Relaxing the Independence Assumption in Sequential Posted Pricing, Prophet Inequality, and Random Bipartite Matching

Ioannis Caragiannis[1(✉)], Nick Gravin[2], Pinyan Lu[2], and Zihe Wang[3]

[1] Department of Computer Science, Aarhus University, Aarhus, Denmark
iannis@cs.au.dk

[2] ITCS, Shanghai University of Finance and Economics, Shanghai, China
{nikolai,lu.piyan}@mail.shufe.edu.cn

[3] Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China
wang.zihe@ruc.edu.cn

Abstract. We reexamine three classical settings of *optimization under uncertainty*, which have been extensively studied in the past, assuming that the several random events involved are mutually independent. Here, we assume that such events are only pair-wise independent; this gives rise to a much richer space of instances. Our aim has been to explore whether positive results are possible even under the more general assumptions. We show that this is indeed the case.

Indicatively, we show that, when applied to pair-wise independent distributions of buyer values, sequential posted pricing mechanisms get at least $\frac{1}{1.299}$ of the revenue they get from mutually independent distributions with the same marginals. We also adapt the well-known prophet inequality to pair-wise independent distributions of prize values to get a 1/3-approximation using a non-standard uniform threshold strategy. Finally, in a stochastic model of generating random bipartite graphs with pair-wise independence on the edges, we show that the expected size of the maximum matching is large but considerably smaller than in Erdős-Renyi random graph models where edges are selected independently. Our techniques include a technical lemma that might find applications in other interesting settings involving pair-wise independence.

Keywords: Posted pricing · Auctions · Prophet inequality · Revenue maximization · Bipartite matching

1 Introduction

Optimization in environments with uncertainty has received much attention in several research areas. It plays a central role in modern EconCS research (see, e.g., [7,10] for early surveys in bayesian mechanism design) and is also pervasive,

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, pp. 131–148, 2022.
https://doi.org/10.1007/978-3-030-94676-0_8

more broadly, in TCS (the rich theory of random graphs [3] is an example). Uncertainty manifests itself in many different ways as the following three settings indicate:

Sequential Posted Pricing: A seller has a single item to sell to n buyers. Each buyer has a random value v_i for the item distributed as $v_i \sim F_i$. The seller knows distributions $\{F_i\}_{i=1}^n$ and approaches buyers one by one in an arbitrary or fixed order. She offers the item to buyer i at a price p_i, which i takes when $v_i \geq p_i$ and pays p_i to the seller. The goal is to find a pricing scheme that maximizes the seller's revenue in expectation.

Optimal Stopping: A gambler plays a series of n games, each game $i \in [n]$ has a prize v_i distributed according to distribution F_i. The order of the games and the distribution of the prize values are known in advance to the gambler. Once the prize v_i for game i is realized, the gambler must decide whether to keep this prize and abandon the remaining games, or to discard this prize and continue playing. The gambler wants to maximize the expected reward.

Random Graph Models: Well-known models for the generation of random graphs assume a fixed set of nodes and produce each edge e between a pair of distinct nodes with a marginal probability p_e. Several graph parameters (e.g., the size of the maximum matching) have important meaning in areas like brain science, networking, or social sciences, and bounding these parameters is the subject of much research in many fields.

A simplifying assumption in most studies of the above settings is that the marginal distributions are mutually independent, i.e., the joint distribution is a product distribution. Under such an assumption, it is well-known that sequential posted pricing yields approximately-optimal revenue in single-parameter settings and generalizes nicely to multi-parameter environments [6]. Also, the optimal stopping strategy for the gambler can be computed by backward induction. A celebrated result, known as the *prophet inequality* [15], suggests that a simple threshold strategy can give an expected reward to the gambler that is at least 50% of the reward that could be achieved by a very powerful prophet, who has access to the maximum realized prize value [10,14]. Finally, in the random graph model where edges among pairs of n fixed nodes are drawn independently with the same probability p, a value of $p = \Omega(\ln n/n)$ is sufficient so that a hamiltonian cycle and, hence, a perfect matching exists, with high probability (see [3] for related results in random graphs). Unfortunately, such results (i.e., tight bounds or good approximations to revenue, gambler reward, or size of the maximum matching) do not hold for arbitrary joint distributions.

On the other hand, a recent line of work [5,9] on the monopoly problem for an additive buyer has proposed an alternative *correlation-robust framework* to study general distributions from the robust optimization perspective (see also [1,2]). In this framework, the algorithm designer knows only marginal distributions $\{F_i\}_{i=1}^n$ of each piece of the input and is given no information about correlation across different pieces in the joint distribution. The evaluation of the algorithm's performance is then taken in the worst-case, over the uncertainty of the problem, i.e., over all possible joint distributions with the specified set of marginal distributions $\{F_i\}_{i=1}^n$. The fact that the joint distribution is not explicitly given has an important practical advantage. Indeed, the representation and sampling complexity for learning correlated multidimensional distributions is exponential in the dimension n. In this respect, learning and operating with information about separate marginals is a much simpler task that does not suffer from the curse of dimensionality. However, the correlation-robust framework does not allow to incorporate any extra information about the distribution beyond the marginals. For example in the monopoly problem, there is no obstacle for the seller to acquire additional information, say, about dependencies between pairs of items by doing more extensive market research. On the other hand, the expected performance guarantees in the correlation robust framework are rather pessimistic compared to the mutually independent case and the worst-case joint distribution of the input often admits strong dependencies between its parts. The latter is not something we usually expect in practice, where it is more likely to see rather weak dependencies and significant variability between any given pair of input's components.

Our goal in this work is to model and study such situations with potential (weak) dependencies between input components. A straightforward approach would be to extend the correlation robust framework as follows: (i) specify the set of marginal distributions for any pair of input components $\{F_{i,j}\}_{i,j\in[n]}$ (e.g., the joint distribution of values for each pair of items $(v_i, v_j) \sim F_{i,j}$) and (ii) evaluate the expectation in the worst-case over all feasible joint distributions that agree with $\{F_{i,j}\}_{i,j\in[n]}$. Unfortunately, not all such pair-wise distributions (even consistent with singleton marginals $\{F_i\}_{i=1}^n$) would admit a feasible joint distribution[1]. Even when there exists a feasible joint distribution the set of pair-wise marginal distributions can sometimes uniquely identify the joint distribution. Moreover, the extra information does not necessarily help when we compare to the worst case distribution π^* for $\{F_i\}_{i=1}^n$. Indeed, one can take the worst-case joint distribution π^* and write down the restriction of π^* to $\{\pi^*_{i,j}\}_{i,j\in[n]}$. Then, the performance for this set of pair-wise marginals would not be better than that for π^*.

