On the Menu Complexity of Consumer Utility Maximization*

Kira Goldner† Divyarthi Mohan‡ Anh Phung§

November 3, 2026

Abstract

We study the structure and menu complexity of mechanisms that maximize consumer utility—the buyers' value for their allocations minus their payments—providing the first menu-complexity lower bounds for this objective. Together, our bounds show that a utility-optimal mechanism may require infinite menu size under both correlated and independent distributions. To study the complexity of multi-buyer mechanisms through a single-buyer notion, we use the ex-ante relaxation to set constraints on expected allocations. We give explicit two-item instances in which every utility-optimal mechanism requires an infinite menu, under both correlated and independent values. For finite correlated distributions, we obtain a tight bound: for every $T \geq 3$, there is a distribution over T types and a pair of ex-ante caps for which every optimum requires T non-null menu options. These results hold for both unit-demand and additive buyers. Our lower bounds also extend to settings with multiple unit-demand buyers. In particular, an infinite menu can be necessary for optimality even with three buyers and two items where values are independent across buyers and items.

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

†Boston University; goldner@bu.edu.

‡Columbia University; divyarthi.m@columbia.edu

§Boston University; anhnp@bu.edu

1 Introduction

Consider social services that may be provided without monetary payments, yet recipients still face costs: waiting for medical treatment, completing paperwork for food assistance, or joining waiting lists for housing. These *ordeals* can help ration demand and elicit private information in order to allocate more efficiently, but they impose costs on recipients without benefiting the provider. Thus, the designer must balance improving the allocation with minimizing the ordeals required to do so.

This tradeoff is captured by *consumer utility maximization*: the objective is the buyers' value for the goods they receive, less the payments they make. In this paper, we study the structure and complexity of mechanisms that maximize this objective. For single-parameter settings, Hartline and Roughgarden [27] completely characterize optimal mechanisms. With multiple items, however, the structure of utility-optimal mechanisms is much less understood.

The corresponding question for revenue has been a central problem in mechanism design since Myerson's characterization of optimal single-item auctions [33]. In the multi-item setting, finding a revenue-optimal mechanism has been a longstanding challenge. Here, optimal mechanisms can be computationally intractable or prohibitively complex even for a single additive buyer with independent item values [12, 13]. In particular, one difficulty concerns *menu complexity* [24]: a truthful mechanism offers the buyer options specifying allocation probabilities and a price, and lets them choose their favorite. Menu complexity counts the distinct options the buyer can receive, excluding the option of receiving nothing and paying nothing. Even with just two independent item values, every revenue-optimal mechanism can require infinitely many such options [13]. Understanding optimal mechanisms therefore includes determining when such complexity is necessary and when simpler mechanisms suffice. Even for just two items and one buyer, a long line of work has addressed these questions by characterizing optimal mechanisms for particular value distributions [30, 32, 31, 34, 18] and giving conditions under which a proposed mechanism is optimal [23, 13].

For consumer utility, recent work gives reason to expect similarly complicated optimal mechanisms. Goldner and Lundy [19] show that an optimal mechanism may leave an item unallocated to reduce the payments needed for truthfulness. They also show that familiar upper bounds from revenue maximization do not directly extend to utility. These observations leave open a basic question about the optimal mechanisms themselves: can utility maximization require arbitrarily large, or even infinite, menus? We investigate this question for two items, under both correlated and independent values.

One major obstacle, however, is that the single-buyer problem that has been so useful for studying optimal revenue mechanisms is trivial for utility: offer every item for free and let the buyer choose their favorite feasible allocation. It is the multi-buyer competition for the same items that might require the designer to steer a buyer toward an item with less competing demand, using menu options with different allocation probabilities and payments to induce this choice. To study this through the single-buyer notion of menu complexity, we impose *ex-ante supply constraints*: item j may be allocated with probability at most α_j, in expectation over the buyer's values and the mechanism's randomization. Allocating an item more often to one type from the prior distribution then leaves less allocation probability for other types, and truthfulness restricts how the designer can make these choices.

These constraints arise naturally from the ex-ante relaxation of a multi-buyer problem [1, 7]. We relax the requirement that allocations be feasible (allocate each item at most once) in every outcome, assign each buyer a cap on their expected allocation of each item, and optimize their mechanisms separately; this gives an upper bound on optimal multi-buyer utility. To obtain a

lower bound, we must choose both a value distribution and caps for which every utility-optimal mechanism requires a large menu. For our constructions, the resulting lower bounds also extend to the multi-buyer problem: adding two deterministic buyers makes our chosen caps optimal and eliminates the relaxation gap. Every optimal joint mechanism then requires the same menu lower bound for the original buyer (Appendix C). This extends our menu-complexity lower bounds to utility-optimal multi-buyer mechanisms.

1.1 Our Results

We show that utility-optimal mechanisms under ex-ante supply constraints can require arbitrarily large, and even infinite, menus with just two items. Our results hold for both unit-demand and additive buyers. Throughout, we require truthfulness, individual rationality, and nonnegative payments.

- **A tight bound for finite correlated distributions.** For every $T \geq 3$, there is a distribution over T types and a pair of ex-ante caps for which every utility-optimal mechanism assigns a distinct non-null option to each type (Theorem 4.1). The required menu size is exactly T, matching the upper bound obtained by assigning one option to each type.
- **An infinite menu for correlated values.** Extending this construction to a distribution with a continuous component gives a single instance for which every utility-optimal mechanism has infinite menu complexity (Corollary 4.5).
- **An infinite menu for independent values.** There are independent distributions over the two item values and a pair of ex-ante caps for which every utility-optimal mechanism has infinite menu complexity (Theorem 5.1). In our construction, the second item's value is uniform on $(0, 1)$.

Our constructions also yield menu-complexity lower bounds for utility-optimal mechanisms with three buyers and two items, via the reduction in Appendix C. Observe that any lower bounds for the two-item setting immediately imply the same bounds for the more general n-item setting.[1]

Our main result on independent values shows that correlation is not necessary for infinite menu complexity, highlighting that instead, the multidimensionality of the values is the reason for complexity. Even when the buyer's two values are drawn independently, maximizing utility can require infinitely many different allocation probabilities and prices.

1.2 Our Approach

Showing that one complicated mechanism is optimal does not establish a menu lower bound; a simpler mechanism might achieve the same utility. We therefore construct instances in which optimality forces specific allocation probabilities and prices. In both constructions, the candidate menu offers lotteries for item 1 and a free option for item 2. A buyer considering an item-1 lottery can always choose item 2 for free instead. We use this comparison to determine the utility that the item-1 options must provide.

[1] Simply extend our instances to have $n - 2$ zero-valued items.

Deferring choice of the ex-ante caps. We must choose both a distribution and supply caps for which a large menu is necessary. Rather than choose the caps first, we assign a nonnegative cost λ_j to allocating item j and maximize expected utility minus expected allocation cost:

$$L_{\lambda_1,\lambda_2}(M) = \mathbb{E}\big[u_M(v_1, v_2) - \lambda_1 x_1^M(v_1, v_2) - \lambda_2 x_2^M(v_1, v_2)\big].$$

These costs are part of the analysis, not payments charged to the buyer. Once we find a maximizing mechanism M^*, we set the caps to its expected allocations. A standard Lagrangian argument (Lemma 3.1) shows that M^* is utility-optimal under these caps and that every utility-optimal mechanism subject to these caps also maximizes L. It therefore suffices to show that every maximizer of L requires a large menu.

Correlated values. We construct the instance in two steps. First, we choose the possible types but leave their probabilities unspecified. We specify properties that we want the optimal menu to satisfy and show that, if these properties hold, the menu must take a particular form. We use correlation to choose $v_2 = \rho(v_1)$, so that the utility from receiving item 2 for free is a strictly convex function of v_1. Each item-1 lottery, in contrast, gives utility that is affine in v_1. We require each type to be indifferent between their assigned lottery and free item 2, and impose indifference between consecutive lottery options. Geometrically, these indifferences pin down two points on each lottery's utility line, determining the option (see Figure 1). Strict convexity then forces distinct lotteries for different types. We characterize the resulting menu.

Second, we choose the probabilities of the types to make this menu uniquely optimal. We choose these probabilities so that attaining an upper bound on L requires precisely the menu properties used in the first step. At a high level, this follows because optimality forces certain IC constraints to bind (creating indifferences). The menu we constructed attains this bound, and every optimum must therefore have the form already determined by those properties. Thus, optimality forces a distinct option for every type. A continuous version of the construction uses the same geometric idea to force infinitely many distinct options.

Independent values. The correlated construction serves as a warm-up, using the relationship between the two item values to force distinct menu options. With independent values, we instead start from the single-item problem, which admits a simple utility-optimal menu even under an ex-ante cap. We show how adding a second item as a free fallback can make the optimal menu for the first item require infinitely many lotteries. The fallback itself is simple: its value is uniform on $(0, 1)$, independently of the first-item value. The buyer's first-item value determines their favorite item-1 lottery, but they now compare its utility to the value of taking item 2 for free. Our task is then to construct a first-item distribution for which this additional choice makes an infinite menu necessary.

We first identify conditions on the first-item distribution under which any menu can be replaced by item-1 lotteries and a single free lottery for item 2 without decreasing L. Splitting an option into single-item lotteries can reduce utility, but these conditions ensure that the savings in expected allocation costs compensate for the loss (Lemma 5.2). The uniform second-item distribution then lets us replace its options with one free lottery that gives at least as much utility without increasing either item's expected allocation (Lemma 5.4). These transformations also preserve the utility provided when $v_2 = 0$.

For menus of this form, the objective depends only on the item-1 utility curve and the item-2 free lottery's allocation probability. Integrating over the second-item value and using integration by parts lets us bound L by optimizing the utility pointwise at each first-item value. We work backwards from an implementable single-item utility curve that is strictly convex on an interval and choose a first-item density so that attaining the bound requires matching this curve on that interval. This density also satisfies the conditions needed for the menu reduction. Our candidate infinite menu attains the bound, and every optimal menu of the restricted form must provide the chosen curve on that interval. Because the reduction preserves utility when $v_2 = 0$, every optimal menu must reproduce the curve there. But a finite menu gives the maximum of finitely many affine functions along this axis, so it cannot produce such a curve.

Beyond unit-demand buyers. Our candidate menus allocate at most one item, but their optimality proofs continue to hold when the buyer is additive and competing mechanisms may allocate both. The bounds use only that each item's allocation probability is at most one. Thus, the same instances give the additive lower bounds (Corollaries 4.6 and 5.7). They also give lower bounds for broader classes containing these valuations, including downward-closed constrained-additive and subadditive valuations. For a fixed downward-closed feasibility constraint, any two individually feasible items give either the additive case, if they can be allocated together, or the unit-demand case, if they cannot.

1.3 Related Work

Consumer utility. Optimal mechanisms have been characterized for money burning and costly signaling in the single-parameter settings [27, 6, 10]. In a related healthcare allocation model, finding utility-maximizing equilibrium assignments and waiting times under a treatment budget is NP-hard [5]. For multidimensional utility, much of the literature studies approximation guarantees relative to optimal social welfare [4, 17, 19, 16, 15]. Goldner et al. [20] instead obtain constant-factor approximations to optimal utility for independent item values, using posted prices and free lotteries via the ex-ante relaxation. We study the structure of optimal mechanisms for these ex-ante-constrained single-buyer problems and show instances in which every utility-optimal mechanism requires an infinite menu, motivating simple, approximately-optimal mechanism design.

Menu complexity and revenue. Optimal mechanisms for revenue maximization can require infinite menus, even for a single additive buyer with two independent item values [24, 13]. For a unit-demand buyer, optimal revenue can require exponentially many lotteries in the number of items, even with independent binary values [9]. However, $(1-\varepsilon)$-approximation is possible with menu size depending only on the number of items and ε, regardless of support sizes for independent subadditive values [3, 29]. This dependence on ε is necessary: even for two bounded independent values, additive error ε can require menu size polynomial in $1/\varepsilon$ [21]. For constant-factor approximation, the better of selling separately and bundling suffices for independent additive and even subadditive buyers [2, 35]. In contrast, under correlation, no fixed menu-size bound guarantees any positive constant approximation, even for two additive items [24]. Our lower bounds establish that exact utility maximization can require infinite menus with two items under both independent and correlated values.

Other settings and measures of complexity. Menu complexity has also been used to characterize optimal mechanisms in more structured settings, such as those with totally ordered outcomes [36] or partially ordered outcomes [14]. Other work changes how complexity is measured: even selling separately has an exponential menu, motivating measures such as symmetric menu complexity, which counts options up to item permutations [29]. Under buy-many mechanisms, which allow the buyer to make repeated purchases, bounded menus suffice for near-optimal revenue even with correlated values [8]. We study the standard setting in which the buyer chooses one option and complexity counts the distinct options they may receive.

Related techniques. Existing menu-complexity bounds for revenue do not directly yield corresponding bounds for consumer utility, because a small menu may extract much less revenue while leaving the buyer more utility or vice versa. Although utility equals welfare minus revenue, changing allocation rules affects both welfare and payments, and ex-ante constraints do not fix this. The ex-ante constraints also introduce a difficulty when modifying menus. Modifications that preserve revenue may reduce utility or cause an item's expected allocation to exceed its ex-ante cap.

