

Greedy Approaches to Online Stochastic Matching

Allan Borodin *

Calum MacRury †

Akash Rakheja ‡

Abstract

Within the context of stochastic probing with commitment, we consider the online stochastic matching problem; that is, the one-sided online bipartite matching problem where edges adjacent to an online node must be probed to determine if they exist based on edge probabilities that become known when an online vertex arrives. If a probed edge exists, it must be used in the matching (if possible). We consider the competitiveness of online algorithms in both the adversarial order model (AOM) and the random order model (ROM). More specifically, we consider a bipartite stochastic graph $G = (U, V, E)$ where U is the set of offline vertices, V is the set of online vertices and G has edge probabilities $(p_e)_{e \in E}$ and edge weights $(w_e)_{e \in E}$. Additionally, G has probing constraints $(\mathcal{C}_v)_{v \in V}$, where \mathcal{C}_v indicates which sequences of edges adjacent to an online vertex v can be probed. We assume that U is known in advance, and that \mathcal{C}_v , together with the edge probabilities and weights adjacent to an online vertex are only revealed when the online vertex arrives. This model generalizes the various settings of the classical bipartite matching problem, and so our main contribution is in making progress towards understanding which classical results extend to the stochastic probing model.

1 Introduction

Stochastic probing problems are part of the larger area of decision making under uncertainty and more specifically, stochastic optimization. Unlike more standard forms of stochastic optimization, it is not just that there is some stochastic uncertainty in the set of inputs, stochastic probing problems involve inputs that cannot be determined without probing (at some cost and/or within some constraint) so as to reveal the inputs. Applications of stochastic probing occur naturally in many settings, such as in matching problems where compatibility (for example, in online dating and kidney exchange applications) or legality (for example, a financial transaction that must be authorized before it can be completed) cannot be determined without some trial or investigation. Amongst other applications, the online bipartite stochastic matching problem notably models online advertising where the probability of an edge can correspond to the probability of a purchase in online stores or to pay-per-click revenue in online searching.

The (*offline*) *stochastic matching* problem was introduced by Chen et al. [6]. In this problem, we are given an adversarially generated *stochastic graph* $G = (V, E)$ with a probability p_e associated with each edge e and a patience (or time-out) parameter ℓ_v associated with each vertex v . An algorithm probes edges in E within the constraint that at most ℓ_v edges are probed incident to any particular vertex $v \in V$. Also, when an edge e is probed, it is guaranteed to exist with probability exactly p_e . If an edge (u, v) is found to exist, it is added to the matching and then u and v are no longer available. The goal is to maximize the expected size of a matching constructed in this way.

*Department of Computer Science, University of Toronto, Toronto, ON, Canada bor@cs.toronto.edu

†Department of Computer Science, University of Toronto, Toronto, ON, Canada cmacrury@cs.toronto.edu

‡Department of Computer Science, University of Toronto, Toronto, ON, Canada rakhejaakash@gmail.com

This problem can be generalized to vertices or edges having weights. We shall refer to this setting as the *known stochastic graph* setting.

Mehta and Panigrahi [20] adapted the offline stochastic matching model to online bipartite matching as originally studied in the classical (non-stochastic) adversarial order online model. That is, they consider the setting where the stochastic graph is unknown and online vertices are determined by an adversary. More specifically, they studied the problem in the case of an unweighted stochastic graph $G = (U, V, E)$ where U is the set of offline vertices and the vertices in V arrive online without knowledge of future online node arrivals. They considered the special case of uniform edge probabilities (i.e, $p_e = p$ for all $e \in E$) and *unit* patience values, that is $\ell_v = 1$ for all $v \in V$. Mehta et al. [21] considered the unweighted online stochastic bipartite setting with arbitrary edge probabilities, attaining a competitive ratio of 0.534, and recently, Huang and Zhang [13] additionally handled the case of arbitrary offline vertex weights, while improving this ratio to 0.572. However, as in [20], both [21] and [13] are restricted to unit patience values, and moreover require edge probabilities which are *vanishingly small*¹. Goyal and Udmani [11] improved on both of these works by showing a 0.596 competitive ratio in the same setting.

In all our results we will assume *commitment*; that is, when an edge is probed and found to exist, it must be included in the matching (if possible without violating the matching constraint). The patience constraint can be viewed as a simple form of a budget constraint for the online vertices. For some, but not all of our settings, we generalize patience and budget constraints by associating a set \mathcal{C}_v of probing sequences for each online node v where \mathcal{C}_v indicates which sequences of edges adjacent to vertex v can be probed.

For random order or adversarial order of online vertex arrivals, results for the online Mehta and Panigrahi model (even for unit patience) generalize the corresponding classical non-stochastic models where edges adjacent to an online node are known upon arrival and do not need to be probed. It follows that any in-approximations in the classical setting apply to the corresponding stochastic setting. Further generalizing the classical settings, when the stochastic graph is unknown, a competitive ratio for the random order model implies that the same ratio is obtained in the stochastic i.i.d. model (for an unknown distribution) as proven in the classical setting by Karande et al. [14].

In a related paper [4], we consider the setting when the stochastic graph is unknown, but there is a known stochastic type graph from which arrivals are drawn independently, thereby generalizing the i.i.d. model introduced by Bansal et al. [3]. Our generalization in [4] allows for independent (but not necessarily identical) distributions; that is, we have an i.d. model.

1.1 Preliminaries

An input to the **(online) stochastic matching problem** is a **(bipartite) stochastic graph**, specified in the following way. Let $G = (U, V, E)$ be a bipartite graph with edge weights $(w_e)_{e \in E}$ and edge probabilities $(p_e)_{e \in E}$. We draw an independent Bernoulli random variable of parameter p_e for each $e \in E$. We refer to this Bernoulli as the **state** of the edge e , and denote it by $\text{st}(e)$. If $\text{st}(e) = 1$, then we say that e is **active**, and otherwise we say that e is **inactive**. For each $v \in V$, suppose that $\partial(v)^{(*)}$ corresponds to the collection of strings (tuples) formed from *distinct* edges of $\partial(v)$. Each $v \in V$ has a set of **online probing constraint** $\mathcal{C}_v \subseteq \partial(v)^{(*)}$ which is **substring-closed**. That is, \mathcal{C}_v has the property that if $e \in \mathcal{C}_v$, then so is any substring of e . Similarly, we say that \mathcal{C}_v is **permutation-closed**, provided if $e \in \mathcal{C}_v$, then so is any permutation string of e . Observe that if \mathcal{C}_v is both substring-closed and permutation-closed, then it corresponds to a downward-closed family

¹Vanishingly small edge probabilities must satisfy $\max_{e \in E} p_e \rightarrow 0$, where the asymptotics are with respect to the size of G .

of subsets of $\partial(v)$. Thus, in particular, our setting encodes the case when v has a patience value ℓ_v , and more generally, when \mathcal{C}_v corresponds to a matroid or budgetary constraint on $\partial(v)$. Note that we will often assume *w.l.o.g.* that $E = U \times V$, as we can always set $p_{u,v} := 0$.

A solution to online stochastic matching is an **online probing algorithm**. An online probing algorithm is initially only aware of the identity of the offline vertices U of G . We think of V , as well as the relevant edges probabilities, weights, and probing constraints, as being generated by an adversary. An ordering on V is then generated either through an adversarial process or uniformly at random. We refer to the former case as the **adversarial order model (AOM)** and the latter case as the **random order model (ROM)**.

Based on whichever ordering is generated on V , the nodes are then presented to the online probing algorithm one by one. When an online node $v \in V$ arrives, the online probing algorithm sees all the adjacent edges and their associated probabilities, as well as \mathcal{C}_v . However, the edge states $(\text{st}(e))_{e \in \partial(v)}$ remain hidden to the algorithm. Instead, the algorithm must perform a **probing operation** on an adjacent edge e to reveal/expose its state, $\text{st}(e)$. Moreover, the online probing algorithm must **respect commitment**. That is, if an edge $e = (u, v)$ is probed and turns out to be active, then e must be added to the current matching, provided u and v are both currently unmatched. The probing constraint \mathcal{C}_v of the online node then restricts which sequences of probes can be made to $\partial(v)$. As in the classical problem, an online probing algorithm must decide on a possible match for an online node v before seeing the next online node. The goal of the online probing algorithm is to return a matching whose expected weight is as large as possible. Since \mathcal{C}_v may be exponentially large in the size of U , in order to discuss the efficiency of an online probing algorithm, we work in the **membership query model**. That is, upon receiving the online vertex $v \in V$, an online probing algorithm may make a **membership query** to any string $e \in \partial(v)^{(*)}$, thus determining in a single operation whether or not $e \in \partial(v)^{(*)}$ is in \mathcal{C}_v .

It is easy to see we cannot hope to obtain a non-trivial competitive ratio against the expected value of an optimum matching of the stochastic graph². Thus, the standard in the literature is to instead benchmark the performance of an online probing algorithm against an *optimum offline probing algorithm*. An **offline probing algorithm** knows $G = (U, V, E)$, but initially the edge states $(\text{st}(e))_{e \in E}$ are hidden. It can adaptively probe the edges of E in any order, but must satisfy the probing constraints $(\mathcal{C}_v)_{v \in V}$ at each step of its execution³, while respecting commitment; that is, if a probed edge $e = (u, v)$ turns out to be active, then e is added to the matching (if possible). The goal of an offline probing algorithm is to construct a matching with optimum weight in expectation. We define the **committal benchmark** as an optimum offline probing algorithm, and use $\text{OPT}(G)$ to denote the expected value of the matching the committal benchmark constructs. We abuse notation slightly, and also use $\text{OPT}(G)$ to refer to the *strategy* of the committal benchmark on G . In Appendix A, we introduce the stronger **non-committal benchmark**, and indicate which of our results hold against it.

1.2 Our Results

We first consider the unknown stochastic matching problem in the most general setting of arbitrary edge weights, and substring-closed probing constraints. Note that since no non-trivial competitive

²Consider a single online vertex with patience 1, and n offline (unweighted) vertices where each edge e has probability $\frac{1}{n}$ of being present. The expectation of an online probing algorithm will be at most $\frac{1}{n}$ while the expected size of an optimal matching (over all instantiations of the edge probabilities) will be $1 - (1 - \frac{1}{n})^n \rightarrow 1 - \frac{1}{e}$. This example clearly shows that no constant ratio is possible if the patience is sublinear (in $n = |U|$).

³Edges $e \in E^{(*)}$ may be probed in order, provided $e^v \in \mathcal{C}_v$ for each $v \in V$, where e^v is the substring of e restricted to edges of $\partial(v)$.

ratio can be proven in the case of adversarial arrivals, we work in the ROM setting. We generalize the matching algorithm of Kesselheim et al. [16] so as to apply to the stochastic probing setting.

Theorem 1.1. *Suppose the adversary presents an edge-weighted stochastic graph $G = (U, V, E)$, with substring-closed probing constraints $(\mathcal{C}_v)_{v \in V}$. If \mathcal{M} is the matching returned by Algorithm 2 when executing on G , then*

$$\mathbb{E}[w(\mathcal{M})] \geq \left(\frac{1}{e} - \frac{1}{|V|} \right) \cdot \text{OPT}(G),$$

provided the vertices of V arrive uniformly at random (u.a.r.). If the constraints $(\mathcal{C}_v)_{v \in V}$ are also permutation-closed, then Algorithm 2 can be implemented efficiently in the membership oracle model.

Upon receiving the online vertices $V_t := \{v_1, \dots, v_t\}$, in order to generalize the matching algorithm of Kesselheim et al. [16], Algorithm 2 would ideally probe the edges of $\partial(v_t)$ suggested by $\text{OPT}(G_t)$, where $G_t := G[U \cup V_t]$ is the **induced stochastic graph**⁴ on $U \cup V_t$. However, since we wish for our algorithms to be efficient in addition to attaining optimum competitive ratios, this strategy is not feasible. Our solution is to instead solve a configuration LP (LP-config) for G_t , which was recently introduced by the authors in [4], and whose optimum value upper bounds $\text{OPT}(G_t)$. Unlike LP-std – the most prevalent LP used in the stochastic matching literature, originally introduced by Bansal et al. [3] (see Appendix B)– LP-config allows us to handle general probing constraints, while not overestimating the performance of $\text{OPT}(G_t)$. Our algorithm differs from the classical algorithm of Kesselheim et al. in that it is randomized.

Since Theorem 1.1 achieves the optimum asymptotic competitive ratio for edge weights, in order to improve this result we next consider the case when the stochastic graph $G = (U, V, E)$ has **(offline) vertex weights** – i.e., there exists $(w_u)_{u \in U}$ such that $w_{u,v} = w_u$ for each $v \in N(u)$. We consider a *greedy* online probing algorithm. That is, upon the arrival of v , the probes to $\partial(v)$ are made in such a way that v gains as much value as possible (in expectation), provided the currently unmatched nodes of U are equal to R . As such, we must follow the probing strategy of the committal benchmark when restricted to $G[\{v\} \cup R]$, which we denote by $\text{OPT}(R, v)$ for convenience.

