A Constant-Factor Approximation to Multidimensional Consumer Utility*

Kira Goldner† Taylor Lundy‡ Thodoris Tsilivis§

November 3, 2026

Abstract

Motivated by social services where consumers pay with non-transferable ordeals, we study mechanisms that maximize consumer utility for multiple unit-demand buyers and heterogeneous items whose values are drawn independently from known prior distributions. Prior work in utility maximization approximates social welfare and shows that the gap between optimal utility and social welfare is logarithmic. We resolve the question of whether simple mechanisms can guarantee a constant-factor approximation to optimal utility itself. Our mechanisms achieve a $(5.67 + \varepsilon)$-approximation for general independent values, improving to $2e/(e-1) < 3.164$ when values are i.i.d. Each buyer chooses their favorite option from posted item prices or free item lotteries, and contention resolution determines which buyers' requests are served; these unit-demand mechanisms are BIC, ex-post individually rational, and computable in polynomial time. Our main technical contribution is a general upper bound on the optimal ex-ante-constrained utility that separates the contributions captured by posted prices and free lotteries. The same tools also give a $(5.67 + \varepsilon)$-approximation for matroid-constrained additive buyers, including unconstrained additive buyers. These results establish a utility counterpart to the theory of simple, approximately revenue-optimal mechanisms.

*The work of K. Goldner and T. Tsilivis was supported by NSF CAREER Award CCF-2441071.

†Boston University; goldner@bu.edu.

‡University of British Columbia; tlundy@cs.ubc.ca.

§Boston University; tsilivis@bu.edu.

1 Introduction

In this paper, we study the design of simple, approximately-optimal mechanisms for consumer utility, also called *consumer surplus* or *residual surplus*: the buyers' value for the goods they receive less the payments they make. This objective is well-motivated by social services, where recipients are often required to pay with their time and effort. For example, patients wait to receive treatment, and applicants for public benefits complete labor-intensive paperwork [8, 32]. For a service provider who cannot or does not wish to charge money, these *ordeals* offer a way to ration demand and elicit information about how much people value the service being provided. However, such ordeals are costly to consumers, and provide no tangible benefit to providers. We study the design of mechanisms for this setting by modeling nonnegative (non-monetary) payments that are burned—that is, not transferred to the seller; our objective is hence the buyers' total utility.

Utility maximization requires a trade-off between maximizing efficiency in the allocation and the payments that are required to truthfully implement such an allocation. For example, for two buyers with similar values, a second-price auction achieves almost no utility, as the payment consumes all of it. Here, a free lottery would be a much better choice. However, if the buyers have a very large gap between their values, the opposite is true; prices are necessary to distinguish the buyer with the high value. In a multi-item setting, the designer also needs to factor in which items each buyer prefers, even if no payment is charged.

The single-dimensional setting is completely resolved, with a Myersonian-like theory by Hartline and Roughgarden [31] giving a closed-form description of the optimal mechanism. Goldner and Lundy [26] initiate the study of multidimensional Bayesian utility maximization, constructing simple mechanisms for n unit-demand buyers and m items with i.i.d. values. Their mechanism, which asks buyers what their favorite item is, and then sells each item separately, gives a tight $\Theta(1 + \log(\max\{1, n/m\}))$-approximation to the benchmark of social welfare. Ezra et al. [21] extend this approximation to general distributions by parameterizing the amount of item-symmetry present. Goldner and Lundy [26] leave the open question: what is a better upper-bound on optimal utility than social welfare, and can simple mechanisms guarantee a constant-factor approximation to the optimal utility?

Revenue maximization has been the impetus for decades of foundational work. It is one of the largest and most influential literatures in algorithmic mechanism design. A landmark achievement of this area, reached through a long sequence of breakthroughs, was to show that simple mechanisms can obtain a constant-fraction of the optimal revenue in multidimensional multi-buyer settings—from unit-demand values [12] to additive values [40] to a unified framework [10] to XOS values [9]—despite the complexity of optimal mechanisms. Our result establishes a parallel milestone for consumer utility, providing a foundation for a corresponding theory of simple, approximately-utility-optimal mechanism design.

Our Results. We develop tools for approximating optimal utility and use them to prove our main result: a constant-factor approximation for unit-demand buyers with general independent values. The symmetric setting introduces these tools, and the matroid extension demonstrates their broader applicability. Our contributions are:

- *A general utility benchmark.* We upper-bound a single buyer's optimal ex-ante-constrained utility by a sum of optimal single-dimensional utilities and a weighted free-lottery term (Theorem 1). This gives an interpretable upper bound on multidimensional utility that accounts

for the allocation constraints, even when social welfare is a much larger benchmark. The bound applies broadly whenever utility is linear in independent item parameters, including constrained-additive valuations.

- *General unit-demand buyers.* For independent non-identical values, we give a $2e/(e-1) + (e/(e-1))^2 + \varepsilon < 5.667 + \varepsilon$ approximation for every $\varepsilon > 0$ (Theorem 3).
- *Symmetric items, unit-demand buyers.* For values drawn i.i.d. across items and buyers, we give a $2e/(e-1) < 3.164$ approximation (Theorem 5). The same guarantee, up to an additive ε in the approximation ratio, holds when values are i.i.d. across items for each buyer but buyers have different distributions (Theorem 6).
- *Matroid-constrained additive buyers.* The same $2e/(e-1) + (e/(e-1))^2 + \varepsilon$ approximation extends to buyers whose values are additive over sets satisfying a known matroid constraint (Theorem 4). This generalizes both unit-demand and additive buyers.

Our unit-demand mechanisms are simple, BIC, ex-post individually rational, and computable in polynomial time under standard assumptions (Section 2.4). The matroid extension retains the general unit-demand approximation and incentive guarantees; its polynomial-time computation requires additional oracle access. In the symmetric unit-demand setting, each buyer is offered the same (possibly randomized) price for every item and chooses their favorite item. In the general unit-demand setting, each buyer is offered either item prices or item-specific probabilities of free allocation and chooses their surplus-maximizing option. In the matroid-constrained setting, each buyer instead chooses a surplus-maximizing independent set at posted item prices, or their preferred free option among independent sets and randomizations between two independent sets. Buyers may always choose not to participate.

The symmetric guarantees stated above are proved in Appendix B. Section 4 presents the pricing-and-lottery construction that extends to general independent values.

Our Approach: A Single-Buyer Upper Bound. We decompose utility into the contribution from each buyer. Typically, the revenue maximization literature has first solved the single-buyer problem as a foothold to the multi-buyer problem. However, this poses a problem for utility maximization, as it is the competition between buyers that requires us to use ordeals. For just a single buyer, we can offer them any item for free.

Instead, we use the *ex ante relaxation* [1, 11, 14] to limit how often each item can be allocated. These probabilities come from a multi-buyer mechanism and determine the probability that item j is sold to buyer i in expectation over *all* of the multi-buyer value profiles $\mathcal{V}$ that could be drawn. The sum of the optimal utility from each single-buyer ex-ante-constrained problem provides an upper bound on the optimal multi-buyer utility.

With this constraint on allocation, even in a single-buyer problem, the buyer's different possible types now compete for each item's limited allocation. Allocating an item to one type leaves less probability for other types, but incentivizing the buyer to choose an alternative requires costly payments. Goldner et al. [27] show that achieving optimal utility under these constraints can require infinitely many menu options, even for a single unit-demand buyer with two independent item values. We therefore focus on simple mechanisms that achieve a constant-factor approximation, and develop an upper bound on the ex-ante-constrained utility to prove this guarantee.

To form our upper bound, we aim to decompose the multi-item utility into single-item terms. We do so by considering each item's *deletion loss*: how much utility is lost in the multi-item mechanism if we remove the item. We must consider that this term may be large (this item is uniquely the buyer's favorite) or small (there are many other good options for the buyer).

To account for both possibilities, we consider the utility from two independent draws from the buyer's value distribution, $\mathbf{V}$ and $\mathbf{W}$. Their average has the same expectation as a single draw, so we can write it as

$$\mathbb{E}[u(\mathbf{V})] = \mathbb{E}[\min\{u(\mathbf{V}), u(\mathbf{W})\}] + \frac{1}{2}\mathbb{E}[|u(\mathbf{V}) - u(\mathbf{W})|].$$

In Section 3, we show that this first term is the portion of utility bounded by free lotteries, and this second term is bounded by deletion loss, which we can then bound using item-pricing. Sections 4 and 5 construct truthful single-buyer multi-item mechanisms that then approximate these terms, and convert them into a multi-buyer mechanism.

Our approach parallels the core-tail analysis of revenue maximization [29, 3], where a cutoff on item value separates items into the outliers (in the tail) and the typical values (in the core), and then simple mechanisms are used to bound each portion. In contrast, rather than divide items based on their realizations, we divide the utility contribution below and above a threshold. This is most similar to a prophet inequality analysis [38, 34]. We then show that simple mechanisms bound each portion.

Cai et al. [10] gave a unifying framework for generating an upper bound on optimal revenue across multiple valuation classes. There are several impediments to applying that framework here, discussed in Goldner and Lundy [26]. First, the CDW framework decomposes revenue into single-buyer contributions and generates upper bounds for each separately. As discussed, that does not work in our setting, as the unconstrained single-buyer problem is trivial. Second, even for the ex-ante-constrained single-buyer problem, the analog of the canonical revenue flows can give bounds no better than welfare. Note also that the copies setting, where each buyer is replaced with single-minded copies, is not necessarily an upper bound on utility [26]. These limitations motivate our upper bound, which incorporates the ex-ante allocation constraints and separates utility into contributions that prices and free lotteries can recover.

Combining Buyers Using Contention Resolution. Our multi-buyer mechanism just runs the single-buyer ex-ante-constrained mechanism for each buyer independently. The only problem, then, is when buyers request the same item. We resolve this contention by making every item equally available to each buyer with a constant probability, so the buyers have no incentive to prefer one item over another due to the competition.

When buyer values are drawn symmetrically, resolving contention uniformly-at-random is truthful and guarantees at least a $(1-1/e)$-fraction of utility for each buyer (Section 4.2). Without item symmetry, independently chosen single-buyer mechanisms might concentrate demand on the same items, making contention impossible to resolve without great utility loss. To avoid this, we move the item-supply constraints into the objective with Lagrangian multipliers that let us control the demand for each item, causing over-allocated items to incur a penalty. These *demand weights* let each buyer's prices and lotteries account for competition from other buyers, while still allowing us to construct their mechanisms separately.

For every fixed choice of demand weights, our single-buyer analysis guarantees that the combined Lagrangian objective is at least a constant fraction of optimal utility. To return this to a

guarantee on the utility of an ex-ante feasible mechanism, we update the weights over multiple rounds, increasing an item's weight when its expected demand exceeds supply and decreasing it when demand falls below supply. A no-regret analysis shows that the average utility is large enough to cover the loss from reducing participation, which restores ex-ante feasibility. The mechanism independently samples a round for each buyer and offers them the prices or free-allocation options computed in that round. Finally, we resolve contention so that every request is served with probability exactly $1 - 1/e$, preserving buyers' choices (Section 5.2).

Beyond Unit Demand. Our approach extends to buyers who value feasible sets additively, where each buyer's feasible sets form a known matroid (Section 6). The main differences from the unit demand case concern the single-buyer mechanisms. Under prices, the buyer can now choose a surplus-maximizing feasible set, while when facing free lotteries, the buyer may choose randomizations over (at most two) feasible sets. The demand-weight updates and contention-resolution arguments extend without any significant changes.

Related Work. The question of single-dimensional Bayesian utility maximization was completely resolved by Hartline and Roughgarden [31]. For multidimensional utility maximization, several works have given logarithmic approximations to the upper bound of social welfare. Fotakis et al. [23] give an $O(\log |\Omega|)$ approximation to welfare, where Ω is the set of feasible outcomes. More recently, Ezra [20] obtains the tight factor $H_n = \sum_{k=1}^{n} 1/k$ against welfare for arbitrary nonnegative valuations over a finite outcome space, and gives reductions from truthful welfare approximations to utility approximations.

Our mechanism is most similar to the Favorites mechanisms of Goldner and Lundy [26]: each buyer chooses their favorite item, and a fixed single-item mechanism allocates among buyers who choose the same item. Our mechanism differs in that, rather than being prior-free, it uses known distributions to approximate optimal utility. However, our favorite-item mechanism (Appendix B) builds on the Favorites approach, showing that appropriately chosen prices and contention resolution do in fact give a constant-factor approximation to optimal utility for item-i.i.d. values, even when buyers have different distributions. In contrast to our mechanism, Ezra et al. [21] design the VCG-with-copies mechanism that randomizes the number of copies of each item, runs VCG on that instance, and rounds the allocation and payment to satisfy the original supply constraints. Their mechanism is also prior-free.

Our work follows the approach of designing simple mechanisms to approximate multidimensional revenue [12, 3, 10] and draws on many tools from this literature. We use the ex ante relaxation to reduce to single-buyer problems, as in Alaei [1], Chawla and Miller [11], Chawla et al. [14], and similarly to Bhawalkar et al. [7]. Our use of single-dimensional mechanisms to bound the utility of a multi-item mechanism parallels the revenue bounds of Chawla et al. [12, 13].

Our use of deletion to bound optimal utility is conceptually related to Eden et al. [15], who in the interdependent values setting, decompose welfare using the value remaining after deleting an agent's signal and the corresponding deletion loss, and bound the resulting contributions using different mechanisms. Related per-signal deletion bounds also appear in Eden et al. [16, 17]. We similarly use per-item deletion to obtain single-dimensional quantities, combining bounds on these deletion losses with a bound on the remaining utility.

There has been a recent surge in the study of utility maximization. This includes utility maximization in market regulation [19], online allocation with predictions [28], budget-feasible

procurement [18], and dynamic or constrained auctions [5, 6]. Related work in additional domains examines waiting times in healthcare [8], verification in resource allocation [35, 37], utility and retention in mobile games [36], and combinatorial pen testing [24]. One particularly related work is Bei and Huang [4], who give a general reduction from welfare algorithms to consumer-utility mechanisms in downward-closed Bayesian environments with finite type distributions. Their mechanism obtains an $\Omega(1/\log(1/\delta))$ fraction of the input algorithm's expected welfare, where δ is the smallest positive probability of any buyer's type. Their result allows general valuations with possible correlation across items; by assuming independence across item values, we obtain a constant-factor approximation to optimal utility.

2 Preliminaries

We consider the setting where the mechanism designer wants to sell one copy of each of $m \geq 1$ heterogeneous items to $n \geq 1$ buyers. Buyer i has a value vector $v_i = (v_{i1}, \ldots, v_{im}) \in \mathbb{R}_+^m$ for the m items. Buyer-item values are drawn independently from known distributions $V_{ij} \sim F_{ij}$ on $\mathbb{R}_+$. We write $\mathbf{V}_i = (V_{i1}, \ldots, V_{im})$ for buyer i's random value vector, and $\mathcal{V} = (\mathbf{V}_1, \ldots, \mathbf{V}_n)$ for the random instance. We assume $0 < \mathbb{E}[V_{ij}] < \infty$ throughout, so that the benchmarks we compare against are positive and finite.

A mechanism $\mathcal{M}$ takes a reported value profile $v = (v_1, \ldots, v_n)$ and specifies a possibly randomized allocation of items and nonnegative payments for each buyer.[1] We write $\mathcal{M} = (x, p)$, where $x_{ij}(v)$ is the probability that buyer i receives item j and $p_i(v) \geq 0$ is their expected payment, averaging over the mechanism's internal randomness.

Throughout the paper, we consider two valuation settings for the buyers:

- **Unit-demand:** Each buyer i has value function $v_i(S) = \max_{j \in S} v_{ij}$ for a nonempty item set $S \subseteq [m]$, with $v_i(\varnothing) = 0$. In this setting, we want to compute a *matching-feasible* allocation: an allocation in which each buyer receives at most one item, and each item is allocated to at most one buyer.

- **Constrained-additive:** Each buyer i has a known nonempty downward-closed family $\mathcal{I}_i$ of feasible item sets and values a feasible set $S \in \mathcal{I}_i$ at $\sum_{j \in S} v_{ij}$, with $v_i(\varnothing) = 0$. An allocation must give each buyer a set in $\mathcal{I}_i$ and allocate each item to at most one buyer.

We provide results for unit-demand buyers and constrained-additive buyers with matroid constraints (defined in Section 2.1).

For the incentive constraints, define buyer i's interim allocation and payment rules as the averages over the other buyers' values:

$$\bar{x}_i(v_i) := \mathbb{E}_{\mathcal{V}_{-i}}[x_i(v_i, \mathcal{V}_{-i})], \qquad \bar{p}_i(v_i) := \mathbb{E}_{\mathcal{V}_{-i}}[p_i(v_i, \mathcal{V}_{-i})].$$

A mechanism is Bayesian incentive compatible (BIC) if, for every true type v_i and alternative report w_i,

$$u_i(v_i) := v_i \cdot \bar{x}_i(v_i) - \bar{p}_i(v_i) \geq v_i \cdot \bar{x}_i(w_i) - \bar{p}_i(w_i).$$

[1] Allowing negative payments substantially changes the problem: unrestricted subsidies make the utility objective unbounded.

A mechanism is interim individually rational if, for every true type v_i,

$$u_i(v_i) \geq 0.$$

Our constructed mechanisms are BIC and satisfy ex-post individual rationality: truthful utility is nonnegative in every realized outcome, a stronger requirement than interim individual rationality. All our mechanisms include a non-participation option with zero probability of allocation and zero payment. Any ties with this non-participation option are always resolved in favor of not participating. Unless specified otherwise, we break all other ties uniformly at random.

For a mechanism $\mathcal{M}$ define its *residual surplus* as the expected total buyer utility $U(\mathcal{M}) := \sum_i \mathbb{E}_{\mathbf{V}_i}[u_i^{\mathcal{M}}(\mathbf{V}_i)]$. Our objective is to design mechanisms that maximize residual surplus $U(\mathcal{M})$, and we will compare against $\mathrm{OPT}^{n,m}$, which denotes the supremum over feasible, BIC, interim individually rational mechanisms with nonnegative payments.

For a real number z, write $z^+ := \max\{z, 0\}$ for its positive part.

2.1 Matroid Constraints

A nonempty downward-closed family $\mathcal{I}_i$ is a *matroid* if it satisfies the *exchange property*: if $A, B \in \mathcal{I}_i$ and $|A| < |B|$, some item $j \in B \setminus A$ has $A \cup \{j\} \in \mathcal{I}_i$. Its feasible sets are called *independent sets*. Unit demand is the special case $\mathcal{I}_i = \{T \subseteq [m] : |T| \leq 1\}$; unconstrained additive buyers correspond to $\mathcal{I}_i = 2^{[m]}$.

The *rank* $r_i(A)$ of an item set $A \subseteq [m]$ is the size of its largest independent subset:

$$r_i(A) := \max\{|T| : T \subseteq A,\ T \in \mathcal{I}_i\}.$$

Without loss of generality, we restrict each buyer's ground set to items j with $\{j\} \in \mathcal{I}_i$. [2]

The *matroid polytope* is

$$\mathcal{X}_i := \left\{ z \in \mathbb{R}^m_{\geq 0} : \sum_{j \in A} z_j \leq r_i(A) \text{ for every } A \subseteq [m] \right\}. \tag{2.1}$$

2.2 The Ex-Ante Relaxation

We formulate here the ex-ante relaxation to upper-bound optimal multi-buyer utility by a sum of single-buyer utilities. Each single-buyer mechanism respects its buyer's feasibility constraint, but we relax item supply: across buyers, each item is allocated at most once in expectation.

For a mechanism $\mathcal{M}$, let $q_{ij}^{\mathrm{alloc}}(\mathcal{M}) := \mathbb{E}_{\mathcal{V}}[x_{ij}^{\mathcal{M}}(\mathcal{V})]$ be buyer i's ex-ante allocation probability for item j and $q_i^{\mathrm{alloc}}(\mathcal{M}) := (q_{i1}^{\mathrm{alloc}}(\mathcal{M}), \ldots, q_{im}^{\mathrm{alloc}}(\mathcal{M}))$ be the vector of these probabilities. Now, for buyer i and item j, define *ex-ante cap* $q_{ij} \in [0, 1]$ as an upper bound on the probability of allocating item j to buyer i. Thus, mechanism $\mathcal{M}$ satisfies a cap q_{ij} if $q_{ij}^{\mathrm{alloc}}(\mathcal{M}) \leq q_{ij}$. Write $\mathbf{q}_i = (q_{i1}, \ldots, q_{im})$ and $Q_i := \sum_j q_{ij}$ for the vector and sum of ex-ante caps.

Let $\mathrm{OPT}^{1,m}(\mathbf{V}_i; \mathbf{q}_i)$ denote the supremum expected utility of a truthful, individually rational single-buyer mechanism with nonnegative payments that allocates to buyer i a set in $\mathcal{I}_i$ and satisfies ex-ante cap q_{ij} for each item j. For independent buyers, consider any feasible, BIC,

[2] By downward closure, any excluded item belongs to no feasible set for that buyer and can be discarded. Buyers with no remaining items receive nothing.

interim individually rational mechanism C with nonnegative payments. Item supply holds in every realization, so $\sum_i q_{ij}^{\text{alloc}}(C) \le 1$ for every item j. Set ex-ante caps $q_{ij} := q_{ij}^{\text{alloc}}(C)$ for every buyer i and item j, and observe that buyer i's interim mechanism satisfies these caps, so $\text{OPT}^{1,m}(\mathbf{V}_i; \mathbf{q}_i)$ upper-bounds their expected utility under C. Summing over buyers and optimizing over caps that satisfy expected item supply gives

$$U(C) \le \sum_{i=1}^{n} \text{OPT}^{1,m}(\mathbf{V}_i; q_i^{\text{alloc}}(C)) \le \sup_{\substack{\mathbf{q}_i \in [0,1]^m \ \forall i \\ \sum_i q_{ij} \le 1 \ \forall j}} \sum_{i=1}^{n} \text{OPT}^{1,m}(\mathbf{V}_i; \mathbf{q}_i).$$

Since this holds for every feasible comparison mechanism C, taking the supremum over all such mechanisms yields the ex-ante relaxation:

$$\text{OPT}^{n,m} \le \sup_{\substack{\mathbf{q}_i \in [0,1]^m \ \forall i \\ \sum_i q_{ij} \le 1 \ \forall j}} \sum_{i=1}^{n} \text{OPT}^{1,m}(\mathbf{V}_i; \mathbf{q}_i). \tag{2.2}$$

2.3 The Single-Dimensional Benchmark for a Single Buyer

Let X be a nonnegative single-dimensional value with finite expectation. A single-dimensional mechanism specifies an allocation probability $a(z) \in [0,1]$ and an expected payment $\tau(z) \ge 0$ for each report $z \ge 0$. For $q \in [0,1]$, define

$$\text{OPT}^{1,1}(X; q) := \max \left\{ \mathbb{E}[X a(X) - \tau(X)] : (a, \tau) \text{ is BIC and IR, } \mathbb{E}[a(X)] \le q \right\}. \tag{2.3}$$

By a slight abuse of notation, we use $\text{OPT}^{1,1}(X; q)$ for both the optimal mechanism and its expected utility. This mechanism can always be implemented by drawing a (possibly infinite) posted price independently of the value from at most 2 (price) options; see Appendix C.1 for this characterization and computational considerations.

2.4 Computational Assumptions

We assume that access to the value distributions permits exact evaluation of the expectations and event probabilities required by our mechanisms. These evaluations involve one or several independent draws from the value distributions and account for the mechanisms' internal randomness and specified tie-breaking rules. For matroid-constrained buyers, we additionally assume exact evaluation of item allocation probabilities under the matroid single-buyer mechanisms.

For the optimization tasks required by our constructions, we assume oracle access to the single-dimensional benchmark $\text{OPT}^{1,1}$, returning its value and an implementing price distribution. For matroid-constrained buyers, we additionally assume access to membership tests for each buyer's matroid and a free-lottery utility oracle (defined in Section D.1.1).

Our algorithms use polynomially many oracle calls and polynomial additional computation.

3 A Single Buyer Ex-Ante Utility Upper Bound

In this section, we develop the single-buyer bounds used throughout the paper. We first upper-bound ex-ante-constrained utility for independent constrained-additive values. We then show how, for unit-demand buyers, posted item prices recover the single-dimensional terms in this bound.

For ease of notation, in this single buyer section we drop buyer subscripts i; subscripts j only refer to items. We then bound the utility of *any* given truthful and individually rational mechanism (x, p) with non-negative payments. As the mechanism is fixed, our analysis is in terms of utility $u(v) := v \cdot x(v) - p(v)$, but note that this is a function of the mechanism (x, p) as well. Let q_j be item j's ex ante item cap, and hence an upper bound on the ex-ante probability of allocating item j in x.

