

The Simple Economics of Approximately Optimal Auctions[†]

Saeed Alaei
 Dept. of Computer Science
 Cornell University
 Ithaca, NY
 saeed.a@gmail.com

Hu Fu
 New England Lab
 Microsoft Research
 Cambridge, MA
 hufu@microsoft.com

Nima Haghpanah
 EECS Dept.
 Northwestern University
 Evanston, IL
 nima.haghpanah@gmail.com

Jason Hartline
 EECS Dept.
 Northwestern University
 Evanston, IL
 hartline@northwestern.edu

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 [1] show that the optimal auction of Myerson [2] 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 implement the marginal revenue maximization mechanism. Our contributions are three fold: we characterize the settings for which marginal revenue maximization is optimal (by identifying an important condition that we call *revenue linearity*), 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). Because the marginal revenue mechanism is optimal for well-studied single-dimensional agents, our generalization immediately extends many approximation results for single-dimensional agents to more general preferences.

Finally, one of the biggest open questions in Bayesian algorithmic mechanism design is in developing methodologies that are not brute-force in size of the agent type space (usually exponential in the dimension for multi-dimensional agents). Our methods identify a subproblem that, e.g., for unit-demand agents with values drawn from product distributions, enables approximation mechanisms that are polynomial in the dimension.

Keywords-Bayesian mechanism design, Approximation, Marginal revenue

I. INTRODUCTION

Marginal revenue plays a fundamental role in microeconomic theory. For example, a monopolist providing a commodity to two markets each with its own concave

revenue (as a function of the supply provided to that market) optimizes her profit by dividing her total supply to equate the marginal revenues across the two markets. Moreover, this central economic principle governs classical auction theory. Myerson [2] characterizes profit maximizing single-item auction as formulaically optimizing the *virtual value* of the winner; Bulow and Roberts [1] reinterpret Myerson’s virtual value as the marginal revenue of a certain concave revenue curve.

Because it is simple and intuitive, the Myerson-Bulow-Roberts approach provides the basis for much of Bayesian auction theory. Unfortunately though, this theory has been limited to settings where agents have linear single-dimensional preferences, i.e., where an agent’s utility is given by her value for service less her payment. Consequently, Bayesian auction theory is often similarly limited. With more general forms of agent preferences, especially multi-dimensionality, e.g., for multi-item auctions, or non-linearity, e.g., risk aversion or budgets, auction theory is complex, less versatile, and often not well understood.

Our main result is to show that hidden under the complexity of optimal mechanism design problems for agents with multi-dimensional and non-linear (henceforth: general) preferences is marginal revenue maximization. The approach of marginal revenue maximization is to express a multi-agent mechanism design problem as a composition simple single-agent mechanism design problems, i.e., from the construction of the appropriate notion of revenue curves. This new approach for general preferences uncovers a condition we refer to as *revenue linearity* that is satisfied by all linear single-dimensional preferences and governs the performance of the marginal revenue mechanism. When the single-agent problems are revenue linear, marginal revenue maximization is optimal and the Myerson-Bulow-Roberts mechanism generalizes exactly. When the single-agent problems are approximately revenue linear, marginal revenue maximization is approximately optimal (though the composition of the single-agent mechanisms to implement marginal revenue maximization requires new techniques). Finally, because our marginal revenue approach is structurally similar to the classical approach, many results for agents with linear single-dimensional preferences approximately and automatically

[†]Full version of this paper is available at <http://arxiv.org/abs/1206.3541>. This work was done in part while all authors were at Northwestern University. The second author was supported by NSF grants CCF-0643934 and AF-0910940 at Cornell University, the remaining by NSF CCF-0830773 at Northwestern University.

extend to general preferences.

A central result to classical auction theory comes from reinterpreting the Myerson-Bulow-Roberts mechanism (i.e., for maximizing marginal revenue) in the special case of symmetric agents. As an example of the benefits of our approach, compare this classical reinterpretation with a similar reinterpretation of our results. In the classical setting there is a single item for sale and agents with i.i.d. values for it; in our setting there is a single item for sale which the seller can configure on one of several ways and agents have i.i.d. values for each configuration, e.g., a car that can be painted red or blue (importantly, the seller sets the configuration and the buyer cannot change it).¹

Selling a car: Classical auction theory says that (a) the optimal way to sell an object (henceforth: a car) to a single agent with value drawn from a uniform distribution on $[0, 1]$ is to post a take-it-or-leave-it price of $1/2$, (b) the optimal way to sell a car to one of multiple agents with uniformly distributed values is to run a second-price auction with reserve price $1/2$, and (c) more generally the optimal way to sell the car to multiple agents with i.i.d. values is to run the second price auction with the same reserve price that would be offered as a take-it-or-leave-it price to one agent (assuming the distribution satisfies some mild assumptions).

Selling a red-or-blue car: Consider selling a car that, on sale, can be painted one of two colors, red or blue.² Our theory says that (a) the optimal way to sell a red-or-blue car to a single agent with values for the different colors each drawn independently and uniformly from $[0, 1]$ is to post a take-it-or-leave-it price of $\sqrt{1/3}$ for either color, (b) the optimal way to sell a red-or-blue car to one of multiple agents each with i.i.d. uniform values for each color is to run the second-price auction with reserve $\sqrt{1/3}$ and allow the winning agent to choose her favorite color on sale, and (c) more generally to sell a red-or-blue car to one of multiple agents each with values drawn i.i.d. (from a distribution that satisfies the same mild assumptions as above) for each color, the second price auction with the reserve price equal to the same price that would be offered to a single agent is (at worst) a 4-approximation to the optimal auction.

It should be noted that reducing the multi-dimensional preference to a single-dimensional preference by always selling the winning agent her favorite color is, though natural and practical, not generally optimal beyond $U[0, 1]$. Even for a single agent with values for both colors distribution

¹The red-or-blue car example is slightly unnatural as a forward auction (i.e., the auctioneer is selling); however, the analogous reverse auction (i.e., the auctioneer is buying) is an important problem in procurement. For instance the government may wish to hire a contractor to build a bridge. Contractors can build different kinds of bridges. From bids of the contractors over the different bridges the auctioneer selects a kind of bridge to procure, which contractor to procure it from, and how much is to be paid. Our results for reverse auctions are analogous to those for forward auctions; interested readers can find the details in the full version of the paper.

