Theory of Computing Systems (2019) 63:1470–1498
https://doi.org/10.1007/s00224-019-09918-y

Online Random Sampling for Budgeted Settings

Alon Eden[1] · **Michal Feldman**[1,2] · **Adi Vardi**[1]

Published online: 30 March 2019

Abstract

We study online multi-unit auctions in which each agent's private type consists of the agent's arrival and departure times, valuation function and budget. Similarly to secretary settings, the different attributes of the agents' types are determined by an adversary, but the arrival process is random. We establish a general framework for devising truthful random sampling mechanisms for online multi-unit settings with budgeted agents. We demonstrate the applicability of our framework by applying it to different objective functions (revenue and liquid welfare), and a range of assumptions about the agents' valuations (additive or general) when selling identical divisible items. Our main result is the design of mechanisms for additive bidders with budget constraints that extract a constant fraction of the optimal revenue (under a standard large market assumption). We also show a mechanism that extracts a constant fraction of the optimal liquid welfare for general valuations.

Keywords Online mechanism · Mechanism design · Budgets · Revenue maximization · Liquid welfare

This article is part of the Topical Collection on *Special Issue on Algorithmic Game Theory (SAGT 2017)*

This work was partially supported by the European Research Council under the European Unions Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 337122.

✉ Alon Eden
alonarden@gmail.com

Michal Feldman
michal.feldman@cs.tau.ac.il

Adi Vardi
adi.vardi@gmail.com

[1] Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel

[2] Microsoft Research, Hertsliya, Israel

1 Introduction

In a typical setting of sales of online ad slots, advertisers arrive at different times, each with her own preferences. The auctioneer (e.g., *cnn.com*) decides about the allocation of the ad slots to the advertisers and how much to charge them. Scenario of this type have inspired the study of mechanism design in online settings, and random sampling has been proposed as a useful approach for the design of truthful mechanisms in online settings [23].

The random sampling framework was first introduced in the context of selling identical items in offline settings [22]. The basic idea of random sampling is to divide the agents into two sets of roughly equal size. Then, sell half of the items to each set at a price calculated according to the counterpart set. Random sampling mechanisms differ from each other in the allocation and pricing functions they apply, but they all operate according to the principle described above. The use of random sampling became widespread due to its desired properties: It is simple, trivially truthful and achieves good guarantees for a wide variety of settings. Our goal in this paper is to generalize the random sampling approach to settings in which agents arrive in an online fashion and have budgets.

Random sampling mechanisms have been previously proposed for online settings without budgets and for offline settings with budgets. Hajiaghayi et al. [23] applied the random sampling approach to online settings without budgets, and devised truthful mechanisms that approximate the auctioneer's revenue and the social welfare up to constant factors. While this was a major progress in the applicability of random sampling, their techniques are restricted to unit-demand valuations and quasi-linear utilities.

Borgs et al. [8] applied the random sampling approach to offline settings with budgets. Budget constraints impose major challenges on auction design since utilities are no longer quasi-linear. For example, the seminal VCG mechanism cannot work in non quasi-linear settings. In fact, it was shown that no truthful mechanism can give a non-trivial approximation to the social welfare in such settings [15].[1] Borgs et al. [8] were able to overcome these challenges in offline settings, and designed a truthful mechanism that gives constant approximation to revenue in offline settings with budgets (under a standard large market assumption).

The scenarios described above address either the budget constraints or the online nature of arrival. It might seem that we already have all the necessary ingredients to address the combination of both. Not surprisingly, however, it is not at all clear how to combine the existing techniques for our setting. First, some of the techniques above are tailored to unit-demand valuations, whereas we are interested in mechanisms for additive (or even general) valuations. Like in other problems in the literature on revenue approximation, the transition from unit-demand valuations to additive valuations requires entirely new techniques (e.g., pricing mechanisms for revenue maximization, [5, 10, 12]). Second, the combination of online arrival and

[1] This impossibility holds even if budgets are public.

budgets imposes new challenges that require new ideas. In what follows, we present our model followed by our results and techniques.

1.1 Model

We consider a setting with a set $N = \{1, \dots, n\}$ of n agents, who arrive in an online fashion. Each agent i has a private type, represented by the following parameters: (i) a_i and d_i: The respective arrival and departure times of agent i (clearly, $d_i \geq a_i$). The interval $[a_i, d_i]$ is referred to as agent i's *time frame*. (ii) $t_i = (v_i, b_i)$: The utility type of the agent, which contains the agent's valuation function $v_i : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ (which maps a given [possibly fractional] amount of items to a non-negative value) and the agent's budget, b_i. We consider a setting with m identical and divisible items. An allocation function determines an allocation x_i for agent i. The utility of agent i for obtaining an allocation x_i within her *real* time frame for a payment of p_i is:

$$u_i(x_i, p_i) = \begin{cases} v_i(x_i) - p_i & p_i \leq b_i \\ -\infty & p_i > b_i \end{cases} . \tag{1}$$

If allocated outside her time frame, the agent incurs the payment but gains no value. We note that the value an agent gets from an allocation within her time frame is the same, no matter at which point in the timeframe the allocation is being made. This fits common situations where the items might expire (e.g., ad slots) and is also the case in the [23]. There are scenarios where it is perfectly reasonable to assume the value might decrease at later stages of the time frame, or that the agent can sell the item acquired after his timeframe ends. We leave extensions of the model to future works.

We consider an online auction with a secretary flavor. An adversary states a vector of time frames $([a_1, d_1], [a_2, d_2], \dots, [a_n, d_n])$ such that $a_i < a_j$ for every $i < j$, and a vector of utility types $(t_1, t_2, \dots, t_n)$. A random permutation is used to match time frames with utility types. That is, a permutation $\pi : N \mapsto N$ is sampled uniformly at random and agent i's type is given by the tuple $\left(a_i, d_i, t_{\pi(i)}\right)$. As in [23],[2] we assume that arrival times are distinct, but the results also extend to non-distinct arrival times if agents cannot bid before their real arrival times.[3]

Agents report their type upon arrival, and can manipulate any component of it. In particular, they can report earlier or later arrival and departure times, and arbitrary utility types. We consider mechanisms that satisfy the following properties: (i) Feasibility: The mechanism does not sell more items than are available. (ii) Ex-ante Individual Rationality: An agent's expected utility from an allocation and payment is non-negative. (iii) Incentive Compatibility: An agent's expected utility is maximized when she reports her true type.

[2] Based on personal communication with the authors, this is essentially what is assumed for the correctness of Mechanism RM_k in Section 6 in [23].

[3] We refer to Appendix B for a description of the tie-breaking rule for this case.

1.2 Previous Techniques and Their Limitations

Hajiaghayi et al. [23] suggested the following scheme for online settings with unit-demand valuations (and no budget constraints): Set the first (roughly) half of the arriving agents as the sampling set, and the remaining ones as the performance set. All the revenue guarantees are obtained from the agents in the performance set; the agents in the sampling set are used merely for learning the valuations in the market. Induce the agents in the sampling set to reveal their true type by applying the VCG mechanism,[4] and use this information to extract revenue from the agents in the performance set. The proposed mechanism is truthful and obtains a constant factor approximation to revenue for unit-demand agents. Borgs et al. [8] devised allocation and payment functions that, when applied to a random sampling scheme, give good approximation to the revenue obtained in offline settings with budgets.

A natural approach for devising mechanisms for our setting is to combine the techniques of Hajiaghayi et al. [23] and Borgs et al. [8]. In particular, use the scheme of the former (to address the online nature of arrival) with the allocation and payments of the latter (to handle budgets). However, there are several obstacles to such an approach. First, the scheme proposed by [23] is restricted to agents with unit-demand valuations. In particular, agents cannot benefit from reporting a later arrival time (thereby being considered in the performance set) since the scheme ensures that the price in the sampling set never exceeds the price in the performance set. This is indeed sufficient for unit-demand agents, but does not extend beyond this class of valuations (in particular, to additive valuations). Second, while VCG can be used to encourage truthfulness in settings without budgets, it is well known that VCG is limited to quasi-linear settings, thus cannot work for settings with budgets.

1.3 Our Results and Techniques

We establish a general framework for devising truthful random sampling mechanisms for online multi-unit settings with budgeted agents. We demonstrate the applicability of our framework by applying it to different objective functions (revenue and liquid welfare), and a range of assumptions about the agents' valuations (additive or general). While we apply our framework to settings with budgets, it is general enough to be applied to other non quasi-linear settings.

Our framework splits the agents into sampling and performance sets of roughly the same size, based on the arrival time. In order to induce agents in the sampling set to report their type truthfully, we invoke random sampling at the time where the last agent in the sampling set arrives (recall VCG cannot be used due to budget constraints). In order to ensure that an agent cannot gain by delaying her arrival time, we use an *impersonation technique*. Namely, we treat an agent in the performance set *as*

[4]While VCG is defined and analyzed for offline settings, it is shown in [23] that it can also be applied in online settings by invoking it at the time where last agent arrives, serving only the agents that haven't departed yet. While this method does not give any revenue guarantees, it is only used to extract truthful information from agents in the sampling set.

if she were in the sampling set. In particular, an agent in the performance set receives at most the allocation she would have received had she been in the sampling set, and at the same price. The impersonation techniques guarantees truthfulness; it remains to design the allocation and payment functions so that they do not lead to a high loss in revenue. We devise such functions for two cases of interest. Our main result is the following:

Theorem: We devise truthful mechanisms for budgeted agents with additive valuations in online settings. These mechanisms give a constant approximation to the optimal revenue, under a large market assumption.

The proposed mechanism combines our general framework with the allocation and payment functions of [8]. We first show that the impersonation technique does not lead to a high loss in revenue. In particular, we have to prove that when the allocation of an agent from the performance set is limited by the allocation she would have gotten in the sampling set, the approximation guarantee is not significantly hurt. As in [8], we apply a large market assumption. Without this assumption, no truthful mechanism can achieve a constant approximation to the optimal revenue.

In addition to the setting above, we demonstrate the applicability of our framework in a different setting with budgets:

Theorem: We devise a truthful mechanism for budgeted agents with general (monotone) valuation functions in online settings. This mechanism gives a constant approximation to the optimal liquid welfare.

This setting imposes new challenges, since the valuations are general and no large market assumption is invoked.

1.4 Additional Related Work

Online auctions have been the subject of a vast body of work. Lavi and Nisan [26] introduced an online model in which agents have decreasing marginal valuations over identical indivisible items. While in their model the agent's value is private, her arrival time is public. A wide variety of additional online auction settings, such as digital goods [6, 7] and combinatorial auctions [2], have been studied under the assumption of private values and public arrival times.

Similarly to our model and the model of [23], online settings in which agents arrive in a random order were also considered in [4, 25]. Friedman et al. [20] considered the case where an agent's arrival time is also part of her private information, and thus can also be manipulated. Additional auction settings were studied under the assumption that an agent's arrival and departure times are private [24, 27, 31]. The reader is referred to [30] for an overview of online mechanism design (Fig. 1).