The representation of truthful mechanisms through convex utility functions connects menu complexity to the geometry of utility. The convex-utility representation, together with integration by parts and optimal-transport duality, has been used to characterize revenue-optimal mechanisms and prove menu-size lower bounds [11, 21]. Our independent-values construction also uses convex utility and integration by parts, but relies on a reduction from arbitrary to structured menus and direct bounds on the Lagrangian objective. Our correlated lower bound connects to revenue-maximization results with costly participation [28, 22], where a convex, type-dependent outside option can force menus as large as the support. We first show that, for our construction, the second item can be offered for free without loss, thereby serving as a private outside option. While costly-participation mechanisms use distinct menu options to extract revenue, ours use them to conserve the limited supply of the outside-option item while maximizing buyer utility. We use binding incentive constraints to force distinct options, an approach also used in lower bounds for partially ordered interests [14]. In both constructions, we choose allocation costs and then supply caps so that every utility-optimal mechanism must satisfy the properties that force a large menu.

1.4 Organization

Section 2 gives the model and menu representation, and Section 3 introduces the modified Lagrangian and its relation to supply caps. Sections 4 and 5 give the correlated and independent constructions, respectively. The appendices contain the deferred optimality proofs, the continuous correlated construction, and the extension to the multi-buyer setting.

2 Preliminaries

We consider a mechanism designer in a Bayesian setting who has two items to sell to a single buyer. The buyer's private type or valuation (v_1, v_2) is drawn from a known prior distribution $\mathcal{D}$ with $\text{support}(\mathcal{D}) \subseteq \mathbb{R}^2_{\geq 0}$, where v_1 and v_2 denote the buyer's values for each of the two items respectively. We consider a unit-demand buyer:[2] their valuation for the bundle of both items $\{1, 2\}$

[2]For the general n-item setting, $v(B) = \max_{j \in B} v_j$ for any bundle $B \subseteq [n]$.

is just $\max\{v_1, v_2\}$. (Our results also extend to the additive buyer setting.[3])

A mechanism M is defined by an allocation rule (x_1^M, x_2^M) and a payment rule p^M. A buyer with type (v_1, v_2) receives items 1 and 2 with probabilities $x_1^M(v_1, v_2)$ and $x_2^M(v_1, v_2)$ respectively at a price $p^M(v_1, v_2)$. For the unit-demand setting, we restrict attention to mechanisms that allocate at most one item, i.e., $x_1^M(v_1, v_2) + x_2^M(v_1, v_2) \leq 1$. The buyer's utility from mechanism M is:

$$u_M(v_1, v_2) := v_1 \cdot x_1^M(v_1, v_2) + v_2 \cdot x_2^M(v_1, v_2) - p^M(v_1, v_2).$$

We consider truthful mechanisms that are incentive compatible (IC) and individually rational (IR) with nonnegative payments.

- **Incentive Compatible (IC).** The buyer's utility is maximized by truthfully reporting their type (v_1, v_2) to the mechanism:

$$u_M(v_1, v_2) \geq v_1 \cdot x_1^M(v_1', v_2') + v_2 \cdot x_2^M(v_1', v_2') - p^M(v_1', v_2').$$

- **Individually Rational (IR).** The buyer's utility from M is always nonnegative:

$$u_M(v_1, v_2) \geq 0.$$

Incentive compatibility of a mechanism M connects the buyer's utility u_M to the allocation rule (x_1^M, x_2^M) through the envelope theorem: the gradient of the utility function equals the allocations. This is a generalization of Myerson's payment identity [33] to the multidimensional setting. For completeness, we restate the version in Hartline [26, Theorem 8.8.1].

Theorem 2.1 (The envelope theorem)**.** *For a buyer with linear utility, an allocation rule x and a utility function u define an incentive-compatible mechanism if and only if*

(i) u is convex; and

(ii) $x(v) = \nabla u(v)$.[4]

Our objective is to maximize the buyer's expected utility. However, this is trivial in the single-buyer setting without any supply constraints: simply offer the items for free. We consider *ex-ante supply constraints*, or caps for short, that restrict the expected allocation of the mechanism. A mechanism M subject to caps or ex-ante supply constraints $(\alpha_1, \alpha_2) \in [0, 1]^2$ must satisfy:

$$\mathbb{E}_{(v_1, v_2) \sim \mathcal{D}} \left[x_j^M(v_1, v_2) \right] \leq \alpha_j \qquad \forall \text{ items } j \in \{1, 2\}.$$

2.1 The Menu Representation of a Mechanism

It is well-known by the taxation principle [26, Section 8.2] that any truthful single-buyer mechanism is equivalent to a menu of (allocation, payment) options, from which the buyer can pick any one option. In particular, a menu option/lottery is a tuple (q_1, q_2, p) where the buyer pays a price $p \geq 0$ and gets the items 1 and 2 with probabilities $q_1 \in [0, 1]$ and $q_2 \in [0, 1]$ respectively.

[3] The buyer's valuation for a bundle of items B is the sum $v(B) = \sum_{j \in B} v_j$.

[4] The gradient of a convex function exists almost everywhere. At a type v where u is not differentiable, the condition is $x(v) \in \partial u(v)$: the allocation must be a subgradient of u at v. Here, $\partial u(v)$ denotes the subdifferential of u at v. Equivalently, for every type w, $u(w) \geq u(v) + x(v) \cdot (w - v)$.

A truthful mechanism M induces a menu as follows:

$$\left\{(x_1^M(v_1,v_2), x_2^M(v_1,v_2), p^M(v_1,v_2)) : (v_1,v_2) \in \text{support}(\mathcal{D})\right\},$$

and conversely, any menu M with the null option $(0,0,0)$ and appropriate tie-breaking specifies an IC and IR mechanism: for each buyer type simply give their utility-maximizing option from the menu.[5] Hence we use menu and truthful mechanism interchangeably.

Menu complexity [24]. The menu complexity or menu size of a mechanism M is the number of distinct (non-null) menu options offered by M. We observe that the menu complexity of any mechanism can be trivially upper bounded by $|\text{support}(\mathcal{D})|$.

Structured Menus. A key step in our proofs is a reduction to menus with certain structural properties. We first define a menu option (q_1, q_2, p) to be either a single-item lottery or a two-item lottery:

- A single-item lottery is one that exclusively sells either the first item or the second one, never both. That is, $q_1 = 0$ or $q_2 = 0$.
- A two-item lottery is any lottery that is not single-item: it allocates each item with positive probability: $q_1, q_2 > 0$.

We call a menu *structured* if it satisfies the following properties:

(S1) It only contains single-item lotteries.

(S2) There is only a single option selling the second item, and it is free.

3 A New Objective: The Modified Lagrangian

We introduce a modified objective that incorporates the ex-ante supply constraints. To construct a lower bound, we must choose ex-ante constraints (alongside a prior) for which every utility-optimal mechanism requires a large menu. By modifying our objective, we defer the choice of the ex-ante constraints, and instead assign a cost to allocating each item—the Lagrangian multipliers $\vec{\lambda}$. Our modified Lagrangian objective is then the buyer's expected utility minus the expected allocation cost. Once we find an optimal mechanism, we set the ex-ante caps to its expected allocations. In this sense, we are offloading the choice of supply constraints to the choice of allocation costs.

More specifically, instead of subjecting the mechanism to hard supply constraints, we penalize it for each unit of supply allocated. The nonnegative costs are λ_1 for item 1 and λ_2 for item 2. We define:

$$L_{\lambda_1,\lambda_2}(M) = \mathbb{E}_{(v_1,v_2)\sim\mathcal{D}}[u_M(v_1,v_2) - \lambda_1 x_1^M(v_1,v_2) - \lambda_2\, x_2^M(v_1,v_2)].$$

This objective comes from a Lagrangian relaxation of the ex-ante supply constraints. Given ex-ante constraints (α_1, α_2), our mechanism design problem maximizes the buyer's expected utility

[5] Without loss remove any (non-null) menu option that is not picked by any type in the support.

over all IC, IR mechanisms with nonnegative payments, subject to $\mathbb{E}[x_i^M(v_1, v_2)] \leq \alpha_i$ for each item i. Introducing the nonnegative multipliers (λ_1, λ_2) for each of these constraints gives the Lagrangian

$$\begin{aligned}\mathcal{L}(M, \vec{\lambda}) &= \mathbb{E}[u_M(v_1, v_2)] - \lambda_1(\mathbb{E}[x_1^M(v_1, v_2)] - \alpha_1) - \lambda_2(\mathbb{E}[x_2^M(v_1, v_2)] - \alpha_2) \\ &= L_{\lambda_1,\lambda_2}(M) + \lambda_1\alpha_1 + \lambda_2\alpha_2,\end{aligned}$$

where the λ's and α's are fixed. Thus, maximizing the Lagrangian $\mathcal{L}(M, \vec{\lambda})$ is equivalent to maximizing the modified Lagrangian objective $L_{\lambda_1,\lambda_2}(M)$.

For fixed costs, this objective depends only on M and the prior $\mathcal{D}$, but not on the ex-ante constraints. The following lemma shows how a maximizer of L_{λ_1,λ_2} determines caps under which it is utility-optimal.

Lemma 3.1 (Lagrangian Caps)**.** *Suppose that M^* attains a finite maximum of L_{λ_1,λ_2} for $\lambda_1, \lambda_2 > 0$, and set $\alpha_i = \mathbb{E}_{(v_1,v_2)\sim\mathcal{D}}[x_i^{M^*}(v_1, v_2)]$ for $i = 1, 2$. Then M^* maximizes expected buyer utility subject to the ex-ante constraints (α_1, α_2). Moreover, every utility-maximizing mechanism subject to these constraints also maximizes L_{λ_1,λ_2}.*

Proof. Consider any mechanism M satisfying these caps. By the optimality of M^* and the feasibility of M,

$$\begin{aligned}\mathbb{E}_{(v_1,v_2)\sim\mathcal{D}}[u_M(v_1, v_2)] &= L_{\lambda_1,\lambda_2}(M) + \lambda_1\,\mathbb{E}[x_1^M(v_1, v_2)] + \lambda_2\,\mathbb{E}[x_2^M(v_1, v_2)] \\ &\leq L_{\lambda_1,\lambda_2}(M^*) + \lambda_1\,\alpha_1 + \lambda_2\,\alpha_2 \\ &= \mathbb{E}[u_{M^*}(v_1, v_2)], \qquad (1)\end{aligned}$$

which means that M^* is also a utility-maximizing mechanism for the prior $\mathcal{D}$ under this particular choice of α.

Now consider any other utility-maximizing mechanism M that satisfies the caps induced by M^*. We have:

$$\begin{aligned}L_{\lambda_1,\lambda_2}(M) &= \mathbb{E}[u_M(v_1, v_2)] - \lambda_1\mathbb{E}[x_1^M(v_1, v_2)] - \lambda_2\,\mathbb{E}[x_2^M(v_1, v_2)] \\ &\geq \mathbb{E}[u_{M^*}(v_1, v_2)] - \lambda_1\alpha_1 - \lambda_2\alpha_2 \\ &= L_{\lambda_1,\lambda_2}(M^*)\end{aligned}$$

and since M^* maximizes L_{λ_1,λ_2}, we must have $L_{\lambda_1,\lambda_2}(M) = L_{\lambda_1,\lambda_2}(M^*)$, so M is another maximizer of L_{λ_1,λ_2}.

Thus, we can conclude that every utility-maximizing mechanism under the caps induced by the mechanism M^* must exhaust both the caps, and obtain the same maximal value of L_{λ_1,λ_2}. □

Remark 1. Our lower bounds require choosing both the prior distribution *and* the ex-ante supply constraints carefully. Neither can be arbitrary; we cannot hope for a stronger statement such as:

- "There exists an independent prior $\mathcal{D}_1 \times \mathcal{D}_2$ where any ex-ante constraints would yield infinite menu size for all optimal mechanisms." (This is false; for all priors, when $\alpha_1 = \alpha_2 = 1$, giving away either item for free is OPT, and has menu size of 2.)
- "There exist ex-ante constraints (α_1, α_2) where no matter what the prior is, the optimal mechanism would require infinitely many options." (This is false; for any choice of α_i's, when $\mathcal{D}_1 \times \mathcal{D}_2$ is a point mass distribution, a trivial free lottery $(\alpha_1, \alpha_2, 0)$ is optimal.[6])

[6] If $\alpha_1 + \alpha_2 > 1$, then either $\{(\alpha_1, 1 - \alpha_1, 0)\}$ or $\{(1 - \alpha_2, \alpha_2, 0)\}$ is the optimal unit-demand menu.

The Lagrangian Caps Lemma will play a critical role in our lower-bound construction. It allows us to focus on constructing a prior distribution by providing a set recipe for the appropriate ex-ante supply constraints via the Lagrangian costs. In the rest of the paper, we normalize $\lambda_1 = 1$ and set $\lambda_2 = c$ where $c \in (0, 1)$ is a constant that we choose appropriately as required by our construction.[7]

4 Lower Bounds for Correlated Items

In this section, we show that even with just two items, the trivial upper bound on the menu complexity is tight when the item values can be correlated. In particular, for any support size $T \geq 3$, there is an instance over T types such that any utility-optimal mechanism has menu size at least T. We also provide an extension to an instance with continuous support that requires infinite menu size.

