Selling in Exclusive Markets: Some Observations on Prior-free Mechanism Design

Anna R. Karlin [1] C. Thach Nguyen [1] Yuval Peres [2]

March 1, 2012

Abstract

We consider prior-free benchmarks in non-matroid settings. In particular, we show that a very desirable benchmark proposed by Hartline and Roughgarden [22] is too strong, in the sense that no truthful mechanism can compete with it even in a very simple non-matroid setting where there are two exclusive markets and the seller can only sell to agents in one of them. On the other hand, we show that there is a mechanism that competes with a symmetrized version of this benchmark. We further investigate the more traditional best fixed price profit benchmark and show that there are mechanisms that compete with it in any downward-closed settings.

1 Introduction

We consider the design of truthful profit maximizing mechanisms for allocation problems in single-parameter settings. In these settings, there is an auctioneer that is offering goods or a service to a set N of n agents. Each agent has their own private value v for receiving service. Depending on the scenario, there are different constraints on which subset of agents can be served simultaneously. The constraints of the particular scenario are expressed using a set system $(N, \mathcal{S})$, where each set $S \in \mathcal{S}$ is a subset of N, and represents a legal allocation, meaning that it is feasible for the set of agents S to receive service simultaneously. Examples of such settings include:

- single item auctions, in which $\mathcal{S}$ consists of all sets of agents of size 1;
- k-unit auctions, in which $\mathcal{S}$ is the collection of all sets containing at most k agents;
- Digital goods auctions, in which $\mathcal{S}$ is the collection of all subsets of N.
- Single-minded combinatorial auctions: In a single-minded combinatorial auction, there is a seller with a collection of items. Each agent, say the ith, is interested in some particular subset T_i of the items. The feasible sets in $\mathcal{S}$ are all subsets of agents whose sets T_i of interest are disjoint.
- Multiple, exclusive markets: In a multiple, exclusive market setting, the agents N are partitioned into a set of markets. The service provider can pick one of the markets and serve any subset of agents in that market, but cannot serve two agents that are in different markets.

[1] University of Washington, email: {karlin,ncthach}@cs.washington.edu. Research supported in part by NSF grant CCF 1016509 and an award from Yahoo! Research.
[2] Microsoft Research, email: peres@microsoft.com.

In an auction for allocating service in any of these settings, the auctioneer takes as input a bid from each of the agents, with b_i the bid of agent i, and chooses as output a feasible subset $S \in \mathcal{S}$ of winning agents, and a payment from each winning agent. We assume the standard quasi-linear utility model: agents bid to maximize their own utility $u_i = v_i x_i - p_i$, where v_i is the true private value of agent i for receiving service, x_i is an indicator of whether or not agent i receives service, and p_i is the payment agent i is required to make to the auctioneer.

We consider the goal of designing truthful auctions that maximize the auctioneer's profit, $\sum_i p_i$.

The most common approach to the problem of profit maximization or *optimal mechanism design*, as it is termed in the economics literature, is in a *Bayesian* setting. This assumes that players' values are drawn from prior distributions *known* to the mechanism designer. The designer then uses his knowledge of the priors to design a mechanism that is optimal, that is, maximizes the auctioneer's expected profit (where the expectation is taken over the draws from the prior distributions).

In a beautiful sweeping result, Roger Myerson [29] showed how to construct the optimal Bayesian auction in any single-parameter environment, under the assumption that the distributions for agents values are independent (though not necessarily identical).

As noted, crucial to the Myerson result is the fact that the auctioneer knows the prior distributions from which the agent's values are drawn. This can be problematic however, since determining priors may not be convenient, reasonable or even possible. This is particularly true in small markets, where the process of collecting information, e.g. via market research techniques, may negatively impact both the incentives of the agents and the performance of the mechanism. Priors may also change over time. Lastly, it is unclear how sensitive the Myerson mechanism is to (the inevitable) errors there might be about the prior.

Motivated by these observations, there has been a great deal of interest in developing mechanisms that are more robust and do not rely on Bayesian assumptions, but rather, work well even in the worst-case. One of the first results of this kind was in the works of Goldberg, Hartline and Wright [16] and Goldberg, Hartline, Karlin, Saks and Wright [15], which showed how to design profit maximizing mechanisms for selling digital goods that work well no matter what the player's valuations are, and without any assumptions about prior distributions. Since there is no such thing as an absolutely optimal truthful auction, the approach taken was a sort of "competitive analysis". This consisted of defining a *profit benchmark* $\mathcal{G}$, a function that maps instances (in this case, the values v_i of each of the agents) to a real number that represents the target profit we wish the truthful auction to try to attain, given those values as input. Given a profit bencharK, the challenge is to design a truthful mechanism that approximates the profit benchmark $\mathcal{G}$ *on every instance*.

In the work on digital goods, the profit benchmark chosen was the *best fixed price* profit (BFP), which is the maximum profit one can obtain by posting a price and selling to any agent willing to accept it. The authors showed that there was a mechanism that can guarantee a constant fraction of the profit of BFP on every instance.[1]

They also considered another natural profit benchark – the maximum possible profit the auctioneer can make given the player's valuations, which is the social welfare or sum of player's valuations. They showed that social welfare is too strong a benchmark: no truthful mechanism can achieve better than a $\log n$ approximation to this benchmark [15]. Moreover, there are mechanisms that get a $\log n$ approximation to social welfare on every instance, but do not do any better than this on "easy" instances, such as when all the values are the same. Thus, competing with this benchmark does not distinguish between "good" and "bad" mechanisms. On the other hand, an auction that is constant-competitive to the best fixed price profit BFP is automatically $\log n$-competitive to social welfare. This and the fact that BFP is very natural gave one the feeling, but not a principled argument, that it was the "right" benchmark.

The work on digital goods motivated a line of research aimed at designing mechanisms with better competitive ratios, as well as on developing new mechanisms that deal with added constraints such

[1] Caveat: There can be no one dominant bidder.

as the scenarios described above (e.g. [1, 2, 5, 7, 11–13, 18, 21], and for more of a survey, see [19, 20]). However, it remained unclear in these more complex setting how to generalize the benchmark. Indeed much of the research approached the benchmark question in a somewhat ad-hoc fashion.

In 2009, Hartline and Roughgarden [22] proposed a framework for approaching the benchmark question that explicitly ties prior-free mechanism design to Bayesian mechanism design. Specifically, they proposed the following recipe for prior free mechanism design.

1. Characterize all i.i.d. Bayesian optimal mechanisms for the environment in question, i.e, given the particular setting, the set N of agents and the set system of feasible allocations $\mathcal{S}$, characterize the optimal mechanism $\mathsf{OPT}_F^{(N,\mathcal{S})}$ for every distribution F.

2. Define the prior-free benchmark to be the profit of the best possible i.i.d. Bayesian mechanism on each instance, i.e.
$$\mathcal{G}(\mathbf{v}) = \sup_F \mathsf{OPT}_F^{(N,\mathcal{S})}(\mathbf{v}).$$

3. Design a prior-free, truthful mechanism $\mathcal{M}^{(N,\mathcal{S})}$ that approximates $\mathcal{G}$ on any instance of the environment, i.e.,
$$\mathcal{M}^{(N,\mathcal{S})}(\mathbf{v}) \geq \frac{1}{c} \cdot \mathcal{G}(\mathbf{v})$$
for any $\mathbf{v}$, where $\mathcal{M}^{(N,\mathcal{S})}(\mathbf{v})$ is the profit of the mechanism $\mathcal{M}^{(N,\mathcal{S})}$ on the given instance. The constant c is called the *competitive ratio* of $\mathcal{M}^{(N,\mathcal{S})}$.

4. Show that $\mathcal{M}^{(N,\mathcal{S})}$ achieves the best possible competitive ratio.

As discussed above, in the single-parameter settings we are discussing, the Baysian optimal mechanisms are well understood: they are the Myerson mechanisms. We will use the notation $\mathsf{Mye}_{\mathbf{F}}(\mathbf{v})$ to represent the profit of the Myerson mechanism tailored to the distribution $\mathbf{F}$ in the given environment $(N, \mathcal{S})$ on the instance $\mathbf{v}$. The resulting benchmark $\sup_F \mathsf{Mye}_{\mathbf{F}}(\mathbf{v})$ was dubbed the *optimal Baysian optimal* benchmark (OBO) [7] as it corresponds to profit one can obtain by choosing the best possible Bayesian optimal mechanism for each given instance.

Hartline and Roughgarden pointed out that, in retrospect, the work on profit maximization in digital goods auctions precisely fits into their framework. For these environments, every i.i.d. Bayesian optimal mechanism sets a (different) fixed price and sells to all agents who are willing to pay that much. Hence, BFP is in fact the benchmark $\mathcal{G}$ in this framework.

To the extent that it is possible to design mechanisms that are competitive to the OBO benchmark, it is extremely appealing. This is because OBO upper bounds the revenue of *every* i.i.d. Bayesian optimal mechanism on *every* instance. Thus, if it is possible to design a mechanism $\mathcal{M}^{(N,\mathcal{S})}$ that is competitive to the OBO profit benchmark on all inputs, this mechanism $\mathcal{M}^{(N,\mathcal{S})}$ is *competitive to all i.i.d. Baysian mechanisms simultaneously*! Such a mechanism achieves the best of both worlds, that is, the i.i.d. Bayesian world and the worst-case world (and everything in between).

Amazingly enough, Hartline and Roughgarden [23] observed that in all single-parameter matroid settings, it *is* possible to design mechanisms that are constant competitive to a benchmark nearly as strong as OBO.(The benchmark was $\sup_{F \text{ regular}} \mathsf{OPT}_F^{(N,\mathcal{S})}(\mathbf{v})$.) By matroid setting, we mean that the set system $(N, \mathcal{S})$ contains the independent sets of a matroid. Numerous interesting settings fall into this framework.

Thus a major remaining challenge was to understand if it is possible to design truthful mechanisms that are constant competitive to the OBO benchmark in other settings, that are not matroids.

This is the question we look at in this paper. To this end, we focus on one of the simplest non-matroid environments: two exclusive markets. This is an environment in which there are two disjoint markets M_1 and M_2. A seller can sell to any subset of agents in one of the markets, but cannot sell to two agents from different markets. Thus, the set system of the environment is

$$\mathcal{S} = \bigcup_{i=1}^{2} \{S : S \subseteq M_i\}.$$

The exclusive market environment is important as it models a number of real world situations. For example, there may be regulations that keep service providers from operating across market segments, or there may be supply constraints that keep sellers from selling in different markets. For example, a vendor truck may only operate in one city on each particular day, an airline with a limited number of aircraft may operate in a certain region during the year, etc.

Results

Our main results are the following:

- We show that even for the very simple setting of two exclusive markets, the benchmark OBO is too strong: no truthful mechanism can achieve a competitive factor better than $\Omega(\log n)$ to OBO. In fact, our proof shows that even if we weaken the benchmark and only consider the optimal Bayesian optimal over monotone hazard rate distributions, i.e. $\sup_{F \text{ MHR}} \mathsf{OPT}_F^{(N,\mathcal{S})}(\mathbf{v})$, no truthful mechanism can achieve a competitive factor better than $\Omega(\log n)$ (see Section 2 for the definition of MHR distributions). This puts to rest the idea of using the OBO benchmark, at least if one strives for constant competitiveness.

- In light of the previous result, we consider a different benchmark, sOBO, which is a symmetrized version of OBO.[2] This benchmark makes much more sense in an asymmetric environment like exclusive markets, and retains many strong properties of OBO. Specifically, it is still the case that a mechanism that competes with with sOBO is competitive to all i.i.d. Bayesian optimal auctions simultaneously. We show that for environments with two exclusive markets, sOBO is within a constant factor of BFP.

  Together with the results outlined in the next bullet point, this implies that for two exclusive market, there are mechanisms that compete with sOBO. While the two exclusive market environment is extremely restrictive, this yields the first prior-free result in a non-matroid environment that is also provably competitive to all i.i.d. Bayesian optimal auctions simultaneously. Moreover, while sOBO has already been considered in the literature [24] (see discussion in Section 6), no previous work has considered mechanisms that achieve *pointwise* guarantees relative to sOBO.

- Finally, using an elegant result of Ha and Hartline [17], we show that BFP can be constant approximated by a prior-free truthful mechanism pointwise in any downward-closed setting. The resulting competitive ratio is 17.5. In any exclusive market environment (for any number of exclusive markets), we give a simpler truthful mechanism that approximates BFP pointwise and achieves a better competitive ratio, 9.

Related Work

In addition to the related work mentioned above, there is another interesting body of work in the design of profit-maximizing mechanisms, and this is the work on prior-independent mechanisms. These results consider the problem of designing mechanisms that work well when the players values are drawn from an *unknown* prior distribution. However, unlike what we are looking at here, they are

[2]This benchmark is considered in [24].

not looking at worse-case bounds or pointwise guarantees. Rather they seek to design mechanisms that work well in expectation over the draws from the distribution. Results of this flavor include [3,6,8].

The most closely related work to our own are the papers of Hartline and Yan [24] and Ha and Hartline [17]. We discuss them in more detail in Section 6.

