A Constant Factor Prophet Inequality for Online Combinatorial Auctions

José Correa   Andrés Cristi

Universidad de Chile

April 2023

Abstract

In online combinatorial auctions m indivisible items are to be allocated to n agents who arrive online. Agents have random valuations for the different subsets of items and the goal is to allocate the items on the fly so as to maximize the total value of the assignment. A prophet inequality in this setting refers to the existence of an online algorithm guaranteed to obtain, in expectation, a certain fraction of the expected value obtained by an optimal solution in hindsight. The study of prophet inequalities for online combinatorial auctions has been an intensive area of research in recent years, and constant factor prophet inequalities are known when the agents' valuation functions are submodular or fractionally subadditive. Despite many efforts, for the more general case of subadditive valuations, the best known prophet inequality has an approximation guarantee of $O(\log \log m)$. In this paper, we prove the existence of a constant factor prophet inequality for the subadditive case, resolving a central open problem in the area.

Our prophet inequality is achieved by a novel, but elementary, sampling idea which we call the *Mirror Lemma*. This lemma is essentially concerned with understanding online algorithms for which the set of items that are allocated and those that are not, distribute equally. The other main ingredient is a nonstandard application of Kakutani's fixed point theorem. Finally, we note that our prophet inequality works against an almighty adversary and even can be implemented in an incentive compatible way.

Keywords: prophet inequality, combinatorial auction, subadditive valuation, mechanism design

1 Introduction

Efficiently distributing a set of valuable items to a set of agents is an old economic question dating back at least to the work of Leon Walras over a century ago. From an economics perspective a major difficulty is that the agents' valuations for the items is unknown. Therefore designing mechanisms, or combinatorial auctions, in which the agents have incentives to reveal their true preferences becomes central. In this space, pricing mechanisms have dominated the scene and understanding their effectiveness has been an active area of research in recent years [2, 1, 4, 9, 10, 16, 5, 23].

While the issue that agents' valuations for the items may be private information is an important difficulty, it is far from being the only. First, even if valuations are public information, the allocation problem is computationally hard. Second, in many situations agents arrive online and the set to be allocated has to be decided on the fly. The latter setting is known as that of online combinatorial auctions and constitutes the main topic of this paper.

Specifically, in online combinatorial auctions (or online combinatorial allocations) we are given a set of m items M and a set of n agents, denoted by N. Each agent $i \in N$ has a valuation function $v_i : 2^M \to \mathbb{R}_+$, which is randomly and independently chosen according to a given distribution $\mathcal{F}_i$ (defined over a set of possible valuation functions). As it is standard, we assume that each possible realization of each v_i is monotone, i.e., for all sets of items $A, B \subset M$ such that $A \subseteq B$ we have that $v_i(A) \le v_i(B)$. Agents arrive sequentially, and upon arrival their valuation function is realized. At the time agent i arrives, the online algorithm has to decide which set of items A_i, among those that are still available $(M \setminus \cup_{j=1}^{i-1} A_j)$, to allocate to the agent. This decision is irrevocable and the goal of the online algorithm is to maximize the social welfare of the allocation, i.e., the sum of the agents' valuations $\sum_{i\in N} v_i(A_i)$.

A prophet inequality in the online combinatorial allocations problem establishes the existence of an online algorithm, say ALG, such that the expected welfare of the resulting allocation is at least a certain fraction α of that of the optimal allocation in the hindsight.[1] That is, if we denote by ALG_i and OPT_i the (random) sets assigned to agent i by the algorithm and by the optimal allocation, a prophet inequality is an inequality of the form:

$$\mathbb{E}\left(\sum_{i\in N} v_i(ALG_i)\right) \ge \alpha \cdot \mathbb{E}\left(\sum_{i\in N} v_i(OPT_i)\right).$$

When there is a single item ($m = 1$), we land in the terrain of the classic prophet inequality of Krengel and Sucheston (and Garling) [21, 22], who proved that the largest possible α is 1/2. Various extensions of the single item prophet inequality to special cases of this general setting have been studied recently. These, for instance, model combinatorial constraints such as matching in a bipartite (or general) graph, hyper-graph matching, matroid constraints, among many others [19, 14, 20, 12, 7, 24].

A central and interesting research direction has focused on finding prophet inequalities for increasingly general classes of valuation functions. These classes not only model several combinatorial constraints, but also allow to express agnts' preferences in richer ways. The most prominent classes, in increasing order of generality, are:

[1] The allocation maximizing the social welfare once all valuation functions have been revealed.

Submodular. A valuation v is submodular if for all $A, B \subseteq M$, $v(A \cup B) \leq v(A) + v(B) - v(A \cap B)$.

Fractionally Subadditive or XOS. A valuation v is XOS if for every $A \subseteq M$ and every fractional covering $\{\lambda_i, T_i\}_{i=1}^k$ of A, i.e., such that $\lambda_i > 0$, $T_i \subseteq M$, and $\sum_{i:j \in T_i} \lambda_i \geq 1$ for all $j \in A$, we have that $v(A) \leq \sum_{i=1}^k \lambda_i v(T_i)$.

Subadditive. A valuation v is subadditive if for all $A, B \subseteq M$, $v(A \cup B) \leq v(A) + v(B)$.

The interest in these classes lies primarily because they model *complement-free* valuations. Interestingly, even in the offline setting, where the valuation functions are deterministic and the goal is to find an allocation of items to agents so as to maximize the sum of the valuations, the problem is $NP-$hard and the best known approximation guarantees are 1/2 for subadditive valuations [18], and $1 - 1/e$ if the valuations are XOS [8, 18].

In the online setting, Feldman, Gravin, and Lucier [17] obtain the best possible prophet inequality when valuations are XOS, and, consequently, also for submodular since 1/2 is the best possible factor even in the single item prophet inequality (see also [12]). The optimal factor improves to $1 - 1/e$ if the arrival order of the agents is random rather than arbitrary [13]. Despite these efforts, the existence of a constant factor prophet inequality for general subadditive valuations has remained unknown. In this context, Feldman et al. [17] obtained a prophet inequality of factor $\log(m)$ which was slightly improved by Zhang [26]. Very recently a major improvement, to a factor $O(\log\log(m))$, was obtained by Dütting, Kesselheim, and Lucier [11].

1.1 Our Results

In this paper we obtain a constant factor prophet inequality for online combinatorial auctions (or more precisely online combinatorial allocations). That is, we prove that there is an online algorithm such that

$$6 \cdot \mathbb{E}\left(\sum_{i \in N} v_i(ALG_i)\right) \geq \mathbb{E}\left(\sum_{i \in N} v_i(OPT_i)\right).$$

This resolves an open problem posed by Feldman et al. [17], which Dütting et al. [11] consider a *vexing* open problem in the area.

Similarly to several results in the area of prophet inequalities, we show that our approach works against an almighty adversary. This adversary can choose the arrival order adaptively, observing all past realizations of the valuations and the random choices of the algorithm in advance.

Our approach, unfortunately, is nonconstructive, so while we prove a 1/6 prophet inequality for subadditive valuations, we do not know how to implement the underlying online algorithm. Furthermore, our algorithm cannot be implemented with posted prices. This is in sharp contrast with the price based algorithms of Cai and Zhao [6] or that of Dütting et al. [11]. On the other hand, our algorithm critically needs access to the valuation of each agent upon their arrival, which is standard in the prophet inequality literature but not in the combinatorial auctions literature. However, we note that this is not really needed and the existence of a prophet inequality for the online combinatorial *allocation* problem easily implies the existence of a prophet inequality for online combinatorial *auctions*, i.e., with an incentive compatible mechanism. Therefore access to the valuation of each agent is not really needed. We expect that closing the gap of what can be achieved by price based mechanisms in the subadditive case will be of much interest in the future, as our result opens the way for the existence of such a mechanism with a constant factor welfare guarantee.

1.2 Main Technical Ingredients

To prove the main result we consider in Section 2 a family of algorithms that make use of *random score generators*. A Random Score Generator (RSG) is just a function that maps a valuation function v to a distribution over $\mathbb{R}_+^M$ of real numbers for each item. Given n RSGs, one for each agent, the algorithm is completely determined. First we sample one valuation function for each buyer v'_i and pass them through the RSGs to obtain scores b'_{ij} for each item and each agent. Then, when agent i arrives, and her valuation v_i is realized, we again pass it through its corresponding RSG to obtain scores $b_{i,j}$. Then we assign to agent i all remaining items for which $b_{i,j} > p'_j = \max_{k\in N} b'_{kj}$. The idea of this algorithm is inspired in the work of Azar, Kleinberg, and Weinberg [3] and that of Rubinstein, Wang, and Weinberg [25] who, for the single item prophet inequality use the maximum of samples from each distribution as a threshold for stopping the sequence.

To establish that this algorithm performs well for some appropriate selection of RSGs, we first prove the *Mirror* Lemma in Section 3. This lemma considers three copies of each agent with independent valuations v_i, v'_i, v''_i, pass them through the RSG and defines the sets $W_i = \{j \in M : b_{i,j} > \max p'_j, p''_j\}$. Then it asserts that

$$\mathbb{E}(ALG) := \mathbb{E}\left(\sum_{i\in N} v_i(ALG_i)\right) \geq \frac{1}{2}\mathbb{E}\left(\sum_{i\in N} v_i(W_i)\right).$$

The key observation to prove the Mirror Lemma is to note that the set of items that the algorithm allocates and those that the algorithm does not allocate, have essentially the same distribution. Noting that the Mirror Lemma holds for any collection of RSGs, the remaining task is to find suitable RSGs for which $\mathbb{E}\left(\sum_{i\in N} v_i(W_i)\right)$ is close to the optimal allocation's welfare.