Observe that if v has unit patience, then $\text{OPT}(R, v)$ reduces to probing the adjacent edge $(u, v) \in R \times \{v\}$ such that the value $w_u \cdot p_{u,v}$ is maximized. Moreover, if v has unlimited patience, then $\text{OPT}(R, v)$ corresponds to probing the adjacent edges of $R \times \{v\}$ in non-increasing order of the associated vertex weights. Building on a result in Purohit et al. [22], Brubach et al. [5] showed how to devise an *efficient* probing strategy for v whose expected value matches $\text{OPT}(R, v)$, no matter the patience constraint. Using this probing strategy, they devised an online probing algorithm which achieves a competitive ratio of $1/2$ for arbitrary patience values. We extend their result to substring-closed probing constraints.

Theorem 1.2. *Suppose the adversary presents a vertex weighted stochastic graph $G = (U, V, E)$, with substring-closed probing constraints $(\mathcal{C}_v)_{v \in V}$. If \mathcal{M} is the matching returned by Algorithm 3 when executing on G , then*

$$\mathbb{E}[w(\mathcal{M})] \geq \frac{1}{2} \cdot \text{OPT}(G),$$

provided the vertices of V arrive in adversarial order. Moreover, Algorithm 3 can be implemented efficiently, provided the constraints $(\mathcal{C}_v)_{v \in V}$ are also permutation-closed.

⁴Given $L \subseteq U, R \subseteq V$, the induced stochastic graph $G[L \cup R]$ is formed by restricting the edges weights and probabilities of G to those edges within $L \times R$. Similarly, each probing constraint \mathcal{C}_v is restricted to those strings whose entries lie entirely in $L \times \{v\}$.

Since Algorithm 3 is deterministic, the $1/2$ competitive ratio is best possible for deterministic algorithms in the adversarial arrival setting. One direction is thus to instead consider what can be done if the online probing algorithm is allowed randomization. In the case of unit patience, this is the previously discussed setting of [20], in which Mehta and Panigrahi showed that $.621 < 1 - \frac{1}{e}$ is a randomized in-approximation with regard to guarantees made against LP-std-unit, even for the unweighted uniform probability setting. Note that for unit patience, LP-std-unit can be viewed as encoding a *relaxed* optimum probing algorithm which need only match the offline nodes once in expectation, and thus upper bounds/relaxes $\text{OPT}(G)$ (see Appendix B for details). This hardness result led Goyal and Udvani [11] to consider a new unit patience LP that is a tighter relaxation of $\text{OPT}(G)$ than LP-std-unit, thereby allowing them to prove a $1 - 1/e$ competitive ratio for the case of **vertex-decomposable**⁵ edge probabilities. However, they also discuss the difficulty of extending this result to the case of general edge probabilities. Our next contribution is thus to consider Algorithm 3 in the ROM setting with unit patience, where we show these difficulties do *not* arise. In fact, we show that a $1 - 1/e$ performance guarantee is provable against LP-std-unit, which shows that the in-approximation of Mehta and Panigrahi does not apply (even for deterministic probing algorithms).

Upon the arrival of vertex v , if v has unit patience, then Algorithm 3 reduces to probing the edge $e = (u, v) \in \partial(v)$ such that $u \in U$ is currently unmatched, and for which $w_u \cdot p_{u,v}$ is maximized. Let us refer to this special case of Algorithm 3 as GREEDYPROBE.

Theorem 1.3. *Suppose the adversary presents a vertex weighted stochastic graph $G = (U, V, E)$, with unit patience values. If \mathcal{M} is the matching returned by GREEDYPROBE when executing on G , then $\mathbb{E}[w(\mathcal{M})] \geq (1 - \frac{1}{e}) \cdot \text{OPT}(G)$, provided the vertices of V arrive in random order.*

Remark 1.4. The guarantee of Theorem 1.3 is proven against LP-std-unit, and together with the 0.621 inapproximation of Mehta and Panigrahi, implies that deterministic probing algorithms in the ROM setting have strictly more power than randomized probing algorithms in the adversarial order model. The analysis of GREEDYPROBE is tight, as an execution of GREEDYPROBE corresponds to the seminal Karp et al. [15] RANKING algorithm for unweighted non-stochastic (i.e., $p_e \in \{0, 1\}$ for all $e \in E$) bipartite matching⁶.

Our final result, Theorem 3.4, makes partial progress towards understanding the ROM setting in the case of general probing constraints. However, unlike the adversarial setting, the complexity of the constraints greatly impacts what we are able to prove. We state and prove our result in Section 3, which we note subsumes Theorem 1.3.

2 Edge Weights

Let us suppose that $G = (U, V, E)$ is a stochastic graph with arbitrary edge weights, probabilities and constraints $(C_v)_{v \in V}$. For each $k \geq 1$ and $\mathbf{e} = (e_1, \dots, e_k) \in E^{(*)}$, define $g(\mathbf{e}) := \prod_{i=1}^k (1 - p_{e_i})$. Notice that $g(\mathbf{e})$ corresponds to the probability that all the edges of \mathbf{e} are inactive, where $g(\lambda) := 1$ for the empty string λ . We also define $\mathbf{e}_{<e_i} := (e_1, \dots, e_{i-1})$ for each $2 \leq i \leq k$, which we denote by $\mathbf{e}_{<i}$ when clear. By convention, $\mathbf{e}_{<1} := \lambda$. Observe then that $\text{val}(\mathbf{e}) := \sum_{i=1}^{|\mathbf{e}|} p_{e_i} w_{e_i} \cdot g(\mathbf{e}_{<i})$

⁵Vertex-decomposable means that there exists probabilities $(p_u)_{u \in U}$ and $(p_v)_{v \in V}$, such that $p_{(u,v)} = p_u \cdot p_v$ for each $(u, v) \in E$.

⁶In the classical (unweighted, non-stochastic) online matching problem, an execution of the randomized RANKING algorithm in the adversarial setting can be coupled with an execution of the deterministic greedy algorithm in the ROM setting – the latter of which is a special case of GREEDYPROBE. The tightness of the ratio $1 - 1/e$ therefore follows since this ratio is tight for the RANKING algorithm.

corresponds to the expected weight of the first active edge of \mathbf{e} if \mathbf{e} is probed in order of its indices. Finally, for each $v \in V$ and $\mathbf{e} \in \mathcal{C}_v$, we introduce a decision variable $x_v(\mathbf{e})$. We can then express the following LP from [4]:

$$\begin{array}{ll} \text{maximize} & \sum_{v \in V} \sum_{\mathbf{e} \in \mathcal{C}_v} \text{val}(\mathbf{e}) \cdot x_v(\mathbf{e}) \end{array} \quad (\text{LP-config})$$

$$\begin{array}{ll} \text{subject to} & \sum_{v \in V} \sum_{\substack{\mathbf{e} \in \mathcal{C}_v: \\ (u,v) \in \mathbf{e}}} p_{u,v} \cdot g(\mathbf{e}_{\prec(u,v)}) \cdot x_v(\mathbf{e}) \leq 1 \quad \forall u \in U \end{array} \quad (2.1)$$

$$\sum_{\mathbf{e} \in \mathcal{C}_v} x_v(\mathbf{e}) = 1 \quad \forall v \in V, \quad (2.2)$$

$$x_v(\mathbf{e}) \geq 0 \quad \forall v \in V, \mathbf{e} \in \mathcal{C}_v \quad (2.3)$$

Theorem 2.1 (Theorem 3.1 in [4]). *For any stochastic graph $G = (U, V, E)$ with substring-closed probing constraints, $\text{OPT}(G) \leq \text{LPOPT}(G)$, where $\text{LPOPT}_{\text{conf}}(G)$ is the optimum value of LP-config.*

In order to prove Theorem 2.1, the natural approach is to define $x_v(\mathbf{e})$ to be the probability that the committal benchmark probes the edges of \mathbf{e} in order, where $v \in V$ and $\mathbf{e} \in \mathcal{C}_v$. However, to our knowledge, this interpretation of the decision variables of LP-config does not seem to yield a proof of Theorem 2.1, for technical reasons which we discuss in detail in Appendix B.

The approach we take in [4] is to instead introduce a **combinatorial relaxation** of the committal benchmark, which is a new stochastic probing problem on G whose optimum solution upper bounds $\text{OPT}(G)$. Specifically, we introduce the **relaxed stochastic matching problem**, a solution to which we refer to as a **relaxed probing algorithm**. A relaxed probing algorithm operates in the same probing framework as an offline probing algorithm, however it does *not* return a (two-sided) matching of G . Instead, it returns a subset of its active probes which form a one-sided matching of V . This one-sided matching \mathcal{N} has the additional property that each offline vertex of U is included in \mathcal{N} at most once *in expectation*. We define the **relaxed benchmark** to be an optimum relaxed probing algorithm, and denote its evaluation on G by $\text{OPT}_{\text{rel}}(G)$. Since each offline probing algorithm is a relaxed probing algorithm, clearly $\text{OPT}(G) \leq \text{OPT}_{\text{rel}}(G)$. On the other hand, by defining $x_v(\mathbf{e})$ to be the probability the relaxed benchmark probes $\mathbf{e} \in \mathcal{C}_v$ of $v \in V$ in order, we can show that $\text{OPT}_{\text{rel}}(G) = \text{LPOPT}(G)$, which implies Theorem 2.1. For completeness, we include the details in Appendix B.

Not only does LP-config relax the committal benchmark, it can be solved efficiently, provided the constraints $(\mathcal{C}_v)_{v \in V}$ are assumed to be closed under substrings and permutations. The approach we take in [4] is to first consider the dual of LP-config. It is not hard to verify that the DP-OPT algorithm of Theorem 3.1 can be used as a (deterministic) polynomial time separation oracle for this LP. This ensures that the dual of LP-config can be solved efficiently, as a consequence of how the ellipsoid algorithm [23, 10] executes. Moreover, by tracking which polynomial number of constraints of the dual of LP-config are queried by this separation oracle, one can reduce the number of decision variables needed in LP-config to a polynomial number. Since this restriction of LP-config has a polynomial number of constraints, it can then be solved efficiently. We omit the details, as a more complete proof is given in [4], and this technique for solving LPs which have an exponential number of variables is well-known in the literature (see [26, 25, 1, 17] for instance).

We now define a *fixed vertex* probing algorithm, called VERTEXPROBE, which is applied to an online vertex s of an arbitrary stochastic graph (potentially distinct from G):

Algorithm 1 VertexProbe

Input: an online vertex s of a stochastic graph, $\partial(s)$, and probabilities $(z(\mathbf{e}))_{\mathbf{e} \in \mathcal{C}_s}$ which satisfy $\sum_{\mathbf{e} \in \mathcal{C}_s} z(\mathbf{e}) = 1$.

Output: an active edge \mathcal{N} of $\partial(s)$.

```
1:  $\mathcal{N} \leftarrow \emptyset$ .
2: Draw  $\mathbf{e}'$  from  $\mathcal{C}_s$  with probability  $z(\mathbf{e}')$ .
3: if  $\mathbf{e}' = \lambda$  then ▷ the empty string is drawn.
4:   return  $\mathcal{N}$ .
5: else
6:   Denote  $\mathbf{e}' = (e'_1, \dots, e'_k)$  for  $k := |\mathbf{e}'| \geq 1$ .
7:   for  $i = 1, \dots, k$  do
8:     Probe the edge  $e'_i$ .
9:     if  $\text{st}(e'_i) = 1$  then
10:      Add  $e'_i$  to  $\mathcal{N}$ , and exit the “for loop”.
11:    end if
12:  end for
13: end if
14: return  $\mathcal{N}$ .
```

Lemma 2.2. *Suppose VERTEXPROBE (Algorithm 1) is passed a fixed online node s of a stochastic graph, and probabilities $(z(\mathbf{e}))_{\mathbf{e} \in \mathcal{C}_s}$ which satisfy $\sum_{\mathbf{e} \in \mathcal{C}_s} z(\mathbf{e}) = 1$. If for each $e \in \partial(s)$,*

$$\tilde{z}_e := \sum_{\substack{\mathbf{e}' \in \mathcal{C}_v: \\ e \in \mathbf{e}'}} g(\mathbf{e}'_{<e}) \cdot z_v(\mathbf{e}'),$$

then $e \in \partial(s)$ is probed with probability \tilde{z}_e , and returned by the algorithm with probability $p_e \cdot \tilde{z}_e$.

Remark 2.3. If VERTEXPROBE outputs the edge $e = (u, s)$ when executing on the fixed node s , then we say that s **commits** to the edge $e = (u, s)$, or that s commits to u .

