows, we study these quantities for each of the model variants, and at the end of the section, we consider the case where the number of arrivals is two, and we present a tight bound for the ratios in Theorem 5.

3.4.1 Full recall case

Note that in this case, the maximal feasible sum of payoffs obtained by the players is simply the expected value of the two best realizations (they can wait until the end of the horizon and pick the best and second best values). Thus, here, the notions of the Price of Stability and Prophet Ratio are equivalent.

By Theorem 2, the sum of payoffs for the worst SPE is given by $2l_n$ and for the best SPE is $2h_n$, and therefore given an instance of the game we have in this setting:

$$\mathrm{PoA}_n^{FR}(F) = \frac{\mathbb{E}(X_{(1:n)} + X_{(2:n)})}{2l_n}, \quad \text{and}$$

$$\mathrm{PoS}_n^{FR}(F) = \mathrm{PR}_n^{FR}(F) = \frac{\mathbb{E}(X_{(1:n)} + X_{(2:n)})}{2h_n}.$$

Notice that if $n = 2$, that is, we have only two arrivals, each value will be eventually picked by a player, so that $\mathrm{PoA}_2^{FR}(F) = \mathrm{PoS}_2^{FR}(F) = \mathrm{PR}_2^{FR}(F) = 1$ for every value distribution F.

On the other hand, if n goes to infinity, then we also have that both the Price of Anarchy and Price of Stability goes to 1. Then, the interesting question is what happens with these ratios when n is finite and greater than 2.

Although we have a general characterization of the ratios for any value of n and distribution F, these quantities are not always easy to compute for any n even if we fix the distribution F.

3.4.2 No recall case

If $n = 2$, picking the two best observations is a feasible strategy since player one can get X_1 and player two X_2, and then $\mathrm{PoS}_2^{NR}(F) = \mathrm{PR}_2^{NR}(F)$. However, for $n \geq 3$, as soon as the support of F has at least three points, picking almost surely the two best observations is no longer feasible and $\mathrm{PR}_n^{NR}(F) > \mathrm{PoS}_n^{NR}(F)$.

Note that in this case, the feasible strategy giving the maximal sum of expected payoffs is the same as the strategy of one player selecting two values among n in the classical online selection problem. The following Lemma gives us a recursive formula for the maximal expected sum of payoffs over feasible strategies.

Lemma 2. *Assume we are in the no recall case with n arrivals following a distribution F with mean m. Let X denote a random variable with law F, and c_n the value of the decision problem in the no recall case with one decision-maker and n arrivals. Then, the expected maximal feasible sum of payoffs s_n satisfies*

1. *$s_1 = m$ if $n = 1$,*
2. *$s_2 = 2m$ if $n = 2$,*

3. $s_n = \mathbb{P}(X \geq s_{n-1} - c_{n-1})\mathbb{E}(X + c_{n-1} \mid X \geq s_{n-1} - c_{n-1}) + s_{n-1}\mathbb{P}(X < s_{n-1} - c_{n-1})$ *if* $n > 2$.

Using Theorem 3 together with Lemma 2 we have:

$$\mathrm{PoA}_n^{NR}(F) = \frac{s_n}{2\alpha_n}, \quad \mathrm{PoS}_n^{NR}(F) = \frac{s_n}{2\beta_n}, \quad \text{and} \quad \mathrm{PR}_n^{NR}(F) = \frac{\mathbb{E}(X_{(1:n)} + X_{(2:n)})}{2\beta_n}.$$

If $n = 1$ the ratios are equal to 1 and when n tends to infinity, we also obtain that the ratios go to 1. Thus, as in the full recall case, the interesting cases are the ones in between.

Recall that from Theorem 4 we have that for every n and every continuous distribution F, $h_n \geq \beta_n$ and therefore $\mathrm{PR}_n^{FR}(F) \geq \mathrm{PR}_n^{NR}(F)$. However, the comparison is not direct for PoA and PoS as the numerators are different in the full recall and the no recall cases.

3.4.3 Two arrivals case

To conclude this section, we study the efficiency of SPE when we fix the number of arrivals at two and we look at the worst-case ratios over F. In the case with full recall, all ratios are 1 and the question is trivial, so we consider the no recall case here. In other words, we consider the game Γ_2^{NR} and we want to study how bad it may be to play the best or worst SPE, in terms of the sum of payoffs obtained, compared with the maximal feasible sum of payoffs.

We obtain the following result, which states that both the Price of Anarchy and Price of Stability are upper bounded by 4/3, and that this bound is tight. The proof can be found in Section 4.3.

Theorem 5. *If $n = 2$, under the no recall case it holds that for every distribution F,*

$$PoS_2^{NR}(F) \leq 4/3 \qquad \text{and} \qquad PoA_2^{NR}(F) \leq 4/3.$$

Furthermore, this bound is tight for both the price of stability and the price of anarchy.

3.5 Example: Uniform distribution

In this section, we apply the results obtained in the former sections to the particular case where the realizations are from a uniform distribution on [0,1].

3.5.1 Computation of SPEP

We start with the computation of the SPEP. To this end, we divide the analysis according to whether we are under the full recall or no recall case. For the former, we compute E_3^{FR} to illustrate how to apply Theorem 1. For the latter, we implement the recursive formulas for α_n, β_n, and α'_n.

Full recall case. We now compute the set of SPEP of the game Γ_3^{FR}, when we have three values arriving from a uniform $[0,1]$ distribution.

Given $1 \geq a \geq b \geq 0$, by Theorem 1 it holds that $l_0(a,b) = h_0(a,b) = (a+b)/2$ and

$$E_0^{FR}(a,b) = \{((a+b)/2, (a+b)/2)\}.$$

Let us now compute a closed-form expression for $E_1^{FR}(a,b)$.

Using the notation introduced in Theorem 1:

$$\begin{aligned} d_1^-(a,b) &= d_1^+(a,b) = \mathbb{E}\left(\frac{a+b}{2}\mathbb{1}_{X<b} + \frac{a+X}{2}\mathbb{1}_{X\geq b}\right) \\ &= b(a+b)/2 + (1-b)(a/2+(1+b)/4) = \frac{1}{2}a + \frac{1}{4} + \frac{1}{4}b^2. \end{aligned}$$

If $X \sim \text{Unif}[0,1]$ and $k \in [0,1]$, we have:

$$\mathbb{E}(X \vee k) = k\mathbb{P}(X \leq k) + \mathbb{E}(X \mid X > k)\mathbb{P}(X > k) = \frac{1+k^2}{2}.$$

Using that $\frac{1+2a+b^2}{4} = \frac{1}{2}\left(a + \frac{1+b^2}{2}\right)$, we deduce that

$$\begin{aligned} l_1(a,b) &= L\left(a, \frac{1+b^2}{2}, \frac{1+2a+b^2}{4}\right) = h_1(a,b) \\ &= H\left(a, \frac{1+b^2}{2}, \frac{1+2a+b^2}{4}\right) = \frac{1+2a+b^2}{4}. \end{aligned}$$

By Theorem 1 we obtain

$$E_1^{FR}(a,b) = \left\{\left(\frac{a}{2} + \frac{1+b^2}{4}, \frac{a}{2} + \frac{1+b^2}{4}\right)\right\}.$$

To compute $E_2^{FR}(a,b)$, we need to compute $\mathbb{E}(X_{(2)} \vee b)$. The expected value of the maximum of n Unif$[0,1]$ random variables and a constant $k \in [0,1]$ is given by:

$$\begin{aligned} \mathbb{E}(X_{(n)} \vee k) &= k\mathbb{P}(X_{(n)} \leq k) + \mathbb{E}(X_{(n)} \mid X_{(n)} > k)\mathbb{P}(X_{(n)} > k) \\ &= k^{n+1} + \frac{n}{n+1}(1-k^{n+1}) \\ &= \frac{n+k^{n+1}}{n+1}, \end{aligned}$$

and therefore $\mathbb{E}(X_{(2)} \vee b) = (2 + b^3)/3$. We have:

$$\begin{aligned} d_2^-(a,b) = d_2^+(a,b) &= \mathbb{E}_X\left(\frac{a \vee X}{2} + \frac{1 + ((a \wedge X) \vee b)^2}{4}\right) \\ &= \frac{1+a^2}{2} + \frac{b^3 - a^3}{6}. \end{aligned}$$

We obtain:

$$l_2(a,b) = \begin{cases} (1+a^2)/2 + (b^3 - a^3)/6 & \text{if } a \le (2+b^3)/3 \\ a/2 + (2+b^3)/6 & \text{if } a > (2+b^3)/3 \end{cases}$$

$$h_2(a,b) = \begin{cases} (1+a^2)/2 + (b^3 - a^3)/6 & \text{if } a \le \nu \\ a/2 + (2+b^3)/6 & \text{if } a > \nu, \end{cases}$$

where $\nu = \max\{(2+b^3)/3, (1+a^2)/2 + (b^3 - a^3)/6\}$.

In particular, $l_2(0,0) = h_2(0,0) = 1/2$, and therefore $E_2^{FR} = E_2^{FR}(0,0) = \{(1/2, 1/2)\}$. Now, we compute E_3^{FR}. We have

$$l_2(a,0) = \begin{cases} (1+a^2)/2 - a^3/6 & \text{if } a \le 2/3 \\ a/2 + 1/3 & \text{if } a > 2/3 \end{cases}$$

$$h_2(a,0) = \begin{cases} (1+a^2)/2 - a^3/6 & \text{if } a \le a^* \\ a/2 + 1/3 & \text{if } a > a^* \end{cases},$$

where $a^* \approx 0.6778$ is the unique root of $a = (1+a^2)/2 - a^3/6$. We deduce that

$$\begin{aligned} l_3(0,0) &= d_3^-(0,0) = \mathbb{E}(l_2(X,0)) = \int_0^{2/3} ((1+a^2)/2 - a^3/6)\mathrm{d}a \\ &+ \int_{2/3}^1 (a/2 + 1/3)\mathrm{d}a = \frac{607}{972} \approx 0.6244, \end{aligned}$$

$$\begin{aligned} h_3(0,0) &= d_3^+(0,0) = \mathbb{E}(h_2(X,0)) = \int_0^{a^*} ((1+a^2)/2 - a^3/6)\mathrm{d}a \\ &+ \int_{a^*}^1 (a/2 + 1/3)\mathrm{d}a \approx 0.6245. \end{aligned}$$

We conclude that

$$E_3^{FR}(0,0) = \{(u,u) : u \in [l_3(0,0), h_3(0,0)]\}.$$

What is played at equilibrium? In the best equilibrium, both players pick X_1 if and only if $X_1 \ge a^*$, whereas, in the worst equilibrium, both players pick

X_1 if and only if $X_1 \geq 2/3$. More generally, if $X_1 \in (2/3, a^*)$, there is a unique continuation common SPEP $d = d_2^+(X_1, 0) = d_2^-(X_1, 0) = \frac{1-X_1^2}{2} - \frac{X_1^3}{6}$ if X_1 is not taken at stage 1 and thus we are in the case $d > X_1 > c$ where $c = 2/3$ is the maximal expected payoff a player can obtain if the other player gets X_1 at the first stage. As mentioned earlier, the game played at the first stage is a coordination game with two pure Nash equilibria and one completely mixed Nash equilibrium. Either both the players choose not to take X_1 and they obtain the best SPEP, or they both choose to try to take X_1 and they obtain the worst SPEP. The mixed equilibrium gives an intermediate SPEP.

