

Sequential Posted Price Mechanisms with Correlated Valuations

Marek Adamczyk
marek.adamczyk@tum.de
Technische Universität München

Allan Borodin
bor@cs.toronto.edu
University of Toronto

Diodato Ferraioli
dferraioli@unisa.it
University of Salerno

Bart de Keijzer
keijzer@cwi.nl
Centrum Wiskunde & Informatica (CWI), Amsterdam

Stefano Leonardi
leonardi@dis.uniroma1.it
Sapienza University of Rome

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 (i.e., everywhere-supported, continuous, symmetric, and normalized (conditional) distributions that satisfy *regularity*, the *MHR condition*, and *affiliation*). 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 ratio of the highest and lowest possible valuation. We prove that a simple generalization of these mechanisms allows to achieve a better revenue performance: 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 ($d = n - 1$). In order to prove this result 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 's. We prove that a constant fraction $(2 - \sqrt{e})/4 \approx 0.088$ of the optimal revenue can be always extracted by this generalized mechanism.

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 how much they are forced to pay. The reason for this is the *revelation principle* (see, e.g., Börgers [2015]), which implies that one may resort to studying only direct revelation mechanisms for many purposes, such as maximizing the social welfare or the auctioneer's revenue. Some of the most celebrated mechanisms follow this approach, such as the Vickrey-Clark-Groves Mechanism [Vickrey, 1961, Clarke, 1971, Groves, 1973] and the Myerson Mechanism [Myerson, 1981].

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 [2004] highlight that this assumption frequently fails because of: (i) a buyer's unwillingness to fully specify their values, (ii) 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 it is known that there is nothing to be gained from this behavior. The failure of direct revelation mechanisms is also confirmed by experimental studies (see, e.g., Kagel et al. [1987]).

This has recently motivated the research about auction mechanisms that are conceptually simple. Among these, the class of *sequential posted price mechanisms* [Chawla et al., 2010] is particularly attractive. First studied by Sandholm and Gilpin [2004] (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 for the service. Moreover, they are individually rational, i.e., informally, participation in such an auction is never harmful to the buyer. Finally, they possess a trivial dominant strategy: indeed, it is always best for a buyer to respond sincerely to any offer that the mechanism makes him. (equivalently, following the definition of Li [2015], we can say that these mechanisms are *obviously strategy proof*).

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 questions related to the existence of sequential posted price mechanisms that achieve a high revenue. That is, we quantify the quality of a mechanism by comparing its expected revenue to that of the *optimal mechanism*, defined as the mechanism that achieves the highest expected revenue among all dominant strategy incentive compatible and ex-post individually rational mechanisms. (In Section 2, we explain and motivate this optimality notion further.)

We assume a standard Bayesian, transferable, quasi-linear utility model and we study a *unit demand, single parameter, multi-unit* setting: there is one service (or type of item) being provided by the auctioneer, any buyer is interested in receiving the service (or item) 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. We focus in this paper 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 recent substantial work on the subject of revenue performance for *simple* mechanisms [Hartline and Roughgarden, 2009, Hart and Nisan, 2012, Babaioff et al., 2014, Rubinstein and Weinberg, 2015, Devanur et al., 2015]. In particular, Babaioff et al. [2014] highlight the importance of understanding what is the relative strength of simple versus complex mechanisms with regard to revenue maximization.

As described above, sequential posted price mechanisms are an example of such a simple class of mechanisms. Sandholm and Gilpin [2004] 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 [2008] 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. [2012] consider *prior-independent* posted price mechanisms with *k*-limited supply for the setting where the only information known about the valuation distribution is that all valuations are independently drawn from the same distribution with support $[0, 1]$. Posted-price mechanisms have also been previously studied in [Kleinberg and Leighton, 2003, Blum and Hartline, 2005, Blum et al., 2004], albeit for a non-Bayesian, on-line setting. In a recent work Feldman et al. [2015a] study “on-line” posted price mechanisms for combinatorial auctions when valuations are independently drawn.

The works of Chawla et al. [2010] and Gupta and Nagarajan [2013] are closest to our present work, although they only consider sequential posted price mechanisms in the independent values setting. In particular, Chawla et al. [2010] 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 *order-oblivious* (i.e., “on-line”) sequential posted price mechanisms in which the order of the buyers is fixed and adversarially determined. They use order-oblivious mechanisms in order to establish some results for the more general multi-parameter case. Yan [2011] builds on this work and strengthens some of the results of Chawla et al. [2010]. Moreover, Kleinberg and Weinberg [2012]

and Feldman et al. [2015b] prove results that imply a strengthening of some of the results of Chawla et al. [2010].

Gupta and Nagarajan [2013] introduce a more abstract stochastic probing problem that includes Bayesian sequential posted price mechanisms as well as the stochastic matching problem introduced by Chen et al. [2009]. Their approximation bounds were later improved by Adamczyk et al. [2014] who in particular matched the approximation of Chawla et al. [2010] for the case where there is a single matroid feasibility constraint.

All previous work only consider the independent setting. In this work, we instead focus on the correlated setting. The lookahead mechanism of Ronen [2001] is a fundamental reference for the correlated setting and resembles our blind offer mechanism but is different in substantial ways, as we will discuss below. Cremer and McLean [1988] made a fundamental contribution to auction theory in the correlated value setting, by characterizing exactly for which valuation distributions it is possible to extract the full optimal social welfare as revenue. They do this for the ex-post IC, interim IR mechanisms and for the dominant strategy IC, interim IR mechanisms. Segal [2003] gives a characterization of optimal ex-post incentive compatible and ex-post individually rational optimal mechanisms. Roughgarden and Talgam-Cohen [2013] study optimal mechanism design in 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 [Dobzinski et al., 2011, Roughgarden and Talgam-Cohen, 2013, Chawla et al., 2014], that develops mechanisms with good approximation guarantees for revenue maximization in the correlated setting. These mechanisms build on the lookahead mechanism of Ronen [2001] and thus they also differ from the mechanisms proposed in this work.

Contributions and outline We define some preliminaries and notation in Section 2. In Section 3.1 we first present a simple sequence of instances. It shows that for unrestricted correlated distributions, the expected revenue of the best sequential posted price mechanism does not approximate within any constant factor the expected revenue of 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.

The maximum revenue that a sequential posted price mechanism can generate on our examples is shown to be characterized by the logarithm of the ratio between the highest and lowest valuations in the support of the distribution. We show in Section 3.2 that this approximation ratio is essentially tight.

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 mechanisms*. An enhanced sequential posted price 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 enhanced sequential posted price mechanism requires that some fraction of buyers reveal their valuation truthfully. Thus, the original property that the bidders not have to reveal their preferences is *partially* sacrificed, in return for a more powerful class of mechanisms and (as we will see) a better revenue performance. For practical implementation, such mechanisms can be slightly adjusted by providing a bidder with a small monetary reward in case he is asked to reveal his valuation¹.

For the enhanced sequential posted price mechanisms, we prove that again there are instances in which the revenue is not within a constant fraction of the optimal revenue. However, we show that this class of mechanisms can extract a fraction $\Theta(1/n)$ of the optimal revenue, i.e., a fraction that is independent 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 their submitted bids; it is possible to offer a price to a buyer only when selected; and the buyer gets the service only if it accepts the

¹We also note that in some realistic scenarios, the valuation of some buyers may be known a priori to the auctioneer (through, for example, the repetition of auctions or accounting and profiling operations). In such settings, though it is not necessary to query such a buyer for his valuation, our results for enhanced sequential posted price mechanisms are also relevant.

offered price), except that the price offered to a bidder i may now depend on the bids submitted by *all* buyers other than i . This generalization sacrifices entirely the property that buyers valuations do not need to be revealed, and 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. Also, blind offer mechanisms 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. Unfortunately it turns out that blind offer mechanisms are not inherently truthful, although this is only for a subtle reason: while it is true that a buyer cannot influence the price he is offered nor the decision of the mechanism to pick him, he can still be incentivized to misreport his valuation in order to *influence the probability that the supply has not run out before the buyer is picked*. However, this is not a big obstruction, as we will show that there is a straightforward way to turn a blind offer mechanism into an incentive compatible one, with only a small loss in approximation factor.

Even though blind offer mechanisms sacrifice some simplicity (and practicality), it is still interesting that a mechanism that allocates items to buyers on-line (i.e., it works regardless of the order in which buyers are served) is able to achieve a constant approximation of the optimal revenue even with correlated valuations. On the contrary, all previously known mechanisms (as, e.g., Ronen [2001] and Chawla et al. [2010])) that achieve a constant approximation with correlated valuations allocate the items to the buyers that maximize the profit. Thus they do not work for every order, but only for very specific ones.

Moreover, our result for blind offer mechanisms has a constructive purpose: it provides the intermediate step in the tentative of establishing revenue approximation bounds for other mechanisms. We will show how blind offer mechanisms serve to this purpose in Section 4 for the construction of good enhanced sequential posted price mechanisms.

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. An exception is Section 4, where 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 or dependence that may exist among the buyers. This stands in contrast with common 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 enhanced sequential posted price mechanisms then states that when the valuation distribution is d -dimensionally dependent, there exists an enhanced sequential posted price mechanism that extracts an $\Omega(1/d)$ fraction of the optimal revenue. The proof of this result consists of three key ingredients:

- An upper bound on the optimal ex-post IC, ex-post IR revenue in terms of the solution of a linear program. This upper bound has the form of a relatively simple expression that is important for the definition and analysis of a blind offer mechanism that we define subsequently. This part of the proof generalizes a linear programming characterization introduced by Gupta and Nagarajan [2013] for the independent distribution setting.
- A proof for the fact that incentive compatible blind offer mechanisms are powerful enough to extract a constant fraction of the optimal revenue of any instance. (Note that this is in itself the main result that we prove for blind offer mechanisms.) This makes crucial use of the linear program mentioned above.
- A *conversion lemma* showing that blind offer mechanisms can be turned into enhanced sequential posted price mechanisms while maintaining a fraction $\Omega(1/d)$ of the revenue of the blind offer mechanism².

While our focus is on proving the (*non*-)existence of simple mechanisms that perform well in terms of revenue, we note the following about the computational complexity of our mechanisms: all the mechanisms that we use in our

²The sequential posted price mechanism constructed in this lemma depends on a parameter q which can be adjusted to trade off the amount of valuation elicitation against the hidden constant in the $\Omega(1/d)$ expression.

positive results run in polynomial time when the valuation distribution is given as a description of the valuation vectors together with their probability mass³.

Additionally, we note that all of our negative results hold for randomized mechanisms. On the other hand, our positive results only require randomization in a limited way. For our positive results for classical sequential posted price mechanisms, only the offered prices need to be randomly chosen, while the order in which the buyers are picked is arbitrary. This makes these positive results hold through for the *order-oblivious* (i.e., on-line) setting in which the mechanism has no control over the agent that is picked in each iteration. Our positive result for blind offer mechanisms only requires randomized pricing in case $k < n$ and works for any ordering in which the buyers are picked, as long as the mechanisms knows the ordering in advance. Our positive result for enhanced sequential posted price mechanisms requires randomized pricing and the assumption that the mechanism can pick a uniformly random ordering of the buyers (i.e., it holds in the *random order model* (ROM) of arrivals).

Some of the proofs have been omitted in the main body of the paper, and can be found in the appendices.

2 Preliminaries and notation

For $a \in \mathbb{N}$, we write $[a]$ to denote the set $\{1, \dots, a\}$. We write $\mathbf{1}[X]$ to denote the indicator function for property X (i.e., it evaluates to 1 if X holds, and to 0 otherwise). When \vec{v} is a vector and a is an arbitrary element, we denote by (a, \vec{v}_{-i}) 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. As mentioned in the introduction, the buyers have valuations for the service offered, which are drawn from a *valuation distribution*, defined as follows.

Definition 1 (Valuation distribution). *A valuation distribution π for n buyers is a probability distribution on $\mathbb{R}_{\geq 0}^n$.*

We will assume throughout this paper that π is discrete, except for in Theorem 2, where we assume that various standard assumptions about continuous valuation distributions hold, such as *regularity*, the *monotone hazard rate* (MHR) condition and *affiliation*. We refer the interested reader to Appendix A for a definition and a brief discussion of these properties.

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

Each of the buyers is interested in receiving the service at most once. The auctioneer runs a *mechanism* with which the buyers interact. In general, a mechanism consists of (i) the strategies available to the buyers, (ii) a function that maps each vector of strategies chosen by the buyers to an outcome. The mechanism, when provided with a strategy profile of the buyers, outputs an outcome that consists of a vector $\vec{x} = (x_1, \dots, x_n)$ and a vector $\vec{p} = (p_1, \dots, p_n)$: vector \vec{x} is the *allocation vector*, i.e., the $(0, 1)$ -vector that indicates to which of the buyers the auctioneer allocates the service, and $\vec{p} = (p_1, \dots, p_n)$ is the vector of *prices* that the auctioneer asks from the buyers. For outcome (\vec{x}, \vec{p}) , the utility of a buyer $i \in [n]$ is $x_i v_i - p_i$. The auctioneer is interested in maximizing the *revenue* $\sum_{i \in [n]} p_i$, and is assumed to have full knowledge of the valuation distribution, but not of the actual valuations of the buyers.

We formalize the above as follows.

Definition 2. *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 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*

³When the valuation distribution is not accessible in such a form, and can instead only be sampled from, then standard sampling techniques can be used in order to obtain an estimate of the distribution. When this estimate is reasonably accurate (which should be the case when the distribution in question does not have extreme outliers), then the mechanisms in our paper can still be used with only a small additional loss in revenue.

$i \in [n]$, 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 $p_i(\vec{s})$. (Whenever we use this notation, the mechanism M will always be clear from context.)

Definition 3. Let (n, π, k) be an instance and M be a randomized direct revelation mechanism for (n, π, k) . Mechanism M is dominant strategy incentive compatible (dominant strategy IC) iff for all $i \in [n]$ and $\vec{s} \in \times_{j \in [n]} \text{supp}(\pi_j)$ and $\vec{v}_i \in \text{supp}(\pi_i)$,

$$x_i(v_i, \vec{s}_{-i})v_i - p_i(v_i, \vec{s}_{-i}) \geq x_i(\vec{s})v_i - p_i(\vec{s}).$$

Mechanism M is ex-post individually rational (ex-post IR) iff for all $i \in [n]$ and $\vec{s} \in \text{supp}(\pi)$,

$$x_i(s)v_i - p_i(s) \geq 0.$$

For convenience we usually will not treat a mechanism explicitly as a probability distribution over deterministic mechanisms, but rather as the result of a randomized procedure that interacts in some way with the buyers. In this case we say that a mechanism is *implemented by* that procedure. *Sequential posted price mechanisms* are the mechanisms that are implemented by a particular such procedure, defined as follows.

