

Truthful Secretaries with Budgets ^{*}

Alon Eden¹, Michal Feldman¹ and Adi Vardi¹

¹Blavatnik School of Computer Science, Tel-Aviv University
alonarden@gmail.com, michal.feldman@cs.tau.ac.il, adivardi@gmail.com

June 30, 2015

Abstract

We study online auction settings in which agents arrive and depart dynamically in a random (secretary) order, and each agent's private type consists of the agent's arrival and departure times, value and budget. We consider multi-unit auctions with additive agents for the allocation of both divisible and indivisible items. For both settings, we devise truthful mechanisms that give a constant approximation with respect to the auctioneer's revenue, under a large market assumption. For divisible items, we devise in addition a truthful mechanism that gives a constant approximation with respect to the liquid welfare — a natural efficiency measure for budgeted settings introduced by Dobzinski and Paes Leme [ICALP'14]. Our techniques provide high-level principles for transforming offline truthful mechanisms into online ones, with or without budget constraints. To the best of our knowledge, this is the first work that addresses the non-trivial challenge of combining online settings with budgeted agents.

1 Introduction

In a typical setting of sales of online ad slots, advertisers arrive at different times, each with her own preferences and budget. The auctioneer (e.g., *cnn.com*) has to make decisions about the allocation of the ad slots to the advertisers and how much to charge them. Such settings give rise to various optimization problems associated with different objective functions. The most common objective functions that have been considered in the literature are (a) social welfare: the sum of the bidders' valuations from their allocations, and (b) the auctioneer's revenue: the total payment the auctioneer derives from the sale.

This scenario inspired the online model studied by Hajiaghayi et al. [20], who used the *random sampling* paradigm to design a truthful mechanism for this setting. The random sampling paradigm was first introduced in the context of selling m identical items in offline settings without budgets [18]. This mechanism divides the agents into two sets, A and B , by tossing a fair coin. Next, it calculates prices p_A and p_B (per item) based on the sets A and B , respectively. Finally, it sells $m/2$ items at price p_A to agents in B , in an arbitrary order, and similarly $m/2$ items at price p_B to agents in A . This mechanism is known to give a constant approximation with respect to the optimal revenue.

Hajiaghayi et al. [20] managed to apply the random sampling paradigm to the online setting, while losing only a constant factor in the approximation of the revenue. Their mechanism is based on the following observations and techniques:

^{*}This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

- A random partition of the n agents into two sets in the offline setting is equivalent to the following procedure in the online setting: Given a set of n agents arriving in a random order, place the first j agents (where j is sampled from the binomial distribution $B(n, 1/2)$) in one set, called the sampling set, and the remaining ones in a second set, called the performance set.
- One can acquire the private types of agents in the sampling set by applying the VCG mechanism on this set.
- Since the random sampling mechanism sells items in an arbitrary order, it is without loss of generality that items are sold to agents in the performance set according to their online order of arrival.

Applying the random sampling paradigm to online settings was a big step toward the applicability of truthful mechanisms. Yet, like many other papers in the mechanism design literature, it ignored the fact that agents are usually restricted by budget constraints. Since budget constraints have a huge effect on agent behavior in real-life scenarios, our goal in this paper is to incorporate budgets into the online setting described above.

The allocation of indivisible items among agents with budget constraints in an offline setting was considered by Borgs et al. [7]. As identified by [7], budget constraints impose many challenges on auction settings. For example, the seminal VCG mechanism [19] fails since the utility function is no longer quasi-linear. Borgs et al. [7] devised a truthful mechanism that adjusted the random sampling mechanism to deal with budget constraints. Their mechanism approximates the revenue within a constant factor under a large market assumption.

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 [12] that no truthful auction among n agents can achieve a better than n approximation to the optimal social welfare.

This problem was recently addressed by Dobzinski and Paes Leme. [12], 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. Dobzinski and Paes Leme considered the allocation of a divisible item to agents with public budgets (and private values), and devised truthful mechanisms that give a constant approximation to the optimal liquid welfare. Lu et al. [26] extended this to agents with *private* budgets and any monotonically non decreasing valuation function. The mechanism devised in [26] is a probabilistic combination of a random sampling mechanism and a modified VCG mechanism.

In the present paper, the main challenge is the combination of two non-trivial challenges; namely the online setting and budgeted agents. From the description above it seems that a lot of the necessary components are already in place. First, [7] and [26] devised random sampling based mechanisms that give a constant approximation with respect to revenue and liquid welfare in budgeted offline settings. Second, [20] presented a technique that adjusts mechanisms based on the random sampling paradigm to online settings. A natural approach for dealing with online budgeted settings would be to combine these techniques to obtain nearly-optimal mechanisms. Not surprisingly, as demonstrated in the sequel, this approach fails badly. Indeed, in order to address the combination of online settings and budget constraints, new ideas are required.

1.1 Our Model

We consider the combination of the online setting without budgets of [20] and the offline settings with budgets of [26, 7]. In our setting, there is a set $N = \{1, \dots, n\}$ of agents. Every agent $i \in N$ has a private type (known only to the agent) that is represented by a tuple (a_i, d_i, v_i, b_i) , where:

- a_i and d_i are the respective arrival and departure time of agent i (clearly, $d_i \geq a_i$). The interval $[a_i, d_i]$ is referred to as agent i 's *time frame*.
- v_i is agent i 's value for receiving a single item within her time frame.
- b_i is the budget of agent i .

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 pairs (value, budget), $([v_1, b_1], [v_2, b_2], \dots, [v_n, b_n])$. A random permutation is used to match these pairs and determine the types of the agents. More formally, let Π_N be the set of all possible permutations on the set N . A permutation $\pi : N \mapsto N$ is sampled uniformly at random from the set Π_N . Then, agent i 's type is given by the tuple $(a_i, d_i, v_{\pi(i)}, b_{\pi(i)})$.

As in standard secretary-like problems, a random order is essential to get any guarantee on the performance of the mechanism. As in [20]¹, 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 times². It should be noted that most of the papers studying online settings with strategic agents consider restrictive inputs or weaker solution concepts (e.g., [24, 15, 20, 21, 25, 28]).

In our model, agents can manipulate any component of their type. In particular, they can report earlier or later arrival and departure times, and arbitrary values and budgets. Upon arrival, agents report their type. Based on these reports, the mechanism determines an allocation and a payment for each agent. The utility of agent i for obtaining an allocation of x_i , within her *real* time frame, for a payment of p_i is:

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

We assume that agents are risk neutral. We consider mechanisms that satisfy the following properties:

- **Feasibility:** The mechanism does not sell more items than are available.
- **Individual Rationality:** An agent's expected utility from an allocation and payment is non-negative.
- **Incentive Compatibility:** An agent's expected utility is maximized when she reports her true type.

Since all the mechanisms devised in this paper are trivially feasible and individually rational, our analysis focuses on incentive compatibility.

1.2 Our Results and Techniques

In Section 2 we establish techniques that constitute high level principles for transforming nearly optimal offline mechanisms into online ones.

In Section 2.1 we present methods for the allocation of divisible and indivisible items to budgeted agents who arrive online. We first describe a simple greedy allocation method for divisible items at a given price

¹Based on personal communication with the authors, this is essentially what is assumed for the correctness of Mechanism RM_k in Section 6 in [20].

²We refer to Appendix D for a description of the tie-breaking rule for this case.

per item p , where agents are ordered randomly. Every agent whose value exceeds p is allocated as many items as she can afford with her budget at a price per item p , limited by the number of remaining items.

In an offline setting, this method can be used for selling indivisible items as well. This is done by translating fractions of items into lottery tickets with the corresponding probabilities, and resolving the market via correlated lotteries upon the "arrival" of the last agent. In an online setting, this method cannot be applied since by the time the last agent arrives, all previous agents may have already departed. To overcome this obstacle, we present an allocation method that resolves the agent's lottery in an online fashion, and loses only a factor 2 in revenue compared to the offline setting.

In Section 2.2 we present a prototype mechanism that transforms an offline mechanism based on the random sampling paradigm into an online one. The prototype mechanism receives as input the set of agents, the number of items to be allocated, a pricing function and an allocation function.

Our starting point is the mechanism devised in [20] for an online setting without budget constraints. This mechanism splits the agents into a sampling set and a performance set. It induces agents in the sampling set to reveal their true values by applying VCG on them, and uses this information to extract revenue from the agents in the performance set. In many settings (such as budgeted settings), VCG cannot be applied. To mitigate this problem, the prototype mechanism applies a random sampling based method on the sampling set.

The above technique ensures that in cases where arrival time is public information, agents in the sampling set do not have incentive to misreport their private information (such as the value and budget in our model). Since we consider models in which arrival time is also private, we need to ensure that agents have no incentive to manipulate their arrival time either. For agents with unit-demand valuations, this issue is handled in [20] by ensuring that the price used in the sampling set never exceeds that in the performance set, and every agent in the sampling set whose value exceeds the price receives it.