To avoid these complications we consider an important special case where pair-wise marginal distributions $\{F_{i,j}\}_{i,j\in[n]}$ are all independent, i.e., $F_{i,j} = F_i \times F_j$ for all $i, j \in [n]$. In other words, we assume that joint distribution is *pair-wise independent*. At first glance, pair-wise independence might appear rather similar to the standard mutually independence assumption. However, there are some important differences which we discuss below. We will also highlight the importance of this robust optimization approach assuming a pair-wise independent joint distribution.

1. *(Statistics vs. probability model).* The idealistic model with mutually independent distributions is a probability model that is not easy to verify in the proper statistical sense. Indeed, the joint distribution has exponential dependency on the number of components, so it would take a super-polynomial number of samples to confirm that the input distribution is close in total

[1] For example, if value of item B is always equal to the values of items A and C, then items A and C cannot be independent.

variation distance to the specific product distribution. On the other hand, the pair-wise independence condition is a statistical condition that can be checked in practice with only polynomially many samples.

2. *(Robustness to weak dependencies of data).* In practice, a multi-dimensional data distribution usually exhibits some form of mutual dependency that might be noticeable at the level of pair-wise marginal distributions. However, these dependencies are often weak and it is still reasonable to approximate each marginal pair distribution $F_{i,j}$ as the product distribution $F_i \times F_j$. In other words, we might want to allow small approximation errors in our model (using a pair-wise independent distribution) to the real joint distribution. The mutually independent model is a specific distribution, and as such can be too far from the most likely distribution matching the data. On the other hand, a robust guarantee for any pair-wise independent distribution is still meaningful even if the pair-wise marginals of the true input distribution are slightly perturbed compared to $F_i \times F_j$.
3. *(Large class of distributions).* To understand the size of the class of pair-wise independent distributions, let us consider the case of finite discrete supports, i.e., each marginal distribution F_i has finite support of size $|F_i|$. In this case the dimension of the simplex of feasible joint distributions is $|F_1| \cdot \ldots \cdot |F_n|$ and the mutually independent distribution is a single point. On the other hand, there are not more than $\sum_{i<j} |F_i| \cdot |F_j|$ linear constraints in the description of a pair-wise independent distribution. The product distribution is pair-wise independent and has positive probability (i.e., the inequality $\mathbf{Pr}[\mathbf{v}] \geq 0$ is not tight) for any point $\mathbf{v}$ in the support. Hence, the dimension of the pair-wise independent distribution is at least $\prod_{i=1}^{n} |F_i| - \sum_{i<j} |F_i| \cdot |F_j|$.

1.1 Our Results and Techniques

In this paper we study the three settings we discussed in the beginning of this section. We show that *any* sequential posted pricing mechanism with a given set of prices $\{p_i\}_{i=1}^n$ has an expected revenue that is at most 1.299 times larger in the case of mutually independent distributions of buyer values compared to the case of pair-wise independent distributions with the same marginals. Our result only requires that prices $\{p_i\}_{i=1}^n$ are offered in Pareto-optimal order, i.e., from higher to lower prices. The main tool we exploit to prove this result is a lemma that is conceptually similar to Lovász Local Lemma (LLL; see [13]). Recall that LLL bounds the probability that none among a series of events happens in terms of the marginal probabilities of these events, provided that they have a certain structure of dependencies. Our LLL-type statement bounds the probability that none among a series of pair-wise independent events happen in terms of their marginal probabilities. We believe that this lemma will find application beyond the scope of the current paper. We give an example that shows that our lemma is essentially tight; this implies that our bound on the revenue of sequential posted pricing is tight as well.

We also present variations of the prophet inequality when prizes have pair-wise independent values. A non-standard uniform threshold strategy yields the following guarantee. The worst expected reward of the gambler among all pair-wise independent prize value distributions with given marginals is at least 1/3 of the best expected prophet's reward over all pair-wise independent distributions with the same marginals. Again, we exploit an alternative expression of our local lemma. Interestingly, we show that uniform threshold strategies cannot yield a guarantee better than 40%, in contrast to the 50% guarantee of the classical prophet inequality [15] (see also [10, 14]) for mutually independent distributions. A non-uniform threshold strategy (exploiting ideas from [6]) is shown to break this barrier and achieve a 41.4% guarantee, at least for continuous pair-wise independent distributions. It is slightly more complicated though and requires additional information on the joint distribution besides the marginals.

Notice that the prophet inequality bounds are not universal like the ones for sequential posted pricing. Specifically, we show that there exists a uniform threshold strategy that achieves constant approximation to the prophet's reward for a mutually independent prize value distribution but achieves very low expected reward for a pair-wise independent distribution with the same marginals. Our results indicate that sequential posted pricing and optimal stopping are two economic settings where positive results are possible by relaxing independence to pair-wise independence. We demonstrate that such results are not possible for second price auctions. Broadening the class of economic problems that are "friendly" to the pair-wise independence assumption is an important direction for future research.

Finally, we consider a stochastic model for generating random bipartite graphs with n nodes in each side of the bipartition, so that each edge exists with some (non-necessarily uniform) probability. We assume that the expected degree of any node is Δ. When edges exist in the graph mutually independently, folklore results (e.g., see [3]) suggest that a perfect matching exists almost certainly, provided that $\Delta = \Omega(\ln n)$. Furthermore, we can show that the expected size of the maximum matching is $n - n \cdot O(\exp(-\Delta))$. In contrast, in the case of pair-wise independence (on the existence of edges), the lower bound we can show is $n - n/\sqrt{\Delta}$, which leaves open the possibility of non-existence of perfect matchings for all interesting range of values for parameter Δ. Our proof is based on a second-moment argument and exploits the fact that the maximum matching has the same size with the minimum vertex cover in a bipartite graph. We also present a non-trivial pair-wise independent distribution over bipartite graphs that shows that our bound is essentially tight. These results indicate that a revision of classical results on random graphs under the pair-wise independence lens might reveal a very interesting new picture.

1.2 Roapmap

The rest of the paper is structured as follows. We present our local-lemma-type statement in Sect. 2. Sequential posted pricing is studied in Sect. 3. Our results for prophet inequalities are presented in Sect. 4. Section 5 is devoted to proving our bounds for random bipartite matchings.

2 A Local-Lemma-Type Probability Statement

We begin by proving an LLL-style probability statement for pair-wise independent distributions. The lemma will be particularly useful in Sects. 3 and 4 but we believe that it will find applications in other settings as well.

Let $\{E_i\}_{i=1}^n$ be a set of random events with $\mathbf{Pr}\left[E_i\right] = q_i$. We are interested in the probability that at least one of the random events happens. If these events are mutually independent, this probability is exactly

$$\mathbf{Pr}\left[\bigvee_{i=1}^n E_i\right] = 1 - \prod_{i=1}^n (1 - q_i).$$

We want to lower bound this probability when the events are only known to be pair-wise independent. In Lovász local lemma (LLL), these random events are either mutually independent or worst-case correlated, and LLL gives a low bound on the probability that none of the events happens. In a sense, the two lemmas both relax the independence assumption to a "local" assumption but in different directions. LLL models situations where the dependencies are only happening locally: an event is mutually independent with all other events except its neighbors. In our lemma, we model the setting where independence is only guaranteed locally: any pair of events are independent with each other but not necessarily globally. We prove that, in any pair-wise independent distribution, the probability is at least a constant fraction of the probability in the mutually independent setting.

