

Bayesian Incentive Compatibility via Fractional Assignments

Xiaohui Bei *

Zhiyi Huang †

Abstract

Very recently, Hartline and Lucier [HL10] studied single-parameter mechanism design problems in the Bayesian setting. They proposed a black-box reduction that converted Bayesian approximation algorithms into Bayesian incentive compatible (BIC) mechanisms and preserved the expected social welfare. It remains an open question if one can find similar reduction in the more important multi-parameter setting. In this paper, we give positive answer to this question when the prior distribution has finite and small support. We propose a black-box reduction for designing BIC multi-parameter mechanisms. The reduction converts any algorithm into an ϵ -BIC mechanism with only marginal loss in social welfare. As a result, for combinatorial auctions with sub-additive agents we get an ϵ -BIC mechanism that achieves constant approximation.

*Institute for Theoretical Computer Science, Tsinghua University. Email: bxh08@mails.tsinghua.edu.cn. Supported in part by the National Natural Science Foundation of China Grant 60553001, the National Basic Research Program of China Grant 2007CB807900, 2007CB807901.

†Computer and Information Science, University of Pennsylvania. Email: hzyi@cis.upenn.edu. This material is based upon work supported by the National Science Foundation under Grant No. 0716172

1 Introduction

In this paper, we consider the problem of designing computationally efficient and truthful mechanism for multi-parameter mechanism design problems in the Bayesian setting.

Suppose a major Internet search service provider wants to sell multiple advertisement slots to a number of companies. From the history of previous transactions, we can estimate a prior distribution of each company's valuation of the advertisement slots. What mechanism shall the search service provider use to obtain good social welfare, or good revenue?

This is a typical multi-parameter mechanism design problem. In general, we consider the scenario in which a principal wants to sell a number of different services to multiple heterogeneous strategic agents subject to some feasibility constraints (e.g. total cost of providing these services must not exceed the budget), so that some desired objective (e.g. social welfare, revenue, residual surplus) is achieved. If we interpret this as simply a combinatorial optimization problem, then there exists approximation algorithms for many of these problems. And the approximation ratios of many of these algorithms are tight subject to certain computational complexity assumptions. However, if we want to design protocols of allocations and setting prices in order to achieve the desired objective in the equilibrium strategic behavior of the agents, we usually have much worse approximation ratio. Therefore, it is natural to ask the following question:

Can we convert any algorithm into a truthful mechanism while preserving the performance, say, social welfare?

Unfortunately, from previous work we learn that this is impossible for some problems. Papadimitriou et al. [PSS08] showed the first significant gap between the performance of deterministic algorithms and deterministic truthful mechanisms via the Combinatorial Public Project problem.

Bayesian setting. The standard game theoretic model for incomplete information is the Bayesian setting, in which the agent valuations are drawn from a publicly known distribution. The standard solution concept in this setting is *Bayesian-Nash Equilibrium*. In a Bayesian-Nash equilibrium, each player maximizes its expected payoff by following the strategy profile given the prior distribution of the agent valuations.

In this paper, we will consider multi-parameter welfare-preserving algorithm/mechanism reductions in the Bayesian setting, and weaken truthfulness constraint from Incentive Compatibility (IC) to Bayesian Incentive Compatibility (BIC), which means truth telling is the equilibrium strategy over random choice of the mechanism as well as the random realization of the other agent valuations. In many real world applications such as online auctions, AdWords auctions, spectrum auctions etc., the availability of data of past transactions make it possible to obtain good estimation of the prior distribution of the agent valuations. Thus, revisiting the algorithm/mechanism reduction problem in the Bayesian setting is of both theoretical and practical importance.

Hartline and Lucier [HL10] studied this problem in the single-parameter setting. They showed a brilliant black-box reduction from any approximation algorithm to BIC mechanism that preserves the performance with respect to social welfare maximization. In this paper, we prove that similar reduction also exists for the realm of multi-parameter mechanism design for social welfare! Moreover, we can also obtain BIC mechanism for revenue or residual surplus via some variants of our black-box reduction.

Our results and technique. Our main result is a black-box reduction that converts algorithms into BIC mechanisms with essentially the same social welfare for arbitrary multi-parameter mech-

anism design problem in the Bayesian setting. More concretely, given an algorithm \mathcal{A} that provides $SW^{\mathcal{A}}$ social welfare, the reduction provides a mechanism that gives $SW^{\mathcal{A}} - \epsilon$ social welfare and is ϵ -BIC. The running time is polynomial in the input size and $1/\epsilon$. This resolves an open problem in [HL10]. The key idea is to decouple the reported valuations and the input valuations for the algorithm \mathcal{A} . When the reported valuations are v_1, v_2, \dots, v_n , we will manipulate the valuations via some carefully designed intermediate algorithms $\mathcal{B}_1, \dots, \mathcal{B}_n$, and use allocation $\mathcal{A}(\mathcal{B}_1(v_1), \dots, \mathcal{B}_n(v_n))$. We prove that there exist intermediate algorithms $\mathcal{B}_1, \dots, \mathcal{B}_n$ so that there are prices that achieve BIC. Under certain conditions, the marginal loss factor in social welfare can be made multiplicative.

As an application of this reduction, we get a $(\frac{1}{2} - \epsilon)$ -approximate and ϵv_{max} -BIC mechanism for social welfare maximization in combinatorial auctions with sub-additive agents. For the more restricted case of fractionally sub-additive agents, we obtain $(1 - \frac{1}{e} - \epsilon)$ -approximate mechanism.

Related work. The problem of maximizing social welfare against strategic agents is one of the oldest and most famous problems in the area of mechanism design. It has been extensively studied by the economists in both Bayesian and prior-free setting without considering computational power constraint. The celebrated VCG mechanism [Cla71, Gro73, Vic61] which guarantees optimal social welfare and incentive compatibility is one of the most exciting results in this domain. However, implementing the VCG mechanism is NP-hard in general. This is one of the reasons that VCG mechanism is rarely used in practice despite of its lovely theoretical features.

In the past decade, computer scientists introduced many novel techniques in the prior-free setting to design computationally efficient mechanisms that provide incentive compatibility and/or good approximation to optimal social welfare for various families of valuation functions.

On the one hand, Dobzinski, Nisan and Schapira [DNS06] proposed poly-time mechanisms which achieved $\Omega(1/\sqrt{n})$ -approximation for general agents and $\Omega(1/\log^2 n)$ -approximation for sub-modular agents. Dobzinski [Dob07] later proposed a truthful mechanism which achieved an improved $\tilde{\Omega}(1/\log n)$ -approximation for a strictly broader class of sub-additive agents.

On the other hand, if we focus on the algorithmic problem of maximizing social welfare assuming all valuations are truthfully revealed, then the algorithm by Dobzinski, Nisan and Schapira [DNS05] gave $\Omega(1/\sqrt{n})$ -approximation for general case and $\Omega(1/\log n)$ -approximation for sub-additive agents. The latter approximation ratio is later improved to $\frac{1}{2}$ for sub-additive agents [Fei06] and $(1 - \frac{1}{e})$ for the more restricted class of fractionally sub-additive agents [Dob07, FV06].

The above results suggest that there exists a gap between the performance of the best poly-time algorithms and that of the best poly-time and incentive compatible mechanism. As an effort to study the relation between designing algorithms and designing truthful mechanisms with good approximation ratio, Lavi and Swamy [LS05] proposed a meta-mechanism that converted strong rounding algorithms for the standard LP of social welfare maximization into IC mechanisms. However, their approach required the rounding algorithm to work for arbitrary valuation functions. This requirement prevents their technique to get good approximation beyond cases of general valuations and additive valuations (via a different linear program). But the more interesting classes of valuations (e.g. sub-additive valuations and sub-modular valuations) lies between these two extremes. Another notable attempt on reducing IC mechanism design to algorithm design is the very recent work by Dughmi and Roughgarden [DR10]. They proved that for any packing problem that admitted an FPTAS, there was an IC mechanism that was also an FPTAS.

Most of the previous effort from computer scientists has focused on the prior-free setting. Until very recently, there has been a few work that brought more and more Bayesian analysis into the field of algorithmic mechanism design. Hartline and Lucier [HL10] gave a black-box reduction that

converted any Bayesian approximation algorithm into a Bayesian incentive compatible mechanism that preserved social welfare in the single parameter domain. Bhattacharya et al. [BGGM10] studied the revenue maximization problem for auctioning heterogeneous items when the valuations of the agents were additive. Their result gave constant approximation in the Bayesian setting even when the agents had public known budget constraints. Chawla et al. [CHMS10] considered the revenue maximization problem in the multi-dimensional multi-unit auctions. They introduced mechanism that gave constant approximation in various settings via sequential posted pricing.

Finally, in concurrent and independent work, Hartline et al. [HKM11] study the relation of algorithm and mechanism in Bayesian setting and propose similar reduction. In the discrete support setting that is considered in this paper, they use essentially the same reduction. However, their work achieves perfectly BIC instead of ϵ -BIC. They also extend the reduction to the more general continuous support setting.

