

Improved Online Correlated Selection ^{*}

Ruiquan Gao [†] Zhongtian He [†] Zhiyi Huang [‡] Zipei Nie [§] Bijun Yuan [‡]
Yan Zhong [‡]

December 2021

Abstract

This paper studies the online correlated selection (OCS) problem. It was introduced by Fahrbach, Huang, Tao, and Zadimoghaddam (2020) to obtain the first edge-weighted online bipartite matching algorithm that breaks the 0.5 barrier. Suppose that we receive a pair of elements in each round and immediately select one of them. Can we select with negative correlation to be more effective than independent random selections? Our contributions are threefold. For semi-OCS, which considers the probability that an element remains unselected after appearing in k rounds, we give an optimal algorithm that minimizes this probability for all k . It leads to 0.536-competitive unweighted and vertex-weighted online bipartite matching algorithms that randomize over only two options in each round, improving the 0.508-competitive ratio by Fahrbach et al. (2020). Further, we develop the first multi-way semi-OCS that allows an arbitrary number of elements with arbitrary masses in each round. As an application, it rounds the BALANCE algorithm in unweighted and vertex-weighted online bipartite matching and is 0.593-competitive. Finally, we study OCS, which further considers the probability that an element is unselected in *an arbitrary subset of rounds*. We prove that the optimal “level of negative correlation” is between 0.167 and 0.25, improving the previous bounds of 0.109 and 1 by Fahrbach et al. (2020). Our OCS gives a 0.519-competitive edge-weighted online bipartite matching algorithm, improving the previous 0.508-competitive ratio by Fahrbach et al. (2020).

^{*}This is the second version on arXiv. Compared to the first version, this one adds a discussion on two concurrent works on the same topic, gives a more accurate description of previous results, and improves the presentation based on the feedbacks by anonymous reviewers. The conference version appears in FOCS 2021.

[†]IIS, Tsinghua University. Email: grq18@mails.tsinghua.edu.cn, hezt18@mails.tsinghua.edu.cn.

[‡]The University of Hong Kong. Email: zhiyi@cs.hku.hk, bisjack@connect.hku.hk, yanzhong.cs@gmail.com.

[§]Nine-Chapter Lab, Huawei. Email: niezipei@huawei.com.

Contents

1	Introduction	1
1.1	Our Contributions	2
1.2	Other Related Works	4
1.3	Concurrent Works	5
2	Preliminaries	5
3	Optimal Semi-OCS	6
3.1	Algorithms	6
3.2	Negative Correlation in Weighted 2-Way Sampling without Replacement	6
3.3	Analysis	7
3.4	Hardness	8
4	Multi-way Semi-OCS	8
4.1	Definitions	8
4.2	Algorithm: Weighted Sampling without Replacement	9
4.3	Analysis	9
5	Improved Algorithms and Hardness for OCS	12
5.1	Definitions	12
5.2	Roadmap	12
5.2.1	Matching-based Approach versus Automata-based Approach on a Path	12
5.2.2	Automata-based Approach	13
5.3	Forest Constructor	15
5.3.1	Warm-up: Good Online Selection Instances	15
5.3.2	General Online Selection Instances, Pseudo-paths, and Pseudo-matchings	16
5.3.3	Forest Constructor for General Instances	17
5.3.4	Properties of the Automata	18
5.3.5	Analysis of Forest Constructor for General Instances: Proof of Lemma 11	20
5.4	Forest OCS	23
5.4.1	Algorithm	23
5.4.2	Analysis: Proof of Lemma 12	24
5.5	Hardness	25
6	Applications in Online Bipartite Matching	26
6.1	Online Bipartite Matching Preliminaries	26
6.2	Semi-OCS and Unweighted and Vertex-weighted Online Bipartite Matching	27
6.3	OCS and Edge-weighted Online Bipartite Matching	28
6.3.1	Online Primal-Dual Algorithm	28
6.3.2	Improved Online Primal-Dual Analysis	28
6.4	A Variant of OCS and Edge-weighted Online Bipartite Matching	32
6.5	Multi-way Semi-OCS and Unweighted and Vertex-weighted Online Bipartite Matching	35
A	Missing Proofs in Section 3	41
A.1	Positive Correlation in 3-Way Sampling without Replacement	41
A.2	Proof of Theorem 4	41

B	Missing Proofs in Section 4	42
B.1	Proof of Lemma 7	42
B.2	Proof of Lemma 8	44
B.3	A Weaker Version of Lemma 7	45
C	Missing Proofs in Section 5	48
C.1	Proof of Lemma 15	48
C.2	Proof of Lemma 16	48
C.3	Proof of Lemma 17	48
C.4	Proof of Lemma 18	49
C.5	Proof of Lemma 19	49
C.6	Proof of Lemma 23	50
C.7	Proof of Lemma 24	51
C.8	Proof of Lemma 25	51
C.9	Proof of Lemma 26	52
C.10	Proof of Lemma 27	53
C.11	Proof of Lemma 28	53
C.12	Proof of Lemma 29	54
C.13	Proof of Lemma 30	56
D	Missing Proofs in Section 6	56
D.1	Proof of Theorem 33	56
D.2	Proof of Theorem 45	57

1 Introduction

Real-life optimization problems often need to make decisions based on the information at hand instead of the full picture in hindsight. Online advertising platforms show advertisements within milliseconds after receiving each user query. Ride hailing applications match riders and drivers without full knowledge of future ride requests. Cloud service providers assign computational tasks to physical servers not knowing what tasks the users may submit later. Due to the broad applications, the design and analysis of online algorithms for these optimization problems are a central topic in computer science and operations research.

Lacking accurate knowledge of the full picture, there is usually no universally good decision for all possible future input in these online optimization problems. As a result, online algorithms need to hedge against different possibilities through randomized decisions. Consider the online bipartite matching problem by Karp, Vazirani, and Vazirani [26] as a running example. We want to find a matching in a bipartite graph and maximize its size. Initially, we only know the left-hand-side of the bipartite graph, a.k.a., the offline vertices. Online vertices on the right-hand-side arrive one at a time. We must immediately and irrevocably match each of them upon arrival. Any deterministic greedy algorithm gives a maximal matching, and therefore its size is at least half of the maximum matching in hindsight. Beating this trivial bound of half, however, necessitates randomization even for bipartite graphs with only two vertices on each side.

Example. Consider vertices 1 and 2 on the left and 3 and 4 on the right. Vertex 3 arrives first with edges to both 1 and 2. A deterministic algorithm then immediately matches 3, e.g., to 1. If vertex 4 only has an edge to 1, however, the algorithm cannot match it even though a perfect matching exists in hindsight. A randomized algorithm that matches 3 to 1 and 2 with equal probability, on the other hand, matches $\frac{3}{2}$ edges in expectation.

Would it suffice to independently randomize over two offline neighbors? Unfortunately, the answer is negative (e.g., Fahrbach, Huang, Tao, and Zadimoghaddam [18]). We need to correlate different rounds' selections to break the $\frac{1}{2}$ barrier. One can introduce correlation through problem specific methods. The RANKING algorithm of Karp et al. [26], for example, samples a random order of the offline vertices at the beginning, and then matches each online vertex to the first unmatched offline neighbor by that order. RANKING and its variants achieve the optimal $1 - \frac{1}{e}$ competitive ratio for unweighted [26] and vertex-weighted online bipartite matching [2], but have difficulties extending to the more general edge-weighted problem (a.k.a., Display Ads) [19] and AdWords [28].

Fahrbach et al. [18], on the other hand, formulate a generic online selection problem and design online correlated selection (OCS) algorithms that lead to the first edge-weighted online bipartite matching algorithm that breaks the $\frac{1}{2}$ barrier. Subsequently, Huang, Zhang, and Zhang [24] break the $\frac{1}{2}$ barrier in AdWords using a similar approach. The online selection problem considers a set of ground elements (e.g., the offline vertices) and a sequence of pairs of these elements (e.g., a pair of offline neighbors for each online vertex). The algorithm immediately selects an element upon receiving each pair. If it independently selects a random element from each pair, then with probability 2^{-k} an element remains unselected in a subset of k pairs involving it. Can we be more effective than independent random selections?