Theorem 4.1. *For all $T \geq 3$, there exists a correlated distribution $\mathcal{D}$ over two items' valuations with support size T and ex-ante constraints (α_1, α_2) for which every utility-optimal mechanism requires at least T non-null menu options.*

We use the modified Lagrangian objective $L = L_{1,c}$ with $c = 1/2$. We first choose the support of the distribution and identify structural properties that uniquely determine a menu with a distinct option for every type (Section 4.1). We then choose probabilities over this support, defining a prior distribution, so that every maximizer of L must satisfy these structural properties (Section 4.2). The Lagrangian Caps Lemma (Lemma 3.1) supplies the corresponding ex-ante constraints. Finally, we extend the construction to a distribution with infinitely many types (Section 4.3).

4.1 Buyer Types and Menu Properties

We now formally define the support of our prior distribution and introduce the desirable properties that, if an optimal menu satisfies them, provide our lower bound. In Lemma 4.2, we show that there is a unique (structured) menu M^* that satisfies these properties.

We start by defining a support of size $T \geq 3$ for our prior over (v_1, v_2). There will be an *outside type* o who only wants the second item, $(v_1^o, v_2^o) = (0, 1)$, and $S = T - 1$ *inside types* who prefer item 1 over item 2. In particular, each inside type $i \in [S]$ has values (v_1^i, v_2^i) where:

- $v_1^i = i/S$ for item 1, and
- $v_2^i = \rho(v_1^i) = i^2/4S^2$ for item 2, where $\rho(v) = v^2/4$ is our choice of correlation.

For a truthful mechanism M, write $M(i)$ for type i's chosen menu option, which must be one of their favorites among all options. We want any optimal menu M to have the following properties:

(C1) The outside type o gets item 2 for free, i.e., $M(o) = (0, 1, 0)$.

(C2) Each inside type receives a lottery for only item 1, i.e., $M(i) = (q_1^i, 0, p^i)$ for all $i \in [S]$.

(C3) The first indifference: each inside type $i \in [S - 1]$ is indifferent between their own lottery $M(i) = (q_1^i, 0, p^i)$ and $M(i + 1) = (q_1^{i+1}, 0, p^{i+1})$. Furthermore, $p^1 = 0$.

[7] The normalization of λ_1 is without loss by setting $c = \lambda_2/\lambda_1$ and scaling the values by λ_1.

(C4) The second indifference: each inside type $i \in [S]$ is indifferent between their own lottery $M(i) = (q_1^i, 0, p^i)$ and the free item 2 option $M(o) = (0, 1, 0)$.

We observe that properties (C1) and (C2) induce a structured menu. These properties determine the outside type's option, leaving only the item-1 allocation probability q_1^i and payment p^i to be specified for each inside type. Moreover, we see that properties (C3) and (C4) provide linear constraints that help us to precisely derive the allocation probabilities and payments of our desired menu M.

Fix an item-1 option $M(i) = (q_1^i, 0, p^i)$. A buyer with first-item value x gets utility $q_1^i x - p^i$ from this option. As x varies, this utility traces out a line with slope q_1^i and intercept $-p^i$. This is the blue line in Figure 1. Thus, to determine the allocation probability and price, it suffices to determine the utility this same option gives at two distinct first-item values.

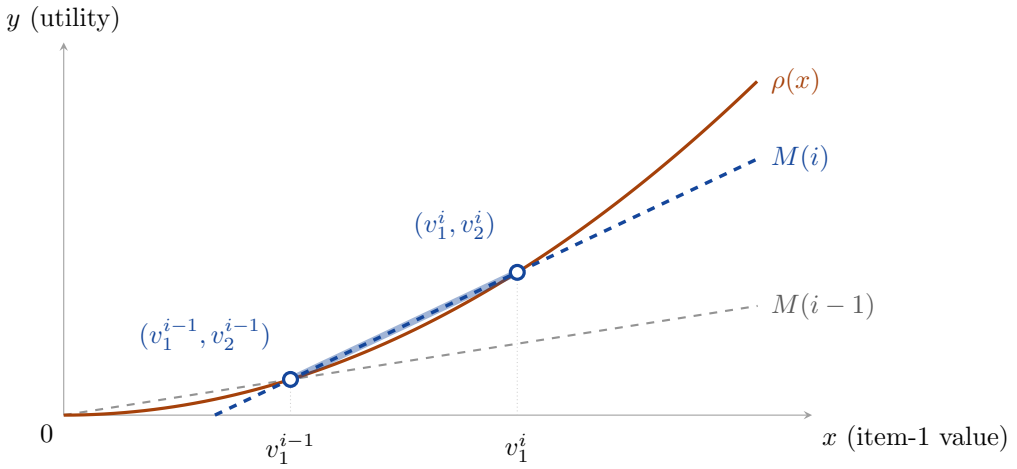


Figure 1: The indifferences of the two types i and $i-1$ determine the option $M(i)$. The blue dashed line is the utility from $M(i)$, and the orange curve is the utility from free item 2. At the left marked point, (C3) makes type $i-1$ indifferent between $M(i)$ and $M(i-1)$, while (C4) equates this utility with $\rho(v_1^{i-1})$, its utility from item 2 for free. At the right marked point, (C4) similarly gives type i utility $\rho(v_1^i)$ from $M(i)$. These two points uniquely determine the blue line, whose slope and intercept determine the allocation probability and price of $M(i)$.

For $i \geq 2$, we use types i and $i-1$. By (C4), type i must get utility v_2^i from $M(i)$, since this is their utility from free item 2. Hence,

$$v_2^i = q_1^i v_1^i - p^i.$$

For type $i-1$, (C3) says that they get the same utility from $M(i)$ as from their own option $M(i-1)$, and (C4) says that this utility is v_2^{i-1}. Thus,

$$v_2^{i-1} = q_1^{i-1} v_1^{i-1} - p^{i-1} = q_1^i v_1^{i-1} - p^i,$$

so the utility line for $M(i)$ must pass through (v_1^{i-1}, v_2^{i-1}) and (v_1^i, v_2^i), two consecutive points on the curve $y = \rho(x)$.

For the first option, $p^1 = 0$ fixes the intercept at zero, so its line passes through the origin and (v_1^1, v_2^1). Writing $(v_1^0, v_2^0) = (0, 0)$ lets us use the same formula for every option: the allocation

probability is the slope between the two points, and the price makes type i's utility equal to v_2^i.

$$q_1^i = \frac{v_2^i - v_2^{i-1}}{v_1^i - v_1^{i-1}} \qquad \text{and} \qquad p^i = q_1^i v_1^i - v_2^i. \tag{2}$$

The origin is only a reference point, not an additional type in the support.

The allocation probabilities are the slopes of consecutive chords of ρ. Since ρ is strictly convex, these slopes are strictly increasing, giving distinct item-1 lotteries. Substituting $(v_1^i, v_2^i) = (i/S, i^2/(4S^2))$ into (2) gives

$$M^* = \{(0,1,0),(0,0,0)\} \cup \left\{ \left(\frac{2i-1}{4S}, 0, \frac{i(i-1)}{4S^2} \right) \right\}_{i \in [S]}. \tag{3}$$

For inside type i, we break ties in favor of their corresponding item-1 option $M^*(i)$.

Lemma 4.2. *The menu M^* defined above in Eq. (3) is the unique IC and IR mechanism with nonnegative payments that satisfies all the constraints (C1)-(C4).*

Proof. The constraint (C1) is easy to check. The outside type o would get positive utility only from the $(0,1,0)$ lottery, so it is their chosen option.

Next we show that any inside type i weakly prefers their own corresponding lottery $M^*(i) = \left(\frac{2i-1}{4S}, 0, \frac{i(i-1)}{4S^2}\right)$ to every other inside option $M^*(j) = \left(\frac{2j-1}{4S}, 0, \frac{j(j-1)}{4S^2}\right)$. The difference between type i's utilities from these options is:

$$(q_1^i v_1^i - p^i) - (q_1^j v_1^i - p^j) = \frac{(i-j)(i-j+1)}{4S^2}. \tag{4}$$

Notice that since $i-j$ and $i-j+1$ are two consecutive integers, their product is nonnegative, so the utility of any inside type i will not increase by choosing another inside option $M^*(j)$.

We now show (C4): the utility of inside type i from getting the free item-2 lottery $(0,1,0)$ is simply $v_2^i = \frac{i^2}{4S^2}$. The utility i gets from their own inside option $M^*(i)$ is exactly the same:

$$q_1^i v_1^i - p^i = \frac{2i-1}{4S} \cdot \frac{i}{S} - \frac{i(i-1)}{4S^2} = \frac{i^2}{4S^2}.$$

By breaking ties in favor of $M^*(i)$, we can ensure that (C2) is satisfied.

Constraint (C3) follows from Equation (4) when $j = i+1$, as the difference in i's utility for their own option $M^*(i)$ and the option $M^*(i+1)$ is 0. The formula for the first option gives $p^1 = 0$.

The uniqueness of M^* follows immediately from Eq. (2). □

Although each inside type is indifferent between several menu options, these choices have different allocation costs with respect to L. Type i's assigned option has cost $q_1^i < c = 1/2$, whereas item 2 has cost c. When $i < S$, the next item-1 option also gives them the same utility but has higher cost q_1^{i+1}. We therefore break ties in favor of $M^*(i)$, which gives type i the same utility at the lowest allocation cost. We next choose probabilities that make this menu optimal among all mechanisms.

4.2 A Prior with a Unique Optimal Menu

In this section, we provide our choice of prior distribution $\mathcal{D}$ with support size T such that M^* is the *unique* truthful mechanism that optimizes our modified Lagrangian L. The following probabilities ensure that every truthful mechanism maximizing L satisfies (C1)–(C4).

Example 1 (Correlated Distribution)**.** Let $T \geq 3$ be any integer and $S = T - 1$. We define a correlated distribution $\mathcal{D}$ over two item valuations as follows:

- $(v_1, v_2) = (0, 1)$ with probability $f_o = 3/4$,
- $(v_1, v_2) = \left(\frac{i}{S}, \frac{i^2}{4S^2}\right)$ with probability $f_i = \frac{2S+1-2i}{4S^2}$, for all $i \in [S]$.

Proposition 4.3. *Let $\mathcal{D}$ be a prior distribution defined in Example 1 and M^* be the structured menu as defined in Eq. (3). Then M^* is the unique truthful mechanism that maximizes the modified Lagrangian L.*

To establish that M^* is the unique maximizer of L among all truthful mechanisms, we bound $L(M)$ by an expression that depends only on the prior. We choose the probabilities so that equality holds exactly when (C1)–(C4) hold. By Lemma 4.2, M^* is the unique mechanism satisfying these conditions.

Adjacent IC constraints first lower-bound allocation costs in terms of utility differences, giving an upper bound on L that depends on the inside types only through their utilities; equality requires (C2) and (C3). After weighting by the prior probabilities and regrouping, the coefficient of each inside type's utility in this bound is $Sf_{i+1} - (S-1)f_i$, with $f_{S+1} = 0$. We choose the probabilities to make these coefficients strictly negative. To attain the bound, each inside type must therefore receive as little utility as IC permits. IC requires each inside type to receive at least the utility it would obtain from the outside type's assigned option. Since the utility coefficients are negative, substituting these lower bounds yields an upper bound on L that is linear in q_1^o, q_2^o, p^o. Finally, the conditions on f_o and f_1 make the remaining bound uniquely maximized when the outside type receives item 2 for free, giving (C1) and making the preceding indifferences precisely (C4). The following lemma states the resulting bound; its proof is in Appendix A.

Lemma 4.4. *Suppose that $f_o > \frac{1}{2}$, $f_1 < \frac{1}{S}$ and $Sf_{i+1} < (S-1)f_i$ for all $1 \leq i < S$. We then have an upper bound on L for all truthful mechanisms M:*

$$L(M) \leq \frac{1}{2}f_o + \sum_{i=1}^{S} f_i \left(v_2^i - \frac{2i-1}{4S}\right),$$

where equality holds if and only if M satisfies all the constraints (C1)-(C4).

With this we are ready to prove Proposition 4.3.

Proof of Proposition 4.3. Observe that the chosen probabilities f_o and f_i's of Example 1 satisfy all the conditions needed for Lemma 4.4:

- $Sf_{i+1} - (S-1)f_i = -\frac{2i-1}{4S^2} < 0$ for all $1 \leq i < S$
- $f_o = \frac{3}{4} > \frac{1}{2}$ for all $S \geq 1$.

- $f_1 = \frac{2S-1}{4S^2} < \frac{1}{S}$

Therefore, by Lemma 4.4, every IC and IR mechanism M with nonnegative payments must satisfy:

$$L(M) \leq \frac{1}{2} f_o + \sum_{i=1}^{S} f_i \left(v_2^i - \frac{2i-1}{4S} \right) = \frac{33S^3 + 2S^2 + 1}{96S^3}.$$

Recall that, by Lemma 4.2, M^* is the unique menu that satisfies all the desirable properties (C1)-(C4). So Lemma 4.4 further implies that $L(M^*)$ equals this upper bound and thus M^* is the unique maximizer of L.

□

4.2.1 Finishing the Proof

We now apply the Lagrangian Caps Lemma 3.1 to obtain the ex-ante constraints for our lower-bound instance.

