The Simple Economics of Approximately Optimal Auctions

Saeed Alaei Hu Fu Nima Haghpanah Jason Hartline Azarakhsh Malekian

First draft: June 14, 2012.

Abstract

The intuition that profit is optimized by maximizing marginal revenue is a guiding principle in microeconomics. In the classical auction theory for agents with quasi-linear utility and single-dimensional preferences, Bulow and Roberts (1989) show that the optimal auction of Myerson (1981) is in fact optimizing marginal revenue. In particular Myerson's virtual values are exactly the derivative of an appropriate revenue curve.

This paper considers mechanism design in environments where the agents have multi-dimensional and non-linear preferences. Understanding good auctions for these environments is considered to be the main challenge in Bayesian optimal mechanism design. In these environments maximizing marginal revenue may not be optimal, and furthermore, there is sometimes no direct way to implementing the marginal revenue maximization mechanism. Our contributions are three fold: we characterize the settings for which marginal revenue maximization is optimal, we give simple procedures for implementing marginal revenue maximization in general, and we show that marginal revenue maximization is approximately optimal. Our approximation factor smoothly degrades in a term that quantifies how far the environment is from an ideal one (i.e., where marginal revenue maximization is optimal).

1 Introduction

Marginal revenue plays a fundamental role in microeconomic theory. For example, a monopolist in a market where the demand induces concave revenues and convex costs (in the amount of the good produced) optimizes her profit by producing an amount where the marginal revenue is equal to the marginal cost. Moreover this central economic principle also governs classical auction theory. Myerson (1981) characterizes profit maximizing single item auction as formulaically optimizing the *virtual value* of the winner; Bulow and Roberts (1989) reinterpret Myerson's virtual value as the marginal revenue of a certain concave revenue curve.

The Myerson-Bulow-Roberts approach applies to environments where the agents have risk-neutral quasi-linear single-dimensional preferences, i.e., where an agent's utility is given by her value for service less her payment. In this paper we generalize this marginal revenue approach to general agent preferences which include multi-dimensionality, risk aversion, and budgets. One of the main difficulties of this generalization is that the marginal revenue approach may not continue to be optimal. We give a structural, and computationally tractable, characterization of the preferences for which optimizing marginal revenue is optimal. We show that the heuristic of marginal revenue maximization is always implementable and is always approximately optimal. Furthermore, the approximation factor of the marginal revenue mechanism degrades smoothly in a specific quantity that describes how far an agent's preference is from satisfying the characterization.

In Myerson-Bulow-Roberts environments an agent has a value from a known range drawn from a known distribution. For this agent a revenue curve can be constructed from the following thought experiment. Given a constraint q on the ex ante probability that the agent can be served, what single-agent *lottery pricing*[1] maximizes expected revenue while serving the agent with the given probability q (where expectation and probability is taken over the randomization in the lottery pricing and the randomization in the agent's value being drawn from the distribution). The *revenue curve* for this agent is given by this expected revenue as a function of the ex ante probability constraint q. The marginal revenue is the derivative of this revenue curve with respect to q. For example, if the agent's value is drawn from uniformly from $[0, 1]$, then the optimal lottery that sells with probability q is to offer the agent the item with probability one at price $1 - q$ and probability zero at price zero, the expected revenue obtained is $R(q) = q - q^2$; the marginal revenue is $R'(q) = 1 - 2q$.

Similar revenue curves can be defined analogously for agents with multi-dimensional or non-quasi-linear preferences. Consider the case of an agent with quasi-linear unit-demand preferences over several distinct items for sale. This agent's preference is given by a value for each item and the agents utility is the value for the item she receives less the amount she is required to pay. The question of optimizing lottery pricings was considered for general distributions by Briest et al. (2010) (optimization) and for product distributions by Chawla et al. (2010b) (approximation); while both of these results assume that there is no constraint on the probability the agent is served, i.e., $q = 1$, it is not hard to adapt them to general ex ante constraint q. The same kind of calculation can be applied for agents with budgets or that are risk-averse, and furthermore, the budgets and risk-preferences can be private to the agent.

Obtaining concise economic characterization of optimal mechanisms with multi-dimensional preferences is one of the big challenge areas in mechanism design. This area is challenging because it is not generally possible to order the possible preferences of the agents based on strength. Notice

[1] A lottery pricing is a menu that assigns a price to every probability of serving the agent.

that in the single-dimensional example given above an agent preference with a higher value is stronger than a preference with a smaller value. In any revenue-optimal auction an agent is at least as likely to win when she has a higher value. In contrast, when an agent has a private value and a private budget, it is not clear whether an agent is stronger with a high value and a low budget or a low value and a high budget. Often which preference is stronger depends on the context, i.e., the other agents and the supply or feasibility constraint on the auction environment. For instance, when there is not much competition, budgets may not be binding; whereas with lots of competition values may not be binding.

Consider a single agent and a lottery pricing. An agent with a given preference will choose an outcome that maximizes her utility. This choice induces a mapping from the space of agent preferences to the probability that the agent is served. From these probabilities we can infer an ordering on the relative strengths of the different possible preferences of the agent. A preference is strictly stronger than another preference if the probability that the agent is served is strictly larger. Given a distribution over the preferences of this agent we can quantify the relative strength of a particular preference as the probability that a random preference from this distribution is stronger than the given preference. We will refer to this relative strength as *quantile*, denoted q, where low quantile is strong and high quantile is weak. For simplicity of exposition assume that strengths being equal is a measure zero event and therefore the distribution of quantile of a random preference is always uniform on $[0, 1]$. Given this mapping from preference to quantile and the mapping from preference to probability of service we can consider the *allocation rule* of the lottery pricing as a function from quantile to probability of service. Notice that by definition this allocation rule is monotone non-increasing in quantile.

As any lottery pricing induces an allocation rule, we can ask the following optimization question. Given an allocation rule as a constraint, find the revenue-optimal lottery pricing that satisfies this allocation constraint. In fact, the revenue curve construction above can be viewed as a special case of this where the constraint of ex ante probability of service being q corresponds to the allocation constraint that preferences stronger than q are served and preferences weaker than q are not served. I.e., the allocation constraint, as a function, steps from one to zero at q. Notice that any allocation rule can be expressed as a convex combination of step functions and, therefore, one approach to solving this optimization is to solve the problem for each step function and then take the convex combination of the solutions. Clearly the revenue of such an approach is given simply by the same convex combination of the revenues associated with each step function via the revenue curve. It can be shown that in expectation this revenue is equal to the expected marginal revenue as given by the revenue curve. Note, however, that the optimal revenue for such an allocation constraint may generally be larger than this quantity.

The Myerson-Bulow-Roberts analysis shows that for the special case of single-dimensional preferences, the convex combination approach is always optimal (for any allocation constraint). We show that it continues to be optimal in some multi-dimensional environments; it is never optimal with non-quasi-linear preferences such as budgets. When it is optimal, additional nice structure can be observed. In particular the mapping from preferences to quantiles is constant across different allocation constraints. This implies that, unlike the general case, relative strength of two types does not depend on context. Because of this the following marginal revenue mechanism is well defined. First, calculate quantiles from preferences as given by the context-free mapping. Second, calculate marginal revenues for each of these quantiles. Third, serve the agents to maximize the total marginal revenue. Prices, for instance, can be calculated via the standard approach: find the weakest quantile a winner must have had to still win, and charge her based on the step function mechanism for that quantile.

The main contribution of this paper is to show that even when it is not optimal this marginal revenue approach is approximately optimal and, furthermore, there is a well defined mechanism that implements marginal revenue maximization. We address the implementation question in two ways. The first approach assumes a structural property on optimal mechanisms for step constraints and gives a way to sample quantiles and a way to calculate the minimum quantile a winner must have to still win that is internally consistent. This structural property on optimal mechanisms for step constraints is satisfied for a number of interesting preferences, e.g., preferences with private value and public budgets. The second approach applies to the general case but is more complex. Both approaches are computationally simple. We address the approximation question in two ways as well. First, we show that the approximation factor smoothly degrades as the single-agent problems become less ideal. For unit-demand preferences we upper bound this degradation, and thus the approximation, by a factor of four. Second, we give upper bounds on the approximation factor when the feasibility constraint imposed by the environment has nice structural properties. In particular, the "correlation gap" approach (cf. Yan, 2011) gives an upper bound of $e/(e-1)$ on the approximation factor for single-item and matroid feasibility constraints; for any downward-closed feasibility constraint we derive an upper bound on the approximation factor that is logarithmic in the number of agents.

These results are significant for several reasons. First, the marginal revenue approach provides a lot of economic intuition for optimal and approximately optimal auctions. Second, the marginal revenue mechanism smoothly trades off performance from ideal environments where it is optimal to non-ideal environments were it is approximately optimal. Third, it gives approaches for approximately optimal mechanism design in general environments that are beyond the reach of previous algorithmic or structural characterizations. Finally, it can be viewed as a reduction from the mechanism design to algorithm design that is compatible with approximation algorithms, e.g., via the approach of Hartline and Lucier (2010).

All of this work applies to general objectives. E.g., when agents have budgets, welfare maximization is a challenging problem as well. We can simply replace "revenue" with any other objective and use the marginal revenue approach with that objective. The only requirement on the objective is that it is additively separable across the agents.

Related Work. In mechanism design problems for agents with non-linear preferences, Matthews (1983, 1984) and Maskin and Riley (1984) developed the approach of optimizing interim allocation rules subject to feasibility. Border (1991, 2007) characterized the necessary and sufficient properties of interim feasibility. Alaei et al. (2012) and Cai et al. (2012) give computationally efficient algorithms for checking and optimizing over interim feasibility in matroid and single-item environments, respectively. Our work can be viewed as approximately extending these latter results to downward-closed environments.

Reductions from multi- to single-agent problems in multi-dimensional mechanism design have been studied by Alaei (2011) and Alaei et al. (2012). Alaei (2011) considers multi-unit combinatorial auctions and provides constant-approximation reductions to the single agent problem. Alaei et al. (2012) defined service constrained environments and provided optimal reductions to single-agent problems in matroid environments. Our reduction is to a simpler problem of maximizing revenue subject to an ex ante constraint, whereas Alaei et al. (2012) reduces the problem to maximizing revenue subject to any allocation constraint.

Our approximation bounds make use of two techniques from the literature. Yan (2011) developed a "correlation gap" approach to prove that sequential pricing is a good approximation to optimal mechanisms in single-dimensional matroid environments. Our work extends the use of correlation gap to multi-dimensional agent problems. Furthermore, our approach generalized beyond the reach of correlation gap based approaches and admits approximation bounds for general downward-closed environments. Chawla et al. (2010a,b) together show that a simple unit-demand item pricing approximates the optimal single-agent mechanism (with no allocation constraint). We extend their technique to single-agent problems with allocation constraints and use this extension to derive 4-approximation mechanisms for unit-demand agents in downward-closed service constrained environments.

