

Bayesian Mechanism Design for Budget-Constrained Agents *

Shuchi Chawla
Univ. of Wisconsin-Madison
Madison, WI, USA
shuchi@cs.wisc.edu

David L. Malec
Univ. of Wisconsin-Madison
Madison, WI, USA
dmalec@cs.wisc.edu

Azarakhsh Malekian
Northwestern University
Evanston, IL, USA
a-malekian@northwestern.edu

ABSTRACT

We study Bayesian mechanism design problems in settings where agents have budgets. Specifically, an agent’s utility for an outcome is given by his value for the outcome minus any payment he makes to the mechanism, as long as the payment is below his budget, and is negative infinity otherwise. This discontinuity in the utility function presents a significant challenge in the design of good mechanisms, and classical “unconstrained” mechanisms fail to work in settings with budgets. The goal of this paper is to develop general reductions from budget-constrained Bayesian MD to unconstrained Bayesian MD with small loss in performance. We consider this question in the context of the two most well-studied objectives in mechanism design—social welfare and revenue—and present constant factor approximations in a number of settings. Some of our results extend to settings where budgets are private and agents need to be incentivized to reveal them truthfully.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Bayesian mechanism design, budgets, revenue, welfare

1. INTRODUCTION

Auction and mechanism design have for the most part focused on agents with quasilinear utility functions: each agent is described by a function that assigns values to possible outcomes, and the agent’s utility from an outcome is her value minus any payment that she makes to the mechanism. This implies, for example, that

*This work was supported in part by NSF awards no. CCF-0830494 and CCF-0643763, and a Sloan Foundation fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC’11, June 5–9, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0261-6/11/06 ...\$10.00.

an agent offered an outcome at a price below her value for the outcome should in the absence of better alternatives immediately accept that outcome. This simple model fails to capture a basic practical issue—agents may not necessarily be able to afford outcomes that they value highly. For example, most people would value a large precious stone such as the Kohinoor diamond at several millions of dollars (for its resale value, if not for personal reasons), but few can afford to pay even a fraction of that amount. Many real-world mechanism design scenarios involve financially constrained agents and values alone fail to capture agents’ preferences. Budget constraints have frequently been observed in FCC spectrum auctions [5, 8], Google’s auction for TV ads [18], and sponsored search auctions, to take a few examples.

From a theoretical viewpoint, the introduction of budget constraints presents a challenge in mechanism design because they make the utility of an agent nonlinear and discontinuous as a function of the agent’s payment—the utility decreases linearly with payment while payment stays below the budget, but drops to negative infinity when the payment crosses the budget. The assumption of linearity in payments (i.e. quasilinearity of utility) underlies much of the theoretical framework for mechanism design. Consequently, standard mechanisms such as the VCG mechanism can no longer be employed in settings involving budgets.

The goal of this paper is to develop connections between budget-constrained mechanism design and the well-developed theory of unconstrained mechanism design. Specifically we ask “when can budget-constrained mechanism design be reduced to unconstrained mechanism design with some small loss in performance?” We consider this question in the context of the two most well-studied objectives in mechanism design—social welfare and revenue. Some of our results assume that the mechanism knows the budgets of the agents, but others hold even when budgets are private and agents need to be incentivized to reveal them truthfully.

Recent work in computer science has begun exploring a theory of mechanism design for budget-constrained agents (see, for example, [1, 4, 11, 10, 2]). Most of this work has focused on *prior-free* or worst-case settings, where the mechanism designer has no information about agents’ preferences. Unsurprisingly, the mechanism designer has very little power in such settings, and numerous impossibility results hold. For example, in the worst-case setting no truthful mechanism can obtain a non-trivial approximation to social welfare [4]. The goal of achieving good social welfare has therefore been abandoned in favor of weaker notions such as Pareto optimality [11]. For the revenue objective while approximations can be achieved in simple enough settings, e.g. multi-unit auctions [4], hardness results hold for more general feasibility constraints even in the absence of budgets. In this paper, we sidestep these impossibility results by considering Bayesian settings where the mech-

anism designer has prior information about the distributions from which agents' private values and private budgets are drawn.

We restrict our attention to direct revelation truthful mechanisms. Our mechanisms are allowed to randomize, and agents' utilities are computed in expectation over the randomness used by the mechanism. As is standard, we assume that both the mechanism and the agents possess a common prior from which values are drawn. While we optimize over the class of Bayesian incentive compatible (BIC) mechanisms, all of the mechanisms we develop are dominant strategy incentive compatible (DSIC) (see, for example, [17] for definitions of these solution concepts).

In addition, we require that our mechanisms satisfy the ex-post individual rationality (EPIR) constraint, namely that the payment of any agent never exceeds her value for the mechanism's outcome. This implies, in particular, that the mechanism cannot charge any agent to whom no item or service is allocated. In contrast, most previous work has enforced the individual rationality constraint only in expectation over the mechanism's randomness as well as the randomness in other agents' values (i.e. interim IR).

It is worth noting here that the EPIR constraint is not without loss in performance. Consider the following example: suppose we are selling a single item to one of n agents, each with a value of v with probability $1/n$ (that is publicly known) and a public budget of v/n with probability $1/n$. Now, under the IIR constraint, the optimal auction asks agents to pay what they bid and offers each agent that pays at least v/n a fair chance at winning the item. Each agent pays v/n , the item is allocated to a random agent, and the mechanism's revenue is v . Under the EPIR constraint, however, a mechanism can only charge the agent that wins the item and can charge this agent no more than v/n . As we can see, the revenue gap between the optimal IIR and the optimal EPIR mechanism gets larger and larger as n grows.

It is well known that over the class of BIC IIR mechanisms, the revenue-optimal as well as welfare-optimal mechanisms are both so-called "all-pay" auctions [15, 19]. In all-pay auctions agents pay the mechanism a certain (distribution dependent) function of their value regardless of the allocation that the mechanism makes. The optimality of all-pay auctions follows by noting that any allocation rule that admits some BIC budget-feasible payment function can be implemented with an all-pay payment rule with worst-case payments that are no larger than those in any other truthful payment rule and are therefore budget-feasible. Unfortunately all-pay auctions have many undesirable properties. In many settings it is simply not feasible to force the agents to pay upfront without knowing the outcome of the mechanism. Moreover all-pay auctions may admit many Bayes-Nash equilibria (BNE), truth-telling being merely one of them. Then the fact that a certain objective is achieved when all the agents report their true types does not necessarily imply that the objective will be achieved in practice if a different BNE gets played out. Therefore, in a departure from previous work, we choose to enforce ex-post individual rationality.

Our results and techniques.

We begin our investigation with the revenue objective and give an exact characterization of the optimal mechanism for a single agent with a public budget. While in the absence of budgets the optimal mechanism is a fixed sale price and therefore deterministic, with budgets the optimal mechanism may need to randomize and offer multiple buying options to the agent. This complicates the design of optimal mechanisms in more general settings involving multiple agents or private budgets. We therefore consider approximations. When budgets are known publicly, we obtain constant factor approximations in nearly all settings where constant

factor approximations are known for unconstrained mechanism design. This includes, for example, all single-parameter settings with a downwards closed feasibility constraint, but also multi-parameter settings with "unit-demand" agents and a matroid feasibility constraint (see, e.g., [6]). Our mechanisms are for the most part direct reductions to unconstrained settings, and are extremely simple.