Lemma 1. *Let $\{E_i\}_{i=1}^n$ be a set of random events. Let $\mathbf{F}_{ind}$ and $\mathbf{F}_\pi$ be a mutually independent and a pair-wise independent distribution over these events, respectively, with $\mathop{\mathbf{Pr}}_{\mathbf{F}_{ind}}\left[E_i\right] = \mathop{\mathbf{Pr}}_{\mathbf{F}_\pi}\left[E_i\right] = q_i$. Then,*

$$\mathop{\mathbf{Pr}}_{\mathbf{F}_\pi}\left[\bigvee_{i=1}^n E_i\right] \geq \frac{\sum_{i=1}^n q_i}{1 + \sum_{i=1}^n q_i}$$

and

$$\mathop{\mathbf{Pr}}_{\mathbf{F}_\pi}\left[\bigvee_{i=1}^n E_i\right] \geq \frac{1}{1.299} \mathop{\mathbf{Pr}}_{\mathbf{F}_{ind}}\left[\bigvee_{i=1}^n E_i\right] = \frac{1}{1.299}\left(1 - \prod_{i=1}^n (1 - q_i)\right).$$

Proof. We prove only the first inequality here; the proof of the second one is omitted. We denote by X_i the indicator random variable for event E_i and define the random variable $X = \sum_{i=1}^n X_i$. Then $\bigvee_{i=1}^n E_i$ (the event we are interested in) is the random event $\{X > 0\}$. By definition, we have $\mathbf{E}\left[X_i\right] = q_i$ and thus $\mathbf{E}\left[X\right] = \sum_{i=1}^n q_i$.

Since the random variables $\{X_i\}_{i=1}^n$ are pair-wise independent, we have

$$\mathbf{Var}\left[X\right] = \sum_{i=1}^n \mathbf{Var}\left[X_i\right] + 2\sum_{i<j} \mathrm{Cov}(X_i, X_j) = \sum_{i=1}^n \mathbf{Var}\left[X_i\right] = \sum_{i=1}^n q_i(1 - q_i).$$

Let $f_k = \mathop{\mathbf{Pr}}_{X \sim \mathbf{F}_\pi}[X = k]$ for all $i \in [n]$. Using this notation and applying Cauchy-Schwartz's inequality, we have

$$\mathop{\mathbf{Pr}}_{X\sim \mathbf{F}_\pi}[X > 0] = \sum_{k=1}^{n} f_k \geq \frac{\left(\sum_{k=1}^n k \cdot f_k\right)^2}{\sum_{k=1}^n k^2 \cdot f_k} = \frac{\mathbf{E}[X]^2}{\mathbf{E}[X^2]}$$
$$= \frac{\mathbf{E}[X]^2}{\mathbf{Var}[X] + \mathbf{E}[X]^2} = \frac{(\sum_{i=1}^n q_i)^2}{\sum_{i=1}^n q_i(1-q_i) + (\sum_{i=1}^n q_i)^2}.$$

This immediately gives us $\mathop{\mathbf{Pr}}_{\mathbf{F}_\pi}[\bigvee_{i=1}^n E_i] = \mathop{\mathbf{Pr}}_{X\sim\mathbf{F}_\pi}[X > 0] \geq \frac{(\sum_{i=1}^n q_i)^2}{\sum_{i=1}^n q_i + (\sum_{i=1}^n q_i)^2} = \frac{\sum_{i=1}^n q_i}{1+\sum_{i=1}^n q_i}$. □

The upper bound of 1.299 in the statement of Lemma 1 is almost tight. Here, we give an example where the gap between $\mathop{\mathbf{Pr}}_{\mathbf{F}_{\text{ind}}}[\bigvee_{i=1}^n E_i]$ and $\mathop{\mathbf{Pr}}_{\mathbf{F}_\pi}[\bigvee_{i=1}^n E_i]$ is at least 1.296. In the example, $q_i = q = 2/(n-1)$ for all $i \in [n]$. For the distribution $\mathbf{F}_{\text{ind}}$, we have $\mathop{\mathbf{Pr}}_{\mathbf{F}_{\text{ind}}}[\bigvee_{i=1}^n E_i] = 1-(1-q)^n$, which approaches $1-e^{-2}$ as n goes to infinity.

Now consider the following probability distribution. With probability $\frac{2n}{3(n-1)}$ a set of exactly three events among $\{E_i\}_{i=1}^n$ happen. These three events are chosen uniformly at random among the $\binom{n}{3}$ possible choices. With the remaining probability of $1-\frac{2n}{3(n-1)}$ no event happens. In this distribution, the probability that at least one event happens approaches $\frac{2}{3}$ as n goes to infinity. The ratio between the two probabilities approaches $(1-e^{-2})/(\frac{2}{3}) = 1.29699$.

It remains to show that this distribution is pair-wise independent. Indeed, for any $i \neq j$, we have:

$$\mathbf{Pr}[E_i \wedge E_j] = \frac{2n}{3(n-1)} \cdot \frac{\binom{n-2}{1}}{\binom{n}{3}} = \frac{4}{(n-1)^2} = q^2,$$
$$\mathbf{Pr}\left[E_i \wedge \overline{E_j}\right] = \mathbf{Pr}\left[\overline{E_i} \wedge E_j\right] = \frac{2n}{3(n-1)} \cdot \frac{\binom{n-2}{2}}{\binom{n}{3}} = \frac{2(n-3)}{(n-1)^2} = q(1-q), \text{ and}$$
$$\mathbf{Pr}\left[\overline{E_i} \wedge \overline{E_j}\right] = 1 - q^2 - 2q(1-q) = (1-q)^2,$$

as desired.

3 Sequential Posted Pricing

In this section, we consider the setting with a seller, who aims to sell a single item to n potential buyers. Buyer $i \in [n]$ has value v_i distributed according to distribution F_i. The seller uses a *sequential posted pricing* mechanism. He considers the buyers one by one according to their index; when considering buyer i, the seller offers her the item at price p_i. We consider Pareto-efficient pricing schemes that satisfy $p_1 \geq p_2 \geq ... \geq p_n$. In this way, the seller does not risk to loose revenue to low price buyers, who will have their chance only after the item is offered to other buyers at a higher price. Our aim is to analyze sequential posted pricing mechanisms assuming that the distributions F_i are pair-wise independent and compare their revenue to the revenue they would have when applied to mutually independent valuations with the same marginals.