Proof of Theorem 4.1. The proof follows immediately from Proposition 4.3 and Lemma 3.1. In particular, for any $T \geq 3$, consider the correlated distribution $\mathcal{D}$ in Example 1 with support size T. We set the ex-ante supply constraints (α_1, α_2) to be the expected allocations of mechanism M^* under this prior:

$$\alpha_1 = \mathbb{E}_{(v_1,v_2)\sim\mathcal{D}}[x_1^{M^*}(v_1, v_2)] = \frac{2S^2+1}{48S^2} \qquad \alpha_2 = \mathbb{E}_{(v_1,v_2)\sim\mathcal{D}}[x_2^{M^*}(v_1, v_2)] = \frac{3}{4}.$$

By Proposition 4.3, the menu M^* is the unique mechanism that maximizes L, and in turn, is the unique utility-optimal mechanism subject to ex-ante supply constraints (α_1, α_2) by Lemma 3.1. The menu complexity of M^* is T by definition. Thus, for distribution $\mathcal{D}$ and ex-ante constraints (α_1, α_2) the menu complexity of any utility-optimal menu is exactly T. □

4.3 Continuous Distributions and Infinite Menu Complexity

A natural extension of our prior to a continuous distribution shows that utility-optimal mechanisms under ex-ante supply constraints can require infinitely many menu options. We replace the finite set of inside types with the curve $(v_1, \rho(v_1))$ for $v_1 \in [0, 1]$. The item-1 utility lines are now tangents to $\rho(v_1) = v_1^2/4$.

The continuous analogue of the preceding upper bound forces inside types to receive only item 1 and to obtain utility $\rho(v_1)$. The envelope theorem then requires allocation probability $v_1/2$.[8] Since a finite menu offers only finitely many allocation probabilities, it cannot attain this bound. Applying the Lagrangian Caps Lemma gives the following result; the prior and proof are in Section A.1.

Corollary 4.5. *There exists a correlated prior distribution and ex-ante supply constraints for which every utility-optimal mechanism has infinite menu complexity.*

[8] Statements about inside allocations in this continuous construction hold almost everywhere under the inside distribution.

Additive Valuations. Finally, observe that in the constructions for both the finite-support lower bound and its infinite-support extension, the proposed menu M^* never allocates both items, while the upper-bound arguments use only that each allocation probability lies in $[0, 1]$, without requiring $q_1 + q_2 \leq 1$. Thus, allowing additive allocations does not improve the optimal objective or alter the equality conditions. The same instances therefore require menu size T for support size T, and an infinite menu for the continuous construction, respectively.

Corollary 4.6. *(i) For an additive buyer, the finite correlated prior $\mathcal{D}$ with support size T and ex-ante constraints (α_1, α_2) constructed in the proof of Theorem 4.1 also require every utility-maximizing mechanism to have menu size at least T.*

(ii) For an additive buyer, the continuous correlated prior $\mathcal{D}$ and ex-ante constraints (α_1, α_2) constructed in the proof of Corollary 4.5 also require every utility-maximizing mechanism to have infinite menu size.

5 Lower Bound for Independent Items

In Section 4, buyer values were allowed to be correlated across items. In this section, we add the assumption that the item values are independent and show that the optimal mechanism can still require a menu of infinite size.

Theorem 5.1. *There exists a prior distribution $\mathcal{D}_1 \times \mathcal{D}_2$ over two independent item valuations and ex-ante constraints (α_1, α_2) where any optimal mechanism has infinite menu complexity.*

This result suggests that multidimensionality rather than carefully chosen correlation is the primary cause of the hardness result. The proof relies on building a class of instances—product distributions and corresponding ex-ante constraints—whose optimal mechanisms have structured menus that are easier to optimize over. In particular, we focus our attention on finding a hard single-dimensional first-item distribution such that, when the second item's value distribution is fixed as uniform over $(0, 1)$, there is an optimal menu that has a simple structure.

Similar to the correlated lower bound, we will use the modified Lagrangian objective $L = L_{1,c}$ to characterize the optimal mechanism. By the Lagrangian Caps Lemma 3.1, we defer choosing ex-ante constraints by optimizing over the modified objective L. However, for this argument to hold, we can no longer freely choose any cost $c \in (0, 1)$; we must choose c to satisfy certain conditions. We leave c as a variable until we determine its value in Appendix B.1.

The proof rests on a simple observation about the buyer's utility function given a truthful mechanism M. The utility of a buyer for a fixed menu option $(\vec{q}, p)$ is an affine function $\vec{q} \cdot \vec{v} - p$ over valuation $\vec{v}$. Given a finite menu M, the buyer chooses the best option, and thus the utility of a buyer $u_M(\vec{v})$ is a maximum over affine functions. Moreover, the restriction of this function to $v_2 = 0$ is a piecewise affine function; see Figure 2. Then if the optimality of some instance requires a utility function that is strictly convex on an interval along this line, no finite menu can implement it.

We build such an instance in the following order:

1. We hope for *structured menus*, those with the following structure: (1) *single-item lotteries*, where each menu option only allocates a single item, and (2) *item 2 is free*, where there is only a single menu option that sells item 2, and that is a free lottery. We find conditions on the priors that ensure there exists an optimal mechanism satisfying this structure without loss of generality (Section 5.1).

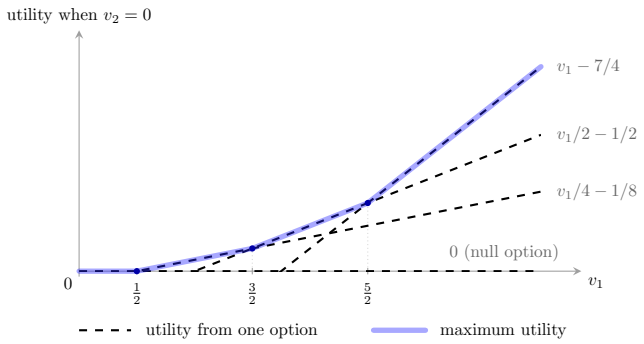


Figure 2: The utility induced by a finite menu is piecewise affine along $v_2 = 0$.

2. With this simplifying assumption, we get a much simpler form of the objective based on only 2 variables, one being the utility curve of the mechanism when item 2 has zero value ($v_2 = 0$). Finally, we can instantiate a prior in the above class where the objective is optimized only when this projected utility curve is strictly convex, implying that no finite menu can induce the optimal utility function (Section 5.2, Section 5.3, Appendix B.1, and Appendix B.2).

5.1 Priors with a Structured Optimal Menu

We first identify conditions on the prior under which a structured menu is optimal. For convenience, we will repeat the simple properties that a structured menu must have:

(S1) There are no two-item lotteries: every option sells either the first item or the second one, never both.

(S2) There is only a single option selling the second item, and it is free.

5.1.1 The Prior for Item 2

Throughout this paper, one can think of item 2 as the fallback option for those not getting the first item. For the sake of simplicity, we choose its prior to be $U(0, 1)$. Under this prior, the usage of item 2 depends only on the utility at the two endpoints of its distribution: when $v_2 = 0$ and $v_2 = 1$. More specifically, for a fixed v_1, by incentive compatibility and the envelope theorem (Theorem 2.1),[9]

$$\mathbb{E}_{v_2 \sim \mathcal{D}_2}\left[x_2^M(v_1, v_2)\right] = \int_0^1 \frac{\partial}{\partial v_2} u_M(v_1, v_2)\, dv_2 = u_M(v_1, 1) - u_M(v_1, 0). \tag{5}$$

[9]Throughout this section and its deferred proofs, derivative identities are understood almost everywhere. Menu utility functions are Lipschitz in each coordinate and hence absolutely continuous along coordinate lines, so their derivatives recover utility differences by integration. We omit the almost-everywhere qualification in subsequent derivative calculations. When equality of continuous utility functions follows almost everywhere on an interval with a strictly positive density, continuity extends that equality to every point of the interval.

5.1.2 Structure (S1): Replacing Two-Item Lotteries with Single-Item Lotteries

We first show that, under some distributional assumptions, we can replace any two-item menu M with a menu M' that contains only single-item lotteries by projecting each option to the $v_1 = 0$ and $v_2 = 0$ axes (by zeroing out each item's allocation). While the resulting mechanism may lose utility, and may even allocate more of a particular item to some buyer types, in aggregate, we show that the resulting menu weakly increases L.

Lemma 5.2. *Suppose that $\mathcal{D}_2 = U(0,1)$ and that f_1 satisfies*

(F1) *f_1 is continuous and piecewise smooth on $(0,\infty)$, with $v_1 f_1(v_1) \to 0$ both as $v_1 \to 0^+$ and as $v_1 \to \infty$, and $\mathcal{D}_1$ has a finite mean;*

(F2) *$f_1(v_1) + f_1'(v_1) \le 0$ whenever $f_1(v_1)$ is differentiable.*

Then, for any truthful mechanism M, there exists another mechanism M' that only uses single-item lotteries, while $L(M) \le L(M')$. Moreover, M' uses weakly less supply of each item than M.

Proof. We build a new menu M' by splitting each option in M into two single-item options at the same price: for every option (q_1, q_2, p) in M, add $(q_1, 0, p)$ and $(0, q_2, p)$ to M'. Since a buyer facing M' can now only choose one single-item menu option, this is equivalent to their best option under M when they only value a single item, having no value for the other. That is, M' projects M onto the two axes $v_1 = 0$ and $v_2 = 0$:

Lemma 5.3. *Given any menu M and the projected menu M', we observe the utility of any type (v_1, v_2) when participating in either:*

(i) $u_{M'}(v_1, v_2) = \max\{u_M(v_1, 0), u_M(0, v_2)\}$

(ii) $u_M(v_1, 0) = u_{M'}(v_1, 0)$ *and* $u_M(0, v_2) = u_{M'}(0, v_2)$.

We use Lemma 5.3 to complete the proof of Lemma 5.2, then prove the projection lemma 5.3.

Let $w = u_M - u_{M'}$ denote the utility lost by the projection. Fix any valuation (v_1, v_2) at which the utilities from both mechanisms are differentiable. The difference of the two objectives $L(M) - L(M')$ at this valuation, using the envelope theorem, is

$$w(v_1, v_2) - \frac{\partial w}{\partial v_1}(v_1, v_2) - c\,\frac{\partial w}{\partial v_2}(v_1, v_2), \tag{6}$$

since the differences in the two allocations are the partial derivatives of w. We integrate the three terms over both supports.

For the second term, fix v_2 and integrate by parts over v_1:

$$\int_0^\infty f_1(v_1)\frac{\partial w}{\partial v_1}(v_1, v_2)\,dv_1 = \Big[f_1(v_1)\,w(v_1, v_2)\Big]_{v_1 \to 0^+}^{v_1 \to \infty} - \int_0^\infty f_1'(v_1)\,w(v_1, v_2)\,dv_1. \tag{7}$$

We observe that $u_M(v_1, v_2) - u_{M'}(v_1, v_2) \le v_1$. Intuitively, since M' breaks each lottery in M into two single-item lotteries by removing the probability of getting one of the items, the loss in utility is at most the utility from that removed item, which is at most the buyer's value for that item. Let (q_1, q_2, p) be the menu option chosen in M by a buyer with type (v_1, v_2). Then the utility of (v_1, v_2) from M' will be at least $q_2 v_2 - p$. Hence $u_{M'}(v_1, v_2) \ge q_2 v_2 - p = u_M(v_1, v_2) - q_1 v_1 \ge u_M(v_1, v_2) - v_1$.

Using this fact, we have $0 \le f_1(v_1)w(v_1, v_2) \le f_1(v_1)v_1$, which tends to 0 at both ends by (F1). There is no boundary term at interior points because f_1 is continuous. Hence

$$\int_0^\infty f_1(v_1)\frac{\partial w}{\partial v_1}(v_1, v_2)\, dv_1 = -\int_0^\infty f_1'(v_1)\, w(v_1, v_2)\, dv_1. \tag{8}$$

For the third term, fix v_1. By (5),

$$\int_0^1 \frac{\partial w}{\partial v_2}(v_1, v_2)\, dv_2 = w(v_1, 1) - w(v_1, 0) = w(v_1, 1), \tag{9}$$

where the last line uses $w(v_1, 0) = 0$ thanks to (ii) in the projection lemma 5.3. Plugging (8) and (9) into the integral of (6), we get

$$L(M) - L(M') = \int_0^\infty \int_0^1 \big(f_1(v_1) + f_1'(v_1)\big)\, w(v_1, v_2)\, dv_2\, dv_1 - c\int_0^\infty f_1(v_1)\, w(v_1, 1)\, dv_1. \tag{10}$$

Both terms are non-positive: $w \ge 0$, and $f_1 + f_1' \le 0$ by (F2).

For the supplies, the same two computations give

$$\begin{aligned}
\mathbb{E}[x_1^M] - \mathbb{E}[x_1^{M'}] &= -\int_0^\infty \int_0^1 f_1'(v_1)\, w(v_1, v_2)\, dv_2\, dv_1 \ \ge 0,\\
\mathbb{E}[x_2^M] - \mathbb{E}[x_2^{M'}] &= \int_0^\infty f_1(v_1)\, w(v_1, 1)\, dv_1 \ \ge 0,
\end{aligned}$$

where the first inequality uses $f_1' \le -f_1 \le 0$, which follows from (F2). □

Finally, the only step missing is our promised proof for the projection lemma 5.3.