For private budgets, the problem becomes much harder and we focus on settings with single-dimensional values. We design a novel mechanism based on "lotteries" that obtains a good approximation whenever each agent's value distribution satisfies the monotone hazard rate (MHR) condition (see Section 2 for a definition). Our mechanism's novelty lies in offering each agent a carefully constructed set of different buying options such that the best option for the agent is to either spend his entire budget or a fraction of the monopoly price for that agent. The MHR assumption is a frequently used assumption in mechanism design literature and many natural distributions satisfy it. In fact the mechanism obtains a good approximation more generally under mild technical conditions on the values and budgets. We believe that our techniques should extend to provide good approximations for arbitrary distributions.

Next we examine the welfare objective. While for revenue, the budget of an agent is a natural upper bound on the contribution of that agent to the revenue and allows us to "cap" values at the budget, for welfare this doesn't work. In fact, a mechanism can generate a non-trivial guarantee on welfare even when budgets are 0. Consider a setting with two unit demand buyers and two items. Consider the following mechanism: the mechanism asks each agent to give a preference list of the items. If the top choices of the buyers are different, then each buyer gets allocated his top choice and welfare is maximized. Otherwise, the mechanism ignores the preferences of the buyers and computes the allocation that maximizes the social welfare ex-ante. Note that this mechanism is truthful. When agents' values for the items are i.i.d., the obtained social welfare from this example is at least $3/4$ of the maximum social welfare we can obtain with no budget constraint. On the other hand, a mechanism that "ignores" values above the budget (i.e. does not distinguish between them in the allocation function) cannot obtain an approximation better than $1/2$. The gap between the two mechanisms increases as the number of agents grows.

We again focus on single-parameter settings and public budgets, but with arbitrary downwards closed feasibility constraints. For these settings, we show a tradeoff between an approximation on budget and an approximation on welfare: for any ϵ , we can get a $1/\epsilon$ approximation to welfare with respect to the welfare that an optimal mechanism can get when budgets are scaled down by a factor of $1 - \epsilon$. This mechanism has an extremely simple form: it replaces every value larger than its budget by its expectation conditioned on being larger than the budget, and runs the VCG mechanism on these modified values. Moreover, if we are willing to sacrifice EPIR in favor of the less restrictive IIR, we can convert this mechanism into a 4 -approximate IIR mechanism (with no approximation on budgets).

Finally, if the value distributions satisfy the MHR condition, we achieve a $2(1 + \epsilon)$ -approximation to welfare via an EPIR mechanism by reducing budget-feasible welfare maximization to budget-feasible revenue maximization.

One nice property of our reductions from budget feasible mechanism design to unconstrained mechanism design is that they are for the most part oblivious to the feasibility constraint imposed on the mechanism. They therefore work for a broad range of feasibility constraints and add minimal complexity to the mechanism design problem.

Related work.

Several works in economics have studied characterizations of optimal BIC IIR budget-feasible mechanisms (e.g., [19, 14, 9, 15]). However, these works are generally weak in the kinds of settings they consider (typically just single-item auctions) and the kinds of value distributions they allow¹. Laffont and Robert [14] considered single item settings where bidders have a private value and public common budget. Che and Gale [9] considered the setting with a single item and a single buyer, but allowed both the value and the budget to be private. Pai and Vohra [19] gave a more general result in which they designed an optimal auction for a single item and multiple buyers with private i.i.d. values and private budgets.

Bhattacharya et al. [3] were the first to study settings beyond single-item auctions and focused on revenue maximization. They considered a setting with heterogeneous items and additive values, and exhibited a (large) constant factor DSIC approximation mechanism as well as an all-pay auction which admits truthtelling as a BNE and in that BNE obtains a 4-approximation. However, these results required the value distributions to satisfy the MHR condition. The mechanisms are LP-based. In contrast most of our mechanisms are easy to compute, work for general distributions, enforce EPIR, and achieve small approximation factors.

In prior-free settings few results are known for revenue maximization. Borgs et al. [4] looked at multi unit auctions for homogeneous goods where agents have private values and budgets and considered worst case competitive ratio (see also [1]). They designed a mechanism based on random sampling that maximizes revenue when the number of bidders is large.

Social welfare maximization has also been considered under budget constraints. Maskin [15] considered the setting of a single item and multiple buyers with public budgets. He defined and showed how to compute the constrained efficient mechanism, the truthful feasible mechanism under budget constraints that maximizes the expected social welfare (however, the result holds only for some distribution functions [19]). In prior-free settings for multi unit homogeneous items, Nisan et al. [11] studied Pareto efficient DSIC mechanisms with budget constraints. They showed that if the budgets are private there is no Pareto optimal incentive compatible mechanism; for public budgets they showed that there exists a unique mechanism based on the *clinching auction*. Chen et al. [10] considered a setting with multiple goods and unit demand buyers and showed how to compute competitive prices that enforce truthfulness under budget constraints if such prices exist. Finally, the work of Alaei et al. [2] stands out in their study of “soft” budget constraints, where buyers pay an increasing interest rate for payments made above their budgets. They showed how to exactly compute the smallest competitive prices in this setting that result in an incentive compatible mechanism with an outcome in the core.

2. NOTATION AND DEFINITIONS

In this work, we consider instances of the Bayesian mechanism design problem where agents have single- and multi-dimensional types; instances are of the form $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$.

In the single-dimensional case, $\mathbf{F} = \prod_i F_i$ is a product distribution; each agent i has a single value $v_i \sim F_i$ for receiving service (and derives value 0 if not served), and an upper limit B_i on how much he or she can pay for service; and \mathcal{S} is a feasibility constraint specifying which sets of agents may be simultaneously served.

In the multi-dimensional case, the seller offers a number of ser-

vices indexed by j to agents, and $\mathbf{F} = \prod_{i,j} F_{ij}$ is again a product distribution; each agent i has a value for receiving service j of $v_{ij} \sim F_{ij}$ and is interested in receiving at most one service; and the agent has a budget limit of B_i . Here, \mathcal{S} is a feasibility constraint over pairs (i, j) .

We also consider settings with private budgets; in that case, we replace \mathbf{B} with a distribution $\mathbf{G} = \prod_i G_i$ and agent i has a budget $B_i \sim G_i$.

We focus on incentive compatible (IC) and individually rational (IR) mechanisms, and further distinguish between Bayesian IC and dominant strategy IC, and interim IR and ex post IR. See [17] for definitions of these concepts.

Let M be a mechanism for the instance \mathcal{I} . We shall denote its expected allocation to each agent by the vector $\mathbf{x}(\mathbf{v}, \mathbf{B})$, and agents’ expected payments by $\mathbf{p}(\mathbf{v}, \mathbf{B})$ (we omit the second parameter when \mathbf{B} is fixed). Then the expected revenue of M is $\mathcal{R}^M = \mathbb{E}_{\mathbf{v}, \mathbf{B}}[\mathbf{p}(\mathbf{v}, \mathbf{B}) \cdot \mathbf{x}(\mathbf{v}, \mathbf{B})]$, and its expected social welfare is $\mathbb{E}_{\mathbf{v}, \mathbf{B}}[\mathbf{v} \cdot \mathbf{x}(\mathbf{v}, \mathbf{B})]$

Given that an agent i has a value v for receiving service, and a budget constraint B , his or her utility from receiving the service with probability x at a price p is $u(v, B) = x \cdot v - p$ if $p \leq B$, and $u(v, B) = -\infty$ otherwise. A mechanism is budget feasible if it never requests an agent to make a payment above his budget.

Virtual values and the monotone hazard rate condition.