²This result generalizes to more colors but with a different reserve price.

uniformly on $[5, 6]$, an analysis of Thanassoulis [3] shows that the optimal pricing does not sell the agent her favorite item subject to a reserve (in fact, it is not even deterministic).

Approach: We focus on *service constrained environments* where in any outcome the mechanism produces, each agent is either considered served or unserved. The designer has a feasibility constraint that governs which subset of agents can be simultaneously served, but other aspects of the outcome, e.g., payments, are unconstrained. This model allows additional unconstrained attributes of the service (e.g., the color of the car in the previous example). We assume that the space of mechanisms is closed under convex combination which allows for randomized mechanisms.

The agents in the mechanism have independently but not necessarily identically distributed preferences (a.k.a., types). We do not place any assumption on the agent preferences other than they are expected utility maximizers. This includes the most challenging preference models in Bayesian mechanism design such as multi-dimensionality, public or private budgets, and risk-aversion (e.g., as given by a concave utility function).

Revenue curves result from the following single-agent mechanism design problem. Consider a single agent with preferences drawn from a known distribution. Via the taxation principle (e.g., Wilson [4]) the outcomes of a mechanism, for all possible preference reports the agent might make, can be viewed as a menu where the agent selects her favorite outcome by making the appropriate report. This menu may contain outcomes that are randomized and for this reason we refer to it as a *lottery pricing*. Ex ante, i.e., in expectation over the distribution of the agent's preference, a lottery pricing induces a probability with which the agent receives an outcome that corresponds to service, and an expected payment, i.e., revenue.

As every lottery pricing induces an ex ante service probability and expected revenue, we can ask the optimization question of identifying the lottery pricing with a given ex ante service probability that has the highest expected revenue. Considering this optimal revenue as a function of the ex ante service probability gives rise to the agent's *revenue curve*. Important in the construction of revenue curves are the lottery pricings, i.e., single-agent mechanisms, that give the optimal revenue for each ex ante service probability. As the space of lottery pricings is closed under convex combination, the revenue curves are always concave. The marginal revenue curve is the derivative of the revenue curve with respect to ex ante service probability.

As discussed in the opening paragraph, the standard economic intuition suggests that a monopolist splitting the sale of a commodity between two markets should do so to equate marginal revenue. There is an intuitive algorithmic reinterpretation of this fact. If we order the consumers of each market by willingness to pay and attribute to each consumer the change in revenue from adding that consumer

(i.e., the marginal revenue), then the total revenue of an allocation is the sum of the marginal revenues of each consumer served. A simple algorithm for optimizing this cumulated marginal revenue is to repeatedly allocate a unit to the market that has the highest marginal revenue at its current allocation (until the good is totally allocated or marginal revenues are non-positive). Clearly this process results in a final allocation that roughly equates where the markets' marginal revenues at the quantities allocated as in the microeconomic interpretation. This allocation is optimal.

The main contribution of this paper is a methodology for constructing multi-agent mechanisms from the simple single-agent lottery pricings that define the revenue curve. The main task of such a construction is to specify a method for combining the single-agent mechanisms into a multi-agent mechanism that is both feasible with respect to the service constraint and obtains good revenue. We refer to the family of mechanisms that take the following form as *marginal revenue mechanisms*.

- 1) Map each agent type (which may lie in an arbitrary type space) to a *quantile* in $[0, 1]$.
- 2) Calculate the marginal revenue of each agent as the derivative of her revenue curve at her quantile.
- 3) Select for service the set of agents that maximize cumulative marginal revenue subject to feasibility.
- 4) Calculate for each agent the appropriate non-service aspects of the outcome, e.g., payments.

Thus far in the discussion only Steps 2 and 3 should be clear. The remaining steps are non-trivial in general and a main issue that we will be resolving.

Results: This paper generalizes the marginal-revenue approach for agents with single-dimensional linear preferences which is due to Bulow and Roberts [1] to general preferences. Our main algorithmic contribution is to generalize Steps 1 and 4 thereby reducing service constrained multi-agent mechanism design problems to a collection of (single agent) ex ante constrained lottery pricing problems. There are a number of challenges in this endeavor. First, revenue equivalence does not hold for general preferences (which is used in the proof of optimality for single-dimensional preferences). Second, there is not a natural ordering on preferences for general preferences (making it difficult to map preferences to quantiles, i.e., Step 1). Third, the set of agents served by the marginal revenue mechanism may be randomized. None of these issues are present for single-dimensional linear preferences. Finally, the reduction focuses attention on this ex ante lottery pricing problem as a fundamental building block of good mechanisms. For general preferences these lottery pricing problems have not previously been considered in the literature.

Orthogonal to the question of implementing the marginal revenue mechanism for general preferences are questions of quantifying its performance. Via the Myerson-Bulow-Roberts analysis it is known that for single-dimensional

linear preferences, the marginal revenue mechanism is optimal. As a first step in understanding the performance of the mechanism more generally we give a new derivation of the optimality for single-dimensional agents that exposes a previously unobserved property of single-dimensional preferences which we refer to as *revenue linearity*. Generally, i.e., beyond single-dimensional preferences, the optimality of the marginal revenue mechanism is implied by revenue linearity. Moreover, if general single-agent lottery pricing problems are α -approximately revenue linear (e.g., bounded from below by a linear function and from above by α times the function), then marginal revenue maximization is an α -approximation to the optimal mechanism.

Revisiting our red-or-blue car examples above, (a) is a description of the optimal unconstrained lottery pricing, (b) is a consequence of the revenue-linearity of types that are uniformly distributed on a multi-dimensional hypercube, and (c) is a consequence of 4-approximate revenue linearity for agents with types drawn from any product distribution.

