

Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers

Saeed Alaei *

Date: 2011-04-29 01:02:28 -0400 (Fri, 29 Apr 2011) , Revision: 118 †

Abstract

For Bayesian combinatorial auctions, we present a general framework for reducing the mechanism design problem for many buyers to the mechanism design problem for one buyer. Our generic reduction works for any separable objective (e.g., welfare, revenue, etc) and any space of valuations (e.g. submodular, additive, etc) and any distribution of valuations as long as valuations of different buyers are distributed independently (not necessarily identically). Roughly speaking, we present two generic n -buyer mechanisms that use 1-buyer mechanisms as black boxes. We show that if we have an α -approximate 1-buyer mechanism for each buyer¹ then our generic n -buyer mechanisms are $\frac{1}{2}\alpha$ -approximation of the optimal n -buyer mechanism. Furthermore, if we have several copies of each item and no buyer ever needs more than $\frac{1}{k}$ of all copies of each item then our generic n -buyer mechanisms are $\gamma_k\alpha$ -approximation of the optimal n -buyer mechanism where $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$. Observe that γ_k is at least $\frac{1}{2}$ and approaches 1 as k increases.

Applications of our main theorem include the following improvements on results from the literature. For each of the following models we construct a 1-buyer mechanism and then apply our generic expansion: For revenue maximization in combinatorial auctions with hard budget constraints, [BGM10] presented a $\frac{1}{4}$ -approximate BIC mechanism for additive/correlated valuations and an $O(1)$ -approximate² sequential posted pricing mechanism for additive/independent valuations. We improve this to a γ_k -approximate BIC mechanism and a $\gamma_k(1 - \frac{1}{e})$ -approximate sequential posted pricing mechanism respectively. For revenue maximization in combinatorial auctions with unit demand buyers, [CHMS10] presented a $\frac{1}{6.75}$ -approximate sequential posted pricing mechanism. We improve this to a $\frac{1}{2}\gamma_k$ approximate sequential posted pricing mechanism. We also present a γ_k -approximate sequential posted pricing mechanism for unit-demand multi-unit auctions(homogeneous) with hard-budget constraints. Furthermore, our sequential posted pricing mechanisms assume no control or prior information about the order in which buyers arrive.

*Dept. of Computer Science, University of Maryland, College Park, MD 20742. saeed@cs.umd.edu. Part of this work was done when the author was visiting Microsoft Research, New England. This work was partially supported by the NSF grant CCF-0728839

†Id: bca.tex 118 2011-04-29 05:02:28Z saeed

¹Note that we can use different 1-buyer mechanisms for different buyers.

² $O(1) = \frac{1}{96}$

1 Introduction

In this paper we consider the problem of designing Bayesian combinatorial auctions for maximizing any separable objective (welfare, revenue, etc) and present a general framework for reducing the mechanism design problem for many buyers to the mechanism design problem for one buyer. Designing mechanisms for objectives other than welfare has inherently been more difficult as the objective of the mechanism is no longer aligned with the incentives of buyers. For example [HKM11] presents a black box reduction from BIC mechanism design to algorithmic design in multidimensional setting, but unfortunately their technique only works for maximizing welfare and there is no trivial way to extend it to other objectives like revenue.

From a high level point of view we face the following two challenges in designing mechanisms for many buyers:

- (i) The decisions made by the mechanism for different buyers should be coordinated because of supply constraints.
- (ii) The decisions made by the mechanism for each buyer has to be optimal (or approximately optimal).

In this paper we address the first challenge by showing that if we know how to make an approximately optimal decision for each buyer separately then we can also make approximately optimal coordinated decisions for all buyers. We show that, by going from uncoordinated decisions to coordinated decisions, we lose a factor of at most $\frac{1}{2}$ in the approximation factor. Furthermore, we show that this loss is no more than $\frac{1}{\sqrt{k+3}}$ when the ratio of the maximum possible demand of any buyer for any item to the supply of that item is no more than $\frac{1}{k}$. Observe that if we had an unlimited supply of items (i.e., $k \rightarrow \infty$) this loss would go down to 0, which is what we would expect, because with unlimited supply the optimal mechanism would treat each buyer independently³. Also observe that this does not depend on the number of buyers.

The paper is organized as follows: In section 2 we explain our model and a summary of our main results. In section 4, we present a toy problem and a near optimal algorithm for it which we use to improve the bound of the generalized prophet inequality for sum of k -choices. The best known bound for this generalization was $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ by [HKS07]. We improve this bound to $1 - \frac{1}{\sqrt{k+3}}$ by constructing a gambler that achieves this bound. This gambler uses the solution of our toy problem as a black box. This toy problem has important applications in Bayesian mechanism design and online stochastic optimization. The algorithm we develop for this toy problem is the main ingredient for our generic construction of n -buyer mechanisms. In section 5, we present two generic n -buyer mechanisms that use 1-buyer mechanisms as black boxes. We show that if we have an α -approximate 1-buyer mechanism for each buyer then our generic n -buyer mechanisms are $\frac{1}{2}\alpha$ -approximation of the optimal n -buyer mechanism. Furthermore, if we have several copies of each item and no buyer ever needs more than $\frac{1}{k}$ of all copies of each item then our generic n -buyer mechanisms are $\gamma_k\alpha$ -approximation of the optimal n -buyer mechanism where $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$. In Appendix A, we present improved n -buyer mechanisms for some of the models recently considered in the literature. For each model we construct a 1-buyer mechanism and then apply our generic expansion to convert it to an n -buyer mechanism. All of the missing proofs are provided in Appendix E.

³Throughout this paper we always assume that valuations of different buyers are independent

2 Model and Main Results

In this section we present a summary of our main results. We start by formally defining our model and the problem.

Definition 1 (Model and Problem Definition). *There are n buyers and m types of items. We use N to denote the set of buyers and M to denote the set of item types. We have k_j copies of each item $j \in M$ and we assume that no buyer needs more than one copy of each item (in Appendix D, we show that the more general model, in which each buyer may demand more than one copy of each item but no more than $\frac{1}{k}$ of all copies of an item, can be reduced to this simpler model). Let $k = \min_j k_j$. We assume that valuations of different buyers are distributed independently (not necessarily identically). However, we make no assumption on how the valuations of an individual buyer for different items or bundles are distributed. Furthermore, we do not make any assumption on the space of valuations of each buyer (e.g., it could be submodular, additive, etc). Moreover, we allow each individual buyer to have feasibility constraints (e.g., budget, matroid, etc). The problem we want to solve is to design an n -buyer Bayesian combinatorial auction mechanism for maximizing a given objective (e.g., welfare, revenue, etc). We assume that the objective that we want to maximize is separable over buyers, i.e., it can be written as the sum of the objective values obtained from each buyer. For example both revenue and welfare are separable objectives.*

For our generic expansion, we need certain kinds of 1-buyer mechanisms that allow us to specify an upper bound on the probability of allocating each item. We call such a 1-buyer mechanism a “primary mechanism” which we formally define next.

Definition 2 (Primary Mechanism). *A primary mechanism is a 1-buyer mechanism with the extra feature that it allows the auctioneer to specify an upper bound on the probability of allocating each item. We use the notation $\mathcal{M}(\bar{q})$ where $\bar{q} \in [0, 1]^m$ to denote a primary mechanism \mathcal{M} when restricted to allocate each item j with probability at most \bar{q}_j . Note that primary mechanisms can be defined for any separable objective and under any set of feasibility constraints just like ordinary mechanisms. We define an α -approximate primary mechanism to consist of a primary mechanism \mathcal{M} and a concave benchmark function $R(\bar{q}) : [0, 1]^m \rightarrow \mathbb{R}_+$. The benchmark function gives an upper bound on the expected objective value of the optimal primary mechanism when it is restricted to allocate each item j with probability at most \bar{q}_j . \mathcal{M} achieves at least α fraction of the benchmark in expectation.*

Note that we require the benchmark functions to be concave. We justify this in Theorem 5 by showing that the expected objective value of the optimal primary mechanism (both for welfare and revenue) is always concave in \bar{q} . The following informal theorem summarizes our main result on generic construction of n -buyer mechanisms using primary mechanisms.

Theorem 1 (Market Expansion). *If we have an α -approximate truthful primary mechanism \mathcal{M}_i for each buyer $i \in N$ and with some further assumptions (explained next) we can construct a generic $\gamma_k \alpha$ -approximate n -buyer truthful mechanism ($\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$) using $\mathcal{M}_1, \dots, \mathcal{M}_n$ as black boxes.*

Note that we do not require the primary mechanisms for different buyers to be of the same kind. This allows the flexibility of having different classes of buyers in the same auction (e.g., one buyer might have additive valuations with a hard budget constraint while another one might be unit demand with no budget constraints). Below, we present a formal explanation of the assumptions and the result of Theorem 1. Suppose for each buyer i we have an α -approximate truthful (in expectation) primary mechanism \mathcal{M}_i , then each one of the following theorems specify a separate

set of assumptions that is sufficient to apply one of our generic expansions (the terms BIC, DSIC, AIP, etc. are defined at the end of this section).

- Theorem 7: Assuming that for any $\bar{q} \in [0, 1]^m$ the exact probability by which each $\mathcal{M}_i(\bar{q})$ allocates each item can be computed (remember that \bar{q} is just an upper bound), and assuming valuations of each buyer can be represented as a weighted rank function of a matroid (see Def. 11), the BIC-expansion mechanism Mech. 1 is an n -buyer BIC mechanism which is a $\gamma_k \alpha$ -approximation of the optimal n -buyer BIC mechanism.
- Theorem 8: Assuming the benchmark functions $R_i(\cdot)$ are submodular (see Def. 12) the DSIC-expansion mechanism Mech. 2 is an n -buyer DSIC mechanism which is a $\gamma_k \alpha$ -approximation of the optimal n -buyer BIC mechanism.
- Theorem 9: Assuming each \mathcal{M}_i is an α -approximate primary IP mechanism that approximates the optimal primary IP mechanism (as opposed to approximating the optimal primary mechanism), and assuming the benchmark functions $R_i(\cdot)$ are submodular, the DSIC-expansion mechanism Mech. 2 is an SIP mechanism that is a $\gamma_k \alpha$ -approximation of the optimal n -buyer AIP mechanism.

We should emphasize that when each primary mechanism approximates the optimal (randomized) primary mechanism, our generic n -buyer mechanism approximates the optimal (randomized) n -buyer mechanism. However, when each primary mechanism approximates the optimal (randomized) primary IP mechanism, our generic n -buyer mechanism approximates the optimal (randomized) AIP mechanism.

To illustrate the applicability of our generic expansions, we consider some of the more popular models from the literature and present an improved n -buyer mechanism for each one. For each model we construct a primary mechanism and then apply one of our generic expansions to convert it to an n -buyer mechanism. The models considered along with the resulting mechanisms are listed in Table 1. We construct the primary mechanisms for these models in Appendix A.

Primary Mech/Expanded by	Setting	Type	Approximation
A.1 Mech. 3 / Mech. 1	additive correlated valuations with polynomial number of types, budget, capacity, revenue or welfare	general(BIC)	γ_k of optimal BIC
A.2 Mech. 4/ Mech. 2	unit demand, single item(multi unit), budget, revenue	SIP	γ_k of optimal AIP
A.3 Mech. 5 / Mech. 2	additive valuations, product distribution, budget, revenue	SIP	$\gamma_k(1 - \frac{1}{e})$ of optimal AIP
A.4 Mech. 6/ Mech. 2	unit demand, product distribution (regular), revenue	SIP	$\frac{1}{2}\gamma_k$ of optimal AIP

Table 1: Summary of mechanisms constructed using our generic expansion.

Next, we present some definitions that will be used throughout the rest of the paper:

Definition 3 (Bayesian Incentive Compatible (BIC)). *A mechanism is BIC iff for every agent truth telling maximizes the expected payoff where the expectation is taken over the types of other agents and random choices of the mechanism.*

Definition 4 (Dominant Strategy Incentive Compatible (DSIC)). *A mechanism is DSIC iff for every agent truth telling maximizes the expected payoff for any possible reports of other agents where the expectation is taken over random choices of the mechanism.*

Definition 5 (Asymmetric Item Pricing (AIP)). *A combinatorial auction mechanism is AIP iff it can be interpreted as a mechanism of the following form: The mechanism collects the reports (types) of all buyers and computes the allocations/payments. However, for each buyer, it also computes an assignment of prices to individual items using only the reports of other buyers and the random choices of the mechanism such that the allocation/payment that was computed for this buyer would correspond to an optimal bundle for this buyer according to these prices. Note that AIP mechanisms could be randomized. Furthermore, to accommodate buyers that have **budget constraints**, an AIP mechanism offers a lottery option for each item. The lottery option works as follow: a buyer can pay a fraction of the price of an item and then receive the item with probability proportional to the paid fraction of the price ⁴. Note that an AIP mechanism with budget constrained buyers does not need to explicitly ask buyers whether they want to choose a lottery option, instead it is only required that the outcome of the mechanism can be interpreted in such a way (i.e., the mechanism may only output the final allocations/payments).*

Definition 6 (Item Pricing (IP)). *A combinatorial auction mechanism is IP iff it is AIP and the assigned item prices are the same from the perspective of every buyer. Note that AIP and IP are the same for primary mechanisms.*

Definition 7 (Sequential Item Pricing (SIP)). ⁵ *A mechanism is SIP iff it is AIP and it can be implemented in the following form: The mechanism visits the buyers in some order which is not controlled by the mechanism and for each buyer it posts prices for items where these prices are computed based on the prior information, the outcomes of previously visited buyers and the random choices of the mechanism. Each buyer buys her optimal bundle based on the prices offered to her.*

Note that the space of SIP mechanisms is a strict subset of the space of AIP mechanisms. For example, when all buyers are unit demand, the VCG mechanism is a deterministic AIP mechanism however it cannot be implemented as an SIP mechanism. Observe that every AIP mechanism is DSIC.