As a warm-up, in Section 4 we show how a very simple family of RSGs is enough to give a constant factor prophet inequality for XOS valuations. Of course, this is not a new result and moreover the resulting constant is sub-optimal.[2] However we present it here to illustarate our approach more clearly. These RSGs are as follows: For a valuation v_i for agent i, sample the valuations of all other agents and compute OPT_i the set assigned to i in an optimal allocation. Then $b_{i,j}$ is just zero if $j \notin OPT_i$ and a uniform [0,1] random variable otherwise.

Our main result is then presented in Section 5. The proof borrows inspiration from a result by Feldman, Fu, Gravin, and Lucier [15, Lemma 1] where our RSGs can take the form of agents' bids in a combinatorial auction[3]. This first step takes the form of a basic inequality that we prove has a nonempty set of solutions. Then, this inequality is put into the context of a certain *Kakutani* map defined from the space of RSGs into itself. For this mapping we can establish the existence of a fixed point. Finally, it is relatively easy to prove that the RSGs corresponding to such fixed point

[2]The optimal constant, due to Feldamn et al. [17], is 1/2

[3]For the reader familiarized with the work of Feldman et al. our Lemma 4 is similar to their Lemma 1 but replacing their *First price auction* by an *All-pay auction*.

lead to an algorithm for which $(6+\varepsilon)\cdot\mathbb{E}\left(\sum_{i\in N} v_i(ALG_i)\right) \geq \mathbb{E}\left(\sum_{i\in N} v_i(OPT_i)\right)$. The application of Kakutani fixed point theorem in Section 5 requires the technical assumption that the set of possible valuation functions is finite. In Section 6 we relax this assumption by loosing an additional ε in the approximation factor.

In Section 7 we describe the optimal online solution using dynamic programming. As a corollary of our main result, we have that the optimal online algorithm attains a factor 6 prophet inequality. We also explain how the optimal DP can be transformed into an incentive compatible mechanism using dynamic bundle prices.

Finally, Section 8 presents a number of extensions. First, we discuss how our algorithm is robust to stronger adversaries. In particular we note that our prophet inequality holds against an almighty adversary as it is also the case in the standard prophet inequality. Second, we dive into the issue of incentive compatibility and the algorithm's access to the valuation of agent i upon her arrival. We show an incentive compatible implementation of our algorithm (and thus not needing to access the actual valuations) by using some form of all-pay auction. This connects our result with the work of Feldman et al. [15] who establish that parallel per-item first-price auctions guarantee a constant Price of Anarchy for offline combinatorial auctions. Finally, we note that the existence of a factor 2 prophet inequality for deterministic valuations using subadditive prices follows from existing results by Feldman, Gravin, and Lucier [16].

2 Model, algorithm and main result

Model. We are given a set of m items M and a set of n agents N. Each buyer $i \in N$ has a valuation function $v_i : 2^M \to \mathbb{R}_+$, which is randomly and independently chosen according to a given distribution $\mathcal{F}_i$ defined over a set of possible valuation functions V_i. As it is standard, we assume that each possible realization $v \in V_i$ is monotone (i.e., $A \subseteq B \Rightarrow v(A) \leq v(B)$), and that for all $A \subseteq M$ and all $i \in N$, $\mathbb{E}(v_i(A)) < \infty$.

The agents arrive sequentially one by one, in a fixed order $\sigma : [n] \to N$.[4] When an agent i arrives, we get to observe her valuation $v_i \sim \mathcal{F}_i$. After observing v_i, we must decide which items we will be allocating to i. We denote the set allocated to i by ALG_i. The expected welfare of the resulting allocation is then

$$\mathbb{E}(ALG) := \mathbb{E}\left(\sum_{i\in N} v_i(ALG_i)\right) = \sum_{i\in N} \mathbb{E}(v_i(ALG_i)).$$

The prophet sees the realizations of all valuation functions in advance and therefore assigns the items optimally. Denoting by OPT_i the random set assigned to agent i by the prophet, the expected welfare of this optimal allocation is:[5]

$$\mathbb{E}(OPT) := \mathbb{E}\left(\max_{\substack{X_1,\dots,X_n \\ \text{partition of } M}} \sum_{i\in N} v_i(X_i)\right) = \mathbb{E}\left(\sum_{i\in N} v_i(OPT_i)\right).$$

[4]For easy of notation and without loss of generality we often assume that the arrival order is just $1,\dots,n$.

[5]Note that the assumption that for all $A \subset M$ and all $i \in N$, $\mathbb{E}(v_i(A)) < \infty$, holds if and only if $\mathbb{E}(OPT) < \infty$. If this is not the case, and there is a set and an agent such that $\mathbb{E}(v_i(A)) = \infty$, then $\mathbb{E}(OPT) = \infty$, but then the trivial online algorithm that simply assigns A to i also recovers infinite expected welfare.

Note that we are not allowed to give an item to more than one agent, so for every $t \in [n]$, the assignment of the algorithm must satisfy

$$ALG_{\sigma(t)} \subseteq M \setminus \cup_{\tau<t} ALG_{\sigma(\tau)}.$$

For the agent $i = \sigma(t)$, we call the latter set the *remaining* items when i arrives and we denote it by R_i. More precisely, for $i \in N$ we denote

$$R_i = M \setminus \cup_{\tau<\sigma^{-1}(i)} ALG_{\sigma(\tau)}.$$

Random Score Generators (RSG). A random score generator for an agent $i \in N$ is a function $D_i : V_i \to \Delta(\mathbb{R}_+^M)$ that takes a valuation function $v \in V_i$ and outputs a *distribution* $D_i(v)$ over $\mathbb{R}_+^M$. A sample $b_i \sim D_i(v)$ from this distribution gives a number $b_{i,j}$ for each item $j \in M$ which we call *scores*. Intuitively, these scores have the role of providing a (random) per-item representation of the valuation v.

Algorithms. We consider a class of algorithms, where we draw a set imaginary agents and their scores —an independent sample of all valuations and corresponding scores—, and whenever a new (real) agent comes, we give her the available items for which her score is strictly larger than the score of each of the imaginary agents.

More precisely ALG works as follows. Given RSGs $(D_i)_{i\in N}$,

1. Sample independent valuations $v_i' \sim \mathcal{F}_i$ and scores $b_i' = (b_{i,j}')_{j\in M} \sim D_i(v_i')$ for each agent $i \in N$. Define for each item $j \in M$ the value $p_j' = \max_{i\in N} b_{i,j}'$.
2. When agent i arrives, observe v_i and draw $b_i \sim D_i(v_i)$. Define $S_i = \{j : b_{i,j} > p_j'\}$.
3. Give agent i the set $ALG_i = R_i \cap S_i$ and go back to the previous step with the next buyer.

Main Result. Our main result establishes the existence of a constant factor prophet inequality for online combinatorial auctions with subadditive valuations.

Theorem 1. *For every $\varepsilon > 0$, if all valuations are subadditive, there are RSGs such that*

$$(6 + \varepsilon) \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT).$$

3 The Mirror Lemma

For the valuations $v_i \sim \mathcal{F}_i$, sample $b_i \sim D_i(v_i)$ for each agent $i \in N$, independently. Additionally, sample two independent copies of these variables, $v_i', v_i'' \sim \mathcal{F}_i$ and $b_i' \sim D_i(v_i')$, $b_i'' \sim D_i(v_i'')$, for each agent $i \in N$. Define $p_j = \max_{i\in N} b_{i,j}$ and p_j', p_j'' analogously, for each item $j \in M$. Furthermore, define for each agent $i \in N$ the set $W_i = \left\{ j : b_{i,j} > \max\{p_j', p_j''\} \right\}$. We have the following result that holds for *any* RSGs $(D_i)_{i\in N}$.

Lemma 1 (Mirror Lemma)**.** *If the valuations are subadditive, then*

$$\mathbb{E}(ALG) \geq \frac{1}{2} \sum_{i \in N} \mathbb{E}(v_i(W_i)).$$

Proof. Recall that for an agent $i \in N$, R_i was defined as the set of remaining items when i arrives and $S_i = \{j : b_{i,j} > p'_j\}$. Therefore, we have that

$$\mathbb{E}(ALG) = \sum_{i \in N} \mathbb{E}(v_i(R_i \cap S_i)).$$

Define $S''_i = \{j : b''_{i,j} > p'_j\}$, and notice that R_i does not depend on v_i, nor on b_i. Therefore, since the distributions of the pairs (v_i, b_i) and (v''_i, b''_i) are identical,

$$\mathbb{E}(v_i(R_i \cap S_i)) = \mathbb{E}(v''_i(R_i \cap S''_i)).$$

Define now $R = \{j : p'_j \geq p_j\}$, which is the set of remaining items at the very end of the sequence of agents. Notice that $R \subseteq R_i$. Define also $W''_i = \left\{j : b''_{i,j} > \max\{p'_j, p_j\}\right\}$, and notice that $W''_i \subseteq S''_i$. With this, and the fact that the valuations are monotone, we obtain that

$$\begin{aligned} \mathbb{E}(v_i(R_i \cap S_i)) &= \mathbb{E}(v''_i(R_i \cap S''_i)) \\ &\geq \mathbb{E}\left(v''_i\left(R \cap W''_i\right)\right). \end{aligned}$$