Fahrbach et al. [18] study two versions of online selection called semi-OCS and OCS. Semi-OCS focuses on the probability that an element is unselected at the end when it is in k pairs. They give a semi-OCS that upper bounds this probability by $2^{-k}(1 - \gamma)^{k-1}$ for $\gamma = 0.109$, and call it a γ -semi-OCS. They also prove that 1-semi-OCS is impossible. OCS further considers the probability

that an element is unselected in *an arbitrary subset* of pairs involving it. When the subset is the union of m consecutive subsequences of the pairs involving the element, with lengths k_1, k_2, \dots, k_m , their OCS bounds the unselected probability by $\prod_{i=1}^m 2^{-k_i} (1 - \gamma)^{k_i - 1}$ also for $\gamma = 0.109$. They call it a γ -OCS. The weaker guarantee of semi-OCS is sufficient for unweighted and vertex-weighted online bipartite matching, while the stronger guarantee of OCS is sufficient for the edge-weighted problem. The main idea is to randomly match the pairs so that two matched pairs share a common element that is not in any pair in between. Their semi-OCS and OCS then select oppositely from the matched pairs with respect to the common element.

Fahrbach et al. [18] explicitly leave two open questions: (1) *What is the best possible γ for which γ -semi-OCS and γ -OCS exist?* (2) *Are there multi-way online selection algorithms that select from multiple elements in each round, with sufficient negative correlation such that the resulting online matching algorithms are better than the two-way counterparts?* We remark that the matching-based approach fails fundamentally in the multi-way extension. If each round has n elements, it can then be matched to $\Omega(n)$ other rounds. For large n , any matching is sparse and the resulting negative correlation is negligible.

1.1 Our Contributions

Semi-OCS and Weighted Sampling without Replacement. This paper gives a complete answer to the first open question for semi-OCS. In fact, we not only show that the optimal γ equals $\frac{1}{2}$ for semi-OCS, but also find that the unselected probability converges to zero much faster than the guarantee of γ -semi-OCS when the element is in $k \geq 3$ rounds.

Informal Theorem 1. *There is a polynomial-time semi-OCS such that an element that is in k pairs is selected with probability at least $1 - 2^{-2^k + 1}$. This is the best possible for all $k \geq 1$.*

In each round the optimal semi-OCS selects the element that appears more in previous rounds, and is unselected thus far, breaking ties randomly. It is the limit case of weighted sampling without replacement, when an element's weight is exponential in the number of previous rounds with the element, and when the base of the exponential tends to infinity. The main lemma in our analysis, which may be of independent interest, shows that the selections of different elements are negatively correlated in weighted sampling without replacement with two elements per round. See Section 3.

This paper further answers the second open question affirmatively for semi-OCS by studying a multi-way online selection problem that allows an arbitrary marginal distribution over the elements in each round. We will refer to the marginal probability of an element as its mass in that round. Consider the probability that an element is unselected at the end when its total mass is y . On the one hand, sampling independently from the marginal distributions (with replacement) bounds the probability by $\exp(-y)$. On the other hand, for an overly idealized algorithm, which samples from the marginal distributions and ensures that each element is sampled at most once, this probability is $\max\{1 - y, 0\}$. For $y \in [0, 1]$, it is $\exp(-y - \frac{y^2}{2} - \frac{y^3}{3} - \dots)$ by the Taylor series of $\ln(1 - y)$. We match the overly idealized bound up to the quadratic term, with a smaller cubic coefficient.

Informal Theorem 2. *There is a polynomial-time multi-way semi-OCS such that an element with total mass y is selected with probability at least $1 - \exp(-y - \frac{y^2}{2} - \frac{4-2\sqrt{3}}{3}y^3)$.*

It may be tempting to conjecture that sampling independently from the marginal distributions *without replacement* already improves the trivial bound of $\exp(-y)$. Unfortunately, this is false. Consider an element e that is in all T rounds each with mass $\epsilon = \frac{y}{T}$, and let there be a distinct

Table 1: A summary of the results in this paper on online correlated selection and their applications in online bipartite matching, with a comparison to those by Fahrbach et al. [18].

	Fahrbach et al. [18]	This Paper
Semi-OCS ¹	$2^{-k}(1 - 0.109)^{k-1}$	2^{-2^k+1}
Multi-way Semi-OCS ²	-	$\exp(-y - \frac{y^2}{2} - \frac{4-2\sqrt{3}}{3}y^3)$
γ -OCS	$0.109 \leq \gamma < 1$	$0.167 \leq \gamma \leq \frac{1}{4}$
Unweighted/Vertex-weighted (2-Way)	0.508	0.536
Unweighted/Vertex-weighted (Multi-way) ³	-	0.593
Edge-weighted	0.508	0.519

¹ The table presents upper bounds on the probability that an element is unselected when it is in k pairs.

² The table presents upper bounds on the probability that an element is unselected when its total mass is y .

³ RANKING by Karp et al. [26] is $1 - \frac{1}{e}$ -competitive which is optimal. Nonetheless, the algorithms in this paper are the first ones other than RANKING whose competitive ratios are beyond the $0.5 + \epsilon$ regime.

element other than e with mass $1 - \epsilon$ in each round. Then element e remains unselected at the end with probability $\exp(-y)$ when T tends to infinity.

This example suggests that if an element has accumulated some mass and is still unselected, we shall give it a higher priority than new elements that are not in any previous rounds. It motivates weighted sampling without replacement where the weight is a function of the total mass of the element in previous rounds. We choose the weight to be the inverse of the upper bound on the unselected probability so that *the expected sampling weight of any element is at most its mass in the round*, an invariant that is the key to our analysis. See Section 4.

OCS and Probabilistic Automata. This paper also contributes to the first open question for OCS by narrowing the gap between the upper and lower bounds on the best possible γ .

Informal Theorem 3. *There is a polynomial-time 0.167-OCS. Further, there is no γ -OCS for any $\gamma > \frac{1}{4}$, even with unlimited computational power.*

The improved OCS also abandons the matching-based approach and instead introduces an automata-based approach. Informally, it picks an element from each round to probe the element’s state. If the element was selected last time, select the other element this time. If the element has not been selected in the last two appearances, select it this time. Finally, only when the element was not selected last time but was selected before that, the OCS selects with fresh randomness. The actual algorithm is more involved. For example, we cannot pick an element to probe independently in each round in general. See Section 5 for detail.

Applications in Online Bipartite Matching. The new results on online correlated selection from this paper lead to better online bipartite matching algorithms. For the unweighted and vertex-weighted problems, we get a 0.536-competitive two-way algorithm, improving the 0.508-competitive algorithm by Fahrbach et al. [18]. We further show that the multi-way semi-OCS can round the (fractional) BALANCE algorithm (e.g., [25, 28]), and be 0.593-competitive. For the edge-weighted problem, the 0.167-OCS gives a 0.512-competitive algorithm. In the process, we refine the reductions from online matching problems to online correlated selection so that the competitive

ratios admit close-formed expressions, and a guarantee strictly weaker than γ -OCS already suffices for the edge-weighted problem. Motivated by the relaxed guarantee, we design a variant of OCS that further improves the edge-weighted competitive ratio. See Section 6.

Informal Theorem 4. *There is a polynomial-time 0.519-competitive algorithm for edge-weighted online bipartite matching.*¹

1.2 Other Related Works

Online Rounding. Online correlated selection is related to online rounding algorithms. It is common to first design online algorithms for an easier fractional online optimization problem, and to round it using online rounding algorithms to solve the original integral problem. Independent rounding is the simplest and most general online rounding; it corresponds to independent random selections in the online selection problem in this paper. For instance, Buchbinder and Naor [6] first design fractional online covering and packing algorithms under the online primal-dual framework, and then round them with independent rounding. More involved online rounding algorithms are usually designed on a problem-by-problem basis in the literature, e.g., for k -server [4], online submodular maximization [8], online edge coloring [10], etc.