Returning to the problem of designing an online probing algorithm for G , let us assume that $n := |V|$, and that the online nodes of V are denoted v_1, \dots, v_n , where the order is generated *u.a.r.*. Denote V_t as the set of first t arrivals of V ; that is, $V_t := \{v_1, \dots, v_t\}$. Moreover, set $G_t := G[U \cup V_t]$, and $\text{LPOPT}_{\text{conf}}(G_t)$ as the value of an optimum solution to LP-config (this is a random variable, as V_t is a random subset of V). The following inequality then holds:

Lemma 2.4. *For each $t \geq 1$, $\mathbb{E}[\text{LPOPT}_{\text{conf}}(G_t)] \geq \frac{t}{n} \text{LPOPT}_{\text{conf}}(G)$.*

In light of this observation, we design an online probing algorithm which makes use of V_t , the currently known nodes, to derive an optimum LP solution with respect to G_t . As such, each time an online node arrives, we must compute an optimum solution for the LP associated to G_t , distinct from the solution computed for that of G_{t-1} .

Algorithm 2 Unknown Stochastic Graph ROM

Input: U and $n := |V|$.

Output: a matching \mathcal{M} from the (unknown) stochastic graph $G = (U, V, E)$ of active edges.

- 1: Set $\mathcal{M} \leftarrow \emptyset$.
 - 2: Set $G_0 = (U, \emptyset, \emptyset)$
 - 3: **for** $t = 1, \dots, n$ **do**
 - 4: Input v_t , with $(w_e)_{e \in \partial(v_t)}$, $(p_e)_{e \in \partial(v_t)}$ and \mathcal{C}_{v_t} .
 - 5: Compute G_t , by updating G_{t-1} to contain v_t (and its relevant information).
 - 6: **if** $t < \lfloor n/e \rfloor$ **then**
 - 7: Pass on v_t .
 - 8: **else**
 - 9: Solve LP-config for G_t and find an optimum solution $(x_v(\mathbf{e}))_{v \in V_t, \mathbf{e} \in \mathcal{C}_v}$.
 - 10: Set $e_t \leftarrow \text{VERTEXPROBE}(v_t, \partial(v_t), (x_v(\mathbf{e}))_{\mathbf{e} \in \mathcal{C}_{v_t}})$.
 - 11: **if** $e_t = (u_t, v_t) \neq \emptyset$ and u_t is unmatched **then**
 - 12: Add e_t to \mathcal{M} .
 - 13: **end if**
 - 14: **end if**
 - 15: **end for**
 - 16: **return** \mathcal{M} .
-

Let us consider the matching \mathcal{M} returned by the algorithm, as well as its weight, which we denote by $w(\mathcal{M})$. Set $\alpha := 1/e$ for clarity, and take $t \geq \lceil \alpha n \rceil$. Define R_t as the *unmatched vertices* of U when vertex v_t arrives. Note that committing to $e_t = (u_t, v_t)$ is necessary, but not sufficient, for v_t to match to u_t . With this notation, we have that $\mathbb{E}[w(\mathcal{M})] = \sum_{t=\lceil \alpha n \rceil}^n \mathbb{E}[w(u_t, v_t) \cdot \mathbf{1}_{[u_t \in R_t]}]$. Moreover, we claim the following:

Lemma 2.5. *For each $t \geq \lceil \alpha n \rceil$, $\mathbb{E}[w(e_t)] \geq LPOPT_{conf}(G)/n$.*

Proof of Lemma 2.5. Set $\alpha := 1/e$ for clarity, and take $t \geq \lceil \alpha n \rceil$. Define $e_t := (u_t, v_t)$, where u_t is the vertex of U which v_t commits to (which is the empty set \emptyset , if no such commitment occurs). For each $u \in U$, denote $C(u, v_t)$ as the event in which v_t commits to u . Let us now condition on the random subset V_t , as well as the random vertex v_t . In this case,

$$\mathbb{E}[w(e_t) \mid V_t, v_t] = \sum_{u \in U} w_{u, v_t} \mathbb{P}[C(u, v_t) \mid V_t, v_t].$$

Observe however that once we condition on V_t and v_t , Algorithm 2 corresponds to executing VERTEXPROBE on the instance $(v_t, \partial(v_t), (x_{v_t}(\mathbf{e}))_{\mathbf{e} \in \mathcal{C}_{v_t}})$, where we recall that $(x_v(\mathbf{e}))_{\mathbf{e} \in \mathcal{C}_v, v \in V_t}$ is an optimum solution to LP-config for $G_t = G[U \cup V_t]$. Thus, Lemma 2.2 implies that $\mathbb{P}[C(u, v_t) \mid V_t, v_t] = p_{u, v_t} \tilde{x}_{u, v_t}$, where

$$\tilde{x}_{u, v_t} := \sum_{\substack{\mathbf{e}' \in \mathcal{C}_{v_t}: \\ e \in \mathbf{e}'}} g(\mathbf{e}'_{<e}) \cdot x_{v_t}(\mathbf{e}'),$$

and so $\mathbb{E}[w(e_t) \mid V_t, v_t] = \sum_{u \in U} w_{u, v_t} p_{u, v_t} \tilde{x}_{u, v_t}$. On the other hand, if we condition on *solely* V_t , then v_t remains distributed uniformly at random amongst the vertices of V_t . Moreover, once we condition on V_t , the graph G_t is determined, and thus so are the values $(x_v(\mathbf{e}))_{v \in V_t, \mathbf{e} \in \mathcal{C}_v}$. These observations together imply that

$$\mathbb{E}[w_{u, v_t} p_{u, v_t} \tilde{x}_{u, v_t} \mid V_t] = \frac{\sum_{v \in V_t} w_{u, v} p_{u, v} \tilde{x}_{u, v}}{t} \quad (2.4)$$

for each $u \in U$ and $\lceil \alpha n \rceil \leq t \leq n$. If we now take expectation over v_t , then using the law of iterated expectations,

$$\begin{aligned} \mathbb{E}[w(e_t) | V_t] &= \mathbb{E}[\mathbb{E}[w(e_t) | V_t, v_t] | V_t] \\ &= \mathbb{E} \left[\sum_{u \in U} w_{u,v_t} p_{u,v_t} \tilde{x}_{u,v_t} | V_t \right] \\ &= \sum_{u \in U} \mathbb{E}[w_{u,v_t} p_{u,v_t} \tilde{x}_{u,v_t} | V_t] \\ &= \sum_{u \in U} \sum_{v \in V_t} \frac{w_{u,v} p_{u,v} \tilde{x}_{u,v}}{t}, \end{aligned}$$

where the final equation follows from (2.4). Observe however that

$$\text{LPOPT}_{\text{conf}}(G_t) = \sum_{v \in V_t} \sum_{u \in U} w_{u,v} p_{u,v} \tilde{x}_{u,v},$$

as $(x_v(\mathbf{e}))_{v \in V_t, \mathbf{e} \in \mathcal{C}_v}$ is an optimum solution to LP-config for G_t . As a result,

$$\mathbb{E}[w(e_t) | V_t] = \frac{\text{LPOPT}_{\text{conf}}(G_t)}{t},$$

and so

$$\mathbb{E}[w(e_t)] = \frac{\mathbb{E}[\text{LPOPT}_{\text{conf}}(G_t)]}{t},$$

after taking taking expectation over V_t . On the other hand, Lemma 2.4 implies that

$$\frac{\mathbb{E}[\text{LPOPT}_{\text{conf}}(G_t)]}{t} \geq \frac{\text{LPOPT}_{\text{conf}}(G)}{n}.$$

Thus,

$$\mathbb{E}[w(e_t)] \geq \frac{\text{LPOPT}_{\text{conf}}(G)}{n},$$

provided $\lceil \alpha n \rceil \leq t \leq n$, thereby completing the proof. \square

Lemma 2.6. *For each $t \geq \lceil \alpha n \rceil$, define $f(t, n) := \lfloor \alpha n \rfloor / (t - 1)$. In this case, $\mathbb{P}[u_t \in R_t | V_t, v_t] \geq f(t, n)$, where $V_t = \{v_1, \dots, v_t\}$ and v_t is the t^{th} arriving node of V .⁷*

Proof of Lemma 2.6. Let us assume that $\lceil \alpha n \rceil \leq t \leq n$ is fixed, and $(x_v^{(t)}(\mathbf{e}))_{v \in V, \mathbf{e} \in \mathcal{C}_v}$ is the optimum solution of LP-config for G_t , as used by Algorithm 2. For each $u \in U$ and $v \in v$, define the **edge variable** $\tilde{x}_{u,v}^{(t)}$, where

$$\tilde{x}_{u,v}^{(t)} := \sum_{\substack{e' \in \mathcal{C}_{v_t}: \\ e \in e'}} g(e'_{<e}) \cdot x_{v_t}^{(t)}(e')$$

We wish to prove that for each $u \in U$,

$$\mathbb{P}[u \in R_t | V_t, v_t] \geq \lfloor \alpha n \rfloor / (t - 1). \quad (2.5)$$

⁷Note that since V_t is a set, conditioning on V_t only reveals which vertices of V encompass the first t arrivals, *not* the order they arrived in. Hence, conditioning on v_t as well reveals strictly more information.

As such, we must condition on (V_t, v_t) throughout the remainder of the proof. To simplify the argument, we abuse notation slightly and remove (V_t, v_t) from the subsequent probability computations, though it is understood to implicitly appear.

Given arriving node v_j for $j = 1, \dots, n$, once again denote $C(u, v_j)$ as the event in which v_j commits to $u \in U$. As R_t denotes the unmatched nodes after the vertices v_1, \dots, v_{t-1} are processed by Algorithm 2, observe that $u \in R_t$ if and only if $\neg C(u, v_j)$ occurs for each $j = 1, \dots, t-1$, and so $\mathbb{P}[u \in R_t] = \mathbb{P}[\bigcap_{j=1}^{t-1} \neg C(u, v_j)]$. We therefore focus on lower bounding $\mathbb{P}[\bigcap_{j=1}^{t-1} \neg C(u, v_j)]$ in order to prove the lemma.

First observe that for $j = 1, \dots, \lfloor \alpha n \rfloor$, the algorithm passes on probing $\partial(v_j)$ by definition, and so (2.5) holds if $t = \lceil \alpha n \rceil$. As such, we may thereby assume $t \geq \lceil \alpha n \rceil + 1$ and focus on lower bounding $\mathbb{P}[\bigcap_{j=\lceil \alpha n \rceil}^{t-1} \neg C(u, v_j)]$. Observe that this event depends only on the probes of the vertices of $V_{t-1} \setminus V_{\lceil \alpha n \rceil}$. We denote $\bar{t} := t - 1 - \lfloor \alpha n \rfloor = t - \lceil \alpha n \rceil$ as the number of vertices within this set.

Let us first consider the vertex v_{t-1} , and the edge variable $\tilde{x}_{u,v}^{(t-1)}$ for each $v \in V_{t-1}$. Observe that after applying Lemma 2.2,

$$\begin{aligned} \mathbb{P}[C(u, v_{t-1})] &= \sum_{v \in V_{t-1}} \mathbb{P}[C(u, v_{t-1}) \mid v_{t-1} = v] \cdot \mathbb{P}[v_{t-1} = v] \\ &= \frac{1}{t-1} \sum_{v \in V_{t-1}} \tilde{x}_{u,v}^{(t-1)} p_{u,v}, \end{aligned}$$

as once we condition on the set V_t and the vertex v_t , we know that v_{t-1} is uniformly distributed amongst V_{t-1} . On the other hand, the values $(\tilde{x}_{u,v}^{(t-1)})_{u \in U, v \in V_{t-1}}$ are derived from a solution to LP-config for G_{t-1} , and so by constraint (2.1),

$$\sum_{v \in V_{t-1}} \tilde{x}_{u,v}^{(t-1)} p_{u,v} \leq 1.$$

We therefore get that $\mathbb{P}[C(u, v_{t-1})] \leq \frac{1}{t-1}$. Similarly, if we fix $1 \leq k \leq \bar{t}$, then we can generalize the above argument by conditioning on the identities of all the vertices preceding v_{t-k} , as well as the probes they make; that is, $(u_{t-1}, v_{t-1}), \dots, (u_{t-(k-1)}, v_{t-(k-1)})$ (in addition to V_t and v_t as always).

In order to simplify the resulting indices, let us reorder the vertices of $V_{t-1} \setminus V_{\lceil \alpha n \rceil}$. Specifically, define $\bar{v}_k := v_{t-k}$, $\bar{u}_k := u_{t-k}$ and $\bar{e}_k := e_{t-k}$ for $k = 1, \dots, \bar{t}$. With this notation, we denote \mathcal{H}_k as encoding the information available based on the vertices $\bar{v}_1, \dots, \bar{v}_k$ and the edges they (potentially) committed to, namely $\bar{e}_1, \dots, \bar{e}_k$ ⁸. By convention, we define \mathcal{H}_0 as encoding the information regarding V_t and v_t .

An analogous computation to the above case then implies that

$$\mathbb{P}[C(u, \bar{v}_k) \mid \mathcal{H}_{k-1}] = \sum_{v \in V_{t-k}} \tilde{x}_{u,v}^{(t-k)} p_{u,v} \mathbb{P}[\bar{v}_k = v] \leq \frac{1}{t-k},$$

for each $k = 1, \dots, \bar{t}$, where $\tilde{x}_{u,v}^{(t-k)}$ is the edge variable for $v \in V_{t-k}$.