2 Preliminaries

Bayesian mechanism design. An agent has a private type t from type space T drawn from distribution F with density function f. She may be assigned outcome w from outcome space W. This outcome encodes what kind of service the agent receives and any payments she must make for the service. In particular the payment specified by an outcome w is denoted by $\text{Payment}(w)$. The agent has a von Neumann–Morgenstern utility function: for type and deterministic outcome w her utility is $u(t,w)$, when w is drawn from a distribution her utility is $\mathbf{E}_w[u(t,w)]$.[2] We will extend the definition of the utility function to distributions over outcomes $\Delta(W)$ linearly. For a random outcome w from a distribution, $\text{Payment}(w)$ will denote the expected payment.

There are n agents indexed $\{1,\ldots,n\}$ and each agent i may have her own distinct type space T_i, utility function u_i, etc. A *direct revelation* mechanism takes as its inputs a profile of types $\mathbf{t} = (t_1,\ldots,t_n)$, and then outputs for each agent i an outcome $\tilde{w}_i(\mathbf{t})$. $\tilde{w}_i : T_1 \times \cdots \times T_n \to \Delta(W_i)$ is the ex post outcome rule of the mechanism. Agent i with type t_i, as the other agents' types are distributed over T_{-i}, faces an *interim outcome rule* $\tilde{w}_i(t_i)$ distributed as $\tilde{w}_i(t_i,\mathbf{t}_{-i})$ with each $t_j \sim F_j$ for $j \neq i$. We say that a mechanism is *Bayesian incentive compatible* if

$$u_i(t_i,\tilde{w}_i(t_i)) \geq u_i(t_i,\tilde{w}_i(t'_i)), \quad \forall i, \forall t_i, t'_i \in T_i. \tag{BIC}$$

A mechanism is *interim individually rational* if

$$u_i(t_i,\tilde{w}_i(t_i)) \geq 0, \quad \forall i, \forall t_i \in T_i. \tag{IIR}$$

The mechanism designer seeks to optimize an objective subject to (BIC), (IIR), and ex post feasibility. We consider the objective of expected revenue, i.e., $\mathbf{E}_{\mathbf{t}}[\sum_i \text{Payment}(\tilde{w}_i(t_i))]$; however, any objective that separates linearly across the agents can be considered. Below we discuss the mechanism's feasibility constraint.

Service constrained environments. In a *serviced constrained environment* the outcome w provided to an agent is distinguished as being a *service* or *non-service* outcomes, respectively with $\text{Alloc}(w) = 1$ or $\text{Alloc}(w) = 0$. There is a feasibility constraint restricting the set of agents that may be simultaneously served; there is no feasibility constraint on how an agent is served. With respect to the feasibility constraint any outcome $w \in W$ with $\text{Alloc}(w) = 1$ is the same. For

[2] This form of utility function allows for encoding of budgets and risk aversion; we do not require quasi-linearity.

example, payments are not constrained by the environment. An agent may have multi-dimensional and non-linear preferences over distinct service and non-service outcomes.

From least rich to most rich, standard service constrained environments are *single-unit environments* where at most one agent can be served, *multi-unit environments* where at most a fixed number of agents can be served, *matroid environments* where the set of agents served must be the independent set of a given matroid, *downward-closed environments* where the set of agents served can be specified by an arbitrary set systems for which subsets of a feasible set are feasible. and *general environments* where the feasible subsets of agents can be given by an arbitrary set system that may not even be downward closed.

Allocation rules. Given an outcome rule $\tilde{w} : T \to \Delta(W)$, we define an *allocation rule* $\tilde{x} : T \to [0, 1]$ by $\tilde{x}(t) = \mathbf{E}[\text{Alloc}(\tilde{w}(t))]$. We call $\tilde{x}$ an interim allocation rule if it is derived from an interim outcome rule $\tilde{w}$. In service constrained environments, an allocation rule contains all the information of an outcome rule that is relevant for feasibility considerations.

An agent's allocation rule induces a natural ordering on his type space based on the probability each type is served: a stronger type has a higher probability of being served than a weaker type. We measure the strength of a type by what we call *quantile* which is the probability that a random type from the distribution F is stronger than it.

Definition 1. An allocation rule $\tilde{x} : T \to [0, 1]$ induces a partial order on T; extend it to a total order by tie-breaking via any total order on T; and denote this total order by $>_{\tilde{x}}$. The *quantile* of $t \in T$ is the probability that a random type $t' \sim F$ is allocated with higher probability, denoted $\text{Quant}(t) = \mathbf{Pr}_{t' \sim F}[t' >_{\tilde{x}} t]$ with inverse $\text{Type}(\cdot)$ which is well defined because the mapping is one-to-one. The *normalized allocation rule* is $x(q) = \tilde{x}(\text{Type}(q))$; the *cumulative allocation rule* is the integral of the normalized allocation rule $X(q) = \int_0^q x(q)\, dq$.

Note that the normalized allocation rule x is bounded on $[0, 1]$ and monotone non-increasing and the cumulative allocation rule X is continuous, concave, and non-decreasing. For discrete type spaces the cumulative allocation rule is piecewise linear.

Definition 2. Allocation rule $\tilde{x}_a$ *dominates* allocation rule $\tilde{x}_b$, denoted $\tilde{x}_a \succeq \tilde{x}_b$, if type space T and distribution F induce cumulative allocation rules X_a and X_b that satisfy $X_a(q) \geq X_b(q)$ for all $q \in [0, 1]$ and $X_a(1) = X_b(1)$.

As a matter of notational consistency rules mapping type to an outcome or allocation such as $\tilde{w}$ and $\tilde{x}$ will always be denoted with a tilde. Rules for quantiles will always be denoted as x without a tilde.

Constrained lottery pricings. Every outcome rule $\tilde{w}(\cdot)$ induces some revenue $\mathbf{E}_{t \sim F, \tilde{w}(t)}[\text{Payment}(\tilde{w}(t))]$ and some normalized allocation rule $x(\cdot)$. Dominance as defined above allows us to consider a normalized allocation rule as a constraint and optimize revenue subject to incentive compatibility and this constraint. Allocation rules as constraints will generally be denoted as y with cumulative allocation rule Y. The suggested optimization is to find the $\tilde{w}$ that induces the x that is dominated by y and obtains the largest expected revenue.

$$
\begin{aligned}
\max_{\tilde{w}}: \quad & \mathbf{E}_{t,\tilde{w}(t)}\left[\text{Payment}(\tilde{w}(t))\right] && \text{(LotP)}\\
\text{s.t.} \quad & \mathbf{E}_{t,\tilde{w}(t)}\left[\text{Alloc}(\tilde{w}(t)) \;\mid\; t \in S\right] \mathbf{Pr}\left[t \in S\right] \leq Y(f(S)), \qquad \forall S \subset T\\
& \mathbf{E}_{t,\tilde{w}(t)}\left[\text{Alloc}(\tilde{w}(t))\right] = Y(1),\\
& \tilde{w} \text{ is IC and IR.}
\end{aligned}
$$

We denote the outcome rule that optimizes this program by $\tilde{w}^* = \text{Outcome}(y)$ and its revenue by $\text{Rev}(y) = \mathbf{E}_{t,\tilde{w}^*(t)}[\text{Payment}(\tilde{w}^*(t))]$. The feasibility constraint is equivalent to the cumulative allocation rule for $\tilde{w}^*$ satisfying $X^*(q) \leq Y(q)$ for all q with equality at $q = 1$. This is because for sets S with cumulative probability $f(S) = q$, the left-hand side is maximized by choosing the q strongest types and for this choice it is exactly $X^*(q)$.

Revenue and marginal revenue. In this paper we focus on the use of marginal revenue in designing approximately optimal auctions. In the classical single item auctions, marginal revenues are the derivatives of the *revenue curve*, a curve of $R(q)$ for $q \in [0, 1]$, where $R(q)$ is the optimal revenue one may obtain if the ex ante probability of serving the agent is exactly q. Since each agent's type space is single dimensional in this setting and admits an ordering, there is a natural mapping from type space to $[0, 1]$, the quantile space. The marginal revenue $R'(q) = \frac{d}{dq}R(q)$ corresponding to a type is then the derivative of R at the quantile q corresponding to the type. Bulow and Roberts (1989) showed that $R'(q)$ is equal to the (ironed) virtual valuation of Myerson (1981) who previously showed that the problem of optimal mechanism design reduces to virtual valuation maximization.

The revenue curve concept generalizes naturally to service constrained environments where $R(q)$ represents the optimal revenue obtainable by any incentive compatible mechanism subject to the constraint that the ex ante probability of serving the agent is exactly q. This optimization problem is equivalent to solving (LotP) for the *q-step constraint* $y^q(\cdot)$ which is a "reverse step function" from 1 to 0 at q. We denote this revenue by $R(q) = \text{Rev}(y^q)$, the outcome rule that achieves it by $\text{Outcome}(y^q) = \tilde{w}^q$, and the induced allocation rule by $\tilde{x}^q$. Notice that the normalized allocation rule x^q may not equal y^q as y^q only provides an upper bound (via its cumulative allocation rule) on x^q. We refer to the resulting mechanism as the *q-step mechanism* and denote it by $\mathcal{M}^q$.

The marginal revenue for a given allocation constraint y can be defined as follows. View allocation constraint y, a monotone non-increasing function, as a convex combination of the q-step allocation constraints (i.e., the step constraints form a basis for any monotone allocation constraint). The coeficient on the q-step function is $-y'(q) = -\frac{d}{dq}y(q)$. The revenue from this convex combination is exactly $-\int_0^1 R(q)y'(q)\,dq$ which is equal to $[R(q)y(q)]_0^1 + \int_0^1 R'(q)y(q)\text{d}q$. Often the additive term $[R(q)y(q)]_0^1 = 0$: for type spaces with bounded utility $R(0) = 0$, for downward closed environments we are free to set $y(1) = 0$, and for non-downward closed environments and type spaces with a zero utility type $R(1) = 0$. We therefore refer to this revenue as the *marginal revenue* of y and denote it by $\text{MR}(y)$.

Clearly, for any allocation constraint y, the optimal revenue $\text{Rev}(y)$ is at least the marginal revenue $\text{MR}(y)$. The space of mechanisms is convex therefore optimal revenues $\text{Rev}(y)$ and $R(q)$ are concave.