To our knowledge, the only general online rounding method other than independent rounding is the online contention resolution schemes initiated by Feldman et al. [20] and further developed by Adamczyk and Włodarczyk [1], Lee and Singla [27], and Dughmi [14, 15]. It has found applications mainly in online problems with stochastic information, such as prophet inequality [17, 20], posted pricing [20], stochastic probing [20], and stochastic matching [17, 21, 22].

There is an important difference between the above usages of online rounding algorithms and the applications of OCS in online bipartite matching. The above online rounding algorithms and the corresponding fractional online algorithms are designed separately; the final competitive ratio is the product of their ratios.² By contrast, this paper and previous works on OCS [18, 24] take an end-to-end approach: the online matching algorithms make fractional decisions based on the guarantee of OCS to directly optimize the expected objective of the rounded matching. For example, the algorithms for vertex-weighted and edge-weighted matching in Section 6 rely on discount functions derived from optimization problems that take the OCS guarantees as parameters. Another example with a similar spirit is the convex rounding technique by Dughmi et al. [16] and Dughmi [13] from the algorithmic game theory literature.

Online Matching. We refer readers to Mehta et al. [29] for a survey on online matching problems. The unweighted, vertex-weighted, and edge-weighted online bipartite matching problems are first studied by Karp et al. [26], Aggarwal et al. [2], and Feldman et al. [19]. Later, Devanur et al. [11] and Devanur et al. [12] simplify the analyses under the online primal-dual framework. In particular, Devanur et al. [12] view the expected maximal edge-weight matched to an offline vertex as an integral of the complementary cumulative distribution function, a key ingredient of the application of OCS in edge-weighted matching. Finally, Buchbinder et al. [7], Cohen and Wajc [9], Gamlath et al. [23], Papadimitriou et al. [30], and Saberi and Wajc [31] also build on negative correlation properties to analyze their online algorithms, although these negative correlation properties and their usage are orthogonal to those in this paper.

¹We assume free disposals, which is standard in the edge-weighted problem under worst-case analysis.

²Using this two-step approach to analyze the applications of the multi-way semi-OCS in this paper in unweighted and vertex-weighted online bipartite matching would lead to a much worse competitive ratio of about 0.514. Using it with the (two-way) semi-OCS and OCS even gives a ratio strictly smaller than 0.5!

1.3 Concurrent Works

Concurrently and independently, Blanc and Charikar [5] and Shin and An [32] also improved the results of Fahrbach et al. [18]. They mainly study multi-way OCS, while this paper focuses on 2-way semi-OCS, 2-way OCS, and multi-way semi-OCS. Hence, these two papers are almost orthogonal to ours. Using a 6-way OCS, Blanc and Charikar [5] obtained a 0.5368-competitive algorithm for edge-weighted online bipartite matching. They also gave similar simplifications for the reduction of online matching problems to OCS. Shin and An [32] gave a method for converting 2-way OCS to 3-way OCS. Applying their method to the OCS of Fahrbach et al. [18] gives a 0.509-competitive algorithm for edge-weighted online bipartite matching. Applying it to the improved 2-way OCS in this paper further improves the ratio to 0.513.

2 Preliminaries

The *online selection problem* considers a set \mathcal{E} of elements and a selection process that proceeds in T rounds. For any round $1 \leq t \leq T$, a pair of elements \mathcal{E}^t arrive and the *online selection algorithm* needs to immediately select an element from \mathcal{E}^t . Let s^t denote the selected element in round t .

For any subset of rounds $T' \subseteq T$, we say that an element e is *unselected* in T' if the algorithm does not select the element in any round in T' , i.e., if $s^t \neq e$ for any $t \in T'$. If $T' = T$, we simply say that element e is unselected. For any $0 \leq t \leq T$, let \mathcal{U}^t denote the set of elements that are unselected in rounds $1, 2, \dots, t$. Let $\mathcal{U} = \mathcal{U}^T$ denote the set of unselected elements at the end.

Semi-OCS considers the probability that an element $e \in \mathcal{E}$ is unselected at the end, and seeks to bound it as a function of the number rounds containing element e .

Definition 1 (γ -semi-OCS, c.f., Fahrbach et al. [18]). *An online selection algorithm is a γ -semi-OCS if for any online selection instance and any element e that appears in k rounds, element e is unselected with probability at most:*

$$2^{-k}(1 - \gamma)^{k-1} .$$

Selecting an element in each round independently and uniformly at random is a 0-semi-OCS. Fahrbach et al. [18] give a 0.109-semi-OCS and prove that there is no 1-semi-OCS.

OCS further considers the probability that an element $e \in \mathcal{E}$ is unselected in an *arbitrary subset of rounds* containing the element. The upper bounds on this probability depend on the structure of the subset of rounds. A consecutive subsequence of the rounds containing element e is a subset of rounds $\{t_1, t_2, \dots, t_k\}$ such that each round t_i contains e , i.e., $e \in \mathcal{E}^{t_i}$ for any $1 \leq i \leq k$, and no round in between contains e , i.e., $e \notin \mathcal{E}^t$ for any $1 \leq i \leq k - 1$ and any $t_i < t < t_{i+1}$.

Definition 2 (γ -OCS, c.f., Fahrbach et al. [18]). *An online selection algorithm is a γ -OCS if for any online selection instance, any element e , and any subset of rounds $T' \subseteq T$ containing e such that T' is the union of m consecutive subsequences of the rounds containing e , with lengths k_1, k_2, \dots, k_m , element e is unselected in T' with probability at most:*

$$\prod_{i=1}^m 2^{-k_i}(1 - \gamma)^{k_i-1} .$$

For example, suppose that rounds 1, 2, 5, 6, 9 are the ones that contain element e , and consider $T' = \{1, 2, 6, 9\}$. Then, T' is the union of two consecutive subsequences 1, 2 and 6, 9, whose lengths are 2. Fahrbach et al. [18] give a 0.109-OCS. Since OCS is stronger than semi-OCS, the impossibility result for semi-OCS implies that there is no 1-OCS.

3 Optimal Semi-OCS

3.1 Algorithms

This paper considers a semi-OCS that remembers the number of rounds involving each element thus far, and selects from each round the element that appears more and is unselected so far, breaking ties uniformly at random and independently in different rounds. See Algorithm 1.

Algorithm 1 Optimal Semi-OCS

State variables: (for each element e)

- The number of previous rounds that contain element e , denoted as k_e .
- Whether element e has been selected in any previous rounds.

For each round t : (suppose $\mathcal{E}^t = \{e, e'\}$)

1. If both e and e' have been selected, select arbitrarily, e.g., $s^t \in \{e, e'\}$ uniformly at random.
 2. If only one of e and e' has been selected, select s^t to be the one that has not been selected.
 3. If neither e nor e' has been selected:
 - If $k_e \neq k_{e'}$, select s^t to be the one with more previous appearances.
 - Otherwise, select $s^t \in \{e, e'\}$ uniformly at random.
-

The main lemma in the analysis of Algorithm 1 will prove that the (un)selections of elements are negatively correlated. It is more instructive to prove this lemma for a broader family of weighted sampling algorithms (Algorithm 2). Algorithm 1 is the special case when we let $w_e^t = 1$ if e appears in previous rounds at least as many times as the other element does, and let $w_e^t = 0$ otherwise.

Algorithm 2 Weighted 2-Way Sampling without Replacement

Parameters:

- $w_e^t \geq 0$, weight of element $e \in \mathcal{E}^t$ in round t

For each round t :

1. If both elements in \mathcal{E}^t have been selected, select arbitrarily, e.g., uniformly at random.
 2. Otherwise, select each unselected element $e \in \mathcal{E}^t$ with probability proportional to w_e^t .
-

3.2 Negative Correlation in Weighted 2-Way Sampling without Replacement

Recall that \mathcal{U}^t denotes the set of unselected elements after the first t rounds. Hence the event that a subset $S \subseteq \mathcal{E}$ of elements are all unselected after the first t rounds can be written as $S \subseteq \mathcal{U}^t$. We shall establish the negative correlation of such events in the next lemma.