Decomposing Utility: Our end goal is to upper-bound the utility of any mechanism using a collection of simple mechanisms. Natural candidates for these simple mechanisms are free lotteries, which incur no cost but cannot screen types, and posted prices, which can save the item for higher-valued types in exchange for incurring a cost. We therefore first ask: when is it better to use a free lottery instead of a price and when is it necessary to allocate extra to high types, treating them differently?

When the types concentrate, there is little benefit to treating them differently, and a lottery can suffice. However, when there is substantial variation, prices can distinguish and benefit the higher types. This is precisely the intuition behind our utility benchmark. We consider two independent draws $\mathbf{V}$ and $\mathbf{W}$, whose average utility has the same expectation $\mathbb{E}_{\mathbf{V}}[u(\mathbf{V})]$. For each pair of draws, we decompose their average utility into the *overlapping* utility present in both draws—the smaller of their two utilities—and the *variation* from there to their average. Observe that if utility is constant across draws, such as in a point mass distribution, then all utility is in the overlap term. More generally, taking the minimum removes the excess utility present in only one draw. In essence, the overlap term is then the portion of utility for which a lottery gives a good approximation; the variation term is the portion for which we need to use prices to differentiate types.

Note that while different from the classic core-tail analysis in revenue, there are multiple parallels. We too divide our bound into a term that one candidate mechanism can cover and the excess, then bounded by the other candidate mechanism. Concentration or lack thereof helps explain the difference between the terms.

Formally, with independent draws $\mathbf{V}$ and $\mathbf{W}$ from the buyer's value distribution, we get the following decomposition

$$\begin{aligned}\mathbb{E}_{\mathbf{V}}[u(\mathbf{V})] &= \frac{\mathbb{E}_{\mathbf{V}}[u(\mathbf{V})] + \mathbb{E}_{\mathbf{W}}[u(\mathbf{W})]}{2} \\ &= \underbrace{\mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\min\{u(\mathbf{V}), u(\mathbf{W})\}\right]}_{\text{overlap term}} + \underbrace{\frac{1}{2}\mathbb{E}_{\mathbf{V},\mathbf{W}}\left[|u(\mathbf{V}) - u(\mathbf{W})|\right]}_{\text{variation term}}.\end{aligned} \tag{3.1}$$

Overlap term: One reason expected utility may be high and difficult to upper-bound is because of correlation between the values and allocation rules: when a buyer's values are higher, their allocation may also be higher. For the first term in our decomposition, the minimum utility of two independent value draws $\mathbf{V}$ and $\mathbf{W}$, we show in Equation (3.6) that this is upper-bounded by replacing the correlation with independence, using the value and allocation from independent draws.

More precisely, we show that this component of utility is well-approximated by a lottery. Using the AM-GM inequality gives that this is upper-bounded in expectation by allocating to a buyer with value $\mathbf{V}$ with the allocation probability of a fresh draw $\mathbf{W}$, eliminating the correlation between value and allocation. Taking expectations, independence lets us upper-bound each component

separately by the expected highest item value multiplied by the expected allocation probability, which is at most the sum of the ex-ante allocation caps across items, Q.

For ex ante caps stemming from a unit-demand mechanism, $Q \leq 1$, and the term equals the expected utility from giving the buyer their favorite item for free with probability Q. Hence the expected minimum utility of two draws is upper-bounded by this free lottery's utility. In the i.i.d. setting with common item caps $q_j = Q/m$, item symmetry ensures that the lottery also respects each item's cap. In the non-i.i.d. setting, we bound the term using item-specific free lotteries, where each item's probability of allocation is chosen in order to balance utility against demand.

Variation term: The variation term measures the expected difference in utility between independent draws, and we show that this component of utility is well-approximated by item prices. By symmetry, half of the expected absolute difference between $u(\mathbf{V})$ and $u(\mathbf{W})$ equals the expected positive part of $u(\mathbf{V}) - u(\mathbf{W})$. To bound this difference, we consider the more extreme effect of removing each item's value entirely, when positive. That is, we replace the coordinates of $\mathbf{W}$ with those of $\mathbf{V}$, one at a time, and observe that the total increase in utility is at most the sum of the positive increases at each step. Since $W_j \geq 0$, increasing coordinate j from 0 to V_j can only give a larger gain than replacing W_j with V_j. This gain is exactly the loss from zeroing out coordinate j in the resulting vector.[3] Independence ensures that each intermediate vector has the same distribution as $\mathbf{V}$, so its expected *deletion loss* is D_j. Summing over coordinates bounds the variation term by D.

Formally we define the total expected deletion loss D of a mechanism as

$$D := \sum_{j=1}^{m} D_j, \qquad D_j := \mathbb{E}_{\mathbf{V}}\big[u(\mathbf{V}) - u(\mathbf{V}_{-j}, 0)\big]. \tag{3.2}$$

The variation term is bounded by this deletion loss:

$$\frac{1}{2}\mathbb{E}_{\mathbf{V},\mathbf{W}}\big[|u(\mathbf{V}) - u(\mathbf{W})|\big] \leq D. \tag{3.3}$$

The full calculation appears in Appendix A.

We bound each of these terms to establish our utility benchmark in Theorem 1. In Lemma 2, we show how to attain expected utility equal to D_j via an implementable single-item mechanism, hence upper-bounding D by the sum of utility from each optimal ex-ante-constrained single-item mechanism. Finally, in Lemma 3, we show that running the per-item mechanisms simultaneously to a unit-demand buyer yields a constant fraction of the sum of the single-item utilities.

Lemma 1. *Consider any truthful, individually rational single-buyer mechanism with nonnegative payments, utility u, and ex-ante probability of allocating item j at most q_j. For any $\ell_1, \ldots, \ell_m \in (0, 1]$,*

$$\mathbb{E}_{\mathbf{V},\mathbf{W}}\big[\min\{u(\mathbf{V}), u(\mathbf{W})\}\big] \leq \mathbb{E}_{\mathbf{V}}\left[\max_j \ell_j V_j\right] \sum_j \frac{q_j}{\ell_j}. \tag{3.4}$$

The parameters ℓ_j describe free item lotteries: choosing item j gives it to the buyer with probability ℓ_j, so their expected utility from the best lottery is $\max_j \ell_j V_j$.

[3] If zero lies outside a coordinate's support, extend the mechanism to the resulting reports by choosing utility-maximizing allocation–payment pairs among its original offers and their limits, including non-participation.

Proof. We bound the overlap term in Equation (3.1). First, observe that utility is of course upper-bounded by welfare as payments are nonnegative. Then, for each item j, we can rewrite its contribution to the *welfare*, $v_j \cdot x_j(v)$ as the *utility* of a free lottery that selects j with probability ℓ_j, namely $\ell_j v_j$, scaled by the weight $x_j(v)/\ell_j$. Summing over items and replacing each lottery's utility by the utility of the buyer's favorite lottery gives an upper bound: the favorite lottery's utility multiplied by the total scaling weight $\sum_j x_j(v)/\ell_j$.

Formally, individual rationality implies

$$0 \le u(v) \le \sum_j v_j \cdot x_j(v) = \sum_j (\ell_j v_j) \cdot \left(\frac{x_j(v)}{\ell_j}\right) \le \left(\max_j \ell_j v_j\right) \left(\sum_j \frac{x_j(v)}{\ell_j}\right). \tag{3.5}$$

As we know from [26], a welfare bound for the full expected utility can be too large to yield a constant approximation. However, by applying Equation (3.5) only to the overlap term, which is smaller than $\mathbb{E}[u(\mathbf{V})]$, we get a small enough bound:

$$\begin{aligned}
&\mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\min\{u(\mathbf{V}), u(\mathbf{W})\}\right] && (3.6)\\
&\le \mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\sqrt{u(\mathbf{V}) \cdot u(\mathbf{W})}\right] && u(\mathbf{V}), u(\mathbf{W}) \ge 0\\
&\le \mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\sqrt{\left(\max_j \ell_j V_j\right)\left(\sum_j \frac{x_j(\mathbf{V})}{\ell_j}\right) \cdot \left(\max_j \ell_j W_j\right)\left(\sum_j \frac{x_j(\mathbf{W})}{\ell_j}\right)}\right] && (3.5)\\
&\le \frac{1}{2}\mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\left(\max_j \ell_j V_j\right)\sum_j \frac{x_j(\mathbf{W})}{\ell_j} + \left(\max_j \ell_j W_j\right)\sum_j \frac{x_j(\mathbf{V})}{\ell_j}\right] && \text{AM-GM}\\
&= \frac{1}{2}\left(\mathbb{E}_{\mathbf{V}}\left[\max_j \ell_j V_j\right]\mathbb{E}_{\mathbf{W}}\left[\sum_j \frac{x_j(\mathbf{W})}{\ell_j}\right] + \mathbb{E}_{\mathbf{W}}\left[\max_j \ell_j W_j\right]\mathbb{E}_{\mathbf{V}}\left[\sum_j \frac{x_j(\mathbf{V})}{\ell_j}\right]\right) && \text{independence}\\
&= \mathbb{E}_{\mathbf{V}}\left[\max_j \ell_j V_j\right]\mathbb{E}_{\mathbf{V}}\left[\sum_j \frac{x_j(\mathbf{V})}{\ell_j}\right]\\
&\le \mathbb{E}_{\mathbf{V}}\left[\max_j \ell_j V_j\right]\sum_j \frac{q_j}{\ell_j}, && \mathbb{E}_{\mathbf{V}}[x_j(\mathbf{V})] \le q_j\\
& && (3.7)
\end{aligned}$$

where the two equalities use that the two vector draws are i.i.d. □

Writing $U := \mathbb{E}_{\mathbf{V}}[u(\mathbf{V})]$, combining the overlap bound with the decomposition (3.1) and the variation bound (3.3) gives

$$U - D \le \mathbb{E}_{\mathbf{V}}\left[\max_j \ell_j V_j\right]\sum_j \frac{q_j}{\ell_j}. \tag{3.8}$$

We next bound the deletion loss by single-item utilities, showing that, for each item j, the deletion loss D_j is the expected utility attained by a single-item mechanism with ex-ante allocation probability at most q_j. Thus, D_j is upper-bounded by the *optimal* single-item ex-ante-constrained utility $\mathrm{OPT}^{1,1}(V_j; q_j)$.

Lemma 2. *For any truthful, individually rational single-buyer mechanism with nonnegative payments and utility u, let D_j be its expected deletion loss for item j, and let q_j be an upper bound on its ex-ante probability of allocating that item. Then, for every item j,*

$$D_j \le \mathrm{OPT}^{1,1}(V_j; q_j).$$

Proof. For a fixed item j, we construct a single-dimensional mechanism that has utility equal to item j's deletion loss D_j. It takes 1-dimensional input V_j by resampling fresh draws to use as values for the other items.

That is, draw fresh samples $Z_k \sim F_k$ independently for every item $k \neq j$; we write $\mathbf{Z}_{-j} := (Z_k)_{k\neq j}$. The agent then reports their bid b_j for item j, and averaging over these draws $\mathbf{Z}_{-j}$, we define the function

$$d_j(b_j) := \mathbb{E}_{\mathbf{Z}_{-j}}\big[u(\mathbf{Z}_{-j}, b_j) - u(\mathbf{Z}_{-j}, 0)\big].$$

For a realized value v_j, $d_j(v_j)$ is the expected utility lost by deleting item j, averaging over the other item values. A single-item mechanism whose truthful utility is $d_j(v_j)$ therefore has expected utility D_j. Indeed, taking expectations over V_j gives

$$\mathbb{E}_{V_j}[d_j(V_j)] = \mathbb{E}_{V_j,\mathbf{Z}_{-j}}[u(\mathbf{Z}_{-j}, V_j) - u(\mathbf{Z}_{-j}, 0)] = \mathbb{E}_{\mathbf{V}}[u(\mathbf{V}) - u(\mathbf{V}_{-j}, 0)] = D_j.$$

We now define the single-item mechanism (a, τ). The allocation rule a averages the original mechanism's allocation probability x_j over the sampled values of the other items: $a(b_j) := \mathbb{E}_{\mathbf{Z}_{-j}}\big[x_j(\mathbf{Z}_{-j}, b_j)\big]$. Then in expectation, it continues to satisfy the ex-ante allocation cap q_j, as

$$\mathbb{E}_{V_j}[a(V_j)] = \mathbb{E}_{V_j,\mathbf{Z}_{-j}}[x_j(\mathbf{Z}_{-j}, V_j)] = \mathbb{E}_{\mathbf{V}}[x_j(\mathbf{V})] \le q_j.$$

We choose the payment so that truthful reporting gives expected utility $d_j(v_j)$, the average deletion loss for v_j: define $\tau(b_j) := b_j a(b_j) - d_j(b_j)$.

Now, for the buyer's true single-dimensional value $v_j \ge 0$ and for any alternate bid $b_j \ge 0$, truthfulness of the original single-buyer mechanism that implements u gives

$$\begin{aligned}
u(\mathbf{Z}_{-j}, v_j) &= v_j\, x_j(\mathbf{Z}_{-j}, v_j) + \sum_{k\neq j} Z_k x_k(\mathbf{Z}_{-j}, v_j) - p(\mathbf{Z}_{-j}, v_j) \\
&\ge v_j\, x_j(\mathbf{Z}_{-j}, b_j) + \sum_{k\neq j} Z_k x_k(\mathbf{Z}_{-j}, b_j) - p(\mathbf{Z}_{-j}, b_j) && \text{IC} \\
&= \left(b_j\, x_j(\mathbf{Z}_{-j}, b_j) + \sum_{k\neq j} Z_k x_k(\mathbf{Z}_{-j}, b_j) - p(\mathbf{Z}_{-j}, b_j) \right) + (v_j - b_j) x_j(\mathbf{Z}_{-j}, b_j) \\
&= u(\mathbf{Z}_{-j}, b_j) + (v_j - b_j) x_j(\mathbf{Z}_{-j}, b_j),
\end{aligned}$$

where the final two equalities add and subtract the reported value's contribution to item j and use the definition of u.

Subtracting $u(\mathbf{Z}_{-j}, 0)$ from both sides and taking expectations over $\mathbf{Z}_{-j}$ gives

$$d_j(v_j) \ge d_j(b_j) + (v_j - b_j) a(b_j). \tag{3.9}$$

Since $d_j(0) = 0$ and $a(b_j) \ge 0$, inequality (3.9) establishes the required incentive and payment properties:

Truthfulness: For true value v_j and alternative report b_j,

$$\begin{aligned} v_j a(v_j) - \tau(v_j) &= d_j(v_j) \\ &\geq d_j(b_j) + (v_j - b_j)a(b_j) \\ &= v_j a(b_j) - \tau(b_j). \end{aligned}$$

Individual rationality: Setting $b_j = 0$ gives

$$d_j(v_j) \geq d_j(0) + v_j a(0) = v_j a(0) \geq 0.$$

Nonnegative payments: Setting $v_j = 0$ and using $d_j(0) = 0$ gives

$$\tau(b_j) = b_j a(b_j) - d_j(b_j) \geq 0.$$

We have thus constructed a single-dimensional mechanism with expected utility D_j and ex-ante probability of allocation at most q_j. Since $\mathrm{OPT}^{1,1}(V_j; q_j)$ is the optimal expected utility among such mechanisms, we conclude that

$$D_j \leq \mathrm{OPT}^{1,1}(V_j; q_j).$$

□

Combining the decomposition with the overlap bound and the bound on deletion loss gives the following single-buyer utility benchmark. None of these bounds requires the mechanism to allocate at most one item: they apply whenever utility is linear in independent type parameters and payments are nonnegative. Thus, we can apply the same tools in the constrained additive setting as well.

Theorem 1 (Single-buyer ex-ante utility bound)**.** *Consider a buyer with independent nonnegative type parameters $\mathbf{V} = (V_1, \ldots, V_m)$ and a feasible outcome set $\mathcal{F} \subseteq [0,1]^m$. A mechanism specifies an expected outcome $x(b) \in \mathcal{F}$ and an expected payment $p(b)$ for each report $b \in \mathbb{R}_+^m$. A buyer of type v who reports b obtains expected utility $v \cdot x(b) - p(b)$.*

For $\mathbf{q} \in [0,1]^m$, let $\mathrm{OPT}_{\mathcal{F}}^{1,m}(\mathbf{V}; \mathbf{q})$ be the supremum expected utility over truthful, individually rational mechanisms with nonnegative payments satisfying $\mathbb{E}_{\mathbf{V}}[x_j(\mathbf{V})] \leq q_j$ for every j. Then, for any $\ell_1, \ldots, \ell_m \in (0,1]$,

$$\mathrm{OPT}_{\mathcal{F}}^{1,m}(\mathbf{V}; \mathbf{q}) \leq \sum_{j=1}^{m} \mathrm{OPT}^{1,1}(V_j; q_j) + \mathbb{E}_{\mathbf{V}}\left[\max_j \ell_j V_j\right] \sum_{j=1}^{m} \frac{q_j}{\ell_j}. \tag{3.10}$$

Recovering the single-item terms for unit-demand buyers. One remaining challenge is that the sum of these single-item utilities is not obviously achievable via an implementable multi-dimensional mechanism. To approximate this bound, we compare mechanisms $\mathcal{M}_j^1$ for the single-item environment (where the buyer is only interested in item j) with a mechanism $\mathcal{M}^{UD}$ for the multi-item environment (where, for unit-demand values, the buyer is only interested in purchasing at most one item, which may or may not be item j). The following lemma shows that by combining the single-item posted price mechanisms into a multi-item posted price mechanism with

unit-demand constraints, not much utility is lost, and the multi-item mechanism continues to satisfy the given ex ante constraints. To recover the single-item utilities in the benchmark using a feasible mechanism, we recall the posted-price representation of $\mathrm{OPT}^{1,1}(V_j; q_j)$ from Section 2.3: it is implemented by posting a randomly drawn price over support at most 2.

Lemma 3. *Let $V_1, \ldots, V_m$ be independent item values, and P_j be an independent (possibly randomized) nonnegative price for item j for which $\Pr_{V_j,P_j}(V_j > P_j) \le q_j$, which for the one-item setting yields expected utility $U(\mathcal{M}_j^1)$. If $\sum_j q_j \le 1$, then the mechanism $\mathcal{M}^{UD}$ that offers these prices together and lets the buyer choose their utility-maximizing option allocates item j with probability at most q_j and has expected utility*

$$U(\mathcal{M}^{UD}) \ge (1 - 1/e) \sum_j U(\mathcal{M}_j^1).$$

To parse this statement, observe that the assumption $\Pr_{V_j,P_j}(V_j > P_j) \le q_j$ corresponds to a posted price for the single-item setting satisfying the ex ante constraint q_j. The assumption that $\sum_j q_j \le 1$ holds when the ex ante constraints q_j are given by the multi-item setting, where we have a feasibility constraint that the mechanism allocates at most one item, which is without loss for unit-demand valuations.

Proof. Let $\mathbf{P} = (P_1, \ldots, P_m)$ be independent (possibly randomized) nonnegative posted prices for each item. Consider first the single-item setting, where only item j is for sale. Define the surplus from posted price P_j as $S_j := (V_j - P_j)^+$. Then the expected utility of posted price P_j is $U(\mathcal{M}_j^1) := \mathbb{E}_{V_j,P_j}[S_j]$.

Now consider the multi-item posted price mechanism $\mathcal{M}^{UD}$ with a unit-demand feasibility constraint. That is, post prices $\mathbf{P}$ and let the buyer choose one item. In this setting, let $q_j^{\mathrm{alloc}}(\mathcal{M}^{UD})$ denote the ex ante probability that mechanism $\mathcal{M}^{UD}$ allocates item j. Observe that it is bounded by the probability of sale in the single-item posted price setting, and thus also the single-item cap q_j:

$$q_j^{\mathrm{alloc}}(\mathcal{M}^{UD}) \le \Pr_{\mathbf{V},\mathbf{S}}\left(S_j > 0 \text{ and } S_j = \max_k S_k\right) \le \Pr_{V_j,P_j}(S_j > 0) \le q_j.$$

Finally, we show that the utility of the multi-item mechanism is a good approximation to the sum of that from the single-item mechanisms. For $t \ge 0$, define the expected number of items whose surplus exceeds t as $L(t) := \sum_j \Pr_{V_j,P_j}(S_j > t)$, which by the single-item ex ante constraints, is at most $\sum_j q_j \le 1$. We bound the utility of the multi-item posted prices as follows.

$$\begin{aligned}
U(\mathcal{M}^{UD}) = \mathbb{E}_{\mathbf{V},\mathbf{S}}\left[\max_j S_j\right] &= \int_0^\infty \left(1 - \prod_j \big(1 - \Pr_{V_j,P_j}(S_j > t)\big)\right) dt && \text{independence} \\
&\ge \int_0^\infty (1 - e^{-L(t)})\, dt && 1 - x \le e^{-x} \\
&\ge (1 - 1/e) \int_0^\infty L(t)\, dt && 1 - e^{-s} \ge (1 - 1/e)s,\ s \in [0,1] \\
&= (1 - 1/e) \sum_j \int_0^\infty \Pr_{V_j,P_j}(S_j > t)\, dt \\
&= (1 - 1/e) \sum_j U(\mathcal{M}_j^1).
\end{aligned}$$

□

4 A Constant-Factor Approximation for i.i.d. Buyers

In this section, we give a constant-factor approximation for unit-demand buyers whose item values are drawn independently from a common distribution F. Thus $V_{ij} \sim F$ independently for all buyers i and items j; write $V \sim F$ for a generic draw. First, we use the bound from Section 3 to construct an approximately optimal single-buyer mechanism (Section 4.1). Posted item prices recover the single-item utility terms, and free allocation recovers the overlap term. We then combine buyers through contention resolution (Section 4.2), first stating the general condition that preserves incentives and then verifying it for uniform contention resolution in the i.i.d. setting. This section gives a simple first application of our tools: symmetry makes the ex-ante caps explicit, which simplifies the choice of appropriate prices and free lotteries. This prepares for our main result for general independent unit-demand values in Section 5. The stronger favorite-item mechanism in Appendix B instead exploits symmetry to offer one common price for all items.

We use Q for a buyer's total ex-ante allocation cap. For i.i.d. item values, symmetrizing a mechanism over item labels preserves expected utility and makes its allocation probabilities equal across items, hence we set a common item cap at $q := Q/m$, with $\mathbf{q} = (q, \dots, q)$. Thus the optimum under a total cap Q equals $\mathrm{OPT}^{1,m}(\mathbf{V}; \mathbf{q})$, using the item-cap notation of Sections 2 and 3. We leave $Q \in [0, 1]$ arbitrary in the single-buyer argument and set $Q = \min\{1, m/n\}$ when combining i.i.d. buyers.

4.1 The Single-Buyer Mechanism

We build directly on Theorem 1. For any total ex-ante cap $Q \in [0, 1]$, with common item cap $q := Q/m$, define the single-buyer mechanism $\mathcal{M}^{1,m}_{\mathrm{iid}}(Q)$ as follows:

1. With probability $e/(2e-1)$, run $\mathcal{M}^{UD}$: Draw each item's price independently according to $\mathrm{OPT}^{1,1}(V; q)$. Let the buyer choose the item with the largest positive utility (if any) and charge its price.

2. With probability $(e-1)/(2e-1)$, run $\mathcal{M}^{\mathrm{free}}$: Let the buyer choose their favorite item and allocate it to them with probability Q free of charge.

To implement $\mathcal{M}^{UD}$, we use the posted-price representation of $\mathrm{OPT}^{1,1}(V; q)$ from Section 2.3. This gives a distribution over at most two prices, which we note is efficiently computable (see Appendix C.1).

Note that $\mathcal{M}^{\mathrm{free}}$ is not simply $\mathcal{M}^{UD}$ with the price set to 0, as it allocates the favorite item with the combined probability Q rather than only q, which still satisfies the per-item cap q in expectation over the buyer's favorite item.

We now show that the utility of our mechanism $\mathcal{M}^{1,m}_{\mathrm{iid}}(Q)$ is a constant-factor approximation to the optimal single-buyer ex-ante-constrained utility: its components bound the terms discussed in Theorem 1.