Optimal revenue and marginal revenue. Consider any procedure that maps a profile of quantiles $\mathbf{q} = (q_1, \ldots, q_n)$ to a feasible allocation. With quantiles drawn uniformly from $[0, 1]$, such a procedure induces a profile of interim allocation rules $(y_1, \ldots, y_n)$. Such a profile of allocation rules is *interim feasible* because there exists a mechanism that induces it. We can view each allocation rule in this profile as a constraint for a single agent problem and optimize single-agent revenue. I.e., for constraint y_i the optimal revenue is $\mathrm{Rev}(y_i)$. The mechanism that optimizes the cumulative revenue from each agent subject to interim feasibility gives the *optimal revenue*.[3]

Consider the procedure for selecting a feasible allocation for a profile of quantiles by mapping each agent i's quantile q_i to marginal revenue $R_i'(q_i)$ and then selecting the feasible set that maximizes the cumulative marginal revenue (i.e., summed over all agents served). As above, this procedure induces a profile of allocation constraints, denoted $(y_1^{MR}, \ldots, y_n^{MR})$, and for each agent i a marginal revenue $\mathrm{MR}(y_i^{MR})$ is generated. We define the expected cumulative marginal revenue of this process as the *optimal marginal revenue*.

Downward closure. In downward closed environments we can relax the equality constraint $\mathbf{E}_{t,\tilde{w}(t)}[\mathrm{Alloc}(\tilde{w}(t))] = Y(1)$ of (LotP); however, for expositional convenience we will instead encorporate downward closure into the outcome space. To designate the fact that an agent can always be excluded from service we will expand the outcome space by duplicating all non-service outcomes and labeling the duplicates as service outcomes. Representing downward closure in this manner means that optimal revenue is monotone with respect to dominance. I.e., if y_a dominates y_b then $\mathrm{Rev}(y_a) \geq \mathrm{Rev}(y_b)$. Likewise, the revenue $R(q)$ is monotone increasing in q.

3 Context freedom, revenue linearity, and orderability

In the classical single-dimensional theory of optimal auctions, the optimal auction is given by a virtual value function for each agent that maps her type to a real number and the optimal auction is the one that selects as winners the feasible set of agents with the highest virtual surplus. This virtual valuation function is in fact the marginal revenue. Of course, the marginal revenue is defined in the quantile space, but with single-dimensional agents there is a natural mapping from type (i.e., value) to the quantile space: in the order of value. Importantly an agent's virtual value is a function only of her type and distribution over type space; it is not dependent on the context (the other agents or the feasibility constraint).

There are a few important consequences of this form of optimal auction. First, the optimal auction deterministically chooses which set of agents to serve. As the original optimal auction design problem allowed for randomized mechanisms, it is interesting that randomization is not necessary. Second, the optimal auction is dominant strategy incentive compatible. Meaning: no matter what the actions of the other agents are, an agent (weakly) prefers truthtelling to any other action. Recall, that our original optimization problem was looking for optimal BIC mechanism. It is interesting that we get the stronger dominant strategy solution concept at no cost to the objective. A final observation about single-dimensional optimal auctions is that the expected revenue is equal to the expected virtual surplus (by definition: equal to the expected sum of marginal revenues).

An agent is described by a type space and distribution over it. In this section we characterize the kinds of agents for which optimal auctions behave the same way as for single dimensional

[3] For example, Alaei et al. (2012) give computationally efficient algorithms for solving this problem in matroid environments.

agents. Our characterization will be based on properties of the single-agent constrained lottery pricing problem. Such a problem is specified by an allocation constraint y; its optimal revenue is $\mathrm{Rev}(y)$ and the outcome rule that gives this revenue is $\mathrm{Outcome}(y)$.

Definition 3. A single-agent problem is *revenue linear* if the optimal revenue $\mathrm{Rev}(y)$ is linear in y with respect to the class of monotone allocation constraints.

In section A we show a simple, fast algorithm that tests if an agent is revenue linear.

Remark 1. *Recall that we use the reverse step functions for space of normalized allocation rules. An easy consequence from linearity w.r.t. this basis is:*

$$\mathrm{Rev}(y) = \mathrm{Rev}\left(-\int_0^1 y^q \,\mathrm{d}y(q)\right) = -\int_0^1 R(q)\,\mathrm{d}y(q) = \mathrm{MR}(y). \tag{1}$$

The second equality follows from linearity and the definition of $R(q) = \mathrm{Rev}(y^q)$*, and the last is the definition of* $\mathrm{MR}(y)$*. Therefore, when we have linearity of* $\mathrm{Rev}(\cdot)$*, the optimal revenue of a normalized allocation rule is equal to the corresponding marginal revenue. It is easy to see that the converse is also true.*

If a characterization holds analogously to the Myerson-Bulow-Roberts characterization of single-dimensional optimal mechanisms, i.e., there is a context-free virtual value function mapping types to real numbers and the revenue optimal mechanism serves the agents with the highest virtual surplus (sum of virtual values), then the virtual valuations induce an ordering on all types, and the optimal auction respects this ordering. The following definition captures this feature.

Definition 4. A single-agent problem is *orderable* if there is an equivalence relation on the types, and there is an ordering on the equivalence classes, such that for any allocation constraint y, the optimal outcome rule $\mathrm{Outcome}(y)$ induces an allocation rule that is greedy by this ordering with ties between types in a same equivalence class broken uniformly at random.

Recall that revenue curves are assumed to be concave because the class of mechanisms is convex. Essentially if they were not concave, we could find the smallest concave function that upper bounds the revenue curve and on an interval (q_1, q_2) where the revenue curve is strictly below this upper bound, we can mix between the two quantiles with the appropriate probability to obtain a mechanism with any ex ante allocation probability $q \in (q_1, q_2)$. This is exactly the ironing procedure that Myerson (1981) used for constructing monotone "ironed" virtual value from non-monotone (original) virtual value: the derivative of a concave function is monotone. The types ironed by this approach have the same marginal revenue (ironed virtual value) and therefore are in the same equivalence class in the above definition of orderability.

In Theorem 2, we show that orderability is a consequence of linearity and therefore the environments where the optimal mechanism is given by a context-free mapping from each agent's type to a "virtual type" (in our case the marginal revenue) and maximization of the "virtual surplus" are exactly the environments where the single-agent problems are revenue linear. When we want to draw attention to this consequence of revenue linearity we will refer to the single agent problem as being *context free*.

Theorem 2. *For any single-agent problem, revenue linearity implies orderability.*

Before showing Theorem 2, we will formalize the consequences of context freedom (from revenue linearity and orderability) alluded to in the discussion above. We define a simple mechanism for maximizing marginal revenue for orderable agents, prove that it is deterministic, feasible, dominant strategy incentive compatible, and obtains the optimal revenue.

Definition 5. The *marginal revenue mechanism* for orderable agents works as follows.

1. Map reported types $\mathbf{t} = (t_1, \ldots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \ldots, q_n)$ via the implied ordering.[4]
2. Calculate the marginal revenue of each agent i as $R_i'(q_i)$.
3. For each agent i, calculate the maximum quantile q_i^* that she could possess and be in the marginal revenue maximizing feasible set (breaking ties consistently).
4. For each agent i, run the single agent mechanism $\mathcal{M}_i^{q_i^*}$

Proposition 3. *The marginal revenue mechanism deterministically selects a feasible set of agents to serve and is dominant strategy incentive compatible.*

Proof. It is clear that the set of winners of the mechanism is decided deterministically. Feasibility follows as the set of agents that select service outcomes is exactly the marginal revenue maximizing set subject to feasibility. To verify the dominant strategy incentive compatibility consider any agent i's perspective. The parameter q_i^* is a function only of the other agents' reports; the agent's outcome is determined by the q_i^*-step mechanism $\mathcal{M}_i^{q_i^*}$ which is incentive compatible for any q_i^*. □

Proposition 4. *In context-free environments, the marginal revenue mechanism obtains the optimal marginal revenue.*

By revenue linearity, from Remark 1 it immediately follows that:

Corollary 5. *In context-free environments, the marginal revenue mechanism obtains the optimal revenue.*

Even though we optimized over the larger class of randomized and Bayesian incentive compatible mechanisms we obtain one that deterministicly selects the set of winners and is dominant strategy incentive compatible.

Corollary 6. *In context-free environments, there is an optimal mechanism (among randomized and Bayesian incentive compatible mechanisms) that deterministically selects the set of agents to serve and is dominant strategy incentive compatible.*

In the remainder of this section we give a proof of Theorem 2, i.e., that revenue linearity implies orderability. By the definition of marginal revenue maximization an agent is (potentially) treated distinctly with quantiles with different marginal revenue and the identically for quantiles with the same marginal revenue. It is also obvious for quantiles in intervals where the revenue is locally linear (therefore: the marginal revenue is constant and the second derivative of the revenue curve is zero) can never be the critical quantiles in Step 3 of the marginal revenue mechanism and therefore

[4] This ordering can be found by calculating the optimal single-agent mechanism for allocation constraint $y(q) = 1 - q$.

the corresponding step mechanisms are never used. These structural properties are not that useful, however, if there is not a fixed mapping from type space to quantile space. Indeed, the main challenge of multi-dimensional mechanism design is that there is no explicit ordering on type space; revenue linearity, however, implies that there is an implicit fixed mapping. Consequently, much of the intuition for optimal mechanisms in single-dimensional environments extends to optimal mechanisms in context-free environments.

The intuition behind this proof comes from connecting our marginal revenue approach to Myerson's theory of ironing for single-dimensional agents. For single dimensional types there is a natural ordering on types and a revenue curve can be defined by considering mechanisms that serve the top q quantile of types by this ordering. In these step mechanisms a type is either served or not served with probability one. This revenue curve may not be concave and must therefore be ironed (by mixing between deterministic step mechanisms) to get a concave revenue curve. The step mechanism that correspondes to mixing between two step mechanisms does not serve agents with probability one. Instead it has three levels: high types are served with probability one, medium types are served with the same constant probability (depending on the probabilities of the convex combination), and low times are served with probability zero. However, again because the revenue curve is linear (second derivative of the revenue $R''(q)$ is zero) on such an interval these mixed step mechanisms are never used. This exact same structure generalizes to context-free environments.