Definition 4. An sequential posted price 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 and study the two generalizations of sequential posted price mechanisms that we mentioned in the introduction.

Note that each buyer in a sequential posted price 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 sequential posted price mechanism. Thus, by the revelation principle (see, e.g., [Börger, 2015]), a sequential posted price mechanism can be straightforwardly converted into a dominant strategy IC and ex-post IR direct revelation mechanism that achieves the same expected revenue.

Our interest lies in analyzing the revenue performance of sequential posted price mechanisms. We do this by comparing the expected revenue of such mechanisms to the maximum expected revenue that can be obtained by a mechanism that satisfies dominant strategy IC and ex-post IR. Thus, for a given instance let OPT be the maximum expected revenue that can be attained by a dominant strategy IC, ex-post IR mechanism and let $REV(C)$ be the maximum expected revenue achievable by some class of mechanisms C . Our goal throughout this paper is to derive instance-independent lower and upper bounds on the ratio $REV(C)/OPT$, when C is the class of sequential posted price mechanisms or one of the generalizations mentioned.

A more general class of mechanisms is formed by the *ex-post incentive compatible*, ex-post individually rational mechanisms.

Definition 5. Let (n, π, k) be an instance and M be a randomized direct revelation mechanism for (n, π, k) . 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 always (i.e., for any valuation vector) a pure equilibrium for the buyers to 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 interested reader to Roughgarden and Talgam-Cohen [2013] for a further discussion of and comparison between various notions of incentive compatibility and individual rationality.

3 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 lies only a constant factor away from the optimal expected revenue that can be achieved by a dominant strategy IC, ex-post IR mechanism. In this section we show that this is unfortunately impossible. In fact, we will show that the non-constant approximation ratios established for the examples in this section are asymptotically optimal in terms of a distributional parameter to be defined in Section 3.2.

3.1 Non-existence of good posted price mechanisms

We next prove the following theorem.

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. We first consider the unit supply setting, i.e., instances of the form $(n, \pi, 1)$. As a first step, we show that a posted price mechanism cannot achieve a constant factor approximation of the expected *optimal social welfare*, defined as:

$$\text{OSW} = \mathbf{E}_{\tilde{v} \sim \pi}[\max\{v_i : i \in [n]\}].$$

Let OR be the optimal revenue that a dominant strategy IC and ex-post IR mechanism can achieve on (n, π, k) . (OR depends on the valuation distribution π , but we assume that the valuation distribution is given, and implicit from context.) It is clear that OSW is an upper bound to OR regardless of π , since a dominant strategy IC and ex-post IR mechanism will not charge (in expectation) any buyer a higher price than its expected valuation.

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 such that $\pi_1(1/a) = 1/m$ for all $a \in [m]$. In this setting, a posted price mechanism will offer the buyer a price p , which the buyer subsequently accepts iff $v_1 \geq p$. After that, the mechanism terminates.

Note that $\text{OSW} = (1/m) \sum_{a=1}^m (1/a)$. The expected revenue of the mechanism is

$$\text{RM} = p \Pr_{v_1 \sim \pi_1}[v_1 \geq p] = p \frac{|\{a : 1/a \geq p\}|}{m} = \frac{|\{a : 1/a \geq 1/p^{-1}\}|}{mp^{-1}} = \frac{p^{-1}}{mp^{-1}} = \frac{1}{m}. \quad (1)$$

Therefore:

$$\lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OSW}} = \lim_{m \rightarrow \infty} \frac{1}{\sum_{a \in [m]} 1/a} = \lim_{m \rightarrow \infty} \frac{1}{H(m)} = 0,$$

where $H(m)$ denotes the m -th harmonic number. So, no posted price mechanism can secure in expectation a revenue that lies a constant factor away from the expected optimal social welfare. (Because our analysis is for an instance with only one buyer, this inapproximability result also holds for instances with independent valuations.)

We extend the above example in a simple way 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.

Fix $m \in \mathbb{N}_{\geq 1}$ and consider a setting with 2 buyers, where the type vector (v_1, v_2) takes values in $\{(1/a, 1/a) : a \in [m]\}$ according to the probability distribution where $\pi((1/a, 1/a)) = 1/m$ for all $a \in [m]$. A mechanism that always gives Buyer 1 the service, and charges Buyer 1 the bid of Buyer 2, is clearly dominant strategy IC and also clearly achieves a revenue equal to the optimal social welfare.

In this two-buyer setting, the value OSW is again $(1/m) \sum_{a=1}^m (1/a)$. By symmetry, we may assume without loss of generality that a posted price mechanism works by first proposing a price p_1 to Buyer 1, and then proposing a price p_2 to Buyer 2 if Buyer 1 rejected the offer. Using arguments similar to (1), we derive that the revenue of this mechanism is:

$$\begin{aligned} \text{RM} &= p_1 \Pr_{(v_1, v_2) \sim \pi}[v_1 \geq p_1] + p_2 \Pr_{(v_1, v_2) \sim \pi}[v_1 < p_1 \cap v_2 \geq p_2] \\ &= \frac{1}{m} + p_2 \Pr_{(v_1, v_2) \sim \pi}[v_1 < p_1 \cap v_2 \geq p_2] \leq \frac{1}{m} + p_2 \Pr_{(v_1, v_2) \sim \pi}[v_2 \geq p_2] = \frac{2}{m}. \end{aligned}$$

Therefore:

$$\lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OR}} = \lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OSW}} \leq \lim_{m \rightarrow \infty} \frac{2}{\sum_{a \in [m]} 1/a} = 0.$$

The above example establishes the non-existence of a good sequential posted price mechanism in the case where the service has to be provided to a single buyer. Suppose now that the service can be provided to two buyers (i.e., $k = 2$), and each buyer gets the service at most once. Consider again two buyers whose values are drawn from the probability distribution π as defined above. As above, by symmetry we may assume that our posted price mechanism first proposes price p_1 to Buyer 1, and then proposes either price p_2 or p'_2 to Buyer 2: p_2 is proposed in case the offer was rejected by Buyer 1, and p'_2 is proposed otherwise. The difference with the previous analysis for the unit supply case is that the mechanism proposes a price to Buyer 2 regardless of whether Buyer 1 accepted the offer or not.

We derive:

$$\begin{aligned} \text{RM} &= p_1 \Pr_{(v_1, v_2) \sim \pi}[v_1 \geq p_1 \cap v_2 < p'_2] + p_2 \Pr_{(v_1, v_2) \sim \pi}[v_1 < p_1 \cap v_2 \geq p_2] \\ &\quad + (p_1 + p'_2) \Pr_{(v_1, v_2) \sim \pi}[v_1 \geq p_1 \cap v_2 \geq p'_2] \\ &\leq \frac{2}{m} + (p_1 + p'_2) \Pr_{(v_1, v_2) \sim \pi}[v_1 \geq p_1 \cap v_2 \geq p'_2] \\ &\leq \frac{2}{m} + 2 \max\{p_1, p'_2\} \Pr_{(v_1, v_2) \sim \pi}[v_1 \geq \max\{p_1, p'_2\}] \leq \frac{4}{m}. \end{aligned}$$

The optimal incentive compatible mechanism works by giving the service to both buyers while charging the bid of Buyer 1 to Buyer 2, and charging the bid of Buyer 2 to Buyer 1. The resulting expected revenue is exactly the expected optimal social welfare: $\text{OR} = \text{OSW} = (1/m) \sum_{a=1}^m (2/a)$. We therefore obtain

$$\lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OR}} = \lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OSW}} \leq \lim_{m \rightarrow \infty} \frac{4}{\sum_{a \in [m]} 2/a} = 0.$$

The above yields an impossibility result for 2-limited supply. By adding to this instance dummy buyers that always have valuation 0, we obtain an impossibility result for k -limited supply, where $k \in \mathbb{N}$. \square

We prove that the above impossibility result holds also in the continuous case even if we assume that all of the following conditions simultaneously hold: 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. We remark that Roughgarden and Talgam-Cohen [2013] 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 the same mechanism that is optimal in the independent value setting. Thus, these conditions make the correlated setting very similar to the independent one with respect to revenue maximization. Yet, our following result shows that, whereas posted price mechanisms can achieve a constant approximation revenue in the independent setting, this result does not extend to the correlated setting.

Theorem 2. *There exists a valuation distribution π with the properties that*

1. π has support $[0, 1]^2$;
2. the expectation $\mathbf{E}_{\vec{v} \sim \pi}[v_i]$ is finite for any $i \in \{1, 2\}$;
3. π is symmetric in all its arguments;
4. π is continuous and nowhere zero on $[0, 1]^2$;
5. the conditional marginal densities $\pi_{i|\vec{v}_{-i}}$ are nowhere zero for any $\vec{v}_{-i} \in [0, 1]^{n-1}$ and any $i \in \{1, 2\}$;
6. π has a monotone hazard rate and is regular;
7. π satisfies affiliation;

such that there does not exist a sequential posted price mechanism for instance $(2, \pi, k)$, for $k \in \{1, 2\}$, that extracts a constant fraction of the expected revenue of the optimal dominant strategy IC, ex-post IR mechanism.

Proof. Consider $c > 1$ and $m \geq (c - \log c)/2$ and set $M = 1 + 1/m$. Let V be a random variable whose value is drawn over the support $[1/m, 1]$ according to the probability density function

$$f_V(x) = \frac{1}{(m-1)x^2}.$$

Let N_1 and N_2 be two random variables whose values are independently drawn over the support $[0, M - v]$ according to the conditional density function

$$f_{N|V}(z | V = v) = \frac{c^{-z} \ln(c)}{Z(v)}$$

with $c > 1$ and $Z(v) = 1 - c^{-(M-v)}$. Finally, let f be the probability density function of the pair $(X, Y) = (V + N_1 - 1/m, V + N_2 - 1/m)$.

Properties 1, 2 and 3 are trivial and can be immediately checked.

For $v \in [0, 1]$, let $f_{X|V}(\cdot | V = v)$ and $f_{Y|V}(\cdot | V = v)$ be respectively the probability density functions of X and Y conditioned on the event that $V = v$. In order to establish the remaining properties, observe that $f_{X|V}(x | V = v) = f_{N|V}(x - v + 1/m | V = v)$ if $x + 1/m \geq v$ and 0 otherwise. Equivalently, $f_{Y|V}(y | V = v) = f_{N|V}(y - v + 1/m | V = v)$ if $y + 1/m \geq v$ and 0 otherwise. Consider now the triple (X, Y, V) . The joint density function of this triple is

$$f_{X,Y,V}(x, y, v) = f_{X|Y,V}(x | Y = y, V = v) \cdot f_{Y|V}(y | V = v) \cdot f_V(v).$$

Note that $f_{X|Y,V}(x | Y = y, V = v) = f_{X|V}(x | V = v)$ if $\min\{x, y\} + 1/m \geq v$ and 0 otherwise. Then

$$f_{X,Y,V}(x, y, v) = f_{N|V}(x - v + 1/m | V = v) \cdot f_{N|V}(y - v + 1/m | V = v) \cdot f_V(v).$$

if $\min\{x, y\} + 1/m \geq v$ and 0 otherwise. Hence, we can compute f as follows:

$$f(x, y) = \int_{1/m}^1 f_{X,Y,V}(x, y, v) dv = \frac{\ln^2(c)}{m-1} \cdot c^{-(x+y)} \cdot \int_{1/m}^{\alpha+1/m} \frac{c^{2v}}{v^2 Z(v)^2} dv,$$

where $\alpha = \min\{1 - 1/m, x, y\}$. Note that the integrated function is continuous and positive in the interval in which it is integrated. Hence, the integral turns out to be non-zero. From this, we observe that $f(x, y)$ is continuous and nowhere zero on $[0, 1]^2$, satisfying Property 4.

Let us now derive the conditional probability density functions. By symmetry it will be sufficient to focus only on $f_{X|Y}$.

$$\begin{aligned} f_{X|Y}(x | Y = y) &= \int_{1/m}^1 f_{X|Y,V}(x | Y = y, V = v) \cdot f_V(v) dv = \frac{\ln(c)}{m-1} \cdot c^{-x} \cdot \int_{1/m}^{\alpha+1/m} \frac{c^v}{v^2 Z(v)} dv \\ &= \frac{m^2 c^M \ln(c)}{m-1} \cdot c^{-x} \cdot \int_0^\alpha \frac{1}{(mz+1)^2 (c^{1-z} - 1)} dz, \end{aligned}$$

where the last equality is obtained by setting $z = v - \frac{1}{m}$. It is now obvious that the conditional probability density functions are continuous and nowhere zero, as desired by Property 5.

Let $\gamma(z) = 1/((mz+1)^2 (c^{1-z} - 1))$, $g(a) = \int_0^a \gamma(z) dz$ with $a \in [0, 1]$ and let $\alpha' = \min\{y, 1 - 1/m\}$. Then

$$f_{X|Y}(x | Y = y) = \frac{m^2 c^M \ln(c)}{m-1} \cdot c^{-x} \cdot \begin{cases} g(x), & \text{if } x < \alpha'; \\ g(\alpha'), & \text{otherwise.} \end{cases}$$

Moreover, we have that

$$\begin{aligned} 1 - F_{X|Y}(x | Y = y) &= \int_x^1 f_{X|Y}(z | Y = y) dz \\ &= \frac{m^2 c^M \ln(c)}{m-1} \cdot \begin{cases} \int_x^{\alpha'} c^{-z} g(z) dz + g(\alpha') \int_{\alpha'}^1 c^{-z} dz, & \text{if } x < \alpha'; \\ g(\alpha') \int_x^1 c^{-z} dz = \frac{g(\alpha')(c^{-x} - c^{-1})}{\ln(c)}, & \text{otherwise.} \end{cases} \end{aligned}$$

Hence, the inverse hazard rate is

$$I(x) = \frac{1 - F_{X|Y}(x | Y = y)}{f_{X|Y}(x | Y = y)} = \begin{cases} \frac{\int_x^{\alpha'} c^{-z} g(z) dz + g(\alpha') \int_{\alpha'}^1 c^{-z} dz}{c^{-x} g(x)}, & \text{if } x < \alpha'; \\ \frac{1 - c^{x-1}}{\ln(c)}, & \text{otherwise.} \end{cases}$$

We prove that $I(x)$ is non-increasing in x in the interval $[0, 1]$ and thus f has the monotone hazard rate and is, as a consequence, regular, as required by Property 6.