Lemma 4. *For i.i.d. item values and every $Q \in [0, 1]$, with $q = Q/m$ and $\mathbf{q} = (q, \dots, q)$, mechanism $\mathcal{M}^{1,m}_{\mathrm{iid}}(Q)$ is BIC, ex-post individually rational, has nonnegative payments, allocates each*

item with ex-ante probability at most q, and satisfies

$$U(\mathcal{M}_{\text{iid}}^{1,m}(Q)) \geq \frac{e-1}{2e-1}\text{OPT}^{1,m}(\mathbf{V};\mathbf{q}).$$

Proof. Apply Theorem 1 with $\ell_j = 1$ and $q_j = q$ for every item. Since $\sum_j q_j = Q$, we obtain

$$\text{OPT}^{1,m}(\mathbf{V};\mathbf{q}) \leq m\text{OPT}^{1,1}(V;q) + Q\mathbb{E}[\max_j V_j] = m\text{OPT}^{1,1}(V;q) + U(\mathcal{M}^{\text{free}}). \tag{4.1}$$

The equality holds because $\mathcal{M}^{\text{free}}$ gives the buyer their favorite item with probability Q, independently of their values. The prices in $\mathcal{M}^{UD}$ attain $\text{OPT}^{1,1}(V;q)$ for each item separately, and as $mq = Q \leq 1$, Lemma 3 gives

$$m\text{OPT}^{1,1}(V;q) \leq \frac{e}{e-1}U(\mathcal{M}^{UD}). \tag{4.2}$$

Combining Equations (4.1) and (4.2) yields

$$\text{OPT}^{1,m}(\mathbf{V};\mathbf{q}) \leq \frac{e}{e-1}U(\mathcal{M}^{UD}) + U(\mathcal{M}^{\text{free}}).$$

Using the probabilities of $\mathcal{M}_{\text{iid}}^{1,m}(Q)$ for running each sub-mechanism, we get

$$U(\mathcal{M}_{\text{iid}}^{1,m}(Q)) = \frac{e}{2e-1}U(\mathcal{M}^{UD}) + \frac{e-1}{2e-1}U(\mathcal{M}^{\text{free}}) \geq \frac{e-1}{2e-1}\text{OPT}^{1,m}(\mathbf{V};\mathbf{q}).$$

Both sub-mechanisms satisfy the item caps and let the buyer choose her utility-maximizing option. Their combination is therefore truthful, ex-post individually rational, and has nonnegative payments. □

We then conclude that the better (utility-maximizing) of these two mechanisms gives a constant-factor approximation.

Corollary 1. *Under the assumptions of Lemma 4,*

$$\max\{U(\mathcal{M}^{UD}), U(\mathcal{M}^{\text{free}})\} \geq \frac{e-1}{2e-1}\text{OPT}^{1,m}(\mathbf{V};\mathbf{q}).$$

All of our guarantees then apply to both the randomized $\mathcal{M}_{\text{iid}}^{1,m}(Q)$ and a replacement that uses only the utility-maximizing sub-mechanism.

Simpler, Stronger, and More General Guarantees. The above mechanism and analysis are provided for their simplicity and because they parallel the ideas in Section 5. We improve on this result, using an even simpler mechanism, with the result holding in the more general setting where each buyer's item values are i.i.d. but distributions may differ across buyers, that is, in the *item-i.i.d. setting*. The mechanism computes the 1-item (randomized) posted price for the buyer's distribution of their value for their favorite item and posts this price for every item. For fully i.i.d. buyers, using this mechanism for every buyer and resolving contention uniformly gives expected utility at least $(1 - 1/e)\text{OPT}^{n,m}/2$ (see Theorem 5). For buyers with different distributions, use the contention rule and computed caps in Appendix B.3 to obtain the guarantee of Theorem 6.

4.2 The Multi-Buyer Mechanism

Simultaneously running $\mathcal{M}_{\text{iid}}^{1,m}(Q)$ for each buyer may result in a single item being allocated to multiple buyers. In this section, we resolve the contention caused by multiple requests for an item, and regain matching-feasibility in the multi-buyer mechanism.

We first state the contention resolution argument used here and in Section 5. Run each single-buyer mechanism independently, treating its allocation as tentative and collecting payment only if the buyer ultimately receives the item. If a buyer is tentatively allocated item j, we call this a "request" to item j. For a free lottery with allocation probability Q, toss the allocation coin after the buyer chooses an item and before resolving contention; convert this to a request only if the coin succeeds. Buyers make their choices without observing the tentative allocations or random outcomes for other buyers.

We follow a common approach of contention resolution in showing that each request is served with constant probability c_i.

Observation 1. *Suppose each buyer i's request is served with probability c_i, regardless of their report or tentative outcome, and the buyer is charged only when served. This multiplies the utility of every report by c_i, preserving incentive compatibility and giving total expected utility $\sum_i c_i U_i$, where U_i is their utility before contention. Serving at most one request per item ensures feasibility; individual rationality and nonnegative payments are inherited from the single-buyer mechanisms.*

Observation 1 does not require identical distributions, only independence.

To identify c_i's for the i.i.d. case, we leverage that by item symmetry, each item has the same distribution of competing requests. Resolving conflicts uniformly-at-random then gives the same c_i for each item.

We now instantiate the argument for i.i.d. buyers. Fix

$$Q := \min\{1, m/n\}, \qquad q := Q/m = \min\{1/m, 1/n\}, \qquad \mathbf{q} = (q, \ldots, q).$$

The multi-buyer mechanism $\mathcal{M}_{\text{iid}}^{n,m}$ is as follows:

1. Run $\mathcal{M}_{\text{iid}}^{1,m}(Q)$ independently for each buyer, treating its allocation as a request. In the free sub-mechanism, a request is made only if the probability-Q allocation coin toss succeeds.

2. For each item with at least one request, serve one requester uniformly at random and charge them the corresponding price, if any.

The allocation caps ensure expected demand of at most one per item. We next show that uniform contention resolution serves each buyer's request with probability at least $1 - 1/e$. By item-symmetry, a buyer faces the same competition no matter which item they select. The given (optimal) ex-ante allocation caps ensure that the expected demand for each item is at most 1. Together, this implies that uniform contention resolution gives each request a constant probability of succeeding, as needed by Observation 1.

Theorem 2. *For i.i.d. values, mechanism $\mathcal{M}_{\text{iid}}^{n,m}$ is BIC, ex-post individually rational, has non-negative payments, respects matching-feasibility, and achieves expected utility*

$$U(\mathcal{M}_{\text{iid}}^{n,m}) \geq \frac{e-1}{2e-1}\left(1 - \frac{1}{e}\right) \text{OPT}^{n,m}.$$

Proof. Symmetry and Lemma 4 give that before contention resolution, utility is at least $\frac{e-1}{2e-1}\text{OPT}^{n,m}$. Since $q \leq 1/n$, the Favorites argument of Goldner and Lundy [26, Theorem 1] gives that each item is available to be served to each buyer with probability at least $1 - 1/e$. The result follows from Observation 1. □

5 Independent Item Values across Multiple Buyers

We now allow every $V_{ij} \sim F_{ij}$ to have a different distribution, retaining independence across items and buyers. The single-buyer upper bounds of Section 3 still apply and will drive the design of the mechanism. We emphasize that we do not have access to the optimal mechanism's ex-ante probability caps[4]. Even if we did, it is not immediately clear how one can efficiently construct single-buyer mechanisms that recover the benchmark's free-allocation term and abide by the desired item caps.

Our goal remains to construct a simple, ex-ante-feasible mechanism with approximately optimal utility that can be implemented using a contention resolution scheme. In Section 4, item symmetry ensures symmetric expected demand across items, making it easy to resolve contention while preserving a constant fraction of utility.

Without item symmetry, single-buyer mechanisms that yield high utility may do so by each allocating the same item, posing an issue for resolving this contention without great loss in utility. This concentration of demand also underlies the worst-case gap between utility and welfare for general distributions [26]. To prevent this, our plan is to set each item's menu option in a way that controls expected demand across items. We do this by moving the supply constraints into the objective with a Lagrangian dual variable, or a *demand weight* that represents the marginal cost to the objective of increasing each item's allocation. We first describe the intuition informally, leaving the precise comparisons and guarantees to the analysis below.

Rather than compare the utility of our mechanism to the optimal, we instead compare this modified Lagrangian objective for our mechanism to the optimal for any fixed dual variables, showing it achieves our desired approximation. With the guarantee holding for all demand weights, the remaining question is then how to set these variables (which instantiate the prices and probabilities in the mechanism) such that demand is not too high on any particular item, and we can resolve contention without losing too much utility.

Borrowing tools from online no-regret optimization, we choose these weights iteratively in simulated rounds using the known value distributions, increasing an item's weight when its expected demand exceeds supply and decreasing it when demand is below supply. The updated weights guide the prices and lottery probabilities in the next round, so each buyer's mechanism accounts for the demand generated by all buyers. A no-regret analysis shows that averaging the resulting mechanisms and reducing participation gives feasible expected demand while preserving our utility approximation.

In Section 5.1, we first adapt the single-buyer analysis of Section 4.1 to incorporate demand weights, constructing posted prices and free lotteries that provide a guarantee in utility and the weighted demand. We then specify the demand-weight update rule and use a no-regret analysis to show that averaging across rounds and reducing participation yields an ex-ante-feasible mechanism

[4]The ex-ante relaxation is typically formulated as a convex program, assuming computational access to the single-buyer benchmarks [1]. This access is nontrivial: even with independent item values, the linear program can be exponential in the number of items.

with approximately optimal utility. Finally, Section 5.2 gives a contention resolution scheme that converts this mechanism into a feasible BIC mechanism while preserving a constant fraction of its utility.

5.1 Demand Weights and Balancing Item Requests

Our goal is to construct single-buyer pricing and lottery mechanisms that collectively obtain high utility while respecting each item's supply of one in expectation. As a first step, we lagrangify the item-supply constraints and move them into the objective. Formally, for fixed nonnegative Lagrange multipliers λ_j, we consider the partial Lagrangian objective of a mechanism $\mathcal{M}$

$$\mathcal{L}_\lambda(\mathcal{M}) = \sum_i U_i(\mathcal{M}) + \sum_j \lambda_j \left(1 - \sum_i q_{ij}^{\text{alloc}}(\mathcal{M})\right). \tag{5.1}$$

We will refer to these multipliers λ_j as *demand weights*. The benefit of working with the partial Lagrangian is that the prices and lotteries we construct for individual buyers will have to account for the demand weights as well as utility. We can therefore choose these weights to convey information about each item's aggregate expected demand when we compute the single-buyer prices and lotteries.

For any fixed demand weights, we will construct single-buyer mechanisms, which when run simultaneously have a Lagrangian objective that is at least a constant fraction of the expected utility of any feasible multi-buyer mechanism. We do this by constructing single-buyer mechanisms with a utility guarantee that incorporates the demand weights (Lemma 5). Notice that these mechanisms do not yet guarantee high utility or ex-ante feasibility, since the item-supply constraints have been relaxed, and the Lagrangian may benefit from underallocated items with large demand weights. We will address both of these issues by running the procedure over multiple rounds with updated demand weights in Section 5.1.1. For now, we focus on the single-buyer comparison for arbitrary fixed demand weights λ_j.

For a fixed buyer i and demand weights λ_j, define the *allocation cost* of a single-buyer mechanism $\mathcal{M}_i$ as $\sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}_i)$. We construct a mixture of posted item prices and free item lotteries whose utility minus its allocation cost is at least a constant fraction of any comparison mechanism's utility, minus that mechanism's allocation cost.

Lemma 5. *For a buyer i with independent item values and any nonnegative demand-weight vector $\lambda = (\lambda_1, \ldots, \lambda_m)$, there exists a BIC, ex-post individually rational single-buyer unit-demand mechanism $\mathcal{M}_i$ with nonnegative payments such that, for every BIC, individually rational single-buyer unit-demand mechanism C_i with nonnegative payments and ex-ante allocation probabilities $q_{ij}^{\text{alloc}}(C_i)$,*

$$U(\mathcal{M}_i) - \sum_j \lambda_j q_{ij}^{\text{alloc}}(\mathcal{M}_i) \geq \frac{U(C_i)}{2 + e/(e-1)} - \sum_j \lambda_j q_{ij}^{\text{alloc}}(C_i).$$

Before proving the lemma, we show that running these single-buyer mechanisms independently in parallel would yield a mechanism whose Lagrangian objective is at least a constant fraction of $\text{OPT}^{n,m}$.

Fix a nonnegative demand weight vector λ and any feasible, BIC, interim individually rational multi-buyer mechanism C with nonnegative payments. For each buyer i, let C_i be the single-buyer mechanism obtained by averaging C over the other buyers' values. These mechanisms

preserve the buyers' expected utilities and ex-ante allocation probabilities, so $\sum_i U(C_i) = U(C)$ and $\sum_i q_j^{\text{alloc}}(C_i) \leq 1$ for every item j. The mechanisms C_i are used only as comparators and need not be computed. We now apply Lemma 5 to each buyer, obtaining mechanisms $\mathcal{M}_i$ whose guarantees hold against the corresponding mechanisms C_i. Let $\mathcal{M}$ run these mechanisms independently, and let the utility of this mechanism be $U(\mathcal{M}) = \sum_i U(\mathcal{M}_i)$. We can now lower-bound the Lagrangian objective of mechanism $\mathcal{M}$ using the guarantee of Lemma 5:

$$\begin{aligned}
\mathcal{L}_\lambda(\mathcal{M}) &= U(\mathcal{M}) + \sum_j \lambda_j \left(1 - \sum_i q_j^{\text{alloc}}(\mathcal{M}_i)\right) && \text{definition (5.1)} \\
&\geq U(\mathcal{M}) + \sum_j \lambda_j \left(\sum_i q_j^{\text{alloc}}(C_i) - \sum_i q_j^{\text{alloc}}(\mathcal{M}_i)\right) && \sum_i q_j^{\text{alloc}}(C_i) \leq 1,\ \lambda_j \geq 0 \\
&= \sum_i \left[U(\mathcal{M}_i) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}_i) + \sum_j \lambda_j q_j^{\text{alloc}}(C_i)\right] && \text{rearranging sums} \\
&\geq \frac{1}{2 + e/(e-1)} \sum_i U(C_i) && \text{Lemma 5} \\
&= \frac{U(C)}{2 + e/(e-1)}. && \text{definition of } C_i
\end{aligned}$$

This bound holds for any mechanism C, and thus taking the supremum over C yields the desired bound

$$\mathcal{L}_\lambda(\mathcal{M}) \geq \frac{\text{OPT}^{n,m}}{2 + e/(e-1)}. \tag{5.2}$$

We reiterate that this guarantee falls short of our eventual goal in two respects. First, the Lagrangian objective of $\mathcal{M}$ is at least a constant fraction of $\text{OPT}^{n,m}$, but it may receive a positive contribution from underallocated items with large demand weights, rather than from utility alone. Second, $\mathcal{M}$ need not be ex-ante feasible, since there is no constraint requiring each item's aggregate expected demand to be at most one. The updates push against both problems: increasing the weights of overallocated items discourages excess demand, while decreasing the weights of underallocated items reduces the contribution of unused supply to the Lagrangian. Both issues will be addressed in Section 5.1.1 through an iterative procedure that uses the mechanisms of Lemma 5 and updates the demand weights based on each round's aggregate expected item demand.

We now provide the proof of Lemma 5.

Proof of Lemma 5. Fix buyer i and suppress their index throughout the proof. Fix an arbitrary comparison mechanism C and let its deletion loss be $D = \sum_j D_j$.

We start again from Equation (3.8), which upper-bounds the comparison mechanism's utility $U(C)$ by its deletion loss D and a weighted free-allocation term. As in Section 4.1, we recover these contributions using posted prices and free lotteries, respectively, but now compare utility after accounting for the allocation costs of both the constructed mechanism and C.

Posted Prices $\mathcal{M}^{UD}$. First, we recover the deletion loss D through posted prices. Specifically, we construct a pricing mechanism $\mathcal{M}^{UD}$ whose utility minus its allocation cost is at least a constant fraction of any comparison mechanism's deletion loss, minus that mechanism's allocation cost:

$$U(\mathcal{M}^{UD}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}^{UD}) \geq \left(1 - \frac{1}{e}\right) D - \sum_j \lambda_j q_j^{\text{alloc}}(C). \tag{5.3}$$

In the i.i.d. setting, symmetry supplied the optimal itemwise caps for each buyer, which we used to compute the posted prices. Here, we must decide how to divide the buyer's ex-ante allocation probability among items, and then use the resulting caps to compute the corresponding prices. This choice must also account for the allocation cost of the pricing mechanism. For these two reasons, we choose the ex-ante caps on which to base our prices by solving the following program:

$$\max_{\substack{q_j \geq 0 \\ \sum_j q_j \leq 1}} \sum_j \left[\mathrm{OPT}^{1,1}(V_j; q_j) - \frac{e}{e-1} \lambda_j q_j \right]. \tag{5.4}$$

The objective is a sum of concave functions of the individual caps q_j, which allows us to approximate the optimum efficiently up to additive error (see Appendix C.2 for details). For the purposes of this proof, we assume access to optimal caps $q_j^\star$, but we account for the approximation error later, when applying Lemma 5 in Theorem 3.

Let $\mathcal{M}^{UD}$ be the mechanism that simultaneously offers the (independently drawn) prices P_j prescribed by $\mathrm{OPT}^{1,1}(V_j; q_j^\star)$ and lets the buyer pick their surplus-maximizing item (or no item). By Lemma 3, mechanism $\mathcal{M}^{UD}$ combines independent prices while preserving a $(1-1/e)$ fraction of their single-item utilities and the allocation caps. Multiplying the demand weight by $e/(e-1)$ ensures that, after applying this utility factor, the penalty is exactly $\lambda_j q_j$. Combining this guarantee with optimality of $\mathbf{q}^\star$ and feasibility of $q^{\mathrm{alloc}}(C)$ in (5.4) yields a comparison between $\mathcal{M}^{UD}$'s utility and C's deletion loss that accounts for both mechanisms' allocation costs:

$$\begin{aligned}
&U(\mathcal{M}^{UD}) - \sum_j \lambda_j q_j^{\mathrm{alloc}}(\mathcal{M}^{UD}) \\
&\quad = \mathbb{E}\left[\max_j (V_j - P_j)^+\right] - \sum_j \lambda_j q_j^{\mathrm{alloc}}(\mathcal{M}^{UD}) && \text{definition of } \mathcal{M}^{UD} \\
&\quad \geq \left(1 - \frac{1}{e}\right) \sum_j \mathrm{OPT}^{1,1}(V_j; q_j^\star) - \sum_j \lambda_j q_j^\star && \text{Lemma 3, } q_j^{\mathrm{alloc}}(\mathcal{M}^{UD}) \leq q_j^\star \\
&\quad \geq \left(1 - \frac{1}{e}\right) \sum_j \mathrm{OPT}^{1,1}(V_j; q_j^{\mathrm{alloc}}(C)) - \sum_j \lambda_j q_j^{\mathrm{alloc}}(C) && \text{optimality of } \mathbf{q}^\star \\
&\quad \geq \left(1 - \frac{1}{e}\right) D - \sum_j \lambda_j q_j^{\mathrm{alloc}}(C). && \text{Lemma 2}
\end{aligned}$$

Free Lottery $\mathcal{M}^{\mathrm{free}}$. For the free-lottery mechanism, unlike in the i.i.d. setting, we must choose the item-specific lottery probabilities without specific ex-ante caps. Our choice must also account for allocation costs, to combine with the guarantee we proved for mechanism $\mathcal{M}^{UD}$. We will construct a free-lottery mechanism $\mathcal{M}^{\mathrm{free}}$ whose utility minus its allocation cost recovers half of $U(C) - D$, minus mechanism C's allocation cost:

$$U(\mathcal{M}^{\mathrm{free}}) - \sum_j \lambda_j q_j^{\mathrm{alloc}}(\mathcal{M}^{\mathrm{free}}) \geq \frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\mathrm{alloc}}(C). \tag{5.5}$$

Our starting point is again Equation (3.8), which gives, for any mechanism C and any lottery probabilities $\ell_j \in (0, 1]$

$$U(C) - D \leq \mathbb{E}\left[\max_j \ell_j V_j\right] \sum_j \frac{q_j^{\mathrm{alloc}}(C)}{\ell_j}.$$

Notice that the expectation is the utility of offering free lotteries with probabilities ℓ_j. The question thus becomes how to set the lottery probabilities ℓ_j.

We choose the lottery probabilities so that whichever lottery the buyer selects, its allocation cost is at most a common parameter $h > 0$. Subject to this restriction, we make each lottery as generous as possible: for item j, we choose the largest probability $\ell_j(h) \leq 1$ satisfying $\lambda_j \ell_j(h) \leq h$. Thus, define

$$\ell_j(h) := \begin{cases} \min\{1, h/\lambda_j\}, & \lambda_j > 0, \\ 1, & \lambda_j = 0. \end{cases}$$

Let mechanism $\mathcal{M}_h^{\text{free}}$ offer these lotteries: the buyer chooses an item maximizing $\ell_j(h)V_j$ and receives it for free with probability $\ell_j(h)$. We now verify that mechanism $\mathcal{M}_h^{\text{free}}$ has allocation cost at most h (as promised), using the fact that the buyer chooses at most one item's lottery:

$$\sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}_h^{\text{free}}) = \sum_j \lambda_j \ell_j(h) \Pr(\text{buyer chooses item } j) \leq h \sum_j \Pr(\text{buyer chooses item } j) \leq h. \tag{5.6}$$

Setting parameter h introduces a tradeoff to consider: increasing it weakly increases the utility of mechanism $\mathcal{M}_h^{\text{free}}$ (since that weakly increases each lottery's allocation probability), but also raises the upper bound we can prove on the mechanism's allocation cost. Naturally, we will choose parameter h so that mechanism $\mathcal{M}_h^{\text{free}}$ provides enough utility relative to this allocation cost.

To guide our choice of parameter h, we first rewrite Equation (3.8) using the lottery probabilities $\ell_j(h)$ of mechanism $\mathcal{M}_h^{\text{free}}$:

$$\begin{aligned}
U(C) - D &\leq \mathbb{E}_{\mathbf{V}}[\max_j \ell_j(h) V_j] \sum_j \frac{q_j^{\text{alloc}}(C)}{\ell_j(h)} && \text{Equation (3.8)} \\
&= U(\mathcal{M}_h^{\text{free}}) \sum_j \frac{q_j^{\text{alloc}}(C)}{\ell_j(h)} && \text{definition of mechanism } \mathcal{M}_h^{\text{free}} \\
&\leq U(\mathcal{M}_h^{\text{free}}) \left(\sum_j q_j^{\text{alloc}}(C) + \frac{\sum_j \lambda_j q_j^{\text{alloc}}(C)}{h} \right) && 1/\ell_j(h) \leq 1 + \lambda_j/h \\
&\leq \frac{U(\mathcal{M}_h^{\text{free}})}{h} \left(h + \sum_j \lambda_j q_j^{\text{alloc}}(C) \right). && \sum_j q_j^{\text{alloc}}(C) \leq 1 \quad (5.7)
\end{aligned}$$

We would like to select h so that $U(\mathcal{M}_h^{\text{free}}) = 2h$. Since its allocation cost is at most h, this leaves utility minus allocation cost at least h. The same choice makes the coefficient $U(\mathcal{M}_h^{\text{free}})/h$ in (5.7) equal to two, giving $U(C) - D \leq 2(h + \sum_j \lambda_j q_j^{\text{alloc}}(C))$. If such a parameter $h > 0$ exists,

the desired bound follows:

$$
\begin{aligned}
U(C) - D &\leq \frac{U(\mathcal{M}_h^{\text{free}})}{h}\left(h + \sum_j \lambda_j q_j^{\text{alloc}}(C)\right) && \text{inequality (5.7)} \\
&= 2\left(h + \sum_j \lambda_j q_j^{\text{alloc}}(C)\right) && \text{choice of } h \\
&= 2\left(U(\mathcal{M}_h^{\text{free}}) - h + \sum_j \lambda_j q_j^{\text{alloc}}(C)\right) && h = U(\mathcal{M}_h^{\text{free}}) - h \\
&\leq 2\left(U(\mathcal{M}_h^{\text{free}}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}_h^{\text{free}}) + \sum_j \lambda_j q_j^{\text{alloc}}(C)\right). && \text{inequality (5.6)} \quad (5.8)
\end{aligned}
$$

If no such parameter h exists, we claim that $U(\mathcal{M}_h^{\text{free}})/h < 2$ for every parameter $h > 0$. Notice that the utility of mechanism $\mathcal{M}_h^{\text{free}}$ is bounded above by $\mathbb{E}[\max_j V_j]$, the ratio $U(\mathcal{M}_h^{\text{free}})/h$ is continuous in h and tends to zero if we let h tend to infinity. Thus, if the ratio exceeded two for any parameter h, continuity would imply existence of a parameter at which it equals two, contradicting the assumption that no parameter satisfies $U(\mathcal{M}_h^{\text{free}}) = 2h$. Having established $U(\mathcal{M}_h^{\text{free}})/h < 2$, we rewrite (5.7) as

$$U(C) - D \leq 2h + 2\sum_j \lambda_j q_j^{\text{alloc}}(C).$$

Taking the limit as parameter h tends to zero gives

$$\frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\text{alloc}}(C) \leq 0.$$

This means that the free lottery mechanism $\mathcal{M}_0^{\text{free}}$ that never allocates satisfies:

$$U(\mathcal{M}_0^{\text{free}}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}_0^{\text{free}}) = 0 \geq \frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\text{alloc}}(C).$$

Thus, selecting mechanism $\mathcal{M}^{\text{free}} := \mathcal{M}_h^{\text{free}}$ when a parameter satisfying $U(\mathcal{M}_h^{\text{free}}) = 2h$ exists, and mechanism $\mathcal{M}^{\text{free}} := \mathcal{M}_0^{\text{free}}$ otherwise, yields the desired inequality (5.5) in both cases. Note again that this selection is independent of the comparison mechanism C. In terms of computation, Appendix C.3 shows how to find a free-lottery mechanism satisfying (5.5) up to arbitrarily small additive error.