In the absence of budget constraints, for revenue maximization Myerson in his seminal work [16] gives a characterization of the optimal mechanism as a “virtual value maximizer”. Specifically, given any distribution function F with density f , Myerson defines a *virtual value* function as follows:

$$\phi(v) = v - \frac{1 - F(v)}{f(v)}$$

We use the following characterization by Myerson of the expected revenue of BIC mechanisms in terms of their virtual surplus.

LEMMA 1. *Consider any BIC mechanism with allocation function \mathbf{x} for a single-parameter problem $\mathcal{I} = (\mathbf{F}, \mathcal{S})$. Then the expected revenue of the mechanism is exactly $\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\sum_i x_i(\mathbf{v})\phi_i(\mathbf{v})]$.*

A distribution is said to be *regular* if $\phi(v)$ is a non-decreasing function of v . When value distributions are regular, a mechanism that allocates to the feasible set that maximizes the total virtual value is BIC and optimal. For a single agent, this mechanism allocates to the agent as long as his value is above the threshold $\phi^{-1}(0)$; we call this threshold the *monopoly price* corresponding to the value distribution.

When value distributions are not regular, Myerson gives an “ironing” procedure that converts a virtual value function into an ironed virtual value function, $\bar{\phi}$, such that maximizing ironed virtual surplus results in a BIC optimal mechanism. We omit the details of the ironing.

Some of our results require a stronger condition on distributions called the monotone hazard rate condition, a common assumption in mechanism design literature. This condition is satisfied by many common distributions such as the uniform, Gaussian, exponential, and power law distributions.

DEFINITION 1. *A distribution F with density f is said to have a **monotone hazard rate** if the function $h(v) = f(v)/(1 - F(v))$ is non-decreasing in v . Distributions satisfying MHR are regular.*

¹E.g., [19] and [15] make the assumption that value distributions have a monotone hazard rate as well as a nondecreasing density function, conditions that few distributions satisfy in combination.

3. MAXIMIZING REVENUE

We first consider the revenue objective, and begin by characterizing the optimal budget feasible mechanism for a single agent setting. The characterization relies on describing the mechanism as a collection of so-called lotteries or randomized pricings. We then consider settings with public budgets. Our general approach towards budget-constrained mechanism design in these settings is to approximate the optimal revenue in two parts: the contribution to optimal revenue by agents whose budget is binding (i.e. their budget is less than their value), and the contribution by agents whose budget is not binding (i.e. their budget is above their value). We present different mechanisms for approximating these two benchmarks. We demonstrate this approach first in the simple setting of single-parameter agents with public budgets and an arbitrary downwards closed feasibility constraint. Then in subsequent sections we extend the approach to settings involving more complicated incentive constraints—multi-dimensional values and private budgets. In private budget settings, instead of asking agents to reveal budgets directly, our mechanism once again relies on collections of lotteries to motivate agents to pay a good fraction of their budgets when their values are high enough.

Single agent settings with public budgets.

Before presenting our general approach, we first consider the most basic version of this problem—namely a setting with one single-parameter agent and a public budget constraint. Even this simple setting, however, reveals the challenges budget constraints introduce to the problem of mechanism design. Without the budget constraint, the optimal mechanism is to offer the item at a fixed price. With budgets, however, the following example shows that a single fixed price can be a factor of 2 from optimal. After the example we proceed to characterize the optimal mechanism.

EXAMPLE 1. Fix $n > 1$. Consider an agent whose value for receiving an item is $v = 1$ with probability $1 - 1/n$, and is $v = n^2$ with probability $1/n$. Let the agent have a budget of $B = n$. Any single fixed price that respects the budget in this setting receives a revenue of at most 1.

We now describe the optimal mechanism. The mechanism offers two options to the agent: either buy the item at price n , or receive the item with a probability of $n/(n + 1)$ at a price of $n/(n + 1)$. This generates an expected revenue of $2n/(n + 1) = 2 - o(1)$.

The optimal mechanism in the above example is what we call a *lottery menu* mechanism. A lottery is a pair (x, p) and offers to the agent at a price p a probability x of winning. A lottery menu is a collection of lotteries that an agent is free to choose from in order to maximize his expected utility. We will now show that for any single agent setting with a public budget, the optimal mechanism is a lottery menu mechanism with at most two options.

Consider a setting \mathcal{I} with a single agent with private value $v \sim F$ and a public budget B . Let ϕ be the virtual value function corresponding to F . For ease of exposition, throughout the following discussion we will assume that F is regular and ϕ is non-decreasing; when F is non-regular, we can merely replace ϕ by $\bar{\phi}$, the ironed virtual value, in the following discussion.

We first note that if $B \geq \phi^{-1}(0)$ then the unconstrained optimal mechanism is already budget feasible. Therefore, for the rest of this section we assume that $B < \phi^{-1}(0)$. Following Lemma 1, our

goal is to solve the following optimization problem.

$$\begin{aligned} \max_x \int x(v)\phi(v)f(v)dv & \quad \text{subject to} \\ \int (x_{\max} - x(v))dv \leq B \cdot x_{\max}, & \quad \text{and,} \\ x(v) \text{ is a non-decreasing function.} & \end{aligned}$$

Here $x_{\max} \leq 1$ is the probability of allocation at the upper end of the support of the value distribution. The first constraint encodes the budget constraint. In particular, the left hand side of the inequality is the expected payment made by the agent at his highest value; the right side is an upper bound on the expected payment under EPIR because the agent can pay a maximum of B when he gets allocated, and 0 otherwise.

Let x^* be the optimal solution to the above optimization problem. We make the following observations (each follows by observing that a function x can be modified in a natural fashion to satisfy them, while maintaining the constraints and improving the objective value). In the following, we denote the inverse virtual value of 0 as $v^* = \phi^{-1}(0)$ to simplify notation.

CLAIM 1. Without loss of generality, we may assume $x_{\max}^* = 1$.

CLAIM 2. Without loss of generality, we may assume that for all $v \geq v^*$, $x^*(v) = 1$.

Following these claims, our optimization problem changes to the following (the monotonicity constraint on x is implicit).

$$\begin{aligned} \max_x \int x(v)\phi(v)f(v)dv & \quad \text{subject to} \\ \int (1 - x(v))dv \leq B & \\ x(v) = 1 \quad \forall v \geq v^* & \end{aligned}$$

This can be simplified to:

$$\begin{aligned} \min_x \int_0^{v^*} x(v)(-\phi(v))f(v)dv & \quad \text{subject to} \\ \int_0^{v^*} (1 - x(v))dv = B & \end{aligned}$$

Note that we replace the inequality in the budget constraint with an equality. This is because if the constraint is not tight, we can feasibly reduce $x(v)$ and thereby reduce the objective function value. For the sake of brevity, we define $B' = v^* - B$, and $g(v) = -\phi(v)f(v)$. The budget constraint then changes to $\int_0^{v^*} x(v)dv = B'$. Note that $v^* \geq B' \geq 0$, and g is nonnegative on $[0, v^*]$. Finally, we define the set of allocations

$$\mathcal{A} = \left\{ \begin{array}{l} \text{increasing } x : [0, v^*] \rightarrow [0, 1] \\ \text{such that } \int_0^{v^*} x(v)dv = B' \end{array} \right\}$$

Then, we can express our objective as

$$\min_{x \in \mathcal{A}} \int_0^{v^*} x(v)g(v)dv.$$

If g is non-increasing on $[0, v^*]$, then we immediately have that the optimal solution is to set $x(v) = 1$ if $v \geq v^* - B'$ ($= B$) and 0 otherwise.