3 Related Work

Until recently, in the CS literature, most of the work on Bayesian mechanism design for objectives other than welfare had been focused on revenue maximization in single dimensional settings. Many of them study mechanisms that have simple implementation and approximate the Myerson’s mechanism [Mye81] (e.g., [BR89, BLP06, BH08, HR09, DRY10]). Following the recent work of [CHMS10] and [BGM10], there has been a trend toward designing approximation mechanisms with simple implementation for maximizing revenue in multi-dimensional settings. [CHMS10] presents several sequential posted pricing mechanisms for various settings with different types of matroid feasibility constraints. Their mechanisms are sequential both in buyers and in items and for multidimensional settings they assume a worst case ordering on the sequence of offers. [CMM11] considers various settings with budget constraints. For selling k units of an item to unit demand buyers [CEDG⁺10] and [Yan11] present sequential posted pricing mechanisms that obtain $1 - \frac{1}{\sqrt{2\pi k}}$ fraction of the revenue of the optimal mechanism. Both of these mechanisms compute a reserve price for each buyer and then serve them in decreasing order of reserve prices. It is crucial for these mechanisms

⁴Note that a rational buyer with no budget constraints never chooses the lottery option. Furthermore, a buyer with submodular valuations and a budget constraint chooses the lottery option on at most one item, i.e., on the last item for which she runs out of budget.

⁵Note that we use the term “*sequential posted pricing*” as a generic term. We use SIP to denote a specific form of sequential posted pricing in which the order is not controlled or even known in advance by the mechanism

to control the order in which buyers are served otherwise their revenue could be arbitrarily bad. That is because the mechanism has less chance to offer an item to a buyer that is served later. Unfortunately, their approach cannot be extended to multiple (heterogenous) items as it may no longer be possible to find an ordering of buyers such that for every item the reserve prices be decreasing. In general, the main problem with mechanisms that are sequential in buyers is that, as the mechanism allocates items to buyers, the probability that it can sell to buyers that arrive later decreases. We show that this problem can be addressed by a randomized strategy that ensures that every buyer will have the same chance of being offered with an item regardless of their arrival order. By using the same randomized strategy we obtain an improved bound on the prophet inequality for the sum of k -choices. Prophet inequalities have been extensively studied in the past (e.g. [HK92]). The best known bound for generalization to sum of k -choices was $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ by [HKS07] which we improve to $1 - \frac{1}{\sqrt{k+3}}$.

4 The Magician’s Problem and Prophet Inequalities

In this section, we present an abstract online stochastic toy problem and also a near-optimal online algorithm for it. This algorithm is the main ingredient for connecting single buyer mechanisms together to form n -buyer mechanisms in section 5. We also use it to construct a gambler and prove a generalized prophet inequality. It also has direct applications in online stochastic optimization (see Appendix C). Roughly speaking, it can be used to construct online rounding algorithms for online stochastic optimization problems in which online decisions are guided by an offline solution for the expected instance. The following is the abstract description of the toy problem:

Definition 8 (The Magician’s Problem). *A magician is presented with a series of boxes one by one. There is a prize hidden in one of the boxes. He has k magic wands that can be used to open the boxes. On each box is written a probability. If a wand is used on a box, it opens, but with at most the written probability the wand breaks. Let q_i denote this probability for the i^{th} box. The magician wishes to maximize the probability of obtaining the prize, but unfortunately the sequence of boxes, written probabilities, and the box in which the prize is hidden have been arranged by a villain and the magician has no prior information about them (not even the number of boxes). It is given that $\sum_i q_i \leq k$ and that the villain cannot make any changes once the game has started.*

Note that the magician could fail to open a box either because he chose to skip the box or because he ran out of magic wands before coming to the box. Therefore if he greedily tries to open all boxes then the probability of having a magic wand left for the last box is minimized and that is where the villain would put the prize. Observe that the villain puts the prize in the box that has the minimum ex-ante probability of being opened. Therefore, in order for the magician to obtain the prize with probability at least γ , he has to devise a strategy that guarantees an ex-ante probability of at least γ for opening each box. Note that the nature of the prize or even whether there is actually a prize does not affect the problem. It is easy to see the following strategy ensures an ex-ante probability of at least $\frac{1}{4}$ for opening each box: For each box randomize and use a wand with probability $\frac{1}{2}$. But can we do better? By using standard concentration bounds, we can construct an algorithm that makes a random decision for each box independently and ensures an ex-ante probability of $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ for opening each box and this is tight if we make an independent decision for each box.⁶ Next we present an algorithm that can ensure an ex-ante probability of

⁶In fact, it can be shown that for any constant c if for each box we independently randomize and choose to open the box with probability $1 - \frac{c}{\sqrt{k}}$ then there are family of instances in which with a constant positive probability (constant in k) we run out of magic wands before getting to the last box.

at least $1 - \frac{1}{\sqrt{k+3}}$ for opening each box. This algorithm takes a parameter γ and tries to ensure a minimum ex-ante probability of γ for opening each box. In Theorem 2, we show that for any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ this algorithm can indeed ensure that the ex-ante probability of opening each box is at least γ . In section 5 we will use this algorithm extensively in constructing n -buyer mechanisms.

Algorithm 1 (γ -Conservative Magician).

- Construct a strategy table y_{ij} using the dynamic programs, that are presented below, in which γ is a parameter that is given in advance. The strategy table can be constructed incrementally as the boxes are revealed. Upon being presented with box i , if j is the number of magic wands broken so far, do the following:

- If $y_i^j = 1$ then open the box i .
- If $y_i^j = 0$ then discard the box.
- Otherwise randomize and open the box with probability y_i^j .

We use Y_i as the indicator random variable which is 1 iff the magician chooses to open the box i . The strategy table can be computed using the following dynamic programs:

$$y_i^j = \begin{cases} 1 & i \geq 1, \quad 0 \leq j < \theta_i \\ (\gamma - \phi_i^{\theta_i-1}) / (\phi_i^{\theta_i} - \phi_i^{\theta_i-1}) & i \geq 1, \quad j = \theta_i \\ 0 & \text{otherwise.} \end{cases} \quad (DP.y)$$

$$\theta_i = \min\{j \mid \phi_i^j \geq \gamma\} \quad (DP.\theta)$$

$$\phi_i^j = \begin{cases} 1 & i = 1, \quad j = 0 \\ y_{i-1}^j q_{i-1} \phi_{i-1}^{j-1} + (1 - y_{i-1}^j q_{i-1}) \phi_{i-1}^j & i \geq 2, \quad j \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (DP.\phi)$$

Note that computing y_i^j only requires the knowledge of q_1, \dots, q_{i-1} so in fact computing y_i^j and making a decision about box i can be done even before seeing the box.

Interpretation of γ -Conservative Magician(Alg. 1) The main idea of the algorithm is the following: Upon arrival of the i^{th} box, we open the box when we have many wands and discard it when we have few wands and we do this in such a way that ensures the ex-ante probability of opening each box at the time the game starts is at least γ . Formally: let S_i be the random variable that represents the number of magic wands broken prior to seeing the i^{th} box. We identify the smallest integer θ_i for which $Pr[S_i \leq \theta_i] \geq \gamma$. Now if we open the i^{th} box only when we have broken no more than θ_i wands and discard it if we have broken more than θ_i wands, we can guarantee that the ex-ante probability of opening the i^{th} box is at least γ . Furthermore, if $Pr[S_i \leq \theta_i]$ is strictly more than γ , then in the event of $S_i = \theta_i$ we can in fact randomize and choose to open the box i with a probability which is strictly less than 1 and is just enough to ensure that the total ex-ante probability of opening the i^{th} box is γ . It can be verified that ϕ_i^j which is computed by the dynamic program is a lower bound for $Pr[S_i \leq j]$. In fact, if q_1, \dots, q_{i-1} are the exact probabilities of breaking a wand for each of first $i - 1$ boxes then $Pr[S_i \leq j] = \phi_i^j$. In order to prove that the above strategy ensures that each box is opened with an ex-ante probability of at least γ , we need to show that $y_i^j = 0$ for all $j \geq k$ and all i . i.e., we need to show that the strategy table of the magician does not instruct him to open a box if he has broken all of his k wands. In Theorem 2 we present a sufficient condition on γ that ensures $y_i^j = 0$ for all $j \geq k$ and all i .

Theorem 2 (γ -Conservative Magician). *For any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$, a γ -conservative magician guarantees that each box is opened with an ex-ante probability at least γ . Furthermore, if q_i are the exact probabilities of breaking a wand then the γ -conservative magician opens each box with an ex-ante probability exactly γ ⁷*

Definition 9 (γ_k). *We define γ_k to be the largest probability such that for all $k' \geq k$, if $\sum_i q_i \leq k'$ then a γ_k -conservative magician can guarantee that each box is opened with probability at least γ_k . By Theorem 2, we know that $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$*

Observe that γ_k is a non-decreasing function in k which is at least $\frac{1}{2}$ (when $k = 1$) and approaches 1 as k gets larger. The next theorem shows that no other magician can do considerably better than the γ -conservative magician. It also shows that the lower bound of $1 - \frac{1}{\sqrt{k+3}}$ on γ_k is almost tight.

Theorem 3 (Optimal Magician). *For any $\epsilon > 0$, even the optimal magician cannot guarantee that each box is opened with probability at least $1 - \frac{k^k}{e^k k!} + \epsilon$ (i.e., no magician can guarantee it). Note that $1 - \frac{k^k}{e^k k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}$.*

Next, we present a direct application of the conservative magician to a generalization of prophet inequalities.

Definition 10 (Sum of k -Choices). *A sequence of n random variables V_1, \dots, V_n are presented to a gambler one by one in an arbitrary order. The gambler knows n and the distribution of each random variable in advance but not the order in which they are presented. Upon being presented with the random variable V_i , the gambler observes the actual draw of V_i and he has to decide whether to keep it or to discard it and this decision cannot be changed later. The gambler wants to select k of the random draws from the sequence and his objective is to maximize the sum of the selected ones. The prophet knows all the actual draws in advance so he chooses the k highest draws. We assume that the order in which the random variables are presented to the gambler is fixed in advance and does not change during the process.*

[HKS07] showed that there is a gambler that guarantees at least $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ fraction of the payoff of the prophet (in expectation) by using a non-decreasing sequence of k stopping rules (thresholds)⁸. In what follows, we construct a gambler that obtains at least $1 - \frac{1}{\sqrt{k+3}}$ fraction of the prophet's payoff in expectation by using a conservative magician as a black box and using only a single threshold.

Theorem 4 (Sum of k -Choices, Prophet Inequality). *The following gambler obtains at least γ_k fraction of the expected payoff of the prophet. To simplify the exposition assume that the distribution of V_i do not have any point mass⁹*

- Find the threshold τ such that $\sum_i Pr[V_i > \tau] = k$. This can be done by a binary search on τ .
- Use a γ_k -conservative magician with k magic wands. Upon seeing each V_i , create a box and write $q_i = Pr[V_i > \tau]$ on the box and present it to the magician. If the magician chooses to open the box and $V_i > \tau$ then select V_i and break the magician's wand.

⁷In particular the fact that the probability of breaking a wand for the i^{th} box is exactly q_i conditioned on any sequence of prior events implies that for each box the event of breaking a wand has to be independent of the sequence of past events and independent of other boxes.

⁸A gambler with stopping rules τ_1, \dots, τ_k works as follows: When presented with V_i , he selects it iff $V_i \geq \tau_{j+1}$ where j is the number of random draws selected so far.

⁹Our theorem holds with slight modifications if we allow point masses.

5 From One Buyer to Many Buyers

In this section, we present two generic approaches for constructing n -buyer mechanisms by using 1-buyer mechanisms as black boxes. We assume the model that was explained in Def. 1. To apply our generic construction, we assume that we have an α -approximate primary mechanism for each buyer. Remember that according to Def. 2, an α -approximate primary mechanism for buyer i consists of a primary mechanism \mathcal{M}_i and a concave benchmark function $R_i(\cdot)$. $R_i(\bar{q}_i)$ gives an upper bound on the expected objective value of the optimal primary mechanism for buyer i when restricted to allocate each item j with probability at most \bar{q}_{ij} . Furthermore, $\mathcal{M}_i(\bar{q}_i)$ obtains in expectation at least α -fraction of $R_i(\bar{q}_i)$. Notice that we require the benchmark functions to be concave. To justify this requirement we present the following theorem.