Offline mechanisms for budgeted agents have been considered, both for revenue maximization and welfare maximization. In the context of revenue maximization, Abrams et al. [1] also considered the model of Borgs et al. [8] where indivisible and identical items are sold to additive agents with hard budget constrains. Abrams establishes the relation between the revenue achieved by an optimal uniform-price mechanism and an optimal mechanism with heterogeneous prices. Maximizing revenue was also considered in Bayesian settings [11] and with the goal of approximating the optimal envy free revenue [13, 17].

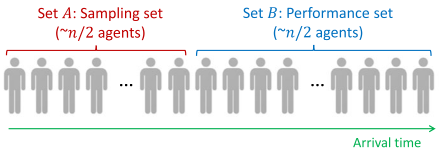

(a) Agents are split into sets A and B.

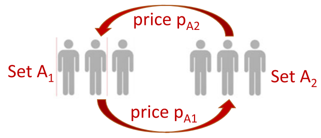

(b) Upon the arrival of the last agent in A, random sampling with $m/2$ items is applied on set A.

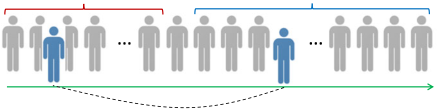

(c) $m/2$ items are sold to agents in B, upon arrival, where each agent in B is treated as if she were in set A.

Fig. 1 The different steps of Online-RS

As stated above, in general budgeted settings it is impossible to approximate the optimal social welfare using a truthful mechanism. Therefore, previous works that seeks to maximizing efficiency ensure that the outcome is Pareto Optimal [14, 18, 21]. In [13] the authors focus on approximating the welfare of the optimal envy free mechanism. The most related works to our model are [8, 15, 23, 28], which are discussed in great detail throughout the paper.

2 Online Implementation of Random Sampling Based Mechanisms

In this section, we present a general framework for adjusting offline random sampling based mechanisms to online settings. Our template mechanism Online-RS is specified

Online-RS $(N, m, \mathcal{P}, \mathit{Allocation})$

Partition:

1. Rename agents according to a permutation $\pi : [n] \to [n]$ chosen uniformly at random.
2. $j \leftarrow Bin(n-1, 1/2) + 1$, $A \leftarrow$ first j arriving agents, $B \leftarrow N \setminus A$, $t_0 \leftarrow \hat{a}_j$.

Sampling phase: (Set A)

1. At time t_0: Partition the agents in A into sets A_1, A_2, uniformly at random. (without loss of generality, $j \in A_2$).
2. $x^1 \leftarrow \mathit{Allocation}(A_1, \mathcal{P}(A_2), \frac{m}{4}, \infty)$, $x^2 \leftarrow \mathit{Allocation}(A_2, \mathcal{P}(A_1), \frac{m}{4}, \infty)$.
3. For every agent $i \in A_1$: if $\hat{d}_i \geq t_0$, agent i is allocated x_i^1 items and pays $\mathcal{P}(A_2)$ per item. Apply the analogous rules for agents in A_2.

Performance phase: (Set B)
$\alpha_1, \alpha_2 \leftarrow \frac{m}{4}$. Upon arrival of agent $i \in B$:

1. with probability $1/2$, $i \in B_1$, agent i is allocated
$x_i \leftarrow \mathit{Allocation}_i\left(\{A_1 \cup i\}, \mathcal{P}(A_2 - j), \frac{m}{4}, \alpha_1\right)$, and pays $\mathcal{P}(A_2 - j)$ per item.
$\alpha_1 \leftarrow \alpha_1 - x_i$.
2. with probability $1/2$, $i \in B_2$, agent i is allocated
$x_i \leftarrow \mathit{Allocation}_i\left(\{A_2 \cup i\} - j, \mathcal{P}(A_1), \frac{m}{4}, \alpha_2\right)$, and pays $\mathcal{P}(A_1)$ per item.
$\alpha_2 \leftarrow \alpha_2 - x_i$.

Fig. 2 A template online random sampling mechanism in Fig. 2. It receives as input a set of agents N, number of items m, an *offline* pricing function $\mathcal{P}$, and an *offline* allocation function Allocation.

The price function $\mathcal{P}$ receives a set of agents S and computes a per-item price. The allocation function $Allocation$ receives a set of agents S, a per-item price p, total number of items k, and a cap per agent ℓ (where the cap per agent limits the number of items an individual agent can get). It outputs an allocation vector, where $Allocation_i$ denotes the number of items allocated to agent i.

Before we give a formal definition of Mechanism Online-RS, we provide a non-formal description.

Step (a) [splitting]: We split the agents into two sets A and B, such that roughly the first $n/2$ arriving agents are placed in set A and the rest in set B (see Fig. 1a).

Step (b) [sampling set]: We apply an (offline) random sampling mechanism on set A. In particular, once the last agent of set A arrives, we divide the agents into two sets, A_1 and A_2, uniformly at random. We set a per-item price p_1 by applying $\mathcal{P}$ on set A_1 and sell $m/4$ items to agents in A_2. We apply an analogous procedure to agents in set A_2 (see Fig. 1b).

Step (c) [performance set]: Upon the arrival of an agent in set B, she is placed in one of sets B_1 or B_2 uniformly at random. An agent i in B_1 is treated as if she were in A_1, with the additional limitation determined by the actual number of remaining items. Therefore, the price is calculated based on A_2. An agent in B_2 is treated analogously, with allocation and payments according to sets A_2 and A_1, respectively (see Fig. 1c).

The repositioning of an agent from set B in set A pushes the last agent in A toward B, thereby altering the set A. This issue is addressed by a careful specification of the sets by which the allocation and price are determined to agents from B. Let j denote the last agent to arrive to A (as defined in Online-RS) and let $N' = N - \{j\}$. The following observation is useful in our analysis.

Observation 1 *Sets A_1, $A_2 - j$, B_1 and B_2 form a uniform partition of the agents in N'.*

Proof Observe that $|N'| = n - 1$ and $|A - j|$ is distributed according to the Binomial distribution $B(n-1, 1/2)$. Since the agents of N' arrive in a random order, the sets $A - j$ and B form a uniform random partition of N'. Moreover, agents in $A - j$ are divided into sets A_1 and $A_2 - j$ uniformly at random, and agents in set B are divided into sets B_1 and B_2 in a similar manner. □

Recall that while a_i and d_i are the real arrival and departure times of agent i, an agent may misreport its type. We denote the reported arrival and departure times of agent i by $\hat{a}_i$, $\hat{d}_i$, respectively. We next provide a formal definition of Mechanism Online-RS.

2.1 Truthfulness and Feasibility

An allocation is said to be *feasible* if the total number of allocated items does not exceed the limitation and the number of items allocated to every individual agent does not exceed the cap per agent. Consider the offline mechanism **M**, that determines allocation based on *Allocation* and payment based on price per item $\mathcal{P}(\cdot)$. An allocation function is said to be *cap monotone* if the expected utility of an agent in mechanism **M** is non-decreasing in the cap per agent.

The main result in this section is that mechanism Online-RS is truthful and feasible as long as mechanism **M** satisfies incentive compatibility (IC),[5] individual rationality (IR), feasibility and cap monotonicity.

Theorem 1 *Mechanism* Online-RS *is truthful and feasible if: (i) Mechanism* **M** *is IC, IR, feasible and cap monotone. (ii) Agents' valuations are monotonically non-decreasing.*

Proof First we prove that mechanism Online-RS is truthful. We consider some agent i and show that her expected utility cannot increase when misreporting her type. Recall that the type of the agent is comprised of her utility type t_i and her time frame $[a_i, d_i]$.

Misreporting t_i Agents in sets A and B cannot change their price or the set of agents they compete with by misreporting their utility type t_i, and the mechanism allocates items only within the reported time frame of the agent. In addition, from the IR and IC properties of **M**, we get that given a fixed price and set of agents, an agent maximizes her utility function when reporting her true type.

Misreporting d_i Agents in B are being allocated upon arrival. Hence, misreporting d_i has no effect. Agents in A are being allocated at time t_0 (the arrival time of agent j). Therefore, by reporting an earlier departure time, agent i might not be allocated. By delaying her departure time, agent i might receive the items outside her time frame.

Misreporting a_i We analyze the following cases.

1. The mechanism places agent i in set B according to both her true arrival time and reported arrival time: Recall that agents in B are being allocated upon arrival. By reporting an earlier arrival time, agent i receives the items outside her time frame. By delaying her arrival time the number of real items will not increase (it might decrease, since other agents might be allocated before agent i). Hence, by the cap monotonicity property of **M**, the agent's expected utility might decrease.
2. The mechanism places agent i in set B according to her true arrival time and in set A according to her reported arrival time: Since agent i is placed in set B according to her true arrival time, $a_i > t_0$. Hence, when agent i is placed in set A she receives the items outside her time frame.

[5] A mechanism is IC if every agent's expected utility is maximized when reporting her true valuation. A mechanism is IR if every agent's expected utility is non-negative when reporting her true valuation.

3. The mechanism places agent i in set A according to both her true arrival time and reported arrival time: Since set A is the same in both cases, so is the expected allocation she receives (recall that the allocation function is offline).
4. The mechanism places agent i in set A according to her true arrival time and in set B according to her reported arrival time: Note that by moving to set B, agent i caused agent j to move to set A. If agent i had been reporting her true arrival time, with probability 1/2 she would have been placed in set A_1 and be allocated according to $Allocation\left(\{A_1 \cup i\}, \mathcal{P}(A_2 - j), \frac{m}{4}, \infty\right)$, and with probability 1/2 she would have been placed in set A_2 and be allocated according to $Allocation\left(\{A_2 \cup i\} - j, \mathcal{P}(A_1), \frac{m}{4}, \infty\right)$.

   By moving to set B, with probability 1/2 the agent is allocated according to $Allocation\left(\{A_1 \cup i\}, \mathcal{P}(A_2 - j), \frac{m}{4}, \alpha_1\right)$, and with probability 1/2 she is allocated according to $Allocation\left(\{A_2 \cup i\} - j, \mathcal{P}(A_1), \frac{m}{4}, \alpha_2\right)$. Since α_1 and α_2 are at most $m/4$, by the cap monotonicity property of **M**, we get that the agent's expected utility only decreased by moving to set B.

Next we show that Online-RS returns a feasible allocation whenever **M** is feasible. First note that since **M** is feasible, agents in set A_1 receive at most $m/4$ items and the same constraint applies to agents in set A_2. Next, note that the cap limitation used by **M** to allocate items to the agents in set B_1 guarantee that they receive at most $m/4$ items and the same guarantee applies to agents in set B_2. □

3 Revenue Maximization for Additive Agents with Budgets

In this section, we study a mechanism for agents with hard budget constraints and additive valuation functions. The utility of agent i for obtaining an allocation x_i within her time frame for a payment of p_i is:

$$u_i(x_i, p_i) = \begin{cases} v_i \cdot x_i - p_i & p_i \leq b_i \\ -\infty & p_i > b_i \end{cases} . \tag{2}$$

Our goal in this section is to devise a truthful mechanism that approximates the optimal revenue.