This technique, however, does not apply to agents with additive valuations. Indeed, there are cases where the price used in the performance set is higher than the price used in the sampling set, and yet an agent in the sampling set may wish to delay her arrival time, in order to receive a larger number of items (as part of the performance set). In order to ensure an agent cannot gain by delaying her arrival time, we treat an agent in the performance set *as if* she was 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 the same price.

We demonstrate the broad applicability of this prototype by applying it to settings with budgeted agents for the allocation of two different types of items (namely, divisible and indivisible) and for two different objective functions (namely, revenue and liquid welfare maximization).

In Section 3 we devise mechanisms for the allocation of identical items (both divisible and indivisible) that approximate the optimal revenue. This is done by applying the above prototype using appropriately chosen pricing and allocation functions. As in [7], without a large market assumption, no truthful mechanism can achieve a constant approximation to the optimal revenue, even for the case of a single item. This is because an agent whose budget is significantly greater than that of the others must receive the item for any value she reports. We establish the following theorem:

Theorem: There exist truthful mechanisms for the allocation of divisible and indivisible items among budgeted agents with additive valuations in online settings that give a constant approximation to the optimal revenue, under a large market assumption.

In Section 4 we devise a mechanism for the allocation of identical divisible items that approximates the optimal liquid welfare.

Theorem: There exists a truthful mechanism for the allocation of divisible items among budgeted agents

with additive valuations in online settings that gives a constant approximation to the optimal liquid welfare.

Under a large market assumption, a similar mechanism to the one presented in Section 3 approximates the optimal liquid welfare to within a constant factor. For liquid welfare, though, we devise a new truthful mechanism that gives a constant approximation for any market (i.e., without employing large market assumption). As in the offline setting described in [26], our mechanism is a probabilistic combination of a random sampling based mechanism and a modified VCG mechanism. However, both mechanisms should be adjusted to the online setting. The random sampling mechanism is adjusted using the prototype mechanism described in Section 2.2.

The modified VCG mechanism should also be adjusted, since in the original one described in [26], all bids are received simultaneously, and this is inherently infeasible in online settings. Indeed, when the last agent arrives, all other agents may have already departed. Therefore, we propose a new modified VCG as follows: The mechanism places the first $n/2$ agents in set A and the rest in set B . Upon the arrival of the last agent in A , it applies the modified VCG mechanism of [26] on the agents in set A , and determines a threshold value, which is greater than the highest bid in A . With probability $1/2$, all the items are allocated to the winner in A ; otherwise, all the items are allocated to the first agent in B that surpasses the threshold.

1.3 Related Work

Online auctions have been the subject of a vast body of work. Lavi and Nisan [24] introduced an online model in which agents have decreasing marginal valuations over identical indivisible items. While in this 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, 4] and combinatorial auctions [2], have been studied under the assumption of private values and public arrival times.

Similarly to our model and the model presented in [20], online settings in which agents arrive in a random order were also considered in [3, 23, 22].

Friedman and Parkes [15] considered the case where an agent's arrival time is also part of her private information, and thus can also be manipulated. Since then, additional auction settings were studied under the assumption that an agent's arrival and departure times are private [21, 28, 20, 25, 15].

The reader is referred to [27] for an overview of online mechanism design.

Offline mechanisms for budgeted agents have long been considered, both in the context of revenue maximization and of welfare maximization. In the context of revenue maximization, Abrams [1] also considered the model of [7] where indivisible and identical items are sold to additive agents with hard budget constraints. She established the relation between the revenue achieved by an optimal mechanism that is allowed to offer multiple prices and the one of an optimal mechanism who is only allowed to sell at a single price. Maximizing revenue was also considered in Bayesian settings [8] and with the goal of approximating the optimal envy free revenue [10, 13].

Considering the welfare maximization objective, as stated before, it is in general impossible to approximate the sum of agents' welfare using a truthful mechanism. To overcome this obstacle, a more relaxed solution concept of Pareto efficiency was considered in [5, 9, 11, 14, 17, 16], and in [10] they focus on approximating the welfare of the optimal envy free mechanism.

The most related works to our model are [20, 7, 12, 26], which are discussed in great detail throughout the paper.

2 High Level Techniques

2.1 Divisible and Indivisible Allocation

In this section we present methods for the allocation of divisible and indivisible items in online settings with budgets. Let S be a set of agents, p a price and k an amount of items. Recall that every agent i is associated with value v_i and budget b_i . We define **Div_Alloc** to be a procedure for the allocation of divisible items:

Div_Alloc (S, p, k)

1. Initialize $k' \leftarrow k$, $x_i \leftarrow 0$ for every $i \in S$.
2. Let π be a permutation of S chosen uniformly at random from the set of all permutations.
3. For $i = 1, 2, \dots, |S|$:
 - If $v_{\pi(i)} \geq p$:

$$x_{\pi(i)} \leftarrow \min\{b_{\pi(i)}/p, k'\}$$

$$k' \leftarrow k' - x_{\pi(i)}$$
4. Return (x, x) .

Figure 1: A procedure for the allocation of divisible items.

Div_Alloc is a greedy allocation procedure for divisible items at a given price per item p . Note that **Div_Alloc** returns two allocation vectors: one for the allocation and a second one for purposes of charging. While in this procedure the two allocations are identical, we keep this format for the sake of uniformity across divisible and indivisible allocations (see below).

Let $\mathcal{M}$ be a mechanism that returns an allocation vector x , such that x_i is the allocation of agent i . We define $\mathbb{E}_{\mathcal{M}}[x_i]$ to be the expected value of x_i , where the expectation is taken over the coin flips of $\mathcal{M}$. We show the following monotonicity property:

Observation 1. *For every two sets of agents S and T , such that $T \subseteq S$, agent $i \in T$, price p and a (possibly fractional) number of items k , $\mathbb{E}_{\text{Div_Alloc}}[x_i^T] \geq \mathbb{E}_{\text{Div_Alloc}}[x_i^S]$ where x^T, x^S are the allocations returned by **Div_Alloc**(T, p, k) and **Div_Alloc**(S, p, k), respectively.*

Proof. Let B_S, B_T be the expected sum of the budgets of all agents with value at least p that are processed before agent i in a random permutation on sets S and T respectively. It is clear that $B_T \leq B_S$. ■

Budget limitations might prevent the allocation of integral items. In offline settings with indivisible items, a common solution for this problem is to sell correlated lottery tickets to the agents. Due to the online nature of our settings, this cannot be done. Therefore, We propose to realize an agent's allocation within her time frame, via uncorrelated lotteries, as described in the procedure below.

This procedure calculates the number of items allocated to each agent, when agents arrive in a random order. The procedure returns two allocation vectors, x and $\tilde{x}$ where x determines the actual allocation, and $\tilde{x}$ determines the allocation agents are charged for. For every agent i , the allocation x_i must be integral (since items are indivisible), while $\tilde{x}_i$ might be fractional.

Indiv_Alloc (S, p, k)

1. Let π be a random permutation of S .
2. Initialize $k' \leftarrow k$, $x_i \leftarrow 0$ and $\tilde{x}_i \leftarrow 0$ for every $i \in S$.
3. For $i = 1, 2, \dots, |S|$:
 - (a) If $v_{\pi(i)} \geq p$:
 - If $k' \leq \frac{b_{\pi(i)}}{p}$:
Set $x_{\pi(i)} \leftarrow k'$ and $\tilde{x}_{\pi(i)} \leftarrow k'$.
 - Else:
Set $\tilde{x}_{\pi(i)} \leftarrow \frac{b_{\pi(i)}}{p}$.
Set $x_{\pi(i)} \leftarrow \begin{cases} \lceil \frac{b_{\pi(i)}}{p} \rceil & \text{with probability } \frac{b_{\pi(i)}}{p} - \lfloor \frac{b_{\pi(i)}}{p} \rfloor. \\ \lfloor \frac{b_{\pi(i)}}{p} \rfloor & \text{with probability } 1 - \frac{b_{\pi(i)}}{p} + \lfloor \frac{b_{\pi(i)}}{p} \rfloor. \end{cases}$
 - (b) $k' \leftarrow k' - x_{\pi(i)}$.
4. Return $(x, \tilde{x})$.

Figure 2: A procedure for the allocation of indivisible items.

The following lemma shows that by using uncorrelated lotteries we lose at most a factor of 2. The proof is deferred to Appendix B.

Lemma 1. *Let $(x, \tilde{x})$ be the tuple returned by **Indiv_Alloc** (S, p, k) and $(x', \tilde{x}')$ the tuple returned by **Div_Alloc** (S, p, k). Then, $E_{\text{Indiv_Alloc}} [\sum_{i \in S} \tilde{x}_i] \geq \sum_{i \in S} \tilde{x}'_i / 2$.*

2.2 Online Implementation of Random Sampling Based Mechanisms

In this section we present a high level technique for the implementation of random sampling mechanisms in online settings. In the sequel, we show that by using suitable allocation and pricing functions, this technique can be used to devise truthful and nearly optimal mechanisms in online budgeted settings. However, as discussed in Section 1.2, this technique can be used for other settings, with or without budgets.