Theorem 5. *Consider any convex program of the form (CP_R) in which $u(\cdot)$ is an arbitrary concave function, $g_j(\cdot)$ are arbitrary convex functions and \mathbb{X} is an arbitrary convex set. Let $R(\bar{q})$ denote the optimal objective value of this program as a function of $\bar{q} = (\bar{q}_1, \dots, \bar{q}_m)$. Then $R(\bar{q})$ is concave. Furthermore, both for welfare and for revenue and under any set of feasibility constraints for the buyer (e.g., matroid constraints, budget constraints, etc), the allocation function of the optimal primary mechanism can be interpreted as the solution of a convex program of this form such that $R(\bar{q})$ be the expected objective value of the optimal primary mechanism when restricted to allocate the items with probabilities at most $\bar{q}_1, \dots, \bar{q}_m$.*

$$\begin{aligned} & \text{maximize:} && u(x) && (CP_R) \\ \forall j : & && g_j(x) \leq \bar{q}_j \\ & && x \in \mathbb{X} \end{aligned}$$

In Appendix A we present several approximation primary mechanisms along with their corresponding benchmarks.

Next, we present two approaches for expanding primary mechanisms to n -buyer mechanisms. The two approaches are quite similar however each one requires slightly different assumptions. The basic idea is the following: In both approaches, we initially solve a convex program using the benchmark functions to compute the approximately optimal probability of allocating a copy of item j to buyer i for each i and j which we denote by \bar{q}_{ij} . The optimal objective value of the convex program is an upper bound on the expected objective value of the optimal n -buyer mechanism. Observe that if for each i we run $\mathcal{M}_i(\bar{q}_i)$ on buyer i independently of other buyers, there is a good chance that we allocate more items than we actually have. The two mechanisms that we present differ in the way they address this issue. Mech. 1 runs each primary mechanism independently and then unallocates the over-allocated items at random while ensuring that each allocated item will remain allocated with probability at least γ . Mech. 2 visits buyers in an arbitrary order (not controlled by the mechanism) and it avoids over-allocating items by not offering some of the items to some of the buyers at random while ensuring that each item can be presented to each buyer with probability at least γ . The following is the convex program:

$$\begin{aligned} & \text{maximize:} && \sum_i R_i(\bar{q}_i) && (CP_{R_N}) \\ \forall j : & && \sum_i \bar{q}_{ij} \leq k_j \\ \forall i, \forall j : & && \bar{q}_{ij} \geq 0 \end{aligned}$$

Theorem 6. *The optimal objective value of (CP_{R_N}) is an upper bound on the expected objective value of the optimal n -buyer BIC mechanism.*

Mechanism 1 (γ -BIC-Expansion).

- (I) *Solve the convex program of (CP_{R_N}) and let \bar{q}_{ij} denote an optimal assignment for it.*
- (II) *For each buyer $i \in 1 \cdots n$: run the corresponding primary mechanism $\mathcal{M}_i(\bar{q}_i)$ on buyer i and let X_{i1}, \dots, X_{im} and P_i denote the random variables for allocation and payment of \mathcal{M}_i (i.e., X_{ij} is an indicator random variable which is 1 iff \mathcal{M}_i allocated a copy of item j to buyer i). Furthermore, let \hat{q}_{ij} be the actual marginal probability of allocating item j to buyer i by $\mathcal{M}_i(\bar{q}_i)$. Note that $\hat{q}_{ij} \leq \bar{q}_{ij}$.*
- (III) *For each item type $j \in 1 \cdots m$:*
 - (a) *Create a new instance of the γ -conservative magician (see Alg. 1) with k_j magic wands. This is the j^{th} magician.*
 - (b) *For each $i \in 1 \cdots n$: create a box corresponding to X_{ij} and write \hat{q}_{ij} on the box and present it to the j^{th} magician. Let Y_{ij} denote the indicator random variable which is 1 iff the magician chooses to open the box. Set $X'_{ij} \leftarrow X_{ij}Y_{ij}$. If $X'_{ij} = 1$ then break the magician's wand.*
- (IV) *For each buyer $i \in 1 \cdots n$: charge buyer i a payment of $P'_i \leftarrow \gamma P_i$ and for each $j \in 1 \cdots m$, allocate a copy of item j to buyer i iff $X'_{ij} = 1$.*

In order for γ -BIC expansion to be a BIC mechanism, the valuations of each buyer should be in the form of a weighted rank function of some matroid. Next, we define this formally:

Definition 11 (Valuations as Matroid Weighted Rank Functions). *Valuations of a buyer for bundles of items can be represented as a weighted rank function of a matroid if there is a matroid whose ground set is the set of items and such that for any bundle S of items:*

- *If S is an independent set of the matroid, then the valuation of the buyer for S is just the sum of her valuations for each item in S .*
- *If S is not an independent set, then the valuation of the buyer for S is equal to her valuation of an independent subset $S' \subset S$ with the maximum valuation.*

*In particular, additive valuations with capacities, unit demand valuations, etc. can be represented as matroid weighted rank functions*¹⁰.

Theorem 7 (γ -BIC-Expansion). *Suppose $\mathcal{M}_1, \dots, \mathcal{M}_n$ are α -approximate truthful (in expectation) primary mechanisms. Assuming that the valuations of each buyer can be represented by a weighted rank function of a matroid and assuming that we can compute the actual marginal probabilities of allocation (i.e. \hat{q}_{ij} in Mech. 1) for each primary mechanism \mathcal{M}_i , then for any parameter $\gamma \in [0, \gamma_k]$ the γ -BIC expansion (Mech. 1) is a BIC mechanism which is a $\gamma \cdot \alpha$ -approximation of the optimal n -buyer BIC mechanism.*

Note that by Theorem 2, $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$ so the above theorem implies that by using a γ -BIC-expansion mechanism with $\gamma = 1 - \frac{1}{\sqrt{k+3}}$, we get at least $(1 - \frac{1}{\sqrt{k+3}})\alpha$ fraction of the expected objective value of the optimal n -buyer BIC mechanism.

¹⁰note that budget constraints are not part of the valuations.

The assumption of being able to compute the actual marginal probability of allocation for each \mathcal{M}_i might be a strong requirement. Our next mechanism does not require neither this assumption nor does it require the valuations to be in the form of matroid weighted rank functions. However it requires a submodularity condition on the benchmark functions. The mechanism is as follows:

Mechanism 2 (γ -DSIC-Expansion).

- (I) Solve the convex program of (CP_{R_N}) and let \bar{q}_{ij} denote an optimal assignment for it.
- (II) For each item $j \in 1 \cdots m$: create an instance of γ -conservative magician (see Alg. 1) with k_j magic wands.
- (III) For each buyer $i \in 1 \cdots n$:
 - (a) For each $j \in 1 \cdots m$: write \bar{q}_{ij} on a box and present it to the j^{th} magician. Let Y_{ij} denote the indicator random variable which is 1 iff the magician opens the box. Set $\bar{q}'_{ij} \leftarrow \bar{q}_{ij} Y_{ij}$.
 - (b) Run the mechanism $\mathcal{M}_i(\bar{q}'_i)$ on buyer i and use its outcome as the final outcome for buyer i . Furthermore, let X_{i1}, \dots, X_{im} denote the indicator random variables for allocation of $\mathcal{M}_i(\bar{q}'_i)$ (i.e., X_{ij} is 1 iff \mathcal{M}_i allocated a copy of item j to buyer i).
 - (c) For each $j \in 1 \cdots m$: if $X_{ij} = 1$ then break the wand of the j^{th} magician.

Definition 12 (Submodularity). A set function $f : \{0, 1\}^m \rightarrow \mathbb{R}$ is submodular iff $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$. However, throughout this paper we use the following extended definition of submodularity: A function $f : [0, 1]^m \rightarrow \mathbb{R}$ is submodular iff the function $\mathbb{S}_f : \{0, 1\}^m \times [0, 1]^m \rightarrow \mathbb{R}$, as defined next, is submodular in its first argument. We define $\mathbb{S}_f(S, x) = f(x')$ in which for each j : if $j \in S$ then $x'_j = x_j$ else $x'_j = 0$.

Theorem 8 (γ -DSIC-Expansion). Suppose $\mathcal{M}_1, \dots, \mathcal{M}_n$ are α -approximate truthful (in expectation) primary mechanisms. If all benchmark functions $R_i(\cdot)$ are submodular then for any parameter $\gamma \in [0, \gamma_k]$ the γ -DSIC expansion (Mech. 2) is a DSIC mechanism which is a $\gamma \cdot \alpha$ -approximation of the optimal n -buyer BIC mechanism.

Observe that to use γ -DSIC expansion we do not need to compute the exact marginal probability of allocation for each \mathcal{M}_i and we do not need to make any assumption on the valuations of buyers. The only requirement is that the benchmark functions be submodular.

An interesting implication of Theorem 8 is the following theorem.

Theorem 9. Suppose $\mathcal{M}_1, \dots, \mathcal{M}_n$ are α -approximate primary IP mechanisms that approximate the optimal primary IP mechanisms for the corresponding buyers. If all benchmark functions $R_i(\cdot)$ are submodular then for any parameter $\gamma \in [0, \gamma_k]$ the γ -DSIC expansion (Mech. 2) is a SIP mechanism which is a $\gamma \cdot \alpha$ -approximation of the optimal n -buyer AIP mechanism.

Observe that the space of primary AIP mechanisms collapses to the space of primary IP mechanisms because a primary mechanism considers only a single buyer. Note that a primary IP mechanism simply computes a distribution of prices for items and draws the prices from that distribution. On the other hand, if we expand primary IP mechanisms to an n -buyer mechanism using the γ -DSIC expansion, the resulting mechanism is SIP but it approximates the optimal AIP mechanism.

Remark 1. The γ -DSIC expansion (Mech. 2) assumes no control or prior information about the order in which buyers are visited. The order specified in the mechanism is arbitrary and could be replaced by any other ordering which does not even need to be known to the mechanism.

6 Conclusion

In this paper, for Bayesian combinatorial auctions, we presented a reduction from n -buyer mechanisms to 1-buyer mechanisms. This shows that the inherent difficulty of designing Bayesian mechanisms in multidimensional settings does not stem from the problem of making coordinated decisions for all buyer, but instead it stems from the difficulty of aligning the behavior of the mechanism with the incentives of each individual buyer even in the absence of other buyers.

7 Acknowledgment

I thank Jason Hartline for many helpful suggestions and comments and Azarakhsh Malekian for a fruitful discussion that lead to the proof of the sand theorem (Theorem 17) and many other suggestions. The current exposition of the magician’s problem (Def. 8) was suggested by Jason Hartline, the original exposition used a gambler with k tokens where the tokens could be used to open the boxes.

References

- [Arm96] Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75, January 1996.
- [BGGM10] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget constrained auctions with heterogeneous items. In *STOC*, pages 379–388, 2010.
- [BH08] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce*, page 49, 2008.
- [BLP06] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and implementation in undominated strategies. In *SODA*, pages 1054–1063, 2006.
- [BR89] Jeremy Bulow and John Roberts. The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–90, October 1989.
- [CEDG⁺10] Tanmoy Chakraborty, Eyal Even-Dar, Sudipto Guha, Yishay Mansour, and S. Muthukrishnan. Approximation schemes for sequential posted pricing in multi-unit auctions. In *WINE*, 2010.
- [CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010.
- [CMM11] Shuchi Chawla, David Malec, and Azarakhsh Malekian. Bayesian mechanism design for budget-constrained agents. In *ACM Conference on Electronic Commerce*, 2011.
- [DRY10] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138, 2010.
- [GKS02] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy 0002, and Aravind Srinivasan. Dependent rounding in bipartite graphs. In *FOCS*, pages 323–332, 2002.

- [HK92] Theodore P. Hill and Robert P. Kertz. A survey of prophet inequalities in optimal stopping theory. In *Contemporary Mathematics*, 1992.
- [HKM11] Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility and matchings. In *SODA*, 2011.
- [HKS07] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, pages 58–65, 2007.
- [HR09] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234, 2009.
- [Mye81] Roger B. Myerson. Optimal auction design. 1981.
- [Yan11] Qiqi Yan. Mechanism design via correlation gap. In *SODA*, 2011.

A Primary Mechanisms

In this section, we present several approximation primary mechanisms. Note that a primary mechanism as defined in Def. 2 only considers a single buyer. Once we construct a primary mechanism, we can use one of the generic expansions of section 5 to convert it to a mechanism for many independent buyers. Except for subsection A.1, for the rest of this section we restrict the space of mechanisms to AIP mechanisms. Note that a primary AIP mechanism is simply an IP mechanism. In other words, a primary AIP mechanism simply chooses a price distribution irrespective of the type of the buyer and then draws the prices of items from that distribution and offers them to the buyer. Since for primary mechanisms AIP and IP are the same we will simply refer to these primary mechanisms as primary IP mechanisms. However, by Theorem 9, the γ -DSIC expansion of α -approximate primary IP mechanisms is a SIP mechanism that is a $\gamma \cdot \alpha$ approximation of the optimal n -buyer AIP mechanism.

