The Simple Economics of Approximately Optimal Auctions*

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

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

Nima Haghpanah
Northwestern University
EECS Department
nima.haghpanah@gmail.com

Jason Hartline
Northwestern University
EECS Department
hartline@northwestern.edu

May 21, 2014

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 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 implement the marginal revenue maximization. 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 approximately extends many 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 the 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.

*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.

1 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 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.

Because it is simple and intuitive, the Myerson-Bulow-Roberts approach provides the basis for most 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 decomposes a multi-agent mechanism design problem as a composition of simple single-agent mechanism design problems, specifically, 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 more generally. 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 the marginal revenue approach is structurally similar to the classical approach, many results from classical auction theory approximately and automatically extend to general preferences.

A central result in classical auction theory is derived from an interpretation the Myerson-Bulow-Roberts mechanism (i.e., for maximizing marginal revenue) in the special case of symmetric agents. Our generalization admits a similar interpretation. 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 in 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

[1] The red-or-blue car example is slightly unnatural as a forward auction (i.e., when 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 Appendix E.

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 a mild assumption).

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 assumption 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 a multi-dimensional preference to a single-dimensional preference by always selling the winning agent her favorite color is very natural and practical; however, it is not generally optimal. For example, when the agent's values for each color is distributed uniformly on $[5, 6]$, the analysis of Thanassoulis (2004) shows that the optimal auction does not sell the agent her favorite color subject to a reserve (in fact, it is not even deterministic). However, many relevant distributions, including the uniform distribution on $[5, 6]$, satisfy the mild assumption sufficient for the four approximation, above.

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 the 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 red-or-blue car example, or the grade or quality of a service). 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 private type drawn from a known distribution. Via the taxation principle (see e.g., Wilson, 1997) the outcomes of a mechanism, for all possible reports the agent might make in the mechanism, 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 type, a lottery pricing induces a probability with which the agent receives an outcome that corresponds to service, and an expected payment, i.e., revenue.

[2] In this example we give the reserve price for $m = 2$ colors; however, with the appropriate reserve price these results hold for any number of colors.

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. As a function of the ex ante service probability this optimal revenue induces a *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 (mechanisms and hence) lottery pricings is closed under convex combination, the revenue curves are always concave. The marginal revenue curve is given by 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 break the allocation to each market into tiny pieces ordered by willingness to pay and attribute to each piece the change in revenue from adding that piece (i.e., the marginal revenue), then the total revenue of an allocation is the sum of the marginal revenues of each piece. A simple algorithm for optimizing this *surplus of marginal revenue* is to repeatedly allocate a tiny amount 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 results in a final allocation where the markets marginal revenues are roughly equal 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.

Definition 1. The family of *marginal revenue mechanisms* take the following form:

1. Map each agent's private 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 the revenue curve at her quantile.
3. Select for service the set of agents that maximize the *surplus of marginal revenue*, i.e., the total marginal revenue of agents served, 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. For the special case of the selling-a-car example; where the agents' values are independently, identically, and uniformly drawn from the $[0, 1]$ interval; the marginal revenue mechanism is instantiated as follows.

For an agent with value drawn uniformly from the $[0, 1]$ interval, the optimal lottery pricing for ex ante service probability $\hat{q}$ is to post a take-it-or-leave-it price of $\hat{v} = 1 - \hat{q}$. The revenue from such pricing is the price times the probability that it is accepted. Therefore, the revenue curve is $R(\hat{q}) = (1 - \hat{q}) \times \hat{q}$ and the marginal revenue curve is its derivative $R'(\hat{q}) = 1 - 2\hat{q}$.

The optimal lottery for ex ante probability $\hat{q}$ serves the agent if her value v is on interval $[1 - \hat{q}, 1]$. This is the strongest $\hat{q}$ measure of the values from the distribution. This motivates, in Step 1, mapping value v to quantile $q = 1 - v$. Composing this mapping from value to quantile with the above mapping from quantile to marginal revenue gives a mapping from value v to marginal revenue as $2v - 1$.

For a single-item auction, in Step 3 the surplus of marginal revenue is maximized by serving nobody if all have negative marginal revenues and, otherwise, by serving the agent with the highest marginal revenue. As the agents are symmetric and marginal revenue is monotone in value, equivalently, the highest-valued agent wins as long as her value is at least 1/2 (the value for which marginal revenue is zero, i.e., solving $2v - 1 = 0$).

The appropriate calculation of payments for Step 4 is the following. All losers have payments equal to zero. The payment of a winner is the minimum value she could declare and still win in Step 3, i.e., it is the maximum of the the second highest agent value and 1/2.

This auction, as claimed in the earlier discussion of the selling-a-car example, is the second-price auction with reserve 1/2. Moreover, the mapping from value to marginal revenue is identical to the virtual values in the derivation of Myerson (1981).

Results. This paper generalizes the marginal-revenue approach for agents with single-dimensional linear preferences (Bulow and Roberts, 1989) to general preferences. Our main algorithmic contribution is to generalize Steps 1 and 4 thereby allowing the construction of service constrained multi-agent mechanisms from single-agent ex ante lottery pricings. There are a number of challenges in this endeavor. First, revenue equivalance does not hold for general preferences (which is used in the proof of optimality for single-dimensional preferences).[3] Second, there is not a natural ordering on types for general preferences (making it difficult to map types to quantiles). 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.

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-dimensinal agents. Our derivation exposes a previously unobserved property of single-dimensional linear preferences which we refer to as *revenue linearity.* Generally, i.e., beyond single-dimensional linear preferences, the optimality of the marginal revenue mechanism is implied by revenue linearity. Moreover, if the single-agent problems are α-approximately revenue linear, then the marginal revenue mechanism is an α approximation to the optimal mechanism.

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

One of the main benefits of considering the marginal revenue mechanism for approximately optimal mechanism design is that, as its structure is similar to optimal mechanisms for single-dimensional environments, many results from the extensive single-dimensional mechanism design literature can be easily generalized. The following are some of the most important consequences.

Algorithmic mechanism design. When weighted optimization is hard we can replace an exact algorithm for weighted maximization (Step 3 of Definition 1) with any approximation

[3]For single-dimensional linear preference agents, the revenue equivalence theorem states that any two auctions with the same allocation in expectation have the same expected revenue.

algorithm using either of the single-dimensional black-box reductions of Hartline and Lucier (2010) and Hartline et al. (2011).

Sequential posted pricing. Sequential posted pricing mechanisms of Chawla et al. (2010a) and Yan (2011) that are approximately optimal for single-dimensional agents are approximately optimal for general agents (in the same service constrained environment) and the same approximation factor is guaranteed. Moreover, these sequential posted pricing bounds give another bound on the approximation factor of the marginal revenue mechanism. The marginal revenue mechanism is in fact optimal within a class of mechanisms that contains the sequential posted pricing mechanisms; therefore, its approximation factor is no worse. As an example, for the single-item service constraint, a sequential posted pricing bound implies an $e/(e-1)$-approximation regardless of approximate revenue linearity of the single-agent problems.

Simple versus optimal. While our marginal revenue mechanism is already generally much simpler than the optimal mechanism, we can get even simpler approximation mechanisms by applying methods developed for single-dimensional preferences to prove that simple mechanisms approximate the marginal revenue mechanism. In particular, in single-dimensional environments maximizing marginal revenue is more complex than simple reserve-price-based mechanisms, i.e., mechanisms that maximize welfare subject to a reserve price. Nonetheless, Hartline and Roughgarden (2009) show that reserve-price-based mechanisms are often approximately optimal. When uniform pricing is approximately optimal, e.g., in generalizations of the red-or-blue car example, these mechanisms extend to general preferences.

Single-sample mechanisms. Approaches above have been for Bayesian optimal mechanism design where the designer optimizes a mechanism given a distribution of preferences. Dhangwatnotai et al. (2010) relax the assumption that the distribution is known and show that a mechanism based on drawing a single sample from the distribution gives a good approximation to the Bayesian optimal mechanism. Again, the single-sample framework extends to general preferences for which uniform pricing is approximately optimal.

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. (2012) and Cai et al. (2012a,b, 2013) gave polynomial-time mechanisms for large important classes of Bayesian mechanism deign problems; the former considers general preferences in service constrained settings (as does this paper) and the latter considers multi-dimensional additive preferences. The two main conclusions of these works is 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 the 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 in 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.) There are a number of important distinctions between our work and these algorithmic results. First, the weights in our derivation have a natural economic interpretation as marginal revenues. 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 multi-agent optimization problem. Third, in most cases, the weights in our derivation depend only on the single-agent problem and not on the multi-agent 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. As an example that breaks this barrier, our approach gives approximately optimal mechanisms for multi-dimensional unit-demand agents with values from product distributions; these mechanisms are easy to compute with a computational complexity that scales linearly with the dimensionality of the type space (i.e., logarithmicly in the size of the type space).

Organization. In Section 2 we review the Myerson-Bulow-Roberts single-dimensional linear agent model 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 3 we formalize the 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 4 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 5 we show that approximate revenue linearity, properly defined, implies approximate optimality. In Section 6 we suggest numerous extensions of results in the single-dimensional mechanism design literature to general preferences that are direct consequences of the marginal revenue mechanism framework.

2 Warmup: 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. In this proof we will introduce many concepts that make our generalization possible (which were not present in previous proofs). 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 linear 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 decomposition implies that the optimal revenue can be expressed in terms of the *surplus of marginal revenue*: the sum of derivatives of the revenue curves of agents served evaluated at points corresponding to the agents' types. The marginal revenue mechanism optimizes this latter term pointwise and, therefore, also in expectation. In the interest of brevity we will keep the discussion informal; many of the proofs in this section are subsumed by generalizations in Section 3 which are given formally.

Model. A single-dimensional linear agent has a private type (a.k.a. valuation) $v \in \mathbb{R}_+$ drawn at random from a probability distribution with cumulative distribution function F and density function f. Let (x, p) denote the outcome of receiving a good or service with probability x and making expected payment p. For such an outcome, an agent with type 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 private type by its strength relative to the distribution. 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 higher-valued types, i.e., an agent with type v has quantile $q = 1-F(v) = V^{-1}(v)$. Importantly, for

v drawn at random 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 her respective inverse demand curve (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 a 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 type, she may choose different outcomes. We may as well index the outcomes in the menu by the quantile corresponding to the type for which the agent would select the outcome, i.e., the 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.,

$$V(q)x(q) - p(q) \geq V(q)x(q') - p(q'), \qquad \forall q, q' \in [0, 1]. \tag{IC}$$

$$V(q)x(q) - p(q) \geq 0, \qquad \forall q \in [0, 1]. \tag{IR}$$

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

The Myerson (1981) 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* $p(\cdot)$ is given precisely as a function of $x(\cdot)$.[4] An important consequence of the latter part of this characterization is *revenue equivalence.* We will make strong usage of both monotonicity and revenue equivalence below, though the specific form of the payment rule will not be important.

Constrained Lottery Pricings. Given 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 $\hat{q}$, find the lottery pricing that serves the agent with probability $\hat{q}$ and maximizes revenue.

Definition 2. The *revenue curve* $R(\hat{q})$ is defined for all $\hat{q} \in [0, 1]$ as the revenue of the *ex ante optimal lottery pricing* with allocation probability $\hat{q}$.

To show that optimal mechanisms are convex combinations of ex ante optimal lottery pricings, we consider a more general lottery pricing problem. Notice that the ex ante lottery pricing problem gives an (equality) constraint on the total probability that the agent is served in expectation over

[4]Notice that quantiles are ordered in the opposite direction as types. Higher-valued types have low quantile and lower-valued types have high quantile. Thus, the allocation rule should be non-increasing in quantile.

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 to 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 $\hat{X}(1)$ at $\hat{q} = 1$.

To see why this constrained lottery pricing problem is the right one to consider, notice the following. First, because any allocation rule is monotone, meaning stronger quantiles receive no lower probability of service than weaker quantiles, the only set of measure $\hat{q}$ for which the constraint $\hat{X}(\hat{q})$ on service probability may be tight is the strongest $\hat{q}$ measure of quantiles, i.e., $[0, \hat{q}]$. For allocation rule $x(\cdot)$ the probability of service to the strongest $\hat{q}$ measure of agents is exactly $X(\hat{q}) = \int_0^{\hat{q}} x(q)\,\mathrm{d}q$. 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, among allocation rules that satisfy X as a constraint, x has the highest probability on stronger (i.e., lower) quantiles. Conversely, the allocation constraint $\hat{X}$ (with corresponding $\hat{x}(q) = \frac{\mathrm{d}}{\mathrm{d}q}\hat{X}(q)$) is met by any allocation rule x that relatively has allocation probability shifted from stronger quantiles to weaker quantiles. Specifically, *$\hat{x}$ majorizes x*.

Definition 3. We say an allocation rule x is *weaker* than another allocation rule $\hat{x}$ if $\hat{x}$ majorizes x in the sense discussed above. $\mathrm{Rev}[\hat{x}]$ is defined for all allocation constraints $\hat{x}$ as the revenue of the *interim optimal lottery pricing* with allocation rule weaker than $\hat{x}$.[5]

Recall the ex ante lottery pricing problem of optimally serving 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 $\hat{q}\,V(\hat{q})$ which is a lower bound on $R(\hat{q})$. 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 among all allocation rules that satisfy the ex ante allocation constraint $\hat{q}$. Of course, the revenue it generates $\hat{q}\,V(\hat{q})$ may not be a concave function of $\hat{q}$ whereas it must be that the revenue curve $R(\cdot)$ is concave. It can be shown, in fact, that $R(\cdot)$ is exactly the concave hull of $\hat{q}\,V(\hat{q})$ and the optimal lottery for any $\hat{q}$ is given by a posted pricing or, if $R(\cdot)$ is linear at $\hat{q}$, the convex combination of two posted pricings (corresponding to the end points of the interval containing $\hat{q}$ on which $R(\cdot)$ is linear). The allocation rule of this convex combination is a convex combination of two reverse step functions and, in the sense described above, relative to posting price $V(\hat{q})$ has service probability shifted from stronger quantiles to weaker quantiles. This specific form (which is not obvious) is not important for our rederivation of the optimal mechanism; what is important is the following proposition (which is obvious from the above discussion).