For every agent i let $\hat{a}_i, \hat{d}_i$ be the reported arrival time and departure time of agent i respectively. Let N be a set of agents and m be an amount of items. Let $\mathcal{P}$ be a function that receives a set of agents as input and returns a price. Let Alloc be a function that receives as input a set of agents, a price per item, and an amount of items, and returns two allocation vectors, x and $\tilde{x}$, where x is the actual allocation and $\tilde{x}$ is used for charging the agents. Note that x_i and $\tilde{x}_i$ might be different (e.g., if Alloc runs lotteries).

Mechanism **RS_Online** works as follows: Toss a fair coin $n - 1$ times, and let j be the number of tails plus one. Place the first j agents in set A and the rest in set B . Upon the arrival of the last agent in set A , apply the random sampling paradigm on the set A . Upon the arrival of each agent in set B , place her into set B_1 or B_2 , by tossing a fair coin. For each agent placed in B_1 , give her allocation as if she was in set A_1 , and charge her a price computed according to set A_2 . Apply an analogous allocation and charging for agents in set B_2 .

RS_Online ($N, m, \mathcal{P}, Alloc$)

Partition:

1. Toss a fair coin $n - 1$ times independently. Let j be the number of tails *plus* 1.
Let $t_0 \leftarrow \hat{a}_j$ (the reported arrival time of the j -th agent).
2. Let A be the set of the first j agents (based on reported arrival time) and let B be $N \setminus A$.

Sampling phase: (Based on set A)

1. At time t_0 , partition the agents in A into two sets A_1 and A_2 by tossing an independent fair coin for each agent. Assume without loss of generality that agent j belongs to A_2 .
2. Let $(x, \tilde{x})$ be the tuple returned by $Alloc(A_1, \mathcal{P}(A_2), \frac{m}{4})$ and $(x', \tilde{x}')$ the tuple returned by $Alloc(A_2, \mathcal{P}(A_1), \frac{m}{4})$.
3. Resolve the allocation and payments of the sampling set:
 - (a) For every agent $i \in A_1$, if $\hat{d}_i \geq t_0$ (i.e., agent i did not depart yet), allocate to agent i x_i items and charge her $\tilde{x}_i \cdot \mathcal{P}(A_2)$.
 - (b) Apply an analogous procedure for agents in A_2 , i.e., for every agent $i \in A_2$, if $\hat{d}_i \geq t_0$, allocate to agent i x'_i items and charge her $\tilde{x}'_i \cdot \mathcal{P}(A_1)$.

Revenue collection phase: (Based on set B)

1. Let $m_1 \leftarrow \frac{m}{4}, m_2 \leftarrow \frac{m}{4}$.
2. For every agent $i \in B$, once she arrives, toss a fair coin.
 - (a) If she turns heads:
 - i. Assign agent i to set B_1 .
 - ii. Let $(x', \tilde{x}')$ be the tuple returned by $Alloc(\{A_1 \cup i\}, \mathcal{P}(A_2 - j), \frac{m}{4})$.
 - iii. Let $x_i \leftarrow \min\{x'_i, m_1\}, \tilde{x}_i \leftarrow \min\{\tilde{x}'_i, m_1\}$.
 - iv. Allocate to agent i x_i items at a price of $\tilde{x}_i \cdot \mathcal{P}(A_2 - j)$.
 - v. Update $m_1 \leftarrow m_1 - x_i$.
 - (b) If the coin turns tails:
 - i. Assign agent i to set B_2 .
 - ii. Let $(x', \tilde{x}')$ be the tuple returned by $Alloc(\{A_2 \cup i\} - j, \mathcal{P}(A_1), \frac{m}{4})$.
 - iii. Let $x_i \leftarrow \min\{x'_i, m_2\}, \tilde{x}_i \leftarrow \min\{\tilde{x}'_i, m_2\}$.
 - iv. Allocate To agent i x_i items at a price of $\tilde{x}_i \cdot \mathcal{P}(A_1)$.
 - v. Update $m_2 \leftarrow m_2 - x_i$.

Figure 3: A random sampling based online mechanism.

Let $N' = N - \{j\}$. The following observation shows that sets $A_1, A_2 - j, B_1$ and B_2 form a random partition of the agents in N' .

Observation 2. Mechanism **RS_Online** places each agent $i \in N'$ randomly and independently in either sets $A_1, A_2 - j, B_1$ or B_2 with equal probability.

Proof. Observe that $|N'| = n - 1$ and $|A - j| \sim B(n - 1, 1/2)$. Since the agents of N' arrive in a random

order, the sets $A - j$ and B form a uniformly-sampled 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. ■

3 Revenue Maximization

In this Section we study mechanisms for the allocation of identical items (both divisible and indivisible) with constant approximation to the optimal revenue under a large market assumption.

Let S be a set of agents, p a price and k a number of items. We define $p_{S,k}^* = \operatorname{argmax}_p \left\{ \min \left(\sum_{i \in S, v_i \geq p} \frac{b_i}{p}, k \right) \cdot p \right\}$, i.e., $p_{S,k}^*$ is the price that maximizes the revenue of selling at most k items to agents in S at a uniform price. Let $OPT(S, k) = \min \left(\sum_{i \in S, v_i \geq p_{S,k}^*} \frac{b_i}{p_{S,k}^*}, k \right) \cdot p_{S,k}^*$, i.e., the optimal revenue from selling at most k items to agent in S at a uniform price. Let $OPT^*(S, k)$ be the optimal revenue from selling at most k items to agents in S at heterogeneous prices. The following theorem from [1] shows that a mechanism that gives a constant approximation to $OPT(S, k)$ also gives a constant approximation with respect to $OPT^*(S, k)$.

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

3.1 The Mechanism

For the allocation of divisible and indivisible items we apply the **RS_Online** mechanism described in the previous section. For divisible items, we define **Rev_Div** = **RS_Online**(N, m, P_{rev}^m , **Div_Alloc**), where $P_{rev}^m(S) = p_{S, \frac{m}{4}}^*$, and **Div_Alloc** is as defined in Section 2.1. For indivisible items, we define **Rev_Indiv** = **RS_Online**(N, m, P_{rev}^m , **Indiv_Alloc**), where P_{rev}^m is as above, and **Indiv_Alloc** is as defined in Section 2.1.

Recall that without a large market assumption, no truthful mechanism can achieve a constant approximation to the revenue. Similarly to [7], we characterize the size of the market by ϵ , described as follows: Let $b_S^{\max} = \max\{b_i\}_{i \in S}$ be the maximum budget of any agent in set S ; $\epsilon(S, k) = \frac{b_S^{\max}}{OPT(S, k)}$. Intuitively, a smaller ϵ implies a larger market. The main result of this section is:

Theorem 2. **Rev_Indiv** is a truthful mechanism that allocates indivisible items and gives a 32-approximation to the optimal revenue as ϵ tends to 0. In addition, **Rev_Div** is a truthful mechanism that allocates divisible items and gives a 16-approximation to the optimal revenue as ϵ tends to 0.

In what follows we prove the theorem for the allocation of indivisible items. It is not difficult to observe that proving 32-approximation for indivisible items directly implies 16-approximation for divisible items. This is because our analysis goes through the allocation of divisible items and then applies Lemma 1 to show that the transition to indivisible items loses at most a factor 2 in the approximation.

3.2 Analysis

Theorem 3. Mechanism **Rev_Indiv** is truthful, i.e., each agent maximizes her utility by reporting her true type.

Proof. Let i be any agent. We prove that the utility of agent i can only decrease if she misreports her type.

Misreporting v_i, b_i : Agents in set A and B cannot change their price or position by misreporting their

value or budget. Since given a fixed price and a position the mechanism allocates the best possible allocation, agents do not have the incentive to misreport their bids.

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 . Therefore, by reporting an earlier departure time, agent i might depart before time t_0 . 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 agent might receive less items (since other agents might be allocated before agent i).
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.
3. The mechanism places agent i in set A according to both her true arrival time and reported arrival time. Since agent i is placed in set A , by misreporting her arrival time she can only affect the arrival order and the time t_0 . Observe that the mechanism allocates the agents in a random order and takes into consideration departed agents (although they do not receive items).
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 . Let $(x, \tilde{x})$ be the tuple returned by $M_1 = \text{Indiv_Alloc} \left(\{A_1 \cup i\}, p_{A_2-j}^*, \frac{m}{4}, \frac{m}{4} \right)$. Let $(x', \tilde{x}')$ be the tuple returned by $M_2 = \text{Indiv_Alloc} \left(\{A_2 \cup i\} - j, p_{A_1}^*, \frac{m}{4}, \frac{m}{4} \right)$. If agent i had been reporting her true arrival time, she would have been placed in set A_1 with probability $\frac{1}{2}$, and in set A_2 with probability $\frac{1}{2}$.