A.1 Correlated Valuations with Capacity and Hard Budget Constraints

In this subsection, we consider a single buyer with correlated valuations for different items and with polynomially bounded number of types. For each possible type t of buyer, let v_{tj} denote her valuation for item j . Also let $f(t)$ denote the probability that the buyer is of type t . Furthermore, suppose that the buyer has a total budget of B and is interested in at most C items. We assume that the only private information of the buyer is her type and everything else is publicly known. Note that this is exactly the setting considered in [BGGM10]. Next, we present a 1-approximate (i.e., optimal) truthful primary mechanism for maximizing revenue in this setting.

Consider the following LP in which x_{tj} is the probability of allocating item j when the buyer has reported type t and p_t is her payment. The optimal objective value of this LP is an upper bound on the revenue of the optimal primary mechanism when it is restricted to allocate each item j with probability at most \bar{q}_j :

$$\begin{aligned}
& \text{maximize:} && \sum_t f(t)p_t && (LP_{rev}) \\
\forall j : & && \sum_t f(t)x_{tj} \leq \bar{q}_j \\
\forall t : & && \sum_j x_{tj} \leq C \\
\forall t, t' : & && \sum_j v_{tj}x_{tj} - p_t \geq \sum_j v_{t'j}x_{t'j} - p_{t'} \\
\forall t, \forall j : & && x_{tj} \in [0, 1] \\
\forall t : & && p_t \in [0, B]
\end{aligned}$$

We construct the primary mechanism as follows:

Mechanism 3.

- Define the benchmark $R(\bar{q})$ to be the optimal objective value of (LP_{rev}) as a function of $\bar{q} = (\bar{q}_1, \dots, \bar{q}_m)$.
- Given $\bar{q}_1, \dots, \bar{q}_m$, solve the linear program of (LP_{rev}) to compute x_{tj} and p_t .
- If the buyer reported her type as t then charge her a payment of p_t and allocate each item j with probability x_{tj} as follows: Use the dependent randomize rounding algorithm of [GKS02] to round each x_{tj} to either 0 or 1 such that if X_j is the result of rounding the x_{tj} then $E[X_j] = x_{tj}$ and such that $\sum_j X_j \leq C$. Then, for each j allocate a copy of item j to the buyer iff $X_j = 1$.

Theorem 10. *The primary mechanism Mech. 3 is a 1-approximate truthful primary mechanism for revenue and it satisfies all requirements of γ -BIC expansion.*

Proof. The proof of truthfulness and optimality of Mech. 3 trivially follows from the (LP_{rev}) . So, we only focus on proving that this mechanism satisfies the requirements of Theorem 7 for γ -BIC expansion. First, observe that the benchmark function, $R(\bar{q})$, is concave (this follows by applying Theorem 5). Second, observe that the valuations of the buyer can be represented as a weighted rank function of a uniform matroid of rank C . Third, notice that given $\bar{q}_1, \dots, \bar{q}_m$, we can compute the exact marginal probabilities of allocation, i.e. $\hat{q}_1, \dots, \hat{q}_m$ as follows: $\hat{q}_j = \sum_t x_{tj}$. So the mechanism Mech. 3 and its associated benchmark satisfy the requirements of Theorem 7 for γ -BIC expansion. \square

Observe that if we replace the objective function of (LP_{rev}) with $\sum_{t,j} f(t)v_{tj}x_{tj}$ we get a 1-approximate truthful primary mechanism for welfare instead of revenue.

A.2 Single-Item with Hard Budget Constraints

In this subsection, we consider a single buyer who is interested in at most one copy of an indivisible item and has a publicly known budget B . Her only private information is her valuation for the item which is drawn from a publicly known distribution with CDF $F(\cdot)$. To avoid complicating the

proofs, we assume that $F(\cdot)$ is continuous and strictly increasing in its domain¹¹. We restrict the space of mechanisms to IP mechanisms. Next, we present a 1-approximate (i.e., optimal) primary IP mechanism for maximizing revenue in this setting. We start by defining the modified CDF function $F^B(\cdot)$ as follows:

$$F^B(v) = \begin{cases} F(v) & v \leq B \\ 1 - (1 - F(v))\frac{B}{v} & v \geq B \end{cases} \quad (F^B)$$

Intuitively, $1 - F^B(p)$ is the probability of allocating the item to the buyer if we offer the item at price p . Note that the buyer only buys if her valuation is more than p which happens with probability $1 - F(p)$. If $p > B$ then she will pay her whole budget and only get the item with probability $\frac{B}{p}$. Observe that if we want to allocate the item with probability q we can offer a price of $F^{B(-1)}(1 - q)$ in which case it can be easily verified that we get a revenue of $q \cdot F^{B(-1)}(1 - q)$ in expectation. Now define the function $H(q) = q \cdot F^{B(-1)}(1 - q)$ and let $\widehat{H}(q)$ denote its concave closure (i.e., the smallest concave function that is an upper bound on $H(q)$ for every q). We address the problem of efficiently computing $\widehat{H}(q)$ later in Lem. 1. Next, we show that the revenue of the optimal primary IP mechanism when restricted to allocate the item with probability at most \bar{q} is no more than the optimal objective value of the following convex program:

$$\begin{aligned} \text{maximize:} \quad & \widehat{H}(q) && (CP_{rev-1}) \\ & q \leq \bar{q} \\ & q \geq 0 \end{aligned}$$

Theorem 11. *The revenue of the optimal primary IP mechanism, when restricted to allocate the item with probability at most \bar{q} , is equal to the optimal objective value of the convex program (CP_{rev-1}) . Furthermore, assuming that q^* is the optimal assignment for the convex program, if $\widehat{H}(q^*) = H(q^*)$ then the optimal primary IP mechanism uses a single price $p = F^{B(-1)}(1 - q^*)$ otherwise, it randomizes between two prices but the probability of allocation is still q^* .*

Proof. First we prove that the expected revenue of the optimal primary IP mechanism which we denote by R^* is upper bounded by $\widehat{H}(q^*)$. Then we construct a distribution over prices that obtains this revenue. Note that any primary IP mechanism can be specified as a distribution over prices. Let \mathcal{P} be the optimal price distribution. So $R^* = E_{p \sim \mathcal{P}}[p(1 - F^B(p))]$. Note that every price p corresponds to an allocation probability $q = 1 - F^B(p)$. So any probability distribution over p can be specified as a probability distribution over q . Let \mathcal{Q} denote the probability distribution over q that corresponds to price distribution \mathcal{P} , then we can write $R^* = E_{q \sim \mathcal{Q}}[q \cdot F^{B(-1)}(1 - q)] = E_{q \sim \mathcal{Q}}[H(q)] \leq E_{q \sim \mathcal{Q}}[\widehat{H}(q)]$ also notice that by Jensen's inequality this is less than or equal to $\widehat{H}(E_{q \sim \mathcal{Q}}[q])$. Note that $E_{q \sim \mathcal{Q}}[q]$ is exactly the probability of allocating the item if we use the price distribution \mathcal{P} so it must be no more than \bar{q} which implies that $E_{q \sim \mathcal{Q}}[q]$ is a feasible assignment for (CP_{rev-1}) and therefore $\widehat{H}(E_{q \sim \mathcal{Q}}[q]) \leq \widehat{H}(q^*)$ which completes the first part of the proof. Next, we construct the optimal price distribution: If $\widehat{H}(q^*) = H(q^*)$ then the optimal price distribution is just a single price $p = F^{B(-1)}(1 - q^*)$. Otherwise, by definition of concave closure, there are two points q^- and q^+ and $\theta \in (0, 1)$ such that $q^* = \theta q^- + (1 - \theta)q^+$ and $\widehat{H}(q^*) = \theta H(q^-) + (1 - \theta)H(q^+)$. In this case the optimal price distribution offers price p^- with probability θ and offers price p^+ with probability $1 - \theta$. \square

¹¹The proofs can be modified to work without this assumption.

We construct the optimal revenue maximizing primary IP mechanism as follows:

Mechanism 4.

- Define the benchmark $R(\bar{q})$ to be the optimal objective value of (CP_{rev-1}) as a function of \bar{q} .
- Given \bar{q} , solve the convex program of (CP_{rev-1}) to compute the optimal q .
- Using the optimal q compute the optimal price as follows: if $\hat{H}(q) = H(q)$ then offer the single price $p = F^{B(-1)}(1 - q)$ otherwise randomize between two prices p^- and p^+ with probabilities θ and $1 - \theta$ as explained in the proof of Theorem 11.

Theorem 12. *Mech. 4 is the optimal revenue maximizing primary IP mechanism. Furthermore, this mechanism satisfies the requirements of DSIC-expansion.*

Proof. The proof of the optimality of the mechanism follows from Theorem 11. Furthermore, the benchmark function, $R(\bar{q})$, is concave (this follows by applying Theorem 5) and it is trivially submodular and so it meets the requirements of DSIC expansion. \square

Now we get back to the problem of efficiently computing $\hat{H}(\cdot)$:

Lemma 1. *A $(1 + \epsilon)$ -approximation of $\hat{H}(\cdot)$ which we denote by $\hat{H}_{1+\epsilon}(\cdot)$ can be constructed using a piece-wise linear function with $\ell = \frac{\log L}{\log(1+\epsilon)}$ pieces and in time $O(\ell \log \ell)$ in which L is the ratio of the maximum valuation to minimum non-zero valuation. Note that we need at least $\log_2 L$ bits just to represent such valuations so this construction is polynomial in input size for constant ϵ .*

Proof. WLOG, assume that all possible non-zero valuations of the buyer are in range $[1, L]$. Let $\ell = \lfloor \frac{\log L}{\log(1+\epsilon)} \rfloor$. For $r = 0 \dots \ell$, consider the prices $p_r = (1 + \epsilon)^{\ell - r}$ and compute the corresponding $q_r = 1 - F^B(p_r)$. Construct $\hat{H}_{1+\epsilon}(\cdot)$ by constructing the convex hull of the points: $(0, 0), (q_1, p_1 q_1), (q_2, p_2 q_2), \dots, (q_\ell, p_\ell q_\ell), (1, 0)$. This can be done in time $O(\ell \log \ell)$. Note that $p = F^{B(-1)}(1 - q)$ is a decreasing function of q so at every point $q \in [q_r, q_{r+1}]$, the corresponding price is $F^{B(-1)}(q) \in [p_{r+1}, p_r]$ but $p_r = (1 + \epsilon)p_{r+1}$ therefore at every point q : $H_{1+\epsilon}(q) \leq \hat{H}(q) \leq (1 + \epsilon)H_{1+\epsilon}(q)$ which completes the proof. \square

Remark 2. *In order to use $\hat{H}_{1+\epsilon}(\cdot)$ in Mech. 4, we need to use $(1 + \epsilon)\hat{H}_{1+\epsilon}(\cdot)$ in the objective function of the (CP_{rev-1}) instead of $\hat{H}(\cdot)$ for computing the benchmark. Furthermore, the mechanism will be a $(1 - \epsilon)$ -approximation of the optimal primary IP mechanism. Also notice that finding p^- and p^+ by using the $\hat{H}_{1+\epsilon}(\cdot)$ is trivial.*

A.3 Additive Independent Valuations with Hard Budget Constraints

In this subsection, we consider m indivisible items and a single buyer with a publicly known budget B who has additive valuations for bundles of items (i.e., her valuation for a bundle of items is the sum of her valuations for individual items in the bundle). We assume that for each item j , the buyer's valuation is drawn independently from a publicly known distribution with CDF $F_j(\cdot)$. To avoid complicating the proofs, we assume that each $F_j(\cdot)$ is continuous and strictly increasing in its domain¹². Next, we present a $(1 - \frac{1}{e})$ -approximate primary IP mechanism for maximizing revenue in this setting. As in subsection A.2, we start by defining the modified CDF function $F_j^B(\cdot)$ for each item j as follows:

¹²The proofs can be modified to work without this assumption.

$$F_j^B(v) = \begin{cases} F_j(v) & v \leq B \\ 1 - (1 - F_j(v)) \frac{B}{v} & v \geq B \end{cases} \quad (F_j^B)$$

Furthermore, for each item j , let $H_j(q) = q \cdot F_j^{B-1}(1 - q)$ and let $\widehat{H}_j(\cdot)$ be its concave closure as in subsection A.2. Similarly, for each j , define $R_j(\bar{q}_j)$ to be the optimal objective value of the following convex program as a function of \bar{q}_j .

$$\begin{aligned} \text{maximize:} \quad & \widehat{H}_j(q) && (CP_{rev-j}) \\ & q \leq \bar{q}_j \\ & q \geq 0 \end{aligned}$$

The next theorem presents an upper bound on the revenue of the optimal primary IP mechanism.