Clearly, $I(x)$ is non-increasing in x in the interval $[\alpha', 1]$ since in this case $I(x) = (1 - c^{x-1})/\ln(c)$. Moreover, $I(x)$ does not have discontinuities for $x = \alpha'$. So, it is sufficient to show that $I(x)$ is non-increasing also in the interval $[0, \alpha']$. To this aim, observe that for $x < \alpha'$,

$$\begin{aligned} \frac{dI(x)}{dx} &= \frac{d}{dx} \frac{\int_x^{\alpha'} c^{-z} g(z) dz + g(\alpha') \int_{\alpha'}^1 c^{-z} dz}{c^{-x} g(x)} \\ &= \left(c^{-x} g(x) \frac{d}{dx} \int_x^{\alpha'} c^{-z} g(z) dz - \int_x^{\alpha'} c^{-z} g(z) dz \frac{d}{dx} c^{-x} g(x) \right. \\ &\quad \left. + c^{-x} g(x) g(\alpha') \frac{d}{dx} \int_{\alpha'}^1 c^{-z} dz - g(\alpha') \int_{\alpha'}^1 c^{-z} dz \frac{d}{dx} c^{-x} g(x) \right) / (c^{-x} g(x))^2. \end{aligned}$$

Observe that, according to the second fundamental theorem of calculus,

$$\frac{d}{dx} \int_x^{\alpha'} c^{-z} g(z) dz = -c^{-x} g(x),$$

whereas

$$\frac{d}{dx} c^{-x} g(x) = c^{-x} \gamma(x) - c^{-x} g(x) \ln(c),$$

and $\frac{d}{dx} \int_{\alpha'}^1 c^{-z} dz = 0$. Then,

$$\frac{dI(x)}{dx} = -1 + \left(\ln(c) - \frac{\gamma(x)}{g(x)} \right) \frac{\int_x^{\alpha'} c^{-z} g(z) dz + g(\alpha') \int_{\alpha'}^1 c^{-z} dz}{c^{-x} g(x)} = -1 + \left(\ln(c) - \frac{\gamma(x)}{g(x)} \right) I(x).$$

The result then follows by showing that $\gamma(x)/g(x) \geq \ln(c)$.

To this aim, let us consider the function $\gamma'(z) = c^z(c^{1-z} - 1)(mz + 1)^2$ for $z \in [0, x]$. Note that

$$\frac{d\gamma'(z)}{dz} = 2mc(mz + 1) \left(1 - \frac{2m + mz \log c + \log c}{c^{1-z}} \right) \leq 2mc(mz + 1) \left(1 - \frac{2m + \log c}{c} \right) \leq 0,$$

where we used the fact that $m \geq (c - \log c)/2$. Thus, $\gamma'(z)$ is non-increasing in its argument and, in particular,

$$\gamma'(z) \geq \gamma'(x) \geq (c^x - 1)(c^{1-x} - 1)(mx + 1)^2.$$

By simple algebraic manipulation, it then follows that $\gamma(z) \leq \gamma(x)c^z/(c^x - 1)$. Then

$$\frac{\gamma(x)}{g(x)} = \frac{\gamma(x)}{\int_0^x \gamma(z) dz} \geq \frac{\gamma(x)}{\int_0^x \frac{c^z}{c^x - 1} \gamma(x) dz} = \frac{c^x - 1}{\int_0^x c^z dz} = \ln(c),$$

as desired.

Set now $C = \ln^2(c)/(m - 1)$ and let

$$h(a) = \int_{1/m}^{\min\{1-1/m, a\}+1/m} \frac{c^{2v}}{v^2 Z(v)^2} dv$$

with $a \in [0, 1]$. Note that integrated function is positive for any $v \in [1/m, 1]$. Hence, the integral increases as the size of the interval in which it is defined increases. In other words, the function $h(a)$ is non-decreasing in a .

Consider now two pairs (x, y) and (x', y') . Moreover, let $\hat{x} = \max\{x, x'\}$ and $\check{x} = \min\{x, x'\}$ and, similarly, define \hat{y} and \check{y} . We show that $f(x, y)f(x', y') \leq f(\hat{x}, \hat{y})f(\check{x}, \check{y})$, satisfying in this way also Property 7.

Indeed,

$$f(x, y)f(x', y') = C^2 c^{-(x+y+x'+y')} h(\min\{x, y\})h(\min\{x', y'\}).$$

If $x \geq x'$ and $y \geq y'$ ($x < x'$ and $y < y'$, respectively), then $(\hat{x}, \hat{y}) = (x, y)$ ((x', y') , resp.) and $(\check{x}, \check{y}) = (x', y')$ ((x, y) , resp.), and the desired result immediately follows. Suppose instead that $(\hat{x}, \hat{y}) = (x, y')$ and $(\check{x}, \check{y}) = (x', y)$. Then

$$f(\hat{x}, \hat{y})f(\check{x}, \check{y}) = C^2 c^{-(x+y+x'+y')} h(\min\{x, y'\})h(\min\{x', y\}).$$

We will prove that in this case $h(\min\{x, y\})h(\min\{x', y'\}) \leq h(\min\{x, y'\})h(\min\{x', y\})$. First observe that on both sides one of the two factors must be $h(\min\{x, y, x', y'\}) = h(\min\{x', y\})$. Suppose without loss of generality, that $\min\{x', y\} = y$. Then it is sufficient to prove that $h(\min\{x', y'\}) \leq h(\min\{x, y'\})$, or, since h is non-decreasing, that $\min\{x', y'\} \leq \min\{x, y'\}$. If $x \leq y'$, then $x' \leq x \leq y'$ by hypothesis and the claim follows. If $y' < x$, then it immediately follows that $\min\{x', y'\} \leq y'$. The case that $(\hat{x}, \hat{y}) = (x', y)$ and $(\check{x}, \check{y}) = (x, y')$ is similar and hence omitted.

Finally, observe that $\lim_{c \rightarrow \infty} f_{N|V}(0 \mid V = v) = 1$ and $\lim_{c \rightarrow \infty} f_{N|V}(z \mid V = v) = 0$ for any $z > 0$. Hence, $\lim_{c \rightarrow \infty} X = \lim_{c \rightarrow \infty} Y = V - 1/m$.

Let us consider the case that the service can be offered to only one buyer. In this setting, the following is a dominant strategy IC and ex-post IR mechanism: it offers to Buyer 1 the service at a price equivalent to the valuation of Buyer 2 minus a fixed constant ϵ . For small enough c , ϵ can be chosen arbitrarily small. Thus, for any ϵ there exists a choice of c such that in expectation this mechanism extracts as revenue all but ϵ of the social welfare. The expected optimal social welfare (and thus the optimal expected revenue) is:

$$\text{OR} = \text{OSW} = \mathbf{E}_{v \sim f_v}[v - 1/m] = \int_{1/m}^1 \frac{v - 1/m}{(m-1)v^2} dv = \frac{\ln(m)}{m-1} - \frac{1}{m} \geq \frac{\ln(m-1)}{m-1}.$$

A posted price mechanism will offer Buyer 1 a price $p_1 \geq 0$, which the buyer subsequently accepts iff $X \geq p_1$. After that, if Buyer 1 rejects, the mechanism offers a price p_2 to Buyer 2. Thus, if $p_1 \in [0, 1 - 1/m]$, then

$$\begin{aligned} p_1 \mathbf{Pr}_{v \sim f_v}[X \geq p_1] &= p_1 \mathbf{Pr}_{v \sim f_v}[v \geq p_1 + 1/m] \\ &= p_1 \int_{p_1+1/m}^1 \frac{1}{(m-1)v^2} dv = \frac{p_1}{m-1} \left(\frac{1}{p_1+1/m} - 1 \right) \leq \frac{1-p_1}{m-1} \leq \frac{1}{m-1}. \end{aligned}$$

Moreover, if $p_1 \geq 1 - 1/m$, then $p_1 \mathbf{Pr}_{v \sim f}[X \geq p_1] = 0$. Hence,

$$\text{RM} = p_1 \mathbf{Pr}_{v \sim f}[v \geq p_1] + p_2 \mathbf{Pr}_{v \sim f}[v < p_1 \cap v \geq p_2] \leq \frac{1}{m-1} + p_2 \mathbf{Pr}_{v \sim f}[v \geq p_2] \leq \frac{2}{m-1}.$$

Therefore:

$$\lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OR}} = \lim_{m \rightarrow \infty} \frac{\text{RM}}{\text{OSW}} \leq \lim_{m \rightarrow \infty} \frac{2}{\ln(m-1)} = 0.$$

The case in which the service can be offered to both buyers is similar and omitted. \square

We note that above theorem can be easily extended to an arbitrary number of buyers, by adding dummy buyers whose valuation is independently drawn over the support $[0, 1]$ according to the probability density function

$$f(x) = \frac{c^{-x} \ln c}{1 - c^{-1}}.$$

Note that, with this extension, the resulting distribution π does not satisfy the symmetry condition anymore, but on the other hand, the symmetry property is not necessary for the optimality of the Myerson mechanism to hold (see Roughgarden and Talgam-Cohen [2013]).

3.2 A revenue guarantee for sequential posted price mechanisms

In the previous section we have demonstrated that it is impossible to have a sequential posted price mechanism extract a constant fraction of the optimal revenue. More precisely, in our example instances it was 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 . We show here that this is indeed the case, asymptotically: For every valuation distribution π , there exists a mechanism that extracts in expectation at least a $\Theta(1/\log(r))$ fraction of the revenue of the optimal revenue. We note that in many realistic scenarios, we do not expect the extremal valuations of the buyers to lie too far from each other, because often the valuation of a buyer is strongly impacted by prior objective knowledge of the value of the service to be auctioned. The results of this section are valuable when that is the case. We start with the unit supply case.

Definition 6. 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 sequential posted price mechanism that, when run on instance $(n, \pi, 1)$, extracts in expectation at least an $\Omega(1/\log(r_\pi))$ fraction of the optimal social welfare (and therefore also of the expected revenue of the optimal dominant strategy IC and ex-post IR auction).

Proof. The proof uses a standard bucketing trick (see, e.g., Babaioff et al. [2007]). Let M be the sequential posted price mechanism that draws a value p uniformly at random from the set $S = \{v_\pi^{\min}2^k : k \in [\lceil \log(r_\pi) \rceil] \cup \{0\}\}$. M offers price p to all the bidders in an arbitrary order, until a bidder accepts.

Let π_{\max} be the probability distribution of the coordinate-wise maximum of π . Note that $|S|$ does not exceed $\log(r_\pi^{\max})$. Therefore the probability that p is the highest possible value (among the values in S) that does not exceed the value drawn from π_{\max} , is equal to $1/\log(r_\pi)$. More formally, let π_S be the probability distribution from which p is drawn; then

$$\Pr_{v^{\max} \sim \pi_{\max}, p \sim \pi_S} [p \leq v^{\max} \cap (\nexists p' \in S : p' > p \wedge p' \leq v^{\max})] \leq \frac{1}{\log(r_\pi)}.$$

Thus, with probability $1/\log(r_\pi)$, the mechanism generates a revenue of exactly $v_\pi^{\min}2^k$, where k is the number such that the value drawn from π_{\max} lies in between $v_\pi^{\min}2^k$ and $v_\pi^{\min}2^{k+1}$. This implies that with probability $1/\log(r_\pi)$ the mechanism generates a revenue that lies a factor of at most $1/2$ away from the optimal social welfare $OPT(\vec{v})$ (i.e., the coordinate-wise maximum valuation):

$$\mathbf{E}_{\vec{v} \sim \pi} [\text{revenue of } M(\vec{v})] \geq \frac{1}{\log(r_\pi)} \frac{1}{2} \mathbf{E}_{\vec{v} \sim \pi} [OPT(\vec{v})] \geq \frac{1}{2 \log(r_\pi)} \mathbf{E}_{\vec{v} \sim \pi} [OPT(\vec{v})]. \quad \square$$

This result can be generalized to yield revenue bounds for the case of k -limited supply, where $k > 1$. We prove a more general variant of the above in Appendix C.

The above result does not always guarantee a good revenue; for example in the extreme case where $v_\pi^{\min} = 0$. However, it is straightforward to generalize the above theorem such that it becomes useful for a much bigger family of probability distributions: let \hat{v} and \check{v} be two any two values in the support of π_{\max} , and let $c(\hat{v}, \check{v}) = \Pr_{v^{\max} \sim \pi_{\max}} [\check{v} \leq v^{\max} \leq \hat{v}]$. Then by replacing the values v_π^{\max} and v_π^{\min} in the above proof by respectively \hat{v} and \check{v} , we obtain a sequential posted price mechanism that extracts in expectation a $c(\hat{v}, \check{v})/(2 \log(\hat{v}/\check{v}))$ fraction of the optimal social welfare. By choosing \hat{v} and \check{v} such that this ratio is maximized, we obtain a mechanism that extracts a significant fraction of the optimal social welfare in any setting where the valuation distribution of a buyer is concentrated in a relatively not too large interval.

A better result can be given for the unlimited supply case.

Definition 7. For a valuation distribution π on \mathbb{R}^n and any $i \in [n]$, let $v_{\pi,i}^{\max}$ and $v_{\pi,i}^{\min}$ be $\max\{v_i : \vec{v} \in \text{supp}(\pi)\}$ and $\min\{v_i : \vec{v} \in \text{supp}(\pi)\}$ respectively. Let $r_{\pi,i} = v_{\pi,i}^{\max}/v_{\pi,i}^{\min}$, be the ratio between the highest and lowest valuation of Buyer i in the support of π .

Proposition 2. Let $n \in \mathbb{N}_{\geq 1}$, and let π be a probability distribution on \mathbb{R}^n . There exists a sequential posted price mechanism that, when run on instance (n, π, n) , extracts in expectation at least an $\Omega(1/\log(\max\{r_{\pi,i} : i \in [n]\}))$ fraction

of the expected revenue of the expected optimal social welfare (and therefore also the expected revenue of the optimal dominant strategy IC and ex-post IR mechanism).

The proof (found in Appendix B) applies similar techniques as those in the proof of Proposition 1. The techniques used for proving these results can be applied to improve the approximation guarantees for the more general k -limited supply setting, for any $k \in [n]$, under special conditions. We give an example of such a result in Appendix C.