No recall case. Now, we turn to reducing the formulas obtained in Theorem 3 for the particular case where the values are uniformly distributed in the interval $[0, 1]$. Recall that $c_1 = \mathbb{E}(X)$, $c_n = \mathbb{E}(X \vee c_{n-1})$ for $n > 1$ and $X \sim F$. To obtain the expressions of β_n and α_n for the uniform case, we use the following technical result.

Lemma 3. *If $\alpha'_n < a < \beta_n$, then $a + c_n \geq 2\beta_n$.*

Proof It is enough to show that $\alpha'_n + c_n \geq 2\beta_n$. We prove that by induction on n.

Notice that $E_1^{NR} = \{(1/4, 1/4)\}$ and thus $\beta_1 = \alpha'_1 = 1/4$. On the other hand, $c_1 = \mathbb{E}(X) = 1/2$, and putting all together we have that

$$\alpha'_1 + c_1 \geq 2\beta_1,$$

and thus the Lemma holds for $n = 1$.

Let us assume now that the inequality holds for n and we prove that it also holds for $n + 1$, that is:

$$\alpha'_{n+1} + c_{n+1} \geq 2\beta_{n+1}.$$

By Theorem 3, we have that

$$\begin{aligned}
2\beta_{n+1} &= 2\alpha'_n\beta_n + \int_{\alpha'_n}^{\beta_n} \max\{a + c_n, 2\beta_n\}\mathrm{d}a + (1 - \beta_n)c_n + \frac{1}{2}(1 - \beta_n^2) \\
&= 2\alpha'_n\beta_n + (\beta_n - \alpha'_n)c_n + \frac{1}{2}(\beta_n^2 - {\alpha'_n}^2) + (1 - \beta_n)c_n + \frac{1}{2}(1 - \beta_n^2) \\
&= 2\alpha'_n\beta_n + c_n(1 - \alpha'_n) + \frac{1}{2}(1 - {\alpha'_n}^2) \\
&\leq (\alpha'_n + c_n)\alpha'_n + c_n(1 - \alpha'_n) + \frac{1}{2}(1 - {\alpha'_n}^2) \\
&= c_n + \frac{1}{2}{\alpha'_n}^2 + \frac{1}{2},
\end{aligned}$$

where the second equality and the inequality follow by the induction hypothesis.

On the other hand, from Theorem 3, we obtain

$$\alpha'_{n+1} = \frac{1}{2}{\alpha'_n}^2 + \frac{1}{2}c_n - \frac{1}{4}c_n^2 + \frac{1}{4}.$$

Then, to show $\alpha'_{n+1} + c_{n+1} \geq 2\beta_{n+1}$ it suffices to show

$$\frac{1}{4}(c_n + 1)^2 \leq c_{n+1},$$

which holds due to $c_{n+1} = \mathbb{E}(X \vee c_n) = (1 + c_n^2)/2$.

Therefore, we have that $2\beta_{n+1} \leq \alpha'_{n+1} + c_{n+1}$, and the proof is completed. □

Using Theorem 3 together with Lemma 3, we obtain that the sum of payoffs for the best SPE for $n > 1$ is

$$\begin{aligned} 2\beta_n &= \int_0^{\alpha'_{n-1}} 2\beta_{n-1}\mathrm{d}a + \int_{\alpha'_{n-1}}^1 (a + c_{n-1})\mathrm{d}a \\ &= 2\beta_{n-1}\alpha'_{n-1} + c_{n-1}(1-\alpha'_{n-1}) + \frac{1}{2}(1-{\alpha'_{n-1}}^2). \end{aligned}$$

Notice that by definition, $\beta_n \geq \alpha_n$ for all n, and therefore, by Lemma 3 we conclude that if $\alpha'_n < a < \alpha_n$ then $a + c_n \geq 2\alpha_n$. Thus, applying Theorem 3, the sum of payoffs for the worst SPE for $n > 1$ is

$$2\alpha_n = \frac{1}{2} + \frac{5}{2}c_{n-1}^2 + 3\beta_{n-1}^2 + c_{n-1} - 6c_{n-1}\beta_{n-1} + \alpha_{n-1}^2 - 4(c_{n-1} - \beta_{n-1})^2 \ln(2).$$

If the realizations are uniformly distributed in the interval $[0, 1]$, the recursive formulas given in Theorem 3 are easy to implement numerically.

In Table 1 we present the values of α'_n, α_n, and β_n for n from 1 to 4 and $F = \mathrm{Unif}[0, 1]$, whereas in Fig. 4 we compare the sum of payoffs for the best and worst SPE for n up to 10.

	α'_n	α_n	β_n
n=1	0.25	0.25	0.25
n=2	0.4688	0.4759	0.4844
n=3	0.5747	0.5803	0.5881
n=4	0.6419	0.6465	0.6533

Table 1: Values of α'_n, α_n, and β_n for one, two, three, and four arrivals and values uniformly distributed in $[0, 1]$.

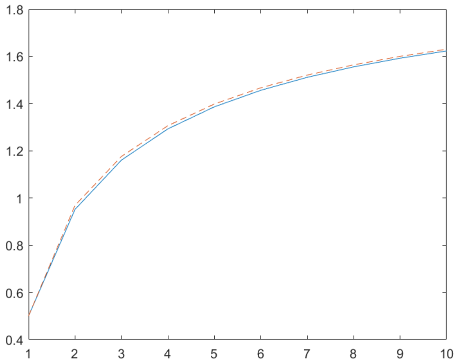

Fig. 4: Sum of payoffs for the best SPE (dashed line) and the worst SPE (continuous line) for up to 10 arrivals, no recall, and values uniformly distributed in $[0, 1]$.

In the game with 3 stages, if $X_1 \in (\alpha_2, \beta_2)$ and if the continuation SPEP after stage 1 if X_1 was not taken are given by $(d, e) = (\alpha_2, \alpha_2)$, we have $d \wedge e < X_1 < c_2$. As mentioned earlier, the game faced by the players at the first stage is an anti-coordination game. There are two pure Nash equilibria, in which only one player tries to take X_1 and gets the smaller payoff X_1, while the other player gets the larger payoff c_2. The same situation occurs if $X_1 \in (\beta_2, c_2)$ and $(d, e) = (\beta_2, \beta_2)$.

3.5.2 Comparison of model variants

We know by Theorem 4 that $\alpha_n \leq \beta_n \leq l_n \leq h_n$ for every n, meaning that having full recall is advantageous to the players. We quantify this difference for small values of n numerically,

In Table 2 we present the values of α_n, β_n, l_n, and h_n for $n \leq 5$ arrivals. The values corresponding to the no recall case were computed using the formulas obtained in the previous section, whereas for the full recall case, for $n = 1, 2, 3$ the values were obtained explicitly and for $n = 4, 5$ numerically via discretization.

From the values of Table 2, we can see that under the full recall case, each player has an advantage compared to the no recall case of between 3% and 5% for $n = 2$, between 6% and 7.6% for $n = 3$, between 6.9% and 8% for $n = 4$ and between 7% and 8% for $n = 5$.

	l_n	h_n	α_n	β_n
n=1	0.25	0.25	0.25	0.25
n=2	0.5	0.5	0.4759	0.4844
n=3	0.6244	0.6245	0.5803	0.5881
n=4	0.6989	0.699	0.6465	0.6533
n=5	0.7484	0.7486	0.6932	0.6991

Table 2: Values of l_n, h_n, α_n, and β_n for up to 5 arrivals and values uniformly distributed in $[0, 1]$.

3.5.3 Efficiency of equilibria

We show in Table 3 the values of the ratios $\mathrm{PoA}_n(F), \mathrm{PoS}_n(F)$ and $\mathrm{PR}_n(F)$ for $X \sim \mathrm{Unif}[0, 1]$ and n up to 5 in both settings. We notice that the ratios are close to 1, i.e. equilibria are close to being efficient in both model variants, being more efficient in the full recall case. We highlight here that these ratios are a measure of the efficiency of equilibria for each game, but not between the two different games. That is to say, it is not correct to say that one setting is better than the other by comparing the PoA or PoS because, for each game, these ratios are computed in a different way (the numerators are different). For such a comparison, we could use the value of PR, but comparing the PR for both problems is equivalent to comparing β_n and h_n, which is a comparison we already made.

	$\text{PoA}_n(\text{F})$		$\text{PoS}_n(\text{F})$		$\text{PR}_n(\text{F})$	
	Full Recall	No Recall	Full Recall	No Recall	Full Recall	No Recall
n=2	1	1.0507	1	1.0323	1	1.0323
n=3	1.000823	1.0299	1.0008	1.0161	1.0008	1.0627
n=4	1.00157	1.0212	1.00143	1.0105	1.00143	1.0714
n=5	1.0021	1.0164	1.00187	1.0077	1.00187	1.0728

Table 3: Efficiency of equilibria with up to 5 arrivals from Unif[0,1] distribution.

Another question we address here is the number of arrivals that gives the worst gaps, in the no recall case for which the numerics are easier. We obtain that, for both the Price of Anarchy and Price of Stability, the ratios reach their maximum when $n = 2$, see Fig. 5a and 5b. This result is somehow intuitive: as we are in the no recall case, there exists a positive probability of getting nothing, and then the smaller the number of arrivals, the more likely this seems to happen. Regarding the Prophet Ratio, we also compute it as a function of n (see Fig. 5c), and we obtained that the maximum is reached when $n = 5$. Here, the result is more surprising.

4 Proofs

In this section, we provide the proofs omitted in Section 3.

4.1 Omitted proofs from Section 3.2

Full recall case. In this section, we prove Theorems 1 and 2. Before that, we need a few preliminary results stating some properties of the game $\Gamma_n^{FR}(a,b)$ defined in Section 3.2. Recall that this game is defined as Γ_n^{FR} with n arrivals but with two initial values $a \geq b$ already present in the market.

At first, to study the SPEP of Γ_n^{FR}, note that this game has the same SPEP as the game which is identical to Γ_n^{FR} and terminates at the first time a player gets a value or at stage n. If the game terminates because one player did get a realization, the payoff of the other player is the value of the one-player continuation problem, and the payoff of the players if none of them did stop strictly before stage n is the expected payoff given by the pair of strategies "bid for the best available value". Indeed, it is easy to check that in both these situations the SPEP in the continuation games are unique and correspond to the payoffs of this auxiliary game. In the following, we assume that Γ_n^{FR} is the auxiliary game. In this game, a history is just a sequence of values $(X_1, ..., X_t)$, and the strategy of player i at time $t < n$ is a measurable map $\sigma_{i,t}$ from histories into the probabilities over $\{\emptyset\} \cup \{1, ..., t\}$. We use the same identification for the game $\Gamma_n^{FR}(a,b)$.