It is important to contrast the simplicity of the marginal revenue approach with recent algorithmic results in Bayesian mechanism design for general agent preferences. Recently, Alaei et al. [5] and Cai et al. [6, 7, 8] gave polynomial time mechanisms for large important classes of Bayesian mechanism design problems; the former considered general preferences in service constrained settings (as does this paper) and the latter considered multi-dimensional additive preferences. The two main conclusions of these works are that (a) optimal mechanisms continue to have weighted maximization at their core, and (b) the appropriate weights (i.e., virtual values) are stochastic and can be solved for as a convex optimization problem, e.g., via ellipsoid method, that takes into account the feasibility constraint and the distribution over types of all agents. (This latter result is simply because the space of mechanisms is convex, any point on the interior of a convex set can be implemented by a convex combination of vertices, and vertices correspond to linear, a.k.a., weighted, optimization.)

Our results are distinct from these algorithmic results in several respects. First, the weights in our derivation have a natural economic interpretation as marginal revenue. Second, the weights in our derivation can be found easily from solutions to the single-agent lottery pricing problems and are not derived from the solution to an additional optimization problem. Third, in most cases, the weights in our derivation depend only on the single-agent problem and not on the feasibility constraint or presence of other agents. Therefore, our approach affords significant structural simplification and interpretation that enables the consequences previously enumerated. Finally, one of the biggest open questions in the above algorithmic work is in developing approaches that are not brute-force in each agent's type space. For example, our approach gives mechanisms that have runtime polynomial in the dimension of the type space (i.e., logarithmic in the

size of the type space) for multi-dimensional unit-demand agents with values from product distributions.

Organization: In Section II we review the Myerson-Bulow-Roberts single-dimensional linear agent model, their approach to Bayesian optimal mechanism design, and give a new proof that the marginal revenue mechanism is revenue optimal. The proof follows from an argument that for single-dimensional linear agents a class of single-agent lottery pricing problems satisfies a natural revenue-linearity property. In Section III we formalize our service constrained model for general preferences and generalize the marginal revenue derivation to general preferences that satisfy the previously identified revenue linearity property. In Section IV we give general methods for implementing the marginal revenue mechanism (e.g. Steps 1 and 4) for general preferences regardless of revenue linearity, and in Section V we show that approximate linearity implies approximate optimality.

II. WARM-UP: SINGLE-DIMENSIONAL LINEAR PREFERENCE

In this section we warm up by giving a new proof that the marginal revenue mechanism is revenue optimal for agents with single-dimensional linear preferences. We will introduce many concepts (which were not present in previous proofs) that make our generalization possible. The basic approach is as follows. We formulate an important class of lottery pricing problems the solution to which define a revenue curve. We show that single-dimensional agents are *revenue linear* in the sense that it is optimal to decompose the allocation to any agent as a convex combination of the solutions to these lottery pricing problems. Finally, we observe that this implies that the optimal revenue can be expressed in terms of the cumulative (over agents served) marginal revenue (given by the derivative of the revenue curve). The marginal revenue mechanism optimizes this latter term point-wise and, therefore, also in expectation. In the interest of brevity we will keep the discussion informal, a formal treatment is given in Section III with proofs deferred to the full version of the paper.

Model: A single-dimensional linear agent has a private type $v \in \mathbb{R}_+$ drawn at random with cumulative distribution function F and density function f . For outcome, let (x, p) denote receiving a good or service with probability x and making expected payment p . For such an outcome, an agent with value v has a linear utility $u = vx - p$.

The geometry of single-dimensional auction theory is more readily apparent when we index an agent's strength relative to the distribution (instead of values). Let $V(q) = F^{-1}(1 - q)$ be the *inverse demand curve*, i.e., $V(\hat{q})$ is the posted price that would be accepted by the \hat{q} measure of highest-valued agents (and rejected by all others). The *quantile* of an agent is the measure of agents with higher values, i.e., for value v the agent's quantile is $q = V^{-1}(v)$. Importantly, for v drawn from the distribution F , $q = V^{-1}(v)$ is

uniform on $[0, 1]$ (therefore, expectations of functions of q are given by integrals with probability density one).

A multi-agent mechanism design problem is given by n such single-dimensional agents each with their respective inverse demand curves (which may be distinct) and a feasibility constraint governing the subsets of agents that can be simultaneously served. E.g., for a single-item auction, the feasibility constraint says that at most one agent can be served; more generally, the feasibility constraint could be given by any set system. In the interim stage, i.e., when an agent knows her own value but not the values of other agents, the mechanism looks to the agent like a single-agent mechanism. It will thus be sufficient for most of the analysis of optimal multi-agent mechanisms to consider the appropriate single-agent problems.

From the perspective of an agent in a single-agent mechanism and as a function of the agent's report, the agent is served with some probability and makes some expected payment. We can view this function as a menu of service probabilities and expected payments where the agent selects her favorite outcome by submitting the corresponding report. Notice that depending on the agent's value for service, she may choose different outcomes. We may as well index the outcomes in the menu by the quantile of the agent that selects the outcome, i.e., agent with quantile q chooses outcome $(x(q), p(q))$. We assume that outcome $(x, p) = (0, 0)$ is in the menu. This relabeling and assumption imply *incentive compatibility* and *individual rationality*, respectively, i.e., for all $q, q' \in [0, 1]$,

$$V(q)x(q) - p(q) \geq V(q)x(q') - p(q'), \quad (\text{IC})$$

$$V(q)x(q) - p(q) \geq 0. \quad (\text{IR})$$

We call such a menu a *lottery pricing*. When the lottery pricing is induced in the interim stage of a multi-agent mechanism, then the constraints above are *Bayesian incentive compatibility* and *interim individual rationality*.

The Myerson [2] characterization of Bayesian incentive compatible mechanisms applies to lottery pricings and implies that the *allocation rule* $x(\cdot)$ is monotone non-increasing and the *payment rule* is given precisely as a function of $x(\cdot)$. An important consequence of the latter part of this characterization is *revenue equivalence*. We will make use of both monotonicity and revenue equivalence below, though the specific form of the payment rule will not be important.