Let us consider the simple case where a uniform price p is offered to all buyers. In this case, we can directly use Lemma 1 to conclude that the expected revenue over any pair-wise independent distribution is at least $\frac{1}{1.299}$ of the revenue of a corresponding mutually independent distribution. Indeed, denote by $E_i = \{v_i \geq p\}$ the event that buyer i accepts the price. As the price is the same for all buyers, the probability guarantee immediately translates into the revenue guarantee. Hence, the revenue of the mechanism when the events E_i are pair-wise independent is at least $\frac{1}{1.299}$ times the revenue of the mechanism when these events are mutually independent but have the same probabilities.

For the general case where prices can be different, we also need to pay attention to who gets the item. To do so, we apply Lemma 1 n times, each time considering the first k buyers for $k = 1, 2, \cdots, n$. Let $\lambda = 1.299$ be the approximation guarantee from Lemma 1. Let X_i be the random variable indicating whether the value of buyer i is at least p_i, and let q_i be the probability that $X_i = 1$. Using the second inequality from Lemma 1, we have

$$\lambda \cdot \mathbf{Pr}\left[\sum_{i=1}^{k} X_i > 0\right] \geq 1 - \prod_{i=1}^{k}(1 - q_i).$$

We multiply this inequality by the price difference $p_k - p_{k+1}$ to get

$$\lambda(p_k - p_{k+1})\mathbf{Pr}\left[\sum_{i=1}^{k} X_i > 0\right] \geq (p_k - p_{k+1})\left(1 - \prod_{i=1}^{k}(1 - q_i)\right),$$

where $p_{n+1} = 0$. After summing these inequalities for $k \in \{1, \ldots, n\}$, we get

$$\lambda \cdot \sum_{k=1}^{n} p_k \cdot \mathbf{Pr}\left[X_k = 1,\ \forall i < k,\ X_i = 0\right] \geq \sum_{k=1}^{n} p_k q_k \prod_{i=1}^{k-1}(1 - q_i) \tag{1}$$

Observe that the LHS of equation (1) is equal to λ times the revenue generated by the sequential posted pricing mechanism, while the RHS of (1) is equal to the revenue the mechanism would have if the valuations of buyers were mutually independent. We summarize this observation in the following statement.

Theorem 1. *A posted pricing mechanism under a pairwise independent distribution of buyer valuations achieves at least* $1/1.299$ *fraction of the revenue under a mutually independent distribution with the same marginals.*

Our lower bound from Sect. 2 directly translates to a posted pricing instance with uniform prices. Thus, a sequential posted pricing mechanism for mutually independent distributions can generate revenue at least 1.296 time more than for pair-wise independent distributions.

We note that such robust properties do not necessary hold in other mechanism design settings. In particular, in a second price auction, the revenue gap in the cases of pair-wise independent and mutually independent buyer valuations can be huge. The reason is that in the mutually independent case, the second largest bid is high with large probability, while in the pairwise independent case, this probability can be very small.

First, consider the following setting. There are n i.i.d. buyers. Each buyer has value 1 with probability $\frac{1}{n-1}$ and value 0 otherwise. The revenue of the second price auction is then the probability that at least two buyers have value 1, i.e., $1-\left(1-\frac{1}{n-1}\right)^n-\frac{n}{n-1}\left(1-\frac{1}{n-1}\right)^{n-1}=1-2\left(1-\frac{1}{n-1}\right)^{n-1}\geq 1-2/e$.

We now construct a pairwise independent distribution in which the value of each buyer is 1 with probability $\frac{1}{n-1}$ and 0 otherwise, and under which the generated revenue is very small. In this distribution, there are two kinds of valuation profiles. The first profile appears with probability $\frac{1}{(n-1)^2}$ and all buyers have value 1. In the second profile, one buyer, selected uniformly at random, has value 1 and the rest of the buyers have value 0. The second price auction has expected revenue of $\frac{1}{(n-1)^2}$ as it gets a revenue of 1 only on the first profile. One can easily verify that the probability distribution is indeed pair-wise independent.

4 Prophet Inequality

Sequential posted pricing is closely related to the *prophet inequality* from optimal stopping theory [15]. In this scenario a gambler plays sequentially a series of n games. Each game $i \in [n]$ has prize v_i distributed according to distribution F_i. The order of the games and the distribution of the prize values are known in advance to the gambler. Once the prize v_i for game i is realized, the gambler must decide whether to keep this prize and abandon the remaining games, or to discard this prize and continue playing. A prophet in this setting knows the realization of all prizes in advance and therefore can stop at the right moment and take the highest prize.

4.1 A Uniform Threshold Policy

It is well-known that the gambler can achieve a 2-approximation of the optimal prize by following a simple *uniform threshold strategy*, which is given by a single threshold $\widehat{v}$ and requires the gambler to accept the first prize i with $v_i \geq \widehat{v}$. The standard assumption in the prophet inequality literature is that the prize distributions $\{F_i\}_{i=1}^n$ for different games are mutually independent.[2] Here, we only assume that the distributions $\{F_i\}_{i=1}^n$ are pair-wise independent.

[2] We note that prophet inequalities for classes of prize distributions with limited correlation have been studied before. The survey of Hill and Kertz [11] discusses early related results in stopping theory while the papers [4,8,12] are representative of recent related work by the EconCS community. However, such results are rarely based on the use of simple threshold strategies as the one we use here.

Theorem 2. *For any set of marginal prize distributions $\{F_i\}_{i=1}^n$, there exists a threshold $\widehat{v}$ such that the expected reward of the uniform threshold strategy for any pair-wise independent joint distribution is at least $1/3$ of the expected value of the maximum prize*[3]*.*

Proof. Let REF denote the expected reward of the prophet, i.e., the expected maximum prize $\mathbf{E}[\max_i v_i]$, and APX denote the expected reward of the gambler with a uniform threshold strategy. We denote by $\widehat{v}$ the uniform threshold (to be defined later). Let $x = \sum_{i=1}^n \mathbf{Pr}[v_i \geq \widehat{v}]$.

We use the same upper bound for REF as in the standard exposition of the prophet inequality (e.g., see [10,14]).

$$\mathsf{REF} = \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathbf{F}}\left[\max_i v_i\right] \leq \widehat{v} + \mathop{\mathbf{E}}_{\mathbf{v}\sim\mathbf{F}}\left[\max_i (v_i - \widehat{v})^+\right] \leq \widehat{v} + \sum_i \mathop{\mathbf{E}}_{v_i\sim F_i}\left[(v_i - \widehat{v})^+\right].$$

We note that this upper bound holds for any joint distribution with the given marginal distributions (not necessary pair-wise independent). In this bound the RHS depends only on the marginal distributions.

