


Sequential Posted Price Mechanisms with Correlated Valuations

Marek Adamczyk¹, Allan Borodin², Diodato Ferraioli³,
Bart de Keijzer¹, and Stefano Leonardi¹

¹ Sapienza University of Rome, Rome, Italy
{adamczyk, dekeijzer, leonardi}@dis.uniroma1.it

² University of Toronto, Toronto, Canada
bor@cs.toronto.edu

³ University of Salerno, Fisciano, SA, Italy
dferraioli@unisa.it

Abstract. We study the revenue performance of sequential posted price mechanisms and some natural extensions, for a general setting where the valuations of the buyers are drawn from a correlated distribution. Sequential posted price mechanisms are conceptually simple mechanisms that work by proposing a “take-it-or-leave-it” offer to each buyer. We apply sequential posted price mechanisms to single-parameter multi-unit settings in which each buyer demands only one item and the mechanism can assign the service to at most k of the buyers. For standard sequential posted price mechanisms, we prove that with the valuation distribution having finite support, no sequential posted price mechanism can extract a constant fraction of the optimal expected revenue, even with unlimited supply. We extend this result to the case of a continuous valuation distribution when various standard assumptions hold simultaneously. In fact, it turns out that the best fraction of the optimal revenue that is extractable by a sequential posted price mechanism is proportional to the ratio of the highest and lowest possible valuation. We prove that for two simple generalizations of these mechanisms, a better revenue performance can be achieved: if the sequential posted price mechanism has for each buyer the option of *either* proposing an offer *or* asking the buyer for its valuation, then a $\Omega(1/\max\{1, d\})$ fraction of the optimal revenue can be extracted, where d denotes the “degree of dependence” of the valuations, ranging from complete independence ($d = 0$) to arbitrary dependence ($d = n - 1$). When we generalize the sequential posted price mechanisms further, such that the mechanism has the ability to make a take-it-or-leave-it offer to the i -th buyer that depends on the valuations of all buyers except i , we prove that a constant fraction $(2 - \sqrt{e})/4 \approx 0.088$ of the optimal revenue can be always extracted.

This work is supported by the EU FET project MULTIPLEX no. 317532, the ERC StG Project PAAI 259515, the Google Research Award for Economics and Market Algorithms, and the Italian MIUR PRIN 2010-2011 project ARS TechnoMedia – Algorithmics for Social Technological Networks.

1 Introduction

A large body of literature in the field of mechanism design focuses on the design of auctions that are optimal with respect to some given objective function, such as maximizing the social welfare or the auctioneer’s revenue. This literature mainly considered direct revelation mechanisms, in which each buyer submits a bid that represents his valuation for getting the service, and the mechanism determines the winners and the payments. The reason for this is the *revelation principle* (see, e.g., [9]), which implies that one may study only direct revelation mechanisms for many purposes. Some of the most celebrated mechanisms follow this approach, such as the VCG mechanism [12, 17, 29] and the Myerson mechanism [23].

A natural assumption behind these mechanisms is that buyers will submit truthfully whenever the utility they take with the truthful bid is at least as high as the utility they may take with a different bid. However, it has often been acknowledged that such an assumption may be too strong in a real world setting. In particular, Sandholm and Gilpin [27] highlight that this assumption usually fails because of: (1) a buyer’s unwillingness to fully specify their values, (2) a buyer’s unwillingness to participate in ill understood, complex, unintuitive auction mechanisms, and (3) irrationality of a buyer, which leads him to underbid even when there is nothing to be gained from this behavior.

This has recently motivated the research about auction mechanisms that are conceptually simple. Among these, the class of *sequential posted price mechanisms* [11] is particularly attractive. First studied by Sandholm and Gilpin [27] (and called “take-it-or-leave-it mechanisms”), these mechanisms work by iteratively selecting a buyer that has not been selected previously, and offering him a price. The buyer may then accept or reject that price. When the buyer accepts, he is allocated the service. Otherwise, the mechanism does not allocate the service to the buyer. In the sequential posted-price mechanism we allow both the choice of buyer and the price offered to that buyer to depend on the decisions of the previously selected buyers (and the prior knowledge about the buyers’ valuations). Also, randomization in the choice of the buyer and in the charged price is allowed. Sequential posted price mechanisms are thus conceptually simple and buyers do not have to reveal their valuations. Moreover, they possess a trivial dominant strategy (i.e., buyers do not have to take strategic decisions) and are individually rational (i.e., participation is never harmful to the buyer).

Sequential posted price mechanisms have been mainly studied for the setting where the valuations of the buyers are each drawn independently from publicly known buyer-specific distributions, called the *independent values* setting. In this paper, we study a much more general setting, and assume that the entire vector of valuations is drawn from one publicly known distribution, which allows for arbitrarily complex dependencies among the valuations of the buyers. This setting is commonly known as the *correlated values* setting. Our goal is to investigate the revenue guarantees of sequential posted price mechanisms in the correlated value setting. We quantify the quality of a mechanism by comparing its expected revenue to that of the *optimal mechanism*, that achieves the highest

expected revenue among all dominant strategy incentive compatible and ex-post individually rational mechanisms (see the definitions below).

We assume a standard Bayesian, transferable, quasi-linear utility model and we study the *unit demand, single parameter, multi-unit* setting: there is one service (or type of item) being provided by the auctioneer, there are n buyers each interested in receiving the service once, and the *valuation* of each buyer consists of a single number that reflects to what extent a buyer would profit from receiving the service provided by the auctioneer. The auctioneer can charge a price to a bidder, so that the utility of a bidder is his valuation (in case he gets the service), minus the charged price. In this paper, our focus is on the k -limited supply setting, where service can be provided to at most k of the buyers. This is an important setting because it is a natural constraint in many realistic scenarios, and it contains two fundamental special cases: the *unit supply* setting (where $k = 1$), and the *unlimited supply* setting where $k = n$.