2 Preliminaries

2.1 Notations.

We use $\{x_i\}_{1 \leq i \leq n}$ to denote an array of size n . We also use the natural extension of this notation for multi-dimensional arrays. We will use bold font \mathbf{x} to denote a vector (x_1, \dots, x_n) . We let $\Delta(S)$ denote the set of distributions over the elements in a set S . For a random variable x , we let $\mathbf{E}[x]$ denote its expectation and let $\sigma[x]$ denote its standard deviation. We use subscripts to represent the random choices over which we consider the expectation and variance. For instance, $\mathbf{E}_{y \sim F}[x]$ is the expectation of x when y is drawn from distribution F . We sometimes use $\mathbf{E}_y[x]$ for short when the distribution F is clear from the context.

2.2 Model and definitions.

In this section, we will formally introduce the model in this paper. We study the general multi-parameter mechanism design problems. In a multi-parameter mechanism design problem, a principal wants to sell a set of services to multiple heterogeneous agents in order to optimize the desired objective (e.g. social welfare, revenue, residual surplus, etc.). A Bayesian multi-parameter mechanism design problem with n agents is defined by a tuple $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$.

- $\mathbf{I} = (I_1, \dots, I_n)$: The set of services that the principal wants to sell to the agents.

Since we can impose arbitrary feasibility constraints on the allocations, we assume without loss of generality that the services are partitioned into n disjoint sets I_1, \dots, I_n so that the services in I_i only aim for agent i , and each agent i is interested in any one of these services.

- $\mathbf{J} \subseteq I_1 \times \dots \times I_n$: The set of feasible allocations.
- $\mathbf{V} = V_1 \times \dots \times V_n$: The space of agent valuations.

We let $V_i \subseteq \mathbb{R}^{I_i}$ denote the set of possible valuations of agent i . We let v_{max} denote the maximal valuation, that is, $v_{max} = \max_{i, v \in V_i, S \in I_i} v(S)$

- $\mathbf{F} = F_1 \times F_2 \times \dots \times F_n$: The joint prior distribution of the agent valuations.

We assume the prior distribution is a product distribution. We let $F_i \in \Delta(V_i)$ denote the prior distribution of the valuation of agent i . In this paper, we only consider distributions with finite and polynomially large support. We will assume without loss of generality that

the support of each distribution F_i is $\{v_i^1, \dots, v_i^\ell\}$. Suppose $v_i \sim F_i$, We will let $f_i(t)$ denote the probability that $v_i = v_i^t$.

For example, in the *combinatorial auction* problem with n agents and m items, we let $[m] = \{1, 2, \dots, m\}$ denote the set of items. The set of services for each agent i is the set of all subsets of items, that is, $I_i = 2^{[m]}$, $1 \leq i \leq n$. The set of feasible allocations is

$$\mathbf{J} = \{(S_1, \dots, S_n) : S_i \in I_i, S_i \cap S_j = \emptyset\} .$$

The set of valuations, V_i , is the set of mappings from subset of items I_i to \mathbb{R}_+ that are monotone ($v_i(S) \leq v_i(T)$ for $S \subseteq T$) and normalized ($v_i(\emptyset) = 0$). We usually assume that the valuations in $\bigcup_i V_i$ satisfies certain properties, e.g. sub-additivity, sub-modularity, etc.

Algorithm. An algorithm for a multi-parameter mechanism design problem $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$ is a protocol (may or may not be randomized) that takes a realization of agent valuations $\mathbf{v} \in \mathbf{V}$ as input, and outputs a feasible allocation $\mathbf{S} \in \mathbf{J}$.

Mechanism. A mechanism is an interactive protocol (may or may not be randomized) between the principal and the agents so that the principal can retrieve information from the agents (presumably via their bids), and determine an allocation of services $\mathbf{S} \in \mathbf{J}$ and a collection of prices $\mathbf{p} = (p_1, \dots, p_n)$. The extra challenge for mechanism design, compared to algorithm design, is to retrieve genuine valuations from the agents and handle their strategic behavior.

For each $1 \leq i \leq n$, we will assume the prior distribution F_i is public known. But the actual realization $v_i \sim F_i$ is *private* information of agent i . Each agent i aims to maximize the quasi-linear utility $v_i(S_i) - p_i$, where S_i is the service it gets and p_i is the price. Thus, the agents may not reveal their genuine valuations if manipulating their bids strategically can increase their utility.

Objectives. We will consider three different objectives: social welfare, revenue, and residual surplus. The expected *social welfare* of a mechanism \mathcal{M} is

$$SW^{\mathcal{M}} = \mathbf{E}_{\mathbf{v} \sim \mathbf{F}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(\mathbf{v})} \left[\sum_{i=1}^n v_i(S_i) \right] .$$

Similarly, we will let $SW^{\mathcal{A}}$ denote the expected social welfare of an algorithm \mathcal{A} .

Definition 1. An algorithm \mathcal{A} is α -approximate in social welfare for a multi-parameter mechanism design problem $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$, if $SW^{\mathcal{A}} \geq \alpha \text{OPT}$.

The expected *revenue* of a mechanism is

$$R^{\mathcal{M}} = \mathbf{E}_{\mathbf{v} \sim \mathbf{F}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(\mathbf{v})} \left[\sum_{i=1}^n p_i \right] .$$

The last objective, *residual surplus*, was recently proposed by Hartline and Roughgarden [HR08] as an alternative objective in the flavour of social welfare. In the residual surplus maximization problem, the principal aims to maximize the sum of the agents' utilities instead of the sum of their valuations. The expected *residual surplus* is

$$RS^{\mathcal{M}} = \mathbf{E}_{\mathbf{v} \sim \mathbf{F}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(\mathbf{v})} \left[\sum_{i=1}^n (v_i(S_i) - p_i) \right] .$$

We will let OPT denote the optimal social welfare, that is, $\text{OPT} = \max_{\mathcal{M}} SW^{\mathcal{M}}$. Since both revenue and residual surplus are upper-bounded by social welfare. We will use OPT as our benchmark for all three objectives.

Solution concepts. Ideally, a mechanism shall provide incentive for the agents to reveal their valuations truthfully. In this section, we will formalize this requirement by introducing the game-theoretical solution concepts that we use in this paper.

Definition 2. A mechanism \mathcal{M} is Bayesian incentive compatible (BIC) if for each agent i and any two valuations $v_i, \tilde{v}_i \in V_i$, we have

$$\mathbf{E}_{\mathbf{v}_{-i}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(v_i, \mathbf{v}_{-i})} [v_i(S_i) - p_i] \geq \mathbf{E}_{\mathbf{v}_{-i}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(\tilde{v}_i, \mathbf{v}_{-i})} [v_i(S_i) - p_i] .$$

Definition 3. A mechanism \mathcal{M} is ϵ -Bayesian Incentive Compatible (ϵ -BIC) if for any agent i and any two valuations $v_i, \tilde{v}_i \in V_i$,

$$\mathbf{E}_{\mathbf{v}_{-i}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(v_i, \mathbf{v}_{-i})} [v_i(S_i) - p_i] \geq \mathbf{E}_{\mathbf{v}_{-i}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(\tilde{v}_i, \mathbf{v}_{-i})} [v_i(S_i) - p_i] - \epsilon .$$

Other than the above constraints of incentive compatibility, the mechanism shall also guarantee that the agents always get non-negative utility. Otherwise, the agents may choose not to participate in the mechanism. This is known as the *individual rationality* constraint.

Definition 4. A mechanism \mathcal{M} is individually rational (IR) if for any realization \mathbf{v} of agent valuations, and any allocation \mathbf{S} and prices \mathbf{p} by the mechanism, we always have that $v_i(S_i) - p_i \geq 0$ for all agent i ,

3 Characterization of BIC mechanisms

In this section, we will introduce a non-trivial characterization of BIC multi-parameter mechanisms via a novel connection between BIC mechanisms and envy-free prices. This characterization inspires our reduction in the next section.

3.1 Fractional assignment problem.

We will first introduce the fractional assignment problems, which will play a critical role in the results of this paper, and a useful lemma about envy-free prices in fractional assignment problems.

In order to distinguish the notations for fractional assignment problems and those for the mechanism design problems, we will use superscripts instead of subscripts to specify different entries of a vector for the fractional assignment problems. For instance, we will use x^s to denote the s^{th} entry of a vector \mathbf{x} .

Let us consider a market with ℓ buyers and m infinitely divisible products. Each buyer s has a non-negative demand α^s , which denotes the maximal amount of products the buyer will buy. Each product t has a non-negative supply β^t , which denotes the available amount of this product in the market. For each buyer s and each product t , we let w^{st} denote the non-negative value of buyer s of product t .

The goal is to set prices for the products and to assign the products to the buyers subject to the demand and supply constraints. Thus, a solution (\mathbf{x}, \mathbf{p}) to the fractional assignment problem

consists of a collection of prices $\mathbf{p} = (p^1, \dots, p^\ell)$ and a feasible allocation $\mathbf{x} = \{x^{st}\}_{1 \leq s \leq \ell, 1 \leq t \leq m}$ in the polytope:

$$\left\{ \mathbf{x} : \forall s, \sum_{t=1}^m x^{st} \leq \alpha^s; \forall t, \sum_{s=1}^{\ell} x^{st} \leq \beta^t; \mathbf{x} \geq 0 \right\},$$

where x^{st} denotes the amount of product t that is assigned to buyer s .