Theorem 13. *The revenue of the optimal primary IP mechanism, when restricted to allocate items with probabilities at most $\bar{q}_1, \dots, \bar{q}_m$, is no more than $\min(\sum_j R_j(\bar{q}_j), B)$.*

Proof. For any j if we were only to sell the item j , by Theorem 11, the maximum revenue we could get using an IP primary mechanism would be $R_j(\bar{q}_j)$. Now observe that if we compute the optimal price distribution for each item separately, we might only get less revenue because the budget is shared among all items and the buyer might not be able to buy some of the items that she would otherwise buy if there were no other items. That means the actual probability of allocating each item j would then be less than \bar{q}_j . So the optimal joint price distribution might sell at lower prices but the extra revenue will only come from buyers of lower type who were excluded by the optimal primary mechanism of each individual item so the overall revenue from each item j cannot be more than $R_j(\bar{q}_j)$. Finally, observe that the expected revenue of the mechanism cannot be more than B so $\min(\sum_j R_j(\bar{q}_j), B)$ is an upper bound on the revenue of the optimal primary IP mechanism. \square

Next, we present a $(1 - \frac{1}{e})$ -approximate primary IP mechanism for maximizing revenue.

Mechanism 5.

- Define the benchmark $R(\bar{q}) = \min(\sum_j R_j(\bar{q}_j), B)$.
- Given \bar{q} , solve the convex program of (CP_{rev-j}) for each item j . Let q_j denote the optimal assignment for the convex program of item j .
- Using the optimal q_1, \dots, q_n compute the prices as follows: for each item j if $\widehat{H}_j(q_j) = H_j(q_j)$ then offer the single price $p_j = F_j^{B(-1)}(1 - q_j)$ otherwise randomize between two prices p_j^- and p_j^+ with probabilities θ_j and $1 - \theta_j$, as explained in the proof of Theorem 11. Note that the randomization is done for each item independently.

Theorem 14. *Mech. 5 obtains at least $1 - \frac{1}{e}$ fraction of the revenue of the optimal primary IP mechanism. Furthermore, this mechanism satisfies the requirements of DSIC-expansion.*

Proof. First, we show that Mech. 5 obtains at least $1 - \frac{1}{e}$ fraction of its benchmark $R(\bar{q})$ which by Theorem 13 is an upper bound on the revenue of the optimal primary IP mechanism. Consider an imaginary replica of the original buyer who has exactly the same valuations as the original buyer but has a separate budget B for each item. We call this imaginary buyer the “jumbo replica”.

Furthermore, suppose that any payment received from the jumbo replica beyond B is lost (i.e., if the jumbo replica pays Z the mechanism receives only $\min(Z, B)$). Observe that for any assignment of prices, the payment received from the original buyer and the payment received from the jumbo replica are exactly the same because if the original buyer has't hit his budget limit then both the original buyer and the jumbo replica will buy the same items and pay the exact same amount. Otherwise, if the original buyer hits his budget limit, then the mechanism receives exactly B from both the original buyer and the jumbo replica. So we only need to show that the revenue received from the jumbo replica by using the price distribution of Mech. 5 is at least $1 - \frac{1}{e}$ of $R(\bar{q})$. Observe that from the view point of the jumbo replica there is no connection between different items so he makes a decision for each item independently. Let Z_j be the random variable that denotes the amount paid by the jumbo replica for item j using the price distribution of Mech. 5. By Theorem 11, we know that $E[Z_j] = R_j(\bar{q}_j)$ and the total revenue received by the mechanism is $Z = \min(\sum_j Z_j, B)$. Notice that Z_1, \dots, Z_m are independent random variables in the range $[0, B]$. By applying the Lem. 2, we can conclude that $E[\min(\sum_j Z_j, B)] \geq (1 - \frac{1}{e}) \min(\sum_j E[Z_j], B) = (1 - \frac{1}{e})R(\bar{q})$ which proves our claim. Next, we show that Mech. 5 satisfies the requirements of DSIC-expansion: observe that all $R_j(\cdot)$ are concave so $R(\bar{q})$ is also concave. Furthermore, $R(\bar{q})$ is submodular according to the Def. 12 because the function $\mathbb{S}_R(S, \bar{q}) = \min(\sum_{j \in S} R_j(\bar{q}_j), B)$ is submodular in S . Therefore, Mech. 5 meets the requirements of DSIC-expansion. \square

Lemma 2. *Let B be an arbitrary positive number and let Z_1, \dots, Z_m be independent random variables such that for each j : $Z_j \in [0, B]$. Then:*

$$E[\min(\sum_j Z_j, B)] \geq (1 - \frac{1}{e^{\frac{1}{\sum_j E[Z_j]}/B}})B \geq (1 - \frac{1}{e}) \min(\sum_j E[Z_j], B)$$

A.4 Unit Demand Independent Valuations

In this subsection, we consider m indivisible items and a single unit demand buyer (i.e., her valuation for a bundle of items is the maximum of her valuations for individual items in the bundle). We assume that for each item j , the buyer's valuation is drawn independently from a publicly known distribution with CDF $F_j(\cdot)$. To avoid complicating the proofs, we assume that each $F_j(\cdot)$ is continuous and strictly increasing in its domain. Furthermore, we require the distributions to be regular. Next, we present a primary mechanism that obtains at least $\frac{1}{2}$ -fraction of the revenue of the optimal primary IP mechanism in this setting. Note that this is the same setting considered in [CHMS10] except for the regularity assumption for distributions.

Our approach is similar to subsection A.2. For each item j , define $H_j(q) = q \cdot F_j^{-1}(1 - q)$. Because $F_j(\cdot)$ is the CDF of a regular distribution the function $H_j(\cdot)$ and its concave closure $\hat{H}_j(\cdot)$ are the same (i.e., $H_j(\cdot)$ is concave itself). This is shown by the following lemma.

Lemma 3. *If $F(\cdot)$ is the CDF of a regular distribution then the function $H(q) = q \cdot F^{-1}(1 - q)$ is concave.*

Proof. To show that $H(q)$ is concave it is enough to show that $\frac{\partial}{\partial q} H(q)$ is non-increasing in q . But $\frac{\partial}{\partial q} H(q) = F^{-1}(1 - q) - \frac{q}{f(F^{-1}(1 - q))}$ in which $f(\cdot)$ is the PDF corresponding to $F(\cdot)$. Now if we substitute $q = 1 - F(p)$ then it is enough to show that the resulting function is non-decreasing in p because q is itself non-increasing in p . However, by this substitution we get $\frac{\partial}{\partial q} H(q) = p - \frac{1 - F(p)}{f(p)}$ which is non-decreasing in p by definition of regularity. \square

Intuitively, $\widehat{H}_j(q_j)$ is the maximum revenue that can be obtained by a mechanism that allocates item j with probability q_j . Next, we show that the revenue of the optimal primary IP mechanism when restricted to allocate the items with probabilities at most $\bar{q}_1, \dots, \bar{q}_m$ is no more than the optimal objective value of the following convex program:

$$\begin{aligned}
& \text{maximize:} && \sum_j \widehat{H}_j(q_j) && (CP_{rev-u}) \\
& \forall j : && q_j \leq \bar{q}_j \\
& && \sum_j q_j \leq 1 \\
& \forall j : && q_j \geq 0
\end{aligned}$$

Theorem 15. *The revenue of the optimal primary IP mechanism, when restricted to allocate the items with probabilities at most $\bar{q}_1, \dots, \bar{q}_m$, is upper bounded by the optimal objective value of the convex program (CP_{rev-u}) .*

Proof. Let q_1^*, \dots, q_m^* be the probabilities of allocating the items to the buyer by the optimal primary IP mechanism. For each item j , the expected revenue that can be obtained if item j is allocated with probability q_j^* is upper bounded by $\widehat{H}_j(q_j^*)$ (the proof of this claim is essentially the same as the proof of Theorem 11). Therefore, the expected revenue of the optimal primary IP mechanism cannot be more than $\sum_j \widehat{H}_j(q_j^*)$. Furthermore, observe that q^* is a feasible solution for the convex program (CP_{rev-u}) because the optimal primary mechanism never allocates more than one item so $\sum_j q_j^* \leq 1$ and also for each j it must be that $q_j^* \leq \bar{q}_j$. Therefore, the revenue of the optimal primary IP mechanism cannot be more than the optimal objective value of the convex program. \square

Next, we construct a $\frac{1}{2}$ -approximate primary IP mechanism as follows:

Mechanism 6.

- Define the benchmark $R(\bar{q})$ to be the optimal objective value of (CP_{rev-u}) as a function of \bar{q} .
- Given \bar{q} , solve the convex program of (CP_{rev-u}) and let q denote the optimal assignment of the convex program.
- For each item j : assign the price $p_j = F_j^{-1}(1 - q_j)$. WLOG, assume that items are labeled such that $p_1 \leq \dots \leq p_m$.
- For each item j , define $r_j = \max(q_j p_j + (1 - q_j)r_{j+1}, r_{j+1})$ and let $r_{m+1} = 0$. Let T be the subset of items defined as follows: $T = \{j | p_j \geq r_j\}$. Only offer the items in T at prices computed in the previous step (i.e., set the price of other items to infinity).

Theorem 16. *Mech. 6 obtains at least $\frac{1}{2}$ fraction of the revenue of the optimal primary IP mechanism. Furthermore, this mechanism satisfies the requirements of DSIC-expansion.*

Proof. First, we show that Mech. 5 obtains at least $\frac{1}{2}$ fraction of its benchmark $R(\bar{q})$ which by Theorem 15 is an upper bound on the revenue of the optimal primary IP mechanism. Observe that in Mech. 6, $R(\bar{q}) = \sum_j q_j p_j$ and notice that q_j is exactly the probability that the valuation of the buyer for item j is at least p_j . Now consider a replica of the buyer which has the exact same valuations as the original buyer but who always buys the item with the lowest price among

the items that are priced below her valuation. We call this imaginary replica the “malevolent replica”. Notice that for any assignment of prices the revenue obtained from the malevolent replica is a lower bound on the revenue obtained from the original buyer. So it is enough to show that the mechanism obtains a revenue of at least $\frac{1}{2} \sum_j q_j p_j$ from the malevolent replica. Observe that r_j as defined in Mech. 6 is exactly the expected revenue obtained from the malevolent replica by offering the items $T \cap \{j, \dots, m\}$. Basically, the item j is offered (i.e., $j \in T$) iff $p_j \geq r_j$ which also implies that $r_j = q_j p_j + (1 - q_j) r_{j+1}$. Otherwise, item j is not offered at all and $r_j = r_{j+1}$. Note that if we only offer the items in $T \cap \{j, \dots, m\}$ and if $j \in T$ then the malevolent replica buys item j with probability q_j and generates a revenue of p_j otherwise with probability $1 - q_j$ a revenue of r_{j+1} is obtained from items in $T \cap \{j + 1, \dots, m\}$. Now observe that the expected revenue obtained from the malevolent replica by Mech. 6 is exactly r_1 . By applying Lem. 4 we conclude that $r_1 \geq \frac{1}{2} \sum_j p_j q_j$ which completes the proof of the first claim. Next, we show that this mechanism satisfies the requirements of DSIC expansion. Observe that by Theorem 5, the optimal objective value of (CP_{rev-u}) is a concave function of $\bar{q}_1, \dots, \bar{q}_m$ so the benchmark function $R(\bar{q})$ is concave. Now it only remains to show that $R(\bar{q})$ is also submodular according to Def. 12. \square

Lemma 4. *Let p_1, \dots, p_m and q_1, \dots, q_m be a two sequences of non-negative real numbers and suppose $\sum_j q_j \leq 1$. For each $j = 1 \dots m$ define $r_j = \max(q_j p_j + (1 - q_j) r_{j+1}, r_{j+1})$ and let $r_{m+1} = 0$. Then $r_1 \geq \frac{1}{2} \sum_j q_j p_j$.*

B The Proof of Theorem 2

In this section, we present the proof of Theorem 2. We prove the theorem in two parts. First we show that, assuming the γ -conservative magician has infinitely many magic wands, the dynamic programming based strategy of Alg. 1 indeed guarantees that each box is opened with probability at least γ . The second part, which is the challenging part, is to show that for any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ the strategy table of the γ -conservative magician never instructs him to open a box if he has broken all of his k magic wands. In other words, we show that if $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ then $y_i^j = 0$ for every $j \geq k$ and every i . We should emphasize that it is not possible to get a bound of better than $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ by using ordinary concentration bounds or martingale inequalities. Furthermore, in Theorem 3 we show that it is not possible to guarantee that each box is opened with probability at least $1 - \frac{1}{\sqrt{2\pi k}} + \epsilon$, so the bound of $1 - \frac{1}{\sqrt{k+3}}$ is almost tight.