Organization of the paper

Section 2 reviews background material and formalizes our setting. In Section 3, we outline the proof that the OBO benchmark is too strong for the two exclusive market setting and, in Section 4, we describe the symmetrized benchmark sOBO and outline the proof that in the same setting, it is well approximated (in expectation) by the BFP benchmark. In Section 5, we present a truthful mechanism that is constant-competitive to BFP in any downward closed environment and the simpler mechanism for the exclusive market setting. We conclude in Section 6 with a discussion in which we put the work in context and highlight some interesting related open problems.

For readability, most of the actual proofs and technical details are deferred to the appendices.

2 Preliminaries

We begin by reviewing a number of preliminaries and formalizing our setting.

Mechanisms

As discussed above, we focus on mechanisms for single-parameter allocation problems, in which there is a service provider offering a set of abstract services (or goods), each targeted to an agent in a set N. Each agent $i \in N$ has, as her private information, a value v_i of being served, i.e., of being allocated the service. We use $\mathbf{v}$ to denote the vector of agent's values. The constraints of the particular scenario are expressed using a set system $\mathcal{S}$, where each set $S \in \mathcal{S}$ is a subset of N, and represents a legal allocation, meaning that it is feasible for the set of agents S to receive service simultaneously. When discussing the Bayesian approach, we assume that the value v_i of player i is drawn independently from a publicly known distribution F_i. We use $\mathbf{F} = \prod_{i \in N} F_i$ to denote the product distribution from which the vector $\mathbf{v}$ is drawn.

In the large part of this paper, we focus on *exclusive markets*, in which there are k markets $M_1, M_2, \ldots M_k$ which are mutually exclusive, i.e., the set of agents $N = \bigcup_{1 \le i \le k} M_i$, and $M_i \cap M_j = \emptyset$ when $i \neq j$. The service provider can serve any number of agents in one of the markets, but cannot serve agents from different markets, i.e. the set system is $\mathcal{S} = \bigcup_{i=1}^{k} \{S : S \subseteq M_i \text{ for some } i\}$. Sections 3 and 4 deal with the case where $k = 2$. Section 5 deals partly with the general exclusive market settings, and partly with the general downward-closed settings. In a general downward-closed settings, the set system satisfies the following condition: if $S \in \mathcal{S}$ and $T \subseteq S$ then $T \in \mathcal{S}$. In particular, all the settings discussed above, i.e., single item auctions, k-unit auctions, digital good auctions, single-minded combinatorial auctions, and exclusive markets are downward-closed settings.

A mechanism asks the agents to bid, i.e., to submit their private values, and, based on their responses, chooses an allocation and computes the agents' payments. The set W of agents that are served will always belong to $\mathcal{S}$ and the vector $\mathbf{p}$ will denote the payments, where p_i is agent i's payment. Together, the allocation and payments constitute the outcome of the mechanism. In the following, we will also denote the outcome $(W, \mathbf{p})$ by a pair of vectors $(\mathbf{x}, \mathbf{p})$ where $\mathbf{x}$ is a vector whose i'th entry x_i is 1 if agent i is served and 0 otherwise. When we need to emphasize the agents' bids that lead to an outcome, we will use the notation $(\mathbf{x}(\mathbf{b}), \mathbf{p}(\mathbf{b}))$ where b_i is the bid of agent i. For randomized mechanisms, x_i will denote the probability that agent i is served, and p_i will denote her expected payment.

Utility and truthfulness

The *utility* of an agent i for an outcome $(\mathbf{x}, \mathbf{p})$ of a mechanism is $u_i = v_i x_i - p_i$. Naturally, we assume that agents act to maximize their utilities. Also, in all of the following, we only consider mechanisms that satisfy the following two natural conditions

- *Individual rationality:* agents have non-negative utilities in every possible outcome.
- *No positive transfers:* the prices are always non-negative; i.e. the mechanisms never "give money" to the agents.

We focus here on truthful mechanisms, i.e., ones in which to maximize their (expected) utilities, agents bid their true values. To review the formal definition, we introduce some standard notation.

Notation. *Let $\mathbf{v}$ be a one-dimensional vector, then $\mathbf{v}_{-i}$ is the vector obtained by removing $\mathbf{v}'s$ ith entry, i.e. $\mathbf{v}_{-i} = (v_1, v_2, \dots v_{i-1}, v_{i+1}, v_{i+2}, \dots, v_n)$, where n is the length of $\mathbf{v}$. Thus, we have three equivalent notations: $\mathbf{v} = (v_i, \mathbf{v}_{-i}) = (v_1, v_2, \dots v_n)$. Similarly, if $\mathbf{F} = F_1 \times F_2 \times \dots \times F_n$, then $\mathbf{F}_{-i}$ is the product distribution $F_1 \times F_2 \times \cdots F_{i-1} \times F_{i+1} \times \cdots \times F_n$.*

Definition 1. *A mechanism is* truthful in expectation *if for any agent i and any bid vector $\mathbf{b}$,*

$$x_i(v_i, \mathbf{b}_{-i})v_i - p_i(v_i, \mathbf{b}_{-i}) \geq x_i(\mathbf{b})v_i - p_i(\mathbf{b}).$$

In single parameter environments, truthful mechanisms are well understood. The following theorem provides a complete characterization:

Theorem 2 ([29]). *A (randomized) mechanism is truthful in expectation if and only if for any agent i and any fixed choice of bids by the other agents $\mathbf{b}_{-i}$,*

- *$x_i(b_i, \mathbf{b}_{-i})$ is monotone non-decreasing in b_i.*
- $p_i(b_i, \mathbf{b}_{-i}) = b_i x_i(b_i, \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z, \mathbf{b}_{-i})dz + p_i(0, \mathbf{b}_{-i})$.

We will always assume that $p_i(0, \mathbf{b}_{-i}) = 0$ and hence, when restricted to deterministic mechanisms, the above theorem becomes:

Corollary 3. *Any deterministic truthful mechanism is specified by a set of functions $t_i(\mathbf{b}_{-i})$ which determine, for each bidder i, and each set of bids $\mathbf{b}_{-i}$, an offer price to bidder i such that i wins and pays t_i if $b_i > t_i$ and loses and pays nothing if $b_i < t_i$. The allocation can be arbitrary when $b_i = t_i$. t_i is sometimes called agent i's* threshold bid.

When we discuss truthful mechanisms, we will assume that $b_i = v_i$ for all i.

Profit and Myerson mechanism

Our goal will be to study the design of truthful mechanisms that maximize the seller's *profit*, defined as $\sum_i p_i$. We will also use the notation $\mathcal{M}^{(N,\mathcal{S})}(\mathbf{v})$ to denote the profit of the mechanism $\mathcal{M}^{(N,\mathcal{S})}$ when N is the set of agents, their valuation profile is $\mathbf{v}$, and the set system is $\mathcal{S}$.

We begin by reviewing the Bayesian approach, where agents' values are drawn independently from publicly known (but not necessarily identical) distributions, and Myerson's optimal mechanism for maximizing the auctioneer's *expected* profit over values drawn from these distributions.

Assume that for each agent i, her value v_i is drawn independently from a publicly known distribution F_i. We will need the following definition:

Definition 4. *The* virtual value *of an agent with value $v \sim F$ is*

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

Myerson proved the following:

Lemma 5 (Myerson's payment identity)**.** *For any truthful mechanism with allocation and payment rules* $\mathbf{x}$ *and* $\mathbf{p}$*, and any valuations* $\mathbf{v}_{-i}$ *of other agents besides* i*, we have*

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\mathbf{v})\right] = \mathbb{E}_{v_i \sim F_i}\left[\phi_{F_i}(v_i) x_i(\mathbf{v})\right].$$

We define the *virtual surplus* of an outcome to be the sum of virtual values of the winners. The following theorem, which is an immediate consequence of Lemma 5 is key to designing optimal truthful mechanisms.

Theorem 6. *The expected profit of a truthful mechanism is equal to the expected virtual surplus of its outcome. Formally, for any mechanism* $\mathcal{M}$ *whose allocation rule is* $\mathbf{x}$ *and which is truthful, we have*

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}\left[\mathcal{M}\left(\mathbf{v}\right)\right] = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}\left[\sum_i x_i(\mathbf{v}) \phi_{F_i}(v_i)\right].$$

By Theorem 6, to maximize the expected profit, we should find a monotone allocation rule that maximizes the expected virtual surplus of the outcome. The monotonicity constraint is important because we would like to design a truthful mechanism. The obvious way to maximize expected virtual surplus is to maximize it point-wise. If this allocation rule turns out to be monotone, we are done. To this end, we review the definition of a regular distribution.

Definition 7. *A probability distribution, given by density* F*, is* regular *if* $\phi_F(v)$ *is monotone non-decreasing.*

An even more restrictive definition is that of an MHR destribution:

Definition 8. *A probability distribution, given by density* F*, satisfies the monotone hazard rate condition (MHR) if* $f(v)/(1 - F(v)$ *is monotone non-decreasing.*

Of course every MHR distribution is regular. Many commonly encountered distributions, such as Gaussian and exponential distributions are MHR. Regular distributions are an even broader class, including many heavy-tailed distributions.

Lemma 9. *If each* ϕ_{F_i} *is monotone non-decreasing (i.e.,* F_i *is regular for all* i*), then the allocation rule that maximizes virtual surplus point-wise is monotone.*

The combination of Theorem 6 and Lemma 9 shows that if the valuations of the agents are drawn from regular distributions then by choosing the allocation that maximizes virtual surplus pointwise we get an optimal (for profit), truthful mechanism. This leads to

Definition 10 (Myerson's mechanism for regular distributions (Mye))**.** *The Myerson mechanism for regular distributions is defined by the following steps:*

1. *Solicit a bid vector* $\mathbf{b}$ *from the agents.*
2. *Compute the virtual value* b'_i *of each agent, where* $b'_i = \phi_{F_i}(b_i)$*, and choose the set of winners* W *such that* W *is feasible and* $\sum_{i \in W} b'_i$ *is maximized.*
3. *Serve the agents in* W *and charge each winner* i *her threshold bid.*

Theorem 11. *If* $\mathbf{F}$ *is a product of regular distributions,* Mye *is truthful and maximizes the auctioneer's expected profit over values drawn from* $\mathbf{F}$*. In other words, for any other truthful mechanism* $\mathcal{M}$*, we have*

$$\mathbb{E}_{\mathbf{v} \sim \mathbf{F}}\left[\mathsf{Mye}\left(\mathbf{v}\right)\right] \geq \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}\left[\mathcal{M}\left(\mathbf{v}\right)\right]$$

Note that the Myerson mechanism can be generalized to deal with irregular distributions as well. In the extension, virtual values and surpluses are replaced by their *ironed* counterparts. We will use the extension in this paper, but do not review it in detail. Readers are referred to [29] for the details.

Benchmarks

Our goal will be to study prior-free mechanism design using a competitive approach, and for this purpose, we need to formally define the benchmarks we will be considering. One technical detail that we must dispense with right away when discussing benchmarks is the fact that no mechanism can obtain a revenue comparable to the $\max_i v_i$ on every instance [15,16]. Thus any benchmark that is lower bounded by $\max_i v_i$ is a priori impossible to compete with. The standard way to resolve this issue in the literature [15, 16], is to first "normalize" the valuation profiles by replacing the largest value of any agent by the second largest value, then define benchmarks and analyze mechanisms on normalized valuation profiles.

Definition 12. *Given a valuation profile* $\mathbf{v}$*, the* normalized valuation profile $\mathbf{v}^{(2)}$ *is obtained by replacing the largest entry by the second largest entry. Formally, let* $i_1 = \text{argmax}_i v_i$ *and* $i_2 = \text{argmax}_{i \neq i_1} v_i$*. Then* $\mathbf{v}^{(2)}$ *is defined by* $v_{i_1}^{(2)} = v_{i_2}$ *and* $v_i^{(2)} = v_i$ *for all* $i \neq i_1$*.*

We now define two of the benchmarks that we will be studying in this paper. A third will be introduced in Section 4. The first benchmark is the optimal Bayesian optimal benchmark introduced by Hartline and Roughgarden [22].

Definition 13 (optimal Bayesian optimal benchmark)**.** *The Optimal Bayesian Optimal Benchmark is a function that maps instances to real numbers and is defined by*

$$\mathsf{OBO}\left(\mathbf{v}\right) = \sup_{F : v_i \in S(F) \forall i} \mathsf{Mye}_{F^n}\left(\mathbf{v}^{(2)}\right)$$

where $S\left(F\right)$ *denotes the support of a distribution* F*.*

We note that this benchmark can be slightly weakened by requiring the class of distributions over which the supremum is taken to satisfy some extra conditions. For example, we may require the distribution F in the above definition to be regular or satisfy the MHR condition, resulting in the profit benchmark OBO_{reg} and OBO_{mhr} respectively.

We also emphasize that OBO provides a *point-wise* upper bound on the profit of any i.i.d. Bayesian optimal mechanism on a normalized valuation profile.

The second benchmark we will consider is the best fixed price profit benchmark [15, 16]. We generalize this benchmark to general downward-closed setting

Definition 14 (Best fixed price profit benchmark)**.** *Given an environment* $(N, \mathcal{S})$*, the best fixed price profit benchmark for this environment is*

$$\mathsf{BFP}^{(N,\mathcal{S})}(\mathbf{v}) = \sup_{r \in R}\left(r \cdot \max_{S \in \mathcal{S} : v_i^{(2)} \geq r \forall i \in S} |S|\right)$$

We will usually suppress the superscript $(N, \mathcal{S})$ and assume that the relevant set system is clear.