Proposition 1. *For single-dimensional linear agents, the ex ante optimal lottery pricings have weaker allocation rules than posted prices and higher revenue.*

[5]From the agent's perspective in a multi-agent mechanism, the allocation constraint $\hat{x}$ is applied at the interim stage of the mechanism, i.e., when the agent knows her own type but considers the types of other agents to be drawn from their respective distributions.

Revenue Linearity. We are now ready to give the new derivation of the marginal revenue mechanism and its optimality.

Definition 4. An agent is *revenue linear* if $\text{Rev}[\cdot]$ is a linear functional, i.e., if the optimal revenue for allocation constraints $\hat{x} = \hat{x}^A + \hat{x}^B$ is $\text{Rev}[\hat{x}] = \text{Rev}[\hat{x}^A] + \text{Rev}[\hat{x}^B]$.

We can derive a lower bound on the optimal revenue for any allocation constraint $\hat{x}$ as follows. 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 from by drawing $\hat{q}$ at random from the distribution $G^{\hat{x}}$ with density $-\hat{x}'(q) = -\frac{\mathrm{d}}{\mathrm{d}q}\hat{x}(q)$ and then posting price $V(\hat{q})$ (with allocation rule $\hat{x}^{\hat{q}}$). The allocation rule of the convex combination is exactly $\hat{x}$; its expected revenue is a lower bound on $\text{Rev}[\hat{x}]$.

We can derive a better lower bound by, for ex ante constraint $\hat{q} \sim G^{\hat{x}}$, offering the ex ante optimal lottery pricing (instead of posting price $V(\hat{q})$). As the allocation rule for each of these lottery pricings 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, x is feasible for $\hat{x}$ and its revenue gives a lower bound on $\text{Rev}[\hat{x}]$. Formally, with $q \sim U[0,1]$,

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

The second equality follows from integration by parts and the third equality from $R(1) = R(0) = 0$. (Minor assumption: if the agent is always served or never served then no revenue is obtained.) This construction motivates the following definition.

Definition 5. The *marginal revenue* for an agent with quantile q is $R'(q) = \frac{\mathrm{d}}{\mathrm{d}q}R(q)$; 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. *For a revenue-linear agent, the optimal revenue for an allocation constraint is equal to its marginal revenue, i.e., for all $\hat{x}$, $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$.*

The revenue linearity of single-dimensional linear agents is a simple consequence of revenue equivalence (Myerson, 1981) 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).

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

Proof. As we have seen above, the marginal revenue of an allocation constraint is a lower bound on its optimal revenue. To show revenue linearity, then, it suffices to upper bound the optimal revenue by the marginal revenue.

For any allocation rule x (or constraint) marginal revenue can be written as

$$\begin{aligned} \mathrm{MR}[x] &= \mathbf{E}\left[-x'(q)\,R(q)\right] &(1)\\ &= \mathbf{E}\left[R'(q)\,x(q)\right], \text{ and} &(2)\\ &= R'(1)X(1) + \mathbf{E}\left[-R''(q)\,X(q)\right]. &(3)\end{aligned}$$

We already saw the derivation of equation (2) from (1), which follows from integration by parts and $R(0) = R(1) = 0$. Equation (3) follows from integrating by parts again and $X(0) = 0$ (by definition). From equation (1), it is apparent that higher revenue curves give higher revenue (as "$-x'(\cdot)$" is non-negative for monotone allocation rule $x(\cdot)$). From equation (3), it is apparent that higher allocation rules, in the sense of majorization, give higher revenue (as "$-R''(\cdot)$" is non-negative for concave revenue curve $R(\cdot)$ and majorization requires equality of $X(1)$).

Let $P(\hat{q})$ denote the expected revenue from posting price $V(\hat{q})$, i.e., $P(\hat{q}) = \hat{q}\,V(\hat{q})$. Suppose we optimize for $\hat{x}$ and get some (possibly less restrictive) allocation rule x, then optimizing for x as a constraint gives the same revenue,

$$\mathrm{Rev}[\hat{x}] = \mathrm{Rev}[x].$$

By revenue equivalence, the revenue of any allocation rule is given by its price-posting revenue curve $P(\cdot)$. Therefore,

$$\mathrm{Rev}[x] = \mathbf{E}\left[-x'(q)\,P(q)\right].$$

As $P(q) \leq R(q)$ for all q, equation (1) implies that the marginal revenue from $P(\cdot)$ is at most that of $R(\cdot)$ for allocation rule $x(\cdot)$:

$$\mathbf{E}\left[-x'(q)\,P(q)\right] \leq \mathbf{E}\left[-x'(q)\,R(q)\right] = \mathrm{MR}[x].$$

As x is majorized by $\hat{x}$, equation (3) implies that the marginal revenue of $x(\cdot)$ is at most that of $\hat{x}(\cdot)$ for revenue curve $R(\cdot)$:

$$\mathrm{MR}[x] = R'(1)\,X(1) + \mathbf{E}\left[-R''(q)\,X(q)\right] \leq R'(1)\,\hat{X}(1) + \mathbf{E}\left[-R''(q)\,\hat{X}(q)\right] = \mathrm{MR}[\hat{x}]. \qquad \square$$

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

Multi-agent Mechanisms. The conclusion of the preceding discussion is that the optimal revenue for any allocation constraint is equal to its marginal revenue.

Definition 6. Any mechanism and distribution over types induces a profile $\boldsymbol{x} = (x_1, \ldots, x_n)$ of interim allocation rules. The *surplus of marginal revenue* is the sum of the marginal revenues of interim allocation rules of each agent $\sum_i \mathrm{MR}[x_i]$.

Multi-agent mechanism design problems reduce to single-agent lottery pricing problems as follows. The following argument is the standard in auction theory. For an agent in the optimal mechanism, her contribution to the revenue is equal to the marginal revenue of her allocation rule (Corollary 4). We thus look for the mechanism that optimizes the surplus of marginal revenue.

Consider relaxing the incentive constraints (namely: monotonicity of the allocation rule) and optimizing marginal revenue pointwise. Specifically, when the agent quantiles are $\boldsymbol{q} = (q_1, \ldots, q_n)$ select the allocation $\boldsymbol{x} = (x_1, \ldots, x_n)$ to maximize the surplus of marginal revenue $\sum_i R_i'(q_i)\, x_i$ subject to feasibility of $\boldsymbol{x}$ (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 previously relaxed incentive constraints are not violated. Notice that since revenue curves are concave, the marginal revenues are monotone non-increasing in quantile, 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 $V_i^{-1}(\cdot)$ 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 and 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. *The marginal revenue mechanism is revenue optimal for single-dimensional linear agents.*

Proof. The optimal mechanism induces some profile $\boldsymbol{x}$ of interm allocation rules. By revenue linearity, the expected revenue of this profile of interim allocation rules is equal to its surplus of marginal revenue. The marginal revenue mechanism selects its outcome to optimize surplus of marginal revenue pointwise for the feasibility constraint. Its expected surplus of marginal revenue is, thus, at least that of the optimal mechanism. □

3 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 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 t and deterministic outcome w her utility is $u(t, w)$, and when w is drawn from a distribution her utility is $\mathbf{E}_w[u(t, w)]$.[6] 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.

Example 1 (A unit-demand quasi-linear-utility agent)**.** The preferences of a *unit-demand quasi-linear-utility* agent are as follows. There are m *alternatives* and the agent's type is given by a vector $(v^1, \ldots, v^m)$ representing her value for each alternative. An outcome is of the form $(p, \pi^1, \ldots, \pi^m)$, where p denotes the payment, and each $\pi^j \in \{0, 1\}$ indicates whether the agent gets the alternative j, with $\sum_j \pi^j \leq 1$. The agent's utility at such an outcome is then given by the linear form $\sum_j v^j \pi^j - p$. When randomizing over such outcomes, we relax the π^j's to be in $[0, 1]$, still with $\sum_j \pi^j \leq 1$. Such a distribution with a price p is called a *lottery.*

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

Example 2 (A single-dimensional public-budget agent)**.** The preferences of a *single-dimensional public-budget* agent are as follows. The agent has a publicly known budget B, and her type is given by her private value v for an item being auctioned. An outcome $w = (x, p)$ indicates by $x \in \{0, 1\}$ whether the agent gets the item, and by p the amount of payment she makes. In contrast to the single-dimensional linear-utility agents of Section 2, this agent's utility is $v \cdot x - p$ only if $p \leq B$, and negative infinity otherwise.

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. The agents types are indepently distributed. A *direct revelation* mechanism takes as its input a profile of types $\boldsymbol{t} = (t_1, \ldots, t_n) \in T_1 \times \cdots \times T_n$ and outputs ex post outcome $\tilde{\boldsymbol{w}}(\boldsymbol{t}) \in \Delta(W_1 \times \cdots \times W_n)$. Agent i's *ex post outcome rule* is denoted by $\tilde{w}_i(\boldsymbol{t})$ and, with the other agents' types drawn from the distribution, her *interim outcome rule* $\tilde{w}_i(t_i)$ is distributed as $\tilde{w}_i(t_i, \boldsymbol{t}_{-i})$ with $t_j \sim F_j$ for each $j \neq i$. We say that a mechanism is *Bayesian incentive compatible* if

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

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., $\mathbf{E}_{\boldsymbol{t}}[\sum_i \text{Payment}(\tilde{w}_i(t_i))]$; 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* the outcome w provided to an agent is distinguished as being either a *service* or a *non-service* outcome, 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 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 an 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 all 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.

Ex Ante Lottery Pricings and 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 ex ante optimal pricing for allocation probabilities $\hat{q} \in [0, 1]$. This is the lottery pricing (i.e., collection of outcomes where the agent is permitted to choose her type-dependent favorite) denoted $\tilde{w}^{\hat{q}}(\cdot)$ that optimizes revenue subject to the constraint that $\mathbf{E}_t[\text{Alloc}(\tilde{w}^{\hat{q}}(t))] = \hat{q}$. 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 2.

Allocation rules. The first challenge in generalizing the marginal revenue approach to general preferences is determining the mapping from types to quantiles. This challenge arises as there is no explicit ordering of an agent's type space T by strength. 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 $\text{Alloc}(\tilde{w}(t)) \geq \text{Alloc}(\tilde{w}(t'))$. This definition induces a mapping from the type space to quantile space; moreover, the distribution of quantiles induced by this mapping and the distribution on types is uniform.[7] 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 heights; the resulting monotone non-increasing piecewise constant function from $[0, 1]$ to $[0, 1]$ is the allocation rule. This is generalized for continuous distributions as follows.

Definition 7. 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 $x(\hat{q}) = \sup\{y : \mathbf{Pr}_{t\sim F}[\text{Alloc}(\tilde{w}(t)) \geq y] \leq \hat{q}\}$.

Optimal Lottery Pricing. With the definition of allocation rules for any lottery pricing above, 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 3 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 $\hat{q}$ *ex ante optimal pricing* is denoted $x^{\hat{q}}$.

Revenue Linearity and Marginal Revenue. Revenue linearity and marginal revenue have the same definitions (Definition 4 and Definition 5) 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 optimal pricings 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 that its revenue is equal to the optimal revenue (Theorem 2). We will desribe the marginal-revenue approach for non-revenue-linear agents by analogy to the single-dimensional case.

Definition 8. The *single-dimensional analog* of a service constrained environment for general agents is the environment with single-dimensional linear agents with the same revenue curves. The

[7] Quantiles are uniformly distributed when ties in allocation probability are measure zero; when there is a measureable probability of ties, quantiles can be defined by drawing uniformly from the interval containing the tie.

optimal marginal revenue for a service constrained environment for general agents is the optimal revenue of the single-dimensional analog (which is equal to its surplus of marginal revenue).

Our approach to multi-agent mechanism design via the single-dimensional analog is to look at the profile of interim allocation rules induced by maximization of surplus of marginal revenue and then to construct a mechanism for general agents that looks to each agent like the convex combination of the ex ante optimal pricings for her allocation rule. For revenue curves $R_1, \ldots, R_n$, draw quantiles $\boldsymbol{q} = (q_1, \ldots, q_n)$ uniformly from $[0,1]^n$, serve to maximize surplus of marginal revenue pointwise as $\sum_i R_i'(q_i)\, x_i$ for feasible $\boldsymbol{x} = (x_1, \ldots, x_n)$. We can interpret the allocation rules induced by this process as allocation constraints for the general environment and denote them by $\hat{\boldsymbol{x}}^{MR} = (\hat{x}_1^{MR}, \ldots, \hat{x}_n^{MR})$. As for single-dimensional linear agents (see Section 2), one way to serve an agent subject to allocation constraint $\hat{x}$ is to draw a quantile $\hat{q}$ from the distribution $G^{\hat{x}}$ with density $-\frac{\mathrm{d}}{\mathrm{d}q}\,\hat{x}(q)$ and run the ex ante optimal pricing for ex ante constraint $\hat{q}$. This approach suggests attempting to implement the general mechanism with outcome rules that correspond to allocation rules of the single-dimensional analog. Denoting the outcome rule for the $\hat{q}$ ex ante optimal pricing for agent i by $\tilde{w}_i^{\hat{q}}(t_i)$. The agent's outcome rule corresponding to constraint $\hat{x}_i^{MR}$ is $\tilde{w}_i^{MR}(t_i) = \int_0^1 \tilde{w}^{\hat{q}}(t_i)\,(-\mathrm{d}\hat{x}_i^{MR}(q))$. There may be multiple ways to implement this profile of outcome rules ex post; however, the direct approach employed for single-dimensional linear agents in Section 2 does not always generalize.