Given a history $h = (x_1, ..., x_t)$ of length t, let $\Gamma_n(h)$ denote the subgame of Γ_{n+t}^{FR} starting at stage t after observing h in which the two players are still present and $E_n(h)$ is the set of SPEP of this game.

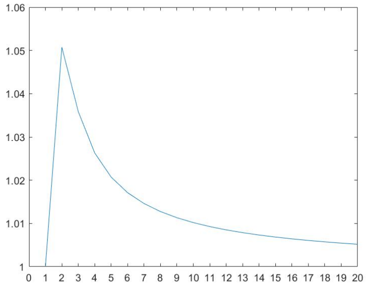

(a) $\mathrm{PoA}_n^{NR}(F)$.

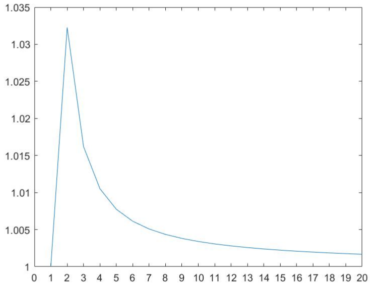

(b) $\mathrm{PoS}_n^{NR}(F)$

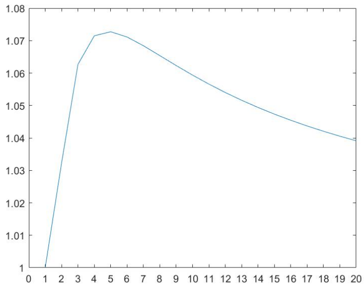

(c) $\mathrm{PR}_n^{NR}(F)$

Fig. 5: Price of anarchy, price of stability, and prophet ratio as a function of n when $F = \mathrm{Unif}[0,1]$ in the no recall case.

In the following, variables $(X, X_1, X_2, ...)$ will denote independent variables with distribution F, and $X_{(t)} = \max(X_1, ..., X_t)$.

Without loss of generality, we assume that $\mathbb{P}(X > 0) > 0$ (otherwise the set of equilibrium payoffs is reduced to $\{(0,0)\}$).

We state now without proof some properties of SPEP which follow easily from usual arguments in dynamic game theory (see, e.g., Section 4.2 in Fudenberg and Tirole 1991). Recall that if $x \in [0,1] \to S(x)$ is a set-valued map, $\mathbb{E}(S(X))$ is defined by

$$\mathbb{E}(S(X)) = \{\mathbb{E}(f(X)) \mid f \text{ measurable such that } \forall x \in [0,1], f(x) \in S(x)\}.$$

Lemma 4. *The following properties hold:*

1. *$E_n = E_n(0,0)$.*
2. *If $h = (x_1, ..., x_t)$ is a history of length $t \geq 2$, then $E_n(h) = E_n(a,b)$ where a and b denote the first and second largest values in h respectively.*

3. *$(x,y) \in E_n(a,b)$ with $a \geq b$ if and only if there exists $(d,e) \in \mathbb{E}(E_{n-1}(a \vee X, med[a,b,X]))$ such that (x,y) is a mixed Nash equilibrium payoff of the finite game $G_n(a,b;\ d,e)$ with payoff matrix:*

$$\begin{array}{c c} & \begin{array}{ccc} a & b & \emptyset \end{array} \\ \begin{array}{c} a \\ b \\ \emptyset \end{array} & \begin{pmatrix} \left(\frac{a+M_b}{2}, \frac{a+M_b}{2}\right) & (a,b) & (a, M_b) \\ (b,a) & \left(\frac{b+M_a}{2}, \frac{b+M_a}{2}\right) & (b, M_a) \\ (M_b, a) & (M_a, b) & (d,e) \end{pmatrix} \end{array}$$

Fig. 6: Payoff matrix of $G_n(a,b;\ d,e)$, where $M_b = \mathbb{E}(X_{(n)} \vee b)$ and $M_a = \mathbb{E}(X_{(n)} \vee a)$.

4. *Similarly, (σ_1, σ_2) is a pair of first-stage strategies of some SPE in $\Gamma_n^{FR}(a,b)$ with payoff (x,y) in $\Gamma_n^{FR}(a,b)$ if and only if there exists $(d,e) \in \mathbb{E}(E_{n-1}(a \vee X, med[a,b,X]))$ such that (σ_1, σ_2) is a mixed Nash equilibrium of the matrix game $G_n(a,b;\ d,e)$ with payoff (x,y).*

The first point follows from the fact that any strategy in any subgame of $\Gamma_n(0,0)$ that bids with a positive probability for one of the two initial observations with value 0 is strictly dominated by the modified strategy which waits until the last stage and bids for the best available realization. Therefore such strategies are not played in any equilibrium in $\Gamma_n(0,0)$ and the result follows easily. The other points can be proven by induction on n using the recursive structure of SPE, the one-shot deviation principle, and measurable selection arguments.

In the following proposition, we prove that if a player weakly (strictly) prefers to pass instead of taking a given that the other player takes a, then he weakly (strictly) prefers to pass instead of taking a if the other player passes.

Proposition 2 *Let $a > b$, $(d,e) \in \mathbb{E}(E_{n-1}(a \vee X, med[a,b,X]))$ and $c = \mathbb{E}(X_{(n)} \vee b)$. If $c \geq a$ then $d \geq a$ and $e \geq a$. Similarly, if $c > a$ then $d > a$ and $e > a$.*

By the symmetry of the game, it is enough to prove the inequality for d, then the same arguments will hold for e.

To prove the Proposition, we need the following two lemmas.

Lemma 5. *If X, Y are integrable random variables and a is a real number such that $\mathbb{E}(X \vee Y) \geq a$ (resp. $> a$), then it holds that $\frac{1}{2}\mathbb{E}(X \vee a) + \frac{1}{2}\mathbb{E}((X \wedge a) \vee Y) \geq a$ (resp. $> a$).*

Proof It is enough to show that $\mathbb{E}(X \vee a + (X \wedge a) \vee Y) \geq \mathbb{E}((X \vee Y) \vee a + (X \vee Y) \wedge a)$, obtaining then the desired inequality by using that $\mathbb{E}(X \vee Y) \geq a$.

We now even prove that $X \vee a + (X \wedge a) \vee Y \geq (X \vee Y) \vee a + (X \vee Y) \wedge a$ almost surely. We consider two cases.

Case 1. Assume that $X \leq a$. In this case, we have

$$X \vee a + (X \wedge a) \vee Y = a + X \vee Y = (X \vee Y) \vee a + (X \vee Y) \wedge a.$$

Case 2. Assume that $X > a$. In this case, we have

$$X \vee a + (X \wedge a) \vee Y = X + a \vee Y.$$

On the other hand,

$$(X \vee Y) \vee a + (X \vee Y) \wedge a = \begin{cases} Y + a & \text{if } Y > X, \\ X + a & \text{if } Y \leq X, \end{cases}$$

and since $X > a$ and $a \vee Y \geq a$, we conclude that

$$X + a \vee Y \geq \begin{cases} a + Y & \text{if } Y > X, \\ X + a & \text{if } Y \leq X, \end{cases}$$

obtaining that $X \vee a + (X \wedge a) \vee Y \geq (X \vee Y) \vee a + (X \vee Y) \wedge a$ almost surely. □

Lemma 6. *Under the same notation and assumptions as in Proposition 2, we have*

$$d \geq \frac{1}{2}\mathbb{E}(X \vee a) + \frac{1}{2}\mathbb{E}(\beta(X)),$$

where X is a random variable with distribution F, $\beta(x) = \mathbb{E}(med[a,b,x] \vee X_{(n-1)})$ and $X_{(n-1)} = \max(X_1, .., X_{n-1})$ for some i.i.d. sequence $X_1, ..., X_{n-1}$ of observations of X.

Proof By assumption, $d = \mathbb{E}(f(X))$ for some measurable function f such that $f(x)$ is the expected payoff of player 1 in some SPE in $\Gamma^{FR}_{n-1}(a \vee x, med[a,b,x])$. We will obtain a lower bound for $f(x)$ by providing a lower bound for the payoff associated with a particular strategy in $\Gamma^{FR}_{n-1}(a \vee x, med[a,b,x])$.

To this end, suppose that Player 1 plays the following strategy:

1. If $x \vee a \geq \beta(x)$, bid for $x \vee a$.
2. If $x \vee a < \beta(x)$, wait until the end and get at least the second best value.

Define $\Omega_1 := \{x \in [0,1] : x \vee a \geq \beta(x)\}$, $\Omega_2 := \{x \in [0,1] : x \vee a < \beta(x)\}$. Since the second best realization in $\Gamma^{FR}_{n-1}(a \vee x, med[a,b,x])$ has expected value $\beta(x)$, we deduce that the payoff of player 1, independently of the strategy of player 2, is at least $\frac{1}{2}(x \vee a + \beta(x))$ if $x \in \Omega_1$ and $\beta(x)$ if $x \in \Omega_2$. We deduce that

$$d = \mathbb{E}(f(X)) \geq \mathbb{E}\left(\frac{1}{2}(X \vee a + \beta(X))\mathbb{1}_{X \in \Omega_1} + \beta(X)\mathbb{1}_{X \in \Omega_2}\right).$$

Note that the last term is at least $\mathbb{E}(\frac{1}{2}(X \vee a + \beta(X))\mathbb{1}_{X \in \Omega_2})$ because $x \vee a < \beta(x)$ for $x \in \Omega_2$. Therefore, we conclude that

$$\begin{aligned} d &\geq \mathbb{E}\left(\frac{1}{2}(X \vee a + \beta(X))\mathbb{1}_{X \in \Omega_1} + \frac{1}{2}(X \vee a + \beta(X))\mathbb{1}_{X \in \Omega_2}\right) \\ &= \mathbb{E}\left(\frac{1}{2}(X \vee a + \beta(X))\right) \end{aligned}$$

and the result follows. □

Proof of Proposition 2 Take X with distribution F independent of some i.i.d sequence $(X_1, ..., X_{n-1})$ of variables distributed according to F, and $Y = X_{(n-1)} \vee b$. Then $\mathbb{E}(X \vee Y) = c \geq a$, and applying Lemma 5 we have

$$\frac{1}{2}\mathbb{E}(X \vee a) + \frac{1}{2}\mathbb{E}((X \wedge a) \vee X_{(n-1)} \vee b) \geq a.$$

On the other hand, by independence and since $med[a, b, x] = (x \wedge a) \vee b$, we have

$$\mathbb{E}((X \wedge a) \vee X_{(n-1)} \vee b) = \mathbb{E}(\beta(X)).$$

By Lemma 6, $d \geq \frac{1}{2}\mathbb{E}(X \vee a) + \frac{1}{2}\mathbb{E}(\beta(X))$. Putting all together we obtain $d \geq a$ and the proof is completed. The case with strict inequality is similar. □

The next result allows us to reduce the analysis of $G_n(a, b;\ d, e)$ to a smaller matrix game.