3 Lowerbound for OBO

In this section, we show that no prior-free, truthful mechanism can be constant competitive to the strong benchmark OBO even in the simple 2-exclusive market setting.

Theorem 15. *For any truthful in expectation mechanism* $\mathcal{M}^{(N,\mathcal{S})}$*, there exists a valuation profile* $\mathbf{v}$ *such that the revenue of* $\mathcal{M}^{(N,\mathcal{S})}$ *is at most* $O\left(\frac{1}{\log n}\mathsf{OBO}\left(\mathbf{v}\right)\right)$*.*

In fact, as the proof shall show, this theorem holds even if we replace OBO by the weaker benchmark OBO_{mhr}.

Before discussing the idea behind the proof, we introduce some notation and terminology. For an instance $(M_1 \cup M_2, \mathcal{S}, \mathbf{v})$, the *surpluses* of the markets are defined by $V_1 = \sum_{i \in M_1} v_i$ and $V_2 = \sum_{i \in M_2} v_i$. Similarly, the *virtual surpluses* of the markets, according to some symmetric product distribution $\mathbf{F} = F^n$ are $W_1 = \sum_{i \in M_1} \max\{\phi_F(v_i), 0\}$ and $W_2 = \sum_{i \in M_2} \max\{\phi_F(v_i), 0\}$, respectively. We will only work with regular distributions in this section (in fact, only MHR), therefore, we assume that Mye sells to the market with higher virtual surplus, breaking ties arbitrarily .

3.1 Proof idea

We prove the lower bound using the probabilistic method. To this end, we will construct a distribution $\mathbf{D}$ over valuation profiles such that the expected revenue of any truthful mechanism is at most a constant while the expected value of OBO is $\Omega(\log n)$. This shows that every prior-free, truthful mechanism fails to compete with OBO on some valuation profile in the support of $\mathbf{D}$.

To understand the required properties of such a distribution $\mathbf{D}$, first observe that Mye obtains the most revenue when the virtual surpluses of the two markets are close to each other. In fact, if we can make their virtual surpluses equal, each winning agent will be charged an amount equal to her value, thus Mye's revenue is the surplus of the winning market (at least those agents with nonnegative virtual values).

The second observation is that we can think of virtual values as "shifted values" – the virtual value of an agent is just its value shifted down. Hence, starting with two markets with unbalanced surpluses, we can exploit these shifts to balance the virtual surpluses of the two markets. In fact, for each $\mathbf{v}$ in the distribution $\mathbf{D}$, we can use a different prior distribution F in order to balance these virtual surpluses. Herein lies the power of the benchmark.

Exponential distributions with different parameters are the easiest distributions to use for this purpose. In fact, if F is an exponential distribution with parameter λ then it is easy to verify that for any v, $v - \phi_F(v) = \lambda$. Hence, if we make sure that the market with larger virtual surplus has more agents than the smaller one, then the larger market will be shifted by a larger amount. Moreover, the precise difference between the shifts of the two markets can be tuned using the parameter of F. To be precise, we will set the number of agents m_1 in the first market to be $2n/3$ and the number of agents m_2 in the second market to be $n/3$. (It is also easy to modify the proof so that both markets have the same number of agents, but we do not bother with that.) The difference between the shifts is then $\lambda n/3$. Assuming that $V_1 > V_2$, choosing $\lambda = 3(V_1 - V_2)/n$ makes the two virtual surpluses equal.

However, we are not done. Recall that Mye does not allocate to agents with negative virtual values. Hence, if λ is bigger than an agent's value, the shift of that agent is not counted. To avoid this, we have to make sure that λ is at most the smallest value of any agent, which implies that we need V_1 and V_2 to be sufficient close to each other. One way to assure this is to choose the value of each agent in M_1 uniformly from a distribution D_1 and the value of each agent in M_2 from a distribution D_2 such that the expectations of V_1 and V_2 are the same and the variances are small. In that case, the distributions of V_1 and V_2 are approximately normal distributions with the same mean and small variance. This can be done to guarantee that V_1 and V_2 are close with at least some constant probability as needed.

Another requirement of $\mathbf{D}$ is that the expectation of BFP should be small compared to $\mathbb{E}[V_1]$, since we know that prior-free truthful mechanisms can extract a constant fraction of BFP of any particular market. Finally, we would like the expected values of V_1 and V_2 to be large, so that Mye makes good profit.

In summary, $\mathbf{D}$ should exhibit the following properties

1. $\mathbf{D} = D_1^{2n/3} \times D_2^{n/3}$

2. $\mathbb{E}_{v \sim D_1}[v] = \frac{1}{2}\mathbb{E}_{v' \sim D_2}[v']$

3. $\mathbb{E}[\mathsf{BFP}(\mathbf{v})]$ is small.

4. $\mathbb{E}[V_1]$ is large compared to $\mathbb{E}[\mathsf{BFP}(\mathbf{v})]$.

The natural choices of D_1 and D_2 to satisfy the last two requirements are uniform distributions over some harmonic sets (variants on the "equal-revenue distribution") – so that BFP is a constant while V_1 and V_2 are on the order of $\log n$. We will show in the formal proof that these natural choices work.

3.2 Formal proof

Following the informal discussion in the previous subsection, we define the distribution $\mathbf{D}$ as follows. Let $n = 3k+2, m_1 = 2k+1$ and $m_2 = k+1$ for some parameter k that will be chosen later. Let $i^* \in M_1$ and $j^* \in M_2$ be two special agents. We set $v_{i^*} = v_{j^*} = 1$. For $i \in M_1 \backslash \{i^*\}$, v_i is drawn independently and uniformly at random from the set

$$T_1 = \left\{ \frac{1}{t+1}, \frac{1}{t+2}, \ldots, \frac{1}{t+2k} \right\}$$

and for $j \in M_2 \backslash \{j^*\}$, v_j is drawn independently and uniformly at random from the set

$$T_2 = \left\{ \frac{a}{t+1}, \frac{a}{t+2}, \ldots, \frac{a}{t+k} \right\}$$

for some parameters a and t that will be described later.

We will choose a and t so that $\mathbb{E}[V_1] = \mathbb{E}[V_2]$. Then, the probability that V_1 and V_2 are close is at least the probability that both random variables are close to their expectations.

Intuitively, this probability should be large by the Central Limit Theorem (CLT). We bound the error term of the CLT in Lemma 16. The lower bound is then completed by bounding the expected revenue of any truthful in expectation mechanism. (Lemma 20).

Thus, the first lemma consists of formalizing the application of the CLT to our setting. This is proven using the Berry-Esseen bound in Appendix A.

Lemma 16. *For all $d > 0$ and $0 \le \delta \le 1/4$, there exists $T > 0$ such that for all $r > T$, there is an $N > 0$ such that for all $m > N$ and $1 \le b \le 2$, if X_i's are i.i.d random variable where each X_i is drawn uniformly from $\left\{ \frac{b}{r+1}, \frac{b}{r+2}, \ldots, \frac{b}{r+m} \right\}$ and $X = \sum_{i=1}^{m} X_i$ then*

$$\mathbf{Pr}[\mathbb{E}[X] \le X \le \mathbb{E}[X] + d] \ge \delta, \quad \textit{and} \quad \mathbf{Pr}[\mathbb{E}[X] - d \le X \le \mathbb{E}[X]] \ge \delta$$

Let T and N be the numbers satisfying Lemma 16 for $d = 1/6$ and $\delta = 1/4$, we set $t = T$, choose $k > N$ be sufficiently large, and set $a = \frac{H_{2n+t} - H_t}{H_{n+t} - H_t} \le 2$. (Note in the following that the only quantities going to infinity are n and k.) With these parameters, we can bound $\mathbb{E}_{\mathbf{v} \sim \mathbf{D}}[\mathsf{OBO}(\mathbf{v})]$ as follows.

Lemma 17. *For a valuation profile $\mathbf{v}$ drawn from $\mathbf{D}$, we have*

$$\mathbf{Pr}_{\mathbf{v}}\left[0 \le V_1 - V_2 \le \frac{k}{2k+t} \text{ and } V_1 = \Omega(\log n) \right] \ge \frac{1}{16}.$$

Proof. Let $V_1' = \sum_{i\in M_1\setminus\{i^*\}} v_i = V_1 - 1$ and $V_2' = \sum_{j\in M_2\setminus\{j^*\}} v_j = V_2 - 1$. Since $\mathbb{E}[V_1'] = \mathbb{E}[V_2'] = \Omega(\log n)$ (since $n = 3k+2$), the above probability is at least

$$\begin{aligned}
&\mathbf{Pr}\left[\mathbb{E}[V_1'] \le V_1' \le \mathbb{E}[V_1'] + \frac{k}{4k+2t}\right] \cdot \mathbf{Pr}\left[\mathbb{E}[V_2'] - \frac{k}{4k+2t} \le V_2' \le \mathbb{E}[V_2']\right] \\
&\ge \mathbf{Pr}\left[\mathbb{E}[V_1'] \le V_1 \le \mathbb{E}[V_1'] + \frac{1}{6}\right] \cdot \mathbf{Pr}\left[\mathbb{E}[V_2'] - \frac{1}{6} \le V_2' \le \mathbb{E}[V_2']\right] \qquad (1)\\
&\ge \frac{1}{4}\cdot\frac{1}{4} = \frac{1}{16} \qquad (2)
\end{aligned}$$

where (1) follows from the fact that t is a constant and k is large, and (2) follows from Lemma 16. □

Lemma 18. *If* $0 \le V_1 - V_2 \le \frac{k}{2k+t}$ *then* $\mathsf{OBO}(\mathbf{v}) \ge V_2$.

Proof. Let $\lambda = \frac{V_1 - V_2}{k} \le \frac{1}{2k+t} \le \min_i v_i$ and let F be the exponential distribution with parameter λ. Then $v_i \in S(F)$ and $\phi_F(v_i) = v_i - \lambda \ge 0$ for all $i \in N$. Furthermore, $\sum_{i\in M_1} \phi_F(v_i) = \sum_{j\in M_2} \phi_F(v_j)$. Hence, for $\mathbf{F} = F^n$, $\mathsf{Mye}_{\mathbf{F}}$ charges the winning market its full surplus, which is at least V_2. □

Combining Lemma 17 and 18 yields

Corollary 19. *The expected value of* $\mathsf{OBO}(\mathbf{v})$ *over* $\mathbf{v}$ *drawn from* $\mathbf{D}$ *is* $\Omega(\log n)$.

Next, we bound the expected revenue of any truthful mechanism $\mathcal{M}^{(N,\mathcal{S})}$ (proof in Appendix A).

Lemma 20. *Let* $\mathcal{M}^{(N,\mathcal{S})} = (\mathbf{x}, \mathbf{p})$ *be a truthful in expectation mechanism and* v_i *is drawn uniformly from the set* $\left\{\frac{1}{r+1}, \frac{1}{r+2}, \ldots, \frac{1}{r+m}\right\}$. *Then for any* $\mathbf{v}_{-i}$, $\mathbb{E}_{v_i}[p_i(\mathbf{v})] \le 1/m$.

Corollary 21. *The expected revenue of any truthful in expectation mechanism over valuation profile drawn from the distribution* $\mathbf{D}$ *is at most* 5.

Proof. The expected amount that $\mathcal{M}^{(N,\mathcal{S})}$ charges i^* and j^* is at most $v_{i*} + v_{j*} = 2$. By Lemma 20, the expected revenue $\mathcal{M}^{(N,\mathcal{S})}$ gets from $M_1\setminus\{i^*\}$ is at most 1. Similarly, the expected revenue $\mathcal{M}^{(N,\mathcal{S})}$ gets from $M_2\setminus\{j^*\}$ is at most $a < 2$. Hence, the expected revenue of $\mathcal{M}^{(N,\mathcal{S})}$ is at most 5. □

Combining Corollaries 19 and 21 yields Theorem 15.

4 Comparing sOBO and BFP

As we saw in Section 3, the OBO benchmark, while a very desirable target, is too strong for non-matroid environments: no prior-free truthful mechanism can obtain a constant fraction of it even in the simple exclusive market setting with two markets. Therefore, we consider alternative benchmarks.

To this end, observe that the requirement that a prior-free benchmark upper bounds the profit of any i.i.d. Bayesian optimal mechanism *pointwise* is perhaps too strong, since this allows the benchmark to choose a different Bayesian optimal mechanism for each input vector of values and placement of the values into the set system (or in our case, partition of the values between the two markets). Indeed, this was the key fact that we took advantage of in the proof of Theorem 15. Moreover, Bayesian optimal mechanisms are only optimal in expectation. Thus, if we have a benchmark that upper bounds all i.i.d. Bayesian optimal mechanisms *in expectation*, and a mechanism that approximates it, then this mechanism still approximates the appropriate Bayesian optimal mechanism in expectation when the agents' values are actually i.i.d.

The second observation is that if the agents' values are actually i.i.d. then all permutations of a valuation profile are equally likely. Therefore, if a benchmark upper bounds a Bayesian optimal mechanism in expectation over random permuations of the valuation profile then the former upper bounds the latter in expectation over valuation profiles drawn from any symmetric product distribution. This motivates us to look at the following benchmark (discussed in [24]).