Let S be a set of agents and k a number of items. We define $P_k(S)$ to be the price that maximizes the revenue of selling at most k items to agents in S at a uniform price: $P_k(S) = \text{argmax}_p \min\left(\sum_{i \in S, v_i \geq p} \frac{b_i}{p}, k\right) \cdot p$. Let $OPT(S, k)$ be the optimal revenue from selling at most k items to agent in S at a uniform price: $OPT(S, k) = \min\left(\sum_{i \in S, v_i \geq P_k(S)} \frac{b_i}{p}, k\right) \cdot P_k(S)$. Let $OPT^*(S, k)$ be the optimal revenue from selling at most k items to agents in S at non-uniform prices (i.e., where the seller can discriminate between the agents).

The following lemma of Abrams [1] implies that it is sufficient to bound the performance of our mechanism with respect to $OPT^*(S, k)$.

Rev-Allocation (S, p, k, ℓ)

1. Initialize $k' \leftarrow k$, $x_i \leftarrow 0$ for every $i \in S$.
2. For each $i \in S$ approached in the order of the random renaming (see first line in Online-RS):
   - If $v_i \geq p$:
     $x_i \leftarrow \min\{b_i/p, k'\}$
     $k' \leftarrow k' - x_i$
3. For every i, if $x_i > \ell$, set $x_i \leftarrow \ell$.
4. Return x.

Fig. 3 The allocation function used for the revenue maximizing mechanism Rev-RS

Lemma 1 [1] *For every set of agents S and k items,* $OPT(S, k) \geq \frac{OPT^*(S,k)}{2}$.

3.1 The Mechanism

Our mechanism, Rev-RS is an implementation of mechanism Online-RS with the following allocation and payment functions. The allocation function, Rev-Allocation (see Fig. 3) approaches the agents sequentially, allocates the items in a greedy manner, and verifies that none of the agents exceeds the cap limitation. The payment function is $P_{m/4}(\cdot)$.

Recall that without a large market assumption, no truthful mechanism can achieve a constant approximation to the revenue [8]. Adopting the parameter used in Borgs et al. [8], we define $\epsilon(S, k) = \frac{\max\{b_i\}_{i \in S}}{OPT(S,k)}$. Intuitively, a smaller ϵ implies a larger market. The main result of this section is:

Theorem 2 Rev-RS *is a truthful mechanism that gives an 8-approximation to the optimal revenue as ϵ tends to 0.*

In Section 3.4 we prove the above theorem.

3.2 A Constant Loss Due to the Impersonation Technique

Our mechanism uses an impersonation technique. Specifically, it calculates the allocation for each agent in set B_1 as if she were in a random permutation in set A_1, and for each agent in B_2 as if she were in $A_2 - j$. According to Observation 1, the sets A_1 and B_1 are equally distributed. Therefore, one can view the impersonation technique as follows. We are given a set $N_1 = A_1 \cup B_1$. Then, the sets are partitioned uniformly at random into sets A_1 and B_1, and each agent in B_1 gets the price and allocation as if she were placed in set A_1. The same holds for $N_2 = (A_2 - j) \cup B_2$. With this interpretation in mind, we prove that we lose at most a factor 2 in the revenue due to the impersonation technique.

Let Offline-Sale(S, p, k) be a mechanism that allocates k items to agents in set S according to the allocation returned by Rev-Allocation(S, p, k, ∞), and charges a price p per item. Impersonation-Sale(S, p, k) (see Fig. 4) is a mechanism that sells according to the impersonation technique described above. *I.e.*, it divides S into sets S_1 and S_2 uniformly at random. It then sells items to agents in set S_1 in a random

Impersonation-Sale (S, p, k)

1. Initialize $\ell \leftarrow k$.
2. For each agent $i \in S$: With probability of $\frac{1}{2}$ place i in set S_1 and with probability of $\frac{1}{2}$ in set S_2.
3. Let $\pi_{S_1} : \{1, 2, \dots, |S_1|\} \rightarrow S_1$ be a random permutation of the agents of S_1.
4. For $i = 1, 2, \dots, |S_1|$:
   (a) Let $j \leftarrow \pi_{S_1}(i)$, and let x be the allocation returned by Rev-Allocation $(S_2 \cup \{j\}, p, k, \ell)$.
   (b) Allocate to agent j x_j items at a price of p per item.
   (c) Update $\ell \leftarrow \ell - x_j$.

Fig. 4 A mechanism defined for the analysis of the impersonation technique

order at a fixed price p per item. The allocation for each agent is calculated as if the agent was in a random permutation in set S_2.

Observation 2 *Let T and S be two sets of agents such that $T \subseteq S$ and let k be a number of items. For every agent $i \in T$,* E[Rev-Allocation$_i(T, p, k, \infty)$] $\geq$ E[Rev-Allocation$_i(S, p, k, \infty)$].

Proof Let i be some agent in $T \cap S$. For a given permutation π determined by the random renaming of agent (the first step of Online-RS), it is clear that the set of agents getting allocated before agent i in Rev-Allocation(T, p, k, ∞) is a subset of the agents getting allocated before agent i in Rev-Allocation(S, p, k, ∞). Since every agent gets as much as her budget allows, we get that agent i has less items available to her when getting allocated in Rev-Allocation(S, p, k, ∞), and the lemma follows. □

Lemma 2 *For every set S, price p and a (possibly fractional) number of items k, the expected revenue from mechanism* Impersonation-Sale(S, p, k) *is at least half of the expected revenue from mechanism* Offline-Sale(S, p, k).

Proof Let r be the expected revenue from mechanism Offline-Sale(S, p, k). For the sake of the analysis, assume that Offline-Sale begins by placing each agent $i \in S$ in S_1 and S_2 uniformly at random. Since each agent has a probability of $\frac{1}{2}$ to be placed in S_1, the expected revenue extracted from agents in S_1 is $\frac{r}{2}$. Therefore, it is sufficient to show that for each partition of S into S_1 and S_2, the expected revenue Impersonation-Sale(S, p, k) extracts from agents in S_1 is at least the expected revenue Offline-Sale(S, p, k) extracts from agents in S_1. There are two cases. If Impersonation-Sale(S, p, k) allocates the entire fraction of the item, then it collects the maximal revenue possible from price p from agents in S_1 and the claim is obviously true.

Otherwise, for each agent $i \in S_1, x_i < \ell$, where x_i is the allocation agent i receives in Impersonation-Sale. Therefore, the allocation of every agent $i \in S_1$ is not bounded by ℓ in Rev-Allocation$(S_2 \cup \{i\}, p, k, \ell)$, *i.e.*, the agent receives the same allocation she would receive in Rev-Allocation$(S_2 \cup \{i\}, p, k, \infty)$. Let $\bar{x}$ be the allocation returned by Rev-Allocation(S, p, k, ∞). According to Observation 2, $\mathrm{E}[x_i] \geq \mathrm{E}[\bar{x}_i]$. Since the revenue extracted out of agent i in Impersonation-Sale is $x_i \cdot p$ while the revenue extracted out of agent i in Offline-Sale is $\bar{x}_i \cdot p$, we get the desired result. □

Therefore, it is sufficient to bound the performance of Offline-Sale(S, p, k).

3.3 Useful Lemmas from Borgs et al. [8]

We use lemmas from Borgs et al. [8] to bound the performance of Offline-Sale (S, p, k). Their mechanism, which we refer to as Offline-Rev-Maximization divides the agents into two sets S_1 and S_2 uniformly at random. Then it calculates the optimal uniform price for selling at most $k/2$ items to agents in sets S_1 and S_2, denoted by p_1 and p_2, respectively. As a final step, the mechanism sells $k/2$ items to agents in S_1 at a price per item p_2, and $k/2$ items to agents in S_2 at a price per item p_1.

The following lemmas follow directly from the proofs in [8].[6] For the sake of completeness, we include the full proofs of these lemmas in Appendix A.

Lemma 3 [8] *Let r_{S_1}, r_{S_2} be the revenue from selling at most $k/2$ items to sets S_1, S_2, respectively, at a price of $P_k(S)$ per item. For any $\delta \in [0, 1]$, the probability that both that both r_{S_1} and r_{S_1} are greater than or equal to $\frac{1-\delta}{2} OPT(S, k)$ is at least* $1 - 2e^{-\delta^2/(4\epsilon(S,k))}$.

Lemma 4 [8] *For any $\delta \in [0, \frac{1}{3}]$, the probability that the revenue of mechanism* Offline-Rev-Maximization *from agents in set S_1 greater than or equal to $\frac{1-3\delta}{2} OPT(S, k)$, is at least* $1 - 2e^{-\delta^2/(4\epsilon(S,k))}$.

The above lemmas hold even when mechanism Offline-Rev-Maximization sell items in an arbitrary order, in particular, according to the online arrival order. Moreover, our mechanism partitions the agents into sets A and B in an equivalent way to mechanism Offline-Rev-Maximization (see Observation 1). Therefore, we can use the above theorems in our analysis.

3.4 Analysis Of Our Mechanism

Theorem 3 *Mechanism* Rev-RS *is a truthful and feasible mechanism.*

Proof It is sufficient to show that Rev-Allocation maintains all the properties stated in Theorem 1.

- IR, IC: Rev-Allocation only allocates items to agents with value per item that exceeds the price. Moreover, each agent gets the maximal amount of available items, subject to its budget constraint. By the additive nature of the agents, the properties follow.
- Feasibility: Trivially follows from the definition of Rev-Allocation.
- Cap Monotonicity: Let $\alpha_1 < \alpha_2$. For a given order of agents in Rev-Allocation, the amount of items allocated to an agent when α_2 is the cap per agent is bigger

[6]Lemmas 3 and 4 follow from the respective proofs of Lemma 5.2 and Theorem 5.1 in [8]. While these are stronger statements than the ones that appeared in [8], they are implied by the proofs in [8].

than the amount of items allocated to an agent when α_1 is the cap per agent. The property again follows by the additive nature of the agents.

□

Recall that $N' = N - j$ and let $OPT = OPT(N, m)$ and $OPT' = OPT(N', m)$.

Lemma 5 $\Pr\left[OPT' \geq OPT/2\right] \geq 1 - 1/n$.

Proof There can be at most a single agent from set N that may contribute more then $OPT/2$ to the optimal revenue. Since agent j is chosen uniformly at random from set N, the probability that such an agent is not in $N - j$ is at most $1/n$. □

Let $N_1 = B_1 \cup A_2 - j$, $N_2 = B_2 \cup A_1$, $OPT_1 = OPT(N_1, m/2)$, $OPT_2 = OPT(N_2, m/2)$, $\epsilon' = \epsilon(N', m)$. The following lemma proves that both N_1 and N_2 contain a significant fraction of N'.