Definition 5. A solution (\mathbf{x}, \mathbf{p}) is envy-free if for any $x^{st} > 0$, then t is a product that maximizes the quasi-linear utility of agent s , and the utility for agent s is non-negative. That is,

$$\forall s, t : x^{st} > 0 \Rightarrow w^{st} - p^t = \max_k \{w^{sk} - p^k\} \geq 0. \quad (1)$$

Definition 6. An allocation \mathbf{x} is market-clearing if all demand constraints and supply constraints hold with equality. That is,

$$\forall 1 \leq s \leq \ell : \sum_{t=1}^m x^{st} = \alpha^s, \quad \forall 1 \leq t \leq m : \sum_{s=1}^{\ell} x^{st} = \beta^t.$$

The social welfare maximization problem for a fractional assignment problem is characterized by the following linear program (P) and its dual (D).

<p>(P) Maximize $\sum_{s=1}^{\ell} \sum_{t=1}^m x^{st} w^{st}$ s.t.</p> <p>$\sum_{t=1}^m x^{st} \leq \alpha^s$ $\forall s$</p> <p>$\sum_{s=1}^{\ell} x^{st} \leq \beta^t$ $\forall t$</p> <p>$x^{st} \geq 0$ $\forall s, t$</p>	<p>(D) Minimize $\sum_{s=1}^{\ell} \alpha^s u^s + \sum_{t=1}^m \beta^t p^t$ s.t.</p> <p>$u^s + p^t \geq w^{st}$ $\forall s, t$</p> <p>$u^s \geq 0$ $\forall s$</p> <p>$p^t \geq 0$ $\forall t$</p>
---	--

We will introduce two useful lemmas about the connection between envy-free prices and social welfare maximization for fractional assignment problems. These lemmas were known in different forms in the economics literature [GS99].

Lemma 1. If there exist envy-free prices \mathbf{p} for a market-clearing allocation \mathbf{x} , then \mathbf{x} maximizes the social welfare, that is, $\mathbf{x} \cdot \mathbf{w} = \max_{\mathbf{z}} \mathbf{z} \cdot \mathbf{w}$.

Proof. Suppose there exist envy-free prices \mathbf{p} for an allocation \mathbf{x} . Let $u^s = \max_t \{w^{st} - p^t\}$. We have that $u^s + p^t \geq w^{st}$ for all s, t . So (\mathbf{u}, \mathbf{p}) is a feasible solution for the dual LP.

Moreover, by definition of envy-freeness, we have

$$\forall s, t : x^{st} > 0 \Rightarrow u^s = w^{st} - p^t.$$

Therefore, we get that

$$\sum_{s=1}^{\ell} \sum_{t=1}^m x^{st} w^{st} = \sum_{s=1}^{\ell} \sum_{t=1}^m x^{st} (u^s + p^t) = \sum_{s=1}^{\ell} \alpha^s u^s + \sum_{t=1}^m \beta^t p^t.$$

The last equality holds because \mathbf{x} is market clearing. Notice that \mathbf{x} is a feasible solution to the primal LP. By duality theorem, we get that the allocation \mathbf{x} maximizes the social welfare for the fractional assignment problem. \square

Lemma 2. *If an allocation \mathbf{x} maximizes the social welfare, then there exist envy-free prices \mathbf{p} for the fractional assignment problem.*

Proof. Suppose the allocation \mathbf{x} maximizes the social welfare. Let (\mathbf{u}, \mathbf{p}) be an optimal solution to the dual LP. By complementary slackness we get that $x^{st} > 0$ only if the corresponding dual constraint is tight, that is, $u^s + p^t = w^{st}$. Therefore, $x^{st} > 0$ implies that $w^{st} - p^t = u^s \geq w^{sk} - p^k$ for all k . Thus \mathbf{p} is a collection of envy-free prices for the allocation \mathbf{x} in this fractional assignment problem. \square

Note that the above proof also gives a poly-time algorithm for finding the welfare maximizing allocation \mathbf{x} and the corresponding envy-free prices \mathbf{p} by solving the primal and dual LPs. Moreover, we also get that the envy-free prices \mathbf{p} satisfy a weak uniqueness in the sense that it must be part of an optimal solution for the dual LP.

Corollary 1. *There exists a poly-time algorithm that computes the welfare-maximizing market-clearing allocation and the envy-free prices.*

3.2 Characterizing BIC via envy-free prices.

We first introduce some notations that will simplify our discussion. Given a mechanism \mathcal{M} for a multi-parameter mechanism design problem $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$, we will consider the expected values and expected prices for each agent when it choose a specific strategy (each strategy corresponds to reporting a specific valuation).

Assuming the other agents report their valuations truthfully, agent i 's expected value of the service it gets, when the genuine valuation is v_i^s and the reported valuation is v_i^t , is

$$w_i^{st} = \mathbf{E}_{\mathbf{v}_{-i}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)] \ .$$

Similarly, we let p_i^t denote the expected price the mechanism would charge to agent i if its reported valuation is v_i^t , that is,

$$p_i^t = \mathbf{E}_{\mathbf{v}_{-i}, (\mathbf{S}, \mathbf{p}) \sim \mathcal{M}(v_i^t, \mathbf{v}_{-i})} [p_i] \ .$$

By the definition of BIC and IR, the mechanism \mathcal{M} is BIC and IR if and only if for any $1 \leq i \leq n$ and $1 \leq s \leq \ell$,

$$w_i^{ss} - p_i^s = \max_t \{w_i^{st} - p_i^t\} \geq 0 \ . \quad (2)$$

The above equation (2) is similar to equation (1) in the definition of envy-freeness in fractional assignment problem. In fact, the key observation is that the above BIC condition is equivalent to the envy-free condition for a set of properly chosen fractional assignment problems.

Induced assignment problems. For each agent i , we will consider the following *induced assignment problem*. We consider ℓ virtual buyers with demands $f_i(1), \dots, f_i(\ell)$ respectively, and ℓ virtual products with supplies $f_i(1), \dots, f_i(\ell)$ respectively. For each virtual buyer s and each virtual product t , let virtual buyer s has value w_i^{st} on virtual product t . We will refer to this fractional assignment problem the *induced assignment problem* of agent i .

Let us consider the *identity allocation* \mathbf{x}_i defined as follows:

$$x_i^{st} = \begin{cases} f_i(s) & , \text{ if } s = t \ , \\ 0 & , \text{ otherwise.} \end{cases}$$

We can easily verify that a collection of prices $\mathbf{p}_i = (p_i^1, \dots, p_i^\ell)$ satisfies constraint (2) if and only if \mathbf{p}_i satisfies the envy-free condition (1) of the induced assignment problem of agent i with respect to the above identity allocation. Hence, we have the following connection between BIC mechanism and the envy-free prices of the induced assignment problems.

Lemma 3 (Characterization Lemma [Roc87]). *A mechanism \mathcal{M} is BIC if and only if in the induced assignment problem of each agent i the identity allocation $\mathbf{x}_i = \{x_i^{st}\}_{1 \leq s, t \leq \ell}$ maximizes the social welfare, and $\mathbf{p}_i = (p_i^1, \dots, p_i^\ell)$ are chosen to be the corresponding envy-free prices.*

Comparing with Myerson's characterization. Suppose the problem falls into the single-parameter domain. Each valuation v_i^s is represented by a single non-negative real number. With a little abuse of notation, we let v_i^s denote this value. Without loss of generality, we assume that $v_i^1 > \dots > v_i^\ell$. We let y_i^t denote the probability that agent i would be served if the reported value was v_i^t . The values \mathbf{w}_i in the fractional assignment problems of agent i are $w_i^{st} = v_i^s y_i^t$ for $1 \leq s, t \leq \ell$. Myerson's famous characterization [Mye81] of truthfulness in single-parameter domain implied that the mechanism is BIC if and only if $y_i^1 \geq \dots \geq y_i^\ell$. Indeed, due to rearrangement inequality, the identity allocation \mathbf{x}_i maximizes the social welfare if and only if $y_i^1 \geq \dots \geq y_i^\ell$. Thus, the characterization lemma implies Myerson's characterization in the single-parameter domain.

4 Reduction for social welfare

Lemma 3 suggests an interesting connection between BIC and envy-free prices for the induced assignment problems. Hence, given an algorithm \mathcal{A} , we will manipulate the allocation by \mathcal{A} based on this connection in order to make it satisfy the condition in Lemma 3.

4.1 Main ideas.

Let us first briefly convey two key ideas in the construction of the welfare-preserving reduction.

The first idea is to decouple the reported agent valuations and the input agent valuations for algorithm \mathcal{A} . More concretely, we will introduce n intermediate algorithm $\mathcal{B}_1, \dots, \mathcal{B}_n$. Each \mathcal{B}_i will take the reported valuation v_i' as input, then output a valuation \tilde{v}_i that may or may not equals v_i' . Then, we will use algorithm \mathcal{A} to compute the allocation \mathbf{S} for agent valuations $\tilde{v}_1, \dots, \tilde{v}_n$, and allocate services according to \mathbf{S} .