Note that $E_{M_1}[x_i] = E_{M_1}[\tilde{x}_i]$ and $E_{M_1}[x'_i] = E_{M_1}[\tilde{x}'_i]$. Therefore, by reporting her true arrival time, with probability $\frac{1}{2}$, agent i receives an expected amount of $E_{M_1}[x_i]$ items at a price of $E_{M_1}[x_i] \cdot p_{A_2-j}$, and with probability $\frac{1}{2}$, agent i receives an expected amount of $E_{M_2}[x'_i]$ items at a price of $E_{M_2}[x'_i] \cdot p_{A_1}$. Since agent i misreported her arrival time, she is placed in set B_1 with probability $\frac{1}{2}$ and placed in set B_2 with probability of $\frac{1}{2}$. Hence, with probability $\frac{1}{2}$, agent i receive an expected amount of $x \leq E_{M_1}[x_i]$ items at an expected price of $x \cdot p_{A_2-j}$, and with probability $\frac{1}{2}$, agent i receives an expected amount of $x \leq E_{M_2}[x'_i]$ items at an expected price of $x \cdot p_{A_1}$ (recall that for set B the allocation is bounded by the remaining amount of the item). It is clear that the expected utility of agent i can only decrease by misreporting her arrival time. ■

Recall that $N' = N - j$. Let $OPT^* = OPT^*(N, m)$, $OPT = OPT(N, m)$, $\epsilon = \epsilon(N, m)$, $OPT' = OPT(N', m)$, $\epsilon' = \epsilon(N', m)$. The following lemma shows that there is no significant difference between OPT' and OPT :

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

Proof. There can be at most a single agent that can contribute more then $OPT/2$ to the optimal revenue from set N . Since agent j is chosen uniformly at random from set N (recall that the agents arrive in a random order), the probability that such an agent (if exists) is not in $N - j$ is $1/n$. ■

Our analysis uses theorems from [7]. Their mechanism, which we will refer to as **Rev_Offline** is based on the random sampling paradigm. The mechanism 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 . For a formal description see Figure 5 at Appendix A.

Let $\bar{\epsilon} = \epsilon(S, k)$. Borgs et al. showed the following:

Theorem 4. ³ [7] *Let r_{S_1} be the revenue from selling at most $k/2$ items to set S_1 at a price of $p_{S,k}^*$. Let r_{S_2} be the revenue from selling at most $k/2$ items to set S_2 at a price of $p_{S,k}^*$. For any $\delta \in [0, 1]$, the probability that both $r_{S_1} \geq \frac{1-\delta}{2} \text{OPT}(S, k)$ and $r_{S_2} \geq \frac{1-\delta}{2} \text{OPT}(S, k)$ is at least $1 - 2e^{-\delta^2/(4\bar{\epsilon})}$.*

Theorem 5. ⁴ [7] *Let r_{S_1} be the revenue of mechanism **Rev_Offline** from agents in set S_1 . Then, for any $\delta \in [0, \frac{1}{3}]$, the probability that $r_{S_1} \geq \frac{1-3\delta}{2} \text{OPT}(S, k)$ is at least $1 - 2e^{-\delta^2/(4\bar{\epsilon})}$.*

Mechanism **Rev_Offline** sells items in an arbitrary order, in particular, according to the online arrival order. Moreover, according to Observation 2 (see Section 2.2), our mechanism partition the agents into sets in an equivalent way to mechanism **Rev_Offline**. Therefore, we can use the theorems from [7] in our analysis.

Let $N_1 = B_1 \cup A_2 - j$ and $N_2 = B_2 \cup A_1$. Let $\text{OPT}_1 = \text{OPT}(N_1, m/2)$, $\epsilon_1 = \epsilon(N_1, m/2)$, $\text{OPT}_2 = \text{OPT}(N_2, m/2)$ and $\epsilon_2 = \epsilon(N_2, m/2)$. The following lemma proves that with high probability both sets receive a significant fraction of N' .

Lemma 3. $\Pr [\text{OPT}_1 \geq \frac{1}{4} \text{OPT}', \text{OPT}_2 \geq \frac{1}{4} \text{OPT}', \epsilon_1 \leq 4\epsilon', \epsilon_2 \leq 4\epsilon'] \geq 1 - 2e^{-1/16\epsilon'}.$

Proof. Using Observation 2, we get that N_1 and N_2 form a uniformly sampled partition of N' . Hence, we can apply Theorem 4 with $\delta = \frac{1}{2}$, $S = N'$, $S_1 = N_1$, $S_2 = N_2$, $k = m$. Therefore, with probability of $1 - 2e^{-1/16\epsilon'}$ we get that $\text{OPT}_1 \geq r_{N_1} \geq \frac{1}{4} \text{OPT}'$, $\text{OPT}_2 \geq r_{N_2} \geq \frac{1}{4} \text{OPT}'$. Since $b_{N_1}^{\max} \leq b_{N'}^{\max}$ and $b_{N_2}^{\max} \leq b_{N'}^{\max}$ we achieve the bounds for ϵ_1 and ϵ_2 stated in the assertion of the lemma (recall that $\epsilon(S, k) = \frac{b_S^{\max}}{\text{OPT}(S, k)}$). ■

Our mechanism, as opposed to offline random sampling based mechanisms, calculates the allocation for each agent in set B_1 as if she was in a random permutation in set A_1 (and applies an analogous procedure for sets B_2 and $A_2 - j$). This modification is vital for the truthfulness of the mechanism.

The following lemma is used to show that we lose only a constant factor due to this modification. Let x be the allocation returned by **Div_Alloc**(S, p, k). **Sell_Div**(S, p, k) (see Figure 6 at Appendix A) is a mechanism that allocates the items according to x and charges a price of p per item. **Sell_Modif**(S, p, k, z) (see Figure 7 at Appendix A) is a mechanism that sells according to the modification described above, i.e., it divides set S into sets S_1 and S_2 such that each agent has a probability of z to be placed in set S_1 and a probability of $1 - z$ to be placed in set S_2 . Then it sells the items to agents in set S_1 in a random 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 .

Lemma 4. *For every set S , price p , fraction of item k and a probability z , the expected revenue from mechanism **Sell_Modif**(S, p, k, z) is at least z times the expected revenue from mechanism **Sell_Div**(S, p, k).*

³Note that this is stronger than the statement of Lemma 5.2 in [7], but this is essentially what they proved.

⁴Note that this result is shown in proof of Theorem 5.1 in [7].

Proof. Let r be the expected revenue from mechanism $\text{Sell_Div}(S, p, k)$. For the sake of the analysis, assume that mechanism $\text{Sell_Div}(S, p, k)$ begins by placing each agent $i \in S$ in S_1 and S_2 with respective probability z and $1 - z$ (as mechanism $\text{Sell_Modif}(S, p, k, z)$ does). Since each agent has a probability of z to be placed in S_1 , the expected revenue extracted from agents in S_1 is $z \cdot r$. Therefore, it is sufficient to show that for each partition of S into S_1 and S_2 , the expected revenue from $\text{Sell_Modif}(S, p, k, z)$ equals to at least the expected revenue of $\text{Sell_Div}(S, p, k)$ from agents in set S_1 . There are two cases. If $\text{Sell_Modif}(S, p, k, z)$ allocated the entire fraction of the item, then it collected the maximal revenue possible from price p and the claim is obviously true.

Otherwise, for each agent $i \in S_1$, $\min\{x'_i, m\} = x'_i$ at step 3a of mechanism Sell_Modif . We prove that for each such an agent i , her expected revenue in $\text{Sell_Modif}(S, p, k)$ (denoted by $r_{i,2}$) equal at least her expected revenue in $\text{Sell_Div}(S, p, k)$ (denoted by $r_{i,1}$). Since $\min\{x'_i, m\} = x'_i$ then $r_{i,2} = \mathbb{E}_{\text{Div_Alloc}}[\tilde{x}_i] \cdot p$ where $\tilde{x}$ is the allocation returned by $\text{Div_Alloc}(\{S_2 \cup i\}, p, k)$. $r_{i,1} = \mathbb{E}_{\text{Div_Alloc}}[\tilde{x}_i] \cdot p$ where $\tilde{x}$ is the allocation returned by $\text{Div_Alloc}(S, p, k)$. Since $\{S_2 \cup i\} \subseteq S$, then according Observation 1 $\mathbb{E}_{\text{Div_Alloc}}[\tilde{x}_i] \geq \mathbb{E}_{\text{Div_Alloc}}[\tilde{x}_i]$. This complete the proof. ■

We are now ready to prove the main technical lemma of this section.