Clearly, the stated bound of $O(1/\log(\max\{r_{\pi,i}: i \in [n]\}))$ is very crude. For most practical settings we expect that it is possible to do a much sharper revenue analysis of the mechanisms in the proofs of the above propositions, by taking the particular valuation distribution into account. Moreover, as suggested above, also for the unlimited supply case it is possible to tweak the proof in a straightforward way in order to achieve a good revenue in cases where the ratios $r_{\pi,i}$ are very large. Finally, note that the mechanisms in the proofs of these two propositions do not take into account any dependence and correlation among the valuations of the buyers. When provided with a particular valuation distribution, a better revenue and sharper analysis may be obtained by taking such dependence into account, and adapting the mechanisms accordingly.

4 Enhanced sequential posted price mechanisms

Our negative results on sequential posted price mechanisms suggest that it is necessary for a mechanism to have a means to retrieve the valuations of some of the buyers in order to improve the revenue performance. We accordingly propose two generalizations of sequential posted price mechanisms, in such a way that they possess the ability to retrieve the valuations of some of the buyers.

Definition 8. *An enhanced sequential posted price mechanism for an instance (n, π, k) is a 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 :*

- Offer the 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 (in which case the buyer pays nothing and does not get the service).

Randomization is allowed.

This generalization is still dominant strategy IC and ex-post IR (that is, the revelation principle allows us to convert them to dominant strategy IC, ex-post IR direct revelation mechanisms). Actually, one can see that this mechanism still enjoys also the stronger property of being obviously strategy-proof [Li, 2015]. Indeed, for enhanced sequential posted price mechanisms, when a buyer gets asked his valuation, he has no incentive to lie, because in this case he does not get the service and he pays nothing⁴.

Next, we analyze the revenue performance of enhanced sequential posted price 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: The next section establishes an $O(1/n)$ bound for enhanced sequential posted price mechanisms. Enhanced sequential posted price mechanisms nonetheless turn out to be more powerful than standard sequential posted price mechanisms. Contrary to what we had for the former ones, enhanced mechanisms can be shown to extract a fraction of the optimal revenue that is independent of the valuation distribution: The $O(1/n)$ bound turns out to be asymptotically tight. Our main positive result for enhanced sequential posted price 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, in Section 4.3 we define the concept of d -dimensional dependence and prove that for an instance (n, π, k) where π is d -dimensionally dependent, there exists an enhanced sequential posted price mechanism that extracts an $\Omega(1/d)$ fraction of the optimal revenue. (This implies the claimed $\Omega(1/n)$ bound by taking $d = (n - 1)$.)

It is instructive to identify the basic reason(s) why, in the case of general correlated distributions, standard and enhanced sequential posted price mechanisms fail to achieve a constant approximation of the optimum revenue. There are two main limitations of (enhanced) sequential posted price mechanisms; namely, (i) such mechanisms do not

⁴A problematic aspect from a practical point of view is that while there is no incentive to a buyer to lie, there is also no incentive to tell the truth. This problem is addressed in Appendix D.

solicit valuations from all buyers, and (ii) such mechanisms are constrained to interact with buyers in a sequential manner. We next show that we can generalize enhanced sequential posted price mechanisms further, such that the first of these two limitations vanishes entirely, and prove that the resulting class of mechanisms is able to approximate the optimum revenue to within a constant. In contrast to previously known approximation results, this shows that it is possible to achieve a constant fraction of the optimum revenue by a mechanism that allocates items sequentially, and moreover, in an on-line manner. Specifically, we consider the following superclass of the enhanced sequential posted price mechanisms.

Definition 9. Let (n, π, k) be an instance and let M be a direct revelation mechanism for (n, π, k) . Mechanism M is a blind offer mechanism iff it can be implemented as follows. Let \vec{b} be the submitted bid vector:

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. Propose Buyer i to offer service at price p_i , where p_i is drawn from a probability distribution that depends only on $\pi_{i, \vec{b}_{-i}}$ (hence the distribution where p_i is drawn from is determined by \vec{b}_{-i} and it does not depend on b_i).
4. Go to Step 2 if there is supply left, i.e., if the number of buyers who have accepted offers does not exceed k .

Randomization is allowed.

In the above definition “offering a price” p_i to Buyer i in the context of a direct revelation mechanism means that we allocate i the service if $v_i \geq p_i$ and charge him a payment of p_i , and otherwise we do not allocate i the service and charge him nothing.

Remark 1. The price offered to a buyer i by a blind offer mechanism is entirely determined by the submitted valuations of the buyers other than i , 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, we note that 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 . This fact can be used to bound the loss in revenue by a factor of c . Indeed, we apply this principle in the proof of Theorem 4 below, in order to obtain a dominant strategy IC blind offer mechanism with a constant factor revenue performance⁵.

It is well known under the name of the *taxation principle* that (when restricted to our setting) a mechanism is dominant strategy IC if and only if it can be implemented by an algorithm that (roughly) works as follows: (i) the buyers are *simultaneously* presented with a payment that does not depend on their own bid; (ii) the items are allocated to the buyers for which the increase in value of some given, fixed, social choice function, is maximized. At first sight, such a mechanism seems to closely resemble the description of blind offer mechanisms. However, we emphasize that there are some significant differences. First, using the taxation principle, prices are set in advance of offering service, whereas in blind offer mechanisms prices are presented sequentially, thus the price offered to the i -th buyer can depend on the decisions taken by the previous buyers. Second, using the taxation principle, the winners are chosen at the end, whereas in blind offer mechanisms the winners are decided on-line, so that it is possible that items are allocated to buyers without maximizing the profit. Hence, there are dominant strategy IC mechanisms that cannot be implemented as blind offer mechanisms (and conversely, there are blind offer mechanisms that are not dominant strategy IC).

A blind offer mechanisms preserves the same conceptually simple structure as standard and enhanced posted price mechanisms, but it bases its proposed price to a buyer i on the set of *all* valuations other than that of i , whereas an

⁵More precisely, for the particular (non-truthful) blind offer mechanism that we propose and analyze in Section 4.2, it will turn out that applying the transformation described here does not result in any additional loss in revenue.

enhanced sequential posted price mechanism bases its proposal to a buyer i only on the valuations that have been *revealed* by buyers in *previous* iterations. This increases the power of blind offer mechanisms: for example, 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, there exist optimal incentive compatible blind offer mechanisms when buyers' valuations are independent. We will prove in Section 4.2 that, even when buyer valuations are *arbitrarily dependent*, blind offer mechanisms can always extract a constant fraction of the optimal revenue, even for the ex-post IC, ex-post IR solution concept. For correlated valuation distributions, other mechanisms that achieve a constant approximation to the optimal revenue have been defined by Ronen [2001], and then by Chawla et al. [2014] and Dobzinski et al. [2011]. However, these mechanisms, as in the taxation principle setting, allocate the items to revenue maximizing buyers. Thus, they are different from blind offer mechanisms in which the allocation is on-line.

4.1 Limitations of enhanced sequential posted price mechanisms

Here we show that enhanced sequential posted price mechanisms cannot extract a constant fraction of the expected revenue of the optimal dominant strategy IC, ex-post IR mechanisms. This is done by constructing a family of instances on which no enhanced sequential posted price mechanism can perform well.

Theorem 3. *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 enhanced sequential posted price 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. We prove this for the case of $k = n$. The proof is easy to adapt for different k .

Let $n \in \mathbb{N}$ and $m = 2^n$. We prove this claim by specifying an instance I_n with n buyers, and proving that $\lim_{n \rightarrow 0} \text{RM}(I_n)/\text{OR}(I_n) = 0$, where $\text{RM}(I_n)$ denotes the largest expected revenue achievable by any enhanced sequential posted price mechanism on I_n , and $\text{OR}(I_n)$ denotes the largest expected revenue achievable by a dominant strategy IC, ex-post IR mechanism.

I_n is defined as follows. Let $\epsilon \in \mathbb{R}_{>0}$ be a number smaller than $1/nm^2$. 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

$$v_{i^*} = \frac{1}{\left(\sum_{j \in [n] \setminus \{i^*\}} c_j\right)_{\text{mod } m} + 1}. \quad (2)$$

Observe that for this distribution it holds that for all $i \in [n]$, the valuation v_i is uniquely determined by the valuations $(v_j)_{j \in [n] \setminus \{i\}}$. The optimal (direct revelation) mechanism can therefore extract the total optimal social welfare as its revenue, as follows: it provides service to every buyer, and sets the payment as follows. Let b_i be the bid, i.e., the reported valuation of Buyer i . Then,

- if $b_j < 1/m$ for all $j \in [n] \setminus \{i\}$, charge i a price of

$$\frac{1}{\left(\sum_{j \in [n] \setminus \{i\}} b_j / \epsilon\right)_{\text{mod } m} + 1};$$

- otherwise, if there is a buyer $j \in [n] \setminus \{i\}$ and a number $c_i \in [m]$ such that

$$b_j = \frac{1}{\left(c_i + \sum_{\ell \in [n] \setminus \{i, j\}} b_\ell / \epsilon\right)_{\text{mod } m} + 1},$$

then charge i the price $c_i \epsilon$;

- otherwise, the mechanism charges i an arbitrary price.

This mechanism is dominant strategy IC because the mechanism's decision to provide service to a buyer does not depend on his bid, and the price that a buyer is charged is not dependent on his own bid. This mechanism is ex-post IR because bidding truthfully always gives the buyer a utility of 0. This mechanism achieves a revenue equal to the optimal social welfare because (by definition of the pricing rule) the price that a buyer is charged is equal to the

valuation of that buyer, if all buyers bid truthfully. Also, note that the situation corresponding to the third point in the above specification of the mechanism will not occur when the buyers bid truthfully, and is only included for the sake of completely specifying the mechanism.

We argue that

$$\text{OR}(I_n) = \mathbf{E}_{\vec{v} \sim \pi} \left[\sum_{i \in [n]} v_i \right] = \sum_{i \in [n]} \mathbf{E}_{\vec{v} \sim \pi} [v_i] = \frac{(n-1)m\epsilon}{2} + \frac{H_m}{m},$$

where the last equality follows because the expected valuation of each of the buyers is $((n-1)/n)(m\epsilon/2) + (1/n)(H_m/m)$. This in turn holds because a buyer is elected as Buyer i^* with probability $1/n$, and Buyer i^* 's marginal distribution is the distribution π' induced by drawing a value from the set $\{1/a : a \in [m]\}$ uniformly at random. The latter distribution has already been encountered in the beginning of Section 3.1, where we concluded that its expected value is H_m/m .

We now proceed to prove an upper bound on $\text{RM}(I_n)$. Let M be an arbitrary enhanced posted price mechanism. Because M is randomized, running M on I_n can be viewed as a probability distribution on a sample space of deterministic enhanced posted price mechanisms that are run on I_n . We analyze the revenue of the mechanism conditioned on three disjoint events that form a partition of this sample space. Consider first the event E_1 that Buyer i^* gets asked for his valuation (when running M on I_n). Conditioned on this event, the mechanism does not attain a revenue of more than $(n-1)m\epsilon$ because the revenue of each buyer in $[n] \setminus \{i^*\}$ is at most $m\epsilon$.

Consider next the event E_2 where Buyer i^* does not get asked for his valuation and Buyer i^* is not the last buyer that is selected. Then a price p_{i^*} is proposed to i^* . Without loss of generality, M draws p_{i^*} from a probability distribution P_{i^*} with finite support, and the choice of distribution P_{i^*} depends on the sequence S of buyers queried prior to i^* together with the responses of the buyers in S . These responses take the form of a reported valuation in case a buyer in S is asked to report his valuation, and the form of an accept/reject decision otherwise. Because i^* is not the last buyer selected, $[n] \setminus (S \cup \{i^*\})$ is non-empty, and there exists a buyer $j \in [n] \setminus (S \cup \{i^*\})$ such that the choice of P_{i^*} does not depend on c_j . By the fact that c_j is drawn independently and uniformly at random from $[m]$ for all $j \in [n] \setminus (S \cup \{i^*\})$ and by (2), the marginal probability distribution of the valuation of Buyer i^* conditioned on E_2 , is π' (which we defined above). Therefore

$$\mathbf{E}_{p_{i^*} \sim P_{i^*}, \vec{v} \sim \pi} [p_{i^*} \mathbf{Pr}[v_{i^*} \leq p_{i^*}]] = \mathbf{E}_{p_{i^*} \sim P_{i^*}, \vec{v} \sim \pi'} [p_{i^*} \mathbf{Pr}[v_{i^*} \leq p_{i^*}]] = \frac{1}{m},$$

where the last equality follows from (1). Thus, the expected revenue of M conditioned on E_2 is at most $1/m + (n-1)m\epsilon$.

In the event E_3 , the mechanism selects i^* last. The expected revenue of M conditioned on this event is at most the expected maximum social welfare: $(n-1)m\epsilon/2 + H_m/m$. The probability of event E_3 occurring is $1/n$, because of the following. For $\ell \in [n]$, let E_3^ℓ be the event that i^* is not the ℓ -th buyer selected by M , and let $E_3^{<\ell}$ be the event that i^* is not among the first $\ell-1$ buyers selected by M . Note that this means that $\mathbf{Pr}[E_3^{<1}] = 1$. Then,

$$\mathbf{Pr}[E_3] = \mathbf{Pr}[E_3^{<n}] = \mathbf{Pr}[E_3^{n-1} \mid E_3^{<n-1}] \mathbf{Pr}[E_3^{<n-1}] = \prod_{\ell \in [n-1]} \mathbf{Pr}[E_3^\ell \mid E_3^{<\ell}].$$

For every $\ell \in [n-1]$, and every set S of $\ell-1$ buyers, it holds that if $i^* \notin S$ and M selects S as the first $\ell-1$ buyers, the probability of selecting Buyer i^* as the ℓ -th buyer is $1/(n-(\ell-1))$, by the definition of π (particularly because Buyer i^* is a buyer picked uniformly at random). Therefore,

$$\mathbf{Pr}[E_3] = \mathbf{Pr}[E_3^{<n}] = \prod_{\ell \in [n-1]} \left(1 - \frac{1}{n-(\ell-1)} \right) = \prod_{\ell \in [n-1]} \frac{n-\ell}{n-\ell+1} = \frac{1}{n}.$$

Thus, we obtain the following upper bound on $\text{RM}(I_n)$:

$$\begin{aligned} \text{RM}(I_n) &\leq \mathbf{Pr}[E_1](n-1)m\epsilon + \mathbf{Pr}[E_2] \left(\frac{1}{m} + (n-1)m\epsilon \right) + \frac{1}{n} \left(\frac{(n-1)m\epsilon}{2} + \frac{H_m}{m} \right) \\ &\leq \frac{1}{m} + 2(n-1)m\epsilon + \frac{m\epsilon}{2} + \frac{H_m}{mn} \leq 3(n-1)m\epsilon + \frac{H_m}{mn} + \frac{1}{m}. \end{aligned}$$

This leads us to conclude that

$$\frac{\text{RM}(I_n)}{\text{OR}(I_n)} \leq \frac{3nm\epsilon + H_m/mn + 1/m}{H_m/m} = \frac{3nm^2\epsilon + H_m/n + 1}{H_m} \leq \frac{H_m/n + 4}{H_m} = \frac{1}{n} + \frac{4}{H_m} \in O\left(\frac{1}{n}\right). \quad \square$$

4.2 Revenue guarantees for blind offer mechanisms

We prove in this section that blind offer mechanisms can always extract a constant fraction of the optimal revenue, without making any assumptions on the valuation distribution. Specifically we prove the following theorem.