Proof of Lemma 5.3. (i) Consider a buyer with valuation (v_1, v_2) in mechanism M'. Suppose, of the menu options that sell only item 1, $(q_1, 0, p)$ maximizes this buyer's utility. Then it gives utility

$$q_1 v_1 - p.$$

This is exactly the same as offering the corresponding option (q_1, q_2, p) in M to a buyer of valuation $(v_1, 0)$. Symmetrically, of the options selling only the second item, the buyer's favorite will yield utility equal to $u_M(0, v_2)$. Hence the buyer's favorite between these will have utility $u_{M'}(v_1, v_2) = \max\{u_{M'}(v_1, 0), u_{M'}(0, v_2)\}$. This correspondence holds for every menu option, so the best item-1 option in M' gives utility $u_M(v_1, 0)$, and the best item-2 option gives utility $u_M(0, v_2)$. The buyer chooses the better of the two, giving the claimed maximum.

(ii) When $v_1 = 0$ or $v_2 = 0$, the utility the buyer would get from M is exactly the same as in M'. Without loss of generality, assume that $v_2 = 0$. Applying the formula from (i), we have:

$$\begin{aligned}
u_{M'}(v_1, 0) &= \max\{u_M(v_1, 0), u_M(0, 0)\}\\
&= u_M(v_1, 0)
\end{aligned}$$

since $u_M(0, 0) = 0 \le u_M(v_1, 0)$ due to the IR constraints on the type $(v_1, 0)$. □

5.1.3 Structure (S2): A Single Free Lottery for Item 2

Now, we replace the multiple item-2 options with a single free option for item 2, and show that, for $\mathcal{D}_2 = U(0,1)$, this only increases our objective L. We observe that transforming the menu preserves u_1 and the utility of the highest item-2-only type.

Lemma 5.4. *Suppose that $\mathcal{D}_2 = U(0,1)$ and M contains only single-item lotteries. Let $d = u_M(0,1)$, and let M' be the menu obtained by keeping all item-1 options in M and replacing all item-2 options with the single free lottery $(0,d,0)$. Then $L(M) \le L(M')$, and M' uses weakly less supply of each item than M. No condition on f_1 is needed.*

Proof. Consider any menu M with single-item lotteries. Denote its two projections by

$$u_1(v_1) = u_M(v_1, 0), \qquad u_2(v_2) = u_M(0, v_2),$$

so that $u_M(v_1, v_2) = \max\{u_1(v_1), u_2(v_2)\}$. Let $d = u_2(1) \in [0,1]$. Our choice of constructing M' from M is by keeping all options for the first item and replacing all options for the second item by the single free lottery $(0,d,0)$. Then $u_{M'}(v_1, v_2) = \max\{u_1(v_1), dv_2\}$.

Recall that $u_2(\cdot)$ is convex with $u_2(0) = 0$, hence for all $v_2 \in [0,1]$ we have

$$u_2(v_2) \le (1 - v_2)\, u_2(0) + v_2\, u_2(1) = d\, v_2.$$

So for every buyer the options for item 1 are the same in both menus, while the option for item 2 is weakly better in M'. Hence the utility weakly increases. Moreover, a buyer who chooses item 2 under M satisfies $u_1(v_1) \le u_M(v_1, v_2) = u_2(v_2) \le dv_2$. With ties broken in favor of the item chosen under M, they continue to choose item 2 under M'. Thus, the expected allocation of item 1 weakly decreases.

It remains to show that the usage of item 2 does not increase. Fix any v_1. Because item 2 has a uniform prior, by (5), the expected allocation of item 2 in M is

$$u_M(v_1, 1) - u_M(v_1, 0) = \max\{u_1(v_1), d\} - u_1(v_1) = (d - u_1(v_1))_+,$$

and the same computation for M' also gives $(d - u_1(v_1))_+$. So the usage of item 2 is unchanged, and $L(M) \le L(M')$. □

Then for $\mathcal{D}_2 = U(0,1)$ and any f_1 satisfying (F1) and (F2), every menu M can be replaced by a menu with structures (S1) and (S2) without decreasing L. The resulting menu is described by two objects that depend on M itself and do not change along the way:

$$u_1(v_1) = u_M(v_1, 0) \qquad \text{and} \qquad d = u_M(0,1) \in [0,1]. \tag{11}$$

5.2 The Modified Lagrangian for Structured Menus

We now compute L for a structured menu. The buyer's favorite item-1 option gives utility $u_1(v_1)$, while the free lottery for item 2 is allocated with probability d, giving utility dv_2. The buyer chooses the better of these options, so

$$u(v_1, v_2) = \max\{u_1(v_1), dv_2\}.$$

We compute L by fixing v_1 and then integrating over $\mathcal{D}_1$. We do this one component at a time.

Utility. For fixed v_1, by dividing utility into the utility the buyer obtains from the item-2 lottery plus the surplus from choosing item 1, we get

$$\mathbb{E}_{v_2\sim\mathcal{D}_2}[u(v_1,v_2)] = \mathbb{E}_{v_2\sim\mathcal{D}_2}[dv_2] + \mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+] = \frac{d}{2} + \mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+].$$

Item 1. By the envelope theorem, the allocation of item 1 is the derivative of the utility in v_1, so its expectation over v_2 is the derivative of $\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+]$. Integrating by parts,

$$\begin{aligned}\int_0^\infty f_1(v_1)\mathbb{E}_{v_2\sim\mathcal{D}_2}[x_1(v_1,v_2)]\,dv_1 &= \Big[f_1(v_1)\,\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+]\Big]_{v_1\to 0^+}^{v_1\to\infty} \\ &\quad - \int_0^\infty f_1'(v_1)\,\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+]\,dv_1.\end{aligned}$$

Since $0 \le \mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+] \le u_1(v_1) \le v_1$, the first term vanishes by (F1), which requires $f_1(v_1)v_1 \to 0$ approaching both endpoints.

Item 2. By Equation (5), the expected allocation of item 2 is $(d-u_1(v_1))_+ = d-\min\{d,u_1(v_1)\}$.

Assembling. Combining the three terms and using $\int f_1 = 1$,

$$L(M) = d\left(\frac{1}{2}-c\right) + \int_0^\infty \Big[(f_1(v_1)+f_1'(v_1))\,\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1)-dv_2)_+] + c\,f_1(v_1)\min\{d,u_1(v_1)\}\Big]\,dv_1. \tag{12}$$

5.3 The Target Utility Curve and Its Menu

Recall that $u_1(v_1)$ is the utility a buyer with value v_1 gets from their favorite item-1 option (or when $v_2 = 0$). To complete the proof, we'll select a target one-dimensional utility curve $r(v_1)$ that is strictly convex (so no finite menu can attain it), and construct priors where L is maximized only when $u_1(v_1) = r(v_1)$ in the chosen range $[0,2]$.

We now choose $r(v_1)$. There are many choices that work, but they would need to satisfy:

- r is implementable for the single-dimensional item-1 setting: $r(0) = 0$, and r is convex with $r'(v_1) \in [0,1]$ (this is the allocation rule). This automatically holds for the utility curve of every menu: the utility is a maximum of affine functions, $r'(v_1)$ is the probability of allocating item 1, and nonnegative prices give utility 0 at the type $(0,0)$.
- r is strictly convex on an interval; here we choose $(0,2)$. This is what no finite menu can reproduce.
- $r'(0) > c$. This will be explained further in Appendix B.1, where this condition is sufficient for the prior to be integrable.

For the sake of this proof, we pick:

$$r(v_1) = \begin{cases} v_1/4 + v_1^2/8, & 0 \le v_1 \le 2, \\ v_1 - 1, & v_1 \ge 2. \end{cases} \tag{13}$$

Lemma 5.5. *The utility curve $r(v_1)$ described above is implementable in single-item settings, i.e., there exists a single-item menu that produces such utility.*

Proof. We will directly find the menu options that realize this curve. On the axis $v_2 = 0$ where this is a one-item problem, item 1's allocation probability is the slope of the utility, and the price follows from the payment identity. This gives the menu options

$$\left(\frac{1+t}{4}, \frac{t^2}{8}\right) \quad (0 \le t \le 2) \qquad \text{and} \qquad (1, 1).$$

□

With this single-item menu in mind, the most obvious extension to two items with the above two structural properties would be:

$$\left(\frac{1+t}{4}, 0, \frac{t^2}{8}\right) \quad (0 \le t \le 2), \qquad (1, 0, 1), \qquad \text{and} \qquad (0, 1, 0). \tag{14}$$

We call this menu M^*. A buyer with valuation (v_1, v_2) gets the utility

$$u^*(v_1, v_2) = \max\{r(v_1), v_2\}.$$

Now that we have designed a utility curve r that no finite menu can implement because it's strictly convex on $[0, 2]$, in Appendix B.1, we continue to back-solve by constructing a distribution $\mathcal{D}_1$ specified by its density f_1. Recall that in any structured menu, d is the probability of receiving item 2 in its single free lottery option $(0, d, 0)$. We choose f_1 so that maximizing L over structured menus with $d = 1$ forces $u_1 = r$ almost everywhere on $(0, 2)$. The construction in Appendix B.1 gives the density (19) with $c = 1/8$ and verifies (F1)–(F4).

The next step (Appendix B.2) is to show that restricting to menus with $d = 1$ does not exclude any other optimal structured menu. For our chosen prior, every structured menu M with free item-2 lottery $(0, d, 0)$ satisfies $L(M) \le dL(M^*)$, even when we allow its item-1 menu to vary. Since $L(M^*) > 0$, every $d < 1$ is strictly worse than M^*. We now combine this fact with the reductions to structured menus to show that M^* is optimal and that every optimal mechanism requires an infinite menu.

5.4 Finishing the Proof

We now combine the reduction to structured menus with the prior construction, which forces every optimal item-1 utility curve to equal r on $(0, 2)$. Since the reduction preserves utility when $v_2 = 0$, every maximizer of L must reproduce this strictly convex curve, which no finite menu can do.

Theorem 5.6. *For $\mathcal{D}_1$ with density (19), $\mathcal{D}_2 = U(0, 1)$, and $c = 1/8$, the menu M^* in (14) maximizes L. Moreover, every maximizer M of L satisfies*

$$u_M(v_1, 0) = \frac{v_1}{4} + \frac{v_1^2}{8} \qquad (0 < v_1 < 2).$$

Consequently, no finite menu maximizes L.

Proof. Consider any menu M. By Lemmas 5.2 and 5.4, we can replace M with a structured menu M' without decreasing L. Let $(0, d, 0)$ be the free option for item 2 in M'. By Lemma B.2,

$$L(M) \leq L(M') \leq dL(M^*) \leq L(M^*).$$

This holds for every menu M. Thus, M^* maximizes L among all menus.

Now, suppose that M also maximizes L. As the transformation from M to M' preserves $u(v_1, 0)$, we have:

$$u_{M'}(v_1, 0) = u_M(v_1, 0).$$

Moreover, as M maximizes L, all inequalities in the sequence above must hold with equality, leading to M' being a structured menu that maximizes L, as well as having $d = 1$. Then, using Lemma B.1, we can claim that

$$u_{M'}(v_1, 0) = r(v_1) = \frac{v_1}{4} + \frac{v_1^2}{8} \qquad \forall\, 0 < v_1 < 2.$$

Notice that if M has finite menu size, $u_M(v_1, 0)$ is the maximum of finitely many affine functions. It therefore cannot agree with this strictly convex quadratic function on the interval $(0, 2)$, implying that no finite menu maximizes L. □

Finally, we return from the modified Lagrangian L to our true objective of utility by reintroducing the ex-ante constraints using the Lagrangian Caps Lemma 3.1. We set the ex-ante caps equal to the expected allocation of M^*, showing that both L and supply terms are the highest possible, and thus implying M^* has the highest utility possible.

Proof of Theorem 5.1. Take the independent priors $\mathcal{D}_1$ and $\mathcal{D}_2$ constructed in Appendix B.1. We choose the ex-ante constraints to be the expected allocations of M^*:

$$\alpha_i = \mathbb{E}_{(v_1,v_2)\sim\mathcal{D}_1\times\mathcal{D}_2}[x_i^{M^*}(v_1, v_2)], \qquad i = 1, 2.$$

Since each allocation probability lies in $[0, 1]$, these are valid ex-ante constraints.

More explicitly,

$$\alpha_1 = \int_0^2 f_1(v_1)\, r(v_1)\, r'(v_1)\, dv_1 + \int_2^\infty f_1(v_1)\, dv_1,$$
$$\alpha_2 = \int_0^2 f_1(v_1)\, (1 - r(v_1))\, dv_1.$$

Indeed, fix $0 < v_1 < 2$. The buyer chooses an item-1 option when $v_2 < r(v_1)$. Since $\mathcal{D}_2 = U(0, 1)$, this happens with probability $r(v_1)$, and the chosen option allocates item 1 with probability $r'(v_1)$. Otherwise, the buyer chooses the free option for item 2 and receives it with probability 1. For $v_1 > 2$, the buyer always receives item 1.

By Theorem 5.6, M^* maximizes L. Lemma 3.1 therefore implies that M^* maximizes expected buyer utility subject to (α_1, α_2). Moreover, every utility-maximizing mechanism subject to these constraints also maximizes L. By Theorem 5.6, no finite menu maximizes L. Thus, every utility-maximizing mechanism under these ex-ante constraints has infinite menu complexity. □

Additive Valuations. Finally, observe that throughout this section, our independent construction never uses the unit-demand assumption (that $q_1 + q_2 \leq 1$). The menu M^* inherently never allocates both items, and the proofs only use that each allocation probability is at most 1. The same instance therefore shows that an additive buyer with two independent items can require infinite menu complexity.