Lemma 1. *For weighted 2-way sampling (Algorithm 2) with any weights, any $0 \leq t \leq T$, and any disjoint subsets of elements $A, B \subseteq \mathcal{E}$:*

$$\Pr[A \cup B \subseteq \mathcal{U}^t] \leq \Pr[A \subseteq \mathcal{U}^t] \Pr[B \subseteq \mathcal{U}^t] .$$

Proof. We shall prove the lemma by induction on t . The base case when $t = 0$ is trivial since $\mathcal{U}^0 = \mathcal{E}$, and thus $\Pr[A \subseteq \mathcal{U}^0] = \Pr[B \subseteq \mathcal{U}^0] = \Pr[A \cup B \subseteq \mathcal{U}^0] = 1$. Next suppose that the lemma holds for round $t - 1$ and consider round t .

Case 1: $\mathcal{E}^t \cap (A \cup B) = \emptyset$, i.e., no element in this round belongs to A or B . Since the selection in round t does not affect the events of concern, the lemma continues to hold after round t by the inductive hypothesis.

Case 2: $|\mathcal{E}^t \cap (A \cup B)| = 1$, i.e., exactly one element in this round belongs to A or B . Denote this element as $e \in \mathcal{E}^t$ and the other element as $e' \in \mathcal{E}^t$. Further suppose without loss of generality that $e \in A$. Since the elements in B are not involved in round t , we have:

$$\Pr[B \subseteq \mathcal{U}^t] = \Pr[B \subseteq \mathcal{U}^{t-1}] .$$

Next consider the elements in A . If e' has been selected in the first $t - 1$ rounds, e would certainly be selected after round t . Hence, to have $A \subseteq \mathcal{U}^t$, we need not only $A \subseteq \mathcal{U}^{t-1}$, but also $e' \in \mathcal{U}^{t-1}$. Further the algorithm must select e' in round t . Putting together we have:

$$\Pr[A \subseteq \mathcal{U}^t] = \Pr[A \cup \{e'\} \subseteq \mathcal{U}^{t-1}] \frac{w_{e'}^t}{w_e^t + w_{e'}^t} .$$

Similarly we have:

$$\Pr[A \cup B \subseteq \mathcal{U}^t] = \Pr[A \cup B \cup \{e'\} \subseteq \mathcal{U}^{t-1}] \frac{w_{e'}^t}{w_e^t + w_{e'}^t} .$$

Cancelling the common term $\frac{w_{e'}^t}{w_e^t + w_{e'}^t}$, the inequality in the lemma is equivalent to:

$$\Pr[A \cup B \cup \{e'\} \subseteq \mathcal{U}^{t-1}] \leq \Pr[A \cup \{e'\} \subseteq \mathcal{U}^{t-1}] \Pr[B \subseteq \mathcal{U}^{t-1}] .$$

This follows by the inductive hypothesis for subsets $A \cup \{e'\}$ and B in round $t - 1$.

Case 3: $\mathcal{E}^t \subseteq (A \cup B)$, i.e., both elements in round t belong to A or B . Since one element in \mathcal{E}^t is selected in round t , we have $\Pr[A \cup B \subseteq \mathcal{U}^t] = 0$. Hence the stated inequality trivially holds. \square

We remark that the lemma no longer holds if we have 3 or more elements in each round. Appendix A.1 provides a counter-example. See also Alexander [3].

3.3 Analysis

Theorem 2. *For any instance and any element that appears in k rounds, the probability that the element is never selected by Algorithm 1 is at most:*

$$2^{-2^k + 1} .$$

Proof. We shall prove the theorem by induction on the number of rounds T in the instance. The base case when $T = 0$ is trivial since k must be 0 in this case. Next suppose that the lemma holds for up to $T - 1$ rounds. Consider an arbitrary instance with T rounds, and any element e that appears k times. Without loss of generality, we may assume that e is in the last round T ; otherwise it follows directly from the inductive hypothesis. Further suppose that the other element in round T is e' . Consider three cases depending on the relation between the number of appearances k_e and $k_{e'}$ before round T . Observe that $k = k_e + 1$.

Case 1: $k_e > k_{e'}$. By the definition of Algorithm 1, element e is selected with certainty after T . Hence the probability of concern is 0, and is trivially smaller than the stated bound.

Case 2: $k_e < k_{e'}$. By the definition of Algorithm 1, element e' is selected with certainty in round T if it is not yet selected previously. Hence, element e is never selected by the algorithm at the end if and only if both e and e' are unselected before round T . This probability is:

$$\begin{aligned}
\Pr[\{e, e'\} \subseteq \mathcal{U}^{T-1}] &\leq \Pr[e \in \mathcal{U}^{T-1}] \Pr[e' \in \mathcal{U}^{T-1}] && \text{(Lemma 1)} \\
&\leq 2^{-2^{k_e}+1} \cdot 2^{-2^{k_{e'}}+1} && \text{(Inductive hypothesis)} \\
&\leq 2^{-2^{k_e}+1} \cdot 2^{-2^{k_e+1}+1} && (k_e < k_{e'}) \\
&\leq 2^{-2^k+1} . && (k = k_e + 1)
\end{aligned}$$

Case 3: $k_e = k_{e'}$. By the definition of Algorithm 1, elements e and e' would be selected with equal probability if neither has been selected before. Therefore, element e is never selected by the algorithm at the end if and only if both e and e' are unselected before round T , and the algorithm selects e' in round T . The latter happens with probability half and is independent with the former. Hence, this probability is:

$$\begin{aligned}
2^{-1} \Pr[\{e, e'\} \subseteq \mathcal{U}^{T-1}] &\leq 2^{-1} \Pr[e \in \mathcal{U}^{T-1}] \Pr[e' \in \mathcal{U}^{T-1}] && \text{(Lemma 1)} \\
&\leq 2^{-1} \cdot 2^{-2^{k_e}+1} \cdot 2^{-2^{k_{e'}}+1} && \text{(Inductive hypothesis)} \\
&\leq 2^{-1} \cdot 2^{-2^{k_e}+1} \cdot 2^{-2^{k_e}+1} && (k_e = k_{e'}) \\
&= 2^{-2^k+1} . && (k = k_e + 1)
\end{aligned}$$

Summarizing the three cases completes the inductive step and thus the proof of the theorem. \square

Since $2^{-2^k+1} \leq 2^{-2^{k+1}} = 2^{-k}(1 - \frac{1}{2})^{k-1}$, Theorem 2 leads to the following corollary in terms of the original definition of semi-OCS.

Corollary 3. *Algorithm 1 is a $\frac{1}{2}$ -semi-OCS.*

We remark that the guarantee of $\frac{1}{2}$ -semi-OCS only gives an $\frac{8}{15} \approx 0.533$ -competitive algorithm for unweighted and vertex-weighted matching, while Theorem 2 leads to least 0.536. That is, the tighter analysis in Theorem 2 indeed results in better competitive ratios in online matching.

3.4 Hardness

Finally, we show that the semi-OCS (Algorithm 1) and its analysis (Theorem 2) are optimal for all k simultaneously. The proof is deferred to Appendix A.2.

Theorem 4. *For any algorithm and any $k \geq 0$, there is an instance and an element that appears in k rounds, such that with probability at least 2^{-2^k+1} the algorithm never selects the element.*

The special case of $k = 2$ further implies a hardness for the original definition of γ -semi-OCS.

Corollary 5. *There is no γ -semi-OCS for $\gamma > \frac{1}{2}$.*

4 Multi-way Semi-OCS

4.1 Definitions

The *multi-way online selection problem* considers a set of elements \mathcal{E} and a selection process that proceeds in T rounds as follows. Each round $1 \leq t \leq T$ is associated with a non-negative vector

Algorithm 3 Multi-way Semi-OCS: Weighted Sampling without Replacement

Parameters:

Non-decreasing weight function $w : [0, +\infty) \rightarrow [1, +\infty)$ such that $w(0) = 1$.

Our result lets $w(y) = \exp\left(y + \frac{y^2}{2} + cy^3\right)$ where $c = \frac{4-2\sqrt{3}}{3}$.