Lemma 7. *Let $a \geq b$ in $[0, 1]$, n be a positive natural number, and $(d, e) \in \mathbb{E}(E_{n-1}(a \vee X, med[a, b, X]))$. The game $G_n(a, b;\ e, d)$ has the same Nash equilibrium payoffs as the game with matrix represented in Fig. 7.*

$$\begin{array}{cc} & \begin{array}{cc} a & \emptyset \end{array} \\ \begin{array}{c} a \\ \emptyset \end{array} & \begin{pmatrix} \left(\frac{a+\mathbb{E}(X_{(n)} \vee b)}{2}, \frac{a+\mathbb{E}(X_{(n)} \vee b)}{2}\right) & \left(a, \mathbb{E}(X_{(n)} \vee b)\right) \\ \left(\mathbb{E}(X_{(n)} \vee b), a\right) & (d, e) \end{pmatrix} \end{array}$$

Fig. 7: Reduced payoff matrix for $G_n(a, b;\ d, e)$.

Proof We first consider the case $a > b$.

For both players, the strategy b is strictly dominated by a mixed strategy consisting of playing a with probability 1/2 and pass with probability 1/2. Due to the symmetry of the payoff matrix, it is enough to show that for one player (let us say player 1) the expected payoff if he plays a with probability 1/2 and passes with probability 1/2 is strictly higher than the expected payoff he obtains if he plays b. Note that because $a > b$ we have that:

1. $\frac{1}{2}\frac{a+\mathbb{E}(X_{(n)} \vee b)}{2} + \frac{1}{2}\mathbb{E}(X_{(n)} \vee b) = \frac{a}{4} + \frac{3}{4}\mathbb{E}(X_{(n)} \vee b) > b$,
2. $\frac{1}{2}a + \frac{1}{2}\mathbb{E}(X_{(n)} \vee a) > \frac{b+\mathbb{E}(X_{(n)} \vee a)}{2}$, and
3. $\frac{1}{2}a + \frac{1}{2}d > \frac{1}{2}b + \frac{1}{2}b = b$,

where the last inequality follows because $d \geq b$. Indeed, in any SPE, player 1 obtains at least the second largest value (otherwise he could deviate to the strategy which does not bid until stage n), which is not smaller than b in $\Gamma_n^{FR}(a, b)$. We conclude that b is a strictly dominated strategy for both players (due to the symmetry of the game), and the result follows.

If $a = b$, we consider two subcases. If $\mathbb{E}(X_{(n)} \vee a) > a$, then Proposition 2 implies that $d > a$ and $e > a$, so that the strategies a and b are strictly dominated by $\emptyset$.

There is a unique Nash equilibrium$(\emptyset,\emptyset)$ with payoff (d,e). Since the same holds for the matrix game given in Fig. 7, the result follows.

If $\mathbb{E}(X_{(n)} \vee a) = a$, the actions a and b are equivalent in the sense that they induce the same payoffs. Eliminating b leads to a matrix game with the same Nash equilibrium payoffs. □

Denoting $c = \mathbb{E}(X_{(n)} \vee b)$, the payoff matrix introduced in Lemma 7 has the particular form presented in Fig. 8. Thus, it is enough to compute the NE of this matrix game where (d,e) are parameters that correspond to the expected continuation equilibrium payoffs for players.

$$\begin{array}{cc} & \begin{array}{cc} a & \emptyset \end{array} \\ \begin{array}{c} a \\ \emptyset \end{array} & \begin{pmatrix} \left(\frac{1}{2}(a+c), \frac{1}{2}(a+c)\right) & (a,c) \\ (c,a) & (d,e) \end{pmatrix} \end{array}$$

Fig. 8: General form of the reduced payoff matrix for $G_n(a,b;\ d,e)$.

We are now ready to prove Theorem 1.

Proof of Theorem 1 Let us first study the game $G_n(a,b;\ d,d)$ with $(d,d) \in \mathbb{E}(E_{n-1}(a \vee X, med[a,b,X]))$, i.e. the case of continuation payoffs which belong to the diagonal in $\mathbb{R}^2$, given in Fig. 9 below where $c = \mathbb{E}(X_{(n)} \vee b)$.

$$\begin{array}{cc} & \begin{array}{cc} a & \emptyset \end{array} \\ \begin{array}{c} a \\ \emptyset \end{array} & \begin{pmatrix} \left(\frac{1}{2}(a+c), \frac{1}{2}(a+c)\right) & (a,c) \\ (c,a) & (d,d) \end{pmatrix} \end{array}$$

Fig. 9: Payoff matrix for $G_n(a,b;\ d,d)$.

Let us compute the mixed Nash equilibria of this game. Using Proposition 2, $c > a \Rightarrow d > a$ and $c \geq a \Rightarrow d \geq a$. We have that:

1. If $a > c \vee d$, then (a,a) is the unique NE with payoff $\left(\frac{1}{2}(a+c), \frac{1}{2}(a+c)\right)$.
2. If $d > a > c$, there are two pure NE (a,a) and $(\emptyset,\emptyset)$ and a symmetric mixed equilibrium in which both agents play a with probability $\frac{2(d-a)}{2d-a-c}$ and pass with probability $\frac{a-c}{2d-a-c}$. Furthermore, the equilibrium payoffs are:
$$\left(\frac{1}{2}(a+c), \frac{1}{2}(a+c)\right), (d,d) \text{ and } \left(\frac{dc-2ac+ad}{2d-a-c}, \frac{dc-2ac+ad}{2d-a-c}\right),$$
respectively.
3. If $a < c$, then $(\emptyset,\emptyset)$ is the unique NE, and (d,d) is the expected payoff.

4. If $d = a > c$ or $d > a = c$, there are two pure NE (a, a) and $(\emptyset, \emptyset)$ with payoffs:
$$\left(\frac{1}{2}(a+c), \frac{1}{2}(a+c)\right) \text{ and } (d, d).$$
5. If $d = a = c$, then any profile is a NE with payoff (d, d).

From the analysis above we conclude that the set of NE payoffs of the game is contained in the diagonal.

We now prove by induction on n that $E_n(a, b)$ is also a subset of the diagonal, that is:
$$E_n(a, b) \subset \{(u, u) : u \in \mathbb{R}_+\}.$$
At first, $E_0(a, b) = \{(\frac{a+b}{2}, \frac{a+b}{2})\}$, so the statement is correct for $n = 0$.

Let us assume that the statement is correct for $E_n(a, b)$ for all pairs (a, b). Then, it implies that $\mathbb{E}(E_n(a \vee X, med[a, b, X]))$ is also a subset of the diagonal. Therefore, by point 3. of Lemma 4 and the above analysis, we deduce that $E_{n+1}(a, b)$ is a subset of the diagonal.

The first statement of the theorem is proved.

Let us now show that $\min \mathrm{P}E_n(a, b) = l_n(a, b)$ and $\max \mathrm{P}E_n(a, b) = h_n(a, b)$, where $\mathrm{P}E_n(a, b)$ is the projection of $E_n(a, b)$ to its first coordinate in $\mathbb{R}$.

Note that from the analysis done above, we have that if $a > c \vee d$ or $a < c$ or $a = c = d$, there is only one NE payoff but in the other cases we have multiple equilibrium payoffs. When playing a SPE, the expected payoff of the players when both pass depends on which equilibrium is played in the following stages of the game. It is known that, in general, we cannot assume that if “the worst” or “the best” equilibrium is played at each stage, that results in the worst or best equilibrium of the game. However, below we show that this is indeed true for the game we are considering and then the result will follow just computing the expected payoff corresponding to the best and worst equilibrium.

To this end, it is enough to prove that the expected payoff of a player is increasing in d.

Define $\Omega = \{(a, c, d) \in \mathbb{R}^3 : (c > a \Rightarrow d > a) \text{ and } (c \geq a \Rightarrow d \geq a)\}$ and consider the multivalued function $\psi : \Omega \rightrightarrows \mathbb{R}$ defined by
$$\psi(a, c, d) = \begin{cases} d & \text{if } c > a \text{ or } a = c = d \\ \{\frac{a+c}{2}, d\} & \text{if } c < a < d \text{ or } c = a < d \text{ or } c < a = d \\ \frac{a+c}{2} & \text{if } c \vee d < a. \end{cases}$$
Note that if $d > a > c$, then
$$d > \frac{dc - 2ac + ad}{2d - a - c} > \frac{1}{2}(a+c),$$
and thus d and $\frac{1}{2}(a + c)$ are respectively the "best" and the "worst" equilibrium payoffs for player 1 in $G_n(a, b;\ d, d)$. Therefore, $\psi(a, c, d)$ represents the "best" and "worst" NE payoffs for player 1 in the game represented in Fig. 9.

We say that the multivalued function is non-decreasing in d if for each a, c, d_1, d_2 such that $d_1 < d_2$ and $(a, c, d_i) \in \Omega$ for $i = 1, 2$, it holds that $\min \psi(a, c, d_1) \leq \min \psi(a, c, d_2)$ and $\max \psi(a, c, d_1) \leq \max \psi(a, c, d_2)$. Let us see that ψ is non-decreasing in d. To show that, fix a, c and take d_1, d_2 such that $d_1 < d_2$ and $(a, c, d_1) \in \Omega$, $(a, c, d_2) \in \Omega$. We have the following cases:

1. If $c > a$, then $\psi(a, c, d_i) = \{d_i\}$ for $i = 1, 2$ and the result is obvious.

2. If $c < a \leq d_1$, then $c < a < d_2$ and $\frac{a+c}{2} < d_1 < d_2$. Thus, $\min \psi(a, c, d_1) = \frac{a+c}{2} = \min \psi(a, c, d_2)$ and $\max \psi(a, c, d_1) = d_1 \leq d_2 = \max \psi(a, c, d_2)$.
3. If $a > c \vee d_2$, then $\psi(a, c, d_1) = \{\frac{a+c}{2}\} = \psi(a, c, d_2)$ and the result is obvious.
4. If $c \vee d_1 < a \leq d_2$, then $\psi(a, c, d_1) = \{\frac{a+c}{2}\}$ and $\psi(a, c, d_2) = \{\frac{a+c}{2}, d_2\}$, obtaining that $\min \psi(a, c, d_1) = \max \psi(a, c, d_1) \leq \min \psi(a, c, d_2) \leq \max \psi(a, c, d_2)$.
5. If $a = c \leq d_1$, then $\psi(a, c, d_1) = \{a\}$ or $\{a, d_1\}$ and $\psi(a, c, d_2) = \{a, d_2\}$, and the result follows.

Thus, we conclude that ψ is non-decreasing in d. The monotonicity property we proved for ψ means that playing a "better" equilibrium in the continuation game gives a "better" equilibrium for $\Gamma_n^{FR}(a, b)$ and playing a "worse" equilibrium one gives a "worse" equilibrium for $\Gamma_n^{FR}(a, b)$.