Corollary 5.7. *For an additive buyer, the independent prior $\mathcal{D}_1 \times \mathcal{D}_2$ and ex-ante constraints (α_1, α_2) constructed in the proof of Theorem 5.1 also require every utility-maximizing mechanism to have infinite menu complexity.*

AI Disclosure

We used ChatGPT extensively in developing this paper. We supplied existing research notes, partial results, attempted approaches, and relevant literature, and guided an iterative investigation by formulating questions, directing proof attempts and extensions, and providing feedback. Within this process, GPT-6 Astra generated draft proofs of most of the results that we built off of in the paper, although many initial proofs were generated by Claude Fable 5.0 and GPT-5.6 Sol.

We wrote this paper and the proofs by completely rewriting AI-generated drafts, which involved understanding and reinterpreting the arguments, drawing out the underlying intuitions, developing new notation, simplifying and restructuring the proofs. We used ChatGPT and Claude in this process for developing results and copyediting. We take full 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] 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.

[3] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 869–877, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4528-6. doi: 10.1145/3055399.3055426. URL http://doi.acm.org/10.1145/3055399.3055426.

[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.

[5] Mark Braverman, Jing Chen, and Sampath Kannan. Optimal provision-after-wait in healthcare. *Mathematics of Operations Research*, 41(1):352–376, 2016.

[6] Surajeet Chakravarty and Todd R. Kaplan. Optimal Allocation without Transfer Payments. *Games and Economic Behavior*, 77(1):1–20, 2013. doi: 10.1016/j.geb.2012.08.006.

[7] 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 Computation*, 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.

[8] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. Menu-size complexity and revenue continuity of buy-many mechanisms. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 2020. URL https://arxiv.org/abs/2003.10636.

[9] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. On the complexity of optimal lottery pricing and randomized mechanisms for a unit-demand buyer. *SIAM Journal on Computing*, 51(3):492–548, 2022.

[10] Daniele Condorelli. What Money Can't Buy: Efficient Mechanism Design with Costly Signals. *Games and Economic Behavior*, 75(2):613–624, 2012. doi: 10.1016/j.geb.2012.02.018.

[11] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 269–286, 2013. URL http://doi.acm.org/10.1145/2482540.2482593.

[12] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1302–1318. SIAM, 2014.

[13] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767, 2017.

[14] Nikhil R Devanur, Kira Goldner, Raghuvansh R Saxena, Ariel Schvartzman, and S Matthew Weinberg. Optimal mechanism design for single-minded agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 193–256, 2020.

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

[16] 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.

[17] Dimitris Fotakis, Dimitris Tsipras, Christos Tzamos, and Emmanouil Zampetakis. Efficient money burning in general domains. *Theory of Computing Systems*, 59:619–640, 2016.

[18] Yiannis Giannakopoulos and Elias Koutsoupias. Selling Two Goods Optimally. In *ICALP (2)*, volume 9135 of *Lecture Notes in Computer Science*, pages 650–662. Springer, 2015.

[19] Kira Goldner and Taylor Lundy. Multidimensional Bayesian utility maximization: Tight approximations to welfare. In *Advances in Neural Information Processing Systems*, 2025. URL https://arxiv.org/abs/2402.12340.

[20] Kira Goldner, Taylor Lundy, and Thodoris Tsilivis. A constant-factor approximation to multidimensional consumer utility, 2026. URL https://arxiv.org/abs/2610.10516.

[21] Yannai A. Gonczarowski. Bounding the menu-size of approximately optimal auctions via optimal-transport duality. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 123–131, 2018. doi: 10.1145/3188745.3188786. URL http://doi.acm.org/10.1145/3188745.3188786.

[22] Yannai A Gonczarowski, Nicole Immorlica, Yingkai Li, and Brendan Lucier. Revenue maximization for buyers with costly participation. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 41–73. SIAM, 2024.

[23] Nima Haghpanah and Jason Hartline. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 757–758. ACM, 2015. Updated version: http://arxiv.org/abs/1404.1341.

[24] 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.

[25] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019. doi: 10.1016/j.jet.2019.07.006. URL https://math.huji.ac.il/hart/papers/m-corr.pdf.

[26] Jason D. Hartline. Mechanism design and approximation, 2017. URL https://jasonhartline.com/MDnA/. Book draft, Chapter 8, July 5, 2017.

[27] 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.

[28] Bruno Jullien. Participation constraints in adverse selection models. *Journal of Economic Theory*, 93(1):1–47, 2000. ISSN 0022-0531. doi: https://doi.org/10.1006/jeth.1999.2641. URL https://www.sciencedirect.com/science/article/pii/S0022053199926418.

[29] Pravesh K. Kothari, Divyarthi Mohan, Ariel Schvartzman, Sahil Singla, and S. Matthew Weinberg. Approximation schemes for a unit-demand buyer with independent items via symmetries. In *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science*, pages 220–232, 2019. doi: 10.1109/FOCS.2019.00023. URL https://arxiv.org/abs/1905.05231.

[30] Jean-Jacques Laffont, Eric Maskin, and Jean-Charles Rochet. Optimal nonlinear pricing with two-dimensional characteristics. *Information, Incentives and Economic Mechanisms*, pages 256–266, 1987.

[31] Alejandro M. Manelli and Daniel R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1 – 35, 2006. URL http://dx.doi.org/10.1016/j.jet.2005.08.007.

[32] R. Preston McAfee and John McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335 – 354, 1988. URL http://dx.doi.org/10.1016/0022-0531(88)90135-4.

[33] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. URL http://dx.doi.org/10.1287/moor.6.1.58.

[34] Gregory Pavlov. Optimal mechanism for selling two goods. *The B.E. Journal of Theoretical Economics*, 11(3), 2011. URL http://dx.doi.org/10.2202/1935-1704.1664.

[35] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394. ACM, 2015.

[36] Raghuvansh R. Saxena, Ariel Schvartzman, and S. Matthew Weinberg. The menu-complexity of one-and-a-half dimensional mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, 2018.

A Missing Proofs from Section 4

Let M be any IC and IR mechanism with nonnegative payments. For each inside type $i \in [S]$, we let (q_1^i, q_2^i, p^i) denote type i's allocation probabilities and payment in M. Similarly, let (q_1^o, q_2^o, p^o) denote the outside type $o = (0, 1)$'s allocation and payment in M. Let u_i and u_o denote the utilities of inside type i and the outside type, respectively.

The following is a helper lemma that will be used in the proof of Lemma 4.4. We upper bound $u_i - q_1^i - cq_2^i$, the contribution of each inside type i to the objective L.

Lemma A.1. *For every $1 \le i \le S$:*

$$u_i - q_1^i - cq_2^i \le u_i - S(u_i - u_{i-1}),$$

where u_i is the utility of type i in M, and when $i = 0$ utility and allocations are both 0. Moreover, equality holds simultaneously for all i only when the mechanism M satisfies (C3): the first indifference constraint and $p^1 = 0$, as well as (C2): all inside types would not receive the second item.

Proof. Consider the IC condition corresponding to the first indifference constraint: agent type $i-1$ should not prefer the ith option to its own.

$$\begin{aligned} u_{i-1} &\ge u_i - (v_1^i - v_1^{i-1})q_1^i - (v_2^i - v_2^{i-1})q_2^i \\ &\ge u_i - \frac{1}{S}q_1^i - \frac{2i-1}{4S^2}q_2^i. \end{aligned}$$

Notice that even in the degenerate case of $i = 1$, the above inequalities are still true thanks to nonnegative payment:

$$\begin{aligned} u_1 - (v_1^1 - v_1^0)q_1^1 - (v_2^1 - v_2^0)q_2^1 &= u_1 - v_1^1 q_1^1 - v_2^1 q_2^1 \\ &= -p^1 \\ &\le 0 = u_0. \end{aligned}$$

Using the IC condition for all types i, we can rewrite the first item allocation as:

$$\begin{aligned} q_1^i &\ge S(u_i - u_{i-1}) - \frac{2i-1}{4S}q_2^i \\ &\ge S(u_i - u_{i-1}) - \frac{1}{2}q_2^i, \end{aligned}$$

since $\frac{2i-1}{4S} < \frac{1}{2}$ for all $i \le S$. Plugging the above formula back into the contribution of agent i to L, we get an upper bound involving only utilities:

$$\begin{aligned} u_i - q_1^i - cq_2^i &\le u_i - \left(S(u_i - u_{i-1}) - \frac{1}{2}q_2^i\right) - \frac{1}{2}q_2^i \\ &\le u_i - S(u_i - u_{i-1}), \end{aligned}$$

where equality holds throughout if and only if $q_2^i = 0$ for every inside type, each type $i \in [S-1]$ is indifferent between its own option and the next higher option, and $p^1 = 0$. These are precisely (C2) and (C3).

□

We now prove Lemma 4.4.

Lemma 4.4. *Suppose that $f_o > \frac{1}{2}$, $f_1 < \frac{1}{S}$ and $Sf_{i+1} < (S-1)f_i$ for all $1 \le i < S$. We then have an upper bound on L for all truthful mechanisms M:*

$$L(M) \le \frac{1}{2} f_o + \sum_{i=1}^{S} f_i \left(v_2^i - \frac{2i-1}{4S} \right),$$

where equality holds if and only if M satisfies all the constraints (C1)-(C4).

Proof. By Lemma A.1, we upper bound L as follows:

$$\begin{aligned} L(M) &= f_o(u_o - q_1^o - cq_2^o) + \sum_{i=1}^{S} f_i(u_i - q_1^i - cq_2^i) \\ &\le f_o(u_o - q_1^o - cq_2^o) + \sum_{i=1}^{S} f_i(u_i - S(u_i - u_{i-1})) \qquad (15) \\ &= f_o(u_o - q_1^o - cq_2^o) + \sum_{i=1}^{S} (Sf_{i+1} - (S-1)f_i)u_i, \end{aligned}$$

where the last line follows simply by regrouping and setting $f_{S+1} = 0$. Recall that inequality (15) becomes equality if and only if M satisfies $(C2)$ and $(C3)$, by Lemma A.1.

By our assumption on the prior, we have that $(Sf_{i+1} - (S-1)f_i) < 0$ for all $i < S$. Moreover, $(S-1)f_S > 0$ because by definition $S = T - 1 \ge 2$ and $f_S > 0$. That is, the coefficients of u_i are negative for all $i \in [S]$. We thus update the bound on L as follows,

$$\begin{aligned} L(M) &\le f_o(u_o - q_1^o - cq_2^o) + \sum_{i=1}^{S} (Sf_{i+1} - (S-1)f_i)u_i \\ &\le f_o(u_o - q_1^o - cq_2^o) + \sum_{i=1}^{S} (Sf_{i+1} - (S-1)f_i)(v_1^i q_1^o + v_2^i q_2^o - p^o), \qquad (16) \end{aligned}$$

by using the IC condition $u_i \ge v_1^i q_1^o + v_2^i q_2^o - p^o$, i's utility from M is at least their utility from getting the outside type's option (q_1^o, q_2^o, p^o). Observe that the final inequality (16) holds as equality if and only if every inside type is indifferent between its assigned option and the outside type assigned option.

Substituting in $u_o = q_2^o - p^o$, the above upper bound on L is just a linear function of the outside type's allocation probabilities and payment,

$$L(M) \le \gamma_1 q_1^o + \gamma_2 q_2^o + \gamma_p p^o,$$

where the coefficients of q_1^o, q_2^o, and p^o are defined as follows:

$$\gamma_1 = -f_o - \sum_{i=1}^{S} f_i(1 - v_1^i)$$

$$\gamma_2 = f_o(1-c) + \sum_{i=1}^{S} f_i \left(v_2^i - \frac{2i-1}{4S} \right) = \frac{1}{2} f_o + \sum_{i=1}^{S} f_i \left(v_2^i - \frac{2i-1}{4S} \right)$$
$$\gamma_p = S f_1 - (f_o + \sum_{i=1}^{S} f_i) = S f_1 - 1,$$

where we use $c = 1/2$ and $f_o + \sum_i f_i = 1$.

First, notice that γ_1 is always negative. Next, since $f_1 < \frac{1}{S}$, γ_p is also always negative. Finally,

$$\begin{aligned}
\gamma_2 &= \frac{1}{2} f_o + \sum_{i=1}^{S} f_i \left(v_2^i - \frac{2i-1}{4S} \right) \\
&\geq \frac{1}{2} f_o - \sum_{i=1}^{S} f_i \left(\frac{2i-1}{4S} \right) && \text{since } v_2^i \geq 0 \text{ for all } i \\
&\geq \frac{1}{2} f_o - \sum_{i=1}^{S} \frac{1}{2} f_i && \text{since } i \leq S \text{ for all } i \\
&\geq \frac{1}{2}(f_o - (1 - f_o)) && \text{the sum of all probabilities is 1} \\
&> 0 && \text{since } f_o > 1/2.
\end{aligned}$$

Thus, for all IC and IR mechanisms M with nonnegative payments, we have

$$\begin{aligned}
L(M) &\leq \gamma_1 q_1^o + \gamma_2 q_2^o + \gamma_p p^o \\
&\leq \gamma_2
\end{aligned} \tag{17}$$

where γ_2 depends only on the prior, and not the mechanism M itself.

