

# Stochastic Vertex Cover with Few Queries

Soheil Behnezhad\*  
Stanford University

Avrim Blum†  
TTIC

Mahsa Derakhshan\*  
Princeton University

## Abstract

We study the minimum vertex cover problem in the following stochastic setting. Let  $G$  be an arbitrary given graph,  $p \in (0, 1]$  a parameter of the problem, and let  $G_p$  be a random subgraph that includes each edge of  $G$  independently with probability  $p$ . We are unaware of the realization  $G_p$ , but can learn if an edge  $e$  exists in  $G_p$  by querying it. The goal is to find an approximate minimum vertex cover (MVC) of  $G_p$  by querying few edges of  $G$  non-adaptively.

This stochastic setting has been studied extensively for various problems such as minimum spanning trees, matroids, shortest paths, and matchings. To our knowledge, however, no non-trivial bound was known for MVC prior to our work. In this work, we present a:

- $(2 + \varepsilon)$ -approximation for **general** graphs which queries  $O(\frac{1}{\varepsilon^3 p})$  edges per vertex, and a
- 1.367-approximation for **bipartite** graphs which queries  $\text{poly}(1/p)$  edges per vertex.

Additionally, we show that at the expense of a triple-exponential dependence on  $p^{-1}$  in the number of queries, the approximation ratio can be improved down to  $(1 + \varepsilon)$  for bipartite graphs.

Our techniques also lead to improved bounds for bipartite stochastic matching. We obtain a 0.731-approximation with nearly-linear in  $1/p$  per-vertex queries. This is the first result to break the prevalent  $(2/3 \sim 0.66)$ -approximation barrier in the  $\text{poly}(1/p)$  query regime, improving algorithms of [Behnezhad *et al.*, SODA'19] and [Assadi and Bernstein, SODA'19].

---

\*Part of the work was done while the first and the third authors were interns at TTIC.

†This work was supported by the National Science Foundation under grant CCF-1733556.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Our Results . . . . .	4
<b>2</b>	<b>Main Techniques</b>	<b>4</b>
<b>3</b>	<b>Preliminaries &amp; Paper Organization</b>	<b>6</b>
<b>4</b>	<b>Tool I: The Half-Stochastic Matching Lemma</b>	<b>7</b>
4.1	Proofs of Claims 4.4 and 4.5 . . . . .	10
4.2	Proof of Claim 4.6 . . . . .	11
<b>5</b>	<b>Tool II: A New Vertex-Independent Matching Lemma</b>	<b>13</b>
5.1	The Vertex-Independent Matching Algorithm . . . . .	13
5.2	Proving Lemma 5.1 via the Vertex-Independent Lemma 5.2 . . . . .	15
5.2.1	Proof of Claim 5.4 . . . . .	16
<b>6</b>	<b>Upper Bounds</b>	<b>19</b>
6.1	Bipartite Graphs: 1.36-Approximation with $\text{poly}(\frac{1}{p})$ Queries . . . . .	19
6.2	General Graphs: $(2 + \varepsilon)$ -Approximation with $O(\frac{1}{p})$ Per-Vertex Queries . . . . .	22
6.3	Bipartite Graphs: $(1 + \varepsilon)$ -Approximation with $O_p(1)$ Per-Vertex Queries . . . . .	25
<b>7</b>	<b>Lower Bounds</b>	<b>27</b>
7.1	Proof of Theorem 7.1 . . . . .	29
7.2	Proof of Theorem 7.2 . . . . .	30
7.3	Proof of Theorem 7.3 . . . . .	30
7.4	Proof of Theorem 7.4 . . . . .	31
7.5	Proof of Theorem 7.5 . . . . .	33
7.6	Why Random Queries Do Not Work . . . . .	33
<b>A</b>	<b>Deferred Proofs</b>	<b>35</b>
A.1	Polynomial-Time Implementation of Algorithm 2 . . . . .	35
A.2	Proof of Claim 5.8 . . . . .	37

# 1 Introduction

We study the following *stochastic vertex cover* problem. Let  $G = (V, E)$  be a given  $n$ -vertex graph,  $p \in (0, 1]$  a parameter of the problem, and let  $G_p \subseteq G$  be a random subgraph that includes each edge in  $E$  independently with probability  $p$ . We are unaware of the realization of  $G_p$ , but can learn if an edge  $e \in E$  is realized in  $G_p$  by *querying* it. The goal is to find an approximate minimum vertex cover (MVC) of  $G_p$  by querying, non-adaptively, few edges in  $G$ .

This stochastic setting has been studied extensively over the last two decades for various problems such as minimum spanning trees and matroids [14, 15], packing problems [22], shortest paths [21], and most relevant to our work, matchings [9, 10, 2, 3, 8, 22, 5, 1, 7, 6, 4]. There has also been quite a lot of related work on “network reliability” in random subgraphs; see the book [11] for some classic results of the 1980’s as well as [18, 16] and the references therein for more recent works. While this is by no means a comprehensive list of all the related works, we are, to our knowledge, the first to consider a *covering* problem in the setting.

It would be useful to overview the known bounds for the matching problem. It was shown by Blum *et al.* [9, 10] that a  $(1/2 - \varepsilon)$ -approximate matching of  $G_p$  can be found by querying  $O_{p,\varepsilon}(1)$  edges of each vertex in  $G$ , where the dependence on  $p$  was exponential. Later, Assadi, Khanna, and Li [2] improved the dependence on  $p$  and obtained the same approximation with  $\text{poly}(1/\varepsilon p)$  queries. Numerous follow up works [3, 22, 7, 1, 6, 4] then improved the approximation ratio. Particularly, the algorithm of Assadi and Bernstein [1] (see also [7]) obtains a  $(2/3 - \varepsilon)$ -approximation with  $\text{poly}(1/\varepsilon p)$  queries. It was already observed in [2] that  $2/3$ -approximation is a barrier for the problem. Recently, Behnezhad, Derakhshan, and Hajiaghayi [6] broke this barrier and showed that one can obtain a  $(1 - \varepsilon)$ -approximation with  $O_{\varepsilon,p}(1)$  per-vertex queries, where the dependence on  $\varepsilon$  and  $p$  is super-polynomial. Determining the best approximation achievable via  $\text{poly}(1/\varepsilon p)$  queries remains an important open question for the stochastic matching problem.

In light of this progress on the approximate matching problem, it is natural to ask whether the same can also be achieved for the dual minimum vertex cover problem. Particularly,

**Question 1.** *Can we find an approximate MVC of  $G_p$  by querying few, preferably  $\text{poly}(1/p)$ , edges of each vertex in the base graph  $G$ ?*

Observe that a vertex cover of  $G$  is also a valid vertex cover of  $G_p$  since  $G_p \subseteq G$ . However, since some of the edges of  $G$  may not belong to  $G_p$ , the MVC of  $G_p$  might be smaller than that of  $G$ . In fact, the MVC of  $G$  may be as large as  $1/p$  times the MVC of  $G_p$  in expectation — an example is when  $G$  is simply a matching. It turns out that to obtain any constant approximation (independent of  $p$ ),  $\Omega(n/p)$  total queries are necessary; see Theorem 7.5. One may wonder if randomly querying the edges of  $G$  may help. However, we show that the set of queried edges must be picked with much more care — see Section 7.6 for why random queries do not work.

Perhaps the simplest known constant approximation for the MVC problem is through *maximal* matchings. The set of endpoints of the edges in a maximal matching is well-known to form a 2-approximate MVC. For this to hold, however, the *maximality* of the matching is essential, and even a  $(1 - \varepsilon)$ -approximate maximum matching that is not maximal is not useful. Unfortunately, none of the works above on the stochastic matching problem yield a maximal matching of  $G_p$ . In fact, we prove a separation (Theorem 7.2) via a simple lower bound, that, unlike approximate matchings, finding a maximal matching of  $G_p$  requires  $\Omega(n \log n)$  total queries. The situation seems even more complicated on the algorithmic side. In fact, we do not know if a maximal matching of  $G_p$  can be found with  $o(n^2)$  queries (note that the whole graph  $G_p$  can be learned with  $O(n^2)$  queries).

## 1.1 Our Results

In this work, we make progress on Question 1 on several fronts.

Our main end results are in the  $\text{poly}(1/p)$  query regime and we prove the following two results for general and bipartite graphs, respectively.

**Result 1** (see Theorem 6.2). *For any graph  $G$ , any  $\varepsilon > 0$ , and any  $p \in (0, 1]$ , there is a poly-time algorithm that finds a  $(2 + \varepsilon)$ -approximate MVC of  $G_p$  via  $O(\frac{1}{\varepsilon^3 p})$  per-vertex queries.*

**Result 2** (see Theorem 6.1). *For any bipartite graph  $G$ , and any  $p \in (0, 1]$ , there is a poly-time algorithm that finds a  $1.367(\approx \frac{e+1}{e})$ -approximate MVC of  $G_p$  via  $\text{poly}(1/p)$  per-vertex queries.*

As discussed,  $\Omega(1/p)$  per-vertex queries are necessary to obtain any constant approximation; see Theorem 7.5. Therefore, Result 1 is asymptotically query-optimal. Note, on the other hand, that for  $p = 1$ , the problem reduces to the non-stochastic MVC problem. This means that, restricting ourselves to polynomial-time algorithms, the approximation ratio achieved in Result 1 is also optimal for general graphs (up to an additive  $\varepsilon$ ) under the Unique Games Conjecture [19].

For bipartite graphs, however, the UGC based lower bound does not hold. Indeed, an optimal MVC can be found in polynomial time. Result 2 asserts that in the stochastic setting, too, one can get around the 2-approximation barrier with just  $\text{poly}(1/p)$  per-vertex queries.

To prove Result 2 we prove a number of tools (overviewed in Section 2) that, in a sense, give a better understanding of matchings in stochastic graphs. Using these tools, we also obtain the guarantee of Corollary 1.1 for stochastic matching in bipartite graphs, improving the previous close to  $2/3$  approximations of [7, 1] (which respectively obtain  $0.656$  and  $(2/3 - \varepsilon)$ -approximations) in the  $\text{poly}(1/p)$  query regime. We note that Corollary 1.1, importantly, is the first result to break the  $2/3$ -approximation barrier of [2] with just  $\text{poly}(1/p)$  queries.

**Corollary 1.1.** *For any bipartite graph  $G$ , and any  $p \in (0, 1]$ , there is a poly-time algorithm that finds a  $0.731(\approx \frac{e}{e+1})$ -approximate matching of  $G_p$  in expectation via  $O(\frac{\log 1/p}{p})$  per-vertex queries.*

Finally, we turn our attention to the regime where super-polynomial-in- $1/p$  queries per-vertex are allowed. We show that in this setting, the approximation guarantee of Result 2 can be improved all the way to  $(1 + \varepsilon)$ .

**Result 3** (see Theorem 6.3). *For any bipartite graph  $G$ , any  $\varepsilon > 0$ , and any  $p \in (0, 1]$ , there is a poly-time algorithm that finds a  $(1 + \varepsilon)$ -approximate MVC of  $G_p$  via  $O_{\varepsilon,p}(1)$  per-vertex queries.*

We note that the dependence of the number of per-vertex queries in Result 3 on  $p$  is in the order  $\exp(\exp(\exp(\text{poly}(1/p))))$ . It remains an important open problem to determine whether a  $(1 + \varepsilon)$ -approximation for bipartite graphs is achievable via  $\text{poly}(1/\varepsilon p)$  queries. (The same is also open for stochastic matching as discussed.)

## 2 Main Techniques

All of our algorithms in this paper for the stochastic MVC problem return a subset  $C \subseteq V$  which is with probability one a vertex cover of  $G_p$ . That is, all the edges of  $G_p$  have at least one endpoint in  $C$  at all times. Let  $Q$  denote the subset of edges in  $G$  that we query and let  $S$  denote the rest of the edges. Observe that since we are unaware of the realization of edges in  $S$ , we have to cover

them all, no matter which ones are realized. Therefore, once we fix the subgraph  $Q$  to be queried, the “best algorithm” is well-defined: Report a MVC of graph  $H := Q_p \cup S$  which includes the realized edges in  $Q$ , but all the edges in  $S$ .<sup>1</sup> Observe that since we require  $Q$  to be sparse, the vast majority of edges will be in  $S$  and are always assumed to be realized. The challenge is to ensure that the extra covering constraints imposed by these edges do not increase the size of our vertex cover by much, compared to the actual minimum vertex cover of  $G_p$ .

The discussion above actually unveils an interesting connection between the stochastic vertex cover problem and the stochastic matching problem particularly in bipartite graphs where by König’s famous theorem MVC and maximum matching have the same size. On the one hand, the stochastic matching problem asks for a subgraph  $Q$ , such that if we **remove** all the rest of the edges  $S := E \setminus Q$  from  $G_p$ , the size of maximum matching in graph  $G_p \setminus S$  remains close to that of  $G_p$ . On the other hand, the stochastic vertex cover problem asks for a subgraph  $Q$  such that if we **add** all the edges  $S := E \setminus Q$  to  $G_p$ , the size of the maximum matching (which equals the size of MVC) in  $G_p \cup S$  remains close to that of  $G_p$ .

Now let us describe how we actually pick subgraph  $Q$  to query, and how we analyze the size of the minimum vertex cover achieved by querying this subgraph.

**The Half-Stochastic Matching Lemma:** This lemma, which we prove in Section 4, is one of the main components of our paper. It mainly provides a partitioning  $Q, S$  of the edge-set  $E$  of  $G$ . Let  $H := Q_p \cup S$  denote a “half-stochastic” graph which includes each edge of  $Q$  independently with probability  $p$  but includes all the edges of  $S$  with probability 1. We use this partitioning in our algorithm for Result 2 in particular. There, we query only the edges in  $Q$  and report the MVC of graph  $H$ , as outlined before. As a result, the partitioning should clearly ensure  $Q$  has a small maximum degree. In addition, the nice property of this partitioning is that the edges  $e \in S$  have a relatively small probability  $\leq \varepsilon^2 p$  of being part of matching  $\mathcal{M}(H)$ , where  $\mathcal{M}$  is a (near) maximum matching algorithm that is also provided by the lemma. Next, we describe some of the challenges that we face in proving this lemma and how we overcome them.

Let us start with the trivial partitioning  $Q_0 = \emptyset, S_0 = E$  and let  $\mathcal{M}$  be an arbitrary (possibly randomized) maximum matching algorithm. The problem with this solution is that an edge  $e \in S_0$  may have a large probability of being part of  $\mathcal{M}(H_0)$ . This occurs if some edges in  $H_0$  are crucial for the matching to be maximum. We can try to put these edges of  $S_0$  violating the probability constraint in  $Q_0$ , and obtain a new partitioning  $Q_1, S_1$ . The problem, however, is that the corresponding graph  $H_1$  has a different distribution than graph  $H_0$ . Thus, it could be that edges in  $S_1$  that previously had a small probability of joining  $\mathcal{M}(H_0)$  now become crucial for the maximum matching of  $H_1$ . This can in fact continue for a super-constant number of iterations, inevitably violating the maximum degree constraint of  $Q$ .

Instead of using an arbitrary maximum matching algorithm, our first idea is to use a special matching algorithm  $\mathcal{M}$  that maximizes the following objective:

$$\Phi := \sum_{e \in E} (\Pr[e \in \mathcal{M}(H)] - \varepsilon \Pr[e \in \mathcal{M}(H)]^2).$$

The first term in the sum, intuitively, ensures that the size of matchings produced by the algorithm is large. The second term, intuitively, is to ensure the edges tend to have small probabilities of joining  $\mathcal{M}(H)$ . A nice “averaging” property of this objective, is that if some matching algorithm  $\mathcal{M}_1$  guarantees an objective of  $\Phi_1$  and another algorithm  $\mathcal{M}_2$  guarantees  $\Phi_2$ , then the algorithm that with probability  $1/2$  picks the output of  $\mathcal{M}_1$  and with probability  $1/2$  the output of  $\mathcal{M}_2$ , has

<sup>1</sup>Since exact MVC is NP-hard for general graphs, we actually end up using a different algorithm for Result 1.

objective strictly larger than  $\frac{\Phi_1 + \Phi_2}{2}$ , unless the vast majority of edges have the same probability of joining  $\mathcal{M}_1$  and  $\mathcal{M}_2$ .

We plug in this new matching algorithm in the aforementioned framework for obtaining a list of partitionings  $(Q_0, S_0), \dots, (Q_k, S_k)$ . But now, we use the averaging property of our special matching algorithm to argue that we reach our desired partitioning for some  $k = \text{poly}(1/\varepsilon p)$ .

**A New Vertex-Independent Matching Algorithm:** Now suppose that we have the partitioning  $(Q, S)$  provided by the lemma discussed above. How should we argue that the MVC of graph  $H = Q_p \cup S$  approximates the MVC of the actual realization  $G_p$ ? One of the key parts of our analysis, which we discuss thoroughly in Section 5, is a new *vertex-independent matching* (VIM) lemma. VIMs were introduced recently in [6] (further refined in [4]) and were shown to be extremely useful for stochastic matchings. Roughly speaking, given a stochastic graph  $G_p$  and an arbitrary matching algorithm  $\mathcal{A}$ , a VIM algorithm  $\mathcal{B}$  has the property that its matching  $\mathcal{B}(G_p)$  approximates  $\mathcal{A}(G_p)$  (in both the total size and marginal probabilities of edges/vertices joining the matching), but in addition, for “most” vertices  $u$  and  $v$  of the graph, whether or not they are matched in  $\mathcal{B}(G_p)$  are independent events. In the works of [6, 4] for example, this independence is satisfied for  $u, v \in V$  if they are at distance at least  $\text{polylog } \Delta$  in  $G$ , where  $\Delta$  is the maximum degree of  $G$ . This requirement on the distance to achieve independence is provably necessary for the approach taken in [6, 4] which is through distributed local algorithms. In contrast, we use a completely different approach to achieve independence in this work. Our new VIM works for bipartite graphs, but unlike prior works, for any two vertices  $u$  and  $v$  (in different partitions) that are non-adjacent, we have independence. This independence in particular holds, even if  $u$  and  $v$  are connected via a path of length 3. This better guarantee on independence is the key, for example, to why we are allowed to break the RS-barrier for stochastic matchings via  $\text{poly}(1/p)$  queries, whereas the previous approaches [6, 4] required a super-polynomial in  $1/p$  queries. It is also used, crucially, in the analysis of the stochastic MVC algorithm we described above for Result 2.

### 3 Preliminaries & Paper Organization

**Notation.** For any graph  $G$ , we use  $\nu(G)$  to denote the size of the minimum vertex cover of  $G$  and use  $\mu(G)$  to denote the size of the maximum matching of  $G$ . A “fractional matching”  $\mathbf{x}$  of a graph  $G = (V, E)$  is an assignment  $\{x_e\}_{e \in E}$  to the edges, where  $x_e \in [0, 1]$  and for each vertex  $v \in V$ ,  $x_v := \sum_{e \ni v} x_e \leq 1$ . We use  $|\mathbf{x}| := \sum_e x_e$  to denote the size of a fractional matching and for any subset  $E' \subseteq E$ , use  $\mathbf{x}(E')$  to denote  $\sum_{e \in E'} x_e$ . For any integer  $k$  we use  $[k]$  to denote set  $\{1, \dots, k\}$ . We say  $S_1, \dots, S_k$  “partitions” set  $S$  if  $S_1 \cup \dots \cup S_k = S$  and  $S_i \cap S_j = \emptyset$  for all  $i, j \in [k]$ .

As in the literature (see e.g. [9, 6]), we say a (random) matching  $M$  provides an  $0 < \alpha \leq 1$  approximation for the stochastic matching problem if  $M \subseteq G_p$  and  $\mathbb{E}|M| \geq \alpha \cdot \mathbb{E}[\mu(G_p)]$ . We say a (random) subset  $C \subseteq V$  is a  $\beta \geq 1$  approximate stochastic minimum vertex cover, if any edge in  $G_p$  has at least an endpoint in  $C$  with probability 1, and  $\mathbb{E}|C| \leq \beta \cdot \mathbb{E}[\nu(G_p)]$ .

We use the following well-known propositions throughout the paper.