Definition 9. The *marginal revenue outcome rule* of an allocation rule x is $\tilde{w}(t) = \int_0^1 \tilde{w}^{\hat{q}}(t)\,(-\mathrm{d}x(q))$. A *marginal revenue mechanism* is one with interim outcome rules equal to the marginal revenue outcome rules corresponding to the optimal marginal revenue.

Implementation with Revenue Linearity. We show now that the marginal revenue mechanism generalizes exactly for general preferences that satisfy revenue linearity. Moreover, we show that 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 is deterministic and *context free*[8] in that it does not depend on the feasibility constraint or other agents in the mechanism. 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 10. 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

[8]Note 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. (2012) and Cai et al. (2012a,b).

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.[9]

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.*

The theorem is proved by the following two lemmas which characterize the structure of optimal lottery pricings.

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

Proof. Since x is the optimal allocation rule subject to $\hat{x}$, we have $\mathrm{Rev}[x] = \mathrm{Rev}[\hat{x}]$. Linearity implies that

$$\mathrm{MR}[\hat{x}] = \int_0^1 x(q)R'(q)\,\mathrm{d}q = \int_0^1 \hat{x}(q)R'(q)\,\mathrm{d}q = \mathrm{MR}[x].$$

Integrating by parts, we have

$$\Big[X(q)R'(q)\Big]_0^1 - \int_0^1 X(q)R''(q)\,\mathrm{d}q = \Big[\hat{X}(q)R'(q)\Big]_0^1 - \int_0^1 \hat{X}(q)R''(q)\,\mathrm{d}q. \tag{4}$$

Note that $\hat{x}$ and x have the same ex ante probability of allocation $\hat{X}(1) = X(1)$; also by definition $X(0) = \hat{X}(0) = 0$. Combining these observations with (4) we have

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

and therefore,

$$\int_0^1 [X(q) - \hat{X}(q)]R''(q)\,\mathrm{d}q = 0. \tag{5}$$

Notice that for any q, $X(q) - \hat{X}(q)$ and $R''(q)$ are non-positive (by domination and concavity, respectively) so their product is non-negative. Therefore, (5) can be satisfied only if $[X(q) - \hat{X}(q)]R''(q) = 0$ for all q. This implies that if $R''(q) < 0$, then we must have $X(q) = \hat{X}(q)$, which completes the proof. $\square$

Lemma 7 in particular implies that for $\hat{q}$ with $R''(\hat{q}) \neq 0$ the $\hat{q}$ ex ante optimal pricing (i.e., with allocation constraint given by the reverse step function $\hat{x}^{\hat{q}}$) has allocation rule $x^{\hat{q}} = \hat{x}^{\hat{q}}$. I.e., the $\hat{q}$ ex ante optimal pricing has only full lotteries (all types are served with either probability one or zero).

For any such $\hat{q}$, define $T_{\hat{q}}$ to be the set of types allocated (with full lotteries) in the optimal allocation subject to $\hat{x}^{\hat{q}}$. The following lemma shows that these sets are nested.

[9] By greedy by the given ordering, we mean process each equivalance class in order and serve the corresponding types with as much probability as possible subject to the allocation constraint. If all equivalance classes are measure zero, then the resulting allocation rule is equal to the allocation constraint.

Lemma 8. *For a revenue-linear single-agent problem, for any $\hat{q}_1 > \hat{q}_2$ and $R''(\hat{q}_1), R''(\hat{q}_2) \neq 0$, we must have $T_{\hat{q}_1} \supseteq T_{\hat{q}_2}$.*

Proof. Assume for contradiction that $T_{\hat{q}_2} \backslash T_{\hat{q}_1} \neq \emptyset$. Let $\alpha = F(T_{\hat{q}_2} \backslash T_{\hat{q}_1}) > 0$. Consider the following allocation constraint

$$\hat{x}(q) = \begin{cases} 1 & q \leq \hat{q}_2 \\ 1/2 & \hat{q}_2 < q \leq \hat{q}_1 \\ 0 & \hat{q}_1 < q. \end{cases}$$

By revenue linearity, the revenue of the optimal auction subject to $\hat{x}$ is $[R(\hat{q}_1) + R(\hat{q}_2)]/2$. Notice that the mechanism that runs $R(\hat{q}_1)$ and $R(\hat{q}_2)$ 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 $\hat{x}$, and achieves the optimal revenue. Yet, we have

$$\hat{X}(\hat{q}_1) = \int_{q=0}^{\hat{q}_1} \hat{x}(q) \, \mathrm{d}q > \int_{q=0}^{\hat{q}_1} x(q) \, \mathrm{d}q = X(\hat{q}_1).$$

This contradicts Lemma 7. □

Proof of Theorem 6. By Lemma 8, all $\hat{q}$ ex ante optimal pricings order the types by the same equivalance classes. By revenue linearity the optimal lottery pricing for an allocation constraint $\hat{x}$ is a convex combination of the $\hat{q}$ ex ante optimal pricings. Therefore, it allocates greedily to types by the same equivalance classes. □

Given orderability and the fact that (by Lemma 7) the optimal $\hat{q}$ ex ante optimal pricings are full lotteries for $\hat{q}$ for which $R(\hat{q})$ is locally non-linear, the marginal revenue mechanism is easy to define.

Definition 11. The *marginal revenue mechanism for orderable agents* works as follows.

1. Map reported types $\boldsymbol{t} = (t_1, \ldots, t_n)$ of agents to quantiles $\boldsymbol{q} = (q_1, \ldots, q_n)$ via the implied ordering.[10]

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.

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

Proposition 9. *The marginal revenue mechanism deterministically selects a feasible set of agents to serve and is dominant strategy incentive compatible.*

Proof. Because ties are broken consistently, critical values cannot fall in intervals where the revenue curve is locally linear (and the marginal revenue curve is locally constant). Therefore, the lottery pricings offered to each agent are full lotteries; each type is deterministically served or not served. 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 $\hat{q}_i$ is a function only of the other agents' reports; the agent's outcome is determined by the $\hat{q}_i$ ex ante optimal pricing which is incentive compatible for any $\hat{q}_i$. □

Proposition 10. *In service constrained environment with revenue-linear agents, the marginal revenue mechanism obtains the optimal marginal revenue (which equals the optimal revenue).*

Testing Revenue Linearity. Revenue linearity is computationally easy to test. From the concavity of $\mathrm{Rev}[\cdot]$ and equality of revenue and marginal revenue for allocation constraints $\hat{x}^{\hat{q}}$ which are a basis for general allocation constraints, it suffices to check the equality of revenue and marginal revenue, i.e., $\mathrm{Rev}[\hat{x}] = \mathrm{MR}[\hat{x}]$, for any allocation constraint $\hat{x}$ with positive derivative ($\hat{x}$ as a convex combination of $\hat{x}^{\hat{q}}$ has positive density on each $\hat{q}$). For example, $\hat{x}(q) = 1 - q$ is such an allocation constraint. Since the theorem facilitates testing the property, we discretize the quantile space to $Q_N = \{0, \frac{1}{N} \cdots, \frac{N-1}{N}, 1\}$ for an arbitrary integer $N > 0$.

Theorem 11. *Let $\hat{x} : Q_N \to [0, 1]$ be any strictly decreasing function, (e.g., $\hat{x}(\hat{q}) = 1 - \hat{q}$). Then a set of single-agent pricings are revenue linear (or, more precisely,* $\mathrm{Rev}[\cdot]$ *is a linear functional for non-increasing functions mapping Q_N to* $[0, 1]$*), if* $\mathrm{Rev}[\hat{x}] = \mathrm{MR}[\hat{x}]$.

Proof. Consider the $N + 2$ reverse step functions that "steps down" from 1 to 0 at a point in Q_N. Any non-increasing function mapping Q_N to $[0, 1]$ is a convex combination of these base functions, and a strictly decreasing function can be written uniquely as such a convex combination. Therefore $\mathrm{Rev}[\hat{x}] = \mathrm{MR}[\hat{x}]$ amounts to saying that $\mathrm{Rev}[\cdot]$ is linear on one interior point in a simplex, and the theorem states that $\mathrm{Rev}[\cdot]$ is linear on the whole simplex. If we shift $\mathrm{Rev}[\cdot]$ by a linear functional such that it is zero on all the base functions, then this theorem follows from the simple fact that, if a concave function g is 0 on all vertices of a simplex *and* one interior point A, then g is uniformly 0 on the simplex. To see this, suppose on point B in the simplex, $g(B) \neq 0$. By concavity, $g(B) > 0$. A can be written as a convex combination of B and vertices of the simplex with a strictly positive coefficient on B. (E.g., connect B and A with a straight line and extend it to intersect at one facet of the simplex formed by $N - 1$ vertices, then A can be written as a convex combination of B and these $N - 1$ vertices, where the coefficient on B in the decomposition is strictly positive.) But the concavity of g implies $g(A) > 0$, a contradiction. □

Example 3 (A multi-dimensional revenue-linear example)**.** The example of the seller who can paint her car red or blue as she sells it to agents with independent and uniform values for each color is revenue linear (proof given in Appendix D). Therefore, the marginal revenue mechanism is optimal and its simple form can be derived from Definition 11 as follows. For a unit-demand agent with values for m variants of a service (i.e., possible colors of the car) distributed uniformly on $[0, 1]^m$, we show that the ex ante optimal mechanism for constraint $\hat{q}$ is to post a price of $\sqrt[m]{1 - \hat{q}}$ for any service.

Notice that such a price will be accepted with probability $\hat{q}$, and therefore the revenue function is $R(\hat{q}) = \hat{q}\sqrt[m]{1-\hat{q}}$, and the marginal revenue function is $R'(\hat{q}) = (1-\hat{q})^{1/m-1}(1-\hat{q}-\hat{q}/m)$. The quantile of each type is $t = (t^1, \ldots, t^m)$ to be $q = 1 - (\max_i t^i)^m$. Notice that both the mapping and the marginal revenue function are monotone. Therefore serving the agent with the highest marginal revenue (Definition 11) means serving the player with the highest value for any kind of service and charging her the minimum she needs to bid to exceed the second-highest value (subject to the reserve of $\sqrt[m]{\frac{1}{m+1}}$ which is where the marginal revenue becomes zero). Revenue-linearity implies that this mechanism is optimal.

4 Implementation

The marginal revenue mechanism for agents with orderable types (Definition 11) 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 $\hat{q}$ ex ante optimal pricings satisfies a natural monotonicity requirement: that the probability that an agent with a given type is served is monotone in the ex ante constraint $\hat{q}$. Key to this construction is a *randomized* mapping from an agent's types to quantiles that is determined by the agent's type space and distribution alone, and is therefore context free, i.e., unaffected by the presence of other agents and the feasibility constraints. Consequently, (a) the resulting mechanism is dominant strategy incentive compatible but, (b) the set of winners is generally a randomized function of the profile of types. The second approach is brute-force but easily computable and completely general. It results in a Bayesian incentive compatible mechanism. Both these mechanisms will differ from the marginal revenue mechanism for orderable types 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.

We conclude this section by describing a relevant class of agents for which the ex ante optimal pricings satisfy the monotonicity property required by the first approach. The example considers single-dimensional agents with a public budget that constrains their maximum payment.

4.1 Monotone ex ante optimal pricings

We consider agents whose ex ante optimal pricings satisfy the following natural monotonicity property.

Definition 12. 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 one or zero. In this case the monotone ex ante optimal pricings assumption would require that the sets of types served for each $\hat{q}$ 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 $\hat{q}$ 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).

Recall that the $\hat{q}$ ex ante optimal pricing, as a function of the agent's type, has an allocation and outcome rule $\tilde{x}^{\hat{q}}$ 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 above, can be interpreted as a cumulative distribution function. Recall that $\hat{q}$ ex ante optimal pricing has probability of service $\mathbf{E}_t[\tilde{x}^{\hat{q}}(t)] = \hat{q}$. Therefore, if t is drawn from the type distribution and then q drawn from G_t then the distribution of q is uniform on $[0, 1]$.

Lemma 12. *If $t \sim F$ and $q \sim G_t$ then q is $U[0, 1]$.*

Definition 13. The *marginal revenue mechanism for agents with monotone ex ante optimal pricings* works as follows.

1. Map reported types $\boldsymbol{t} = (t_1, \ldots, t_n)$ of agents to quantiles $\boldsymbol{q} = (q_1, \ldots, q_n)$ 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 $\hat{q}_i$ that she could possess and be in the marginal revenue maximizing feasible set (breaking ties consistently).

4. For each agent i, offer the $\hat{q}_i$ ex ante optimal pricing conditioned so that i is served if $q_i \leq \hat{q}_i$ and not served otherwise.

The last step of the marginal revenue mechanism warrants an explanation. In the $\hat{q}_i$ 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 $\tilde{x}_i^{\hat{q}_i}(t_i) = \mathbf{Pr}_{q_i}[q_i \leq \hat{q}_i] = G_{t_i}(\hat{q}_i)$ by our choice of q_i. When we offer agent i the $\hat{q}_i$ ex ante optimal pricing we must draw an outcome from the distribution given by $\tilde{w}_i^{\hat{q}_i}(t_i)$. Some of these outcomes are service outcomes, some of these are non-service outcomes. If $q_i \leq \hat{q}_i$ then we draw an outcome from the distribution $\tilde{w}_i^{\hat{q}_i}(t_i)$ conditioned on service; if $q_i > \hat{q}_i$ then we draw an outcome conditioned on no-service. Notice that, 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.

Proposition 13. *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 $\hat{q}_i$ is a (randomized) function only of the other agents' reports; the agent's outcome is determined by the $\hat{q}_i$ ex ante optimal pricing which is incentive compatible for any $\hat{q}_i$. □

Theorem 14. *The marginal revenue mechanism for agents with monotone ex ante optimal pricings implements marginal revenue maximization (Definition 9).*

Proof. From each agent i's perspective, the other agents' quantiles are distributed independently and uniformly on $[0, 1]$ (Lemma 12). Therefore, this agent faces a distribution over ex ante optimal pricings that is identical to the distribution of "critical quantiles" in the maximization of marginal revenue, i.e., with density $\frac{\mathrm{d}}{\mathrm{d}\hat{q}} G_i^{MR}(\hat{q})$. □

4.2 General ex ante optimal pricings

For general agents for whom the ex ante optimal pricings do not satisfy the monotonicity condition (Definition 12), we give in Appendix A an simple procedure to implement the marginal revenue mechanism (recall Definition 9). This mechanism is given by Definition 25 in Appendix A. The key to the proof of Theorem 15 is a variation of the technique of vector majorization (Hardy et al., 1929).