If we revisit the values $\tilde{\mathbf{w}}_i$ in the induced assignment problem of agent i after this manipulation, we will get that for any $1 \leq s, t \leq \ell$,

$$\tilde{w}_i^{st} = \mathbf{E}_{\mathbf{v}_{-i}, \tilde{\mathbf{v}} \sim \mathcal{B}(v_i^t, \mathbf{v}_{-i}), \mathbf{S} \sim \mathcal{A}(\tilde{\mathbf{v}})} [v_i^s(S_i)] \quad .$$

By Lemma 3, we need to choose \mathcal{B}_i 's carefully, so that the identity allocations in the manipulated assignment problems are welfare-maximizing allocations. However, from the above equation we can see that by using \mathcal{B}_i to manipulate the i^{th} valuation, we may change not only the structure of the induced assignment problem of agent i , but the structure of the induced assignment problems of other agents as well. Hence, we need to handle such correlation among the induced assignment problems when we choose intermediate algorithms $\mathcal{B}_1, \dots, \mathcal{B}_n$.

The idea that handles this correlation is to impose an extra constraint on each intermediate algorithm \mathcal{B}_i : if the input valuation v_i' is drawn from the distribution F_i , then the output valuation \tilde{v}_i also follows the same distribution, that is, for all $1 \leq i \leq n$ and $1 \leq t \leq \ell$,

$$\Pr_{v_i' \sim F_i, \tilde{v}_i \sim \mathcal{B}_i(v_i')} [\tilde{v}_i = v_i^t] = f_i(t) \quad . \quad (3)$$




Figure 1: High-level picture of the reduction for social welfare. \mathcal{B}_i 's are intermediate algorithms for manipulating the input of algorithm \mathcal{A} . \tilde{v}_i 's are the reported valuations. v'_i 's are the manipulated input valuations for algorithm \mathcal{A} . S is the final allocation.

With this extra constraint, the values $\tilde{\mathbf{w}}_i$ after the manipulation in the induced assignment problem of agent i becomes

$$\begin{aligned} \tilde{w}_i^{st} &= \mathbf{E}_{\mathbf{v}_{-i} \sim F_{-i}, \tilde{\mathbf{v}} \sim \mathcal{B}(v_i^t, \mathbf{v}_{-i}), \mathbf{S} \sim \mathcal{A}(\tilde{\mathbf{v}})} [v_i^s(S_i)] \\ &= \mathbf{E}_{\tilde{\mathbf{v}}_{-i} \sim F_{-i}, \tilde{v}_i \sim \mathcal{B}_i(v_i^t), \mathbf{S} \sim \mathcal{A}(\tilde{\mathbf{v}})} [v_i^s(S_i)] \\ &= \mathbf{E}_{\mathbf{v}_{-i} \sim F_{-i}, \tilde{v}_i \sim \mathcal{B}_i(v_i^t), \mathbf{S} \sim \mathcal{A}(\tilde{v}_i, \mathbf{v}_{-i})} [v_i^s(S_i)] . \end{aligned}$$

Thus, from the Bayesian viewpoint of agent i , the intermediate algorithms \mathcal{B}_{-i} of other agents are transparent. This property enables us to manipulate the structure of each assignment problem separately.

4.2 Black-box reduction.

Given an algorithm \mathcal{A} , the black-box reduction for social welfare will convert algorithm \mathcal{A} into the following mechanism $\mathcal{M}_{\mathcal{A}}$:

1. For each agent i , $1 \leq i \leq n$ **(Pre-computation)**
 - (a) Estimate the values $\mathbf{w}_i = \{w_i^{st}\}_{1 \leq s, t \leq \ell}$ of the induced assignment problem of i with respect to algorithm \mathcal{A} . Let $\hat{\mathbf{w}}_i = \{\hat{w}_i^{st}\}_{1 \leq s, t \leq \ell}$ denote the estimated values.
 - (b) Find the social welfare maximizing allocation $\mathbf{x}_i = \{x_i^{st}\}_{1 \leq s, t \leq \ell}$ and the corresponding envy-free prices $\mathbf{p}_i = (p_i^1, \dots, p_i^\ell)$ for the induced assignment problem of agent i with estimated values.
2. Manipulate the valuations with intermediate algorithms \mathcal{B}_i , $1 \leq i \leq n$, as follows: **(Decoupling)**

Suppose the reported valuation of agent i is $v'_i = v_i^s$, $1 \leq i \leq n$. Let $\tilde{v}_i = \mathcal{B}_i(v'_i) = v_i^t$ with probability $x_i^{st}/f_i(s)$ for $1 \leq t \leq \ell$.
3. Allocate services according to $\mathcal{A}(\tilde{\mathbf{v}})$. **(Allocation)**
 - (a) Let $\mathbf{S} = (S_1, \dots, S_n)$ denote the allocation by algorithm \mathcal{A} with input $\tilde{\mathbf{v}}$.
 - (b) For each agent i , suppose the reported valuation is $v'_i = v_i^s$ and the manipulated valuation is $\tilde{v}_i = v_i^t$, charge agent i with price $p_i^t v_i^s(S_i) / \hat{w}_i^{st}$.

The following theorem states that this reduction produces BIC while preserving the performance with respect to social welfare.

Theorem 1. Suppose \mathcal{A} is an algorithm for a multi-parameter mechanism design problem $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$.

1. If the estimated values $\hat{\mathbf{w}}_i$ are accurate, then mechanism $\mathcal{M}_{\mathcal{A}}$ is BIC, IR, and guarantees at least $SW^{\mathcal{A}}$ of social welfare.
2. If the estimated values $\hat{\mathbf{w}}_i$ satisfy that for any $1 \leq s, t \leq \ell$, $\hat{w}_i^{st} \in [(1 - \epsilon)w_i^{st}, (1 + \epsilon)w_i^{st}]$, then mechanism $\mathcal{M}_{\mathcal{A}}$ is $4\epsilon v_{\max}$ -BIC, IR, and guarantees at least $(1 - 2\epsilon) \cdot SW^{\mathcal{A}}$ of social welfare.
3. If the estimated values $\hat{\mathbf{w}}_i$ satisfy that for any $1 \leq s, t \leq \ell$, $\hat{w}_i^{st} \in [w_i^{st} - \epsilon, w_i^{st} + \epsilon]$, then mechanism $\mathcal{M}_{\mathcal{A}}$ is 4ϵ -BIC, IR, and guarantees at least $SW^{\mathcal{A}} - 2n\epsilon$ of social welfare.

Let us illustrate the proof of part 1. The proofs of the other two parts are tedious and simple calculations along the same line. We will omit these proofs in this extended abstract.

Proof. We consider the case when the estimated values $\hat{\mathbf{w}}_i$ are accurate, that is, $\hat{w}_i^{st} = w_i^{st}$ for all $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$.

Individual rationality. By our choice of envy-free prices, we have that $p_i^t \leq w_i^{st}$ for all $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$. Thus, we always guarantee

$$p_i^t \frac{v_i^s(S_i)}{w_i^{st}} \leq v_i^s(S_i) .$$

So the mechanism $\mathcal{M}_{\mathcal{A}}$ that we get from the reduction always provides non-negative utilities for the agents. Essentially the same proof also shows IR for part 2 and 3.

Bayesian incentive compatibility. We will first show that the intermediate algorithms in the decoupling step of the reduction satisfy constraint (3). Let \mathbf{x}_i denote the social welfare maximizing allocation that the reduction finds for the induced assignment problem of agent i for $1 \leq i \leq n$. Note that these social welfare maximizing allocations are market-clearing. We have that if the reported valuation v'_i follows the distribution F_i , then the distribution of the manipulated valuation \tilde{v}_i satisfies that

$$\Pr [\tilde{v}_i = v_i^t] = \sum_{s=1}^{\ell} \Pr [v'_i = v_i^s] \Pr [\tilde{v}_i = v_i^t : v'_i = v_i^s] = \sum_{s=1}^{\ell} f_i(s) \frac{x_i^{st}}{f_i(s)} = \sum_{s=1}^{\ell} x_i^{st} = f_i(t) .$$

Indeed, the intermediate algorithms satisfy constraint (3). Thus, for each $1 \leq i \leq n$ the intermediate algorithm \mathcal{B}_i only changes the structure of induced assignment problem of agent i and leaves the induced assignment problems of other agents untouched.

Next, we will verify that in each of the manipulated assignment problem, the identity allocation maximizes the social welfare and the prices are the corresponding envy-free prices.

For each agent i , we let $\tilde{\mathbf{w}}_i = \{\tilde{w}_i^{st}\}_{1 \leq s, t \leq \ell}$ and $\tilde{\mathbf{p}}_i = (\tilde{p}_i^1, \dots, \tilde{p}_i^{\ell})$ denote the values and the expected prices of the virtual products respectively in the manipulated assignment problem of agent i . We have that for any $1 \leq r, s \leq \ell$,

$$\begin{aligned} \tilde{w}_i^{rs} &= \sum_{t=1}^{\ell} \Pr [\tilde{v}_i = v_i^t] \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^r(S_i)] = \sum_{t=1}^{\ell} \frac{x_i^{st}}{f_i(s)} w_i^{rt} , \\ \tilde{p}_i^s &= \sum_{t=1}^{\ell} \Pr [\tilde{v}_i = v_i^t] \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} \left[p_i^t \frac{v_i^s(S_i)}{w_i^{rt}} \right] = \sum_{t=1}^{\ell} \frac{x_i^{st}}{f_i(s)} p_i^t . \end{aligned}$$

Thus, in the manipulated assignment problem of agent i , the utility of the virtual buyer r of the virtual product s , $1 \leq r, s \leq \ell$, is

$$\tilde{w}_i^{rs} - \tilde{p}_i^s = \sum_{t=1}^{\ell} \frac{x_i^{st}}{f_i(s)} (w_i^{rt} - p_i^t) \leq \sum_{t=1}^{\ell} \frac{x_i^{st}}{f_i(s)} \max_k \{w_i^{rk} - p_i^k\} = \max_k \{w_i^{rk} - p_i^k\} .$$

Since \mathbf{p}_i are chosen to be the envy-free prices, we have that $x_i^{rt} > 0$ only if $w_i^{rt} - p_i^t = \max_k \{w_i^{rk} - p_i^k\}$. Hence, when agent i reports its valuation truthfully, that is, $r = s$, the above inequality holds with equality. So the \tilde{p}_i are envy-free prices with respect to the identity allocation $\tilde{\mathbf{x}}_i$ of the manipulated assignment problem of agent i . By Lemma 1 we know the allocation $\tilde{\mathbf{x}}_i$ maximizes the social welfare. Thus, mechanism $\mathcal{M}_{\mathcal{A}}$ is BIC according to Lemma 3.