Finally, notice that $(p_j)_{j \in M}$ and $(p'_j)_{j \in M}$ are independent and identically distributed, so we can interchange them inside the expectation. So if we define $\tilde{R} = \{j : p_j \geq p'_j\}$, we have that

$$\mathbb{E}\left(v''_i\left(R \cap W''_i\right)\right) = \mathbb{E}\left(v''_i\left(\tilde{R} \cap W''_i\right)\right).$$

Therefore, by the fact that $R \cup \tilde{R} = M$ and by the subadditivity of v_i, we conclude that

$$\begin{aligned} \mathbb{E}(v_i(R_i \cap S_i)) &\geq \frac{1}{2}\left(\mathbb{E}\left(v''_i\left(R \cap W''_i\right)\right) + \mathbb{E}\left(v''_i\left(\tilde{R} \cap W''_i\right)\right)\right) \\ &\geq \frac{1}{2}\mathbb{E}\left(v''_i(W''_i)\right). \end{aligned}$$

Summing over agents and switching the roles of $(v, b), (v', b')$ and (v'', b'') we obtain the statement of the lemma. ☐

4 Constant approximation for XOS valuations

We show here how to use the Mirror Lemma to obtain a constant approximation for XOS valuations. We construct an RSG as follows: given v_i for an agent $i \in N$, sample independent valuations $v_{i'}^{(i)} \sim \mathcal{F}_{i'}$ for each $i' \neq i$, calculate the optimal allocation under the valuations $(v_i, v_{-i}^{(i)})$, and let $OPT_i^{(i)}$ be the set of item that i receives. Sample independent Uniform$[0, 1]$ variables $U_{i,j}$ for each item $j \in M$, and let

$$b_{i,j} = \begin{cases} U_{i,j} & \text{if } j \in OPT_i^{(i)} \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

In the next lemma we prove that if we use this RSG, each item in $OPT_i^{(i)}$ is also in W_i with constant probability. Essentially, this is because any item j is in expectation in at most one of the sets $(OPT_i^{(i)})_{i \in N}$, so p'_j and p''_j cannot be much larger than uniform variables.

Lemma 2. *If we use the RSGs defined in Equation* (1)*, then, for all $i \in N$ and $A \subseteq M$,*

$$\mathbb{P}\left(j \in W_i \,\middle|\, OPT_i^{(i)} = A\right) \geq 1/12, \quad \forall j \in A.$$

Proof. Fix the set $A \subseteq M$ and an item $j \in A$. For each $i \in N$, denote by $OPT_i'^{(i)}$ and $OPT_i''^{(i)}$ the sets corresponding to (v', b') and (v'', b''), respectively, as in Equation (1). Now define the random variables

$$X_j = \left|\left\{i : j \in OPT_i'^{(i)}\right\}\right|, \text{ and } Y_j = \left|\left\{i : j \in OPT_i''^{(i)}\right\}\right|.$$

Since OPT allocates each item to only one agent, they satisfy $\mathbb{E}(X_j) = \mathbb{E}(Y_j) = 1$. By Markov's inequality, we have that

$$\mathbb{P}(X_j \leq 1) = \mathbb{P}(Y_j \leq 1) \geq 1/2.$$

Therefore, since $OPT_i^{(i)}$ is independent of X_j and Y_j, we have that

$$\begin{aligned}
&\mathbb{P}\left(j \in W_i \,\middle|\, OPT_i^{(i)} = A\right) \\
&\geq \mathbb{P}\left(j \in W_i \,\middle|\, OPT_i^{(i)} = A, X_j \leq 1, Y_j \leq 1\right) \cdot \mathbb{P}(X_j \leq 1, Y_j \leq 1) \\
&\geq \frac{1}{4} \cdot \mathbb{P}\left(j \in W_i \,\middle|\, OPT_i^{(i)} = A, X_j \leq 1, Y_j \leq 1\right).
\end{aligned}$$

Finally, recall that $j \in W_i$ if and only if $b_{i,j} > \max\{p'_j, p''_j\}$. Notice that, conditional on $OPT_i^{(i)} = A, X_j \leq 1$ and $Y_j \leq 1$, we have that $b_{i,j} \sim$Uniform$[0,1]$, and also that p'_j and p''_j are both dominated by independent Uniform$[0,1]$ random variables. Therefore, conditional on the mentioned event, $j \in W_i$ with probability at least $1/3$. □

Lemma 2 implies immediately a 24-approximation for XOS valuations. In fact, notice that conditional on $OPT_i^{(i)} = A$, the set W_i is independent of v_i. Thus,

$$\begin{aligned}
&\mathbb{E}(v_i(W_i)) \\
&= \sum_{A \subseteq M} \sum_{B \subseteq M} \mathbb{E}(v_i(B) | W_i = B, OPT_i^{(i)} = A) \cdot \mathbb{P}(W_i = B | OPT_i^{(i)} = A) \cdot \mathbb{P}(OPT_i^{(i)} = A) \\
&= \sum_{A \subseteq M} \sum_{B \subseteq M} \mathbb{E}(v_i(B) | OPT_i^{(i)} = A) \cdot \mathbb{P}(W_i = B | OPT_i^{(i)} = A) \cdot \mathbb{P}(OPT_i^{(i)} = A) \\
&= \sum_{A \subseteq M} \mathbb{E}\left(\sum_{B \subseteq M} v_i(B) \cdot \mathbb{P}(W_i = B | OPT_i^{(i)} = A) \,\middle|\, OPT_i^{(i)} = A\right) \cdot \mathbb{P}(OPT_i^{(i)} = A) \\
&\geq \sum_{A \subseteq M} \frac{1}{12} \cdot \mathbb{E}(v_i(A) | OPT_i^{(i)} = A) \cdot \mathbb{P}(OPT_i^{(i)} = A) \\
&= \frac{1}{12} \cdot \mathbb{E}(v_i(OPT_i)).
\end{aligned}$$

For the inequality we applied the property that defines XOS valuations, together with the fact that Lemma 2 guarantees that the collection $\{(B, 12 \cdot \mathbb{P}(W_i = B | OPT_i^{(i)} = A))\}_{B \subseteq M}$ is a fractional covering of A.

The constant in Lemma 2 can be tightened to be 1/3, yielding together with Lemma 1 a 6-approximation for XOS valuations. However, the per-item guarantee the lemma provides is insufficient to prove a constant approximation for general subadditive valuations. In the next section, we push this beyond the XOS case and use the Mirror Lemma with more sophisticated RSGs to obtain a 6-approximation for the general subadditive case.

5 Proof of the main result: A fixed point approach

In this section we prove our main result for general subadditive valuations. Throughout the section we assume that, for all $i \in N$, V_i is a finite set. In Section 6 we discuss how to relax this condition.

Theorem 1. *For every $\varepsilon > 0$, if all valuations are subadditive, there are RSGs such that*

$$(6 + \varepsilon) \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT).$$

The main ingredient for the proof of the theorem, besides the Mirror Lemma, is the following existence result, whose proof relies on a fixed-point argument. For ease of notation, for a vector $f \in \mathbb{R}^M$ and a set $S \subseteq M$, we write $f(S) = \sum_{j \in S} f_j$.

Lemma 3. *For every $\varepsilon > 0$ there are RSGs that guarantee that, almost surely,*

$$\mathbb{E}\Big(v_i\left(W_i\right) - b_i\left(M\right) \;\Big|\; v_i\Big) \geq \max_{X \subseteq M} \left\{ \frac{1}{3} v_i(X) - \mathbb{E}\left(p'(X)\right) - \varepsilon \cdot |X| \right\}.$$

Before proving this lemma, we show first how it implies Theorem 1.

Proof of Theorem 1. We take the distributions guaranteed to exist by Lemma 3, using some $\varepsilon' > 0$ that we will set later. Let OPT_i be the set of items that agent i gets in the optimal allocation, under valuations $(v_i)_{i \in N}$. By Lemma 1,

$$\begin{aligned}
\mathbb{E}(ALG) &\geq \frac{1}{2} \sum_{i \in N} \mathbb{E}(v_i(W_i)) \\
&= \frac{1}{2} \sum_{i \in N} \mathbb{E}(v_i(W_i) - b_i(M)) + \sum_{i \in N} \mathbb{E}(b_i(M)) \\
&= \frac{1}{2} \sum_{i \in N} \mathbb{E}\Big(\mathbb{E}\big(v_i(W_i) - b_i(M) \;\big|\; v_i\big)\Big) + \sum_{i \in N} \mathbb{E}(b_i(M)) \\
&\geq \frac{1}{2} \sum_{i \in N} \mathbb{E}\left(\frac{1}{3} v_i(OPT_i) - p'(OPT_i) - \varepsilon' \cdot |OPT_i|\right) + \sum_{i \in N} \mathbb{E}(b_i(M)),
\end{aligned}$$

where in the last line we used Lemma 3 taking $X = OPT_i$. Now, since OPT_i is a partition of M, we have that

$$\begin{aligned}\mathbb{E}(ALG) &\geq \frac{1}{6}\cdot\mathbb{E}(OPT) - \varepsilon'\cdot|M| - \mathbb{E}(p'(M)) + \mathbb{E}\left(\sum_{i\in N} b_i(M)\right)\\ &\geq \frac{1}{6}\cdot\mathbb{E}(OPT) - \varepsilon'\cdot|M|,\end{aligned}$$

where we used the fact that $\mathbb{E}(p'_j) = \mathbb{E}(p_j) = \mathbb{E}(\max_{i\in N} b_{i,j}) \leq \mathbb{E}(\sum_{i\in N} b_{i,j})$. Finally, we can take ε' to be as small as we want. In particular, we can set it to be $\varepsilon' \leq \frac{\varepsilon\cdot\mathbb{E}(OPT)}{6\cdot(6+\varepsilon)\cdot|M|}$, which concludes the proof of the theorem. □

As mentioned earlier, to prove Lemma 3 we use a fixed-point approach. We will define a mapping whose fixed points are exactly the RSGs that satisfy the condition of the lemma. Before diving into that, let us define a discretized space for the scores and prove a simpler existence result. Let $v_{\max} = \max_{v\in V_i, i\in N} v(M)$. For a given $\varepsilon > 0$, let $B_\varepsilon = \{s\cdot\varepsilon : s\in\mathbb{N} \text{ and } s\cdot\varepsilon \leq v_{\max}\}^M$.