Theorem 4. *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. Moreover, if $k = n$, then there is a blind offer mechanism for which the expected revenue equals the full 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

$$\max \sum_{i \in [n]} \sum_{\vec{v}_{-i} \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_{-i}) \sum_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})} \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i] v_i y_i(v_i, \vec{v}_{-i}) \quad (3)$$

subject to the constraints

$$\forall i \in [n], \vec{v}_{-i} \in \text{supp}(\pi_{-i}): \sum_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})} y_i(v_i, \vec{v}_{-i}) \leq 1, \quad (4)$$

$$\forall \vec{v} \in \text{supp}(\pi): \sum_{i \in [n]} \sum_{v'_i \in \text{supp}(\pi_{i, \vec{v}_{-i}}): v'_i \leq v_i} y_i(v'_i, \vec{v}_{-i}) \leq k, \quad (5)$$

$$\forall i \in [n], \vec{v} \in \text{supp}(\pi): y_i(\vec{v}) \geq 0. \quad (6)$$

The next lemma states that the solution to this linear program forms an upper bound on the revenue of the optimal mechanism, and that the solution to the above linear program is integral in case $k = n$.

Lemma 1. *For any instance (n, π, k) , the linear program (3-6) upper bounds the maximum expected revenue achievable by an ex-post IC, ex-post IR mechanism. Moreover, when $k = n$ the optimal solution to (3-6) is to set $y_i(v_i, \vec{v}_{-i})$ to 1 for the value v_i that maximizes $v_i \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i]$ (for all $i \in [n], \vec{v} \in \text{supp}(\pi)$).*

Proof. Integrality for $k = n$ is the easiest to prove among these two claims, so we do that first. Note that in case $k = n$, we can safely remove the constraints (5) from the linear program, because when $k = n$ these constraints are implied by (4) and (6). The linear program that remains tells us how to optimize a sum of convex combinations of values. That is, it effectively tells us to pick for each $i \in [n]$ and $\vec{v}_{-i} \in \text{supp}(\pi_{-i})$ a convex combination of the values $\{v_i \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i]\}_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})}$. The optimal solution is therefore to put weight 1 on the maximum values in these sets, i.e., to set $y_i(v_i, \vec{v}_{-i})$ to 1 for the value v_i that maximizes $v_i \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i]$.

It remains to prove the first claim. To this aim, let us first introduce some specialized notation: Let σ now be a probability distribution on a finite subset of $\mathbb{R}_{\geq 0}$ and $x \in \text{supp}(\sigma)$, we write $\text{prec}_\sigma(x)$ to denote $\max \text{supp}(\sigma) \cap [0, x)$ if $\text{supp}(\sigma) \cap [0, x)$ is non-empty. Otherwise, if $\text{supp}(\sigma) \cap [0, x) = \emptyset$, we define $\text{prec}_\sigma(x) = 0$. Similarly, we write $\text{succ}_\sigma(x)$ to denote $\min \text{supp}(\sigma) \cap (x, \infty]$. (We leave $\text{succ}_\sigma(x)$ undefined if $\text{supp}(\sigma) \cap (x, \infty]$ is empty.)

Suppose now that A is an optimal dominant strategy IC, ex-post IR mechanism. For $\vec{v} \in \text{supp}(\pi)$, denote by $x(\vec{v})$ the expected allocation vector output by A when the buyers report valuation vector \vec{v} (so that for $i \in [n]$, the value $x_i(\vec{v})$ is the probability that i gets allocated service, when the buyers report \vec{v}) and denote by $p(\vec{v})$ the vector of expected prices charged by A when the buyers report \vec{v} . Ex-post incentive compatibility states that

$$\forall i \in [n], \vec{v}_{-i} \in \text{supp}(\pi_{-i}), (v_i, v'_i) \in \text{supp}(\pi_i)^2: v_i x_i(v_i, \vec{v}_{-i}) - p_i(v_i, \vec{v}_{-i}) \geq v_i x_i(v'_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}),$$

and ex-post individual rationality states that

$$\forall i \in [n], \vec{v} \in \text{supp}(\pi): v_i x_i(v_i, \vec{v}_{-i}) - p_i(v_i, \vec{v}_{-i}) \geq 0.$$

The next lemma states that, in A , the allocation probability for a buyer is non-decreasing in his reported valuation.

Lemma 2. *For all $i \in [n]$, all $\vec{v}_{-i} \in \pi_{-i}$, and all $v_i, v'_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})$, with $v_i < v'_i$, it holds that*

$$x_i(v_i, \vec{v}_{-i}) \leq x_i(v'_i, \vec{v}_{-i}).$$

Proof. By way of contradiction, we assume that $\epsilon = x_i(v_i, \vec{v}_{-i}) - x_i(v'_i, \vec{v}_{-i}) > 0$. By ex-post incentive compatibility it holds that

$$\begin{aligned} v_i x_i(v_i, \vec{v}_{-i}) - p_i(v_i, \vec{v}_{-i}) &\geq v_i x_i(v'_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}), \\ v'_i x_i(v'_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}) &\geq v'_i x_i(v_i, \vec{v}_{-i}) - p_i(v_i, \vec{v}_{-i}). \end{aligned}$$

We now rewrite these inequalities as

$$\begin{aligned} v_i x_i(v_i, \vec{v}_{-i}) - v_i x_i(v'_i, \vec{v}_{-i}) &\geq p_i(v_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}), \\ v'_i x_i(v_i, \vec{v}_{-i}) - v'_i x_i(v'_i, \vec{v}_{-i}) &\leq p_i(v_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}). \end{aligned}$$

This results in the following pair of inequalities.

$$\begin{aligned} v_i \epsilon &\geq p_i(v_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}), \\ v'_i \epsilon &\leq p_i(v_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}). \end{aligned}$$

The two inequalities contradict each other, because $v'_i > v_i$ and we assumed $\epsilon > 0$. \square

The next lemma upper bounds the prices charged by A .

Lemma 3. For all $i \in [n]$, all $\vec{v}_{-i} \in \text{supp}(\vec{v}_{-i})$ and all $v_i \in \text{supp}(\pi_{i,\vec{v}_{-i}})$, it holds that

$$p_i(v_i, \vec{v}_{-i}) \leq v_i x_i(v_i, \vec{v}_{-i}) - \sum_{v'_i \in \text{supp}(\pi_{i,\vec{v}_{-i}}): v'_i < v_i} (\text{succ}_{\pi_{i,\vec{v}_{-i}}}(v_i) - v_i) x_i(v'_i, \vec{v}_{-i}). \quad (7)$$

Proof. For $v'_i \in \text{supp}(\pi_{i,\vec{v}_{-i}})$ the ex-post IC constraint for \vec{v}_{-i}, v'_i , $\text{prec}_{\pi_{i,\vec{v}_{-i}}}(v_i)$ can be written as

$$v'_i (x_i(v'_i, \vec{v}_{-i}) - x_i(v''_i, \vec{v}_{-i})) \geq p_i(v'_i, \vec{v}_{-i}) - p_i(v''_i, \vec{v}_{-i}),$$

where $v''_i = \text{prec}_{\pi_{i,\vec{v}_{-i}}}(v_i)$. Summing the above over all $v'_i \in \text{supp}(\pi_{i,\vec{v}_{-i}}), v'_i < v_i$ yields (7). \square

The optimal revenue among all ex-post IC, ex-post IR mechanisms (and thus the expected revenue of A) can be written as the following linear program, where $(x(\vec{v}))_{\vec{v} \in \text{supp}(\pi)}$ and $(p(\vec{v}))_{\vec{v} \in \text{supp}(\pi)}$ are the variables:

$$\max \left\{ \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} p_i(\vec{v}) \quad \middle| \right. \quad (8)$$

$$\forall i \in [n], \vec{v} \in \text{supp}(\pi): v_i x_i(\vec{v}) - p_i(\vec{v}) \geq 0 \quad (9)$$

$$\forall i \in [n], (v_i, v'_i) \in \text{supp}(\pi_i)^2, \vec{v}_{-i} \in \text{supp}(\pi_{-i,v_i}): v_i x_i(v_i, \vec{v}_{-i}) - p_i(v_i, \vec{v}_{-i}) \geq v_i x_i(v'_i, \vec{v}_{-i}) - p_i(v'_i, \vec{v}_{-i}) \quad (10)$$

$$\forall \vec{v} \in \text{supp}(\pi): \sum_i x_i(\vec{v}) \leq k \quad (11)$$

$$\forall i \in [n], \vec{v} \in \text{supp}(\pi): 0 \leq x_i(\vec{v}) \leq 1 \quad (12)$$

In the above linear program, (9) are the ex-post IR constraints, (10) are the ex-post IC constraints, and (11) expresses that the service cannot be provided to more than k buyers.

By Lemma 2, it is possible to add to the above linear program the constraints $x_i(v_i, \vec{v}_{-i}) \geq x_i(\text{prec}_{\pi_{i,\vec{v}_{-i}}}(v_i))$ for $i \in [n], \vec{v}_{-i} \in \text{supp}(\pi_{-i}), v_i \in \text{supp}(\pi_{i,\vec{v}_{-i}})$. Moreover, by Lemma 3, we can replace the objective function by

$$\sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \left(v_i x_i(v_i, \vec{v}_{-i}) - \sum_{v'_i \in \text{supp}(\pi_{i,\vec{v}_{-i}}): v'_i < v_i} (\text{succ}_{\pi_{i,\vec{v}_{-i}}}(v_i) - v_i) x_i(v'_i, \vec{v}_{-i}) \right)$$

and remove the constraints (9) and (10). This results in the following linear program that upper bounds the optimal revenue among the ex-post IC, ex-post IR mechanisms:

$$\max \left\{ \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \left(v_i x_i(v_i, \vec{v}_{-i}) - \sum_{v'_i \in \text{supp}(\pi_{i,\vec{v}_{-i}}): v'_i < v_i} (\text{succ}_{\pi_{i,\vec{v}_{-i}}}(v_i) - v_i) x_i(v'_i, \vec{v}_{-i}) \right) \quad \middle| \right. \quad (13)$$

$$\forall i \in [n], \vec{v} \in \text{supp}(\pi): x_i(\vec{v}) \geq x_i(\text{prec}_{\pi_i, \vec{v}_{-i}}(v_i), \vec{v}_{-i}) \quad (14)$$

$$\forall \vec{v} \in \text{supp}(\pi): \sum_i x_i(\vec{v}) \leq k \quad (15)$$

$$\forall i \in [n], \vec{v} \in \text{supp}(\pi): 0 \leq x_i(\vec{v}) \leq 1 \quad (16)$$

We will show that the linear program (3–6) is equivalent to the above. Set $y_i(\vec{v}) = x_i(\vec{v}) - x_i(\text{prec}_{\pi_i, \vec{v}_{-i}}(v_i), \vec{v}_{-i})$ for all $i \in [n], \vec{v} \in \text{supp}(\pi)$, and observe that the constraints (14), (15) and (16) are then equivalent to (4), (5) and (6) respectively. Moreover, with this correspondence between y and x , we can rewrite the objective function as follows:

$$\begin{aligned} & \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \left(v_i x_i(v_i, \vec{v}_{-i}) - \sum_{v'_i \in \text{supp}(\pi_i, \vec{v}_{-i}): v'_i < v_i} (\text{succ}_{\pi_i, \vec{v}_{-i}}(v_i) - v_i) x_i(v'_i, \vec{v}_{-i}) \right) \\ &= \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})} \pi_{i, \vec{v}_{-i}}(v_i) v_i x_i(v_i, \vec{v}_{-i}) \\ &\quad - \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})} \pi_{i, \vec{v}_{-i}}(v_i) \sum_{v'_i \in \text{supp}(\pi_i, \vec{v}_{-i}): v'_i < v_i} (\text{succ}_{\pi_i, \vec{v}_{-i}}(v_i) - v_i) x_i(v'_i, \vec{v}_{-i}) \\ &= \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})} \pi_{i, \vec{v}_{-i}}(v_i) v_i x_i(v_i, \vec{v}_{-i}) \\ &\quad - \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})} x_i(v_i, \vec{v}_{-i}) (\text{succ}_{\pi_i, \vec{v}_{-i}}(v_i) - v_i) \Pr_{v'_i \sim \pi_i, \vec{v}_{-i}}[v'_i > v_i] \\ &= \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})} x_i(v_i, \vec{v}_{-i}) (\pi_{i, \vec{v}_{-i}}(v_i) v_i - (\text{succ}_{\pi_i, \vec{v}_{-i}}(v_i) - v_i) \Pr_{v'_i \sim \pi_i, \vec{v}_{-i}}[v'_i > v_i]) \\ &= \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})} x_i(v_i, \vec{v}_{-i}) (\Pr_{v'_i \sim \pi_i, \vec{v}_{-i}}[v'_i \geq v_i] v_i - (\text{succ}_{\pi_i, \vec{v}_{-i}}(v_i)) \Pr_{v'_i \sim \pi_i, \vec{v}_{-i}}[v'_i > v_i]) \\ &= \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})} (x_i(v_i, \vec{v}_{-i}) - x_i(\text{prec}_{\pi_i, \vec{v}_{-i}}(v_i), \vec{v}_{-i})) v_i \Pr_{v'_i \sim \pi_i, \vec{v}_{-i}}[v'_i \geq v_i] \\ &= \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}). \end{aligned}$$

This completes the proof. \square

We can now proceed to prove our main result about blind offer mechanisms. We first handle the case of unlimited supply. We use the following blind offer mechanism.