The combined single-buyer mechanism. Let mechanism $\mathcal{M}$ run pricing mechanism $\mathcal{M}^{UD}$ with probability $e/(3e-2)$ and free-lottery mechanism $\mathcal{M}^{\text{free}}$ with probability $2(e-1)/(3e-2)$, independently of the buyer's report. By linearity of utility and allocation cost under this mixture,

the bounds (5.3) and (5.5) give the desired bound:

$$\begin{aligned} U(\mathcal{M}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}) &\geq \frac{e}{3e-2}\left[\left(1-\frac{1}{e}\right)D - \sum_j \lambda_j q_j^{\text{alloc}}(C)\right] \\ &\quad + \frac{2(e-1)}{3e-2}\left[\frac{U(C)-D}{2} - \sum_j \lambda_j q_j^{\text{alloc}}(C)\right] \\ &= \frac{U(C)}{2+e/(e-1)} - \sum_j \lambda_j q_j^{\text{alloc}}(C). \end{aligned} \tag{5.9}$$

Mechanism $\mathcal{M}$ has nonnegative payments and is ex-post individually rational, since both sub-mechanisms have these properties. The buyer chooses their preferred offer in each sub-mechanism, so the mixture is BIC as well.

□

To recap, our single-buyer mechanism $\mathcal{M}_{\text{het}}^{1,m}$ is defined as follows: Fix $\lambda \geq 0$. Let $\mathbf{q}^\star$ maximize the cap program (5.4) in the proof of Lemma 5. Then

1. With probability $e/(3e-2)$, run $\mathcal{M}^{UD}$: Draw each item j's price independently according to $\text{OPT}^{1,1}(V_j; q_j^\star)$. Let the buyer choose the item with the largest positive utility (if any) and charge its price.

2. With probability $2(e-1)/(3e-2)$, run $\mathcal{M}^{\text{free}}$: For a given $h > 0$, let $\ell_j(h) = \min\{1, h/\lambda_j\}$ be the probability of allocating item j if the buyer chooses it, with $\ell_j(h) = 1$ when $\lambda_j = 0$. The mechanism lets the buyer choose an item j maximizing their expected utility $\ell_j(h)V_j$ and allocates it to them with probability $\ell_j(h)$ free of charge. Choose $h > 0$ satisfying $\mathbb{E}[\max_j \ell_j(h)V_j] = 2h$, if one exists; otherwise allocate nothing.

We can then conclude that the better of these two mechanisms gives a constant-factor approximation (measured in utility minus allocation cost).

Corollary 2. *Under the assumptions of Lemma 5, for every comparison mechanism C,*

$$\max_{\mathcal{M}\in\{\mathcal{M}^{UD},\mathcal{M}^{\text{free}}\}} \left\{ U(\mathcal{M}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}) \right\} \geq \frac{U(C)}{2+e/(e-1)} - \sum_j \lambda_j q_j^{\text{alloc}}(C).$$

5.1.1 Choosing λ's.

We now construct an ex-ante-feasible, approximately-optimal multi-buyer mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ by choosing the demand weights iteratively. We initialize each item's demand weight λ_j at zero. In each round of the T rounds, we compute single-buyer mechanisms using the current weights, then increase (or decrease) an item's weight if its expected demand exceeds (or falls below) supply.

Our updates follow the projected online gradient descent method of Zinkevich [41]. Its no-regret guarantee controls the supply term in the Lagrangian on average over the rounds, allowing us to turn the Lagrangian guarantee into a utility guarantee. In particular, the average utility remains large enough to cover the utility lost when we reduce participation to restore ex-ante

feasibility. We implement this average by independently choosing one of the T rounds uniformly at random for each buyer. We use the prices or lotteries computed for that buyer in that round. To restore ex-ante feasibility, we sometimes make no offer, with the same probability for every buyer.

Formally, initialize demand weights to $\lambda^1 = 0$ and set a step size $\eta = S/(n\sqrt{T})$ (where $S = \sum_{i,j} \mathbb{E}[V_{ij}]$ is a crude welfare upper bound on the expected utility of any multidimensional mechanism). We will run the following procedure for T rounds. In round t, apply Lemma 5 with demand weights λ^t to compute a mechanism $\mathcal{M}_i^t := \mathcal{M}_{\text{het},i}^{1,m}(\lambda^t)$ for each buyer i. Let $r_j^t := \sum_i q_j^{\text{alloc}}(\mathcal{M}_i^t)$ be the expected number of requests for item j in this round, and update its demand weight according to

$$\lambda_j^{t+1} = \min \left\{ \sum_{i,k} \mathbb{E}[V_{ik}], \max\left\{0, \lambda_j^t + \eta(r_j^t - 1)\right\} \right\}. \tag{5.10}$$

After the T rounds, define the average expected number of requests $\overline{r}_j := T^{-1} \sum_{t=1}^{T} r_j^t$. Mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ independently draws a uniform round t_i for each buyer i and runs for this buyer $\mathcal{M}_i^{t_i}$ with probability $1/\max\{1, \max_j \overline{r}_j\}$ or otherwise makes no offer to that buyer. The following lemma describes the properties and utility guarantee of mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$. Its ε loss accounts for running only finitely many demand-weight updates and for approximately computing the single-buyer mechanisms whose existence is guaranteed by Lemma 5. The proof of the following lemma and exact implementation details of $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ are deferred to Appendix C.4.

Lemma 6. *Using the demand weights $\lambda^1, \ldots, \lambda^T$ generated by the procedure above, for every $\varepsilon > 0$, after $T = O(n^4 m^4/\varepsilon^2)$ rounds, $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ is ex-ante feasible and the mechanism's total expected utility satisfies*

$$U(\overline{\mathcal{M}}_{\text{het}}^{n,m}) \geq \left(\frac{1}{2 + e/(e-1)} - \varepsilon \right) \text{OPT}^{n,m}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\text{OPT}^{1,1}$ and exact evaluation of the required expectations and item allocation probabilities.

5.2 Resolving Contention without Changing Buyers' Choices

Let $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ be the mechanism constructed by the demand-weight procedure in Section 5.1.1. By Lemma 6, its requests satisfy each item feasibility constraint in expectation, but several buyers may still request the same item. In Section 4.2, item symmetry ensured that every item a buyer might request had the same probability of being served. Here, competition can differ across items. To preserve the buyer's choices, we use a contention resolution scheme that serves each request with at least a constant probability, then reduce its probability so that every request is served with exactly the same probability. Since choosing a lower-surplus item cannot improve a buyer's chance of service, the buyer still maximizes expected utility by choosing their surplus-maximizing item.

We treat each allocation under mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ as a request for the contention-resolution stage. Thus, $q_{ij}^{\text{alloc}} := q_{ij}^{\text{alloc}}(\overline{\mathcal{M}}_{\text{het}}^{n,m})$ is the probability that buyer i receives item j under mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$, but in the contention-resolution stage, q_{ij}^{alloc} is instead the probability buyer i requests item

j (with the scheme determining whether the request is served or not). For each item j, requests are independent across buyers and $\sum_i q_{ij}^{\text{alloc}} \le 1$ (at most one item is allocated in expectation in mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$). We use the fair contention rule of Feige and Vondrak [22] to select which agent might be served an item j. If there are no requests for the item, no buyer is selected, if there is one request that buyer is selected, and if there are more requests, each requesting buyer i is selected with probability

$$\frac{1}{\sum_k q_{kj}^{\text{alloc}}} \left(\frac{\sum_{k \in A_j \setminus \{i\}} q_{kj}^{\text{alloc}}}{|A_j| - 1} + \frac{\sum_{k \notin A_j} q_{kj}^{\text{alloc}}}{|A_j|} \right),$$

where A_j is the set of buyers requesting item j (if $\sum_k q_{kj}^{\text{alloc}} = 0$, a requester is selected uniformly instead). This rule guarantees at least a constant probability of a buyer being selected, conditional on them requesting the item

$$c_j = \Pr\big(i \text{ selected for } j \mid i \text{ requests } j\big) = \begin{cases} 1, & \text{if } \sum_k q_{kj}^{\text{alloc}} = 0, \\ \dfrac{1 - \prod_k (1 - q_{kj}^{\text{alloc}})}{\sum_k q_{kj}^{\text{alloc}}}, & \text{otherwise,} \end{cases} \tag{5.11}$$

where $c_j \ge 1 - \frac{1}{e}$.

If the rule selects buyer i for item j, serve the request with probability $(1 - 1/e)/c_j$, independently of the earlier randomness and charge the corresponding single-buyer payment only if the item is served. This guarantees that every possible request (including one induced by a deviation) is served with probability exactly $1 - 1/e$, independently of the buyer's report. This implies that it is in the best interest of each buyer to request their surplus-maximizing item. See Appendix C.5 for the full contention resolution argument and computational guarantees.

By Observation 1, the resulting mechanism $\mathcal{M}_{\text{het}}^{n,m}$ is BIC, ex-post individually rational, and matching-feasible, has nonnegative payments, and satisfies

$$U(\mathcal{M}_{\text{het}}^{n,m}) = \left(1 - \frac{1}{e}\right) U(\overline{\mathcal{M}}_{\text{het}}^{n,m}). \tag{5.12}$$

Combining this utility guarantee with Lemma 6 gives the following approximation.

Theorem 3. *In the independent-values setting, for every $\varepsilon > 0$, the demand-weight procedure followed by the contention rule above gives a BIC, ex-post individually rational, matching-feasible mechanism $\mathcal{M}_{\text{het}}^{n,m}$ with nonnegative payments. Its expected utility satisfies*

$$U(\mathcal{M}_{\text{het}}^{n,m}) \ge \frac{\text{OPT}^{n,m}}{2e/(e-1) + (e/(e-1))^2 + \varepsilon}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\text{OPT}^{1,1}$ and exact evaluation of the required expectations and item allocation probabilities.

Proof. It suffices to consider $0 < \varepsilon \le 1$; for larger ε, use the mechanism obtained with $\varepsilon = 1$. We choose the accuracy in Lemma 6 so that, after the utility loss due to contention, the additive error in the approximation factor is ε.

Set the error parameter θ of Lemma 6 to

$$\theta := \frac{1}{2 + e/(e-1)} - \frac{1}{2 + e/(e-1) + (1-1/e)\varepsilon} > 0.$$

By Lemma 6, running the demand-weight procedure with accuracy θ gives mechanism $\overline{\mathcal{M}}^{n,m}_{\text{het}}$ with utility

$$U(\overline{\mathcal{M}}^{n,m}_{\text{het}}) \geq \left(\frac{1}{2 + e/(e-1)} - \theta\right) \text{OPT}^{n,m} = \frac{\text{OPT}^{n,m}}{2 + e/(e-1) + (1-1/e)\varepsilon}.$$

By the contention resolution argument (5.12),

$$U(\mathcal{M}^{n,m}_{\text{het}}) \geq \frac{1-1/e}{2 + e/(e-1) + (1-1/e)\varepsilon} \text{OPT}^{n,m} = \frac{\text{OPT}^{n,m}}{2e/(e-1) + (e/(e-1))^2 + \varepsilon}.$$

Because we selected $\theta = \Theta(\varepsilon)$ for $0 < \varepsilon \leq 1$ and Lemma 6 requires $T = O(n^4 m^4/\varepsilon^2)$ rounds for its additive error guarantee, the construction uses polynomially many rounds and single-buyer optimization calls in $n, m, 1/\varepsilon$. The contention rule and the computation of its selection probabilities also take polynomial time under the stated access assumptions. □

6 Matroid-Constrained Additive Buyers

We extend the mechanism of Section 5 to matroid-constrained additive buyers, retaining the same approximation guarantee (Theorem 4). This setting includes both unit-demand buyers ($\mathcal{I}_i = \{T : |T| \leq 1\}$) and unconstrained additive buyers ($\mathcal{I}_i = 2^{[m]}$). The design and analysis follow the same steps: posted prices and free lotteries provide a single-buyer guarantee, demand-weight updates balance expected item requests while maintaining high utility, and itemwise contention resolution makes the allocation feasible ex-post. The only substantial differences are in how we design the single-buyer mechanisms and analyze them. We describe these adaptations below; the full construction and analysis appear in Appendix D.

In the posted-price part of the single-buyer mechanism, we offer item prices and let the buyer choose an independent set maximizing total surplus. As before, we first optimize the single-item posted prices, now subject to their allocation probabilities satisfying the matroid constraints in expectation. We then compare the expected utility from offering these prices with the sum of the single-item utilities. Using the matroid correlation-gap bound of Yan [39, Lemma 4.1], we extend Lemma 3 and recover the same $1 - 1/e$ fraction of the sum of the single-item utilities.

For the free-lottery part of the single-buyer mechanism, we still search for a lottery mechanism with low expected allocation cost and high expected utility by varying a parameter h that bounds the cost under demand weights λ. This mechanism has a simple implementation: the buyer selects their preferred option among independent sets and randomizations between two independent sets with expected allocation cost at most h. As in Section 5.1, we choose parameter h so that the mechanism's expected utility equals $2h$, if such a parameter exists, or otherwise allocate nothing (for this we assume access to a *free-lottery utility oracle* that evaluates the mechanism's expected utility for a given h). The resulting mechanism satisfies the same free-lottery bound as in the unit-demand setting.

Since we have constructed pricing and free-lottery mechanisms with the same guarantees as in the unit-demand case, combining them with the same mixing probabilities as in Section 5.1

gives the same guarantee as Lemma 5, with a matroid constraint replacing the unit-demand constraint. The demand-weight and contention-resolution arguments require only minor adaptations for requested sets, without changing their guarantees. Specifically, we define mechanism $\overline{\mathcal{M}}^{n,m}_{\text{mat}}$ by using these single-buyer mechanisms in the demand-weight procedure of Section 5.1.1, with the same update rule, averaging, and participation reduction. Mechanism $\overline{\mathcal{M}}^{n,m}_{\text{mat}}$ allocates an independent set to each buyer, is ex-ante feasible for item supply, and has the same approximate utility guarantee as Lemma 6. To turn this into the ex-post feasible mechanism $\mathcal{M}^{n,m}_{\text{mat}}$, we apply the fair contention rule of Feige and Vondrak [22] item by item to the buyers' requested independent sets. As in Section 5.2, we adjust each request's service probability to exactly $1-1/e$ and charge posted prices only for delivered items. Conditional on a buyer's entire requested set, each requested item is delivered with probability $1-1/e$, including under deviations. Since values and payments add across delivered items, contention resolution multiplies the expected utility of every option by $1-1/e$, preserving the buyer's utility-maximizing choices. Each item is allocated at most once, and downward closure ensures that every buyer's delivered set remains feasible. This gives the same approximation guarantee as Theorem 3, now for matroid-constrained buyers.

Theorem 4. *In the matroid-constrained setting, for every $\varepsilon > 0$, the demand-weight procedure followed by the contention rule above gives a BIC, ex-post individually rational, feasible mechanism $\mathcal{M}^{n,m}_{\text{mat}}$ with nonnegative payments. Its expected utility satisfies*

$$U(\mathcal{M}^{n,m}_{\text{mat}}) \geq \frac{\text{OPT}^{n,m}}{2e/(e-1) + (e/(e-1))^2 + \varepsilon}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\text{OPT}^{1,1}$, matroid membership tests, a free-lottery utility oracle, and exact evaluation of the required expectations and item allocation probabilities.

Throughout the matroid extension, ties among utility-maximizing options are resolved using a fixed deterministic ordering, with non-participation preferred whenever it is optimal. The assumed allocation probabilities use this same rule.

AI Disclosure

We used a variety of AI models in developing this paper including frontier models from both OpenAI and Anthropic. This work built on five years of our preliminary research, including results, technical approaches, and intuition about what had and had not worked. We supplied this background and guided an iterative investigation: we formulated the original research questions, posed further questions as the work developed, suggested techniques, and directed which approaches to pursue. Within this process, GPT-6 Astra generated a draft paper and proofs of most of the results. The first proof of the result for independent values was obtained using GPT-5.6.

We wrote the paper and proofs presented here by completely rewriting the AI-generated drafts. This involved reinterpreting and simplifying the arguments, developing new notation, breaking the proofs into smaller conceptual steps, and drawing out the underlying intuition. GPT-6 Astra contributed substantially to discovering results and generating proofs within that process. We take responsibility for the paper's content.

References

[1] S. Alaei. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 512–521, Oct 2011. doi: 10.1109/FOCS.2011.90.

[2] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, page 17, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450314152. doi: 10.1145/2229012.2229017. URL https://doi.org/10.1145/2229012.2229017.

[3] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 21–30. IEEE, 2014.

[4] Xiaohui Bei and Zhiyi Huang. Bayesian Incentive Compatibility via Fractional Assignments. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 720–733. Society for Industrial and Applied Mathematics, 2011. doi: 10.1137/1.9781611973082.57. URL https://doi.org/10.1137/1.9781611973082.57.

[5] Batya Berzack, Rotem Oshman, and Inbal Talgam-Cohen. Dynamic Rental Games with Stagewise Individual Rationality. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC '25, page 785, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400719431. doi: 10.1145/3736252.3742624. URL https://doi.org/10.1145/3736252.3742624.

[6] Batya Berzack, Rotem Oshman, and Inbal Talgam-Cohen. Optimal Auction Design for Constrained Buyers, 2026. URL https://arxiv.org/abs/2606.30776.

[7] Kshipra Bhawalkar, Marios Mertzanidis, Divyarthi Mohan, and Alexandros Psomas. Mechanism Design via the Interim Relaxation, 2024. URL https://arxiv.org/abs/2407.12699.

[8] Mark Braverman, Jing Chen, and Sampath Kannan. Optimal Provision-After-Wait in Healthcare. *Mathematics of Operations Research*, 41(1):352–376, 2016.

[9] Yang Cai and Mingfei Zhao. Simple Mechanisms for Subadditive Buyers via Duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 170–183, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4528-6. doi: 10.1145/3055399.3055465. URL http://doi.acm.org/10.1145/3055399.3055465.

[10] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 926–939, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4132-5. doi: 10.1145/2897518.2897645. URL http://doi.acm.org/10.1145/2897518.2897645.

[11] Shuchi Chawla and J. Benjamin Miller. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computa-*

tion, EC '16, pages 579–596, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0. doi: 10.1145/2940716.2940756. URL http://doi.acm.org/10.1145/2940716.2940756.

[12] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010. URL http://dl.acm.org/citation.cfm?id=1806733.

[13] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The Power of Randomness in Bayesian Optimal Mechanism Design. *Games and Economic Behavior*, 91:297 – 317, 2015. ISSN 0899-8256. doi: http://dx.doi.org/10.1016/j.geb.2012.08.010. URL http://www.sciencedirect.com/science/article/pii/S0899825612001303.

[14] Shuchi Chawla, Rojin Rezvan, Yifeng Teng, and Christos Tzamos. Buy-Many Mechanisms for Many Unit-Demand Buyers. In *International Conference on Web and Internet Economics*, pages 21–38. Springer, 2023.

[15] Alon Eden, Michal Feldman, Amos Fiat, Kira Goldner, and Anna R. Karlin. Combinatorial Auctions with Interdependent Valuations: SOS to the Rescue. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, Phoenix, AZ, USA, 2019. ACM. ISBN 978-1-4503-6792-9. doi: 10.1145/3328526.3329759. URL https://arxiv.org/abs/1903.08384.

[16] Alon Eden, Kira Goldner, and Shuran Zheng. Private Interdependent Valuations. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2920–2939. SIAM, 2022.

[17] Alon Eden, Michal Feldman, Kira Goldner, Simon Mauras, and Divyarthi Mohan. Constant Approximation for Private Interdependent Valuations. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 148–163. IEEE, 2023.

[18] Alon Eden, Kira Goldner, Eldar Kerner, and Thodoris Tsilivis. From Welfare to Utility: Generalized Objectives in Budget-Feasible Procurement. In *Forty-third International Conference on Machine Learning*, 2026.

[19] Meryem Essaidi, Kira Goldner, and S. Matthew Weinberg. To Regulate or Not to Regulate: Using Revenue Maximization Tools to Maximize Consumer Utility. In Guido Schäfer and Carmine Ventre, editors, *Proceedings of the 17th International Symposium on Algorithmic Game Theory (SAGT)*, pages 315–332, Cham, 2024. ISBN 978-3-031-71033-9. URL https://arxiv.org/abs/2002.06326.

[20] Tomer Ezra. Money Burning Mechanism Design: From Welfare to Surplus, 2026. URL https://arxiv.org/abs/2608.01693v2. Version 2, September 24, 2026.

[21] Tomer Ezra, Daniel Schoepflin, and Ariel Shaulker. Multi-Parameter Mechanisms for Consumer Surplus Maximization. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 483–494, 2025.

[22] Uriel Feige and Jan Vondrak. Approximation Algorithms for Allocation Problems: Improving the Factor of 1 - 1/e . In *2006 47th Annual IEEE Conference on Foundations of Computer Science*, pages 667–676, Los Alamitos, CA, USA, October 2006. IEEE Computer Society. doi: 10.1109/FOCS.2006.14.

[23] Dimitris Fotakis, Dimitris Tsipras, Christos Tzamos, and Emmanouil Zampetakis. Efficient Money Burning in General Domains. *Theory of Computing Systems*, 59:619–640, 2016.

[24] Aadityan Ganesh and Jason Hartline. Combinatorial Pen Testing (Or Consumer Surplus of Deferred-Acceptance Auctions). In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, pages 52:1–52:22. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025. doi: 10.4230/LIPIcs.ITCS.2025.52.

[25] Michel X. Goemans. Lecture notes on matroid intersection. Massachusetts Institute of Technology, 18.453: Combinatorial Optimization, Lecture 6, May 2017. URL `https://math.mit.edu/~goemans/18453S17/matroid-intersect-notes.pdf`.

[26] Kira Goldner and Taylor Lundy. Multidimensional Bayesian Utility Maximization: Tight Approximations to Welfare. In *Proceedings of the Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[27] Kira Goldner, Divyarthi Mohan, and Anh Phung. On the Menu Complexity of Consumer Utility Maximization. Unpublished manuscript, 2026.

[28] Kira Goldner, Divyarthi Mohan, and Thodoris Tsilivis. Knowing Who, Not How Much: Learning-Augmented Mechanisms for Consumer Utility Maximization. In *Forty-third International Conference on Machine Learning*, 2026. URL `https://arxiv.org/abs/2607.00175`.

[29] Sergiu Hart and Noam Nisan. The Menu-Size Complexity of Auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 565–566, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1962-1. doi: 10.1145/2482540.2482544. URL `http://doi.acm.org/10.1145/2482540.2482544`.

[30] Jason D. Hartline. Mechanism Design and Approximation, October 2013. URL `https://jasonhartline.com/MDnA/`. Book draft.

[31] Jason D. Hartline and Tim Roughgarden. Optimal Mechanism Design and Money Burning. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 75–84, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-047-0. doi: http://doi.acm.org/10.1145/1374376.1374390.

[32] Colleen Heflin, William Clay Fannin, and Leonard Lopoo. Local Control, Discretion, and Administrative Burden: SNAP Interview Waivers and Caseloads During the COVID-19 Pandemic. *American Review of Public Administration*, 2023.

[33] D. Kashaev and R. Santiago. A Simple Optimal Contention Resolution Scheme for Uniform Matroids. *Theoretical Computer Science*, 940:81–96, 2023. URL `https://arxiv.org/abs/2105.11992`.

[34] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 123–136. ACM, 2012. doi: 10.1145/2213977.2213991. URL `https://doi.org/10.1145/2213977.2213991`.

[35] Taylor Lundy, Alexander Wei, Hu Fu, Scott Duke Kominers, and Kevin Leyton-Brown. Allocation for social good: auditing mechanisms for utility maximization. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 785–803, 2019.

[36] Taylor Lundy, Narun Raman, Hu Fu, and Kevin Leyton-Brown. Pay to (not) play: Monetizing impatience in mobile games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9856–9864, 2024.

[37] Rohit Patel and Can Urgun. Costly verification and money burning, March 2025. URL `https://curgun.scholar.princeton.edu/document/66`. Working paper.

[38] Ester Samuel-Cahn. Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *The Annals of Probability*, pages 1213–1216, 1984.

[39] Qiqi Yan. Mechanism Design via Correlation Gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 710–719. Society for Industrial and Applied Mathematics, 2011. doi: 10.1137/1.9781611973082.56. URL `https://arxiv.org/abs/1008.1843`.

[40] Andrew Chi-Chih Yao. An n-to-1 Bidder Reduction for Multi-Item Auctions and Its Applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109. SIAM, 2015. doi: 10.1137/1.9781611973730.8. URL `http://arxiv.org/abs/1406.3278`.