Observe now that in each step, we condition on strictly more information; that is, $\mathcal{H}_{k-1} \subseteq \mathcal{H}_k$ for each $k = 2, \dots, \bar{t}$. On the other hand, observe that if we condition on \mathcal{H}_{k-1} for $1 \leq k \leq \bar{t}-1$, then the event $C(u, \bar{v}_j)$ can be determined from \mathcal{H}_{k-1} for each $1 \leq j \leq k-1$.

⁸Formally, \mathcal{H}_k is the sigma-algebra generated from V_t, v_t and $\bar{e}_1, \dots, \bar{e}_k$.

Using these observations, for $1 \leq k \leq \bar{t}$, the following recursion holds:

$$\begin{aligned} \mathbb{P}[\cap_{j=1}^k \neg C(u, \bar{v}_j)] &= \mathbb{E} \left[\mathbb{E} \left[\prod_{j=1}^k \mathbf{1}_{[\neg C(u, \bar{v}_j)]} \mid \mathcal{H}_{k-1} \right] \right] \\ &= \mathbb{E} \left[\prod_{j=1}^{k-1} \mathbf{1}_{[\neg C(u, \bar{v}_j)]} \mathbb{P}[\neg C(u, \bar{v}_k) \mid \mathcal{H}_{k-1}] \right] \\ &\geq \left(1 - \frac{1}{t-k} \right) \mathbb{P}[\cap_{j=1}^{k-1} \neg C(u, \bar{v}_j)] \end{aligned}$$

It follows that if $k = \bar{t} = t - \lceil \alpha n \rceil$, then applying the above recursion implies that

$$\mathbb{P}[\cap_{j=\lceil \alpha n \rceil}^{t-1} \neg C(u, v_j)] \geq \prod_{k=1}^{t-\lceil \alpha n \rceil} \left(1 - \frac{1}{t-k} \right).$$

Thus, after cancelling the pairwise products,

$$\mathbb{P}[u \in R_t] = \mathbb{P}[\cap_{j=\alpha n}^{t-1} \neg C(u, v_j)] \geq \frac{\lfloor \alpha n \rfloor}{t-1},$$

and so (2.5) holds for all $t \geq \lceil \alpha n \rceil$, thereby completing the argument. \square

With these lemmas, together with the efficient solvability of LP-config, the proof of Theorem 1.1 follows easily:

Proof of Theorem 1.1. Clearly, Algorithm 2 can be implemented efficiently, since LP-config is efficiently solvable, provided it involves a stochastic graph whose probing constraints are permutation-closed. Thus, we focus on proving the algorithm attains the desired asymptotic competitive ratio.

Let us consider the matching \mathcal{M} returned by the algorithm, as well as its weight, which we denote by $w(\mathcal{M})$. Set $\alpha := 1/e$ for clarity, and take $t \geq \lceil \alpha n \rceil$, where we define R_t to be the *unmatched vertices* of U when vertex v_t arrives. Moreover, define e_t as the edge v_t commits to, which is the empty-set by definition if no such commitment is made. Observe that

$$\mathbb{E}[w(\mathcal{M})] = \sum_{t=\lceil \alpha n \rceil}^n \mathbb{E}[w(u_t, v_t) \cdot \mathbf{1}_{[u_t \in R_t]}]. \quad (2.6)$$

Fix $\lceil \alpha n \rceil \leq t \leq n$, and first observe that $w(u_t, v_t)$ and $\{u_t \in R_t\}$ are conditionally independent given (V_t, v_t) , as the probes involving $\partial(v_t)$ are independent from those of v_1, \dots, v_{t-1} . Thus,

$$\mathbb{E}[w(u_t, v_t) \cdot \mathbf{1}_{[u_t \in R_t]} \mid V_t, v_t] = \mathbb{E}[w(u_t, v_t) \mid V_t, v_t] \cdot \mathbb{P}[u_t \in R_t \mid V_t, v_t].$$

Moreover, Lemma 2.6 implies that

$$\mathbb{E}[w(u_t, v_t) \mid V_t, v_t] \cdot \mathbb{P}[u_t \in R_t \mid V_t, v_t] \geq \mathbb{E}[w(u_t, v_t) \mid V_t, v_t] f(t, n),$$

and so $\mathbb{E}[w(u_t, v_t) \mathbf{1}_{[u_t \in R_t]} \mid V_t, v_t] \geq \mathbb{E}[w(u_t, v_t) \mid V_t, v_t] f(t, n)$. Thus, by the law of iterated expectations⁹

$$\begin{aligned} \mathbb{E}[w(u_t, v_t) \cdot \mathbf{1}_{[u_t \in R_t]}] &= \mathbb{E}[\mathbb{E}[w(u_t, v_t) \cdot \mathbf{1}_{[u_t \in R_t]} \mid V_t, v_t]] \\ &\geq \mathbb{E}[\mathbb{E}[w(u_t, v_t) \mid V_t, v_t] f(t, n)] = f(t, n) \mathbb{E}[w(u_t, v_t)]. \end{aligned}$$

⁹ $\mathbb{E}[w(u_t, v_t) \cdot \mathbf{1}_{[u_t \in R_t]} \mid V_t, v_t]$ is a random variable which depends on V_t and v_t , and so the outer expectation is over the randomness in V_t and v_t .

As a result, using (2.6), we get that

$$\mathbb{E}[w(\mathcal{M})] = \sum_{t=\lceil \alpha n \rceil}^n \mathbb{E}[w(u_t, v_t) \mathbf{1}_{[u_t \in R_t]}] \geq \sum_{t=\lceil \alpha n \rceil}^n f(t, n) \mathbb{E}[w(u_t, v_t)].$$

We may thus conclude that

$$\mathbb{E}[w(\mathcal{M})] \geq \text{LPOPT}_{\text{conf}}(G) \sum_{t=\lceil \alpha n \rceil}^n \frac{f(t, n)}{n},$$

after applying Lemma 2.5. As $\sum_{t=\lceil \alpha n \rceil}^n f(t, n)/n \geq (1/e - 1/n)$, the result holds. \square

3 Vertex Weights

Suppose that $G = (U, V, E)$ is a vertex weighted stochastic graph with weights $(w_u)_{u \in U}$. Note that it will be convenient to denote $w_{u,v} := w_u$ provided $(u, v) \in \partial(v)$ for $v \in V$. Let us now fix $s \in V$, and recall that $\text{val}(\mathbf{e})$ is the expected weight of the edge matched, provided the edges of \mathbf{e} are probed in order, where $\mathbf{e} \in \mathcal{C}_s$. Observe then the following claim, which builds upon the work of Brubach et al. [5], and before that, Purohit et al. [22]:

Theorem 3.1. *There exists a dynamical programming (DP) based algorithm DP-OPT, which given access to $G[\{s\} \cup U]$, computes a tuple $\mathbf{e}' \in \mathcal{C}_s$, such that $\text{OPT}(s, U) = \text{val}(\mathbf{e}')$. Moreover, if \mathcal{C}_s is closed under substrings and permutations, then DP-OPT is efficient, assuming access to a membership oracle for \mathcal{C}_s .*

Proof of Theorem 3.1. It will also be convenient to denote $w_{u,s} := w_u$ for each $u \in U$ such that $(u, s) \in \partial(s)$.

We first must show that there exists some $\mathbf{e}' \in \mathcal{C}_s$ such that $\text{val}(\mathbf{e}') = \text{OPT}(s, U)$, where

$$\text{val}(\mathbf{e}) := \sum_{i=1}^{|\mathbf{e}|} p_{e_i} w_{e_i} \prod_{j=1}^{i-1} (1 - p_{e_j}), \quad (3.1)$$

for $\mathbf{e} \in \mathcal{C}_s$, and $\text{OPT}(s, U)$ is the value of the committal benchmark on $G[\{s\} \cup U]$. Since the committal benchmark must respect commitment – i.e., match the first edge to s which it reveals to be active – it is clear that \mathbf{e}' exists.

Let us now additionally assume that \mathcal{C}_s is also closed under permutations. Our goal is to show that \mathbf{e}' can be computed efficiently. Now, for any $\mathbf{e} \in \mathcal{C}_s$, let \mathbf{e}^r be the rearrangement of \mathbf{e} , based on the non-increasing order of the weights $(w_e)_{e \in \mathbf{e}}$. Since \mathcal{C}_s is closed under permutations, we know that \mathbf{e}^r is also in \mathcal{C}_s . Moreover, $\text{val}(\mathbf{e}^r) \geq \text{val}(\mathbf{e})$. Hence, let us order the edges of $\partial(s)$ as e_1, \dots, e_m , such that $w_{e_1} \geq \dots \geq w_{e_m}$, where $m := |\partial(s)|$. Observe then that it suffices to maximize (3.1) over those strings within \mathcal{C}_s which respect this ordering on $\partial(s)$. Stated differently, let us denote \mathcal{I}_s as the family of subsets of $\partial(s)$ induced by \mathcal{C}_s , and define the set function $f : 2^{\partial(s)} \rightarrow [0, \infty)$, where $f(B) := \text{val}(\mathbf{b})$ for $B = \{b_1, \dots, b_{|B|}\} \subseteq \partial(s)$, such that $\mathbf{b} = (b_1, \dots, b_{|B|})$ and $w_{b_1} \geq \dots \geq w_{b_{|B|}}$. Our goal is then to efficiently maximize f over the set-system $(\partial(s), \mathcal{I}_s)$. Observe that since \mathcal{C}_s is both substring-closed and permutation-closed, \mathcal{I}_s is downward closed. Moreover, clearly we can simulate oracle access to \mathcal{I}_s , based on our oracle access to \mathcal{C}_s .

For each $i = 0, \dots, m - 1$, denote $\partial(s)^{>i} := \{e_{i+1}, \dots, e_m\}$, and $\partial(s)^{>m} := \emptyset$. Moreover, define the family of subsets $\mathcal{I}_s^{>i} := \{B \subseteq \partial(s)^{>i} : B \cup \{e_i\} \in \mathcal{I}_s\}$ for each $2 \leq i \leq m$, and $\mathcal{I}_s^{>0} := \mathcal{I}_s$. Observe then that $(\partial(s)^{>i}, \mathcal{I}_s^{>i})$ is a downward-closed set system, as \mathcal{I}_s is downward-closed. Moreover, we may simulate oracle access to $\mathcal{I}_s^{>i}$ based on our oracle access to \mathcal{I}_s .

Denote $\text{OPT}(\mathcal{I}_s^{>i})$ as the maximum value of f over constraints $\mathcal{I}_s^{>i}$. Observe then the following recursion:

$$\text{OPT}(\mathcal{I}_s) := \max_{i \in [m]} (p_{e_i} \cdot w_{e_i} + (1 - p_{e_i}) \cdot \text{OPT}(\mathcal{I}_s^{>i})) \quad (3.2)$$

Hence, given access to the values $\text{OPT}(\mathcal{I}_s^{>1}), \dots, \text{OPT}(\mathcal{I}_s^{>m})$, we can compute $\text{OPT}(\mathcal{I}_s)$ efficiently. In fact, it is clear that we can use (3.2) to recover an optimum solution to f , and so the proof follows by an inductive argument on $|\partial(s)|$. We can define DP-OPT to be a memoization based implementation of (3.2). \square

Given $R \subseteq U$, consider the induced stochastic graph, $G[\{s\} \cup R]$ for $R \subseteq U$ which has probing constraint $\mathcal{C}_s^R \subseteq \mathcal{C}_v$, constructed by restricting \mathcal{C}_s to those strings whose entries all lie in $R \times \{s\}$. Moreover, denote the output of executing DP-OPT on $G[\{s\} \cup R]$ by $\text{DP-OPT}(s, R)$. Consider now the following online probing algorithm, where we assume the online vertices of G arrive in an adversarially chosen unknown order v_1, \dots, v_n , where $n := |V|$.

Algorithm 3 Greedy-DP

Input: offline vertices U with vertex weights $(w_u)_{u \in U}$.

Output: a matching \mathcal{M} of active edges of the unknown stochastic graph $G = (U, V, E)$.

```

1:  $\mathcal{M} \leftarrow \emptyset$ .
2:  $R \leftarrow U$ .
3: for  $t = 1, \dots, n$  do
4:   Let  $v_t$  be the current online arrival node, with constraint  $\mathcal{C}_{v_t}$  and edges probabilities
       $(p_e)_{e \in \partial(v_t)}$ .
5:   Set  $e \leftarrow \text{DP-OPT}(v_t, R)$ 
6:   for  $i = 1, \dots, |e|$  do
7:     Probe  $e_i$ .
8:     if  $\text{st}(e_i) = 1$  then
9:       Add  $e_i$  to  $\mathcal{M}$ , and update  $R \leftarrow R \setminus \{u_i\}$ , where  $e_i = (u_i, v_t)$ .
10:    end if
11:  end for
12: end for
13: return  $\mathcal{M}$ .
```

In general, the behaviour of $\text{OPT}(s, R)$ can change very much, even for minor changes to R . For instance, if $R = U$, then $\text{OPT}(s, U)$ may probe (u, s) first – thus giving it highest priority – whereas by removing $u^* \in U$ from U (where $u^* \neq u$), $\text{OPT}(s, U \setminus \{u^*\})$ may not probe (u, v) at all:

Example 3.2. Let $G = (U, V, E)$ be a bipartite graph with $U = \{u_1, u_2, u_3, u_4\}$, $V = \{v\}$ and $\ell_v = 2$. Set $p_{u_1, v} = 1/3$, $p_{u_2, v} = 1$, $p_{u_3, v} = 1/2$, $p_{u_4, v} = 2/3$. Fix $\varepsilon > 0$, and let the weights of offline vertices be $w_{u_1} = 1 + \varepsilon$, $w_{u_2} = 1 + \varepsilon/2$, $w_{u_3} = w_{u_4} = 1$. We assume that ε is sufficiently small – concretely, $\varepsilon \leq 1/12$. If $R_1 := U$, then $\text{OPT}(v, R_1)$ probes (u_1, v) and then (u_2, v) in order. On the other hand, if $R_2 = R_1 \setminus \{v_2\}$, then $\text{OPT}(v, R_2)$ does *not* probe (u_1, v) . Specifically, $\text{OPT}(v, R_2)$ probes (u_3, v) and then (u_4, v) .