State variables: (for each element e)

- Cumulative mass y_e^t of element e up to any round t .
- Whether element e has been selected in any previous rounds.

For each round t :

1. If all elements in \mathcal{E}^t have been selected, select arbitrarily, e.g., uniformly at random.
 2. Otherwise, select an unselected $e \in \mathcal{E}^t$ with probability proportional to $x_e^t \cdot w(y_e^{t-1})$.
-

$\mathbf{x}^t = (x_e^t)_{e \in \mathcal{E}}$ such that $\sum_{e \in \mathcal{E}} x_e^t = 1$. We shall refer to x_e^t as the *mass* of element e in round t . The vectors are unknown at the beginning and are revealed to an *multi-way online selection algorithm* at the corresponding rounds. Let $\mathcal{E}^t = \{e : x_e^t > 0\}$ be the set of elements with positive masses in round t , i.e., those that may be selected in the round. Upon observing the mass vector \mathbf{x}^t for round t , the algorithm selects an element from \mathcal{E}^t .

We may interpret x_e^t as the probability of selecting element e in the round *if none of the elements have appeared in previous rounds*, although in general the correlation introduced by the multi-way online selection algorithms will complicate the selection probabilities. For any $0 \leq t \leq T$, let $y_e^t = \sum_{t' \leq t} x_e^{t'}$ be the cumulative mass of element e in the first t rounds. Let $y_e = y_e^T$ be its total mass in the instance for brevity.

Definition 3 (*p*-Multi-way Semi-OCS). *A multi-way online selection algorithm is a p-multi-way semi-OCS for a non-increasing function $p : [0, +\infty) \rightarrow [0, 1]$ if for any multi-way online selection instance and any element e , e is unselected with probability at most $p(y_e)$.*

4.2 Algorithm: Weighted Sampling without Replacement

We consider weighted sampling without replacement, which is parameterized by a weight function $w : [0, +\infty) \rightarrow [1, +\infty)$ with $w(0) = 1$. In each round t , the sampling weight of an element $e \in \mathcal{E}^t$ equals 0 if the element has already been selected in the previous rounds, and equals $x_e^t w(y_e^{t-1})$ otherwise. See Algorithm 3.

We remark that the optimal (2-way) semi-OCS in Section 3 can be interpreted as the limit case when $w(y) = W^y$ and W tends to infinity.

4.3 Analysis

Theorem 6. *Weighted Sampling without Replacement (Algorithm 3) with weight function:*

$$w(y) = \exp\left(y + \frac{y^2}{2} + cy^3\right) \tag{1}$$

where $c = \frac{4-2\sqrt{3}}{3} \approx 0.179$ is a *p*-multi-way semi-OCS for:

$$p(y) = \frac{1}{w(y)} = \exp\left(-y - \frac{y^2}{2} - cy^3\right).$$

Consider an overly idealized algorithm which selects each element e in round t with probability exactly x_e^t and never selects any element more than once. It would be a p^* -multiway semi-OCS for $p^*(y) = \max\{1 - y, 0\}$. By the Taylor series of $\log(1 - y)$ for $0 \leq y < 1$, it can be written as:

$$p^*(y) = \exp\left(-\sum_{i=1}^{\infty} \frac{y^i}{i}\right).$$

The guarantee of Theorem 6 matches the overly idealized bound up to the quadratic term and has a smaller coefficient for the cubic term.

First, we prove some properties about the weight function w in Eqn. (1).

Lemma 7. *For any $0 \leq x < 1$ and any $y \geq 0$:*

$$\frac{w(y+x)}{w(y)} \leq \frac{x}{1-x}w(y) + 1.$$

The proof of Lemma 7 involves tedious calculations and computer-aided numerical verifications that are not insightful. Hence, we defer it to Appendix B.1; see also Appendix B.3 for a proof that does not use computer-aided numerical verifications for a weaker version of the lemma. We further introduce a generalized version of Lemma 7 whose proof is also deferred to Appendix B.2.

Lemma 8. *For any $k \geq 1$, any $x_i, y_i \geq 0$ for $1 \leq i \leq k$ such that $\sum_{i=1}^k x_i \in [0, 1]$:*

$$\frac{1 - \sum_{i=1}^k x_i}{\sum_{i=1}^k x_i w(y_i) + 1 - \sum_{i=1}^k x_i} \leq \prod_{i=1}^k \frac{w(y_i)}{w(y_i + x_i)}.$$

With these two lemmas, we bound the unselected probability for any subset of elements, which implies Theorem 6 as a special case.

Theorem 9. *Weighted Sampling without Replacement (Algorithm 3) with weight function*

$$w(y) = \exp\left(y + \frac{y^2}{2} + cy^3\right)$$

with $c = \frac{4-2\sqrt{3}}{3}$ ensures that any subset of elements $\mathcal{E}' \subseteq \mathcal{E}$ are unselected with probability at most:

$$\prod_{e \in \mathcal{E}'} p(y_e),$$

where $p(y) = \frac{1}{w(y)}$.

Proof. Recall that \mathcal{U}^t denotes the set of unselected elements after round t . Hence, $\mathcal{E}' \subseteq \mathcal{U}^t$ is the event that the elements in \mathcal{E}' are unselected in the first t rounds. We shall prove by induction on $0 \leq t \leq T$ that:

$$\Pr[\mathcal{E}' \subseteq \mathcal{U}^t] \leq \prod_{e \in \mathcal{E}'} p(y_e^t),$$

which implies Theorem 9 as a special case when $t = T$.

The base case when $t = 0$ holds vacuously because both sides of the inequality equal 1.

Next suppose that it holds for $t-1$ rounds for some $t > 0$, and consider the case of t rounds. Let \bar{X}_e^t be the indicator of whether element e is unselected after round t , and define $\bar{X}_{\mathcal{E}'}^t = \prod_{e \in \mathcal{E}'} \bar{X}_e^t$ for any $\mathcal{E}' \subseteq \mathcal{E}$. Finally, we write \bar{X}^t for $(\bar{X}_e^t)_{e \in \mathcal{E}}$.

$$\begin{aligned} \Pr[\mathcal{E}' \subseteq \mathcal{U}^t] &= \mathbf{E} \bar{X}_{\mathcal{E}'}^t \\ &= \mathbf{E}_{\bar{X}^{t-1}} \left[\bar{X}_{\mathcal{E}'}^{t-1} \left(1 - \frac{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1}}{\sum_{e \in \mathcal{E}} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1}} \right) \right]. \end{aligned}$$

Here we artificially define $\frac{0}{0} = 0$ for ease of presentation. Readers may verify that our argument stays true with this caveat.

Next, multiply $\bar{X}_{\mathcal{E}'}^{t-1}$ with both the numerator and denominator in the above fraction. Using that $Y^2 = Y$ for $Y \in \{0, 1\}$, we have:

$$\begin{aligned} \Pr[\mathcal{E}' \subseteq \mathcal{U}^t] &= \mathbf{E}_{\bar{X}^{t-1}} \left[\bar{X}_{\mathcal{E}'}^{t-1} \left(1 - \frac{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1}}{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1} + \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1}} \right) \right] \\ &= \mathbf{E}_{\bar{X}^{t-1}} \left[\frac{\bar{X}_{\mathcal{E}'}^{t-1} \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1}}{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1} + \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \bar{X}_e^{t-1}} \right]. \end{aligned}$$

By the concavity of $f(x, y) = \frac{xy}{x+y}$, it follows from Jensen's inequality that:

$$\Pr[\mathcal{E}' \subseteq \mathcal{U}^t] \leq \frac{\mathbf{E} \bar{X}_{\mathcal{E}'}^{t-1} \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \mathbf{E} \bar{X}_{\mathcal{E}' \cup \{e\}}^{t-1}}{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t \mathbf{E} \bar{X}_{\mathcal{E}'}^{t-1} + \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \mathbf{E} \bar{X}_{\mathcal{E}' \cup \{e\}}^{t-1}}.$$