Lemma 4. *For every subadditive and monotone valuation function $v : 2^M \to \mathbb{R}_+$ with $v(M) \leq v_{\max}$, and every $\varepsilon > 0$; if p' and p'' are random i.i.d. vectors in $\mathbb{R}_+^M$, there exists a vector $f \in B_\varepsilon$ such that*

$$\mathbb{E}\Big(v\left(\{j : f_j > \max\{p'_j, p''_j\}\}\right)\Big) - f(M) \geq \frac{1}{3}v(X) - \mathbb{E}\left(p'(X)\right) - \varepsilon\cdot|X|, \qquad \forall X\subseteq M. \quad (2)$$

Proof. Note that only the right-hand side of Equation (2) depends on X, so it is enough to find an $f\in B_\varepsilon$ that satisfies Equation (2) for the set X^* that maximizes its right-hand side. Let p''' be an i.i.d. copy of p'. We define first a random vector $\hat{f}$ as follows:

$$\hat{f}_j = \begin{cases}\lfloor p'''_j/\varepsilon\rfloor\cdot\varepsilon + \varepsilon & \text{if } j\in X^*\\ 0 & \text{otherwise.}\end{cases}$$

With this definition, $\hat{f}_j > p'''_j$ for all $j\in X^*$, so by the monotonicity of v,

$$\mathbb{E}\Big(v\left(\left\{j : \hat{f}_j > \max\{p'_j, p''_j\}\right\}\right)\Big) \geq \mathbb{E}\Big(v\left(\{j : p'''_j \geq \max\{p'_j, p''_j\}\}\cap X^*\right)\Big).$$

Since p', p'', p''' are i.i.d., we can interchange them inside the expectation, i.e.,

$$\begin{aligned}\mathbb{E}\Big(v\left(\{j : p'''_j \geq \max\{p'_j, p''_j\}\}\cap X^*\right)\Big) &= \mathbb{E}\Big(v\left(\{j : p''_j \geq \max\{p'_j, p'''_j\}\}\cap X^*\right)\Big)\\ &= \mathbb{E}\Big(v\left(\{j : p'_j \geq \max\{p'''_j, p''_j\}\}\cap X^*\right)\Big).\end{aligned}$$

Also, note that

$$M = \Big\{j : p'''_j \geq \max\{p'_j, p''_j\} \text{ or } p''_j \geq \max\{p'_j, p'''_j\} \text{ or } p'_j \geq \max\{p'''_j, p''_j\}\Big\},$$

so by the subadditivity of v, we obtain that

$$\mathbb{E}\Big(v\left(\left\{j : \hat{f}_j > \max\{p'_j, p''_j\}\right\}\right)\Big) \geq \frac{1}{3}v(X^*).$$

Because $\hat{f}_j \le p'''_j + \varepsilon$ for all $j \in X^*$, we conclude that

$$\mathbb{E}\Big(v\left(\left\{j : \hat{f}_j > \max\{p'_j, p''_j\}\right\}\right) - \hat{f}(M)\Big) \ge \frac{1}{3}v(X^*) - \mathbb{E}\left(p'(X^*)\right) - \varepsilon \cdot |X^*|. \tag{3}$$

To finish the proof of the lemma we have to make sure $\hat{f} \in B_\varepsilon$, so we need to control the event that $\hat{f}_j > v_{\max}$ for some $j \in M$. But if this happens, notice that $v(M) - \hat{f}(M) < 0$. Thus, if we replace $\hat{f}$ with the vector 0 in this event, we only increase the expectation on the left-hand side of Equation (3). Finally, since a random vector in B_ε satisfies the inequality in expectation, there must exist a deterministic vector in B_ε that also satisfies it. □

Definition of the Mapping. Denote by $\mathcal{L} = \bigtimes_{i \in N} \Delta(B_\varepsilon)^{V_i}$ the space of the RSGs with scores in B_ε. We define the set-valued function $\psi : \mathcal{L} \to 2^{\mathcal{L}}$ in the following way. Given an RSG $D = (D_i)_{i \in N} \in \mathcal{L}$, we sample for each $i \in N$, $v'_i, v''_i \sim \mathcal{F}_i$ independently, $b'_i \sim D_i(v'_i)$, and $b''_i \sim D_i(v''_i)$, also independently, and we denote $p'_j = \max_{i \in N} b'_{i,j}$, and $p''_j = \max_{i \in N} b''_{i,j}$. We define $\psi(D) \subseteq \mathcal{L}$ as the set of RSGs $G = (G_i)_{i \in N} \in \mathcal{L}$ that satisfy that: for all $i \in N$, and for all $v \in V_i$, if $f \sim G_i(v)$ and is independent of p' and p'', then

$$\mathbb{E}\Big(v\left(\{j : f_j > \max\{p'_j, p''_j\}\}\right) - f(M)\Big) \ge \max_{X \subseteq M} \left\{\frac{1}{3}v(X) - \mathbb{E}\left(p'(X)\right) - \varepsilon \cdot |X|\right\}. \tag{4}$$

Lemma 5. *$\mathcal{L}$ is a non-empty, convex, and compact subset of a euclidean space; for every $D \in \mathcal{L}$, the set $\psi(D)$ is non-empty, closed, and convex; and the function ψ has a closed graph.*

Proof. First, we establish that $\mathcal{L}$ is non-empty, convex, and compact. To this end note that the set B_ε is finite and has size $|B_\varepsilon| = \lfloor v_{\max}/\varepsilon \rfloor^{|M|}$. Therefore, we can represent $\mathcal{L}$ as a subset of $[0,1]^\ell$, with $\ell = \lfloor v_{\max}/\varepsilon \rfloor^{|M|} \cdot \sum_{i \in N} |V_i|$ as

$$\mathcal{L} = \left\{x = (x_{b,i,v})_{b \in B_\varepsilon, i \in N, v \in V_i} \in [0,1]^\ell : \sum_{b \in B_\varepsilon} x_{b,i,v} = 1,\ \forall i \in N, v \in V_i\right\},$$

where $x_{b,i,v}$ is the probability that the scores of an agent $i \in N$ are $b \in B_\varepsilon$ when her valuation is $v \in V_i$. It is immediate from this that $\mathcal{L}$ is a non-empty, convex and compact subset of $\mathbb{R}^\ell$. Throughout the rest of this proof we will use this representation of $\mathcal{L}$.

Second, note that $\psi(D)$ is non-empty. In fact, this is immediate from Lemma 4.

Third, we show that $\psi(D)$ is closed and convex. For $i \in N, v \in V_i$, let $q_{i,v} \in [0,1]$ denote the probability that $v_i = v$. For a given $x \in \mathcal{L}$, we denote by $\pi(x) \in \Delta(B_\varepsilon)$ the resulting distribution of p' (and of p''), i.e., for $f \in B_\varepsilon$,

$$\pi_f(x) = \sum_{(v_i)_{i \in N} \in \bigtimes_{i \in N} V_i} \left(\prod_{i \in N} q_{i,v_i}\right) \sum_{(b_i)_{i \in N} \in B_\varepsilon^N : \max_{i \in N} b_{i,j} = f_j} \left(\prod_{i \in N} x_{b_i,i,v_i}\right).$$

Now, using this notation, for a pair $x, y \in \mathcal{L}$, we have that $y \in \psi(x)$ if and only if, for all $i \in N, v \in V_i, X \subseteq M$,

$$\sum_{f \in B_\varepsilon} \sum_{p' \in B_\varepsilon, p'' \in B_\varepsilon} y_{f,i,v} \cdot \pi_{p'}(x) \cdot \pi_{p''}(x) \cdot \Big(v\left(\{j : f_j > \max\{p'_j, p''_j\}\}\right) - f(M) \Big)$$
$$\geq \frac{1}{3} v(X) - \sum_{p' \in B_\varepsilon} \pi_{p'}(x) \cdot p'(X) - \varepsilon \cdot |X|. \tag{5}$$

For fixed x, these are finitely many linear constraints on y, so $\psi(x)$ is a convex and compact set.