Related Work. There has been substantial work [5, 14, 19, 20, 26] on *simple* mechanisms. Babaioff et al. [5] highlight the importance of understanding the strength of simple versus complex mechanisms for revenue maximization.

As described above, sequential posted price mechanisms are an example of such a simple class of mechanisms. Sandholm and Gilpin [27] have been the first ones to study sequential posted price mechanisms. They give experimental results for the case in which values are independently drawn from the uniform distribution in $[0, 1]$. Moreover, they consider the case where multiple offers can be made to a bidder, and study the equilibria that arise from this. Blumrosen and Holenstein [8] compare fixed price (called symmetric auctions), sequential posted price (called discriminatory auctions) and the optimal mechanism for valuations drawn from a wide class of i.i.d. distributions. Babaioff et al. [3] consider *prior-independent* posted price mechanisms with k -limited supply for the setting where the only information known is that all valuations are independently drawn from the same distribution with support $[0, 1]$. Posted-price mechanisms have also been previously studied in [6, 7, 21], albeit for a non-Bayesian, on-line setting. In a recent work Feldman et al. [16] study on-line posted price mechanisms for combinatorial auctions when valuations are independently drawn.

The works of Chawla et al. [11] and Gupta and Nagarajan [18] are closest to our present work, although they only consider sequential posted price mechanisms in the independent values setting. In particular, Chawla et al. [11] prove that such mechanisms can extract a constant factor of the optimal revenue for single and multiple parameter settings under various constraints on the allocations. They also consider on-line (called *order-oblivious* in [11]) sequential posted price mechanisms in which the order of the buyers is fixed and adversarially determined. They use on-line mechanisms to establish results for the more general multi-parameter case. Yan [30], and Kleinberg and Weinberg [22] build on this work and strengthen some of the results of Chawla et al. [11].

Gupta and Nagarajan [18] introduce a more abstract stochastic probing problem that includes Bayesian sequential posted price mechanisms. Their approximation

bounds were later improved by Adamczyk et al. [1] who in particular matched the approximation of Chawla et al. [11] for single matroid settings.

All previous work only consider the independent setting. In this work we instead focus on the correlated setting. The lookahead mechanism of Ronen [24] is a fundamental reference for the correlated setting. It also resembles some of the mechanisms considered in this work. However, as we will indicate, it turns out to be different in substantial ways. Cremer and McLean [13] made a fundamental contribution to auction theory in the correlated value setting, by exactly characterizing for which valuation distributions it is possible to extract the full optimal social welfare as revenue. Segal [28] gives a characterization of optimal ex-post incentive compatible and ex-post individually rational optimal mechanisms. Roughgarden and Talgam-Cohen [25] study the even more general *interdependent* setting. They show how to extend the Myerson mechanism to this setting for various assumptions on the valuation distribution. There is now a substantial literature [10, 15, 25] that develops mechanisms with good approximation guarantees for revenue maximization in the correlated setting. These mechanisms build on the lookahead mechanism of Ronen [24] and thus they also differ from the mechanisms proposed in this work.

Contributions and Outline. We first define some preliminaries and notation. In Sect. 2 we give a simple sequence of instances which demonstrate that for (unrestricted) correlated distributions, sequential posted price (SPP) mechanisms cannot obtain a constant approximation with respect to the revenue obtained by the optimal dominant strategy incentive compatible and ex-post individually rational mechanism. This holds for any value of k (i.e., the size of the supply). We extend this impossibility result by proving that a constant approximation is impossible to achieve even when we assume that the valuation distribution is continuous and satisfies all of the following conditions simultaneously: the valuation distribution is supported everywhere, is entirely symmetric, satisfies *regularity*, satisfies the *monotone hazard rate* condition, satisfies *affiliation*, all the induced marginal distributions have finite expectation, and all the conditional marginal distributions are non-zero everywhere.

Given these negative results, we consider a generalization of sequential posted price mechanisms that are more suitable for settings with limited dependence among the buyers' valuations: *enhanced sequential posted price (ESPP) mechanisms*. An ESPP mechanism works by iteratively selecting a buyer that has not been selected previously. The auctioneer can either offer the selected buyer a price or ask him to report his valuation. As in sequential posted price mechanisms, if the buyer is offered a price, then he may accept or reject that price. When the buyer accepts, he is allocated the service. Otherwise, the mechanism does not allocate the service to the buyer. If instead, the buyer is asked to report his valuation, then the mechanism does not allocate him the service. Note that the ESPP mechanism requires that some fraction of buyers reveal their valuation truthfully. Thus, the property that the bidders not have to reveal their preferences is *partially* sacrificed, for a more powerful class of mechanisms and (as we will see) a better revenue performance. For the ESPP mechanisms, again

there are instances in which the revenue is not within a constant fraction of the optimal revenue. However, these mechanisms can extract a fraction $\Theta(1/n)$ of the optimal revenue, regardless of the valuation distribution.