Social welfare. The expected social welfare for this mechanism is $\sum_{i=1}^n \sum_{s=1}^{\ell} \sum_{t=1}^{\ell} x_i^{st} w_i^{st}$. Since for any $1 \leq i \leq n$ the allocation \mathbf{x}_i maximizes the social welfare for the induced assignment problem of agent i , the social welfare of \mathbf{x}_i is at least as large as that of the identity allocation, that is,

$$\forall i : \sum_{s=1}^{\ell} \sum_{t=1}^{\ell} x_i^{st} w_i^{st} \geq \sum_{s=1}^{\ell} f_i(s) w_i^{ss} = \mathbf{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{S} \sim \mathcal{A}(\mathbf{v})} [v_i(S_i)] .$$

Thus, we have that

$$SW^{\mathcal{M}_{\mathcal{A}}} = \sum_{i=1}^n \sum_{s,t=1}^{\ell} x_i^{st} w_i^{st} \geq \sum_{i=1}^n \mathbf{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{S} \sim \mathcal{A}(\mathbf{v})} [v_i(S_i)] = \mathbf{E}_{\mathbf{v} \sim \mathbf{F}, \mathbf{S} \sim \mathcal{A}(\mathbf{v})} \left[\sum_{i=1}^n v_i(S_i) \right] = SW^{\mathcal{A}} .$$

□

4.3 Estimating values by sampling.

There is still one technical issue that we need to settle in the reduction. In this section, we will briefly discuss how to use the standard sampling technique to obtain good estimated values of $\mathbf{w}_i = \{w_i^{st}\}_{1 \leq s, t \leq \ell}$ for the induced assignment problem of agent i for $1 \leq i \leq n$.

By definition, w_i^{st} is the expectation of a random variable $v_i^s(S_i)$, where S_i is the allocated service given by \mathcal{A} over random realization of the valuations \mathbf{v}_{-i} of other agents and random coin flips of the algorithm. Hence, if the standard deviation of $v_i^s(S_i)$ is not too large compared to its mean (no more than a polynomial factor), then we can draw polynomially many independent samples and take the average value as our estimated value. More concretely, the sampling algorithm proceeds as follows.

1. Draw $N = 4c^2 \log(n\ell^2/\epsilon)/\epsilon^2$ independent samples of $\mathbf{v} \sim \mathbf{F}$ conditioned on that the valuation of agent i is v_i^t , where

$$c = \frac{\sigma_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)]}{\mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)]} .$$

Let $\mathbf{v}^1, \dots, \mathbf{v}^N$ denote these N sample.

2. Use algorithm \mathcal{A} to compute an allocation $\mathbf{S}^k \sim \mathcal{A}(\mathbf{v}^k)$ for each sample \mathbf{v}^k , $1 \leq k \leq N$.
3. Let \hat{w}_i^{st} be the average of $v_i^s(S_i^k)$, $1 \leq k \leq N$.

Lemma 4. *The estimated values \hat{w}_i , $1 \leq i \leq n$, by the above sampling procedure satisfy for any $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$,*

$$\hat{w}_i^{st} \in [(1 - \epsilon)w_i^{st}, (1 + \epsilon)w_i^{st}]$$

with probability at least $1 - \epsilon$.

Proof. We shall have that

$$\begin{aligned} \mathbf{E} [\hat{w}_i^{st}] &= \mathbf{E}_{\mathbf{v}_{-i}, S \sim \mathcal{A}(\mathbf{v}_{-i}^t, \mathbf{v}_{-i})} [v_i^s(S_i)] = w_i^{st} , \\ \sigma [\hat{w}_i^{st}] &= \frac{1}{\sqrt{N}} \sigma_{\mathbf{v}_{-i}, S \sim \mathcal{A}(\mathbf{v}_{-i}^t, \mathbf{v}_{-i})} [v_i^s(S_i)] = \frac{c}{\sqrt{N}} \mathbf{E} [\hat{w}_i^{st}] = \frac{c}{\sqrt{N}} w_i^{st} . \end{aligned}$$

By Chernoff-Hoeffding bound, we get

$$\begin{aligned} \Pr [|\hat{w}_i^{st} - w_i^{st}| > \epsilon w_i^{st}] &= \Pr \left[|\hat{w}_i^{st} - \mathbf{E} [\hat{w}_i^{st}]| > \frac{\epsilon \sqrt{N}}{c} \sigma [\hat{w}_i^{st}] \right] \\ &= \Pr \left[|\hat{w}_i^{st} - \mathbf{E} [\hat{w}_i^{st}]| > 2 \sqrt{\log(n\ell^2/\epsilon)} \sigma [\hat{w}_i^{st}] \right] \\ &\leq e^{-\log(n\ell^2/\epsilon)} = \frac{\epsilon}{n\ell^2} . \end{aligned}$$

Since we only need to estimate $n\ell^2$ values, by union bound we get that with probability at least $1 - \epsilon$ the estimated value \hat{w}_i^{st} is within ϵ relative error compared to w_i^{st} for all $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$. \square

Thus, if the allocation algorithm \mathcal{A} admits $SW^{\mathcal{A}}$ social welfare and the ratio c is only polynomially large, then by part 2 of Theorem 1 we get that mechanism $\mathcal{M}_{\mathcal{A}}$ gives $(1 - 3\epsilon) \cdot SW^{\mathcal{A}}$ social welfare and is $4\epsilon v_{\max}$ -BIC. The running time is polynomial in the input size and $1/\epsilon$, assuming a black-box call to algorithm \mathcal{A} can be done in a single step. In other words, we get a FPTAS reduction.

The next lemma gives an alternative bound of the sampling error with respect to absolute error.

Lemma 5. *If we draw $N' = 4 \log(n\ell^2/\epsilon)/\epsilon^2$ independent samples, then with probability at least $1 - \epsilon$ the estimated values $\hat{w}_i^{st} \in [w_i^{st} - \epsilon v_{\max}, w_i^{st} + \epsilon v_{\max}]$ for all $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$.*

Proof. In this case, we have

$$\begin{aligned} \mathbf{E} [\hat{w}_i^{st}] &= \mathbf{E}_{\mathbf{v}_{-i}, S \sim \mathcal{A}(\mathbf{v}_{-i}^t, \mathbf{v}_{-i})} [v_i^s(S_i)] = w_i^{st} , \\ \sigma [\hat{w}_i^{st}] &= \frac{1}{\sqrt{N'}} \sigma_{\mathbf{v}_{-i}, S \sim \mathcal{A}(\mathbf{v}_{-i}^t, \mathbf{v}_{-i})} [v_i^s(S_i)] \leq \frac{1}{\sqrt{N'}} \max_{S_i} v_i^s(S_i) \leq \frac{1}{\sqrt{N'}} v_{\max} . \end{aligned}$$

By Chernoff bound we get that

$$\begin{aligned} \Pr [|\hat{w}_i^{st} - w_i^{st}| > \epsilon v_{\max}] &\leq \Pr \left[|\hat{w}_i^{st} - \mathbf{E} [\hat{w}_i^{st}]| > \frac{\epsilon}{\sqrt{N'}} \sigma [\hat{w}_i^{st}] \right] \\ &= \Pr \left[|\hat{w}_i^{st} - \mathbf{E} [\hat{w}_i^{st}]| > 2 \sqrt{\log(n\ell^2/\epsilon)} \sigma [\hat{w}_i^{st}] \right] \\ &\leq e^{-\log(n\ell^2/\epsilon)} = \frac{\epsilon}{n\ell^2} . \end{aligned}$$

By union bound, we have $\hat{w}_i^{st} \in [w_i^{st} - \epsilon v_{\max}, w_i^{st} + \epsilon v_{\max}]$ for all $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$. \square

Suppose the ratio $v_{\max}/SW^{\mathcal{A}}$ is upper bounded by a polynomial of the input size. Then, if we choose $\epsilon = \delta SW^{\mathcal{A}}/2nv_{\max}$ in the above lemma, we will get that

$$|\hat{w}_i^{st} - w_i^{st}| < \delta SW^{\mathcal{A}}/2n .$$

By part 3 of Theorem 1 we obtain that mechanism $\mathcal{M}_{\mathcal{A}}$ provides at least $(1 - \delta)SW^{\mathcal{A}}$ of social welfare and is 4ϵ -BIC and IR. The running time is polynomial in the input size and $1/\delta$.