Definition 22 (symmetrized optimal Bayesian optimal benchmark)**.** *Given a single parameter environment defined by a set of agents N and a set system $\mathcal{S}$, the symmetrized optimal Bayesian optimal benchmark is defined by*

$$\mathsf{sOBO}\left(\mathbf{v}\right) = \sup_{\mathbf{F}=F^n:\mathbf{v}\in S(\mathbf{F})} \mathbb{E}_{\pi\in\Pi_n}\left[\mathsf{Mye}_{\mathbf{F}}\left(\pi\left(\mathbf{v}^{(2)}\right)\right)\right].$$

where Π_n denotes the set of all permutations of the vector $(1, 2, \ldots, n)$ and $\pi(\mathbf{v})$ denotes the vector $(v_{\pi_1}, v_{\pi_2}, \ldots, v_{\pi_n})$.

Similarly to above, we can also restrict the class of the distributions in this definition to obtain weaker benchmarks sOBO_{reg} and sOBO_{mhr} respectively.

We can see that sOBO upper bounds the best i.i.d. Bayesian optimal mechanism point-wise, as long as the values of the agents are randomly permuted before that optimal mechanism is applied.

The main theorem of this section is that for two exclusive markets, sOBO and BFP are equivalent up to a constant factor.

Theorem 23. *For the two exclusive market environment, where the first market has m_1 agents and the second market has $m_2 = n - m_1$ agents,*

$$\mathsf{sOBO}\left(\mathbf{v}\right) = O\left(\mathsf{BFP}\left(\mathbf{v}^{(2)}\right)\right)$$

for any valuation profile $\mathbf{v}$.

As there are mechanisms that compete with BFP (even in much more general environments), this suggests that sOBO is a promising and strong benchmark for prior-free truthful mechanism design. In particular, it works beautifully for two exclusive markets.

Proof idea

We outline the structure of the proof that for two exclusive markets $\mathsf{sOBO} = O\left(\mathsf{BFP}\left(\mathbf{v}^2\right)\right)$ for any valuation profile $\mathbf{v}$. All the details of this proof can be found in Appendix B.

To this end, we first recall that the optimal mechanism for general distributions works exactly like that for regular distributions (see Section 2 for more details), except that it uses ironed virtual values and surpluses [29]. For the purpose of this section and Appendix B, we only need two easy-to-verify properties of the ironed virtual value function $\bar{\phi}_F$: (i) $\bar{\phi}_F$ is monotone and (ii) for any value u in the domain of F, $\bar{\phi}_F(u) \leq u$.

To prove Theorem 23, consider a valuation profile $\mathbf{v}$ and a permutation $\mathbf{u}$ of it where $u_1 \geq u_2 \geq \cdots \geq u_n$. Let $a = \max_{1\leq i\leq n} iu_i$. Then $\mathsf{BFP}\left(\mathbf{v}\right) \geq a/2$. Therefore, it suffices to show that for any distribution $\mathbf{F} = F^n$

$$\mathbb{E}_{\pi}\left[\mathsf{Mye}_{\mathbf{F}}\left(\pi(\mathbf{u})\right)\right] = O\left(a\right). \tag{3}$$

Let $w_i = \bar{\phi}_F(u_i)$. Then $w_i \leq u_i \leq a/i$ and $w_1 \geq w_2 \geq \cdots \geq w_n$. For a particular permutation π, let $W_1(\pi)$ and $W_2(\pi)$ be the ironed virtual surpluses of the two markets M_1 and M_2 respectively, and $\Delta(\pi) = |W_1 - W_2|$. (We will generally drop the argument of π and assume it is understood when we write W_1, W_2 and Δ.)

Consider the case where $w_n \geq 0$. (The case $w_n < 0$ is a simple extension.) In this case, $\mathsf{Mye}_{\mathbf{F}}$ serves all agents in the market with the higher ironed virtual surplus. By Lemma 2, it charges each winning agent whose ironed virtual value is at most Δ an amount of $\bar{\phi}_F^{-1}(0) \leq u_n$. This means that the total revenue $\mathsf{Mye}_{\mathbf{F}}$ earns from such agents is at most $nu_n \leq a$. Hence, we only need to bound the expected revenue $\mathsf{Mye}_{\mathbf{F}}$ earns from the winning agents whose virtual values are greater than Δ.

To this end, for each i, we bound $\mathbf{Pr}[w_i > \Delta]$ over the random permutation π. If we can show that this probability is at most $1/i^\delta$ for some constant $\delta > 0$, then the expected revenue $\mathsf{Mye}_{\mathbf{F}}$ earns from the agents whose virtual values are greater than Δ is at most $\sum_{i=1}^n w_i/i^\delta \leq \sum_{i=1}^n u_i/i^\delta \leq \sum_{i=1}^n a/i^{\delta+1} = O(a)$ as required. Thus, the proof boils down to proving that $\mathbf{Pr}[w_i > \Delta] < 1/i^\delta$ for some $\delta > 0$.

We use an idea due to Paul Erdős [9]. To illustrate his clever idea, consider a different but related model in which the size of the markets are not fixed and instead of randomly permuting the valuation profile, we independently put each value into each market with probability 1/2. Consider a particular value u_i and fix the placements of all values u_j where $j > i$. Each placement of the values $u_1, u_2, \ldots, u_i$ is characterized by the subset S of them which are put into the market with the larger virtual surplus. We say that a set S is *bad* if it corresponds to a placement in which $\Delta < w_i$. Now consider two sets S_1 and S_2 such that $S_1 \subset S_2$. If S_1 is bad than S_2 cannot be bad, because when going from S_1 to S_2, Δ increases by at least $2w_i$. Hence, the set of placements where $\Delta < w_i$ corresponds to a collection of subsets of $\{u_1, u_2, \ldots, u_i\}$ where no subset contains another. In other words, this set of placements is an anti-chain in the partial order defined by set inclusion. The size of this collection is, by Sperner's lemma [30], at most $\binom{i}{i/2}$. Hence, the probability in question is at most $\binom{i}{i/2}/2^n = O(1/\sqrt{i})$.

The problem with applying this argument to the fixed-size market setting is that once all the values smaller than u_i are placed, there are a fixed number of "slots" left for the remaining elements in each market, and *all* placements of the remaining elements form an anti-chain in this partial order. Hence, the fact that we only count the bad placements does not yield any reduction in the probability of the bad event. To work around this problem, we first observe that for any fixed i^*, the revenue $\mathsf{Mye}_{\mathbf{F}}$ gets from the agents whose values are in the set $\{u_{i^*+1}, u_{i^*+2} \ldots u_n\}$ is at most nu_{i^*} (the exact value of i^* will be decided later). Hence, if $i^* = \Omega(n)$, the revenue $\mathsf{Mye}_{\mathbf{F}}$ gets from these agents is at most $O(a)$.

Now consider some $i \leq i^*$. If $\Delta \geq w_i - w_{i^*}$ then the agent with value u_i also pays at most u_{i^*}, hence the expected revenue from such agents is also bounded by $O(a)$. Hence, instead of bounding $\mathbf{Pr}[\Delta < w_i]$, it suffices to bound $\mathbf{Pr}[\Delta < w_i - w_{i^*}]$, which we do by fixing the placements of the values between u_i and u_{i^*}, and considering all possible placements of the remaining elements. We define a different partial order in which bad placements form an antichain, and prove a simple variant of Sperner's lemma, using the same structure as the one used in the proof due to Lubell [27] to complete our proof.

In a bit more detail, we partition each placement P, where $P = \{j : \pi_j \in M_1\}$, into three sets S, R and T where $S = P \cap \{1, 2, \ldots, i\}$, $R = P \cap \{i+1, i+2, \ldots, i^*\}$ and $T = P \cap \{i^*+1, i^*+2, \ldots, n\}$.

For a fixed R, define the following partial order:

Definition 24. $(S_1, R, T_1) \prec_i (S_2, R, T_2)$ *if and only if* $S_1 \subset S_2$ *and* $T_2 \subset T_1$.

We say that a placement (S, R, T) is *bad* if $\Delta(S, R, T) < w_i - w_i^*$. It is straightforward to prove the following lemma:

Lemma 25. *If two placements* (S_1, R, T_1) *and* (S_2, R, T_2) *are both bad then neither* $(S_1, R, T_1) \prec_i (S_2, R, T_2)$ *nor* $(S_2, R, T_2) \prec_i (S_1, R, T_1)$.

Figure 4 gives a "proof by picture" of this lemma.

For each fixed R, let $\mathcal{A}_R$ be the set of bad placements (S, R, T). We prove the following simple analogue of Sperner's lemma in this setting.

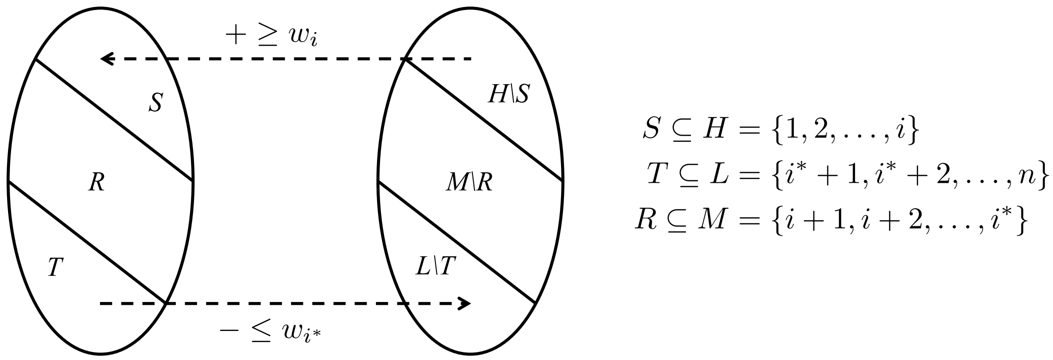

Figure 1: A proof by picture for Lemma 25. Starting with a bad configuration (S, R, T), if we construct a new configuration $(S', R, T') \succ (S, R, T)$ by moving some elements of $H \backslash S$ into S and moving the same number of elements of $L \backslash T$ into T then the change in $V_1 - V_2$ is more than twice $v_i - v_{i^*}$. Hence the new configuration cannot be bad.

Lemma 26. $|\mathcal{A}_R| \leq \max_{s+t=m_1-r} \binom{n-i^*}{t}\binom{i}{s}$ *where* $r = |R|$.

Let $\mathcal{A}$ be the set of all bad placements. Then Lemma 26 implies

$$|\mathcal{A}| = \sum_R |\mathcal{A}_R| \leq \sum_{r=0}^{p} \left(\binom{i^* - i}{r} \cdot \max_{s+t=m_1-r} \binom{i}{s}\binom{n - i^*}{t} \right) = O\left(\frac{1}{i^{\delta/2}}\right)\binom{n}{m_i}, \tag{4}$$

where p is the minimum of the size of the smaller market and $i^* - i$, and δ is some small constant. Readers are referred to Appendix B for detailed definitions of these quantities and the proof of the last equality.

Combining this with the fact that

$$\mathbf{Pr}_\pi \left[\Delta(\pi) < w_i - w_i^*\right] \leq \frac{|\mathcal{A}|}{\text{number of placements}} = \frac{|\mathcal{A}|}{\binom{n}{m_1}} \tag{5}$$

is sufficient to complete the proof of Theorem 23. As mentioned above, all the details can be found in Appendix B.

5 Approximating BFP

We next present a mechanism that is competitive to BFP in exclusive markets and another one that is competitive to BFP in all downward-closed environments, albeit with a worse competitive ratio. Both mechanisms are based on a truthful profit extractor for BFP in downward-closed environments. A *profit extractor* for a profit benchmark B is a truthful mechanism that takes an extra parameter T and obtains revenue at least T whenever $T \leq B(\mathbf{v})$. Given a profit extractor, the immediate idea to design a competitive mechanism is to first find a way to estimate BFP $(\mathbf{v})$ truthfully and then run the profit extractor with this estimation being the target. (Of course, the combination of these two steps may break truthfulness, and has to be handled carefully.)

We first describe the profit extractor in Subsection 5.1. Then, in Subsections 5.2 and 5.3, we describe the competitive mechanisms in exclusive markets and downward-closed environments respectively. As a corollary, we obtain

Corollary 27. *For two exclusive markets, there is a prior-free mechanism that is competitive to* sOBO.

5.1 A profit extractor for BFP in downward-closed environments

As in [15], the profit extractor $\mathsf{ProfitExtract}_T$ is constructed using the cost sharing mechanism of Moulin and Shenker [28]. It works as follows:

1. Solicit a bid vector $\mathbf{b}$ from the agents.
2. Choose the largest feasible set S that can "share" the target T, i.e., the largest feasible set S such that $b_i \geq T/|S|$ for all $i \in S$, as the winning set, breaking ties according to a fixed pre-specified ordering on the sets.
3. Charge each agent in S an amount of $T/|S|$.

Lemma 28. *For any downward-closed environment and any $T \geq 0$, $\mathsf{ProfitExtract}_T$ is truthful and the allocation is feasible. Moreover, $\mathsf{ProfitExtract}_T(\mathbf{v}) \geq T$ whenever $\mathsf{BFP}(\mathbf{v}) \geq T$.*

Proof. Feasibility is immediate. To see that the mechanism is truthful, fix the bids of all agents except for agent i, and observe that this determines a threshold bid for i. Finally, since the mechanism is truthful, the second part of the lemma is obvious. □

5.2 A competitive mechanism based on random sampling for exclusive markets

Using ProfitExtract, we design the Iterative Sampling Cost Sharing (ICS) mechanism (Figure 2) for exclusive markets. Our mechanism is inspired by the Sampling Cost Sharing (SCS) mechanism [12, 15] for digital good settings. In a digital good environment, SCS works as follows:

1. Randomly divide the set of agents into a sample set S and a market M,
2. Compute the best fixed price profit f and f' in S and M respectively,
3. Try to extract f from M and f' from S simultaneously.