Theorem 15. *For service constrained environments, there is a simple Bayesian incentive compatible implmentation of the marginal revenue mechanism.*

4.3 Example: single-dimensional agents with public budgets

In this section we exhibit a class of single-agent problems with non-linear utilities that has monotone ex ante optimal pricings (Definition 12). Consider an agent with a single-dimensional value for receiving a good but has a public budget that limits the payment she could make. 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 ex ante optimal pricings.

The following proposition is a consequence of techniques developed by Laffont and Robert (1996) and Pai and Vohra (2008); for completeness we provide a proof in Appendix B whose steps largely resemble the ones in these two references.

Proposition 16. *For regular distribution F with non-decreasing density, budget B, and $\hat{q} \leq 1 - F(B)$, the $\hat{q}$ ex ante optimal pricing offers a single take-it-or-leave-it lottery for price B that serves with probability π, where π is the solution to the equation $\hat{q} = \pi[1 - F(B/\pi)]$. This lottery is bought by the agent when her value is at least B/π which happens with probability $1 - F(B/\pi)$.*

For $\hat{q} > 1 - F(B)$, it is easy to see that the budget does not bind and the $\hat{q}$ ex ante optimal pricing is the same as when there is no budget.

Notice that the allocation rule of the mechanism satisfying Proposition 16 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. If π increases, B/π decreases and more types are served and with a higher probability; thus, the ex ante probability of service is increased. Analogously, if we increase the ex ante probability of service, we enlarge the set of types served and their probability of service. We conclude with the following consequence.

Theorem 17. *An agent with value drawn from a regular distribution with non-increasing density has monotone ex ante optimal pricings.*

Proof. The only case not argued by the text above is when $\hat{q} \geq 1 - F(B)$. In this case, the budget is not binding and the ex ante optimal pricing posts price p that satisfies $\hat{q} = 1 - F(p)$ and serves agents willing to pay this price with probability one. The ex ante optimal pricings are monotone over these quantiles as well. □

Example 4 (Implementation with public budgets)**.** The following procedure implements the marginal revenue mechanism in a single item auction for bidders each with a publicly known budget B and value drawn uniformly from $[0, 1]$. The auction is easier to describe separately for the two cases when $B < 1/2$ and $B \geq 1/2$.

For $B < 1/2$, if no bidder bids above B, the item is not sold; if only one bidder bids above B, she wins the item and makes a payment of B; if at least two bidders bid above B, the winner of the item will be decided among these bidders by a random procedure described shortly. The winner always makes a payment of B.

For $B \geq 1/2$, if no bidder bids above B, a second price auction is run with a reserve price of $1/2$; if one bidder bids above B, she wins the item and makes the same payment as in a second price auction with reserve $1/2$; if at least two bidders bid above B, one of them is decided to be the winner by a random procedure, but all bidders that bid above a randomly chosen threshold also makes a payment of B.

Now we describe the random procedure used to determine the winner in both cases. Note again that only bidders who bid at least B will enter this procedure. Each such bidder i draws a random number r_i uniformly from $[0, 1]$, and her quantile q_i will be $\max\{r_i, B/v_i\} - B$. Whichever bidder i^* having the smallest quantile is declared the winner. The threshold above which other bidders make the payment is $B/(B + q_{i^*})$.

Appendix B gives the derivation showing that this is the instantiation of Definition 13.

Note the role played by the random mapping in this example. When multiple bidders bid above B, the highest bidder is not guaranteed to win the item. Her higher value helps her obtain a lower quantile by posing a smaller B/v_i, but with positive probability she may lose to a lower bidder.

5 Approximation

In previous sections, we have shown that for any service constrained environment the marginal revenue mechanism can be implemented. In Section 3 we have also shown that for revenue linear agents, it obtains the optimal revenue. In this section, we show that, quite generally, the optimal marginal revenue is a good approximation to the optimal revenue.

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 of each agent: if for all allocation constraints $\hat{x}$, the marginal revenue $\mathrm{MR}[\hat{x}]$ is a good approximation to the optimal revenue $\mathrm{Rev}[\hat{x}]$, then the marginal revenue mechanism is a good approximation to the optimal mechanism. The second approach will derive approximation bounds from the feasibility constraint. With no feasibility constraint, marginal revenue maximization is optimal; for matroid environments, it remains a $1-1/e$ approximation; and for general downward-closed environments with n quasi-linear-preference agents, it gives an $O(\log n)$ approximation.

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

5.1 Agent-based Approximation

If, for all allocation constraints, 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, what we refer to as, ex ante pseudo pricings. An ex ante pseudo pricing respects an ex ante constraint but may not be optimal. If for every ex ante service probability $\hat{q}$ the $\hat{q}$ ex ante pseudo pricing approximates the linear upper bound, then for all allocation constraints $\hat{x}$, the marginal revenue $\mathrm{MR}[\hat{x}]$ approximates the optimal revenue $\mathrm{Rev}[\hat{x}]$. Furthermore, these ex ante pseudo pricings can be directly optimized over and the same approximation factor is obtained. Such an approach might be desirable if the ex ante pseudo pricings are better behaved than the (optimal) ex ante optimal pricings, e.g., if they are easy to compute, respect an ordering on types (à la Definition 10), or are monotone (à la Definition 12). This approach is formalized by the following sequence of definitions and propositions.

Proposition 18. *If for any agent i and allocation constraint $\hat{x}_i$, the marginal revenue $\mathrm{MR}[\hat{x}_i]$ is at least an α approximation to the optimal revenue $\mathrm{Rev}[\hat{x}_i]$, then the marginal revenue mechanism in the multi-agent setting is an α approximation to the optimal mechanism.*

Definition 14. An *linear revenue bound*, UB, is a function mapping an allocation constraint to a revenue, which is

1. linear in the allocation constraint, i.e., for all allocation constraints $\hat{x} = \hat{x}^A + \hat{x}^B$, $\mathrm{UB}(\hat{x}) = \mathrm{UB}(\hat{x}^A) + \mathrm{UB}(\hat{x}^B)$; and
2. an upper bound on revenue for all allocation constraints, i.e., $\forall \hat{x},\ \mathrm{UB}(\hat{x}) \geq \mathrm{Rev}[\hat{x}]$, and

Definition 15. A *ex ante pseudo pricing* is one that respects an ex ante service probability constraint but is not necessarily revenue optimal for such a constraint. The revenue of a $\hat{q}$ ex ante pseudo pricing is denoted $\tilde{R}(\hat{q})$; and the *pseudo marginal revenue* for allocation constraint $\hat{x}$ is $\mathrm{PMR}[\hat{x}] = \mathbf{E}[\tilde{R}'(\hat{q})\hat{x}(\hat{q})]$.

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 $\hat{q}$ ex ante pseudo pricing be the revenue-optimal lottery pricing in the closure that serves with ex ante probability $\hat{q}$. This construction is analogous to the ironing method of Myerson (1981).

Proposition 19. *For a given linear revenue bound* UB*, if for all $\hat{q} \in [0, 1]$ the $\hat{q}$ ex ante pseudo pricing α approximates the bound on the $\hat{q}$ ex ante constrained revenue $\mathrm{UB}(\hat{x}^{\hat{q}})$, 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 pseudo marginal revenue. □

Definition 16. The *pseudo marginal revenue mechanism* is the one that maximizes pseudo marginal revenue via any of the approaches of Definition 11, Definition 13, or Definition 25 that applies.

Pseudo ex ante optimal pricing for downward-closed unit-demand agents. We illustrate the methodology proposed above for the example of downward-closed service-constrained environments and unit-demand agents. Recall for unit-demand agents a service outcome is one of m alternatives. An agent's type is described by the vector $(v^1, \ldots, v^m)$, her valuations for each of the m alternatives, and her utility for obtaining alternative j with payment p is simply $v^j - p$. The agent's type is drawn from a product distribution over the distinct alternatives.

Chawla et al. (2010a,b) show, for a single-unit demand agent with values for distinct alternatives from a product distribution and no feasibility constraint, that individually pricing alternatives is a four approximation to optimal lottery pricing. Our approach in this section will be to extend this result to settings with ex ante and interim allocation constraints. Our generalization preserves the approximation bound of four and exposes the approximate linearity condition required by Proposition 18.

Consider the syntactically-related 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 of agent i is drawn independently from F_i. As described earlier (Section 2), 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 where the alternatives correspond to the identity of the winning representative. We refer to this original environment as the *unit-demand environment* and the revenue of the optimal lottery pricing as the *optimal unit-demand revenue.* Chawla et al. (2010a,b) considered quantifying the performance of optimal unit-demand lottery pricings relative to the optimal representative revenue. The approach of Chawla et al. (2010a) is to set individual prices for each alternative in the unit-demand environment so as to mimic the outcome of the optimal auction for the representative environment. As the optimal auction in the representative environment orders representatives by virtual values, a natural approach to pricing the alternatives in the unit-demand environment is to set a uniform virtual price, i.e., the price for each alternative has the same virtual value (with respect to the distribution from which the agent's value for that alternative is drawn).[11] The prices in value space are generally distinct when the agent's value distributions for the alternatives are non-identical. 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 the single-agent problem of serving an agent with independent values for m alternatives subject to an allocation constraint $\hat{x}$. In particular, twice the optimal representative revenue is a linear revenue bound (Definition 14), and for any allocation constraint it upper bounds the optimal (unit-demand) revenue. We define a class of ex ante pseudo pricings where the $\hat{q}$ ex ante pseudo pricing is given by a uniform virtual pricing that sells with probability $\hat{q}$. Since the virtual values are weakly increasing in the representative agents' values, the sets of types served by these ex ante pseudo pricings respect an ordering on types (Definition 10). Therefore, the pseudo marginal revenue mechanism can be implemented via the marginal revenue mechanism for orderable agents (Definition 11). Finally, we show that for all $\hat{q}$ the $\hat{q}$ ex ante pseudo pricing is a four approximation to the linear upper bound given by twice the optimal representative revenue. This result, with Proposition 19, implies that the pseudo marginal revenue mechanism is a

[11] 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 his representatives we will refer to a representative's marginal revenue as his virtual value.

four approximation to the optimal revenue for any allocation constraint. The proof of Theorem 20, below, is a non-trivial but straightforward extension of Chawla et al. (2010a,b) and we include it in Appendix C.

Definition 17. The $\hat{q}$ ex ante pseudo pricing for a unit-demand agent with values for alternatives drawn independently from $F^1, \ldots, F^m$ is given by the pricing that sets a uniform virtual price for the alternatives such that the probability that the agent buys any alternative is equal to $\hat{q}$. (If this class does not have a concave pseudo revenue curve we take its closure with respect to convex combination to make it concave; if this class does not have a monotone non-decreasing pseudo revenue curve $\tilde{R}(\cdot)$ we invoke downward closure to make it monotone.)[12]

Theorem 20. *In downward-closed (service constrained) environments with unit-demand agents, both the pseudo marginal revenue mechanism and the marginal revenue mechanism give 4-approximations to the optimal revenue.*

5.2 Feasibility-based Approximation

We now show that feasibility constraints imply approximation bounds. As a first trivial observation, if there is no feasibility constraint (e.g., for digital good environments) then marginal revenue maximization is optimal. With no feasibility constraint, each agent can be considered separately. For any agent i, suppose the revenue optimal mechanism serves with probability $\hat{q}_i$, by definition the revenue it obtains is equal to that of the $\hat{q}_i$ ex ante optimal pricing. The optimal revenue $\sum_i \mathrm{Rev}[\hat{x}_i^{\hat{q}}]$ is equal to the marginal revenue $\sum_i \mathrm{MR}[\hat{x}_i^{\hat{q}}] = \sum_i R_i(\hat{q}_i)$. This observation approximately generalizes as follows. The marginal revenue mechanism is a an $e/(e-1)$ approximation in service-constraind matroid environments and an $O(\log n)$ bound for downward-closed environments on n quasi-linear-utility agents.

Matroid environments, by single-dimensional-agent reduction. Marginal revenue maximization is an $e/(e-1)$ approximation for service-constrained matroid environments, i.e., when the feasibility constraint is induced by independent sets of a matroid set system. Multi-unit environments, where at most a fixed number k of the agents can be simultaneously served, are a special case of matroid environments (corresponding to the k-uniform matroid). For the k=1 unit environment, which corresponds to a single-item auction, the bound remains $e/(e-1)$; for general k the bound improves to $\sqrt{2\pi k}/(\sqrt{2\pi k}-1)$, as simplified by Stirling's approximation. These results follow by reduction to the correlation gap theorem of Yan (2011).

Our approach is to reduce the question of approximation of the optimal mechanism by the marginal revenue mechanism to a question of approximation in the single-dimensional analog environment (recall Definition 8). In particular, we consider relaxing the feasibility constraint to hold *ex ante* instead of ex post. Such a relaxation potentially enables a higher revenue to be obtained. The single-dimensional-agent approximation question is to quantify the extent to which the optimal mechanism for the ex post feasibility constraint approximates the optimal mechanism for the ex ante feasibility constraint.

[12]For showing approximate linearity for downward-closed environments it is expedient to incorporate the downward closure into the outcome space by duplicating each non-service outcome and relabeling the duplicate outcome as a service outcome. This transformation is allowed for downward closed environments because we are always allowed to withhold service to an agent who would otherwise be served and this withholding will not violate the feasibility constraint. Of course, with such a transformation the revenue curves are non-decreasing.