Lemma 5. *For every $\delta \in [0, \frac{1}{3}]$, the probability that the expected revenue obtained from mechanism $\text{Rev_Indiv}(N, m)$ is greater than $\frac{(1-3\delta)OPT'}{16}$ is at least $(1 - 2e^{-1/16\epsilon'}) (1 - 4e^{-\delta^2/(16\epsilon')})$.*

Proof. We calculate the expected revenue only for 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 from mechanism $\text{Sell_Div}(B_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4})$. Note that each agent $i \in S_1$ is randomly and independently placed to either B_1 or $A_2 - j$ with equal probability. Hence, we can apply Theorem 5 with $S_1 = B_1$, $S_2 = A_2 - j$ and get that the probability that $r \geq \frac{1-3\delta}{2}OPT_1 \geq \frac{1-3\delta}{8}OPT'$ is at least $1 - 2e^{-\delta^2/(4\epsilon_1)} \geq 1 - 2e^{-\delta^2/(16\epsilon')}$. Note that the expected revenue from mechanism $\text{Sell_Div}(B_1 \cup A_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4})$ is at least the expected revenue from mechanism $\text{Sell_Div}(B_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4})$, since the agents in set $B_1 \cup A_1$ with value at least $p_{A_2-j, \frac{m}{4}}^*$ are a superset of the ones in set B_1 . Hence, the probability that the expected revenue from mechanism $\text{Sell_Div}(B_1 \cup A_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4})$ is at least $\frac{1-3\delta}{8}OPT'$ is at least $1 - 2e^{-\delta^2/(16\epsilon')}$.

By applying Lemma 4 with $S = B_1 \cup A_1$, $S_1 = B_1$, $S_2 = A_1$, $p = p_{A_2-j, \frac{m}{4}}^*$, $k = \frac{m}{4}$ and $z = \frac{1}{2}$ we get the following. The probability that the expected revenue from $\text{Sell_Modif}(B_1 \cup A_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4}, \frac{1}{2})$ is greater than $\frac{1}{2} \cdot \frac{1-3\delta}{8}OPT' = \frac{1-3\delta}{16}OPT'$ is at least $1 - 2e^{-\delta^2/(16\epsilon')}$. By a symmetric argument we get that the probability that the expected revenue from mechanism $\text{Sell_Modif}(B_2 \cup A_2 - j, p_{A_1, \frac{m}{4}}^*, \frac{m}{4}, \frac{1}{2})$ 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 from mechanism $\text{Sell_Modif}(B_1 \cup A_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4}, \frac{1}{2})$ and mechanism $\text{Sell_Modif}(B_2 \cup A_2 - j, p_{A_1, \frac{m}{4}}^*, \frac{m}{4}, \frac{1}{2})$ is greater than $\frac{1-3\delta}{8}OPT'$ is at least $1 - 4e^{-\delta^2/(16\epsilon')}$. Note that we can use Lemma 4 although we enforce that $OPT_1 \geq \frac{1}{4}OPT'$ 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 this conditions are symmetric.

One can easily observe that the expected revenue from mechanism $\text{Rev_Div}(N, m)$ is at least the expected revenue from mechanism $\text{Sell_Modif}(B_1 \cup A_1, p_{A_2-j, \frac{m}{4}}^*, \frac{m}{4}, \frac{1}{2})$ and mechanism $\text{Sell_Modif}(B_2 \cup A_2 - j, p_{A_1, \frac{m}{4}}^*, \frac{m}{4}, \frac{1}{2})$. To bound the expected revenue from mechanism $\text{Rev_Indiv}(N, m)$ we apply Lemma 1 and lose an additional factor of 2 in the revenue. Combining the above with Lemma 3, we get the desired result. ■

We now use Theorem 1 and Lemma 2, to bound the expected revenue from mechanism **Rev_Indiv** with respect to OPT^* .

Theorem 6. *For every $\delta \in [0, \frac{1}{3}]$, the probability that the expected revenue obtained from mechanism **Rev_Indiv**(N, m) is greater than $\frac{(1-3\delta)OPT^*}{64}$ is at least $(1 - 1/n) (1 - 2e^{-1/32\epsilon}) (1 - 4e^{-\delta^2/(32\epsilon)})$.*

Proof. According to Lemma 2, $\Pr[OPT' \geq OPT/2] \geq 1 - 1/n$. Combining $OPT' \geq OPT/2$ and $\epsilon' \leq 2\epsilon$ with the bound from Lemma 5 and Theorem 1 yields the desire result. ■

We are now ready to show the main result of this section.

Proof of Theorem 2: By definition, when ϵ tends to 0, OPT' approaches to OPT . Therefore, the bound from Lemma 5 becomes: for every $\delta \in [0, \frac{1}{3}]$, the probability that the expected revenue obtained from mechanism **Rev_Indiv**(N, m) is greater than $\frac{(1-3\delta)OPT^*}{32}$ is at least $(1 - 2e^{-1/32\epsilon}) (1 - 4e^{-\delta^2/(32\epsilon)})$. Since ϵ tends to 0, for every δ , $(1 - 2e^{-1/32\epsilon}) (1 - 4e^{-\delta^2/(32\epsilon)})$ tends to 1. ■

4 Liquid Welfare Maximization

In the liquid welfare setting the seller has a single divisible item⁵. Let $\tilde{v}_i(x_i) = \min(v_i x_i, b_i)$. The liquid welfare from an allocation vector x is $\tilde{W}(x) = \sum_i \tilde{v}_i(x_i)$. Let OPT denote the optimal liquid welfare. For an agent $i \in N$ we define $\bar{v}_i = \min\{v_i, b_i\}$ to be the *liquid value* of the agent. We refer to an agent whose liquid value is a constant fraction of OPT as a “dominant” agent.

Our mechanism, which we refer to as **Liquid_Div**, is composed of two mechanisms. It runs mechanism **MVCG**(N) with probability μ and mechanism **RS_Liquid** with probability $1 - \mu$, where the exact value of μ will be determined in the analysis.

We define **RS_Liquid** = **RS_Online**($N, 1, \bar{P}, \text{Div_Alloc}$), where **Div_Alloc** is as defined in Section 2.1, and $\bar{P}$ is described as follows: Let S be a set of n' agents such that $v_{i_1} \geq v_{i_2} \dots \geq v_{i_{n'}}$. Let k be the maximum integer such that $\sum_{j=1}^k b_{i_j} \leq v_{i_k}$. The function $\bar{P}(S)$ is the market clearing price⁶ defined as follows: $\bar{P}(S) = \max\{\sum_{j=1}^k b_{i_j}, v_{i_{k+1}}\}$.

Mechanism **MVCG**(N), which is defined in Figure 4, achieves a constant approximation to the optimal liquid welfare if there is a single dominant agent, and Mechanism **RS_Liquid** achieves a constant approximation in all other cases. The main result of this section is:

Theorem 7. *Mechanism **Liquid_Div** is a truthful mechanism for the allocation of a divisible item that gives a constant approximation to the optimal liquid welfare.*

⁵By normalization, this is equivalent to m divisible items.

⁶Market clearing price is a term taken from the general equilibrium theory. In the special case of a single type of good, market clearing price is the maximum price satisfying that each agent can be assigned an optimal fraction of the good such that there is no deficiency or surplus.

MVCG (N)

1. Let A be the set of the first $\frac{n}{2}$ agents (based on reported arrival time) and let B be $N \setminus A$. Let $t_{0.5} = \hat{a}_{\frac{n}{2}}$ (the reported arrival time of the $\frac{n}{2}$ -th agent).
2. At time $t_{0.5}$: Order the agents in set A according to their liquid value such that $\bar{v}_{i_1} \geq \bar{v}_{i_2} \dots \geq \bar{v}_{i_{\frac{n}{2}}}$. Set $p_2 = \gamma \cdot \bar{v}_{i_2}$ and $p_1 = \gamma \cdot \bar{v}_{i_1}$. Toss a fair coin.
 - If it turns heads, then if $\bar{v}_{i_1} \geq p_2$ and $d_{i_1} \geq t_{0.5}$, allocate the entire item to agent i_1 at a price p_2 .
 - If the coin turns tail, then wait for the first agent $j \in B$ such that $\bar{v}_j \geq p_1$ (if exists) and allocate to her the entire item at a price p_1 .

Figure 4: Online modified VCG mechanism.

The value of $\gamma > 1$ is determined in the analysis.

4.1 Analysis

Theorem 8. *The mechanism is truthful for any $\gamma > 1$, i.e., each agent maximizes her utility by reporting her true type.*

The proof of the above theorem is deferred to Appendix C

We now bound the performance guarantee of the mechanism. Some parts of the analysis resemble the analysis from [26]. We bound the performance guarantee for several cases. For some of them we prefer to bound the expected revenue instead of bounding the expected liquid welfare. The relation between revenue and liquid welfare is stated in Lemma 3.4 from [26].