Finally, we prove ψ has a closed graph, i.e, that the set $\{(x, y) : x \in \mathcal{L},\ y \in \psi(x)\}$ is closed. This follows immediately from the continuity of Equation (5) in x and y. Indeed, let $(x^{(k)}, y^{(k)})_{k \in \mathbb{N}}$ be a convergent sequence in the graph of ψ that converges to $(x^{(\infty)}, y^{(\infty)})$. Since the sequence is in the graph of ψ, we have that $y^{(k)} \in \psi(x^{(k)})$ for all $k \in \mathbb{N}$. Also, since $\pi_f(x)$ is a continuous function for every $f \in B_\varepsilon$, both the left-hand side and the right-hand side of Equation (5) are continuous functions of x and y. Thus, the pair $(x^{(\infty)}, y^{(\infty)})$ satisfies Equation (5) for all $i \in N, v \in V_i, X \subseteq M$, and then $y^{(\infty)} \in \psi(x^{(\infty)})$. This means that $(x^{(\infty)}, y^{(\infty)})$ is in the graph of ψ, and therefore, ψ has a closed graph. □

Proof of Lemma 3. Recall that Kakutani's fixed-point theorem states that if $\mathcal{L}$ is a non-empty, compact, and convex subset of a euclidean space, and $\psi : \mathcal{L} \to 2^{\mathcal{L}}$ has non-empty and convex values, and a closed graph, then ψ has a fixed point. Lemma 5 gives exactly these conditions, and therefore, there exists an RSG, D, such that $D \in \psi(D)$, which is exactly what we want. □

6 General valuations

So far we assumed that the supports of the distributions of the valuation functions are finite. By losing an additional ε we can remove this assumption and actually prove our main result for general valuations so long as their expectation is finite.

First, if the supports are not finite but the valuations are uniformly bounded by a constant $v_{\max}$, we can discretize them and obtain finite supports. In fact, we only need the assumption of finite support for the existence result in Lemma 3, so let's apply it for the discretized valuations defined for each $i \in N, X \subseteq M$, as

$$\tilde{v}_i(X) = \lfloor v_i(X)/\varepsilon \rfloor \cdot \varepsilon.$$

Applying Lemma 3 we obtain RSGs such that, almost surely,

$$\begin{aligned}
\mathbb{E}\Big(v_i(W_i) - b_i(M) \;\Big|\; v_i \Big) &\geq \mathbb{E}\Big(\tilde{v}_i(W_i) - b_i(M) \;\Big|\; v_i \Big) \\
&= \mathbb{E}\Big(\tilde{v}_i(W_i) - b_i(M) \;\Big|\; \tilde{v}_i \Big) \\
&\geq \max_{X \subseteq M} \left\{ \frac{1}{3} \tilde{v}_i(X) - \mathbb{E}\left(p'(X)\right) - \varepsilon \cdot |X| \right\} \\
&\geq \max_{X \subseteq M} \left\{ \frac{1}{3} v_i(X) - \mathbb{E}\left(p'(X)\right) - 2\varepsilon \cdot |X| \right\}.
\end{aligned}$$

We thus obtain an algorithm such that $(6 + O(\varepsilon)) \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT)$.

Second, we can assume valuations are uniformly bounded by a sufficiently large constant $v_{\max}$ by truncating them. In fact, since $\mathbb{E}(OPT) < \infty$, we know that

$$\lim_{t\to\infty} \mathbb{E}\big(OPT \cdot \mathbb{1}_{\{OPT>t\}}\big) = 0.$$

Thus, there exists a constant $v_{\max}$ such that

$$\mathbb{E}\big(OPT \cdot \mathbb{1}_{\{OPT\leq v_{\max}\}}\big) \geq (1-\varepsilon) \cdot \mathbb{E}(OPT).$$

We already know that there must exist RSGs that guarantee a $6 + O(\varepsilon)$-approximation for the truncated valuations $\bar{v}_i \sim v_i | v_i(M) \leq v_{\max}$ for each $i \in N$. Then, if we apply this algorithm to the original valuations, in the case $v_i(M) > v_{\max}$ for some $i \in N$, necessarily $OPT > v_{\max}$, so the expected welfare must be at least

$$\frac{1}{6+O(\varepsilon)} \mathbb{E}\big(OPT \cdot \mathbb{1}_{\{OPT\leq v_{\max}\}}\big) \geq \frac{1}{6+O(\varepsilon)} \mathbb{E}(OPT).$$

7 The Optimal Dynamic Program and its DSIC Implementation

If the arrival order of the buyers is fixed and given, the optimal allocation in each step can be obtained by dynamic programming. Assume the agents arrive in the order $1, \ldots, n$ and denote by $V_i(X)$ the expected welfare that the algorithm recovers from agents i or larger, when the leftover set of items is $X \subset M$. By the principle of dynamic programming we have that

$$\begin{aligned}
V_n(X) &= \mathbb{E}(v_n(X)) && \text{for all } X \subset M,\\
V_i(X) &= \mathbb{E}\left(\max_{Y\subseteq X} v_i(Y) + V_{i+1}(X \setminus Y)\right) && \text{for all } X \subset M.
\end{aligned}$$

$V_1(M)$ is then the maximum expected social welfare attainable if decisions are made online. Therefore, for any online algorithm ALG we have that $V_1(M) \geq \mathbb{E}(ALG)$. Applying Theorem 1 and taking $\varepsilon \to 0$ we obtain the following result.

Corollary 1. *If the arrival order of the agents is fixed and given, there is an online algorithm ALG such that*

$$6 \cdot \mathbb{E}(ALG) \geq \mathbb{E}(OPT).$$

Implementation Using Dynamic Bundle Pricing. In our result we assume we can observe not only the distributions of the valuations, but also the valuations themselves (when the corresponding agent arrives), and therefore our prophet inequality establishes a bound for the performance of an online algorithm with respect to the optimal allocation. However, our result also implies the existence of a prophet inequality where the online algorithm does not have access to the realizations of the valuations and, moreover, by a dominant strategies incentive compatible (DSIC) mechanism. Indeed, we can implement the optimal online solution given by the dynamic program by offering each buyer personalized prices for each bundle of items.

Consider the prices $p_{i,R}(X)$, which is the price of set X, offered to agent i, if the remaining set upon her arrival is R.

$$p_{i,R}(X) = V_{i+1}(R) - V_{i+1}(R \setminus X) \quad \text{for all } X \subset R \subset M.$$

These prices represent the externality caused by agent i if she buys set X. Defining $V_{n+1}(R) = 0$ for all $R \subset M$, the prices are nonnegative (since $V_i(\cdot)$ is $p_{i,R}(\phi) = 0$.

Proposition 1. *If the agents sequentially select their preferred set when faced to prices $p_{i,R}(X)$, the expected welfare of the resulting allocation equals $V_1(M)$.*

Proof. By backwards induction assume that if the set of remaining items is X and $i+1$ arrives, the expected welfare of the pricing mechanism from agents $\{i+1, \ldots, n\}$ is $V_{i+1}(X)$.

Then, when agent i arrives and set R are the remaining items she solves:

$$\begin{aligned} \max_{X \subset R} \{v_i(X) - p_{i,R}(X)\} &= \max_{X \subset R} \{v_i(X) - V_{i+1}(R) + V_{i+1}(R \setminus X)\} \\ &= - V_{i+1}(R) + \max_{X \subset R} \{v_i(X) + V_{i+1}(R \setminus X)\} . \end{aligned}$$

Thus, i buys the set $X^* = \arg\max_{X \subset R} v_i(X) + V_{i+1}(R \setminus X)$. This is the same set that the optimal online algorithm allocates to agent i, given the set of remaining items R. Therefore, what the pricing mechanism obtains from agents $\{i, \ldots, n\}$ is $\mathbb{E}(v_i(X^*) + V_{i+1}(R \setminus X^*))$ which equals $V_i(R)$. □

8 Concluding remarks

The Almighty Adversary. Like in the classical Prophet Inequality setting, throughout the paper we assumed the agents come in an adversarial fixed order. Recent works have investigated which (or to what extent) approximation results can be extended to more powerful adversaries. The most powerful adversary, the *almighty adversary* can choose the arrival order adaptively, observing all realizations of the valuations and the random choices of the algorithm in advance. Our result can be extended to work against the almighty adversary. This requires to slightly adapt the Mirror Lemma. In fact, notice that if the order can depend on the realizations of the valuations, then the set R_i is no longer independent of v_i, so we cannot just replace v_i with v_i'' in the proof without modifying R_i. However, if we define

$$R^{(-i)} = \left\{ j : p'_j \geq b_{i',j}, \ \forall i' \neq i \right\},$$

we get that $R \subseteq R^{(-i)} \subseteq R_i$. And since $R^{(-i)}$ is independent of v_i, we can add an intermediate step where we replace R_i with $R^{(-i)}$, then we can replace v_i with v_i'' and then replace $R^{(-i)}$ with R and continue with the next steps of the proof. Note that $\mathbb{E}(v_i(W_i))$ does not depend on the arrival order, so once the Mirror Lemma is established, the rest of the proof of Theorem 1 carries over.

Computational Complexity. Our result is purely existential: we prove there exists an algorithm that achieves the constant factor Prophet Inequality. An obvious question is whether such policy can be computed. Since we must find a probability distribution, it is not even immediately obvious whether we can compute it in finite time when the supports of the valuations are finite. We can answer this affirmatively by further discretizing the space of RSGs $\mathcal{L}$. Let $\hat{\mathcal{L}} \subseteq \mathcal{L}$ be the set of RSGs whose probabilities are multiples of a very small $\varepsilon' > 0$. Since there exist RSGs in $\mathcal{L}$ that satisfy the condition of Lemma 3, and both sides of the inequality are continuous in the RSG, there must exist one in $\hat{\mathcal{L}}$ that satisfies the constraint approximately. We can make ε' sufficiently small so we only lose an extra ε in the approximation factor. Searching over all possible options gives an algorithm that terminates in finite time. An open question is whether we can compute efficiently this policy, or any policy that yields a constant approximation factor.