This result seems to suggest that to achieve a constant approximation of the optimal revenue it is *necessary* to collect all the bids truthfully. Consistent with this hypothesis, we prove that a constant fraction of the optimal revenue can be extracted by dominant strategy IC *blind offer mechanisms*: these mechanisms inherit all the limitations of sequential posted price mechanisms (i.e., buyers are considered sequentially in an order independent of any bids; buyers are only offered a price when selected; and the buyer gets the service only if he accepts the offered price), except that the price offered to a bidder i may now depend on the bids submitted by all players other than i . This generalization sacrifices entirely the property that buyers valuations need not be revealed. Blind offer mechanisms are thus necessarily direct revelation mechanisms. However, this comes with the reward of a revenue that is only a constant factor away from optimal. In conclusion, blind offer mechanisms achieve a constant approximation of the optimal revenue, largely preserve the conceptual simplicity of sequential posted price mechanisms, and are easy to grasp for the buyers participating in the auction. In particular, buyers have a conceptually simple and practical strategy: to accept the price if and only if it is not above their valuation, regardless of how the prices are computed. We stress that, even if blind offer mechanisms sacrifice some simplicity (and practicality), we still find it theoretically interesting that a mechanism that allocates items to buyers *in any order* and thus not necessarily in an order that maximizes profit, say as in [24], is able to achieve a constant approximation of the optimal revenue even with correlated valuations. Moreover, blind offer mechanisms provide the intermediate step en-route to establishing revenue approximation bounds for other mechanisms. We will show how blind offer mechanisms serve this purpose in Sect. 3.

We highlight that our positive results do not make any assumptions on the marginal valuation distributions of the buyers nor the type of correlation among the buyers. However, in Sect. 3 we consider the case in which the degree of dependence among the buyers is limited. In particular, we introduce the notion of *d-dimensionally dependent distributions*. This notion informally requires that for each buyer i there is a set S_i of d other buyers such that the distribution of i 's valuation when conditioning on the vector of other buyers' valuations can likewise be obtained by only conditioning on the valuations of S_i . Thus, this notion induces a hierarchy of n classes of valuation distributions with increasing degrees of dependence among the buyers: for $d = 0$ the buyers have independent valuations, while the other extreme $d = n - 1$ implies that the valuations may be dependent in arbitrarily complex ways. Note that d -dimensional dependence does not require that the marginal valuation distributions of the buyers themselves satisfy any particular property, and neither does it require anything from the type of correlation that may exist among the buyers. This stands in contrast with commonly made assumptions such as *symmetry*, *affiliation*, the *monotone-hazard rate assumption*,

and *regularity*, that are often encountered in the auction theory and mechanism design literature.

Our main positive result for ESPP mechanisms then states that if the valuation distribution is d -dimensionally dependent, there exists an ESPP mechanism that extracts an $\Omega(1/d)$ fraction of the optimal revenue. The proof of this result consists of three key ingredients: (i) An upper bound on the optimal ex-post IC, ex-post IR revenue in terms of the solution of a linear program. This part of the proof generalizes a linear programming characterization introduced by Gupta and Nagarajan [18] for the independent distribution setting. (ii) A proof that incentive compatible blind offer mechanisms are powerful enough to extract a constant fraction of the optimal revenue of any instance. This makes crucial use of the linear program mentioned above. (iii) A conversion lemma showing that blind offer mechanisms can be turned into ESPP mechanisms while maintaining a fraction $\Omega(1/d)$ of the revenue of the blind offer mechanism.

Many proofs and various important parts of the discussion have been omitted from this version of our paper, due to space constraints. We refer the reader to [2] for full proofs and a complete discussion of our work and results.

Preliminaries. For $a \in \mathbb{N}$, $[a]$ denotes the set $\{1, \dots, a\}$. For a vector \vec{v} and an arbitrary element a , let (a, \vec{v}_{-i}) be the vector obtained by replacing v_i with a .

We face a setting where an auctioneer provides a service to n buyers, and is able to serve at most k of the buyers. The buyers have valuations for the service offered, which are drawn from a *valuation distribution* π , i.e., a probability distribution on $\mathbb{R}_{\geq 0}^n$. We will assume throughout this paper that π is discrete, except where otherwise stated.

We will use the following notation for conditional and marginal probability distributions. Let π be a discrete finite probability distribution on \mathbb{R}^n , let $i \in [n]$, $S \subset [n]$ and $\vec{v} \in \mathbb{R}^n$. For an arbitrary probability distribution π , denote by $\text{supp}(\pi)$ the support of π , by \vec{v}_S the vector obtained by removing from \vec{v} the coordinates in $[n] \setminus S$, by π_S the distribution induced by drawing a vector from π and removing the coordinates corresponding to index set $[n] \setminus S$, by $\pi_{\vec{v}_S}$ the distribution of π conditioned on the event that \vec{v}_S is the vector of values on the coordinates corresponding to index set S , and by π_{i, \vec{v}_S} the marginal distribution of the coordinate of $\pi_{\vec{v}_S}$ that corresponds to buyer i . In the subscripts we sometimes write i instead of $\{i\}$ and $-i$ instead of $[n] \setminus \{i\}$.

An *instance* is a triple (n, π, k) , where n is the number of participating buyers, π is the valuation distribution, and $k \in \mathbb{N}_{\geq 1}$ is the supply, i.e., the number of services that the auctioneer may allocate to the buyers. A *deterministic mechanism* f is a function from $\times_{i \in [n]} \Sigma_i$ to $\{0, 1\}^n \times \mathbb{R}_{\geq 0}^n$, for any choice of *strategy sets* $\Sigma_i, i \in [n]$. When $\Sigma_i = \text{supp}(\pi_i)$ for all $i \in [n]$, mechanism f is called a *deterministic direct revelation mechanism*. A *randomized mechanism* M is a probability distribution over deterministic mechanisms. For $i \in [n]$ and $\vec{s} \in \times_{j \in [n]} \Sigma_j$, we will denote i 's *expected allocation* $\mathbf{E}_{f \sim M}[f(\vec{s})_i]$ by $x_i(\vec{s})$ and i 's *expected payment* $\mathbf{E}_{f \sim M}[f(\vec{s})_{n+i}]$ by $p_i(\vec{s})$. For $i \in [n]$ and $\vec{s} \in \times_{j \in [n]} \Sigma_j$, the *expected utility* of buyer i is $x_i(\vec{s})v_i - p_i(\vec{s})$. The auctioneer is interested in maximizing the *revenue*

$\sum_{i \in [n]} p_i(\vec{s})$, and is assumed to have full knowledge of the valuation distribution, but not of the actual valuations of the buyers.