Definition 18. A profile of ex ante service probabilities $\hat{\boldsymbol{q}} = (\hat{q}_1, \ldots, \hat{q}_n)$ is *ex ante feasible* if there is a distribution over feasible subsets of agents such that for each i, $\hat{q}_i$ is the (ex ante) probability agent i is in the subset. The *ex ante optimal mechanism* is the one that maximizes $\sum_i R_i(\hat{q}_i)$ subject to ex ante feasibility of $\hat{\boldsymbol{q}}$.

Proposition 21. *The ex ante optimal revenues for a general service constrained environment and its single-dimensional analog are equal and an upper bound on the (ex post feasible) optimal revenues (which may not be equal). If the optimal mechanism is a β-approximation to the ex ante optimal revenue in the single-dimensional analog environment, then the marginal revenue mechanism is a β-approximation to the optimal revenue in the original environment.*

Proof. The ex ante optimal revenue is defined only in terms of revenue curves and feasibility for the service constrained environment; therefore, a general environment and its single-dimensional analog have the same ex ante optimal revenue. By Definition 8 the (ex post feasible) optimal revenue in the single-dimensional analog is equal to the (ex post feasible) optimal marginal revenue of the original environment. To show the reduction, then, it suffices to observe that the ex ante optimal revenue is an upper bound on the optimal revenue in the original environment. As every ex post feasible mechanism is ex ante feasible (i.e., the latter is a relaxation of the former), the observation holds. $\square$

The following single-dimensional agent theorem is an immediate consequence of results of Yan (2011); his results, in fact, gave a specific (ex post feasible but non-optimal) mechanism that satisfies the claimed bound. Of course, then, the (ex post feasible) optimal mechanism satisfies the bound too. We obtain our desired result for general agents as a corollary of this theorem and Proposition 21.

Theorem 22. *For matriod environments with single-dimensional linear agents, the (ex post feasible) optimal mechanism is an $e/(e-1)$ approximation to the ex ante optimal mechanism; in any k-unit environment the bound improves to $\sqrt{2\pi k}/(\sqrt{2\pi k}-1)$.*

Corollary 23. *In any service constrained matroid environment, the marginal revenue mechanism is an $e/(e-1)$ approximation to the optimal mechanism; in any service constrained k-unit environment the bound improves to $\sqrt{2\pi k}/(\sqrt{2\pi k}-1)$.*

Downward-closed environments. In this section we show that in downward-closed environments and for a large class of agent preferences, the optimal marginal revenue is a logarithmic approximation, in the number of agents, to the optimal revenue. For example, this class includes quasi-linear preferences. In contrast to Section 5.1 where we gave a four approximation for unit-demand preferences with a product distribution (over alternatives), the results here apply, for example, to agents with correlated value distributions over alternatives and to quasi-linear preferences beyond unit demand.

To show this result we will incorporate the downward closure of the environment in the single-agent lottery pricing problems. Specifically, it is without loss of generality for downward-closed environments to duplicate every non-service outcome and label the duplicate a service outcome. This transformation implies that revenue is monotone in the allocation constraint, i.e., weaker constraints give no lower revenue.

A summary of the construction in the proof is as follows. If we consider allocation constraints with a minimum probability of 2^{-K} for allocating to any type, then the allocation constraint can be partitioned into K pieces such that the highest and lowest probabilities of allocation in each piece are within a factor of two of each other. If the single-agent lottery pricing problems satisfy a natural scalability property then the revenue of each piece can be approximated by a $\hat{q}$ ex ante optimal pricing scaled appropriately so that it is dominated by the original allocation constraint. The optimal revenue, then, is at most an $O(K)$ multiple of the revenue of the best such scaled ex ante optimal pricing. By downward closure, the optimal marginal revenue exceeds this revenue and is thus an $O(K)$ approximation. We obtain a logarithmic approximation by observing that attention can be restricted to allocation constraints for which $K \approx \log n$.

Recall that the revenue operator $\mathrm{Rev}[\cdot]$ is concave in its argument and therefore, for any $\gamma \in [0,1]$, $\mathrm{Rev}[\gamma\hat{x}] \geq \gamma \,\mathrm{Rev}[\hat{x}] + (1-\gamma)\,\mathrm{Rev}[0]$, where $\mathrm{Rev}[0] = R(0)$ is the optimal revenue when the agent is never served. We assume for simplicity of exposition that $\mathrm{Rev}[0] = R(0) = 0$, i.e., that an agent who is not served generates no revenue. The revenue scalability property we need is the opposite of this inequality, which, if the property holds, must therefore be an equality. In fact, revenue scalability can be viewed as a very permissive relaxation of revenue linearity.

Definition 19. An agent is *revenue scalable* if for any $\gamma \in [0,1]$ and any allocation constraint $\hat{x}$, the optimal revenue for $\gamma\hat{x}$ is equal to γ times the optimal revenue for $\hat{x}$. I.e.,

$$\mathrm{Rev}[\gamma\hat{x}] = \gamma\,\mathrm{Rev}[\hat{x}].$$

For example, as we will show, quasi-linear agents satisfy revenue scalability, but are not generally revenue linear. Moreover, if individual rationality is assumed, which usually implies that the utility and payment of an agent for non-service outcomes is zero, then even non-quasi-linear agents are revenue scalable. These observations are formalized in the following lemma.

Lemma 24. *Both (a) quasi-linear agents with no value for non-service outcomes and (b) agents with no utility and payment for any non-service outcome are revenue scalable.*

Proof. A key property of agents that are quasi-linear or have no utility and payment for non-service outcomes is that their utility and payment for any non-service outcome can be arbitrarily scaled upward. If an agent's utility and payment for a non-service outcome is zero, then scaling it upwards is trivial; if an agent is quasi-linear then his value for a non-service outcome is (minus) his payment and quasi-linearity requires that scaled payments translate to scaled utility. Thus, it suffices to show that agents with scalable utility and payment for non-service outcomes are revenue scalable.

Consider any allocation constraint $\hat{x}$ and the optimal lottery pricing for the scaled constraint $\gamma\hat{x}$. Denote by L the set of priced lotteries. As $\gamma\hat{x}(\hat{q}) \leq \gamma$ for all $\hat{q}$, the probability of a service outcome in any of the lotteries of L is at most γ. The theorem holds if we can define an alternative set of priced lotteries L' that meets the constraint $\hat{x}$ where the utility and payment of any type for any lottery is scaled upward a $1/\gamma \geq 1$ multiple. This is achieved by scaling the probability of any service outcome in any lottery upwards by a $1/\gamma$ multiple (without changing its payment), scaling the remaining probability of non-service outcomes down (so that the total probability is one), and scaling the utility and payment for non-service outcomes so that it is $1/\gamma$ multiple of that for the original lottery (which is possible by the assumption on scalability for non-service outcomes). Let $\gamma x(\cdot)$ denote the optimal allocation rule for constraint $\gamma\hat{x}$; the allocation rule from this construction is x and it is feasible for constraint $\hat{x}$. ☐

We now show that the marginal revenue approximates the optimal revenue for revenue-scalable agents in downward-closed service-constrained environments.

Lemma 25. *For a revenue-scalable agent, any allocation constraint with minimum allocation probability* $\hat{x}(1) \geq 2^{-K}$ *has revenue* $\mathrm{Rev}[\hat{x}]$ *at most* $2K\,\mathrm{MR}[\hat{x}]$.

Proof. Let $R^* = \mathrm{Rev}[\hat{x}]$ be the optimal revenue under allocation constraint $\hat{x}$. Let $x \preceq \hat{x}$ (where notation $x \preceq \hat{x}$ denotes allocation rules whose cumulative allocation rules satisfy $X(\hat{q}) \leq \hat{X}(\hat{q})$ for $\hat{q} \in [0,1]$; importantly, $X(1) = \hat{X}(1)$ is not required) be the allocation of optimal mechanism subject to $\hat{x}$. Therefore, $\mathrm{Rev}[x] = \mathrm{Rev}[\hat{x}]$. If we prove the claim for x, the proof for $\hat{x}$ follows because $\mathrm{Rev}[\hat{x}] = \mathrm{Rev}[x] \leq 2K\,\mathrm{MR}[x] \leq 2K\,\mathrm{MR}[\hat{x}]$, where the last inequality follows by definition of dominance and concavity of the revenue function. Therefore, in the rest of the proof we can assume without loss of generality that the optimal allocation subject to $\hat{x}$ is $\hat{x}$ itself.

Define a sequence of quantiles $0 = q_0 \leq q_1 \leq \cdots \leq q_K = 1$ such that $\hat{x}(q_{j-1}) \leq 2\hat{x}(q_j)$, for $j \in \{1, \ldots, K\}$. Define R_j^* to be the expected revenue from types that are mapped to a quantile in $[q_{j-1}, q_j]$, where the quantile of a type t is the probability that a type drawn at random has a higher probability of service than that of t (as per Definition 7 in Section 3). 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 allocation rules $z_j(\cdot)$ for all j, such that $z_j \preceq \hat{x}$ (where notation $x \preceq \hat{x}$ denotes allocation rules whose cumulative allocation rules satisfy $X(\hat{q}) \leq \hat{X}(\hat{q})$ for $\hat{q} \in [0,1]$; importantly, $X(1) = \hat{X}(1)$ is not required), and also $R_j^* \leq 2\,\mathrm{MR}[z_j]$. In particular, for j^* we will have $z_{j^*} \preceq \hat{x}$, and $2\,\mathrm{MR}[z_{j^*}] \geq R_{j^*}^* \geq R^*/K$, which will imply that

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

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

The main technical component of the proof is to show that, for z_j defined above, $R_j^* \leq 2\,\mathrm{MR}[z_j]$. By construction of z_j, and recalling that $\hat{x}(q_j) \leq 2\hat{x}(q_{j+1})$,

$$\begin{aligned} 2\,\mathrm{MR}[z_j] &= 2\int_0^1 z_j(q) R'(q)\,\mathrm{d}q \\ &= 2\hat{x}(q_{j+1}) R(q_{j+1} - q_j) \\ &\geq \hat{x}(q_j) R(q_{j+1} - q_j) \end{aligned}$$

It is therefore sufficient to show that $\hat{x}(q_j) R(q_{j+1} - q_j) \geq R_j^*$. 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 $\hat{x}$ with probability at most $\hat{x}(q_j)$. Now define L to be the set of lotteries chosen by types in $[q_j, q_{j+1}]$, and offer only these lotteries to the agent.[13] Notice that types with quantiles in $[q_j, q_{j+1}]$ choose the same lottery in L as they did in $\hat{x}$ (whereas other types that used to choose a lottery either switch to some lottery in L or no longer choose one if none in L give them non-negative utility). 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 $q_{j+1} - q_j$.[14] Call this new set of lotteries L'. Notice that the revenue from L' is at least

[13] Recall that, by the taxation principle, any incentive compatible mechanism consists of a set of lotteries, from which the agent chooses the one maximizing her utility.

[14] This requires continuity of the type space. We assume continuity for simplicity, but the proof can be easily generalized to handle discrete types.

R_j^*. Now recall that all the lotteries in L, and therefore L', allocate with probability at most $\hat{x}(q_j)$. Revenue scalability implies that the revenue of L' is at most $\hat{x}(q_j)R(q_{j+1} - q_j)$.

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

Theorem 26. *In downward-closed revenue-scalable environments with n agents, 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 serves that agent, regardless of his type and without charging him. This alternative mechanism is obviously incentive compatible, and its revenue is half of the optimal. Let $x_1, \ldots, x_n$ be the allocation rules for the alternative mechanism. Notice also that by construction, for each i and any $q \in [0, 1]$ we have $x_i(q) \geq 1/2n$. Therefore we can invoke Lemma 25 with $K = \log 2n$ to conclude that the revenue of the alternative mechanism is at most

$$2 \log n \sum\nolimits_i \mathrm{MR}_i[x_i]. \qquad \square$$

6 Single Dimensional Extension Theorems

The marginal revenue approach allows natural generalizations of techniques developed for single-dimensional linear agent environments. We will focus here on results for the approximation of optimal mechanisms by simple mechanisms. In such a study we are not free to arbitrarily design the simple mechanism. Instead, we show that performance guarantees for simple mechanisms are often obtainable by relating them to marginal revenue mechanisms.

Consider a variant of the red-or-blue car example from the introduction. There are n agents, k cars, and m possible colors. The social surplus maximizing mechanism (a.k.a. VCG; see Vickrey, 1961; Clarke, 1971; and Groves, 1973) selects the k agents whose values for their favorite color are the highest, serves these agents, and paints each car as the agent prefers. We consider this mechanism simple and practical, and we compare its revenue against the optimal revenue (cf. Hartline and Roughgarden, 2009). Shortly we will give conditions under which this mechanism is approximately optimal.

One feature of the VCG mechanism in service constrained environments is that, ex post, i.e., after agents make reports to the mechanism, each agent faces a *uniform* price over the alternatives. This price is equal to the favorite-color value among the other agents. Thus, in the interim stage each agent faces a distribution over uniform prices. We show that the VCG mechanism has near-optimal revenue in two steps. First, we show that the VCG revenue is close to the revenue of the optimal mechanism that only offers agents uniform prices. Second, we show that the latter revenue is close to the optimal revenue by any mechanism. An important observation is that the intermediate revenue in between these two steps is the optimal pseudo marginal revenue with uniform ex ante pseudo pricings (cf. Definition 15).

For the first step of the argument, the gap between the VCG revenue and the optimal pseudo marginal revenue is governed by the single-dimensional theory. Both mechanisms opperate on the type space given by projection of each unit-demand agent's multi-dimensional type into the single-dimensional space given by his value for his favorite alternative. In particular, the VCG revenue for the unit-demand agents is equal to its revenue under this single-dimensional projection, and the optimal pseudo marginal revenue for the unit-demand agents is equal to the optimal revenue for the projection. For the second step in our argument, by the theory of agent-based approximation we developed in Section 5 (e.g., Proposition 19), we need only analyze how good the uniform pricings are, as ex ante pseudo pricings.