If g is not non-increasing, we “iron” the function g to produce a non-increasing function \hat{g} with the property that any non-decreasing function x that is constant over intervals where \hat{g} is constant has the same integral with respect to \hat{g} as with respect to g . Let $\tilde{\mathcal{A}}$ be the

subset of \mathcal{A} containing all functions x that are constant over intervals where \hat{g} is constant. We obtain the following lemma. (The details of the ironing procedure and the proof of the following lemma can be found in the full version of the paper [7].)

LEMMA 2. *For all $x \in \mathcal{A}$, there exists a $\tilde{x} \in \tilde{\mathcal{A}}$, such that $\int_0^{v^*} x(v)g(v)dv \geq \int_0^{v^*} \tilde{x}(v)g(v)dv$.*

The lemma lets us confine our optimization to the set $\tilde{\mathcal{A}}$:

$$\min_{x \in \tilde{\mathcal{A}}} \int_0^{v^*} x(v)g(v)dv = \min_{x \in \tilde{\mathcal{A}}} \int_0^{v^*} x(v)g(v)dv$$

Finally, we define \mathcal{A}^* to be a subset of $\tilde{\mathcal{A}}$ in which functions x take on at most three different values – 0, 1, and an intermediate value. The final part of our proof is to show that the optimal solution lies in this set.

THEOREM 3. *For any single agent setting $\mathcal{I} = (F, B)$, there is an optimal mechanism with allocation rule in the set \mathcal{A}^* .*

PROOF. Recall that the optimal solution x^* lies in the set $\tilde{\mathcal{A}}$. Suppose for contradiction that this function takes on two different intermediate values, $x^*(v_1) = y$ and $x^*(v_2) = z$, between 0 and 1 with $y < z$. Then, since \hat{g} is non-increasing and x^* is non-decreasing, we must have $\hat{g}(v_1) > \hat{g}(v_2)$. Now we can improve our objective function value by increasing x^* between v_2 and the value at which it becomes 1, and decreasing x^* between the value at which it becomes strictly positive and v_1 , while maintaining the budget constraint. This contradicts the optimality of x^* . \square

Single parameter setting with public budgets.

We now consider single parameter settings with multiple agents. Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ be an instance of single-parameter budget-constrained revenue maximization. Define the truncated distributions \hat{F}_i as follows.

$$\hat{F}_i(v) = \begin{cases} F_i(v) & \text{if } v < B_i; \text{ and} \\ 1 & \text{if } v \geq B_i. \end{cases} \quad (1)$$

Let $\hat{\mathcal{I}} = (\hat{\mathbf{F}}, \mathcal{S})$ be the modified setting where we replace \mathbf{F} with $\hat{\mathbf{F}}$ — note that for each i , the support of \hat{F}_i ends at or before B_i , and so we may remove the budgets since they place no constraint on the instance $\hat{\mathcal{I}}$.

A mechanism for $\hat{\mathcal{I}}$ naturally extends to \mathcal{I} , while satisfying budget feasibility and obtaining the same revenue. Our general technique will be to relate the revenue of a mechanism for \mathcal{I} to that of a mechanism for $\hat{\mathcal{I}}$. In general, the latter can be quite small, and so we introduce the following quantity to bound this loss. Define the set \mathcal{B} as

$$\mathcal{B} = \operatorname{argmax}_{S \subseteq \mathcal{S}} \{ \sum_{i \in S} B_i \mid \forall i \in S, v_i \geq B_i \}. \quad (2)$$

Our basic approach is to design a BIC mechanism \hat{M} for the setting $\hat{\mathcal{I}}$ based on the original mechanism M such that we have

$$\mathcal{R}^M \leq \mathcal{R}^{\hat{M}} + \mathbb{E} \left[\sum_{i \in \mathcal{B}} B_i \right]. \quad (3)$$

Then, the first term on the right is bounded above by the revenue of the optimal mechanism for $\hat{\mathcal{I}}$. We further demonstrate in each case that we can bound the expectation $\mathbb{E} \left[\sum_{i \in \mathcal{B}} B_i \right]$ by another mechanism for $\hat{\mathcal{I}}$.

We define the mechanism \hat{M} in terms of its expected allocation and payment. Let $\mathbf{x}(\mathbf{v})$ and $\mathbf{p}(\mathbf{v})$ be the expected allocations and

expected payments for M , respectively. Define the expected allocation and expected payment rules for \hat{M} as follows. For each agent i in the setting $\hat{\mathcal{I}}$ with valuation \hat{v}_i , draw a corresponding v_i consistent with $\hat{v}_i = \min(v_i, B_i)$; in this case that simply means $v_i = \hat{v}_i$ if $\hat{v}_i < B_i$, and $v_i \sim F_i(v \mid v \geq B_i)$ otherwise. Then \hat{M} 's expected allocation and payment are given by

$$\begin{aligned} \hat{x}_i(\hat{\mathbf{v}}) &= x_i(\mathbf{v}_{-i}, \hat{v}_i); \text{ and} \\ \hat{p}_i(\hat{\mathbf{v}}) &= p_i(\mathbf{v}_{-i}, \hat{v}_i), \end{aligned}$$

respectively.

LEMMA 4. *\hat{M} is a feasible BIC mechanism for $\hat{\mathcal{I}}$.*

PROOF. We first note that from the point of view of a single agent i , the expected allocation and price function of \hat{M} behave as though other agents' values are the same as before. Therefore, the expected allocation is still an increasing function of value and the payments satisfy BIC. We will now argue that the expected allocation function can be implemented in a way that the resulting outcome is a randomization over feasible outcomes. To do so, we first compute $x_i(\mathbf{v})$, as well as $\hat{x}_i(\hat{\mathbf{v}})$ for all i . Starting with the allocation returned by $\mathbf{x}(\mathbf{v})$, for every agent i in this allocation, with probability $\hat{x}_i(\hat{\mathbf{v}})/x_i(\mathbf{v})$, we serve this agent, and with the remaining probability we remove her from the allocated set. Since \mathcal{S} is a downward closed feasibility constraint, feasibility is maintained, and we achieve the target allocation probabilities. We remark here that our goal is to merely exhibit that \hat{M} is feasible and not to actually compute it. \square

We now prove the bound (3) on \mathcal{R}^M .

LEMMA 5. *Given any mechanism M for $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, where \mathcal{S} is downward-closed, if we define the mechanism \hat{M} for $\hat{\mathcal{I}}$ as above, then (3) holds.*

PROOF. In order to prove the statement, we couple the values \mathbf{v} that \hat{M} draws for each $\hat{\mathbf{v}}$ with the \mathbf{v} in the other expectations. So fix some corresponding pair of value vectors \mathbf{v} and $\hat{\mathbf{v}}$; consider the contribution of each agent i to the revenue of M . Split the agents into two sets L and H , defined by

$$L = \{i \mid v_i \leq B_i\}; \text{ and} \quad H = \{i \mid v_i \geq B_i\}.$$

Recall that for all $i \in L$, we have that $v_i = \hat{v}_i$, and so $\hat{p}_i(\hat{\mathbf{v}}) = p_i(\mathbf{v}_{-i}, \hat{v}_i) = p_i(\mathbf{v})$. Furthermore, since M faces the downward-closed feasibility constraint \mathcal{S} , any subset of H that M serves is one of the sets \mathcal{B} maximizes over. Since M can never charge any agent more than their budget, we can see that

$$\begin{aligned} \mathcal{R}^M(\mathbf{v}) &= \sum_{i \in L} \mathcal{R}_i^M(\mathbf{v}) + \sum_{i \in H} \mathcal{R}_i^M(\mathbf{v}) \\ &\leq \sum_{i \in L} \mathcal{R}_i^{\hat{M}}(\hat{\mathbf{v}}) + \sum_{i \in \mathcal{B}} B_i \\ &\leq \mathcal{R}^{\hat{M}}(\hat{\mathbf{v}}) + \sum_{i \in \mathcal{B}} B_i. \end{aligned}$$

Taking expectations on both sides (according to the previously mentioned coupling) proves the claim. \square

Note that $\mathcal{R}^{\hat{M}}$ can be easily achieved by simply running the (unconstrained) revenue-optimal mechanism over $\hat{\mathcal{I}}$. It remains to be shown that we can, in fact, upper bound $\mathbb{E}[\sum_{i \in \mathcal{B}} B_i]$ also by the revenue of the same (unconstrained) revenue-optimal mechanism over $\hat{\mathcal{I}}$.