**Proposition 3.1** (König’s Theorem). *In any bipartite graph  $G$ ,  $\nu(G) = \mu(G)$ .*

**Proposition 3.2** (Chebyshev’s inequality). *Let  $X$  be a random variable with finite expected value  $\mathbb{E}[X]$  and finite non-zero variance  $\text{Var}[X]$ . For any  $\lambda > 0$ ,*

$$\Pr[|X - \mathbb{E}X| \geq \lambda] \leq \frac{\text{Var}[X]}{\lambda^2}.$$

**Paper organization.** In Sections 4 and 5 we prove two of the main tools introduced in this paper, particularly the “half-stochastic matching lemma” and the “new vertex-independent matching lemma”. In Section 6 we present our algorithms for the stochastic MVC problem and also the improved result for the stochastic matching problem. Finally, in Section 7 we prove several lower bounds for both the stochastic vertex cover problem and the stochastic matching problem.

## 4 Tool I: The Half-Stochastic Matching Lemma

The Half-Stochastic Matching Lemma, constructively, gives a partitioning  $(Q, S)$  of the edge-set  $E$  of graph  $G$ . This partitioning is accompanied with a special near-maximum matching algorithm  $\mathcal{M}$  that operates on the “half-stochastic” random subgraph  $H$  of  $G$  in which each edge of  $Q$  is stochastic (i.e. realized with probability  $p$ ) and each edge of  $S$  appears with probability one.

Although this lemma seems to be about matchings only, it actually plays an important role in the stochastic vertex cover algorithm for Theorem 6.1 in both the algorithm in deciding which edges to query, and the analysis of the approximation ratio achieved by this algorithm.

**Lemma 4.1** (Half-Stochastic Matching Lemma). *Let  $G = (V, E)$  be a (possibly non-bipartite) graph and let  $\varepsilon \in [0, 1]$  and  $p \in [0, 1]$  be two given parameters. There is a partitioning of  $E$  into subsets  $Q$  and  $S$  such that:*

(i) *The maximum degree in  $Q$  is  $O(\frac{1}{\varepsilon^{11p^6}})$ .*

*Moreover,  $Q$  and  $S$  are such that there exists a randomized matching algorithm  $\mathcal{M}$ , where by letting  $H := Q_p \cup S$  denote a random graph which includes all edges in  $S$ , but includes each edge of  $Q$  independently with probability  $p$ , we get:*

(ii)  *$\mathbb{E}|\mathcal{M}(H)| \geq (1 - 2\varepsilon) \cdot \mathbb{E}[\mu(H)]$ . That is, the matching algorithm  $\mathcal{M}$  should find an approximate maximum matching of  $H$  in expectation.*

(iii) *For any edge  $e \in S$ ,  $\Pr[e \in \mathcal{M}(H)] \leq \varepsilon^2 p$ .*

*We note that the probabilistic statements above are with respect to both the inherent randomization in graph  $H$ , and also the randomization used in algorithm  $\mathcal{M}$ .*

We now turn to present the proof of Lemma 4.1. To do so, we have to output a triplet  $(Q, S, \mathcal{M})$  where  $Q$  and  $S$  form a partitioning of  $E$ , and  $\mathcal{M}$  is a matching algorithm.

For some  $k = O_{\varepsilon, p}(1)$  which we specify later, we construct a list  $(Q_0, S_0), (Q_1, S_1), \dots, (Q_k, S_k)$  of partitionings. We will argue that if  $k$  is large enough, there is one of the partitions,  $(Q_i, S_i)$ , which satisfies all the properties required by Lemma 4.1. This will be our partitioning  $(Q, S)$ .

**The construction.** The base partitioning is simply  $Q_0 = \emptyset, S_0 = E$  and for each  $i \geq 0$ ,  $(Q_{i+1}, S_{i+1})$  is constructed from the previous partitioning  $(Q_i, S_i)$ . To describe the construction, let us define for each partitioning  $i$  a random graph  $H_i$  which includes each edge of  $Q_i$  independently with probability  $p$ , and includes every edge of  $S_i$  with probability 1. We will also soon formalize a special randomized matching algorithm  $\mathcal{M}_i$  that we run on graph  $H_i$ . Having  $\mathcal{M}_i$ , for each  $i \geq 0$  we define

$$D_i := \{e \in S_i \mid \Pr[e \in \mathcal{M}_i(H_i)] > \varepsilon^2 p\}, \quad (1)$$



and construct  $(Q_{i+1}, S_{i+1})$  in the following way:

$$Q_{i+1} \leftarrow Q_i \cup D_i \quad \text{and} \quad S_{i+1} \leftarrow S_i \setminus D_i. \quad (2)$$

Observe that in the construction above, each partitioning  $(Q_{i+1}, S_{i+1})$  is obtained from the previous partitioning  $(Q_i, S_i)$  by “moving” the edges of  $D_i$  (which are all by definition in  $S_i$ ) from  $S_i$  to partition  $Q_i$ . For this reason, we have

$$Q_0 \subset Q_1 \subset \dots \subset Q_k \quad \text{and} \quad S_0 \supset S_1 \supset \dots \supset S_k.$$

Let us emphasize that graphs  $H_0, H_1, \dots, H_k$  have **different** distributions. Intuitively, since  $Q_i$  grows as  $i$  increases, and graph  $H_i$  includes only  $p$  fraction of the edges in  $Q_i$  but all the edges in  $S_i$ , random graph  $H_i$  tends to get smaller and smaller by increasing  $i$ . In fact it would be useful to consider a coupling  $(H_0, H_1, \dots, H_k)$  as follows: On each edge  $e$  we draw an independent  $p$ -Bernoulli random variable  $X_e$  and use this to define  $H_i$  for all  $i$  as:

$$H_i = \{e \mid (e \in Q_i \text{ and } X_e = 1) \text{ or } (e \in S_i)\}. \quad (3)$$

This coupling is useful because in each outcome of the joint distribution  $(H_0, \dots, H_k)$ , each graph  $H_i$  is a subgraph of the previous graph  $H_{i-1}$ .

Let us now finalize our construction by formalizing algorithm  $\mathcal{M}_i$ .

**The Matching  $\mathcal{M}_i$ :** For any  $i$  and any (possibly randomized) matching algorithm  $\mathcal{M}'$ , we define

$$\Phi_i(\mathcal{M}') := \sum_{e \in E} \left( \Pr_{\mathcal{M}', H_i} [e \in \mathcal{M}'(H_i)] - \varepsilon \Pr_{\mathcal{M}', H_i} [e \in \mathcal{M}'(H_i)]^2 \right), \quad (4)$$

where let us emphasize that the probabilities are taken over both the possible randomization in algorithm  $\mathcal{M}'$ , and the randomization in graph  $H_i$  (regarding the realization of edges belonging to  $Q_i$ ). Having this definition, we now simply let  $\mathcal{M}_i$  be the algorithm maximizing  $\Phi_i$ , i.e.:

$$\mathcal{M}_i := \arg \max_{\mathcal{M}'} \Phi_i(\mathcal{M}'),$$

and we use  $\Phi_i := \Phi_i(\mathcal{M}_i)$  to simply denote the optimal objective value obtained by this algorithm.

Let us, for now, not concern ourselves with how the objective function (4) can be maximized in polynomial time and assume that this algorithm  $\mathcal{M}_i$  is simply given. We will later address this issue and obtain a polynomial-time algorithm once it becomes clear how we use  $\mathcal{M}_i$ .

**The intuition behind objective (4).** To see the intuition behind why we define the objective function (4) this way, first note by linearity of expectation that

$$\Phi_i = \mathbb{E}|\mathcal{M}_i(H_i)| - \varepsilon \sum_{e \in E} \Pr[e \in \mathcal{M}_i(H_i)]^2. \quad (5)$$

The intuition behind the first term is clear: We want the expected size of the matching to be large; Observation 4.3 below formalizes this by showing that the matching algorithm  $\mathcal{M}_i$  maximizing  $\Phi_i$  must be a  $(1 - \varepsilon)$ -approximate matching. The second term, on the other hand, ensures that the marginal probabilities of edges appearing in the matching tend to be small. This is useful for Property (iii) of Lemma 4.1 which requires small marginals for all edges in  $S$ .

**Observation 4.2.** For any  $i$ ,  $\mathbb{E}|\mathcal{M}_i(H_i)| \geq \Phi_i$ .



*Proof.* By Equation (5),  $\mathbb{E}|\mathcal{M}_i(H_i)| = \Phi_i + \varepsilon \sum_{e \in E} \Pr[e \in \mathcal{M}_i(H_i)]^2 \geq \Phi_i$ . □

**Observation 4.3.**  $\mathbb{E}|\mathcal{M}_i(H_i)| \geq (1 - \varepsilon) \cdot \mathbb{E}[\mu(H_i)]$ .

*Proof.* Consider a deterministic algorithm  $\mathcal{M}'$  that picks an arbitrary maximum matching  $M$  of its random input  $H_i$ . Since  $\mathbb{E}|M| = \mathbb{E}[\mu(H_i)]$ , we have

$$\Phi_i(\mathcal{M}') = \mathbb{E}[\mu(H_i)] - \varepsilon \sum_{e \in E} \Pr[e \in M]^2 \geq \mathbb{E}[\mu(H_i)] - \varepsilon \sum_{e \in E} \Pr[e \in M] = (1 - \varepsilon) \cdot \mathbb{E}[\mu(H_i)].$$

Since  $\mathcal{M}_i$  maximizes  $\Phi_i = \Phi_i(\mathcal{M}_i)$ , we get  $\Phi_i \geq \Phi_i(\mathcal{M}') \geq (1 - \varepsilon) \cdot \mathbb{E}[\mu(H_i)]$ . Combined with  $\mathbb{E}|\mathcal{M}_i(H_i)| \geq \Phi_i$  due to Observation 4.2, this implies  $\mathbb{E}|\mathcal{M}_i(H_i)| \geq (1 - \varepsilon) \cdot \mathbb{E}[\mu(H_i)]$ . □

We now turn to prove that one of the partitionings  $(Q_i, S_i)$  must satisfy the properties required by Lemma 4.1. The next set of claims are used for this purpose.

**Claim 4.4.** *It holds that*

$$\mu(G) \geq \Phi_0 \geq \Phi_1 \geq \dots \geq \Phi_k \geq (1 - \varepsilon) \cdot p \cdot \mu(G).$$

**Claim 4.5.** *There is an interval  $I = \{s, \dots, s + k'\}$  in  $[k]$  with  $|I| \geq \frac{1}{300}\varepsilon^6 p^3 k$  such that*

$$|\Phi_i - \Phi_j| \leq 0.01\varepsilon^6 p^3 \mu(G) \quad \text{for all } i < j \text{ in } I.$$

**Claim 4.6.** *Let  $I$  be as defined in Claim 4.5. Either there is some  $i \in I$  where  $\mathbb{E}|D_i \cap \mathcal{M}_i(H_i)| < \varepsilon p \mu(G)$ , or otherwise for any  $i, j \in I$  with  $i < j$ , it holds that*

$$\mathbb{E}|D_i \cap \mathcal{M}_j(H_j)| \geq 0.25\varepsilon^3 p^2 \cdot \mu(G).$$

Claims 4.4 and 4.5 are proved in Section 4.1 and Claim 4.6 is proved in Section 4.2. Claim 4.6 is, in particular, the key part of the proof. It is proved by showing that if the condition of Claim 4.6 is not satisfied, then the randomized matching algorithm that with probability 0.5 picks the output of  $\mathcal{M}_j(H_j)$  and otherwise the output of  $\mathcal{M}_i(H_i)$ , should obtain a larger objective than  $\Phi_i$  which we show is a contradiction.

Having proved these claims, we now turn to prove Lemma 4.1.

*Proof of Lemma 4.1.* First, we set  $k = 5 \cdot \frac{300}{\varepsilon^9 p^5}$ . Since for any  $i$ , each edge in  $D_i$  has probability at least  $\varepsilon^2 p$  of being in matching  $\mathcal{M}_i(H_i)$ , and that the probabilities around each vertex sum up to at most one, there are at most  $1/\varepsilon^2 p$  edges connected to each vertex in  $D_i$ . This implies that for any  $j \in [k]$ ,  $Q_j$  has maximum degree at most  $k \cdot \frac{1}{\varepsilon^2 p} = O(\frac{1}{\varepsilon^{11} p^6})$ , satisfying Property (i). Now we prove there exists some  $(Q_i, S_i)$  satisfying Properties (ii) and (iii) as well.

Let  $I$  be as provided by Claim 4.5. There are two possible cases:

**Case 1** — There is some  $i \in I$  where  $\mathbb{E}|D_i \cap \mathcal{M}_i(H_i)| < \varepsilon p \mu(G)$ :

In this case, we can let  $Q \leftarrow Q_i, S \leftarrow S_i$ , which implies graph  $H$  of Lemma 4.1 has the same distribution as  $H_i$ . We now let matching algorithm  $\mathcal{M}$ , required by Lemma 4.1, to be the same as matching algorithm  $\mathcal{M}_i$ , except that we exclude the edges of  $D_i$  from the matching. That is, we let  $\mathcal{M}(H) = \mathcal{M}_i(H_i) \setminus D_i$ .

Since we exclude the edges in  $D_i$  from the matching, we get that for all edges  $e \in S$ ,  $\Pr[e \in \mathcal{M}(H)] \leq \varepsilon^2 p$  satisfying Property (iii). On the other hand,

$$\mathbb{E}|\mathcal{M}(H)| = \mathbb{E}|\mathcal{M}_i(H_i) \setminus D_i| = \mathbb{E}|\mathcal{M}_i(H_i)| - \mathbb{E}|D_i \cap \mathcal{M}_i(H_i)|$$

$$\begin{aligned}
&> \mathbb{E}|\mathcal{M}_i(H_i)| - \varepsilon p \mu(G) && \text{(By the assumption of Case 1.)} \\
&\geq (1 - \varepsilon) \mathbb{E}[\mu(H_i)] - \varepsilon p \mu(G) && \text{(By Observation 4.3.)} \\
&\geq (1 - 2\varepsilon) \mathbb{E}[\mu(H_i)] && \text{(Since } \mathbb{E}[\mu(H_i)] \geq p \mu(G).) \\
&= (1 - 2\varepsilon) \mathbb{E}[\mu(H)]. && \text{(Since } H \text{ and } H_i \text{ have the same distribution.)}
\end{aligned}$$

This proves  $\mathcal{M}$  is a  $(1 - 2\varepsilon)$ -approximate matching algorithm, satisfying Property (ii).

**Case 2** — For all  $i \in I$ ,  $\mathbb{E}|D_i \cap \mathcal{M}_i(H_i)| \geq \varepsilon p \mu(G)$ :

In this case, by Claim 4.6, we have  $\mathbb{E}|D_i \cap \mathcal{M}_j(H_j)| \geq 0.25\varepsilon^3 p^2 \cdot \mu(G)$  for all  $i < j$  in  $I$ . Let us denote  $I = \{a_1, a_2, \dots, a_\ell\}$  where  $a_1 < \dots < a_\ell$ . Letting  $j = a_\ell$ , we thus get

$$\mathbb{E}|D_{a_i} \cap \mathcal{M}_{a_\ell}(H_{a_\ell})| \geq 0.25\varepsilon^3 p^2 \cdot \mu(G) \quad \text{for all } i \in \{1, \dots, \ell - 1\}.$$

On the other hand, observe from construction (2) that sets  $D_{a_1}, \dots, D_{a_{\ell-1}}$  are all pairwise disjoint. This implies

$$\mathbb{E}|\mathcal{M}_{a_\ell}(H_{a_\ell})| \geq \sum_{i=1}^{\ell-1} \mathbb{E}|D_{a_i} \cap \mathcal{M}_{a_\ell}(H_{a_\ell})| \geq (\ell - 1) \cdot 0.25\varepsilon^3 p^2 \cdot \mu(G). \quad (6)$$

Recall from Claim 4.5 that  $\ell = |I| \geq \frac{1}{300}\varepsilon^6 p^3 k$ . Since we set  $k = 5 \cdot \frac{300}{\varepsilon^9 p^5}$ , we get  $\ell \geq \frac{5}{\varepsilon^3 p^2}$  which combined with (6) implies  $\mathbb{E}|\mathcal{M}_{a_\ell}(H_{a_\ell})| > \mu(G)$  which is a contradiction since  $H_{a_\ell}$  is a subgraph of  $G$  and cannot have a larger matching than  $\mu(G)$ . This contradiction implies that if we set  $k$  large enough, this second case essentially does not happen. As a result, we always end up at Case 1, which we just showed how it proves Lemma 4.1.

The proof of Lemma 4.1 is thus complete.  $\square$

Finally, we remark that our techniques also lead to a partitioning with the same guarantee as in Lemma 4.1 that can be found in polynomial time. We defer the details of this polynomial-time implementation to Appendix A.1.

## 4.1 Proofs of Claims 4.4 and 4.5

*Proof of Claim 4.4.* As discussed, in the coupling of Eq 3,  $H_i$  is always a subgraph of  $H_{i-1}$ . As a result, matching algorithm  $\mathcal{M}_i$  is also applicable on graph  $H_{i-1}$ , implying that  $\Phi_{i-1} \geq \Phi_i$  for all  $i$ .

To see why  $\mu(G) \geq \Phi_0$ , note from Observation 4.2 that  $\mathbb{E}|\mathcal{M}_0(H_0)| \geq \Phi_0$ . On the other hand, no matter what matching algorithm we use for  $\mathcal{M}_0$ , we have  $\mathbb{E}|\mathcal{M}_0(H_0)| \leq \mu(G)$  as the output must be a matching in  $H_0$  and thus  $G$ . Combining the two bounds gives  $\mu(G) \geq \Phi_0$ .

Finally, to see why  $\Phi_k \geq (1 - \varepsilon) \cdot p \cdot \mu(G)$ , fix a maximum matching  $M$  of  $G$  which has to have size  $\mu(G)$ . Every edge  $e \in M$  either is in  $Q_k$  or  $S_k$ ; in either case,  $e \in H_k$  with probability at least  $p$ . We thus have  $\mathbb{E}[\mu(H_k)] \geq \mathbb{E}|M \cap H_k| \geq p|M| = p \cdot \mu(G)$ . Now consider a choice for  $\mathcal{M}_k$  which deterministically picks a maximum matching  $M_k$  of  $H_k$ . This proves

$$\begin{aligned}
\Phi_k &\geq \sum_{e \in E} \Pr[e \in M_k] - \varepsilon \Pr[e \in M_k]^2 \geq \sum_{e \in E} \Pr[e \in M_k] - \varepsilon \Pr[e \in M_k] \\
&= (1 - \varepsilon) \sum_{e \in E} \Pr[e \in M_k] \geq (1 - \varepsilon) \cdot \mathbb{E}|M_k| = (1 - \varepsilon) \cdot \mathbb{E}[\mu(H_k)] \geq (1 - \varepsilon) \cdot p \cdot \mu(G),
\end{aligned}$$

completing the proof.  $\square$

*Proof of Claim 4.5.* Let us define  $I_j$  for any  $j \in \{1, \dots, \lceil 100/\varepsilon^6 p^3 \rceil\}$  as follows

$$I_j := \{i : (1 - 0.01j\varepsilon^6 p^3)\mu(G) < \Phi_i \leq (1 - 0.01(j-1)\varepsilon^6 p^3)\mu(G)\}.$$

Recall from Claim 4.4 that  $\mu(G) = \Phi_0 \geq \Phi_1 \geq \dots \geq \Phi_k \geq (1-\varepsilon)p\mu(G)$ . Thus,  $I_1, I_2, \dots$  partition  $[k]$  into consecutive intervals where for all elements  $i, j$  in the same interval,  $|\Phi_i - \Phi_j| \leq 0.01\varepsilon^6 p^3 \mu(G)$ .