5 Reductions for revenue and residual surplus

In the reduction for social welfare in the previous section, we only consider market-clearing allocations in the induced assignment problems. This is because for any agent i , we want to make sure that the intermediate algorithm \mathcal{B}_i is transparent to all agents except agent i . If we restrict ourselves to market-clearing allocations, we do not know any way to get reasonable bounds on revenue and residual surplus.

However, if we focus on an important sub-class of multi-parameter mechanism design problems that includes the combinatorial auction problem and its special cases, then we have some flexibility in choosing the allocations for the induced assignment problem and obtain theoretical bounds on revenue and residual surplus. More concretely, we will consider mechanism design problems that are *downward-closed*. We let ϕ denote the null service so that allocating ϕ to an agent implies that agent is not served, that is, $v_i(\phi) = 0$ for all $1 \leq i \leq n$.

Definition 7. A multi-parameter mechanism design problem $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$ is downward-closed if for any feasible allocation $\mathbf{S} = (S_1, \dots, S_n) \in \mathbf{J}$ and any $1 \leq i \leq n$, the allocation $(S_1, \dots, S_{i-1}, \phi, S_{i+1}, \dots, S_n)$ is also feasible.

We let $\delta = \min\{f_i(s) : 1 \leq i \leq n, 1 \leq s \leq \ell, f_i(s) > 0\}$ denote the granularity of the prior distributions. We will prove the following result.

Theorem 2. For any algorithm \mathcal{A} , there is a mechanism that is IR, BIC, and provides at least $\Omega(SW^{\mathcal{A}}/\log(1/\delta))$ of revenue (residual surplus).

5.1 Meta-reduction.

We will first introduce a meta-reduction scheme based on algorithms that compute envy-free solutions for fractional assignment problems. Suppose \mathcal{C} is an algorithm that computes envy-free solutions (\mathbf{x}, \mathbf{p}) for any given fractional assignment problem. Let \mathcal{A} be an algorithm for a multi-parameter mechanism design problem $\langle \mathbf{I}, \mathbf{J}, \mathbf{V}, \mathbf{F} \rangle$. We will convert algorithm \mathcal{A} into to a mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$:

1. For each agent i (Pre-computation)
 - (a) Estimate the values $\mathbf{w}_i = \{w_i^{st}\}_{1 \leq s, t \leq \ell}$ for the induced assignment problem of agent i with respect to \mathcal{A} . Let $\hat{\mathbf{w}}_i = \{\hat{w}_i^{st}\}_{1 \leq s, t \leq \ell}$ denote the estimated values.
 - (b) Use \mathcal{C} to solve the induced assignment problems with estimated values. Let $(\mathbf{x}_i, \mathbf{p}_i)$ denote the solution by \mathcal{C} for the induce assignment problem of agent i .
 - (c) Let $y_i^t = f_i(t) - \sum_{s=1}^{\ell} x_i^{st}$ denote the unallocated supply of virtual product t in solution $(\mathbf{x}_i, \mathbf{p}_i)$ for all $1 \leq i \leq n$ and $1 \leq t \leq \ell$.
 - (d) Let $y_i = \sum_{t=1}^{\ell} y_i^t$ denote the total amount of unallocated virtual products in $(\mathbf{x}_i, \mathbf{p}_i)$ for all $1 \leq i \leq n$.

2. Manipulate the valuations with intermediate algorithm \mathcal{B}_i , $1 \leq i \leq n$, as follows: **(Decoupling)**
 - (a) Suppose the reported valuation of agent i is $v'_i = v_i^s$.
 - (b) Let $\tilde{v}_i = \mathcal{B}_i(v'_i) = v_i^t$ with probability $x_{st}^i/f_i(s)$ for $1 \leq t \leq \ell$.
 - (c) With probability $1 - \sum_t x_{st}^i/f_i(s)$, the manipulated valuation \tilde{v}_i is unspecified in the previous step. In this case, let $\tilde{v}_i = v_i^t$ with probability y_i^t/y_i for $1 \leq t \leq \ell$.
3. Allocate services as follows: **(Allocation)**
 - (a) Compute a tentative allocation $\tilde{\mathbf{S}} = (\tilde{S}_1, \dots, \tilde{S}_n) = \mathcal{A}(\tilde{\mathbf{v}})$.
 - (b) For each agent i , let $S_i = \tilde{S}_i$ if the manipulated valuation \tilde{v}_i is specified in step 2b). Let $S_i = \phi$ otherwise. Allocate services according to $\tilde{\mathbf{S}}$.
 - (c) For each agent i , suppose the reported valuation is $v'_i = v_i^s$ and the manipulated valuation is $\tilde{v}_i = v_i^t$, charge agent i with price $p_i^t v_i^s(S_i)/\hat{w}_i^{st}$.

The following theorem states the above meta-reduction scheme converts algorithms into IR and BIC mechanisms.

Theorem 3. *Suppose the algorithm \mathcal{C} always provides envy-free solutions.*

1. *If the estimated values $\hat{\mathbf{w}}_i$ are accurate, then mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ is IR and BIC.*
2. *If the estimated values $\hat{\mathbf{w}}_i$ satisfy that for any $1 \leq s, t \leq \ell$, $\hat{w}_i^{st} \in [(1 - \epsilon)w_i^{st}, (1 + \epsilon)w_i^{st}]$, then $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ is IR and $4\epsilon v_{\max}$ -BIC.*
3. *If the estimated values $\hat{\mathbf{w}}_i$ satisfy that for any $1 \leq s, t \leq \ell$, $\hat{w}_i^{st} \in [w_i^{st} - \epsilon, w_i^{st} + \epsilon]$, then $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ is IR and 4ϵ -BIC.*

Proof. Let us outline the proof for part 1. Proofs of the other two parts are calculations along the same line.

Note that $p_i^t \leq w_i^{st}$ for all $1 \leq i \leq n$ and $1 \leq s, t \leq \ell$. The mechanism is IR because for any $1 \leq i \leq n$ and $1 \leq s \leq \ell$ the utility for an agent i with valuation v_i^s in any realization is

$$v_i^s(S_i) - p_i^t \frac{v_i^s(S_i)}{w_i^{st}} \geq 0 \quad .$$

Next, we will show that mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ is BIC. We first verify that the intermediate algorithms \mathcal{B}_i , $1 \leq i \leq n$, satisfy the constraint (3). For any agent i , if its valuation v_i is drawn from distribution F_i , then the probability that the manipulated valuation $\tilde{v}_i = \mathcal{B}_i(v_i) = v_i^t$ is

$$\begin{aligned} \sum_{s=1}^{\ell} f_i(s) \left[\frac{x_{st}^i}{f_i(s)} + \left(1 - \sum_{r=1}^{\ell} \frac{x_{sr}^i}{f_i(s)} \right) \frac{y_i^t}{y_i} \right] &= \sum_{s=1}^{\ell} x_{st}^i + \left(\sum_{s=1}^{\ell} f_i(s) - \sum_{s=1}^{\ell} \sum_{r=1}^{\ell} x_{sr}^i \right) \frac{y_i^t}{y_i} \\ &= \sum_{s=1}^{\ell} x_{st}^i + \left(\sum_{r=1}^{\ell} f_i(r) - \sum_{r=1}^{\ell} \sum_{s=1}^{\ell} x_{sr}^i \right) \frac{y_i^t}{y_i} \\ &= \sum_{s=1}^{\ell} x_{st}^i + \sum_{r=1}^{\ell} \left(f_i(r) - \sum_{s=1}^{\ell} x_{sr}^i \right) \frac{y_i^t}{y_i} \\ &= \sum_{s=1}^{\ell} x_{st}^i + \sum_{r=1}^{\ell} y_i^r \frac{y_i^t}{y_i} = \sum_{s=1}^{\ell} x_{st}^i + y_i^t = f_i(t) \quad . \end{aligned}$$

Thus, we get that for each agent i , the intermediate algorithms \mathcal{B}_j , $1 \leq j \leq n$ and $j \neq i$, are transparent to it. So the expected value of agent i of the service it gets, when its genuine valuation is $v_i = v_i^s$ and the manipulate valuation, is $\tilde{v}_i = v_i^t$ is exactly

$$w_i^{st} = \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)] \quad .$$

Hence, the expected value of agent i of the service it gets, when its genuine valuation is $v_i = v_i^s$ and the reported valuation is $v_i' = v_i^t$, is

$$\tilde{w}_i^{st} = \sum_{r=1}^{\ell} \frac{x_i^{tr}}{f_i(t)} w_i^{sr} \quad .$$

And the expected price for agent i when the reported valuation is $v_i' = v_i^t$ is

$$\tilde{p}_i^t = \sum_{r=1}^{\ell} \frac{x_i^{tr}}{f_i(t)} \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^r, \mathbf{v}_{-i})} \left[p_i^r \frac{v_i^t(S_i)}{w_i^{tr}} \right] = \sum_{r=1}^{\ell} \frac{x_i^{tr}}{f_i(t)} p_i^r \quad .$$

Thus, the the expected utility of agent i , when its genuine valuation is $v_i = v_i^s$ and its reported valuation is $v_i' = v_i^t$, is

$$\tilde{w}_i^{st} - \tilde{p}_i^t = \sum_{r=1}^{\ell} \frac{x_i^{tr}}{f_i(t)} (w_i^{sr} - p_i^r) \leq \sum_{r=1}^{\ell} \frac{x_i^{tr}}{f_i(t)} \max_k \{w_i^{sk} - p_i^k\} = \max_k \{w_i^{sk} - p_i^k\} \quad .$$

Since \mathbf{p}_i are chosen to be the envy-free prices, we have that $x_i^{sr} > 0$ only if $w_i^{sr} - p_i^r = \max_k \{w_i^{sk} - p_i^k\}$. Hence, when agent i reports its valuation truthfully, that is, $s = t$, the above inequality is tight. Moreover, the above utility is always non-negative. So mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ is BIC. \square

Moreover, the revenue and residual surplus of mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ is related to the social welfare and revenue of the induced assignment problems as stated in following proposition.