LEMMA 6. *There exists a mechanism $M_{\mathcal{B}}$ for the setting $\widehat{\mathcal{I}}$ such that $\mathbb{E}_{v \sim \mathbf{F}} [\sum_{i \in \mathcal{B}} B_i] \leq \mathcal{R}^{M_{\mathcal{B}}}$.*

PROOF. We define the mechanism $M_{\mathcal{B}}$ as implementing the allocation rule \mathcal{B} . Note that membership of i in \mathcal{B} is monotone in v_i , and that the truthful payment for $i \in \mathcal{B}$ is precisely B_i , since this is the minimum value required for allocation. Thus, we can immediately see that

$$\mathcal{R}^{M_{\mathcal{B}}} = \mathbb{E}_{v \sim \widehat{\mathbf{F}}} \left[\sum_{i \in \mathcal{B}} B_i \right] = \mathbb{E}_{v \sim \mathbf{F}} \left[\sum_{i \in \mathcal{B}} B_i \right],$$

as desired. \square

By combining the results of Lemmas 5 and 6, we get the following theorem.

THEOREM 7. *Given a single parameter setting $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, the optimal mechanism \mathcal{M} for the modified setting $\widehat{\mathcal{I}} = (\widehat{\mathbf{F}}, \mathcal{S})$ gives a 2-approximation to the optimal revenue for \mathcal{I} .*

Multi-parameter setting with public budgets.

We next consider settings where a seller offers multiple kinds of service and agents have different preferences over them. Agents are unit-demand and want any one of the services; the seller faces a general downward closed feasibility constraint. As before, we use the tuple $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ to denote an instance of this problem; throughout, i indexes agents and j indexes services. Let \mathcal{S} be a downward-closed feasibility constraint over (i, j) pairs, and furthermore assume each agent i has a budget B_i .

Ideally, we would like to follow the same approach as in the previous section. We use the same basic benchmark, defining $\widehat{\mathbf{F}}$ and \mathcal{B} analogously to (1) and (2) for the instance \mathcal{I} . We can't apply the same reduction from M to \widehat{M} directly, however, because truncating each of a multi-parameter agent's values to their budget affects the agents' preferences across different items, a concern we did not have before.

Instead, we make use of a reduction of Chawla et al. [6] from multi-parameter Bayesian MD to single-parameter Bayesian MD to first bring the problem into a single parameter domain and then apply the approach from the previous section. We describe the reduction and our mechanism in the full version of the paper [7], obtaining the following theorem.

THEOREM 8. *Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ be an instance with multi-parameter, unit-demand agents and \mathcal{S} being a matroid or simpler feasibility constraint. Then, there exists a polynomial time computable mechanism for $\widehat{\mathcal{I}}$ that is budget feasible and DSIC for \mathcal{I} and obtains a constant fraction of the revenue of the optimal budget-feasible mechanism for \mathcal{I} .*

Private budgets.

We next consider settings where budgets are part of agents' private types, but where the mechanism designer knows the distributions from which budgets are drawn. We assume that values and budgets are drawn from independent distributions. We focus on settings where agents' values are single-dimensional.

Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ denote an instance of this setting. We follow a similar approach as for public budgets. Switching from public to private budgets, however, adds new complexity; in particular it becomes tricky to achieve our benchmark $\mathbb{E}[\sum_{i \in \mathcal{B}} B_i]$. In this section, we present an approximately optimal mechanism for the case

when each distribution in \mathbf{F} satisfies the MHR condition (see Definition 1 in Section 2). Our analysis uses the MHR condition in a very mild way and in fact holds for any setting where for every agent the probability that his value exceeds his monopoly price is lower bounded by a constant. Even when this condition is not satisfied, we can obtain a good approximation through a slight modification of our mechanism under a technical condition on values and budgets. These details are discussed in the full version of the paper [7]. We believe that the general idea behind our mechanism can be extended to obtain good approximations for arbitrary distributions.

We focus on settings $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ where \mathcal{S} is a matroid set system, and each distribution in \mathbf{F} satisfies the MHR condition. We begin with some definitions. Given a pair of value and budget vectors, we consider the "extractable value" of an agent i to be $\min(v_i, B_i)$; we modify our definition of \mathcal{B} to reflect this:

$$\mathcal{B} = \operatorname{argmax}_{S \in \mathcal{S}} \sum_{i \in S} \min(v_i, B_i).$$

Similarly to the public budgets case, our approach is to split the revenue of an arbitrary mechanism into two terms, which (loosely speaking) look like revenue in a truncated value setting, and the sum of the budgets in \mathcal{B} ; we then demonstrate a lottery menu mechanism whose revenue upper bounds both of these terms.

Our proposed mechanism (which we denote $M^{\mathcal{L}}$) is based on lottery menus $\mathcal{L}(p, \bar{p})$ parameterized by a minimum price p and a maximum price \bar{p} . (Recall the discussion of lotteries at the beginning of this section.) There are two cases. If $p \geq \bar{p}/3$, then $\mathcal{L}(p, \bar{p})$ contains the single fixed price of p ; otherwise, it contains, for all $2p/\bar{p} \leq \alpha \leq 2/3$, a lottery that with probability $(1/3 + \alpha)$ allows the agent to purchase service at a price of $\alpha\bar{p}/2$.

Note that the probability of allocation in the above lottery system rises faster as a function of α than the price of the lottery. Effectively this ensures that the agent is willing to buy the most expensive lottery that he can afford. So, in particular, if all lotteries bring positive utility then the agent spends his entire budget, the maximum amount that any mechanism can hope to achieve from the agent. This powerful idea is what enables our approximation.

Note the following properties of the lottery menu $\mathcal{L}(p, \bar{p})$.

LEMMA 9. *When an agent with $\min(v, B) \geq p$ is offered the menu $\mathcal{L}(p, \bar{p})$, he purchases an option yielding expected revenue at least $p/3$.*

PROOF. Note that in either case, the lottery system $\mathcal{L}(p, \bar{p})$ always contains an option with price exactly p , and that this is the lowest priced option. Furthermore, for $\min(v, B) \geq p$, this option always gives non-negative utility. Thus our claim follows immediately from the fact that every lottery assigns the item with probability at least $1/3$. \square

LEMMA 10. *When an agent with $v \geq \bar{p}$ is offered the menu $\mathcal{L}(p, \bar{p})$, he purchases an option yielding expected revenue at least $\min(\bar{p}, B)/3$.*

PROOF. In the first case (when $\mathcal{L}(p, \bar{p})$ contains the single price of p), this follows trivially; in the other case, consider the utility of an agent with value v when purchasing the lottery with parameter α , which we denote $u_{\alpha}(v)$. We can see that

$$u_{\alpha}(v) = (1/3 + \alpha)(v - \alpha\bar{p}/2), \text{ and so} \\ \frac{\partial u_{\alpha}(v)}{\partial \alpha} = v - (\alpha + 1/6)\bar{p} \geq 0$$

by our assumption. So we can see that an agent will purchase the lottery with the highest α value they can afford; since the lottery for

$\alpha = 2/3$ assigns service with probability 1 at a price of $\bar{p}/3$, and every lottery provides service with probability at least one third, we can see that an agent will purchase a lottery yielding revenue at least $\min(\bar{p}, B)/3$. \square

Our mechanism $M^{\mathcal{L}}$ serves the set \mathcal{B} . For each i , let T_i be the threshold corresponding to inclusion in \mathcal{B} , i.e. $T_i = \min\{v' : i \in \mathcal{B} \text{ for } ((\mathbf{v}_{-i}, v'), (\mathbf{B}_{-i}, v'))\}$. Our mechanism offers agent i the lottery system $\mathcal{L}(T_i, \phi_i^{-1}(0))$, where $\phi_i^{-1}(0)$ is the monopoly price for i .