[41] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning (ICML)*, pages 928–936, 2003.

A Bounding Variation by Deletion Loss

We give the calculation establishing (3.3), using the utility function and deletion loss defined in Section 3.

To bound the variation term, we replace the coordinates of $\mathbf{W}$ with those of $\mathbf{V}$, one at a time. Let $H^{(j)} := (V_1, \dots, V_j, W_{j+1}, \dots, W_m)$ be the vector after the first j replacements, so $H^{(0)} = \mathbf{W}$ and $H^{(m)} = \mathbf{V}$. The changes in utility telescope, and discarding negative changes gives an upper bound:

$$\begin{aligned}
\frac{1}{2}\mathbb{E}_{\mathbf{V},\mathbf{W}}\big[\big|u(\mathbf{V}) - u(\mathbf{W})\big|\big] &= \mathbb{E}_{\mathbf{V},\mathbf{W}}\big[(u(\mathbf{V}) - u(\mathbf{W}))^+\big] && \text{symmetry} \\
&= \mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\Big(u(H^{(m)}) - u(H^{(0)})\Big)^+\right] && \text{def. of } H^{(j)} \\
&= \mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\left(\sum_{j=1}^{m}\big(u(H^{(j)}) - u(H^{(j-1)})\big)\right)^+\right] && \text{telescoping} \\
&\le \mathbb{E}_{\mathbf{V},\mathbf{W}}\left[\sum_{j=1}^{m}\big(u(H^{(j)}) - u(H^{(j-1)})\big)^+\right] && (\textstyle\sum_j a_j)^+ \le \sum_j a_j^+ \\
&\le \sum_{j=1}^{m}\mathbb{E}_{\mathbf{V}}\big[u(\mathbf{V}) - u(\mathbf{V}_{-j}, 0)\big] && \text{monotonicity of } u \\
&= D, && \text{def. of } D \qquad \text{(A.1)}
\end{aligned}$$

where the final inequality follows because u is coordinatewise nondecreasing and $H^{(j)}$ has the same distribution as $\mathbf{V}$.

B Stronger Favorite-Item Guarantees

We prove a factor-2 single-buyer guarantee for the *favorite-item* mechanism. This gives a single-buyer mechanism that is simpler (a single price for all items) and achieves a better approximation guarantee than $\mathcal{M}^{1,m}_{\text{iid}}(Q)$ from Section 4.1. For i.i.d. buyers, combining this mechanism with uniform contention resolution gives a better guarantee than the one in Theorem 2. We also extend the result to the item-i.i.d. setting, where values are identically distributed across items for each buyer, but not necessarily across buyers.

Throughout Appendix B, Q denotes a buyer's total allocation cap, $q = Q/m$ the common item cap, and $\mathbf{q} = (q, \dots, q)$ the vector of item caps. By item symmetry, the optimum under total cap Q equals $\text{OPT}^{1,m}(\mathbf{V}; \mathbf{q})$, using the notation of the main body.

B.1 A Factor-2 Approximation through the Favorite-Item Mechanism

The simplified mechanism offers one random price, common to all items, according to the *favorite-item mechanism* $\text{OPT}^{1,1}(M; Q)$, where $M := \max_j V_j$. The buyer selects a favorite item if its value strictly exceeds the price. By symmetry, each item's allocation probability is at most Q/m. To prove Lemma 7, we use the prices from $\text{OPT}^{1,1}(V; q)$ to construct a mechanism for the favorite

value M, recovering the single-item utility terms of Theorem 1 without the loss from combining independent item prices. The following lemma states the resulting approximation guarantee.

Lemma 7. *Let $M := \max_j V_j$. For i.i.d. item values and every $Q \in [0,1]$, with $q = Q/m$ and $\mathbf{q} = (q, \dots, q)$, the favorite-item mechanism $\mathrm{OPT}^{1,1}(M; Q)$ is truthful, ex-post individually rational, has nonnegative payments, allocates each item with ex-ante probability at most q, and satisfies*

$$\mathrm{OPT}^{1,1}(M; Q) \geq \frac{1}{2}\mathrm{OPT}^{1,m}(\mathbf{V}; \mathbf{q}).$$

Fix a single buyer with i.i.d. item values, a total cap $Q \in [0,1]$, and a comparison mechanism with utility U and ex-ante allocation probability at most Q. As before, it is without loss of generality to consider mechanisms with buyer item ex-ante probabilities at most Q/m. From Equation (3.8) and lemma 2 with $\ell_j = 1$ we get

$$U \leq D + Q\mathbb{E}[M] \leq m\mathrm{OPT}^{1,1}(V; Q/m) + Q\mathbb{E}[M]. \tag{B.1}$$

We first recover the deletion bound through a mechanism for the favorite value M.

Lemma 8. *Let $V, V_1, \dots, V_m$ be i.i.d. nonnegative random variables drawn from F, and let $M := \max_j V_j$. For every $Q \in [0,1]$,*

$$m\mathrm{OPT}^{1,1}(V; Q/m) \leq (1+Q)\mathrm{OPT}^{1,1}\left(M; \frac{Q}{1+Q}\right). \tag{B.2}$$

Proof of Lemma 8. Fix an optimal single-dimensional mechanism for V with ex-ante allocation probability at most Q/m, whose expected utility is $\mathrm{OPT}^{1,1}(V; Q/m)$. Since this is a single-dimensional mechanism, we can view it as a distribution over posted-price mechanisms (for details see [30]). For a single-dimensional value X, a posted price t is accepted with probability $\Pr(X > t)$ and provides expected utility $\mathbb{E}[(X-t)^+]$.

Let ω_k and t_k denote the weights and prices in the distribution over posted prices representing $\mathrm{OPT}^{1,1}(V; Q/m)$. For each component k, let $\pi_k^V := \Pr(V > t_k)$ denote the probability that its posted price is accepted. Feasibility requires that the weights define a valid distribution and that the ex-ante allocation probability under V is at most Q/m, that is,

$$\sum_k \omega_k \leq 1, \qquad \sum_k \omega_k \pi_k^V \leq Q/m.$$

We will construct a valid distribution over posted-price mechanisms for the maximum value M that obtains comparable utility to that of the mechanism for V, and satisfies an ex-ante probability constraint of $Q/(1+Q)$. We use the same prices t_k and let $\pi_k^M := \Pr(M > t_k)$ denote their acceptance probabilities under M. For each component k with $\pi_k^V > 0$, we select its probability weight

$$\widetilde{\omega}_k := \frac{m\omega_k \pi_k^V}{(1+Q)\pi_k^M}.$$

We first verify that the weights $\widetilde{\omega}_k$ define a valid distribution. Independence gives

$$\pi_k^M = 1 - (1-\pi_k^V)^m \geq 1 - e^{-m\pi_k^V} \geq 1 - \frac{1}{1+m\pi_k^V} = \frac{m\pi_k^V}{1+m\pi_k^V}, \tag{B.3}$$

where both inequalities follow from $e^x \geq 1+x$, applied with $x = -\pi_k^V$ and $x = m\pi_k^V$, respectively. We can therefore upper-bound the sum of the new probability weights:

$$\sum_k \widetilde{\omega}_k = \sum_k \frac{m\omega_k \pi_k^V}{(1+Q)\pi_k^M} \leq \frac{\sum_k \omega_k (1+m\pi_k^V)}{1+Q} \leq \frac{1+Q}{1+Q} = 1,$$

where the first inequality is by (B.3), while the second follows from the feasibility of the original distribution, namely $\sum_k \omega_k \leq 1$ and $\sum_k \omega_k \pi_k^V \leq Q/m$. Thus, the weights define a valid distribution (with any remaining probability assigned to making no offer).

The new mechanism's ex-ante allocation probability is

$$\sum_k \widetilde{\omega}_k \pi_k^M = \sum_k \frac{m\omega_k \pi_k^V}{(1+Q)\pi_k^M} \pi_k^M = \frac{m}{1+Q} \sum_k \omega_k \pi_k^V \leq \frac{Q}{1+Q},$$

where the inequality follows from the ex-ante allocation probability of the original mechanism, $\sum_k \omega_k \pi_k^V \leq Q/m$.

To compare utilities after reweighting, we show that an accepted offer gives at least as much expected surplus for M as for V. Fix a price t with $\Pr(V > t) > 0$. For $s \geq t$ with $\Pr(V > s) > 0$,

$$\frac{\Pr(M > s)}{\Pr(V > s)} = \frac{1 - F(s)^m}{1 - F(s)} = \sum_{\ell=0}^{m-1} F(s)^\ell \geq \sum_{\ell=0}^{m-1} F(t)^\ell = \frac{\Pr(M > t)}{\Pr(V > t)}.$$

Rearranging the tail inequality and integrating over $s \geq t$ yields

$$\frac{\mathbb{E}[(M-t)^+]}{\Pr(M > t)} \geq \frac{\mathbb{E}[(V-t)^+]}{\Pr(V > t)}. \tag{B.4}$$

We lower bound the expected utility U_M of the new mechanism:

$$\begin{aligned} U_M &:= \sum_k \widetilde{\omega}_k \mathbb{E}[(M-t_k)^+] = \frac{m}{1+Q} \sum_k \omega_k \frac{\pi_k^V}{\pi_k^M} \mathbb{E}[(M-t_k)^+] \\ &\geq \frac{m}{1+Q} \sum_k \omega_k \mathbb{E}[(V-t_k)^+] = \frac{m}{1+Q} \mathrm{OPT}^{1,1}(V; Q/m), \end{aligned}$$

where the inequality is due to (B.4).

The constructed mechanism is therefore feasible for $\mathrm{OPT}^{1,1}(M; Q/(1+Q))$. By optimality and the preceding utility bound,

$$\mathrm{OPT}^{1,1}(M; Q/(1+Q)) \geq U_M \geq \frac{m}{1+Q} \mathrm{OPT}^{1,1}(V; Q/m),$$

which proves the lemma. □

We can now apply this to inequality (B.1) to prove Lemma 7:

Proof of Lemma 7. The case $Q = 0$ is immediate, so assume $Q > 0$. Consider first $0 < Q \leq 1/2$. We construct a single-buyer mechanism for M with ex-ante probability at most Q by mixing the two mechanisms corresponding to the terms of (B.1). With probability $(1+Q)/2$, run $\mathrm{OPT}^{1,1}(M; Q/(1+Q))$; with probability $Q/2$, allocate the favorite item for free; and with the

remaining probability $(1-2Q)/2$, offer nothing. The mixture's ex-ante allocation probability is at most

$$\frac{1+Q}{2}\frac{Q}{1+Q}+\frac{Q}{2}=Q.$$

By Lemma 8, its expected utility is at least

$$\frac{1+Q}{2}\mathrm{OPT}^{1,1}\left(M;\frac{Q}{1+Q}\right)+\frac{Q}{2}\mathbb{E}[M]\geq\frac{1}{2}\big(m\mathrm{OPT}^{1,1}(V;Q/m)+Q\mathbb{E}[M]\big)\geq\frac{U}{2},\tag{B.5}$$

where the final inequality is (B.1). The mixture preserves incentive compatibility, individual rationality, and nonnegative payments, and is therefore feasible for $\mathrm{OPT}^{1,1}(M;Q)$. Hence $U\leq 2\mathrm{OPT}^{1,1}(M;Q)$ in this case.

For $1/2<Q\leq 1$, nonnegative payments and unit-demand feasibility imply $U\leq\mathbb{E}[M]$, while giving the favorite item for free with probability Q shows that $\mathrm{OPT}^{1,1}(M;Q)\geq Q\mathbb{E}[M]\geq U/2$.

Since the comparison mechanism was arbitrary, taking the supremum over comparison mechanisms proves the lemma:

$$\mathrm{OPT}^{1,m}(\mathbf{V};\mathbf{q})\leq 2\mathrm{OPT}^{1,1}(M;Q).\tag{B.6}$$

□

B.2 Fully i.i.d. Buyers

Using Lemma 7, we can immediately improve on the approximation factor of Theorem 2 by a factor of $(2e-1)/(2(e-1))$.

Theorem 5. *For i.i.d. values, the mechanism that independently runs* $\mathrm{OPT}^{1,1}(M;\min\{1,m/n\})$ *for each buyer and resolves contention uniformly at random is BIC, ex-post individually rational, has nonnegative payments, respects matching-feasibility, and achieves expected utility at least*

$$\frac{1-1/e}{2}\mathrm{OPT}^{n,m}.$$

This follows by applying Lemma 7 instead of Lemma 4 in the analysis of Section 4.2.

B.3 Different Distributions across Buyers

We now extend the guarantee to the *item-i.i.d.* setting: buyer i has i.i.d. item values from F_i, but these distributions may differ across buyers, although they remain independent. Write $M_i:=\max_j V_{ij}$. Both single-buyer guarantees (Section 4.1, Lemma 7) apply separately to each buyer, with total cap Q_i and item caps $q_{ij}=Q_i/m$. The multi-buyer argument changes in two respects.

The first issue pertains to computing appropriate ex-ante total probability caps Q_i. If there are at least as many items as buyers ($m\geq n$), setting $Q_i=1$ for every buyer remains feasible and optimal, so we restrict to cases with fewer items ($m<n$). Without symmetry across buyers, the common cap $Q=m/n$ may not be optimal, and we therefore need buyer-specific caps Q_i.

Ideally, we would choose these caps by optimizing the ex-ante relaxation itself. Even though its objective is concave in the allocation caps, we cannot efficiently optimize it directly because we cannot efficiently evaluate feasible solutions and compute their gradients. Instead, we use Lemma 7 to upper-bound the relevant single-buyer ex-ante objectives by the simpler 1-dimensional objective $2\mathrm{OPT}^{1,1}(M_i;Q_i)$, and optimize

$$\mathcal{R} := \max_{\substack{0 \le Q_i \le 1 \\ \sum_i Q_i \le m}} \sum_{i=1}^{n} \mathrm{OPT}^{1,1}(M_i; Q_i). \tag{B.7}$$

Fix any feasible, BIC, interim individually rational multi-buyer mechanism C with nonnegative payments. For each buyer i, let C_i be the single-buyer mechanism obtained by averaging C over the other buyers' values, and let Q_i^C be its ex-ante allocation probability. Since $Q_i^C \le 1$ and $\sum_i Q_i^C \le m$, applying Lemma 7 to each buyer gives $U(C) \le 2\mathcal{R}$. Taking the supremum over all such mechanisms C yields

$$\mathrm{OPT}^{n,m} \le 2\mathcal{R}. \tag{B.8}$$

The advantage of $\mathcal{R}$ is that each $\mathrm{OPT}^{1,1}(M_i; Q_i)$ can be evaluated efficiently. Moreover, the objective is the sum of concave functions, each depending only on one buyer's cap. We can therefore use a discretized greedy water-filling procedure to find approximately optimal caps. In particular, we can obtain caps whose objective value is within a multiplicative factor $1+\varepsilon$ of $\mathcal{R}$ in time $\mathrm{poly}(n, m, 1/\varepsilon)$ (see Appendix C.2 and Lemma 10 for the details).

The second issue is contention resolution. Unlike the fully i.i.d. setting, some buyers may contribute much more utility than others (consider for example, a buyer whose item values are always much higher than everyone else's) so treating all requests uniformly may not retain the same utility guarantee. Fortunately, the contention construction of Section 5.2 applies here as well: it serves every possible request with probability $1 - 1/e$ while preserving buyers' choices.

The following theorem extends the favorite-item guarantee to item-i.i.d. values.

Theorem 6. *For item-i.i.d. values and every $\varepsilon > 0$, we can compute buyer allocation caps $Q_i \in [0,1]$ with $\sum_i Q_i \le m$ and a contention-resolution scheme such that independently running $\mathrm{OPT}^{1,1}(M_i; Q_i)$ for each buyer, where $M_i := \max_j V_{ij}$, and resolving the resulting requests using this scheme gives a mechanism that is BIC, ex-post individually rational, has nonnegative payments, respects matching-feasibility, and achieves expected utility at least*

$$\frac{\mathrm{OPT}^{n,m}}{2e/(e-1)+\varepsilon}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\mathrm{OPT}^{1,1}$ and exact evaluation of the required expectations and item allocation probabilities.

Proof of Theorem 6. It suffices to consider $0 < \varepsilon \le 1$; for larger ε, use the mechanism obtained with $\varepsilon = 1$. If $m \ge n$, set $Q_i = 1$ for every buyer—these are optimal ex-ante caps. Otherwise, choose $\delta := \varepsilon(1-1/e)/2$ and use the item-i.i.d. cap procedure in Lemma 10 to compute feasible caps satisfying $\sum_i \mathrm{OPT}^{1,1}(M_i; Q_i) \ge \mathcal{R}/(1+\delta)$. For the following bound, use the same definition of δ in both cases.

Independently run $\mathrm{OPT}^{1,1}(M_i; Q_i)$ for each buyer, breaking item ties uniformly. Item symmetry bounds buyer i's probability of requesting any fixed item by Q_i/m, so the requests are independent across buyers and have total expected demand at most $\sum_i Q_i/m \le 1$ for every item. The contention construction of Section 5.2 therefore applies directly: it preserves the stated incentive, payment, and feasibility properties and retains a $1-1/e$ fraction of the single-buyer utilities. Expected utility is thus at least

$$\left(1 - \frac{1}{e}\right) \sum_i \mathrm{OPT}^{1,1}(M_i; Q_i) \ge \frac{1-1/e}{1+\delta}\mathcal{R} \ge \frac{1-1/e}{2(1+\delta)}\mathrm{OPT}^{n,m} \ge \frac{\mathrm{OPT}^{n,m}}{2e/(e-1)+\varepsilon},$$

where the inequalities use the computed caps' objective guarantee, the ex-ante upper bound (B.8), and the choice of δ, respectively. The computational guarantees follow from the cap procedure, the single-dimensional optimization, and contention probability calculations in Appendix C. $\square$

C Approximate Optimization and Further Implementation Details

This appendix provides the optimization and implementation details omitted throughout the paper. In Appendix C.1, we discuss the single-dimensional single-buyer benchmark $\mathrm{OPT}^{1,1}$ in further detail. Then, in Appendix C.2 we provide a greedy algorithm that yields approximately optimal ex-ante caps for both the item-i.i.d. optimization problem (formulation (B.7)) and the unit-demand cap optimization problem (formulation (5.4)). Appendix C.3 explains how to compute the free-lottery parameter up to additive error. Next, in Appendix C.4, we prove Lemma 6, which guarantees that our non-iid unit-demand mechanism $\overline{\mathcal{M}}^{n,m}_{\mathrm{het}}$ is both ex-ante feasible and approximately optimal. Finally, in Appendix C.5, we describe and analyze the contention-resolution scheme for unit-demand buyers. The matroid extension is addressed in Appendix D.2.5.

C.1 Solving the Single-Dimensional Optimization Problem

The mechanism that attains the single-dimensional benchmark $\mathrm{OPT}^{1,1}$ is a fundamental building block of the constructions presented throughout the paper. As such, existence of a simple mechanism that attains utility $\mathrm{OPT}^{1,1}$ as well as implementation thereof must be explicitly addressed.

We start by characterizing the optimal solution of $\mathrm{OPT}^{1,1}$. A standard idea in mechanism design is to reexpress any single-dimensional mechanism as a distribution over posted price mechanisms (see for example Hartline [30, Chapter 3]). We apply this to a mechanism that attains $\mathrm{OPT}^{1,1}$. Let μ be a distribution over finite nonnegative posted prices and define the allocation and payment functions for a specific posted price $t \geq 0$ to be

$$a_t(z) := \mathbf{1}\{z > t\}, \qquad p_t(z) := t \cdot \mathbf{1}\{z > t\}.$$

With this notation, we can rewrite the single-dimensional benchmark as follows

$$\begin{aligned} \mathrm{OPT}^{1,1}(X; q) = \max_{\mu \geq 0} \quad & \int \mathbb{E}_X[X a_t(X) - p_t(X)]\, d\mu(t) \\ \text{s.t.} \quad & \int \mathbb{E}_X[a_t(X)]\, d\mu(t) \leq q, \\ & \int d\mu(t) \leq 1. \end{aligned}$$

Note that in the formulation, we allow μ to have total mass less than 1, interpreting the missing probability mass as the probability of making no offer (or equivalently, posting an infinite price).

This ex-ante constrained mechanism design problem for a utility maximization objective admits the same optimal solution structure as its revenue maximization counterpart (see Hartline [30, Chapter 3]), with the optimal solution being a distribution μ over at most two prices. To sketch why that is true, consider an optimal distribution $\mu^\star$ that uses more than two prices (but finitely many). We can pick any three prices with positive probability from $\mu^\star$ and redistribute probability

among them while preserving both the sum of their probabilities and their contribution to the ex-ante allocation probability constraint. Any such redistribution can proceed in either direction, reversing which probability weights increase and decrease. Given that the objective is linear in the probabilities of $\mu^\star$, this implies that there always exists one direction of the redistribution that weakly increases the objective. Continuing in that direction until a probability weight becomes zero eliminates a price. We can keep applying this idea until we eventually reach a distribution over at most two prices. We refer to Hartline [30, Chapter 3] which provides a more precise geometric argument.[5]

Following oracle-based approaches in Bayesian revenue maximization [12, 2], we assume oracle access to both the optimal objective value and an implementing mechanism for these single-dimensional optimization programs. For an approximate implementation, we refer to the methodology of Alaei [1, Section 6.1, Lemma 2], which optimizes over a discretized set of prices for a revenue objective. We note that for their methodology to apply to a utility objective, the rounding argument must be modified to control the resulting rounding loss in expected utility.

C.2 Discretized Greedy Cap Optimization

The item-i.i.d. cap program (formulation (B.7)) and the unit-demand cap program (formulation (5.4)) can be approximated by the same greedy discretized procedure. Both maximize a sum of concave functions, each depending on one cap, subject to a bound on the sum of the caps. We first describe the procedure and its guarantees in Lemma 9, then derive the accuracy guarantees for the two applications.

A greedy algorithm for a concave cap-allocation problem. Let $\phi_1, \ldots, \phi_d$ be concave functions on $[0,1]$ with $\phi_i(0) = 0$, and define the concave objective $\Phi(q) := \sum_{i=1}^{d} \phi_i(q_i)$ for $q \in [0,1]^d$. For an integer total cap $C \in \{1, \ldots, d\}$, we want to approximately solve the following optimization problem

$$\begin{aligned} \max_{q} \quad & \Phi(q) \\ \text{s.t.} \quad & \sum_i q_i \le C \\ & 0 \le q_i \le 1 \quad \forall i \in [d]. \end{aligned} \tag{C.1}$$

To efficiently approximate this program's optimal solution, we restrict ex-ante probability caps to the discretized values $\{0, 1/K, \ldots, 1\}$, where $K \ge 1$ is an integer parameter to be chosen later. We build the caps incrementally: at each step, we compare the gains from increasing each cap by $1/K$ and increment the cap with the largest gain in its corresponding function ϕ_i. We now present the greedy algorithm formally.

Fix K and define the increment budget $B := CK$. Let k_i be an integer that records the increments assigned to coordinate i by the algorithm. The algorithm executes the following steps:

1. For every i, set $k_i \leftarrow 0$.

[5] Although stated for revenue, the same convex-hull argument applies to expected utility, which is also linear in the price mixture. Unlike our sketch, it does not require assuming that the optimal price distribution has finite support.

2. Repeat at most B times:

   (a) If no coordinate i with $k_i < K$ has $\phi_i((k_i+1)/K) - \phi_i(k_i/K) > 0$, break. Otherwise, choose such a coordinate that maximizes this gain and set $k_i \leftarrow k_i + 1$.

3. Return the cap vector q^{G} with ${q_i}^{\mathrm{G}} = k_i/K$ for every i.

Note that the algorithm stops when the largest available gain is nonpositive, because no remaining increment could improve the objective.

The following lemma establishes that the greedy algorithm produces feasible ex-ante caps and that the objective value using these caps approximates the value obtained by rounding each optimal cap upward to the next discretized value. The effect of the rounding will then be handled separately for the two following applications of the lemma.

Lemma 9. *Let $\phi_1, \dots, \phi_d$ be concave functions on $[0,1]$ with $\phi_i(0) = 0$. Let $q^\star$ be an optimal solution to optimization problem* (C.1)*, define the rounded caps by ${q_i}^{\uparrow} := \lceil K q_i^\star \rceil / K$, and let q^{G} be the caps returned by the preceding greedy algorithm with increment size $1/K$ and increment budget $B := CK$. The caps q^{G} are feasible for the optimization program and satisfy*

$$\Phi(q^{\uparrow}) \le \left(1 + \frac{d}{B}\right) \Phi(q^{\mathrm{G}}).$$

The algorithm uses $\mathrm{poly}(d, B)$ *evaluations of the functions ϕ_i and polynomial additional computation in d, B.*

Proof. To argue feasibility of caps q^{G}, note that the algorithm maintains $k_i \le K$, so ${q_i}^{\mathrm{G}} = k_i/K \le 1$. Moreover, at most $B = CK$ increments are assigned (across the vector q^{G}), giving $\sum_i {q_i}^{\mathrm{G}} \le C$.