While this behaviour isn't problematic in the case of adversarial arrivals, we must restrict our attention to executions of Algorithm 3 which are less *adaptive* for our primal-dual proof to work in the case of ROM arrivals.

Given a vertex $v \in V$, and an ordering π_v on $\partial(v)$, if $R \subseteq U$, then define $\pi_v(R)$ to be the longest¹⁰ string constructible by iteratively appending the edges of $R \times \{v\}$ via π_v , subject to respecting constraint \mathcal{C}_v^R . We say that v is **rankable**, provided there exists a choice of π_v which depends *solely* on $(p_e)_{e \in \partial(v)}$, $(w_e)_{e \in \partial(v)}$ and \mathcal{C}_v , such that for *every* $R \subseteq U$, the strings $\text{DP-OPT}(v, R)$ and $\pi_v(R)$ are equal. Crucially, if v is rankable, then when vertex v arrives while executing Algorithm 3, one can compute the ranking π_v on $\partial(v)$ and probe the adjacent edges of $R \times \{v\}$ based on this order, subject to not violating the constraint \mathcal{C}_v^R . By following this probing strategy, the optimality of DP-OPT ensures that the expected weight of the match made to v will be $\text{OPT}(v, R)$. We consider three (non-exhaustive) examples of rankability:

Proposition 3.3. *Let $G = (U, V, E)$ be a stochastic graph, and suppose that $v \in V$. If v satisfies either of the following conditions, then v is rankable:*

1. v has unit patience or unlimited patience; that is, $\ell_v \in \{1, |U|\}$.
2. v has patience ℓ_v , and for each $u_1, u_2 \in U$, if $p_{u_1, v} \leq p_{u_2, v}$ then $w_{u_1} \leq w_{u_2}$.
3. G is unweighted, and v has a budget¹¹ B_v with edge probing costs $(c_{u, v})_{u \in U}$, and for each $u_1, u_2 \in U$, if $p_{u_1, v} \leq p_{u_2, v}$ then $c_{u_1, v} \geq c_{u_2, v}$.

We refer to the stochastic graph G as **rankable**, provided all of its vertices are themselves rankable. We emphasize that distinct vertices of V may each use their own separate rankings of their adjacent edges.

Theorem 3.4. *Suppose Algorithm 3 returns the matching \mathcal{M} when executing on the rankable stochastic graph $G = (U, V, E)$ with prefix-closed constraints $(\mathcal{C}_v)_{v \in V}$. In this case,*

$$\mathbb{E}[w(\mathcal{M})] \geq \left(1 - \frac{1}{e}\right) \cdot \text{OPT}(G),$$

provided the vertices of V arrive u.a.r.. Algorithm 3 can be implemented efficiently, provided the constraints $(\mathcal{C}_v)_{v \in V}$ are also permutation-closed.

Remark 3.5. Proposition 3.3 includes the unit patience setting of Theorem 1.3, as well as when G is unweighted and has arbitrary patience values, and so we focus on proving Theorem 3.4 for the remainder of the section.

In order to show that Algorithm 3 attains the claimed competitive ratios, we upper bound $\text{OPT}(G)$ using an LP relaxation which accounts for arbitrary probing constraints. For each $u \in U$ and $v \in V$, let $x_{u, v}$ be a decision variable corresponding to the probability that $\text{OPT}(G)$ probes the

¹⁰Given e' after processing e_1, \dots, e_i via ordering π_v , append e_{i+1} if $(e', e_{i+1}) \in \mathcal{C}_v^R$, else move to e_{i+2} .

¹¹In the case of a budget B_v and edge probing costs $(c_v)_{v \in V}$, any subset of $\partial(v)$ may be probed, provided its cumulative cost does not exceed B_v .

edge (u, v) .

$$\text{maximize} \quad \sum_{u \in U} \sum_{v \in V} w_u \cdot p_{u,v} \cdot x_{u,v} \quad (\text{LP-DP})$$

$$\text{subject to} \quad \sum_{v \in V} p_{u,v} \cdot x_{u,v} \leq 1 \quad \forall u \in U \quad (3.3)$$

$$\sum_{u \in R} w_u \cdot p_{u,v} \cdot x_{u,v} \leq \text{OPT}(v, R) \quad \forall v \in V, R \subseteq U \quad (3.4)$$

$$x_{u,v} \geq 0 \quad \forall u \in U, v \in V \quad (3.5)$$

Denote $\text{LPOPT}_{\text{DP}}(G)$ as the optimum value of this LP. Constraint (3.3) can be viewed as ensuring that the expected number of matches made to $u \in U$ is at most 1. Similarly, (3.4) can be interpreted as ensuring that expected stochastic reward of v , suggested by the solution $(x_{u,v})_{u \in U, v \in V}$, is actually attainable by the committal benchmark. Thus, $\text{OPT}(G) \leq \text{LPOPT}_{\text{DP}}(G)$ (a formal proof specific to patience values is proven in [5]).

3.0.1 Defining the Primal-Dual Charging Schemes

In order to prove Theorems 1.2 and 3.4, we employ primal-dual charging arguments based on the dual of LP-DP. For each $u \in U$, define the variable α_u . Moreover, for each $R \subseteq U$ and $v \in V$, define the variable $\phi_{v,R}$ (these latter variables correspond to constraint (3.4)).

$$\text{minimize} \quad \sum_{u \in U} \alpha_u + \sum_{v \in V} \sum_{R \subseteq U} \text{OPT}(v, R) \cdot \phi_{v,R} \quad (\text{LP-dual-DP})$$

$$\text{subject to} \quad p_{u,v} \cdot \alpha_u + \sum_{\substack{R \subseteq U: \\ u \in R}} w_u \cdot p_{u,v} \cdot \phi_{v,R} \geq w_u \cdot p_{u,v} \quad \forall u \in U, v \in V \quad (3.6)$$

$$\alpha_u \geq 0 \quad \forall u \in U \quad (3.7)$$

$$\phi_{v,R} \geq 0 \quad \forall v \in V, R \subseteq U \quad (3.8)$$

The dual-fitting argument used to prove Theorem 3.4 has an initial set-up which proceeds as in Devanur et al. [8]. Specifically, let $F := 1 - 1/e$ and define $g : [0, 1] \rightarrow [0, 1]$ where $g(z) := \exp(z - 1)$ for $z \in [0, 1]$. For each $v \in V$, draw $Y_v \in [0, 1]$ independently and uniformly at random. We assume that the vertices of V are presented to Algorithm 3 in a non-decreasing order, based on the values of $(Y_v)_{v \in V}$.

We now describe how the charging assignments are made while Algorithm 3 executes on G . Firstly, we initialize a dual solution $((\alpha_u)_{u \in U}, (\phi_{v,R})_{v \in V, R \subseteq U})$ where all the variables are set equal to 0. Let us now take $v \in V, u \in U$, and $R \subseteq U$, where $u \in R$. If R consists of the unmatched vertices of v when it arrives at time Y_v , then suppose that Algorithm 3 matches v to u while making its probes to a subset of the edges of $R \times \{v\}$. In this case, we **charge** $w_u \cdot (1 - g(Y_v))/F$ to α_u and $w_u \cdot g(Y_v)/(F \cdot \text{OPT}(v, R))$ to $\phi_{v,R}$. Observe that each subset $R \subseteq U$ is charged at most once, as is each $u \in U$. Thus, by definition,

$$\mathbb{E}[w(\mathcal{M})] = F \cdot \left(\sum_{u \in U} \mathbb{E}[\alpha_u] + \sum_{v \in V} \sum_{R \subseteq U} \text{OPT}(v, R) \cdot \mathbb{E}[\phi_{v,R}] \right), \quad (3.9)$$

where the expectation is over the random variables $(Y_v)_{v \in V}$ and $(\text{st}(e))_{e \in E}$. If we now set $\alpha_u^* := \mathbb{E}[\alpha_u]$ and $\phi_{v,R}^* := \mathbb{E}[\phi_{v,R}]$ for $u \in U, v \in V$ and $R \subseteq U$, then (3.9) implies the following lemma:

Lemma 3.6. *Suppose $G = (U, V, E)$ is a stochastic graph for which Algorithm 3 returns the matching \mathcal{M} when presented V based on $(Y_v)_{v \in V}$ generated u.a.r. from $[0, 1]$. In this case, if the variables $((\alpha_u^*)_{u \in U}, (\phi_{v,R}^*)_{v \in V, R \subseteq U})$ are defined through the above charging scheme, then*

$$\mathbb{E}[w(\mathcal{M})] = F \cdot \left(\sum_{u \in U} \alpha_u^* + \sum_{v \in V} \sum_{R \subseteq U} OPT(v, R) \cdot \phi_{v,R}^* \right).$$

We also make the following claim regarding the feasibility of the variables $((\alpha_u^*)_{u \in U}, (\phi_{v,R}^*)_{v \in V, R \subseteq U})$:

Lemma 3.7. *If $G = (U, V, E)$ is a rankable stochastic graph whose online nodes are presented to Algorithm 3 based on $(Y_v)_{v \in V}$ generated u.a.r. from $[0, 1]$, then the solution $((\alpha_u^*)_{u \in U}, (\phi_{v,R}^*)_{v \in V, R \subseteq U})$ is a feasible solution to LP-dual-DP.*

Since LP-DP is a relaxation of the committal benchmark, Theorem 3.4 follows from Lemmas 3.6 and 3.7 in conjunction with weak duality. On the other hand, if we redefine $g(z) := 1/2$ and $F := 1/2$, then analogous versions of Lemmas 3.6 and 3.7 hold, even when the values $(Y_v)_{v \in V}$ are generated adversarially and G is not rankable. With these analogous lemmas, Theorem 1.2 follows in the same way. We focus on the ROM setting for the remainder of the section, as the analogous version of Lemma 3.7 for the adversarial setting follows similarly, and is in fact easier to prove.

3.0.2 Proving Dual Feasibility: Lemma 3.7

Let us suppose that the variables $((\alpha_u)_{u \in U}, (\phi_{v,R})_{v \in V, R \subseteq U})$ are defined as in the charging scheme of Section 3.0.1. In order to prove Lemma 3.7, we must show that for each fixed $u_0 \in U$ and $v_0 \in V$, we have that

$$\mathbb{E}[p_{u_0, v_0} \cdot \alpha_{u_0} + w_{u_0} \cdot p_{u_0, v_0} \sum_{\substack{R \subseteq U: \\ u_0 \in R}} \phi_{v_0, R}] \geq w_{u_0} \cdot p_{u_0, v_0}. \quad (3.10)$$

Our strategy for proving (3.10) first involves the same trick used by Devanur et al. [8]. Specifically, we define the stochastic graph $\tilde{G} := (U, \tilde{V}, \tilde{E})$, where $\tilde{V} := V \setminus \{v_0\}$ and $\tilde{E} := G[U \cup \tilde{V}]$. We wish to compare the execution of the algorithm on the instance \tilde{G} to its execution on the instance G . It will be convenient to couple the randomness between these two executions by making the following assumptions:

1. For each $e \in \tilde{E}$, e is active in \tilde{G} if and only if it is active in G .
2. The same random variables, $(Y_v)_{v \in \tilde{V}}$, are used in both executions.

If we now focus on the execution of \tilde{G} , then define the random variable \tilde{Y}_c where $\tilde{Y}_c := Y_{v_c}$ if u_0 is matched to some $v_c \in \tilde{V}$, and $\tilde{Y}_c := 1$ if u_0 remains unmatched after the execution on \tilde{G} . We refer to the random variable \tilde{Y}_c as the **critical time** of vertex u_0 with respect to v_0 . We claim the following lower bound on α_{u_0} in terms of the critical time \tilde{Y}_c . We emphasize that this is the only part of the proof of Theorem 3.4 which requires the rankability of G .