Incentive Compatible Implementation and PoA for Simultaneous Auctions. Our result and the discussion in Section 7 establish the existence of a factor 6 prophet inequality which is implementable using a DSIC mechanism. However we do this by analyzing the optimal online algorithm and we do not have a reduction that transforms an online algorithm into a DSIC mechanism. Thus, even if we could efficiently find the RSGs that attain the $1/(6+\varepsilon)$-approximation, we do not immediately obtain from that an incentive compatible mechanism.

We explain here how to modify our algorithm to obtain an incentive compatible mechanism. The analysis of this mechanism is related to the result of Feldman et al. [15], where we replace the first-price auction with a modified all-pay auction, that we call *undisputed-winner all-pay auctions with hidden reserve prices.* This auction has a Price of Anarchy of $(6+\varepsilon)$, and thanks to the Mirror Lemma, can be implemented online.

An undisputed-winner all-pay auction with hidden reserve prices is a simultaneous auction where we set a random reserve price p'_j for each item $j \in M$. Then, each agent $i \in N$ chooses a bid $b_{i,j}$ for each item $j \in M$ (without observing the realized prices). An agent $i \in N$ gets each item $j \in M$ for which she is the undisputed winner, i.e., $b_{i,j} > p'_j \geq b_{i',j}$ for all $i' \neq i$. All agents pay all their bids, i.e., each agent $i \in N$ pays $b_i(M) = \sum_{j \in M} b_{i,j}$. Colloquially, we *burn* every item for which no agent bids more than the reserve price, and also every item for which two or more agents bid more than the reserve price. We can modify the Mirror Lemma and Lemma 3 to show that if (i) the supports of the valuations are finite, and (ii) we restrict the space of possible bids to be multiples of a given $\varepsilon > 0$, then there are (random) reserve prices for which there is a Bayesian Nash Equilibrium, and the expected welfare is at least $1/(6 + O(\varepsilon))$ times the expected optimal welfare. The purpose of defining this rather esoteric class of auctions is that we can emulate them on the fly in our Online Combinatorial Auction setting. When an agent $i \in N$ arrives we can —without disclosing what is the set of remaining items— ask her a vector of bids, simulate the bids of the agents that have not arrived yet, and give i the items for which she is the undisputed winner. Notice that because of the undisputed-winner condition, we never attempt to allocate an item to more than one agent. Therefore, from an agent's point of view, this procedure is equivalent to the simultaneous auction, so the distributions of the set of items she gets and her bids are the same as in the simultaneous auction. Thus, by linearity of expectation, the expected welfare is the same as in the simultaneous auction. We defer the details of this construction to the appendix.

Posted Prices. A natural class of mechanisms is that of per-item posted prices, where we simply set a price p_j for each item $j \in M$, and then allow each agent to buy her most preferred subset of the remaining items. The 2-approximation for XOS valuations [17] and the $O(\log \log |M|)$-approximation for subadditive valuations [11] from previous work are in this class. Even though the scores in our algorithm have a certain flavor of prices, we do not know whether it can be transformed into a posted-prices mechanism. When designing a posted-prices mechanism for the Online Combinatorial Auction setting there are two main challenges. The first is the coordination of the agents: since each agent just takes her most preferred set, we might sell too early items that we should reserve for an agent with a high valuation later in the sequence. The second is the uncertainty in the valuations: we do not know which agents will have a high valuation, and which subsets of items will be the more valuable ones. The first challenge appears even if we try to use a posted-prices mechanism for deterministic valuations. However, a result by Feldman et al. [16, Theorem 3.2] nicely resolves this issue by using posted prices per bundles. For completeness,[6] we present here a slight adaptation of their result attaining a 2-approximation of the optimal welfare for deterministic subadditive valuations. For each subset $X \subseteq M$, post the price

$$p(X) = \frac{1}{2}\sum_{i \in N} v_i(OPT_i \cap X).$$

Let S be the set of sold items. By the subadditivity we have that the revenue is at least

$$\begin{aligned}
\text{Revenue} &= \sum_{i \in N} p(ALG_i) \\
&= \frac{1}{2}\sum_{i \in N}\sum_{i' \in N} v_{i'}(ALG_i \cap OPT_{i'}) \\
&\geq \frac{1}{2}\sum_{i' \in N} v_{i'}(S \cap OPT_{i'}).
\end{aligned}$$

Now, let $R = M \setminus S$ be the set of remaining items. All items in R are available when each of the agents arrives, so we have that the sum of utilities of the agents is at least

$$\begin{aligned}
\text{Utility} &\geq \sum_{i \in N} \max_{X \subseteq R} \{v_i(X) - p(X)\} \\
&\geq \sum_{i \in N} \Big(v_i(OPT_i \cap R) - p(OPT_i \cap R)\Big) \\
&= \sum_{i \in N} \left(v_i(OPT_i \cap R) - \frac{1}{2}\sum_{i' \in N} v_{i'}(OPT_i \cap OPT_{i'} \cap R)\right) \\
&= \sum_{i \in N} \left(v_i(OPT_i \cap R) - \frac{1}{2} v_i(OPT_i \cap R)\right) \\
&= \frac{1}{2}\sum_{i \in N} v_i(OPT_i \cap R).
\end{aligned}$$

The fourth line follows from the fact that $(OPT_i)_{i \in N}$ is a partition of M. Adding the revenue and the utility, since the valuations are subadditive and $S \cup R = M$, we get that the welfare is at least $\frac{1}{2}\sum_{i \in N} v_i(OPT_i)$.

[6] What we present is actually a special case of the result by Feldman et al. [16], which works for fully general monotone valuations keeping the same 2-approximation.

We believe this evidence, together with our main result, suggests there should exist a posted-prices mechanism for our setting.

Acknowledgements. We are grateful to the Simons Institute for the Theory of Computing and the organizers of the Data-Driven Decision Processes program, where large parts of this work were conceived. We are grateful to Paul Dütting for introducing the problem to us and for stimulating discussions. We also thank Bobby Kleinberg for pointing us to [16, Theorem 3.2] and Thomas Kesselheim for pointing us to [15]. This work was partially supported by the Center for Mathematical Modeling at the University of Chile (ANID FB210005), Grant Anillo Information and Computation in Market Design (ANID ACT210005) and by a Meta Research PhD Fellowship.

References

[1] S. Assadi, T. Kesselheim, S. Singla. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. SODA 2021.

[2] S. Assadi and S. Singla. Improved truthful mechanisms for combinatorial auctions with submodular bidders. FOCS 2019.

[3] P. Azar, R. Kleinberg, S.M. Weinberg. Prophet inequalities with limited information. SODA 2014.

[4] M. Babaioff, B. Lucier, N. Nisan, R. Paes Leme. On the efficiency of the walrasian mechanism. EC 2014.

[5] E. Baldwin, P. Klemperer. Understanding preferences:"demand types", and the existence of equilibrium with indivisibilities. *Econometrica* 87(3):867–932, 2019.

[6] Y. Cai, M. Zhao. Simple mechanisms for subadditive buyers via duality. STOC 2017.

[7] J. Correa, A. Cristi, A. Fielbaum, T. Pollner, S.M. Weinberg. Optimal item pricing in online combinatorial auctions. IPCO 2022.

[8] S. Dobzinski, M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. SODA 2006.

[9] S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. STOC 2005.

[10] S. Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. STOC 2016.

[11] P. Dütting, T. Kesselheim, B. Lucier. An O(log log m) Prophet Inequality for Subadditive Combinatorial Auctions. *SIAM J. on Computing*, to appear (Preliminary version in FOCS 2020).

[12] P. Dutting, M. Feldman, T. Kesselheim, B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. on Computing*, 49(3):540–582, 2020.

[13] S. Ehsani, M. Hajiaghayi, T. Kesselheim, S. Singla Prophet secretary for combinatorial auctions and matroids. SODA 2018.

[14] T. Ezra, M. Feldman, N. Gravin, Z. Tang. Online stochastic max-weight matching: prophet inequality for vertex and edge arrival models. EC 2020.

[15] M. Feldman, H. Fu, N. Gravin, B. Lucier. Simultaneous Auctions are (almost) Efficient. STOC 2013.

[16] M. Feldman, N. Gravin, B. Lucier. Combinatorial walrasian equilibrium. *SIAM J. on Computing*, 45(1):29–48, 2016.

[17] M. Feldman, N. Gravin, B. Lucier. Combinatorial Auctions via Posted Prices. SODA 2014.

[18] U. Feige. On maximizing welfare when utility functions are subadditive. STOC 2006.

[19] N. Gravin, H. Wang. Prophet inequality for bipartite matching: merits of being simple and non adaptive. EC 2019.

[20] R. Kleinberg, S.M. Weinberg. Matroid prophet inequalities. STOC 2012.

[21] U. Krengel, L. Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83:745–747, 1977.

[22] U. Krengel, L. Sucheston. On semiamarts, amarts, and processes with finite value. *Adv. in Probability* 4:197–266, 1978.

[23] R. Paes Leme, S. Wong. Computing Walrasian equilibria: Fast algorithms and structural properties. *Mathematical Programming* 179(1):343–384 2020

[24] A. Rubinstein, S. Singla. Combinatorial prophet inequalities. SODA 2017.

[25] A. Rubinstein, J.Z. Wang, S.M. Weinberg. Optimal Single-choice prophet inequalities from samples. ITCS 2020.