By the inductive hypothesis and the monotonicity of $f(x, y) = \frac{xy}{x+y}$, we further get that:

$$\begin{aligned} \Pr[\mathcal{E}' \subseteq \mathcal{U}^t] &\leq \frac{\prod_{e \in \mathcal{E}'} p(y_e^{t-1}) \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \prod_{e' \in \mathcal{E}' \cup \{e\}} p(y_{e'}^{t-1})}{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t \prod_{e \in \mathcal{E}'} p(y_e^{t-1}) + \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t \prod_{e' \in \mathcal{E}' \cup \{e\}} p(y_{e'}^{t-1})} \\ &= \prod_{e \in \mathcal{E}'} p(y_e^{t-1}) \frac{\sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t p(y_e^{t-1})}{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t + \sum_{e \notin \mathcal{E}'} w(y_e^{t-1}) x_e^t p(y_e^{t-1})} \\ &= \prod_{e \in \mathcal{E}'} p(y_e^{t-1}) \frac{\sum_{e \notin \mathcal{E}'} x_e^t}{\sum_{e \in \mathcal{E}'} w(y_e^{t-1}) x_e^t + \sum_{e \notin \mathcal{E}'} x_e^t}. \end{aligned}$$

Next combine the above with Lemmas 7 and 8:

$$\begin{aligned} \Pr[\mathcal{E}' \subseteq \mathcal{U}^t] &\leq \frac{1 - \sum_{e \in \mathcal{E}'} x_e^t}{\sum_{e \in \mathcal{E}'} x_e^t w(y_e^{t-1}) + 1 - \sum_{e \in \mathcal{E}'} x_e^t} \prod_{e \in \mathcal{E}'} p(y_e^{t-1}) && (\sum_{e \in \mathcal{E}} x_e^t = 1) \\ &\leq \prod_{e \in \mathcal{E}'} \frac{w(y_e^{t-1})}{w(y_e^{t-1}) + x_e^t} \prod_{e \in \mathcal{E}'} p(y_e^{t-1}) && (\text{Lemmas 7 and 8}) \\ &= \prod_{e \in \mathcal{E}'} p(y_e^{t-1} + x_e^t). && (p(y) = \frac{1}{w(y)}) \end{aligned}$$

□




Figure 1: Example of ex-ante dependence graph

5 Improved Algorithms and Hardness for OCS

5.1 Definitions

Recall that an online selection algorithm is a γ -OCS, if for any (2-way) online selection instance, any element e , and any disjoint consecutive subsequences of the rounds involving e with lengths k_1, k_2, \dots, k_m respectively, the probability that e is unselected in these rounds is at most:

$$\prod_{\ell=1}^m 2^{-k_\ell} (1 - \gamma)^{k_\ell - 1}.$$

Definition 4 (Ex-ante Dependence Graph, c.f., Fahrbach et al. [18]). *The ex-ante dependence graph $G^{\text{ex-ante}} = (V, E^{\text{ex-ante}})$ is a directed graph defined with respect to an online selection instance. We shall refer to its vertices and edges as nodes and arcs to make a distinction with those in online matching problems. The nodes correspond to rounds:*

$$V = \{1, 2, \dots, T\}.$$

The arcs correspond to neighboring appearances of an element (indicated by the subscript³):

$$E^{\text{ex-ante}} = \left\{ (t, t')_e : t < t'; e \in \mathcal{E}^t; e \in \mathcal{E}^{t'}; \forall t < t'' < t', e \notin \mathcal{E}^{t''} \right\}.$$

See Figure 1 for an illustrative example of the ex-ante dependence graph.

5.2 Roadmap

5.2.1 Matching-based Approach versus Automata-based Approach on a Path

This subsection reviews the matching-based approach of Fahrbach et al. [18] and its limitation, and explains the automata-based approach in this paper. As a running example, consider an instance with the same two elements *head* (H) and *tail* (T) in every round, and thus the ex-ante dependence graph is a directed path (more precisely, two identical parallel directed paths).

Matching-based Approach. Fahrbach et al. [18] propose to select a matching from the ex-ante graph, and then to select elements in each pair of matched nodes and in each isolated node with independent random bits; each pair of matched nodes shall select the opposite elements. We shall select the matching such that 1) the selections of different arcs are negatively dependent (including independent), and 2) the probability of selecting each arc is as high as possible. If we could select the arcs each with probability at least β with negative dependence, we would obtain a β -OCS because

³There could be parallel arcs in the ex-ante dependence graph, e.g., when rounds t and $t' = t + 1$ have the same two elements. The subscript helps distinguish such parallel arcs.

of the following argument. For any disjoint consecutive subsequences of lengths k_1, k_2, \dots, k_m , they contain $\sum_{i=1}^m (k_i - 1)$ arcs that could have been selected into the matching. If we select at least one of them into the matching, the opposite selections in its two nodes ensure selecting both elements. By the aforementioned properties, the matching has none of these arcs with probability at most $(1 - \beta)^{\sum_{i=1}^m (k_i - 1)}$. Even in that case, we still have $\sum_{i=1}^m k_i$ independent selections in these rounds; the probability of not selecting a given element in them is at most $2^{-\sum_{i=1}^m k_i}$.

For example, Fahrbach et al. [18] let each node independently pick an incident arc, and then select an arc into the matching if both nodes pick it. This selects each arc with probability $\frac{1}{4}$ in the special case when the ex-ante graph is a directed path. It is possible to improve in the special case. For instance, we could let each arc independently sample a number uniformly from $[0, 1]$ and select an arc if its number is bigger than its neighbors'. This selects each arc with probability $\frac{1}{3}$. To our best effort, however, we cannot find any matching-based algorithm that selects each arc with probability more than $\frac{\sqrt{5}-1}{2} \approx 0.382$. Further, some of these ideas that improve the $\frac{1}{4}$ bound by the algorithm of Fahrbach et al. [18] fail to generalize beyond the special case.

Automata-based Approach. This paper introduces a different approach that selects elements using a probabilistic automaton. We shall refer to both this automaton and its transition function as σ^* . It has five states q_O, q_H, q_{H^2}, q_T , and q_{T^2} . The original state q_O is both the initial state of the automaton and the state it resets to after selecting the same element in two consecutive rounds. State q_H (resp., q_T) means that the automaton selects H (resp., T) in the previous round but not twice in a roll; from this state the automaton selects T (resp., H) with a higher chance, and the margin $\beta \in [0, 1]$ will be optimized to be $\beta = \sqrt{2} - 1$ in our analysis. State q_{H^2} (resp., q_{T^2}) means that the automaton selects H (resp., T) in the last two rounds; from this state the automaton will select T (resp., H) with certainty and resets to the original state O. Below is the transition function σ^* that takes a state as input and returns a state and an element from $\{H, T\}$ (see also Figure 2):

$$\sigma^*(q_O) = \begin{cases} (q_H, H) & \text{w.p. } \frac{1}{2} \\ (q_T, T) & \text{w.p. } \frac{1}{2} \end{cases}, \quad \sigma^*(q_H) = \begin{cases} (q_{H^2}, H) & \text{w.p. } \frac{1-\beta}{2} \\ (q_T, T) & \text{w.p. } \frac{1+\beta}{2} \end{cases}, \quad \sigma^*(q_T) = \begin{cases} (q_H, H) & \text{w.p. } \frac{1+\beta}{2} \\ (q_{T^2}, T) & \text{w.p. } \frac{1-\beta}{2} \end{cases},$$

$$\sigma^*(q_{H^2}) = (q_O, T), \quad \sigma^*(q_{T^2}) = (q_O, H).$$

We find that using this automaton to select elements in different rounds is a $(\sqrt{2} - 1)$ -OCS in the special case. Readers will find the proof of a stronger claim in Subsection 5.4. This is strictly better than our best effort using the matching-based approach. More importantly, it generalizes to arbitrary online selection instances using the techniques in the rest of the section.

5.2.2 Automata-based Approach

This subsection outlines how to generalize the automata-based approach to general online selection instances and obtain an improvement over the 0.109-OCS of Fahrbach et al. [18].