To prove the inequality of the lemma, we first argue that greedy produces the best feasible solution with caps restricted to the discretized values $\{0, 1/K, \dots, 1\}$. We then show that rounding the continuous optimal caps $q^\star$ upward uses at most d additional increments beyond the B-increment budget, and bound the gain of each additional increment by that of greedy's last increment.

Let $\Phi_{\mathrm{G}} := \Phi(q^{\mathrm{G}})$ denote the objective value under the caps q^{G} returned by the greedy algorithm. We restrict the optimization program (C.1) to the discretized caps $\{0, 1/K, \dots, 1\}$ and define the best solution in this restricted space

$$\Phi_K^\star := \max_{\substack{k_i \in \{0,\dots,K\} \\ \sum_i k_i \le B}} \Phi(k_1/K, \dots, k_d/K).$$

We first show that greedy solves the discretized problem optimally, that is, $\Phi_{\mathrm{G}} = \Phi_K^\star$. Let γ_t be the gain selected at iteration $t = 1, \dots, B$, setting it to zero after the algorithm stops. For each function ϕ_i, the successive increment gains $\phi_i(k/K) - \phi_i((k-1)/K)$ are nonincreasing in k by concavity. Combining this with the fact that greedy always chooses the largest gain among the next increments of all coordinates, the gain it selects cannot increase from one iteration to the next. Thus, the sequence γ_t is nonnegative and nonincreasing. Now, take an optimal solution to the discretized problem, express its objective value as the sum of the gains from at most B increments, and order its gains nonincreasingly (to compare them with the ones of the greedy algorithm). By a greedy exchange argument, we can argue that the greedy algorithm attains at least the same objective value, so $\Phi_{\mathrm{G}} = \Phi_K^\star$.

We next compare the greedy objective Φ_{G} with $\Phi_\uparrow := \Phi(q^\uparrow)$, the objective of the rounded optimal caps ${q_i}^\uparrow = \lceil Kq_i^\star \rceil / K$ (where recall $q^\star$ are optimal caps of program (C.1)). The rounded vector $q^\uparrow$ remains in $[0,1]^d$, but may exceed the increment budget B. Each coordinate i requires $K{q_i}^\uparrow$ increments of size $1/K$, so the total number of increments is

$$\sum_{i=1}^{d} K{q_i}^\uparrow = \sum_{i=1}^{d} \lceil Kq_i^\star \rceil \le K \sum_{i=1}^{d} q_i^\star + d \le B + d, \tag{C.2}$$

where the first inequality uses $\lceil x \rceil \le x + 1$, and the second uses feasibility of $q^\star$ and $B = CK$.

To bound the benefit of d additional increments, we define an *extended greedy run*, used only as an analytical tool. It follows the greedy algorithm's selection rules, assigning at most K increments to each coordinate, but omits the total increment limit of B and stops when no positive gain remains.

The extended greedy run makes the same choices as the greedy algorithm for its first B iterations. We retain the notation γ_t for its gains beyond B, again setting them to zero after it stops, and note that the gains γ_t of the extended sequence remain nonnegative and nonincreasing. We can thus bound the objective $\Phi_\uparrow$

$$\begin{aligned}
\Phi_\uparrow &\le \sum_{t=1}^{B+d} \gamma_t = \sum_{t=1}^{B} \gamma_t + \sum_{t=B+1}^{B+d} \gamma_t \\
&\le \sum_{t=1}^{B} \gamma_t + d\gamma_B \\
&\le \sum_{t=1}^{B} \gamma_t + \frac{d}{B} \sum_{t=1}^{B} \gamma_t = \left(1 + \frac{d}{B}\right) \sum_{t=1}^{B} \gamma_t = \left(1 + \frac{d}{B}\right) \Phi_{\mathrm{G}},
\end{aligned} \tag{C.3}$$

where the first inequality is because the rounded caps use at most $B + d$ increments so by the previous greedy exchange argument their objective is bounded by the first $B + d$ gains of the extended run, the second inequality follows from γ_t being nonincreasing, and the third uses that γ_B is at most the average of the first B gains (again since γ_t is nonincreasing), and the final equality holds because the greedy algorithm and the extended run make the same choices until the greedy algorithm stops.

In terms of computation, we compute each coordinate's next available gain once and update only when it is selected. The procedure therefore uses at most $B + d$ evaluations of the functions $\phi_i\colon d$ to compute the initial gains and at most one new evaluation per increment. Selecting the largest available gain at each iteration requires only polynomial additional time in d and B. □

Item-i.i.d. optimization. We now discuss how the greedy algorithm and Lemma 9 give us an approximate solution to the item-i.i.d. cap program (B.7). Recall that we require an approximate solution only for the case where $m < n$.

With $d = n$, $C = m$, and $\phi_i(q) = \mathrm{OPT}^{1,1}(M_i; q)$, the general concave program (C.1) coincides exactly with the item-i.i.d. cap program (B.7), whose optimal value is $\mathcal{R}$. For $0 < \varepsilon \le 1$, choose $K := \lceil n/(m\varepsilon) \rceil$ for the discretization and apply the greedy algorithm. The following lemma states the guarantees for the returned caps and the greedy algorithm's running time.

Lemma 10. *For $m < n$ and $0 < \varepsilon \le 1$, the greedy algorithm with $d = n$, $C = m$, $\phi_i(q) = \mathrm{OPT}^{1,1}(M_i; q)$, and $K = \lceil n/(m\varepsilon) \rceil$ returns ex-ante feasible caps q^{G} satisfying*

$$\sum_i \mathrm{OPT}^{1,1}(M_i; {q_i}^{\mathrm{G}}) \ge \frac{\mathcal{R}}{1+\varepsilon}.$$

The algorithm uses $\mathrm{poly}(n, 1/\varepsilon)$ *oracle calls to* $\mathrm{OPT}^{1,1}$ *and polynomial additional computation in* $n, 1/\varepsilon$.

Proof. With the selected parameters, the optimization program has optimal value $\Phi^\star = \mathcal{R}$ and the greedy algorithm objective is $\Phi_{\mathrm{G}} = \sum_i \mathrm{OPT}^{1,1}(M_i; {q_i}^{\mathrm{G}})$.

The returned caps q^{G} are feasible by Lemma 9. We directly bound the expected utility of the produced greedy caps:

$$\begin{aligned}
\mathcal{R} = \Phi^\star &\le \Phi_\uparrow && {q_i}^\uparrow \ge q_i^\star,\ \mathrm{OPT}^{1,1} \text{ nondecreasing in its cap} \\
&\le \left(1 + \frac{n}{mK}\right) \Phi_{\mathrm{G}} && \text{Lemma 9} \\
&\le (1+\varepsilon) \sum_i \mathrm{OPT}^{1,1}(M_i; {q_i}^{\mathrm{G}}), && \text{choice of } K
\end{aligned}$$

The computational guarantees follow from Lemma 9 with $B = mK = O(n/\varepsilon)$ and $d = n$. □

Unit-demand cap optimization. We now discuss how the greedy algorithm and Lemma 9 give us a solution to the unit-demand cap program (5.4) with an additive approximation guarantee.

With $d = m$, $C = 1$, and $\phi_j(q) := \mathrm{OPT}^{1,1}(V_j; q) - \frac{e}{e-1}\lambda_j q$, the general concave program (C.1) coincides exactly with the unit-demand cap program (5.4). Denote the optimal value by $\mathcal{R}_\lambda := \Phi^\star$ and its optimal caps $q_j^\star$. For $0 < \delta \le 1$, choose $K := \lceil m(1 + e/(e-1))/\delta \rceil$ and apply the greedy algorithm. Unlike the previous application, rounding upward may decrease the objective because it increases the demand-weight terms. Combining a bound on this rounding loss with Lemma 9 gives the additive approximation guarantee of Lemma 11.

Solving the program up to an additive error lets us compute a mechanism whose utility guarantee matches that of Lemma 5 up to an additive loss, which in turn carries over to the guarantees of Lemma 6. In proving Lemma 6 we normalize values, payments, and demand weights by dividing by $\sum_{i,j} \mathbb{E}[V_{ij}]$. Since normalization affects the additive approximation guarantee, the following lemma and its proof are under the same normalization.

Lemma 11. *Suppose $\sum_j \mathbb{E}[V_j] \le 1$ and $0 \le \lambda_j \le 1$ for every item j. For $0 < \delta \le 1$, the greedy algorithm with $d = m$, $C = 1$, $\phi_j(q) = \mathrm{OPT}^{1,1}(V_j; q) - \frac{e}{e-1}\lambda_j q$, and $K = \lceil m(1 + e/(e-1))/\delta \rceil$ returns ex-ante feasible caps q^{G} satisfying*

$$\sum_j \left[\mathrm{OPT}^{1,1}(V_j; {q_j}^{\mathrm{G}}) - \frac{e}{e-1} \lambda_j {q_j}^{\mathrm{G}} \right] \ge \mathcal{R}_\lambda - \delta.$$

The algorithm uses $\mathrm{poly}(m, 1/\delta)$ *oracle calls to* $\mathrm{OPT}^{1,1}$ *and polynomial additional computation in* $m, 1/\delta$.

Proof. With the selected parameters, the optimization program has optimal value $\Phi^\star = \mathcal{R}_\lambda$ and the greedy algorithm objective is $\Phi_{\mathrm{G}} = \sum_j [\mathrm{OPT}^{1,1}(V_j; {q_j}^{\mathrm{G}}) - \frac{e}{e-1}\lambda_j {q_j}^{\mathrm{G}}]$.

The returned caps q^{G} are feasible by Lemma 9. The rounding used in Lemma 9 increases item j's demand-weight term by at most $\frac{e}{(e-1)K}\lambda_j$. Its utility term cannot decrease, by monotonicity of $\mathrm{OPT}^{1,1}(V_j;\cdot)$. Thus the rounded objective loses at most the sum of these demand-weight term increases:

$$\Phi_\uparrow \geq \mathcal{R}_\lambda - \frac{e}{(e-1)K}\sum_j \lambda_j.$$

Meanwhile, Lemma 9, with $d = m$, $C = 1$, and $B = K$, bounds the rounded objective by

$$\Phi_\uparrow \leq \left(1 + \frac{m}{K}\right)\Phi_{\mathrm{G}}.$$

Combining these two inequalities proves the desired inequality

$$\begin{aligned}
\mathcal{R}_\lambda - \Phi_{\mathrm{G}} &\leq \frac{m}{K}\Phi_{\mathrm{G}} + \frac{e}{(e-1)K}\sum_j \lambda_j \\
&\leq \frac{m(1+e/(e-1))}{K} && \Phi_{\mathrm{G}} \leq \sum_j \mathbb{E}[V_j] \leq 1,\ \sum_j \lambda_j \leq m \\
&\leq \delta. && \text{choice of } K
\end{aligned}$$

The computational guarantees follow from Lemma 9 with $B = K = O(m/\delta)$ and $d = m$. □

C.3 Computing the Free-Lottery Parameter

We work in the normalized units used in Appendix C.4, with additive error $\gamma \in (0, 1]$. The construction of mechanism $\mathcal{M}^{\mathrm{free}}$ in Section 5.1 seeks a parameter $h > 0$ with $U(\mathcal{M}_h^{\mathrm{free}}) = 2h$. The ratio $U(\mathcal{M}_h^{\mathrm{free}})/h$ is continuous and nonincreasing, since each coefficient $\ell_j(h)/h$ is nonincreasing. In normalized units, bisect $[a, b] = [0, 1]$, moving the lower endpoint to the midpoint c if $U(\mathcal{M}_c^{\mathrm{free}}) > 2c$ and the upper endpoint otherwise, until $b - a \leq \gamma$. By construction, the final upper endpoint satisfies $U(\mathcal{M}_b^{\mathrm{free}}) \leq 2b$: this holds initially by normalization and is preserved whenever the upper endpoint is updated. For the lower bound, consider two cases.

Case 1: $a > 0$. The lower endpoint was set to a tested midpoint satisfying $U(\mathcal{M}_a^{\mathrm{free}}) > 2a$. Monotonicity and $b - a \leq \gamma$ give

$$U(\mathcal{M}_b^{\mathrm{free}}) \geq U(\mathcal{M}_a^{\mathrm{free}}) > 2a \geq 2b - 2\gamma.$$

Case 2: $a = 0$. Termination gives $b \leq \gamma$, so nonnegative utility implies

$$U(\mathcal{M}_b^{\mathrm{free}}) \geq 0 \geq 2b - 2\gamma.$$

Thus $2b - 2\gamma \leq U(\mathcal{M}_b^{\mathrm{free}}) \leq 2b$ after $O(1 + \log(1/\gamma))$ exact evaluations, even if no root exists. By (5.7) and using these bounds and the fact that the allocation cost at most b, we obtain

$$\begin{aligned}
\frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\mathrm{alloc}}(C) \leq \frac{U(\mathcal{M}_b^{\mathrm{free}})}{2} &\leq U(\mathcal{M}_b^{\mathrm{free}}) - b + \gamma \\
&\leq U(\mathcal{M}_b^{\mathrm{free}}) - \sum_j \lambda_j q_j^{\mathrm{alloc}}(\mathcal{M}_b^{\mathrm{free}}) + \gamma.
\end{aligned}$$

Rearranging yields

$$U(\mathcal{M}_b^{\text{free}}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}_b^{\text{free}}) \geq \frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\text{alloc}}(C) - \gamma. \tag{C.4}$$

This additive guarantee holds whether or not an exact root exists.

C.4 Proof of the Demand-Weight Balancing Lemma

In this subsection, we prove Lemma 6, which describes the properties and the utility guarantee of mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$.

Lemma 6. *Using the demand weights $\lambda^1, \ldots, \lambda^T$ generated by the procedure above, for every $\varepsilon > 0$, after $T = O(n^4 m^4/\varepsilon^2)$ rounds, $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ is ex-ante feasible and the mechanism's total expected utility satisfies*

$$U(\overline{\mathcal{M}}_{\text{het}}^{n,m}) \geq \left(\frac{1}{2 + e/(e-1)} - \varepsilon\right) \text{OPT}^{n,m}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\text{OPT}^{1,1}$ and exact evaluation of the required expectations and item allocation probabilities.

Proof of Lemma 6. It suffices to consider $0 < \varepsilon < 1/(2 + e/(e-1))$, since otherwise the lemma's utility lower bound is nonpositive (and is trivially true).

For the proof of Lemma 6, it is convenient to normalize expected total utility to lie in $[0, 1]$ (note that this preserves incentives and multiplicative approximation ratios). To this end, let $S := \sum_{i,j} \mathbb{E}[V_{ij}]$, a loose upper bound on $\text{OPT}^{n,m}$ that counts the expected welfare of allocating every item to every buyer. Divide values, payments, and demand weights by S, retaining the same notation for the normalized quantities. In normalized units, the step size in the construction of $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ is $\eta = 1/(n\sqrt{T})$, and the demand-weight update rule (5.10) becomes

$$\lambda_j^{t+1} := \min\left\{1, \max\left\{0, \lambda_j^t + \eta(r_j^t - 1)\right\}\right\}.$$

Recall the construction of $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ in Section 5.1.1: starting from demand weights $\lambda^1 = 0$, we compute single-buyer mechanisms $\mathcal{M}_i^t$ in each of T rounds and update the demand weights based on the expected item requests of that round. The final mechanism independently samples a round for each buyer and either offers the corresponding mechanism $\mathcal{M}_i^t$ or makes no offer, with probabilities chosen to ensure ex-ante feasibility of $\overline{\mathcal{M}}_{\text{het}}^{n,m}$. To track the performance of mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$, let $U^t := \sum_i U(\mathcal{M}_i^t)$ and $r_j^t := \sum_i q_j^{\text{alloc}}(\mathcal{M}_i^t)$ denote the total expected utility and expected number of requests for item j in round t and define their averages across the T rounds as

$$\overline{U} := \frac{1}{T}\sum_{t=1}^{T} U^t, \qquad \overline{r}_j := \frac{1}{T}\sum_{t=1}^{T} r_j^t.$$

Now, if mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ ran all sampled mechanisms, it would generate expected demand $\overline{r}_j$ for item j, which may exceed one. For this reason, the mechanism instead offers each buyer's sampled mechanism with probability $1/\max\{1, \max_k \overline{r}_k\}$ and otherwise makes no offer. The resulting expected number of requests for each item j is $\overline{r}_j/\max\{1, \max_k \overline{r}_k\} \leq 1$, ensuring ex-ante feasibility.

It remains to show that making no offer with the prescribed probability does not remove too much utility from mechanism $\overline{\mathcal{M}}_{\text{het}}^{n,m}$. We first prove that the utility loss because of the probability of no offer is at most the excess expected item requests that would arise if every sampled mechanism were run with probability 1:

$$\begin{aligned}
U(\overline{\mathcal{M}}_{\text{het}}^{n,m}) &= \frac{\overline{U}}{\max\{1, \max_j \overline{r}_j\}} && \text{definition of } \overline{\mathcal{M}}_{\text{het}}^{n,m} \\
&= \overline{U} - \overline{U}\left(1 - \frac{1}{\max\{1, \max_j \overline{r}_j\}}\right) \\
&\geq \overline{U} - \left(1 - \frac{1}{\max\{1, \max_j \overline{r}_j\}}\right) && \overline{U} \leq 1 \text{ by normalization} \\
&\geq \overline{U} - \big(\max\{1, \max_j \overline{r}_j\} - 1\big) && 1 - 1/z \leq z - 1 \text{ for } z \geq 1 \\
&= \overline{U} - \max_j (\overline{r}_j - 1)^+ \\
&\geq \overline{U} - \sum_j (\overline{r}_j - 1)^+. && \text{(C.5)}
\end{aligned}$$

Before proceeding further with the utility lower bound, we address how errors due to approximate computation come into play here. To be precise, each round in the construction of $\overline{\mathcal{M}}_{\text{het}}^{n,m}$ uses the single-buyer construction of Lemma 5, which in turn requires solving the unit-demand cap program (5.4) for each buyer and computing a free lottery. Lemma 11 shows how to solve this program with additive error $\delta \in (0, 1]$ (in normalized units), while in Appendix C.3 we discuss how to compute the free-lottery mechanism with additive error $\gamma > 0$.

Taking this into account, the following lemma shows that with an appropriate choice of the step size η, the right-hand side of (C.5) is at least a constant fraction of $\text{OPT}^{n,m}$, up to additive errors from the approximate solution of the unit-demand cap program (5.4), the free-lottery mechanism computation in each round and from the projected online gradient descent used across the T rounds. We state the lemma here, but we will prove it after completing the proof of Lemma 6.

Lemma 12. *Let $\delta \in (0, 1]$ and $\gamma \in (0, 1]$ be the additive errors in cap optimization and free-lottery computation, respectively, for each buyer at each round. With step size $\eta = 1/(n\sqrt{T})$, after T demand-weight update rounds, we get*

$$\overline{U} - \sum_j (\overline{r}_j - 1)^+ \geq \frac{\text{OPT}^{n,m}}{2 + e/(e-1)} - \frac{nm}{\sqrt{T}} - n(\delta + \gamma).$$

We now choose the number of rounds T and the additive errors δ and γ of the approximate computation for each buyer and each round:

$$T := \left\lceil \frac{4n^4 m^4}{\varepsilon^2} \right\rceil, \qquad \delta + \gamma \leq \frac{\varepsilon}{2n^2 m}.$$

These choices of parameters balance the two error terms in Lemma 12, and so the final guar-

antee follows:

$$
\begin{aligned}
U(\overline{\mathcal{M}}^{n,m}_{\text{het}}) &\geq \overline{U} - \sum_j (\overline{r}_j - 1)^+ && \text{(C.5)}\\
&\geq \frac{\text{OPT}^{n,m}}{2 + e/(e-1)} - \frac{nm}{\sqrt{T}} - n(\delta + \gamma) && \text{Lemma 12}\\
&\geq \frac{\text{OPT}^{n,m}}{2 + e/(e-1)} - \frac{\varepsilon}{2nm} - \frac{\varepsilon}{2nm} && \text{choices of } T,\ \delta, \text{ and } \gamma\\
&\geq \frac{\text{OPT}^{n,m}}{2 + e/(e-1)} - \varepsilon \text{OPT}^{n,m} && \text{OPT}^{n,m} \geq 1/(nm)\\
&= \left(\frac{1}{2 + e/(e-1)} - \varepsilon\right) \text{OPT}^{n,m},
\end{aligned}
$$

where $\text{OPT}^{n,m} \geq 1/(nm)$ because allocating the item for free to the buyer in the buyer-item pair with the largest expected value is feasible and yields utility at least $1/(nm)$ in normalized units (since $\sum_{i,j} \mathbb{E}[V_{ij}] = 1$).

We emphasize again that multiplying values and payments by the normalizing factor S restores the original units without changing request probabilities, incentive properties or multiplicative approximation factors. □

It remains to prove Lemma 12.

Proof of Lemma 12. The proof is split into two parts: lower bounding the average utility $\overline{U}$ and upper bounding the excess expected requests term $\sum_j (\overline{r}_j - 1)^+$.

Average utility bound: To lower bound the average utility $\overline{U}$, we essentially apply Lemma 5 to the mechanisms produced at each round. However, note that Lemma 5 used the optimal solution of the unit-demand cap program (5.4) while we instead propose a greedy algorithm that approximately solves the same program up to additive error (see Lemma 11). For this reason, we also need to address how this error of the approximate solution of program (5.4) propagates to the guarantee of Lemma 5.

To be precise, by Lemma 11, we compute approximately optimal ex-ante caps whose objective value is within δ of the optimum (in normalized units). In the construction proving Lemma 5, this affects only the pricing branch, weakening inequality (5.3) by at most $(1 - 1/e)\delta$. Free-lottery computation with error γ weakens (5.5) by at most γ (see Section C.3). Mixing the two bounds gives, for every buyer i and round t,

$$
\begin{aligned}
U(\mathcal{M}^t_i) - \sum_j \lambda^t_j q^{\text{alloc}}_j(\mathcal{M}^t_i) + \sum_j \lambda^t_j q^{\text{alloc}}_j(C_i) &\geq \frac{U(C_i) - \delta - 2\gamma}{2 + e/(e-1)}\\
&\geq \frac{U(C_i)}{2 + e/(e-1)} - (\delta + \gamma).
\end{aligned}
$$

Following the same steps used to derive inequality (5.2), fix any feasible, BIC, interim individually rational multi-buyer mechanism C with nonnegative payments. For each buyer i, let C_i be the single-buyer mechanism obtained by averaging C over the other buyers' values. These mechanisms preserve the buyers' expected utilities and ex-ante allocation probabilities, so $\sum_i U(C_i) = U(C)$

and $\sum_i q_j^{\text{alloc}}(C_i) \le 1$ for every item j. Using the preceding bound with the additive error, we get

$$\begin{aligned} U^t + \lambda^t \cdot (\mathbf{1} - r^t) &\ge \sum_i \left[U(\mathcal{M}_i^t) - \sum_j \lambda_j^t q_j^{\text{alloc}}(\mathcal{M}_i^t) + \sum_j \lambda_j^t q_j^{\text{alloc}}(C_i) \right] \\ &\ge \frac{\sum_i U(C_i)}{2 + e/(e-1)} - n(\delta + \gamma) \\ &= \frac{U(C)}{2 + e/(e-1)} - n(\delta + \gamma). \end{aligned}$$

Since this bound holds for every comparison mechanism C, taking the supremum over C yields

$$U^t + \lambda^t \cdot (\mathbf{1} - r^t) \ge \frac{\text{OPT}^{n,m}}{2 + e/(e-1)} - n(\delta + \gamma).$$

Averaging the inequality over the T rounds yields the following bound on average utility $\overline{U}$:

$$\overline{U} - \frac{1}{T} \sum_{t=1}^{T} \lambda^t \cdot (r^t - \mathbf{1}) \ge \frac{\text{OPT}^{n,m}}{2 + e/(e-1)} - n(\delta + \gamma), \tag{C.6}$$

where $\mathbf{1}$ denotes the all-1s vector.