In order to relate the revenue of a mechanism M for the setting \mathcal{I} to that of our proposed mechanism $M^{\mathcal{L}}$, we break the revenue of M into two parts – that derived from agents in \mathcal{B} , and that derived from agents not in \mathcal{B} ; we denote these quantities by $\mathcal{R}^{M \cap \mathcal{B}}$ and $\mathcal{R}^{M \setminus \mathcal{B}}$, respectively. We bound the two terms separately.

LEMMA 11. $\mathcal{R}^{M \setminus \mathcal{B}} \leq 3\mathcal{R}^{M^{\mathcal{L}}}$

PROOF. Let S denote the set of agents served by M . Note that by our definition of \mathcal{B} , it will always be a maximal independent set in \mathcal{S} ; as such, for every pair of vectors (\mathbf{v}, \mathbf{B}) , we can get a 1-1 function $g : S \setminus \mathcal{B} \rightarrow \mathcal{B}$ such that for all $i \in S \setminus \mathcal{B}$, $\mathcal{B} \setminus \{g(i)\} \cup \{i\} \in \mathcal{S}$. Note that by EPIR, we must have that the revenue M derives from each agent i is no more than $\min(v_i, B_i)$, and by the definition of T_i , we get that

$$\begin{aligned} \mathcal{R}^{M \setminus \mathcal{B}}(\mathbf{v}, \mathbf{B}) &\leq \sum_{i \in S \setminus \mathcal{B}} \min(v_i, B_i) \\ &\leq \sum_{i \in S \setminus \mathcal{B}} T_{g(i)} \leq \sum_{i \in \mathcal{B}} T_i \\ &\leq 3\mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}, \mathbf{B}) \end{aligned}$$

Here the second inequality follows from noting that in order to be in the set \mathcal{B} , the agent $g(i)$ must have an extractable value no smaller than that of i . The last inequality follows from applying Lemma 9 to $\mathcal{L}(T_i, \phi_i^{-1}(0))$. Taking expectation over (\mathbf{v}, \mathbf{B}) completes the proof. \square

In order to prove our next revenue bound, we need the following property of MHR variables (Lemma 4.1 in [12]):

LEMMA 12. *For v distributed according to some F satisfying the MHR, the probability that the value v exceeds the monopoly price $\phi^{-1}(0)$ is at least $1/e$.*

LEMMA 13. $\mathcal{R}^{M \cap \mathcal{B}} \leq 3e\mathcal{R}^{M^{\mathcal{L}}}$

PROOF. Consider some agent i , and fix $(\mathbf{v}_{-i}, \mathbf{B}_{-i})$; note that this fixes T_i as well. Fix B_i . Recall that an agent i contributes to $\mathcal{R}_i^{M \cap \mathcal{B}}$ only when $\min(v_i, B_i) \geq T_i$. We split the analysis into two cases; in each case we consider the optimal revenue a mechanism could derive from agents with $v_i \geq T_i$ if allowed to ignore the budgets constraints.

- Case 1: $T_i \geq \phi_i^{-1}(0)$. In this case, the maximum revenue that can be obtained from agent i conditioned on $v_i \geq T_i$ and ignoring feasibility constraints is T_i and can be obtained by offering a fixed price of T_i . Our mechanism on the other hand offers a single option of buying service at a fixed price of T_i . Therefore,

$$\begin{aligned} \mathbb{E}_{v_i} \left[\mathcal{R}_i^{M \cap \mathcal{B}}(\mathbf{v}, \mathbf{B}) \mid v_i \geq T_i \right] \\ \leq \mathbb{E}_{v_i} \left[\mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}, \mathbf{B}) \mid v_i \geq T_i \right] \end{aligned}$$

- Case 2: $T_i < \phi_i^{-1}(0)$. The maximum revenue that can be obtained from agent i conditioned on $v_i \geq T_i$ and ignoring feasibility constraints is at most $\phi_i^{-1}(0)$ and can be obtained by offering a fixed price of $\phi_i^{-1}(0)$. Considering the budget constraint we conclude that $\mathbb{E}_{v_i} [\mathcal{R}_i^{M \cap \mathcal{B}}(\mathbf{v}, \mathbf{B}) \mid v_i \geq T_i] \leq \min(\phi_i^{-1}(0), B_i)$. On the other hand, applying Lemma 10 to $\mathcal{L}(T_i, \phi_i^{-1}(0))$, we get that for $v_i \geq \phi_i^{-1}(0)$, $\mathcal{R}_i^{M^{\mathcal{L}}} \geq \min(\phi_i^{-1}(0), B_i)/3$. Lemma 12 implies that the event $v_i \geq \phi_i^{-1}(0)$ happens with probability at least $1/e$. So we get

$$\begin{aligned} \mathbb{E}_{v_i} [\mathcal{R}_i^{M \cap \mathcal{B}}(\mathbf{v}, \mathbf{B}) \mid v_i \geq T_i] \\ \leq (1/3e) \mathbb{E}_{v_i} [\mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}, \mathbf{B}) \mid v_i \geq T_i] \end{aligned}$$

We have $\mathbb{E}_{v_i} [\mathcal{R}_i^{M \cap \mathcal{B}}(\mathbf{v}, \mathbf{B})] \leq 3e \mathbb{E}_{v_i} [\mathcal{R}^{M^{\mathcal{L}}}(\mathbf{v}, \mathbf{B})]$ in either case; taking expectations over $(\mathbf{v}_{-i}, \mathbf{B})$, and summing over i yields our claim. \square

Combining the above two lemmas immediately gives the following theorem.

THEOREM 14. *For any setting $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ where \mathcal{S} is a matroid constraint and each distribution in \mathbf{F} satisfies MHR, the mechanism $M^{\mathcal{L}}$ is a $3(1+e)$ approximation to the optimal revenue.*

In fact, we can use a similar lottery pricing technique to get a constant approximation even in the absence of the MHR assumption; however, we still need a technical assumption relating the distributions of agents' values and budgets. This once again ensures that there is a good probability that agents' values are high enough for the lottery system to extract a constant fraction of their budget. We state the theorem here, but defer the proof of this result to the full version of the paper [7].

THEOREM 15. *Suppose that every agent's median value is no smaller than a constant fraction of her maximum budget. Then we can construct a budget-feasible mechanism that is DSIC with respect to both values and budgets, and obtains a constant fraction of the revenue of the optimal such mechanism.*

4. MAXIMIZING WELFARE

In this section we focus on the welfare objective. In particular, the seller's goal is to maximize the total value of the allocation in expectation. Once again we assume that budgets are known publicly.