[26] H. Zhang. Improved prophet inequalities for combinatorial welfare maximization with (approximately) subadditive agents. *Journal of Computer and System Sciences* 123:143–146, 2022 (Preliminary version in ESA 2020).

A Incentive compatible implementation

We describe here how to implement our result if the valuations are private information and agents behave selfishly. For that we first introduce a certain class of simultaneous auctions.

Undisputed-winner all-pay auctions with hidden reserve prices. In this class of auctions we set a reserve price p'_j for each item $j \in M$. These prices are possibly random, in which case only the distribution is known to the agents. For a prespecified set $\mathcal{B} \subseteq \mathbb{R}^M_+$, each agent $i \in N$ submits item-bids $b_i = (b_{i,j})_{j \in M} \in \mathcal{B}$. Each agent $i \in N$ gets those items $j \in M$ such that $b_{i,j} > p'_j \geq b_{i',j}$ for all $i' \neq i$. For such an item, we say i is the undisputed winner of item j. Each agent $i \in N$ pays the sum of all her bids, i.e., agent i pays $b_i(M) = \sum_{j \in M} b_{i,j}$ (regardless of which items she gets).

For each agent $i \in N$ and a vector $f_i \in \mathcal{B}$, let $A_i(f_i)$ denote the set of items she gets if she bids f_i, given the other agents' bids, i.e., $A_i(f_i) = \{j \in M : f_{i,j} > p'_j \geq b_{i',j}, \forall i' \neq i\}$. We assume i chooses $f_i \in \mathcal{B}$ so as to maximize her expected utility $E(v_i(A_i(f_i)) - f_i(M)|v_i)$. For a given $\varepsilon > 0$, we take $\mathcal{B} = B_\varepsilon = \{s \cdot \varepsilon : s \in \mathbb{N} \text{ and } s \cdot \varepsilon \leq v_{\max}\}^M$. As in Section 5 we define $\mathcal{L}$ as the space of RSGs with scores in B_ε.

Theorem 2. *There are random reserve prices* p' *and a Bayesian Nash Equilibrium such that the welfare of the resulting allocation satisfies*

$$(6 + \varepsilon) \cdot \sum_{i \in N} \mathbb{E}(A_i(b_i)) \geq \mathbb{E}(OPT).$$

Online Implementation. Given (a distribution of) prices p' and a corresponding Bayesian Nash Equilibrium, we can implement the allocation on the fly. Notice that the distributions of the bids in the equilibrium *are* RSGs. Denote them by $(D_i)_{i \in N}$. For ease of notation let's identify the set of agents N with $[n]$ and assume they arrive in the corresponding order. We proceed as follows: when agent i arrives, without disclosing what are the prices and the remaining items, we ask her a vector of bids $b_i \in \mathcal{B}$. For each agent $i' > i$ (the ones that have not arrived yet), we simulate (draw a fresh sample of) her valuation $v^{(i)}_{i'} \sim \mathcal{F}_{i'}$ and the corresponding equilibrium bids $b^{(i)}_{i'} \sim D_{i'}(v^{(i)}_{i'})$. Note that at this point we already know the true bids $b_{i'}$ of previous agents $i' < i$. We give agent i the set

$$ALG_i = \left\{ j \in M : b_{i,j} > p'_j \geq \max\left\{ \max_{i' < i} b_{i',j}, \max_{i' > i} b^{(i)}_{i',j} \right\} \right\}.$$

Notice that this is feasible because, by definition, the sets ALG_i are disjoint. We can prove inductively that for every agent $i \in N$, submitting bids $b_i \sim D_i(v_i)$ is an equilibrium, i.e., if all the other agents follow these strategies, these bids maximize each agent's utility. In fact, take an agent $i \in N$ and assume all previous agents $i' < i$ submit bids $b_{i'} \sim D_{i'}(v_{i'})$. From the perspective of agent i the distribution of the set she will get as a function of her bid is exactly the same as in the equilibrium of the simultaneous auction, and therefore, it is optimal to play the same strategy as in the simultaneous auction, i.e., $b_i \sim D_i(v_i)$. Therefore, the equilibrium bids for the simultaneous auction are also an equilibrium for the online implementation. Finally, notice that this implies that the distribution of the utility and the payment of each agent $i \in N$ is the same as in the equilibrium of the simultaneous auction, and therefore, by linearity of expectation, the resulting welfare is the same as in the equilibrium of the simultaneous auction.[7]

We now prove there are prices p' and a corresponding equilibrium for the simultaneous auction whose expected welfare is at least $1/(6+\varepsilon)$ of the expected optimal welfare. We start by proving there are random prices that distribute exactly as the maximum bids of the resulting equilibrium. This is an application of Kakutani's Fixed-Point Theorem, similar to Lemma 3.

Lemma 6. *There are RSGs $(D_i)_{i\in N}$ in $\mathcal{L}$ such that if we sample for each $i \in N$ independently $v'_i \sim \mathcal{F}_i$ and $b'_i \sim D_i(v'_i)$, and take as reserve prices $p'_j = \max_{i\in N} b'_{i,j}$, then there is a Bayesian Nash Equilibrium of the simultaneous auction where each agent $i \in N$ submits bids $b_i \sim D_i(v_i)$.*

Proof. We define a similar mapping as for Lemma 3. We take the same space $\mathcal{L} = \bigtimes_{i\in N} \Delta(B_\varepsilon)^{V_i}$ which is the space of the RSGs with scores in B_ε. We define the set-valued function $\hat{\psi} : \mathcal{L} \to 2^{\mathcal{L}}$ in the following way. Given an RSG $D = (D_i)_{i\in N} \in \mathcal{L}$, we sample for each $i \in N$, $v'_i, v''_i \sim \mathcal{F}_i$ independently, $b'_i \sim D_i(v'_i)$, and $b''_i \sim D_i(v''_i)$, also independently, and we denote $p'_j = \max_{i\in N} b'_{i,j}$, and $p''_j = \max_{i\in N} b''_{i,j}$. We define $\psi(D) \subseteq \mathcal{L}$ as the set of RSGs $G = (G_i)_{i\in N} \in \mathcal{L}$ that satisfy that: for all $i \in N$, and for all $v \in V_i$, a random vector $f \sim G_i(v)$ satisfies

$$\mathbb{E}\Big(v\left(\{j : f_j > \max\{p'_j, p''_j\}\}\right) - f\left(M\right)\Big) \geq \max_{g\in B_\varepsilon} \mathbb{E}\Big(v\left(\{j : g_j > \max\{p'_j, p''_j\}\}\right) - g\left(M\right)\Big). \qquad (6)$$

We show this mapping has a fixed point (which is are by definition the RSGs we are looking for) using Kakutani's Theorem. To use the theorem, we require that $\mathcal{L}$ is a non-empty, convex and compact subset of an euclidean space, that $\hat{\psi}(D)$ is non-empty, convex and closed, and that $\hat{\psi}$ is upper-hemicontinuous.

First, we prove $\mathcal{L}$ is a non-empty, convex and compact subset of an euclidean space. As in the proof of Lemma 5, this is immediate of considering the representation of $\mathcal{L}$ as

$$\mathcal{L} = \left\{ x = (x_{b,i,v})_{b\in B_\varepsilon, i\in N, v\in V_i} \in [0,1]^\ell : \sum_{b\in B_\varepsilon} x_{b,i,v} = 1,\ \forall i \in N, v \in V_i \right\},$$

where $x_{b,i,v}$ is the probability that the scores of an agent $i \in N$ are $b \in B_\varepsilon$ when her valuation is $v \in V_i$.

Second, $\hat{\psi}(D)$ by definition non-empty: we can just take f to be deterministically equal to the vector $g \in B_\varepsilon$ that attains the maximum of the right-hand side of Equation (6) (recall that B_ε is a finite set).

Third, we show that $\hat{\psi}(D)$ is closed and convex. Just as in Lemma 5, for $i \in N, v \in V_i$, let $q_{i,v} \in [0,1]$ denote the probability that $v_i = v$. For a given $x \in \mathcal{L}$, we denote by $\pi(x) \in \Delta(B_\varepsilon)$ the

[7] The reader might have noticed that the distribution of the resulting allocation as a whole will not necessarily be the same as in the simultaneous auction, as we might "burn" items because they are claimed by the current agent and one of the "simulated copies" of the future agents, but not by the real future agents. However, by the linearity-of-expectation argument, this does not affect the expected welfare of the allocation. Separately, the contribution of each agent has the same distribution in both the simultaneous auction and in the online implementation.

resulting distribution of p' (and of p''), i.e., for $f \in B_\varepsilon$,

$$\pi_f(x) = \sum_{(v_i)_{i\in N} \in \times_{i\in N} V_i} \left(\prod_{i\in N} q_{i,v_i} \right) \sum_{(b_i)_{i\in N} \in B_\varepsilon^N : \max_{i\in N} b_{i,j} = f_j} \left(\prod_{i\in N} x_{b_i,i,v_i} \right).$$

Now, using this notation, for a pair $x, y \in \mathcal{L}$, we have that $y \in \psi(x)$ if and only if, for all $i \in N, v \in V_i, g \in B_\varepsilon$,

$$\begin{aligned} &\sum_{f\in B_\varepsilon} \sum_{p'\in B_\varepsilon, p''\in B_\varepsilon} y_{f,i,v} \cdot \pi_{p'}(x) \cdot \pi_{p''}(x) \cdot \Big(v\left(\{j : f_j > \max\{p'_j, p''_j\}\}\right) - f(M) \Big) \\ &\geq \sum_{p'\in B_\varepsilon, p''\in B_\varepsilon} \pi_{p'}(x) \cdot \pi_{p''}(x) \cdot \Big(v\left(\{j : g_j > \max\{p'_j, p''_j\}\}\right) - g(M) \Big). \end{aligned} \tag{7}$$

For fixed x, these are finitely many linear constraints on y, so $\hat{\psi}(x)$ is a convex and compact set.