Definition 10. Consider an instance (n, π, n) . For $i \in [n]$ and $\vec{v}_{-i} \in \text{supp}(\pi_{-i})$, fix $\hat{p}_{i, \vec{v}_{-i}}$ to be any value in the set $\arg_p \max \{ p \Pr_{v'_i \sim \pi_i, \vec{v}_{-i}}[p \leq v_i] : p \in \mathbb{R} \}$. Define M_π^n to be the following blind offer mechanism for allocating service to n buyers when the valuations of these buyers are drawn from π .

Let \vec{b} be the submitted vector of bids. For $i \in [n]$, if $\vec{b}_{-i} \in \text{supp}(\pi_{-i})$ and $b_i \geq \hat{p}_{i, \vec{b}_{-i}}$, then M_π^n gives service to i and charges i the price $\hat{p}_{i, \vec{b}_{-i}}$. If $b_i < \hat{p}_{i, \vec{b}_{-i}}$, then the price charged to i is 0, and i is not given service. Otherwise, if $\vec{b}_{-i} \notin \text{supp}(\pi_{-i})$, the price charged to i is 0 and i is not allocated service.

Lemma 4. For instance (n, π, n) , Mechanism M_π^n extracts the maximum revenue among the class of ex-post IC, ex-post IR mechanisms.

Proof. Denote by $p_i(\vec{v})$ the price charged to Buyer $i \in [n]$ when the buyers have valuation vector $\vec{v} \in \text{supp}(\pi)$. We can write the expected revenue of M_π^n as follows:

$$\mathbf{E}_{\vec{v} \sim \pi} \left[\sum_{i \in [n]} p_i(\vec{v}) \right] = \sum_{i \in [n]} \mathbf{E}_{\vec{v} \sim \pi} [p_i(\vec{v})] = \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) p_i(\vec{v})$$

$$= \sum_{i \in [n]} \sum_{\vec{v}_{-i} \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_{-i}) \Pr_{v_i \sim \pi_{i, \vec{v}_{-i}}} [v_i \geq \hat{p}_{i, \vec{v}_{-i}}] \hat{p}_{i, \vec{v}_{-i}}.$$

Lastly, by Lemma 1 and the objective function (3) of the linear program, we conclude that the latter expression is equal to the solution of the linear program, which is an upper bound on the optimal revenue among all ex-post IC, ex-post IR mechanisms by Theorem 1. \square

For the case of k -limited supply where $k < n$, things are somewhat more complicated. Indeed, there does not seem to exist a blind offer mechanism as simple and elegant as M_π^n . However, we are still able to define a blind offer mechanism that extracts at least a constant fraction of the optimal revenue.

Definition 11. Let (n, π, k) be an arbitrary instance. Let $(y_i^*(\vec{v}))_{i \in [n]}$ be the optimal solution to the linear program (3–6) 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 once k buyers have accepted an offer, or when iteration $n + 1$ is reached.

Lemma 5. On instance (n, π, k) , the expected revenue of blind offer mechanism M_π^k is at least a $(2 - \sqrt{e})/4 \approx 0.088$ fraction of the expected revenue of the optimal ex-post IC, ex-post IR mechanism. Moreover, there exists a dominant strategy IC blind offer mechanism of which the expected revenue is at least a $(2 - \sqrt{e})/4 \approx 0.088$ fraction of the expected revenue of the optimal ex-post IC, ex-post IR mechanism.

Proof. We will show that the expected revenue of M_π^k is at least

$$\frac{2 - \sqrt{e}}{4} \sum_{i \in [n]} \sum_{\vec{v}_{-i} \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_{-i}) \sum_{v_i \in \text{supp}(\pi_{i, \vec{v}_{-i}})} \Pr_{v'_i \sim \pi_{i, \vec{v}_{-i}}} [v'_i \geq v_i] v_i y_i^*(v_i, \vec{v}_{-i}),$$

which is by Theorem 1 and the objective function (3) 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 (as described in Definition 11). 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 the “do nothing”-option during iteration i , so that $D_{i, \vec{v}_{-i}}$ is a probability distribution on the set

$\{V\} \cup \{v_i : y_i^*(v_i, \vec{v}_{-i}) > 0\}$. Let us formulate an initial lower bound on the expected revenue of M_π^k .

$$\begin{aligned}
& \mathbf{E}_{\vec{v} \sim \pi} [\text{revenue of } M_\pi^k(\vec{v})] \\
&= \mathbf{E}_{\substack{\vec{v} \sim \pi, \\ p_1 \sim D_{1, \vec{v}_{-1}}, \\ \vdots \\ p_n \sim D_{n, \vec{v}_{-n}}}} \left[\sum_{i \in [n]} p_i \mathbf{1}[p_i \leq v_i] \mathbf{1}[|\{j \in [i-1] : p_j \leq v_j\}| < k] \right] \\
&= \sum_{i \in [n]} \mathbf{E}_{\substack{\vec{v} \sim \pi, \\ p_1 \sim D_{1, \vec{v}_{-1}}, \\ \vdots \\ p_n \sim D_{n, \vec{v}_{-n}}}} \left[p_i \mathbf{1}[p_i \leq v_i] \mathbf{1}[|\{j \in [i-1] : p_j \leq v_j\}| < k] \right] \\
&= \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{\substack{p_i \in \text{supp}(D_{i, \vec{v}_{-i}}) \\ \vdots \\ p_i \leq v_i}} p_i \mathbf{1}[p_i \leq v_i] \mathbf{1}[|\{j \in [i-1] : p_j \leq v_j\}| < k] \prod_{j \in [i]} D_{j, \vec{v}_{-j}}(p_j) \\
&= \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{\substack{p_i \in \text{supp}(D_{i, \vec{v}_{-i}}) \\ \vdots \\ p_i \leq v_i}} p_i D_{i, \vec{v}_{-i}}(p_i) \sum_{\substack{p_{i-1} \in \text{supp}(D_{i-1, \vec{v}_{-(i-1)}}) \\ \vdots \\ |\{j \in [i-1] : p_j \leq v_j\}| < k}} \prod_{j \in [i-1]} D_{j, \vec{v}_{-j}}(p_j) \\
&= \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{\substack{p_i \in \text{supp}(D_{i, \vec{v}_{-i}}) \\ \vdots \\ p_i \leq v_i}} \frac{p_i y_i^*(p_i, \vec{v}_{-i})}{2} \mathbf{Pr}_{\substack{p_1 \sim D_{1, \vec{v}_{-1}} \\ \vdots \\ p_{i-1} \sim D_{i-1, \vec{v}_{-(i-1)}}}} [|\{j \in [i-1] : p_j \leq v_j\}| < k] \\
&\geq \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{\substack{p_i \in \text{supp}(D_{i, \vec{v}_{-i}}) \\ \vdots \\ p_i \leq v_i}} \frac{p_i y_i^*(p_i, \vec{v}_{-i})}{2} \mathbf{Pr}_{\substack{p_1 \sim D_{1, \vec{v}_{-1}} \\ \vdots \\ p_{n-1} \sim D_{n-1, \vec{v}_{-(n-1)}}}} [|\{j \in [n-1] : p_j \leq v_j\}| < k].
\end{aligned} \tag{17}$$

For the second equality, we applied linearity of expectation; the third equality follows from the definition of expected value; to obtain the fourth equality we eliminate the indicator functions by removing the appropriate terms from the summation; in the fifth equality we substitute $D_{i, \vec{v}_{-i}}(p_i, \vec{v}_{-i})$ and $D_{j, \vec{v}_{-j}}(p_j, \vec{v}_{-j})$ by concrete probabilities. For the last inequality we lower bounded the last probability in the expression by replacing i by n .

For $\vec{v} \in \text{supp}(\pi)$ and $i \in [n-1]$, let us denote by $z_i^{\vec{v}}$ the probability that a price drawn from $D_{i, \vec{v}_{-i}}$ does not exceed v_i , i.e.,

$$z_i^{\vec{v}} = \sum_{p_i \in \text{supp}(D_{i, \vec{v}_{-i}}) : p_i \leq v_i} D_{i, \vec{v}_{-i}}(p_i) = \sum_{v'_i \in \text{supp}(\pi_{i, \vec{v}_{-i}}) : v'_i \leq v_i} \frac{y_i^*(v'_i, \vec{v}_{-i})}{2},$$

and let $X_i^{\vec{v}}$ denote the random variable that takes the value 1 with probability $z_i^{\vec{v}}$ and the value 0 with probability $1 - z_i^{\vec{v}}$. Then the final probability in the derivation above, i.e.,

$$\begin{aligned}
& \mathbf{Pr}_{\substack{p_1 \sim D_{1, \vec{v}_{-1}} \\ \vdots \\ p_{n-1} \sim D_{n-1, \vec{v}_{-(n-1)}}}} [|\{j \in [n-1] : p_j \leq v_j\}| < k]
\end{aligned}$$

can be written as

$$1 - \mathbf{Pr} \left[\sum_{i \in [n-1]} X_i^{\vec{v}} \geq k \right].$$

Let $\mu = \mathbf{E} \left[\sum_{i \in [n-1]} X_i^{\vec{v}} \right]$. Next, we use a Chernoff bound:

Theorem 5 (Chernoff bound (as in [Motwani and Raghavan, 1995])). *Let X_1, \dots, X_n be independent random (0, 1)-variables such that, for $i \in [n]$, $\mathbf{Pr}[X_i = 1] = p_i$ where $p_i \in [0, 1]$. Then, for $X = \sum_{i \in [n]} X_i$, $\mu = \mathbf{E}[X] = \sum_{i \in [n]} p_i$ and*

any $\delta > 0$,

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu.$$

This implies that the expression above is bounded as follows.

$$1 - \Pr \left[\sum_{i \in [n-1]} X_i^{\vec{v}} \geq \left(1 + \left(\frac{k}{\mu} - 1 \right) \right) \mu \right] \geq 1 - \left(\frac{e^{k/\mu-1}}{(k/\mu)^{k/\mu}} \right)^\mu.$$

By the definition of $z_i^{\vec{v}}$ and constraint (5) of the linear program, it holds that $\mu = \sum_{i \in [n-1]} z_i^{\vec{v}} \leq k/2$. We can lower bound the expression above by replacing μ by $k/2$. To see this, we first rewrite it as follows:

$$1 - \left(\frac{e^{k/\mu-1}}{(k/\mu)^{k/\mu}} \right)^\mu = 1 - \frac{e^{k-\mu+k \ln(\mu)}}{k^k}.$$

The derivative of the exponent of e (with respect to μ) is positive for $\mu \in [0, k]$, which means that the exponent of e is increasing in μ on $[0, k]$. Thus, replacing μ by its upper bound $k/2$ increases the exponent and therefore decreases the expression above. Therefore:

$$1 - \Pr \left[\sum_{i \in [n-1]} X_i^{\vec{v}} \geq k \right] \geq 1 - \left(\frac{e}{4} \right)^{k/2} \geq 1 - \left(\frac{e}{4} \right)^{1/2} = \frac{2 - \sqrt{e}}{2}. \quad (18)$$

Continuing from (17), we obtain

$$\begin{aligned} \mathbf{E}_{\vec{v} \sim \pi}[\text{revenue of } M_\pi^k(\vec{v})] &\geq \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{p_i \in \text{supp}(D_{i, \vec{v}_i}): p_i \leq v_i} p_i \frac{y_i^*(p_i, \vec{v}_i)}{2} \frac{2 - \sqrt{e}}{2} \\ &= \frac{2 - \sqrt{e}}{4} \sum_{i \in [n]} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{p_i \in \text{supp}(D_{i, \vec{v}_i}): p_i \leq v_i} p_i y_i^*(p_i, \vec{v}_i) \\ &= \frac{2 - \sqrt{e}}{4} \sum_{i \in [n]} \sum_{\vec{v}_i \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_i) \sum_{v_i, p_i \in \text{supp}(\pi_{i, \vec{v}_i}): p_i \leq v_i} \pi_{i, \vec{v}_i}(v_i) p_i y_i^*(p_i, \vec{v}_i) \\ &= \frac{2 - \sqrt{e}}{4} \sum_{i \in [n]} \sum_{\vec{v}_i \in \text{supp}(\pi_{-i})} \pi_{-i}(\vec{v}_i) \sum_{p_i \in \text{supp}(\pi_{i, \vec{v}_i}): p_i \leq v_i} \Pr_{\vec{v}_i \sim \pi_{i, \vec{v}_i}}[v_i \geq p_i] p_i y_i^*(p_i, \vec{v}_i), \end{aligned}$$

which proves the first of the two claims. For the second claim, observe that (18) states a lower bound of $(2 - \sqrt{e})/2$ on the probability that all buyers get selected by the mechanism. Therefore, we can combine (18) with the principle explained in Remark 1, which allows us to transform M_π^k into a dominant strategy IC blind offer mechanism \hat{M}_π^k . Observe now that (17) is still a lower bound on the revenue of \hat{M}_π^k so that the revenue analysis of M_π^k also holds for \hat{M}_π^k . \square

Theorem 4 now follows by combining Lemmas 4 and 5. We note that the approximation factor of the theorem is certainly not tight and can be improved with additional work. For example, it is possible to show that for $k = 1$ the revenue of M_π^k is in fact at least $1/4$ of the optimal revenue. Moreover, recall that mechanism M_π^k works by scaling the probabilities $y_i(\vec{v})$ down by $1/2$. By making this scaling factor dependent on k and choosing it appropriately, we can improve the approximation factor further. We emphasize that the focus and purpose of the above result is merely to show that a constant factor of the optimal revenue (independent of the supply k) is achievable.

4.3 Revenue guarantees for enhanced sequential posted price mechanisms

Finally, in this section we evaluate the revenue guarantees of the enhanced sequential posted price 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, we have the following definition.

Definition 12. 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 where there is a "true" valuation v for the item, an expert buyer that keeps this valuation, and remaining buyers who roughly follow the opinion of the expert buyer, so that their valuation is that of the expert buyer, 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 exact conditional distribution of any other buyer. Our definition of 1-dimensional dependence is satisfied if such an expert buyer exists, even if the auctioneer does not know which buyer is the expert. Moreover, our definition of d -dimensional dependence is much more flexible than that: In the example above, even if each bidder picks another bidder as his own "personal" expert and adds noise to the valuation of his expert, the distribution would remain 1-dimensionally dependent.