Noting that each of the exclusive markets is similar to a digital good environments, we attempt to simulate SCS in one of them. Ideally, we would like to simulate SCS in the market with the highest BFP. However, the sampling may mislead us into using the wrong market. The trick to circumvent this problem is to sample all the markets, compute the BFP in each sample and attempt to extract these profits in the decreasing order. This ensures that in the right market, we still have a constant probability of doing well. Since we use the sampled BFP to determine the order of the markets to be considered, the sampled agents cannot be allowed to win, otherwise the mechanism may not be truthful. Therefore, we can only simulate "half" of SCS. This leads to a factor of 2 loss in the competitive ratio, as Theorem 30 shows.

Lemma 29. ICS *is truthful.*

Proof. Consider an agent $j \in M_i$. If $j \in S_i$, changing her bid does not help j. If $j \in M_i'$, changing her bid does not change the order in which the effective markets are considered in step 3, hence does not change the chance of M_i' being considered. Finally, if $j \in M_i'$ and M_i' ends up being considered in Step 3, changing her bid does not help j because of the truthfulness of $\mathsf{ProfitExtract}_{f_i}$. □

Lemma 30. *For any $\mathbf{v}$, the expected revenue of the Iterative Cost Sharing mechanism, where the expectation is taken over the randomness of the mechanism, is at least $1/8$ times $\mathsf{BFP}\left(\mathbf{v}^{(2)}\right)$ if the latter is obtained by serving at least 2 agents.*

Iteratıve Cost Sharing (ICS) Mechanism

1. Ignore all markets of size 1.
2. Partition each remaining market M_i into a random sample S_i and an effective market $M'_i = M_i \setminus S_i$ by independently putting each agent into S_i with probability 1/2.
3. Compute the best fixed price profit f_i of each S_i, i.e.

$$f_i = \max_{j \in S_i} v_j \, |\{j' \in S_i : v_{j'} \geq v_j\}|$$

4. Go through the markets in decreasing order of f_i and run $\mathsf{ProfitExtract}_{f_i}$ on M'_i until some agents are served.

Figure 2: A truthful mechanism that approximates BFP for k exclusive markets when the best fixed price profit is obtained by serving at least 2 agents.

Proof. Let M_i be the market with the largest best fixed price profit (that includes at least 2 agents). Because markets are considered in order of decreasing f_j, the expected profit ICS gets is at least the expected profit it gets from market M_i, which is $\mathbb{E}[f_i | f'_i \geq f_j] \mathbf{Pr}[f'_i \geq f_j]$ where f'_i is the best fixed price profit of M_i.

Fiat *et al* [12, 15] analyzed a mechanism that simultaneously tried to extract f_i from M'_i and f'_i from S_i and showed that it obtained a profit of at least $\frac{1}{4}\mathsf{BFP}(M_i)$. Our lemma is a direct corollary of this result. □

Again, ICS only approximates BFP if the latter is obtained by serving at least 2 agents. Therefore, it alone is not enough to approximate BFP in all cases[3]. However, this problem is easy to fix. Observe that if BFP is obtained by serving one agent, then the second price auction (SPA), which serves only the agent with the highest bid and charges her the second highest bid, obtains the same revenue on normalized valuation profiles. Therefore, we obtain a competitive mechanism by randomly choosing either ICS or the SPA to run with probability α and $1 - \alpha$ respectively. By choosing $\alpha = 8/9$, we obtained a mechanism that approximates BFP within a factor of 9.

5.3 A competitive mechanism based on consensus estimate for downward-closed environments

For general downward-closed environments, random sampling based mechanisms are difficult to analyze. Thus, we use a different technique, namely consensus estimate with cross checking, developed first by Goldberg and Hartline [14] and then, in much greater generality, by Ha and Hartline [17]. We first review this technique and its application to digital good setting.

Definition 31. *For any given $v > 0$ and $\rho > 1$, a function g is a ρ-consensus estimate of v if*

- *g is a consensus: for any w such that $v/\rho \leq w \leq v$, we have $g(w) = g(v)$.*
- *$g(v)$ is a non-trivial lower bound of v, i.e. $v/c < g(v) \leq v$, for some $c > 1$.*

Goldberg and Hartline [14] prove the following result

Theorem 32. *There is a distribution $\mathcal{G}$, parameterized by a parameter c, over functions such that*

[3]Even if the valuation profile is normalized, BFP can still be achieved by serving only one agent if the agents with the highest values are in different markets.

Downward-closed Consensus Revenue Estimate (DCCORE) **Mechanism**

1. Solicit bid vector **b** from the agents.

2. For each pair of agents i and j, compute

$$T_{-ij} = \sup_r \max_{S \text{ feasible}, i,j \notin S} r \cdot |\{\ell \in S : b_\ell \geq r\}|.$$

3. Draw a function g from the distribution $\mathcal{G}$ in Theorem 32 with $\rho = 3$ and c appropriately chosen.

4. For each agent i, if $g(T_{-ij}) = g(T_{-ik}) = T$ for all j and k, run $\mathsf{ProfitExtract}_T$ on i, i.e., serve i and charge her a price p if and only if $\mathsf{ProfitExtract}_T$ does so.

Figure 3: A truthful mechanism that is competitive to BFP when BFP serves at least 3 agents.

- *for any v and ρ, a function g drawn from G is a ρ-consensus estimate of v with probability at least $1 - \log_c \rho$, and*
- *$g(v) \geq v/c$ for all v.*

With this, Goldberg and Hartline design a competitive mechanism for the digital good settings as follows. Assume that the best fixed price profit T is obtained by serving at least k agents. If we remove the bid of an agent i and find the best fixed price profit when serving the remaining agents, we get some value T_{-i}. All the T_{-i} are within a factor of $k/(k-1)$ of each others. Now draw a function g from the distribution in Theorem 32, and for each i, offer him the same price as $\mathsf{ProfitExtract}_{g(T_{-i})}$ does. This mechanism is truthful because the price it offers each agent is independent of her bid. On the other hand, if g is a consensus estimate of T, which happens with a constant probability with an appropriate choice of c, then $g(T_{-i}) = g(T_{-j}) = \Omega(T)$ for all i and j. Therefore, in this case, the mechanism obtains a profit of $\Omega(T)$.

The reason this mechanism does not extend to all downward-closed environments is feasibility. Specifically, when g is not a consensus estimate, the mechanism may run different profit extractors on different agents, therefore offering acceptable prices to agents who do not constitute a feasible set. (This was not a problem in digital good settings because any subset of agents is feasible). To circumvent this problem, we have to make sure that the same profit extractor is used for all agents.

Ha and Hartline [17] suggest the following elegant solution. Instead of computing T_{-i}, we compute T_{-ij}, the BFP of the reduced valuation profile when both i and j are removed, for any pair of agents i and j. We say that an agent i *reaches a consensus* if for all $j, \ell \neq i$, $g(T_{-ij}) = g(T_{-i\ell})$. The mechanism only considers serving an agent who reaches a consensus. The nice property of this approach is that if i and j both reach a consensus, they reach the same consensus (which is $g(T_{-ij})$). Hence, the mechanism runs the same profit extractor on all the agents (or no profit extractor on those agents that do not reach consensus). The full mechanism DCCORE is presented in Figure 3.

It is clear that DCCORE is truthful, as the price offered to an agent (and whether she is offered a price to begin with) is independent of her bid. Moreover, it always serves a feasible set, since it actually runs the same profit extractor on the agents who reach a consensus. We argue that if BFP is obtained by serving at least 3 agents, then DCCORE obtains a constant fraction of it.

Theorem 33. *If* BFP *is obtained by serving at least* 3 *agents, then* $\mathsf{DCCORE}(\mathbf{v}) \geq \frac{2}{31}\mathsf{BFP}(\mathbf{v})$.

Proof. Since BFP is obtained by serving at least 3 agents, $T_{-ij} \geq \frac{1}{3}\mathsf{BFP}(\mathbf{v})$ for all i and j. By Theorem 32, with probability at least $1 - \log_c 3$, $g(T_{-ij}) = T \geq \frac{1}{c}\mathsf{BFP}(\mathbf{v})$ for all pairs of agents i, j, i.e. every agent reaches the same consensus. When this happens, the mechanism obtains a profit of exactly T. Therefore, the expected revenue of the mechanism is at least $\frac{1-\log_c 3}{c}\mathsf{BFP}(\mathbf{v})$. Choosing c to maximize $\frac{1-\log_c 3}{c}$ yields the theorem. □

Finally, note that DCCORE is only competitive to BFP when the latter is obtained by serving at least 3 agents. Similar to the previous subsection, if BFP is obtained by serving at most 2 agents, then the second price auction (SPA) that serves at most one agent obtains at least 1/2 of BFP (note that BFP is defined on normalized valuation profiles, so the second highest bid and the highest bid are the same). Hence, by running DCCORE with probability 31/35 and SPA with probability 4/35, we gets a truthful mechanism that is competitive to BFP with the competitive ratio of 17.5.

6 Conclusion

In this paper, we have revisited the question of benchmarks that one could try to compete with when designing truthful, prior-free mechanisms for maximizing profit. Our interest is in truthful prior-free mechanisms that can compete with a benchmark on each possible input, i.e. are maximally robust. The setting we have focused on is a fundamental and natural one that has not received much attention in the literature: exclusive markets. Our first result is that the strong benchmark proposed by Hartline and Roughgarden [22], OBO, is not achievable, even in what is arguably the simplest possible non-matroid setting – two exclusive markets. It can and has been argued [25] that OBO does not make sense in environments in which the agents do not play the same role in the set system. However, in general matroid settings, agents do not play the same role in the set system either and nonetheless the weaker benchmark OBO_{reg} is approximable by a truthful mechanism to within a constant factor. Thus, it was of interest to know whether this result generalized. It does not, even against the much weaker benchmark OBO_{mhr}.

In light of this negative result, we turned our attention to a symmetrized version of the same benchmark. The symmetrized optimal Bayesian optimal benchmark sOBO retains the nice connection to the Bayesian optimal mechanism design. In particular, any truthful mechanism $\mathcal{M}^{(N,\mathcal{S})}$ that is constant competitive with sOBO is automatically constant competitive to the i.i.d. Bayesian optimal benchmark no matter what the prior distribution is. Thus, in settings in which it is reasonable to assume that agents values are i.i.d. from some distribution, such an $\mathcal{M}^{(N,\mathcal{S})}$ is to within a constant of the Bayesian optimal, and does not depend on knowledge of the distribution. But such an $\mathcal{M}^{(N,\mathcal{S})}$ also has a robustness that the Bayesian mechanisms do not have: it has a guaranteed profit level (a constant fraction of the benchmark) no matter where the input comes from.

We were able to show that in the 2-exclusive market setting sOBO is equivalent to a much simpler benchmark, the best fixed price benchmark BFP, and that BFP is approximable to within a constant factor in this setting by a truthful mechanism. This gives a proof of concept that sOBO might be an achievable benchmark. However, even just considering BFP as a benchmark on its own right is interesting and, using an elegant idea due to Ha and Hartline, we showed that BFP is approximable in any downward-closed set system.

It is worth comparing these results to the recent work of Hartline and Yan [24] and Ha and Hartline [17]. These two papers consider a different benchmark called the Envy-Free-Optimal benchmark or EFO, which is the maximum profit that can be achieved using prices that are envy-free, i.e., such that no agent envies the outcome of another. The definition of EFO, like the definition of sOBO involves computing an expected value over permutations of the agents within the set system.[4] Hartline and Yan (and Ha and Hartline) present prior-free mechanisms whose expected profit, over random permutations of elements within the set system, is guaranteed to be Ω(EFO). They also prove that for set systems that are matroids, EFO is Ω(sOBO). The key distinction between their results and ours in terms of mechanisms is that the guarantees on the prior-free mechanism are only proven in expectation over the random permutations. In other words, they are not proving that for a fixed set

[4]Indeed, envy-freedom does not make sense unless the positions of the agents in the set system are symmetric.

of values and for a fixed placement of agents within the set system, their mechanism achieves a certain guarantee, but rather only over random permutations. Thus, while their results are very strong and the EFO benchmark is natural, interesting, and in some settings stronger than BFP (though not in the 2-exclusive market setting), our mechanisms are more robust (in the worst-case sense).

The biggest open question remaining in this paper is to understand the relationship between sOBO (or sOBO_{reg}) and BFP or, more generally, obtain a better characterization of sOBO and/or sOBO_{reg} for other settings. At this point there is strong evidence that BFP is not $\Omega(\mathsf{sOBO})$, even for matroids [25]. However, we conjecture that BFP is $\Omega(\mathsf{sOBO})$ for any exclusive market setting. (We only proved this for the case of two markets.) Moreover, it is possible that for all downward closed set systems BFP is $\Omega(\mathsf{sOBO}_{reg})$.[5] We simply do not know.