Finally, again analogously to the proof of Lemma 5, the upper-hemicontinuity of $\hat{\psi}$ follows because $\pi_f(x)$ is continuous on x, and therefore, both sides of Equation (7) are continuous functions of x and y. □

The RSGs provided by Lemma 6 are exactly those we use for Theorem 2. To prove the approximation result we require two more lemmas. The first is the analogue of Lemma 1 (the Mirror Lemma).

Lemma 7. *Let $(D_i)_{i\in N}$ be the RSGs from Lemma 6 and p' the corresponding reserve prices. Let p'' be an i.i.d. copy of p'. For each $i \in N$, let $b_i \sim D_i(v_i)$ be the equilibrium bids. We have that for all $i \in N$, all $v \in V_i$, and all $f_i \in \mathcal{B}$,*

$$\mathbb{E}(v(A_i(f_i))) \geq \frac{1}{2} \cdot \mathbb{E}(v(\{j \in M : f_{i,j} > \max\{p'_j, p''_j\}\})).$$

Proof. For each $i \in N$ sample independent valuations $v''_i \sim \mathcal{F}_i$, and sample $b''_i \sim D_i(v''_i)$. Define for each $j \in M$ $p''_j = \max_{i\in N} b''_{i,j}$. We have that

$$\begin{aligned} E(v(A_i(f_i))) &= \mathbb{E}\left(v\left(\{j \in M : f_{i,j} > p'_j \geq b_{i',j},\ \forall i' \neq i\}\right)\right) \\ &= \mathbb{E}\left(v\left(\{j \in M : f_{i,j} > p'_j \geq b''_{i',j},\ \forall i' \neq i\}\right)\right) \\ &\geq \mathbb{E}\left(v\left(\{j \in M : f_{i,j} > p'_j \geq p''_j\}\right)\right), \end{aligned}$$

where the last inequality follows from the monotonicity of v. Now, since p' and p'' are i.i.d.,

$$\begin{aligned} &\mathbb{E}\left(v\left(\{j \in M : f_{i,j} > p'_j \geq p''_j\}\right)\right) \\ &= \mathbb{E}\left(v\left(\{j \in M : f_{i,j} > p''_j \geq p'_j\}\right)\right) \\ &= \frac{1}{2} \cdot \mathbb{E}\Big(v\left(\{j \in M : f_{i,j} > p'_j \geq p''_j\}\right) + v\left(\{j \in M : f_{i,j} > p''_j \geq p'_j\}\right)\Big) \\ &\geq \frac{1}{2} \cdot \mathbb{E}\Big(v\left(\{j \in M : f_{i,j} > \max\{p'_j, p''_j\}\}\right)\Big), \end{aligned}$$

where the last inequality follows from the subadditivity of v. □

The second lemma is the analogue of Lemma 4. Together with Lemma 7, it provides a lower bound on the utility of an agent in the equilibrium.

Lemma 8. *For a given $\varepsilon > 0$, let p' and p'' be a pair of i.i.d. vectors in B_ε. For every monotone valuation function $v : 2^M \to \mathbb{R}_+$ with $v(M) \le v_{\max}$, there exists a vector $f \in B_\varepsilon$ such that, for all $X \subseteq M$,*

$$\frac{1}{2}\mathbb{E}(v(\{j : f_j > \max\{p'_j, p''_j\}\})) - f(M) \ge \frac{1}{6}v(X) - \mathbb{E}(p'(X)) - \varepsilon \cdot |X|.$$

Proof. Let p''' be an i.i.d. copy of p' (and also independent of p''). Let X^* be the set that maximizes the right-hand side of the inequality of the lemma. We can take a random vector $\hat{f}$ defined as

$$\hat{f}_j = \begin{cases} p'''_j + \varepsilon & \text{if } j \in X^* \\ 0 & \text{otherwise,} \end{cases}$$

for each $j \in M$. We have that

$$\mathbb{E}\left(\frac{1}{2}v\left(\left\{j : \hat{f}_j > \max\{p'_j, p''_j\}\right\}\right) - \hat{f}(M)\right)$$
$$\ge \mathbb{E}\left(\frac{1}{2}v\left(X^* \cap \{j : p'''_j \ge \max\{p'_j, p''_j\}\}\right) - p'''(X^*) - \varepsilon \cdot |X^*|\right).$$

Since p', p'' and p''' are i.i.d., we can interchange them inside the expectation. Because of this and by the subadditivity of v,

$$\mathbb{E}\left(v\left(X^* \cap \{j : p'''_j \ge \max\{p'_j, p''_j\}\}\right)\right) \ge \frac{1}{3} \cdot v(X^*).$$

Therefore, we have that

$$\mathbb{E}\left(\frac{1}{2}v\left(\left\{j : \hat{f}_j > \max\{p'_j, p''_j\}\right\}\right) - \hat{f}(M)\right) \ge \frac{1}{6} \cdot v\left(X^*\right) - \mathbb{E}(p'''(X^*)) - \varepsilon \cdot |X^*|$$
$$= \max_{X \subseteq M}\left\{\frac{1}{6} \cdot v\left(X\right) - \mathbb{E}(p'(X)) - \varepsilon \cdot |X|\right\}.$$

Finally, the existence of a random vector $\hat{f} \in B_\varepsilon + \varepsilon$ that satisfies the inequality in expectation indicates that there is a vector $f \in B_\varepsilon + \varepsilon$ that also satisfies it. Now, imagine that $f_j > v_{\max}$ for some $j \in M$. If this is the case, the left-hand side is negative, which contradicts the fact that the right-hand side is at least 0, by taking $X = \emptyset$. Therefore, $f \in B_\varepsilon$. □

Proof of Theorem 2. As mentioned earlier, we take the RSGs $(D_i)_{i\in N}$ provided by Lemma 6. If we sample for each $i \in N$ independent valuations $v'_i \sim \mathcal{F}_i$ and bids $b'_i \sim D_i(v'_i)$, and set reserve prices $p'_j = \max_{i\in N} b'_{i,j}$, the lemma states there is a Bayesian Nash Equilibrium where agents bid $b_i \sim D_i(v_i)$. In this equilibrium the expected welfare is

$$\sum_{i\in N}\mathbb{E}(v_i(A_i(b_i))) = \sum_{i\in N}\mathbb{E}(v_i(A_i(b_i)) - b_i(M)) + \sum_{i\in N}\mathbb{E}(b_i(M))$$
$$= \sum_{i\in N}\mathbb{E}\left(\mathbb{E}\left(v_i(A_i(b_i)) - b_i(M) \,|\, v_i\right)\right) + \sum_{i\in N}\mathbb{E}(b_i(M)).$$

Now, notice that

$$\mathbb{E}\left(v_i(A_i(b_i)) - b_i(M) \,|\, v_i\right)$$

is the expected utility of agent i, when her valuation is v_i. Since this is an equilibrium, b_i maximizes this expression, given v_i. Let p'' be an i.i.d. copy of p', and let f_i be the vector guaranteed by Lemma 8. We have that

$$\begin{aligned}
\mathbb{E}\left(v_i(A_i(b_i)) - b_i(M) \,|\, v_i\right) &\geq \mathbb{E}\left(v_i(A_i(f_i)) - f_i(M) \,|\, v_i\right) \\
&\geq \mathbb{E}\left(\frac{1}{2}v_i(\{j : f_{i,j} > \max\{p'_j, p''_j\}\}) - f_i(M) \,\middle|\, v_i\right) \\
&\geq \max_{X \subseteq M}\left\{\frac{1}{6}\cdot v_i(X) - \mathbb{E}(p'(X)) - \varepsilon \cdot |X|\right\}.
\end{aligned}$$

In the second line we applied Lemma 7 and in the third line Lemma 8. Thus, we have that the expected welfare is at least

$$\begin{aligned}
&\sum_{i\in N}\mathbb{E}\left(\max_{X\subseteq M}\left\{\frac{1}{6}\cdot v_i(X) - \mathbb{E}(p'(X)) - \varepsilon\cdot|X|\right\}\right) + \sum_{i\in N}\mathbb{E}(b_i(M)) \\
&\geq \sum_{i\in N}\mathbb{E}\left(\frac{1}{6}\cdot v_i(OPT_i) - p'(OPT_i) - \varepsilon\cdot|OPT_i|\right) + \sum_{i\in N}\mathbb{E}(b_i(M)) \\
&= \frac{1}{6}\cdot\mathbb{E}(OPT) - \varepsilon\cdot|M| - \mathbb{E}(p'(M)) + \mathbb{E}\left(\sum_{i\in N} b_i(M)\right),
\end{aligned}$$

where in the second line we just replaced X with OPT_i, and in the third line we used linearity of expectation and the fact that $(OPT_i)_{i\in N}$ is a partition of M. We conclude by noticing that for every $j \in M$,

$$\mathbb{E}(p'_j) = \mathbb{E}(\max_{i\in M} b'_{i,j}) \leq \mathbb{E}(\sum_{i\in N} b'_{i,j}) = \mathbb{E}(\sum_{i\in N} b_{i,j}),$$

and taking a sufficiently small $\varepsilon > 0$. □