Precisely, define $L : \Omega \to \mathbb{R}$ and $H : \Omega \to \mathbb{R}$ by:

$$L(a, c, d) = \begin{cases} d & \text{if } a < c, \\ \frac{1}{2}(a + c) & \text{if } a \geq c \end{cases} \quad \textit{and} \quad H(a, c, d) = \begin{cases} \frac{1}{2}(a + c) & \text{if } a > c \vee d, \\ d & \text{if } a \leq c \vee d \end{cases}$$

$L(x, y, z)$ and $H(x, y, z)$ represent the lowest and highest equilibrium expected payoffs of player 1 in $G_n(a, b;\ d, d)$. As we aforementioned, we are interested in computing the expected payoff corresponding to the best and the worst NE of the game $\Gamma_n^{FR}(a, b)$, which correspond to the extreme values of the set $\mathbb{E}_n(a, b)$ denoted $l_n(a, b)$ and $h_n(a, b)$. Using the monotony of ψ (and thus of L and H) with respect to d, $l_n(a, b)$ is the lowest equilibrium payoff of $G_n(a, b;\ d, d)$, i.e. $L(a, \mathbb{E}(X_{(n)} \vee b), d)$, when d is the expected continuation payoff obtained by playing the lowest equilibrium in every continuation game, that is:

$$d = \min \mathbb{E}(E_{n-1}(a \vee X, med[a, b, X])) = \mathbb{E}_X \left[l_{n-1}(a \vee X, \text{med}[a, b, X])\right].$$

where $l_0(a, b) = \frac{a+b}{2}$. This completes the proof for l_n and analogous analysis applies for $h_n(a, b)$. □

Remark 1. *Note that we did not use any assumption on the distribution F except that its support is contained in $[0, 1]$, therefore Theorem 1 holds when considering discrete distributions.*

Using Theorem 1, we now prove Theorem 2.

Proof of Theorem 2 By Theorem 1, we know that $E_n^{FR} \subset \{(u, u) : l_n \leq u \leq h_n\}$.

Furthermore, $E_n^{FR} = \left\{\int_{[0,1]} f(x) \mathrm{d}F(x) : f(x) \in \mathbb{E}_{n-1}(x, 0)\right\}$, and then (l_n, l_n) and (h_n, h_n) belong to E_n^{FR} (we just take $f(x) = (l_{n-1}(x, 0), l_{n-1}(x, 0))$ and $f(x) = (h_{n-1}(x, 0), h_{n-1}(x, 0))$, respectively). To obtain the result, it is then enough to prove that the set E_n is convex. To this end, let us define the function $\mu : [0, 1] \to \mathbb{R}$ by

$$\mu(\alpha) = \int_{[0,\alpha)} l_{n-1}(x, 0) \mathrm{d}F(x) + \int_{[\alpha,1]} h_{n-1}(x, 0) \mathrm{d}F(x).$$

Notice that μ is continuous since F is atomless, $\mu(0) = h_n$, and $\mu(1) = l_n$, then using the intermediate value theorem all values between l_n and h_n are taken by μ. But $(\mu(\alpha), \mu(\alpha))$ belongs to E_n for every α by choosing $f(x)$ equal to $(l_{n-1}(x, 0), l_{n-1}(x, 0))$ on $[0, \alpha)$ and to $(h_{n-1}(x, 0), h_{n-1}(x, 0))$ on $[\alpha, 1]$. Therefore E_n is convex. □

No recall case. We pass now to the no recall case and the goal is to prove Theorem 3.

Recall that the game $\Gamma_n^{NR}(a)$ is defined as Γ_n^{NR}, but there are $n+1$ stages: at the first stage, the players can bid for a value a, and the values for the next n stages are randomly drawn as in Γ_n^{NR}.

At first, in order to study the SPEP of Γ_n^{NR}, note that this game has the same SPEP as the game which is identical to Γ_n^{NR} but terminates at the first time a player picks a value (or after stage n). If the game terminates because one player did get a realization, the payoff of the other player is the value of the one-player continuation problem. As for the full recall case, it is easy to check that after a player picks a value, the SPEP in the continuation game are unique and correspond to the payoffs of this auxiliary game. In the following, we assume that Γ_n^{NR} is the auxiliary game. In this game, a history is just a sequence of values $(X_1, ..., X_t)$, and the strategy of player i at time $t < n$ is a measurable map $\sigma_{i,t}$ from histories into the probabilities over $\{\emptyset\} \cup \{1, ..., t\}$. We use the same identification for the game $\Gamma_n^{NR}(a)$.

Given a history $h = (x_1, ..., x_t)$ of length t, let $\Gamma_n^{NR}(h)$ denote the subgame of Γ_{n+t}^{NR} starting at stage t after observing h in which the two players are still present and $E_n^{NR}(h)$ the set of SPEP of this game.

In the following, variables $(X, X_1, X_2, ...)$ will denote independent variables with distribution F, and $X_{(t)} = \max(X_1, ..., X_t)$.

Without loss of generality, we assume that $\mathbb{P}(X > 0) > 0$ (otherwise the set of equilibrium payoffs is reduced to $\{(0, 0)\}$).

As for the full recall case, we state without proof some properties of SPEP which follow easily from usual arguments in dynamic game theory.

Lemma 8. *The following properties hold:*

1. $E_0^{NR} = \{(0,0)\}$ *and for* $n \geq 1$, $E_n^{NR} = E_n^{NR}(0) = \mathbb{E}(E_{n-1}^{NR}(X))$.
2. *If* $h = (x_1, ..., x_t)$ *is a history of length* $t \geq 1$, *then* $E_n^{NR}(h) = E_n^{NR}(x_t)$.
3. $(x, y) \in E_n^{NR}(a)$ *if and only if there exists* $(d, e) \in E_n^{NR}$ *such that* (x, y) *is a mixed Nash equilibrium payoff of the finite game* $G_n(a;\ d, e)$ *with payoff matrix:*

$$\begin{array}{c} \\ a \\ \emptyset \end{array} \begin{array}{c} \begin{array}{cc} \quad\quad\quad a \quad\quad\quad & \quad\quad \emptyset \end{array} \\ \begin{pmatrix} \left(\frac{1}{2}(a+c_n), \frac{1}{2}(a+c_n)\right) & (a, c_n) \\ (c_n, a) & (d, e) \end{pmatrix} \end{array}$$

Fig. 10: Payoff matrix of $G_n(a;\ d, e)$.

where c_n *denotes the value of the decision problem with* n *stages in a standard prophet setting.*

4. *Similarly, (σ_1, σ_2) is a pair of first-stage strategies of some SPE in $\Gamma_n^{NR}(a)$ with payoff (x, y) if and only if there exists $(d, e) \in E_n^{NR}$ such that (σ_1, σ_2) is a mixed Nash equilibrium of the matrix game $G_n(a;\ d, e)$ with payoff (x, y).*

Before concluding the section with the proof of Theorem 3, we show Proposition 1.

Proof of Proposition 1. The fact that E_n^{NR} is symmetric with respect to the diagonal is a direct consequence of the fact that the game is symmetric.

To prove convexity, we can use the properties of the expectation of a set-valued map, also called the Aumann integral. Theorem 8.6.3 in Aubin and Frankowska (2009) implies that the expectation of a set-valued map with non-empty closed values and a compact graph from $[0, 1]$ to $\mathbb{R}^2$ with respect to an atomless measure on $[0, 1]$ is a non-empty convex compact set.

First, $E_0^{NR} = \{(0, 0)\}$ is non-empty compact convex. Then, $a \to E_0^{NR}(a)$ is a set-valued map with non-empty closed values and compact graph using the classical properties of Nash equilibrium payoffs of matrix games. It follows that E_1^{NR} is non-empty compact convex. Let us assume that E_n^{NR} is non-empty compact convex. Let $NEPG_n(a;\ d, e)$ denote the set of Nash equilibrium payoffs of $G_n(a;\ d, e)$, then $(a, d, e) \in [0, 1] \times E_n^{NR} \to NEPG_n(a;\ d, e)$ is a set-valued map with non-empty closed values and a compact graph, and thus $a \to E_n^{NR}(a)$ is a set-valued map with non-empty closed values and a compact graph. We conclude that E_{n+1}^{NR} is a non-empty compact convex set. $\square$

Below, we prove the theorem.

Proof of Theorem 3. Given $a \in [0, 1]$ and n a natural number, we consider the game $\Gamma_n^{NR}(a)$ defined in Section 3.2. We first analyze the game $G_n(a;\ d, e)$ described in Fig. 10 with $(d, e) \in E_n^{NR}$.

One first remark regarding this game is that what a player gets if he stays alone in the game, that is c_n, is at least what he gets if both stay. In other words, $c_n \geq d$ and $c_n \geq e$.

Now, we use Fig. 10 and the remark above to study the NE of the game $G_n(a;\ d, e)$, depending on the relation between the parameters a, c_n, d, e. It is easy to check that the following hold:

1. If $a > c_n$, (a, a) is the unique NE with payoff $\left(\frac{1}{2}(a + c_n), \frac{1}{2}(a + c_n)\right)$.
2. If $a = c_n$, there is a unique NE payoff (c_n, c_n).
3. If $c_n > a > d \vee e$ there are two pure NE $(a, \emptyset)$ and $(\emptyset, a)$, and a mixed equilibrium; with payoffs $(a, c_n), (c_n, a)$ and (γ_n^1, γ_n^2), respectively, with $\gamma_n^1 = \frac{2ac_n - d(c_n + a)}{c_n + a - 2d}$, $\gamma_n^2 = \frac{2ac_n - e(c_n + a)}{c_n + a - 2e}$.
4. If $c_n > a = d > e$, the NE payoffs are (c_n, a) and (a, λ) with $\lambda \in [\gamma_n^2, c_n]$.
5. If $c_n > a = e > d$, the NE payoffs are (a, c_n) and (λ, a) with $\lambda \in [\gamma_n^1, c_n]$.
6. If $c_n > a = d = e$, the NE payoffs are (a, λ) and (λ, a) with $\lambda \in [a, c_n]$.
7. If $d \wedge e > a$, $(\emptyset, \emptyset)$ is the unique NE with payoff (d, e).
8. If $d > a > e$, $(\emptyset, a)$ is the unique NE with payoff (c_n, a).
9. If $e > a > d$, $(a, \emptyset)$ is the unique NE with payoff (a, c_n).
10. If $d > a = e$, the NE payoffs are (d, e) and (λ, a) with $\lambda \in [d, c_n]$.
11. If $e > a = d$, the NE payoffs are (d, e) and (a, λ) with $\lambda \in [e, c_n]$.

In particular, we are interested in computing the sum of the expected payoffs corresponding to the best and the worst SPEs that is, $\max\{x+y : (x,y) \in E_{n+1}^{NR}\}$ and $\min\{x+y : (x,y) \in E_{n+1}^{NR}\}$, respectively; as well as the worst payoff a player can get at equilibrium, that is $\min\{\min\{x,y\} : (x,y) \in E_n^{NR}\}$.