Theorem 10. *There is a polynomial-time 0.167-OCS for the 2-way online selection problem.*

We next explain the ingredients and how to combine them to prove Theorem 10. The sequel subsections will substantiate them, with the proofs of some lemmas deferred to Appendix C.

A main challenge in generalizing the automata-based approach to general instances is deciding from which in-neighbor each node shall inherit the state of the automata. In other words, we need




Figure 2: The probabilistic automaton σ^* that selects an element in each round in the special case whose ex-ante dependence graph is a directed path, and in our β -tree OCS. The transitions are labeled by the selections from $\{H, T\}$, and by the probabilities of transitions.

to select an in-arc for each node to form a directed binary forest.⁴ We find that the naïve approach of independently and randomly selecting an in-arc for each node does not work unless the instance satisfies additional properties (see Subsection 5.3.1), because we need the directed binary forest to satisfy another property defined below.

Definition 5 (Good Forest). A good forest $G^{\text{forest}} = (V, E^{\text{forest}})$ with respect to an online selection instance is a subgraph of the ex-ante dependence graph $G^{\text{ex-ante}} = (V, E^{\text{ex-ante}})$ such that:

1. $G^{\text{forest}} = (V, E^{\text{forest}})$ is a directed binary forest;
2. For any node p with two children c and c' in G^{forest} , the corresponding rounds have no common element, i.e., $\mathcal{E}^p \cap \mathcal{E}^c \cap \mathcal{E}^{c'} = \emptyset$.

In the following definitions, for any subset of nodes $U \subseteq V$, let E_U^{forest} denote the subset of arcs induced by U in the forest G^{forest} . Further for any element e and any subset of nodes $U \subseteq V$ involving element e , let $E_{U,e}^{\text{ex-ante}}$ denote the subset of arcs induced by U and with subscript e :

$$E_{U,e}^{\text{ex-ante}} \stackrel{\text{def}}{=} \left\{ (t, t')_e \in E^{\text{ex-ante}} : t \in U; t' \in U \right\}.$$

Definition 6 (Forest Constructor). A forest constructor takes an online selection instance as input and returns good forest $G^{\text{forest}} = (V, E^{\text{forest}})$. On receiving the elements \mathcal{E}^t of round t , it immediately decides whether each in-arc of t belongs to E^{forest} . It is an α -forest constructor if for any element e , any subset of nodes $U \subseteq V$ involving e , and any $\beta \in [0, 1]$:

$$\mathbf{E}(1 - \beta)^{|E_U^{\text{forest}}|} \leq (1 - \alpha\beta)^{|E_{U,e}^{\text{ex-ante}}|}. \quad (2)$$

The expectation is over the randomness of the forest constructor.

The next lemma is our main result regarding forest constructors. Subsection 5.3 presents the algorithm that proves this lemma.

⁴That is, each node has at most one in-arc from its parent, and at most two out-arcs to its children. The latter is true for the ex-ante dependence graph itself, and therefore also for all its subgraphs.

Algorithm 4 $\frac{1}{2}$ -Forest constructor for good online selection instances

For each round t : (suppose that $\mathcal{E}^t = \{e_1, e_2\}$)

1. For $i \in \{1, 2\}$, let t_i be the most recent round that involves e_i (if exists).
 2. Draw $j \in \{1, 2\}$ uniformly at random, and include arc $(t_j, t)_{e_j}$ into E^{forest} (if t_j is defined).
-

Lemma 11. *There is a polynomial-time 0.404-forest constructor.*

Definition 7 (Forest OCS). *A forest OCS takes both an online selection instance and a good forest $G^{\text{forest}} = (V, E^{\text{forest}})$ as input. At each round t , it observes the elements \mathcal{E}^t in the round and whether each in-arc of t is in E^{forest} , and then selects an element from \mathcal{E}^t . It is a β -forest OCS if for any element e and any subset of nodes $U \subseteq V$ involving e , the probability that e is never selected in the corresponding rounds is at most:*

$$2^{-|U|} (1 - \beta)^{|E_U^{\text{forest}}|}.$$

Our main result regarding forest OCS is the next lemma, whose proof is in Subsection 5.4.

Lemma 12. *There is a polynomial-time $(\sqrt{2} - 1)$ -forest OCS.*

The next lemma combines the two ingredients to get an OCS, and implies Theorem 10 as a corollary using Lemmas 11 and 12.

Lemma 13. *Suppose that there is a polynomial-time α -forest constructor and a polynomial-time β -forest OCS. Together they form a polynomial-time $\alpha\beta$ -OCS.*

Proof. The OCS combines the α -forest constructor and the β -forest OCS as follows. On receiving the elements \mathcal{E}^t of a round t , it calls the forest constructor to determine whether each in-arc of t is in E^{forest} . Then, it puts this information together with the elements \mathcal{E}^t and calls the forest OCS to select an element from \mathcal{E}^t .

For any element e , and any disjoint consecutive subsequences of the rounds involving e , let k_1, k_2, \dots, k_m be the lengths of these subsequences, and let U be the subset of nodes that correspond to these rounds. By the guarantee of the α -forest constructor and the β -forest OCS, the probability that element e is never selected in these rounds is at most:

$$\mathbf{E} \left[2^{-|U|} (1 - \beta)^{|E_U^{\text{forest}}|} \right] \leq 2^{-|U|} (1 - \alpha\beta)^{|E_{U,e}^{\text{ex-ante}}|}.$$

The lemma then follows by $|U| = \sum_{\ell=1}^m k_\ell$ and $|E_{U,e}^{\text{ex-ante}}| = \sum_{\ell=1}^m (k_\ell - 1)$. \square

5.3 Forest Constructor

5.3.1 Warm-up: Good Online Selection Instances

We say that an online selection instance is good if its ex-ante graph satisfies the second requirement of good forests. That is, for any node p with two out-neighbors c and c' in $G^{\text{ex-ante}}$, the corresponding rounds have no common element, i.e., $\mathcal{E}^p \cap \mathcal{E}^c \cap \mathcal{E}^{c'} = \emptyset$. Such good instances admit a simple forest constructor that for each node keeps one of its in-arcs independently and uniformly at random (Algorithm 4). The simple forest constructor and its analysis are instructive, and motivate the forest constructor for general instances, so we include them as a warm-up.

Lemma 14. *Algorithm 4 is a $\frac{1}{2}$ -forest constructor for good online selection instances.*

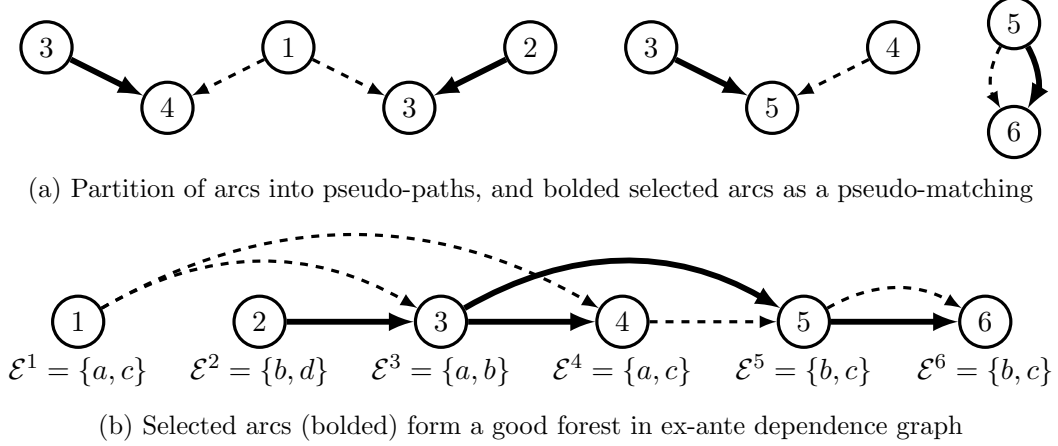


Figure 3: An illustrative example of the forest constructor for general instances