Proposition 1. *The expected revenue (residual surplus) of the mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$ equals the sum of the revenue (residual surplus) of the manipulated assignment problems.*

By choosing proper allocation algorithm \mathcal{C} , we can obtain theoretical bounds for the revenue or residual surplus in the manipulated induced assignment problems and thus theoretical bounds for mechanism $\mathcal{M}_{\mathcal{A}}^{\mathcal{C}}$.

5.2 Assignment algorithms.

In this section, we will introduce two algorithms for computing envy-free solutions for the induced assignment problems. These two algorithms provides theoretical bounds for revenue and residual surplus.

Revenue. The first algorithm provides revenue that is a $\Omega(1/\log(1/\delta))$ fraction of $SW^{\mathcal{A}}$, the social welfare by algorithm \mathcal{A} . The idea is to introduce proper reserve prices to the induced assignment problems by redundant virtual buyers. This is inspired by the techniques by Guruswami et al. [GHK⁺05]. For the induced assignment problem of agent i , $1 \leq i \leq n$, the assignment algorithm \mathcal{C}_R for revenue maximization proceeds as follows:

1. Find the social welfare maximizing allocation $\mathbf{x}_i = \{x_i^{st}\}_{1 \leq s, t \leq \ell}$.

2. Suppose u_{max} is the maximal valuation among the virtual buyer-product pair (s, t) with non-zero x_i^{st} , that is,

$$u_{max} = \max\{w_i^{st} : 1 \leq s, t \leq \ell, x_i^{st} > 0\} .$$

3. Recall that $\delta = \min\{f_i(t) : 1 \leq i \leq n, 1 \leq t \leq \ell, f_i(t) > 0\}$ denotes the granularity of the prior distribution. For $1 \leq k \leq \log(1/\delta)$:

- (a) Consider the following variant of the induced assignment instance of agent i :
For each virtual product $1 \leq t \leq \ell$, add a dummy virtual buyer with demand $1 + \delta$ and value $u_k = u_{max}/2^k$ for virtual product t and value 0 for other virtual products.
- (b) Find social welfare maximizing allocation \mathbf{x}_{ik} and envy-free prices \mathbf{p}_{ik} for this variant.
- (c) Let $(\hat{\mathbf{x}}_{ik}, \hat{\mathbf{p}}_{ik})$ be the projection of $(\mathbf{x}_{ik}, \mathbf{p}_{ik})$ on the original induced assignment problem of agent i , that is, for any $1 \leq s, t \leq \ell$, $\hat{x}_{ik}^{st} = x_{ik}^{st}$, $\hat{p}_{ik}^t = p_{ik}^t$.

4. Return the $(\hat{\mathbf{x}}_{ik}, \hat{\mathbf{p}}_{ik})$, $1 \leq k \leq \log(1/\delta)$, with best revenue.

Lemma 6. *Assignment algorithm \mathcal{C}_R always return an envy-free solution (\mathbf{x}, \mathbf{p}) . The revenue is at least a $\Omega(1/\log(1/\delta))$ fraction of the optimal social welfare of the assignment problem.*

Proof. The envy-freeness follows from the fact that $(\hat{\mathbf{x}}_{ik}, \hat{\mathbf{p}}_{ik})$, $1 \leq k \leq \log(1/\delta)$, are projections of envy-free solutions and thus are also envy-free.

Now we consider the revenue by \mathcal{C}_R . We let r_k denote the revenue by solution $(\hat{\mathbf{x}}_{ik}, \hat{\mathbf{p}}_{ik})$. Note that in $(\hat{\mathbf{x}}_{ik}, \hat{\mathbf{p}}_{ik})$, all prices are at least u_k and the amount of virtual products that are sold is at least $\sum_{s,t:w_i^{st} \geq u_k} x_i^{st}$. Hence, we have

$$r_k \geq w_k \sum_{s,t:w_i^{st} \geq u_k} x_i^{st} .$$

Note that if we extend the definition of u_k and let $u_k = u_{max}/2^k$ for all non-negative integer k , then we have

$$\begin{aligned} \sum_{k=1}^{\infty} u_k \sum_{s,t:w_i^{st} \geq u_k} x_i^{st} &= \sum_{k=1}^{\infty} (u_{k-1} - u_k) \sum_{s,t:w_i^{st} \geq u_k} x_i^{st} \\ &= \sum_{s=1}^{\ell} \sum_{t=1}^{\ell} x_i^{st} \sum_{k:w_i^{st} \geq u_k} (u_{k-1} - u_k) \\ &= \sum_{s=1}^{\ell} \sum_{t=1}^{\ell} x_i^{st} \max_k \{u_{k-1} : w_i^{st} \geq u_k\} \\ &\geq \sum_{s,t} x_i^{st} w_i^{st} . \end{aligned} \tag{4}$$

On the other hand, the contribution of the tail is small compared to the social welfare.

$$\sum_{k=\log(1/\delta)+1}^{\infty} u_k \sum_{s,t:w_i^{st} \geq u_k} x_i^{st} \leq \sum_{k=\log(1/\delta)+1}^{\infty} w_k \leq \frac{\delta w_{max}}{2} \leq \frac{\sum_{s,t} x_i^{st} w_i^{st}}{2} . \tag{5}$$

The last inequality holds because allocating the most valuable virtual product the one of the virtual buyer is a feasible allocation. Hence, consider the difference of the above formulas, (4) – (5), and we get that

$$\sum_{k=1}^{\log(1/\delta)} r_k \geq \sum_{k=1}^{\log(1/\delta)} u_k \sum_{s,t: w_{st} \geq u_k} x_i^{st} \geq \frac{\sum_{s,t} x_i^{st} w_i^{st}}{2} .$$

Thus, by pigeon-hole-principle at least one of the assignment $(\hat{x}_{ik}, \hat{p}_{ik})$ provides revenue that is at least a $1/2 \log(1/\delta)$ fraction of the social welfare. \square

The above lemma leads to the following results for revenue maximization.

Proposition 2. *Suppose the social welfare given by allocation algorithm \mathcal{A} is $SW^{\mathcal{A}}$, the mechanism $\mathcal{M}_{\mathcal{A}}^{C_R}$ guarantees at least $\Omega(SW^{\mathcal{A}}/\log(1/\delta))$ of revenue.*

Complementary lower bound. The approximation ratio with respect to $SW^{\mathcal{A}}$ is tight due to the following example. Consider the auction problem with only one agent and one item. Suppose with probability $1/2^k$ the agent has value 2^k for the item for $k = 1, 2, \dots, \log(1/\delta)$. And with probability δ , the agent has value 0 for the item. In this example, the granularity of the prior distribution is δ . The optimal social welfare is $\sum_{k=1}^{\log(1/\delta)} \frac{1}{2^k} 2^k = \log(1/\delta)$. But no BIC mechanism can achieve revenue better than 1.

Residual surplus. We turn to the residual surplus maximization problem. Note that revenue and residual surplus are symmetric in the induced assignment problems. We will use the following assignment algorithm \mathcal{C}_{RS} based on the same idea we use for the revenue maximization algorithm.

The residual surplus maximizing envy-free algorithm \mathcal{C}_{RS} is as follows:

1. Find the social welfare maximizing allocation $\mathbf{x}_i = \{x_i^{st}\}_{1 \leq s, t \leq \ell}$.
2. Suppose u_{max} is the maximal valuation among the virtual buyer-product pair (s, t) with non-zero x_i^{st} , that is,
$$u_{max} = \max\{w_i^{st} : 1 \leq s, t \leq \ell, x_i^{st} > 0\} .$$
3. Recall that $\delta = \min\{f_i(t) : 1 \leq i \leq n, 1 \leq t \leq \ell, f_i(t) > 0\}$ denotes the granularity of the prior distribution. For $1 \leq k \leq \log(1/\delta)$:
 - (a) Consider the following variant of the induced assignment instance of agent i :
For each virtual buyer $1 \leq t \leq \ell$, add a dummy virtual product with demand $1 + \delta$ and value $u_k = u_{max}/2^k$ for virtual buyer t and value 0 for other virtual buyer.
 - (b) Find social welfare maximizing allocation \mathbf{x}_{ik} and envy-free prices \mathbf{p}_{ik} for this variant.
 - (c) Let $(\hat{x}_{ik}, \hat{p}_{ik})$ be the projection of $(\mathbf{x}_{ik}, \mathbf{p}_{ik})$ on the original induced assignment problem of agent i , that is, for any $1 \leq s, t \leq \ell$,

$$\hat{x}_{ik}^{st} = x_{ik}^{st} \quad , \quad \hat{p}_{ik}^t = p_{ik}^t .$$

4. Return the $(\hat{x}_{ik}, \hat{p}_{ik})$, $1 \leq k \leq \log(1/\delta)$, with best revenue.

The proofs of the following lemma and theorem is almost identical to the revenue maximization part so we omit the proofs here.