Lemma 6. [26] *The liquid welfare produced by any truthful and budget feasible mechanism is at least the revenue of the auctioneer.*

We introduce the following definitions. Recall that j is the last agent in set A (according to her reported arrival time), where $j \sim B(n-1, 1/2) + 1$. Let $N' = N - j$. Let $p^* = \bar{P}(N')$ be the market clearing price of the total set of agents without agent j . Let OPT and OPT' be the optimal liquid welfare for sets N and N' respectively. Let $V = \{i \in N | v_i \geq p^*\}$ and $V' = \{i \in N' | v_i \geq p^*\}$. For every set of agents $T \subset N'$ we define $B_T = \sum_{i \in T} b_i$, $B'_T = \sum_{i \in T \cap V'} b_i$. By definition, $B'_V \geq p^*$ (Recall that the market clearing price is the maximal price such that there is no deficiency or surplus of any good. Hence, it is clear that the budget of the agent from set V' is at least p^*).

The following lemma proves some relations regarding the market clearing price. In particular, it proves that the market clearing price is a good approximation for the optimal liquid welfare. This Lemma was stated and proved by Lu and Xiao in an earlier version of [26], and removed in the latest version. We state and prove it here for the sake of completeness.

Lemma 7 (From a previous version of [26]⁷). *Let S and $T \subseteq S$ be two sets of agents and let W be the optimal liquid welfare from the set S . Then:*

1. $\bar{P}(S) \geq \frac{1}{2}W$.
2. $\bar{P}(T) \leq \bar{P}(S)$.

⁷The previous version can be found in <http://arxiv.org/pdf/1407.8325v2.pdf>.

Proof. Let $S = \{1, \dots, n\}$, $v_1 \geq v_2 \geq \dots \geq v_n$ and let k be the maximum integer such that $\sum_{i=1}^k b_i \leq v_k$. By the definition of the market clearing price we have that $\bar{P}(S) = \max\{\sum_{i=1}^k b_i, v_{k+1}\}$. Let $x^* = \{x_i^*\}_{i \in S}$ be an allocation that maximizes the social welfare of the agents in S . First we prove the first part of the lemma:

$$W = \sum_{i \in S} \min\{v_i x_i^*, b_i\} \leq \sum_{i=1}^k b_i + \sum_{i=k+1}^n v_i x_i^* \leq \sum_{i=1}^k b_i + v_{k+1} \sum_{i=k+1}^n x_i^* \leq \sum_{i=1}^k b_i + v_{k+1} \leq 2\bar{P}(S),$$

where the first inequality is derived from the definition of the min function, the third inequality is due to the fact that only one item is allocated and the fourth inequality results from the definition of clearing market price.

Now we prove the second part of the lemma. Let $T = \{i_1, i_2, \dots, i_{n'}\}$ such that $v_{i_1} \geq v_{i_2} \geq \dots \geq v_{i_{n'}}$. Let k' be the maximum $i \in \{1, \dots, k\}$ such that $k' \in T$. Let k_T the maximum integer such that $\sum_{j=1}^{k_T} b_{i_j} \leq v_{i_{k_T}}$ (i.e., $\bar{P}(T) = \max\{\sum_{j=1}^{k_T} b_{i_j}, v_{i_{k_T}+1}\}$). Since $k' \leq k$, according to the definition of k , $\sum_{j=1}^{k'} b_j \leq v_{k'}$. Combining with the fact that T is a subset of S we have that either $i_{k_T} = k'$ or $i_{k_T} > k$. For the first case we have that $\bar{P}(T) = \max\{\sum_{j=1}^{k_T} b_{i_j}, v_{i_{k_T}+1}\} \leq \max\{\sum_{j=1}^k b_j, v_{k+1}\} = \bar{P}(S)$ since $i_{k_T} = k' \leq k$ and $i_{k_T+1} > k$ (according to the definition of k'). On the other hand, if $i_{k_T} > k$ then $\bar{P}(T) = \max\{\sum_{j=1}^{k_T} b_{i_j}, v_{i_{k_T}+1}\} \leq v_{i_{k_T}} \leq v_{k+1} \leq \max\{\sum_{j=1}^k b_j, v_{k+1}\} = \bar{P}(S)$. When the first inequality results from the definition of the clearing market price.

This complete the proof of the lemma. $\blacksquare$

As a corollary of the Lemma 7, we get that $p^* \geq \frac{1}{2}OPT'$, $p_{A_1} \leq p^*$ and $p_{A_2-j} \leq p^*$. The following lemma shows that there is not a significant difference between OPT' and OPT .

Lemma 8. $\Pr \left[OPT' \geq \left(1 - \frac{1}{k+1}\right)OPT \right] \geq 1 - \frac{k}{n}$.

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 $\tilde{v}_\ell(x_\ell^*) \leq \frac{OPT}{k+1}$; otherwise, $\sum_{i \in (I_k + \ell)} \tilde{v}_i(x_i^*) > (k+1)\frac{OPT}{k+1} = OPT$ 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

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

Using Observation 2, we get that every agent of V' is in each of the sets A_1 , $A_2 - j$, B_1 and B_2 with an equal and independent probability. Let $Q = A_1 \cup B_1 \cup B_2$. Hence, each agent of V' is placed in set Q with a probability of $\frac{3}{4}$ and in set $A_2 - j$ with a probability of $\frac{1}{4}$.

For simplicity, we normalize $p^* = 1$. The following lemma shows that if all the agents in set V' have a relatively small budget, both Q and $A_2 - j$ get a significant fraction of the budget of the agents in set V' .

Lemma 9. If $\max_{i \in V'} b_i \leq \frac{\gamma}{10}$ then $\Pr \left[\min\{B'_Q, B'_{A_2-j}\} \geq \frac{1}{15} \right] \geq 1 - \frac{135}{242}\gamma$.

Proof. Recall that $B_{V'} \geq p^* = 1$. Assume that $B_{V'} = 1$ (if $B_{V'} > 1$, we can decrease some of the budgets, and the claim only gets stronger). Let $\mathbb{I}_i$ be a random indicator such that $\mathbb{I}_i = 1$ if agent i belongs to set Q and $\mathbb{I}_i = 0$ otherwise. Therefore, $B'_Q = \sum_{i \in V'} b_i \mathbb{I}_i$, $E[B'_Q] = \frac{3}{4}$ and

$$\begin{aligned} Var[B'_Q] &= \sum_{i \in V'} Var[b_i \mathbb{I}_i] = \sum_{i \in V'} \left(E[(b_i \mathbb{I}_i)^2] - E[b_i \mathbb{I}_i]^2 \right) \\ &= \sum_{i \in V'} \frac{3}{16} b_i^2 \leq \frac{1}{(\gamma/10)} \frac{3}{16} \cdot (\gamma/10)^2 = \frac{3}{160}\gamma, \end{aligned}$$

where the inequality holds since setting $b_i = \gamma/10$ for every $i \in V'$ maximizes $\sum_{i \in V'} b_i^2$ under the constraints $\max_{i \in V'} b_i \leq \gamma/10$ and $\sum_{i \in V'} b_i = 1$. Using Chebyshev's inequality we get:

$$\Pr \left[\left| B'_Q - \frac{3}{4} \right| \geq 11/60 \right] \leq \frac{\text{Var}[B'_Q]}{\left(\frac{11}{60}\right)^2} \leq \frac{135}{242}\gamma.$$

When $|B'_Q - \frac{3}{4}| \leq \frac{11}{60}$ then $\frac{17}{30} \leq B'_Q \leq \frac{14}{15}$. Since $B'_Q + B'_{A_2-j} = B'_V = 1$, we get that $B'_{A_2-j} \geq \frac{1}{15}$, which completes the proof. $\blacksquare$

In the following theorem we use the relation between mechanisms **Sell_Div** and **Sell_Modif** (see Figures 6 and 7) presented in Section 3. Recall that **Sell_Div**(S, p, k) is a mechanism that allocates the item according to **Div_Alloc**(S, p, k) and charges a price of p per the entire item. **Sell_Modif**(S, p, k, z) is a mechanism that sells according to the following modification. It divides a set S into sets S_1 and S_2 such that each agent has a probability of z to be placed in set S_1 and a probability of $1 - z$ to be placed in set S_2 . It then sells the items to agents in set S_1 in a random order at a price p per the entire item. The allocation for each agent is calculated as if the agent was in a random permutation in set S_2 .

Note that for $n < 100$, devising an 100-approximation truthful mechanism is trivial — allocate the entire item to a randomly chosen agent. Therefore, from now on we assume that $n \geq 100$. Now we are ready to prove our main result of this section:

Theorem 9. *Setting $\mu = 1/10$ and $\gamma = 1.0001$, mechanism **Liquid_Div** gives a 2443-approximation to the optimal liquid welfare.*

Proof. We bound the expected welfare for the event that $OPT' \geq \frac{9}{10}OPT$ occurs. In this case $OPT \leq \frac{10}{9}OPT' \leq \frac{20}{9}p^* = \frac{20}{9}$, where the second inequality stems from Lemma 7.