Lemma 6 $\Pr\left[OPT_1 \geq \frac{1}{4}OPT', OPT_2 \geq \frac{1}{4}OPT'\right] \geq 1 - 2e^{-1/16\epsilon'}$.

Proof According to Observation 1, N_1 and N_2 form a uniform partition of N'. Hence, by applying Lemma 3 with $\delta = \frac{1}{2}$, $S = N'$, $S_1 = N_1$, $S_2 = N_2$, $k = m$ we get that $OPT_1 \geq r_{N_1} \geq \frac{1}{4}OPT'$ and $OPT_2 \geq r_{N_2} \geq \frac{1}{4}OPT'$ with probability $1 - 2e^{-1/16\epsilon'}$. □

The following is the main technical lemma of this section.

Lemma 7 *For every $\delta \in [0, \frac{1}{3}]$, the probability that the expected revenue obtained from mechanism* Rev-RS(N, m) *is greater than $\frac{(1-3\delta)OPT'}{8}$ is at least $\left(1 - 2e^{-1/16\epsilon'}\right)\left(1 - 4e^{-\delta^2/(16\epsilon')}\right)$.*

Proof Let $\epsilon_1 = \epsilon(N_1, m/2)$, $\epsilon_2 = \epsilon(N_2, m/2)$. We calculate the expected revenue under the event that $OPT_1 \geq \frac{1}{4}OPT'$ and $OPT_2 \geq \frac{1}{4}OPT'$ (and therefore also $\epsilon_1 \leq 4\epsilon'$, $\epsilon_2 \leq 4\epsilon'$). Let r be the expected revenue obtained by mechanism Offline-Sale$(B_1, P_{\frac{m}{4}}(A_2 - j), \frac{m}{4})$. Since each agent $i \in N_1$ is randomly and independently placed in either B_1 or $A_2 - j$ with equal probability, we can apply Lemma 4 with $S_1 = B_1$, $S_2 = A_2 - j$ and get that the probability that $r \geq \frac{1-3\delta}{2}OPT_1$ is at least $1 - 2e^{-\delta^2/(4\epsilon_1)} \geq 1 - 2e^{-\delta^2/(16\epsilon')}$. Since $OPT_1 \geq \frac{1}{4}OPT'$, this implies that the probability that $r \geq \frac{1-3\delta}{8}OPT'$ is also at least $1 - 2e^{-\delta^2/(16\epsilon')}$. Since $B_1 \cup A_1 \supseteq B_1$ we have the same guarantees on the expected revenue obtained by mechanism Offline-Sale$(B_1 \cup A_1, P_{\frac{m}{4}}(A_2 - j), \frac{m}{4})$.

By applying Lemma 2 with $S = B_1 \cup A_1$, $S_1 = B_1$, $S_2 = A_1$, $p = P_{\frac{m}{4}}(A_2 - j)$ and $k = \frac{m}{4}$, we get that the probability that the expected revenue from Impersonation-Sale$(B_1 \cup A_1, P_{\frac{m}{4}}(A_2 - j), \frac{m}{4})$ is greater than $\frac{1-3\delta}{16}OPT'$ is at least $1 - 2e^{-\delta^2/(16\epsilon')}$. From a symmetric argument we get that the probability that the expected revenue from mechanism Impersonation-Sale$(B_2 \cup A_2 - j, P_{\frac{m}{4}}(A_1), \frac{m}{4})$ is greater than $\frac{1-3\delta}{16}OPT'$ is at least $1 - 2e^{-\delta^2/(16\epsilon')}$. Using a union bound we get that the probability that the expected revenue obtained from both mechanisms is greater than $\frac{1-3\delta}{8} OPT'$ is at least $1 - 4e^{-\delta^2/(16\epsilon')}$. Note that Lemma 2 can be used despite the additional restrictions of $OPT_1 \geq \frac{1}{4} OPT'$ and $OPT_2 \geq \frac{1}{4} OPT'$ (which implies restrictions on $N_1 = B_1 \cup A_{2-j}$ and $N_2 = B_2 \cup A_1$), since these conditions are symmetric.

We complete the proof by observing that the revenue obtained from agents of B_1 in Rev-RS is exactly the revenue obtained from mechanism Impersonation-Sale$(B_1 \cup A_1, P_{\frac{m}{4}}(A_2 - j), \frac{m}{4})$. Indeed, in Rev-RS every agent in $A_1 \cup B_1$ is placed in A_1 or B_1 with probability $1/2$, and the agents in B_1 arrive in a random order (the random permutation that determines the players types). Every incoming agent $i \in B_1$ is allocated and charged according to Rev-Allocation$(A_1 \cup \{i\}, P_{\frac{m}{4}}(A_2 - j), \frac{m}{4}, \alpha)$, where α is the number of remaining items, the same as in Impersonation-Sale. Similarly, the revenue obtained from the agents of B_2 in Rev-RS is exactly the revenue obtained from mechanism Impersonation-Sale$(B_2 \cup A_2 - j, P_{\frac{m}{4}}(A_1), \frac{m}{4})$. Combining the last statement with Lemma 6 yields the desired result. □

Let $OPT^* = OPT^*(N, m)$ and $\epsilon = \epsilon(N, m)$. We now use Lemma 1 and Lemma 5, to bound the expected revenue from mechanism Rev-RS with respect to OPT^*.

Corollary 1 *For every $\delta \in [0, \frac{1}{3}]$, the probability that the expected revenue obtained from mechanism* Rev-RS(N, m) *is greater than $\frac{(1-3\delta)OPT^*}{32}$ is at least $(1 - 1/n)\left(1 - 2e^{-1/32\epsilon}\right)\left(1 - 4e^{-\delta^2/(32\epsilon)}\right)$.*

The proof of Theorem 2 follows.

Proof of Theorem 2 By definition, when ϵ tends to 0, OPT' approaches OPT. Therefore, the bound from Lemma 7 becomes: for every $\delta \in [0, \frac{1}{3}]$, the probability that the expected revenue obtained from mechanism Rev-RS(N, m) is greater than $\frac{(1-3\delta)OPT^*}{16}$ is at least $\left(1 - 2e^{-1/32\epsilon}\right)\left(1 - 4e^{-\delta^2/(32\epsilon)}\right)$. When ϵ tends to 0, for every δ this probability tends to 1. Therefore, the current analysis gives a 16-approximation.

We can further improve the approximation factor to be arbitrarily close to 8. This improvement can be achieved by applying Lemma 3 using an arbitrarily small constant δ in the proof of Lemma 6. By doing so, we get that OPT_1 and OPT_2 can be arbitrarily close to $OPT'/2$ (and therefore, ϵ_1 and ϵ_2 are bounded from above by a value arbitrarily close to $2\epsilon'$) with probability that still tends to 1 as ϵ tends to 0. This improves the guarantee of Lemma 3 by a constant arbitrarily close to 2, and yields an 8-approximation as ϵ tends to 0. □

4 Liquid Welfare Maximization for Budgeted Agents

In this section, we design a mechanism for agents with hard budget constraints and monotonically non-decreasing valuation functions. Our goal is to devise a truthful mechanism that approximates the optimal liquid welfare, defined below.

Budget constraints impose challenges also with respect to the social welfare objective function. It is well known that if budgets are private, no truthful mechanism can give a non-trivial approximation (better than n) to the social welfare, even for the case of a single item. This is because any such mechanism must allocate the item to an agent with a high value relative to all other agents, for any budget she reports. To make things worse, even if budgets are public, it has been shown in [15] that no truthful auction among n agents can achieve a better than n approximation to the optimal social welfare.

This problem was addressed by Dobzinski and Paes Leme [15], who introduced the liquid welfare objective function. A natural interpretation of the social welfare function is the maximum revenue an omniscient auctioneer can extract. Applying this interpretation to budgeted settings gives rise to the liquid welfare objective function, defined as follows: For every agent, consider the minimum between her value for the allocation (i.e., her willingness to pay) and her budget (i.e., her ability to pay), and take the sum of this minimum over all agents (we formally define this below). Using this interpretation of welfare in budgeted settings has become commonly used and led to many positive results [3, 9, 15, 16, 19, 28, 29] that are unattainable when one considers the standard social welfare objective.

In our settings, the items are identical and divisible. Therefore, it is without loss of generality to consider a single divisible item. The valuation function of an agent i, $v_i : [0, 1] \to \mathbb{R}^+$ is a non-decreasing valuation function of the agent, and b_i is the agents' budget. The *liquid welfare* of an agent i from allocation x_i is defined as the minimum between her valuation for x_i and her budget, and is denoted by $\overline{v}i(x_i) = \min(v_i(x_i), b_i)$.

We refer to $\overline{v}i(1)$ as the liquid value of agent i (i.e., the liquid welfare of the agent from receiving the entire item). The liquid welfare of an allocation vector x is $\overline{W}(x) = \sum_i \overline{v}i(x_i)$. Let $OPT = \max_x \overline{W}(x)$ be the optimal liquid welfare. We use the following lemma from Lu and Xiao [28].

Lemma 8 [28] *For every truthful and individually rational mechanism, the liquid welfare is greater than or equal to the revenue.*

4.1 High-Level Overview of the Mechanism

We refer to an agent whose liquid value is a constant fraction of OPT as a *dominant* agent. A large market can be interpreted as a market without dominant agents. Since we do not assume a large market, we device several mechanisms, each of then deals with a different "state" of the market.

A Single Dominant Agent In this case, it is sufficient to sell the entire item to the dominant agent. In an offline setting without budgets we can use the second price auction. Nevertheless, Lu and Xiao [28] showed that this is not truthful in an offline settings with budget constraints. They overcome the problem by setting the price slightly above the second highest liquid value. Our mechanism, SP-One-Dominant (see Section 4.2) is an adaptation of their mechanism to online settings.

At Least Two Dominant Agents In this case, one cannot use the mechanism from the previous section since it requires a gap between the two agents with the highest liquid values. Nevertheless, since there are at least two dominant agents, if one uses random sampling, then with a constant probability there will be a dominant agent in both the sampling and the performance sets. In this case, if we sell the item using a price that is computed according to the liquid welfare of the dominant agent in the sampling set, we get a high liquid welfare. Our mechanism, RS-Many-Dominant (see Section 4.3) is an implementation of mechanism Online-RS with an allocation function that allocates the entire item to a single agent and with a suitable pricing function.

No Dominant Agent Our mechanism, RS-No-Dominant (see Section 4.4) is an implementation of mechanism Online-RS with suitable pricing and allocation functions. Since all agents are relatively small compared to OPT, random sampling techniques guarantee that the agents are (roughly) equally divided between the sampling and performance sets. It remains to prove that we only lose a constant factor in the performance due to the impersonation technique. The main challenge here is the fact that agents' valuations are general (as opposed to the additive valuations considered in Section 3).