Constrained Lottery Pricings: Given such a lottery pricing and a distribution over the agent's value, an ex ante expected payment $\mathbf{E}_q[p(q)]$ and ex ante probability of service $\mathbf{E}_q[x(q)]$ are induced. The single-agent lottery pricing problem that forms the basis for the marginal revenue mechanism is the following. Given an ex ante constraint \hat{q} on the probability with which the agent is served, find the lottery pricing that serves the agent with probability \hat{q} and maximizes revenue.

Definition 1. The *revenue curve* $R(\hat{q})$ is defined for all $\hat{q} \in [0, 1]$ as the optimal lottery pricing revenue for ex ante constraint \hat{q} .

In order to show that convex combinations of optimal ex ante constrained lottery pricings are optimal in general, we need to consider a more general lottery pricing problem. Notice that the ex ante constrained problem gives an (equality) constraint on the total probability that the agent is served over all quantiles she may have. To get more fine-grained control over the lottery pricing we additionally allow upper bounds to be specified on the total probability of allocation for subsets of quantiles. Consider the following lottery pricing problem: Given a monotone concave function $\hat{X}(q)$, find the optimal lottery pricing where the ex ante probability of allocating to any \hat{q} measure of quantiles is at most $\hat{X}(\hat{q})$ for all $\hat{q} \in [0, 1]$ and exactly equal to $\hat{X}(\hat{q})$ at $\hat{q} = 1$.

To see why this constrained lottery pricing problem is the right one to consider, notice the following. Because any allocation rule is monotone, meaning stronger quantiles receive no lower probability of service than weaker quantiles, the sets of measure \hat{q} for which the constraint of service probability at most $\hat{X}(\hat{q})$ is binding correspond exactly to the strongest \hat{q} measure of quantiles. For allocation rule $x(\cdot)$ the probability of service to the strongest \hat{q} measure of agents is exactly $X(\hat{q}) = \int_0^{\hat{q}} x(q) dq$. We refer to $X(\cdot)$ as the *cumulative allocation rule*. Thus, the allocation constraint is exactly, $X(\hat{q}) \leq \hat{X}(\hat{q})$ for all $\hat{q} \in [0, 1]$ (with equality for $\hat{q} = 1$). Of course we can view the cumulative allocation rule X of x as a constraint and observe that x satisfies the constraint with equality. Moreover, x is the allocation rule that satisfies X as a constraint that has the highest probability on stronger (i.e., lower) quantiles. Therefore, for any constraint \hat{X} (with corresponding $\hat{x}(q) = \frac{d}{dq} \hat{X}(q)$) is met by allocation rule x that relatively has allocation probability shifted from stronger quantiles to weaker quantiles. Specifically, \hat{x} *majorizes* x .

Definition 2. $\text{Rev}[\hat{x}]$ is the optimal revenue of any lottery pricing that satisfies the allocation constraint \hat{x} (via its cumulative allocation rule \hat{X}).

Recall our ex ante constrained lottery pricing where we wish to serve the agent with ex ante probability \hat{q} . A *posted price* is parameterized by a single price and is a simple example of a lottery pricing (i.e., one that is deterministic), the two menu items are to be served and pay the price or not to be served and pay nothing. The agent prefers service when her value exceeds the price and, otherwise, she prefers no service. For an agent with inverse demand curve $V(\cdot)$, the posted price that serves with probability \hat{q} is $V(\hat{q})$. It gives expected revenue $\hat{q} \cdot V(\hat{q})$ (which is at most $R(\hat{q})$). Its allocation rule $\hat{x}^{\hat{q}}$ is the reverse step function that is one on quantiles $[0, \hat{q}]$ and then zero on $(\hat{q}, 1]$. This rule

has the most service probability on strong quantiles of all allocation rules that satisfy the ex ante allocation constraint \hat{q} . Of course, the revenue it generates $\hat{q} \cdot V(\hat{q})$ may not be a concave function of \hat{q} and it must be that the revenue curve $R(\cdot)$ is concave. In can be shown, in fact, that $R(\cdot)$ is exactly the concave hull of $\hat{q} \cdot V(\hat{q})$ and the optimal lottery for any \hat{q} is given by a posted pricing or if $R(\cdot)$ is linear at \hat{q} equal to the convex combination of two posted pricings (corresponding to the boundary of the interval containing \hat{q} on which $R(\cdot)$ is linear). The allocation rule of this convex combination is a convex combination of the appropriate two reverse step functions and, in the sense described above, has service probability shifted from stronger quantiles to weaker quantiles. This specific form (which is not obvious) is not important for our re-derivation of the optimal mechanism, what is important (and obvious) is the following.

Proposition 1. For any ex ante constraint, the optimal lottery has weaker a allocation rule higher revenue than price posting.

Revenue Linearity: We are now ready to give the new derivation of the marginal revenue mechanism and its optimality. We start with the central definition.

Definition 3. The single agent lottery pricing problems are *revenue linear* if $\text{Rev}[\cdot]$ is linear. I.e., the optimal revenue for constraint $\hat{x} = \hat{x}^A + \hat{x}^B$ is $\text{Rev}[\hat{x}] = \text{Rev}[\hat{x}^A] + \text{Rev}[\hat{x}^B]$.

Now consider the following two lower bounds on the optimal revenue for any allocation constraint \hat{x} . The constraint \hat{x} is a monotone non-increasing function. As reverse step functions provide a basis for such functions, we can view \hat{x} as a convex combination of reverse step functions. This convex combination can be sampled by drawing \hat{q} at random from the distribution $G^{\hat{x}}$ with density $-\hat{x}'(q) = \frac{d}{dq} \hat{x}(q)$ and then posting price $V(\hat{q})$ (and allocation rule $\hat{x}^{\hat{q}}$). The allocation rule of the convex combination is exactly \hat{x} , its expected revenue lower bounds the optimal revenue subject to the constraint \hat{x} . A second approach is to use, instead of the posted pricing $V(\hat{q})$, the optimal lottery pricing for ex ante constraint \hat{q} . As the allocation rule for each of these mechanisms is weaker than the corresponding posted pricing allocation rule, the convex combination of the allocation rules (denote it by x) is weaker than the allocation constraint \hat{x} . Therefore, it is feasible for \hat{x} and its revenue gives a lower bound on the optimal revenue for \hat{x} . Formally,

$$\begin{aligned} \text{Rev}[\hat{x}] &\geq \mathbf{E}_{\hat{q} \sim G^{\hat{x}}} [-\hat{x}'(\hat{q})R(\hat{q})] \\ &= [-\hat{x}(\hat{q})R(\hat{q})]_0^1 + \mathbf{E}_q [R'(q)\hat{x}(q)] \\ &= \mathbf{E}_q [R'(q)\hat{x}(q)]. \end{aligned}$$

The second equality follows from the definition of expectation and integration by parts, and the third equality follows from $R(1) = R(0) = 0$ (minor assumption: if we always

serve or never serve the agent we obtain no revenue). This construction motivates the following definition.