Regarding the conjecture that BFP is $\Omega(\mathsf{sOBO})$ for any exclusive market setting, the approach we took for 2 markets was to show that the probability that the sum of nonnegative ironed virtual values in the 2 markets differs by w_i is $O(1/i^\delta)$ for some $\delta > 0$. A useful step towards generalizing this proof to any number of markets would be to resolve the following interesting "voting problem": Suppose there are k candidates in an election (the markets) and n voters (the agents). Each voter has a weight w_i (the maximum of that agent's ironed virtual value and 0) and votes for candidate j with probability p_j (the fraction of agents in market j). We assume the weights are ordered so that $w_1 \geq w_2 \ldots \geq w_n$. The winning candidate is selected using weighted majority voting. What is the best upper bound that can be shown on the probability that voter i can swing the election, i.e. by changing his vote, he can change the winner of the election (ties do not count)? We are able to obtain bounds when i is large enough compared to k, but the small i case is the more interesting case and more important for our application. This problem seems of interest in its own right.

More generally, our community is still a long way from having resolved the problem of prior-free mechanism design for profit maximization in all single parameter environments, let alone the more challenging multi-parameter settings.

Acknowledgement

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions.

References

[1] Saeed Alaei, Azarakhsh Malekian, and Aravind Srinivasan. On random sampling auctions for digital goods. In *ACM Conference on Electronic Commerce*, pages 187–196, 2009.

[2] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *Proc. of the 46th IEEE Symp. on Foundations of Computer Science*, pages 605–614, 2005.

[3] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3, 2003.

[4] Andrew C. Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):122–136, 1941.

[5] K. Deshmukh, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Truthful and competitive double auctions. In *Proc. 10th European Symp. on Algorithms*. Springer, 2002.

[6] N. Devanur, J. Hartline, A. Karlin, and T. Nguyen. Prior-independent multi-parameter mechanism design. In *Workshop on Internet and Network Economics (WINE)*, 2011.

[5] It is known [23] that VCG with a single reserve price for all agents is a 2 approximation to Myerson when the priors F are MHR. Thus, another interesting benchmark to try to compete with is $\sup_r VCG_r$, where r is the reserve price (see [23].)

[7] Nikhil R. Devanur and Jason D. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *ACM Conference on Electronic Commerce*, pages 41–50, 2009.

[8] P. Dhangwatnotai, T. Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce*, 2010.

[9] Paul Erdős. On a lemma of littlewood and offord. *Bulletin Amer. Math. Soc.*, 52:898 – 902, 1945.

[10] Carl-Gustav Esseen. On the liapunoff limit of error in the theory of probability. *Arkiv for matematik, astronomi och fysik*, A28:1–19, 1942.

[11] U. Feige, A. Flaxman, J. Hartline, and R. Kleinberg. On the Competitive Ratio of the Random Sampling Auction. In *Proc. 1st Workshop on Internet and Network Economics*, pages 878–886, 2005.

[12] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Generalized competitive auctions. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 72–81. ACM Press, 2002.

[13] A. V. Goldberg and J. D. Hartline. Competitive auctions for multiple digital goods. In *Proc. 9th European Symp. on Algorithms*. Springer, 2001.

[14] A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. In *Proc. 14th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2003.

[15] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55:242–269, 2006.

[16] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.

[17] Bach Q. Ha and Jason D. Hartline. Consensus estimates, cross checking, and profit extraction. In *Proc. 23rd ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2012.

[18] M. Hajiaghayi, D. Parkes, and R. Kleinberg. Adaptive limited-supply online auctions. In *Proc. 6th ACM Conf. on Electronic Commerce*, 2004.

[19] J. Hartline. Approximation in economic design. 2012.

[20] J. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.

[21] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 175–182, 2005.

[22] J. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing*, 2008.

[23] J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*, 2009.

[24] Jason Hartline and Qiqi Yan. Envy, truth, and profit. In *ACM Conference on Electronic Commerce*, pages 243–252, 2011.

[25] Jason D. Hartline. Personal communication, 2011.

[26] Victor Korolev and Irina Shevtsova. An improvement of the berry-esseen inequality with applications to poisson and mixed poisson random sums. *Scandinavian Actuarial Journal*, pages 1–25, 2010.

[27] D. Lubell. A short proof of sperners theorem. *J. Combinatorial Theory*, 1:299, 1966.

[28] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Journal of Economic Theory*, 18:511–533, 2001.

[29] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[30] E. Sperner. Ein satz uber untermengen einer endlichen menge. *Math. Zeitschrift*, 27:544 – 548, 1928.

A Proofs in Section 3

We first state the Berry-Esseen bound:

Theorem 34 (Berry-Esseen bound [4,10] with parameter estimation by Korolev and Shevtsova [26])**.** *Let $X_1, X_2, \dots X_n$ be i.i.d random variables with $\mathbb{E}[X_1] = 0$, $\mathbb{E}\left[X_1^2\right] = \sigma^2 > 0$ and $\mathbb{E}\left[|X_1|^3\right] = \rho < \infty$. Also, let $Y_n = \frac{\sum_{i=1}^n X_i}{n}$ and F_n be the cumulative distribution function of $\frac{Y_n\sqrt{n}}{\sigma}$. Then*

$$|F_n(x) - \Phi(x)| \le \frac{C\rho}{\sigma^3\sqrt{n}}$$

where Φ is the cumulative distribution function of the standard normal distribution and $C < 1/2$ is a constant.

We will use the following fact in the proof of Lemma 16.

Fact 35. *If $f : \mathbb{R} \mapsto \mathbb{R}$ is a non-negative monotone decreasing function then for any integers $\ell < h$,*

$$\int_{\ell}^{h} f(x)dx \le \sum_{i=\ell}^{h} f(i) \le \int_{\ell-1}^{h} f(x)dx.$$

Proof of Lemma 16. Let $\mathbb{E}[X_i] = \mu$, $Var(X_i) = \sigma^2$. Then

$$\mu = \mathbb{E}[X_i] = \frac{b(H_{m+r} - H_r)}{m} \le \frac{bH_{m+r}}{m} \tag{6}$$

where $H_k = \sum_{i=1}^k \frac{1}{j}$ is the k-th harmonic number.

Also, let $Y_i = X_i - \mu$. Then $\mathbb{E}[Y_i] = 0$ and $\mathbb{E}\left[Y_i^2\right] = Var(X_i) = \sigma^2$. Finally, let $Y = \sum_{i=1}^m Y_i$ and $Z = \frac{Y\sqrt{m}}{\sigma m} = \frac{Y}{\sigma\sqrt{m}}$. We have

$$\begin{aligned}\mathbf{Pr}\left[\mathbb{E}[X] \le X \le \mathbb{E}[X] + d\right] &= \mathbf{Pr}\left[0 \le Y \le d\right] \\ &= \mathbf{Pr}\left[0 \le Z \le \frac{d}{\sigma\sqrt{m}}\right] \\ &= \mathbf{Pr}\left[Z \le \frac{d}{\sigma\sqrt{m}}\right] - \mathbf{Pr}\left[Z < 0\right] \end{aligned} \tag{7}$$

By Theorem 34, we have

$$\mathbf{Pr}\left[Z \le \frac{d}{\sigma\sqrt{m}}\right] \ge \Phi\left(\frac{d}{\sigma\sqrt{m}}\right) - \frac{C\rho}{\sigma^3\sqrt{m}} \quad \text{and} \quad \mathbf{Pr}\left[Z < 0\right] \le \Phi(0) + \frac{C\rho}{\sigma^3\sqrt{m}}.$$

where $\rho = \mathbb{E}\left[|Y_i|^3\right]$ and C is a constant less than 1/2. Therefore,

$$\mathbf{Pr}\left[\mathbb{E}\left[X\right] \le X \le \mathbb{E}\left[X\right] + d\right] \ge \Phi\left(\frac{d}{\sigma\sqrt{m}}\right) - \frac{2C\rho}{\sigma^3\sqrt{m}} - \Phi(0) \tag{8}$$

To bound this quantity, we need to bound σ and ρ. First, we have

$$\begin{aligned}
\sigma^2 = Var(X_i) &= \frac{1}{m}\sum_{j=r+1}^{r+m}\frac{b^2}{j^2} - \mu^2 \\
&\le \frac{1}{m}\sum_{j=r+1}^{r+m}\frac{b^2}{j^2} \\
&\le \frac{1}{m}\int_r^{r+m}\frac{b^2 dx}{x^2} \qquad (9) \\
&= \frac{1}{m}\left(\frac{b^2}{r} - \frac{b^2}{r+m}\right) \\
&\le \frac{b^2}{mr}, \qquad (10)
\end{aligned}$$

where we used the second part of Fact 35 in (9). Furthermore, we have

$$\begin{aligned}
\sigma^2 &= \frac{1}{m}\sum_{j=r+1}^{r+m}\frac{b^2}{j^2} - \mu^2 \\
&\ge \frac{1}{m}\int_{r+1}^{r+m}\frac{b^2 dx}{x^2} - \mu^2 \qquad (11) \\
&= \frac{b^2}{m(r+1)} - \frac{b^2}{m(r+m)} - \mu^2 \qquad (12) \\
&\ge \frac{b^2}{m(r+1)} - \frac{b^2}{m(r+m)} - \frac{b^2 H_{m+r}^2}{m^2} \qquad (13) \\
&\ge \frac{b^2}{2mr} \qquad (14)
\end{aligned}$$

when $r \ge 3$ and m is sufficiently large compared to r so that $m \ge 8H_{m+r}^2 r$. We used the first part of Fact 35 in (11) and substituted the last inequality of (6) into (12) to get (13).

Next, we bound ρ. Let $q = \lfloor \frac{m}{H_{m+r}-H_r} \rfloor$. Since $\frac{b}{j} \ge \frac{b(H_{m+r}-H_r)}{m} = \mu$ if and only if $j \le \frac{m}{H_{m+r}-H_r}$, we have

$$\rho = \mathbb{E}\left[|Y_i|^3\right] = \frac{1}{m}\left(\sum_{j=r+1}^{q}\left(\frac{b}{j}-\mu\right)^3 + \sum_{j=q+1}^{m+r}\left(\mu-\frac{b}{j}\right)^3\right)$$
$$\leq \frac{1}{m}\left(\sum_{j=r+1}^{q}\left(\frac{b^3}{j^3}+\frac{3b}{j}\mu^2\right) + \sum_{j=q+1}^{m+r}\left(\mu^3+\frac{3b^2}{j^2}\mu\right)\right)$$
$$\leq \frac{1}{m}\left(\sum_{j=r+1}^{q}\frac{b^3}{j^3} + 3b(H_q-H_r)\mu^2 + (m+r-q)\mu^3 + 3\mu\sum_{j=q+1}^{m+r}\frac{b^2}{j^2}\right)$$
$$\leq \frac{1}{m}\left(\int_r^q \frac{b^3 dx}{x^3} + 3b(H_q-H_r)\mu^2 + (m+r-q)\mu^3 + 3\mu\int_q^{m+r}\frac{b^2 dx}{x^2}\right) \tag{15}$$
$$\leq \frac{1}{m}\left(\left(\frac{1}{2r^2}-\frac{1}{2q^2}\right)b^3 + 3bH_q\mu^2 + m\mu^3 + 3\mu b^2\left(\frac{1}{q}-\frac{1}{m+r}\right)\right) \tag{16}$$
$$\leq \frac{1}{m}\left(\frac{b^3}{2r^2} + \frac{3b^3H_qH_{m+r}^2}{m^2} + \frac{b^3H_{m+r}^3}{m^2} + \frac{3b^3H_{m+r}}{mq}\right) \tag{17}$$
$$\leq \frac{b^3}{mr^2} \tag{18}$$

when m is sufficiently large compared to r. We used the second part of Fact 35 in (15), and the fact that $r \leq q$ in (16). Moreover, we substituted the second part of (6) into (16) to obtain (17). Finally, (18) follows from the fact that all the terms inside the parentheses in (17), except for the first one, tend to 0 when m/r tends to ∞.

Subtituting (10), (14) and (18) into (8) yields

$$\mathbf{Pr}\left[\mathbb{E}[X] \leq X \leq \mathbb{E}[X]+d\right] \geq \Phi\left(\frac{d}{\sigma\sqrt{m}}\right) - \frac{2C\rho}{\sigma^3\sqrt{m}} - \Phi(0)$$
$$\geq \Phi\left(\frac{d}{\sqrt{\frac{b^2}{mr}}\sqrt{m}}\right) - \frac{\frac{b^3}{mr^2}}{\frac{b^3}{(2mr)^{3/2}}\sqrt{m}} - \Phi(0)$$
$$= \Phi(d\sqrt{r}/b) - \frac{2^{3/2}}{\sqrt{r}} - \Phi(0) \tag{19}$$

Similarly, we have

$$\mathbf{Pr}\left[\mathbb{E}[X]-d \leq X \leq \mathbb{E}[X]\right] \geq \Phi(0) - \Phi(-d\sqrt{r}/b) - \frac{2^{3/2}}{\sqrt{r}}. \tag{20}$$

Since the right hand sides of both (19) and (20) tend to 1/2 when r tends to ∞, for any $\delta \leq 1/4$, there is a T such that both of the above probabilities are at least δ when $r \geq T$. □

Proof of Lemma 20. Fix $\mathbf{v}_{-i}$. If the mechanism is truthful ,this determines the threshold bid p_i (< 1) of agents i. v_i is greater than this threshold bid with probability at most $(1/p_i - r)/m$. The lemma follows. Slightly more complex but still straightforward arguments show that the same lemma holds for mechanisms that are truthful in expectation. □

B Proofs in Section 4

In this appendix, we give the detailed proof of Theorem 23. For the sake of readability, some discussion in Section 4 will be repeated.