Mechanism M is *dominant strategy incentive compatible (dominant strategy IC)* iff for all $i \in [n]$ and $\vec{v} \in \times_{j \in [n]} \text{supp}(\pi_j)$ and $\vec{v} \in \text{supp}(\pi)$, $x_i(v_i, \vec{v}_{-i})v_i - p_i(v_i, \vec{v}_{-i}) \geq x_i(\vec{v})v_i - p_i(\vec{v})$. Mechanism M is *ex-post individually rational (ex-post IR)* iff for all $i \in [n]$ and $\vec{v} \in \text{supp}(\pi)$, $x_i(v)v_i - p_i(v) \geq 0$. For convenience we usually will not treat a mechanism as a probability distribution over outcomes, but rather as the result of a randomized procedure that interacts with the buyers. In this case we say that a mechanism is *implemented by* that procedure.

A *sequential posted price (SPP) mechanism* for an instance (n, π, k) is any mechanism that is implementable by iteratively selecting a buyer $i \in [n]$ that has not been selected in a previous iteration, and proposing a price p_i for the service, which the buyer may accept or reject. If i accepts, he gets the service and pays p_i , resulting in a utility of $v_i - p_i$ for i . If i rejects, he pays nothing and does not get the service, resulting in a utility of 0 for i . Once the number of buyers that have accepted an offer equals k , the process terminates. Randomization in the selection of the buyers and prices is allowed. We will initially be concerned with only sequential posted price mechanisms. Later in the paper we define the two generalizations of SPP mechanisms that we mentioned in the introduction.

Our focus in this paper is on the maximum expected revenue of the SPP mechanisms, and some of its generalizations. Note that each buyer in a SPP mechanism has an obvious dominant strategy: he will accept whenever the price offered to him does not exceed his valuation, and he will reject otherwise. Also, a buyer always ends up with a non-negative utility when participating in a SPP mechanism. Thus, by the revelation principle (see, e.g., [9]), a SPP mechanism can be converted into a dominant strategy IC and ex-post IR direct revelation mechanism with the same expected revenue. Therefore, we compare the maximum expected revenue $REV(M)$ achieved by an SPP mechanism M to OPT , where OPT is defined as the maximum expected revenue that can be obtained by a mechanism that is dominant strategy IC and ex-post IR.

A more general solution concept is formed by the *ex-post incentive compatible*, ex-post individually rational mechanisms. Specifically, let (n, π, k) be an instance and M be a randomized direct revelation mechanism for that instance. Mechanism M is *ex-post incentive compatible (ex-post IC)* iff for all $i \in [n]$, $s_i \in \text{supp}(\pi_i)$ and $\vec{v} \in \text{supp}(\pi)$, $x_i(\vec{v})v_i - p_i(\vec{v}) \geq x_i(s_i, \vec{v}_{-i})v_i - p_i(s_i, \vec{v}_{-i})$. In other words, a mechanism is *ex-post IC* if it is a pure equilibrium for the buyers to always report their valuation. In this work we sometimes compare the expected revenue of our (dominant strategy IC and ex-post IR) mechanisms to the maximum expected revenue of the more general class of ex-post IC, ex-post IR mechanisms. This strengthens our positive results. We refer the reader to [25] for a further discussion of and comparison between various solution concepts.

2 Sequential Posted Price Mechanisms

We are interested in designing a posted price mechanism that, for any given n and valuation distribution π , achieves an expected revenue that is a constant

approximation of the optimal expected revenue achievable by a dominant strategy IC, ex-post IR mechanism. Theorem 1 shows that this is impossible.

Theorem 1. *For all $n \in \mathbb{N}_{\geq 2}$, there exists a valuation distribution π such that for all $k \in [n]$ there does not exist a sequential posted price mechanism for instance (n, π, k) that extracts a constant fraction of the expected revenue of the optimal dominant strategy IC, ex-post IR mechanism.*

Proof sketch. Fix $m \in \mathbb{N}_{\geq 1}$ arbitrarily, and consider the case where $n = 1$ and the valuation v_1 of the single buyer is taken from $\{1/a : a \in [m]\}$ distributed so that $\pi_1(1/a) = 1/m$ for all $a \in [m]$. In this setting, an SPP mechanism will offer the buyer a price p , which the buyer accepts iff $v_1 \geq p$. After that, the mechanism terminates. We show that this mechanism achieve only a fraction $\frac{1}{H(m)}$ of the social welfare. We then extend this example to a setting where the expected revenue of the optimal dominant strategy IC, ex-post IR mechanism is equal to the expected optimal social welfare. \square

The above impossibility result holds also in the continuous case, even if a large set of popular assumptions hold simultaneously, namely, the valuation distribution π has support $[0, 1]^n$; the expectation $\mathbf{E}_{\vec{v} \sim \pi}[v_i]$ is finite for any $i \in [n]$; π is symmetric in all its arguments; π is continuous and nowhere zero on $[0, 1]^n$; the conditional marginal densities $\pi_{i|\vec{v}_{-i}}$ are nowhere zero for any $\vec{v}_{-i} \in [0, 1]^{n-1}$ and any $i \in [n]$; π has a monotone hazard rate and is regular; π satisfies affiliation.