We bound the performance of the mechanism for three different cases. Assume by renaming that 1 and 2 are the agents with the highest liquid value and the second highest liquid value in set N respectively.

Case A (no “dominant” agent): $\max_{i \in V} b_i \leq \gamma/10$.

We bound the expected revenue from mechanism **RS_Liquid**($Q, p_{A_2-j}, \frac{1}{4}$); by Lemma 6 it is sufficient. It is clear that $\max_{i \in V'} b_i \leq \gamma/10$. We calculate the expected revenue for the event that $\min\{B'_Q, B'_{A_2-j}\} \geq 1/15$ occurs. In this case, by Lemma 7, $p_{A_2-j} = \bar{P}(A_2 - j) \geq B'_{A_2-j}/2 \geq \frac{1}{30}$. First we calculate the expected revenue from **Sell_Div**($Q, p_{A_2-j}, \frac{1}{4}$). There are two possible outcomes. If the entire $\frac{1}{4}$ fraction was sold, then the revenue is at least $\frac{1}{4 \cdot 30} = \frac{1}{120}$. Otherwise, the entire budget of the agents with value at least p_{A_2-j} was exhausted. According to Lemma 7, $p_{A_2-j} \leq p^*$. Therefore, the budget of agents in set $Q \cap V'$ was exhausted and the revenue is at least $B'_Q \geq \frac{1}{15}$. Combining the two cases we get that the expected revenue from mechanism **RS_Liquid**($Q, p_{A_2-j}, \frac{1}{4}$) is at least $\min\{\frac{1}{120}, \frac{1}{15}\} = \frac{1}{120}$. Applying Lemma 4 with $S = Q$ and $z = 1/3$ we get that the expected revenue from mechanism **Sell_Modif**($Q, p_{A_2-j}, \frac{1}{4}, \frac{1}{3}$) is at least $\frac{1}{3 \cdot 120} = \frac{1}{360}$.

Observe that the expected revenue obtained from agents in set B_1 by running mechanism **RS_Liquid**(N) is at least the expected revenue from **Sell_Modif**($Q, p_{A_2-j}, \frac{1}{4}, \frac{1}{3}$), since when running mechanism **RS_Liquid**(N) the agents in set B_1 are competing only against agents from set A_1 and when running **Sell_Modif**($Q, p_{A_2-j}, \frac{1}{4}, \frac{1}{3}$) they have to compete against agents from set $A_1 \cup B_2$. We get that the expected revenue is at least $(1 - \mu)(1 - \frac{145}{242}\gamma)/360 \geq 1/1000$.

Case B (More than one “dominant” agent): $\max_{i \in V} b_i > \gamma/10$ and $\bar{v}_1 < \gamma\bar{v}_2$.

We bound the expected revenue from mechanism **RS_Liquid**. Since $\max_{i \in V} b_i \geq \gamma/10$, $\bar{v}_1 \geq \min\{1, \gamma/10\} =$

$\gamma/10$ (recall that p^* is normalized to 1). Moreover, $\bar{v}_2 > \bar{v}_1/\gamma \geq 1/10$. Let i' be the agent with the highest liquid value in set V . Clearly, $\bar{v}_{i'} \geq \gamma/10$.

Next, we bound the probability that agent j is not one of the agents $i', 1, 2$. Recall that the number of agents is at least 100. We assume that $OPT' \geq \frac{9}{10}OPT$, and therefore, j is not one of the 9 agents with the highest welfare in the optimal allocation. It is clear that the worst case is when agents $i', 1, 2$ are not among the 9 agents with the highest welfare. Hence, this probability is at least $\frac{88}{91}$.

Let $Q_1 = A_1 \cup B_1$ and $Q_2 = A_2 - j \cup B_2$. Note that at least one of the agents 1, 2 is not agent i' . Assume without loss of generality that agent 2 is different than agent i' (the other case is similar).

We inspect the case where either $2 \in A_1 \wedge i' \in Q_2$ or $2 \in A_2 - j \wedge i' \in Q_1$. This happens with probability $1/4$ (recall that each agent in set N' has the same independent probability of being placed in either $A_1, B_1, A_2 - j$ or B_2).

Assume without loss of generality that $2 \in A_1$ and $i' \in Q_2$ (the other case is symmetric). Using the same argument as in the previous section, we have that the expected revenue from mechanism **Sell_Div**($Q_2, p_{A_1}, 1/4$) is at least $\min\{\frac{1}{80}, \frac{\gamma}{10}\} = \frac{1}{80}$. By using Lemma 4 with $S = Q_2$ and $z = 1/2$ we get that the expected revenue out of agents in B_2 in **RS_Liquid** is at least $\frac{1}{160}$. Therefore, the expected revenue in this case is at least $(1 - \mu)\frac{88}{91 \cdot 4 \cdot 160} \geq \frac{13}{10000}$.

Case C (one “dominant” agent): $\max_{i \in V} b_i > \gamma/10$ and $\bar{v}_1 < \gamma\bar{v}_2$.

We bound the expected liquid welfare from mechanism **MVCG**. With probability $1/4$ agent 1 is placed in set B and agent 2 is placed in set A , and with probability $1/2$ the item is sold in the second phase of mechanism **MVCG**. Hence, the expected liquid welfare is at least $\frac{\gamma}{2 \cdot 4 \cdot 10} \mu = \frac{1}{80} \mu \gamma = \frac{1250125}{1000000000}$.

We conclude that the expected liquid welfare is at least $\min\{\frac{1}{1000}, \frac{13}{10000}, \frac{1250125}{1000000000}\} = \frac{1}{1000}$. Recall that in the event that $OPT' \geq \frac{9}{10}OPT$, $OPT \leq \frac{20}{9}$. According to Lemma 8, $\Pr[OPT' \geq \frac{9}{10}OPT] \geq 91/100$. Therefore, the approximation ratio is at most $\frac{\frac{20}{9}}{\frac{1}{1000} \cdot \frac{91}{100}} < 2443$. ■

References

- [1] Z. Abrams. Revenue maximization when bidders have budgets. In *SODA*, pages 1074–1082. ACM, 2006.
- [2] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *STOC*, pages 503–510. ACM, 2003.
- [3] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA*, pages 434–443. SIAM, 2007.
- [4] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *SODA*, pages 964–970. Society for Industrial and Applied Mathematics, 2002.
- [5] S. Bhattacharya, V. Conitzer, K. Munagala, and L. Xia. Incentive compatible budget elicitation in multi-unit auctions. In *SODA*, pages 554–572. Society for Industrial and Applied Mathematics, 2010.
- [6] A. Blum and J. D. Hartline. Near-optimal online auctions. In *SODA*, pages 1156–1163. Society for Industrial and Applied Mathematics, 2005.
- [7] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *EC*, pages 44–51. ACM, 2005.

- [8] S. Chawla, D. L. Malec, and A. Malekian. Bayesian mechanism design for budget-constrained agents. In *EC*, pages 253–262. ACM, 2011.
- [9] R. Colini-Baldeschi, M. Henzinger, S. Leonardi, and M. Starnberger. On multiple keyword sponsored search auctions with budgets. In *ICALP*, pages 1–12. Springer, 2012.
- [10] N. R. Devanur, B. Q. Ha, and J. D. Hartline. Prior-free auctions for budgeted agents. In *EC*, pages 287–304. ACM, 2013.
- [11] S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.
- [12] S. Dobzinski and R. P. Leme. Efficiency guarantees in auctions with budgets. In *ICALP*, pages 392–404. Springer, 2014.
- [13] M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *EC*, pages 532–549. ACM, 2012.
- [14] A. Fiat, S. Leonardi, J. Saia, and P. Sankowski. Single valued combinatorial auctions with budgets. In *EC*, pages 223–232. ACM, 2011.
- [15] E. J. Friedman and D. C. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *EC*, pages 240–241. ACM, 2003.
- [16] G. Goel, V. Mirrokni, and R. P. Leme. Clinching auctions with online supply. In *SODA*, pages 605–619. SIAM, 2013.
- [17] G. Goel, V. Mirrokni, and R. Paes Leme. Polyhedral clinching auctions and the adwords polytope. In *STOC*, pages 107–122. ACM, 2012.
- [18] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. volume 55, pages 242–269. Elsevier, 2006.
- [19] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.
- [20] M. T. Hajiaghayi, R. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *EC*, pages 71–80. ACM, 2004.
- [21] M. T. Hajiaghayi, R. D. Kleinberg, M. Mahdian, and D. C. Parkes. Online auctions with re-usable goods. In *EC*, pages 165–174. ACM, 2005.
- [22] T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Algorithms–ESA 2013*, pages 589–600. Springer, 2013.
- [23] R. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 630–631. Society for Industrial and Applied Mathematics, 2005.
- [24] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *EC*, pages 233–241. ACM, 2000.