Lemma 7. *Assignment algorithm \mathcal{C}_{RS} always return an envy-free solution (\mathbf{x}, \mathbf{p}) . The residual surplus is at least a $\Omega(1/\log(1/\delta))$ fraction of the optimal social welfare of the assignment problem.*

Proposition 3. *Suppose the social welfare given by allocation algorithm \mathcal{A} is $SW^{\mathcal{A}}$, the mechanism $\mathcal{M}_{\mathcal{A}}^{CRS}$ guarantees at least $\Omega(SW^{\mathcal{A}}/\log(1/\delta))$ of residual surplus.*

6 Application in combinatorial auctions

In this section we will briefly illustrates how to use the reduction for social welfare in this paper to derive a combinatorial auction mechanism that matches the best algorithmic result.

Combinatorial auctions. In the combinatorial auctions, we consider a principal with m items (exactly one copy of each item) and n agents. Each agent has a private valuation $v_i \sim F_i$ for subsets of items. The goal is to design a protocol to allocate the items and to charge prices to the agents.

We will show the following corollaries of our reduction for social welfare.

Corollary 2. *For combinatorial auctions with sub-additive (or fractionally sub-additive) agents where the prior distributions have finite and poly-size supports, there is a $(\frac{1}{2} - \epsilon)$ -approximate (or $(1 - \frac{1}{e} - \epsilon)$ -approximate respectively), ϵv_{max} -BIC, and IR mechanism for social welfare maximization. The running time is polynomial in the input size and $1/\epsilon$.*

Algorithm. We will consider a variant of the LP-based algorithms by Feige [Fei06] and use the reduction for social welfare to convert it into an IR and ϵv_{max} -BIC mechanism. More concretely, we will consider the Bayesian version of the standard social welfare maximization linear program (CA):

$$\begin{aligned}
& \text{Maximize} && \sum_i \sum_t \sum_S f_i(t) v_i^t(S) x_{i,t,S} && \text{s.t.} \\
& \sum_i \sum_t \sum_{S:j \in S} f_i(t) x_{i,t,S} && \leq 1 && \forall j \\
& \sum_S x_{i,t,S} && \leq 1 && \forall i, t \\
& x_{i,t,S} && \geq 0 && \forall i, t, S
\end{aligned}$$

In this LP, $x_{i,t,S}$ denote the probability that agent i is allocated with a subset of items S conditioned on its valuation is v_i^t . This LP can be solved in polynomial time by the standard primal dual technique via demand queries. See [DNS05] for more details. We let LP^* denote the optimal value of this LP. Moreover, for any basic feasible optimal solution of the above LP, there are at most $nm\ell$ non-zero entries since there are only $nm\ell$ non-trivial constraints. Hence, we have the following lemma:

Lemma 8. *In poly-time we can find an optimal solution \mathbf{x}^* to (CA) with at most $nm\ell$ non-zero entries.*

Next, we will filter this solution \mathbf{x}^* by removing insignificant entries. We let $\hat{x}_{i,t,S} = x_{i,t,S}^* < \epsilon/nm\ell$. Note that $LP^* \geq f_i(t)v_i^t(S)$ for any i, t , and S since always allocating subset S to agent i is a feasible allocation. We get that $\hat{\mathbf{x}}$ is a feasible solution to (CA) with value at least $(1 - \epsilon)LP^*$.

Then, we will use the rounding algorithms by Feige [Fei06] to get a $\frac{1}{2}$ -rounding for sub-additive agents and a $(1 - \frac{1}{e})$ -rounding for fractionally sub-additive agents:

1. Allocate a tentative subset of items \tilde{S}_i to each agent i , $1 \leq i \leq n$, according to the reported valuation $v'_i = v_i^t$ and $\hat{x}_{i,t,\tilde{S}_i}$.
2. Resolve conflicts properly by choosing $S_i \subseteq \tilde{S}_i$ so that $\mathbf{S} = (S_1, \dots, S_n)$ is a feasible allocation.

By extending Feige's original proof, we can show that there is a randomized algorithm for choosing \mathbf{S} such that for sub-additive agents, we have:

$$\mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S}} [v_i(S_i)] \geq \frac{1}{2} v_i(\tilde{S}_i) . \quad (6)$$

And for fractionally sub-additive agents, we have:

$$\mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S}} [v_i(S_i)] \geq \left(1 - \frac{1}{e}\right) v_i(\tilde{S}_i) . \quad (7)$$

We will omit the proof in this extended abstract. We denote the above algorithm as \mathcal{A} . Then, \mathcal{A} provides $(\frac{1}{2} - \epsilon)$ -approximation for sub-additive agents and $(1 - \frac{1}{e} - \epsilon)$ -approximation for fractionally sub-additive agents.

Estimating values. By Theorem 1 and 2, we only need to show how to estimate the values \mathbf{w}_i , $1 \leq i \leq n$, for the induced assignment problem of agent i efficiently. Further, by Lemma 4, we can efficiently estimate the values $\mathbf{w}_i = \{w_i^{st}\}_{1 \leq s, t \leq \ell}$, $1 \leq i \leq n$, if the following lemma holds.

Lemma 9. For any $1 \leq i \leq n$, and any $1 \leq s, t \leq \ell$,

$$\frac{\sigma_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)]}{\mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)]} \leq \sqrt{\frac{4nm\ell}{\epsilon}} .$$

Proof. By inequalities (6) and (7), we get that conditioned on \tilde{S}_i being chosen as the tentative set,

$$\mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i) : \tilde{S}_i] \geq \frac{1}{2} v_i^s(\tilde{S}_i) .$$

We also have that

$$\sigma_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i) : \tilde{S}_i] \leq \max \left\{ v_i^s(S_i) : \tilde{S}_i \right\} \leq v_i^s(\tilde{S}_i) .$$

Hence,

$$\begin{aligned} \sigma_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)]^2 &= \sum_{\tilde{S}_i} \hat{x}_{i,t,\tilde{S}_i} \sigma_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i) : \tilde{S}_i]^2 \\ &\leq \sum_{\tilde{S}_i} \hat{x}_{i,t,\tilde{S}_i} v_i^s(\tilde{S}_i)^2 \\ &\leq \frac{1}{\min \left\{ \hat{x}_{i,t,\tilde{S}_i} > 0 \right\}} \left(\sum_{i,t,\tilde{S}_i} \hat{x}_{i,t,\tilde{S}_i} v_i^s(\tilde{S}_i) \right)^2 \\ &\leq \frac{nm\ell}{\epsilon} \left(\sum_{\tilde{S}_i} \hat{x}_{i,t,\tilde{S}_i} \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i) : \tilde{S}_i] \right)^2 \\ &\leq \frac{4nm\ell}{\epsilon} \mathbf{E}_{\mathbf{v}_{-i}, \mathbf{S} \sim \mathcal{A}(v_i^t, \mathbf{v}_{-i})} [v_i^s(S_i)]^2 . \end{aligned}$$

□

References

- [BGGM10] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget constrained auctions with heterogeneous items. In *ACM 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 2010.
- [CHMS10] S. Chawla, J.D. Hartline, D. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *ACM 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 2010.
- [Cla71] E.H. Clarke. Multipart pricing of public goods. *Public choice*, 11(1):17–33, 1971.
- [DNS05] S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *ACM 37th Annual ACM Symposium on Theory of Computing (STOC)*, page 618. ACM, 2005.
- [DNS06] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *ACM 38th Annual ACM Symposium on Theory of Computing (STOC)*, page 652. ACM, 2006.
- [Dob07] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103, 2007.
- [DR10] S. Dughmi and T. Roughgarden. Black-box randomized reductions in algorithmic mechanism design. In *51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2010.
- [Fei06] U. Feige. On maximizing welfare when utility functions are subadditive. In *ACM 38th Annual ACM Symposium on Theory of Computing (STOC)*, page 50. ACM, 2006.
- [FV06] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1-1/e$. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 667–676, 2006.
- [GHK⁺05] V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Annual ACM-SIAM Symposium on Discrete algorithms*, page 1173. Society for Industrial and Applied Mathematics, 2005.
- [Gro73] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, 41(4):617–631, 1973.
- [GS99] F. Gul and E. Stacchetti. Walrasian Equilibrium with Gross Substitutes* 1. *Journal of Economic Theory*, 87(1):95–124, 1999.
- [HKM11] J. Hartline, R. Kleinberg, and A. Malekian. Bayesian incentive compatibility via matchings. In *Annual ACM-SIAM Symposium on Discrete algorithms*, to appear, 2011.
- [HL10] J.D. Hartline and B. Lucier. Bayesian algorithmic mechanism design. In *ACM 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 2010.
- [HR08] J.D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *ACM 40th Annual ACM Symposium on Theory of Computing (STOC)*, 2008.

- [LS05] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 595–604, 2005.
- [Mye81] R.B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [PSS08] C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 250–259, 2008.
- [Roc87] J.C. Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200, 1987.
- [Vic61] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of finance*, 16(1):8–37, 1961.