Since there are only  $\lceil 100/\varepsilon^6 p^3 \rceil$  intervals and  $\sum_j |I_j| = k$  (as every  $i \in [k]$  belongs to exactly one of the intervals) there is at least one interval  $I$  with  $|I| \geq \frac{k}{\lceil 100/\varepsilon^6 p^3 \rceil + 1} > \frac{1}{300} \cdot \varepsilon^6 p^3 k$ . This interval  $I$  satisfies the required property of the claim by its definition, and has the desired size.  $\square$

## 4.2 Proof of Claim 4.6

*Proof of Claim 4.6.* Suppose for the sake of contradiction that

$$\mathbb{E}|D_i \cap \mathcal{M}_i(H_i)| \geq \varepsilon p \mu(G) \quad \text{for all } i \in I, \quad (7)$$

and that there are  $i, j \in I$  such that  $i < j$  and

$$\mathbb{E}|D_i \cap \mathcal{M}_j(H_j)| < 0.25\varepsilon^3 p^2 \cdot \mu(G). \quad (8)$$

Consider a matching  $\mathcal{M}_{i,j}(H_i)$  which with probability  $1/2$  returns the output of  $\mathcal{M}_i(H_i)$  and with probability  $1/2$  returns the output of  $\mathcal{M}_j(H_j)$ . Since  $i < j$ , by the coupling (3),  $H_j$  is a subgraph of  $H_i$  and thus any matching in  $H_j$  is a matching in  $H_i$ . As a result,  $\mathcal{M}_{i,j}$  is a valid matching algorithm for graph  $H_i$ . We prove that under (7) and (8), algorithm  $\mathcal{M}_{i,j}$  should satisfy  $\Phi_i(\mathcal{M}_{i,j}) > \Phi_i(\mathcal{M}_i)$  which contradicts the assumption that  $\mathcal{M}_i$  maximizes  $\Phi_i(\mathcal{M}_i)$ .

From the definition of objective  $\Phi_i(\mathcal{M}_{i,j})$  we have

$$\begin{aligned} \Phi_i(\mathcal{M}_{i,j}) &= \sum_{e \in E} \Pr[e \in \mathcal{M}_{i,j}(H_i)] - \varepsilon \Pr[e \in \mathcal{M}_{i,j}(H_i)]^2 \\ &= \sum_{e \in E} \left( \frac{\Pr[e \in \mathcal{M}_i(H_i)] + \Pr[e \in \mathcal{M}_j(H_j)]}{2} \right) - \varepsilon \left( \frac{\Pr[e \in \mathcal{M}_i(H_i)] + \Pr[e \in \mathcal{M}_j(H_j)]}{2} \right)^2. \end{aligned}$$

Let us for simplicity of notation use  $p_i(e) := \Pr[e \in \mathcal{M}_i(H_i)]$  and  $p_j(e) := \Pr[e \in \mathcal{M}_j(H_j)]$ . The equality above therefore can be expressed as

$$\Phi_i(\mathcal{M}_{i,j}) = \sum_{e \in E} \left( \frac{p_i(e) + p_j(e)}{2} \right) - \varepsilon \left( \frac{p_i(e) + p_j(e)}{2} \right)^2. \quad (9)$$

Basic mathematical calculations give that for any  $e$ ,

$$\begin{aligned} \left( \frac{p_i(e) + p_j(e)}{2} \right)^2 &= \frac{p_i(e)^2 + p_j(e)^2 + 2p_i(e)p_j(e)}{4} \\ &= \frac{p_i(e)^2}{2} + \frac{p_j(e)^2}{2} - \frac{p_i(e)^2}{4} - \frac{p_j(e)^2}{4} + \frac{2p_i(e)p_j(e)}{4} \\ &= \frac{p_i(e)^2}{2} + \frac{p_j(e)^2}{2} - \left( \frac{p_i(e) - p_j(e)}{2} \right)^2. \end{aligned} \quad (10)$$

Replacing (10) back into (9) gives

$$\Phi_i(\mathcal{M}_{i,j}) = \sum_{e \in E} \left( \frac{p_i(e) + p_j(e)}{2} \right) - \varepsilon \left( \frac{p_i(e)^2}{2} + \frac{p_j(e)^2}{2} - \left( \frac{p_i(e) - p_j(e)}{2} \right)^2 \right)$$

$$\begin{aligned}
&= \sum_{e \in E} \left( \frac{p_i(e) - \varepsilon p_i(e)^2}{2} \right) + \left( \frac{p_j(e) - \varepsilon p_j(e)^2}{2} \right) + \varepsilon \left( \frac{p_i(e) - p_j(e)}{2} \right)^2. \\
&\quad \text{(By simply moving the terms in the previous line.)} \\
&= \frac{\Phi_i}{2} + \frac{\Phi_j}{2} + \varepsilon \sum_{e \in E} \left( \frac{p_i(e) - p_j(e)}{2} \right)^2 \quad \text{(See below.)} \\
&\geq \Phi_i - 0.01\varepsilon^6 p^3 \mu(G) + \varepsilon \sum_{e \in E} \left( \frac{p_i(e) - p_j(e)}{2} \right)^2 \\
&\quad \text{(Since } |\Phi_i - \Phi_j| \leq 0.01\varepsilon^6 p^3 \mu(G) \text{ by Claim 4.5.)}
\end{aligned}$$

The third equality above, simply comes from the definition (4) for  $\Phi_i$ , which implies  $\Phi_i = \sum_e \Pr[e \in \mathcal{M}_i(H_i)] - \varepsilon \Pr[e \in \mathcal{M}_i(H_i)]^2 = \sum_e p_i(e) - \varepsilon p_i(e)^2$ , and from the same bound applied on  $\Phi_j$ .

Now define subset  $D'_i := \{e \in D_i \mid p_j(e) < 0.5\varepsilon^2 p\}$  of  $D_i$ . Using this subset only instead of the set  $E$  of edges in the inequality above gives

$$\begin{aligned}
\Phi_i(\mathcal{M}_{i,j}) &\geq \Phi_i - 0.01\varepsilon^6 p^3 \mu(G) + \varepsilon \sum_{e \in D'_i} \left( \frac{p_i(e) - p_j(e)}{2} \right)^2 \\
&\geq \Phi_i - 0.01\varepsilon^6 p^3 \mu(G) + \varepsilon \sum_{e \in D'_i} \left( \frac{\varepsilon^2 p - 0.5\varepsilon^2 p}{2} \right)^2 \\
&\quad \text{(Since for any } e \in D'_i, p_j(e) < 0.5\varepsilon^2 p \text{ by definition of } D'_i \text{ and } p_i(e) \geq \varepsilon^2 p \text{ by definition of } D_i.) \\
&= \Phi_i - 0.01\varepsilon^6 p^3 \mu(G) + \varepsilon \sum_{e \in D'_i} \frac{\varepsilon^4 p^2}{16} \\
&= \Phi_i - 0.01\varepsilon^6 p^3 \mu(G) + \frac{\varepsilon^5 p^2}{16} |D'_i|.
\end{aligned}$$

To obtain the claimed contradiction, we will prove that

$$|D'_i| \geq 0.5\varepsilon p \mu(G), \quad (11)$$

which combined by inequality above proves

$$\Phi_i(\mathcal{M}_{i,j}) \geq \Phi_i - 0.01\varepsilon^6 p^3 \mu(G) + \frac{\varepsilon^6 p^3}{32} \mu(G) > \Phi_i + 0.01\varepsilon^6 p^3 \mu(G), \quad (12)$$

which contradicts  $\Phi_i$  being the maximum objective achievable.

To complete the proof, it thus only remains to prove (11). We have

$$\mathbb{E}|D_i \cap \mathcal{M}_j(H_j)| = \sum_{e \in D_i} p_j(e) = \sum_{e \in D_i \setminus D'_i} p_j(e) + \sum_{e \in D'_i} p_j(e) \geq \sum_{e \in D_i \setminus D'_i} p_j(e) \geq (|D_i| - |D'_i|)0.5\varepsilon^2 p.$$

Also note that

$$\varepsilon p \mu(G) \stackrel{(7)}{\leq} \mathbb{E}|D_i \cap \mathcal{M}_i(H_i)| = \sum_{e \in D_i} p_i(e) \leq \sum_{e \in D_i} 1 = |D_i|.$$

Combining the two bounds above gives

$$\mathbb{E}|D_i \cap \mathcal{M}_j(H_j)| \geq (\varepsilon p \mu(G) - |D'_i|)0.5\varepsilon^2 p = 0.5\varepsilon^3 p^2 \mu(G) - 0.5\varepsilon^2 p |D'_i|,$$

which combined with the bound  $\mathbb{E}|D_i \cap \mathcal{M}_j(H_j)| < 0.25\varepsilon^3 p^2 \cdot \mu(G)$  of (8) gives

$$0.25\varepsilon^3 p^2 \mu(G) > 0.5\varepsilon^3 p^2 \mu(G) - 0.5\varepsilon^2 p |D'_i|.$$

By moving the terms, we get

$$|D'_i| > \frac{0.25\varepsilon^3 p^2 \mu(G)}{0.5\varepsilon^2 p} = 0.5\varepsilon p \mu(G).$$

This is the desired bound of inequality (11), which as discussed above completes the proof.  $\square$

## 5 Tool II: A New Vertex-Independent Matching Lemma

The notion of “vertex-independent matchings” for stochastic graphs was introduced first in [6]. In this section, we present a new vertex-independent matching lemma for bipartite graphs, that unlike the previous ones [6, 4], which required the vertices to be far apart in the graph to have independence, guarantees independence for any pair of non-adjacent nodes, even if there is a short path of length 3 between them. This stronger guarantee on the independence is the key to improve per-vertex queries from  $O_p(1)$  down to  $\text{poly}(1/p)$ .

We use the vertex-independent lemma to prove the following which use for our bipartite MVC approximate algorithm.

**Lemma 5.1.** *Let  $G = (V, E)$  be a bipartite graph, let realization  $G_p = (V, E_p)$  be a random subgraph of  $G$  that includes each of its edges independently with probability  $p$ . Let  $(Q, S)$  be a partitioning of  $E$  and denote  $Q_p := Q \cap E_p$  and  $S_p := S \cap E_p$ . Suppose also that we are given a (possibly randomized) matching algorithm  $\mathcal{M}$ , and a fractional matching  $\mathbf{q}$  on  $E$  such that:*

1. *For any edge  $e \in Q$ ,  $q_e = \Pr_{Q_p, \mathcal{M}}[e \in \mathcal{M}(Q_p)]$ .*
2. *For any edge  $e \in S$ ,  $q_e \leq \varepsilon^5 p$  for some  $\varepsilon > 0$ .*

*Then  $\mathbb{E}[\mu(G_p)] \geq (1 - 6\varepsilon) \frac{e}{e+1} |\mathbf{q}|$ .*

We first present the vertex-independent matching algorithm in Section 5.1 and use it to prove Lemma 5.1 in Section 5.2.

### 5.1 The Vertex-Independent Matching Algorithm

In this section we present our vertex-independent matching algorithm which satisfies the following:

**Lemma 5.2 (Bipartite Vertex-Independent Matching Lemma).** *Let  $\Gamma(A, B, E_\Gamma)$  be a bipartite graph, let  $\Gamma_p$  be a random subgraph of  $\Gamma$  that contains any of its edges independently with probability  $p$ , let  $\mathcal{A}$  be an arbitrary matching algorithm, possibly randomized, and let  $M_{\mathcal{A}}$  be the matching obtained by running  $\mathcal{A}$  on  $\Gamma_p$ . There is a randomized algorithm (Algorithm 1) for constructing a matching  $M_{\mathcal{B}}$  of  $\Gamma_p$  such that:*

- (i)  $\mathbb{E}|M_{\mathcal{B}}| \geq (1 - \frac{1}{e}) \cdot \mathbb{E}|M_{\mathcal{A}}|$ .
- (ii) *For any vertex  $v \in A$ ,  $\Pr[v \in M_{\mathcal{B}}] \leq \Pr[v \text{ proposes}] = \Pr[v \in M_{\mathcal{A}}]$ .  
(See Algorithm 1 for how the vertices on the  $A$  side “propose”.)*
- (iii) *For any vertex  $u \in B$ ,  $\Pr[u \in M_{\mathcal{B}}] \leq \Pr[u \in M_{\mathcal{A}}]$ .*

- (iv) For any non-adjacent  $v \in A, u \in B$  (i.e.  $(u, v) \notin E$ ), whether  $v$  proposes (see Algorithm 1) is independent of event  $u \in M_B$ .

We emphasize that  $M_A$  (resp.  $M_B$ ) has two sources of randomization, one in the randomization of graph  $\Gamma_p$ , and one the possible randomization in algorithm  $\mathcal{A}$  (resp.  $\mathcal{B}$ ). The probabilistic statements above are with regards to both.

*Proof.* We start by describing the algorithm for constructing matching  $M_B$ .

For any vertex  $v \in A$ , let us use  $R_v$  to denote the realization status of edges connected to  $v$  in graph  $\Gamma_p$ . That is,  $R_v$  reveals which edges connected to  $v$  are realized and which ones are not realized, but crucially does not reveal any information about the realization of the rest of the edges. Using this information, for any edge  $e = (v, u)$  with  $v \in A, u \in B$ , we define  $p_e := \Pr[e \in M_A \mid R_v]$ . That is, in defining  $p_e$  for any edge  $e = (v, u)$  we only need to know which edges connected to  $v$  are realized in  $\Gamma_p$ , and are essentially unaware of realization of the rest of the edges in  $\Gamma_p$ . Similarly, for any vertex  $v \in A$  we denote  $p_v := \sum_{e \ni v} p_e = \Pr[v \in M_A \mid R_v]$ .

Observe that for any  $v \in A$ ,  $0 \leq p_v \leq 1$  since  $p_v$  corresponds to the probability that  $v$  is matched in  $M_A$  conditioned on the realization of its edges. Importantly, however, this does not hold for vertices of the other partition, and  $\sum_{e \ni u} p_e$  may, in fact, exceed one for  $u \in B$ .

Having defined  $p_e$  and  $p_v$  as above, the claimed Algorithm 1 in Lemma 5.2 can be formalized:

**Algorithm 1.** The algorithm for constructing matching  $M_B$  on  $\Gamma_p$ .

```

1 for any vertex  $v \in A$  do
2   Vertex  $v$  either proposes to exactly one neighbor  $u$ , or does not propose at all. This is
   decided by a random procedure, where each neighbor  $u$  has probability exactly  $p_{(u,v)}$ 
   of being proposed to by  $v$ , and there is a probability  $1 - p_v$  that  $v$  does not propose.
3 for any vertex  $u \in B$  do
4   Among the vertices who sent proposals to vertex  $u$ , if any,  $u$  chooses an arbitrary
   winner  $v$  and we add  $(u, v)$  to matching  $M'_B$ .
5 return  $M_B$ .

```

We now prove the properties of Lemma 5.2.

**Property (i).** Fix an arbitrary vertex  $u \in B$ , and let  $\{v_1, \dots, v_d\}$  be its neighbors in  $A$ . Let  $Y_i$  be the indicator random variable for the event that  $v_i$  proposes to  $u$ . First, observe that

$$\mathbb{E}[Y_i] = \mathbb{E}[p_{(v_i, u)}] = \mathbb{E}[\Pr[(u, v_i) \in \mathcal{A}(\Gamma_p) \mid R_{v_i}]] = \Pr[(u, v_i) \in \mathcal{A}(\Gamma_p)].$$

As a result,

$$\sum_i \mathbb{E}[Y_i] = \sum_i \Pr[(u, v_i) \in \mathcal{A}(\Gamma_p)] = \Pr[u \in \mathcal{A}(\Gamma_p)]. \quad (13)$$

Moreover, observe that  $R_{v_1}, \dots, R_{v_d}$  are mutually independent since the edges of  $v_1, \dots, v_d$  are all disjoint. Thus, the proposals of  $v_1, \dots, v_d$  are also mutually independent and so are random variables  $Y_1, \dots, Y_d$ . Now observe that  $u$  remains unmatched in  $M_B$  if and only if none of its neighbors proposes to it; combined with the independence discussed, this implies

$$\Pr[u \in M_B] = 1 - \prod_i \Pr[Y_i = 0] = 1 - \prod_i (1 - \mathbb{E}[Y_i]).$$

Fixing the sum  $\sum_i \mathbb{E}[Y_i]$  to be  $S$ ,  $1 - \prod_i (1 - \mathbb{E}[Y_i])$  is minimized for  $\mathbb{E}[Y_1] = \dots = \mathbb{E}[Y_d] = \frac{S}{d}$ . Thus

$$\Pr[u \in M_{\mathcal{B}}] \geq 1 - \prod_{i=1}^d \left(1 - \frac{S}{d}\right) = 1 - \left(1 - \frac{S}{d}\right)^d \geq (1 - 1/e)S \stackrel{(13)}{=} (1 - 1/e) \Pr[u \in \mathcal{A}(\Gamma_p)]. \quad (14)$$

Now, by linearity of expectation over all choices of  $u \in B$ , we get

$$\mathbb{E}|M_{\mathcal{B}}| = \sum_{u \in B} \Pr[u \in M_{\mathcal{B}}] \stackrel{(14)}{\geq} (1 - 1/e) \sum_{u \in B} \Pr[u \in \mathcal{A}(\Gamma_p)] = (1 - 1/e) \mathbb{E}|M_{\mathcal{A}}|,$$

completing the proof.

**Property (ii).** If a vertex  $v \in A$  does not propose, it remains unmatched in  $M_{\mathcal{B}}$ . Thus:

$$\Pr[v \in M_{\mathcal{B}}] \leq \Pr[v \text{ proposes}] = \mathbb{E}[p_v] = \Pr[v \in M_{\mathcal{A}}].$$

**Property (iii).** For any vertex  $u \in B$ ,

$$\Pr[u \in M_{\mathcal{B}}] \leq \Pr[u \text{ receives a proposal}] \leq \sum_{v \in N(u)} \mathbb{E}[p_{(v,u)}] = \Pr[u \in M_{\mathcal{A}}].$$

**Property (iv).** Observe that in order to determine  $u \in M_{\mathcal{B}}$ , it suffices to reveal the proposals of its neighbors in  $A$ . If one of them proposes to  $u$  then  $u$  is matched and otherwise it is not. Now since  $(u, v) \notin E$  as assumed by Property (iv), the proposals of  $v$  remains completely unknown and independent of  $u \in M_{\mathcal{B}}$ .  $\square$

## 5.2 Proving Lemma 5.1 via the Vertex-Independent Lemma 5.2

To prove Lemma 5.1, we prove two different bounds on the expected size of  $\mu(G_p)$ . The first one is easy to prove and is as follows:

**Claim 5.3.**  $\mathbb{E}[\mu(G_p)] \geq \mathbf{q}(Q)$ .

*Proof.* As assumed in Lemma 5.1,  $q_e = \Pr[e \in \mathcal{M}(Q_p)]$  for any  $e \in Q$ . By linearity of expectation, this implies  $\mathbb{E}|\mathcal{M}(Q_p)| = \mathbf{q}(Q)$ . Since  $Q_p \subseteq E_p$ , any edge in  $\mathcal{M}(Q_p)$  appears in  $G_p$  and thus the same lower bound also holds for  $\mathbb{E}[\mu(G_p)]$  completing the proof.  $\square$

The second bound is the main part of the proof, and reads as follows:

**Claim 5.4.**  $\mathbb{E}[\mu(G_p)] \geq (1 - 6\varepsilon) \left( \frac{e-1}{e} \cdot \mathbf{q}(Q) + \mathbf{q}(S) \right)$ .

Let us first see how the combination of Claims 5.3 and 5.4 proves Lemma 5.1. Observe that since each edge of  $G$  is either in  $Q$  or  $S$ ,  $|\mathbf{q}| = \mathbf{q}(Q) + \mathbf{q}(S)$ . Now consider two cases:

**Case 1** —  $\mathbf{q}(Q) \geq \frac{e}{e+1} |\mathbf{q}|$ : In this case, Claim 5.3 already implies Lemma 5.1.

**Case 2** —  $\mathbf{q}(Q) < \frac{e}{e+1} |\mathbf{q}|$ : In this case, by Claim 5.4, we have

$$\mathbb{E}[\mu(G_p)] \geq (1 - 6\varepsilon) \left( \frac{e-1}{e} \cdot \mathbf{q}(Q) + \mathbf{q}(S) \right).$$



The right hand side is minimized, when as much of the weight of  $\mathbf{q}$  comes from  $Q$  instead of  $S$ . But, by the assumption of Case 2,  $\mathbf{q}(Q) < \frac{e}{e+1}|\mathbf{q}|$ . Thus:

$$\mathbb{E}_{G_p}[\mu(G_p)] \geq (1 - 6\varepsilon) \left( \frac{e-1}{e} \cdot \frac{e}{e+1} \cdot |\mathbf{q}| + \frac{1}{e+1} \cdot |\mathbf{q}| \right) = (1 - 6\varepsilon) \frac{e}{e+1} |\mathbf{q}|,$$

which is the desired bound of Lemma 5.1.