Lemma 7. *For a linear single-agent problem, let x be the optimal allocation rule subject to some constraint y. Then, for any q such that $R''(q) \neq 0$ we have $X(q) = Y(q)$.*

Proof. Since x is the optimal allocation rule subject to y, then we have, by definition, $\mathrm{Rev}(x) = \mathrm{Rev}(y)$. Linearity implies that

$$\int_0^1 x(q)R'(q)\,\mathrm{d}q = \mathrm{MR}(x) = \mathrm{Rev}(x) = \mathrm{Rev}(y) = \mathrm{MR}(y) = \int_0^1 y(q)R'(q)\,\mathrm{d}q.$$

Integrating by parts, we have

$$\Big[X(q)R'(q)\Big]_0^1 - \int_0^1 X(q)R''(q)\,\mathrm{d}q = \Big[Y(q)R'(q)\Big]_0^1 - \int_0^1 Y(q)R''(q)\,\mathrm{d}q. \tag{2}$$

Since y and x have the same ex ante probability of allocation, $Y(1) = X(1)$; by definition $X(0) = Y(0) = 0$. Substituting these into (2), we have

$$\int_0^1 X(q)R''(q)\,\mathrm{d}q = \int_0^1 Y(q)R''(q)\,\mathrm{d}q,$$

or

$$\int_0^1 [X(q) - Y(q)]R''(q)\,\mathrm{d}q = 0. \tag{3}$$

Notice that for any q, $X(q) - Y(q)$ and $R''(q)$ are non-positive (by domination and concavity, respectively) so their product is non-negative. Therefore, (3) can be satisfied only if $[X(q) - Y(q)]R''(q) = 0$ for all q. This implies that if $R''(q) < 0$, then we must have $X(q) = Y(q)$, which completes the proof. □

Lemma 7 in particular implies that for q with $R''(q) \neq 0$ the q-step mechanism (for step constraint y^q) has allocation rule $x^q = y^q$. I.e., the q-step mechanism has full lotteries only (no partial lotteries).

For any such q, define T_q to be the set of types allocated (with full lotteries) in the optimal allocation subject to y^q. The following lemma shows that these sets are nested.

Lemma 8. *For a revenue-linear single-agent problem, for any $q_1 > q_2$ and $R''(q_1), R''(q_2) \neq 0$, we must have $T_{q_1} \supseteq T_{q_2}$.*

Proof. Assume for contradiction that $T_{q_2} \backslash T_{q_1} \neq \emptyset$. Let $\alpha = F(T_{q_2} \backslash T_{q_1}) > 0$. Consider the following allocation constraint

$$y(q) = \begin{cases} 1 & q \leq q_2 \\ 1/2 & q_1 \leq q \leq q_2 \\ 0 & q_1 \leq q. \end{cases}$$

By revenue linearity, the revenue of the optimal auction subject to y is $[R(q_1) + R(q_2)]/2$. Notice that the mechanism that runs $R(q_1)$ and $R(q_2)$ each with probability $1/2$ achieves this revenue. The allocation rule x of this mechanism is

$$x(q) = \begin{cases} 1 & q \leq q_2 - \alpha \\ 1/2 & q_2 - \alpha \leq q \leq q_1 + \alpha \\ 0 & q_1 + \alpha \leq q. \end{cases}$$

Notice that this allocation rule is dominated by y, and achieves the optimal revenue. Yet, we have

$$Y(q_1) = \int_{q=0}^{q_1} y(q)dq > \int_{q=0}^{q_1} x(q)dq = X(q_1).$$

This contradicts Lemma 7. □

In summary, in context-free environments, there is a natural and simple implementation of the marginal revenue mechanism, which is deterministic and dominant strategy incentive compatible, and yet is optimal among randomized and Bayesian incentive compatible mechanisms. Generally, when revenue linearity and orderability do not hold, both the implementation of marginal revenue maximization and the revenue guarantee of marginal revenue maximization are less clear. In the following sections we address these problems respectively, and we show that a large class of environments inherit some of nice properties of context free environments. We will show that both the implementation simplicity and the revenue guarantees degrade rather smoothly as we move away from context freedom.

4 Implementation

The marginal revenue mechanism (Definition 5) for agents with orderable types does not extend to general agents. In this section we give two approaches for defining the marginal revenue mechanism more generally. The first approach assumes that the parameterized family of q-step mechanisms satisfy a natural monotonicity requirement: that the probability that an agent with a given type is served is monotone in q (the mechanism's parameter). Like the marginal revenue mechanism for orderable agents, this mechanism is dominant strategy incentive compatible. Unlike the marginal revenue mechanism for orderable agents, this mechanism does not deterministically select a set of agents to serve. The second approach is brute-force but easily computable and completely general. It results in a Bayesian incentive compatible mechanism. These two mechanisms will differ from the marginal revenue mechanism only in the first (mapping types to quantiles) and last (serving each agent if her quantile is at most her critical quantile) steps; these changes can be mix-and-matched for different agents in the same mechanism.

Marginal revenue maximization for given revenue curves and feasibility constraint induces a profile of normalized interim allocation rules via the following simulation: Draw agent quantiles uniformly from $[0,1]$; calculate the marginal revenues for each agent; serve the set of agents to maximize the marginal revenue. This simulation gives rise to the profile of normalized interim allocation rules that maximize marginal revenue in expectation. Denote these interim allocation rules by $y_1^{MR}, \ldots, y_n^{MR}$. Any real mechanism that maximizes marginal revenue should look to each agent i like sampling a q-step mechanism with density $-\frac{\mathrm{d}}{\mathrm{d}q} y_i^{MR}(q)$. The outcome rule of this convex combination for agent i is given by: $\tilde{w}_i^{MR}(t_i) = -\int_0^1 \tilde{w}^q(t_i)\,\mathrm{d}y_i^{MR}(q)$ where $\tilde{w}^q$ is the outcome rule for the q-step mechanism. Our goal in this section is to the find multi-agent mechanism that induces these interim outcome rules $\tilde{w}_1^{MR}, \ldots, \tilde{w}_n^{MR}$ and is ex post feasible.

We conclude this section by describing a relevant class of agents for which the step mechanisms satisfy the monotonicity property required by the first approach of this section. The example is one of a single-dimensional agent with a public budget.

4.1 Monotone Step Mechanisms

We define below a simple extension of the marginal revenue mechanism for orderable types for the case where the single-agent step mechanisms satisfy a natural monotonicity property. The resulting mechanism is based on a randomized mapping from types to quantiles that is independent across the agents.

Definition 6. An agent has *monotone step mechanisms* if, given her type, the probability she wins in the q-step mechanism $\mathcal{M}^q$ is monotone non-decreasing in q.

Suppose that the q-step mechanisms $\mathcal{M}^q$ for an agent each consist of a menu of full lotteries. I.e., for any type of the agent she will choose a lottery that either serves her with probability 1 or zero. In this case the monotone step mechanisms assumption would require that the sets of types served for each q are nested. There is a simple deterministic mapping from types to quantiles in this case: set the quantile of a type to be the minimum q such that the q-step mechanism serves the type. Below, we generalize this selection procedure to the case with partial lotteries.

Recall that the q-step mechanism $\mathcal{M}^q$ has allocation rule $\tilde{x}^q$ that maps types to probability of service. Fix the type of the agent as t and consider the function $G_t(q) = \tilde{x}^q(t)$ which, by the monotonicity condition above, can be interpreted as a cumulative distribution function. Notice that $\mathcal{M}^q$ has ex ante probability of service $\mathbf{E}_t[\tilde{x}^q(t)] = q$. Therefore, if t is drawn from the type distribution and then q is drawn from G_t then the distribution F of q is uniform on $[0,1]$.

Lemma 9. *If $t \sim F$ and $q \sim G_t$ then q is $U[0,1]$.*

Definition 7. The *marginal revenue mechanism* for agents with monotone step mechanisms works as follows.

1. Map reported types $\mathbf{t} = (t_1, \ldots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \ldots, q_n)$ by sampling q_i from the distribution with cumulative distribution function $G_{t_i}(q) = \tilde{x}_i^q(t_i)$.
2. Calculate the marginal revenue of each agent i as $R_i'(q_i)$.
3. For each agent i, calculate the maximum quantile q_i^* that she could possess and be in the marginal revenue maximizing feasible set (breaking ties consistently).
4. For each agent i, offer the mechanism $\mathcal{M}_i^{q_i^*}$ conditioned so that i is served if $q_i \leq q_i^*$ and not served otherwise.

The last step of the marginal revenue mechanism warrants an explanation. In the q_i^*-step mechanism $\mathcal{M}_i^{q_i^*}$, the outcome that i would obtain with type t_i may be a partial lottery, i.e., it may probabilistically serve i or not. The probability that i is served is $\tilde{x}_i^{q_i^*}(t_i) = \mathbf{Pr}_{q_i}[q_i \leq q_i^*]$ by our choice of q_i. When we offer agent i the mechanism $\mathcal{M}_i^{q_i^*}$ we must draw an outcome from the distribution given by $\tilde{w}_i^{q_i^*}(t_i)$. Some of these outcomes are service outcomes, some of these are non-service outcomes. If $q_i \leq q_i^*$ then we draw an outcome from the distribution $\tilde{w}_i^{q_i^*}(t_i)$ conditioned on service; if $q_i > q_i^*$ then we draw an outcome conditioned on no-service. Notice that it may not be feasible to serve all agents who receive non-trivial partial lotteries; this method of coordinating which of these agents are served has the right distribution and always gives a feasible outcome.

Proposition 10. *The marginal revenue mechanism for agents with monotone step mechanisms is feasible and dominant strategy incentive compatible.*

Proof. Feasibility follows as the set of agents that select service outcomes is exactly the marginal revenue maximizing set subject to feasibility. To verify the dominant strategy incentive compatibility consider any agent i's perspective. The parameter q_i^* is a function only of the other agents' reports; the agent's outcome is determined by the q_i^*-step mechanism $\mathcal{M}_i^{q_i^*}$ which is incentive compatible for any q_i^*. □

Theorem 11. *The marginal revenue mechanism for agents with monotone step mechanisms implements marginal revenue maximization.*

Proof. From each agent i's perspective, the other agents' quantiles are distributed independently and uniformly on $[0, 1]$. Therefore, this agent faces a distribution over step mechanisms that is identical to the distribution of "critical quantiles" in the maximization of marginal revenue, i.e., with density $-\frac{\mathrm{d}}{\mathrm{d}q} x^{MR}(q)$. □

4.2 General Step Mechanisms

We now give a general and, in a manner of fashion, brute-force procedure for implementing marginal revenue maximization with general agents, i.e., without monotone step mechanisms. Consider an agent with interim allocation constraint y^{MR} (from marginal revenue maximization) and mix over q-step mechanisms with probability density given by $-\frac{\mathrm{d}}{\mathrm{d}q} y^{MR}(q)$ to get outcome rule $\tilde{w}^{MR}$. Recall each step mechanism is derived by optimizing revenue subject to a step function constraint on the allocation rule. The resulting outcome rule induces a normalized allocation rule that may not be a step function. Therefore, x^{MR} is dominated by but not necessarily equal to y^{MR}. A natural approach to deriving an ordering on types and thus mapping types to quantiles would be to use the mapping given by $\text{Quant}(\cdot)$ (defined in Section 2) for x^{MR}. This approach would result in allocation rule y^{MR} not the desired allocation rule x^{MR}. Instead the map from types to quantiles needs to be randomized to make it "worse" and equal to the desired allocation rule x^{MR}.