Roughgarden and Talgam-Cohen [25] showed that when all these assumptions are simultaneously satisfied, the optimal ex-post IC and ex-post IR mechanism is the Myerson mechanism that is optimal also in the independent value setting. Thus, these conditions make the correlated setting in some sense similar to the independent one with respect to revenue maximization. Yet our result show that, whereas SPP mechanism can achieve a constant approximation revenue for independent distributions, this does not hold for correlated ones.

A Revenue Guarantee for Sequential Posted Price Mechanisms. More precisely, in our lower bound instances constructed in the proof of Theorem 1, it is the case that the expected revenue extracted by every posted price mechanism is a $\Theta(1/\log(r))$ fraction of the optimal expected revenue, where r is the ratio between the highest valuation and the lowest valuation in the support of the valuation distribution. A natural question that arises is whether this is the worst possible instance in terms of revenue extracted, as a function of r . It turns out that this is indeed the case, asymptotically. The proofs use a standard bucketing technique (see, e.g., [4]) and can be found in the full paper [2].

We start with the unit supply case. For a valuation distribution π on \mathbb{R}^n , let v_π^{\max} and v_π^{\min} be $\max\{v_i : v \in \text{supp}(\pi), i \in [n]\}$ and $\min\{\max\{v_i : i \in [n]\} : v \in \text{supp}(\pi)\}$ respectively. Let $r_\pi = v_\pi^{\max}/v_\pi^{\min}$ be the ratio between the highest and lowest coordinate-wise maximum valuation in the support of π .

Proposition 1. *Let $n \in \mathbb{N}_{\geq 1}$, and let π be a probability distribution on \mathbb{R}^n . For the unit supply case there exists a SPP mechanism that, when run on instance*

$(n, \pi, 1)$, extracts in expectation at least an $\Omega(1/\log(r_\pi))$ fraction of the expected revenue of the optimal social welfare (and therefore also of the expected revenue of the optimal dominant strategy IC and ex-post IR auction).

This result can be generalized to yield revenue bounds for the case of k -limited supply, where $k > 1$. The above result does not always guarantee a good revenue; for example in the extreme case where $v_\pi^{\min} = 0$. However, it is easy to strengthen the above theorem such that it becomes useful for a wide class of distributions.

3 Enhanced Sequential Posted Price Mechanisms

We propose a generalization of sequential posted price mechanisms, in such a way that they possess the ability to retrieve valuations of some buyers.

Specifically, an *enhanced sequential posted price (ESPP) mechanism* for an instance (n, π, k) is a randomized mechanism that can be implemented by iteratively selecting a buyer $i \in [n]$ that has not been selected in a previous iteration, and performing exactly one of the following actions on buyer i :

- Propose service at price p_i to buyer i , which the buyer may accept or reject. If i accepts, he gets the service and pays p_i , resulting in a utility of $v_i - p_i$ for i . If i rejects, he pays nothing and does not get the service, resulting in a utility of 0 for i .
- Ask i for his valuation. (Buyer i pays nothing and does not get service.)

This generalization is still dominant strategy IC and ex-post IR.

Next we analyze the revenue performance of ESPP mechanisms. For this class of mechanisms we prove that, it is unfortunately still the case that no constant fraction of the optimal revenue can be extracted. Specifically, the next theorem establishes an $O(1/n)$ bound for ESPP mechanisms.

Theorem 2. *For all $n \in \mathbb{N}_{\geq 2}$, there exists a valuation distribution π such that for all $k \in [n]$ there does not exist a ESPP mechanism for instance (n, π, k) that extracts more than a $O(1/n)$ fraction of the expected revenue of the optimal dominant strategy IC, ex-post IR mechanism.*

Proof sketch. Let $n \in \mathbb{N}$ and $m = 2^n$. We specify an instance I_n with n buyers, and prove that $\lim_{n \rightarrow \infty} RM(I_n)/OR(I_n) = 0$, where $RM(I_n)$ is the largest expected revenue achievable by any ESPP mechanism on I_n , and $OR(I_n)$ is the largest expected revenue achievable by a dominant strategy IC, ex-post IR mechanism. I_n is defined as follows. Fix ϵ such that $0 < \epsilon < 1/nm^2$. The valuation distribution π is the one induced by the following process: (i) Draw a buyer i^* from the set $[n]$ uniformly at random; (ii) Draw numbers $\{c_j : j \in [n] \setminus \{i^*\}\}$ independently from $[m]$ uniformly at random; (iii) For all $j \in [n] \setminus \{i^*\}$, set $v_j = c_j \epsilon$; (iv) Set $v_{i^*} = ((\sum_{j \in [n] \setminus \{i^*\}} c_j) \bmod m + 1)^{-1}$. \square

However, ESPP mechanisms turn out to be more powerful than the standard sequential posted price. Indeed, contrary to SPP mechanisms, the ESPP mechanisms can be shown to extract a fraction of the optimal revenue that is independent of the valuation distribution. More precisely, the $O(1/n)$ bound turns out

to be asymptotically tight. Our main positive result for ESPP mechanisms is that when dependence of the valuation among the buyers is limited, then a constant fraction of the optimal revenue can be extracted. Specifically, we will define the concept of d -dimensional dependence and prove that for a d -dimensionally dependent instance, there is an ESPP mechanism that extracts an $\Omega(1/d)$ fraction of the optimal revenue.