Equality in the last inequality of Equation (17) holds if and only if the outside type is assigned option $(0, 1, 0)$, giving (C1). Combining (C1) with the equality constraint in Equation (16) gives exactly (C4), while equality in Equation (15) is equivalent to (C2) and (C3). Thus, $L(M) = \gamma_2$ if and only if M satisfies $(C1) - (C4)$. □

A.1 Extending to infinite menu complexity

In this section, we present the deferred proof for Corollary 4.5:

Corollary 4.5. *There exists a correlated prior distribution and ex-ante supply constraints for which every utility-optimal mechanism has infinite menu complexity.*

Proof. The proof is the continuous analogue of the proof for Theorem 4.1. We follow the same steps, while highlighting the necessary modifications. We first characterize the maximizers of the modified Lagrangian $L_{1,c}$ with the cost vector $(1, c) = (1, \frac{1}{2})$, and then apply Lemma 3.1 to obtain the ex-ante caps under which every utility-maximizing mechanism requires an infinite menu.

The continuous prior and menu properties. We first introduce the continuous analogue of the prior $\mathcal{D}$:

- The outside type has valuation $(0, 1)$ and probability $f_o = \frac{3}{4}$.

- The correlation function stays the same $\rho(v_1) = \frac{v_1^2}{4}$. However, the support of v_1 now ranges continuously over $(0, 1)$. For simplicity, as any inside type can be uniquely defined by its first-item value, we abuse notation and identify the type using v_1 instead of $\left(v_1, \frac{v_1^2}{4}\right)$.
- For the density of the inside component, we represent it by the function

$$f(v_1) = \frac{1 - v_1}{2},$$

where the density of the inside types integrates to $\int_0^1 f(v)dv = \frac{1}{4} = 1 - f_o$.

For an inside type $(v_1, \rho(v_1))$, let $(q_1(v_1), q_2(v_1), p(v_1))$ denote its assigned option, and write $u(v_1) = u_M(v_1, \rho(v_1))$.

We now move on to the constraints (C1)-(C4) that we want the optimal mechanism to have. The continuous construction retains analogues of (C1), (C2) and (C4). We will still show that the upper bound on L is realized only when the above three are all true. However, (C3) - the indifferences among adjacent types - is replaced by a consequence of the envelope theorem. More precisely, our new set of constraints is:[10]

(C1) The outside type gets item 2 for free $(0, 1, 0)$.

(C2) Inside types receive no allocation of item 2. More precisely, an inside type with value v_1 is assigned the menu option $(q_1(v_1), 0, p(v_1))$.

(C4) Every inside type is indifferent between its own assigned option and the one assigned to the outside type.

(C3') The alternative for (C3) is:

$$u'(v_1) = q_1(v_1) + \rho'(v_1)q_2(v_1) = q_1(v_1) + \frac{v_1}{2}q_2(v_1).$$

This is always true from IC via the envelope theorem. Fix an interior point $a \in (0, 1)$ at which u is differentiable. IC implies:

$$u(v_1) - (v_1 q_1(a) + \rho(v_1)q_2(a) - p(a)) \geq 0,$$

with equality at $v_1 = a$. Differentiate with respect to v_1, holding a fixed, and evaluate at $v_1 = a$; the derivative is zero at this minimizer.

Finally, we also have to find the intended optimal menu M^* that satisfies all the above constraints, as well as being IC and IR and having nonnegative payments. The resulting menu is:

$$M^* = \{(0, 1, 0), (0, 0, 0)\} \cup \left\{\left(\frac{v_1}{2}, 0, \frac{v_1^2}{4}\right)\right\}_{v_1 \in (0,1]}, \tag{18}$$

with the same tie-breaking rule: the outside type will get the free item-2 lottery, while every inside type v_1 will get their corresponding inside lottery.

Observe that as $T \to \infty$, the M^* for the finite construction converges to this menu. One can verify that this menu M^* satisfies all the stated properties.

[10] Throughout this proof, identities involving derivatives and allocation rules along the inside curve are understood to hold Lebesgue-almost everywhere on $(0, 1)$; we omit this qualification when there is no ambiguity. These exceptional sets have zero probability under the inside distribution and therefore do not affect expected utility or allocations. Statements about the outside type, which has positive probability, hold pointwise. Equality of continuous utility functions almost everywhere along the inside curve extends to the entire curve by continuity.

Maximizing L requires an infinite menu. The first step of the proof involves finding the upper bound on L that uses only utility. Using (C3'), we have that for every type v_1 with valuation $\left(v_1, \frac{v_1^2}{4}\right)$, its contribution to L can be bounded by

$$\begin{aligned} u(v_1) - q_1(v_1) - cq_2(v_1) &= u(v_1) - u'(v_1) - \frac{1 - v_1}{2} q_2(v_1) \\ &\leq u(v_1) - u'(v_1), \end{aligned}$$

where, for $0 < v_1 < 1$, equality holds if and only if $q_2(v_1) = 0$ (i.e., (C2) holds).

Taking the integral of this contribution over all v_1, we have:

$$\begin{aligned} L(M) &= f_o\big(u_o - q_1^o - cq_2^o\big) + \int_0^1 f(v_1)\big(u(v_1) - q_1(v_1) - cq_2(v_1)\big) dv_1 \\ &\leq f_o\big(u_o - q_1^o - cq_2^o\big) + \int_0^1 f(v_1)\big(u(v_1) - u'(v_1)\big) dv_1. \end{aligned}$$

The next step involves handling the integral term in the preceding bound using integration by parts (we use a telescoping sum in the finite proof):

$$\int_0^1 f(v_1)\big(u(v_1) - u'(v_1)\big)\, dv_1 = \int_0^1 \big(f(v_1) + f'(v_1)\big) u(v_1)\, dv_1 = -\frac{1}{2} \int_0^1 v_1 u(v_1)\, dv_1,$$

using $f(1) = 0$, $u(0) = 0$ and $f(v_1) + f'(v_1) = -\frac{v_1}{2}$ for all $0 \leq v_1 \leq 1$.

We proceed to using the IC constraints comparing each inside type's assigned option with the outside type's option:

$$u(v_1) \geq v_1 q_1^o + \rho(v_1) q_2^o - p^o,$$

and using this for the upper bound on L, we get:

$$L(M) \leq \gamma_1 q_1^o + \gamma_2 q_2^o + \gamma_p p^o,$$

where from our choice of prior,

$$\gamma_1 = -\frac{11}{12}, \qquad \gamma_2 = \frac{11}{32}, \qquad \gamma_p = -\frac{1}{2}.$$

Thus, we obtain the upper bound $L(M) \leq \frac{11}{32}$. In fact, $L(M^*)$ is also $\frac{11}{32}$, so M^* is a maximizer of L. Furthermore, equality requires that the optimal mechanism satisfies (C1),(C2) and (C4). In particular, any optimal mechanism satisfies $u(v_1) = \rho(v_1) = v_1^2/4$ and $q_2(v_1) = 0$, and also using the envelope theorem 2.1 we have $u'(v_1) = q_1(v_1)$. Hence, any optimal mechanism must satisfy the allocation identity $q_1(v_1) = \frac{v_1}{2}$ on $(0, 1)$, which is impossible for a menu with finitely many lotteries. Hence, no finite menu can maximize L.

Finishing the proof with the Lagrangian Caps Lemma. We set the ex-ante constraints to be exactly the expected allocations of M^* with the above prior:

$$\alpha_1 = \frac{1}{24}, \qquad \alpha_2 = \frac{3}{4}.$$

By the Lagrangian Caps Lemma 3.1, every utility-maximizing mechanism under these caps also maximizes L. However, as no menu with finitely many options can maximize L, every utility-maximizing mechanism must have infinite menu complexity. □

B Missing Proofs from Section 5

B.1 Choosing the prior that makes our curve optimal

Now that we have designed a utility curve r that no finite menu can implement as it's strictly convex on $[0, 2]$, we continue to backsolve by constructing a distribution $\mathcal{D}_1$, specified by its density f_1. Recall that in any structured menu, d is also the probability of receiving item 2 in its single free lottery option $(0, d, 0)$. We first choose f_1 to make r optimal among structured menus with $d = 1$.

Lemma B.1. *Suppose that $\mathcal{D}_2 = U(0, 1)$ and $\mathcal{D}_1$ satisfies the constraints (F1) and (F2) in Lemma 5.2, together with the following:*

(F3) *For $0 < v_1 < 2$, $f_1(v_1) > 0$ and*

$$(f_1(v_1) + f_1'(v_1))r(v_1) + cf_1(v_1) = 0;$$

(F4) *For $v_1 > 2$,*

$$f_1(v_1) + f_1'(v_1) = 0.$$

Then M^* maximizes L among all structured menus satisfying (S1) and (S2) with $d = 1$. Moreover, every maximizing menu M in this class satisfies

$$u_M(v_1, 0) = r(v_1) \qquad \forall\, 0 < v_1 < 2.$$

Proof. For convenience, we will restate the derived formula of L from (12), but now with $d = 1$:

$$L(M) = \frac{1}{2} - c + \int_0^\infty \Big[(f_1(v_1) + f_1'(v_1))\, \mathbb{E}_{v_2 \sim \mathcal{D}_2}[(u_1(v_1) - v_2)_+] + c\, f_1(v_1) \min\{1, u_1(v_1)\}\Big]\, dv_1.$$

Observe that the derivative of u_1 (the item-1 allocation along $v_2 = 0$) no longer appears; the integrand at v_1 depends only on the utility $u_1(v_1)$. Thus, if setting $u_1(v_1) = r(v_1)$ maximizes this integrand for all v_1, we have shown that M^* indeed maximizes L among structured menus with $d = 1$.

Fix v_1, and consider the integrand of L now as a function of $u_1(v_1)$ instead of v_1. This function is actually different based on whether $u_1(v_1) \geq 1$ or not. Since $\mathcal{D}_2 = U(0, 1)$, the integrand equals

$$\begin{cases} \dfrac{f_1(v_1) + f_1'(v_1)}{2} u_1(v_1)^2 + cf_1(v_1)u_1(v_1), & 0 \leq u_1(v_1) \leq 1, \\ (f_1(v_1) + f_1'(v_1)) \left(u_1(v_1) - \dfrac{1}{2}\right) + cf_1(v_1), & u_1(v_1) \geq 1. \end{cases}$$

Because of this discrepancy, we will move on with two cases according to the value of the target utility curve $r(v_1)$. For $0 < r(v_1) < 1$, which is equivalent to $0 < v_1 < 2$, we will show that $u_1(v_1) = r(v_1)$ is the unique maximizer of this function. For the other case of $r(v_1) > 1$, equivalently, $2 < v_1$, it already suffices to show the weaker statement that $r(v_1)$ is a maximizer. Note that the single point $v_1 = 2$ can be disregarded as it does not affect the integral.

Case 1: $0 < v_1 < 2$. We first compare the value of the integrand at $r(v_1)$ with its value at every other $u_1(v_1) \in [0, 1]$. The integrand in this case is a quadratic function of $u_1(v_1)$, whose derivative

at $r(v_1)$ is 0 thanks to (F3). Furthermore, this function is strictly concave, as the above derivation also gives:

$$f_1(v_1) + f_1'(v_1) = -\frac{cf_1(v_1)}{r(v_1)} < 0,$$

where we get $f_1(v_1) > 0$ also from (F3), and we will later set c to be positive. Therefore, the quadratic expression is uniquely maximized at $u_1(v_1) = r(v_1)$ on $[0, 1]$.

However, we still need to consider when $u_1(v_1) \geq 1$. Notice that in this range, the integrand is simply affine with negative slope, so its maximum on this range is attained at $u_1(v_1) = 1$. However, the two cases of the integrand agree when $u_1(v_1) = 1$, and the quadratic version of the integrand at 1 is already shown to be less than the value at $r(v_1)$. Thus, when $0 < v_1 < 2$, setting $u_1(v_1) = r(v_1)$ is still the unique maximizer over all possible values of $u_1(v_1)$.

Case 2: $2 < v_1$. Using (F4), we evaluate the integrand at all possible values of $u_1(v_1)$. For every $u_1(v_1) \geq 1$, the integrand equals the constant $cf_1(v_1)$. For $0 \leq u_1(v_1) < 1$, the integrand becomes:

$$\frac{f_1(v_1) + f_1'(v_1)}{2} u_1(v_1)^2 + cf_1(v_1)u_1(v_1) = cf_1(v_1)u_1(v_1) < cf_1(v_1).$$

Thus, every utility at least 1 is a maximizer of the integrand, including $r(v_1)$. □

Solving for f_1. We have finally found all the constraints we need f_1 to satisfy:

(F1) f_1 is continuous and piecewise smooth on $(0, \infty)$, with $v_1 f_1(v_1) \to 0$ both as $v_1 \to 0^+$ and as $v_1 \to \infty$, and $\mathcal{D}_1$ has a finite mean;

(F2) $f_1(v_1) + f_1'(v_1) \leq 0$ whenever $f_1(v_1)$ is differentiable.