Part 1: We show that if the γ -conservative magician has infinitely many magic wands then the strategy table of Alg. 1 indeed ensures that each box is opened with probability at least γ . Below, we repeat the dynamic program for computing the strategy table:

$$y_i^j = \begin{cases} 1 & i \geq 1, \quad 0 \leq j < \theta_i \\ (\gamma - \phi_i^{\theta_i-1}) / (\phi_i^{\theta_i} - \phi_i^{\theta_i-1}) & i \geq 1, \quad j = \theta_i \\ 0 & \text{otherwise.} \end{cases} \quad (DP.y)$$

$$\theta_i = \min\{j \mid \phi_i^j \geq \gamma\} \quad (DP.\theta)$$

$$\phi_i^j = \begin{cases} 1 & i = 1, \quad j = 0 \\ y_{i-1}^j q_{i-1} \phi_{i-1}^{j-1} + (1 - y_{i-1}^j q_{i-1}) \phi_{i-1}^j & i \geq 2, \quad j \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (DP.\phi)$$

Part 1(a): First we prove that $Pr[S_i \leq j] \geq \phi_i^j$ by induction on i . The base case $i = 1$ is trivial. For $i > 1$:

$$\begin{aligned}
Pr[S_i \leq j] &\geq Pr[S_{i-1} \leq j-1] + Pr[S_{i-1} = j](1 - y_{i-1}^j q_{i-1}) \\
&= Pr[S_{i-1} \leq j-1] y_{i-1}^j q_{i-1} + Pr[S_{i-1} \leq j](1 - y_{i-1}^j q_{i-1}) \\
&\geq \phi_{i-1}^{j-1} y_{i-1}^j q_{i-1} + \phi_{i-1}^j (1 - y_{i-1}^j q_{i-1}) && \text{by induction hypothesis} \\
&= \phi_i^j && \text{by (DP.\phi)} \\
& && (1.a)
\end{aligned}$$

Observe that all of the above inequalities are met with equality if each q_i is the exact probability of breaking a magic wand for box i and not just an upper bound.

Part 1(b): Next, we show that $Pr[Y_i = 1] \geq \gamma$:

$$\begin{aligned}
Pr[Y_i = 1] &= \sum_{j=0}^{\theta_i} Pr[Y_i = 1 | S_i = j] Pr[S_i = j] \\
&= \sum_{j=0}^{\theta_i} y_i^j Pr[S_i = j] \\
&= Pr[S_i \leq \theta_i - 1] + y_i^{\theta_i} Pr[S_i = \theta_i] && \text{because } y_i^j = 1 \text{ for } j < \theta_i \\
&= (1 - y_i^{\theta_i}) Pr[S_i \leq \theta_i - 1] + y_i^{\theta_i} Pr[S_i \leq \theta_i] \\
&\geq (1 - y_i^{\theta_i}) \phi_i^{\theta_i - 1} + y_i^{\theta_i} \phi_i^{\theta_i} && \text{by (1.a)} \\
&= \phi_i^{\theta_i - 1} + y_i^{\theta_i} (\phi_i^{\theta_i} - \phi_i^{\theta_i - 1}) \\
&= \gamma && \text{by substituting } y_i^{\theta_i} \text{ from (DP.y)}
\end{aligned}$$

Observe that all of the above inequalities are met with equality if each q_i is the exact probability of breaking a magic wand for box i and not just an upper bound.

Part 2: Next, we show that when $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$, the γ -conservative magician never tries to open a box after breaking k magic wands. In other words, we show that $y_i^j = 0$ for every $j \geq k$ and every i , in the strategy table that he computes using the dynamic program (DP.y).

Instead of proving the claim directly, we present a stochastic process on an infinite tape with one unit of infinitely divisible sand and a barrier. Then we establish the connection between this process and the strategy table of the magician. Furthermore, we prove a theorem about this stochastic process and use it to prove second part of the theorem. The process is as follows.

Definition 13 (The Sand/Barrier Process).

- A parameter $\gamma \in (0, 1)$ and a sequence of probabilities q_1, \dots, q_n are given as the input, such that $\sum_i q_i \leq k$.
- There is a tape of infinite length, one unit of infinitely divisible sand, and a barrier. Initially, the barrier is at position 1 on the tape and all the sand is at position 0 (on the left of the barrier). During the process, we gradually move the sand and the barrier to right, distributing the sand over the tape, but never crossing the barrier. The barrier is moved one position to the right whenever the amount of sand on the barrier increases to more than $1 - \gamma$. The

process runs in n rounds. For each round i let λ_i denote the position of the barrier at the beginning of round i , and for each j let s_i^j denote the amount of sand at position j at the beginning of round i . Initially, $\lambda_1 = 1$ and $s_1^0 = 1$ and $s_1^j = 0$ for all $j > 0$. During each round i we do the following.

- We select a fraction of the sand from each position to the left of the barrier such that the total amount of selected sand is exactly γ . Note that we can always do this because the amount of sand that is on the barrier is no more than $1 - \gamma$, so the rest of the sand must be to the left of the barrier. Let $y_i^j \in [0, 1]$ denote the fraction of the sand on position j that gets selected. We start selecting all the sand greedily from left to right until the total amount selected is γ . In other words, for some index θ_i we have $y_i^j = 1$ for all $j < \theta_i$ and $y_i^j = 0$ for all $j > \theta_i$, such that $\sum_j y_i^j s_i^j = \gamma$.
- We then move q_i fraction of the selected sand as follows: for each j we move q_i fraction of the sand that was selected from position j to $j + 1$, we do this simultaneously for all positions (i.e. the amount of sand that is moved to $j + 1$ is exactly $y_i^j s_i^j q_i$).
- If the total amount of sand on the barrier is more than $1 - \gamma$ then we move the barrier one position to the right (i.e., to position $\lambda_i + 1$).

Let $\phi_i^j = \sum_{r=0}^j s_i^r$ be the total amount of sand in positions $0, \dots, j$ at the beginning of round i . Observe that by the above process we have:

$$\begin{aligned} \phi_i^j &= \phi_{i-1}^{j-1} + (1 - y_{i-1}^j q_{i-1})(\phi_{i-1}^j - \phi_{i-1}^{j-1}) \\ &= y_{i-1}^j q_{i-1} \phi_{i-1}^{j-1} + (1 - y_{i-1}^j q_{i-1}) \phi_{i-1}^j \end{aligned}$$

Notice that the state of the process defined above can be computed using exactly the same dynamic program that was used to compute the strategy table of the γ -conservative magician. Furthermore, in order to show that $y_i^j = 0$ for every $j \geq k$ and every i , it is enough to show that the barrier is never moved past position k on the tape. Next, we present a theorem which is the main step of the proof and is also interesting on its own.

Theorem 17 (Sand/Barrier). *Throughout the process defined in Def. 13, the average distance of the sand from the barrier is strictly less than $\frac{1}{1-\gamma}$. In particular, at the beginning of round i this distance is strictly less than $\frac{1-\gamma^{\lambda_i}}{1-\gamma}$, and this is true for any sequence of probabilities q_1, \dots, q_n , regardless of how big $\sum_i q_i$ is.*

We use the above theorem to derive a sufficient condition on γ that ensures the barrier is never moved past position k . At the beginning of round i , let d_i denote the average distance of the sand from the origin and d'_i denote the average distance of the sand from the barrier. Observe that $\lambda_i = d_i + d'_i$. Furthermore, notice that $d_i = \gamma q_{i-1} + d_{i-1}$, in other words, the average distance of the sand from the origin is increased exactly by γq_{i-1} during round $i - 1$ (because the amount of selected sand is exactly γ and q_{i-1} fraction of the selected sand is moved one position to the right). So we get the following inequality.

$$\lambda_i = d_i + d'_i < \sum_{r=1}^{i-1} q_r \gamma + \frac{1 - \gamma^{\lambda_i}}{1 - \gamma} \leq k\gamma + \frac{1 - \gamma^{\lambda_i}}{1 - \gamma}$$

Next, we present another inequality that contradicts the above inequality for $\lambda_i = k + 1$, therefore if we can show that the following inequality always holds then the barrier can never move past position k because if it moves to position $k + 1$ then the above inequality and the following inequality contradict each other.

$$k + 1 \geq k\gamma + \frac{1 - \gamma^{k+1}}{1 - \gamma} \quad (\Lambda)$$

Instead of the above inequality, we can consider the stronger inequality $k + 1 \geq k\gamma + \frac{1}{1-\gamma}$ which is quadratic in γ and can be solved to get a bound of $\gamma \leq 1 - \frac{1}{1/2 + \sqrt{k+1/4}}$. This bound is in fact a weaker condition than $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ when $k \geq 7$. It can also be verified that for $k < 7$ and $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ the inequality (Λ) holds. That completes the proof of Theorem 2.

Next, we present the proof of the sand theorem:

Proof of Theorem 17. First, we show that throughout the process of Def. 13, the following invariant holds:

$$\forall i, \forall j \in [1, \lambda_i - 1] : \quad \phi_i^{j-1} < \gamma \phi_i^j \quad (\phi_{ineq})$$

We prove this by induction on i . The base case of the induction is trivial. For $i \geq 2$:

$$\begin{aligned} \phi_i^j &= \phi_{i-1}^{j-1} + (1 - y_{i-1}^j q_{i-1})(\phi_{i-1}^j - \phi_{i-1}^{j-1}) \\ &\geq \phi_{i-1}^{j-1} + (1 - y_{i-1}^{j-1} q_{i-1})(\phi_{i-1}^j - \phi_{i-1}^{j-1}) && \text{because } y_i^j \text{ is non-increasing in } j \\ &= y_{i-1}^{j-1} q_{i-1} \phi_{i-1}^{j-1} + (1 - y_{i-1}^{j-1} q_{i-1}) \phi_{i-1}^j \\ &> y_{i-1}^{j-1} q_{i-1} \frac{1}{\gamma} \phi_{i-1}^{j-2} + (1 - y_{i-1}^{j-1} q_{i-1}) \frac{1}{\gamma} \phi_{i-1}^{j-1} && \text{by induction hypothesis} \\ &= \frac{1}{\gamma} \phi_i^{j-1} \end{aligned}$$

We also need to consider the case in which the barrier is moved. Suppose the barrier is moved at the end of round i so $\lambda_{i+1} = \lambda_i + 1$, so we must show that the invariant also holds for position $j = \lambda_i$. Notice that the barrier is only moved when there is more than $1 - \gamma$ sand on it. So the total sand to the left of the barrier, just before it was moved, must have been strictly less than γ which means $\phi_{i+1}^{\lambda_i-1} < \gamma$, furthermore $\phi_{i+1}^{\lambda_i} = 1$ because the barrier just moved to position $\lambda_i + 1$. So $\phi_{i+1}^{\lambda_i-1} < \gamma \phi_{i+1}^{\lambda_i}$.

Now, we prove the main claim of the theorem. Let d_i' denote the average distance of the sand

from the barrier at the beginning of round i , then:

$$\begin{aligned}
d'_i &= \sum_{j=0}^{\lambda_i-1} \phi_i^j && \text{sand at position } j \text{ is counted exactly } \lambda_i - j \text{ times} \\
&< \sum_{j=0}^{\lambda_i-1} \gamma^{\lambda_i-1-j} \phi_i^{\lambda_i-1} && \text{by } (\phi_{ineq}) \\
&\leq \sum_{j=0}^{\lambda_i-1} \gamma^{\lambda_i-1-j} && \text{because } \phi_i^{\lambda_i-1} \leq 1 \\
&= \frac{1 - \gamma^{\lambda_i}}{1 - \gamma}
\end{aligned}$$

□

C Online Stochastic Optimization

In this section, we briefly discuss applications of Alg. 1 and Theorem 2 to online stochastic optimization problems in general. We present an abstract rounding problem that arises as a part of many online stochastic optimization problems in which we optimize for the expected instance.

Definition 14. *Let $X_1 \dots, X_n$ be independent non-identical Bernoulli random variables (with known distributions) that are realized in some unknown order during an online process. Suppose we have optimized the process such that $\sum_i E[X_i] \leq k$, however we need to ensure that under any realization of these random variables the constraint is not violated. In order to achieve this we may have to round some of the X_i that are realized as 1 to 0 and we have to do this online. In other words, if X'_i is the result of rounding X_i then we should set X'_i at the time X_i is realized. The objective is to maximize the $Pr[X'_i = 1 | X_i = 1]$ for every i .*

As an application of the above rounding problem, consider the following problem: Suppose we have k units of a resource and each X_i is a possible request for one unit of resource but the sequence of X_i may arrive in an unknown order. Furthermore, suppose we have optimized the system so that the expected number of requests is no more than k , i.e. $\sum_i E[X_i] \leq k$. We wish to maximize the probability of satisfying each request.

Observe that for any $\gamma \leq \gamma_k$ the following algorithm guarantees that $Pr[X'_i = 1 | X_i = 1] \geq \gamma_k$ for any i which means every request regardless of its arrival time can be satisfied with probability γ .

Algorithm 2 (γ -Conservative Rounding).

- Create a γ -conservative magician with k magic wands.
- Upon the realization of X_i , write $E[X_i]$ on a box and present it to the magician. Let Y_i be the indicator random variable which is 1 iff the magician chooses to open the box. Set $X'_i \leftarrow Y_i X_i$. Furthermore, if $X'_i = 1$, break the wand of the magician.