We will split the gambler's reward APX into two parts: (i) the first part, APX_1, is the guaranteed contribution of $\widehat{v}$ if some reward is taken and (ii) the second part, APX_2, is the extra contribution of $v_i - \widehat{v}$ when i is chosen. To bound APX_1, we use Lemma 1:

$$\mathsf{APX}_1 = \mathbf{Pr}_\pi\left[\max_i v_i \geq \widehat{v}\right] \cdot \widehat{v} \geq \frac{\sum_{i=1}^n \mathbf{Pr}[v_i \geq \widehat{v}]}{1 + \sum_{i=1}^n \mathbf{Pr}[v_i \geq \widehat{v}]} \cdot \widehat{v} = \frac{x}{1+x} \cdot \widehat{v}, \tag{2}$$

where $x \stackrel{\text{def}}{=} \sum_{i=1}^n \mathbf{Pr}[v_i \geq \widehat{v}]$. In general, the notation $\mathbf{Pr}_\pi[\cdot]$ is used when pair-wise independent prizes are considered.

To bound APX_2, we define the event $\mathcal{E}_{i,v} \stackrel{\text{def}}{=} \{\mathbf{v}\,|v_i = v;\ \forall j \neq i, v_j < \widehat{v}\}$ for every $v \geq \widehat{v}$, i.e., the reward in game i is $v_i = v$ while all the remaining prizes are below the threshold. The crucial property of any joint pair-wise independent distribution π is that $\mathbf{Pr}_\pi[\mathcal{E}_{i,v}] \geq (1-x)\mathbf{Pr}_{F_i}[v_i = v]$ which we show below. Indeed, by definition

$$\mathbf{Pr}_\pi[\mathcal{E}_{i,v}] = \mathbf{Pr}_{F_i}[v_i = v] \cdot \mathbf{Pr}_\pi\left[\bigcap_{j\neq i}[v_j < \widehat{v}] \middle| v_i = v\right].$$

By the union bound, we have

$$\mathop{\mathbf{Pr}}_\pi\left[\bigcap_{j\neq i}[v_j < \widehat{v}] \middle| v_i = v\right] \geq 1 - \sum_{j\neq i} \mathop{\mathbf{Pr}}_\pi[v_j \geq \widehat{v}|v_i = v].$$

[3] We remark that, while the expectation for the threshold strategy is taken in the *worst case* over any pair-wise independent distribution, the expectation for the prophet is taken in the *best case* over any distribution with the given marginals.

Due to pair-wise independence, $\mathbf{Pr}_{\pi}\left[v_j \geq \widehat{v} | v_i = v\right] = \mathbf{Pr}_{F_j}\left[v_j \geq \widehat{v}\right]$. By definition of x, we know that $\sum_{j\neq i} \mathbf{Pr}\left[v_j \geq \widehat{v}\right] \leq x$. Hence,

$$\mathbf{Pr}\left[\mathcal{E}_{i,v}\right] \geq \mathbf{Pr}\left[v_i = v\right] \cdot \left(1 - \sum_{j\neq i} \mathbf{Pr}\left[v_j \geq \widehat{v}\right]\right) \geq (1-x)\mathbf{Pr}\left[v_i = v\right].$$

When $\mathcal{E}_{i,v_i}$ happens, we get the additional contribution of $v_i - \widehat{v}$. As all random events $\{\mathcal{E}_{i,v_i}\}_{i\in[n],v_i\geq\widehat{v}}$ are disjoint, we have

$$\mathsf{APX}_2 \geq (1-x)\sum_{i=1}^{n} \int_{v_i\geq\widehat{v}} (v_i - \widehat{v})\,\mathrm{d}F_i(v_i) = (1-x)\sum_{i=1}^{n} \mathbf{E}\left[(v_i-\widehat{v})^{+}\right]. \qquad (3)$$

Since $\sum_{i=1}^{n} \mathbf{E}\left[(v_i-\widehat{v})^{+}\right]$ is a continuous function of $\widehat{v}$ decreasing to zero when $\widehat{v} \to \infty$, we can choose[4] the threshold $\widehat{v}$ so that

$$\widehat{v} = 2 \cdot \sum_{i=1}^{n} \mathbf{E}\left[(v_i-\widehat{v})^{+}\right].$$

Then, by the definition of REF, we have

$$\widehat{v} \geq \frac{2}{3}\cdot\mathsf{REF}, \quad \text{and} \quad \sum_{i=1}^{n} \mathbf{E}\left[(v_i-\widehat{v})^{+}\right] \geq \frac{1}{3}\cdot\mathsf{REF}.$$

If $x \geq 1$, then the lower bound (3) is trivial and we only use (2) to get

$$\mathsf{APX} \geq \mathsf{APX}_1 \geq \frac{1}{2}\cdot\widehat{v} \geq \frac{1}{3}\cdot\mathsf{REF}.$$

Otherwise, if $0 \leq x \leq 1$, we combine (2) and (3) to get

$$\begin{aligned}\mathsf{APX} = \mathsf{APX}_1 + \mathsf{APX}_2 &\geq \left(\frac{2}{3}\frac{x}{1+x} + \frac{1}{3}(1-x)\right)\mathsf{REF}\\ &\geq \left(\frac{2}{3}\frac{x}{2} + \frac{1}{3}(1-x)\right)\mathsf{REF} = \frac{1}{3}\cdot\mathsf{REF}.\end{aligned}$$

This completes the proof. □

[4] If the distributions were continuous, we could choose the threshold $\widehat{v}$ so that $x = \frac{\sqrt{5}-1}{2}$. For this value of x we have $\frac{x}{1+x} = 1 - x = \frac{3-\sqrt{5}}{2} = 0.382$. Then, we could get a lower bound on APX by combining (2) and (3) as follows:

$$\mathsf{APX} = \mathsf{APX}_1 + \mathsf{APX}_2 \geq 0.382\left(\widehat{v} + \sum_{i=1}^{n} \mathbf{E}\left[(v_i-\widehat{v})^{+}\right]\right) \geq 0.382\cdot\mathsf{REF}.$$

In our proof, we assume that distributions can be discontinuous and, thus, we may not be able to set $\widehat{v}$ to get a particular value of x.

We now present limitations of uniform threshold strategies for pair-wise independent distributions. These limitations come in contrast with the case of mutually independent distributions, where some of such policies can give at least 50% of the prophet's value as reward.

Theorem 3. *Uniform threshold strategies cannot guarantee more than* 40% *of prophet's value for some pair-wise independent distributions of prize values.*

Proof. In the proof, we use the following distribution. There are $n+2$ items. Item 1 has a deterministic value of 1. Item $n+2$ has value n with probability $\frac{1}{n}$ and value 0 otherwise. The values of items $2, ..., n+1$ have identical marginal distributions: value 2 with probability $\frac{1}{n}$ and 0 with probability $\frac{n-1}{n}$. The joint distribution π is constructed as follows. When item $n+2$ has value n:

- With probability $\frac{1}{n} - \frac{1}{n^2}$, all items $2, ..., n+1$ have value 0.
- With probability $1 - \frac{1}{n}$, exactly one item among $2, ..., n+1$ has value 2 and the remaining items have value 0. The high-value item is selected uniformly at random among the items $2, ..., n+1$.
- With probability $\frac{1}{n^2}$, items $2, ..., n+1$ have all values 2.