Proposition 3.8. *If G is rankable, then $\alpha_{u_0} \geq \frac{w_{u_0}}{F} (1 - g(\tilde{Y}_c))$.*

Proof of Proposition 3.8. For each $v \in V$, denote $R_v^{\text{af}}(G)$ as the unmatched (remaining) vertices of U right after v is processed (attempts its probes) in the execution on G . We emphasize that if a probe of v yields an active edge, thus matching v , then this match is excluded from $R_v^{\text{af}}(G)$. Similarly, define $R_v^{\text{af}}(\tilde{G})$ in the same way for the execution on \tilde{G} (where v is now restricted to \tilde{V}).

Now, since G is rankable and the constraints $(\mathcal{C}_v)_{v \in V}$ are substring-closed, we can use the coupling between the two executions to inductively prove that

$$R_v^{\text{af}}(G) \subseteq R_v^{\text{af}}(\tilde{G}), \quad (3.11)$$

for each $v \in \tilde{V}$ ¹². Now, since $g(1) = 1$ (by assumption), there is nothing to prove if $\tilde{Y}_c = 1$. Thus, we may assume that $\tilde{Y}_c < 1$, and as a consequence, that there exists some vertex $v_c \in V$ which matches to u_0 at time \tilde{Y}_c in the execution on \tilde{G} .

On the other hand, by assumption we know that $u_0 \notin R_{v_c}^{\text{af}}(\tilde{G})$ and thus by (3.11), that $u_0 \notin R_{v_c}^{\text{af}}(G)$. As such, there exists some $v' \in V$ which probes (u_0, v') and succeeds in matching to u_0 at time $Y_{v'} \leq \tilde{Y}_c$. Thus, since g is monotone,

$$\alpha_{u_0} \geq \frac{w_{u_0}}{F} (1 - g(Y_{v'})) \mathbf{1}_{[\tilde{Y}_c < 1]} \geq \frac{w_{u_0}}{F} (1 - g(\tilde{Y}_c)),$$

and so the claim holds. □

By taking the appropriate conditional expectation, we can also lower bound the random variables $(\phi_{v_0, R})_{\substack{R \subseteq U: \\ u_0 \in R}}$.

Proposition 3.9.

$$\sum_{\substack{R \subseteq U: \\ u_0 \in R}} \mathbb{E}[\phi_{v_0, R} \mid (Y_v)_{v \in \tilde{V}}, (st(e))_{e \in \tilde{E}}] \geq \frac{1}{F} \int_0^{\tilde{Y}_c} g(z) dz.$$

Proof of Proposition 3.9. We first define R_{v_0} as the unmatched vertices of U when v_0 arrives (this is a random subset of U). We also once again use \mathcal{M} to denote the matching returned by Algorithm 3 when executing on G . If we now take a *fixed* subset $R \subseteq U$, then the charging assignment to $\phi_{v_0, R}$ ensures that

$$\phi_{v_0, R} = w(\mathcal{M}(v_0)) \cdot \frac{g(Y_{v_0})}{F \cdot \text{OPT}(v_0, R)} \cdot \mathbf{1}_{[R_{v_0} = R]},$$

where $w(\mathcal{M}(v_0))$ corresponds to the weight of the vertex matched to v_0 (which is zero if v_0 remains unmatched after the execution on G). In order to make use of this relation, let us first condition on the values of $(Y_v)_{v \in V}$, as well as the states of the edges of \tilde{E} ; that is, $(st(e))_{e \in \tilde{E}}$. Observe that once we condition on this information, we can determine $g(Y_{v_0})$, as well as R_{v_0} . As such,

$$\mathbb{E}[\phi_{v_0, R} \mid (Y_v)_{v \in V}, (st(e))_{e \in \tilde{E}}] = \frac{g(Y_{v_0})}{F \cdot \text{OPT}(v_0, R)} \mathbb{E}[w(\mathcal{M}(v_0)) \mid (Y_v)_{v \in V}, (st(e))_{e \in \tilde{E}}] \cdot \mathbf{1}_{[R_{v_0} = R]}.$$

On the other hand, the only randomness which remains in the conditional expectation involving $w(\mathcal{M}(v_0))$ is over the states of the edges adjacent to v_0 . Observe now that since Algorithm 3 behaves optimally on $G[\{v_0\} \cup R_{v_0}]$, we get that

$$\mathbb{E}[w(\mathcal{M}(v_0)) \mid (Y_v)_{v \in V}, (st(e))_{e \in \tilde{E}}] = \text{OPT}(v_0, R_{v_0}), \quad (3.12)$$

and so for the *fixed* subset $R \subseteq U$,

$$\mathbb{E}[w(\mathcal{M}(v_0)) \mid (Y_v)_{v \in V}, (st(e))_{e \in \tilde{E}}] \cdot \mathbf{1}_{[R_{v_0} = R]} = \text{OPT}(v_0, R) \cdot \mathbf{1}_{[R_{v_0} = R]}$$

¹²Example 3.2 shows why (3.11) will not hold if G is not rankable.

after multiplying each side of (3.12) by the indicator random variable $\mathbf{1}_{[R_{v_0}=R]}$. Thus,

$$\mathbb{E}[\phi_{v_0,R} \mid (Y_v)_{v \in V}, (\text{st}(e))_{e \in \tilde{E}}] = \frac{g(Y_{v_0})}{F} \mathbf{1}_{[R_{v_0}=R]},$$

after cancellation. We therefore get that

$$\sum_{\substack{R \subseteq U: \\ u_0 \in R}} \mathbb{E}[\phi_{v_0,R} \mid (Y_v)_{v \in V}, (\text{st}(e))_{e \in \tilde{E}}] = \frac{g(Y_{v_0})}{F} \sum_{\substack{R \subseteq U: \\ u_0 \in R}} \mathbf{1}_{[R_{v_0}=R]}.$$

Let us now focus on the case when v_0 arrives before the critical time; that is, $0 \leq Y_{v_0} < \tilde{Y}_c$. Up until the arrival of v_0 , the executions of the algorithm on \tilde{G} and G proceed identically, thanks to the coupling between the executions. As such, u_0 must be available when v_0 arrives. We interpret this observation in the above notation as saying the following:

$$\mathbf{1}_{[Y_{v_0} < \tilde{Y}_c]} \leq \sum_{\substack{R \subseteq U: \\ u_0 \in R}} \mathbf{1}_{[R_{v_0}=R]}.$$

As a result,

$$\sum_{\substack{R \subseteq U: \\ u_0 \in R}} \mathbb{E}[\phi_{v_0,R} \mid (Y_v)_{v \in V}, (\text{st}(e))_{e \in \tilde{E}}] \geq \frac{g(Y_{v_0})}{F} \mathbf{1}_{[Y_{v_0} < \tilde{Y}_c]}.$$

Now, if we take expectation over Y_{v_0} , while still conditioning on the random variables $(Y_v)_{v \in \tilde{V}}$, then we get that

$$\mathbb{E}[g(Y_{v_0}) \cdot \mathbf{1}_{[Y_{v_0} < \tilde{Y}_c]} \mid (Y_v)_{v \in \tilde{V}}, (\text{st}(e))_{e \in \tilde{E}}] = \int_0^{\tilde{Y}_c} g(z) dz,$$

as Y_{v_0} is drawn uniformly from $[0, 1]$, independently from $(Y_v)_{v \in \tilde{V}}$ and $(\text{st}(e))_{e \in \tilde{E}}$. Thus, after applying the law of iterated expectations,

$$\sum_{\substack{R \subseteq U: \\ u_0 \in R}} \mathbb{E}[\phi_{v_0,R} \mid (Y_v)_{v \in \tilde{V}}, (\text{st}(e))_{e \in \tilde{E}}] \geq \frac{1}{F} \int_0^{\tilde{Y}_c} g(z) dz,$$

and so the claim holds. □

With Propositions 3.8 and 3.9, the proof of Lemma 3.7 follows easily:

Proof of Lemma 3.7. We first observe that by taking the appropriate conditional expectation, Proposition 3.8 ensures that

$$\mathbb{E}[\alpha_{u_0} \mid (Y_v)_{v \in \tilde{V}}, (\text{st}(e))_{e \in \tilde{E}}] \geq \frac{w_{u_0}}{F} \cdot (1 - g(\tilde{Y}_c)),$$

where the right-hand side follows since \tilde{Y}_c is entirely determined from $(Y_v)_{v \in \tilde{V}}$ and $(\text{st}(e))_{e \in \tilde{E}}$. Thus, combined with Proposition 3.9,

$$\mathbb{E}[p_{u_0, v_0} \cdot \alpha_{u_0} + w_{u_0} \cdot p_{u_0, v_0} \cdot \sum_{\substack{R \subseteq U: \\ u_0 \in R}} \phi_{v,R} \mid (Y_v)_{v \in \tilde{V}}, (\text{st}(e))_{e \in \tilde{E}}],$$

is lower bounded by

$$\frac{w_{u_0} \cdot p_{u_0, v_0}}{F} \cdot (1 - g(\tilde{Y}_c)) + \frac{w_{u_0} p_{u_0, v_0}}{F} \int_0^{\tilde{Y}_c} g(z) dz.$$

However, $g(z) := \exp(z - 1)$ for $z \in [0, 1]$ by assumption, so

$$(1 - g(\tilde{Y}_c)) + \int_0^{\tilde{Y}_c} g(z) dz = \left(1 - \frac{1}{e}\right),$$

no matter the value of the critical time \tilde{Y}_c . As such, since $F := 1 - 1/e$, we may apply the law of iterated expectations and conclude that

$$\mathbb{E}[p_{u_0, v_0} \cdot \alpha_{u_0} + w_{u_0} \cdot p_{u_0, v_0} \cdot \sum_{\substack{R \subseteq U: \\ u_0 \in R}} \phi_{v, R}] \geq w_{u_0} \cdot p_{u_0, v_0}.$$

As the vertices $u_0 \in U$ and $v_0 \in V$ were chosen arbitrarily, the proposed dual solution of Lemma 3.7 is feasible, and so the proof is complete. \square

4 Conclusion and Open Problems

We considered the online stochastic bipartite matching with commitment in a number of different settings establishing several competitive bounds against the committal benchmark. In Appendix A, we indicate when our results hold against a stronger non-committal benchmark.

In the case of vertex-weighted stochastic graphs, adversarial arrivals, and general probing constraints, we provide a deterministic algorithm that achieves a $\frac{1}{2}$ competitive ratio. This is an optimal competitive ratio for deterministic algorithms and adversarial arrivals.

In the case of the random order model, we provide two results. First, for edge weighted stochastic graphs, and general probing constraints, we provide a randomized algorithm that achieves the optimal asymptotic competitive ratio of $\frac{1}{e}$. For vertex weighted results we provide an algorithm that achieves a $1 - 1/e$ competitive ratio whenever the input graph is “rankable”. Our rankable assumption subsumes most of the stochastic graph settings studied in previous works.

Our work leaves open a number of challenging open problems. For context, we note that currently, even for the classical (i.e., non-stochastic) setting, $1 - \frac{1}{e}$ is the best known ratio for deterministic algorithms operating on unweighted or vertex weighted graphs with random vertex arrivals. The best known ROM in-approximation of 0.823 (due to Manshadi et al. [19]) comes from the classical i.i.d. unweighted graph setting for a known distribution and applies to randomized as well as deterministic algorithms.

- What is the best ratio that a deterministic or randomized online algorithm can obtain for *all* stochastic graphs in the ROM setting? That is, what competitive ratio can be achieved without the rankable assumption? Is there an online probing algorithm which can surpass the $1 - 1/e$ “barrier”? In [4], we show that $1 - 1/e$ is a hardness result for *non-adaptive* online probing algorithms, even for the unweighted unit patience setting when the stochastic graph is known to the algorithm.
- Is there a provable difference between what an optimal online algorithm can obtain against the committal benchmark versus the non-committal benchmark? Specifically, does Algorithm 3 achieve a competitive ratio of $1 - 1/e$ against the non-committal benchmark which holds for

all rankable stochastic graphs or for all stochastic graphs? The hardness result of Proposition A.1 suggests that Algorithm 3 does *not* attain a competitive ratio of $1 - 1/e$, even for rankable stochastic graphs.

- What is the best ratio that a randomized online algorithm can obtain for stochastic graphs in the adversarial arrival model? The Mehta and Panigrahi [20] 0.621 inapproximation shows that randomized probing algorithms (even for unweighted graphs and unit patience) cannot achieve a $1 - 1/e$ performance guarantee against LP-std-unit, however the work of Goyal and Udwani [11] suggests that this is because LP-std-unit is too loose a relaxation of the committal benchmark.
- Is there a online stochastic matching problem in which the optimum competitive ratio provably worse than the optimal ratio for the corresponding classical setting? Note that in the classical setting the benchmark is the weight of an offline optimal matching.
- Can our $1 - \frac{1}{e}$ competitive ratio be improved by a randomized algorithm in the vertex-weighted ROM setting? Here we note that in the classical ROM setting, the RANKING algorithm achieves a 0.696 ratio for unweighted graphs (due to Mahdian and Yan [18]) and a 0.6534 ratio (due to Huang et al. [12]) for vertex weighted graphs. Thus, randomization seems to significantly help in the classical ROM setting.