D When Buyers Need More Than 1 Copy of Each Item

In this section, we show that the more general model in which each buyer may demand more than one copy of each item but no more than $\frac{1}{k}$ of all copies of an item, can be reduced to the simpler

model of Def. 1 in which we have at least k copies of every item and no one demands more than 1 copy of each item.

Definition 15 (Multi-Demand Market Transformation). *Let k_j denote the number of copies of item j . Define $c_j = \lfloor \frac{k_j}{k} \rfloor$ and divide the copies of item j almost equally into c_j bins (i.e., each bin will contain either c_j or $c_j + 1$ copies). Create a new item type for each bin and treat the copies of the same item from different bins as different items.*

Theorem 18. *Any mechanism that does not allocate more than $\frac{1}{k}$ of all copies of an item in the general (multi-demand) model can be converted to a mechanism that does not allocate more than one copy of each item in the transformed market, such that from the view point of buyers the actual allocations and payments are the same as the original mechanism.*

Proof. To prove the theorem, it is enough to show that we can convert any allocation of the original mechanism in the general market to an allocation in the transformed market such that no buyer receives more than one copy of the same item from the same bin. We proceed as follows: For every j , we create a list L_j of bins of item j . L_j is initially sorted in decreasing order of the size of the bins. Let x_{ij} be the number of copies of item j allocated to buyer i by the original mechanism. We define the allocation of the transformed mechanism as follows: for each buyer i we do the following x_{ij} times: Allocate one copy from the bin that is in front of the list L_j and then move the bin to the end of the list. It is easy to see that no two items from the same bin are allocated to the same buyer. \square

Note that by Theorem 18, any mechanism in the original market is equivalent to a mechanism in the transformed market with the exact same allocations/payments from the perspective of buyers. So WLOG, we can work with the transformed market and only consider the mechanisms in this market. However, note that to use our generic expansions in the transformed market, the underlying primary mechanisms should be capable of handling correlated valuations because different copies of the same item which are labeled as different items are perfect substitutes from the view point of a buyer. Among the primary mechanisms presented in Appendix A, only Mech. 3 (subsection A.1) can handle correlated valuations.

E Missing Proofs

Proof of Theorem 3. Suppose we create n boxes and in each box, independently, we put \$1 with probability $\frac{k}{n}$. If the magician opens a box containing a \$1 then he gets the \$1 but we break his wand (i.e., $q_i = \frac{k}{n}$). Observe that the expected total prize is k dollars but because we put a dollar in each box independently, there are some instances in which there are more than k non-empty boxes but the magician cannot win more than k dollars at any instance. Let X_i be the indicator random variable which is 1 iff there is a dollar in box i . The expected total prize is $E[\sum_i X_i] = k$ but the expected prize that the magician can win is at most $E[\min(\sum_i X_i, k)]$. It can be shown that $E[\min(\sum_i X_i, k)] = (1 - \frac{k^k}{e^k k!})k$ asymptotically as n goes to infinity. In fact, for any positive ϵ , there is a large enough n such that $E[\min(\sum_i X_i, k)] < (1 - \frac{k^k}{e^k k!} + \epsilon)k$. On the other hand, observe that if a magician can guarantee that every box is opened with probability at least $\gamma = 1 - \frac{k^k}{e^k k!} + \epsilon$ then in expectation he will win at least $\sum_i \gamma E[X_i] = (1 - \frac{k^k}{e^k k!} + \epsilon)k$ therefore no magician can make such a guarantee. \square

Proof of Theorem 4. First, we compute an upper bound on the expected payoff of the prophet. Let q_i be the probability that the prophet chooses V_i (i.e. the probability that V_i is among the k

highest draws). Now let $u_i(q_i)$ denote the maximum possible contribution of random variable V_i to the expected payoff of the prophet if V_i is selected with probability q_i . Let $F_i(\cdot)$ and $f_i(\cdot)$ be the CDF and PDF of V_i and let v_i^* be such that $\Pr[V_i \geq v_i^*] = q_i$. Then $u_i(q_i) = \int_{v_i^*}^{\infty} v f_i(v) dv$. By substituting $v_i^* = F_i^{-1}(1 - q_i)$ and doing a change of variables in the integral and applying the chain rule we get $u_i(q_i) = \int_0^{q_i} F_i^{-1}(1 - q) dq$. Observe that $\frac{d}{dq_i} u_i(q_i) = F_i^{-1}(1 - q_i)$ is a non-increasing function so $u_i(q_i)$ is a concave function. Furthermore, $\sum_i q_i \leq k$ because the prophet cannot choose more than k random draws. So the optimal objective value of the following convex program is an upper bound on the payoff of the prophet:

$$\text{maximize:} \quad \sum_i u_i(q_i) \quad (U)$$

$$\sum_i q_i \leq k \quad (\tau)$$

$$\forall i : \quad q_i \geq 0 \quad (\mu_i)$$

Now let $L(q, \tau, \mu) = -\sum_i u_i(q_i) + \tau(\sum_i q_i - k) - \sum_i \mu_i q_i$ be the Lagrangian. By KKT stationarity condition, at the optimal assignment $\frac{\partial}{\partial q_i} L(q, \tau, \mu) = 0$. On the other hand, $\frac{\partial}{\partial q_i} L(q, \tau, \mu) = -F_i^{-1}(1 - q_i) + \tau - \mu_i$. Assuming that $q_i > 0$ by complementary slackness $\mu_i = 0$ which then implies that $q_i = 1 - F_i(\tau)$. Furthermore, we can also show that $\sum_i \Pr[V_i > \tau] = k$ because the first constraint is tight. Observe that the contribution of each V_i to the objective value of the convex program is exactly $E[V_i | V_i > \tau] \Pr[V_i > \tau]$. Now, by using a γ_k -conservative magician we can ensure that each box is opened with probability at least γ_k which implies the contribution of each V_i to the expected payoff of the gambler is $E[V_i | V_i > \tau] \Pr[V_i > \tau] \gamma_k$ which proves that the expected payoff of the gambler is at least γ_k fraction of optimal objective value of the convex program which was itself and upper bound on the expected payoff of the prophet. \square

Proof of Theorem 5. The proof has two parts. In the first part, we show that $R(\bar{q})$ is concave. In the second part, we show how to construct a convex program of this form whose solution corresponds to the allocation function of the optimal primary mechanism.

$$\text{maximize:} \quad u(x) \quad (CP_R)$$

$$\forall j : \quad g_j(x) \leq \bar{q}_j \quad (\alpha_j)$$

$$x \in \mathbb{X}$$

Part 1: Define the Lagrange dual function $D(\alpha, \bar{q})$ as follows:

$$D(\alpha, \bar{q}) = \max_{x \in \mathbb{X}} \left(u(x) - \sum_j \alpha_j (g_j(x) - \bar{q}_j) \right)$$

Notice that we can decompose the dual function to two components and write it as $D(\alpha, \bar{q}) = D'(\alpha) + \sum_j \alpha_j \bar{q}_j$ in which $D'(\alpha) = \max_{x \in \mathbb{X}} (u(x) - \sum_j \alpha_j g_j(x))$ is independent of \bar{q} . Now, by strong duality we can write $R(\bar{q})$ as follows:

$$\begin{aligned} R(\bar{q}) &= \min_{\alpha \geq 0} D(\alpha, \bar{q}) \\ &= \min_{\alpha \geq 0} (D'(\alpha) + \sum_j \alpha_j \bar{q}_j) \end{aligned}$$

Notice that, in the last line, the expression inside the parenthesis is a linear function of \bar{q} and we are taking the minimum of an infinite number of linear functions of \bar{q} so the resulting function is a concave function of \bar{q} .

Part 2: We show how to construct such convex program under any feasibility constraints, both for revenue and for welfare. We show that $u(x)$ and all $g_j(x)$ are linear functions in x and the set \mathbb{X} is a convex set.

- \mathbb{X} is the set of all truthful and feasible allocation functions. Note that any truthful primary mechanism can be fully specified by its allocation function $x : \mathbb{R}^{2^m-1} \rightarrow [0, 1]^{2^m}$ that maps every possible type of the buyer to a probability distribution over bundles of items. Also, the expected payment of any truthful mechanism can be specified uniquely (up to an additive constant) by just specifying the allocation function. The buyer's type is a (2^m-1) -dimensional vector that specifies her valuation for every possible bundle of items. Furthermore, observe that we can encode all the feasibility and IC constraints into \mathbb{X} . In other words, \mathbb{X} is the set of all allocation functions for which the corresponding mechanisms satisfy the IC constraints and feasibility constraints (i.e., matroid constraints, budget constraints, etc). Observe that \mathbb{X} is a convex set because any convex combination of two truthful feasible allocation functions is itself a truthful feasible allocation function (remember that an allocation function maps the buyer's types to probability distributions over bundles of items, so any infeasible bundle will still have 0 probability in any convex combination of two feasible allocations).
- $g_j(x)$ is the marginal probability of allocating item j by the mechanism that uses allocation function x . In a Bayesian setting we assume that the seller knows the probability distribution of buyer's types. Let $f : \mathbb{R}_+^{2^m-1} \rightarrow \mathbb{R}_+$ be the PDF of this distribution. Also let $I_j \in \{0, 1\}^{2^m}$ be a vector that has a 1 at every position that corresponds to a bundle containing item j . Then $g_j(x)$ can be specified as follows:

$$g_j(x) = \int_{\mathbb{R}_+^{2^m-1}} I_j \cdot x(t) f(t) dt$$

Observe that $g_j(x)$ is linear in x .

- $u(x)$ is the expected objective value of the mechanism that uses the allocation function x . Both for welfare and for revenue, $u(x)$ can be written in the following form:

$$u(x) = \int_{\mathbb{R}_+^{2^m-1}} \phi(t) \cdot x(t) f(t) dt$$

For welfare maximization we define $\phi(t) = t$. For revenue maximization we define $\phi(t)$ to be the virtual valuation as defined below (see [Arm96] for details):

$$\phi(t) = t - \frac{t}{f(t)} \int_1^\infty r^{2^m-2} f(rt) dr$$

Observe that, both for welfare and for revenue, $u(x)$ is linear in x .

□

Proof of Theorem 6. Let \mathcal{M}_N^* be the optimal n -buyer mechanism and for each i and j let q_{ij}^* be the probability that \mathcal{M}_N^* allocates a copy of item j to buyer i . We show that there exist primary mechanisms $\mathcal{M}_1^*, \dots, \mathcal{M}_n^*$ such that for each i the expected objective value that \mathcal{M}_i^* obtains from buyer i is exactly the same as the expected objective value that \mathcal{M}_N^* obtains from buyer i and \mathcal{M}_i^* allocates each item j to buyer i with probability exactly q_{ij}^* . Then we can argue that the expected objective value obtained by each \mathcal{M}_i^* is upper bounded by the benchmark $R_i(q_i^*)$ and therefore the expected total objective value obtained by \mathcal{M}_N^* is no more than $\sum_i R_i(q_i^*)$. On the other hand, q^* is a feasible solution for (CP_{R_N}) therefore $\sum_i R_i(q_i^*)$ is no more than the optimal objective value of the convex program (CP_{R_N}) . Next, we show how to construct each \mathcal{M}_i^* . For each i , we construct \mathcal{M}_i^* by using \mathcal{M}_N^* on buyer i and $n - 1$ dummy buyers where for each $i' \neq i$, the valuations of dummy buyer i' are drawn independent of other (dummy) buyers and are distributed according to the prior distribution of buyer i' . Observe that \mathcal{M}_i^* as we just described is a primary mechanism for buyer i and allocates each item j to buyer i with the exact same probabilities as \mathcal{M}_N^* . Notice that it is crucial that the valuations of different buyers be independent, otherwise we wouldn't be able to draw the valuations of the dummy buyers from the correct distribution without knowing the valuations of buyer i first. □

Proof of Theorem 7. First we show that the γ -BIC expansion obtains in expectation at least $\gamma \cdot \alpha$ fraction of the expected objective value of the optimal n -buyer mechanism. Notice that the expected objective value of the mechanism is $E[\sum_i P_i] = \gamma E[\sum_i P_i] \geq \gamma \cdot \alpha \sum_i R_i(\bar{q}_i)$ and by Theorem 6 $\sum_i R_i(\bar{q}_i)$ is an upper bound on the objective value of the optimal n -buyer mechanism which proves our claim.