- [25] R. Lavi and N. Nisan. Online ascending auctions for gradually expiring items. In *SODA*, pages 1146–1155. Society for Industrial and Applied Mathematics, 2005.
- [26] P. Lu and T. Xiao. Improved efficiency guarantees in auctions with budgets. 2014.
- [27] D. C. Parkes. Online mechanisms. 2007.
- [28] D. C. Parkes and S. P. Singh. An mdp-based approach to online mechanism design. In *Advances in neural information processing systems*, page None, 2003.

A Omitted Mechanisms

Rev_Offline (S, k)

1. Partition the agents in set S into two sets S_1 and S_2 , by tossing an independent fair coin for each agent. Set $p_1 \leftarrow p_{S_1, m/2}^*$ and $p_2 \leftarrow p_{S_2, m/2}^*$.
2. Let x be the allocation returned by **Div_Alloc**($S_1, p_2, k/2$) and x' be the allocation returned by **Div_Alloc**($S_2, p_1, k/2$).
3. For every agent $i \in S_1$: Allocate to agent i x_i items at a price of $x_i \cdot p_2$.
4. For every agent $i \in S_2$: Allocate to agent i x'_i items at a price of $x'_i \cdot p_1$.
5. Run a correlated lottery between the agents who received fractions of items to resolve the real allocation.

Figure 5: The offline mechanism presented in [7] for selling indivisible items and maximizing revenue.

Sell_Div (S, p, k)

1. Let x be the allocation returned by **Div_Alloc** (S, p, k).
2. For each agent $i \in S$ allocate x_i items at price of $x_i \cdot p$.

Figure 6: A mechanism for selling divisible items at a given price.

Sell_Modif (S, p, k, z)

1. Initialize $m \leftarrow k$.
2. For each agent $i \in S$: With probability of z place i in set S_1 and with probability of $1 - z$ in set S_2 .
3. For each agent $i \in S_1$ (for an arbitrary order of the agents):
 - (a) Let x' be the allocation returned by **Div_Alloc** ($\{S_2 \cup i\}, p, k$). Let $x_i \leftarrow \min\{x'_i, m\}$. Allocate to agent i x_i items at a price of $x_i \cdot p$.
 - (b) Update $m \leftarrow m - x_i$.

Figure 7: A mechanism defined solely for the analysis of our mechanisms.

B Proof of Lemma 1

We consider only the case where each agent in S has a budget smaller than p (it is clear that for the integral part of $\frac{b_i}{p}$, mechanisms **Indiv_Alloc** and **Div_Alloc** will allocate the same amount of items). Let $S' = \{i \in S : v_i \geq p\}$ (agents in $S \setminus S'$ have no effect on the allocation). Let $B_{S'} = \sum_{i \in S'} b_i$. We analyze the following cases:

Case A: $k = 1$ and $B_{S'} \geq p$. It is clear that in this case $\sum_{i \in S} \tilde{x}'_i = 1$ (mechanism **Div_Alloc** sells the entire item). $E_{\text{Indiv_Alloc}}[\sum_{i \in S} \tilde{x}_i]$ is the expected amount of items charged by **Indiv_Alloc** before the item is sold. Let j be the size of set S and assume by renaming that $1, \dots, j$ is the order in which mechanism **Indiv_Alloc** iterates the agents in set S' . Every time an agent $i \in S'$ is accessed before the item is sold, she is charged for a $\tilde{x}_i = b_i/p$ fraction, and gets the item with probability $\tilde{x}_i$.

Let P_i be the probability that agent $i \in S'$ is accessed before the item is sold. Clearly, $X = \sum_{i=1}^j P_i \cdot \tilde{x}_i$ is the expected number of items the agents of S' are charged for before the item is sold.

For the sake of the analysis, consider the following. Let $r = \min\{t \in S' : \sum_{i=1}^t \tilde{x}_i \geq 1\}$. Set $y_i = \tilde{x}_i$ for every $i < r$, $y_r = 1 - \sum_{i=1}^{r-1} y_i$ and $y_i = 0$ for every $i > r$. Let $z \sim U[0, 1]$ be a uniformly sampled number between 0 and 1. Let $\tilde{j} = \min\{t \in S' : \sum_{i=1}^t y_i \geq z\}$, and let $X' = \sum_{i=1}^{\tilde{j}} y_i$. Since $E[z] = 1/2$, we have that $E[X'] \geq 1/2$. Let P'_i be the probability that $i \leq \tilde{j}$. Therefore, $E[X'] = \sum_{i=1}^j P'_i \cdot y_i$.

To complete this case we show by induction that $P_i \geq P'_i$ for every $i \in S'$. Clearly $P_1 = P'_1 = 1$. For $i > 1$, by definition $P_i = P_{i-1} \cdot (1 - \tilde{x}_{i-1})$. On the other hand,

$$P'_i = P'_{i-1} \cdot \left(1 - \frac{y_{i-1}}{\sum_{t=i-1}^j y_t}\right) \leq P_{i-1} \cdot (1 - y_{i-1}) = P_{i-1} \cdot (1 - \tilde{x}_{i-1}) = P_i,$$

where the inequality stems from the induction hypothesis. Since $\tilde{x}_i \geq y_i$ for any $i \in S'$, we conclude that $X \geq E[X'] \geq 1/2$.

Case B: $k \geq 1$ and $B_{S'} \leq p$ (agents can afford at most a single item). By a similar analysis to that of the **Case A**, we get that $E_{\text{Indiv_Alloc}}[\sum_{i \in S} \tilde{x}_i] \geq \sum_{i \in S} \tilde{x}'_i/2$.

Case C: $k \geq 2$ and $\sum_{i \in S} \tilde{x}'_i = k' > 1$. For simplicity we assume that k' is integral, but a similar argument can be applied to the fractional case. Recall that mechanism **Indiv_Alloc** sells the items sequentially. Let $B_{S', \ell}$ be the sum of the budgets of the agents in set S' , after mechanism **Indiv_Alloc** sells item $\ell - 1$ (for $\ell = 1$ clearly $B_{S', \ell} = B_{S'}$). We consider the following cases:

1. Mechanism **Indiv_Alloc** sold at least $k' - 1$ items and $B_{S', k-1} \geq p$. Therefore, before selling each of the items $1, \dots, k'$, we have that the sum of the budgets of agent in S' is at least p . Hence, we can apply the analysis of **Case A** and conclude that $E_{\text{Indiv_Alloc}}[\sum_{i \in S} \tilde{x}_i] \geq k'/2$.
2. Mechanism **Indiv_Alloc** sold less than $k' - 1$ items or $B_{S', k-1} < p$. In this case, agents in S were charged for at least $k' - 1$ items (i.e. $\sum_{i \in S} \tilde{x}_i \geq k' - 1$). Since $k' \geq 2$, $k' - 1 \geq k'/2$. This complete the proof. ■

C Proof of Theorem 8

Since the mechanism is a probabilistic combination of two mechanism, it is sufficient to show that each of the mechanisms is truthful.

RS_Liquid Mechanism

One can easily verify that mechanism **RS_Liquid** is almost identical to mechanism **Rev_Indiv**. The only differences between the mechanisms are the prices and the allocation method which does not use lotteries. Hence, the proof is similar to the proof of Theorem 3.

MVCG Mechanism

Most of the cases have a similar proof as in mechanism **Rev_Indiv**. We describe only the cases with a different proof. Let i be any agent. We prove that the utility of agent i can only decrease if she misreports 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 allocates the best possible allocation, agents do not have the incentive to misreport their bids.

For agent $i \in A$, if $\bar{v}_i \geq \gamma \cdot \bar{v}_{i_2}$ then the agent has a probability of $\frac{1}{2}$ to receive the item at a price of $\gamma \cdot \bar{v}_{i_2}$. Agent i cannot affect the price or the probability to win the item by misreporting $\bar{v}_i$. If $\bar{v}_i < \gamma \cdot \bar{v}_{i_2}$, then to win the item agent i has to bid at least $\gamma \cdot \bar{v}_{i_2}$ and to exceed her budget or value.

Misreporting a_i : We analyze the following case (analogous to case 4 of misreporting a_i in **Rev_Indiv**). 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 $\bar{v}_{i_1} \geq \bar{v}_{i_2} \dots \geq \bar{v}_{i_{\frac{n}{2}}}$ be the agents in set A ordered according to their liquid value, assuming that agent i reported her real arrival time. If $\bar{v}_i < \gamma \cdot \bar{v}_{i_2}$, agent i will not receive the item whether she is in set A or in set B . If $\bar{v}_i \geq \gamma \cdot \bar{v}_{i_2}$, then the agent receives the item at price $\gamma \cdot \bar{v}_{i_2}$ 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).

D 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 **Rev_Indiv**, **Rev_Div**, **MVCG** and **RS_Liquid** 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 all the other cases is similar to the proof of Theorem 3). 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.