Our mechanism, LW, is a probabilistic combination of the above three Mechanisms. It applies the three aforementioned mechanisms with constant probabilities λ_1, λ_2 and λ_3, respectively. These constants will be determined later. To prove that LW achieves a constant factor approximation, we show that each of the mechanisms achieves a constant factor approximation with respect to its corresponding scenario.

Let $N' = N - \{j\}$. Let $OPT' = OPT(N')$. To simplify the analysis, we normalize OPT' to equal 1. While OPT' depends on the choice of j, the fraction of OPT' guaranteed by our mechanism holds for every j, therefore this normalization is valid.

The following lemma shows that we do not lose much by disregarding a random agent j in the analysis.

Lemma 9 $\Pr\left[OPT' \geq \left(1 - \frac{1}{k+1}\right) OPT\right] \geq 1 - \frac{k}{n}$ *for every* $k \leq n$.

Proof Let x^* be an optimal allocation. Let I_k be the set of the k agents with the highest liquid welfare under x^*. Note that for every agent $\ell \in (N \setminus I_k)$, we have that $\overline{v}\ell(x^*_\ell) \leq \frac{OPT}{k+1}$; otherwise, $\sum_{i \in (I_k + \ell)} \overline{v}i(x^*_i) > (k+1)\frac{OPT}{k+1} = OPT$ for some $\ell \in (N \setminus I_k)$. Since agent j is chosen uniformly at random from set N (recall that the agents arrive in a random order), the probability that $j \in I_k$ is $\frac{k}{n}$. In case that $j \in (N \setminus I_k)$, we get that

$$OPT' \geq \sum_{i \in N-j} \overline{v}i(x^*_j) \geq OPT - \frac{OPT}{k+1} = \left(1 - \frac{1}{k+1}\right) OPT.$$

□

As stated above, each of the mechanisms described above is guaranteed to perform well for a certain input type. Let $i_{(1)}$ and $i_{(2)}$ be the agents in N' with the respective highest and second highest liquid value. Let $\beta = \frac{1}{200}$ and $\gamma = \frac{200}{201}$. We distinguish between three types of inputs:

1. No-Dominant: $\max_{i \in N'} \overline{v}i(1) \leq \beta$.
2. Many-Dominant: $\overline{v}i_{(1)}(1) \geq \beta$ and $\overline{v}i_{(2)}(1) > \gamma \cdot \overline{v}1(1)$.
3. One-Dominant: input: $\overline{v}i_{(1)}(1) \geq \beta$ and $\overline{v}i_{(2)}(1) \leq \gamma \cdot \overline{v}1(1)$.

In Section 4.2, we present Mechanism SP-One-Dominant which yields a constant approximation for a One-Dominant input (see Lemma 11). In Section 4.3, we present Mechanism RS-Many-Dominant which yields a constant approximation for a Many-Dominant input (see Lemma 14). In Section 4.4, we present Mechanism RS-No-Dominant which yields a constant approximation to the optimal liquid welfare for a No-Dominant input (see Lemma 19). Finally, in Section 4.5, we analyze mechanism LW and conclude it obtains a constant factor approximation to the optimal liquid welfare.

4.2 Mechanism SP-One-Dominant

Mechanism SP-One-Dominant (see Fig. 5) achieves a good approximation in the case of only one dominant agent.

Lemma 10 *Mechanism* SP-One-Dominant *is truthful and feasible.*

Proof It is clear that the mechanism is feasible, since it never allocates more than one item. As for truthfulness, most cases have a similar proof as in mechanism to the proof of truthfulness of Online-RS (Theorem 1). We describe only the cases where the proofs differ. Consider an agent i. We prove that agent i's utility can only decrease by misreporting her type.

Misreporting v_i, b_i Agents in set B cannot change their price or position by misreporting their value or budget. Since given a fixed price and a position the mechanism

SP-One-Dominant (N)

1. $A \leftarrow$ first $\frac{n}{2}$ arriving agents, $B \leftarrow N \setminus A$, $t_{0.5} \leftarrow \hat{a}_{\frac{n}{2}}$.
2. At time $t_{0.5}$:
   (a) Order the agents in set A according to their liquid value such that $\overline{v}_{i_1}(1) \geq \overline{v}_{i_2}(1) \geq \ldots \geq \overline{v}_{i_{\frac{n}{2}}}(1)$.
   (b) $p_1 \leftarrow \frac{1}{\gamma} \cdot \overline{v}_{i_1}(1)$, $p_2 \leftarrow \frac{1}{\gamma} \cdot \overline{v}_{i_2}(1)$.
   (c) Toss a fair coin.
   – If it turns heads (sampling set):
     If $\overline{v}_{i_1}(1) \geq p_2$ and $d_{i_1} \geq t_{0.5}$ then allocate the entire item to agent i_1 at a price p_2.
   – If the coin turns tails (performance set):
     Wait for the first agent $j \in B$ such that $\overline{v}_j(1) \geq p_1$ (if exists) and allocate to her the entire item at a price p_1.

Fig. 5 Mechanism SP-One-Dominant

allocates the best possible allocation, agents do not have the incentive to misreport their bids.

For an agent $i \in A$, if $\overline{v}i(1) \geq \frac{1}{\gamma} \cdot \overline{v}i_2(1)$ then the agent has a probability of $\frac{1}{2}$ to receive the item at a price of $\frac{1}{\gamma} \cdot \overline{v}i_2(1)$. Agent i cannot affect the price or the probability to win the item by misreporting $\overline{v}i(1)$. If $\overline{v}i(1) < \frac{1}{\gamma} \cdot \overline{v}i_2(1)$, then in order to get the item, agent i has to exceed her budget or value.

Misreporting a_i We analyze the following case (analogous to case 4 of misreporting a_i in Theorem 1).

The mechanism places agent i in set A according to her true arrival time and in set B according to her reported arrival time. Let $\overline{v}i_1(1) \geq \overline{v}i_2(1) \geq \ldots \geq \overline{v}i_{\frac{n}{2}}(1)$ be the agents in set A ordered according to their liquid value, assuming that agent i reported her real arrival time. If $\overline{v}i(1) < \frac{1}{\gamma} \cdot \overline{v}i_2(2)$, agent i will not receive the good whether she is in set A or in set B. If $\overline{v}i(1) \geq \frac{1}{\gamma} \cdot \overline{v}i_2(1)$, then the agent receives the item at price $\frac{1}{\gamma} \cdot \overline{v}i_2(1)$ with probability $1/2$ whether she is in set A or in set B (note that in this case, if agent i is placed in set B then agent i_2 becomes the agent with the highest liquid value in set A). □

Lemma 11 *For a One-Dominant input, the expected liquid welfare in mechanism SP-One-Dominant is at least* $\frac{1}{1600}$.

Proof We inspect the case where agent $i_{(2)} \in A$ and agent $i_{(1)} \in B$. In this case, with probability $1/2$, the item is allocated to agent $i_{(1)}$ and the liquid welfare is at least $1/200$. Since this case happens with probability $1/4$, we get that the expected liquid welfare is at least $\frac{1}{1600}$. □

4.3 Mechanism RS-Many-Dominant

Mechanism RS-Many-Dominant described below achieves a good approximation in the case of at least two dominant agents. We first describe mechanism RS-Many-Dominant′ which is an instance of mechanism Online-RS and sells $m = 4$ items. It uses the pricing function $\gamma \cdot OPT$ and an allocation function Allocate-TIOLI (see Fig. 6). The allocation function orders the agents according to the random renaming, and sells the entire good to the first agent with a liquid value at least p. Mechanism RS-Many-Dominant$infeas$ is not feasible since we have a single divisible item for sale, while RS-Many-Dominant$infeas$ sells an item to each of

Allocate-TIOLI $(S, p, k = 1, \ell)$

- Initialize $x_i \leftarrow 0$ for every agent $i \in S$.
- If $\ell = 1$:
  1. For each $i \in S$ approached in the order of the random renaming
     - If $\overline{v}_{(i)}(1) \geq p$ then $x_{(i)} \leftarrow 1$ and exit the loop.
- Return x.

Fig. 6 Allocation function Allocate-TIOLI the sets A_1, A_2, B_1 and B_2.[7] To overcome this problem, we devise mechanism RS-Many-Dominant which chooses one of the sets uniformly at random and sells the item in this set. Thus it is feasible and obtains $\frac{1}{4}$ of the liquid welfare guaranteed by RS-Many-Dominant$infeas$.

Lemma 12 *Mechanism RS-Many-Dominant is truthful and feasible.*

Proof We first analyze mechanism RS-Many-Dominant'. By Theorem 1, it is sufficient to verify the following properties for Allocate-TIOLI:

- IR and IC: The agents are ordered randomly, and each agent is offered a take-it-or-leave-it price. Therefore, it is trivially IR and IC.
- Cap Monotonicity: ℓ is either 0 or 1. If $\ell = 0$ then no agent is allocated and her utility is 0; otherwise, her utility is non-negative.
- Feasibility: The allocation function is feasible since only one item is allocated and only if there exists an actual item.

Now we analyze mechanism RS-Many-Dominant. Inspecting the proof of Theorem 1, one can notice that if the agent gets the allocation with probability $1/4$, her expected utility is still maximized when reporting her true value. Hence, mechanism RS-Many-Dominant is truthful and feasible. □

The following lemma shows that whenever we have a dominant agent in the performance set, then the item is sold with a high probability to an agent in this set. Let X_{B_1} and X_{B_2} be the events that at least one of the agents $i_{(1)}$ and $i_{(2)}$ is in set B_1 and B_2, respectively.

Lemma 13 *Whenever X_{B_1} occurs, RS-Many-Dominant$infeas$ sells an item to an agent in B_1 with probability at least $1/2$; Similarly, whenever X_{B_2} occurs, RS-Many-Dominant$infeas$ sells an item to an agent in B_2 with probability at least $1/2$.*

Proof We prove the claim for the case where event X_{B_1} occurs (the proof for the case that X_{B_2} happens is symmetric). Let p be the price offered to set B_1, let $N_1(p) = \{i' \in N_1 : \overline{v}i'(1) \geq p\}$ be the set of agents that might get allocated in Allocate-TIOLI if no agent is allocated prior to them, and let $S = B_1 \cap N_1(p)$ and $\bar{S} = (N_1 \setminus B_1) \cap N_1(p)$. Let $r : N \to \{1, \ldots, n\}$ be a function that for every agent i returns its order in the random renaming (first step of Online-RS). Notice that the item is sold to an agent in B_1 whenever $N_1(p)$ is non-empty and

$$\min_{i' \in S} r(i') < \min_{i' \in \bar{S}} r(i'). \tag{3}$$

That is, a potential agent in B_1 is offered the item before other potential agents in $N_1 \setminus B_1$. Let $\ell \in \{i_{(1)}, i_{(2)}\}$ be the agent in B_1. If $N_1(p) - \ell$ is empty, (3) holds

[7] Selling an entire item to each set is needed, since selling a fraction of item to a dominant agent does not give any guarantee on the liquid welfare.

with probability 1. Otherwise, let E be the event $r(\ell) > \min_{i' \in (N_1(p)-\ell)} r(i')$. The probability that (3) holds is:

$$\Pr\left[\min_{i' \in (S-\ell)} r(i') < \min_{i' \in \bar{S}} r(i') \mid E\right] \cdot \Pr[E] + \Pr[\neg E] \leq$$
$$\Pr\left[\min_{i' \in (S-\ell)} r(i') < \min_{i' \in \bar{S}} r(i') \mid E\right].$$

Since the agents in $B_1 - \ell$ and $N_1 \setminus (B_1 - \ell)$ are a uniformly sampled partition of $N_1 - \ell$, by symmetry,

$$\Pr\left[\min_{i' \in (S-\ell)} r(i') < \min_{i' \in \bar{S}} r(i') \mid E\right] = \frac{1}{2}.$$

Therefore, (3) holds with probability greater than $1/2$. □

Lemma 14 *For a* Many-Dominant *input, the expected liquid welfare in mechanism* RS-Many-Dominant *is at least* $\frac{1}{6480}$.