Definition 4. The *marginal revenue* for an allocation constraint \hat{x} is $\text{MR}[\hat{x}] = \mathbf{E}_q[R'(q)\hat{x}(q)]$.

The definition of revenue linearity and the definition of the revenue curve (as the optimal revenue subject to the ex ante constraint \hat{q}) immediately imply the following theorem.

Theorem 2. *If the single-agent lottery pricings are revenue linear then the optimal revenue for an allocation constraint is equal to its marginal revenue, i.e., for all \hat{x} , $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$.*

To show that the marginal revenue is equal to the optimal revenue, we must only prove revenue linearity. The proof of this theorem is a simple consequence of revenue equivalence and the fact that the optimal revenue for ex ante constraint \hat{q} exceeds the posted pricing revenue from $V(\hat{q})$ but has a weaker allocation rule (Proposition 1). We defer it to the full version of the paper.

Theorem 3. *An agent with single-dimensional linear utility is revenue linear.*

Corollary 4. *For single-dimensional agents the optimal revenue for an allocation constraint is equal to its marginal revenue, i.e., for all \hat{x} , $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$.*

Multi-agent Mechanisms: We now look at the problem of optimizing expected revenue over agents in a multi-agent mechanism. The following is the standard argument from auction theory. For each agent, revenue is given by marginal revenue (Corollary 4). Relax incentive constraints (namely: monotonicity of the allocation rule) and optimize marginal revenue pointwise. Meaning, when the agent quantiles are $\mathbf{q} = (q_1, \dots, q_n)$ select the allocation $\mathbf{x} = (x_1, \dots, x_n)$ to maximize the cumulative marginal revenue $\sum_i R'_i(q_i) \cdot x_i$ subject to feasibility of \mathbf{x} (e.g., for a single-item auction, serve the agent with the highest positive marginal revenue, or none if the marginal revenues are all negative). Now check that the incentive constraints hold: Notice that since revenue curves are concave, marginal revenues are monotone non-increasing, for any agent a stronger (lower) quantile corresponds to a weakly higher marginal revenue, and so the induced allocation rule is monotone. Furthermore, as these allocations optimize marginal revenue pointwise for all profiles of agent quantiles, they certainly also maximize marginal revenue in expectation over the agent quantiles.

Comparing the above construction with the marginal revenue mechanism framework described in the introduction, the missing Steps 1 and 4 are simple. For Step 1, the mapping from value to quantile is given by $V_i^{-1}(\cdot) = 1 - F_i(\cdot)$ for each agent i as described above. For Step 4, the appropriate payments can be calculated pointwise as follows: Agents that are not served pay nothing; an agent i that is

served pays the value $V_i(\hat{q}_i)$ corresponding to her critical quantile \hat{q}_i , i.e., the quantile after which she would no longer be served (via the payment identity).

Theorem 5. *For single-dimensional linear agents, the marginal revenue mechanism is incentive compatible and revenue optimal.*

III. MULTI-DIMENSIONAL AND NONLINEAR PREFERENCES

Bayesian mechanism design: An agent has a private type t from type space T drawn from distribution F with density function f . The agent may be assigned outcome w from outcome space W . This outcome encodes what 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 t and deterministic outcome w her utility is $u(t, w)$, and when w is drawn from a distribution her utility is $\mathbf{E}_w[u(t, w)]$.³ 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 $\{1, \dots, n\}$ and each agent i may have her own distinct type space T_i , utility function u_i , etc. In this paper we only consider settings where different agents' types are drawn independently from their respective distributions. A *direct revelation* mechanism takes as its input a profile of types $\mathbf{t} = (t_1, \dots, t_n)$, and then outputs for each agent i an outcome $\tilde{w}_i(\mathbf{t})$. The ex post outcome rule of the mechanism is $\tilde{w}_i : T_1 \times \dots \times T_n \rightarrow \Delta(W_i)$. 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 $t_j \sim F_j$ for each $j \neq i$. We say that a mechanism is *Bayesian incentive compatible* if

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

A mechanism is *interim individually rational* if

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

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., $\mathbf{E}_t[\sum_i \text{Payment}(\tilde{w}_i(t_i))]$; 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 *service constrained environment* a mechanism produces a (potentially) randomized outcome which can be viewed as a joint distribution over deterministic outcomes provided to each agent. A deterministic outcome for an agent is distinguished as a *service* or *non-service* outcome with $\text{Alloc}(w) = 1$ or

³This form of utility function allows for encoding of budgets and risk aversion; we do not require quasi-linearity.

$\text{Alloc}(w) = 0$, respectively. There is a feasibility constraint restricting the set of agents that may be simultaneously served. A randomization over feasible subsets is feasible. For a randomized outcome w an agent's probability of service is $\text{Alloc}(w) \in [0, 1]$. 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 example, payments are part of the outcome but 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.

Revenue Curves: The only aspect of the marginal revenue approach that translates identically from single-dimensional preferences to general preferences is the definition of the \hat{q} ex ante optimal lottery pricing. This is the lottery pricing (i.e., collection of outcomes where the agent is permitted to choose her type-dependent favorite) denoted $\tilde{w}^{\hat{q}}(\cdot)$ with the constraint that $\mathbf{E}_t[\text{Alloc}(\tilde{w}^{\hat{q}}(t))] = \hat{q}$ that optimizes revenue. For the optimal $\tilde{w}^{\hat{q}}(\cdot)$, the revenue curve for the agent is then given by $R(\hat{q}) = \mathbf{E}_t[\text{Payment}(\tilde{w}^{\hat{q}}(t))]$ as per Definition 1.