Next, we prove that the γ -BIC expansion is indeed BIC. Consider any arbitrary buyer i . Let $T_i = \{j | X_{ij} = 1\}$ be the set of items initially allocated to buyer i by $\mathcal{M}_i(\bar{q}_i)$ and let $T'_i = \{j | X'_{ij} = 1\}$ be the final set of items allocated to buyer i by the γ -BIC expansion Mech. 1. By the assumption of the theorem, valuations of buyer i can be interpreted as a weighted rank function of a matroid so WLOG we assume that T_i is an independent set of this matroid (otherwise we can instead allocate an independent subset of T_i with the highest valuation). Observe that the valuation of buyer i for any subset of T_i is additive and T'_i is a random subset of T_i . Therefore, if we show that each item $j \in T_i$ is also in T'_i with probability γ then the expected valuation of buyer i for T'_i is exactly γ times her valuation for T_i . We prove this by applying Theorem 2 to each instance of the γ -conservative magician. Observe that for each j , the mechanism presents n boxes to the j^{th} instance of γ -conservative magician with magic wand breakage probabilities $\hat{q}_{1j}, \dots, \hat{q}_{mj}$. Furthermore, for each j : $\sum_i \hat{q}_{ij} \leq \sum_i \bar{q}_{ij} \leq k_j$ so by applying Theorem 2 to the j^{th} magician we conclude that each box is opened with probability exactly γ . Therefore $Pr[X'_{ij} = 1 | X_{ij} = 1] = Pr[Y_{ij} = 1 | X_{ij} = 1] = Pr[Y_{ij} = 1] = \gamma$ because Y_{ij} and X_{ij} are independent (because buyers are independent). So T'_i is a random subset of T_i that includes each item from T_i with probability exactly γ . Notice that since each \mathcal{M}_i is BIC and all allocations and payments of \mathcal{M}_i are always scaled by the same constant γ , the resulting mechanism is also BIC. Notice that each time we allocate a copy of item j , we break the magic wand of the j^{th} γ -conservative magician therefore we never allocate more items than we actually have. That completes the proof. □

Proof of Theorem 8. First we show that the γ -DSIC expansion is indeed DSIC. Notice that for each buyer i , the γ -DSIC expansion selects the items that can be offered to buyer i prior to running \mathcal{M}_i and once it runs \mathcal{M}_i it uses the allocation/payment of \mathcal{M}_i as the final allocation/payment for buyer i . Therefore the γ -DSIC expansion is DSIC for each buyer i because $\mathcal{M}_i(\bar{q}'_i)$ is truthful for every possible \bar{q}'_i . Furthermore, this mechanism also preserves all ex-post properties of \mathcal{M}_i . For example

if \mathcal{M}_i is ex-post individually rational then the γ -DSIC expansion will also be ex-post individually rational for buyer i .

Next, we show that the expected objective value of the γ -DSIC expansion is at least $\gamma \cdot \alpha$ fraction of the expected objective value of the optimal n -buyer BIC mechanism. We start by showing that each item is available to each \mathcal{M}_i with probability at least γ , i.e., we show that $Pr[Y_{ij} = 1] \geq \gamma$. Observe that for each j , a sequence of n boxes with magic wand breakage probability upper bounds $\bar{q}_{1j}, \dots, \bar{q}_{nj}$ are presented to the j^{th} magician. By applying Theorem 2, when $\gamma \leq \gamma_k$, this magician guarantees that each box is opened with probability at least γ which means $Pr[Y_{ij} = 1] \geq \gamma$. Next, we show that if each item is available to each \mathcal{M}_i with probability at least γ then the expected objective value obtained from each \mathcal{M}_i is at least $\gamma \cdot \alpha R_i(\bar{q}_i)$ which implies that the total expected objective value of the γ -DSIC expansion is at least $\gamma \cdot \alpha \sum_i R_i(\bar{q}_i)$ which proves our claim because, by Theorem 6, $\sum_i R_i(\bar{q}_i)$ is an upper bound on the expected objective value of the optimal BIC mechanism. Let $T_i = \{j | Y_{ij} = 1\}$ denote the random subset of items available to \mathcal{M}_i and let $M_r = \{1, \dots, r\}$ denote the first r item types (remember that M was the set of all item types). Also let $G_i(T) = R_i(q_i'')$ in which if $j \in T$ then $q_{ij}'' = \bar{q}_{ij}$ else $q_{ij}'' = 0$. Note that by the assumption of the theorem the benchmark functions are submodular which by Def. 12 implies that $G_i(\cdot)$ as defined above is submodular. Next, we show that the expected objective value of the γ -DSIC-expansion obtained from buyer i , which we denote by $E[\mathcal{M}_i(\bar{q}_i)]$, is at least $\gamma \cdot \alpha R_i(\bar{q}_i)$ and that completes the proof:

$$\begin{aligned}
E[\mathcal{M}_i(\bar{q}_i)] &\geq E[\alpha G_i(T_i)] \\
&= \alpha E\left[\sum_{r=1}^m G_i(T_i \cap M_r) - \sum_{r=0}^{m-1} G_i(T_i \cap M_r)\right] \\
&= \alpha E\left[\sum_{r=1}^m G_i(T_i \cap M_r) - G_i(T_i \cap M_{r-1})\right] \\
&= \alpha \sum_{r=1}^m Pr[r \in T_i][G_i(\{r\} \cup (T_i \cap M_{r-1})) - G_i(T_i \cap M_{r-1})] \\
&\geq \alpha \sum_{r=1}^m Pr[r \in T_i][G_i(\{r\} \cup M_{r-1}) - G_i(M_{r-1})] && \text{By submodularity of } G_i(\cdot) \\
&\geq \alpha \cdot \gamma \sum_{r=1}^m (G_i(M_r) - G_i(M_{r-1})) \\
&= \alpha \cdot \gamma G_i(M_m) \\
&= \alpha \cdot \gamma R_i(\bar{q}_i)
\end{aligned}$$

□

Proof of Theorem 9. The proof of this theorem is similar to the proof of Theorem 6. Let \mathcal{M}_N^* be the optimal n -buyer AIP mechanism and for each i and j let q_{ij}^* be the probability that \mathcal{M}_N^* allocates a copy of item j to buyer i . We show that there exist primary IP mechanisms $\mathcal{M}_1^*, \dots, \mathcal{M}_n^*$ such that for each i the expected objective value that \mathcal{M}_i^* obtains from buyer i is exactly the same as the expected objective value that \mathcal{M}_N^* obtains from buyer i , such that \mathcal{M}_i^* allocates each item j to buyer i with probability exactly q_{ij}^* . Then we can argue that the expected objective value obtained by each \mathcal{M}_i^* is upper bounded by the benchmark $R_i(q_i^*)$ and therefore the expected total objective value obtained by \mathcal{M}_N^* is no more than $\sum_i R_i(q_i^*)$. On the other hand, q^* is a feasible

solution for (CP_{R_N}) therefore $\sum_i R_i(q_i^*)$ is no more than the optimal objective value of the convex program (CP_{R_N}) . Next, we show how to construct each \mathcal{M}_i^* . For each i , we construct the primary IP mechanism \mathcal{M}_i^* by using \mathcal{M}_N^* on buyer i and $n-1$ dummy buyers where for each $i' \neq i$, the valuations of dummy buyer i' are drawn independent of other (dummy) buyers and are distributed according to the prior distribution of buyer i' . Observe that \mathcal{M}_i^* is a primary AIP mechanism for buyer i and allocates each item j to buyer i with the exact same probabilities as \mathcal{M}_N^* . Notice that it is crucial that the valuations of different buyers be independent, otherwise we wouldn't be able to draw the valuations of the dummy buyers from the correct distribution without knowing the valuations of buyer i first. Now observe that for a single buyer the space of AIP mechanisms collapse to the space of IP mechanisms so each \mathcal{M}_i is in fact a primary IP mechanism because the dummy buyers are randomly created by the mechanism and do not really exist. \square

Proof of Lem. 2. Let $\mu = \sum_j E[Z_j]$. Define the random variables $W_j = \max(W_{j-1} - Z_j, 0)$ and $W_0 = B$. Observe that for each j , $W_j = \max(B - \sum_{r=1}^j Z_r, 0)$ so $\min(\sum_{r=1}^j Z_r, B) + W_j = B$. Therefore $E[\min(\sum_{r=1}^j Z_r, B)] + E[W_j] = B$ and to prove the theorem it is enough to show that $E[W_m] \leq \frac{1}{e^{\mu/B}} \cdot B$. To show this we will prove the following inequality:

$$W_j \leq \left(1 - \frac{E[Z_j]}{B}\right) W_{j-1} \quad (W_j)$$

Assuming that (W_j) is true, we can conclude the following which proves the claim.

$$\begin{aligned} W_m &\leq B \cdot \prod_{j=1}^m \left(1 - \frac{E[Z_j]}{B}\right) \\ &\leq B \cdot \frac{1}{e^{\mu/B}} \end{aligned}$$

The last inequality follows from the fact that $\sum_j \frac{E[Z_j]}{B} = \frac{\mu}{B}$, therefore the right hand side takes its maximum when for all j : $\frac{E[Z_j]}{B} = \frac{\mu}{mB}$ and $m \rightarrow \infty$. Furthermore, to prove the second inequality, we can use the fact that $(1 - x^a) \geq (1 - x)a$ for any $a \leq 1$ and conclude that $(1 - \frac{1}{e^{\mu/B}})B \geq (1 - \frac{1}{e^{\frac{\min(\mu, B)}{B}}})B \geq (1 - \frac{1}{e}) \frac{\min(\mu, B)}{B} B = (1 - \frac{1}{e}) \min(\mu, B)$. Now it only remains to prove the inequality (W_j) :

$$\begin{aligned} E[W_j] &= E[\max(W_{j-1} - Z_j, 0)] \\ &\leq E[\max(W_{j-1} - Z_j \frac{W_{j-1}}{B}, 0)] && \text{because } \frac{W_{j-1}}{B} \leq 1 \\ &= E[W_{j-1} - Z_j \frac{W_{j-1}}{B}] && \text{because } \frac{Z_j}{B} \leq 1 \\ &= E[W_{j-1}] - \frac{1}{B} E[Z_j W_{j-1}] \\ &\leq E[W_{j-1}] - \frac{1}{B} E[Z_j] E[W_{j-1}] && \text{because } Z_j \text{ and } W_{j-1} \text{ are independent.} \\ &= \left(1 - \frac{E[Z_j]}{B}\right) E[W_{j-1}] \end{aligned}$$

That completes the proof. \square

Proof of Lem. 4. To prove the claim, it is enough to show that $\frac{r_1}{\sum_j q_j p_j} \geq \frac{1}{2}$. WLOG, we may assume that $\sum_j p_j q_j = 1$ since we can scale p_1, \dots, p_m by a constant $c = \frac{1}{\sum_j q_j p_j}$ and this will also scale r_1, \dots, r_m by the same constant c so the ratio $\frac{r_1}{\sum_j q_j p_j}$ will not be affected. Now consider the following LP and observe that q_j, p_j , and r_j , as defined in the statement of the lemma, form a feasible assignment for this LP. If we show that the optimal objective value of the LP is bounded below by $\frac{1}{2}$, then that means for any feasible assignment the objective value will be at least $\frac{1}{2}$, and therefore $\frac{r_1}{\sum_j q_j p_j} \geq \frac{1}{2}$. In the following LP, p_j , and r_j are variables and everything else is constant.

$$\begin{aligned}
& \text{minimize:} && r_1 \\
& \forall j \in \{1 \dots m\} : && r_j \geq q_j p_j + (1 - q_j) r_{j+1} && (\alpha_j) \\
& \forall j \in \{1 \dots m\} : && r_j \geq r_{j+1} && (\beta_j) \\
& && \sum_j q_j p_j \geq 1 && (\gamma) \\
& && p_j \geq 0, \quad r_j \geq 0
\end{aligned}$$

To prove that the optimal objective value of the above LP is bounded below by $\frac{1}{2}$, we construct a feasible solution for its dual that obtains an objective value of $\frac{1}{2}$. The following is the dual LP:

$$\begin{aligned}
& \text{maximize:} && \gamma \\
& \forall j \in \{1 \dots m\} : && \gamma \leq \alpha_j && (p_j) \\
& && \alpha_1 + \beta_1 \leq 1 && (r_1) \\
& \forall j \in \{2 \dots m\} : && \alpha_j + \beta_j \leq (1 - q_{j-1}) \alpha_{j-1} + \beta_{j-1} && (r_j) \\
& && 0 \leq (1 - q_m) \alpha_m + \beta_m && (r_{m+1}) \\
& && \alpha_j \geq 0, \quad \beta_j \geq 0, \quad \gamma \geq 0
\end{aligned}$$

Now, suppose that for all j we set $\alpha_j = \gamma$ and $\beta_j = \beta_{j-1} - q_{j-1} \gamma$ except that for $j = 1$ we set $\beta_1 = 1 - \gamma$. From this assignment, we get $\beta_j = 1 - \gamma - \gamma \sum_{k=1}^{j-1} q_k$. Observe that we get a feasible solution as long as all β_j resulting from this assignment are non-negative. Furthermore, it is easy to see that $\beta_j \geq 1 - \gamma - \gamma \sum_{k=1}^m q_k \geq 1 - 2\gamma$ because $\sum_j q_j \leq 1$. Therefore, by setting $\gamma = \frac{1}{2}$, all β_j are non-negative and we always get a feasible solution for the dual LP with an objective value of $\frac{1}{2}$ which completes the proof. \square