Consider a valuation profile $\mathbf{v}$ and a permutation $\mathbf{u}$ of it where $u_1 = u_2 \geq \cdots \geq u_n$. Let $a = \max_{1\leq i\leq n} iu_i$. Then $\mathsf{BFP}(\mathbf{v}) \geq a/2$. Therefore, it suffices to show that for any distribution $\mathbf{F} = F^n$

$$\mathbb{E}_\pi \left[\mathsf{Mye}_{\mathbf{F}}\left(\pi(\mathbf{u})\right)\right] = O\left(a\right). \tag{21}$$

To this end, first note that we only need to concern ourselves with the revenue Mye gets from a fraction of agents, those with high values. Formally, let $i^* = \Omega\left(n\right)$ be some index which is independent of the valuation profile and will be fixed later. Then for each agent whose value is u_i for $i > i^*$, the revenue Mye gets from her is at most $u_i \leq u_{i^*}$. Hence, the revenue Mye gets from all such agents is at most $(n - i^*)u_{i^*} = O\left(i^*\right)u_{i^*} = O\left(a\right)$. This shows that we only need to bound the expected revenue Mye gets from the agents whose values are u_i for some $i \leq i^*$.

Moreover, even an agent with value u_i for $i \leq i^*$ may not pay as much as u_{i^*} in some cases. For example, if the virtual surpluses of the two markets are very different from each other, almost all winning agents would pay only the reserve price $\bar{\phi}^{-1}(0)$. In the following, we discuss the case when the winning agents pay more than u_i^*.

To this end, let $w_i = \bar{\phi}(u_i)$, then $w_i \leq u_i \leq a/i$ and $w_1 \geq w_2 \geq \cdots \geq w_n$. Also, let $W_1(\pi)$ and $W_2(\pi)$ be the ironed virtual surpluses of the two markets M_1 and M_2 respectively, and $\Delta(\pi) = |W_1(\pi) - W_2(\pi)|$ (W_1, W_2 and Δ are random variables dependent on π).

Assume that $w_n \geq 0$ (we defer the case where $w_n < 0$ to the end of this section). In this case, $\mathsf{Mye}_{\mathbf{F}}$ serves all agents in the market with the higher ironed virtual surpluse, and the reserve price ($\bar{\phi}^{-1}(0)$) is not binding. Thefore, a winning agent with value u_i pays $\bar{\phi}^{-1}(w_i - \Delta(\pi))$. If $w_i - \Delta(\pi) \leq w_{i^*}$ then this payment is at most u_{i^*}. Otherwise, it is at most u_i. Therefore, the expected total payment of the winning agents whose values are at least u_{i^*} is at most

$$\begin{aligned}
\sum_{i=1}^{i^*}\left(u_{i^*}\mathbf{Pr}_\pi\left[w_i - \Delta(\pi) \leq w_{i^*}\right] + u_i\mathbf{Pr}_\pi\left[w_i - \Delta(\pi) > w_{i^*}\right]\right) &\leq i^*u_i^* + \sum_{i=1}^{i^*} u_i\mathbf{Pr}_\pi\left[w_i - \Delta(\pi) > w_{i^*}\right] \\
&\leq a + \sum_{i=1}^{i^*} u_i\mathbf{Pr}_\pi\left[w_i - \Delta(\pi) > w_{i^*}\right] \\
&\leq a + \sum_{i=1}^{i^*} \frac{a}{i}\mathbf{Pr}_\pi\left[w_i - \Delta(\pi) > w_{i^*}\right]
\end{aligned}$$

If we can prove that for all $i \leq i^*$, $\mathbf{Pr}_\pi\left[\Delta(\pi) < w_i - w_{i^*}\right]$ is at most $1/i^\delta$ for some δ then the second term of the above sum is also bounded by $O\left(a\right)$. Therefore the proof of Theorem 23 boils down to proving the following lemma

Lemma 36. *Let $i^* = cn$ where $c < 1/3$ is a constant. There exists a constant $\delta > 0$ such that for all $i \leq i^*$, we have*

$$\mathbf{Pr}_\pi\left[\Delta(\pi) < w_i - w_{i^*}\right] \leq O\left(i^{-\delta}\right)$$

The large part of the remainder of this section is devoted to proving this lemma. In the following, we will assume that $m_1 \leq m_2$. To prove Lemma 36, note that each permutation π defines a *placement* of values into markets, which is fully characterized by the set $P = \{j : \pi_j \in M_1\}$. In addition, $W_1(\pi), W_2(\pi)$ and $\Delta(\pi)$ only depend on P; so from now on, we will write them as $W_1(P)$, $W_2(P)$ and $\Delta(P)$ respectively. Furthemore, there are exactly $m_1!m_2!$ permutations that define each placement. Therefore, the probability we would like to bound is equal to

$$\mathbf{Pr}_P\left[\Delta(P) < w_i - w_{i^*}\right]$$

To bound this probability, as discussed in Section 4, we follow the structure of Lubell's proof [27] of Sperner's Lemma [30], which exploits the relationship between chains and anti-chains corresponding to some partial order. We first define these notions.

Definition 37. *Given a partial order $\prec$, a set $Q = \{a_1, a_2, \ldots, a_m\}$ is a* chain *(coresponding to $\prec$) if $a_1 \prec a_2 \prec \cdots \prec a_m$. Q is a* maximal chain *if it is not a proper subset of any other chain.*

Definition 38. *Given a partial order $\prec$, a set $Q = \{a_1, a_2, \ldots, a_m\}$ is an* anti-chain *(corresponding to $\prec$) if for any two elements a_i and a_j of Q, neither $a_i \prec a_j$ nor $a_j \prec a_i$.*

The following observation is obvious from the definitions

Observation 39. *Suppose Q and Q' are a chain and an anti-chain corresponding to the same partial order. Then $|Q \cap Q'| \leq 1$.*

Following Lubell's proof, we define an appropriate partial order on the placements, then show that the event $\Delta(P) < w_i - w_{i^*}$ is an anti-chain corresponding to this order, and finally use the relationship between chains and anti-chains to bound the probability of this event. To this end, we partition each placement P into three sets S, R and T where $S = P \cap \{1, 2, \ldots, i\}$, $R = P \cap \{i+1, i+2, \ldots, i^*\}$ and $T = P \cap \{i^*+1, i^*+2, \ldots, n\}$. In the following, we also use the tuple (S, R, T) to denote the placement characterized by P.

Definition 40. $(S_1, R, T_1) \prec_i (S_2, R, T_2)$ *if and only if $S_1 \subset S_2$ and $T_2 \subset T_1$.*

Let $\xi(P)$ (correspondingly, $\xi(\pi)$ and $\xi(S, R, T)$) be the event $\Delta(P) < w_i - w_i^*$ (correspondingly, $\Delta(\pi) < w_i - w_i^*$ and $\Delta(S, R, T) < w_i - w_i^*$), we define

Definition 41. *A placement (S, R, T) is* bad *if $\xi(S, R, T)$ happens.*

The following lemma relates these two definitions

Lemma 42. *If two placements (S_1, R, T_1) and (S_2, R, T_2) are both bad then neither $(S_1, R, T_1) \prec_i (S_2, R, T_2)$ nor $(S_2, R, T_2) \prec_i (S_1, R, T_1)$.*

Proof. (Recall that Figure 4 has a "proof by picture" of this lemma.) Suppose that $(S_1, R, T_1) \prec_i (S_2, R, T_2)$. Then

$$S_2 = S_1 \cup \{j_1, j_2, \ldots j_k\}, \quad \text{and} \quad T_1 = T_2 \cup \{h_1, h_2, \ldots h_k\}.$$

Let $P_1 = S_1 \cup R \cup T_1$ and $P_2 = S_2 \cup R \cup T_2$, we have

$$\begin{aligned}
(W_1(P_2) - W_2(P_2)) - (W_1(P_1) - W_2(P_1)) &= (W_1(P_2) - W_1(P_1)) + (W_2(P_1) - W_2(P_2)) \\
&= 2\left(\sum_{t=1}^{k} w_{j_t} - \sum_{t=1}^{k} w_{h_t}\right) \\
&= 2\sum_{t=1}^{k}(w_{j_t} - w_{h_t}) \\
&\geq 2k(w_i - w_{i^*}) \qquad (22) \\
&\geq 2(w_i - w_{i^*}) \qquad (23)
\end{aligned}$$

where (22) follows from the fact that $j_t \leq i$ and $h_t \geq i^*$ for all t. Since P_1 is bad, we have $W_1(P_1) - W_2(P_1) > w_{i^*} - w_i$. Hence, $W_1(P_2) - W_2(P_2) > w_{i^*} - w_i + 2(w_i - w_i^*) = w_i - w_{i^*}$, contradicting the fact that (S_2, R, T_2), i.e. P_2, is bad. Hence, $(S_1, T_1) \not\prec_i (S_2, T_2)$.

Similarly, we have $(S_2, T_2) \not\prec_i (S_1, T_1)$. The lemma follows. □

Let $\mathcal{A}$ be the set of bad placements by Definition 41, then

$$\mathbf{Pr}_P\left[\xi(P)\right] \leq \frac{|\mathcal{A}|}{\text{number of placements}} = \frac{|\mathcal{A}|}{\binom{n}{m_1}} \qquad (24)$$

For each R, let $\mathcal{A}_R$ be the set of bad placements (S,R,T). Then $|\mathcal{A}| = \sum_R |\mathcal{A}_R|$. Moreover, by Lemma 42, each $\mathcal{A}_R$ is an anti-chain with regard to $\prec_i$. We will exploit this fact to bound $|\mathcal{A}_R|$, following the structure of Lubell's proof [27] of Sperner's Lemma [30].

For each fixed (S,R,T), let $\mathcal{C}(S,R,T)$ be the set of maximal chains that contain (S,R,T) and $\mathcal{C}_R$ be the set of all maximal chains whose elements have the form (S',R,T'). Since both (S_1,R,T_1) and (S_2,R,T_2) belong to $\mathcal{A}$, by Observation 39, no two sets $\mathcal{C}(S_1,R,T_1)$ and $\mathcal{C}(S_2,R,T_2)$ intersect, for otherwise they contain a common chain and that chain intersects $\mathcal{A}_R$ at two points (S_1,R,T_1) and (S_2,R,T_2). Hence,

$$\sum_{(S,R,T)\in\mathcal{A}_R} |\mathcal{C}(S,R,T)| \leq |\mathcal{C}_R|,$$

which implies

$$|\mathcal{A}_R| \leq \max_{(S,T)} \frac{|\mathcal{C}_R|}{|\mathcal{C}(S,R,T)|} \tag{25}$$

Lemma 43 and Corollary 44 compute the denominator and numerator of the right hand side of this inequality.

Lemma 43. *Let $s=|S|, r=|R|$, and $t=|T|$, and $q=\min\{i, m_1-r\}$. The number of maximal chains containing (S,R,T) is*

$$\frac{s!t!(i-s)!}{(m_1-r-q)!(i-q)!}\cdot\frac{(n-t-i^*)!}{(m_2+r-i^*)!}$$

Proof. Consider constructing a maximal chain containing (S,R,T). There are two directions to go. Along one direction, we remove agents from S and add agents to T until either S is empty or there is no more element from $\{i^*+1,\ldots,n\}$ to add to T. Along the other direction, we add agents to S and remove agents from T until either T is empty or there is no more agent from $\{1,2,\ldots,i\}$ to add to S.

Since $n-i* > n/2 > m_1$, along the first direction, S will be empty before we run out of agents to add to T. Therefore, the number of ways to go along the first direction is

$$\begin{aligned} &s!(n-i^*-t)(n-i^*-t-1)\cdots(n-i^*-t-s+1)\\ &=\frac{s!(n-i^*-t)!}{(n-i^*-s-t)!}=\frac{s!(n-t-i^*)}{(m_2+r-i^*)!} \end{aligned} \tag{26}$$

Next, consider the second direction. If $m_1-r=t+s<i$ then T would be empty before we run out of elements to add to S. Hence, in this case, the number of ways to go along the second direction is

$$\begin{aligned} &t!(i-s)(i-s-1)\cdots(i-s-t+1)=\frac{t!(i-s)!}{(i-s-t)!}\\ &=\frac{t!(i-s)!}{(i-m_1+r)!}=\frac{t!(i-s)!}{(i-q)!(m_1-r-q)!} \end{aligned} \tag{27}$$

On the other hand, if $m_1-r=t+s\geq i$ then we would run out of agents to add to S before T is empty. Hence, the number of ways to go along the second direction is

$$\begin{aligned} &(i-s)!t(t-1)\cdots(t+s-i+1)=\frac{(i-s)!t!}{(t+s-i)!}\\ &=\frac{(i-s)!t!}{(m_1-r-i)!}=\frac{(i-s)!t!}{(m_1-r-q)!(i-q)!} \end{aligned} \tag{28}$$

Combining (26), (27) and (28) yields the lemma. ☐

Corollary 44. *The size of $\mathcal{C}_R$ is*

$$\frac{(n-i^*)!i!}{(m_2+r-i^*)!(m_1-r-q)!(i-q)!}$$

Proof. There are $\binom{n-i^*}{m_1-r}$ ways to construct a placement where $S=\emptyset$. For each of these placements, by Lemma 43, there are

$$\frac{0!(m_1-r)!i!(n-m_1+r-i^*)!}{(m_1-r-q)!(i-q)!(m_2+r-i^*)!)}$$

maximal chains that contain it. Hence, the number of maximal chains whose elements are of the form (S',R,T') is

$$\binom{n-i^*}{m_1-r}\cdot\frac{0!(m_1-r)!i!(n-m_1+r-i^*)!}{(m_1-r-q)!(i-q)!(m_2+r-i^*)!)}=\frac{(n-i^*)!i!}{(m_2+r-i^*)!(m_1-r-q)!(i-q)!}.$$

□

Substituting Lemma 43 and Corollary 44 into (25) yields

$$|\mathcal{A}_R|\leq\max_{s+t=m_1-r}\binom{n-i^*}{t}\binom{i}{s}.$$

where $r=|R|$. (Notice by the way that, to our knowledge, this fact cannot be obtained as a 1-line corollary of Sperner's Lemma since in 2 incomparable elements (S,R,T) and (S',R,T'), it could be that S is contained in S' as long as T and T' are unrelated in the partial order defined by inclusion of subsets of $\{i^*+1,\dots,n\}$ (or vice versa).)