When item $n+2$ has value 0,

- With probability $\frac{n-1}{2n}$, items $2, ..., n+1$ have all value 0.
- With probability $\frac{1}{n}$, exactly one item (selected uniformly at random) among $2, ..., n+1$ has value 2 and the remaining items have value 0.
- With probability $\frac{n-1}{2n}$, exactly two items among $2, ..., n+1$ have value 2 and the remaining items have value 0. The two items are chosen uniformly at random among the $\binom{n}{2}$ possible pairs.

It is straightforward to verify that this is a pair-wise independent distribution. The expected value of the prophet is

$$n \cdot \frac{1}{n} + \left(1 - \frac{1}{n}\right) \cdot \left(1 \cdot \frac{n-1}{2n} + 2 \cdot \frac{n+1}{2n}\right) = \frac{5}{2} - \frac{1}{n} - \frac{1}{2n^2}.$$

Using a threshold that is smaller than 1, the reward of the gambler is (deterministically) 1. Using a threshold higher than 2, her expected reward is $n \cdot \frac{1}{n} = 1$, too. Now, assume that a threshold in $(1, 2]$ is used. Then, the expected reward of the gambler is

$$\frac{1}{n}\left(n\left(\frac{1}{n} - \frac{1}{n^2}\right) + 2\left(1 - \frac{1}{n} + \frac{1}{n^2}\right)\right) + \left(1 - \frac{1}{n}\right)\left(2 \cdot \frac{n+1}{2n}\right) = 1 + \frac{3}{n} - \frac{4}{n^2} + \frac{2}{n^3}.$$

Hence, as n approaches infinity, the reward of the prophet approaches 5/2, while the reward of the gambler approaches 1 when using any uniform threshold strategy. The theorem follows. □

We have proved that there exists a threshold such that prophet inequality still holds with a slightly worse constant. However, unlike the case of sequential posted pricing where a constant gap holds for any choice of prices, this is not true for all choices of thresholds in the prophet inequality setting. We give an example where a certain threshold strategy achieves constant fraction of the maximum welfare in the mutually independent case, but it gets almost zero fraction in the pairwise independent case.

Our example has four items. The values of the first three items are 0 and 1, equally likely; the fourth item has large deterministic value $V > 1$. Assuming mutually independent values, the expected gain of the gambler when she uses a uniform threshold strategy with the threshold 1 is $\frac{7+V}{8}$. Now, consider the following pair-wise independent distribution, in which the gambler always gets a value of 1 when she uses 1 as a uniform threshold. The first three items have values $(1, 1, 1), (0, 1, 0), (1, 0, 0), (0, 0, 1)$ with equal probabilities. Our claim follows for large values of V.

4.2 Non-uniform Threshold Strategies

We now demonstrate that non-uniform threshold strategies can be more powerful than uniform ones. We adapt a technique from [6]. The gambler uses different thresholds τ_1, τ_2, ..., τ_n and n coins where the probability of the i-th coin toss to be heads is q_i. At step i, if the award has not been given before and the prize v_i exceeds the thresholds τ_i, the gambler tosses the i-th coin and gets the prize if it comes heads. In our analysis, we assume that the prize values follow continuous distributions.

Theorem 4. *For any set of continuous marginal prize distributions $\{F_i\}_{i=1}^n$, there exist thresholds $(\tau_i)_{i\in[n]}$ and probabilities $(q_i)_{i\in[n]}$ so that the expected reward of the gambler's strategy is at least $\sqrt{2} - 1 \approx 41.4\%$ of the expected value of the maximum prize.*

Proof. For $i = 1, 2, ..., n$, let $p_i = \mathbf{Pr}[v_i \geq v_j, \forall j \in [n]]$ and define τ_i to be such that $\mathbf{Pr}[v_i \geq \tau_i] = p_i$. Then, $\mathbf{E}[v_i \cdot \mathbb{1}\{v_i \geq v_j, \forall j \in [n]\}] \leq \mathbf{E}[v_i \cdot \mathbb{1}\{v_i \geq \tau_i\}]$ and

$$\mathbf{E}\left[\max_i v_i\right] \leq \sum_i \mathbf{E}\left[v_i \cdot \mathbb{1}\left\{v_i \geq \tau_i\right\}\right]. \tag{4}$$

For $i = 1, 2, ..., n$, let R_i be the event that no award has been given at steps $1, 2, ..., i-1$, P_i the event that $v_i \geq \tau_i$ (i.e., $\mathbf{Pr}[P_i] = p_i$), and Q_i the event that the random coin toss at step i comes heads (i.e., $\mathbf{Pr}[Q_i] = q_i$). For $i \geq 2$, we have

$$\begin{aligned}\mathbf{Pr}[R_i|v_i = v] &= \mathbf{Pr}\left[R_{i-1} \wedge \overline{P_{i-1} \wedge Q_{i-1}} \;\middle|\; v_i = v\right] \\ &\geq \mathbf{Pr}[R_{i-1} \mid v_i = v] + \mathbf{Pr}\left[\overline{P_{i-1} \wedge Q_{i-1}} \;\middle|\; v_i = v\right] - 1 \\ &\geq \mathbf{Pr}[R_{i-1} \mid v_i = v] - p_{i-1}q_{i-1}. \end{aligned} \tag{5}$$

The first inequality uses the property $\mathbf{Pr}[A \wedge B] = \mathbf{Pr}[A] + \mathbf{Pr}[B] - \mathbf{Pr}[A \vee B] \geq \mathbf{Pr}[A] + \mathbf{Pr}[B] - 1$. The second inequality follows since the events

P_{i-1} and Q_{i-1} and $\{v_i = v\}$ are independent. Summing the inequalities (5) for $i = 2, ..., n$ together with the obvious fact $\mathbf{Pr}[R_1 | v_i = v] = 1$, we get

$$\mathbf{Pr}[R_i | v_i = v] \geq 1 - \sum_{j<i} p_j q_j. \tag{6}$$

The expected award APX of the gambler is

$$\begin{aligned}
\mathsf{APX} &= \sum_{i=1}^{n} \int_{\tau_i}^{\infty} v \mathbf{Pr}[R_i | v_i = v] \, q_i \, \mathrm{d}F_i(v_i) \\
&\geq \sum_{i=1}^{n} q_i (1 - \sum_{j<i} p_j q_j) \int_{\tau_i}^{\infty} v \, \mathrm{d}F_i(v_i) \\
&\geq \min_i \left\{ q_i \left(1 - \sum_{j<i} p_j q_j \right) \right\} \cdot \sum_i \mathbf{E}[v_i \, \mathbb{1}\{v_i \geq \tau_i\}] \\
&\geq \min_i \left\{ q_i \left(1 - \sum_{j<i} p_j q_j \right) \right\} \cdot \mathbf{E}\left[\max_i v_i \right].
\end{aligned} \tag{7}$$

The first inequality follows by (6). The second one is obvious and the third one follows by (4).