For uniform distributions over discrete type spaces, domination of allocation rules is equivalent to vector majorization. Hardy et al. (1929) show a transformation can be given between one and the other via a doubly stochastic matrix. This implies an algorithm for mapping types to quantiles based on solving a quadratic sized linear program: solve for the doubly stochastic matrix that gives x^{MR} from y^{MR}, map an agent type to quantile via $\text{Quant}(\cdot)$ induced by x^{MR}, and randomly map this quantile to a distribution over quantiles by sampling according to the probabilities in the quantile's row in the matrix.

As we do not need to construct the matrix, but only to sample from a specific row there is a much simpler construction. This construction further generalizes to non-uniform distributions. Instead of permutation matrices this approach is based on applying a sequence of interval resamplings. An interval resampling is given by an interval $[a, b]$ and if the quantile is in $[a, b]$ it is resampled uniformly from $[a, b]$, otherwise it is unchanged. For allocation rule x, resampling quantiles from $[a, b]$ has the effect of replacing the interval of the cumulative allocation rule X with the line segment connecting $(a, X(a))$ to $(b, X(b))$. The construction below, for type spaces of size $m = |T|$, calculates the requisite sequence of at most m interval resamplings in linear time.

Definition 8. For allocation constraint y and dominated allocation rule x satisfying $Y(1) = X(1)$ on m discrete types, the *interval resampling sequence construction* starts with $x^{(0)} = y$ and calculates $x^{(j+1)}$ from $x^{(j)}$ while $x^{(j)} \neq x$ as follows.

1. Find the highest quantile q where $x(q) \neq x^{(j)}(q)$.
2. Let $q' > q$ be the quantile at which the line tangent to X at q with slope $x(q)$ crosses $X^{(j)}$.[5]
3. The jth resampling interval is $[q, q']$.
4. Let $x^{(j+1)}$ be $x^{(j)}$ averaged on $[q, q']$.

Proposition 12. *The interval sampling sequence construction gives a sequence of at most m intervals such that the composition of y with the sequence of resamplings applied to* $\text{Quant}(\cdot)$ *is equal to x.*

Proof. The proof is by induction on j where the jth step assumes the first $j - 1$ types, in order of $\text{Quant}(\cdot)$, satisfy $x^{(j-1)}(\text{Quant}(t)) = x(\text{Quant}(t))$. Consider step j. The assumption that $Y(1) = X(1)$ ensures that the intersection of the tangent happens at a $q' \leq 1$. The line segment connecting interval $[q, q']$ of $X^{(j)}$ has slope equal to $x(q)$, by definition. Therefore, the jth step in the construction leaves $x^{(j)}(\text{Quant}(t)) = x(\text{Quant}(t))$ for the jth type. The procedure is linear time as both y and x are, without loss of generality, piece-wise constant with m pieces, and in each step q and q' are increasing and at least one piece from y or x is processed. □

[5] For discrete type, this intersection may happen at a quantile q' that does not correspond to the boundary between two types. When this happens split the type into two types each occurring with the same total probability and with q' separating them.

The final ingredient in the construction of the marginal revenue mechanism for agents with general types is in converting the allocation rule back into an outcome rule. This can be done exactly as in Alaei et al. (2012): if an agent with type t is served by the allocation rule, sample from service outcomes of $\tilde{w}^{MR}(t)$, otherwise sample from non-service outcomes of $\tilde{w}^{MR}(t)$.

Definition 9. The *marginal revenue mechanism* for general agents works as follows.

1. Map reported types $\mathbf{t} = (t_1, \ldots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \ldots, q_n)$ by, for each agent, composing the interval resampling transformation with $\text{Quant}(\cdot)$.
2. Calculate the marginal revenue of each agent i as $R_i'(q_i)$.
3. Calculate the set of agents to be served by marginal revenue maximization.
4. Calculate outcomes for each agent i as:
   - sample $w_i \sim \tilde{w}_i^{MR}(t_i)$ conditioned on $\text{Alloc}(w_i) = 1$ if i is to be served, or
   - sample $w_i \sim \tilde{w}_i^{MR}(t_i)$ conditioned on $\text{Alloc}(w_i) = 0$ if i is not to be served.

Theorem 13. *The marginal revenue mechanism for general agents implements marginal revenue maximization and is Bayesian incentive compatible.*

Note that instead of calculating outcome rules by mixing over step mechanisms we could, from the allocation constraint y^{MR} for an agent, calculate the optimal mechanism subject to that constraint, i.e., with outcome rule $\text{Outcome}(y^{MR})$ and revenue $\text{Rev}(y^{MR})$. The construction above can be invoked with this outcome rule in place of $\tilde{w}^{MR}$ without modification.

The assumption in the interval resampling sequence construction that the allocation constraint y and the desired allocation x have cumulative allocations satisfying $Y(1) = X(1)$ can be removed in downward closed settings. With this constraint removed the line tangent to X at q may be strictly below Y at 1. If this happens we set $q' = 1$ and probabilistically reject a quantile falling in this interval. The probability of rejection is set so that the slope of the resampled allocation rule is $x(q)$.

Proposition 14. *In downward closed settings, a normalized allocation rule x dominated by an allocation rule y, where $Y(1) > X(1)$, can still be implemented by interval resamplings given in Definition 8, and a marginal revenue mechanism for x can run similarly as in Definition 9.*

We use $y \succeq x$ to refer to this weaker definition of dominance in downward closed settings.

4.3 Example: single-dimensional agents with public budgets

In this section we exhibit a class of single agent problems that has monotone step mechanisms. Consider an agent with a single-dimensional value for receiving a good and a public budget. Her utility is her value for receiving the good minus her payment as long as her payment is at most her budget. We show that under standard conditions on the agent's valuation distribution, this single agent problem has monotone step mechanisms. Of course, this budget setting is not context-free therefore, the marginal revenue mechanism for agents with monotone step mechanisms is necessary for implementation of marginal revenue maximization.

The following proposition is a consequence of results of Laffont and Robert (1996) and Pai and Vohra (2008); for completeness we provide proof of this special case of their result in section B. Our theorem follows.

Proposition 15. *For regular distribution F with non-decreasing density, budget B, and quantile $q \leq 1 - F(B)$, the q-step mechanism offers a single take-it-or-leave-it lottery for price B that serves with probability π. This lottery is bought by the agent when her value is at least B/π which happens with probability $q = 1 - F(B/\pi)$.*

Notice that the allocation rule of the mechanism satisfying Proposition 15 is a function that steps from 0 to π at value B/π. The required payment B can be viewed as "the area above the allocation curve" which is given by a rectangle with width B/π and height π. The ex ante probability of service is the probability that an agent has value exceeding B/π. Clearly as this probability increases, B/π decreases, and π must increase to keep the area above the curve the same at B. Therefore increasing q increases the set of types served and the probability of service. We conclude with the following consequence.

Theorem 16. *An agent with value drawn from a regular distribution with non-increasing density has monotone step mechanisms.*

Proof. The only case not argued on the text above is when $q \geq 1 - F(B)$. In this case, the budget is not binding and the step mechanism posts price p that satisfies $q = 1 - F(p)$ and serves agents willing to pay this price with probability one. The step mechanisms are monotone over these quantiles as well. □

5 Approximation

In previous sections, we have shown that for any collection of agents the marginal revenue mechanism can be implemented. We know that for context-free agents, this extracts the optimal revenue. In this section, we show that this revenue is a good approximation to the optimal revenue quite generally.

We will give two approaches for approximation bounds. The first kind of bound is based on the single-agent problem, i.e., the distribution and type space. If we can show that for all allocation constraints, the marginal revenue is a good approximation to the optimal revenue, then the marginal revenue mechanism is a good approximation to the optimal mechanism. The second approach will derive bounds from the feasibility constraint. Clearly, with no feasibility constraint, marginal revenue maximization is optimal. We will show that for matroid environments, it gives a $1 - 1/e$ approximation, and for general downward-closed environments, it gives a $O(\log n)$ approximation.

Of course, if we are in an environment where our agent-based arguments imply an α-approximation and our feasibility-based arguments imply a β-approximation, the marginal revenue mechanism is in fact a $\min(\alpha, \beta)$-approximation. In context-free environments $\alpha = 1$ (and marginal revenue is optimal); the approximation smoothly degrades in α as the environment becomes less context free until it reaches the approximation bound β given by the feasibility constraint.

5.1 Agent-based Approximation

If, for all allocation rules, the marginal revenue is close to the optimal revenue, then marginal revenue maximization is approximately optimal. One approach to deriving such a bound is to give a linear upper bound on the optimal revenue and a lower bound through a class of pseudo step mechanisms. A pseudo step mechanism respects a step constraint but may not be optimal. If for every quantile q the pseudo q-step mechanism approximates the linear upper bound, then marginal revenue maximization approximates the optimal revenue for all allocation constraints. Furthermore, these pseudo step mechanisms can be directly optimized over and the same approximation factor is obtained. Such an approach might be desirable if the pseudo step mechanisms are better-behaved than the (optimal) step mechanisms, e.g., if they are easy to compute, respect an ordering on types, or are step monotone.

This approach is formalized by the following sequence of definitions and propositions.

Proposition 17. *If for any agent i and allocation rule x_i, the marginal revenue* $\mathrm{MR}(x_i)$ *is at least an α fraction of the optimal revenue* $\mathrm{Rev}(x_i)$*, then the marginal revenue mechanism in the multi-agent setting is an α approximation to the optimal mechanism.*

Definition 10. A *linear revenue bound*, UB, is

- an upper bound on revenue for all allocation constraints, i.e., $\forall y,\ \mathrm{UB}(y) \geq \mathrm{Rev}(y)$, and
- linear in the allocation constraint, i.e., for all allocation constraints y_a and y_b and $\gamma \in [0,1]$, $\mathrm{UB}(\gamma y_a + (1-\gamma)y_b) = \gamma\,\mathrm{UB}(y_a) + (1-\gamma)\,\mathrm{UB}(y_b)$.

Definition 11. A *pseudo step mechanism* is one that respects a step constraint but is not necessarily revenue optimal for such a constraint. The revenue of the pseudo q-step mechanism $\tilde{\mathcal{M}}^q$ is denoted $\tilde{R}(q)$; and the *pseudo marginal revenue* for allocation constraint y is $\mathrm{PMR}(y) = \mathbf{E}[\tilde{R}'(q)y(q)]$.

We can assume without loss of generality that the pseudo marginal revenue $\tilde{R}$ is concave. If it is not we could always redefine the class by taking its closure with respect to convex combination and letting the pseudo q-step mechanism be the revenue-optimal mechanism in the class that serves with ex ante probability q.