Let $p=\min\{i^*-i,m_1\}$. We have $r\leq p$. Hence,

$$|\mathcal{A}|=\sum_R|\mathcal{A}_R|\leq\sum_{r=0}^{p}\left(\binom{i^*-i}{r}\cdot\max_{s+t=m_1-r}\binom{i}{s}\binom{n-i^*}{t}\right)\tag{29}$$

For brevity, let $g(r)=\max_{s+t=m_1-r}\binom{i}{s}\binom{n-i^*}{t}$, and $f(r)=\binom{i-i^*}{r}g(r)$.

(29) gives us the numerator of the right hand side of (24). To bound it, we rewrite the denominator as a similar sum

$$\binom{n}{m_1}=\sum_{r=0}^{i^*-i}\left(\binom{i^*-i}{r}\binom{n-i^*+i}{m_1-r}\right)\geq\sum_{r=0}^{p}\left(\binom{i^*-i}{r}\binom{n-i^*+i}{m_1-r}\right)\tag{30}$$

(29) and (30) suggest that to bound $|\mathcal{A}|/\binom{n}{m_1}$, we can attempt to bound $g(r)/\binom{n-i^*+i}{m_1-r}$ for all r. This is achieved by applying the following lemma for $a=i,b=n-i^*$ and $m=m_1-r$.

Lemma 45. *Let a,b,m be positive integers such that $am>(a+b)$ and $bm>(a+b)$, and $\alpha=m/(a+b)$. We have*

$$\frac{\max_x\binom{a}{x}\binom{b}{m-x}}{\binom{a+b}{m}}=O\left(\sqrt{\frac{a+b}{\alpha(1-\alpha)ab}}\right)$$

We will use a few simple facts to prove this claim.

Lemma 46. *Let a,b,x,y be positive integers. Then*

$$\binom{a}{x}\binom{b}{y}\leq\binom{a}{x+1}\binom{b}{y-1}\iff(b+1)(x+1)\leq(a+1)y.$$

Proof. We have

$$R = \frac{\binom{a}{x}\binom{b}{y}}{\binom{a}{x+1}\binom{b}{y-1}} = \frac{(x+1)(b-y+1)}{(a-x)y}.$$

Simple calculations show that $R \le 1$ if and only if $(x+1)(b+1) \le y(a+1)$. □

Corollary 47. *Let a, b, m be positive integers where $m \le \min(a,b)$, then*

$$x \in \operatorname{argmax}_x \binom{a}{x}\binom{b}{m-x} \implies \frac{am}{a+b} - 1 \le x \le \frac{am}{a+b} + 1.$$

Lemma 48 (Stirling's Approximation)**.** *For any $a \ge x > 0$,*

$$\binom{a}{x} = \Theta\left(\sqrt{\frac{a}{x(a-x)}} \cdot \frac{a^a}{x^x(a-x)^{a-x}}\right).$$

Proof of Lemma 45. By Corollary 47, there is x' such that $\frac{am}{a+b} - 1 \le x' \le \frac{am}{a+b} + 1$ and for any x:

$$\binom{a}{x}\binom{b}{m-x} \le \binom{a}{x'}\binom{b}{m-x'}.$$

Since $am > (a+b)$, we have $x' > 0$. Therefore, by Lemma 48

$$\begin{aligned}\frac{\binom{a}{x'}\binom{b}{m-x'}}{\binom{a+b}{m}} =& O\left(\sqrt{\frac{abm(a+b-m)}{x'(a-x')(m-x')(b-m+x')(a+b)}}\right)\times \\ &\times \frac{a^a b^b m^m (a+b-m)^{a+b-m}}{(x')^{x'}(a-x')^{a-x'}(m-x')^{m-x'}(b-m+x')^{b-m+x'}(a+b)^{a+b}}.\end{aligned}$$

Denote the second term of the right hand side $t(x')$. By thinking of $t(x')$ as a function over real numbers, we can show that $t(x') \le t\left(\frac{am}{a+b}\right) = 1$. On the other hand, for $\frac{am}{a+b} - 1 \le x' \le \frac{am}{a+b} + 1$, the first term $= O\left(\sqrt{\frac{a+b}{\alpha(1-\alpha)ab}}\right)$. Hence the lemma holds. □

As we discussed, Lemma 45 allows us to bound $g(r)/\binom{n-i^*+i}{m_1-r}$ in some cases. However, because of the condition that $am > (a+b)$ and the appearance of $\alpha = m/(a+b)$ in the denominator of the right hand side, this lemma is only useful for $m_1 - r$ being sufficiently large. If m_1 and m_2 are within a constant of each other, we can choose $i^* = m_1/2$ and guarantee this, but for m_1 and m_2 vastly different from each other, this approach does not work.

To overcome this problem, we will consider two cases where m_1 is very small compared to m_2 (hence, n) and where m_1 is large enough compared to n. The trick is to define "small" and "large" in terms of i. In particular, if $m_1 i^\delta \le n$ for some $\delta > 0$, we have an easy bound using Lemma 49. For the other case, we will separate the right hand side of (29) into two smaller sums among which one can be bounded using Lemma 45. We will then discuss how to bound the remaining sub-sum.

In the following, we fix $\delta = 1/4$. Lemma 49 and Corollary 50 bounds $\mathbf{Pr}_P[\xi(P)]$ when $m_1 i^\delta \le n$ as discussed.

Lemma 49. $\mathbf{Pr}_P[\xi(P)] \le \frac{2m_1}{n-1-m_1}$.

Proof. Observe that $\xi(P)$ implies that either $W_1 > W_2$ or $i \in M_2$ and $W_1 > W_2 - w_i$. Hence, the probability of $\xi(P)$ is at most the sum of the probabilities of the latter events. We will show that the probability of each of them is at most $\frac{m_1}{n-1-m_1}$.

First, we bound the probability that $W_1 > W_2$. To this end, note that a random placement can be constructed through the following process

1. Divide the set of values $w_1, w_2, \ldots w_n$ into n/m_1 groups, each of size m_1, except for one group whose size may be smaller.

2. Choose a random group among the groups of size m_1 to place into M_1.

After Step 1, if there is more than one group with the maximum virtual surplus then $W_1 \leq W_2$. If there is such a group, then the probability that $W_1 > W_2$ is at most the probability that it is placed into M_1, which is at most $\frac{1}{n/m_1 - 1} = \frac{m_1}{n-m_1} < \frac{m_1}{n-1-m_1}$.

Next, the probability that $W_1 > W_2 - w_i$ and $i \in M_2$ is bounded by the probability that $W_1 > W_2 - w_i$ given that $i \in M_2$. Similar to above, we can construct a random placement given that $i \in M_2$ by first dividing the remaining $n-1$ agents into $(n-1)/m_1$ groups, and then placing one of the groups whose sizes are m_1 to M_1. Again, $W_1 > W_2 - w_i$ only if there is a unique group with the maximum virtual surplus and that group is assigned to M_1. The probability of this happening is at most $\frac{1}{(n-1)/m_1 - 1} = \frac{m_1}{n-1-m_1}$. This completes the proof. □

Corollary 50. *If $m_1 i^\delta \leq n$ then $\mathbf{Pr}_P[\xi(P)] \leq O\left(\frac{1}{i^\delta}\right)$.*

Now suppose $m_1 i^\delta > n$. We rewrite the right hand side of (29) as follows

$$\sum_{r=0}^{p} f(r) = \sum_{r=0}^{r^*} f(r) + \sum_{r=r^*+1}^{p} f(r) \tag{31}$$

where $r^* = m_1(1 - i^{2\delta-1})$. The first sub-sum contains the terms where r is "small", hence $m_1 - r$ is "large", so we can bound it using Lemma 45.

Lemma 51. *Suppose $m_1 i^\delta > n$. Then $\sum_{r=0}^{r^*} f(r) \leq O\left(i^{-\delta/2}\right)\binom{n}{m_1}$.*

Proof. Fix some r and let $m' = m_1 - r$. Since $m'i \geq m_1 i^{2\delta-1} \cdot i = m_1 i^{2\delta} > m_1 i^\delta > n \geq n - i^* + i$, Lemma 45 applies with $a = i$, $b = n - i^*$, $m = m'$ and $\alpha = m'/(i+b)$. By Lemma 45, we have

$$\begin{aligned}
\frac{\max_{s+t=m'} \binom{i}{s}\binom{n-i^*}{t}}{\binom{n-i^*+i}{m'}} &= O\left(\sqrt{\frac{i+b}{\alpha(1-\alpha)ib}}\right) = O\left(\sqrt{\frac{i+b}{\alpha ib}}\right) \\
&= O\left(\sqrt{\frac{(i+b)^2}{m'ib}}\right) = O\left(\sqrt{\frac{(i+b)^2}{nbi^\delta}}\right) = O\left(\sqrt{\frac{1}{i^\delta}}\right)
\end{aligned}$$

□

On the other hand, to bound the second sub-sum, we note that its summands are decreasing with r (see Lemma 52). Therefore, it can be bounded by the first summand times the number of summands, which turns out to be good enough. In particular, we have

Lemma 52. *If $p \geq r_2 = r_1 + 1 \geq m_1/2 + 1$ then $f(r_2) \leq f(r_1)$.*

Proof. Let $(s', t') \in \operatorname{argmax}_{s,t:s+t=m_1-r_2} \binom{i}{s}\binom{n-i^*}{t}$. We claim that

$$\binom{i^*-i}{r_2}\binom{i}{s'}\binom{n-i^*}{t'} \leq \binom{i^*-i}{r_2-1}\binom{i}{s'}\binom{n-i^*}{t'+1}.$$

By Lemma 46, it suffices to show that

$$(i^* - i + 1)(t' + 1) \leq (n - i^* + 1)r_2$$

To this end, note that the left hand side is at most $i^*(r_1 + 1) \leq (n - i^*)r_2$.

Hence, we have

$$\binom{i^*-i}{r_2}\binom{i}{s'}\binom{n-i^*}{t'} \le \binom{i^*-i}{r_1}\binom{i}{s'}\binom{n-i^*}{t'+r_2-r_1}$$
$$\le \binom{i^*-i}{r_1}\max_{s+t=m_1-r_1}\binom{i}{s}\binom{n-i^*}{t}.$$

□

Lemma 53. $\sum_{r=r^*+1}^{p} f(r) = O\left(i^{2\delta-1}\right)\binom{n}{m_1}$.

Proof. By Lemma 52, we have

$$\begin{aligned}
\sum_{r=r^*+1}^{p} f(r) &\le \frac{p-r^*+1}{p-m_1/2+1}\sum_{r=m_1/2}^{p} f(r) \\
&\le \frac{p-r^*+1}{p-m_1/2+1}\sum_{r=m_1/2}^{p}\left(\binom{i^*-i}{r}\binom{n-i^*+i}{m-r}\right) \\
&\le \frac{p-r^*+1}{p-m_1/2+1}\binom{n}{m_1} \\
&\le \frac{m_1-r^*+1}{m_1-m_1/2+1}\binom{n}{m_1} \\
&= \frac{m_1 i^{2\delta-1}+1}{m_1/2+1}\binom{n}{m_1} \\
&= O\left(i^{2\delta-1}\right)\binom{n}{m_1}
\end{aligned}$$

as required. □

Lemmas 51 and 53 allow us to bound the right hand side of (29) by $O\left(\frac{1}{i^{\delta/2}}\right)\binom{n}{m_i}$, hence imply Theorem 23.

The case where $w_n < 0$

Consider the case where $w_n < 0$. If $w_1 < 0$, Mye_F's revenue is 0 and we are done. Suppose $w_1 \ge 0$ and let n' be the index where $w'_n \ge 0$ and $w_{n'+1} < 0$. Now fix π_i for all $i > n'$. For each such fixing, we have a reduced two market environment. sOBO on the original environment is at most the maximum of sOBO over these reduced environment. By the result of the previous two subsections, sOBO in any such reduced environment is at most a constant times BFP in that environment, which is at most a constant times a. Hence, sOBO is at most a constant times a, which is a constant times BFP.