We will now define the q_i's appropriately so that $\min_i \left\{ q_i \left(1 - \sum_{j<i} p_j q_j \right) \right\} \geq \sqrt{2} - 1$. The theorem will then follow by (7). Let $\alpha = \frac{\sqrt{2}-1}{2}$ and $\beta = (1+\sqrt{2})^2$ and define the function $g : [0, 1] \to \mathbb{R}_{\geq 0}$ with $g(x) = \sqrt{\frac{\alpha}{\beta - x}}$. It can be verified by tedious calculations that

$$g(x) \left(1 - \int_0^x g(t) \mathrm{d}t \right) = 2\alpha = \sqrt{2} - 1$$

for every $x \in [0, 1]$. Now, let $q_i = g\left(\sum_{j<i} p_j \right)$ and observe that $\sum_{j<i} p_j q_j \leq \int_0^{\sum_{j<i} p_j} g(t) \mathrm{d}t$ (as $g(t)$ is a decreasing function and the integral in the right hand side is larger than its Riemann sum for the partition into the intervals of lengths $(p_j)_{j<i}$). Hence, for every $i \in [n]$, we have

$$q_i \left(1 - \sum_{j<i} p_j q_j \right) \geq g\left(\sum_{j<i} p_j \right) \left(1 - \int_0^{\sum_{j<i} p_j} g(t) \mathrm{d}t \right) = \sqrt{2} - 1$$

as desired. □

5 Matchings in Random Bipartite Graphs

In this section, we consider a stochastic graph model for bipartite graphs, extending the classical Erdős-Renyi model. In particular, the stochastic model

$\mathbf{G}(L, R, \{p_e\}_{e \in L \times R})$ is a distribution over bipartite graphs $G = (L, R, E)$ with $E \subseteq R \times L$, such that the marginal probability of each edge $e \in R \times L$ to appear in G is equal to $\mathbf{Pr}_{G \sim \mathbf{G}}[e \in E(G)] = p_e$. We are interested in the case of stochastically Δ-regular n-vertex models, which generate bipartite graphs with $|L| = |R| = n$ and average degree Δ, i.e.,

$$\mathop{\mathbf{E}}_{G \sim \mathbf{G}}[\deg(u)] = \sum_{e:\{u\} \times R} p_e = \Delta \quad \text{and} \quad \mathop{\mathbf{E}}_{G \sim \mathbf{G}}[\deg(v)] = \sum_{e:L \times \{v\}} p_e = \Delta$$

for every vertex $u \in L$ and $v \in R$, respectively.

Note that there might be many Δ-regular n-vertex models with fixed marginal probabilities. The most well studied case is the adaptation of the Erdős-Renyi model $\mathbf{G}_{\text{ind}}$, where the events $e \in E(G)$ are mutually independent for $e \in L \times R$ with $p_e = p$ for all $e \in L \times R$. Here, we focus on models $\mathbf{G}_\pi$ where these events are pair-wise independent. Our aim is to prove bounds on the expected size $\mu(G)$ of the maximum matching of graph $G \sim \mathbf{G}_\pi$. It is well-known that for the model $\mathbf{G}_{\text{ind}}$, the expected size of the maximum matching is $n - O(\exp(-\Delta))$ and, hence, perfect matchings exist with high probability when Δ becomes (super)logarithmic. Such results are not possible in the more general pair-wise independent case; still, the expected size of the maximum matching is quite large.

Theorem 5. *Let $\mathbf{G}_\pi$ be a stochastic Δ-regular n-vertex model with marginals $\{p_e\}_{e \in L \times R}$ such that the events $\{e \in E(G)\}_{e \in L \times R}$ for $G \sim \mathbf{G}_\pi$ are pair-wise independent. Then, the expected size of the maximum matching of a randomly generated graph $G \sim \mathbf{G}_\pi$ is at least $\mathbf{E}_{G \sim \mathbf{G}_\pi}[\mu(G)] \geq n - n/\sqrt{\Delta}$.*

Proof. The main idea of the proof is to look at the aggregate distribution of the vertex degrees of the whole graph G. On the one hand, the pairwise independence condition allows us to calculate precisely the expectation and variance of the degree of any particular vertex. On the other hand, the non-existence of a large matching $\mu(G)$ in a realized graph $G \sim \mathbf{G}_\pi$ implies a large deviation of degrees of many vertices from their mean Δ. This allows us to get the desired bound on the following random variable $f(G)$, where $G \sim \mathbf{G}_\pi$.

$$f(G) \overset{\text{def}}{=} \sum_{v \in L \cup R} (d_v - \Delta)^2, \quad \text{where } d_v \text{ is the degree of each vertex } v \text{ in } G.$$

In any graph $G \sim \mathbf{G}_\pi$, let $E_v \overset{\text{def}}{=} \{e \in E(G) \big| e \text{ is incident to } v\}$ for each $v \in L \cup R$. The variance of a vertex degree d_v is equal to

$$\begin{aligned}
\mathop{\mathbf{E}}_{G \sim \mathbf{G}_\pi}\left[(d_v - \Delta)^2\right] = \mathop{\mathbf{Var}}_{G \sim \mathbf{G}_\pi}[d_v] &= \mathop{\mathbf{Var}}_{G \sim \mathbf{G}_\pi}\left[\sum_{e \in E_v} \mathbb{1}\{e \in E(G)\}\right] \\
&= \sum_{e \in E_v} \mathop{\mathbf{Var}}_{G \sim \mathbf{G}_\pi}[\mathbb{1}\{e \in E(G)\}] = \sum_{e \in E_v} p_e \cdot (1 - p_e) \\
&= \sum_{e \in E_v} p_e - \sum_{e \in E_v} p_e^2 \leq \Delta,
\end{aligned}$$

where the first equality is due to the definition of variance and the fact that $\mathbf{E}[d_v] = \Delta$, the third equality is due to the property of variance and the fact that random variables $\mathbb{1}\{e \in E(G)\}$ are pairwise independent for all $e \in E_v$, and the fourth equality follows since $e \in E(G)$ is a Bernoulli random variable. Therefore,

$$\mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}[f(G)] = \mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}\left[\sum_{v\in L\cup R}(d_v - \Delta)^2\right] = \sum_{v\in L\cup R}\mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}\left[(d_v-\Delta)^2\right] \le 2\cdot n\cdot\Delta. \tag{8}$$

We also observe that if a realized graph $G \sim \mathbf{G}_\pi$ has a small maximum matching $\mu(G) < n$, then many vertex degrees in G must significantly deviate from Δ. To this end, we first establish the following lemma (its proof is omitted).