Proof We first analyze the expected revenue obtained by mechanism RS-Many-Dominant$infeas$. We analyze the case where one of the following occurs:

– Agent $i_{(1)} \in B_1$ and agent $i_{(2)} \in (A_2 - j)$; or
– Agent $i_{(2)} \in B_1$ and agent $i_{(1)} \in (A_2 - j)$

In these cases, $\overline{v}i_{(2)}(1) > \gamma \cdot \overline{v}i_{(1)}(1) \geq \mathcal{P}(A_2 - j) \geq \gamma \cdot \overline{v}i_{(2)}(1) \geq \gamma^2 \cdot \beta \geq \frac{200}{201^2} > \frac{1}{202.5}$. By Lemma 13 the item is sold with probability of at least $1/2$. Therefore, we get that the expected revenue in this case is at least $\frac{1}{405}$. Similarly, consider the case where one of the following occurs:

– $i_{(1)} \in B_1$ and $i_{(2)} \in (A_2 - j)$; or
– $i_{(2)} \in B_1$ and $i_{(1)} \in (A_2 - j)$.

In this case, the expected revenue is at least $\frac{1}{405}$ as well.

Since the two events are disjoint, and the probability of each one of them is $1/8$, we get that the expected revenue obtained by mechanism RS-Many-Dominant' is at least $\frac{1}{1620}$. Since the expected revenue obtained by mechanism RS-Many-Dominant is a quarter of the expected revenue obtained by RS-Many-Dominant', we get that the revenue in RS-Many-Dominant is at least $\frac{1}{6480}$, and by Lemma 8, we get that the expected liquid welfare is at least $\frac{1}{6480}$ as well. □

4.4 Mechanism RS-No-Dominant

Mechanism RS-No-Dominant is an instance of mechanism Online-RS that achieves a good approximation in the case of no dominant agent. It uses the pricing function $P(S) = \frac{OPT(S)}{4}$, where given a set $S \subseteq N$ of agents, $OPT(S) = \max_x \sum_{i \in S} \overline{v}i(x_i)$ (the optimal liquid welfare of set S). The allocation function is Allocate-Demand (see Fig. 7).

Allocate-Demand (S, p, k, ℓ)

1. Initialize $k' \leftarrow k$, $x_i \leftarrow 0$ for every agent $i \in S$.
2. For each $i \in S$ approached in the order of the random renaming
   - $x_i \leftarrow D_i(p, k')$
   - $k' \leftarrow k' - x_i$
3. For every i, set $x_i \leftarrow D_i(p, \min\{x_{(i)}, \ell\})$.
4. Return x.

Fig. 7 The allocation function for mechanism RS-No-Dominant

We define the demand of agent i for a price p and fraction k of the item:

$$D_i(p, k) = \operatorname{argmax}_{x \leq k}\{u_i(x)\}.$$

In case of multiple solutions for $D_i(p, k)$, we choose the largest x. The allocation function Allocate-Demand orders the agents according to the random renaming, and sells the fraction of the item to the agents. Then the allocation of each agent is bounded according to the cap per agent.

Lemma 15 *Mechanism RS-No-Dominant is truthful and feasible.*

Proof Feasibility holds trivially. By Theorem 1, it is sufficient to show that Allocate-Demand satisfies the following properties:

- IR, IC and Cap Monotonicity: Allocate-Demand allocates the utility maximizing fraction of the good to an agent under the constraint that the allocation does not exceed the total number of items and the cap per agent; Therefore, the IR and IC properties hold.
- Cap Monotonicity: Note that, when an agent receives her utility maximizing fraction she does not exceeds her budget. Hence, the utility of an agent i from receiving a fraction x of the item at a total price p is exactly $v_i(x_i) - p$. Recall that Allocate-Demand orders the agents according to the random permutation sampled by mechanism Online-RS, and sells the contrived items to the agents. Then the allocation of each agent is bounded according to the actual number of items. Hence, it is clear that the allocation is monotone in the number of actual items.

□

Note that we analyze Mechanism RS-No-Dominant only for a No-Dominant input. The next lemma shows that in this case, with a high probability sets A_1 and $(A_2 - j) \cup B_2$ are relatively large. This is crucial since we collect revenue from set B_2 based on the number of items allocated to the contrived set A_1, and the price calculated according to set $A_2 - j$.

Lemma 16 *For a No-Dominant input,*

$$\Pr\left[OPT(A_1) > \frac{3}{16} \wedge OPT\left((A_2 - j) \cup B_2\right) > \frac{3}{8}\right] \geq 0.74.$$

Proof To prove this lemma, we use the following variant of the multiplicative Chernoff bound. Let $\chi_1, \ldots, \chi_k$ be k independent random variables such that $\chi_i \in [0, B]$ for some $B > 0$. Let $\mathrm{E}[\chi_i] = \mu_i$ for every i, and let $\mu = \sum_i \mu_i$. For every $\delta \in (0, 1]$, the following holds:

$$\Pr\left[\sum_{i=1}^{n} \chi_i < (1-\delta)\mu\right] < \exp\left(-\frac{\delta^2 \cdot \mu}{2B}\right). \quad (4)$$

Let x^* be an optimal allocation vector for N'. For every $i \in N'$, let X_i be a random variable defined by:

$$X_i = \begin{cases} \overline{v}i(x_i^*) & i \in (A_2 - j) \cup B_2 \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

and let Y_i be a random variable defined by:

$$Y_i = \begin{cases} \overline{v}i(x_i^*) & i \in A_1 \\ 0 & \text{otherwise} \end{cases}. \quad (6)$$

By Observation 1, we have that $\mathrm{E}\left[\sum_i X_i\right] = \frac{1}{2}$ and $\mathrm{E}\left[\sum_i Y_i\right] = \frac{1}{4}$. Using (4) and a union bound, we get that

$$\Pr\left[\sum_i X_i \geq \frac{3}{8} \wedge \sum_i Y_i \geq \frac{3}{16}\right] \geq 0.74.$$

Since $\{x_i^*\}_{i\in(A_2-j)\cup B_2}$ is some feasible allocation for the agents in $(A_2 - j) \cup B_2$, we have that $\sum_i X_i \leq OPT((A_2-j)\cup B_2)$. Similarly, $\sum_i Y_i \leq OPT(A_1)$. Combining this with the above equation, we get the desired result. □

The following definition and claims are based on the ones presented in [28]. We define:

$$\overline{D}_i(p) = \operatorname{argmax}_{x\leq 1}\{\overline{v}i(x) - p \cdot x\}.$$

In case of multiple solutions for $\overline{D}_i(p)$ we choose the largest one. $\overline{D}_i(p)$ is useful in our analysis as it gives a lower bound on the demand of agent i when there is a sufficiently large fraction of the item.

Lemma 17 [28]

- *For any $p > 0$ and $x < \overline{D}_i(p)$, $\overline{D}_i(p) \leq \frac{b_i}{p}$ and $v_i(x) < b_i$.*
- *For every $p > 0$ and for every $k \geq \frac{b_i}{p}$, $D_i(p, k) \geq \overline{D}_i(p)$.*

Mechanism RS-No-Dominant uses an impersonation technique. In Section 3, we showed that we lose only a constant fraction of the revenue due to this technique. Therefore, we were able to analyze the performance of the mechanism without the impersonation. In this section, agent's valuation functions are general (monotonically non-decreasing). Hence, decoupling the proof is much more complicate, and we use the following lemma which calculates the performance of a mechanism with the impersonation.

Mechanism Impersonation-Allocate-Demand(S, p, k) (see Fig. 8) takes a set of agents S, and partitions them into two sets S_1 and S_2 uniformly at random. Then it

Impersonation-Allocate-Demand (S, p, k)

1. Initialize $\ell \leftarrow k$.
2. For each agent $i \in S$: With probability of $\frac{1}{2}$ place i in set S_1 and with probability of $\frac{1}{2}$ in set S_2.
3. Let $\pi_{S_1} : \{1, 2, \ldots, |S_1|\} \rightarrow S_1$ be a random permutation of the agents of S_1.
4. For $i = 1, 2, \ldots, |S_1|$:
   (a) Let $j \leftarrow \pi_{S_1}(i)$, and let x be the allocation vector returned by Allocate-Demand $(S_2 \cup \{j\}, p, k, \ell)$.
   (b) Allocate to agent j x_j items at a price of $x_j \cdot p$.
   (c) Update $\ell \leftarrow \ell - x_j$.

Fig. 8 A mechanism defined solely for the analysis of the liquid welfare maximization mechanism

sells a fraction of the good k to agents in set S_1 in a random order at a price p for the entire good. The allocation to each agent is calculated as if the agent were in a random permutation in set S_2. Note that mechanism Impersonation-Allocate-Demand(S, p, k) uses the impersonating technique. For a set S of bidders, we define $\overline{v}S(1) = \max_{i \in S} \overline{v}i(1)$.

Lemma 18 *Let $p \leq 1$ and let $0 < OPT(S) < 1$. The expected liquid welfare out the agents in S_1 in procedure* Impersonation-Allocate-Demand(S, p, k) *is at least*

$$\min\left(\frac{OPT(S) - p}{2}, \frac{k \cdot p - \overline{v}S(1)}{2}\right).$$

Proof We begin by analyzing the performance of Allocate-Demand(S, p, k, ∞). Let α_i be the amount of good available for agent i in Allocate-Demand(S, p, k, ∞), let $\overline{x}$ and $\widetilde{x}$ be the allocation vectors returns by Allocate-Demand(S, p, k, ∞) and Impersonation-Allocate-Demand(S, p, k), respectively. Note that for every $i \in S$ and k, $D_i(p, k) \leq \frac{\overline{v}S(1)}{p}$, otherwise, agent i would exceed her budget. We distinguish between the following 2 cases:

1. $\alpha_i \geq \frac{\overline{v}S(1)}{p}$ for every agent i: In this case, every agent gets her utility maximizing fraction of the good, and at most a k-fraction is sold. The liquid welfare achieved by Allocate-Demand(S, p, k, ∞) is at least:

$$\begin{aligned}
\sum_{i \in S} \overline{v}i(\overline{x}_i) &\geq \sum_{i \in S} \overline{v}i\left(\overline{D}_i(p)\right) \\
&\geq \sum_{i \in S} \overline{v}i\left(\overline{D}_i(p)\right) - \overline{D}_i(p) \cdot p \\
&\geq \sum_{i \in S} \overline{v}i(x_i^*) - x_i^* \cdot p \\
&\geq OPT(S) - p,
\end{aligned}$$

where the first inequality is due to the second property of Lemma 17, and the third inequality follows from the definition of $\overline{D}_i(p)$.