Proposition 18. *For a given linear revenue bound, if for all q the pseudo q-step mechanism α-approximates the bound on the q-step constraint, then the pseudo marginal revenue α-approximates the optimal revenue for all allocation constraints.*

Proof. This proposition follows from linearity of both the revenue bound and the pseudo marginal revenue. □

Definition 12. The *pseudo marginal revenue mechanism* is the one that maximizes pseudo marginal revenue via any of the approaches of Definition 5, Definition 7, or Definition 9 that applies.

Pseudo step mechanisms for downward-closed unit-demand agents. Consider the single-agent problem of serving an agent one of m types of services (or none) subject to an allocation constraint y. For clarity with respect to the literature we will refer to this agent as being unit-demand and the types of services as being distinct items. Without an allocation constraint, this problem has seen extensive study. Briest et al. (2010) show that the optimal lottery pricing can be calculated by a linear program that has size equal to the number of distinct types of the agents. When the agent's value for the items are independently distributed, Cai and Daskalakis (2011) give a dynamic program for approximating the optimal item pricing to within a $(1+\epsilon)$ factor for any ϵ in time polynomial in the number of items. These results are distinct in two ways. First, the first result is optimal with respect to randomized mechanisms whereas the second is (nearly) optimal with respect to deterministic mechanisms. Second, the first result would require time exponential in the number of items for a product distribution, while the second result is polynomial (but requires a product distribution). It is not known whether the optimal lottery pricing for product distributions can be calculated arbitrarily closely in polynomial time in the number of items. Attempting to address this question, a combination of the work of Chawla et al. (2010a,b) shows that for product distributions, item pricing is a 4-approximation to lottery pricing. Furthermore, there is an item pricing that is very simple to describe that satisfied this bound. We generalize this theory to single-agent problems with a supply constraint y and use it to give a 4-approximate class of pseudo step mechanisms for unit-demand, downward-closed agents.

Consider the problem of selling a single item to one of m single-dimensional agents with values drawn from a product distribution. I.e., the value v_i for agent i is drawn independently from F_i. As described earlier, the optimal auction for this single-dimensional problem is well understood. Agent values are mapped to virtual values (equivalent to each agent's marginal revenue), and the agent with the highest positive virtual value is selected as the winner of the auction. We refer to this auction environment as the single-dimensional *representative environment*, the revenue obtained by the optimal auction as the *optimal representative revenue,* and the agents participating in the auction as *representatives*.

Notice that if these representatives were all colluding together the problem would be identical to our original single-agent unit-demand problem. We refer to this environment as the *unit-demand environment* and the revenue of the optimal lottery pricing as the *optimal unit-demand revenue*. The approach of Chawla et al. (2010a) is to try to mimic the outcome of the optimal auction for the representative environment to obtain an approximately optimal pricing in the unit-demand environment. As the optimal auction in the representative environment orders representatives by virtual value, a natural approach to pricing the items in the unit-demand environment is to set a uniform virtual price, i.e., the price for each item has the same virtual value (with respect to the distribution from which the agent's value for that item is drawn).[6] Chawla et al. (2010a) show that the unit-demand revenue of such a pricing is a 2-approximation to the optimal representative revenue; Chawla et al. (2010b) show that the optimal unit-demand revenue (e.g., from lottery pricings) is at most twice the optimal representative revenue. Combining these two results, *uniform virtual pricing* is a 4-approximation to the optimal unit-demand revenue.

We generalize the approach above to single-agent problem of serving an agent with independent values for m items subject to an allocation constraint y. In particular we show that twice the optimal representative revenue is a linear revenue bound (Definition 10), i.e., it is linear in the allocation constraint and for any allocation constraint it upper bounds the optimal (unit-demand) revenue. We define a class of pseudo step mechanisms where the pseudo q-step mechanism is given by a uniform virtual pricing that sells with probability q. Since the virtual values are weakly increasing in the representative agents' values, the sets of types served by these pseudo step mechanisms are nested. Therefore, the pseudo marginal revenue mechanism can be implemented via the marginal revenue mechanism for orderable agents (Definition 5). Finally, we show that for all q the pseudo q-step mechanism is a 4-approximation to the linear upper bound given by twice the optimal representative revenue. This result with Proposition 18 implies that the pseudo marginal revenue mechanism is a 4-approximation to the optimal revenue. The proof of Theorem 19, below, is a rather straightforward extension of Chawla et al. (2010a,b) and we include it in section C.

Definition 13. The pseudo q-step mechanism for a unit-demand agent with value for each item

[6] As mentioned above, a representative's virtual value is equal to their marginal revenue. For clarity of discussion and to disambiguate the marginal revenue of the unit demand agent versus that of her representatives we will refer to the representative marginal revenue as virtual value.

drawn independently from $F_1, \ldots, F_n$ is given by the pricing that sets a uniform virtual price for the items such that the probability that the agent buys any item is equal to q. (If this class does not have a monotone non-decreasing pseudo revenue curve $\tilde{R}(\cdot)$ we invoke downward closure to make it monotone; if this class does not have a concave pseudo revenue curve we take its closure with respect to convex combination to make it concave.)

Theorem 19. *In downward-closed (service constrained) environments with unit-demand agents, the pseudo marginal revenue mechanism is a 4-approximation to the optimal revenue.*

Corollary 20. *In downward-closed (service constrained) environments with unit-demand agents, the marginal revenue mechanism is a 4-approximation to the optimal revenue.*

5.2 Feasibility-based Approximation

We now show that the feasibility constraint implies an approximation bound as well. As a first simple bound, if there is no feasibility constraint (e.g., for digital goods) then marginal revenue maximization is optimal. Below we give a $e/(e-1)$ approximation bound for matroid environments and an $O(\log n)$ bound for downward-closed environments.

Matroid environments. Marginal revenue maximization is an $e/(e-1)$ approximation when the feasibility constraint is induced by independent sets of a matroid set system. This same approximation factor governs single-item auctions. For k-unit environments we obtain a $1/(1-(2\pi k)^{-1/2})$-approximation. These results follow from the correlation gap approach of Yan (2011).

Theorem 21. *In a matroid environment the optimal marginal revenue is a $e/(e-1)$-approximation to the optimal revenue; for k-unit environments it is a $1/(1-(2\pi k)^{-1/2})$-approximation.*

Proof sketch. Suppose the optimal mechanism serves agent i with ex ante probability q_i. Relax the feasibility constraints and consider maximizing revenue subject to ex ante probability q_i for each agent i. This revenue is only greater and it is precisely $\sum_i R_i(q_i)$. Sort agents by $R_i(q_i)/q_i$ and run the greedy matroid algorithm: if it is possible to serve i when i is visited by the algorithm, then offer i the q_i-step mechanism $\mathcal{M}^{q_i}$ (by definition, the optimal mechanism that serves with ex ante probability q_i). A corollary of the main theorem of Yan (2011) is that this is an $e/(e-1)$-approximation for general matroids and an $1/(1-(2\pi k)^{-1/2})$-approximation for k-uniform matroids. This greedy-based mechanism's revenue is given by its marginal revenue, and therefore the marginal revenue maximizer is only better. □

Downward-closed environments. In this section we show that in downward-closed environments the optimal marginal revenue is a logarithmic approximation, in the number of agents, to the optimal revenue.

The intuition for the proof is as follows. If we consider allocation constraints with a minimum probability of allocating to any type of 2^{-K}, then the allocation constraint can be partitioned into K pieces with the highest and lowest probability of allocation in each piece being within a factor of two of each other. The revenue each piece can be approximated by a q-step constraint scaled appropriately so that it is dominated by the original allocation constraint. The total revenue is then at most an $O(K)$-fraction of the revenue of the best such scaled step constraint.

We start by proving the following lemma.

Lemma 22. *Any allocation constraint with minimum probability $y(1) \geq 2^{-K}$ has revenue at most $2K\,\mathrm{MR}(y)$.*

Proof. Let $R^* = \mathrm{Rev}(y)$ be the optimal revenue for allocation constraint y. Define sequence of quantiles $0 = q_0 \leq q_1 \leq \cdots \leq q_K = 1$ such that $x(q_{j-1}) \leq 2x(q_j)$. Define R^*_j to be the expected revenue from types that are mapped to a quantile in $[q_{j-1}, q_j]$ by $\mathrm{Quant}(\cdot)$ as specified in Definition 1. Therefore, the revenue of the mechanism is $R^* = \sum_{j=1}^{K} R^*_j$. Then there must exist j^* such that $R^* \leq K R^*_{j^*}$. In what follows, we define normalized allocation rules $z_j(\cdot)$ for all j, such that $z_j \preceq x$, and also $R^*_j \leq 2\,\mathrm{MR}(z_j)$. In particular, for j^* we have $z_{j^*} \preceq x$, and $2\,\mathrm{MR}(z_{j^*}) \geq R^*_{j^*} \geq R^*/K$. This implies that

$$2 \max_{z \preceq x} \mathrm{MR}(z) \geq 2\,\mathrm{MR}(z_{j^*}) \geq R^*/K.$$

Define function $z_j(\cdot)$ to be $z_j(q) = x(q_{i+1})$ if $q \leq q_{i+1} - q_i$, and 0 otherwise. Notice that for any q, we have $z_j(q) \leq x(q)$, and therefore $z_j \preceq x$, by the definition of dominance in downward-closed environments.

We now show that for z_j defined above, $R^*_j \leq 2\,\mathrm{MR}(z_j)$. By construction of z_j,

$$2\,\mathrm{MR}(z_j) = 2\int_0^1 z_j(q) R'(q)\mathrm{d}q = 2x(q_{j+1})R(q_{j+1} - q_j) \geq x(q_j)R(q_{j+1} - q_j)$$

It is therefore sufficient to show that $R^*_j \leq x(q_j)R(q_{j+1} - q_j)$. Recall that R^*_j is the revenue from types that are mapped to quantiles in $[q_j, q_{j+1}]$. Any type in $[q_j, q_{j+1}]$ is allocated in x with probability at most $x(q_j)$. Now define the set of lotteries L to be the lotteries chosen by types in $[q_j, q_{j+1}]$. Notice that types in $[q_j, q_{j+1}]$ choose the same lottery in L as they did in x. As a result, the measure of the types that choose some lottery in L is at least $q_{j+1} - q_j$. Now remove lotteries from L, from the one with lowest price, until the measure of types that choose some lottery is exactly $q_{j+1} - q_j$. Call this new set of lotteries L'. Notice that the revenue from L' is at least R^*_j. Now recall that all the lotteries in L, and therefore L', allocate with probability at most $x(q_j)$. So it is feasible to define a new set of lotteries L'' to be the lotteries in L' scaled up by $1/x(q_j)$. The revenue from L'' is therefore at least $R^*_j/x(q_j)$. Since a fraction $q_{j+1} - q_j$ buy some lottery in L'', by definition the revenue that we get is at most $R(q_{j+1} - q_j)$. We conclude that $R^*_j/x(q_j) \leq R(q_{j+1} - q_j)$.