We consider a concrete simple case before developing the general theory. In the above car-selling example, let k be 1, and each agent's value for each color be i.i.d. (i.i.d. among agents and i.i.d. across the alternatives). Since each agent's values for the alternatives are i.i.d., a uniform price for an agent is also a uniform virtual price (see Definition 17). Thus, Theorem 20 implies that the optimal pseudo marginal revenue (with uniform ex ante pseudop pricing) is a four approximation to the optimal revenue. This constitutes the second step of our planned argument. For the first step, the standard single-dimensional theory. If the distribution of the i.i.d. unit-demand agent's favorite-alternative value satisfies the regularity condition of Myerson (1981), then the theorem of Bulow and Klemperer (1996) implies that, for its single-dimensional analog, the second-price auction is an $\frac{n}{n-1}$ approximation to the single-dimensional optimal revenue, which is in turn equal to the optimal pseudo marginal revenue for the unit-demand agents. Combining the two steps, we have shown that the VCG mechanism for unit-demand agents is a $\frac{4n}{n-1}$ approximation. This discussion is formalized and generalized belowq.

Definition 20. A unit-demand agent is β *uniformly priceable* if, for any allocation constraint $\hat{x}$, a distribution over uniform pricings gives a β approximation to the optimal lottery pricing. The *uniform ex ante pseudo pricings* are the optimal of these pricings for ex ante constraints.

Definition 21. The *favorite-alternative single-dimensional analog* of a unit-demand service constrained environment is given by projecting the values of each unit-demand agent to the value of his favorite alternative. The *favorite-alternative extension* of a single-dimensional mechanism is a mechanism for unit-demand agents that simulates the given single-dimensional mechanism on reported values for favorite alternatives (ignoring the other values). It serves the winners of the simulation their favorite alternatives at the prices of the simulation.

Proposition 27. *For any service constrained environment, unit-demand β-uniformly-priceable agents, and α-approximation mechanism $\mathcal{M}$ for the favorite-alternative single-dimensional analog environment; the favorite-alternative extension of $\mathcal{M}$ is an $\alpha\beta$ approximation for the original environment.*

Proof. By construction, the revenue of the favorite-alternative extension of $\mathcal{M}$ in the original environment is equal to the revenue of $\mathcal{M}$ for the favorite-alternative single-dimensional analog environment. By assumption of the proposition, this revenue is an α-approximation to the optimal revenue for the favorite-alternative single-dimensional analog. This single-dimensional optimal revenue is equal to the optimal pseudo marginal revenue (with uniform ex ante pseudo pricings) in the unit-demand environment. By the assumption that the unit-demand agents are β uniformly priceable, Proposition 19 implies that this optimal pseudo marginal revenue is a β approximation to the unit-demand optimal revenue. These bounds combine to imply that the revenue of the favorite-alternative extension is an $\alpha\beta$ approximation to the optimal revenue. □

To draw single-dimensional extension theorems as consequences to Proposition 27, we first claim that unit-demand agents with values for each alternative independently drawn from (not necessarily identical) regular distributions are eight uniformly priceable. After this, we apply tools from the single-dimensional theory to provide approximation mechanisms for the favorite-alternative single-dimensional analog, and draw immediate corollaries.

Uniform Priceability. As described above, the i.i.d. special case of Theorem 20 implies that any unit-demand agent with i.i.d. values for distinct alternatives is four uniformly priceable. This result approximately generalizes to non-i.i.d. distributions that satisfy the regularity condition of Myerson (1981) as follows.

Definition 22. A distribution specified by distribution function F and density function f is *regular* if $v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing in v. A single-dimensional linear agent is regular if his value is drawn from a regular distribution.[15]

Lemma 28. *A unit-demand agent with values for alternatives drawn independently from (not necessarily identical) regular distributions is eight uniformly priceable.*

Proof Sketch. The proof will follow the template given by Proposition 19 with the following main ingredients.

- Twice the optimal revenue of the representative environment (where the unit-demand agent is replaced by single-dimensional representatives for each alternative) is a linear upper bound on the optimal revenue for any allocation constraint (by Lemma 34).
- Uniform pricing in the representative environment with regular distributions gives a four approximation to the optimal representative revenue; the argument is as follows. Hartline and Roughgarden (2009) show that the second-price auction with a uniform (a.k.a., anonymous) reserve price is a four approximation to the optimal revenue. In fact, this result can be strengthened (a) using a prophet-inequality-like proof to give the same bound for uniform pricing and (b) to allow an ex ante constraint on the probability that any representative is served. These extensions follow from a relatively straightforward modification of Lemma 35 which we omit.
- Uniform pricing in the original environment has the same revenue as uniform pricing in the representative environment. □

Below we will make use of the following slight strengthening of the regularity condition of Definition 22.

Definition 23. A unit-demand agent is *favorite-alternative regular* if the distribution of the agent's value for favorite alternative is regular; a unit-demand agent is *individual-alternative regular* if the agent's value for each alternative is regular; a unit-demand agent is *regular* if he is both favorite- and individual-alternative regular.[16]

Note that Lemma 28 requires only individual-alternative regularity.

[15]Single-dimensional regularity is equivalent (a) to $P(\hat{q}) = R(\hat{q})$ for all $\hat{q}$ (see the proof of Theorem 3), and (b) to the revenue-optimal allocation rule for $\hat{x}$ being $\hat{x}$ itself (see Lemma 7). These properties of regular distributions enable approximation of optimal mechanisms by simple mechanisms in single-dimensional environments.

[16]Neither favorite-alternative nor individual-alternative regularity imply the other.

Monopoly and Anonymous Reserve Pricing. For a single-dimensional single-agent problem, the *monopoly price* is the price that optimizes revenue. For single-dimensional, i.i.d., regular, matroid environments the surplus maximizing mechanism (a.k.a. VCG) with the monopoly reserve price (for the distribution) is revenue optimal. Hartline and Roughgarden (2009) approximately extend this result to non-identical distributions. They show that with regular single-dimensional agents, the revenue of the surplus maximizing mechanism with monopoly reserves is a two approximation to the optimal revenue. The following is a corollary of the above development and their theorems.

Corollary 29. *For independent unit-demand favorite-alternative-regular β-linearly-priceable agents and matroid service-constrained environments, the surplus maximizing mechanism with monopoly reserves (for distributions of favorite alternatives) is a 2β approximation to the optimal revenue. For regular unit-demand agents, $\beta = 8$.*

For single-item environments a similar approximation bound holds for an anonymous reserve price, i.e., one that is the same across the distinct agents. Hartline and Roughgarden (2009) show that with regular single-dimensional linear agents in single-item environments, the revenue of the second-price auction with an appropriate anonymous reserve is a four approximation to the optimal revenue. From this result we obtain the following corollary.

Corollary 30. *For independent unit-demand favorite-alternative-regular β-linearly-priceable agents and single-item service-constrained environments, the surplus maximizing mechanism with a suitably choosen anonymous reserve price is a 4β approximation to the optimal revenue. For regular unit-demand agents, $\beta = 8$.*

Market Expansion. Bulow and Klemperer (1996) show that the revenue of the single-item second-price auction for n i.i.d. regular agents is at least the revenue of the optimal auction for $n-1$ agents. An interpretation of this result is that the revenue loss of running the (surplus-optimal) second-price auction instead of the revenue-optimal auction can be made up by recruiting one more agent to the auction. This result generalizes to matroid environments, see e.g., Dughmi et al. (2009), where the revenue of the surplus maximizing mechanism is at least the revenue of the optimal auction after removing a *base* of the matroid.[17] The corollary, below, extends the ($k = 1$) multi-unit result described informally in the beginning of this section.

Corollary 31. *For i.i.d. unit-demand favorite-alternative-regular β-linearly-priceable agents and matroid service-constrained environments, the surplus maximizing mechanism is a β approximation to the optimal revenue with any base of the matroid removed. For regular unit-demand agents, $\beta = 8$.*

Prior-Independent Mechanisms. Dhangwatnotai et al. (2010) show that, in regular single-dimensional matroid environments, the surplus maximizing auction where each agent faces a reserve price randomly drawn from his value distribution is a four approximation to the optimal auction. If the agents' values are identically distributed then the approximation factor improves to two. Moreover, as long as there are at least two agents with values drawn from the same distribution, this approximation result can be obtained by a prior-independent mechanism, i.e., one that is not

[17] A base of a matroid is a feasible set with maximum cardinality.

parameterized by the prior-distribution. We summarize the consequences of the i.i.d. result in general service-constrained matroid environments as follows.

Corollary 32. *For i.i.d. unit-demand favorite-alternative-regular β-linearly-priceable agents and matroid service-constrained environments, there is a prior-independent mechanism that is a 2β approximation to the optimal revenue. For regular unit-demand agents (whose values for alternatives are drawn independently from not necessarily identical distributions), $\beta = 8$.*

These results are meant as examples of single-dimensional results with automatic extensions to unit-demand service constrained environments. Many other single-dimensional results also can be extended.

References

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*.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *The American Economic Review*, 86(1):180–194.

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

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

Cai, Y., Daskalakis, C., and Weinberg, M. (2012b). Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Foundations of Computer Science*.

Cai, Y., Daskalakis, C., and Weinberg, M. (2013). Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *SODA*.

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 *EC*, pages 149–158.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11:17–33.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.

Dughmi, S., Roughgarden, T., and Sundararajan, M. (2009). Revenue submodularity. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 243–252. ACM.

Groves, T. (1973). Incentives in teams. *Econometrica*, 41:617–631.

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.

Hartline, J. D., Kleinberg, R., and Malekian, A. (2011). Bayesian incentive compatibility and matchings. In *SODA*.

Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.

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

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.

Rochet, J. (1985). The taxation principle and multi-time hamilton-jacobi equations. *Journal of Mathematical Economics*, 14(2):113 – 128.

Thanassoulis, J. (2004). Haggling over substitutes. *J. Economic Theory*, 117(2):217–245.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37.

Wilson, R. (1997). *Nonlinear Pricing.* Oxford University Press.

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

A Proofs from Section 4.2

We now give a procedure for implementing the marginal revenue mechanism (Definition 9) with general agents. Recall that in the marginal revenue mechanism, each agent faces a distribution over ex ante optimal pricings, where the distribution is given by marginal revenue maximization over single-dimensional analog agents having the same revneue curves. This maximization over single-dimensional analogs gives rise to an allocation constraint $\hat{x}^{MR}$, and then the $\hat{q}$ ex ante optimal pricing occurs with probability $-\frac{\mathrm{d}}{\mathrm{d}\hat{q}}\hat{x}^{MR}(\hat{q})$. We show that this mixture of ex ante optimal pricings is implementable within the marginal revenue mechanisms family (Definition 1), and the randomized mapping from type to quantile (Step 1) in this implementation is efficiently computable.

What properties are needed for such a mapping? First, for each agent, we need the quantile to be uniformly distributed over $[0, 1]$. This way, the distribution over marginal revenues faced by each agent is as if the competing agents are single-dimensional linear agents with the same revenue curves. This guarantees that, if we map an agent's type to a quantile q, the probability that she revenue wins a service in the marginal revenue mechanism is equal to $\hat{x}^{MR}(q)$. In other words, this property would designate an allocation probability $\hat{x}^{MR}(q)$ to each quantile q, and therefore in order to get the desired allocation rules for types, we need only to come up with appropriate mappings of types to quantiles. Secondly, we would like the allocation rules obtained by this procedure to match the allocation rule given by the previously described mixture over ex ante optimal pricings.

To be specific, recall that each ex ante optimal pricing is derived from optimizing revenue subject to a step function constraint. The resulting normalized allocation rule may not be a step function and is in general weaker. When these ex ante optimal pricings are composed into the mixture, the resulting allocation rule, which we denote by x^{MR}, is dominated by and not necessarily equal to $\hat{x}^{MR}$. It is the allocation rule x^{MR}, and not $\hat{x}^{MR}$, that we would like to produce.

Recall from the discussion of Definition 9 that our goal is to implement the outcome rule $\tilde{w}^{MR}$. If we order the types according to $\mathrm{Alloc}(\tilde{w}^{MR}(\cdot))$, we get a natural mapping from types to quantiles: $\mathrm{Quant}(t) \triangleq \mathbf{Pr}_{t' \sim F}[\mathrm{Alloc}(\tilde{w}^{MR}(t')) \geq \mathrm{Alloc}(\tilde{w}^{MR}(t))]$. This mapping will have the first property[18], i.e., $\mathrm{Quant}(t)$ will be uniformly on $[0,1]$, but it does not have the second property. This is because by definition the probability of type t winning in the marginal revenue mechanism with this mapping is $\hat{x}^{MR}(\mathrm{Quant}(t))$. Overall, we get the allocation rule $\hat{x}^{MR}$ and not the weaker x^{MR}. If we could keep the first property, then the problem reduces to the following: given two non-increasing functions $x^{MR}, \hat{x}^{MR} : [0,1] \to [0,1]$, such that x^{MR} is weaker than $\hat{x}^{MR}$ (in the sense that $\hat{X}^{MR} \geq X^{MR}$ pointwise), is there a randomized function $g : [0,1] \to [0,1]$, such that $\mathbf{E}[\hat{x}^{MR}(g(q)] = x^{MR}(q)$ for every $q \in [0,1]$, and $g(q)$ is uniform on $[0,1]$ when q is uniform on $[0,1]$? This is a problem addressed by the theory of *majorization* (see, e.g. Hardy et al., 1929), and has a general solution. In our context, we give a particularly simple interval resampling procedure that gives this mapping g, which is to be composed with $\mathrm{Quant}(\cdot)$ for the eventual randomized mapping from types to quantiles.

Definition 24. For allocation constraint $\hat{x}$ and dominated allocation rule x satisfying $\hat{X}(1) = X(1)$ on m discrete types, the *interval resampling sequence construction* starts with $x^{(0)} = \hat{x}$ 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 $X^{(j)}$.[19]
3. The jth resampling interval is $[q, q']$.
4. Let $x^{(j+1)}$ be $x^{(j)}$ averaged on $[q, q']$.