### 5.2.1 Proof of Claim 5.4

In order to argue that  $G_p$  has a matching of our desired expected size, we construct a fractional matching  $\mathbf{x}$  on it. Since the graph is bipartite, any fractional matching can be turned into an integral matching of at least the same size. As a result, it suffices to argue that fractional matching  $\mathbf{x}$  has our desired size in expectation.

To construct fractional matching  $\mathbf{x}$ , we first use the vertex-independent matching Algorithm 1 of Lemma 5.2 to construct an integral matching  $M_B$  on  $Q_p$  (the parameters that we feed into Lemma 5.2 are formalized below). We then use  $M_B$  to define  $\mathbf{y} : E \rightarrow \mathbb{R}_+$  which will be very close to our final fractional matching  $\mathbf{x}$ , except that for a small fraction of vertices, the value of  $y_v := \sum_{e \ni v} y_e$  may exceed one due to deviations in our random process. We then scale down  $\mathbf{y}$  to obtain our fractional matching  $\mathbf{x}$  and finally argue that  $\mathbf{x}$  is large enough.

For brevity, we use  $(v \text{ prop})$  and  $(\overline{v \text{ prop}})$  to indicate respectively the events that a vertex  $v \in A$  proposes and does not propose in Algorithm 1 for constructing  $M_B$ .

The formal definition of  $\mathbf{y}$ , given matching  $M_B$  is given below:

$$y_e \leftarrow \begin{cases} 1 & \text{if } e \in M_B \text{ (this implies } e \in Q_p) \\ q_e / (p \Pr[\overline{v \text{ prop}}] \Pr[u \notin M_B]) & \text{if } e \in S_p, u \notin M_B, \text{ and } \overline{v \text{ prop}} \\ 0 & \text{otherwise,} \end{cases} \quad \forall e = (u, v) \in E, \quad v \in A, u \in B. \quad (15)$$

As discussed,  $\mathbf{y}$  is not necessarily a valid fractional matching since for some vertices  $v$ ,  $y_v$  may be larger than 1 due to some low probability (but still likely to occur) events. To resolve this, we define the final, always valid, fractional matching  $\mathbf{x}$  as follows:

$$x_e \leftarrow \begin{cases} y_e / (1 + \varepsilon) & \text{if } y_v \leq 1 + \varepsilon \text{ and } y_u \leq 1 + \varepsilon, \\ 0 & \text{otherwise,} \end{cases} \quad \forall e = (u, v) \in E. \quad (16)$$

**The Matching  $M_B$ :** As discussed, we use Lemma 5.2 to construct matching  $M_B$ . In order to use this lemma, we have to specify: (i) what graph we run the matching algorithm of this lemma on, and (ii) what algorithm  $\mathcal{A}$  we feed into the lemma. For the first question, the graph  $\Gamma$  on which we run the lemma is simply the subgraph  $Q$  of  $G$ , and we let  $\Gamma_p$  correspond to the realized edges  $Q_p$ . This ensures that the reported matching  $M_B$  of Lemma 5.2 satisfies  $M_B \subseteq Q_p$ . For the second question, we first define an auxiliary matching  $M'_A := \mathcal{M}(Q_p)$  and let  $M_A$  include each edge of  $M'_A$  independently with probability  $1 - \varepsilon$ . This down sampling step is rather technical and is there to just ensure  $\Pr[w \in M_A] \leq 1 - \varepsilon$  for any vertex  $w$ .

Defining  $M_A$  this way, as a corollary of Lemma 5.2 we get:

**Corollary 5.5.** *The matching  $M_B$  constructed as above on subgraph  $Q_p$  satisfies:*

$$(i) \quad \mathbb{E}|M_B| \geq (1 - \frac{1}{e})\mathbb{E}|M_A| = (1 - \varepsilon)(1 - \frac{1}{e}) \cdot \mathbf{q}(Q).$$

- (ii) For any vertex  $v \in A$ ,  $\Pr[v \in M_{\mathcal{B}}] \leq \Pr[v \text{ prop}] = \Pr[v \in M_{\mathcal{A}}] = (1 - \varepsilon) \Pr[v \in \mathcal{M}(Q_p)]$ .
- (iii) For any vertex  $u \in B$ ,  $\Pr[u \in M_{\mathcal{B}}] \leq \Pr[u \in M_{\mathcal{A}}] = (1 - \varepsilon) \Pr[u \in \mathcal{M}(Q_p)]$ .
- (iv) For any  $v \in A, u \in B$  with  $(u, v) \in S$ , events  $(v \text{ prop})$  and  $(u \in M_{\mathcal{B}})$  are independent.

*Proof.* The first three are simply followed by the properties of Lemma 5.2 combined with the definition of  $M_{\mathcal{A}}$  above. The last property holds since  $(u, v) \in S$  implies  $(u, v) \notin Q$  (since each there are no parallel edges and each edge belongs to exactly one of  $Q$  and  $S$ ) which combined with Property (iv) of Lemma 5.2 implies the stated independence.  $\square$

Having defined  $M_{\mathcal{B}}$ , the definitions of  $\mathbf{y}$  and  $\mathbf{x}$  are complete. We now turn to analyze their size. We first analyze the size of  $\mathbf{y}$  and then prove the size of  $\mathbf{x}$  is actually very close to that of  $\mathbf{y}$  by bounding the probability of deviations leading to vertices to have  $y_v > 1 + \varepsilon$ .

**The Expected Size of  $\mathbf{y}$ :** We have  $\mathbb{E}|\mathbf{y}| = \sum_{e \in E} \mathbb{E}[y_e] = \sum_{e \in Q} \mathbb{E}[y_e] + \sum_{e \in S} \mathbb{E}[y_e]$ . For edges  $e \in Q$  we have  $y_e = 1$  iff  $e \in M_{\mathcal{B}}$ , which combined with  $M_{\mathcal{B}} \subseteq Q_p$  implies  $\sum_{e \in Q} \mathbb{E}[y_e] = \mathbb{E}|M_{\mathcal{B}}|$ . On the other hand, by definition of  $\mathbf{y}$  on edges in  $S$ , we have

$$\mathbb{E}|\mathbf{y}| = \mathbb{E}|M_{\mathcal{B}}| + \sum_{e \in S} \Pr[e \in S_p, u \notin M_{\mathcal{B}}, \overline{v \text{ prop}}] \cdot \frac{q_e}{p \Pr[u \notin M_{\mathcal{B}}] \Pr[\overline{v \text{ prop}}]}.$$

For any  $e \in S$ ,  $e \in S_p$  iff  $e$  is realized which is independent of how matching  $M_{\mathcal{B}}$  is constructed. This holds because in constructing matching  $M_{\mathcal{B}}$  we do not look at the realized edges  $S_p$  of  $S$ . On the other hand, by Corollary 5.5 Property (iv),  $u \in M_{\mathcal{B}}$  and  $\overline{v \text{ prop}}$  are also independent, hence

$$\begin{aligned} \mathbb{E}|\mathbf{y}| &= \mathbb{E}|M_{\mathcal{B}}| + \sum_{e \in S} q_e \\ &\geq (1 - \varepsilon)(1 - 1/e) \cdot \mathbf{q}(Q) + \sum_{e \in S} q_e && \text{(By Corollary 5.5 Property (i))} \\ &\geq (1 - \varepsilon) \left( (1 - 1/e) \cdot \mathbf{q}(Q) + \mathbf{q}(S) \right). \end{aligned} \tag{17}$$

Note that  $\mathbf{y}$ , as proved above, is as large as the lower bound required by Claim 5.4. However, we want to show that  $\mathbb{E}|\mathbf{x}|$  is also this large (up to  $1 - \Theta(\varepsilon)$  factors). This is what we prove next.

**The Expected Size of  $\mathbf{x}$ :** Take an edge  $e = (v, u) \in E$ . Observe from the construction of  $\mathbf{x}$  that either  $x_e = y_e/(1 + \varepsilon)$ , or  $x_e = 0$  if  $y_v > 1 + \varepsilon$  or  $y_u > 1 + \varepsilon$ . This implies that

$$\mathbb{E}|\mathbf{x}| = \sum_{e \in E} \mathbb{E}[x_e] = \sum_{e=(u,v) \in E} \mathbb{E}[y_e/(1 + \varepsilon) \mid y_v \leq 1 + \varepsilon, y_u \leq 1 + \varepsilon].$$

Now an edge  $e \in E$  either belongs to  $Q$  or  $S$ . For edges  $e = (u, v) \in Q$ , it is not hard to see that we have  $x_e = y_e/(1 + \varepsilon)$  with probability one, and thus

$$\mathbb{E}[x_e] = \mathbb{E}[y_e]/(1 + \varepsilon) \quad \forall e \in Q. \tag{18}$$

The reason behind this, is that if  $e = (u, v) \in Q$  then by construction of  $\mathbf{y}$ ,  $y_e = \mathbf{1}(e \in M_{\mathcal{B}})$ . Further, if it occurs that  $y_e = 1$ , then both of its endpoints are matched in  $M_{\mathcal{B}}$  and hence we have  $y_{e'} = 0$  for all other edges  $e'$  connected to either  $u$  or  $v$ , again by construction of  $\mathbf{y}$ . As a result,  $y_e = 1$  implies  $y_u = y_v = 1$  and thus  $x_e = y_e/(1 + \varepsilon)$ ; otherwise  $x_e = y_e = 0$ , proving (18).

The situation, however, is more complicated for edges  $e \in S$  as in this case  $\mathbf{y}$  on the endpoints of  $e$  may exceed one. We show, however, that this occurs with a small enough probability that we can still argue that it does not affect the expected value of  $x_e$  by much, and  $\mathbb{E}[x_e] \approx \mathbb{E}[y_e]$ .

**Claim 5.6.** For any vertex  $u \in B$ ,  $\mathbb{E}[y_u \mid u \notin M_B] \leq q_u$ .

*Proof.* Let  $v_1, \dots, v_d$  be the neighbors of  $u$  in  $S$ , and let  $e_i = (u, v_i)$ . We have

$$\mathbb{E}[y_u \mid u \notin M_B] = \sum_{i=1}^d \frac{q_{e_i}}{p \Pr[u \notin M_B] \Pr[v_i \text{ prop}]} \cdot \Pr[e_i \in S_p, u \notin M_B, \overline{v_i \text{ prop}} \mid u \notin M_B].$$

By independence of events in the probability (justified before), we can simplify this to

$$\mathbb{E}[y_u \mid u \notin M_B] = \sum_{i=1}^d \frac{q_{e_i}}{\Pr[u \notin M_B]} = \frac{\sum_{i=1}^d q_{e_i}}{\Pr[u \notin M_B]}.$$

Now let  $q_u^S$  denote the sum of fractional value written on edges of  $u$  in  $S$  and let  $q_u^Q$  denote the same but on edges of  $u$  in  $Q$ . The nominator equals  $q_u^S$ . Also from Corollary 5.5 Property (iii), thus  $\Pr[u \in M_B] \leq (1 - \varepsilon)q_u^Q$ , we get  $\Pr[u \notin M_B] \geq 1 - (1 - \varepsilon)q_u^Q \geq 1 - q_u^Q$ . Combining them gives

$$\mathbb{E}[y_u \mid u \notin M_B] \leq \frac{q_u^S}{1 - q_u^Q} = \frac{q_u \cdot q_u^S}{q_u(1 - q_u^Q)} = \frac{q_u \cdot q_u^S}{q_u - q_u \cdot q_u^Q} \leq \frac{q_u \cdot q_u^S}{q_u - q_u^Q} = \frac{q_u \cdot q_u^S}{q_u^S} = q_u,$$

as desired.  $\square$

**Claim 5.7.** For any vertex  $u \in B$ ,  $\mathbb{E}[y_u \mid y_u \leq 1 + \varepsilon, u \notin M_B] \geq \mathbb{E}[y_u] - \varepsilon q_u$ .

*Proof.* As proved in Claim 5.6,  $\mathbb{E}[y_u \mid u \notin M_B] \leq q_u$ . We prove the desired inequality of the claim via a concentration bound on random variable  $y'_u := (y_u \mid u \notin M_B)$ .

Let  $v_1, \dots, v_d$  be the neighbors of  $u$  in  $S$ , and let  $e_i = (u, v_i)$ . Let us for simplicity define random variable  $y'_{e_i} := (y_{e_i} \mid u \notin M_B)$ . Since we have conditioned on  $u \notin M_B$ , we get by definition of  $\mathbf{y}$  that  $y'_u = \sum_{i=1}^d y'_{e_i}$ . Now we argue that  $y'_{e_1}, \dots, y'_{e_d}$  are actually independent. To see this, observe that once we condition on  $u \notin M_B$ , the only random process that determines each  $y'_{e_i}$  is whether edge  $e_i$  is realized (i.e.  $e_i \in S_p$ ) and if vertex  $v_i$  proposes. Both events are independent of  $u \notin M_B$  since, recall, to determine  $u \notin M_B$  we only need to know the proposals of its neighbors in  $Q$ , and  $v_1, \dots, v_d$  are all neighbors of  $u$  in  $S$ .

This independence allows us to prove a concentration bound on  $y'_u$  and prove the claim. We do this via the second moment method. We have

$$\text{Var}[y'_u] = \sum_{i=1}^d \text{Var}[y'_{e_i}] = \sum_{i=1}^d (\mathbb{E}[(y'_{e_i})^2] - \mathbb{E}[y'_{e_i}]^2) \leq \sum_{i=1}^d \mathbb{E}[(y'_{e_i})^2] \stackrel{\text{see below}}{\leq} \tau \cdot \mathbb{E}[y'_u] \stackrel{\text{Claim 5.6}}{\leq} \tau \cdot q_u,$$

where  $\tau$  is the maximum possible outcome of  $y'_{e_i}$  for any  $i \in [d]$ .

Plugging this into Chebyshev's inequality, we get

$$\Pr[y'_u > \mathbb{E}[y'_u] + \delta] \leq \frac{\text{Var}[y'_u]}{\delta^2} \leq \frac{\tau q_u}{\delta^2}.$$

Finally, for each edge  $e_i$ , by construction of  $\mathbf{y}$ ,  $y_{e_i} \leq q_{e_i} / (p \Pr[v_i \text{ prop}] \Pr[u \notin M_B]) \leq q_{e_i} / p\varepsilon^2$  where the latter follows from Corollary 5.5. Combined with  $q_{e_i} \leq \varepsilon^5 p$  by Lemma 5.1 and  $e_i \in S$ , we get  $\tau \leq \varepsilon^3$ . We thus get  $\Pr[y'_u > 1 + \varepsilon] \leq \Pr[y'_u > \mathbb{E}[y'_u] + \varepsilon] \leq \frac{\varepsilon^3}{\varepsilon^2} q_u$ . This concentration of  $y'_u$  implies  $\mathbb{E}[y_u \mid y_u > 1 + \varepsilon, u \notin M_B] \leq \varepsilon q_u$  and thus the stated bound of the claim.  $\square$

Similarly, for each vertex  $v \in A$  we can first bound the expected value of  $y_v$  conditioned on  $(\overline{v \text{ prop}})$  and then prove this is concentrated, and overall argue  $\mathbf{x}$  remains close to  $\mathbf{y}$ . Namely:

**Claim 5.8.** *For any  $v \in A$ ,  $\mathbb{E}[y_v \mid y_v \leq 1 + \varepsilon, \overline{v \text{ prop}}] \geq \mathbb{E}[y_v] - \varepsilon q_v$ .*

Since the proof is similar to Claim 5.7 for vertices in  $B$  and there are only minor differences, we defer it to Appendix A.

To complete the proof that  $\mathbb{E}[\mathbf{x}] \approx \mathbb{E}[\mathbf{y}]$ , note that

$$\begin{aligned} \mathbb{E}[\mathbf{x}] &= \frac{1}{2} \sum_{w \in A \cup B} \mathbb{E}[x_w] = \frac{1}{2} \sum_{w \in A \cup B} \frac{\mathbb{E}[y_w \mid y_w \leq 1 + \varepsilon]}{1 + \varepsilon} \stackrel{\text{Claims 5.7, 5.8}}{\geq} \left( \frac{1}{2} \sum_{w \in A \cup B} \frac{\mathbb{E}[y_w]}{1 + \varepsilon} \right) - 3\varepsilon |\mathbf{q}| \\ &\geq \frac{1}{1 + \varepsilon} |\mathbf{y}| - 3\varepsilon |\mathbf{q}| \stackrel{(17)}{\geq} (1 - \varepsilon) \left( (1 - \varepsilon) \left( (1 - 1/e) \cdot \mathbf{q}(Q) + \mathbf{q}(S) \right) \right) - 3\varepsilon |\mathbf{q}| \\ &\geq (1 - 6\varepsilon) \left( (1 - 1/e) \cdot \mathbf{q}(Q) + \mathbf{q}(S) \right). \end{aligned}$$

Since as discussed  $\mathbb{E}[\mu(G_p)] \geq \mathbb{E}[\mathbf{x}]$ , this completes the proof of Claim 5.4.

## 6 Upper Bounds

In this section, we present our main algorithms for the stochastic vertex cover problem.

### 6.1 Bipartite Graphs: 1.36-Approximation with $\text{poly}(\frac{1}{p})$ Queries

Our main result in this section is the following stochastic vertex cover result for bipartite graphs:

**Theorem 6.1.** *For any  $p \in (0, 1]$ , any bipartite graph  $G = (V, E)$  has a subgraph  $Q$  of maximum degree  $O(1/p^6)$  where querying only the edges in  $Q$  suffices to find  $C \subseteq V$  such that:*

1.  $C$  is a vertex cover of  $G_p$  with probability 1.
2. The expected size of  $C$  is at most  $1.367 (\approx \frac{e+1}{e})$  times the size of  $\mathbb{E}[\nu(G_p)]$ .
3. Both  $Q$  and  $C$  can be found in polynomial time.

To prove Theorem 6.1, we first use the Half-Stochastic Matching Lemma 4.1 to obtain a partitioning  $(Q, S)$  of  $E$ . We then query the edges in  $Q$ , and report the MVC of  $H = Q_p \cup S$  as the vertex cover for  $G_p$ . We finally use the algorithm  $\mathcal{M}$  provided by this lemma, as well as Lemma 5.1 to analyze the approximation ratio of this algorithm.

The formal algorithm is as follows:

**Algorithm 2.** The algorithm for Theorem 6.1.

- 1 Let  $Q, S$  be the partitioning found by Lemma 4.1 for the following parameters:  $G$  is the given base graph,  $p$  is the realization probability, and let  $\varepsilon > 0$  be a sufficiently small constant which adjusts how close the approximation will be to  $\frac{e+1}{e}$ .
- 2 Query the edges in  $Q$  and let  $Q_p$  be the edges in  $Q$  that are realized.
- 3 Return a minimum vertex cover  $C$  of graph  $H = Q_p \cup S$ .

Intuitively, what we do in Algorithm 2 is to query only the edges in  $Q$ , assume that the rest of the edges in  $S$  are all realized, and report a MVC of the resulting graph  $H = Q_p \cup S$ . Note that the vast majority of the edges in  $G$  belong to  $S$  since  $Q$  has only a constant maximum degree. As

a result, we have to argue that the extra constraints imposed by assuming that all these edges are realized, do not increase the vertex cover size by much.

Before analyzing the size of the vertex cover  $C$ , let us explain why it is always a valid vertex cover of  $G_p$  and analyze the query-complexity of Algorithm 2. These will be simple consequences of the properties provided by Lemma 4.1.

**Query-Complexity and Validity of the Vertex Cover:** By Property (i) of Lemma 4.1 the maximum degree in  $Q$  is at most  $O(1/\varepsilon^{11}p^6)$ ; thus, we query at most  $O(1/p^6)$  edges per vertex. The validity of the vertex cover is also easy to confirm. By definition of  $H$ , an edge  $e \in E$  is not in  $H$  if and only if it belongs to  $Q$  and is not realized. Therefore, all realized edges must belong to  $H$  implying that  $G_p$  is a subgraph of  $H$ . As a result, a vertex cover of  $H$  covers all the edges in  $G_p$ .