To complete the proof, recall that for downward-closed environments revenue curves are monotone non-decreasing so marginal revenues are non-negative. Therefore, by the definition of marginal revenue and dominance, $\mathrm{MR}(y) \geq \mathrm{MR}(z_j)$ for all j. ☐

Theorem 23. *In n-agent downward-closed environments, the optimal marginal revenue is a $4\log n$-approximation to the optimal revenue.*

Proof. Consider an alternative mechanism that runs the optimal mechanism with probability 1/2, and otherwise picks an agent at random and outputs an arbitrary outcome that services that agent, regardless of his type and without charging him. Notice that the revenue of the alternative mechanism is half the revenue of the optimal revenue. Let $x_1, \ldots, x_n$ be the allocation rules for the alternative mechanism. Notice also that by construction of the alternative mechanism, for each i and $q \in [0, 1]$ we have $x_i(q) \geq 1/2n$. Therefore we can invoke Lemma 22 with $K = \log 2n$ to conculde that that the revenue of the alternative mechanism is at most

$$2\log n \sum\nolimits_i \mathrm{MR}_i(x_i).$$ ☐

References

Alaei, S. (2011). Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 512–521.

Alaei, S., Fu, H., Haghpanah, N., Hartline, J., and Malekian, A. (2012). Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*.

Border, K. (2007). Reduced form auctions revisited. *Economic Theory*, 31(1):167–181.

Border, K. C. (1991). Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4):1175–1187.

Briest, P., Chawla, S., Kleinberg, R., and Weinberg, S. M. (2010). Pricing randomized allocations. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 585–597.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–1090.

Cai, Y. and Daskalakis, C. (2011). Extreme-value theorems for optimal multidimensional pricing. In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 522–531.

Cai, Y., Daskalakis, C., and Weinberg, M. (2012). An algorithmic characterization of multi-dimensional mechanisms. In *Symposium on Theory of Computation*.

Chawla, S., Hartline, J., Malec, D., and Sivan, B. (2010a). Multi-parameter mechanism design and sequential posted pricing. In *42nd ACM Symposium on Theory of Computing*, pages 311–320.

Chawla, S., Malec, D., and Sivan, B. (2010b). The power of randomness in bayesian optimal mechanism design. In *ACM Conference on Electronic Commerce*, pages 149–158.

Hardy, G., Littlewood, J., and Pólya, G. (1929). Some simple inequalities satisfied by convex functions. *Messenger of Math*, 58:145–152.

Hartline, J. and Lucier, B. (2010). Bayesian algorithmic mechanism design. In *STOC*, pages 301–310.

Laffont, J.-J. and Robert, J. (1996). Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186.

Maskin, E. and Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, 52(6):1473–1518.

Matthews, S. (1983). Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory*, 30(2):370–400.

Matthews, S. A. (1984). On the implementability of reduced form auctions. *Econometrica*, 52(6):pp. 1519–1522.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73.

Pai, M. and Vohra, R. (2008). Optimal auctions with financially constrained bidders. Discussion papers, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

Yan, Q. (2011). Mechanism design via correlation gap. In *Symposium on Discrete Algorithms*.

A Testing Context Freedom

In this section we present a simple algorithm that tests whether a given type space and a distribution over it is context free. By definition we need only to check the linearity of $\mathrm{Rev}(\cdot)$.

Recall that $\mathrm{Rev}(\cdot)$ is defined on normalized allocation functions, which are nondecreasing functions defined on $[0, 1]$. After discretizing the quantile space to pieces of size $1/N$, each such function is a convex combination of N reverse step functions. The coefficients of these step functions in such a linear representation sum up to 1. Therefore, $\mathrm{Rev}(\cdot)$ can be equivalently seen as *a concave function defined on the* $(N-1)$*-simplex:* $\Delta^N = \{(x_1, \ldots, x_N) \in \mathbb{R}_+^N \mid \sum x_i \leq 1\}$. In what follows we will present a general algorithm that tests the linearity of such a function f.

Let e_i be the vector whose coordinates are 0 except the i-th one being 1 (the vertices of Δ^N), and let P be the vector $(1/N, \cdots, 1/N)$.

Definition 14. The *concave linearity test algorithm* runs as follows: test if $f(P) = \sum_{i=1}^N \frac{1}{N} f(e_i)$; returns Yes if so, and No otherwise.

Theorem 24. *Given that f is concave, the concave linearity test algorithm returns Yes if and only if f is linear.*

We note that the guaranteed concavity of f is important for the validity of the algorithm.

Proof. If f is linear, it is obvious that the algorithm returns Yes. So we only need to prove the other direction. We define a function g:

$$g(x_1, \ldots, x_N) = f(x_1, \ldots, x_N) - \sum_{i=1}^N x_i f(e_i).$$

Then $g(e_i) = 0$ for each e_i. Notice that g is the difference between a concave function f and a linear function, and therefore g is concave. Also, given that g is 0 on every vertex of the polytope Δ^n, we have that $g \geq 0$ on Δ^n by basic properties of concave functions. It is easy to see that f is linear if and only if g is 0 everywhere on Δ^n.

We show the correctness of the algorithm by induction on n. When n is 2, g is defined on the line segment $\{(x_1, x_2) \mid x_1 + x_2 = 1, x_1, x_2 \geq 0\}$, and by basic analysis we know that if $g(1/2, 1/2) = 0$ and g is concave, then g is constantly 0, which shows that f is linear. Now suppose the case of $N-1$ has been proved, we are going to prove the case of N. Suppose the algorithm returns Yes, i.e., $g(P) = 0$, we need to prove that g is constantly 0 on Δ^N. We first consider the value of g on the point $(1/(N-1), \cdots, 1/(N-1), 0)$. This point is colinear with e_N and P, and it lies in $\Delta^{N-1} \times \{0\} \subset \Delta^N$. In geometric terms, it is the point where the line that connects e_N and P

meets the face $\{\mathbf{x} \in \Delta^n \mid \sum_{i=1}^{N-1} x_i = 1\}$. As we have argued, g is nonnegative everywhere on Δ^N. On the other hand, if $g(1/(N-1), \cdots, 1/(N-1), 0) > 0$, by the concavity of g, we will have

$$\begin{aligned} g(P) = g\left(\frac{1}{N}, \cdots, \frac{1}{N}\right) &= g\left(\frac{1}{N}(0, \cdots, 0, 1) + \frac{N-1}{N}\left(\frac{1}{N-1}, \cdots, \frac{1}{N-1}, 0\right)\right) \\ &\geq \frac{1}{N} g(e_N) + \frac{N-1}{N} g\left(\frac{1}{N-1}, \cdots, \frac{1}{N-1}, 0\right) > 0, \end{aligned}$$

which would contradict that $g(P) = 0$. Therefore, we have $g(1/(N-1), \cdots, 1/(N-1), 0) = 0$. By induction, we have that g is constantly 0 on $\Delta^{N-1} \times \{0\}$. By the same reasoning, g is constantly 0 on all points where at least one coordinate is 0. Now consider any point $\mathbf{x} = (x_1, \ldots, x_N) \in \Delta^N$ where $x_i > 0$ for all i. Connect $\mathbf{x}$ and P and elongate the line in both directions till they intersect the faces of Δ^n at points $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$. [7] Then $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$ both have some coordinate that is 0, and therefore $g(\tilde{\mathbf{x}}) = g(\hat{\mathbf{x}}) = g(P) = 0$. Since $g(P)$ lies in the interior of the line segment connecting $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$, by concavity this implies that g is 0 for all points lying on this segment, and hence $g(\mathbf{x}) = 0$. Since $\mathbf{x}$ was an arbitrary point in Δ^N, we have shown that g is constantly 0 on Δ^N, and hence f is linear. □

B Proof of Proposition 15

The technique of this proof largely comes from Laffont and Robert (1996) and can be viewed as a consequence of that work. We remark that the condition of concavity of f (or monotonicity of f), which was not used in the original paper of Laffont and Robert (1996), was in fact needed for their characterization, as correctly pointed out by Pai and Vohra (2008).

Proof of Proposition 15. Let $x(v)$ be the allocation probability for type v. Without loss of generality, we assume that the highest possible valuation is 1. The standard truthful condition (monotonicity of the allocation rule and the revenue equivalence) still holds here. In particular, for $v > v'$, if $x(v) > x(v')$, then the payment of v is also strictly larger than that of v'. Therefore, if the budget constraint is binding (as we assumed), then there is a $\bar{v}$ such that the allocation probability is a constant for all types above $\bar{v}$, and the payment for all these types is B. The theorem then states that, in an optimal auction, the allocation rule for types smaller than $\bar{v}$ is constantly 0.

Truthfulness requires that the payment of type v is $vx(v) - \int_0^v x(z)\,\mathrm{d}z$. We therefore would like to maximize the objective function

$$\max \int_0^{\bar{v}} f(v)x(v)\varphi(v)\,\mathrm{d}v + [1 - F(\bar{v})]\bar{v}x(\bar{v}), \tag{4}$$

where $\varphi(v)$ is the standard virtual valuation function $v - \frac{1-F(v)}{f(v)}$.

We would like the budget constraint to bind, and so we have the constraint

$$\bar{v}x(\bar{v}) - \int_0^{\bar{v}} x(v)\,\mathrm{d}v = B. \tag{5}$$

[7] Formally, let $\tilde{c}$ be $\sup\{c : \mathbf{x} + c(P - \mathbf{x}) \in \Delta^N\}$, and let $\tilde{\mathbf{x}}$ be $\mathbf{x} + \tilde{c}(P - \mathbf{x})$. Similarly, let $\hat{c}$ be $\inf\{c : \mathbf{x} + c(P - \mathbf{x}) \in \Delta^N\}$, and let $\hat{\mathbf{x}}$ be $\mathbf{x} + \hat{c}(P - \mathbf{x})$. Since Δ^N is a closed set, $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$ are well defined points that are in Δ^N.