Allocation rules: Our first challenge, then, in generalizing the marginal revenue approach to general preferences is that we cannot make an upfront transformation from the type space T of an agent to a $[0, 1]$ quantile space ordered by the strength of the agent. E.g., if the type is multi-dimensional then it is unclear which is stronger, a higher value in one dimension and lower in another or vice versa. In fact, which is stronger often depends on the context, e.g., the competition from other agents.

Our approach is based on two observations. First, relative to a mechanism and for a particular agent, the relevant part of the mechanism is the (interim) outcome rule $\tilde{w}(\cdot)$. For a given outcome rule $\tilde{w}(\cdot)$ an ordering on types by strength can be defined. Simply, a type that is more likely to be served is stronger than a type that is less likely to be served. I.e., t is stronger than t' relative to $\tilde{w}(\cdot)$ if $\text{Alloc}(\tilde{w}(t)) \geq \text{Alloc}(\tilde{w}(t'))$. Second, (by the above mapping) any outcome rule $\tilde{w}(\cdot)$ induces an allocation rule $x(\cdot)$ that maps quantile to service probability. This allocation rule has a simple intuition in discrete type spaces: For each type $t \in T$ make a rectangle of width equal to the probability of the type $f(t)$ and height equal to the service probability

of the type $\text{Alloc}(\tilde{w}(t))$. Sort the types in decreasing order of height; the resulting monotone non-increasing piecewise constraint function from $[0, 1]$ to $[0, 1]$ is the allocation rule. This is generalized for continuous distributions as follows.

Definition 5. For an agent with $t \in T$ drawn from distribution F and outcome rule $\tilde{w}(\cdot)$, the *allocation rule* mapping quantiles to service probabilities is given by $x(\hat{q}) = \inf\{y : \Pr_{t \sim F}[\text{Alloc}(\tilde{w}(t)) \geq y] \leq \hat{q}\}$.

Optimal Lottery Pricing: With the definition of allocation rules for any lottery pricing in hand, allocation constrained lottery pricings generalize naturally. Even though the order on types may change from one lottery pricing to another, we can still ask for the lottery pricing with the optimal revenue subject to a constraint on its allocation rule. The optimal lottery pricing for allocation constraint \hat{x} with cumulative allocation constraint \hat{X} is given by the outcome rule $\tilde{w}(\cdot)$ that optimizes expected revenue subject to its corresponding allocation rule x with cumulative allocation rule X satisfying $X(\hat{q}) \leq \hat{X}(\hat{q})$ for $\hat{q} \in [0, 1]$ with equality at $\hat{q} = 1$. As per Definition 2 the optimal revenue for allocation constraint \hat{x} is denoted $\text{Rev}[\hat{x}]$.

We will generally denote by x the optimal allocation rule for constraint \hat{x} . The ex ante constraint on total service probability by \hat{q} is given by the reverse step function at \hat{q} denoted $\hat{x}^{\hat{q}}$; the corresponding allocation rule of the \hat{q} optimal lottery pricing is denoted $x^{\hat{q}}$. Therefore, $R(\hat{q}) = \text{Rev}[\hat{x}^{\hat{q}}]$.

Revenue Linearity and Marginal Revenue: Revenue linearity and marginal revenue have the same definitions (Definition 3 and Definition 4) as for single-dimensional preferences. The marginal revenue of an allocation constraint is $\text{MR}[\hat{x}] = \mathbf{E}_q[R'(q)\hat{x}(q)]$. By its construction as the revenue of the appropriate convex combination of ex ante constrained mechanisms it is a lower bound on the optimal revenue, i.e., $\text{Rev}[\hat{x}] \geq \text{MR}[\hat{x}]$. Again by its construction, revenue linearity would imply it is equal to the optimal revenue.

Definition 6. The *optimal marginal revenue* for a service constrained environment with general agent preferences is the expected revenue (equal to expected cumulative marginal revenue) of the *single-dimensional analog* with each agent replaced by a single dimensional agent with the same revenue curve.

The framework thus defined affords two very natural questions. First, as for general preferences optimal revenue may be strictly larger than marginal revenue, does the optimal marginal revenue approximate the optimal revenue? Second, as the implementation of the marginal revenue mechanism for single-dimensional preferences does not directly extend to general preferences (e.g., Steps 1 and 4), can we implement the marginal revenue mechanisms? In the remainder of this section we will focus on the revenue-linear special case, where the optimal revenue is the optimal cumulative

marginal revenue, and we will answer the implementation question. Non-revenue-linear environments are considered in the next sections.

Implementation with Revenue Linearity: We show now that the marginal revenue mechanism generalizes exactly for general preferences that satisfy revenue linearity. Moreover, in this case the marginal revenue mechanism inherits all of the nice properties of the marginal revenue mechanism for single-dimensional preferences. Namely, it deterministically selects the set of agents to serve, it is dominant strategy incentive compatible (truthful reporting is a best response for any actions of the other agents), and the mapping from types to quantiles to marginal revenues *context free*,⁴ it does not depend on the feasibility constraint or other agents in the mechanism, and deterministic. The mechanism, however, is optimal among the larger class of randomized and Bayesian incentive compatible mechanisms. As motivation for this result, we will show subsequently that there are multi-dimensional preferences that are revenue linear, e.g., when multi-dimensional values are uniformly distributed on a hypercube.

The main challenge of implementing the marginal revenue mechanism is in specifying Step 1, i.e., the mapping from types to quantiles, and Step 4, i.e., selecting the appropriate outcomes for the set of agents that are served. If, however, each agent’s types are orderable by the following definition, then both steps are essentially identical to the single-dimensional case.

Definition 7. 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 \hat{x} , the optimal outcome rule \tilde{w} induces an allocation rule that is greedy by this ordering with ties between types in a same equivalence class broken uniformly at random.⁵

Orderability may look like a stringent and unlikely condition to hold generally. We note that it holds for single-dimensional agents and we show now, more generally, that it is a consequence of revenue linearity.