Impersonation-Allocate-Demand places every agent from S in set S_1 with independent probability of $\frac{1}{2}$. Therefore, we have that:

$$\mathrm{E}_{\text{Allocate-Demand}(S,p,k,\infty)}\left[\sum_{i\in S_1}\overline{v}i(\overline{x}_i)\right]\geq\frac{OPT(S)-p}{2}.$$

Let x be the output vector of Allocate-Demand$(S_2 \cup \{j\}, p, k, \ell)$ in Impersonation-Allocate-Demand. We show by induction on the ordering of the agents in S_1 (by π_{S_1}) that $x_j = \overline{x}_j$ for every agent j. For the first agent j in the ordering, since $\ell = k$, it is clear that in Allocate-Demand$(S_2 \cup \{j\}, p, k, \infty)$ every agent in $S_2 \cup \{j\}$ got her utility maximizing share, and therefore $x_j = \overline{x}_j$. Let us assume the claim is true for all agents in S_1 that are ordered before agent j, and let $S_1^{<j}$ and $S_1^{>j}$ denote the respective sets of agents in S_1 before and after agent j in the ordering. We have that:

$$\begin{aligned}\ell = k-\sum_{i\in S_1^{<j}}\overline{x}_i &\geq k-\sum_{i\in S_1^{<j}}\overline{x}_i-\sum_{i\in S_1^{>j}}\overline{x}_i\\ &\geq \sum_{i\in S}\overline{x}_i-\sum_{i\in S_1^{<j}}\overline{x}_i-\sum_{i\in S_1^{>j}}\overline{x}_i=\sum_{i\in(S_2\cup\{j\})}\overline{x}_i.\end{aligned}$$

We get that every agent in $S_2 \cup \{j\}$ can get her utility maximizing share, and the overall fraction allocated is at most ℓ. Therefore, it is clear that $x_j = \overline{x}_j$ for this agent j, and

$$\mathrm{E}_{\text{Impersonation-Allocate-Demand}(S,p,k,z)}\left[\sum_{i\in S_1}\overline{v}i(\widetilde{x}_i)\right]\geq\frac{OPT(S)-p}{2}.$$

2. There exists an agent i such that $\alpha_i > \frac{\overline{v}S(1)}{p}$: let j denote the first such agent. We perform the following transformation on the allocation returned by Allocate-Demand(S, p, k, ∞):

$$y_i=\begin{cases}\overline{x}_i\cdot\frac{k-\frac{\overline{v}S(1)}{p}}{k-\alpha_j} & \pi(i)<\pi(j)\\ 0 & \text{otherwise}\end{cases}.$$

We have that:

$$\sum_i y_i=\sum_{i:\pi(i)<\pi(j)}y_i=\sum_{i:\pi(i)<\pi(j)}\overline{x}_i\cdot\frac{k-\frac{\overline{v}S(1)}{p}}{k-\alpha_j}=\left(k-\alpha_j\right)\cdot\frac{k-\frac{\overline{v}S(1)}{p}}{k-\alpha_j}=k-\frac{\overline{v}S(1)}{p}.$$

Since every agent in S is placed with independent probability $\frac{1}{2}$ in S_1 in Impersonation-Allocate-Demand, we have that $\mathrm{E}\left[\sum_{i\in S_1} y_i\right] = \frac{k-\frac{\overline{v}S(1)}{p}}{2}$. We now show that $\sum_{i\in S_1}\widetilde{x}_i \geq \sum_{i\in S_1} y_i$. Let $\widetilde{k}_i$ be the amount of good left for agent $i \in S_1$ when allocated in Impersonation-Allocate-Demand(S, p, k). If there is some agent i for which $\widetilde{k}_i < \frac{\overline{v}S(1)}{p}$, then $\sum_{i\in S_1}\widetilde{x}_i > k - \frac{\overline{v}S(1)}{p} = \sum_i y_i \geq \sum_{i\in S_1} y_i$. Otherwise, for every $i \in S_1$, $\widetilde{x}_i = \overline{x}_i$, and since $\overline{x}_i \geq y_i$, the claim follows. We get that the expected revenue in Impersonation-Allocate-Demand(S, p, k) is at least $p \cdot \mathrm{E}\left[\sum_{i \in S_1} \widetilde{x}_i\right] \geq p \cdot \frac{k - \frac{\overline{v}S(1)}{p}}{2} = \frac{pk - \overline{v}S(1)}{2}$. By Lemma 8, we get that the expected liquid welfare is at least $\frac{pk-\overline{v}S(1)}{2}$ as well.

It follows that for every random r_i's, either the expected liquid welfare is at least $\frac{OPT(S)-p}{2}$, or at least $\frac{kp-\overline{v}S(1)}{2}$; Therefore, the expected liquid welfare is at least the minimum between them. □

We conclude that:

Lemma 19 *For a No-Dominant input, mechanism RS-No-Dominant achieves an expected liquid welfare of at least* 1/400.

Proof We analyze the case where $OPT(A_1) > \frac{3}{16}$ and $OPT\left((A_2 - j) \cup B_2\right) > \frac{3}{8}$. In this case, $\frac{3}{64} \leq P(A_1) \leq \frac{1}{4}$. According to Lemma 18, the expected liquid welfare of the agents in A_2 in procedure Impersonation-Allocate-Demand $\left((A_2 - j) \cup B_2, P(A_1), \frac{1}{4}, \frac{1}{2}\right)$ in this case is at least:

$$\min\left(\frac{OPT\left((A_2 - j) \cup B_2\right) - \mathcal{P}(A_1)}{2}, \frac{\mathcal{P}(A_1)/4 - 1/200}{2}\right) \geq$$
$$\min\left(\frac{3/8 - 1/4}{2}, \frac{3/256 - 1/200}{2}\right) \geq 1/298.$$

Notice that in RS-No-Dominant(N, m), the expected liquid welfare of the agents in B_2 is the same as in Impersonation-Allocate-Demand $\left((A_2 - j) \cup B_2, \mathcal{P}(A_1), \frac{1}{4}, \frac{1}{2}\right)$, since every agent in $(A_2 - j) \cup B_2$ is either in $A_2 - j$ or in B_2 with probability 1/2 (Observation 1), and the agents in B_2 arrive in random order. By Lemma 16, we get that the expected liquid welfare is at least $0.74/298 > 1/400$ □

4.5 Approximation Ratio of LW

Given all the above components, we can devise a constant approximation mechanism to the optimal liquid welfare in our setting.

Theorem 4 *Mechanism LW is a truthful and feasible mechanism that achieves an expected liquid welfare of* $OPT/10412$.

Proof Truthfulness and feasibility follow since the mechanism is a probabilistic combination of 3 truthful and feasible mechanisms. It remains to establish the approximation guarantee. According to Lemma 9, $OPT' \geq 0.9 \cdot OPT$ with probability of at least 91/100. Therefore, $OPT \leq \frac{100}{91} \cdot \frac{10}{9} \mathrm{E}[OPT'] = \frac{1000}{819}$. Combining the facts that mechanism LW invokes each of the mechanisms RS-No-Dominant, RS-Many-Dominant and SP-One-Dominant with a constant probability, in each of the three possible scenarios, one of them achieves a constant fraction of OPT'. Setting

$\lambda_1 = 0.19$, $\lambda_2 = 0.76$ and $\lambda_3 = 0.05$, we get that the expected liquid welfare of our mechanism is at least

$$\min\left\{\frac{\lambda_1}{1600}, \frac{\lambda_2}{6480}, \frac{\lambda_3}{400}\right\} \geq \frac{1}{8527} \geq \frac{1}{8527} \cdot \frac{819}{1000} OPT \geq OPT/10412.$$

□

Appendix A: Proofs of the Theorems from Borgs et al.

Recall that $P_k(S)$ is the optimal price for selling k items to a set S of agents and that $\overline{\epsilon} = \epsilon(S,k)$. Let $r_S^\infty(p)$ denote the revenue achieved by selling an unlimited number of items (digital goods) to set S of agents at price per item p. In order to show the proofs of the lemmas in Section 3.3, we need the following lemma from [8]:

Lemma 20 [8] *For any* $\delta > 0$ *and for any* $p \geq P_k(S)$, *the probability that* $\left|r_{S_1}^\infty(p) - r_{S_2}^\infty(p)\right| \leq \delta \cdot OPT(S,k)$ *is at least* $1 - 2e^{-\delta^2/4\overline{\epsilon}}$.

We now turn to prove the lemmas stated in Section 3.3.

Proof of Lemma 3 Since the revenue obtained from selling items at a given price per item is either bounded by the agents' budgets or by the number of items for sale, we get that $OPT(S,k) = \min\left\{P_k(S) \cdot k, r_S^\infty(P_k(S))\right\}$. Therefore, we get

$$OPT(S,k) \leq P_k(S) \cdot k, \tag{7}$$

and

$$OPT(S,k) \leq r_S^\infty(P_k(S)). \tag{8}$$

When unrestricted by the number of items, we have that

$$r_{S_1}^\infty(P_k(S)) + r_{S_2}^\infty(P_k(S)) = r_S^\infty(P_k(S)) \geq OPT(S,k). \tag{9}$$

Whenever $\left|r_{S_1}^\infty(P_k(S)) - r_{S_2}^\infty(P_k(S))\right| \leq \delta \cdot OPT(S,k)$, then

$$r_{S_2}^\infty(P_k(S)) \leq r_{S_1}^\infty(P_k(S)) + \delta \cdot OPT(S,k), \tag{10}$$

and

$$r_{S_1}^\infty(P_k(S)) \leq r_{S_2}^\infty(P_k(S)) + \delta \cdot OPT(S,k). \tag{11}$$

Combining (10) with (9) yields $r_{S_1}^\infty(P_k(S)) \geq \frac{1-\delta}{2} OPT(S,k)$. Similarly, combining (11) with (9) yields $r_{S_2}^\infty(P_k(S)) \geq \frac{1-\delta}{2} OPT(S,k)$. By (7) we clearly have $P_k(S) \cdot \frac{k}{2} \geq OPT(S,k)/2$. Since both $r_{S_1} = \min\left\{r_{S_1}^\infty(P_k(S)), P_k(S) \cdot \frac{k}{2}\right\}$ and $r_{S_2} = \min\left\{r_{S_2}^\infty(P_k(S)), P_k(S) \cdot \frac{k}{2}\right\}$ we get that whenever $\left|r_{S_1}^\infty(P_k(S)) - r_{S_2}^\infty(P_k(S))\right| \leq \delta \cdot OPT(S,k)$ it holds that $r_{S_1} \geq \frac{1-\delta}{2} OPT(S,k)$ and $r_{S_2} \geq \frac{1-\delta}{2} OPT(S,k)$. The proof is now concluded by Lemma 20. □

In order to prove Lemma 4, we also need the following lemma from [8].[8]

Lemma 21 [8] *For all* $\delta, p > 0$ *both* $r^{\infty}_{S_1}(p) \geq \min\left\{r^{\infty}_{S_2}(p) - \delta \cdot OPT(S,k), \frac{1-\delta}{2} \cdot OPT(S,k)\right\}$ *and* $r^{\infty}_{S_2}(p) \geq \min\left\{r^{\infty}_{S_1}(p) - \delta \cdot OPT(S,k), \frac{1-\delta}{2} \cdot OPT(S,k)\right\}$ *with probability at least* $1 - 2e^{-\delta^2/4\epsilon}$.