References

- [1] Marek Adamczyk, Fabrizio Grandoni, Stefano Leonardi, and Michal Włodarczyk. When the optimum is also blind: a new perspective on universal optimization. In *ICALP*, 2017.
- [2] Arash Asadpour and Hamid Nazerzadeh. Maximizing stochastic monotone submodular functions. *Management Science*, 62(8):2374–2391, 2016.
- [3] Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, 2012.
- [4] Allan Borodin, Calum MacRury, and Akash Rakheja. Prophet inequality matching meets probing with commitment. *CoRR*, abs/2102.04325, 2021.
- [5] Brian Brubach, Nathaniel Grammel, and Aravind Srinivasan. Vertex-weighted online stochastic matching with patience constraints, 2019.
- [6] Ning Chen, Nicole Immorlica, Anna R. Karlin, Mohammad Mahdian, and Atri Rudra. Approximating matches made in heaven. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part I*, ICALP '09, pages 266–278, 2009.
- [7] Kevin P. Costello, Prasad Tetali, and Pushkar Tripathi. Stochastic matching with commitment. In Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming*, pages 822–833, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [8] Nikhil R. Devanur, Kamal Jain, and Robert D. Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 101–107, Philadelphia, PA, USA, 2013. Society for Industrial and Applied Mathematics.

- [9] Hao Fu, Jian Li, and Pan Xu. A PTAS for a Class of Stochastic Dynamic Programs. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 56:1–56:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [10] Bernd Gärtner and Jiri Matousek. *Understanding and using linear programming*. Universitext. Springer, 2007.
- [11] Vineet Goyal and R. Udvani. Online matching with stochastic rewards: Optimal competitive ratio via path based formulation. *Proceedings of the 21st ACM Conference on Economics and Computation*, 2020.
- [12] Zhiyi Huang, Zhihao Tang, Xiaowei Wu, and Yuhao Zhang. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. *ACM Transactions on Algorithms*, 15, 04 2018.
- [13] Zhiyi Huang and Qiankun Zhang. Online primal dual meets online matching with stochastic rewards: Configuration lp to the rescue. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 1153–1164, New York, NY, USA, 2020. Association for Computing Machinery.
- [14] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 587–596, 2011.
- [15] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for online bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 352–358, 1990.
- [16] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In Hans L. Bodlaender and Giuseppe F. Italiano, editors, *Algorithms – ESA 2013*, pages 589–600, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [17] Euiwoong Lee and Sahil Singla. Optimal Online Contention Resolution Schemes via Ex-Ante Prophet Inequalities. In Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 57:1–57:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [18] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: An approach based on strongly factor-revealing lps. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 597–606, New York, NY, USA, 2011. ACM.
- [19] Vahideh H. Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. *Math. Oper. Res.*, 37(4):559–573, 2012.
- [20] Aranyak Mehta and Debmalya Panigrahi. Online matching with stochastic rewards. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 728–737. IEEE Computer Society, 2012.

- [21] Aranyak Mehta, Bo Waggoner, and Morteza Zadimoghaddam. Online stochastic matching with unequal probabilities. In *SODA*, pages 1388–1404, 2015.
- [22] Manish Purohit, Sreenivas Gollapudi, and Manish Raghavan. Hiring under uncertainty. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5181–5189. PMLR, 09–15 June 2019.
- [23] D. Seese. Groetschel, m., l. lovasz, a. schrijver: Geometric algorithms and combinatorial optimization. (algorithms and combinatorics. eds.: R. l. graham, b. korte, l. lovasz. vol. 2), springer-verlag 1988, xii, 362 pp., 23 figs., dm 148,-. isbn 3-540-13624-x. *Biometrical Journal*, 32(8):930–930, 1990.
- [24] Danny Segev and Sahil Singla. Efficient approximation schemes for stochastic probing and prophet problems. *CoRR*, abs/2007.13121, 2020.
- [25] Jan Vondrák, Chandra Chekuri, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 783–792, New York, NY, USA, 2011. Association for Computing Machinery.
- [26] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, USA, 1st edition, 2011.

A The Non-committal Benchmark

In this section, we introduce the **non-committal benchmark**. This benchmark must still adaptively probe edges subject to probing constraints, and its goal is the same as the committal benchmark, but it does not need to respect commitment. More precisely, if $P_a \subseteq E$ corresponds to the active probes made by the benchmark, then it returns a matching $\mathcal{M} \subseteq P_a$ of maximum weight. We denote $\text{OPT}_{\text{non}}(G)$ as the expected weight of the matching that the non-committal benchmark constructs, and abuse notation slightly by also using $\text{OPT}_{\text{non}}(G)$ to refer to the *strategy* of the non-committal benchmark on G . Observe that in the case of unlimited patience, $\text{OPT}_{\text{non}}(G)$ may probe all the edges of G , and thus corresponds to the expected weight of the optimum matching of the stochastic graph. Clearly, for any set of probing constraints, the non-committal benchmark is no weaker than the committal benchmark; that is, $\text{OPT}_{\text{non}}(G) \geq \text{OPT}(G)$ for any stochastic graph G . We first show that these values are separated by a ratio of at least 0.856269, even for a single online node.

Proposition A.1. *There exists a stochastic graph G with a single online node v , such that $\text{OPT}(G) = 0.856269 \cdot \text{OPT}_{\text{non}}(G)$.*

Proof of Proposition A.1. Suppose G has a single online node v with patience $\ell_v = 2$, and that there are 3 offline nodes $U = \{u_1, u_2, u_3\}$. For each $i \in \{1, 2, 3\}$, we denote the weight of (u_i, v) by w_i and assume that the edge (u_i, v) is active with probability p_i . We make the following assumptions on these weights and probabilities:

1. $w_1 < w_2 < w_3$.
2. $p_1 > p_2 > p_3$.

3. $w_1 \cdot p_1 \geq w_2 \cdot p_2 > w_3 \cdot p_3$.
4. $p_2 \cdot w_2 - p_3 \cdot w_3 \geq p_1 \cdot w_1 \cdot (p_2 - p_3)$

Clearly, there exists a choice of weights and probabilities which satisfy these constraints. For instance, take $w_1 = 3$, $w_2 = 4$, $w_3 = 98$, $p_1 = 0.8$, $p_2 = 0.6$, and $p_3 = 0.01$. Based on these assumptions let us now consider the value of $\text{OPT}(G)$. Observe that

$$p_2 \cdot w_2 + (1 - p_2) \cdot p_1 \cdot w_1 \geq p_3 \cdot w_3 + (1 - p_3) \cdot p_1 \cdot w_1 \geq p_3 \cdot w_3 + (1 - p_3) \cdot p_2 \cdot w_2,$$

where the first inequality follows from (4), and the second follows from (3). As a result, it is clear to see that the committal benchmark corresponds to probing (u_2, v) and then (u_1, v) (if necessary); thus, $\text{OPT}(G) = p_2 \cdot w_2 + (1 - p_2) \cdot p_1 \cdot w_1$. On the other hand, let us consider $\text{OPT}_{\text{non}}(G)$, the value of the non-committal benchmark on G . Consider the following *non-committal* probing algorithm:

- Probe (u_2, v) , and if $\text{st}(u_2, v) = 1$, probe (u_3, v) .
- Else if $\text{st}(u_2, v) = 0$, probe (u_1, v) .
- Return the edge of highest weight which is active (if any).

Clearly, this probing algorithm uses adaptivity to decide whether to reveal (u_3, v) or (u_1, v) in its second probe. Specifically, if it discovers that (u_2, v) is active, then it knows that it will return an edge with weight at least w_2 . As such, it only makes sense for (u_3, v) to be its next probe, as $w_3 > w_2 > w_1$. On the other hand, if (u_2, v) is discovered to be inactive, it makes sense to prioritize probing the edge (u_1, v) over (u_3, v) , as the *expected* reward is higher; namely, $w_1 \cdot p_1 > w_3 \cdot p_3$. Thus, it is clear that the expected weight of the edge returned is

$$p_2 \cdot p_3 \cdot w_3 + p_2 \cdot (1 - p_3) \cdot w_2 + (1 - p_2) \cdot p_1 \cdot w_1.$$

Observe however that

$$\begin{aligned} p_2 \cdot p_3 \cdot w_3 + p_2 \cdot (1 - p_3) \cdot w_2 + (1 - p_2) \cdot p_1 \cdot w_1 &> p_2 \cdot w_2 + (1 - p_2) \cdot p_1 \cdot w_1 \\ &= \text{OPT}(G), \end{aligned}$$

where the final inequality follows since $w_3 > w_2$. As a result, it is clear that this strategy corresponds to the non-committal benchmark, and so $\text{OPT}_{\text{non}}(G) > \text{OPT}(G)$. In fact, for the specific choice when $w_1 = 3$, $w_2 = 4$, $w_3 = 98$, $p_1 = 0.8$, $p_2 = 0.6$, and $p_3 = 0.01$, it holds that $\text{OPT}(G) = 0.856269 \cdot \text{OPT}_{\text{non}}(G)$. \square

This example slightly improves upon the negative result of [7], in which Costello et al. show that the ratio between $\text{OPT}(G)$ and $\text{OPT}_{\text{non}}(G)$ is at most 0.898 (albeit for unweighted graphs).

We remark that restricted to a single online node, the non-committal probing problem is a special case of the *adaptive* version of **PROBLEMMAX**, a stochastic probing problem which is studied in [2, 9, 24]. Similarly, the committal probing problem is a special case of the *non-adaptive* version of **PROBLEMMAX**, which is also considered in [2, 9, 24]. Thus, we can view 0.856269 as an upper bound (negative result) on the **adaptivity gap** of **PROBLEMMAX**. This is in contrast to the lower bound of $1 - 1/e$ on the adaptivity gap of **PROBLEMMAX** proven by Asadpour et al. [2].

We now provide a number of settings in which the competitive ratios of Algorithm 3 holds against the non-committal benchmark.

Theorem A.2. *Let G be a vertex-weighted stochastic graph such that for each $v \in V$ and $R \subseteq U$, $OPT(v, R) = OPT_{\text{non}}(v, R)$. In this case, Algorithm 3 attains a performance guarantee of $1/2$ against $OPT_{\text{non}}(G)$ assuming adversarial arrivals, and $1-1/e$ against $OPT_{\text{non}}(G)$ assuming random order arrivals and the rankability of G .*

Remark A.3. If G is unweighted or is vertex-weighted yet has unit/unlimited patience, then Theorem A.2 applies.

To prove Theorem A.2, consider the following LP, where each edge $e \in E$ of the edge weighted stochastic graph $G = (U, V, E)$ is associated with two variables, namely x_e and z_e . We interpret the former variable as the probability that the non-committal benchmark probes the edge e , whereas the latter variable corresponds to the probability that e is included in the matching constructed by the non-committal benchmark. Note that for convenience, we assume that $E = U \times V$.

$$\text{maximize} \quad \sum_{u \in U, v \in V} w_{u,v} \cdot z_{u,v} \quad (\text{LP-DP-non})$$

$$\text{subject to} \quad \sum_{v \in V} z_{u,v} \leq 1 \quad \forall u \in U \quad (\text{A.1})$$

$$\sum_{u \in R} w_{u,v} \cdot z_{u,v} \leq OPT_{\text{non}}(v, R) \quad \forall v \in V, R \subseteq U \quad (\text{A.2})$$

$$z_{u,v} \leq p_{u,v} \cdot x_{u,v} \quad \forall u \in U, v \in V \quad (\text{A.3})$$

$$x_{u,v}, z_{u,v} \geq 0 \quad \forall u \in U, v \in V \quad (\text{A.4})$$

We denote $LPOPT_{\text{DP-non}}(G)$ as the value of an optimum solution to LP-DP-non.