Due to Proposition 1, we have that

$$\max\{x+y : (x,y) \in E_{n+1}^{NR}\} = 2\max\{x : (x,x) \in E_{n+1}^{NR}\},$$

and

$$\min\{x+y : (x,y) \in E_{n+1}^{NR}\} = 2\min\{x : (x,x) \in E_{n+1}^{NR}\}.$$

Therefore, defining $\alpha_n := \min\{x : (x,x) \in E_n^{NR}\}$, $\beta_n := \max\{x : (x,x) \in E_n^{NR}\}$ and $\alpha_n' = \min\{\min\{x,y\} : (x,y) \in E_n^{NR}\}$, it follows that it is enough to compute $2\alpha_{n+1}$, $2\beta_{n+1}$ and α_{n+1}'.

Let us start by proving point c) of the theorem. To this end, we compute $2\alpha_{n+1}$ as a function of $\alpha_n', \alpha_n, \beta_n$ and c_n.

Define $2\alpha_n(a)$ as the minimal sum of payoffs in any NE of the family of games $G_n(a;\ d,e)$ when (d,e) ranges through $E_n^{NR}(a)$, so that $\alpha_{n+1} = \int_{[0,1]} \alpha_n(a)\mathrm{d}F(a)$. Notice that if $c_n > a > d \vee e$, the mixed equilibrium gives a worse sum of payoffs than the pure. We claim that:

$$2\alpha_n(a) = \begin{cases} 2\alpha_n & \text{if } a < \alpha_n', \\ \min\{2\alpha_n, a+c_n\} & \text{if } \alpha_n' < a < \alpha_n, \\ 2a & \text{if } \alpha_n < a < \beta_n, \\ 2\frac{2ac_n - \beta_n(a+c_n)}{c_n+a-2\beta_n} & \text{if } \beta_n < a < c_n, \\ a+c_n & \text{if } a > c_n. \end{cases}$$

The proof of the previous equality follows directly from the characterization of NE payoffs of $G_n(a;\ d,e)$ (i.e. the list of 11 cases detailed at the beginning of the present proof). To illustrate the method, let us prove that $\alpha_n(a) = 2a$ when $\alpha_n < a < \beta_n$. The other cases, and the corresponding equalities for $\alpha_n'(a)$ and $\beta_n(a)$ defined below, can be proven using similar arguments.

First, as $c_n \geq \beta_n$, we have $c_n > a$, and thus the 2 first cases in the list are irrelevant. Second, it is straightforward to check that $\gamma_n^1 \geq a$ (resp. $\gamma_n^2 \geq a$) with equality if and only if $d = a$ (resp. $e = a$), and therefore the sum of payoffs in a NE in cases 3,4, and 5 is larger than or equal to $2a$. In the cases from 7 to 11, the sum of NE payoffs is always larger than or equal to $2a$ using the assumptions on (d,e) and that $c_n > a$. The last remaining case is case 6, where $d = e = a < c_n$. Due to Proposition 1, the set E_n^{NR} contains the points (α_n, α_n) and (β_n, β_n). By convexity, it also contains the point (a,a) as $\alpha_n < a < \beta_n$. We deduce that the choice $d = e = a$ is possible and the smallest sum of NE payoffs in this case is $2a$ by taking $\lambda = a$. We conclude that $\alpha_n(a) = 2a$ because any sum of NE payoffs in $G_n(a;\ d,e)$ for $(d,e) \in E_n^{NR}$ is always larger than or equal to $2a$, and this value is attained when $d = e = a$. Note that choosing $(d,e) = (\alpha_n, \alpha_n)$, the smallest sum of NE payoffs is strictly larger compared to the one obtained by taking $(d,e) = (a,a)$, despite the fact that $\alpha_n < a$. This contrasts with the results obtained in the model with recall.

Since F is atomless, it is sufficient to conclude that for $n \geq 1$

$$\begin{aligned} 2\alpha_{n+1} &= \int_{c_n}^{1} (a+c_n)\mathrm{d}F(a) + \int_{\beta_n}^{c_n} \psi_n(a)\mathrm{d}F(a) \\ &+ \int_{\alpha_n}^{\beta_n} 2a\mathrm{d}F(a) + \int_{\alpha_n'}^{\alpha_n} \xi_n(a)\mathrm{d}F(a) + \int_0^{\alpha_n'} 2\alpha_n \mathrm{d}F(a), \end{aligned} \tag{4}$$

where $\psi_n(a) = \frac{4ac_n - 2\beta_n(a+c_n)}{c_n+a-2\beta_n}$ and $\xi_n(a) = \min\{2\alpha_n, a+c_n\}$, which is the third statement of the theorem.

To prove part a) of the theorem, we compute a recursive formula for α'_{n+1}. Define $\alpha'_n(a)$ as the minimal NE payoff of player 1 in the family of games $G_n(a;\ d, e)$ when (d, e) ranges through $E_n^{NR}(a)$. It is clear from the previous characterization that $\alpha'_{n+1} = \int_{[0,1]} \alpha'_n(a)\mathrm{d}F(a)$. Using our previous analysis, notice that:

$$\alpha'_n(a) = \begin{cases} \alpha'_n & \text{if } a < \alpha'_n, \\ a & \text{if } \alpha'_n < a < c_n, \\ \frac{a+c_n}{2} & \text{if } a > c_n. \end{cases}$$

Since F is atomless, it is sufficient to conclude that for $n \geq 1$

$$\begin{aligned} \alpha'_{n+1} &= \int_0^{\alpha'_n} \alpha'_n \mathrm{d}F(a) + \int_{\alpha'_n}^{c_n} a\mathrm{d}F(a) + \int_{c_n}^1 \frac{a+c_n}{2}\mathrm{d}F(a) \\ &= \alpha'_n(F(\alpha'_n) - F(0)) + (c_nF(c_n) - \alpha'_nF(\alpha'_n)) - \int_{\alpha'_n}^{c_n} F(a)\mathrm{d}a \\ &+ \frac{c_n}{2}(F(1) - F(c_n)) + \frac{1}{2}(1F(1) - c_nF(c_n)) - \frac{1}{2}\int_{c_n}^1 F(a)\mathrm{d}a. \\ &= \frac{(c_n+1)}{2} - \int_{\alpha'_n}^{c_n} F(a)\mathrm{d}a - \frac{1}{2}\int_{c_n}^1 F(a)\mathrm{d}a, \end{aligned} \tag{5}$$

and the first statement of Theorem 3 is proved.

Next, we compute $2\beta_{n+1}$ as a function of β_n, α'_n, and c_n. Define $2\beta_n(a)$ as the maximal sum of payoffs in any NE on the family of games $G_n(a;\ d, e)$ when (d, e) ranges through $E_n^{NR}(a)$, so that $\beta_{n+1} = \int_{[0,1]} \beta_n(a)\mathrm{d}F(a)$. We have:

$$2\beta_n(a) = \begin{cases} 2\beta_n & \text{if } a < \alpha'_n, \\ \max\{a+c_n, 2\beta_n\} & \text{if } \alpha'_n < a < \beta_n, \\ a + c_n & \text{if } a > \beta_n. \end{cases}$$

Since F is atomless, it is sufficient to conclude that for $n \geq 1$

$$2\beta_{n+1} = \int_0^{\alpha'_n} 2\beta_n\mathrm{d}F(a) + \int_{\alpha'_n}^{\beta_n} \max\{a+c_n, 2\beta_n\}\mathrm{d}F(a) + \int_{\beta_n}^1 (a+c_n)\mathrm{d}F(a), \tag{6}$$

and the second statement of the theorem is obtained.

Putting together all the foregoing analysis, we obtain the desired result, concluding the proof. □

4.2 Omitted proofs from Section 3.3

Proof of Lemma 1. We assume that one player, let us say player 1, bids in the first stage if $a \vee X_1 \geq c_n$, and passes otherwise, where X_1 is the realization of the first random variable arrived. Let us show that player 1 obtains an expected payoff, namely γ, of at least $(a+c_{n+1})/2$, independently of what player 2 does, and therefore player 1 will obtain a payoff of at least $(a+c_{n+1})/2$ playing any SPE.

We divide the proof into two cases:

Case 1. Assume that $a \geq c_n$. Note that in this case, player 1 bids for $M = a \vee X_1$, and we have that

$$\gamma = \mathbb{E}_{X_1}\left((a \vee X_1)\mathbb{P}(A) + \frac{1}{2}\left((a \vee X_1) + \mathbb{E}\left((a \wedge X_1) \vee X_{(n)}\right)\right)\mathbb{P}(A^c)\right),$$

where A is the event *player 2 does not bid for M.*

Notice that $\left((a \vee X_1) + \mathbb{E}\left((a \wedge X_1) \vee X_{(n)}\right)\right)$ is lower bounded by $a + c_n$ and by $a + X_1$ and therefore we have that

$$\begin{aligned}
\gamma &\geq \mathbb{E}_{X_1}\left((a \vee X_1)\mathbb{P}(A) + \frac{1}{2}\left((a + c_n)\mathbf{1}_{\{X_1 < c_n\}} + (a + X_1)\mathbf{1}_{\{X_1 \geq c_n\}}\right)\mathbb{P}(A^c)\right) \\
&= \frac{a}{2} + \frac{1}{2}\left(\mathbb{P}(A)(2(a \vee X_1) - a) + \mathbb{P}(A^c)\left(c_n\mathbf{1}_{\{X_1 < c_n\}} + X_1\mathbf{1}_{\{X_1 \geq c_n\}}\right)\right) \\
&\geq \frac{a}{2} + \frac{1}{2}\left(c_n\mathbf{1}_{\{X_1 < c_n\}} + X_1\mathbf{1}_{\{X_1 \geq c_n\}}\right) = \frac{a}{2} + \frac{c_{n+1}}{2},
\end{aligned}$$

where the second inequality holds because

$$2(a \vee X_1) - a \geq c_n\mathbf{1}_{\{X_1 < c_n\}} + X_1\mathbf{1}_{\{X_1 \geq c_n\}}$$

due to $a > c_n$, and the last equality because $c_{n+1} = \mathbb{E}(X \vee c_n)$. We conclude that $\gamma \geq (a + c_{n+1})/2$ if $a \geq c_n$.