Excess expected requests bound: We next bound the excess expected requests using the demand-weight updates. Following the online gradient descent analysis of Zinkevich [41], we compare the average demand-weight term $T^{-1} \sum_{t=1}^{T} \lambda^t \cdot (\mathbf{1} - r^t)$ with $T^{-1} \sum_{t=1}^{T} \lambda^{\text{fix}} \cdot (\mathbf{1} - r^t)$ for a fixed vector $\lambda^{\text{fix}} \in [0,1]^m$. Choosing this vector λ^{fix} to assign weight one to items whose average expected requests exceed one, and zero otherwise, makes the latter term equal to the negative of the excess expected requests $\sum_j (\overline{r}_j - 1)^+$. To compare the two average demand-weight terms, we track how the squared distance between λ^t and λ^{fix} changes in each round's update:

$$\begin{aligned} \|\lambda^{t+1} - \lambda^{\text{fix}}\|_2^2 &\le \|\lambda^t + \eta(r^t - \mathbf{1}) - \lambda^{\text{fix}}\|_2^2 \\ &= \|\lambda^t - \lambda^{\text{fix}}\|_2^2 - 2\eta(\lambda^t - \lambda^{\text{fix}}) \cdot (\mathbf{1} - r^t) + \eta^2 \|r^t - \mathbf{1}\|_2^2, \end{aligned}$$

where $\mathbf{1}$ denotes the all-1s vector, and the inequality is because the clipping of $\lambda^t + \eta(r^t - \mathbf{1})$ in the definition of λ^{t+1} moves the expression inside $[0,1]^m$ (if required), and λ^{fix} also lies in $[0,1]^m$, and therefore clipping can only reduce the distance to λ^{fix}.

Rearranging and dividing by $2\eta > 0$ yields

$$(\lambda^t - \lambda^{\text{fix}}) \cdot (\mathbf{1} - r^t) \le \frac{\|\lambda^t - \lambda^{\text{fix}}\|_2^2 - \|\lambda^{t+1} - \lambda^{\text{fix}}\|_2^2}{2\eta} + \frac{\eta}{2} \|r^t - \mathbf{1}\|_2^2. \tag{C.7}$$

We now sum inequality (C.7) over rounds and divide by T:

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\lambda^t\cdot(\mathbf{1}-r^t) &= \lambda^{\text{fix}}\cdot(\mathbf{1}-\overline{r}) + \frac{1}{T}\sum_{t=1}^{T}(\lambda^t-\lambda^{\text{fix}})\cdot(\mathbf{1}-r^t)\\
&\le \lambda^{\text{fix}}\cdot(\mathbf{1}-\overline{r}) + \frac{1}{2\eta T}\sum_{t=1}^{T}\left(\|\lambda^t-\lambda^{\text{fix}}\|_2^2 - \|\lambda^{t+1}-\lambda^{\text{fix}}\|_2^2\right) + \frac{\eta}{2T}\sum_{t=1}^{T}\|r^t-\mathbf{1}\|_2^2\\
&= \lambda^{\text{fix}}\cdot(\mathbf{1}-\overline{r}) + \frac{\|\lambda^1-\lambda^{\text{fix}}\|_2^2 - \|\lambda^{T+1}-\lambda^{\text{fix}}\|_2^2}{2\eta T} + \frac{\eta}{2T}\sum_{t=1}^{T}\|r^t-\mathbf{1}\|_2^2\\
&\le \lambda^{\text{fix}}\cdot(\mathbf{1}-\overline{r}) + \frac{\sum_{j=1}^{m}(\lambda_j^{\text{fix}})^2}{2\eta T} + \frac{\eta}{2T}\,Tmn^2\\
&\le \lambda^{\text{fix}}\cdot(\mathbf{1}-\overline{r}) + \frac{nm}{\sqrt{T}}, \qquad \text{(C.8)}
\end{aligned}$$

where the first equality adds and subtracts the fixed-weight term and uses $\overline{r} = T^{-1}\sum_t r^t$, the first inequality applies (C.7), the second equality telescopes terms, the second inequality uses $\lambda^1 = 0$, drops the nonpositive final-distance term and uses $|r_j^t - 1| \le n$ (since r_j^t is the sum of n request probabilities), and the last inequality uses $0 \le \lambda_j^{\text{fix}} \le 1$ and $\eta = 1/(n\sqrt{T})$.

Since this bound holds for every $\lambda^{\text{fix}} \in [0,1]^m$, choose $\lambda_j^{\text{fix}} = 1$ if $\overline{r}_j > 1$ and $\lambda_j^{\text{fix}} = 0$ otherwise. For this choice, $\lambda^{\text{fix}}\cdot(\mathbf{1}-\overline{r}) = -\sum_j(\overline{r}_j-1)^+$, so rearranging inequality (C.8) gives

$$\sum_j(\overline{r}_j-1)^+ - \frac{1}{T}\sum_{t=1}^{T}\lambda^t\cdot(r^t-\mathbf{1}) \le \frac{nm}{\sqrt{T}}. \qquad \text{(C.9)}$$

Combining (C.9) and (C.6) proves the lemma

$$\overline{U} - \sum_j(\overline{r}_j-1)^+ \ge \frac{\text{OPT}^{n,m}}{2+e/(e-1)} - n(\delta+\gamma) - \frac{nm}{\sqrt{T}}.$$

□

C.5 Contention Resolution and Its Implementation

This subsection describes the contention-resolution scheme used in the paper and explains how to compute buyers' request probabilities and selection probabilities c_j. We describe the construction for unit-demand buyers; the matroid extension is addressed in Appendix D.2.5. All probabilities are computed before buyers choose their options or draw their rounds. A buyer's request probabilities depend only on that buyer's value distributions and single-buyer mechanisms; contention probabilities use the final request probabilities of all buyers for the same item.

Request probabilities. For each buyer i, item j, and round t, we want to compute $q_j^{\text{alloc}}(\mathcal{M}_i^t)$, the probability that buyer i requests item j. Recall that mechanism $\mathcal{M}_i^t$ runs the pricing mechanism $\mathcal{M}_{i,t}^{UD}$ with probability $e/(3e-2)$ and the free-lottery mechanism $\mathcal{M}_{i,t}^{\text{free}}$ with probability $2(e-1)/(3e-2)$. Conditioning on which mechanism is drawn gives

$$q_j^{\text{alloc}}(\mathcal{M}_i^t) = \frac{e}{3e-2}\,q_j^{\text{alloc}}(\mathcal{M}_{i,t}^{UD}) + \frac{2(e-1)}{3e-2}\,q_j^{\text{alloc}}(\mathcal{M}_{i,t}^{\text{free}}).$$

Thus it suffices to compute the request probability under each of these two mechanisms.

In the pricing mechanism $\mathcal{M}_{i,t}^{UD}$, each item's price P_k is drawn independently. Suppose item j's price is drawn from p_1, p_2 with probabilities $\theta, 1-\theta$. We compute its request probability by conditioning on its drawn price and averaging over the two choices:

$$q_j^{\text{alloc}}(\mathcal{M}_{i,t}^{UD}) = \theta \Pr(j \text{ requested} \mid P_j = p_1) + (1-\theta) \Pr(j \text{ requested} \mid P_j = p_2).$$

If a drawn price is infinite, item j is not offered and the corresponding conditional request probability is zero. For a finite price p, the buyer requests j if it gives the largest positive surplus. We compute this probability by conditioning on item j's value and averaging over its value distribution:

$$\begin{aligned}\Pr(j \text{ requested} \mid P_j = p) &= \Pr\left(V_{ij} > p,\ j = \arg\max_k \{V_{ik} - P_k\} \,\middle|\, P_j = p\right) \\ &= \int \mathbf{1}\{v > p\} \Pr(v - p > V_{ik} - P_k \text{ for all } k \neq j)\ dF_{ij}(v) \\ &= \int \mathbf{1}\{v > p\} \prod_{k \neq j} \Pr(v - p > V_{ik} - P_k)\ dF_{ij}(v).\end{aligned}$$

For simplicity, the formula above is written assuming that positive-surplus ties occur with probability zero (and should be adjusted to account for the uniform random tie-breaking rule if ties have positive probability).

In the free-lottery mechanism $\mathcal{M}_{i,t}^{\text{free}}$, the buyer chooses among lotteries, where the lottery for item k has expected utility $\ell_k(h)V_{ik}$. Conditional on choosing the lottery for item j, that item is delivered with probability $\ell_j(h)$. Thus the request probability is

$$q_j^{\text{alloc}}(\mathcal{M}_{i,t}^{\text{free}}) = \ell_j(h) \Pr\left(V_{ij} > 0,\ j = \arg\max_k \{\ell_k(h)V_{ik}\}\right).$$

Here too, we omit ties for simplicity.

Our distribution-access assumptions permit exact evaluation of these probabilities, including the specified tie-breaking rule.

Taking into account the randomization over the T rounds and the probability of making no offer, define q_{ij}^{alloc} to be the probability that buyer i requests item j before contention resolution:

$$q_{ij}^{\text{alloc}} := q_{ij}^{\text{alloc}}(\overline{\mathcal{M}}_{\text{het}}^{n,m}) = \frac{1}{T \max\{1, \max_k \overline{r}_k\}} \sum_{t=1}^{T} q_j^{\text{alloc}}(\mathcal{M}_i^t).$$

The contention resolution rule. After the buyers make their requests, we resolve them by first selecting requesters using the fair rule of Feige and Vondrak [22], then serving the selected requests with a reduced probability specified below. Other contention resolution schemes also work for our mechanisms, for example [33].

For each item, the rule of Feige and Vondrak [22] selects no buyer if there are no requests, selects the sole requester if there is one, and otherwise selects each requester i with probability

$$\frac{1}{\sum_k q_{kj}^{\text{alloc}}} \left(\frac{\sum_{k \in A_j \setminus \{i\}} q_{kj}^{\text{alloc}}}{|A_j| - 1} + \frac{\sum_{k \notin A_j} q_{kj}^{\text{alloc}}}{|A_j|} \right),$$

where A_j is the set of buyers requesting item j; if $\sum_k q_{kj}^{\text{alloc}} = 0$, select a requester uniformly instead.

Under this selection rule, Lemma 1 of Feige and Vondrak [22] gives that every buyer has the same probability of selection conditional on requesting item j:

$$c_j = \Pr\big(i \text{ selected for } j \mid i \text{ requests } j\big) = \begin{cases} 1, & \text{if } \sum_k q_{kj}^{\text{alloc}} = 0, \\ \dfrac{1 - \prod_k (1 - q_{kj}^{\text{alloc}})}{\sum_k q_{kj}^{\text{alloc}}}, & \text{otherwise.} \end{cases} \tag{C.10}$$

A technical detail is that when $q_{ij}^{\text{alloc}} = 0$, the conditional probability above is undefined. Nevertheless, the same averaging argument as in the proof of Feige and Vondrak [22, Lemma 1] gives selection probability c_j for a request induced by a deviation, using the original request probabilities; if $\sum_k q_{kj}^{\text{alloc}} = 0$, any deviating buyer who requests item j is the sole requester and is thus selected with probability 1 (since when checking BIC, all other buyers report truthfully and request this item with probability zero).

If $\sum_k q_{kj}^{\text{alloc}} = 0$, then $c_j = 1$ by definition. Otherwise, we lower-bound the selection probability as follows:

$$c_j \geq \frac{1 - \prod_k (1 - q_{kj}^{\text{alloc}})}{\sum_k q_{kj}^{\text{alloc}}} \geq \frac{1 - e^{-\sum_k q_{kj}^{\text{alloc}}}}{\sum_k q_{kj}^{\text{alloc}}} \geq 1 - \frac{1}{e},$$

where the second inequality follows from $(1 - x) \leq e^{-x}$ and the third from ex-ante feasibility, that is $\sum_k q_{kj}^{\text{alloc}} \leq 1$.

Having selected at most one requester per item, we further reduce service probabilities to preserve buyers' choices. A buyer selected under the fair rule discussed above is then allocated and charged with probability $(1 - 1/e)/c_j$. This makes it so that any possible buyer request, including one induced by a deviation, results in allocation and (possible) payment with probability exactly $1 - 1/e$.

This implementation uses exact request probabilities to equalize service probabilities and preserve BIC. The approximation errors analyzed above concern cap optimization and free-lottery parameter selection; we do not establish an implementation using approximate request probabilities.

D Full Construction for Matroid-Constrained Additive Buyers

We give the full construction and analysis of the matroid-constrained mechanism in Section D.1, with supporting proofs and computational details in Section D.2.

D.1 The Full Construction and Analysis

We extend the mechanism of Section 5 to matroid-constrained additive buyers. The design and analysis follow the same steps: posted prices and free lotteries provide a single buyer guarantee, demand-weight updates balance expected item requests, and itemwise contention resolution makes the allocation feasible ex-post. In fact, the single-buyer comparison has the same guarantee, meaning that we can reuse the same demand-weight update rule and analysis, while additivity and downward closure allow us to reuse the same itemwise contention-resolution argument.

The only substantial differences are in how we design the single-buyer mechanism and prove its guarantees. In the posted-price part of the mechanism, the buyer chooses an independent set maximizing total surplus rather than a single item; the pricing guarantee uses Lemma 14, a matroid analogue of Lemma 3. In the free-lottery part, the buyer chooses a lottery over at most two independent sets instead of a single-item lottery. We develop these adaptations in Section D.1.1, then apply the multi-buyer arguments of Sections 5.1.1 and 5.2 to obtain the guarantee in Theorem 4. Appendix D.2 gives the remaining menu, budget-selection, balancing, and conditioning details.

D.1.1 The Single-Buyer Extension

For nonnegative demand weights λ_j, relax the item-supply constraints while retaining each buyer's matroid constraint. The partial Lagrangian of a mechanism $\mathcal{M}$ is

$$\mathcal{L}_\lambda(\mathcal{M}) := \sum_i U_i(\mathcal{M}) + \sum_j \lambda_j \left(1 - \sum_i q_{ij}^{\text{alloc}}(\mathcal{M})\right). \tag{D.1}$$

The following lemma states the guarantee we can prove for a single-buyer mechanism. It is the same guarantee as Lemma 5, with a matroid constraint replacing the unit-demand constraint (although the mechanism and analysis differ slightly).

Lemma 13. *For a buyer i whose feasible sets form a known matroid $\mathcal{I}_i$, and any nonnegative demand-weight vector λ, there exists a BIC, ex-post individually rational single-buyer mechanism $\mathcal{M}_i$ allocating a set in $\mathcal{I}_i$ with nonnegative payments, such that every BIC, individually rational comparison mechanism C_i allocating a set in $\mathcal{I}_i$ with nonnegative payments satisfies*

$$U(\mathcal{M}_i) - \sum_j \lambda_j q_{ij}^{\text{alloc}}(\mathcal{M}_i) \geq \frac{U(C_i)}{2 + e/(e-1)} - \sum_j \lambda_j q_{ij}^{\text{alloc}}(C_i). \tag{D.2}$$

Proof of Lemma 13. Fix buyer i and suppress their index. Fix a comparison mechanism C, and let D denote its deletion loss. As in Section 5.1, posted prices recover D, while free lotteries recover the remaining term $U(C) - D$.

Posted prices $\mathcal{M}^{\mathcal{I}}$. Replace the unit-demand cap constraint in (5.4) by $q \in \mathcal{X}$, leaving the objective unchanged:

$$\max_{q \in \mathcal{X}} \sum_j \left[\text{OPT}^{1,1}(V_j; q_j) - \frac{e}{e-1}\lambda_j q_j\right]. \tag{D.3}$$

As in the proof of Lemma 5, we use optimal caps $q^\star$ for this argument. Appendix D.2.2 gives an efficient approximation, whose error we account for in the final guarantee.

Draw the prices prescribed by $\text{OPT}^{1,1}(V_j; q_j^\star)$ independently and let the buyer choose an independent set maximizing total surplus.[6] The following lemma extends Lemma 3 to matroid constraints. Its utility bound follows from the matroid correlation-gap bound of Yan [39, Lemma 4.1]; Appendix D.2 gives the precise application.

[6] We omit items with nonpositive surplus before applying the tie-breaking convention of Section 2. Removing a zero-surplus item preserves feasibility and utility, and ensures that item j is allocated only when $V_j - P_j > 0$, as needed for its allocation cap.

Lemma 14. *Let $V_1, \dots, V_m$ be independent item values and let P_j be nonnegative, possibly randomized prices drawn independently of the values and one another, satisfying $\Pr(V_j > P_j) \le q_j$. Let $U(\mathcal{M}_j^1)$ be the expected utility from offering price P_j for item j alone. If $q \in \mathcal{X}$, then the mechanism $\mathcal{M}^{\mathcal{I}}$ that offers these prices together and lets the buyer choose a surplus-maximizing independent set of positive-surplus items allocates item j with probability at most q_j and has expected utility*

$$U(\mathcal{M}^{\mathcal{I}}) \ge (1 - 1/e) \sum_j U(\mathcal{M}_j^1).$$

We now establish the matroid analogue of the unit-demand bound (5.3). The calculation is the same: we only replace Lemma 3 by Lemma 14 and check that the comparison mechanism's allocation probabilities lie in $\mathcal{X}$ rather than sum to at most one.

$$\begin{aligned}
&U(\mathcal{M}^{\mathcal{I}}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}^{\mathcal{I}}) \\
&\quad \ge \left(1 - \frac{1}{e}\right) \sum_j \text{OPT}^{1,1}(V_j; q_j^\star) - \sum_j \lambda_j q_j^\star && \text{Lemma 14, } q_j^{\text{alloc}}(\mathcal{M}^{\mathcal{I}}) \le q_j^\star \\
&\quad \ge \left(1 - \frac{1}{e}\right) \sum_j \text{OPT}^{1,1}(V_j; q_j^{\text{alloc}}(C)) - \sum_j \lambda_j q_j^{\text{alloc}}(C) && \text{optimality of caps } q^\star,\ q^{\text{alloc}}(C) \in \mathcal{X} \\
&\quad \ge \left(1 - \frac{1}{e}\right) D - \sum_j \lambda_j q_j^{\text{alloc}}(C). && \text{Lemma 2}
\end{aligned} \tag{D.4}$$

Free lotteries $\mathcal{M}^{\text{free}}$. For the matroid setting, a free lottery will draw an independent set rather than a single item. As in the unit-demand setting, we choose a parameter $h > 0$ and allow lotteries with item allocation probabilities q_j^{alloc} only if their expected allocation cost $\sum_j \lambda_j q_j^{\text{alloc}}$ is at most h. Since the matroid polytope $\mathcal{X}$ consists exactly of the allocation-probability vectors of lotteries over independent sets, imposing the expected allocation cost bound h restricts these vectors to

$$\mathcal{X}_h := \left\{ q^{\text{alloc}} \in \mathcal{X} : \sum_j \lambda_j q_j^{\text{alloc}} \le h \right\}. \tag{D.5}$$

We now define the free-lottery mechanism $\mathcal{M}_h^{\text{free}}$. Write $\Lambda(T) := \sum_{j \in T} \lambda_j$. We give the buyer the demand weights λ and budget h and let them choose one of the following options, free of charge:

1. Receive any independent set T with $\Lambda(T) \le h$ deterministically.
2. Choose any two independent sets S, T with $\Lambda(S) < h < \Lambda(T)$; receive T with probability $(h - \Lambda(S))/(\Lambda(T) - \Lambda(S))$ and S otherwise.

The buyer submits their chosen option, and the mechanism verifies membership in $\mathcal{I}$ (using the membership oracle) and the allocation cost inequalities before implementing it. In this mechanism, the buyer is responsible for computing a utility-maximizing option (similarly to how the pricing mechanism leaves the choice of a surplus-maximizing independent set to the buyer).

Mechanism $\mathcal{M}_h^{\text{free}}$ makes the menu as generous as possible subject to the allocation cost bound: for every value vector v, it attains the maximum utility among all free lotteries over independent

sets with expected allocation cost at most h (see Appendix D.2.3 for the argument). Formally, denote the expected utility of a buyer with value vector v in mechanism $\mathcal{M}_h^{\text{free}}$ by $u_h^{\text{free}}(v)$, and its expectation over the buyer's values by $H(h)$:

$$u_h^{\text{free}}(v) := \max_{q^{\text{alloc}} \in \mathcal{X}_h} v \cdot q^{\text{alloc}}, \qquad H(h) := \mathbb{E}[u_h^{\text{free}}(\mathbf{V})] = U(\mathcal{M}_h^{\text{free}}). \tag{D.6}$$

As in the unit-demand setting, we will choose parameter h so that mechanism $\mathcal{M}_h^{\text{free}}$'s utility is large enough relative to its allocation cost (which is at most h).

To compare mechanism C with the free-lottery mechanism $\mathcal{M}_h^{\text{free}}$, we want to use the fact that mechanism $\mathcal{M}_h^{\text{free}}$ maximizes utility among free lotteries with expected allocation cost at most h. To do this, we first modify mechanism C's allocation lottery to satisfy the same allocation cost bound. We do this simply by running that lottery less often if its cost exceeds h, and allocating nothing otherwise; its expected value decreases by the same factor as its cost.

Formally, fix a value vector v and consider the allocation lottery of mechanism C. If its expected cost $\lambda \cdot x(v)$ exceeds h, run it with probability $h/(\lambda \cdot x(v))$ and allocate nothing otherwise; if its cost is at most h, leave it unchanged. Call the resulting lottery $C_h(v)$. It still draws independent sets and has expected cost at most h, so maximality of mechanism $\mathcal{M}_h^{\text{free}}$ and nonnegative payments of mechanism C give

$$u_h^{\text{free}}(v) \geq \frac{v \cdot x(v)}{\max\{1, \lambda \cdot x(v)/h\}} \geq \frac{u(v)}{\max\{1, \lambda \cdot x(v)/h\}} \geq \frac{u(v)}{1 + \lambda \cdot x(v)/h}. \tag{D.7}$$

Using this inequality, we adapt the utility upper bound from Section 3. For independent draws $\mathbf{V}$ and $\mathbf{W}$ from the buyer's value distribution,

$$\begin{aligned}
U(C) - D &\leq \mathbb{E}\left[\sqrt{u(\mathbf{V})u(\mathbf{W})}\right] && \text{(3.1), (3.3)} \\
&\leq \frac{1}{2}\mathbb{E}\left[u_h^{\text{free}}(\mathbf{V})\left(1 + \frac{\lambda \cdot x(\mathbf{W})}{h}\right)\right] && \text{(D.7), AM–GM} \\
&\quad + \frac{1}{2}\mathbb{E}\left[u_h^{\text{free}}(\mathbf{W})\left(1 + \frac{\lambda \cdot x(\mathbf{V})}{h}\right)\right] \\
&= H(h)\mathbb{E}\left[1 + \frac{\lambda \cdot x(\mathbf{V})}{h}\right] && \text{independent, identically distributed draws} \\
&= H(h)\left(1 + \frac{\sum_j \lambda_j q_j^{\text{alloc}}(C)}{h}\right). && \text{definition of } q_j^{\text{alloc}}(C)
\end{aligned} \tag{D.8}$$

Note that bound (D.8) is exactly the same as (5.7) in Section 5.1.

For $0 < h < h'$, the same allocation-cost reduction used to construct lottery $C_h(v)$, applied to the lottery chosen under mechanism $\mathcal{M}_{h'}^{\text{free}}$, gives

$$\frac{h}{h'}H(h') \leq H(h) \leq H(h').$$

Thus H is continuous for $h > 0$, and $H(h)/h$ is nonincreasing and tends to zero as $h \to \infty$, since $H(h) \leq \sum_j \mathbb{E}[V_j] < \infty$. The same crossing and no-crossing arguments as in Section 5.1 therefore

apply: choose $h > 0$ with $H(h) = 2h$ if one exists, and otherwise allocate nothing. The resulting mechanism $\mathcal{M}^{\text{free}}$ satisfies

$$U(\mathcal{M}^{\text{free}}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}^{\text{free}}) \geq \frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\text{alloc}}(C). \tag{D.9}$$

In terms of computation, a difference compared to Section 5.1 is that evaluating the free lottery's expected utility $H(h)$ is no longer straightforward. The assumed *free-lottery utility oracle* evaluates $H(h)$ for any buyer, demand-weight vector, and budget $h > 0$.

The combined single-buyer mechanism. Let mechanism $\mathcal{M}$ run pricing mechanism $\mathcal{M}^{\mathcal{I}}$ with probability $e/(3e-2)$ and free-lottery mechanism $\mathcal{M}^{\text{free}}$ with probability $2(e-1)/(3e-2)$, independently of the buyer's report. By linearity of utility and allocation cost under this mixture, bounds (D.4) and (D.9) give

$$\begin{aligned} U(\mathcal{M}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}) &\geq \frac{e}{3e-2}\left[\left(1 - \frac{1}{e}\right) D - \sum_j \lambda_j q_j^{\text{alloc}}(C)\right] \\ &\quad + \frac{2(e-1)}{3e-2}\left[\frac{U(C) - D}{2} - \sum_j \lambda_j q_j^{\text{alloc}}(C)\right] \\ &= \frac{U(C)}{2 + e/(e-1)} - \sum_j \lambda_j q_j^{\text{alloc}}(C). \end{aligned} \tag{D.10}$$

This proves (D.2). Both sub-mechanisms allocate an independent set in every realization, so mechanism $\mathcal{M}$ does as well. Mechanism $\mathcal{M}$ has nonnegative payments and is ex-post individually rational: the pricing mechanism allocates only positive-surplus items, and the free-lottery mechanism charges nothing. The buyer chooses their preferred offer from a fixed menu in each sub-mechanism, so each is BIC; their mixture, drawn independently of the buyer's report, is therefore BIC as well. □

Corollary 3. *Under the assumptions of Lemma 13, for every comparison mechanism C that allocates a set in $\mathcal{I}$,*

$$\max_{\mathcal{M} \in \{\mathcal{M}^{\mathcal{I}}, \mathcal{M}^{\text{free}}\}} \left\{ U(\mathcal{M}) - \sum_j \lambda_j q_j^{\text{alloc}}(\mathcal{M}) \right\} \geq \frac{U(C)}{2 + e/(e-1)} - \sum_j \lambda_j q_j^{\text{alloc}}(C).$$

D.1.2 The Multi-Buyer Mechanism

As before, we have constructed a single-buyer mechanism with a utility guarantee for any fixed demand-weight vector λ (but now this holds for mechanisms which allocate independent sets in the buyer's matroid family $\mathcal{I}_i$). The remaining tools from Section 5 apply identically to the matroid setting, in part because Lemma 13 has the same guarantee for the matroid setting as Lemma 5 for the unit-demand setting.