Lemma 2. *Let d_v be the degree in G of each vertex $v \in L \cup R$. Then, $\forall\delta \ge 0$*

$$\sum_{v\in L\cup R}(d_v-\delta)^2 \ge \frac{2\delta^2(n-\mu(G))^2}{n}.$$

We can now combine Lemma 2 for $\delta = \Delta$ with (8) to get

$$2n\Delta \ge \mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}[f(G)] = \mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}\left[\sum_{v\in L\cup R}(d_v-\Delta)^2\right] \ge \mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}\left[\frac{2\Delta^2(n-\mu(G))^2}{n}\right] \ge \frac{2\Delta^2}{n}\mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}[n-\mu(G)]^2$$

Thus, $\mathop{\mathbf{E}}_{G\sim\mathbf{G}_\pi}[n-\mu(G)] \le \frac{n}{\sqrt{\Delta}}$ and the theorem follows. □

5.1 A Tight Upper Bound

We now show that our bound in Theorem 5 is tight for a wide range of values of parameter Δ (compared to n). We do so using the following stochastic model $\mathbf{G}_\pi$. In our construction we assume that $n - \Delta = \Omega(n)$ and $\Delta \ge 2$ is an integer[5].

1. With probability $1-\alpha$ (where α is a parameter which we will specify later), we select uniformly at random a Δ-regular bipartite graph with $|L| = |R| = n$ vertices. Denote by $\mathbf{D}_{reg}$ the uniform probability distribution over these graphs.
2. With the remaining probability α, we select uniformly at random a subset $A \subset L$ of size $|A| = \frac{n}{2}\left(1 - \frac{1}{c\sqrt{\Delta}}\right)$, where c is the closest to 1 number such that $|A|$ is an integer. Note that c would get arbitrary close to 1 as n goes

[5] Our construction can be extended to cover the case of non-integer Δ with some minor adjustments.

to infinity. Similarly, we select uniformly at random a subset $B \subset R$ of size $|B| = \frac{n}{2}\left(1 + \frac{1}{c\sqrt{\Delta}}\right) = n - |A|$. Next, we describe two distributions $\mathbf{D}(x_1)$ and $\mathbf{D}(x_2)$, each parametrized by a selection probability. In each distribution, we draw edges between the sets A and B and between the sets $L\backslash A$ and $R\backslash B$ i.i.d. with probability x_1 in $\mathbf{D}(x_1)$ and with probability x_2 in $\mathbf{D}(x_2)$. In particular:
(a) with probability 0.5, we generate a bipartite graph $G \sim \mathbf{D}(x_1)$;
(b) with probability 0.5, we generate a bipartite graph $G \sim \mathbf{D}(x_2)$.
We choose x_1 and x_2 so that the expected degree of the graph G drawn from the mixture of $\mathbf{D}(x_1)$ and $\mathbf{D}(x_2)$ is exactly Δ. In particular, we set $x_1 \stackrel{\text{def}}{=} x(1-\delta)$ and $x_2 \stackrel{\text{def}}{=} x(1+\delta)$, where $x \stackrel{\text{def}}{=} \frac{\Delta \cdot n}{2 \cdot |A| \cdot |B|}$ and $\delta^2 \stackrel{\text{def}}{=} \frac{1}{(n-1)c^2\Delta - n}$.

We choose probability

$$\alpha \stackrel{\text{def}}{=} \frac{(n-\Delta)\left((n-1)c^2\Delta - n\right)}{n(n-1)c^2\Delta - n^2 + \Delta n^2 - 2(n-1)c^2\Delta^2}. \tag{9}$$

Theorem 6. *The model $\mathbf{G}_\pi$ is pairwise independent over the set of edges, has probability $p = \frac{\Delta}{n}$ for every edge to be realised, and generates bipartite graphs with a maximum matching of expected size $\mathbf{E}_{G\sim\mathbf{G}_\pi}[\mu(G)] \le n\left(1 - \Omega\left(\frac{1}{\sqrt{\Delta}}\right)\right)$ as long as $n - \Delta = \Omega(n)$ and $\Delta \ge 2$.*

The formal proof of Theorem 6 is omitted due to lack of space.

Acknowledgements. We thank Yuan Zhou for helpful discussions on the expected size of maximum matchings in random bipartite graphs. This work was partially supported by Science and Technology Innovation 2030 - "New Generation of Artificial Intelligence" Major Project No. 2018AAA0100903; COST Action 16228 "European Network for Game Theory"; Innovation Program of Shanghai Municipal Education Commission; Program for Innovative Research Team of Shanghai University of Finance and Economics (IRTSHUFE) and the Fundamental Research Funds for the Central Universities; National Natural Science Foundation of China (Grant No. 61806121); Beijing Outstanding Young Scientist Program (No. BJJWZYJH012019100020098) Intelligent Social Governance Platform; Major Innovation & Planning Interdisciplinary Platform for the "Double-First Class" Initiative, Renmin University of China.

References

1. Babaioff, M., Feldman, M., Gonczarowski, Y.A., Lucier, B., Talgam-Cohen, I.: Escaping cannibalization? Correlation-robust pricing for a unit-demand buyer. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC), p. 191 (2020)
2. Bei, X., Gravin, N., Lu, P., Tang, Z.G.: Correlation-robust analysis of single item auction. In: Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 193–208 (2019)
3. Bollobás, B.: Random Graphs, 2nd edn. Cambridge University Press, Cambridge (2001)

4. Cai, Y., Oikonomou, A.: On simple mechanisms for dependent items. In: Proceedings of the 22nd ACM Conference on Economics and Computation (EC), pp. 242–262 (2021)
5. Carroll, G.: Robustness and separation in multidimensional screening. Econometrica **85**(2), 453–488 (2017)
6. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), pp. 311–320 (2010)
7. Chawla, S., Sivan, B.: Bayesian algorithmic mechanism design. SIGecom Exchanges **13**(1), 5–49 (2014)
8. Correa, J.R., Dütting, P., Fischer, F.A., Schewior, K., Ziliotto, B.: Unknown I.I.D. prophets: better bounds, streaming algorithms, and a new impossibility (extended abstract). In: Proceedings of the 12th Innovations in Theoretical Computer Science Conference (ITCS), pp. 86:1–86:1 (2021)
9. Gravin, N., Lu, P.: Separation in correlation-robust monopolist problem with budget. In: Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA), pp. 2069–2080 (2018)
10. Hartline, J.D.: Mechanism Design and Approximation. Graduate Textbook (in preparation) (2020)
11. Hill, T., Kertz, R.: A survey of prophet inequalities in optimal stopping theory. Contemp. Math. **125**, 191–207 (1992)
12. Immorlica, N., Singla, S., Waggoner, B.: Prophet inequalities with linear correlations and augmentations. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC), pp. 159–185 (2020)
13. Motwani, R., Raghavan, P.: Randomized Algorithms. Cambridge University Press, Cambridge (1995)
14. Roughgarden, T.: Twenty Lectures on Algorithmic Game Theory. Cambridge University Press, Cambridge (2016)
15. Samuel-Cahn, E.: Comparison of threshold stop rules and maximum for independent nonnegative random variables. Ann. Probab. **12**, 1213–1216 (1984)