**The Approximation Ratio:** We use Lemma 5.1, as well as the properties of the partitioning provided by Lemma 4.1, to show that the expected size of the MVC in graph  $H$ , which is reported by Algorithm 2 as the output, is not larger than (almost)  $\frac{e+1}{e}$  times the expected size of the MVC in the actual realization  $G_p$ , namely that for any arbitrarily small constant  $\delta'$  (affecting  $\varepsilon$  in Algorithm 2) it holds that:

$$\mathbb{E}_H[\nu(H)] \leq (1 + \delta') \frac{e+1}{e} \cdot \mathbb{E}_{G_p}[\nu(G_p)].$$

Since the graph is bipartite, by König's theorem, the size of maximum matching and MVC are the same. As such, it suffices to prove

$$\mathbb{E}_H[\mu(H)] \leq (1 + \delta') \frac{e+1}{e} \cdot \mathbb{E}_{G_p}[\mu(G_p)]. \quad (19)$$

In order to show this, we define a fractional matching  $\mathbf{q}$  on  $G$  with size  $|\mathbf{q}| \geq (1 - 2\varepsilon)\mathbb{E}[\mu(H)]$ . We then show that this fractional matching  $\mathbf{q}$  satisfies the required properties of Lemma 5.1 and as a result, implies  $\mathbb{E}[\mu(G_p)]$  has the desired size of (19).

**The Fractional Matching  $\mathbf{q}$ :** Consider (random) matching  $\mathcal{M}(H)$  and recall that  $\mathcal{M}$  is the matching algorithm provided by Lemma 4.1. For each edge  $e \in E$ , we simply let  $q_e \leftarrow \Pr[e \in \mathcal{M}(H)]$ . Clearly  $\mathbf{q}$  is a valid fractional matching since the probabilities around each vertex correspond to its probability of being matched and thus do not exceed one. Moreover,  $|\mathbf{q}| = \mathbb{E}[\mathcal{M}(H)]$  by linearity of expectation, combined with Lemma 4.1 Property (ii) that  $\mathbb{E}[\mathcal{M}(H)] \geq (1 - 2\varepsilon)\mathbb{E}[\mu(H)]$ , we get  $|\mathbf{q}| \geq (1 - 2\varepsilon)\mathbb{E}[\mu(H)]$ ; hence  $|\mathbf{q}|$  has the claimed size too. It remains to prove the two assumptions of Lemma 5.1 are also satisfied by  $\mathbf{q}$ . The first one requires us to give a matching algorithm  $\mathcal{M}'$  where  $q_e = \Pr[e \in \mathcal{M}'(Q_p)]$  for all  $e \in Q$ . Letting  $\mathcal{M}'(Q_p) := \mathcal{M}(H) \cap Q_p$  suffices for this purpose, since recall that any edge in  $\mathcal{M}(H) \cap Q$  is already in  $Q_p$  by definition of graph  $H$  in Lemma 4.1 and thus  $\Pr[e \in \mathcal{M}'(Q_p)] = q_e$ . For the second assumption we need  $q_e \leq \delta^5 p$  for all  $e \in S$  (here we used  $\delta$  instead of  $\varepsilon$  to avoid confusion with parameter  $\varepsilon$  that we feed into Lemma 4.1). Indeed, by Property (iii), we have  $q_e \leq \varepsilon^2 p$  and it suffices to let  $\delta = \varepsilon^{2/5}$ . We can, now, apply Lemma 5.1 to obtain:

$$\mathbb{E}[\mu(G_p)] \geq (1 - 4\delta) \frac{e}{e+1} |\mathbf{q}| \geq (1 - 4\delta)(1 - 2\varepsilon) \frac{e}{e+1} \mathbb{E}[\mu(H)].$$

Since we can let  $\varepsilon$  (and thus  $\delta$ ) be any desirably small constant, this proves (19) and our claim that our reported vertex cover has size at most (arbitrarily close to)  $\frac{e+1}{e} \mathbb{E}[\nu(G_p)]$ .

**Stochastic Matchings:** Finally, we note that the same tools we used for this problem also lead to Corollary 1.1. To prove it, we in fact, provide a novel analysis for a well-known Monte Carlo

algorithm for the stochastic matching problem that is very different from the algorithm we use for Theorem 6.1. This is why the number of queries in Corollary 1.1 and Theorem 6.1 are different. But the new analysis, is also based on Lemma 5.1 similar to above.

We analyze the following algorithm proposed first in [7] and further analyzed in [6, 4].

**Algorithm 3.** A Monte-Carlo algorithm for stochastic matching [7].

- 1 For large enough  $R = O(\frac{\log 1/p}{p})$ , take independent realizations  $G_1, \dots, G_R$  of  $G$ .
- 2 For some deterministic maximum matching algorithm  $\text{MM}(\cdot)$ , query subgraph  $Q := \text{MM}(G_1) \cup \dots \cup \text{MM}(G_R)$  of  $G$  and report the maximum matching in  $Q_p$ .

The following lemma is implied in [7] for Algorithm 3:

**Claim 6.1** ([7]). *For any desirably small constant  $\varepsilon > 0$  (affecting the hidden constants in  $R$ ), there is a partitioning  $C, N$  of  $Q$ , and a fractional matching  $\mathbf{q}$  on  $Q$  such that (i)  $\mathbb{E}|\mathbf{q}| \geq (1 - \varepsilon)\mathbb{E}[\mu(G_p)]$ , (ii) for each edge  $e \in N$ ,  $q_e \leq \varepsilon^5 p$ , (iii) for each  $e \in C$ ,  $q_e = \Pr[e \in \text{MM}(G_p)]$ .*

Here we briefly describe the intuition behind Claim 6.1. For the complete proof see [7].

*Proof sketch of Claim 6.1.* Partition the edges of graph  $G$  into two subsets: *crucial* edges  $e$  defined as those with  $\Pr[e \in \text{MM}(G_p)] \geq \tau$  for a small enough parameter  $\tau = \varepsilon^{O(1)}p$ , and *non-crucial* edges which include all the rest of edges  $e$  with  $\Pr[e \in \text{MM}(G_p)] < \tau$ .

Let subgraphs  $C$  and  $N$  of Claim 6.1 be respectively the set of crucial and non-crucial edges of  $G$  that belong to  $Q$ . We define  $\mathbf{q}$  such that the size of  $\mathbf{q}$  on  $C$  is  $(1 - \varepsilon)$  times the expected contribution of crucial edges to  $\text{MM}(G_p)$  and, similarly, the size of  $\mathbf{q}$  on  $N$  is  $(1 - \varepsilon)$  times the expected contribution of non-crucial edges to  $\text{MM}(G_p)$ . This way, we guarantee property (i).

To define  $\mathbf{q}$  on  $C$ , for any edge  $e \in C$  we let  $q_e = \Pr[e \in \text{MM}(G_p)]$ . It can be easily confirmed that if the parameter  $R$  of Algorithm 3 is larger than  $\log(1/\varepsilon)/\tau$ , then each crucial edge is added to  $Q$  with probability at least  $1 - \varepsilon$ . As such, fractional matching  $\mathbf{q}$  on  $C$  has size at least  $(1 - \varepsilon)$  fraction of the expected contribution of crucial edges to  $\text{MM}(G_p)$ .

On the flip side, however,  $Q$  includes only a small fraction of non-crucial edges. Hence, to maintain the property that  $\mathbf{q}$  on  $N$  is almost as large as the expected contribution of non-crucial edges to  $\text{MM}(G_p)$ , the value of  $q_e$  must be much larger than  $\Pr[e \in \text{MM}(G_p)]$  for  $e \in N$ . To do this, suppose that we define  $q_e$  to be the *fraction* of matchings  $\text{MM}(G_1), \dots, \text{MM}(G_R)$  that include  $e$ . Observe that each edge  $e$  in the graph, crucial or non-crucial, is expected to appear in exactly  $\Pr[e \in \text{MM}(G_p)]$  fraction of matchings  $\text{MM}(G_1), \dots, \text{MM}(G_R)$  since each one includes  $e$  with probability exactly  $\Pr[e \in \text{MM}(G_p)]$ . Hence, by defining  $\mathbf{q}$  on  $N$  this way, we get that the expected size of  $\mathbf{q}$  on  $N$  is at least the expected contribution of  $N$  to  $\text{MM}(G_p)$ . Unfortunately, however,  $\mathbf{q}$  may violate fractional matching constraints with this construction. Namely, that  $q_v$  for a vertex  $v$  may exceed one. The next important observation is that this violation cannot be too large, since the fraction of matchings  $\text{MM}(G_1), \dots, \text{MM}(G_R)$  in which a vertex  $v$  is matched via a non-crucial edge is sufficiently concentrated around the probability that  $v$  is matched via a non-crucial edge in  $\text{MM}(G_p)$ . The final fractional matching is obtained by slightly modifying this fractional matching (particularly by discarding vertices that deviate too much and multiplying the rest of the values by some  $(1 - \varepsilon)$  factor) so that no constraints are violated.  $\square$

Having it, we can now plug  $\mathbf{q}$  into Lemma 5.1 and obtain that  $\mathbb{E}[\mu(Q_p)] \geq (1 - 5\varepsilon)(\frac{e}{e+1})\mathbb{E}[\mu(G_p)]$ , thereby proving Corollary 1.1.



## 6.2 General Graphs: $(2 + \varepsilon)$ -Approximation with $O(\frac{1}{p})$ Per-Vertex Queries

In this section, we prove the following result:

**Theorem 6.2.** *For any  $\varepsilon > 0$  and  $p \in (0, 1]$ , any (general) graph  $G = (V, E)$  has a subgraph  $Q$  of maximum degree  $O(\frac{1}{\varepsilon^3 p})$  where querying only the edges in  $Q$  suffices to find  $C \subseteq V$  such that:*

1.  $C$  is a vertex cover of  $G_p$  with probability 1.
2. The size of  $C$  is in expectation at most  $(2 + \varepsilon)$  times the minimum vertex cover of  $G_p$ .
3. It is possible to find  $Q$  and  $C$  in polynomial time.

We start with a subroutine for constructing a fractional matching on a given graph, and then describe our algorithm which proves Theorem 6.2.

**A fractional matching subroutine:** Consider a simple and well-known fractional matching algorithm which starts with a zero-size fractional matching and gradually increases the fractional values on the edges all at the same (additive) rate. Once the fractional value around a vertex reaches one, we mark this vertex as *inactive* and stop increasing the fractional value of its edges. We will use a slightly different variant of this algorithm in Algorithm 4 where the vertices may be made inactive sooner; i.e., once they reach a given budget; this variant is formalized below.

Let  $G$  be a graph, and for each vertex  $v$ , let  $b(v) \in (0, 1]$  be a given *budget*. Initially, every vertex  $v$  is *active* and for each edge  $e$  we set  $x_e := 0$ . The algorithm proceeds in at most  $n$  steps. In each step, for any edge  $e$  whose both endpoints are active, we increase  $x_e$  for all the edges in the same rate until for a vertex we have  $x_v := \sum_{e \ni v} x_e = b(v)$ . When this event happens for a vertex  $v$  it becomes inactive, which implies that for its edges  $e$ ,  $x_e$  will no longer change. Algorithm 4 is the pseudo-code of this process.

It is clear that throughout the algorithm we have  $x_v \leq b(v) \leq 1$  for every vertex  $v$ . Hence, at every point in the algorithm, vector  $\mathbf{x} := (x_e)_{e \in E}$  is a fractional matching of  $G$ . Let us define  $\mathbf{x}^{(t)}$  to be equal to  $\mathbf{x}$  from an iteration of the algorithm after which for at least one edge we have  $x_e > t$ . (If  $x_e > t$  never happens then  $\mathbf{x}^{(t)}$  is from the last iteration of the algorithm.) We can intuitively think of the algorithm above as a continuous process over a time interval of  $[0, 1]$  that gradually increases the fractional matching on all edges with active endpoints, all at the same additive rate, until every vertex becomes inactive. The value of  $\mathbf{x}^{(t)}$  can thus be interpreted as the fractional matching constructed by time  $t$  of this process.

### Algorithm 4. Filling( $G = (V, E), b : V \rightarrow (0, 1]$ )

- 1 For any edge  $e \in E$ , set  $x_e := 0$ .
- 2 **repeat**
- 3     Call a vertex *inactive* iff  $x_v := \sum_{e \ni v} x_e = b(v)$ , and *active* otherwise.
- 4     Call an edge active iff both its endpoints are active, and inactive otherwise.
- 5     Pick the minimum parameter  $\delta \in (0, 1)$  such that setting  $x_e \leftarrow x_e + \delta$  for all the active edges, results in at least one new inactive vertex.
- 6     Set  $x_e \leftarrow x_e + \delta$  for all active edges.
- 7 **until** All the vertices are inactive
- 8 **return** fractional matching  $\mathbf{x} \leftarrow (x_e)_{e \in E}$ .

**Our stochastic vertex cover algorithm.** We use Algorithm 5 to decide which edges to query and which vertices to put in the vertex cover. Here we give an informal overview of this algorithm.

The algorithm starts by running Algorithm 4 on the static graph  $G$  with a budget of 1 per



vertex, to obtain a fractional matching  $\mathbf{x}$ . We will, in fact, only need the fractional matching  $\mathbf{x}'$  constructed by this algorithm up to time  $t := \Theta(\varepsilon^3 p)$ , i.e.  $\mathbf{x}' := \mathbf{x}^{(t)}$ . Once we have  $\mathbf{x}'$  and before we query any edge, we commit the vertices with fractional matching value 1 in  $\mathbf{x}'$  to be in our final vertex cover. Let  $F := \{v \in V : x'_v = 1\}$  denote the set of these vertices and let  $Q$  be the edges in  $E$  that do not have an endpoint in  $F$ . Note that any edge in  $E \setminus Q$  is already covered by  $F$ ; hence we do not need to query them. We, thus, only query the edges in  $Q$ . Let  $Q_p$  be the subset of edges in  $Q$  that turn out to be realized. At least one of the endpoints of each edge in  $Q_p$  should be added to the vertex cover. To decide which ones join the vertex cover, we again run Algorithm 4, but this time on subgraph  $Q_p$  and we set the budget of each vertex  $v$  to be  $1 - x'_v$ . Let  $\mathbf{y}$  be the resulting fractional matching on  $Q_p$ . We report the set  $C := \{v \in V : x'_v + y_v = 1\}$  as the vertex cover. (Note from the definition that  $F \subseteq C$ , hence satisfying our earlier claim that we “commit”  $F$  to be in the final vertex cover.)

**Algorithm 5.** The algorithm for Theorem 6.2.

- 1 Let  $\mathbf{x} := \text{Filling}(G, b)$  where  $b(v) = 1$  for each  $v \in V$ .
- 2 Fix  $t = \Theta(\varepsilon^3 p)$  and let  $\mathbf{x}' := \mathbf{x}^{(t)}$ .
- 3 Let  $F := \{v \in V : x'_v = 1\}$ .
- 4 Query edges  $Q := \{e = (u, v) \in E : u \notin F \text{ and } v \notin F\}$  with no endpoint in  $F$ .
- 5 Let  $Q_p$  be the realized edges in  $Q$ .
- 6 Run  $\mathbf{y} := \text{Filling}(Q_p, b')$  where  $b'(v) = 1 - x'_v$  for each  $v \in V$ .
- 7 Report  $C := \{v \in V : x'_v + y_v = 1\}$  as the vertex cover of  $G_p$ .

**Validity of the vertex cover.** The proof of why the set  $C$  reported by Algorithm 5 is always a valid vertex cover of  $G_p$  is simple. We start by formalizing our earlier claim that all vertices in  $F$  also appear in  $C$ . To see this, observe that any vertex  $v$  in  $F$  by definition has  $x'_v = 1$  which also implies  $y_v = 0$  since  $b'(v) = 1 - x'_v = 0$ ; this in turn implies  $x'_v + y_v = 1$  and thus by definition  $v \in C$ . Now take an edge  $e$  of the realization  $G_p$ . If  $e$  has an endpoint in  $F$ , then this endpoint is in  $C$ , covering  $e$ . Therefore, let us take an edge  $e = (u, v)$  in  $G_p$  whose both endpoints are in  $V \setminus F$ . Observe that by definition we have  $e \in Q_p$ . Now once we run  $\text{Filling}(Q_p, b')$ , we increase the fractional value on  $e$  until one of its endpoints  $v$  reaches its budget  $b'(v)$ . Since we have set  $b'(v) = 1 - x'_v$  and  $y_v = b'(v)$ , we have  $x'_v + y_v = 1 - b'(v) + b'(v) = 1$  and thus  $v$  belongs to  $C$ , covering  $e$ .

**Analysis of the number of queries.** Observe that in Algorithm 5, we only query the edges in  $Q$ . Thus, it suffices to prove that the maximum degree in  $Q$  is at most  $O(1/\varepsilon^3 p)$  to prove that Algorithm 5 queries at most  $O(1/\varepsilon^3 p)$  edges per vertex. Consider an iteration  $i$  of the algorithm when  $\mathbf{x} = \mathbf{x}^t$ , and let  $\mathbf{x}'$  be the fractional matching from the next iteration (if any). By definition of  $\mathbf{x}^t$  and  $Q$ , at iteration  $i$ , any edge  $e \in Q$  has two active endpoints and thus is active itself. Moreover, by Algorithm 4, at any iteration, all the edges that have been active in the previous iteration have the same fractional value; thus, we have  $x'_e > t$  for any  $e \in Q$ . Moreover, as a result of  $\mathbf{x}'$  being a valid fractional matching, any vertex has at most  $1/t = \Theta(1/\varepsilon^3 p)$  edges in  $Q$ .

**Running time of the Algorithm.** The algorithm is clearly polynomial time as in each iteration of  $\text{Filling}(G, b)$ , at least one vertex becomes inactive and this can happen at most  $n$  times.

**Analysis of the approximation ratio.** The more challenging part is to prove that this vertex cover  $C$  reported by Algorithm 5 is in expectation at most  $(2 + \varepsilon)$  times larger than the minimum vertex cover of  $G_p$ . To prove this, in Lemma 6.2, we show that it is possible to construct a fractional matching  $\mathbf{y}$  of graph  $G_p$  such that  $(2 + \varepsilon)\mathbb{E}[|\mathbf{y}|] \geq \mathbb{E}[|C|]$ . By weak duality, the minimum vertex

cover of a graph is larger than any of its fractional matchings; hence,  $C$  is a  $(2 + \varepsilon)$ -approximate minimum vertex cover.

**Lemma 6.2.** *There exists a fractional matching  $\mathbf{y}$  of  $G_p$  such that  $(2 + \varepsilon)\mathbb{E}[|\mathbf{y}|] \geq \mathbb{E}[|C|]$ .*

*Proof.* We start by giving Algorithm 6 that constructs a fractional matching of  $G_p$ .

**Algorithm 6.** The algorithm for Lemma 6.2.

- 1 For any edge  $e \in S_p$  we set  $y_e \leftarrow \frac{x_e^{(t)}}{(1+\alpha)p}$  for  $\alpha = 0.25\varepsilon$ .
- 2 If a vertex  $v$  does not satisfy  $\sum_{e \ni v} y_e \geq x_v^{(t)}$  then, for any edge  $e \ni v$  set  $y_e \leftarrow 0$ .
- 3 For any  $e \in Q_p$  set  $y_e \leftarrow x'_e$  where  $\mathbf{x}' \leftarrow \text{Filling}(Q_p, b)$  with  $b(v) = 1 - x_v^{(t)}$  for any  $v \in V \setminus F$ .
- 4 Report  $\mathbf{y}$  as the fractional matching of  $G_p$ .

This algorithm consists of two stages. In the first stage, we construct a fractional matching on edges in  $S_p$  and in the second stage, we add the edges of  $Q_p$  to this matching. The first stage starts by setting

$$x_e := \frac{x_e^{(t)}}{(1 + \alpha)p}$$

for any edge  $e \in S$  that is realized, where  $\alpha = 0.25\varepsilon$ . Let us call a vertex *bad* iff it satisfies  $x_v := \sum_{e \ni v} x_e > x_v^{(t)}$ . For any bad vertex  $v$  we decrease  $x_e$  of all the edges  $e \ni v$  to zero.