It is natural to identify the basic reason(s) why, in the case of general correlated distributions, standard and enhanced sequential posted price mechanisms may fail to achieve a constant approximation of the optimum revenue. There are two main limitations of these mechanisms: (i) such mechanisms do not solicit bids or values from all buyers, and (ii) such mechanisms award items in a sequential manner. Although it is crucial to retrieve the valuation of *all* (but one of the) buyers, we show that it is possible to achieve a constant fraction of the optimum revenue by a mechanism that allocates items sequentially in an on-line manner, in contrast to previously known approximation results.

Randomized mechanism M is a *blind offer mechanism* iff it can be implemented as follows. Let (n, π, k) be an instance and let \vec{b} be the submitted bid vector. Then,

1. Terminate if $\vec{b} \notin \text{supp}(\pi)$.
2. Either terminate or select a buyer i from the set of buyers that have not yet been selected, such that the choice of i does not depend on \vec{b} .
3. Offer buyer i the service at price p_i , where p_i is drawn from a probability distribution that depends only on π_i, \vec{b}_{-i} (hence the distribution of p_i is determined by \vec{b}_{-i} and in particular does not depend on b_i).
4. Restart if the number of buyers who have accepted offers does not exceed k .

Note that the price offered to a buyer is entirely determined by the valuations of the remaining buyers, and is independent of what is reported by buyer i himself. Also the iteration in which a buyer is picked cannot be influenced by his bid. Nonetheless, blind offer mechanisms are in general not incentive compatible due to the fact that a bidder may be incentivized to misreport his bid in order to increase the probability of supply not running out before he is picked. However, blind offer mechanisms can easily be made incentive compatible as follows: let M be a non-IC blind offer mechanism, let \vec{b} be a bid vector and let $z_i(\vec{b})$ be the probability that M picks bidder i before supply has run out. When a bidder is picked, we adapt M by *skipping* that bidder with a probability $p_i(\vec{b})$ that is chosen in a way such that $z_i(\vec{b})p_i(\vec{b}) = \min\{z_i(b_i, \vec{b}_{-i}) : b_i \in \text{supp}(\pi_i)\}$. This is a blind offer mechanism in which buyer i has no incentive to lie, because now the probability that i is made an offer is independent of his bid. Doing this iteratively for all buyers yields a dominant strategy IC mechanism M' . Note that the act of *skipping* a bidder can be implemented by offering a price that is so high that a bidder will never accept it, thus M' is still a blind offer mechanism. Moreover, if the probability that any bidder in M is made an offer is lower bounded by a constant c , then in M' the probability that any bidder is offered a price is at least c . We apply this principle in the proof of Theorem 3 below in order to obtain a dominant strategy IC mechanism with a constant factor revenue performance.

It is not hard to see that the classical Myerson mechanism for the *independent* single-item setting belongs to the class of blind offer mechanisms. Thus blind offer mechanisms are optimal when buyers' valuations are independent. We will prove next that when buyer valuations are *correlated*, blind offer mechanisms can always extract a constant fraction of the optimal revenue, even against the ex-post IC, ex-post IR solution concept. Other mechanisms that achieve a constant approximation to the optimal revenue have been defined by Ronen [24], and then by Chawla et al. [10] and Dobzinski et al. [15]. However, these mechanisms allocate the items to profit-maximizing buyers. Thus, they are different from blind offer mechanisms in which the allocation is on-line.

Theorem 3. *For every instance (n, π, k) , there is a dominant strategy IC blind offer mechanism for which the expected revenue is at least a $(2 - \sqrt{e})/4 \approx 0.088$ fraction of the maximum expected revenue that can be extracted by an ex-post IC, ex-post IR mechanism.*

We need to establish some intermediate results in order to build up to a proof for the above theorem. First, we derive an upper bound on the revenue of the optimal ex-post IC, ex-post IR mechanism. For a given instance (n, π, k) , consider the linear program with variables $(y_i(\vec{v}))_{i \in [n], \vec{v} \in \text{supp}(\pi)}$ where the objective is $\max \sum_{i \in [n]} \sum_{\vec{v}_{-i} \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_{-i}) \sum_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})} \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i] v_i y_i(v_i, \vec{v}_{-i})$ subject to the constraints $\forall i \in [n], \vec{v}_{-i} \in \text{supp}(\pi_{-i}): \sum_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})} y_i(\vec{v}) \leq 1; \forall \vec{v} \in \text{supp}(\pi): \sum_{i \in [n]} \sum_{v'_i \in \text{supp}(\pi_{i, \vec{v}_{-i}}): v'_i \leq v_i} y_i(v'_i, \vec{v}_{-i}) \leq k; \vec{v} \in \text{supp}(\pi): y_i(\vec{v}) \geq 0 \forall i \in [n]$. The next lemma states that the solution to this linear program forms an upper bound on the revenue of the optimal mechanism.

Lemma 1. *For any instance (n, π, k) , above linear program upper bounds the maximum expected revenue achievable by an ex-post IC, ex-post IR mechanism.*

Proof sketch. We first prove that a monotonicity constraint holds on the set of possible allocations that a ex-post IC, ex-post IR mechanism can output. Moreover, we show that the prices charged by the mechanism cannot exceed a certain upper bound given in terms of allocation probabilities. Then, we formulate a new linear program whose optimal value equals the revenue of the optimal ex-post IC, ex-post IR mechanism. We finally rewrite this new linear program into the one given above. This proof adapts the approach introduced in [18]. \square

We can now proceed to prove our main result about blind offer mechanisms. Let (n, π, k) be an arbitrary instance. Let $(y_i^*(\vec{v}))_{i \in [n]}$ be the optimal solution to the linear program given above corresponding to this instance. Let M_π^k be the blind offer mechanism that does the following: let \vec{v} be the vector of submitted valuations. Iterate over the set of buyers such that in iteration i , buyer i is picked. In iteration i , select one of the following options: offer service to buyer i at a price p for which it holds that $y_i^*(p, \vec{b}_{-i}) > 0$, or skip buyer i . The probabilities with which these options are chosen are as follows: Price p

is offered with probability $y_i^*(p, \vec{b}_{-i})/2$, and buyer i is skipped with probability $1 - \sum_{p' \in \text{supp}(\pi_{i, \vec{b}_{-i}})} y_i^*(p, \vec{b}_{-i})/2$. The mechanism terminates if k buyers have accepted an offer, or at iteration $n + 1$.