Case 2. Assume that $a < c_n$. In this case, we have that

$$\begin{aligned}
\gamma &\geq \mathbb{E}_{X_1}\left(\left(X_1 \wedge \frac{1}{2}\left(X_1 + \mathbb{E}\left(a \vee X_{(n)}\right)\right)\right)\mathbf{1}_{\{X_1 \geq c_n\}}\right) \\
&+ \mathbb{E}_{X_1}\left(\left(\frac{(a \vee X_1) + c_n}{2} \wedge \mathbb{E}\left(X_{(n)}\right)\right)\mathbf{1}_{\{X_1 < c_n\}}\right) \\
&= \mathbb{E}_{X_1}\left(\left(\frac{X_1}{2} + \frac{1}{2}\left(X_1 \wedge \mathbb{E}\left(a \vee X_{(n)}\right)\right)\right)\mathbf{1}_{\{X_1 \geq c_n\}}\right) \\
&+ \mathbb{E}_{X_1}\left(\left(\frac{c_n}{2} + \frac{1}{2}\left((a \vee X_1) \wedge \left(2\mathbb{E}\left(X_{(n)}\right) - c_n\right)\right)\right)\mathbf{1}_{\{X_1 < c_n\}}\right) \\
&= \frac{c_{n+1}}{2} + \mathbb{E}_{X_1}\left(\frac{1}{2}\left(X_1 \wedge \mathbb{E}\left(a \vee X_{(n)}\right)\right)\mathbf{1}_{\{X_1 \geq c_n\}}\right) \\
&+ \mathbb{E}_{X_1}\left(\frac{1}{2}\left((a \vee X_1) \wedge \left(2\mathbb{E}\left(X_{(n)}\right) - c_n\right)\right)\mathbf{1}_{\{X_1 < c_n\}}\right) \\
&\geq \frac{a + c_{n+1}}{2},
\end{aligned}$$

where the second equality holds because $c_{n+1} = \mathbb{E}(X \vee c_n)$ and the last inequality because both $X_1 \wedge \mathbb{E}\left(a \vee X_{(n)}\right)$ and $(a \vee X_1) \wedge \left(2\mathbb{E}\left(X_{(n)}\right) - c_n\right)$ are at least a when $a < c_n$. Thus, we conclude that $\gamma \geq (a + c_{n+1})/2$ if $a < c_n$.

Putting all together we obtain that $\gamma_n^{FR} \geq (a+c_{n+1})/2$ and the proof is complete. □

Proof of Theorem 4 Note first that $v_1 + v_2 \leq 2\beta_n$ since the set of SPE is symmetric with respect to the diagonal and convex, and β_n denotes the highest symmetric SPE payoff in the no recall case. Indeed, assuming that $v_1 + v_2 > 2\beta_n$, then by symmetry (v_2, v_1) is also a SPE payoff. Using convexity $(\frac{1}{2}(v_1 + v_2), \frac{1}{2}(v_1 + v_2))$ is a symmetric SPE payoff so that $\frac{1}{2}(v_1 + v_2) > \beta_n$, contradicting the definition of β_n. Therefore, to prove Theorem 4 it is enough to show that $l_n \geq \beta_n$ for all n, where l_n denotes the lowest SPE payoff in the full recall case. To this end, we will prove by induction on n that $l_n(a,0) \geq \beta_n(a)$ for all n and for all a, where $l_n(a,0)$ is defined as in the statement of Theorem 1 and $\beta_n(a)$ is defined by

$$\beta_n(a) = \begin{cases} \beta_n & \text{if } a < \alpha'_n \\ \max\{\beta_n, \frac{a+c_n}{2}\} & \text{if } a \in [\alpha'_n, \beta_n] \\ \frac{a+c_n}{2} & \text{if } a > \beta_n, \end{cases}$$

where c_n is the value of the decision problem in the no recall case with one decision-maker and n arrivals.

First, notice that $l_1(a,0) = (a + \mathbb{E}(X))/2$, $c_1 = \mathbb{E}(X)$ and $\beta_1 = \mathbb{E}(X)$, and thus $l_1(a,0) \geq \beta_1$ and $l_1(a,0) \geq (a+c_1)/2$ for all a, concluding that $l_1(a,0) \geq \beta_1(a)$ for all a.

We assume now that $l_n(a,0) \geq \beta_n(a)$ for all a and we prove that $l_{n+1}(a,0) \geq \beta_{n+1}(a)$ for all a. We divide the proof into two cases depending on whether a is greater than $\mathbb{E}\left(X_{(n+1)}\right)$ or not.

Case 1. Assume that $a > \mathbb{E}\left(X_{(n+1)}\right)$. In this case, we have that $a > \beta_{n+1}$ because $\mathbb{E}\left(X_{(n+1)}\right) \geq c_{n+1} \geq \beta_{n+1}$. Thus, by the definition of $\beta_{n+1}(a)$, we have that

$$\beta_{n+1}(a) = \frac{a + c_{n+1}}{2}. \tag{7}$$

On the other hand, it holds that

$$l_{n+1}(a,0) = \frac{a + \mathbb{E}\left(X_{(n+1)}\right)}{2} \geq \frac{a + c_{n+1}}{2}, \tag{8}$$

where the equality follows from using the definition of $l_{n+1}(a,0)$ when $a > \mathbb{E}\left(X_{(n+1)}\right)$ and the inequality holds because $\mathbb{E}\left(X_{(n+1)}\right) \geq c_{n+1}$.

Therefore, we obtain $l_{n+1}(a,0) \geq \beta_{n+1}(a)$ for all $a > \mathbb{E}\left(X_{(n+1)}\right)$ combining (7) and (8).

Case 2. Assume that $a \leq \mathbb{E}\left(X_{(n+1)}\right)$. To prove that $l_{n+1}(a,0) \geq \beta_{n+1}(a)$, we show that $l_{n+1}(a,0) \geq \beta_{n+1}$ and that $l_{n+1}(a,0) \geq (a + c_{n+1})/2$. The latter holds by Lemma 1. On the other hand,

$$l_{n+1}(a,0) = \mathbb{E}_X(l_n(a \vee X, a \wedge X)) \geq \mathbb{E}_X(l_n(X,0)) \geq \mathbb{E}_X(\beta_n(X)) = \beta_{n+1},$$

where the first equality follows from the definition of $l_{n+1}(a,0)$ for $a \leq \mathbb{E}\left(X_{(n+1)}\right)$, the first inequality holds due to the monotonicity of the function l_n in both components, the second inequality follows from the induction hypothesis, and the last equality from the definition of β_{n+1}.

Therefore, we conclude that $l_{n+1}(a,0) \geq \beta_{n+1}(a)$ for all $a \leq \mathbb{E}\left(X_{(n+1)}\right)$.

Putting all together we obtain that $l_n \geq \beta_n$ for all n and the desired result follows. □

4.3 Omitted proofs from Section 3.4

The main result in Section 3.4 we prove here is Theorem 5 and then we work on the competitive selection problem with no recall and two arrivals.

Before going to the proof, we obtain an expression for the price of anarchy and price of stability if the random variables are distributed according to F, namely $\mathrm{PoA}_2^{NR}(F)$ and $\mathrm{PoS}_2^{NR}(F)$ respectively.

Regarding the price of stability, we need to compute

$$\frac{\mathbb{E}\left(X_{(1:2)}+X_{(2:2)}\right)}{2\beta_2},$$

where $2\beta_2$ is obtained from equation (6) by taking $n=1$. That is:

$$2\beta_2 = \int_0^{\alpha_1'} 2\beta_1 \mathrm{d}F(a) + \int_{\alpha_1'}^{\beta_1} \max\{a+c_1, 2\beta_1\}\mathrm{d}F(a) + \int_{\beta_1}^{1} (a+c_1)\mathrm{d}F(a).$$

Now, due to $E_1^{NR} = \left\{\left(\frac{\mathbb{E}(X)}{2}, \frac{\mathbb{E}(X)}{2}\right)\right\}$, we have that $\alpha_1' = \beta_1 = \frac{\mathbb{E}(X)}{2}$. Noting that $c_1 = \mathbb{E}(X)$ and putting all together we have

$$\begin{aligned}
2\beta_2 &= \int_0^{\mathbb{E}(X)/2} \mathbb{E}(X)\mathrm{d}F(a) + \int_{\mathbb{E}(X)/2}^{\mathbb{E}(X)/2} \max\{a+\mathbb{E}(X), \mathbb{E}(X)\}\mathrm{d}F(a)\\
&+ \int_{\mathbb{E}(X)/2}^{1} (a+\mathbb{E}(X))\mathrm{d}F(a) = \int_0^1 \mathbb{E}(X)\mathrm{d}F(a) + \int_{\mathbb{E}(X)/2}^{1} a\mathrm{d}F(a)\\
&= \mathbb{E}(X) + \mathbb{P}(X \geq \mathbb{E}(X)/2)\mathbb{E}(X \mid X \geq \mathbb{E}(X)/2).
\end{aligned}$$

On the other hand, $\mathbb{E}\left(X_{(1:2)}+X_{(2:2)}\right) = 2\mathbb{E}(X)$, and therefore

$$\begin{aligned}
\frac{1}{\mathrm{PoS}_2^{NR}(F)} &= \frac{\mathbb{E}(X) + \mathbb{P}(X \geq \mathbb{E}(X)/2)\mathbb{E}(X \mid X \geq \mathbb{E}(X)/2)}{2\mathbb{E}(X)}\\
&= 1/2 + \frac{\mathbb{P}(X \geq \mathbb{E}(X)/2)\mathbb{E}(X \mid X \geq \mathbb{E}(X)/2)}{2\mathbb{E}(X)}.
\end{aligned} \tag{9}$$

Regarding the price of anarchy, we need to compute

$$\frac{\mathbb{E}\left(X_{(1:n)}+X_{(2:n)}\right)}{2\alpha_2},$$

where $2\alpha_2$ follows from equation (4) by taking $n=1$. After some algebra, we obtain

$$\begin{aligned}
2\alpha_2 &= \mathbb{E}(X) + \int_{\mathbb{E}(X)/2}^{\mathbb{E}(X)} \left(2\mathbb{E}(X) - \frac{\mathbb{E}(X)^2}{a}\right)\mathrm{d}F(a) + \int_{\mathbb{E}(X)}^{1} a\mathrm{d}F(a) \qquad (10)\\
&= 2\mathbb{E}(X) - \int_0^{\mathbb{E}(X)/2} a\mathrm{d}F(a) - \int_{\mathbb{E}(X)/2}^{\mathbb{E}(X)} \left(a - 2\mathbb{E}(X) + \frac{\mathbb{E}(X)^2}{a}\right)\mathrm{d}F(a),
\end{aligned}$$

and thus

$$\frac{1}{\mathrm{PoA}_2^{NR}(F)} = 1 - \frac{1}{2\mathbb{E}(X)} \int_0^{\mathbb{E}(X)/2} a \mathrm{d}F(a) \\ - \frac{1}{2\mathbb{E}(X)} \int_{\mathbb{E}(X)/2}^{\mathbb{E}(X)} \left(a - 2\mathbb{E}(X) + \frac{\mathbb{E}(X)^2}{a}\right) \mathrm{d}F(a). \tag{11}$$

Also, we can write the inverse of the price of anarchy as follows:

$$\frac{1}{\mathrm{PoS}_2^{NR}(F)} - \frac{1}{2\mathbb{E}(X)} \int_{\mathbb{E}(X)/2}^{\mathbb{E}(X)} \left(a - 2\mathbb{E}(X) + \frac{\mathbb{E}(X)^2}{a}\right) \mathrm{d}F(a). \tag{12}$$

We now prove Theorem 5 using the formulas obtained above.