The selling probability constraint is written as

$$\int_0^{\bar{v}} x(v)\,\mathrm{d}v + [1 - F(\bar{v})]x(\bar{v}) = q. \tag{6}$$

Then we have that

$$\forall v, \quad x(v) \geq 0; \tag{7}$$
$$x(\bar{v}) \leq 1. \tag{8}$$

We consider the first-order conditions for the above program. We use δ for the Lagrangian variable for the budget condition (5); λ for the selling probability constraint (6); Π_v for condition (7) for each v ($\Pi_v \leq 0$); η for condition (8) ($\eta \geq 0$). The first order condition would give

$$f(v)\left[\varphi(v) + \lambda - \frac{\delta}{f(v)}\right] + \Pi_v = 0, \quad \forall v < \bar{v}; \tag{9}$$

$$[1 - F(\bar{v})]\left[\bar{v} + \lambda + \frac{\bar{v}\delta}{1 - F(\bar{v})}\right] + \Pi_{\bar{v}} + \eta = 0. \tag{10}$$

By complementary slackness, for any v such that $x(v) > 0$, we have $\Pi_v = 0$. (In particular, $\Pi_{\bar{v}} = 0$.) We argue that δ is negative. Assume there is a $v < \bar{v}$ such that $x(v) > 0$. Then we have

$$\varphi(v) + \lambda - \frac{\delta}{f(v)} = 0;$$
$$\bar{v} + \lambda + \frac{\bar{v}\delta}{1 - F(\bar{v})} + \frac{\eta}{1 - F(\bar{v})} = 0.$$

We can therefore solve for δ:

$$\delta = \left[\varphi(v) - \bar{v} - \frac{\eta}{1 - F(\bar{v})}\right] \Big/ \left[\frac{\bar{v}}{1 - F(\bar{v})} + \frac{1}{f(v)}\right] < 0. \tag{11}$$

Now, if for two different $v, v' < \bar{v}$ such that their allocation probabilities are both strictly positive, then $\Pi_v = \Pi_{v'} = 0$, and we will have

$$\varphi(v) - \frac{\delta}{f(v)} = \varphi(v') - \frac{\delta}{f(v')},$$

or

$$\varphi(v) - \varphi(v') = \delta\left(\frac{1}{f(v)} - \frac{1}{f(v')}\right). \tag{12}$$

Suppose $v < v'$, then $f(v) \geq f(v')$ by our assumption. Since the distribution is regular, we have $\varphi(v) \leq \varphi(v')$. Additionally, we know that $\delta < 0$, and so (12) can hold only if $f(v) = f(v')$, but then the equation says $f(v)(v - v') + F(v) - F(v') = 0$, which cannot be true since $F(v) < F(v')$. Therefore (12) cannot hold under our assumptions.

So far we have shown that in the optimal solution to the above linear program, there can be at most one value $v < \bar{v}$ such that $x(v) > 0$. But then lowering $x(v)$ to 0 affects neither the objective function nor the constraints, and so we obtain a monotone allocation rule [8]. Therefore the solution to the program gives rise to a truthful mechanism, which satisfies Proposition 15.

□

C Proof of Theorem 19

Theorem 19 is a consequence of the two lemmas below and Proposition 18.

Lemma 25. *Twice the optimal representative revenue is a linear upper bound on the optimal unit-demand revenue.*

Proof. Linearity follows simply from the revenue linearity of single-dimensional agents. Consider the distribution of the maximum virtual value (or zero if the maximum virtual value is negative) in the representative environment. Index this distribution by quantile as $\psi_{\max}(q)$. The optimal revenue for any allocation constraint y is $\mathbf{E}_q[\psi_{\max}(q)y(q)]$ which is linear in y; this follows from the proof that the optimal revenue in single-dimensional environments is the virtual surplus maximizer.

We now show that twice the optimal representative revenue upper bounds the optimal unit-demand revenue. To do this we will give two auctions for the representative environment with the allocation constraint y and show that the sum of these auctions' revenue upper bounds the optimal unit-demand revenue for the same constraint. Of course, optimal representative revenue upper bounds each of these auctions revenue.

A mechanism for the unit-demand problem is simply a lottery pricing, i.e., it is a set of lotteries L with each $\ell \in L$ taking the form of $(p^\ell, \pi_1^\ell, \ldots, \pi_m^\ell)$ with $\sum_j \pi_j^\ell \leq 1$. The semantics of a lottery ℓ is that the agent pays the price p^ℓ and then is allocated an item j at random with probability π_j^ℓ; the semantics of the collection of lotteries L is that the agent, upon drawing her type from the distribution, chooses the lottery $\ell \in L$ that maximizes her utility (or none).

Given any collection of lotteries L that satisfies the allocation constraint y we define two auctions for the representative environment that have combined revenue at least that of the collection of lotteries in the unit-demand environment.

The L *mimicking auction* considers the profile of values $\mathbf{v} = (v_1, \ldots, v_m)$ of the representatives and the lottery $\ell \in L$ that would have been selected by the unit-demand agent with these values. It serves the representative j with the highest value with probability π_j^ℓ and charges her $p^\ell - \sum_{j' \neq j} \pi_{j'}^\ell v_{j'} + \mu(\mathbf{v}^{(2)})$ where $\mu(\mathbf{v}^{(2)})$ is the expected utility of the unit-demand agent with valuation profile $\mathbf{v}^{(2)}$ which is $\mathbf{v}$ with v_j replaced with $\max_{j' \neq j} v_{j'}$. Notice that the utility of the winning representative j in this auction is exactly the same as the unit-demand agent less an amount that is a function only of the values of the other representatives, $\mathbf{v}_{-j}$. As the utility of the unit-demand agent is monotone in her value for each item, the utility each representative has for winning is negative when she is not the highest valued representative and positive when she is. Therefore, this auction is incentive compatible, has revenue at least $p_j^\ell - \sum_{j' \neq j} \pi_{j'}^\ell v_{j'}$ on valuation profile $\mathbf{v}$ where j is the highest valued representative, and satisfies allocation constraint y. For a given valuation profile, call the second term in the winning agent's payment the *deficit* of the L mimicking auction.

The motivation for the next auction is that we want to obtain back the deficit lost by the L mimicking auction. Notice that the procedure that charges the highest valued representative the

[8] As a standard practice, we have relaxed the monotonicity condition in the formation of the linear progrma, and only observe that the optimal solution satisfies the monotonicity condition under the assumptions on the valuation distribution.

second highest value and serves with probability $\sum_j \pi_j^\ell$ satisfies the allocation constraint y and more than balances the deficit; however, it may not be incentive compatible.

The *allocation constrained second-price auction* sells to the highest valued representative at the second highest representative's value so as to maximize revenue subject to the allocation constraint y that any representative is served. Consider the distribution of the second order statistic of values and let $\nu_{(2)}(q)$ be the value that the q quantile of this random variable takes on. The optimal revenue obtainable via a second price auction with allocation constraint y is $\mathbf{E}_q[\nu_{(2)}(q)y(q)]$. To obtain this revenue, conditioning on the second highest value being v, with probability $y(\nu_{(2)}^{-1}(v))$ we serve the highest valued representative and charge her v. This auction is incentive compatible and revenue optimal (in expectation) among all second-price procedures that meet the allocation constraint. Therefore, it more than covers the expected deficit of the L mimicking auction.

We have given two incentive compatible auctions for the representative environment with combined expected revenue exceeding the revenue of the lottery pricing L. Therefore, twice the optimal representative revenue is at least the optimal unit-demand revenue. □

Lemma 26. *The pseudo revenue curve $\tilde{R}(\cdot)$ from uniform virtual pricings for a unit-demand agent 2-approximates the optimal representative revenue curve (as a function of q for any q-step constraint).*

Proof. Denote the optimal representative revenue for the q-step constraint as a function of q by the revenue curve $\mathrm{ORR}(q)$. Consider the outcome of the optimal auction for the representative environment with ex ante service constraint q. It sets a uniform virtual price (denoted $\psi(q)$) and serves the agent with the highest virtual value strictly bigger than $\psi(q)$ with probability one. If the probability that the largest virtual value is equal to $\psi(q)$ is strictly positive (which might happen if any virtual value function is constant on an interval, e.g., from ironing), it probabilistically accepts or rejects the maximum virtual value when it is equal to $\psi(q)$ so as to serve with the desired ex ante probability q. The optimal representative revenue can thus be calculated and bounded as follows. Let $(\psi_1, \ldots, \psi_m)$ denote the profile of virtual values of the representatives.

$$\begin{aligned}\mathrm{ORR}(q) &= q \cdot \psi(q) + \mathbf{E}\left[\max_i(\psi_i - \psi(q))^+\right] \\ &\leq q \cdot \psi(q) + \sum\nolimits_i \mathbf{E}\left[(\psi_i - \psi(q))^+\right].\end{aligned}$$

Above, the notation $(\psi_i - \psi(q))^+$ is short-hand for $\max(0, \psi_i - \psi(q))$.

Now we show a lower bound on $\tilde{R}(q)$ for q that does not require probabilistic acceptance in the optimal representative auction described above; denote by $Q \subset [0,1]$ all such quantiles. Let $\mathcal{E}_i$ denote the event that $\psi_j < \psi(q)$ for all $j \neq i$; our lower bound on the pseudo q-step mechanisms revenue will ignore contributions to the virtual surplus from the case that more than one representative has virtual value at least $\psi(q)$.

$$\begin{aligned}\tilde{R}(q) &\geq q \cdot \psi(q) + \sum\nolimits_i \mathbf{E}\left[(\psi_i - \psi(q))^+ \ \mid \ \mathcal{E}_i\right] \cdot \mathbf{Pr}\left[\mathcal{E}_i\right] \\ &\geq q \cdot \psi(q) + (1-q) \cdot \sum\nolimits_i \mathbf{E}\left[(\psi_i - \psi(q))^+\right].\end{aligned}$$

To extend this lower bound on $\tilde{R}(q)$ from $q \in Q$ to all $q \in [0,1]$, consider inserting a virtual value $\psi' = \psi(q) + \epsilon$ with measure zero in the distribution. The q' that corresponds to serving this virtual value or higher has revenue bounded by the formula above but $\psi' \approx \psi(q)$. Keeping the virtual value constant and varying q in the formula interpolates a line between the two revenues. As the pseudo step mechanisms are closed under convex combination, this line gives a lower bound on the pseudo q-step mechanism. Therefore, the bound above on $\tilde{R}(q)$ holds for all q.

To bound $\mathrm{ORR}(q)$ in terms of $\tilde{R}(q)$ we consider two cases. When $q \leq 1/2$ these terms can be directly bounded as the first terms in both bounds are the same and the second terms are within a factor of two of each other (by assumption $1 - q \geq 1/2$). When $q \geq 1/2$ we can compare the bound on $\mathrm{ORR}(q = 1)$ to the bound on $\tilde{R}(q = 1/2)$; these bounds are within a factor of two of each other. Monotonicity (via downward closure) of $\mathrm{ORR}(q)$ and $\tilde{R}(q)$ then implies that they are within a factor of two for any $q \in [1/2, 1]$. □