We first note that we cannot use the approach of the previous section as a roadmap. Even with public budgets, truncating values to the corresponding budgets does not work for the social welfare objective. In particular, the following example shows it is possible for a budget feasible mechanism to distinguish between values above the budget without exceeding the budget in payments.

EXAMPLE 2. *Consider an n agent single-item auction, where agents have i. i. d. values for the item. Each agent has a budget of 1. Each agent's value is 1 with probability $1 - 1/n$ and n with probability $1/n$. Then a mechanism that simply truncates values to budgets cannot distinguish between the agents and gets a social welfare of at most 2. On the other hand, consider a mechanism that orders agents in an arbitrary order and offers two options to each agent in turn while the item is unallocated: getting the item for free with probability $1/n$ and nothing otherwise, or purchasing the item at a price of 1. Then, an agent picks the first option if and only if her value is below $n/(n-1)$, and otherwise picks the second option.*

In particular, an agent with value n always picks the second option, and an agent with value 1 always picks the first option. For large n , with probability approaching $1 - 1/e$ at least one agent has value n , and with probability at least $1/e$ the item is unsold before the first agent with value n is made an offer. The mechanism's expected welfare is therefore at least $1/e(1 - 1/e)n = \Omega(n)$.

Note that the precise choice of budgets in the above example was critical: if budgets were any lower, the proposed mechanism would have been infeasible; and if they were any higher, truncation would have still allowed for distinguishing between agents with low and high values. This suggests considering bicriteria approximations where we compete against an optimal mechanism that faces smaller budgets. We first demonstrate a mechanism achieving an approximation of this sort; we then show that our mechanism also gives a good approximation if we relax the EPIR constraint to an IIR constraint, instead of relaxing budgets. Of course, our ultimate goal is to provide a good approximation for the social welfare objective via an EPIR budget feasible mechanism. While we are unable to do so in general, the final section presents a constant factor approximation for settings where the distributions F_i for every agent i satisfy the MHR condition (Definition 1 in Section 2).

A bicriteria approximation.

Consider a setting \mathcal{I} with budgets \mathbf{B} . Let OPT' denote the EPIR mechanism that is welfare-maximizing and feasible for budgets $(1 - \epsilon)\mathbf{B}$ (i.e. where each budget is scaled down by a factor of $1 - \epsilon$). We claim that we can approximate the welfare of this mechanism while maintaining budget feasibility with respect to the original budgets \mathbf{B} .

THEOREM 16. *For a given instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, let \mathcal{I}' be the instance $(\mathbf{F}, \mathcal{S}, (1 - \epsilon)\mathbf{B})$ where each agent's budget is scaled down by a factor of $1 - \epsilon$. Let OPT' denote the welfare optimal budget feasible mechanism for \mathcal{I}' . Then, there exists an easy to compute ex post IR mechanism (namely, the VCG mechanism over a modified instance) that is budget feasible for \mathcal{I} and obtains at least an ϵ fraction of the social welfare of OPT' .*

PROOF. We first use OPT' to construct a new mechanism M . M proceeds as follows. It elicits values from agents. For all agents i with $v_i \geq B_i$, it resamples agent i 's value from the distribution F_i restricted to the set $[B_i, \infty)$. Other values are left unchanged. It then runs the mechanism OPT' on the resampled values. It is easy to see that M is budget feasible for \mathcal{I} .

We claim that the social welfare of M is at least ϵ times the social welfare of OPT' . To prove the claim, consider a single agent i , and let x_i^1 denote the probability of allocation for this agent in OPT' when her value is B_i , and x_i^2 denote the probability of allocation for this agent in OPT' when her value is v_i^{\max} (the agent's maximum possible value). Note that the expected payment that the agent makes at v_i^{\max} is at least $(x_i^2 - x_i^1)B_i$ plus the payment she makes at B_i . Then, EPIR implies that $(x_i^2 - x_i^1)B_i$ is at most the budget $(1 - \epsilon)B_i$ times x_i^2 . This implies $x_i^2 < x_i^1/\epsilon$. Now, noting that the value distributions for agents are unaltered by resampling, it holds that for $v_i \geq B_i$, the probability of allocation for agent i at v_i under M is equal to the expected probability of allocation for the agent under OPT' conditioned on the agent's value being in the range $[B_i, \infty)$. Since the probability of allocation under OPT' for this range is always between x_i^1 and x_i^2 , the expected probability of allocation is at least $x_i^1 \geq \epsilon x_i^2$. So compared to those under OPT' , the probabilities of allocation under M are at most a factor of ϵ smaller. Therefore, the expected social welfare of M is also at most a factor of ϵ smaller than that of OPT' .

Our goal will then be to approximate the social welfare of M . Note that for every agent i , M treats values above B_i identically. We can therefore consider the following optimization problem: for an instance \mathcal{I} , construct a DSIC EPIR mechanism that maximizes social welfare subject to the additional constraint that for every agent i the mechanism's (distribution over) allocation should be identical across value vectors that differ only in agent i 's value and where agent i 's value is $\geq B_i$. For any such mechanism, agent i 's expected contribution to social welfare from value vectors with $v_i \geq B_i$ conditioned on being allocated is \bar{v}_i , where $\bar{v}_i = E[v_i | v_i \geq B_i]$. Therefore, the following mechanism maximizes welfare over the above class of mechanisms: for every agent i with $v_i \geq B_i$, replace v_i by \bar{v}_i ; other values remain unmodified; run the VCG mechanism over the modified value vector; charge every agent the minimum of the payment returned by the VCG mechanism and their budget. It is easy to verify that this mechanism is DSIC, ex post IR, budget feasible for the original budgets B_i , and obtains expected social welfare at least that of M . Therefore, it satisfies the claim in the theorem. \square

An interim IR mechanism.

Next we note that it is in fact easy to remove the approximation on budget in the above theorem if we are willing to give up on EPIR. In particular, consider an optimal mechanism OPT on the instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$. Then the above theorem implies the existence of a mechanism V that is budget feasible for $\mathcal{I}' = (\mathbf{F}, \mathcal{S}, 2\mathbf{B})$ and obtains half the welfare of OPT (taking $\epsilon = 1/2$). Now consider the mechanism V' described as follows. V' simulates V on the given value vector. Then for every agent i it charges i half the payment charged by V and with probability $1/2$ makes an allocation to i if V makes an allocation to i . Agent i 's expected utility from any strategy under V' is exactly half her expected utility from the same strategy under V . Therefore, V' is DSIC. Moreover, it is budget feasible for the original budgets \mathbf{B} since it always charges half the payments in V . Its expected social welfare is exactly half that of V . We therefore get the following theorem.

THEOREM 17. *For a given instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, let OPT denote the welfare optimal EPIR budget feasible mechanism for \mathcal{I} . Then, there exists an easy to compute IIR mechanism that is budget feasible for \mathcal{I} and obtains at least a quarter of the social welfare of OPT .*

An ex-post IR mechanism for MHR distributions.

As previously remarked, our ultimate goal is to provide a good approximation for the social welfare objective via an EPIR budget feasible mechanism. We now show that under an MHR condition on distributions, we can achieve precisely this goal. In particular, we present a constant factor approximation for settings where the distributions F_i for every agent i satisfy the MHR condition (Definition 1 in Section 2).

Under the MHR condition, we can exhibit a close relationship between the welfare and revenue of any mechanism. Using this relationship along with results from the previous section, we can come up with a budget feasible approximately-revenue-maximizing mechanism that also provides an approximation to social welfare. The MHR condition is quite crucial to our approach. In fact our solution consists of two mechanisms, one of which charges no payments, and the other of which truncates values to their corresponding budgets – approaches that don't work for the example we considered above.