Theorem 6. *For any single-agent problem, revenue linearity implies orderability. Moreover, for an ex ante constraint \hat{q} for which $R(\hat{q})$ is locally linear, the optimal lottery pricing is a full lottery.*⁶

Given the properties above, the marginal revenue mech-

⁴Note that this contrasts with recent algorithmic work in multi-dimensional optimal mechanism design where the optimal mechanism is characterized by mapping types stochastically to “virtual values” and this mapping is solved for from the feasibility constraint and the distributions of all agents types. See Alaei et al. [5] and Cai et al. [6, 7].

⁵By greedy by the given ordering, we mean process each equivalence class in order and serve the corresponding types with as much probability as possible subject to the allocation constraint.

⁶A full lottery is one where each type is either served or not served with probability one.

anism is easy to define in the revenue-linear settings.

Definition 8. The *marginal revenue mechanism* for orderable agents works as follows.

- 1) Map reported types $\mathbf{t} = (t_1, \dots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \dots, q_n)$ via the implied ordering.⁷
- 2) Calculate the marginal revenue of each agent i as $R'_i(q_i)$.
- 3) For each agent i , calculate the maximum quantile \hat{q}_i that she could possess and be in the marginal revenue maximizing feasible set (breaking ties consistently).
- 4) Offer each agent i the \hat{q}_i ex ante optimal pricing.

Proposition 7. *The marginal revenue mechanism deterministically selects a feasible set of agents to serve and is dominant strategy incentive compatible.*

Proposition 8. *In service constrained environments with revenue-linear agents, the marginal revenue mechanism obtains the optimal marginal revenue (which equals the optimal revenue).*

As an example of multi-dimensional single-agent problem that is revenue linear, we show in the full version of the paper the following theorem. From this theorem, Proposition 8 enables the derivation of the optimal auction for the red-or-blue car example described in Section I.

Theorem 9. *A unit-demand agent with values for m variants of a service distributed uniformly on $[0, 1]^m$ is revenue-linear.*

IV. IMPLEMENTATION

In the full version of the paper, we give a proof for the following theorem, the key to which is a variation of the technique of vector majorization of Hardy et al. [9].

Theorem 10. *For service constrained environments the interim allocation rules of the marginal revenue mechanism can ex post implemented by a Bayesian incentive compatible mechanism.*

The mechanism given by Theorem 10 gives polynomial time reduction to the single-agent ex ante pricing problems. However, unlike the revenue linear case, it is only Bayesian incentive compatible and the mapping from types to quantiles to marginal revenues is not context free (i.e., it depends on the feasibility constraint and the competition from other agents).

In the following we give another construction for the special case where the parameterized family of \hat{q} ex ante optimal pricings satisfy a natural monotonicity property (Definition 9, below). The advantage of this approach is that it gives a mechanism that is dominant strategy incentive compatible and, moreover, the mapping from an agent’s

⁷This ordering can be found by calculating the optimal single-agent mechanism for allocation constraint $\hat{x}(q) = 1 - q$.

type to quantile is a randomized function determined by her type space and distribution alone and is unaffected by the environment (the other bidders and the feasibility constraints), i.e., it is context free.

Definition 9. An agent has *monotone ex ante optimal pricings* if, given her type, the probability she wins in the \hat{q} ex ante optimal pricing is monotone non-decreasing in \hat{q} .

Suppose that the \hat{q} ex ante optimal pricing for an agent each consists 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 ex ante optimal pricings assumption would require that the sets of types served for each \hat{q} be nested. There is a simple deterministic mapping from types to quantiles in this case: set the quantile of a type to be the minimum \hat{q} such that the \hat{q} ex ante optimal pricing serves the type. Below, we generalize this selection procedure to the case of partial lotteries (where types may be probabilistically served).

Denote the allocation and outcome rules (as functions of an agent's type) of the \hat{q} ex ante optimal pricing by $\tilde{x}^{\hat{q}}$ and $\tilde{w}^{\hat{q}}$, respectively. Fix the type of the agent as t and consider the function $G_t(\hat{q}) = \tilde{x}^{\hat{q}}(t)$ which, by the monotonicity condition, can be interpreted as a cumulative distribution function. Note that $\tilde{x}^{\hat{q}}$ has ex ante probability of service $\mathbf{E}_t[\tilde{x}^{\hat{q}}(t)] = \hat{q}$. Hence if t is drawn from the type distribution and then q is drawn from G_t then q is uniformly distributed on $[0, 1]$.

Lemma 11. If $t \sim F$ and $q \sim G_t$ then q is $U[0, 1]$.

Definition 10. The *marginal revenue mechanism* for agents with monotone step mechanisms works as follows.

- 1) Map reported types $\mathbf{t} = (t_1, \dots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \dots, q_n)$ by sampling q_i from the distribution with cumulative distribution function $G_{t_i}(q) = \tilde{x}_i^{\hat{q}}(t_i)$.
- 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 to be in the marginal revenue maximizing feasible set (breaking ties consistently).
- 4) For each agent i , offer the q_i^* ex ante optimal pricing 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^* ex ante optimal pricing, 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 $\tilde{x}_i^{q_i^*}(t_i) = \Pr_{q_i}[q_i \leq q_i^*]$ by our choice of q_i . When we offer agent i the q_i^* ex ante optimal pricing 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 non-service. While it may not be feasible to serve all agents who receive non-trivial partial lottery, this method coordinates across the partial lotteries which agents to serve to maintain the right distribution on agent outcomes and ensure feasibility.

Theorem 12. *The marginal revenue mechanism for agents with monotone step mechanisms is ex post feasible, dominant strategy incentive compatible, and implements marginal revenue maximization.*

In the full version of the paper we show that a single-dimensional agent with a publicly known budget (and some standard assumptions on the value distribution) satisfies the monotonicity condition, and for this kind of preference and a single-unit environment we give an interpretation of the marginal revenue mechanism, above.

V. APPROXIMATION

In previous sections, we have shown that for any collection of agents the marginal revenue mechanism can be implemented and for revenue-linear agents that it is optimal. In this section, we show that the marginal revenue mechanism gives a good approximation to the optimal revenue quite generally.