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 notion can be used to rule out the extreme kind of dependence defined in the proof of Theorem 3: 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 -dimensional 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, and for that model a more straightforward procedure than the one in the proof below exists for obtaining the same revenue guarantee (see Appendix F). However, the notion of d -dimensional dependence is both more natural (for our setting) and much more general. In fact, we show in Appendix E that there exist distributions on \mathbb{R}^n that are 1-dimensionally dependent, but are not a Markov random field of degree less than $n/2$. In a sense, our definition of d -dimensional dependence resembles the limited dependence condition under which the *Lovasz Local Lemma* holds.

The main result that we will prove in this section is the following.

Theorem 6. For every instance (n, π, k) where π is d dimensionally dependent, there exists an enhanced sequential posted price 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. Moreover, if $k = n$, then there exists an enhanced posted price mechanism of which the expected revenue is at least a $1/(4d)$ fraction of the maximum expected revenue that can be extracted by an ex-post IC, ex-post IR mechanism.

A corollary of this theorem is that the bound of Theorem 3 is asymptotically tight.

Corollary 1. For every instance (n, π, k) there exists an enhanced sequential posted price mechanism of which the expected revenue is at least a $\Omega(1/n)$ fraction of the maximum expected revenue that can be extracted by an ex-post IC, ex-post IR mechanism.

For proving our main result on enhanced sequential posted price mechanisms, we make use of our insights about blind offer mechanisms. The next lemma shows how we can convert blind offer mechanisms into enhanced sequential posted price mechanisms while losing only a factor of $1/4d$ of the revenue, if the valuation distribution is d -dimensionally dependent.

Lemma 6. Let $\alpha \in [0, 1]$ and let (n, π, k) be an instance where π is d -dimensionally dependent. If there exists a blind offer mechanism for (n, π, k) 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 exists an enhanced sequential posted price 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.

Proof. Let M be 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. Let $p_i^M(\vec{v})$ be the expected price paid to Mechanism M by Buyer $i \in [n]$ when \vec{v} are the reported valuations. Let $q \in [0, 1]$ and consider the following enhanced sequential posted

price mechanism M_q : Mechanism M_q first partitions $[n]$ into two sets A and $B = [n] \setminus A$. It does so by placing each buyer independently with probability q in set A , and placing him in set B otherwise. Then, the mechanism retrieves the vector \vec{v}_A by asking the buyers in A for their valuations. The existence of M implies the existence of a blind offer mechanism M_B that only makes offers to buyers in B such that the expected price $p_i^{M_B}(\vec{v})$ paid to M_B by a buyer in B is at least $p_i^M(\vec{v})$ (this can be achieved by doing the same as M , but refraining from offering to buyers in A). Mechanism M_q offers each buyer $i \in B$ a price that is determined by simulating M_B as follows: make the same decisions as M_B would, except for that an offer of M_B is skipped if $\pi_{i,\vec{v}_{-i}} \neq \pi_{i,\vec{v}_A}$.

Let P be the distribution (induced by Mechanism M_q) on the set of partitions of $[n]$ into 2 sets. For $i \in [n]$, let $S_i \subseteq [n] \setminus \{i\}$ be the set of d buyers such that $\pi_{i,\vec{v}_{S_i}} = \pi_{i,\vec{v}_{-i}}$ for all $\vec{v} \in \mathbb{R}^n$. For $T \subseteq [n]$, let $p_i(T, \vec{v})$ be the expected price paid to M_q by a buyer $i \in T$, conditioned on the event that $B = T$ and $S_i \subseteq A$. Note that $p_i(T, \vec{v}) \geq p_i^{M_B}(\vec{v}) \geq p_i^M(\vec{v})$. Therefore, the expected revenue of M_q is

$$\begin{aligned} \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{i \in [n]} \Pr_{[A,B] \sim P}[i \in B \cap S_i \subseteq A] p_i(B, \vec{v}) &= \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{i \in [n]} (1-q)q^d p_i(B, \vec{v}) \\ &\geq (1-q)q^d \sum_{\vec{v} \in \text{supp}(\pi)} \pi(\vec{v}) \sum_{i \in [n]} p_i^M(\vec{v}). \end{aligned}$$

The last (double) summation is at least α times the expected revenue of the optimal dominant strategy IC, ex post IR mechanism, by definition of M . Therefore, this mechanism extracts at least a $(1-q)q^d \alpha$ fraction of the optimal revenue. For $d = 0$ it is optimal to set $q = 0$, which results in an enhanced sequential posted price mechanism of which the revenue is α -approximately optimal. For $d = 1$ it is optimal to set $q = 1/2$, which results in a $(\alpha/4)$ -approximately optimal enhanced sequential posted price mechanism. For $d \geq 2$ setting $q = 1 - 1/d$ will achieve the desired approximation ratio, since $\lim_{d \rightarrow \infty} (1 - 1/d)^d = 1/e$. Moreover, $(1 - 1/d)^d$ is increasing in d , and equals $1/4$ for $d = 2$. \square

We note that in the above proof it is easy to see that we can decrease the fraction q of buyers being probed for their valuation at the cost of worsening the approximation guarantee. Theorem 6 now follows directly by combining Theorem 4 and Lemma 6 above.

5 Open problems

Besides improving any of the approximation bounds that we established in the present paper, there are many other interesting further research directions. For example, it would be interesting to investigate the revenue guarantees under the additional constraint that the sequential posted price mechanism be *order-oblivious*: i.e., the mechanism has no control over which buyers to pick, and should perform well for any possible ordering of the buyers. We are also interested in resolving some questions regarding the use of randomization in our enhanced posted price mechanism that extracts $O(1/d)$ of the optimal revenue: In the current proof, is it necessary to pick buyers uniformly at random, or does there exist a *deterministic* enhanced sequential posted price mechanism that extracts a constant factor of the optimal revenue?

An obvious and interesting research direction is to investigate more general auction problems. In particular, to what extent can extended SPP 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 buyers may have, for example, a demand of more than one item, or there may be a matroid feasibility constraint on the set of buyers or on the set of items that may be allocated.

Acknowledgments

Part of the work has been done when Marek Adamczyk, Allan Borodin, Diodato Ferraioli, and Bart de Keijzer were all at Sapienza University of Rome. A shorter conference version of this paper appeared in the Proceedings of the 11th Conference on Web and Internet Economics (WINE), 2015 (Lecture Notes in Computer Science 9470, Springer, pages 1–15). We thank Joanna Drummond, Brendan Lucier, Tim Roughgarden and anonymous referees for their constructive comments which we have incorporated. 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.

References

- M. Adamczyk, M. Sviridenko, and J. Ward. Submodular stochastic probing on matroids. In *Proceedings of the 31st International Symposium on Theoretical Aspects of Computer Science (STACS)*. Schloss Dagstuhl, 2014.
- M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2007.
- M. Babaioff, S. Dughmi, R. Kleinberg, and A. Slivkins. Dynamic pricing with limited supply. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*. ACM, 2012.
- M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014.
- A. Blum and J. D. Hartline. Near-optimal online auctions. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2005.
- A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. *Theoretical Computer Science*, 324: 137–146, 2004.
- L. Blumrosen and T. Holenstein. Posted prices vs. negotiations: An asymptotic analysis. In *Proceedings of the 9th ACM Conference on Economics and Computation (EC)*. ACM, 2008.
- T. Börgers. *An Introduction to the Theory of Mechanism Design*. Oxford University Press, 2015.
- S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*. ACM, 2010.
- S. Chawla, H. Fu, and A. Karlin. Approximate revenue maximization in interdependent value settings. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 277–294. ACM, 2014.
- N. Chen, N. Immorlica, A. R. Karlin, M. Mahdian, and A. Rudra. Approximating matches made in heaven. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*. Springer, 2009.
- E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- J. Cremer and R. P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.
- N. R. Devanur, J. Morgenstern, V. Syrgkanis, and S. M. Weinberg. Simple auctions with simple strategies. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 305–322. ACM, 2015.
- S. Dobzinski, H. Fu, and R. D. Kleinberg. Optimal auctions with correlated bidders are easy. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*. ACM, 2011.
- M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2015a.
- M. Feldman, O. Svensson, and R. Zenklusen. Online contention resolution schemes. *CoRR/ArXiv preprint (arXiv:1508.00142)*, 2015b.
- T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.
- A. Gupta and V. Nagarajan. A stochastic probing problem with applications. In *Proceedings of the 16th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 2013.
- S. Hart and N. Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*, 2012.

- J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. *SIGecom Exchanges*, 8(1), 2009.
- J. Kagel, R. M. Harstad, and D. Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, pages 1275–1304, 1987.
- R. Kleinberg and T. Leighton. The value of knowing a demand curve: bounds on regret for online posted-price auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2003.
- R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th ACM Symposium on Theory of Computing (STOC)*. ACM, 2012.
- S. Li. Obviously strategy-proof mechanisms. Available at SSRN 2560028, 2015.
- R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- A. Ronen. On approximating optimal auctions. In *Proceedings of the 2nd ACM Conference on Economics and Computation (EC)*. ACM, 2001.
- T. Roughgarden and I. Talgam-Cohen. Optimal and near-optimal mechanism design with interdependent values. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*. ACM, 2013.
- A. Rubinstein and S. M. Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 377–394. ACM, 2015.
- T. Sandholm and A. Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *Agent-Mediated Electronic Commerce V*. 2004.
- I. Segal. Optimal pricing mechanisms with unknown demand. *The American Economic Review*, 93(3):509–529, 2003.
- W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- Q. Yan. Mechanism design via correlation gap. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2011.

APPENDIX

A Continuous distribution properties

Let π be a valuation distribution on $[0, a_i]^n$, with $a_i \in \mathbb{R}_{\geq 0}$ for $i \in [n]$, with density f that is continuous and nowhere zero. Distribution π is said to satisfy *affiliation* iff for every two valuation vectors $\vec{v}, \vec{w} \in \text{supp}(\pi)$ it holds that $f(\vec{v} \wedge \vec{w})f(\vec{v} \vee \vec{w}) \geq f(\vec{v})f(\vec{w})$, where $\vec{v} \wedge \vec{w}$ is the component-wise minimum and $\vec{v} \vee \vec{w}$ is the component-wise maximum. For $i \in [n]$ and $\vec{v}_{-i} \in \text{supp}(\pi_{-i})$ the *conditional marginal density function* $f_i(\cdot | \vec{v}_{-i})$ is defined as

$$f_i(v_i | \vec{v}_{-i}) = \frac{f(v_i, \vec{v}_{-i})}{\int_0^{a_i} f(t, \vec{v}_{-i}) dt},$$

the *conditional revenue curve* $B_i(\cdot | \vec{v}_{-i})$ is defined as

$$B_i(v_i | \vec{v}_{-i}) = v_i \int_{v_i}^{a_i} f_i(t | \vec{v}_{-i}) dt,$$

and the *conditional virtual value* $\phi_i(\cdot | \vec{v}_{-i})$ is defined as

$$\phi_i(v_i | \vec{v}_{-i}) = -\frac{\frac{d}{dv_i} B_i(v_i | \vec{v}_{-i})}{f_i(v_i | \vec{v}_{-i})}.$$

Denote by $F_i(\cdot | \vec{v}_{-i})$ the cumulative distribution function corresponding to $f_i(\cdot | \vec{v}_{-i})$. Distribution π satisfies *regularity* if $\phi_i(\cdot | \vec{v}_{-i})$ is non-decreasing for all $i \in [n]$ and $\vec{v}_{-i} \in \text{supp}(\pi_{-i})$ and it satisfies the *monotone hazard rate* condition if $(1 - F_i(v_i | \vec{v}_{-i})) / (f_i(v_i | \vec{v}_{-i}))$ is non-increasing in v_i for all $i \in [n]$ and $\vec{v}_{-i} \in \text{supp}(\pi_{-i})$.

A discussion and justification for the above notions is outside of the scope of this paper, and we refer the interested reader to [Roughgarden and Talgam-Cohen, 2013].

Note that Roughgarden and Talgam-Cohen [2013] proved that for any distribution π that satisfies regularity and affiliation the Myerson Mechanism is ex-post IC, ex-post IR and optimal among all ex-post IC and ex-post IR mechanisms.

B Proof of Proposition 2

Proof. Let M be the sequential posted price mechanism that draws for each $i \in [n]$ a value p_i uniformly at random from the set $S_i = \{v_{\pi,i}^{\min} 2^k : k \in [\lceil \log(r_{\pi,i}) - 1 \rceil] \cup \{0\}\}$. M proposes prices to the buyers in an arbitrary order, and offers price p_i to Buyer i .

For $i \in [n]$, let π_i be the probability distribution of the i th coordinate of π . Note that $|S_i|$ does not exceed $\log(r_{\pi,i})$. Therefore the probability that p_i is the highest possible value (among the values in S_i) that does not exceed the value drawn from π_i , is at least $1/\log(r_{\pi,i})$. More formally, let π_{S_i} be the uniform distribution on S_i ; then

$$\Pr_{v_i \sim \pi_i, p_i \sim \pi_{S_i}} [p_i \leq v_i \cap (\nexists p'_i \in S_i : p'_i > p_i \wedge p'_i \leq v_i)] \leq \frac{1}{\log(r_{\pi,i})}.$$

Thus, with probability $1/\log(r_{\pi,i})$, the mechanism extracts from Buyer i a revenue of exactly $v_{i,\min} 2^k$, where k is the number such that the value drawn from π_i lies in between $v_{i,\min} 2^i$ and $v_{i,\min} 2^{i+1}$. This implies that with probability $1/\log(r_{\pi,i})$ the mechanism extracts from Buyer i a revenue that lies a factor of at most $1/2$ away from v_i . This leads to the conclusion that

$$\begin{aligned} \mathbf{E}_{\vec{v} \sim \pi} [\text{revenue of } M(\vec{v})] &\geq \sum_{i \in [n]} \frac{1}{\log(r_{\pi,i})} \frac{1}{2} \mathbf{E}_{v_i \sim \pi_i} [v_i] \\ &\geq \frac{1}{2 \log(\max\{r_{\pi,i} : i \in [n]\})} \sum_{i \in [n]} \mathbf{E}_{v_i \sim \pi_i} [v_i] \\ &= \frac{1}{2 \log(\max\{r_{\pi,i} : i \in [n]\})} \mathbf{E}_{\vec{v} \sim \pi} [OPT(\vec{v})], \end{aligned}$$

where $OPT(\vec{v}) = \sum_{i \in [n]} v_i$ denotes the optimal social welfare when the buyers have valuation vector \vec{v} . □

C A sequential posted price mechanism for the k -limited setting

Definition 13. For a valuation distribution π on \mathbb{R}^n , let $v_\pi^{\max(k)}$ and $v_\pi^{\min(k)}$ be respectively the maximum k -th largest and minimum k -th smallest valuation among the valuation vectors in $\text{supp}(\pi)$. Let $r_\pi^{(k)} = v_\pi^{\max(k)} / v_\pi^{\min(k)}$ be the ratio between these values.