Let $v_i^* = \phi_i^{-1}(0)$ denote the monopoly price for the distribution F_i . We then get the following bound on social welfare, which we are able to approximate.

LEMMA 18. *For any instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, if all the distributions F_i satisfy the MHR condition, then for any non-decreasing allocation function $x(v)$, we have that*

$$\begin{aligned} & \int_{\mathbf{v}} \left(\sum_i v_i x_i(\mathbf{v}) \right) d\mathbf{F}(\mathbf{v}) \\ & \leq \int_{\mathbf{v}} \left(\sum_i (\phi_i(v_i) + 2v_i^*) x_i(\mathbf{v}) \right) d\mathbf{F}(\mathbf{v}). \end{aligned}$$

In order to prove the Lemma, we require some new definitions and claims. Consider a single agent with MHR distribution F , virtual value function ϕ , and monopoly price v^* . Let ϕ^+ and ϕ^- be the positive and negative portions of ϕ respectively; i.e. for all v , $\phi^+(v), \phi^-(v) \geq 0$ and $\phi(v) = \phi^+(v) - \phi^-(v)$. We can then claim the following (the first is a restatement of Lemma 3.1 in [13]).

LEMMA 19. *For a distribution F satisfying the MHR, all values v satisfy $v \leq v^* + \phi^+(v) = v^* + \phi(v) + \phi^-(v)$.*

LEMMA 20. *For any monotone allocation function $x(\cdot)$,*

$$\int \phi^-(v) x(v) dF(v) \leq \int v^* x(v) dF(v).$$

PROOF. We begin by recalling that by Lemma 1, the expected revenue of any BIC mechanism is equal to its expected virtual surplus. Now consider a mechanism for a single agent with value distribution F that always serves the agent. Clearly the revenue of this mechanism is 0. So we get

$$\int_v (\phi^+(v) - \phi^-(v)) dF(v) = 0,$$

which implies that

$$\int_v \phi^+(v) dF(v) = \int_v \phi^-(v) dF(v).$$

Second, the revenue from offering the agent the monopoly price v^* is precisely $v^*(1 - F(v^*))$. Therefore,

$$\begin{aligned} \int_{v^*}^{\infty} (\phi^+(v) - \phi^-(v)) dF(v) &= \int \phi^+(v) dF(v) \\ &= v^* \int_{v^*}^{\infty} dF(v), \end{aligned}$$

where the first equality follows from regularity of F .

Note that the regularity of F implies that ϕ^+ and ϕ^- are identically 0 below and above v^* , respectively. Now, from the above two equalities, we can see that if x is monotone non-decreasing, then,

$$\begin{aligned} & \int \phi^-(v) x(v) dF(v) \leq \int \phi^-(v) x(v^*) dF(v) \\ & = x(v^*) \int \phi^+(v) dF(v) = x(v^*) v^* \int_{v^*}^{\infty} dF(v) \\ & \leq \int v^* x(v) dF(v); \end{aligned}$$

the claim follows. \square

The proof of Lemma 18 follows immediately by combining Lemmas 19 and 20. The Lemma gives us the following approximation.

THEOREM 21. *Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ be an instance where all distributions F_i satisfy the MHR condition. Then, one of the following mechanisms obtains a $2(1 + e)$ -approximation to the social welfare of a welfare-optimal budget-feasible mechanism for \mathcal{I} . Both of these mechanisms are DSIC, EPIR and budget feasible.*

- Mechanism 1: Always allocate to the set S_1^* and charge zero payments, where $S_1^* = \operatorname{argmax}_{S \in \mathcal{S}} \sum_{i \in S} v_i^*$.
- Mechanism 2: Elicit values from agents; for all i with $v_i > B_i$, replace v_i by B_i ; run Myerson's mechanism on the resulting instance.

PROOF. We begin by noting that an immediate consequence of Lemma 12 is that for a distribution F satisfying the MHR, $E_F[v] \geq v^*/e$.

Now, consider some budget feasible mechanism M for the instance \mathcal{I} . Then by Theorem 7, the optimal mechanism for the truncated distributions (1) obtains revenue, and therefore also social welfare, no less than a $1/2$ fraction of the expected revenue $\int \sum_i \phi_i(v_i) x_i(\mathbf{v}) d\mathbf{F}(\mathbf{v})$.

Moreover, $\int_{\mathbf{v}} \sum_i v_i^* x_i(\mathbf{v}) d\mathbf{F}(\mathbf{v})$ is upper bounded by $\sum_{i \in S^*} v_i^*$, where $S^* = \operatorname{argmax}_{S \in \mathcal{S}} \sum_{i \in S} v_i^*$. Then a mechanism which always allocates to the set S^* and charges no payments is budget feasible, DSIC, and obtains welfare $\sum_{i \in S^*} E_{F_i}[v_i] \geq 1/e \sum_{i \in S^*} v_i^*$, where the inequality follows from Lemma 12. The original claim follows. \square

Acknowledgments

We thank Saeed Alaei for many helpful discussions.

5. REFERENCES

- [1] Z. Abrams. Revenue maximization when bidders have budgets. In *Proceedings of the seventeenth annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '06, pages 1074–1082, New York, NY, USA, 2006. ACM.
- [2] S. Alaei, K. Jain, and A. Malekian. Competitive equilibrium for unit-demand buyers with non quasi-linear utilities. In *CORR*, 2010.
- [3] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget constrained auctions with heterogeneous items. In *STOC*, pages 379–388, 2010.
- [4] C. Borgs, J. T. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *ACM Conference on Electronic Commerce*, pages 44–51, 2005.
- [5] S. Brusco and G. Lopomo. Simultaneous ascending auctions with complementarities and known budget constraints. *Economic Theory*, 38(1):105–124, January 2009.
- [6] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010.
- [7] S. Chawla, D. L. Malec, and A. Malekian. Bayesian mechanism design for budget-constrained agents. *CoRR*, abs/1103.6280, 2011.
- [8] Y.-K. Che and I. Gale. Standard auctions with financially constrained bidders. *Review of Economic Studies*, 65(1):1–21, January 1998.
- [9] Y.-K. Che and I. Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory*, 92(2):198–233, June 2000.
- [10] N. Chen, X. Deng, and A. Ghosh. Competitive equilibria in matching markets with budgets. *CoRR*, 2010.

- [11] S. Dobzinski, N. Nisan, and R. Lavi. Multi-unit auctions with budget limits. In *Proc. of the 49th Annual Symposium on Foundations of Computer Science (FOCS)*, 2008.
- [12] J. Hartline, V. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *Proceeding of the 17th international conference on World Wide Web, WWW '08*, pages 189–198, New York, NY, USA, 2008. ACM.
- [13] J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*, 2009.
- [14] J.-J. Laffont and J. Robert. Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186, 1996.
- [15] E. S. Maskin. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44(4-6):667–681, 2000.
- [16] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [17] N. Nisan. Introduction to mechanism design. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9, pages 209–242. Cambridge Press, 2007.
- [18] N. Nisan, J. Bayer, D. Chandra, T. Franji, R. Gardner, Y. Matias, N. Rhodes, M. Seltzer, D. Tom, H. Varian, and D. Zigmond. Google’s auction for tv ads. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part II, ICALP '09*, pages 309–327, Berlin, Heidelberg, 2009. Springer-Verlag.
- [19] M. M. Pai and R. Vohra. Optimal auctions with financially constrained bidders. Discussion papers, Northwestern University, Aug. 2008.