A. Agent-based Approximation

Proposition 13. *If for any agent i and allocation constraint \hat{x}_i , the marginal revenue $\text{MR}(\hat{x}_i)$ is at least an α fraction of the optimal revenue $\text{Rev}(x_i)$, then the marginal revenue mechanism in the multi-agent setting is an α -approximation to the optimal mechanism.*

For many single-agent problems of interest neither the optimal revenue nor the marginal revenue are easy to characterize; therefore, to instantiate Proposition 13 we look for upper bounds on the optimal revenue, lower bounds on the marginal revenue, and their ratio in worst case over allocation constraints. This endeavor is simplified by the following immediate consequence of linearity.

Proposition 14. *Given an upper bound on the optimal revenue and a lower bound on the marginal revenue that are both linear in the allocation constraint, i.e., $\hat{x} = \hat{x}^A + \hat{x}^B$ implies that $B(\hat{x}) = B(\hat{x}^A) + B(\hat{x}^B)$, if the lower bound is an α fraction of the upper bound for any ex ante constraint \hat{q} , then the marginal revenue is a α -approximation to the optimal revenue for all allocation constraints.*

If, in addition, our lower bound for ex ante constraint \hat{q} comes from an approximation algorithm for the \hat{q} lottery pricing problem, i.e., it is the revenue of a lottery pricing that serves with probability \hat{q} , then we can define the marginal revenue mechanism from this approximation algorithm for the lottery pricing problem. Such an approach might be

desirable if the approximations are better behaved than the optimal ex ante lottery pricings, e.g., if they are easy to compute, respect an ordering on types, or satisfy the monotonicity condition of Definition 9 (all of which make implementation of the marginal revenue mechanism easier).

Multi-dimensional Agents with Unit-demand Product Distributions: Consider lottery pricing for an agent who desires one of several items (i.e., unit-demand) with values for the items drawn from a product distribution. The unconstrained version of this lottery pricing problem has seen recent attention in the literature. Chawla et al. [10] showed that the optimal lottery pricing is upper bounded by twice the optimal auction revenue for the *representative environment*. The representative environment is one where the unit-demand agent is replaced by single-dimensional representative agents bidding for a single item. As we saw in Section II, this single-dimensional auction problem is solved by optimizing marginal revenue and its expected revenue is equal to its cumulative marginal revenue. Chawla et al. [11] showed that, for the original unit-demand environment, a simple item pricing based on equalizing the marginal revenues of the price posted is a two approximation to the optimal representative revenue. Thus, this item pricing is a four approximation to the optimal lottery pricing. We extend these results to general constraints as follows.

Theorem 15. *For any unit-demand agent with values drawn from a product distribution and any allocation constraint, twice the optimal revenue for the representative environment upper bounds the revenue of the optimal lottery for the unit-demand agent.*

Theorem 16. *For any unit-demand agent with values drawn from a product distribution and any downward-closed ex ante allocation constraint,⁸ item pricing to equalize the marginal revenues of the prices posted gives a two approximation to the optimal auction revenue for the representative environment.*

Linearity of the revenue in the representative environment (by Theorem 3 and Theorem 5) then implies by Proposition 14 that maximizing marginal revenue using the item pricings of Theorem 16 as approximate solutions to the ex ante optimal pricing problem gives a four approximation to the optimal revenue. Of course, marginal revenue maximization with the optimal ex ante pricing can be no worse.

Corollary 17. *In downward-closed service-constrained environments with unit-demand agents, the marginal revenue mechanism is a four approximation to the optimal mechanism (and so is the marginal revenue mechanism defined from the ex ante pricing approximations).*

⁸A downward closed ex ante constraint \hat{q} allows the service probability to be at most \hat{q} .

B. Feasibility-based Approximation

Feasibility constraints imply approximation bounds for the marginal revenue mechanism. Below, Theorem 18 gives results for matroid settings follow from the correlation gap approach of Yan [12], and results for downward-closed settings which require novel reasoning about the marginal revenue mechanism. Proofs of these results are in the full version of the paper.

Theorem 18. *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. In downward-closed environments for n quasi-linear agents, the optimal marginal revenue is a $4 \log n$ -approximation to the optimal revenue.*

REFERENCES

- [1] J. Bulow and J. Roberts, “The simple economics of optimal auctions,” *Journal of Political Economy*, vol. 97, no. 5, pp. 1060–1090, October 1989.
- [2] R. Myerson, “Optimal auction design,” *Mathematics of Operations Research*, vol. 6, no. 1, pp. pp. 58–73, 1981.
- [3] J. Thanassoulis, “Haggling over substitutes,” *J. Economic Theory*, vol. 117, no. 2, pp. 217–245, 2004.
- [4] R. Wilson, *Nonlinear Pricing*. Oxford University Press, 1997.
- [5] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian, “Bayesian optimal auctions via multi-to single-agent reduction,” in *ACM Conference on Electronic Commerce*, 2012.
- [6] Y. Cai, C. Daskalakis, and M. Weinberg, “An algorithmic characterization of multi-dimensional mechanisms,” in *Symposium on Theory of Computation*, 2012.
- [7] —, “Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization,” in *Foundations of Computer Science*, 2012.
- [8] —, “Reducing revenue to welfare maximization: Approximation algorithms and other generalizations,” in *SODA*, 2013.
- [9] G. Hardy, J. Littlewood, and G. Pólya, “Some simple inequalities satisfied by convex functions,” *Messenger of Math*, vol. 58, pp. 145–152, 1929.
- [10] S. Chawla, D. Malec, and B. Sivan, “The power of randomness in bayesian optimal mechanism design,” in *EC*, 2010, pp. 149–158.
- [11] S. Chawla, J. Hartline, D. Malec, and B. Sivan, “Multi-parameter mechanism design and sequential posted pricing,” in *42nd ACM Symposium on Theory of Computing*, 2010, pp. 311–320.
- [12] Q. Yan, “Mechanism design via correlation gap,” in *Symposium on Discrete Algorithms*, 2011.