Proof sketch (of Theorem 3). We will show that the expected revenue of M_π^k is at least $\frac{2-\sqrt{e}}{4} \cdot \sum_{i \in [n]} \sum_{\vec{v}_{-i} \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_{-i}) \sum_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})} \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i] v_i y_i^*(v_i, \vec{v}_{-i})$, which, by Lemma 1 and the LP above, is a $(2 - \sqrt{e})/4$ fraction of the expected revenue of the optimal ex-post IC, ex-post IR mechanism.

For a vector of valuations $\vec{v} \in \text{supp}(\pi)$ and a buyer $i \in [n]$, denote by $D_{i, \vec{v}_{-i}}$ the probability distribution from which mechanism $M_\pi^k(\vec{v})$ draws a price that is offered to buyer i , in case iteration $i \in [n]$ is reached. We let V be a number that exceeds $\max\{v_i : i \in [n], \vec{v} \in \text{supp}(\pi)\}$ and represent by V the option where $M_\pi^k(\vec{v})$ chooses to skip buyer i , so that $D_{i, \vec{v}_{-i}}$ is a probability distribution on the set $\{V\} \cup \{v_i : y_i^*(v_i, \vec{v}_{-i}) > 0\}$. Then, $\mathbf{E}_{\vec{v} \sim \pi}[\text{revenue of } M_\pi^k(\vec{v})] \geq \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{\substack{p_i \in \text{supp}(D_{i, \vec{v}_{-i}}) \\ : p_i \leq v_i}} \frac{p_i y_i^*(p_i, \vec{v}_{-i})}{2} \Pr_{p_i \sim D_{i, \vec{v}_{-i}}} [|\{j \in [n-1] : p_j \leq v_j\}| < k]$. Then, by applying a Chernoff bound, we can prove that $\Pr_{v_i \sim D_{i, \vec{v}_{-i}}} [|\{j \in [n-1] : p_j \leq v_j\}| < k] \geq 1 - (\frac{e}{4})^{k/2} \geq 1 - (\frac{e}{4})^{1/2} = \frac{2-\sqrt{e}}{2}$. Hence, we have a lower bound of $(2 - \sqrt{e})/2$ on the probability that all players get selected. The theorem follows by combining this with the principle explained above that allows us to transform M_π^k into a dominant strategy IC blind offer mechanism. \square

Revenue Guarantees for ESPP Mechanisms. Finally, in this section we evaluate the revenue guarantees of the ESPP mechanisms in the presence of a form of limited dependence that we will call *d-dimensional dependence*, for $d \in \mathbb{N}$. These are probability distributions for which it holds that the valuation distribution of a buyer conditioned on the valuations of the rest of the buyers can be retrieved by only looking at the valuations of a certain subset of d buyers. Formally, a probability distribution π on \mathbb{R}^n is *d-dimensionally dependent* iff for all $i \in [n]$ there is a subset $S_i \subseteq [n] \setminus \{i\}$, $|S_i| = d$, such that for all $\vec{v}_{-i} \in \text{supp}(\pi_{-i})$ it holds that $\pi_{i, \vec{v}_{S_i}} = \pi_{i, \vec{v}_{-i}}$. Note that if $d = 0$, then π is a product of n independent probability distributions on \mathbb{R} . On the other hand, the set of $(n-1)$ -dimensionally dependent probability distributions on \mathbb{R}^n equals the set of all probability distributions on \mathbb{R}^n . This notion is useful in practice for settings where it is expected that a buyer's valuation distribution has a reasonably close relationship with the valuation of a few other buyers. As an example of one of these practical settings consider the case that there exists a true objective valuation v for the item or service, an expert buyer that knows this valuation precisely, and remaining buyers whose valuation is influenced by independent noise. It is then sufficient to know the valuation of a single buyer, namely the expert one, in order to retrieve the conditional distribution of any other buyer.

In general, d -dimensional dependence is relevant to many practical settings in which it is not necessary to have complete information about the valuations of all the other buyers in order to say something useful about the valuation of

a particular buyer. This rules out the extreme kind of dependence defined in the proof of Theorem 2; there the distributions are not $(n - 2)$ -dimensionally dependent, because for each buyer i it holds that the valuations of all buyers $[n] \setminus \{i\}$ are necessary in order to extract the valuation distribution of i conditioned on the others' valuations.

It is important to realize that the class of d -dimensionally dependent distributions is a strict superset of the class of *Markov random fields of degree d* . A Markov random field of degree d is a popular model to capture the notion of limited dependence. Anyway, d -dimensionally dependent distributions are more general: we show in [2] that there are distributions on \mathbb{R}^n that are 1-dimensionally dependent, but are not a Markov random field of degree less than $n/2$.

Theorem 4. *For every instance (n, π, k) where π is d dimensionally dependent, there exists an ESPP mechanism of which the expected revenue is at least a $(2 - \sqrt{e})/(16d) \geq 1/(46d) \in \Omega(1/d)$ fraction of the maximum expected revenue that can be extracted by an ex-post IC, ex-post IR mechanism.*

As a corollary we have that the bound of Theorem 2 is asymptotically tight.