In the second stage of the algorithm, we construct a fractional matching on edges in  $Q_p$ . Consider  $\mathbf{x}' := \text{Filling}(Q_p, b)$  with  $b(v) = 1 - x_v^{(t)}$  for any  $v \in Q$ . Clearly, combining  $\mathbf{x}'$  with  $\mathbf{y}$  gives us a valid fractional matching since for any vertex  $v \in V$  we have  $\sum_{v \in V} x'_v + y_v < 1$ . Hence, for any edge in  $Q_p$  we set  $y_e := x'_e$ . To complete the proof of this lemma, it suffices to take vertex cover  $C$  outputted by Algorithm 5 and prove

$$(1 + 0.5\varepsilon)\mathbb{E}\left[\sum_{v \in C} y_v\right] \geq |C|, \quad (20)$$

as it gives us

$$(2 + \varepsilon)\mathbb{E}[|X|] = 2(1 + 0.5\varepsilon)\mathbb{E}[|X|] \geq (1 + 0.5\varepsilon)\mathbb{E}\left[\sum_{v \in V} y_v\right] \geq (1 + 0.5\varepsilon)\mathbb{E}\left[\sum_{v \in C} y_v\right] \geq |C|.$$

For any vertex  $v \in C$ , let  $z_v := \sum_{e \ni v, e \in S}$  and let  $w_v := \sum_{e \ni v, e \in Q}$ . By Algorithm 5, any vertex  $v \in C$  satisfies  $w_v + x_v^{(t)} = 1$ . In the rest of the proof we focus on proving

$$\sum_{v \in C} x_v^{(t)} \leq \sum_{v \in C} (1 + 0.5\varepsilon)\mathbb{E}[z_v], \quad (21)$$

since it results in (20) as follows:

$$|C| = \sum_{v \in C} (w_v + x_v^{(t)}) \stackrel{(21)}{\leq} \sum_{v \in C} w_v + (1 + 0.5\varepsilon)\mathbb{E}[z_v] \leq \sum_{v \in C} (1 + 0.5\varepsilon)\mathbb{E}[y_v] = (1 + 0.5\varepsilon)\mathbb{E}\left[\sum_{v \in C} y_v\right].$$

To prove (21), let us start by noting that if we set  $z_e = \frac{x_e^{(t)}}{(1+\alpha)p}$  for any realized edge in  $S$ , we get  $\sum_{v \in C} x_v^{(t)} = (1 + \alpha) \sum_{v \in C} \mathbb{E}[z_v]$  and proves (21). However, in our algorithm we decrease  $z_e$  to zero

for the edges around any bad vertex and get  $z_v = 0$  for any such vertex. Thus, using Chebyshev's inequality, we will show that for any  $v$ , probability of being bad is small and as a result reducing the fractional value of the edges around these vertices does not affect the expected size of our matching significantly. Since for any  $v \in S$ , we have  $\mathbb{E}[z_v] = x_v^{(t)}/(1 + \alpha)$ , we get

$$\Pr[z_v > x_v^{(t)}] = \Pr[z_v > \mathbb{E}[z_v](1 + \alpha)] \leq \Pr[z_v - \mathbb{E}[z_v] > \alpha].$$

Since any edge gets value of  $x_e^{(t)}/(1 + \alpha)p$  with probability  $p$  (if realized) and zero otherwise, for any edge  $e \ni v$ , we have:

$$\text{Var}[z_e] = \mathbb{E}[z_e^2] - \mathbb{E}[z_e]^2 = p \left( x_e^{(t)}/(1 + \alpha)p \right)^2 - \left( x_e^{(t)}/(1 + \alpha) \right)^2 \leq \frac{\left( x_e^{(t)} \right)^2 (1 - p)}{p} \leq \frac{t^2}{p}.$$

Since edges are realized independently, this gives us

$$\text{Var}[z_v] = \sum_{e \ni v} \text{Var}[z_e] \leq \frac{1}{t} \frac{t^2}{p} = \frac{t}{p}.$$

Using Chebyshev's inequality, we have

$$\Pr[z_v - \mathbb{E}[z_v] > \alpha] \leq \frac{\text{Var}[z_v]}{\alpha^2} \leq \frac{t}{p\alpha^2}.$$

Now let us investigate the expected size of the fractional matching after decreasing  $z_e$  to zero for any edge  $e \ni v$  adjacent to a bad vertex  $v$ . We have

$$\mathbb{E} \left[ \sum_{v \in V} z_v \right] \geq \frac{1}{1 + \alpha} \sum_{v \in V} x_v^{(t)} \left( 1 - \frac{t}{p\alpha^2} \right). \quad (22)$$

By setting  $t = 0.25\alpha^2\varepsilon p = \Theta(\varepsilon^3 p)$ , we get

$$\mathbb{E} \left[ \sum_{v \in V} z_v \right] \geq \frac{1 - 0.25\varepsilon}{1 + 0.25\varepsilon} \sum_{v \in V} x_v^{(t)} \geq \frac{1}{1 + 0.5\varepsilon} \sum_{v \in V} x_v^{(t)}.$$

This gives us (21) and completes the proof.  $\square$

### 6.3 Bipartite Graphs: $(1 + \varepsilon)$ -Approximation with $O_p(1)$ Per-Vertex Queries

In this section, we will prove the following result:

**Theorem 6.3.** *For any constant  $\varepsilon > 0$  and constant  $p \in (0, 1]$ , any bipartite graph  $G = (V, E)$  has a constant degree subgraph  $Q$  where querying only the edges in  $Q$  suffices to find  $C \subseteq V$  such that*

1.  $C$  is a vertex cover of  $G_p$  with probability 1.
2. The size of  $C$  is in expectation at most  $(1 + \varepsilon)$  times the minimum vertex cover of  $G_p$ .
3. It is possible to find  $Q$  and  $C$  in polynomial time.

We start by the following lemma. While it is a folklore result, to be self-contained we also provide a proof for it.

**Claim 6.3.** Given  $M_1$  and  $M_2$ , two random matchings of graph  $G$  with  $\mathbb{E}[|M_2|] \geq \mathbb{E}[|M_1|]$ , let  $D$  denote their symmetric difference. For any  $\varepsilon \in (0, 1)$  graph  $D$  contains at least  $\mathbb{E}[|M_2|] - \mathbb{E}[|M_1|] - \frac{\varepsilon \mathbb{E}[|M_2|]}{2}$  maximal paths of length at most  $4/\varepsilon + 1$  who start and end with edges in  $M_2$ .

*Proof.* Any maximal path of graph  $D$  that starts and ends with edges in  $M_2$  is also called an augmenting paths. Since  $D$  is the symmetric difference of  $\mathbb{E}[|M_1|]$  and  $\mathbb{E}[|M_2|]$  it has  $\mathbb{E}[|M_2|] - \mathbb{E}[|M_1|]$  more edges from  $M_2$  compared to  $M_1$ . As a result, it contains at least  $\mathbb{E}[|M_2|] - \mathbb{E}[|M_1|]$  augmenting paths. Also, note that any augmenting paths longer than  $4/\varepsilon + 1$ , contains at least  $2/\varepsilon$  edges from  $M_2$ . Hence, there are most  $\mathbb{E}[|M_2|]/(2/\varepsilon)$  augmenting paths of length greater than  $4/\varepsilon + 1$ . As such,  $D$  contains at least

$$\mathbb{E}[|M_2|] - \mathbb{E}[|M_1|] - \frac{\varepsilon \mathbb{E}[|M_2|]}{2}$$

augmenting paths of length at most  $4/\varepsilon + 1$ .  $\square$

*Proof of Theorem 6.3.* We provide a reduction to approximate stochastic matchings. Suppose that we have a stochastic matching algorithm  $\mu_\delta$  that provides a  $(1 - \delta)$ -approximation via  $f(\delta, p)$  per-vertex queries – such algorithm exists as proved in [6] for

$$f(\delta, p) = \exp \left( \exp \left( \exp \left( O(\delta^{-1}) \right) \times \log \log p^{-1} \right) \right),$$

and takes polynomial time to run for constant  $\delta$  and  $p$ . We give a  $(1 + \varepsilon)$ -approximate stochastic minimum vertex cover algorithm that queries  $f(\frac{\varepsilon p^{2/\varepsilon+2}}{4}, p) = O_{\varepsilon, p}(1)$  edges per vertex.

**The Reduction:** For  $\delta = \frac{\varepsilon p^{2/\varepsilon+2}}{4}$ , let  $Q$  denote the set of edges queried by algorithm  $\mu_\delta$  and let  $S = E \setminus Q$ . We claim that querying set  $Q$  and picking a minimum vertex cover of  $(Q_P \cup S)$  is a  $(1 + \varepsilon)$ -approximate vertex cover algorithm. Let  $\nu_\varepsilon$  be the described algorithm. First, this clearly gives us a valid vertex cover of  $G_p$  as it covers all the realized edges of  $Q$  and all the edges that are not queried (i.e., edges in  $S$ ). Thus, to complete the proof we need to show

$$\mathbb{E}[\nu(Q_P \cup S)] \leq (1 + \varepsilon) \cdot \mathbb{E}[\nu(G_p)]. \quad (23)$$

For the sake of contradiction, we assume that this inequality does not hold and show that it implies  $\mathbb{E}[|\mu(Q_p)|] < (1 - \delta)\mathbb{E}[|\mu(G_p)|]$ , contradicting that  $\mu_\delta$  is a  $(1 - \delta)$ -approximate stochastic matching algorithm. Note that since  $G$  is bipartite, the size of its maximum matching and minimum vertex cover are equal by König's theorem. This implies that if (23) does not hold, then we have

$$\mathbb{E}[\mu(Q_P \cup S)] > (1 + \varepsilon)\mathbb{E}[\mu(G_p)]. \quad (24)$$

We will show that in this case, matching  $\mu(Q_p)$  can be augmented using edges in  $S_p$  to a matching whose size is larger than that of  $\mu(G_p)$  in expectation, which is a contradiction. Let  $D$  denote the symmetric difference of  $M_1 := \text{MM}(Q_p)$  and  $M_2 := \text{MM}(Q_P \cup S)$  where  $\text{MM}(\cdot)$  returns an arbitrary maximum matching. Namely,  $D$  contains an edge  $e$  if it is in exactly one of these matchings. By Claim 6.3, graph  $D$  in expectation contains at least  $\mathbb{E}[|M_2|] - \mathbb{E}[|M_1|] - \frac{\varepsilon \mathbb{E}[|M_1|]}{2}$  maximal paths of length at most  $\frac{4}{\varepsilon} + 1$  which start and end with edges in  $M_2$  (i.e., augmenting paths), where

$$\begin{aligned} \mathbb{E}[|M_2|] - \mathbb{E}[|M_1|] - \frac{\varepsilon \mathbb{E}[|M_1|]}{2} &\geq \mathbb{E}[|\mu(Q_P \cup S)|] - \mathbb{E}[|\mu(Q_p)|] - \frac{\varepsilon \mathbb{E}[|\mu(Q_p)|]}{2} \\ &\geq \mathbb{E}[|\mu(Q_P \cup S)|] - \mathbb{E}[|\mu(G_p)|] - \frac{\varepsilon \mathbb{E}[|\mu(G_p)|]}{2} \end{aligned}$$

$$\begin{aligned}
&\stackrel{(24)}{>} (1 + \varepsilon)\mathbb{E}[|\mu(G_p)|] - \mathbb{E}[|\mu(G_p)|] - \frac{\varepsilon\mathbb{E}[|\mu(G_p)|]}{2} \\
&> \frac{\varepsilon\mathbb{E}[|\mu(G_p)|]}{2}.
\end{aligned}$$

Consider one of these augmenting paths  $P$ . If all the edges of  $P$  are in  $G_p$ , flipping the membership of its edges in  $M_1$  increases the size of this matching by one. Note that any of these augmenting paths has at most  $\frac{2}{\varepsilon} + 2$  edges from  $S$  and these edges are realized (are in  $G_p$ ) independently with probability  $p$ . As a result, each one of these paths is in  $G_p$  with probability at least  $p^{\frac{2}{\varepsilon}+2}$ . Since we have at least  $\frac{\varepsilon\mathbb{E}[|\mu(G_p)|]}{2}$  of these paths, applying all of them on  $M_1$  results in increasing its expected size by at least

$$p^{\frac{2}{\varepsilon}+2} \cdot \frac{\varepsilon\mathbb{E}[|\mu(G_p)|]}{2} = 2\delta\mathbb{E}[|\mu(G_p)|].$$

This means that the resulting matching has size at least

$$\mathbb{E}|M_1| + 2\delta\mathbb{E}[|\mu(G_p)|] \geq (1 - \delta)\mathbb{E}[|\mu(G_p)|] + 2\delta\mathbb{E}[|\mu(G_p)|] > \mathbb{E}[|\mu(G_p)|],$$

which is a contradiction by the fact that  $\mathbb{E}[|\mu(G_p)|]$  is an upper bound for  $\mathbb{E}|M_1|$ .  $\square$

## 7 Lower Bounds

In this section, we prove several lower bounds for both the stochastic vertex cover problem and also the stochastic matching problem. Below we state these results as a series of theorems, and give their proofs later in the section.

**Theorem 7.1.** *For any constant  $p < 1$ , there exist an  $n$ -vertex bipartite stochastic graph  $G$  with realization probability  $p$  such that finding an exact minimum vertex cover or an exact maximum matching of  $G$  with any constant probability requires querying  $\Omega(\frac{n^2}{\log^2 n})$  edges of this graph.*

**Theorem 7.2.** *Finding a maximal matching of  $n$ -vertex stochastic graphs with a constant realization probability  $p \in (0, 1)$  requires  $\Omega(n \log_b n)$  total queries for  $b = \frac{1}{1-p}$ .*

**Theorem 7.3.** *There are absolute constants  $p_0, \varepsilon_0 \in (0, 1)$  such that for any  $p \leq p_0$  and  $\varepsilon \leq \varepsilon_0$ , finding a  $(1 - \varepsilon)$ -approximate maximum matching for a bipartite stochastic graph  $G_p$ , requires querying a subgraph of maximum degree  $\Omega(\frac{1}{\varepsilon p})$ .*

**Theorem 7.4.** *There are absolute constants  $p_0, \varepsilon_0 \in (0, 1)$  such that for any  $p \leq p_0$  and  $\varepsilon \leq \varepsilon_0$ , finding a  $(1 + \varepsilon)$ -approximate minimum vertex cover for a bipartite stochastic graph  $G_p$ , requires querying a subgraph of maximum degree  $\Omega(\frac{1}{\varepsilon p})$ .*

**Theorem 7.5.** *Finding a constant approximation of minimum vertex cover of  $n$ -vertex stochastic bipartite graphs with realization probability  $p$  requires  $\Omega(\frac{n}{p})$  total queries.*

A graph that is particularly useful for our lower bounds is illustrated in Figure 1 and defined formally in Definition 7.1.

**Definition 7.1** ( $S(d, s, N)$ -graphs – Figure 1). *For positive integers  $d, s$ , and  $N$ , an  $S(d, s, N)$ -graph is defined on  $2N(s+1)$  vertices. Of these nodes,  $2N$  form an induced  $d$ -regular bipartite graph  $B$ , with  $N$  nodes in each partition. In addition, each vertex of  $B$  is also connected to  $s$  vertices outside  $B$ , each with degree exactly 1. We use  $S$  to denote the set of vertices outside  $B$ .*

In proving our theorems we use, two essential properties of  $S(d, s, N)$ -graphs which we state below as Lemma 7.2 and Lemma 7.3.




Figure 1: An example of graph  $S(d, s, N)$  formalized in Definition 7.1 for  $d = 3$ ,  $s = 5$ , and  $N = 6$ .

**Lemma 7.2.** *Given a stochastic  $S(d, s, N)$ -graph  $G$  with realization probability  $p$ , consider its subgraphs  $B$  and  $S$  from Definition 7.1 and let  $H$  be the induced subgraph of vertices in  $B$  that do not have any realized edges to vertices in  $S$ . If  $M$  is a  $(1 - \varepsilon)$ -approximate maximum matching of  $G_p$ , it contains a matching of expected size at least  $|\mu(H_p)| - 2N\varepsilon$  from  $H$ .*

*Proof.* Let  $M_{\text{apx}}$  be an arbitrary  $(1 - \varepsilon)$ -approximate maximum matching of  $G_p$ . To prove this lemma, we will show

$$|M_{\text{apx}} \cap H| \geq |\mu(H_p)| - 2N\varepsilon.$$

We start by giving a lower bound for  $|M_{\text{apx}}|$ . We do so by constructing a matching of graph  $G_p$  which we denote by  $M_1$ . In  $M_1$ , we match any vertex  $v \in B$  that has a realized edge to a vertex in  $S$  (i.e., any vertex in  $B$  that is not in  $H$ ) using one of these edges. We also add a maximum matching of subgraph  $H_p$  to  $M_1$ . Let  $V_H$  and  $V_B$  respectively denote the vertex sets of graphs  $H$  and  $B$ . Observe that we have  $|M_1| = |V_B - V_H| + |\mu(H_p)|$ . Knowing that  $M_{\text{apx}}$  is a  $(1 - \varepsilon)$ -approximate maximum matching, gives us:

$$|M_{\text{apx}}| \geq (1 - \varepsilon)(|V_B| - |V_H| + |\mu(H_p)|) \geq |V_B| - |V_H| + |\mu(H_p)| - \varepsilon 2N.$$

We complete the proof by noting that the expected size of the maximum matching on  $G_p \setminus H_p$  is upper bounded by  $|V_B| - |V_H|$  since any edge in  $G_p \setminus H_p$  has an endpoint in  $V_B \setminus V_H$  and  $|V_B \setminus V_H| = |V_B| - |V_H|$ . This completes the proof since it indicates that the remaining  $|\mu(H_p)| - \varepsilon 2N$  edges of matching  $M_{\text{apx}}$  come from  $H_p$ .  $\square$

**Lemma 7.3.** *Given a stochastic  $S(d, s, N)$ -graph  $G$  with realization probability  $p$ , consider its subgraphs  $B$  and  $S$  from Definition 7.1 and let  $H$  be the induced subgraph of vertices in  $B$  that do not have any realized edges to vertices in  $S$ . Any  $(1 + \varepsilon)$ -approximate minimum vertex cover of  $G_p$ , in expectation includes at most  $|\nu(H_p)| + 2N\varepsilon$  vertices from  $H$ .*

*Proof.* Let  $\nu_{\text{apx}}$  be an arbitrary  $(1 + \varepsilon)$ -approximate minimum vertex cover of  $G_p$ . To prove this lemma, we will show

$$|\nu_{\text{apx}} \cap H_p| \leq \nu(H_p) + 2N\varepsilon.$$

We start by giving an upper bound for  $|\nu_{\text{apx}}|$ . We do so by constructing a vertex cover of graph  $G_p$  which we denote by  $\nu_1$ . This vertex cover includes any vertex  $v \in B$  that has a realized edge to a vertex in  $S$  (i.e., any vertex in  $B$  that is not in  $H$ ). This covers all the edges that are not in  $H$ . Thus, we complete  $\nu_1$  by adding a minimum vertex cover of  $H_p$  to it. Let  $V_H$  and  $V_B$  respectively



denote the vertex sets of graphs  $H$  and  $B$ . Observe that we have  $|\nu_1| = |V_B - V_H| + |\nu(H_p)|$ . As a result of  $\nu_{apx}$  being a  $(1 + \varepsilon)$ -approximate minimum vertex cover we get:

$$|\nu_{apx}| \leq (1 + \varepsilon)(|V_B| - |V_H| + |\nu(H_p)|) \leq |V_B| - |V_H| + |\nu(H_p)| + \varepsilon 2N.$$

We complete the proof by noting that  $|V_B| - |V_H|$  is a lower bound for the number of vertices in  $V_B \setminus V_H$  that are in vertex cover since any vertex in  $V_B \setminus V_H$  is connected to at least a degree-one vertex. This concludes the proof as it indicates that at most  $|\mu(H_p)| + \varepsilon 2N$  vertices of  $\nu_{apx}$  can come from  $H_p$ .  $\square$

## 7.1 Proof of Theorem 7.1

We start the proof by the following Lemma.

**Lemma 7.4.** *For any constant  $p < 1$ , there exist an  $n$ -vertex bipartite stochastic graph  $G$  with realization probability  $p$  and a constant  $c \geq 0$  such that finding an exact maximum matching or an exact minimum vertex cover of  $G$  with probability at least  $c$  requires querying  $\Omega(\frac{n^2}{\log^2 n})$  edges of this graph.*

*Proof.* To prove this lemma, we start by an  $n$ -vertex graph  $G$ , a  $S(N, s, N)$ -graph for  $s = \log_{1-p} 1/N$ . We then prove the existence of a constant  $c$  for which the statement of the lemma holds. Particularly, we will show that it is not possible to find an exact maximum matching/minimum vertex cover of  $G_p$  with probability at least  $c$  using only  $o(n^2/\log^2 n)$  queries.