Proof Assuming $\left|r^{\infty}_{S_1}(P_k(S)) - r^{\infty}_{S_2}(P_k(S))\right| \leq \delta \cdot OPT(S,k)$ we get that both $r^{\infty}_{S_2}(p) \geq r^{\infty}_{S_1}(p) - \delta \cdot OPT(S,k)$ and $r^{\infty}_{S_1}(p) \geq r^{\infty}_{S_2}(p) - \delta \cdot OPT(S,k)$ with probability greater than $1 - 2e^{-\delta^2/4\epsilon}$ for $p \geq P_k(S)$.

Recall that in the proof of Lemma 3, we showed that both $r^{\infty}_{S_1}(P_k(S)) \geq \frac{1-\delta}{2} OPT(S,k)$ and $r^{\infty}_{S_2}(P_k(S)) \geq \frac{1-\delta}{2} OPT(S,k)$ whenever $\left|r^{\infty}_{S_1}(P_k(S)) - r^{\infty}_{S_2}(P_k(S))\right| \leq \delta \cdot OPT(S,k)$. Since with an unlimited supply of items, $r^{\infty}_{S}(p_1) \geq r^{\infty}_{S}(p_2)$ whenever $p_1 < p_2$ (only more agents exhaust their budget), we get that for every $p \leq P_k(S)$ both $r^{\infty}_{S_1}(p) \geq \frac{1-\delta}{2} OPT(S,k)$ and $r^{\infty}_{S_2}(p) \geq \frac{1-\delta}{2} OPT(S,k)$.

To conclude, we recall that according to Lemma 20, $\left|r^{\infty}_{S_1}(P_k(S)) - r^{\infty}_{S_2}(P_k(S))\right| \leq \delta \cdot OPT(S,k)$ occurs with probability at least $1 - 2e^{-\delta^2/4\epsilon}$. □

We now proceed to show the proof of Lemma 4.

Proof of Lemma 4 We first assume that $\left|r^{\infty}_{S_1}(P_k(S)) - r^{\infty}_{S_2}(P_k(S))\right| \leq \delta \cdot OPT(S,k)$. By definition, $OPT(S_1, k/2) \geq r_{S_1}$. In this case, as shown in the proof of Lemma 3, $OPT(S_1, k/2) \geq \frac{1-\delta}{2} OPT(S,k)$. Therefore, it is clear that both

$$P_{k/2}(S_1) \cdot \frac{k}{2} \geq \frac{1-\delta}{2} OPT(S,k), \tag{12}$$

and

$$r^{\infty}_{S_1}(P_{k/2}(S_1)) \geq \frac{1-\delta}{2} OPT(S,k). \tag{13}$$

Combining the last inequality with Lemma 21, we get that

$$r^{\infty}_{S_2}(P_{k/2}(S_1)) \geq \frac{1-3\delta}{2} OPT(S,k). \tag{14}$$

Since in Offline-Rev-Maximization either the agents of S_2 exhaust their budget, or $k/2$ items are sold, we get that the revenue of mechanism Offline-Rev-Maximization from agents in set S_2 is not smaller than: $\min\left\{r^{\infty}_{S_2}(P_{k/2}(S_1)), P_{k/2}(S_1) \cdot \frac{k}{2}\right\} \geq \frac{1-3\delta}{2} OPT(S,k)$. To conclude, we note that according to Lemma 20,

[8]This lemma is stated quite differently in [8].

$\left| r_{S_1}^{\infty}(P_k(S)) - r_{S_2}^{\infty}(P_k(S)) \right| \leq \delta \cdot OPT(S,k)$ occurs with probability at least $1 - 2e^{-\delta^2/4\epsilon}$. □

Appendix B: Tie Breaking

In this section, we show how to perform tie breaking when an agents is allowed to report the same arrival time as another agent, but cannot report an earlier arrival time than her real arrival time. The tie breaking rule we present is as follows. Whenever an agent i arrives at time a_i, the mechanism chooses a uniformly at random value $\tilde{a}_i \sim [0, 1]$ and sets agent i's arrival time to be $\langle a_i, \tilde{a}_i \rangle$. The mechanism orders the agents according to the following rule: Agent i precedes agent j if $a_i < a_j$ or if $a_i = a_j$ and $\tilde{a}_i < \tilde{a}_j$.

We change mechanisms Online-RS and SP-One-Dominant according to the above rule. We claim that these mechanisms are truthful. The only problematic case is when an agent reports an earlier arrival time such that she is placed in set A instead of being placed in set B (the proof of all the other cases is similar to the proof of Theorem 1). Since an agent cannot affect the random value assigned to her, and since we specifically prevent an agent from reporting an earlier arrival time, this is impossible.

References

1. Abrams, Z.: Revenue maximization when bidders have budgets. In: SODA, pp. 1074–1082. ACM (2006)
2. Awerbuch, B., Azar, Y., Meyerson, A.: Reducing truth-telling online mechanisms to online optimization. In: STOC, pp. 503–510. ACM (2003)
3. Azar, Y., Feldman, M., Gravin, N., Roytman, A.: Liquid price of anarchy. In: Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L'aquila, Italy, September 12–14, 2017, Proceedings, pp. 3–15 (2017)
4. Babaioff, M., Immorlica, N., Kleinberg, R.: Matroids, secretary problems, and online mechanisms. In: SODA, pp. 434–443. SIAM (2007)
5. Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M.: A simple and approximately optimal mechanism for an additive buyer. In: FOCS (2014)
6. Bar-Yossef, Z., Hildrum, K., Wu, F.: Incentive-compatible online auctions for digital goods. In: SODA, pp. 964–970 (2002). Society for industrial and applied mathematics
7. Blum, A., Hartline, J.D.: Near-optimal online auctions. In: SODA, pp. 1156–1163 (2005). Society for industrial and applied mathematics
8. Borgs, C., Chayes, J., Immorlica, N., Mahdian, M., Saberi, A.: Multi-unit auctions with budget-constrained bidders. In: EC, pp. 44–51. ACM (2005)
9. Caragiannis, I., Voudouris, A.A.: The efficiency of resource allocation mechanisms for budget-constrained users. In: Proceedings of the 2018 ACM Conference On Economics and Computation, Ithaca, NY, USA, June 18–22, 2018, pp. 681–698 (2018)
10. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: STOC (2010)
11. Chawla, S., Malec, D.L., Malekian, A.: Bayesian mechanism design for budget-constrained agents. In: EC, pp. 253–262. ACM (2011)
12. Chawla, S., Malec, D.L., Sivan, B.: The power of randomness in bayesian optimal mechanism design. Games Econom. Behav. **91**, 297–317 (2015)
13. Devanur, N.R., Ha, B.Q., Hartline, J.D.: Prior-free auctions for budgeted agents. In: EC, pp. 287–304. ACM (2013)

14. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits. Games Econom. Behav. **74**(2), 486–503 (2012)
15. Dobzinski, S., Leme, R.P.: Efficiency guarantees in auctions with budgets. In: ICALP, pp. 392–404. Springer (2014)
16. Dughmi, S., Eden, A., Feldman, M., Fiat, A., Leonardi, S.: Lottery pricing equilibria. In: Proceedings of the 2016 ACM Conference On Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016, pp. 401–418 (2016)
17. Feldman, M., Fiat, A., Leonardi, S., Sankowski, P.: Revenue maximizing envy-free multi-unit auctions with budgets. In: EC, pp. 532–549. ACM (2012)
18. Fiat, A., Leonardi, S., Saia, J., Sankowski, P.: Single valued combinatorial auctions with budgets. In: EC, pp. 223–232. ACM (2011)
19. Fotakis, D., Lotidis, K., Podimata, C.: A bridge between liquid and social welfare in combinatorial auctions with submodular bidders. arXiv:1809.01803 (2018)
20. Friedman, E.J., Parkes, D.C.: Pricing Wifi at Starbucks: issues in online mechanism design. In: EC, pp. 240–241. ACM (2003)
21. Goel, G., Mirrokni, V., Leme, R.P.: Clinching auctions with online supply. In: SODA, pp. 605–619. SIAM (2013)
22. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A.: Competitive auctions. Games Econom. Behav. **55**(2), 242–269 (2006)
23. Hajiaghayi, M.T., Kleinberg, R., Parkes, D.C.: Adaptive limited-supply online auctions. In: EC, pp. 71–80. ACM (2004)
24. Hajiaghayi, M.T., Kleinberg, R.D., Mahdian, M., Parkes, D.C.: Online auctions with re-usable goods. In: EC, pp. 165–174. ACM (2005)
25. Kleinberg, R.: A multiple-choice secretary algorithm with applications to online auctions. In: SODA, pp. 630–631 (2005). Society for industrial and applied mathematics
26. Lavi, R., Nisan, N.: Competitive analysis of incentive compatible on-line auctions. In: EC, pp. 233–241. ACM (2000)
27. Lavi, R., Nisan, N.: Online ascending auctions for gradually expiring items. In: SODA, pp. 1146–1155. SIAM (2005)
28. Lu, P., Xiao, T.: Improved efficiency guarantees in auctions with budgets. In: EC 15, Portland, OR, USA, June 15–19, 2015, pp. 397–413 (2015)
29. Lu, P., Xiao, T.: Liquid welfare maximization in auctions with multiple items. In: Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L'aquila, Italy, September 12–14, 2017, Proceedings, pp. 41–52 (2017)
30. Parkes, D.C.: Online mechanisms (2007)
31. Parkes, D.C., Singh, S.P.: An Mdp-based approach to online mechanism design. In: NIPS (2003)