Lemma A.4. *For any stochastic graph $G = (U, V, E)$ with substring-closed probing constraints $(\mathcal{C}_v)_{v \in V}$, $OPT_{\text{non}}(G) \leq LPOPT_{\text{DP-non}}(G)$.*

Proof. Let us suppose that \mathcal{M} is the matching returned by the non-committal benchmark when it executes on $G = (U, V, E)$. If we fix $u \in U$ and $v \in V$, then define $x_{u,v}$ as the probability the non-committal benchmark probes the edge (u, v) , and $z_{u,v}$ as the probability that it includes e in \mathcal{M} . Observe then that

$$OPT_{\text{non}}(G) = \mathbb{E}[w(\mathcal{M})] = \sum_{u \in U, v \in V} w_{u,v} \cdot z_{u,v}.$$

Thus, we need only show that $(x_{u,v}, z_{u,v})_{u \in U, v \in V}$ is a feasible solution to LP-DP-non, as this will ensure that $OPT_{\text{non}}(G) = \sum_{u \in U, v \in V} w_{u,v} \cdot z_{u,v} \leq LPOPT_{\text{DP-non}}(G)$. If we first fix $u \in U$ and $v \in V$, then observe that in order for (u, v) to be included in \mathcal{M} , (u, v) must be probed *and* (u, v) must be active. On the other hand, these two events occur independently of each other. As such, $z_{u,v} \leq p_{u,v} \cdot x_{u,v}$. Now, each $u \in U$ is matched at most once by the non-committal benchmark, thus $\sum_{v \in V} z_{u,v} \leq 1$. Finally, fix $v \in V$, and denote $\mathcal{M}(v)$ as the edge matched to v (which is \emptyset by convention if v remains unmatched), and denote $w(\mathcal{M}(v))$ as the weight of the edge v is matched to (which is 0 provided v remains unmatched). Observe first that $\sum_{u \in U} w_{u,v} \cdot z_{u,v} = \mathbb{E}[w(\mathcal{M}(v))]$. Moreover, executing the non-committal benchmark on G induces¹³ a probing strategy on $G[\{v\} \cup U]$, which we denote by \mathcal{B}_v . However, observe that since the non-committal benchmark decides upon

¹³The strategy \mathcal{B}_v can be defined formally by first drawing (simulated) independent copies of the edge states which are not adjacent to v , say $\tilde{\text{st}}(e)_{e \in E: v \notin e}$. By executing the non-committal benchmark on G with $\tilde{\text{st}}(e)_{e \in E: v \notin e}$ and $\text{st}(e)_{e \in \partial(v)}$, we get the desired strategy on $G[\{v\} \cup U]$.

which edges to match after all its probes are made, so does \mathcal{B}_v . Specifically, the match it makes to v is determined once all its probes to $U \times \{v\}$ are made. Clearly, the expected value of this match is equal to $\mathbb{E}[w(\mathcal{M}(v))]$ and can be no larger than $\text{OPT}_{\text{non}}(v, U)$. Thus,

$$\sum_{u \in U} w_{u,v} \cdot z_{u,v} = \mathbb{E}[w(\mathcal{M}(v))] \leq \text{OPT}_{\text{non}}(v, U).$$

More generally, if we fix $R \subseteq U$, then

$$\sum_{u \in U} w_{u,v} \cdot z_{u,v} = \mathbb{E}[w(\mathcal{M}(v)) \cdot \mathbf{1}_{[\mathcal{M}(v) \in R \times \{v\}]}] \leq \text{OPT}_{\text{non}}(v, R).$$

To see this, consider a modification of \mathcal{B}_v , say $\mathcal{B}_v(R)$, which matches v to $u \in U$ if and only if \mathcal{B}_v matches v to u and $(u, v) \in R \times \{v\}$.

This shows that all the constraints of LP-DP-non hold for $(x_{u,v}, z_{u,v})_{u \in U, v \in V}$, and so the proof is complete. \square

Proof of Theorem A.2. Let us suppose that G is vertex-weighted.

Consider a modified version of LP-DP in which the right-hand side of constraint (3.4) is replaced by the analogous expression for the non-committal benchmark (i.e., $\text{OPT}_{\text{non}}(v, R)$). It is not hard to show that this modified LP is a reformulation of LP-DP-non. Thus, because of the assumptions on G , $\text{LPOPT}_{\text{DP}}(G) = \text{LPOPT}_{\text{DP-non}}(G)$. Since the performance guarantees of Algorithm 3 in Theorems 1.2 and 3.4 are proven against $\text{LPOPT}_{\text{DP}}(G)$, Theorem A.2 follows thanks to Lemma A.4. \square

B Relaxing the Committal Benchmark

Suppose that we are given an arbitrary stochastic graph $G = (U, V, E)$. Let us restate LP-config for convenience:

$$\begin{array}{ll} \text{maximize} & \sum_{v \in V} \sum_{e \in \mathcal{C}_v} \text{val}(e) \cdot x_v(e) \end{array} \quad (\text{LP-config})$$

$$\begin{array}{ll} \text{subject to} & \sum_{v \in V} \sum_{\substack{e \in \mathcal{C}_v: \\ (u,v) \in e}} p_{u,v} \cdot g(e_{<(u,v)}) \cdot x_v(e) \leq 1 \quad \forall u \in U \end{array} \quad (\text{B.1})$$

$$\sum_{e \in \mathcal{C}_v} x_v(e) = 1 \quad \forall v \in V, \quad (\text{B.2})$$

$$x_v(e) \geq 0 \quad \forall v \in V, e \in \mathcal{C}_v \quad (\text{B.3})$$

We contrast LP-config with LP-std, which is defined only when G has patience values $(\ell_v)_{v \in V}$:

$$\begin{array}{ll} \text{maximize} & \sum_{e \in E} w_e \cdot p_e \cdot x_e \end{array} \quad (\text{LP-std})$$

$$\begin{array}{ll} \text{subject to} & \sum_{e \in \partial(u)} p_e \cdot x_e \leq 1 \quad \forall u \in U \end{array} \quad (\text{B.4})$$

$$\sum_{e \in \partial(v)} p_e \cdot x_e \leq 1 \quad \forall v \in V \quad (\text{B.5})$$

$$\sum_{e \in \partial(v)} x_e \leq \ell_v \quad \forall v \in V \quad (\text{B.6})$$

$$0 \leq x_e \leq 1 \quad \forall e \in E. \quad (\text{B.7})$$

Observe that LP-config and LP-std are the same LP in the case of unit patience:

$$\begin{array}{ll} \text{maximize} & \sum_{v \in V} \sum_{e \in \partial(v)} w_e \cdot p_e \cdot x_e \end{array} \quad (\text{LP-std-unit})$$

$$\begin{array}{ll} \text{subject to} & \sum_{e \in \partial(u)} p_e \cdot x_e \leq 1 \quad \forall u \in U \end{array} \quad (\text{B.8})$$

$$\sum_{e \in \partial(v)} x_e \leq 1 \quad \forall v \in V \quad (\text{B.9})$$

$$x_e \geq 0 \quad \forall e \in E \quad (\text{B.10})$$

For completeness, we now present the essential ideas used in the proof of Theorem 2.1, which shows that LP-config relaxes the committal benchmark.

Let us suppose that hypothetically we could make the following assumption regarding the committal benchmark:

(P_1) If $e = (u, v)$ is probed and $\text{st}(e) = 1$, then e is included in the matching, provided v is currently unmatched.

(P_2) For each $v \in V$, the edge probes involving $\partial(v)$ are made independently of the edge states $(\text{st}(e))_{e \in \partial(v)}$.

Observe then that (P_1) and (P_2) would imply that the expected weight of the edge assigned to v is $\sum_{e \in \mathcal{C}_v} \text{val}(e) \cdot x_v(e)$. Moreover, the left-hand side of (2.1) would correspond to the probability that $u \in U$ is matched, so $(x_v(e))_{v \in V, e \in \mathcal{C}_v}$ would be a feasible solution to LP-config, and so we could upper bound $\text{OPT}(G)$ by $\text{LPOPT}(G)$. Now, if we knew that the committal benchmark adhered to some adaptive vertex ordering π on V (i.e., it chooses $v_{\pi(i)}$ based on $v_{\pi(1)}, \dots, v_{\pi(i-1)}$, and probes the edges of $v_{\pi(i)}$ before moving to $v_{\pi(i+1)}$), then it is clear that we could assume (P_1) and (P_2) simultaneously¹⁴ *w.l.o.g.*. However, clearly a probing algorithm with this restriction is in general less powerful than the committal benchmark. As such, the natural interpretation of the variables of LP-config does not seem to easily lend itself to a proof of Theorem 2.1.

In order to get around these issues, we first discuss the relaxed stochastic matching problem defined in Section 2 in more detail. A solution to this problem is a relaxed probing algorithm. A relaxed probing algorithm operates on the stochastic graph $G = (U, V, E)$ in the same framework as an offline probing algorithm. That is, initially the edge states $(\text{st}(e))_{e \in E}$ are unknown to the

¹⁴It is clear that we may assume the committal benchmark satisfies (P_1) *w.l.o.g.*, but not (P_2).

algorithm, and it must adaptively probe the edges of G , while respecting the probing constraints of the online nodes of G . Its output is then a subset of its probes which yielded active edges, which we denote by \mathcal{N} . The goal of the relaxed probing algorithm is to maximize the expected weight of \mathcal{N} , while ensuring that the following properties are satisfied:

1. Each $v \in V$ appears in at most one edge of \mathcal{N} .
2. For each $u \in U$, the *expected* number of edges of \mathcal{N} which contain u is at most one.

We refer to \mathcal{N} as a **one-sided matching** of the online nodes. We abuse terminology slightly, and say that a vertex of G is matched by \mathcal{N} , provided it is included in an edge of \mathcal{N} . A relaxed probing algorithm must **respect commitment**. That is, it has the property that if a probe to $e = (u, v)$ yields an active edge, then the edge is included in \mathcal{N} (provided v is currently not in \mathcal{N}). Observe that this requires the relaxed probing algorithm to include e , even if u is already adjacent to some element of \mathcal{N} .

We define the **relaxed benchmark** as an optimum relaxed probing algorithm on G , and denote $\text{OPT}_{\text{rel}}(G)$ as the expected value of its output when executing on G . Observe that by definition, $\text{OPT}(G) \leq \text{OPT}_{\text{rel}}(G)$, where $\text{OPT}(G)$ is the value of the committal benchmark on G .

Finally, we say that a relaxed probing algorithm is **non-adaptive**, provided its edge probes are statistically independent from the edge states of G ; that is, the random variables $(\text{st}(e))_{e \in E}$. We emphasize that this is equivalent to (randomly) specifying an ordering π on a subset of E , which satisfies the constraints $(\mathcal{C}_v)_{v \in V}$. The edges specified by π are then probed in order, and an active edge is added to the matching, provided its *online* node is unmatched.

Unlike the committal benchmark, $\text{OPT}_{\text{rel}}(G)$ can be attained by a non-adaptive relaxed probing algorithm.

Lemma B.1 (Lemma 3.2 in [4]). *There exists a relaxed probing algorithm which is non-adaptive and attains value $\text{OPT}_{\text{rel}}(G)$ in expectation.*

We refer the reader to [4] for a complete proof of Lemma B.1, and instead show how it allows us to prove Theorem 2.1. In fact, we prove that LP-config encodes the value of the relaxed benchmark exactly, thus implying Theorem 2.1 since $\text{OPT}(G) \leq \text{OPT}_{\text{rel}}(G)$.

Theorem B.2 (Theorem 3.3 in [4]). *For any stochastic graph G , an optimum solution to LP-config has value equal to $\text{OPT}_{\text{rel}}(G)$, the value of the relaxed benchmark on G .*

Proof. Suppose we are presented a solution $(x_v(\mathbf{e}))_{v \in V, \mathbf{e} \in \mathcal{C}_v}$ to LP-config. We can then define the following relaxed probing algorithm:

1. $\mathcal{N} \leftarrow \emptyset$.
2. For each $v \in V$, set $e \leftarrow \text{VERTEXPROBE}(v, \partial(v), (x_v(\mathbf{e}))_{\mathbf{e} \in \mathcal{C}_v})$. If $e \neq \emptyset$, then add $e = (u, v)$ to \mathcal{N} , provided v is currently unmatched.
3. Return \mathcal{N} .

Using Lemma 2.2, it is clear that $\mathbb{E}[w(\mathcal{N})] = \sum_{v \in V} \sum_{\mathbf{e} \in \mathcal{C}_v} \text{val}(\mathbf{e}) \cdot x_v(\mathbf{e})$. Moreover, each vertex $u \in U$ is matched by \mathcal{N} at most once in expectation, as a consequence of (B.1).

In order to complete the proof, it remains to show that if \mathcal{A} is an optimum relaxed probing algorithm, then there exists a solution to LP-config whose value is equal to $\mathbb{E}[w(\mathcal{A}(G))]$ (where $\mathcal{A}(G)$ is the one-sided matching returned by \mathcal{A}). In fact, by Lemma B.1, we may assume that \mathcal{A} is

non-adaptive. Observe then that for each $v \in V$ and $\mathbf{e} = (e_1, \dots, e_k) \in \mathcal{C}_v$ with $k := |\mathbf{e}| \geq 1$ we can define

$$x_v(\mathbf{e}) := \mathbb{P}[\mathcal{A} \text{ probes the edges } (e_i)_{i=1}^k \text{ in order}],$$

where $x_v(\lambda)$ corresponds to the probability no edge adjacent to v is probed. Setting $\mathcal{N} = \mathcal{A}(G)$ for convenience, observe that if $w(\mathcal{N}(v))$ corresponds to the weight of the edge assigned to v (which is 0 if no assignment is made), then

$$\mathbb{E}[w(\mathcal{N}(v))] = \sum_{\mathbf{e} \in \mathcal{C}_v} \text{val}(\mathbf{e}) \cdot x_v(\mathbf{e}),$$

as \mathcal{A} is non-adaptive. Moreover, for each $u \in U$,

$$\sum_{v \in V} \sum_{\substack{\mathbf{e} \in \mathcal{C}_v: \\ (u,v) \in \mathbf{e}}} p_{u,v} \cdot g(\mathbf{e}_{<(u,v)}) \cdot x_v(\mathbf{e}) \leq 1$$

by once again using the non-adaptivity of \mathcal{A} . The proof is therefore complete. □