Consider subgraphs  $B$  and  $S$  of  $G$  from Definition 7.1 and let  $H$  be the induced subgraph of vertices in  $B$  that do not have any realized edges to vertices in  $S$ . Based on Lemma 7.2, any maximum matching of  $G_p$  should include a maximum matching of  $H_p$ . Also, based on Lemma 7.3, any minimum vertex cover of  $G_p$ , includes exactly  $|\nu(H_p)|$  vertices from  $H_p$ . Thus, to be able to find an exact minimum vertex cover/maximum matching of  $G_p$  we need to find a minimum vertex cover/maximum matching of  $H_p$ . Thus, we will focus on graph  $H_p$ .

Since  $s = \Theta(\log N)$ , we have  $n = \Theta(N \log N)$  and  $N = \Omega(n/\log n)$  (based on Definition 7.1). This implies that if at most  $o(n^2/\log^2 n)$  edges are queried from  $G$ , then any edge  $e$  chosen uniformly at random from  $B$  (defined in Definition 7.1) is queried with probability  $o(1)$  (since  $B$  has  $N^2 = \Omega(n^2/\log^2 n)$  edges). Therefore, it suffices to show that if any random edge from  $B$  is queried with probability  $o(1)$ , then it is not possible to find an exact maximum matching/minimum vertex cover of  $H_p$  with probability at least  $c$  for a constant  $c$ .

We claim that with a constant probability subgraph  $H_p$  has only one vertex in each part. Let us denote this event by  $E$  and compute its probability. Since each vertex  $v \in B$  is in  $H$  independently with probability  $(1 - p)^s = 1/N$ , we have

$$\Pr[E] = (N(1/N)(1 - 1/N)^{N-1})^2 \geq 1/e^2.$$

When event  $E$  happens; i.e.,  $H$  has only one vertex in each part, to be able to find its maximum matching, if edge  $(u, v)$  is realized it also should be queried. This edge is realized with probability  $p$ , however, as discussed above, probability of this edge being queried is  $o(1)$ . This implies that with probability at least  $1/e^2(p - o(1))$ , which is a constant, the queried edges do not contain a maximum matching of  $G_p$ . Similarly, if edge  $(u, v)$  is not queried, to have a valid vertex cover, we need to put one of its endpoints in the vertex cover, however, this vertex cover is not minimum if edge  $(u, v)$  is not realized. This event happens with probability  $1/e^2(1 - p - o(1))$  which is again a constant. As a result, setting  $c = \min(1/e^2(1 - p - o(1)), 1/e^2(p - o(1)))$  completes the proof.  $\square$



Given Lemma 7.4, we are now ready to prove Theorem 7.1.

*Proof of Theorem 7.1.* We use proof by contradiction. We first assume the existence of a constant  $\alpha$  such that given any  $n$ -vertex stochastic graph  $G$ , it is possible to find an exact minimum vertex cover/maximum matching of this graph with probability at least  $\alpha$  using only  $o(n^2/\log^2 n)$  queries. We then show that this results in a contradiction. Based on Lemma 7.4, for any  $N$  there exists an  $N$ -vertex bipartite stochastic graph  $G'$  with realization probability  $p$  and a constant  $c \geq 0$  such that finding an exact minimum vertex cover or an exact maximum matching of  $G'$  with probability at least  $c$  requires  $\Omega(N^2/\log^2 N)$  queries. Define graph  $G$  to include  $s = \lceil \log_{1-c}(1-\alpha) \rceil$  copies of  $G'$  with  $N = n/s$ . Since  $n = \Theta(N)$ , querying  $o(\frac{n^2}{\log^2 n})$  edges of  $G$ , means queries  $o(\frac{N^2}{\log^2 N})$  edges of each subgraph. Observe that finding an exact maximum matching/minimum vertex cover of  $G$  means finding an exact maximum matching/minimum vertex cover of all these  $s$  subgraphs. This implies that the queried edges of  $G$  contain an exact maximum matching/minimum vertex cover of  $G_p$  with probability at most  $1 - (1-c)^s \leq 1 - (1-\alpha) \leq \alpha$ . We conclude the proof of this theorem by noting that this is a contradiction with the initial assumption.  $\square$

## 7.2 Proof of Theorem 7.2

*Proof of Theorem 7.2.* Consider an  $n$ -vertex clique  $G$  with realization probability  $p$ . It is known that for any constant  $p \in (0, 1)$ , the expected size of the maximum independent set in  $G_p$  is  $\Theta(\log_b n)$  for  $b = 1/(1-p)$ . (See Theorem 7.3 in Book [13].) This is useful for giving a lower bound for the size of any maximal matching  $M$  of  $G_p$  since vertices that are not in  $M$  form an independent set. This implies  $2|M| \geq n - |\text{MIS}(G_p)|$  where  $\text{MIS}(G_p)$  is the maximum independent set of  $G_p$ . Let  $Q$  be the subgraph we choose to query. If  $Q_p$  contains a maximal matching then its maximum independent set should not be larger than that of  $G_p$ . Therefore, to complete the proof, it suffices to show that if

$$|Q_p| = o(n \log_b n) < \frac{n \log_b n}{4}, \quad (25)$$

then  $\mathbb{E}[|\text{MIS}(Q_p)|] = \omega(\log_b n)$ . To prove this, we will show that the expected number of singleton vertices in  $Q_p$  is  $\omega(\log_b n)$ . Note that based on (25), at least half of vertices have degree at most  $\frac{\log_b n}{2}$  in  $Q$ . For any such vertex  $v$  we have:

$$\Pr[v \text{ is singleton in } Q_p] \geq (1-p)^{\frac{\log_b n}{2}} = (1/b)^{\frac{\log_b n}{2}} = \left(\frac{1}{n}\right)^{\frac{1}{2}}.$$

For any constant  $p \in (0, 1)$ , this gives us:

$$\mathbb{E}[|\text{MIS}(Q_p)|] \geq \mathbb{E}[\text{number of singleton vertices in } Q_p] \geq \frac{n}{2} \cdot \left(\frac{1}{n}\right)^{\frac{1}{2}} = \frac{\sqrt{n}}{2} = \omega(\log_b n),$$

and completes the proof.  $\square$

## 7.3 Proof of Theorem 7.3

*Proof of Theorem 7.3.* Consider the bipartite graph  $G = S(N, s, N)$  for  $s = \lfloor \log_{1-p} 10\varepsilon \rfloor$  (given that  $10\varepsilon \leq 1-p$  holds for small enough values of  $\varepsilon$  and  $p$ ), and let  $n$  be the total number of vertices in  $G$ . To prove this theorem, we show that finding a  $(1-\varepsilon)$ -approximate maximum matching of  $G_p$  requires  $\Omega(\frac{1}{\varepsilon p})$  queries for a vertex. Consider subgraphs  $B$  and  $S$  of  $G$  from Definition 7.1 and let  $H$  be the induced subgraph of vertices in  $B$  that do not have any realized edges to vertices in

$S$ . Based on Lemma 7.2, any  $(1 - \varepsilon)$ -approximate maximum matching of  $G_p$ , includes a matching of expected size at least  $|\mu(H_p)| - 2N\varepsilon$  from  $H$ . Let  $X_1$  and  $X_2$  with be the set of vertices in two parts of the bipartite graph  $H$  with  $|X_1| \leq |X_2|$ . Note that the expected number of vertices in each part is

$$(1 - p)^{\lfloor \log_{1-p} 10\varepsilon \rfloor} N \geq 10\varepsilon N;$$

thus,  $\mathbb{E}[|X_1|] \geq 5\varepsilon N$ . It is well-known that such random graph  $H_p$  that has an edge between any two vertices in  $X_1$  and  $X_2$  with a constant probability  $p$  has a matching of size  $|X_1|$  with high probability. However, for the sake of this proof we only use the fact that a matching of size  $|X_1|$  exists with probability at least  $\frac{4}{5}$ . This gives us  $\mathbb{E}[|\mu(H_p)|] \geq 4\varepsilon N$ . As a result, to find a matching of size at least  $|\mu(H_p)| - 2\varepsilon N$  in  $H_p$ , any randomly chosen vertex  $v \in X_1$  should have at least one queried edge in  $H_p$  with probability at least  $\frac{1}{2}$ . Let  $g_v$  be the expected number of edges queried for vertex  $v$ . Since the other endpoint of any of these queried edges is in  $H_p$  with probability

$$(1 - p)^{\lfloor \log_{1-p} 10\varepsilon \rfloor} \leq 10\varepsilon/(1 - p),$$

and  $e$  is realized with probability  $p$ , we should have  $\frac{10\varepsilon p g_v}{1 - p} \geq 1/2$  which for any  $p \leq 0.99$  gives us  $g_v \geq \frac{1-p}{20\varepsilon p} = \Omega(\frac{1}{\varepsilon p})$  and concludes the proof.  $\square$

## 7.4 Proof of Theorem 7.4

*Proof of Theorem 7.4.* Consider the bipartite graph  $G = S(d, s, N)$  for

$$s = \lfloor \log_{1-p} 96\varepsilon \rfloor \quad \text{and} \quad d = 1/(2^{15}\varepsilon p),$$

and let  $n$  be the total number of vertices in  $G$ . To prove this theorem, we show that it is not possible to find a  $(1 + \varepsilon)$ -approximate minimum vertex cover of  $G_p$  using only  $o(1/\varepsilon p)$  queries per vertex. We use proof by contradiction. We assume the existence of a  $(1 + \varepsilon)$ -approximate minimum vertex  $\nu_{apx}$ , achieved by querying only  $o(1/\varepsilon p)$  edges per vertex and show that it results in a contradiction. Consider subgraphs  $B$  and  $S$  of  $G$  from Definition 7.1, and let  $H$  be the induced subgraph of vertices in  $B$  that do not have any realized edges to vertices in  $S$ . Based on Lemma 7.3, any  $(1 + \varepsilon)$ -approximate minimum vertex cover of  $G_p$ , includes at most  $|\nu(H_p)| + 2N\varepsilon$  vertices from  $H$ , which means

$$|\nu_{apx} \cap H| \leq |\nu(H_p)| + 2N\varepsilon.$$

Let us define a subgraph  $H'$  of  $H$  to contain all the edges of  $H$  that are not queried. To be a valid vertex cover,  $\nu_{apx}$  should include a minimum vertex cover of  $H'$ ; i.e.,

$$|\nu(H')| \leq |\nu_{apx} \cap H|.$$

Thus, to achieve a contradiction, it suffices to prove that the following equation does not hold if any vertex has  $d' = 1/(2^{16}\varepsilon p) + 1$  edges in  $B$  that are not queried:

$$|\nu(H')| \leq |\nu(H_p)| + 2N\varepsilon. \quad (26)$$

Let us assume w.o.l.g., that exactly  $d'$  edges are not queried for any vertex in  $B$  as it only decreases the size of the minimum vertex cover. We start by giving a lower bound for  $|\nu(H')|$ . We do so by constructing a fractional matching  $M$  on  $H'$ . Since by weak duality, the minimum vertex cover of a graph is at least as large as any fractional matching of the graph, this will also be a lower bound for  $|\nu(H')|$ . Define  $x := (d' - 1)(1 - p)^s + 1$ . For any edge  $e$ , let  $m_e$  be the fractional value we assign

to edge  $e$  in matching  $M$ . We set  $m_e := 1/x$  for any edge  $e \in H'_p$  iff both its endpoints have degree at most  $x$  in  $H'_p$  and zero otherwise.  $M$  is obviously a valid fractional matching since the sum of values assigned to the edges around any vertex is at most one. Also, for any edge  $e = (u, v) \in B$ , we have

$$\mathbb{E}[m_e] = \Pr[e \in H'] \cdot \Pr[d_{v,H'} \leq x \mid e \in H'] \cdot \Pr[d_{u,H'} \leq x \mid e \in H'], \quad (27)$$

where  $d_{v,H'}$  is the degree of vertex  $v$  in graph  $H'$ . To compute  $\Pr[e \in H']$ , note that each endpoint of this edge is in  $H$  with probability  $(1-p)^s$ . Combining this with  $s = \lfloor \log_{1-p} 96\varepsilon \rfloor$  gives us

$$(96\varepsilon)^2 \leq \Pr[e \in H']. \quad (28)$$

Consider any vertex  $v \in H'$  and any of its edges  $e \in H'$ , and let  $d_{v,e,H'}$  denote the degree of vertex  $v$  in  $H' \setminus \{e\}$ . Since  $v$  has  $d' - 1$  edges in  $B \setminus \{e\}$ , we have  $\mathbb{E}[d_{v,e,H'}] = (d' - 1)(1-p)^s = x - 1$ . Moreover, we claim that as a result of  $d_{v,e,H'}$  being sum of independent Bernoulli random variables, we have  $\Pr[d_{v,e,H'} > x - 1] < 0.75$ . This can be achieved using a simple application of Chernoff bound, and implies

$$\Pr[d_{v,H'} \leq x \mid e \in H'] \geq \frac{1}{4}.$$

Combining this with (28) and (27) gives us:

$$|M| \geq Nd' \frac{1}{x} (96\varepsilon)^2 \frac{1}{16} \geq \frac{N96^2\varepsilon^2}{2^{16}\varepsilon p x 16}.$$

Observe that for a small enough  $p$ , we have

$$x = (d' - 1)(1-p)^s + 1 \leq \frac{96}{2^{16}p(1-p)} + 1 \leq \frac{2 \cdot 96}{2^{16}p(1-p)}.$$

This implies:

$$|M| \geq \frac{N96^2\varepsilon^2}{16 \cdot 2^{16}\varepsilon p(1-p)} \frac{2^{16}p(1-p)}{2 \cdot 96} \geq \frac{N96\varepsilon}{32}. \quad (29)$$

After finding a lower bound for  $|M|$  which is also a lower bound for  $|\nu(H')|$  we need to find an upper bound for  $|\nu(H_p)|$ . Recall that  $H$  is the induced subgraph of vertices in  $B$  that do not have any realized edges to vertices in  $S$ . To give an upper bound for  $|\nu(H_p)|$ , we start by computing the expected number of edges in this graph which we denote by  $|H_p|$ . Each edge  $e \in B$  is in  $H_p$  with probability  $p(1-p)^2s \leq p96^2\varepsilon^2/(1-p)^2$ . Since  $B$  has  $N/(2^{15}\varepsilon p)$  edges, for a large enough  $(1-p)$  we get

$$|H_p| \leq \frac{N96^2\varepsilon}{2^{15}(1-p)^2} \leq \frac{2N96^2\varepsilon}{2^{15}}.$$

This is also an upper bound for  $|\nu(H_p)|$  since minimum vertex cover is not larger than the total number of edges. Based on (26), to complete the proof we need to show that

$$\frac{N96\varepsilon}{32} \geq \frac{2N96^2\varepsilon}{2^{15}} + 2N\varepsilon.$$

If we simplify both sides of the equation we get

$$\frac{N96\varepsilon}{32} = 3N\varepsilon \geq 2.6N\varepsilon \geq \frac{2N96^2\varepsilon}{2^{15}} \geq +2N\varepsilon,$$

which concludes the proof.  $\square$

## 7.5 Proof of Theorem 7.5

*Proof of Theorem 7.5.* Suppose we want to get a  $c$ -approximate vertex cover for a  $c > 1$ . For any  $p < \frac{1}{3c}$ , we construct a graph  $G$  for which finding such an approximation needs at least  $\Omega(n/p)$  queries. Graph  $G$  is an arbitrary  $d$ -regular bipartite graph with  $d = \frac{1}{3cp}$ . (Here we assume for simplicity that  $d$  is an integer.) In  $G_p$ , any vertex  $v$  is a singleton with probability at least  $1 - dp$  which means  $dn/2$  is an upper bound for the expected size of the vertex cover.

On the other hand, if we do not query an edge, we have to put one of its endpoints in the vertex cover. We claim that to get a  $c$ -approximate at least half of the edges ( $dn/4$ ) should be queried. Since any vertex can cover at most  $d$  edges, not querying  $dn/4$  edges results in a vertex cover of size at least  $n/4$ . Thus, if at most  $dn/4$  edges are queried, the approximation factor of the algorithm would be at least  $\frac{n/4}{dn/2} = \frac{1}{2dp} = \frac{3c}{2} \geq c$ . To conclude, to be able to get a  $c$ -approximate at least  $\frac{n}{6cp} = \Omega(n/cp)$  (half the number of edges) queries are needed.  $\square$

## 7.6 Why Random Queries Do Not Work

One simple approach that, in the first sight, might seem appealing for finding approximate vertex covers of a stochastic graph  $G$  is simply picking a random subset of edges of any vertex to form a subgraph  $Q$ , query these edges and find the vertex cover based on that. In this section, we will show that this approach, in fact, does not have a good performance for some instances of the problem. Namely, we show that there exists a graph  $G$  such that querying a subgraph  $Q$  of  $G$  obtained via sampling  $s = o(n)$  edges per any vertex does not give us better than a  $(\frac{1}{p} - o(1))$ -approximate vertex cover of  $G_p$ . Note that to find a valid vertex cover, any edge in  $G$  that is not queried should be covered; i.e., edges in  $S := G \setminus Q$ . Therefore, the final vertex cover found via this sampling technique has size at least  $|\nu(S \cup Q_p)|$  in expectation.

We construct an  $n$ -vertex bipartite graph  $G$  as follows. For a number  $N$  with  $N = o(n)$  and  $N = \omega(s)$ , graph  $G$  contains a complete bipartite graph  $H$  of  $N$  vertices and a matching  $M$  of size  $(n - N)/2$ . Moreover, for any  $e = (u, v) \in M$ , vertex  $u$  has edges to all the vertices in one part of  $H$  and  $v$  has edges to all the vertices in the other part of  $H$ . Note that such a number  $N$  exists since  $s = o(n)$ . Let us first give an upper bound for  $|\nu(G_p)|$  by constructing a vertex cover  $C$  of this graph.  $C$  contains all the vertices of graph  $H$  and an endpoint of any edge in  $M$  that is realized (an endpoint of any edge in  $M \cup G_p$ ). This set clearly covers all the edges of  $G_p$ . Since any edge in  $M$  is realized with probability  $p$ , we have

$$\mathbb{E}[|\nu(G_p)|] \leq \mathbb{E}[|C|] = N + \frac{p(n - N)}{2} \leq pn/2 + N/2 = pn/2 + o(n). \quad (30)$$

Having established this upper bound, the next step is to give a lower bound for  $|\nu(S \cup Q_p)|$ . Consider an edge  $e = (u, v) \in M$ , and let us compute  $\Pr[e \in Q]$ . Both endpoints of this edge have degree  $N/2 + 1$  in graph  $G$ . Since each vertex randomly chooses  $s$  edges, for the probability of  $e$  being sampled we have:

$$\Pr[e \in Q] \leq \frac{2s}{N/2 + 1} = o(1).$$

As a result the expected number of edges sampled from matching  $M$  is at most  $(1 - o(1))|M|$ . Since any edge that is not queried (sampled) should be covered, any valid vertex cover of  $S \cup Q_p$  should contain at least one endpoint of edges in  $M$  that are not sampled. Hence, we have

$$\mathbb{E}[|\nu(S \cup Q_p)|] \geq (1 - o(1))|M| = (1 - o(1))(n - N)/2 = n/2 - o(n).$$

Combining this with (30) gives us:

$$\frac{\mathbb{E}[|\nu(S \cup Q_p)|]}{\mathbb{E}[|\nu(G_p)|]} \geq \frac{n/2 - o(n)}{pn/2 + o(n)} = \frac{1}{p} - o(1).$$

This proves our claim as it means that by randomly sampling  $o(n)$  edges for any vertex in graph  $G$ , one cannot construct a valid vertex cover of  $G_p$  with an approximation ratio smaller than  $(1/p - o(1))$ .