Theorem 4 follows by combining Theorem 3 with the following lemma.

Lemma 2. *Let $\alpha \in [0, 1]$ and let (n, π, k) be an instance such that π is d -dimensionally dependent. If there is a blind offer mechanism that extracts in expectation at least an α fraction of the expected revenue of the optimal dominant strategy IC, ex-post IR mechanism, then there is an ESPP mechanism that extracts in expectation at least a $\alpha / \max\{4d, 1\}$ fraction of the expected revenue of the optimal ex-post IC, ex-post IR mechanism.*

4 Open Problems

Besides improving approximation bounds established in the present paper, there are many other interesting further research directions. For example, it would be interesting to investigate revenue guarantees under the additional constraint that the sequential posted price mechanism be *on-line*, i.e., the mechanism has no control over which buyers to pick, and should perform well for any possible ordering. We are also interested in the role of randomization in our ESPP mechanism that extracts $O(1/d)$ of the optimal revenue: in the current proof buyers are picked uniformly at random. Does there exist a deterministic ESPP mechanism that attains the same revenue guarantee, or is randomness a necessity?

An obvious and interesting research direction is to investigate more general auction problems. In particular, to what extent can ESPP mechanisms be applied to auctions having non-identical items? Additionally, can such mechanisms be applied to more complex allocation constraints or specific valuation functions for the buyers? The agents may have, for example, a demand of more than one item, or there may be a matroid feasibility constraint.

Acknowledgments. We thank Joanna Drummond, Brendan Lucier, Tim Roughgarden and anonymous referees for their constructive comments.

References

1. Adamczyk, M., Sviridenko, M., Ward, J.: Submodular stochastic probing on matroids. In: STACS 2014 (2014)
2. Adamczyk, M., Borodin, A., Ferraioli, D., de Keijzer, B., Leonardi, S.: Sequential posted price mechanisms with correlated valuations. CoRR, abs/1503.02200 (2015)
3. Babaioff, M., Dughmi, S., Kleinberg, R., Slivkins, A.: Dynamic pricing with limited supply. In: EC 2012 (2012)
4. Babaioff, M., Immorlica, N., Kleinberg, R.: Matroids, secretary problems, and online mechanisms. In: SODA 2007 (2007)
5. Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M.: A simple and approximately optimal mechanism for an additive buyer. In: FOCS 2014 (2014)
6. Blum, A., Hartline, J.D.: Near-optimal online auctions. In: SODA 2005 (2005)
7. Blum, A., Kumar, V., Rudra, A., Wu, F.: Online learning in online auctions. *Theor. Comput. Sci.* **324**, 137–146 (2004)
8. Blumrosen, L., Holenstein, T.: Posted prices vs. negotiations: an asymptotic analysis. In: EC 2008 (2008)
9. Börgers, T.: *An Introduction to the Theory of Mechanism Design*. Oxford University Press, Oxford (2015)
10. Chawla, S., Fu, H., Karlin, A.: Approximate revenue maximization in interdependent value settings. In: EC 2014, pp. 277–294 (2014)
11. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: STOC 2010 (2010)
12. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **11**(1), 17–33 (1971)
13. Cremer, J., McLean, R.P.: Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica* **56**(6), 1247–1257 (1988)
14. Devanur, N.R., Morgenstern, J., Syrgkanis, V., Weinberg, S.M.: Simple auctions with simple strategies. In: EC 2015, pp. 305–322 (2015)
15. Dobzinski, S., Fu, H., Kleinberg, R.D.: Optimal auctions with correlated bidders are easy. In: STOC 2011 (2011)
16. Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: SODA 2015 (2015)
17. Groves, T.: Incentives in teams. *Econometrica* **41**, 617–631 (1973)
18. Gupta, A., Nagarajan, V.: A stochastic probing problem with applications. In: Goemans, M., Correa, J. (eds.) IPCO 2013. LNCS, vol. 7801, pp. 205–216. Springer, Heidelberg (2013)
19. Hart, S., Nisan, N.: Approximate revenue maximization with multiple items. In: EC 2012 (2012)
20. Hartline, J.D., Roughgarden, T.: Simple versus optimal mechanisms. *SIGecom Exch.* **8**(1), 5:1–5:3 (2009)
21. Kleinberg, R., Leighton, T.: The value of knowing a demand curve: bounds on regret for online posted-price auctions. In: FOCS 2003 (2003)
22. Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: STOC 2012 (2012)
23. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
24. Ronen, A.: On approximating optimal auctions. In: EC 2001 (2001)
25. Roughgarden, T., Talgam-Cohen, I.: Optimal and near-optimal mechanism design with interdependent values. In: EC 2013 (2013)
26. Rubinstein, A., Weinberg, S.M.: Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In: EC 2015, pp. 377–394 (2015)

27. Sandholm, T.W., Gilpin, A.: Sequences of take-it-or-leave-it offers: near-optimal auctions without full valuation revelation. In: Faratin, P., Parkes, D.C., Rodríguez-Aguilar, J.-A., Walsh, W.E. (eds.) AMEC 2003. LNCS (LNAI), vol. 3048, pp. 73–91. Springer, Heidelberg (2004)
28. Segal, I.: Optimal pricing mechanisms with unknown demand. *Am. Econ. Rev.* **93**(3), 509–529 (2003)
29. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Financ.* **16**(1), 8–37 (1961)
30. Yan, Q.: Mechanism design via correlation gap. In: SODA 2011 (2011)