Proposition 33. *The interval sampling sequence construction gives a sequence of at most m intervals such that the composition of $\hat{x}$ with the sequence of resamplings applied to* $\mathrm{Quant}(\cdot)$ *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 $\mathrm{Quant}(\cdot)$, satisfy $x^{(j-1)}(\mathrm{Quant}(t)) = x(\mathrm{Quant}(t))$. Consider step j. The assumption that $\hat{X}(1) = X(1)$ 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 $x^{(j)}(\mathrm{Quant}(t)) = x(\mathrm{Quant}(t))$ for the jth type. The procedure is linear time as both $\hat{x}$ 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 $\hat{x}$ or x is processed. □

[18] As before, we break ties appropriately.

[19] 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 the boundary between them at q'.

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

1. Map reported types $\boldsymbol{t} = (t_1, \ldots, t_n)$ of agents to quantiles $\boldsymbol{q} = (q_1, \ldots, q_n)$ by, for each agent, composing the interval resampling transformation with $\mathrm{Quant}(\cdot)$.
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 $\mathrm{Alloc}(w_i) = 1$ if i is to be served, or
   - sample $w_i \sim \tilde{w}_i^{MR}(t_i)$ conditioned on $\mathrm{Alloc}(w_i) = 0$ if i is not to be served.

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

B Proofs from Section 4.3

The technique for the proof of Proposition 16 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, equivalently, the 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 16. For this proof we will only use allocations for types (instead of quantiles), and to simplify notation we let $x(v)$ be the allocation probability for type v. Without loss of generality, we assume that the highest valuation in the support of F is 1. The standard incentive compatibility condition for single-dimensional linear preferences (monotonicity of the allocation rule and the payment identity) still holds. 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 proposition then states that, in $\hat{q}$ ex ante optimal pricing, the allocation for types smaller than $\bar{v}$ is constantly 0.

Payment identity states 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

$$\max \int_0^{\bar{v}} f(v)x(v)\phi(v)\,\mathrm{d}v + [1 - F(\bar{v})]\bar{v}x(\bar{v}), \tag{6}$$

where $\phi(v)$ is the standard virtual valuation function $v - \frac{1-F(v)}{f(v)}$, subject to the constraints:

$$\bar{v}x(\bar{v}) - \int_0^{\bar{v}} x(v)\,\mathrm{d}v = B, \tag{7}$$

$$\int_0^{\bar{v}} x(v)\,\mathrm{d}v + [1 - F(\bar{v})]x(\bar{v}) = \hat{q}, \tag{8}$$

$$\forall v, \quad x(v) \geq 0, \tag{9}$$

$$x(\bar{v}) \leq 1. \tag{10}$$

We consider the first-order conditions for the above program. We use δ for the Lagrangian variable for the budget condition (7); λ for the ex ante selling probability constraint (8); Π_v for condition (9) for each v ($\Pi_v \leq 0$); η for condition (10) ($\eta \geq 0$). The first order condition gives

$$f(v)\left[\phi(v) + \lambda - \frac{\delta}{f(v)}\right] + \Pi_v = 0, \quad \forall v < \bar{v}; \tag{11}$$

$$[1 - F(\bar{v})]\left[\bar{v} + \lambda + \frac{\bar{v}\delta}{1 - F(\bar{v})}\right] + \Pi_{\bar{v}} + \eta = 0. \tag{12}$$

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

$$\phi(v) + \lambda - \frac{\delta}{f(v)} = 0;$$

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

We can therefore solve for δ:

$$\delta = \left[\phi(v) - \bar{v} - \frac{\eta}{1 - F(\bar{v})}\right] \Big/ \left[\frac{\bar{v}}{1 - F(\bar{v})} + \frac{1}{f(v)}\right] < 0. \tag{13}$$

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

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

or

$$\phi(v) - \phi(v') = \delta\left(\frac{1}{f(v)} - \frac{1}{f(v')}\right). \tag{14}$$

Suppose $v < v'$, then $f(v) \geq f(v')$ by our assumption. Since the distribution is regular, we have $\phi(v) \leq \phi(v')$. Additionally, we know that $\delta < 0$, and so (14) 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 (14) 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. [20] Therefore the solution to the program gives rise to an incentive compatible mechanism, which satisfies Proposition 16. $\square$

Derivation of Example 4. We first derive the $\hat{q}$ ex ante optimal pricings for $\hat{q} < 1 - F(B) = 1 - B$. By Proposition 16, a lottery that costs B is offered, and, when bought, it sells the item with probability $\pi = B + \hat{q}$. A type with value at least $B/(B + \hat{q})$ will buy the lottery (and hence wins with probability $B + \hat{q}$). For $\hat{q} > 1 - B$, the budget does not bind and the item is sold at a price of $1 - \hat{q}$; all types with $v \geq 1 - \hat{q}$ wins the item with certainty. This immediately shows us the shape of $G_v(\hat{q})$, the probability of allocation as a function of $\hat{q}$ for a fixed type v. From the perspective of a given type $v \geq B$, $G_v(\hat{q})$ jumps starts at $\hat{q} = \frac{B}{v} - B$, increases linearly with $\hat{q}$ and saturates at $\hat{q} = 1 - B$. This is depicted in Figure 1a. For $v < B$, the budget never binds, and the corresponding $G_v(\hat{q})$ is the familiar step function (Figure 1b).

Calculating the revenue curve is straightforward. For $\hat{q} < 1 - B$, $R(\hat{q})$ is $B \cdot (1 - \frac{B}{B+\hat{q}})$. Its derivative, i.e., the marginal revenue, $\frac{B^2}{(B+\hat{q})^2}$, is strictly positive. (Note that, for $B < 1/2$, this is different from the linear preference case. There, the marginal revenue would be negative for $\hat{q} > 1/2$.) For $\hat{q} \geq 1 - B$, $R(\hat{q})$ is $\hat{q}(1 - \hat{q})$. Its derivative is positive for $\hat{q} < 1/2$ and negative for $\hat{q} > 1/2$.

By Step 1 of Definition 13, whenever $v' < B$, its quantile $1 - v'$ will be larger with probability 1 than the quantile of a type $v \geq B$, which is distributed between $\frac{B}{v} - B$ and $1 - B$. Therefore, such smaller v''s are only considered when there is no bidder bidding above B; when this happens, since the budget does not bind, the auction is the optimal one in the linear preference case, i.e., a second price auction with reserve price $\frac{1}{2}$. When there are bidders bidding above B, the way q_i is computed in Example 4 is simply sampling by $G_v(\hat{q})$ as stipulated in Definition 13. When i^* is the sole bidder bidding above B, she should pay the ex ante optimal pricing at her critical quantile. When $B < 1/2$, this critical quantile is $1 - B$, and she pays B; when $B > 1/2$, the critical quantile is $\min\{1/2, 1 - \max_{i \neq i^*} v_i\}$, and she pays $\max_{i \neq i^*}\{1/2, v_i\}$. When there are other bidders bidding above B, the critical quantile for i^* is always smaller than $1 - B$, and she will pay B in the corresponding ex ante optimal pricing. A losing bidder i bidding above B faces a critical quantile q_{i^*}. We see from Figure 1a that if q_{i^*} is larger than the quantile $\frac{B}{v_i} - B$ at which $G_{v_i}(\hat{q})$ jump starts, she will need to make a payment of B. This happens for $v_i \geq B/(q_{i^*} + B)$.

C Proofs for unit-demand approximation

Theorem 20 is a consequence of the two lemmas below and Proposition 19.

Lemma 34. *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 linear agents. We consider the collection of representatives as a whole (or, say, a single market), and we can ask

[20] 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.

what is the optimal revenue from this market given an ex ante selling probability $\hat{q}$ or an allocation constraint $\hat{x}$. Both terms are easy to find. 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 $\psi_{\max}(\hat{q})$. The optimal revenue for any allocation constraint $\hat{x}$ is $\mathbf{E}_{\hat{q}}[\psi_{\max}(\hat{q})\hat{x}(\hat{q})]$ which is linear in $\hat{x}$; this follows from the proof that the optimal revenue in single-dimensional environments is the virtual surplus maximizer.

We now show that, under any allocation constraint, 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 $\hat{x}$ and show that the sum of these auctions' revenue upper bounds the optimal unit-demand revenue for the same constraint. Of course, the optimal representative revenue in turn 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 a lottery for each type t taking the form of $(p(t), \pi^1(t), \ldots, \pi^m(t))$ with $\sum_j \pi^j(t) \leq 1$. The semantics of a lottery is that the agent pays the price $p(t)$ and then is allocated an alternative j at random with probability $\pi^j(t)$; the semantics of the collection of lotteries L is that the agent, upon drawing her type t from the distribution, chooses the lottery $(p(t), \pi^1(t), \ldots, \pi^m(t))$ that corresponds to her type.

Given any collection of lotteries L that satisfies the allocation constraint $\hat{x}$ 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 $\boldsymbol{v} = (v^1, \ldots, v^m)$ of the representatives and the lottery that would have been selected by the unit-demand agent with these values. It serves the representative j with the highest value with probability $\pi^j(t)$ and charges her (no matter whether we serve her or not) $p(t) - \sum_{j' \neq j} \pi^{j'}(t) v^{j'} + \mu(\boldsymbol{v}^{(2)})$ where $\mu(\boldsymbol{v}^{(2)})$ is the expected utility of the unit-demand agent with valuation profile $\boldsymbol{v}^{(2)}$ which is $\boldsymbol{v}$ with v_j replaced with $\max_{j' \neq j} v^{j'}$.

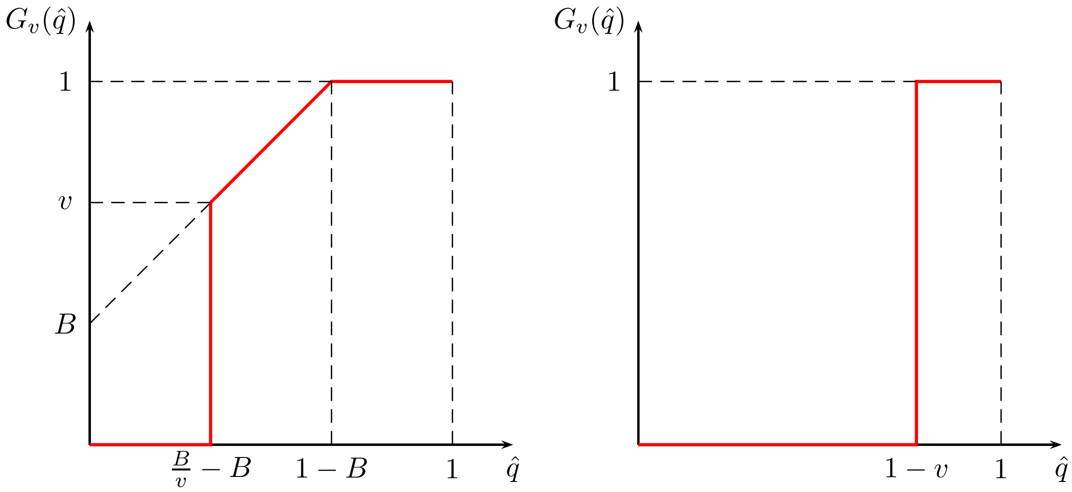

(a) $v \geq B$

(b) $v < B$

Figure 1: The allocation rules for $\hat{q}$ ex ante optimal pricings in Example 4, for a fixed type v.

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, $\boldsymbol{v}^{-j}$. As the utility of the unit-demand agent is monotone in her value for each alternative, the utility each representative has for winning is positive when she is the highest valued representative and negative if she is not (and were to misreport and pretend she were). Therefore, this auction is incentive compatible, has revenue at least $p(t) - \sum_{j' \neq j} \pi^{j'}(t) v^{j'}$ on valuation profile $\boldsymbol{v}$ where j is the highest valued representative, and satisfies allocation constraint $\hat{x}$. For a given valuation profile, call the second term in the winning agent's payment, $\sum_{j' \neq j} \pi^{j'}(t) v^{j'}$, 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 second highest value and serves with probability $\sum_j \pi^j(t)$ satisfies the allocation constraint $\hat{x}$ 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 $\hat{x}$ 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 $\hat{x}$ is $\mathbf{E}_q[\nu_{(2)}(q)\hat{x}(q)]$. To obtain this revenue, conditioning on the second highest value being v, with probability $\hat{x}(\nu_{(2)}^{-1}(v))$ we serve the highest valued representative and charge her v (only when we serve her). 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 35. *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 $\hat{q}$ for any $\hat{q}$-step constraint).*

Proof. The proof closely follows the standard prophet inequality proofs (see for example Chawla et al. (2010b)). As in the proof to the previous lemma, we may view the representatives as one entity and consider its optimal revenue under ex ante constraint on serving. Denote the optimal representative revenue for the $\hat{q}$-step constraint as a function of $\hat{q}$ by the revenue curve $\mathrm{ORR}(\hat{q})$. Consider the outcome of the optimal auction for the representative environment with ex ante service constraint $\hat{q}$. It sets a uniform virtual price (denoted $\psi(\hat{q})$) and serves the agent with the highest virtual value strictly bigger than $\psi(\hat{q})$ with probability one. If the probability that the largest virtual value is equal to $\psi(\hat{q})$ 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 $\psi(\hat{q})$ so as to serve with the desired ex ante probability $\hat{q}$. 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.