Proof of Theorem 5. To prove that $\mathrm{PoS}_2^{NR}(F) \leq 4/3$, let us consider the second term in the rhs of (9) and notice that

$$\begin{aligned} \mathbb{P}(X \geq \mathbb{E}(X)/2)\mathbb{E}(X \mid X \geq \mathbb{E}(X)/2) &= \mathbb{E}(X) - \mathbb{P}(X < \mathbb{E}(X)/2)\mathbb{E}(X \mid X < \mathbb{E}(X)/2) \\ &> \mathbb{E}(X) - \mathbb{P}(X < \mathbb{E}(X)/2)\frac{\mathbb{E}(X)}{2} \\ &\geq \mathbb{E}(X) - \frac{\mathbb{E}(X)}{2} = \frac{\mathbb{E}(X)}{2}. \end{aligned}$$

Then,

$$\frac{1}{\mathrm{PoS}_2^{NR}(F)} \geq \frac{1}{2} + \frac{\mathbb{E}(X)}{2}\frac{1}{2\mathbb{E}(X)} = \frac{3}{4}.$$

Regarding the price of anarchy, by equation (11) and using that $a - 2\mathbb{E}(X) + \frac{\mathbb{E}(X)^2}{a} \leq \frac{\mathbb{E}(X)}{2}$ if $a \in [\mathbb{E}(X)/2, \mathbb{E}(X)]$, we have

$$\frac{1}{\mathrm{PoA}_2^{NR}(F)} \geq 1 - \frac{\mathbb{P}(X < \mathbb{E}(X)/2)\mathbb{E}(X \mid X < \mathbb{E}(X)/2)}{2\mathbb{E}(X)} - \frac{1}{4}\mathbb{P}(\mathbb{E}(X)/2 \leq X \leq \mathbb{E}(X)).$$

But, $\mathbb{E}(X \mid X < \mathbb{E}(X)/2) \leq \mathbb{E}(X)/2$ and thus

$$\begin{aligned} \frac{1}{\mathrm{PoA}_2^{NR}(F)} &\geq 1 - \frac{1}{4}\mathbb{P}(X < \mathbb{E}(X)/2) - \frac{1}{4}\mathbb{P}(\mathbb{E}(X)/2 \leq X \leq \mathbb{E}(X)) \\ &= 1 - \frac{1}{4}\mathbb{P}(X \leq \mathbb{E}(X)) \geq \frac{3}{4}, \end{aligned}$$

obtaining the desired inequalities.

To prove the tightness of the bound, let us take two small positive real numbers $\varepsilon > 0$ and $\eta > 0$, and consider the random variable $X_{\varepsilon,\eta} = (1-\eta) \cdot X + \eta \cdot$ Unif[0, 1], where

$$X = \begin{cases} \varepsilon - \varepsilon^2 & \text{with probability } 1 - \varepsilon, \\ 1 & \text{with probability } \varepsilon. \end{cases}$$

Note that when η goes to 0,

$$\frac{1}{\mathrm{PoS}_2^{NR}(F_{\varepsilon,\eta})} \to \frac{1}{2} + \frac{\mathbb{P}(X \geq \mathbb{E}(X)/2)\mathbb{E}(X \mid X \geq \mathbb{E}(X)/2)}{2\mathbb{E}(X)},$$

where $F_{\varepsilon,\eta}$ represents the c.d.f. of $X_{\varepsilon,\eta}$, and therefore it is enough to prove that

$$\frac{1}{2}+\frac{\mathbb{P}(X\geq\mathbb{E}(X)/2)\mathbb{E}(X\mid X\geq\mathbb{E}(X)/2)}{2\mathbb{E}(X)}\to 3/4$$

when ε goes to 0.

The c.d.f. of X is given by

$$F_\varepsilon(x)=\begin{cases}0 & \text{if } x<\varepsilon-\varepsilon^2,\\ 1-\varepsilon & \text{if } x\in[\varepsilon-\varepsilon^2,1),\\ 1 & \text{if } x\geq 1;\end{cases}$$

and the expected value is $\mathbb{E}(X)=(\varepsilon-\varepsilon^2)(1-\varepsilon)+\varepsilon=\varepsilon(1+(1-\varepsilon)^2)$.

Therefore, after some algebra, it follows that

$$\begin{aligned}\frac{1}{2}+\frac{\mathbb{P}(X\geq\mathbb{E}(X)/2)\mathbb{E}(X\mid X\geq\mathbb{E}(X)/2)}{2\mathbb{E}(X)} &= \frac{1}{2}+\frac{\varepsilon}{2\varepsilon(1+(1-\varepsilon)^2)}\\ &= \frac{1}{2}+\frac{1}{2(1+(1-\varepsilon)^2)},\end{aligned}$$

which converges to $3/4$ when $\varepsilon\to 0$, and therefore we obtain price of stability $4/3$.

On the other hand, note that by definition, for each distribution F, it holds that $\mathrm{PoS}_2(F)\leq\mathrm{PoA}_2(F)$, and thus $\lim_{\varepsilon,\eta\to 0}\mathrm{PoA}_2(F_{\varepsilon,\eta})\geq 4/3$, where $4/3$ is the upper bound we have already proven. Therefore, the bound is tight also for the price of anarchy. □

5 Future directions

This work constitutes a first step in the study of equilibria of prophet inequality problems with competition and recall, and it opens many avenues for future research.

A possible extension is to consider a sequence of arrivals with independent but not identical distributions. Here we believe that the natural recursive approach used in our paper can be extended, at the cost of writing much more complex expressions.

A second very natural extension is to consider a larger number of players. Here the extension is qualitatively more complex, and as soon as there are at least 3 players, under full recall it may be a best reply for a player to bid for the second best available value (and let the other players compete for the first best). The set of equilibrium payoffs will then likely no longer be a subset of the diagonal, i.e. we expect to have non symmetric equilibrium payoffs, even under the full recall assumption.

It is also natural to consider the case of players having different utility functions, i.e. when the utility of a player receiving the realization x can depend on the player (for instance, it could be x for player 1 and x^2 for player 2). How would this affect the set of equilibria ? For instance, starting from the current situation where both players utility functions are the same, could it be always better at equilibrium for player 1 if the utility function of player 2 becomes different ? A related paper in this direction is the one by Garnaev (2006).

Finally, different tie-breaking rules can also be considered. The interesting case where players have a fixed priority given by a predetermined order has already been considered in Ezra et al (2021b).

Declarations

Acknowledgments The authors are grateful to Victor Verdugo and the two anonymous referees for their input that helped us improve the paper.

Ethical Approval Ethical approval not applicable to this article.

Competing interests The authors have no competing interests to declare.

Authors' contributions All authors contributed equally to this work.

Funding The authors gratefully acknowledge funding from ANITI ANR-3IA Artificial and Natural Intelligence Toulouse Institute, grant ANR-19-PI3A-0004, and from the ANR under the Investments for the Future program, grant ANR-17-EURE- 0010. J. Renault also acknowledges the support of ANR MaSDOL-19-CE23-0017-01. One page abstract of this work appeared in SAGT 2022.

Availability of data and materials This declaration is not applicable to this article.

References

Abdelaziz FB, Krichen S (2007) Optimal stopping problems by two or more decision makers: a survey. Computational Management Science 4(2):89–111

Aubin JP, Frankowska H (2009) Set-valued analysis. Springer Science & Business Media

Baston V, Garnaev A (2005) Competition for staff between two departments. Game Theory and Applications 10:13–20

Chawla S, Hartline JD, Malec DL, et al (2010) Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the forty-second ACM symposium on Theory of computing, pp 311–320

Correa J, Foncea P, Hoeksma R, et al (2017) Posted price mechanisms for a random stream of customers. In: Proceedings of the 2017 ACM Conference on Economics and Computation, pp 169–186

Correa J, Foncea P, Pizarro D, et al (2019) From pricing to prophets, and back! Operations Research Letters 47(1):25–29

Ezra T, Feldman M, Kupfer R (2021a) On a competitive secretary problem with deferred selections. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021, pp 175–181

Ezra T, Feldman M, Kupfer R (2021b) Prophet inequality with competing agents. In: International Symposium on Algorithmic Game Theory, Springer, pp 112–123

Ferguson TS (1989) Who solved the secretary problem? Statistical science 4(3):282–289

Freeman P (1983) The secretary problem and its extensions: A review. International Statistical Review/Revue Internationale de Statistique pp 189–206

Fudenberg D, Tirole J (1991) Game theory, 1991. Cambridge, Massachusetts 393(12):80

Garnaev AY (2006) A game-theoretical model of competition for staff between two departments. Banach Center Publications 71:137–145

Gilbert JP, Mosteller F (1966) Recognizing the maximum of a sequence. Journal of the American Statistical Association 61(313):35–73

Hajiaghayi MT, Kleinberg R, Sandholm T (2007) Automated online mechanism design and prophet inequalities. In: AAAI, pp 58–65

Hill TP, Kertz RP (1982) Comparisons of stop rule and supremum expectations of iid random variables. The Annals of Probability pp 336–345

Immorlica N, Kleinberg R, Mahdian M (2006) Secretary problems with competing employers. In: International Workshop on Internet and Network Economics, Springer, pp 389–400

Karlin A, Lei E (2015) On a competitive secretary problem. In: Twenty-Ninth AAAI Conference on Artificial Intelligence

Kertz RP (1986) Stop rule and supremum expectations of iid random variables: a complete comparison by conjugate duality. Journal of multivariate analysis 19(1):88–112

Kesselheim T, Psomas A, Vardi S (2023) On hiring secretaries with stochastic departures. Operations Research

Krengel U, Sucheston L (1978) On semiamarts, amarts, and processes with finite value. Probability on Banach spaces 4:197–266

Lindley DV (1961) Dynamic programming and decision theory. Journal of the Royal Statistical Society: Series C (Applied Statistics) 10(1):39–51

Mazalov VV, Banin MV (2003) N-person best-choice game with voting. Game Theory and Applications 9:45–53

Radzik T, Szajowski K (1990) Sequential games with random priority. Sequential Analysis 9(4):361–377

Ramsey DM, Szajowski K (2006) Correlated equilibria in competitive staff selection problem. Banach Center Publications 71:253–265

Roughgarden T, Tardos É (2002) How bad is selfish routing? Journal of the ACM (JACM) 49(2):236–259

Sakaguchi M (2005) Optimal stopping games where players have weighted privilege. Advances in Dynamic Games: Applications to Economics, Finance, Optimization, and Stochastic Control pp 285–294

Sakaguchi M, Mazalov VV (2004) A non-zero-sum no-information best-choice game. Mathematical Methods of Operations Research 60:437–451

Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. the Annals of Probability pp 1213–1216

Sweet T (1994) Optimizing a single uncertain selection by recall of observations. Journal of applied probability 31(3):660–672

Yang MC (1974) Recognizing the maximum of a random sequence based on relative rank with backward solicitation. Journal of Applied Probability 11(3):504–512