(F3) For $0 < v_1 < 2$, $f_1(v_1) > 0$, as well as $(f_1(v_1) + f_1'(v_1))\, r(v_1) + c\, f_1(v_1) = 0$;

(F4) For $v_1 > 2$, $f_1(v_1) + f_1'(v_1) = 0$.

We start with (F3), which is essentially a differential equation. That means, for $0 < v_1 < 2$:

$$\frac{f_1'(v_1)}{f_1(v_1)} = -1 - \frac{c}{r(v_1)} = -1 - \frac{8c}{v_1(v_1+2)} = -1 - 4c\left(\frac{1}{v_1} - \frac{1}{v_1+2}\right),$$

whose solutions are $f_1(v_1) = Ce^{-v_1}\big((v_1+2)/v_1\big)^{4c}$ for a constant $C > 0$.

Furthermore, working with (F4) gives $f_1(v_1) \propto e^{-v_1}$ on $(2, \infty)$, and continuity at 2 fixes the constant. Thus (F3), (F4) and continuity determine f_1 up to the normalizing constant C:

$$f_1(v_1) = Ce^{-v_1} \begin{cases} \big((v_1+2)/v_1\big)^{4c}, & 0 < v_1 \leq 2, \\ 2^{4c}, & v_1 > 2. \end{cases}$$

Choosing the cost c. Conditions (F3) and (F4) now hold. We choose c to satisfy (F1).

First, observe that the exponential tail of f_1 already guarantees the part of (F1) concerning the limit as $v_1 \to \infty$. Thus, we only need to handle the case when v_1 is near zero, where $f_1(v_1) \approx C\,(2/v_1)^{4c}$. This is integrable, and satisfies $v_1 f_1(v_1) \to 0$ as required by (F1), exactly when $4c < 1$. Note that with our choice of r, we have:

$$r'(0) = \frac{1}{4},$$

meaning the condition $4c < 1$ is exactly

$$c < r'(0).$$

We take $c = 1/8$, which gives the final distribution $\mathcal{D}_1$ with density:

$$f_1(v_1) = Ce^{-v_1} \begin{cases} \sqrt{(v_1+2)/v_1}, & 0 < v_1 \le 2, \\ \sqrt{2}, & v_1 > 2, \end{cases} \tag{19}$$

where $C > 0$ is chosen so that $\int_0^\infty f_1 = 1$.

Checking the conditions. Finally, we verify that this choice of distribution satisfies all the constraints (F1)-(F4):

- (F1): f_1 is continuous and piecewise smooth on $(0, \infty)$; $v_1 f_1(v_1) \approx C\sqrt{2v_1} \to 0$ at zero, and $v_1 f_1(v_1) \to 0$ exponentially at infinity. Also, the mean of $\mathcal{D}_1$,

  $$\mathbb{E}[v_1] = C \int_0^2 \sqrt{v_1(v_1+2)}\, e^{-v_1}\, dv_1 + C\sqrt{2} \int_2^\infty v_1 e^{-v_1}\, dv_1,$$

  is finite.

- (F2): by construction,

  $$f_1(v_1) + f_1'(v_1) = \begin{cases} -\dfrac{c\, f_1(v_1)}{r(v_1)}, & 0 < v_1 < 2, \\ 0, & v_1 > 2, \end{cases} \tag{20}$$

  which is non-positive.

- (F3) and (F4) hold by construction.

B.2 The optimal free lottery for item 2 has $d = 1$

In Appendix B.1, we construct a prior f_1 under which M^* maximizes L among structured menus with $d = 1$. This section shows that this restriction on d is without loss, as we will prove that for every structured menu M with $0 \le d < 1$, it does strictly worse than M^*: $L(M) < L(M^*)$.

Lemma B.2. *For the distributions and costs chosen in Appendix B.1, every structured menu M satisfying (S1) and (S2), with free item-2 lottery $(0, d, 0)$, satisfies*

$$L(M) \le dL(M^*).$$

Consequently, if $d < 1$, $L(M) < L(M^*)$.

Proof. We start with $L(M^*)$: substituting $d = 1$ and $u_1 = r$ into (12) and using properties (F3) and (F4) of the chosen item-1 prior gives

$$L(M^*) = \frac{1}{2} - c + \frac{c}{2}\int_0^2 f_1(v_1) r(v_1)\, dv_1 + c\int_2^\infty f_1(v_1)\, dv_1 \geq \frac{1}{2} - c = \frac{3}{8}. \tag{21}$$

We first handle the edge case of $d = 0$. Substituting $d = 0$ into (12) and simplifying gives

$$L(M) = \int_0^\infty \left(f_1(v_1) + f_1'(v_1)\right) u_1(v_1)\, dv_1 \leq 0,$$

by (F2) and $u_1(v_1) \geq 0$.

With the edge case out of the way, we proceed by fixing any d in the range $(0, 1]$. For convenience, we restate the integrand of $L(M)$:

$$\left(f_1(v_1) + f_1'(v_1)\right)\left(\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1) - dv_2)_+]\right) + cf_1(v_1)\min\left\{d, u_1(v_1)\right\},$$

where by our choice of f_1, we know that:

$$f_1(v_1) + f_1'(v_1) = \begin{cases} -\dfrac{cf_1(v_1)}{r(v_1)}, & 0 < v_1 < 2, \\ 0, & v_1 > 2. \end{cases}$$

We will now bound this integrand separately in the two cases: $0 < v_1 < 2$ and $2 < v_1$. The second case is significantly easier: the integrand becomes $cf_1(v_1)\min\left\{d, u_1(v_1)\right\}$, which is at most $cdf_1(v_1)$. We now work with v_1 in the range of $(0, 2)$, where the formula for the integrand is:

$$\begin{aligned} &\left(f_1(v_1) + f_1'(v_1)\right)\left(\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1) - dv_2)_+]\right) + cf_1(v_1)\min\left\{d, u_1(v_1)\right\} \\ &\quad = -\frac{cf_1(v_1)}{r(v_1)}\left(\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1) - dv_2)_+]\right) + cf_1(v_1)\min\left\{d, u_1(v_1)\right\} \\ &\quad = cf_1(v_1)\left[\min\left\{d, u_1(v_1)\right\} - \frac{\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1) - dv_2)_+]}{r(v_1)}\right]. \end{aligned} \tag{22}$$

As $\mathcal{D}_2 = U(0, 1)$, we evaluate the expectation terms in the above equation separately for $u_1(v_1) \leq d$ and $u_1(v_1) \geq d$ to get:

$$\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1) - dv_2)_+] = \frac{\min\{d, u_1(v_1)\}^2}{2d} + (u_1(v_1) - d)_+.$$

Substituting this expression into (22) and completing the square gives

$$\begin{aligned} &cf_1(v_1)\left[\min\left\{d, u_1(v_1)\right\} - \frac{\mathbb{E}_{v_2\sim\mathcal{D}_2}[(u_1(v_1) - dv_2)_+]}{r(v_1)}\right] \\ &\quad = \frac{cf_1(v_1)}{r(v_1)}\left[r(v_1)\min\left\{d, u_1(v_1)\right\} - \frac{\min\{d, u_1(v_1)\}^2}{2d} - (u_1(v_1) - d)_+\right] \\ &\quad = \frac{cf_1(v_1)}{r(v_1)}\left[\frac{d\, r(v_1)^2}{2} - \frac{\left(\min\{d, u_1(v_1)\} - d\, r(v_1)\right)^2}{2d} - (u_1(v_1) - d)_+\right] \end{aligned}$$

$$= cf_1(v_1)\left[\frac{d\,r(v_1)}{2} - \frac{\big(\min\{d, u_1(v_1)\} - d\,r(v_1)\big)^2}{2d\,r(v_1)} - \frac{(u_1(v_1)-d)_+}{r(v_1)}\right].$$

Observe that both terms getting subtracted from $\frac{cdf_1(v_1)r(v_1)}{2}$ are nonnegative, so we have found an upper bound for the integrand when $0 < v_1 < 2$.

In conclusion, we have found that the integrand is bounded by $\frac{cdf_1(v_1)r(v_1)}{2}$ when $0 < v_1 < 2$, and by $cdf_1(v_1)$ when $2 < v_1$. Using these two bounds in the expression for $L(M)$ gives

$$L(M) \le d\left[\frac{1}{2} - c + \frac{c}{2}\int_0^2 f_1(v_1)r(v_1)\,dv_1 + c\int_2^\infty f_1(v_1)\,dv_1\right] = dL(M^*).$$

As $L(M^*) > 0$, whenever $d < 1$, we have $L(M) \le dL(M^*) < L(M^*)$.

□

C Menu Complexity with Multiple Buyers

We extend our menu-complexity lower bounds to three buyers with ex-post supply constraints. The reduction introduces two deterministic buyers whose values realize the allocation costs in L_{λ_1,λ_2}. Here, we use the definition of multi-buyer menu size from Hart and Nisan [25, Appendix A.4].

For a mechanism M with n buyers and two items, write $v^i = (v_1^i, v_2^i)$ for buyer i's report and v^{-i} for the other buyers' reports. Let

$$x^{M,i}(v^i, v^{-i}) = \big(x_1^{M,i}(v^i, v^{-i}), x_2^{M,i}(v^i, v^{-i})\big)$$

denote buyer i's allocation probabilities, averaging only over the mechanism's randomization.

Definition 1 (Multi-buyer menu size)**.** The mechanism M has menu size at most K if fixing the other buyers' reports leaves each buyer at most K distinct nonzero allocation vectors. Formally, for every $i \in [n]$ and $v^{-i} \in (\mathbb{R}^2_{\ge 0})^{n-1}$,

$$\big|\big\{x^{M,i}(v^i, v^{-i}) : v^i \in \mathbb{R}^2_{\ge 0}\big\} \setminus \{(0,0)\}\big| \le K.$$

In our construction, all buyers other than buyer 1 have deterministic types, denoted by $\bar{v}^{-1}$. Thus, averaging over the other buyers' types is the same as fixing their reports to $\bar{v}^{-1}$. Consequently, buyer 1's interim mechanism has exactly the same allocation probabilities and payments as its ex-post mechanism at these fixed reports. Since the multi-buyer menu size takes the maximum over buyers and fixed reports of the others, it is at least the menu size of this induced single-buyer mechanism. In particular, if the induced single buyer mechanism requires infinitely many options, then the joint mechanism has infinite menu complexity.

Proposition C.1. *Fix a prior $\mathcal{D}$ and nonnegative costs λ_1, λ_2. Suppose M^* attains a finite maximum of L_{λ_1,λ_2}, and set $\alpha_j = \mathbb{E}[x_j^{M^*}]$ for $j = 1, 2$. If every maximizer of L_{λ_1,λ_2} requires K distinct non-null menu options, then there is an instance with three independent unit-demand buyers and two single-copy items in which every mechanism that maximizes the total expected buyer utility also has menu size of at least K. This holds under either Bayesian or dominant-strategy incentive compatibility, with interim IR and nonnegative burned payments.*

Proof. Keep the original buyer with prior $\mathcal{D}$. We interpret the costs λ_1 and λ_2 as additional competition for the items: add a buyer who values only item 1, at λ_1, and another one who values only item 2, at λ_2.

Consider the following multi-buyer mechanism: run M^* for the original buyer and give each remaining item to its designated buyer for free. The added buyers then receive their items with probabilities $1-\alpha_1$ and $1-\alpha_2$, respectively. This mechanism is feasible and dominant-strategy incentive compatible: the original buyer faces M^*, and the added buyers' reports are ignored. Its total expected utility is

$$\mathbb{E}[u_{M^*}] + \lambda_1(1-\alpha_1) + \lambda_2(1-\alpha_2) = L_{\lambda_1,\lambda_2}(M^*) + \lambda_1 + \lambda_2.$$

Now consider any Bayesian incentive-compatible joint mechanism, and let M be the original buyer's induced interim mechanism. Since the other buyers have deterministic types, M is an IC, IR single-buyer mechanism. Write $a_j = \mathbb{E}[x_j^M]$. Supply feasibility and nonnegative payments bound the added buyers' utilities by $\lambda_1(1-a_1)$ and $\lambda_2(1-a_2)$. Thus the total expected utility satisfies

$$\begin{aligned} U_{\text{joint}} &\le \mathbb{E}[u_M] + \lambda_1(1-a_1) + \lambda_2(1-a_2) \\ &= \lambda_1 + \lambda_2 + L_{\lambda_1,\lambda_2}(M) \le \lambda_1 + \lambda_2 + L_{\lambda_1,\lambda_2}(M^*). \end{aligned}$$

The last line follows from the optimality of M^* for the modified Lagrangian L_{λ_1,λ_2}. Our proposed mechanism attains this bound, so every utility-optimal joint mechanism must induce a maximizer of L_{λ_1,λ_2} for the original buyer. By assumption, this induced mechanism must have at least K non-null options. Therefore, every optimal joint mechanism has menu complexity at least K. □

Applying Proposition C.1 to the construction in the proof of Proposition 4.3 with $(\lambda_1, \lambda_2) = (1, 1/2)$ shows that for every $T \ge 3$, there is an instance with three unit-demand buyers and two items in which every utility-optimal mechanism has menu size at least T. Applying the proposition to Theorem 5.6, with $(\lambda_1, \lambda_2) = (1, 1/8)$, gives an instance with three unit-demand buyers and two items in which every utility-optimal mechanism has infinite menu complexity, even when item values are independent across buyers and items.