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

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

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

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

The second inequality followed because $\mathbf{Pr}[\mathcal{E}_i]$, the probability of the event that $\psi_j < \psi(\hat{q})$ for all $j \neq i$ is not less than the probability that $\psi_j < \psi(\hat{q})$ for all j, which is $(1-\hat{q})$. To extend this lower bound on $\tilde{R}(\hat{q})$ from $\hat{q} \in Q$ to all $\hat{q} \in [0,1]$, consider inserting a virtual value $\psi' = \psi(\hat{q}) + \epsilon$ with measure zero in the distribution. The $\hat{q}'$ that corresponds to serving this virtual value or higher has revenue bounded by the formula above but $\psi' \approx \psi(\hat{q})$. Keeping the virtual value constant and varying $\hat{q}$ in the formula interpolates a line between the two revenues. As the ex ante pseudo pricings are closed under convex combination, this line gives a lower bound on the $\hat{q}$ ex ante pseudo pricing. Therefore, the bound above on $\tilde{R}(\hat{q})$ holds for all $\hat{q}$.

To bound $\mathrm{ORR}(\hat{q})$ in terms of $\tilde{R}(\hat{q})$ we consider two cases. When $\hat{q} \leq 1/2$ 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 $1 - \hat{q} \geq 1/2$). To show the claim for $\hat{q} > 1/2$ notice that

$$\begin{aligned}
\mathrm{ORR}(1) &= \mathbf{E}\left[\max_i(\psi_i)_+\right] \\
&= \hat{v} + \mathbf{E}\left[\max_i(\psi_i)_+ - \hat{v}\right] \\
&\leq \hat{v} + \mathbf{E}\left[(\max_i(\psi_i)_+ - \hat{v})_+\right] \\
&= \hat{v} + \mathbf{E}\left[(\max_i \psi_i - \hat{v})_+\right] \\
&\leq \hat{v} + \sum\nolimits_i \mathbf{E}\left[(\psi_i - \hat{v})_+\right],
\end{aligned}$$

for any $\hat{v}$. As a result, by setting $\hat{v} = \psi(1/2)$ we get

$$\begin{aligned}
\mathrm{ORR}(1) &\leq \psi(1/2) + \sum\nolimits_i \mathbf{E}\left[(\psi_i - \psi(1/2))_+\right] \\
&\leq 2\tilde{R}(1/2).
\end{aligned}$$

From monotonicity of ORR and $\tilde{R}$ we then conclude that for any $\hat{q} > 1/2$, $\mathrm{ORR}(\hat{q}) \leq \mathrm{ORR}(1) \leq 2\tilde{R}(1/2) \leq 2\tilde{R}(\hat{q})$.

□

D Revenue Linearity for Unit Demand Valuations Uniform on Hypercubes

In this section we show that unit-demand quasi-linear-utility agents whose values for m alternatives are i.i.d. drawn from $U[0,1]$ are revenue linear. Recall from Appendix C that an incentive compatible mechanism offers a menu of lotteries to the agent. Each lottery takes the form of $(p(t), \pi^1(t), \ldots, \pi^m(t))$, where $\sum_j \pi^j(t) \leq 1$, with p denoting the price of the lottery and π_j the probability with which alternative j is allocated to the agent. We sometimes write π as the vector $(\pi^1, \ldots, \pi^m)$. In this section we abuse the notation and use u to denote a mapping that maps a type $t \in T = [0,1]^m$ to the expected utility of this type in an incentive compatible mechanism. We use the following lemma first noted by Rochet (1985).

Lemma 36. *For a quasi-linear-utility agent, a utility function u corresponds to an incentive compatible mechanism if and only if it is convex. In this case, $p(t) = \nabla u \cdot t - u(t)$, and $\pi(t) = \nabla u(t)$.*

In the above lemma $\nabla u(t)$ is the gradient of the function u. Since selling any alternative accounts as a service, by Lemma 36 the allocation of a type t is $||\nabla u(t)||_1$, the L_1 norm of the vector $\nabla u(t)$. Let W be the space of convex utility functions u, and c the cost of producing an alternative. Using Lemma 36, we can reformulate the problem of revenue maximization under allocation constraint $\hat{x}$ as follows:

$$\begin{aligned}
\text{maximize} & \int_T [\nabla u(t) \cdot t - u(t)] f(t)\, \mathrm{d}t - c \int_T \vec{1} \cdot \nabla u(t) f(t)\, \mathrm{d}t \\
\text{s.t.} \quad & u \in W \\
& \forall S \subseteq T, \int_S ||\nabla u(t)||_1\, \mathrm{d}t \leq \hat{X}(f(S)).
\end{aligned}$$

Recall from Section 2 the definition of the cumulative allocation constraint $\hat{X}$. Note also that the second constraint automatically guarantees the feasibility constraint: for all but a measure zero set of types, $||\nabla u(t)||_1 \leq 1$. By our assumption, $f(t)$ is 1 everywhere on $[0,1]^m$.

For any $t \in T$, define a scaling function $r_t : [0,1] \to T$ as $r_t(\alpha) = \alpha t$. Then $r_t(0) = \vec{0}$, and $r_t(1) = t$, for any t. We now use the gradient theorem and write

$$\forall t, u(t) - u(0) = \int_0^1 \nabla u(r(\alpha)) \cdot r_t'(\alpha)\, \mathrm{d}\alpha.$$

In a revenue optimal mechanism, $u(0) = 0$. Also, by definition of r, $r'(\alpha) = t$. Therefore,

$$u(t) = \int_0^1 \nabla u(\alpha t) \cdot t\, \mathrm{d}\alpha, \qquad \forall t \in T.$$

Using this, we can rewrite the objective function as

$$\begin{aligned}
& \int_T \left[\nabla u(t) \cdot (t - c\vec{1}) - \int_0^1 \nabla u(\alpha t) \cdot t\, \mathrm{d}\alpha \right] \mathrm{d}t \\
= & \int_T \nabla u(t) \cdot (t - c\vec{1})\, \mathrm{d}t - \int_0^1 \int_T \nabla u(\alpha t) \cdot t\, \mathrm{d}t\, \mathrm{d}\alpha.
\end{aligned}$$

In the second term, change variables by defining $v = \alpha t \in [0,1]^m$. Notice that $t = v/\alpha$, and $\mathrm{d}v^j = \alpha\, \mathrm{d}t^j$ for any $1 \leq j \leq m$. Therefore $\mathrm{d}v = \alpha^m\, \mathrm{d}t$. Define T_α to be the set of $t \in T$ such that

$\max_j t^j \leq \alpha$. The objective is now rewritten as

$$\begin{aligned}
&\int_T \nabla u(t) \cdot (t - c\vec{1})\,\mathrm{d}t - \frac{1}{\alpha^m}\int_0^1 \int_{v \in T_\alpha} \nabla u(v) \cdot (v/\alpha)\,\mathrm{d}v\,\mathrm{d}\alpha \\
&= \int_T \nabla u(t) \cdot (t - c\vec{1})\,\mathrm{d}t - \int_{v \in T} \nabla u(v) \cdot v \int_{\alpha = \max_j v^j}^1 \frac{1}{\alpha^{m+1}}\,\mathrm{d}\alpha\,\mathrm{d}v \\
&= \int_T \nabla u(t) \cdot (t - c\vec{1})\,\mathrm{d}t - \int_{v \in T} \nabla u(v) \cdot v \left[\frac{1}{m(\max_j v^j)^m} - \frac{1}{m}\right] \mathrm{d}v \\
&= \int_T \nabla u(t) \cdot \left[t\left(\frac{m+1}{m} - \frac{1}{m(\max_j t^j)^m}\right) - c\vec{1}\right] \mathrm{d}t.
\end{aligned}$$

Now, if we relax the convexity constraint, the optimization problem is expressed solely in terms of the gradient of u. Next we argue that the optimal solution to this optimization problem takes a particularly simple form. First note that the function $t^j(\frac{m+1}{m} - \frac{1}{mt^j})$ is increasing in t^j. Consider any feasible solution ∇u to the program and its alteration $\nabla \tilde{u}$ in the following manner: at any type t where j^* is $\arg\max_j t^j$, let $\nabla^{j^*} \tilde{u}$ be $\sum_j \nabla^j u$, and $\nabla^j \tilde{u}$ be 0 for all $j \neq j^*$. Since this alteration keeps the L_1-norm of ∇u, ∇u^* still satisfies all the constraints (except that we are relaxing the convexity constraint for now). But the objective function is pointwise better for $\nabla \tilde{u}$ than for ∇u. Therefore, it suffices to consider solution gradients whose coordinates at each type t are all zero except the one coordinate where the valuation is maximized (ties can be broken arbitrarily). But then the problem degenerates, and the optimal utility function is given by a simple greedy procedure, which grows, at each type, in the direction of the maximum valued alternative as much as allowed by the allocation constraint $\hat{x}$. Formally, the optimal utility function is given by

$$u^*(t) = \begin{cases} 0, & \max_j t^j \leq \hat{t}^c \\ \int_{\alpha = \hat{t}^c}^{\max_j t^j} \hat{x}(1 - \alpha^m)\,\mathrm{d}\alpha, & \max_j t^j > \hat{t}^c, \end{cases}$$

where $\hat{t}^c$ solves

$$\hat{t}^c \left(\frac{m+1}{m} - \frac{1}{m(\hat{t}^c)^m}\right) = c.$$

In particular, $\hat{t}^0 = \sqrt[m]{\frac{1}{m+1}}$. This utility function u specified above is convex and linear in $\hat{x}$. By Lemma 36, it is easy to see that u^* being linear implies that $\mathrm{Rev}[\cdot]$ is also linear (noting that integral is a linear functional).

To summarize, we have shown that the ex ante optimal mechanism for constraint $\hat{q}$ is to post a price of $\sqrt[m]{1 - \hat{q}}$ for any service. The quantile of each type $t = (t^1, \ldots, t^m)$ is $q = 1 - (\max_i t^i)^m$ (see Figure 2).

E Reverse Auctions

Reverse auctions can naturally be modeled as service constrained environments. Different agents, here sellers, have different costs for providing different services, and the auctioneer has possibly

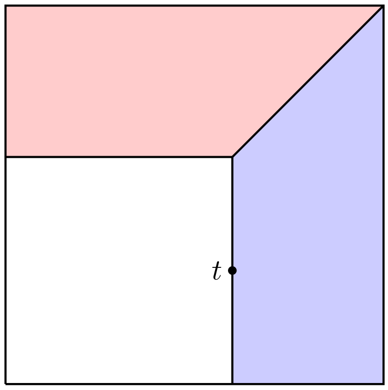

Figure 2: The quantile of type $t = (t^1, t^2)$ with $t^1 \geq t^2$ is $q = 1 - (t^1)^2$.

different values for different services, and wishes to acquire at most one service, and to do so in order to maximize the value for the service acquired minus the payment for the service. Notice that the goal of maximizing value minus payment is equivalent to minimizing payment minus value. Such an objective can be modeled as a forward auction in which the seller has possibly different costs for selling items. In Appendix D we solve the forward auction problem with uniform values and uniform costs, which implies the following results for the reverse auction problem.

More formally, we can transform a reverse auction problem into a forward auction as follows. Consider a single seller that can provide m services where each service i costs c_i. Assume that the cost of each service is drawn uniformly at random from the interval $[0, 1]$, and assume that the value of each service for the auctioneer is 1 (the analysis generalizes to arbitrary distributions and valuations, but the general analysis is not required here). Let $\pi(c) = (\pi^1(c), \ldots, \pi^m(c))$ be the vector of the probabilities of purchasing each service when the cost vector is c, and $p(c)$ the payment made to the seller by the auctioneer. Now the objective is to maximize

$$\int_{c \sim U[0,1]^m} \vec{1} \cdot \pi(c) - p(c) \, \mathrm{d}c.$$

We can also write the incentive compatibility constraint as

$$p(c) - c \cdot \pi(c) \geq p(c') - c \cdot \pi(c')$$

for all cost vectors c and c'. Now define functions $\bar{\pi}$ and $\bar{p}$ to be

$$\begin{aligned} \bar{\pi}(c) &= \pi(\vec{1} - c) \\ \bar{p}(c) &= \vec{1} \cdot \pi(\vec{1} - c) - p(\vec{1} - c). \end{aligned}$$

Using the above notation we can rewrite the objective to be

$$\int_{c \sim U[0,1]^m} \bar{p}(c) \, \mathrm{d}c. \tag{15}$$

Also,

$$p(c) - c \cdot \pi(c) = (\vec{1} - c) \cdot \bar{\pi}(\vec{1} - c) - \bar{p}(\vec{1} - c).$$

Therefore, the incentive compatibility constraint is equal to

$$c \cdot \bar{\pi}(c) - \bar{p}(c) \geq c \cdot \bar{\pi}(c') - \bar{p}(c'), \qquad \forall c, c'. \tag{16}$$

Now notice that the optimization problem given by (15) and (16) is equal to the standard formulation of a forward auction. We can therefore solve the reverse auction problems by transforming them into forward auction problems, solving the problem using our framework, and then transforming the solution back to the reverse auction setting.

In a reverse auction problem, classical auction theory says that (a) the optimal way to buy an object (henceforth: a bridge) with value 1 from a single agent with cost drawn from a uniform distribution on $[0, 1]$ is to offer a take-it-or-leave-it payment of $1/2$, (b) the optimal way to buy a bridge with value $1/2$ from one of multiple agents with uniformly distributed costs is to run a second-price reverse auction with reserve price $1/2$, in which the agent with the lowest cost (if it is less than $1/2$) constructs the bridge and is payed the minimum of the second lowest cost and $1/2$. The above interpretation of the marginal revenue mechanism in i.i.d. settings is one of the most important result in classical auction theory. Our theory generalizes this to multi-dimensional preferences as follows. Consider instead buying a bridge that can be built using technology 1 or technology 2. It says that (a) the optimal way to buy a bridge with value 1 from a single agent with costs for the different technologies each drawn independently and uniformly from $[0, 1]$ is to offer a take-it-or-leave-it payment of $1 - \sqrt{1/3}$ for either technology, (b) the optimal way to buy a bridge with value 1 from one of multiple agents each with i.i.d. uniform costs for each technology is to run the second-price reverse auction with reserve $1 - \sqrt{1/3}$ and allow the winning agent to choose her favorite technology to build the bridge.