For completeness, we define the multi-buyer mechanism $\overline{\mathcal{M}}^{n,m}_{\text{mat}}$. As in the unit-demand mechanism $\overline{\mathcal{M}}^{n,m}_{\text{het}}$, initialize demand weights to $\lambda^1 = 0$ and set step size $\eta = S/(n\sqrt{T})$ (where here $S = \sum_i \sum_{j:\{j\} \in \mathcal{I}_i} \mathbb{E}[V_{ij}] > 0$).[7] In each round t, use Lemma 13 with demand weights λ^t to com-

[7] If $S = 0$ then trivially, $\text{OPT}^{n,m} = 0$.

pute a single-buyer mechanism $\mathcal{M}_i^t$ for each buyer i. Let $r_j^t := \sum_i q_j^{\text{alloc}}(\mathcal{M}_i^t)$ be the expected number of requests for item j, and update the demand weights using rule

$$\lambda_j^{t+1} = \min\left\{\sum_i \sum_{k:\{k\}\in\mathcal{I}_i} \mathbb{E}[V_{ik}], \max\left\{0, \lambda_j^t + \eta(r_j^t - 1)\right\}\right\}.$$

The update rule differs compared to equation (5.10) only because our upper bound of S differs for the matroid constrained setting.

After T rounds, let $\bar{r}_j := T^{-1}\sum_{t=1}^T r_j^t$. Mechanism $\overline{\mathcal{M}}_{\text{mat}}^{n,m}$ independently draws a uniform round t_i for each buyer i and runs mechanism $\mathcal{M}_i^{t_i}$ with probability $1/\max\{1, \max_j \bar{r}_j\}$, otherwise making no offer to that buyer.

The following matroid counterpart of Lemma 6 gives the guarantee of mechanism $\overline{\mathcal{M}}_{\text{mat}}^{n,m}$, including errors from approximate computation. Its proof is given in Appendix D.2.4.

Lemma 15. *Using the demand weights $\lambda^1, \ldots, \lambda^T$ generated by the procedure above, for every $\varepsilon > 0$, after $T = O(n^4m^4/\varepsilon^2)$ rounds, $\overline{\mathcal{M}}_{\text{mat}}^{n,m}$ allocates a set in $\mathcal{I}_i$ to each buyer i and is ex-ante feasible. The mechanism's total expected utility satisfies*

$$U(\overline{\mathcal{M}}_{\text{mat}}^{n,m}) \geq \left(\frac{1}{2 + e/(e-1)} - \varepsilon\right)\text{OPT}^{n,m}. \tag{D.11}$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\text{OPT}^{1,1}$, matroid membership tests, a free-lottery utility oracle, and exact evaluation of the required expectations and item allocation probabilities.

We again apply the fair contention rule of Feige and Vondrak [22] item by item to the buyers' requested independent sets, converting mechanism $\overline{\mathcal{M}}_{\text{mat}}^{n,m}$ into the feasible mechanism $\mathcal{M}_{\text{mat}}^{n,m}$. As in Section 5.2, we adjust each request's service probability to exactly $1 - 1/e$ and charge posted prices only for delivered items.

To see why the same rule preserves buyers' choices, fix a buyer's entire requested set (including the set drawn by the free-lottery mechanism). Other buyers' requests remain independent of this conditioning, so each requested item is delivered with probability $1-1/e$, even under deviations, by the same contention argument as in Section 5.2. Since values and payments add across delivered items, linearity of expectation multiplies every option's expected utility by $1 - 1/e$, preserving utility-maximizing choices and BIC. Delivered priced items have nonnegative surplus, and free allocations charge nothing, preserving ex-post individual rationality and nonnegative payments. Each item is allocated at most once, and every delivered set remains independent by downward closure. Consequently, $U(\mathcal{M}_{\text{mat}}^{n,m}) = (1 - 1/e)U(\overline{\mathcal{M}}_{\text{mat}}^{n,m})$.

The following theorem gives the same approximation guarantee as Theorem 3 for matroid-constrained buyers; its proof is given in Appendix D.2.5.

Theorem 4. *In the matroid-constrained setting, for every $\varepsilon > 0$, the demand-weight procedure followed by the contention rule above gives a BIC, ex-post individually rational, feasible mechanism $\mathcal{M}_{\text{mat}}^{n,m}$ with nonnegative payments. Its expected utility satisfies*

$$U(\mathcal{M}_{\text{mat}}^{n,m}) \geq \frac{\text{OPT}^{n,m}}{2e/(e-1) + (e/(e-1))^2 + \varepsilon}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\text{OPT}^{1,1}$, matroid membership tests, a free-lottery utility oracle, and exact evaluation of the required expectations and item allocation probabilities.

D.2 Supporting Details for the Matroid-Constrained Construction

D.2.1 Combining Independent Prices

A lottery over independent sets is a distribution $\mathcal{D}$ on the family $\mathcal{I}$ of independent sets. Its *marginal vector* z records each item's inclusion probability: $z_j = \Pr_{T\sim\mathcal{D}}(j \in T)$. These individual probabilities do not specify which items are allocated together. The matroid polytope $\mathcal{X}$ in (2.1) consists exactly of the marginal vectors of such lotteries: every lottery over independent sets has its marginal vector in $\mathcal{X}$, and every vector in $\mathcal{X}$ is realized by such a lottery.

Lemma 14. *Let $V_1, \ldots, V_m$ be independent item values and let P_j be nonnegative, possibly randomized prices drawn independently of the values and one another, satisfying $\Pr(V_j > P_j) \le q_j$. Let $U(\mathcal{M}_j^1)$ be the expected utility from offering price P_j for item j alone. If $q \in \mathcal{X}$, then the mechanism $\mathcal{M}^{\mathcal{I}}$ that offers these prices together and lets the buyer choose a surplus-maximizing independent set of positive-surplus items allocates item j with probability at most q_j and has expected utility*

$$U(\mathcal{M}^{\mathcal{I}}) \ge (1 - 1/e) \sum_j U(\mathcal{M}_j^1).$$

Proof of Lemma 14. Let $\mathbf{P} = (P_1, \ldots, P_m)$ be the independent nonnegative posted prices. Consider first the single-item setting, where only item j is for sale. Define the surplus from posted price P_j as $S_j := (V_j - P_j)^+$. The expected utility of the single-item mechanism $\mathcal{M}_j^1$ is $U(\mathcal{M}_j^1) = \mathbb{E}_{V_j,P_j}[S_j]$.

Now consider the multi-item posted-price mechanism $\mathcal{M}^{\mathcal{I}}$ with the matroid constraint $\mathcal{I}$. The buyer chooses an independent set maximizing total surplus. Since mechanism $\mathcal{M}^{\mathcal{I}}$ allocates only positive-surplus items, its ex-ante allocation probability $q_j^{\text{alloc}}(\mathcal{M}^{\mathcal{I}})$ is bounded by the probability of sale in the single-item setting, and hence by the single-item cap q_j:

$$q_j^{\text{alloc}}(\mathcal{M}^{\mathcal{I}}) \le \Pr_{V_j,P_j}(S_j > 0) = \Pr_{V_j,P_j}(V_j > P_j) \le q_j.$$

Finally, we compare the utility of mechanism $\mathcal{M}^{\mathcal{I}}$ with the sum of the single-item utilities, one surplus threshold at a time, as in Lemma 3. At each threshold, the single-item benchmark counts all items whose surplus exceeds that threshold, while the matroid mechanism can select only an independent subset. We will show that the expected size of the largest independent subset of those items is at least a $1 - 1/e$ fraction of the expected number of items whose surplus exceeds that threshold. Integrating this comparison over thresholds will show that the utility of mechanism $\mathcal{M}^{\mathcal{I}}$ is at least the same fraction of the sum of the single-item utilities.

Fix a surplus threshold $t \ge 0$. Let $R_t := \{j : S_j > t\}$ be the random set of items whose surplus exceeds this threshold, and let $p_j(t) := \Pr_{V_j,P_j}(S_j > t)$ be the probability that item j belongs to random set R_t. Its rank $r(R_t)$ is the largest number of items in this set that can be allocated together, so we seek a lower bound on the expected rank $\mathbb{E}[r(R_t)]$.

To bound the expected rank of random set R_t, we compare it with a distribution $\mathcal{D}_t$ over independent sets having the same item inclusion probabilities. Such a distribution exists because $0 \le p(t) \le q \in \mathcal{X}$ and the polytope $\mathcal{X}$ is downward closed. Since the surpluses are independent, random set R_t includes items independently, so the correlation-gap bound of Yan [39, Lemma 4.1],

with all item weights set to one, gives

$$\begin{aligned}\mathbb{E}[r(R_t)] &\geq (1-1/e)\mathbb{E}_{A\sim\mathcal{D}_t}[r(A)] \\ &= (1-1/e)\mathbb{E}_{A\sim\mathcal{D}_t}[|A|] \\ &= (1-1/e)\mathbb{E}_{A\sim\mathcal{D}_t}\left[\sum_j \mathbf{1}\{j\in A\}\right] = (1-1/e)\sum_j p_j(t),\end{aligned} \tag{D.12}$$

where the first equality is because every set A drawn from distribution $\mathcal{D}_t$ is feasible, so its rank equals its size.

For every threshold $t \geq 0$, the buyer's surplus-maximizing independent set contains exactly $r(R_t)$ items with surplus above t (if not, the matroid exchange property would allow adding an item from R_t, possibly replacing an item with surplus at most t, and increasing total surplus). Integrating the expected number of chosen items with surplus above t over all thresholds gives the buyer's total surplus in expectation. Thus mechanism $\mathcal{M}^{\mathcal{I}}$ has utility

$$\begin{aligned}U(\mathcal{M}^{\mathcal{I}}) &= \mathbb{E}\left[\max_{T\in\mathcal{I}}\sum_{j\in T} S_j\right] = \int_0^\infty \mathbb{E}[r(R_t)]\,dt \\ &\geq (1-1/e)\int_0^\infty \sum_j p_j(t)\,dt \qquad \text{(D.12)} \\ &= (1-1/e)\sum_j \int_0^\infty \Pr_{V_j,P_j}(S_j > t)\,dt \\ &= (1-1/e)\sum_j U(\mathcal{M}_j^1).\end{aligned}$$

□

D.2.2 Computing the Caps

We now discuss how the greedy algorithm of Appendix C.2 gives us a solution to the matroid cap program (D.3) with an additive approximation guarantee. The objective is the same as in the unit-demand cap program (5.4), treated in Appendix C.2, but with the unit-demand constraint $\sum_j q_j \leq 1$ replaced by membership in the matroid polytope $\mathcal{X}$.

Let $d \leq m$ be the number of items available to the fixed buyer and write $\phi_j(q) := \mathrm{OPT}^{1,1}(V_j; q) - \frac{e}{e-1}\lambda_j q$ and $\Phi(q) := \sum_j \phi_j(q_j)$. Denote the optimal value of (D.3) by $\mathcal{R}_\lambda := \Phi^\star$ and its optimal caps by $q^\star$. As in Appendix C.2, we work in normalized units with $\sum_j \mathbb{E}[V_j] \leq 1$ and $0 \leq \lambda_j \leq 1$ for every item j.

For matroid-constrained buyers, we also assume efficient membership testing in each buyer's matroid polytope $\mathcal{X}_i$.[8] For $0 < \delta \leq 1$, choose $K := \lceil d(1 + e/(e-1))/\delta \rceil$ and run the incremental greedy algorithm of Appendix C.2, replacing the total-cap constraint by membership in $\mathcal{X}$: allow an increment only if the resulting cap vector remains in $\mathcal{X}$, and stop when no positive-gain feasible increment remains. Write ${q_j}^{\mathrm{G}} = k_j/K$ for the returned caps. The rounding argument must now

[8] This does not require stronger oracle access: membership in the matroid polytope can be tested in polynomial time using an independence oracle for the original matroid.

preserve all matroid rank constraints. Combining a bound on the loss from shrinking and rounding the caps with optimality of greedy on the feasible grid gives the following additive approximation guarantee.

Lemma 16. *Suppose* $\sum_j \mathbb{E}[V_j] \le 1$ *and* $0 \le \lambda_j \le 1$ *for every item* j*. For* $0 < \delta \le 1$*, the preceding greedy algorithm with* $K = \lceil d(1 + e/(e-1))/\delta \rceil$ *returns feasible caps* $q^{\mathrm{G}} \in \mathcal{X}$ *satisfying*

$$\sum_j \left[\mathrm{OPT}^{1,1}(V_j; {q_j}^{\mathrm{G}}) - \frac{e}{e-1} \lambda_j {q_j}^{\mathrm{G}} \right] \ge \mathcal{R}_\lambda - \delta.$$

The algorithm uses $\mathrm{poly}(m, 1/\delta)$ *oracle calls to* $\mathrm{OPT}^{1,1}$ *and matroid membership tests, and polynomial additional computation in* $m, 1/\delta$*.*

Proof. With the selected parameters, the optimization program has optimal value $\Phi^\star = \mathcal{R}_\lambda$ and the greedy algorithm objective is $\Phi_{\mathrm{G}} := \Phi(q^{\mathrm{G}})$. The returned caps q^{G} are feasible because the algorithm preserves membership in $\mathcal{X}$ at every increment. We first show that greedy solves the discretized problem optimally, then bound the loss from shrinking and rounding the optimal caps.

To prove optimality on the grid, make K copies of each item, representing its cap increments of $1/K$. Selecting k_j copies of item j represents cap $q_j = k_j/K$, and the weight of its kth copy is the corresponding gain in the objective $\phi_j(k/K) - \phi_j((k-1)/K)$. Declare a collection with copy counts k_j feasible exactly when $k/K \in \mathcal{X}$, or equivalently

$$\sum_{j \in A} k_j \le K r(A) \quad \text{for every item set } A.$$

These feasible collections form a matroid by the matroid union theorem [25, Section 6.5]. To apply the theorem, first consider collections containing at most one copy of each original item, with the corresponding original items forming an independent set; these collections themselves form a matroid. The theorem then shows that collections satisfying our cap constraints are exactly those that can be partitioned into K such independent collections, and that they too form a matroid.

The standard greedy algorithm maximizes total weight over the independent sets of a matroid. Our incremental cap algorithm coincides with this algorithm on the matroid defined above, with copy weights equal to cap-increment gains. By concavity, successive copy weights are nonincreasing, so we may take the selected copies of each item to be its first k_j copies, breaking equal-weight ties by copy index. The weights of the first k_j copies of item j sum to

$$\sum_{k=1}^{k_j} [\phi_j(k/K) - \phi_j((k-1)/K)] = \phi_j(k_j/K),$$

since $\phi_j(0) = 0$. Thus the total selected weight is $\Phi(k/K)$, and our algorithm maximizes Φ on the feasible grid.

It remains to compare the best feasible grid solution with the continuous optimum. We construct feasible grid caps from the optimal caps $q^\star$ by shrinking them before rounding upward: ${q_j}^\uparrow := \lceil (K-d) q_j^\star \rceil / K$. For every nonempty set A of available items,

$$\sum_{j \in A} {q_j}^\uparrow \le \left(1 - \frac{d}{K}\right) \sum_{j \in A} q_j^\star + \frac{|A|}{K} \le \left(1 - \frac{d}{K}\right) r(A) + \frac{d}{K} r(A) = r(A).$$

The second inequality uses $|A| \le d$ and $r(A) \ge 1$, since every available item can be allocated individually. Thus the rounded caps $q^{\uparrow}$ are feasible. Monotonicity and concavity of $\mathrm{OPT}^{1,1}$, together with $\lambda_j \le 1$, give

$$\Phi(q^{\uparrow}) \ge \left(1 - \frac{d}{K}\right) \mathcal{R}_\lambda - \frac{ed}{(e-1)K}.$$

Meanwhile, optimality of greedy on the feasible grid gives $\Phi_{\mathrm{G}} \ge \Phi(q^{\uparrow})$. Combining these two inequalities proves the desired inequality:

$$\begin{aligned} \mathcal{R}_\lambda - \Phi_{\mathrm{G}} &\le \frac{d}{K}\mathcal{R}_\lambda + \frac{ed}{(e-1)K} \\ &\le \frac{d(1 + e/(e-1))}{K} && \mathcal{R}_\lambda \le \sum_j \mathbb{E}[V_j] \le 1 \\ &\le \delta. && \text{choice of } K \end{aligned}$$

□

D.2.3 Optimality of the Free-Lottery Mechanism

Fix demand weights λ, a value vector v, and a budget $h > 0$. We show that mechanism $\mathcal{M}_h^{\mathrm{free}}$ attains the maximum utility among all free lotteries over independent sets with expected allocation cost $\sum_j \lambda_j q_j^{\mathrm{alloc}}$ at most h. It suffices to show that some utility-maximizing lottery either allocates one independent set deterministically or randomizes between two independent sets with allocation costs below and above h, respectively.

Write p_T for the probability that a lottery allocates independent set $T \in \mathcal{I}$, and recall that $\Lambda(T) = \sum_{j \in T} \lambda_j$ is its allocation cost. Since the lottery charges nothing, maximizing the buyer's expected utility subject to the allocation cost bound is the finite linear program

$$\begin{aligned} \max_{(p_T)_{T \in \mathcal{I}}} \quad & \sum_{T \in \mathcal{I}} p_T \sum_{j \in T} v_j \\ \text{s.t.} \quad & \sum_{T \in \mathcal{I}} p_T = 1, \qquad \sum_{T \in \mathcal{I}} p_T \Lambda(T) \le h, \\ & p_T \ge 0 \quad \text{for every } T \in \mathcal{I}. \end{aligned}$$

An optimal lottery exists, so we choose one assigning positive probability to the fewest sets. Suppose that this lottery uses at least three sets. We can always redistribute probability among three sets while preserving both total probability and expected allocation cost, and since utility is linear in these probabilities, there exists one redistribution direction that does not decrease utility. Following this direction until a probability becomes zero gives an optimal lottery using fewer sets, which is a contradiction. Thus an optimal lottery must use at most two sets.

Having established that at most two sets suffice, optimality of mechanism $\mathcal{M}_h^{\mathrm{free}}$ follows easily: the buyer (who maximizes their utility) chooses their preferred independent set deterministically or their preferred pair under the utility-maximizing randomization with expected allocation cost at most h.

D.2.4 Proof of the Demand-Weight Balancing Lemma

Lemma 15. *Using the demand weights $\lambda^1, \dots, \lambda^T$ generated by the procedure above, for every $\varepsilon > 0$, after $T = O(n^4 m^4/\varepsilon^2)$ rounds, $\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}}$ allocates a set in $\mathcal{I}_i$ to each buyer i and is ex-ante feasible. The mechanism's total expected utility satisfies*

$$U(\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}}) \geq \left(\frac{1}{2 + e/(e-1)} - \varepsilon\right) \mathrm{OPT}^{n,m}. \tag{D.11}$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\mathrm{OPT}^{1,1}$, matroid membership tests, a free-lottery utility oracle, and exact evaluation of the required expectations and item allocation probabilities.

Proof of Lemma 15. We apply the proof of Lemma 6 in Appendix C.4, only checking the points affected by matroid constraints. Use the same normalization by S and the same notation $U^t, r_j^t, \overline{U}, \overline{r}_j$ for round utilities, item requests, and their averages. The bounds used there still hold: $U^t \leq 1$, since utility is at most total allocated value; $0 \leq r_j^t \leq n$, since each buyer requests a given item at most once; and $\mathrm{OPT}^{n,m} \geq 1/(nm)$, since allocating the highest-expected-value buyer–item pair for free is feasible and every available item can be allocated individually. None of these bounds requires a buyer to request at most one item in total.

Each sampled single-buyer mechanism allocates an independent set, and making no offer instead preserves this property. The same participation reduction therefore ensures ex-ante item supply and gives, as in (C.5),

$$U(\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}}) = \frac{\overline{U}}{\max\{1, \max_j \overline{r}_j\}} \geq \overline{U} - \sum_j (\overline{r}_j - 1)^+.$$

As in the proof of Lemma 12, cap-optimization error δ and budget-search error γ contribute at most $\delta + \gamma$ to the single-buyer guarantee; the budget search uses exact evaluations from the free-lottery utility oracle. Averaging a feasible comparison mechanism over the other buyers' values still gives single-buyer mechanisms that allocate independent sets and preserve utilities and allocation probabilities. Consequently, the same round-by-round comparison and projected-update analysis give the same bound as Lemma 12 for the matroid construction:

$$\overline{U} - \sum_j (\overline{r}_j - 1)^+ \geq \frac{\mathrm{OPT}^{n,m}}{2 + e/(e-1)} - \frac{nm}{\sqrt{T}} - n(\delta + \gamma).$$

Choosing the same number of rounds and error tolerances as in the proof of Lemma 6,

$$T := \left\lceil \frac{4n^4 m^4}{\varepsilon^2} \right\rceil, \qquad \delta + \gamma \leq \frac{\varepsilon}{2n^2 m},$$

gives the claimed approximation:

$$U(\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}}) \geq \frac{\mathrm{OPT}^{n,m}}{2 + e/(e-1)} - \frac{\varepsilon}{nm} \geq \left(\frac{1}{2 + e/(e-1)} - \varepsilon\right) \mathrm{OPT}^{n,m}.$$

To compute the single-buyer mechanisms in each round, use the matroid cap algorithm of Appendix D.2.2. For budget bisection, use the interval $[0, 1]$, since $H(h) \leq \sum_j \mathbb{E}[V_j] \leq 1$ in normalized units. Under the stated access assumptions, the mechanism can be computed in time polynomial in $n, m, 1/\varepsilon$. □

D.2.5 Proof of the Matroid Approximation Theorem

Theorem 4. *In the matroid-constrained setting, for every $\varepsilon > 0$, the demand-weight procedure followed by the contention rule above gives a BIC, ex-post individually rational, feasible mechanism $\mathcal{M}^{n,m}_{\mathrm{mat}}$ with nonnegative payments. Its expected utility satisfies*

$$U(\mathcal{M}^{n,m}_{\mathrm{mat}}) \geq \frac{\mathrm{OPT}^{n,m}}{2e/(e-1) + (e/(e-1))^2 + \varepsilon}.$$

The mechanism can be computed using polynomially many oracle calls and polynomial additional computation in $n, m, 1/\varepsilon$, given oracle access to $\mathrm{OPT}^{1,1}$, matroid membership tests, a free-lottery utility oracle, and exact evaluation of the required expectations and item allocation probabilities.

Proof of Theorem 4. The contention-resolution argument in Section D.1.2 establishes feasibility, BIC, ex-post individual rationality, nonnegative payments, and $U(\mathcal{M}^{n,m}_{\mathrm{mat}}) = (1-1/e)U(\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}})$.

It suffices to consider $0 < \varepsilon \leq 1$; for larger ε, use the mechanism obtained with $\varepsilon = 1$. Choose the accuracy in Lemma 15 to be

$$\theta := \frac{1}{2 + e/(e-1)} - \frac{1}{2 + e/(e-1) + (1-1/e)\varepsilon} > 0.$$

Then Lemma 15 and the utility scaling proved above give

$$U(\mathcal{M}^{n,m}_{\mathrm{mat}}) = (1-1/e)U(\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}}) \geq \frac{(1-1/e)\mathrm{OPT}^{n,m}}{2 + e/(e-1) + (1-1/e)\varepsilon} = \frac{\mathrm{OPT}^{n,m}}{2e/(e-1) + (e/(e-1))^2 + \varepsilon}.$$

For $0 < \varepsilon \leq 1$, the accuracy $\theta = \Theta(\varepsilon)$, so the computation of mechanism $\overline{\mathcal{M}}^{n,m}_{\mathrm{mat}}$ is polynomial in $n, m, 1/\varepsilon$ under the access assumptions in Lemma 15. As in Section C.5, we assume efficient, exact evaluation of the required expectations and probabilities, including each buyer's item request probabilities under the matroid single-buyer mechanisms. Averaging over rounds, scaling participation, and computing the contention probabilities c_j then proceed exactly as in Section C.5. □