Theorem 7. Let $n \in \mathbb{N}_{\geq 1}$, and let π be a probability distribution on \mathbb{R}^n . For any $k \in [n]$, there exists a sequential posted price mechanism that, when run on instance (n, π, k) , extracts in expectation at least an $\Omega((1/\log(r_\pi^{(k)})) \cdot (v_\pi^{\max(k)} / v_\pi^{\max(1)}))$ fraction of the expected revenue of the expected optimal social welfare (and therefore also of the expected revenue of the optimal dominant strategy IC, ex-post IR mechanism).

Proof. Let M be the sequential posted price mechanism that draws a value p uniformly at random from the set $S = \{v_\pi^{\min(k)} 2^j : j \in [\lceil \log(r_\pi^{(k)}) - 1 \rceil] \cup \{0\}\}$. M offers price p to all the buyers in an arbitrary order, until k buyers accept.

Let $\pi_{\max(k)}$ be the probability distribution of the k -th highest value of π . Note that $|S|$ does not exceed $\log(r_\pi^{(k)})$. Therefore the probability that p is the highest possible value (among the values in S) that does not exceed the value drawn from $\pi_{\max(k)}$, is equal to $1/\log(r_\pi^{(k)})$. More formally, let π_S be the probability distribution from which p is drawn; then

$$\Pr_{v^{\max(k)} \sim \pi_{\max(k)}, p \sim \pi_S} [p \leq v^{\max(k)} \cap (\nexists p' \in S : p' > p \wedge p' \leq v^{\max(k)})] \leq \frac{1}{\log(r_\pi^{(k)})}.$$

Thus, with probability $1/\log(r_\pi^{(k)})$, the mechanism extracts from each winner a revenue of exactly $v_\pi^{\min(k)} 2^j$, where j is the number such that the value drawn from $\pi_{\max(k)}$ lies in between $v_\pi^{\min(k)} 2^j$ and $v_\pi^{\min(k)} 2^{j+1}$. This implies that with probability $1/\log(r_\pi^{(k)})$ the mechanism extracts from Buyer i a revenue that lies a factor of $O(v_\pi^{\max(k)} / v_\pi^{\max(1)})$ away from $v_\pi^{\max(1)}$. This leads to the conclusion that

$$\mathbf{E}_{\vec{v} \sim \pi} [\text{revenue of } M(\vec{v})] \geq \Omega\left(\frac{1}{\log(r_\pi^{(k)})} \cdot \frac{v_\pi^{\max(k)}}{v_\pi^{\max(1)}}\right) \sum_{i \in W_M} v_\pi^{\max(1)},$$

where W_M denotes the set of buyers for which Mechanism M allocates the service. The theorem then follows since $\sum_{i \in W_M} v_\pi^{\max(1)} = \sum_{i \in W_{OPT}} v_\pi^{\max(1)} \geq OPT = \sum_{i \in W_{OPT}} v_i$, where W_{OPT} denotes the set of buyers at which the optimal mechanism allocates the service, and OPT is the social welfare achieved by the optimal mechanism. \square

We say that an instance (n, π, k) is k -well-separated if for any $\vec{v} \in \text{supp}(\pi)$ the k -th coordinate-wise maximum $v_\pi^{\max(k)}$ is achieved only by a single buyer, i.e., the set $\{i : v_i = v_\pi^{\max(k)}, \vec{v} \in \text{supp}(\pi)\}$ is a singleton. Then we can prove the following proposition.

Proposition 3. Let $n \in \mathbb{N}_{\geq 1}$, and let π be a discrete probability distribution on \mathbb{R}^n . For any $k \in [n]$, if the instance (n, π, k) is k -well-separated, then there exists a sequential posted price mechanism that, when run on instance (n, π, k) , extracts in expectation at least an $\Omega((1/\log(r_\pi^{(k)})) \cdot \max_{i \in [n]} \log(v_\pi^{\max(k)} / v_{\pi,i}^{\max}))$ fraction of the expected optimal social welfare (and therefore also of the expected revenue of the optimal dominant strategy IC and ex-post IR mechanism).

Proof. Since π is discrete, let δ be the smallest ratio larger than 1 between two valuation in $\vec{v} \in \text{supp}(\pi)$, i.e., $\delta = \min_{i,j} \{v_i/v_j > 1 : \vec{v} \in \text{supp}(\pi)\}$. Consider $\epsilon \leq \delta$ and let M be the sequential posted price mechanism that draws a value p uniformly at random from the set $S = \{v_\pi^{\min(k)} \epsilon^j : j \in [\lceil \log_\epsilon(r_\pi^{(k)}) - 1 \rceil] \cup \{0\}\}$. Moreover, M draws for each $i \in [n]$ a value p_i uniformly at random from the set $S_i = \{v_{\pi,i}^{\min} 2^\ell : \ell \in [\lceil \log \frac{p}{v_{\pi,i}^{\min}} \rceil, \lceil \log \frac{p}{v_{\pi,i}^{\min}} - 1 \rceil]\}$. M proposes prices to the buyers in an arbitrary order, and offers price p_i to Buyer i .

Let $\pi_{\max(k)}$ be the probability distribution of the k -th coordinate-wise maximum of π . Note that $|S|$ does not exceed $\log_\epsilon(r_\pi^{(k)})$. Therefore the probability that p is the highest possible value (among the values in S) that does not exceed the value drawn from $\pi_{\max(k)}$, is equal to $1/\log_\epsilon(r_\pi^{(k)})$. More formally, let π_S be the probability distribution from which p is drawn; then

$$\Pr_{v^{\max(k)} \sim \pi_{\max(k)}, p \sim \pi_S} [p \leq v^{\max(k)} \cap (\nexists p' \in S : p' > p \wedge p' \leq v^{\max(k)})] \leq \frac{1}{\log_\epsilon(r_\pi^{(k)})}.$$

Thus, with probability $1/\log_\epsilon(r_\pi^{(k)})$, the mechanism selects $p = v_\pi^{\min(k)} \epsilon^j$, where j is the number such that the value drawn from $\pi_{\max(k)}$ lies in between $v_\pi^{\min(k)} \epsilon^j$ and $v_\pi^{\min(k)} \epsilon^{j+1}$. When this event occurs, since the instance is k -well

separated and by our choice of ϵ , to the set W_M of buyers whose valuation is at least p has size exactly k and corresponds of the set W_{OPT} of buyers with the k highest valuation in $\vec{v} \in \text{supp}(\pi)$. Hence, with probability $1/\log_\epsilon(r_\pi^{(k)})$ Mechanism M extracts revenue only from buyers in W_{OPT} .

Now, for any $i \in W_{OPT}$, let π_i be the probability distribution of the i -th coordinate of π . Note that $|S_i|$ does not exceed $\log_\epsilon \frac{v_{\pi,i}^{\max}}{v_{\pi,i}^{\max(k)}}$. Therefore the probability that p_i is the highest possible value (among the values in S_i) that does not exceed the value drawn from π_i , is at least $\log_\epsilon \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}}$. More formally, let π_{S_i} be the uniform distribution on S_i ; then

$$\Pr_{v_i \sim \pi_i, p_i \sim \pi_{S_i}} [p_i \leq v_i \cap (\nexists p'_i \in S_i : p'_i > p_i \wedge p'_i \leq v_i)] \leq \log_\epsilon \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}}.$$

Thus, with probability $\frac{1}{\log_\epsilon(r_\pi^{(k)})} \cdot \log_\epsilon \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}}$, the mechanism extracts from Buyer $i \in W_{OPT}$ a revenue of exactly $v_{i,\min} 2^\ell$, where ℓ is the number such that the value drawn from π_i lies in between $v_{i,\min} 2^\ell$ and $v_{i,\min} 2^{\ell+1}$. This implies that with probability $\frac{1}{\log_\epsilon(r_\pi^{(k)})} \cdot \log_\epsilon \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}}$ the mechanism extracts from Buyer $i \in W_{OPT}$ a revenue that lies a factor of at most $1/2$ away from v_i . This leads to the conclusion that

$$\begin{aligned} \mathbf{E}_{\vec{v} \sim \pi} [\text{revenue of } M(\vec{v})] &\geq \sum_{i \in W_{OPT}} \frac{1}{\log_\epsilon(r_\pi^{(k)})} \cdot \log_\epsilon \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}} \frac{1}{2} \mathbf{E}_{v_i \sim \pi_i} [v_i] \\ &= \Omega \left(\frac{1}{\log_\epsilon(r_\pi^{(k)})} \cdot \max_{i \in [n]} \log \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}} \right) \sum_{i \in [n]} \mathbf{E}_{v_i \sim \pi_i} [v_i] \\ &= \Omega \left(\frac{1}{\log_\epsilon(r_\pi^{(k)})} \cdot \max_{i \in [n]} \log \frac{v_{\pi,i}^{\max(k)}}{v_{\pi,i}^{\max}} \right) \mathbf{E}_{\vec{v} \sim \pi} [OPT(\vec{v})], \end{aligned}$$

where $OPT(\vec{v}) = \sum_{i \in W_{OPT}} v_i$ denotes the optimal social welfare when the buyers have valuation vector v . \square

D Addressing some practical problems of enhanced sequential posted price mechanisms

A problematic aspect is that while there is no incentive for a buyer to lie, there is also no incentive to tell the truth. Therefore, incentive compatibility is only achieved in weakly dominant strategies. We note that in the literature many (or perhaps most) truthful mechanisms are only incentive compatible in the weak sense. Such mechanisms are of theoretical interest, and may possibly be turned into more practically satisfactory mechanisms.

In the case of enhanced SPP mechanisms, the lack of a strong incentive to be truthful only applies to those buyers who are asked for their values, knowing they will not be allocated the item. Such a buyer may not cooperate at all, or stating a false value. The first problem can be resolved by compensating the buyer with some fixed small amount of money that the auctioneer obtains from the buyers who pay for the service. Having insured some level of cooperation, how do we incentivize these buyers to be truthful?

Here is an example of such an adaptation of our enhanced SPP mechanisms that creates the proper strong incentive. Suppose now that we have provided an incentive for every buyer to reveal a valuation. At the start of the auction, using a cryptographic protocol (or just a normal sealed envelope), we ask each of the buyers for a sealed commitment of their value. Furthermore, for buyers being offered a price, with some (say small) probability, the buyer must reveal their private valuation in order to be allowed the item. If the revealed value is larger than the offered price, then the buyer is punished and the item is not allocated to him. Now this is strongly incentive compatible if we assume buyers are risk averse so that they will not over-bid their valuation. There is clearly no monetary reason for a buyer to under-bid.

E On d -dimensional dependence versus Markov random fields of degree d

This section is intended for readers who are interested in the relative generality of d -dimensionally dependence compared to Markov random fields of degree d . We assume that reader is familiar with the definition of Markov random fields. For convenience we will state a weaker notion here.

Definition 14. Given a undirected graph $G = ([n], E)$, a probability distribution π on \mathbb{R}^n is a local Markov random field with respect to G if the following property, named local Markov property, holds: for all $i \in [n]$, π_i is independent of $\pi_{[n] \setminus (\{i\} \cup \Gamma(i))}$ when conditioning on all coordinates in $\Gamma(i)$. Here, $\Gamma(i)$ denotes the neighborhood of i in G . The degree of a Markov random field is defined as the maximum degree of a vertex in its associated graph.

(In a true Markov random field, two additional technical conditions needs to be satisfied, called the *pairwise Markov property* and the *global Markov property*.) We will give an example of a 1-dimensionally dependent distribution that is not a local Markov random field with respect to any graph G in which all vertices have a degree less than $(n - 2)/2$.

Consider a distribution π on $\{0, 1\}^{n+2}$. A vector v drawn from π is formed according to the following random process: we are given $2n$ distinct probability distributions on $\{0, 1\}$. We name these distributions $\pi^{i,0}$ and $\pi^{i,1}$, for $i \in [n]$. These distributions are such that both 0 and 1 occur with positive probability. Let v' be a value drawn from yet another distribution π' on $\{0, 1\}$ where again both 0 and 1 have positive probability. The final generated vector is then $(v^{1,v'}, v^{2,v'}, \dots, v^{n,v'}, v', v')$, where $v^{i,v'}$ is drawn from $\pi^{i,v'}$.

π is clearly 1-dimensionally dependent, since for $i \in [n]$ the conditional marginal distribution $\pi_{i, \vec{v}_{-i}}$ is determined by only the value v' , which is the value of the $(n + 1)$ -th coordinate. Also, the value of the $(n + 1)$ -th coordinate is entirely determined by the $(n + 2)$ -th coordinate, and vice versa.

We can also easily see that π is not a Markov random field with respect to any graph in which all vertices have a degree less than $n/2$. Let G be a graph such that π is a Markov random field. Suppose for contradiction that there exists an $i \in [n]$ for which it holds that $\Gamma(i) \subseteq [n]$. Then the local Markov property would be violated. Therefore, each vertex in $[n]$ is connected to either vertex $n + 1$ or $n + 2$. Hence, we conclude that either vertex $n + 1$ or $n + 2$ has at least $n/2$ vertices attached to it.

F A blind offer mechanism that $O(1/d)$ -approximates the optimum revenue for Markov random fields of degree d

Consider the following procedure. Given Markov random fields of degree d (in which nodes represent buyers), at each round, mark a uniformly random node, then remove the node and all its neighbors and repeat until no node is left. Note that, each node is marked with probability at least $1/(d + 1)$. Consider now the following blind offer mechanism: ask the unmarked buyers to reveal their valuations, and run a sequential posted price mechanism on the marked buyers. The marked buyers' values are independent conditioned on the other values of the others. Then, one can use the results of Chawla et al. [2014] and Yan [2011] to extract a fraction $1 - 1/e$ of the revenue extractable from them⁶.

⁶We acknowledge an anonymous reviewer for pointing out this simple algorithm.