## References

- [1] Sepehr Assadi and Aaron Bernstein. Towards a Unified Theory of Sparsification for Matching Problems. In *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019, January 8-9, 2019 - San Diego, CA, USA*, pages 11:1–11:20, 2019.
- [2] Sepehr Assadi, Sanjeev Khanna, and Yang Li. The Stochastic Matching Problem with (Very) Few Queries. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 43–60, 2016.
- [3] Sepehr Assadi, Sanjeev Khanna, and Yang Li. The Stochastic Matching Problem: Beating Half with a Non-Adaptive Algorithm. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 99–116, 2017.
- [4] Soheil Behnezhad and Mahsa Derakhshan. Stochastic weight matching:  $(1 - \epsilon)$ -approximation. In *Foundations of Computer Science (FOCS 20)*, to appear, 2020.
- [5] Soheil Behnezhad, Mahsa Derakhshan, Alireza Farhadi, MohammadTaghi Hajiaghayi, and Nima Reyhani. Stochastic Matching on Uniformly Sparse Graphs. In *Algorithmic Game Theory - 12th International Symposium, SAGT 2019, Athens, Greece, September 30 - October 3, 2019, Proceedings*, pages 357–373, 2019.
- [6] Soheil Behnezhad, Mahsa Derakhshan, and MohammadTaghi Hajiaghayi. Stochastic Matching with Few Queries:  $(1 - \epsilon)$  Approximation. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, to appear, 2020.
- [7] Soheil Behnezhad, Alireza Farhadi, MohammadTaghi Hajiaghayi, and Nima Reyhani. Stochastic Matching with Few Queries: New Algorithms and Tools. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2855–2874, 2019.
- [8] Soheil Behnezhad and Nima Reyhani. Almost Optimal Stochastic Weighted Matching with Few Queries. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 235–249, 2018.
- [9] Avrim Blum, John P. Dickerson, Nika Haghtalab, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is Almost Bliss: Near-Optimal Stochastic Matching With Few Queries. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 325–342, 2015.
- [10] Avrim Blum, John P. Dickerson, Nika Haghtalab, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance Is Almost Bliss: Near-Optimal Stochastic Matching with Few Queries. *Operations Research*, 68(1):16–34, 2020.

- [11] Charles J Colbourn. *The combinatorics of network reliability*. Oxford University Press, Inc., 1987.
- [12] Devdatt P Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25), 1996.
- [13] Alan Frieze and Michał Karoński. *Introduction to random graphs*. Cambridge University Press, 2016.
- [14] Michel X. Goemans and Jan Vondrák. Covering minimum spanning trees of random subgraphs. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004*, pages 934–941, 2004.
- [15] Michel X. Goemans and Jan Vondrák. Covering minimum spanning trees of random subgraphs. *Random Struct. Algorithms*, 29(3):257–276, 2006.
- [16] Heng Guo and Mark Jerrum. A polynomial-time approximation algorithm for all-terminal network reliability. *SIAM J. Comput.*, 48(3):964–978, 2019.
- [17] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- [18] David R. Karger. A phase transition and a quadratic time unbiased estimator for network reliability. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 485–495, 2020.
- [19] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within  $2 - \varepsilon$ . *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.
- [20] Alam Khursheed and KM Lai Saxena. Positive dependence in multivariate distributions. *Communications in Statistics-Theory and Methods*, 10(12):1183–1196, 1981.
- [21] Jan Vondrák. Shortest-path metric approximation for random subgraphs. *Random Struct. Algorithms*, 30(1-2):95–104, 2007.
- [22] Yutaro Yamaguchi and Takanori Maehara. Stochastic Packing Integer Programs with Few Queries. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 293–310, 2018.

## A Deferred Proofs

### A.1 Polynomial-Time Implementation of Algorithm 2

Algorithm 2 which proves Theorem 6.1 relies on a partitioning of Lemma 4.1. The described proof for this lemma, as stated above, is not through a polynomial-time construction. In this section, we address this issue and explain how the guarantee of Lemma 4.1 can also be achieved in polynomial (randomized) time, which leads to Algorithm 2 running in polynomial-time.

The reason that our algorithm for finding the partitioning  $(Q, S)$  of Lemma 4.1 is not polynomial-time, is that we assume each matching algorithm  $\mathcal{M}_i$  maximizes the quadratic objective (4) which



is not clear how to do in polynomial time. Here, however, we show how it is possible to get around this barrier and achieve the same guarantee in polynomial time too.

Suppose that we define partitionings  $(Q_i, S_i)$  as before, but instead of using a matching algorithm  $\mathcal{M}_i$  for each which maximizes  $\Phi_i(\mathcal{M}_i)$ , we use simply an arbitrary maximum matching algorithm  $\mathcal{M}'_i$ . What can go wrong? We used the assumption that  $\mathcal{M}_i$  maximizes  $\Phi_i(\mathcal{M}_i)$  in Claim 4.4 which led to Claim 4.5, proving existence of  $\Omega(\varepsilon^6 p^3 k)$  partitionings where every pair of them achieve the same objective up to an additive  $O(\varepsilon^6 p^3 \mu(G))$  difference. It was also used crucially in the proof of Claim 4.6, where we constructed an algorithm  $\mathcal{M}_{i,j}$  which returns either the output of  $\mathcal{M}_i(H_i)$  or that of  $\mathcal{M}_j(H_j)$  each with probability  $1/2$ . There we argued that if the desired bound of Claim 4.6 does not hold, then  $\Phi_i(\mathcal{M}_{i,j}) > \Phi_i(\mathcal{M}_i) + 0.01\varepsilon^6 p^3 \mu(G)$  (see (12)), contradicting the assumption that  $\Phi_i(\mathcal{M}_i)$  is the maximum possible achievable objective. However, in order for this to lead to a contradiction, we do not necessarily need  $\Phi_i(\mathcal{M}_i)$  to have the maximum possible value. Rather, it is sufficient to merely guarantee  $\Phi_i(\mathcal{M}_{i,j}) < \Phi_i(\mathcal{M}_i) + 0.01\varepsilon^6 p^3 \mu(G)$  for all  $i < j$ . Also in order to guarantee Claim 4.5, it suffices to have, say,  $\Phi_i(\mathcal{M}_i) \geq \Phi_j(\mathcal{M}_j) \pm O((\varepsilon p)^{10})$  for all  $i < j$ .

To achieve this guarantee, we first present in Claim A.1 a polynomial-time randomized algorithm for estimating  $\Phi_i(\mathcal{M}')$  with  $O(1)$  additive error, for any given polynomial-time algorithm  $\mathcal{M}'$ .

**Claim A.1.** *Given any matching algorithm  $\mathcal{M}'$ , it is possible to estimate the value of  $\Phi_i(\mathcal{M}')$  (for any  $i$ ) with  $O(1)$  additive error in polynomial time, with high probability.*

*Proof.* For each edge  $e$ , let  $p_e := \Pr[e \in \mathcal{M}'(H_i)]$  and recall from (4) that  $\Phi_i(\mathcal{M}') = \sum_{e \in E} p_e - \varepsilon p_e^2$ . We do not know the value of  $p_e$  but can estimate it by random sampling. Take  $t = n^6$  independent outputs  $M_1, \dots, M_t$  of the random matching  $\mathcal{M}'(H_i)$ . Let  $X_e := \sum_{i=1}^t \mathbf{1}(e \in M_i)$  be the number of these matchings that include  $e$ . Since  $\mathbb{E}X_e = t \cdot p_e$ , a simple application of Chernoff bound gives

$$\Pr[|X_e - tp_e| \geq \sqrt{2t \ln n}] \leq 2e^{-4 \ln n} \leq 2n^{-4}.$$

Therefore, by letting  $q_e := \frac{1}{t}X_e$  and a union bound over the less than  $n^2/2$  choices of  $e$ , we get that with probability at least  $1 - n^{-2}$ , for any edge  $e$  it holds that  $|q_e - p_e| < \frac{\sqrt{2t \ln n}}{t} < \frac{1}{n^2}$ .

Using this estimator  $q_e$  instead of  $p_e$ , with probability  $1 - n^{-2}$  we get the following estimator with constant additive error

$$\sum_{e \in E} q_e - \varepsilon q_e^2 = \Phi_i(\mathcal{M}') \pm O\left(n^2 \cdot \frac{1}{n^2}\right) = \Phi_i(\mathcal{M}') \pm O(1),$$

completing the proof.  $\square$

Having this estimator, we then use an arbitrary maximum matching algorithm  $\mathcal{M}'_i$  for each partitioning  $(Q_i, S_i)$  and based on that construct the next partitioning  $(Q_{i+1}, S_{i+1})$ . Then for a margin  $\delta = \Theta((\varepsilon p)^{10} \mu(G))$ , if it happens for some  $j > i$ , that our estimator predicts  $\Phi_j(\mathcal{M}'_j) > \Phi_i(\mathcal{M}'_i) + \delta$  or  $\Phi_i(\mathcal{M}'_{i,j}) > \Phi_i(\mathcal{M}'_i) + \delta$  (where  $\Phi_i(\mathcal{M}'_{i,j})$  returns the matching  $\Phi_i(\mathcal{M}'_i)$  with probability  $1/2$  and  $\Phi_j(\mathcal{M}'_j)$  otherwise), we simply use that algorithm instead of  $\mathcal{M}'_i$  for  $i$ . Note that using this new algorithm may cause  $D_i$  to change and so we may need to re-construct the the partitionings  $(Q_{i+1}, S_{i+1}), \dots, (Q_k, S_k)$ .

Finally, we argue why this process stops after polynomially many iterations. Every time that we change the matching algorithm of a partitioning  $(Q_i, S_i)$ , its objective  $\Phi_i$  (and not just its estimation) increases by  $O((\varepsilon p)^{10} \mu(G))$ . On the other hand, by definition (4), the value of  $\Phi_i$  for any  $i$  is upper bounded by the maximum matching  $\mu(G)$  of  $G$ . Hence, this  $\Theta((\varepsilon p)^{10} \mu(G))$  increase



in  $\Phi_i$  can only happen for at most  $O((\varepsilon p)^{-10})$  steps for any fixed partitioning  $(Q_i, S_i)$ . This, in turn, implies that the total number of changing the matching algorithm for any of the  $k$  partitionings is bounded by  $k^{O(1/\varepsilon^{10}p^{10})} = O_{\varepsilon,p}(1)$ . As a result, the whole process takes  $O_{\varepsilon,p}(1) \cdot \text{poly}(n)$  time.

## A.2 Proof of Claim 5.8

In this section, we prove Claim 5.8. We start with the notation we use in the proof.

**Notation:** We fix an arbitrary vertex  $v \in A$  and prove Claim 5.8 for it. We use  $u_1, \dots, u_d$  to denote the neighbors of  $v$  in  $S$ , and denote  $e_j = (u_j, v)$ . We also let  $A' = \{v_1, \dots, v_{n'}\} = A \setminus \{v\}$  be the set of all vertices in  $A$  excluding  $v$ .

We start with an auxiliary claim that will be helpful both in bounding the expected value of random variable  $(y_v \mid \overline{v \text{ prop}})$  and also proving a concentration bound for it.

**Claim A.2.** *It holds that  $\frac{1}{\Pr[\overline{v \text{ prop}}]} \sum_{i=1}^d q_{e_i} \leq q_v$ .*

*Proof.* Observe from Corollary 5.5 Property (ii) that  $\Pr[v \text{ prop}] = (1 - \varepsilon) \Pr[v \in \mathcal{M}(Q_p)]$ . Thus,

$$\Pr[\overline{v \text{ prop}}] = 1 - \Pr[v \text{ prop}] = 1 - (1 - \varepsilon) \Pr[v \in \mathcal{M}(Q_p)] \geq 1 - \Pr[v \in \mathcal{M}(Q_p)].$$

As a result, we get

$$\frac{\sum_{i=1}^d q_{e_i}}{\Pr[\overline{v \text{ prop}}]} \leq \frac{\sum_{i=1}^d q_{e_i}}{1 - \Pr[v \in \mathcal{M}(Q_p)]}.$$

Let  $q_v^S$  denote the sum of  $q_e$ 's written on edges  $e \in S$  connected to  $v$  and let  $q_v^Q$  denote the same but on edges of  $v$  in  $Q$ . Observe that the nominator of the fraction above is exactly  $q_v^S$  and the denominator is  $1 - q_v^Q$  by the first assumption of Lemma 5.1. Combined with  $q_v^S + q_v^Q = q_v$  and  $q_v \leq 1$  (since  $\mathbf{q}$  is a valid fractional matching by Lemma 5.1) we get

$$\frac{\sum_{i=1}^d q_{e_i}}{1 - \Pr[v \in \mathcal{M}(Q_p)]} = \frac{q_v^S}{1 - q_v^Q} = \frac{q_v \cdot q_v^S}{q_v(1 - q_v^Q)} = \frac{q_v \cdot q_v^S}{q_v - q_v \cdot q_v^Q} \leq \frac{q_v \cdot q_v^S}{q_v - q_v^Q} = \frac{q_v \cdot q_v^S}{q_v^S} = q_v, \quad (31)$$

which is the stated bound.  $\square$

Let us first bound the expected value of  $y_v$  conditioned on event  $(\overline{v \text{ prop}})$ .

**Claim A.3.** *Let  $v$  be as above. Then  $\mathbb{E}[y_v \mid \overline{v \text{ prop}}] \leq q_v$ .*

*Proof.* We have

$$\mathbb{E}[y_v \mid \overline{v \text{ prop}}] = \sum_{i=1}^d \frac{q_{e_i}}{p \Pr[u_i \notin M_{\mathcal{B}}] \Pr[\overline{v \text{ prop}}]} \cdot \Pr[e_i \in S_p, u \notin M_{\mathcal{B}}, \overline{v \text{ prop}} \mid \overline{v \text{ prop}}].$$

Event  $e_i \in S_p$  is independent of  $u \notin M_{\mathcal{B}}, \overline{v \text{ prop}}$  as discussed before and  $\overline{v \text{ prop}}$  and  $u \notin M_{\mathcal{B}}$  are also independent by Corollary 5.5 Property (iv). This means

$$\mathbb{E}[y_v \mid \overline{v \text{ prop}}] = \sum_{i=1}^d \frac{q_{e_i}}{p \Pr[u_i \notin M_{\mathcal{B}}] \Pr[\overline{v \text{ prop}}]} \cdot p \Pr[u_i \notin M_{\mathcal{B}}] = \frac{1}{\Pr[\overline{v \text{ prop}}]} \sum_{i=1}^d q_{e_i}.$$

Applying Claim A.2 on the RHS concludes the claim.  $\square$

**Claim A.4.** For any pair of edges  $e_i = (v, u_i)$  and  $e_j = (v, u_j)$  with  $e_i \neq e_j$  and  $\{u_i, u_j\} \subseteq \{u_1, \dots, u_d\}$ , let  $y'_{e_i} = (y_{e_i} \mid \overline{v \text{ prop}})$  and  $y'_{e_j} = (y_{e_j} \mid \overline{v \text{ prop}})$ . We have  $\text{Cov}(y'_{e_i}, y'_{e_j}) \leq 0$ .

*Proof.* To prove this claim, we will use some known facts about negatively associated (NA) random variables. By definition, a set of random variables are NA, if any two monotone nondecreasing functions  $f$  and  $g$  defined on disjoint subsets of them satisfy  $\mathbb{E}[f \cdot g] \leq \mathbb{E}[f] \cdot \mathbb{E}[g]$ . Below are three facts about negative association based on [20, 17, 12].

- (1) Any set of Bernoulli random variables whose sum is upper-bounded by one are NA.
- (2) If  $A$  is a set of NA random variables,  $B$  is a set of NA random variables with  $A$  and  $B$  independent of each other,  $A \cup B$  is also a set of NA random variables.
- (3) Let  $X = \{x_1, \dots, x_m\}$  be a set of NA random variables. If  $f_1, \dots, f_k$  are a set of monotone nondecreasing functions defined on disjoint subsets of  $X$ , then  $f_1, \dots, f_k$  are NA.
- (4) Let  $\{x_1, \dots, x_m\}$  be a set of NA random variables. Then, for any  $i \neq j$ ,  $\text{Cov}(x_i, x_j) \leq 0$ .

We will start by showing that random variables  $y'_{e_1}, \dots, y'_{e_d}$  are NA where for any  $i \in [d]$ , we define  $y'_{e_i} = (y_{e_i} \mid \overline{v \text{ prop}})$ . For the rest of the proof, we will omit the condition  $\overline{v \text{ prop}}$  from all the statements for simplicity.

For any pair of vertices  $i \in A/\{v\}$  and  $j \in B$ , let us define Bernoulli random variable  $x_{i,j}$  to be equal to one iff vertex  $i$  sends a proposal to vertex  $j$ . Note that for any  $i$ , we have  $\sum_{j \in B} x_{i,j} \leq 1$ . Also, since vertices send their proposals independently from each other, invoking the first two facts above implies  $\{x_{i,j} : \forall i \in A/\{v\}, j \in B\}$  is a set of NA Bernoulli random variables. Now, for any vertex  $u_j \in \{u_1, \dots, u_d\}$ , we define random variable  $z_j$  to be equal to one if  $u_j \in M_B$ . Note that we have  $z_j = 1$  iff  $j$  receives at least one proposal from  $A/\{v\}$  and so it is a monotone nondecreasing function of  $\{x_{i,j} : i \in A/\{v\}\}$ . Since for any  $z_j$  and  $z'_j$  these subsets are disjoint, invoking the third fact implies  $z_1, \dots, z_d$  are NA. Finally, for any edge, we know  $y'_{e_i}$  is a monotone nondecreasing function of  $z_i$  and whether  $e_i$  is realized. Based on the third fact, this implies negative association of  $y'_{e_1}, \dots, y'_{e_d}$ . Thus, by the fourth property  $\text{Cov}(y'_{e_i}, y'_{e_j}) \leq 0$  for any  $i \neq j$ .  $\square$

We are now ready to prove Claim 5.8.

*Proof.* As proved in Claim 5.6,  $\mathbb{E}[y_v \mid \overline{v \text{ prop}}] \leq q_v$ . We prove the desired inequality of the claim via a concentration bound on random variable  $y'_v := (y_v \mid \overline{v \text{ prop}})$ . Let us for simplicity also define random variable  $y'_{e_i} := (y_{e_i} \mid \overline{v \text{ prop}})$  and observe that  $y'_v = \sum_{i=1}^d y'_{e_i}$ .

By Claim A.4 we know  $\text{Cov}(y'_{e_i}, y'_{e_j}) \leq 0$  for any  $i \neq j$ . We can thus bound the variance of  $y'_v$  as follows:

$$\begin{aligned} \text{Var}[y'_v] &= \sum_{i=1}^d \text{Var}[y'_{e_i}] + 2 \sum_{1 \leq i < j \leq d} \text{Cov}(y'_{e_i}, y'_{e_j}) \\ &\leq \sum_{i=1}^d \mathbb{E}[(y'_{e_i})^2] - \mathbb{E}[y'_{e_i}]^2 \leq \sum_{i=1}^d \mathbb{E}[(y'_{e_i})^2] \stackrel{\text{see below}}{\leq} \tau \cdot \mathbb{E}[y'_v] \stackrel{\text{Claim A.3}}{\leq} \tau \cdot q_v, \end{aligned}$$

where  $\tau$  is the maximum possible outcome of  $y'_{e_i}$  for any  $i \in [d]$ .

Plugging this into Chebyshev's inequality, we get

$$\Pr[y'_v > \mathbb{E}[y'_v] + \delta] \leq \frac{\text{Var}[y'_v]}{\delta^2} \leq \frac{\tau q_v}{\delta^2}.$$

Finally, for each edge  $e_i$ , by construction of  $\mathbf{y}$ ,  $y'_{e_i} \leq y_{e_i} \leq q_{e_i}/(p \Pr[\overline{v \text{ prop}}] \Pr[u_i \notin M_{\mathcal{B}}]) \leq q_{e_i}/p\varepsilon^2$  where the latter follows from Corollary 5.5. Combined with  $q_{e_i} \leq \varepsilon^5 p$  by Lemma 5.1 and  $e_i \in S$ , we get  $\tau \leq \varepsilon^3$ . We thus get  $\Pr[y'_v > 1 + \varepsilon] \leq \Pr[y'_v > \mathbb{E}[y'_v] + \varepsilon] \leq \frac{\varepsilon^3}{\varepsilon^2} q_v$ . This concentration of  $y'_v$  implies  $\mathbb{E}[y_v \mid y_v \leq 1 + \varepsilon, \overline{v \text{ prop}}] \leq \varepsilon q_v$  and thus the stated bound of the claim.  $\square$