Proof. For any arc $a \in E_{U,e}^{\text{ex-ante}}$, let $X_a \in \{0, 1\}$ be the indicator of whether arc a is included into E^{forest} . Since every arc $a \in E_{U,e}^{\text{ex-ante}}$ with $X_a = 1$ belongs to E_U^{forest} , we have:

$$(1 - \beta)^{|E_U^{\text{forest}}|} \leq (1 - \beta)^{\sum_{a \in E_{U,e}^{\text{ex-ante}}} X_a}.$$

Further, since $E_{U,e}^{\text{ex-ante}}$ by definition is a subset of arcs connecting neighboring appearances of element e , they have distinct in- and out-nodes. Hence, by the definition of Algorithm 4, X_a 's are independently and uniformly distributed over $\{0, 1\}$ for all arcs $a \in E_{U,e}^{\text{ex-ante}}$. We get that

$$\mathbf{E}(1 - \beta)^{|E_U^{\text{forest}}|} \leq \prod_{a \in E_{U,e}^{\text{ex-ante}}} \mathbf{E}(1 - \beta)^{X_a} = \prod_{a \in E_{U,e}^{\text{ex-ante}}} \left(1 - \frac{\beta}{2}\right) = \left(1 - \frac{\beta}{2}\right)^{|E_{U,e}^{\text{ex-ante}}|}.$$

□

For good instances, the second property of good forests always holds regardless of which arcs the forest constructor selects. To satisfy the first property, i.e., to form a directed binary forest, consider a partition of arcs according to their destinations, into groups with one or two arcs each. Then, selecting a directed binary forest is equivalent to selecting at most one arc from each group. The above forest constructor indeed independently and randomly selects an arc from each group. Finally, the independent selections of arcs in $E_{U,e}^{\text{ex-ante}}$ help show that it is a $\frac{1}{2}$ -forest constructor.

5.3.2 General Online Selection Instances, Pseudo-paths, and Pseudo-matchings

For general instances, a forest constructor needs to also ensure the second property of good forests. In other words, for any pairs of arcs a and a' with the same origin such that the rounds corresponding to their incident nodes share a common element, a forest constructor must not select both a and a' into the forest. For example, consider an instance with the ex-ante dependence graph in Figure 1. The directed forest property requires that, e.g., arcs $(1, 3)_a$ and $(2, 3)_b$ cannot be both selected, arcs $(1, 4)_c$ and $(3, 4)_a$ cannot be both selected, etc., due to having the same destinations. The second property of good forests further requires that, e.g., arcs $(1, 3)_a$ and $(1, 4)_c$ cannot be both selected. Dropping their directions, the above arcs $(3, 4)_a, (1, 4)_c, (1, 3)_a, (2, 3)_b$ form an undirected path $3 - 4 - 1 - 3 - 2$. More importantly, we can succinctly describe the aforementioned requirements of good forests as not selecting neighboring arcs with respect to the path. Driven by this observation, we define pseudo-paths and the pseudo-matchings below.

Definition 8 (Pseudo-path). *Given any online selection instance and its ex-ante dependence graph, a pseudo-path is a maximal ordered subset of arcs $P = ((t_i, t'_i)_{e_i})_{1 \leq i \leq \ell}$ such that for any $1 \leq i < \ell$:*

- *Either $t'_i = t'_{i+1}$, i.e., the i -th and $(i+1)$ -th arcs in P have the same destination;*
- *Or $t_i = t_{i+1}$, i.e., the i -th and $(i+1)$ -th arcs in P have the same origin, and rounds $t_i = t_{i+1}$, t'_i , and t'_{i+1} have a common element.*

Lemma 15. *The pseudo-paths partition the arcs of the ex-ante graph.⁵*

Figure 3a shows the partition of arcs into pseudo-paths in the aforementioned example whose ex-ante dependence graph is Figure 1. We defer other structural properties about pseudo-paths to sequel subsections where we use them to design and analyze the forest constructor.

Definition 9 (Pseudo-matching). *For any pseudo-path $P = ((t_i, t'_i)_{e_i})_{1 \leq i \leq \ell}$, a subset of its arcs M is a pseudo-matching if it has no adjacent arcs with respect to P , i.e., for any $1 \leq i < \ell$, either $(t_i, t'_i)_{e_i} \notin M$ or $(t_{i+1}, t'_{i+1})_{e_{i+1}} \notin M$.*

We remark that a pseudo-matching may not be a matching of the ex-ante dependence graph. For example, arcs $(3, 4)_a$ and $(2, 3)_b$ form a pseudo-matching of the left-most pseudo-path in Figure 3a even though they share node 3.

Lemma 16. *A subgraph of the ex-ante dependence graph is a good forest if and only if it is a union of pseudo-matchings, one for each pseudo-path.*

Therefore, a forest constructor needs to pick a pseudo-matching from each pseudo-path. Further, it must do so in an online fashion. On observing the elements \mathcal{E}^t of round t , it either appends t 's in-arcs to an existing pseudo-path and lets them start a new pseudo-path on their own, according to the definition of pseudo-paths. It also immediately decides if to include each in-arc into the pseudo-matching. To make it an α -forest constructor for the largest possible α , we want to select as many arcs into the pseudo-matchings as possible, and at the same time to keep the selections sufficiently independent so that an analysis similar to Lemma 14 applies. The latter refutes selecting either all odd arcs or all even arcs from each pseudo-path with equal probability.

5.3.3 Forest Constructor for General Instances


To explain our forest constructor for general instances, we need a structural lemma about the arrival order of arcs in any pseudo-path. We notice that arcs usually arrive in pairs since the in-arcs of t both arrive in round t . We will artificially break ties to be consistent with the lemma.

Lemma 17. *For any pseudo-path $P = ((t_i, t'_i)_{e_i})_{1 \leq i \leq \ell}$ and any $0 \leq t \leq T$, the subset of arcs that arrive in the first t rounds is a sub-pseudo-path, i.e., either it is an empty set, or there exists $1 \leq i_{\min} \leq i_{\max} \leq \ell$ such that the arrived arcs are $((t_i, t'_i)_{e_i})_{i_{\min} \leq i \leq i_{\max}}$.*

That is, after the first arc of a pseudo-path arrives, the arrival of any arc of the pseudo-path appends to the existing sub-pseudo-path; we never need to merge two pseudo-paths. Our argument lets i_0 denote the index of the earliest arc. We say that the arcs with indices $i \geq i_0$ are on the positive end, and those with indices $i < i_0$ are on the negative end.⁶ Finally, changing the roles of

⁵We consider two pseudo-paths with the same subset of arcs but in opposite orders as the same pseudo-path.

⁶The algorithm does not need to know the index i_0 upfront. Instead, it could let the earliest arc have index 0; arcs on the positive ends have indices 0, 1, 2, etc., and those on the negative end have indices $-1, -2, -3$, etc. Nonetheless, the choice of indices in the main text admits cleaner notations in the analysis.



(a) Automaton σ^+ for the positive end

(b) Automaton σ^- for the negative end

Figure 4: The probabilistic automata in our forest constructor (Algorithm 5). The transitions are labeled by the binary decisions and the probabilities of the transitions.

the positive and negative ends will not affect our conclusion, so we will without loss of generality let the second earliest arc of any pseudo-path be on the positive end.

The forest constructor uses a probabilistic automaton σ^+ and its inverse σ^- . The automata have states q_U , q_{U^2} , and q_M . Intuitively, state q_U means that automaton σ^+ leaves the last arc unmatched, but matches the arc before that; state q_{U^2} means that automaton σ^+ leaves the last two arcs unmatched; and state q_M means that automaton σ^+ matches the last arc. The transition functions, which we denote also as σ^+ and σ^- abusing notations, take a state as input and returns the next state and also a binary decision. They are parameterized by $p \in [0, 1]$, the transition probability from q_U to q_M in automaton σ^+ . We will let $p = 0.6616$ in the analysis to optimize the result. Formally, the transition functions are (see also Figure 4):

$$\sigma^+(q_U) = \begin{cases} (q_U, M) & \text{w.p. } p \\ (q_{U^2}, U) & \text{w.p. } 1 - p \end{cases}, \quad \sigma^+(q_{U^2}) = (q_M, M), \quad \sigma^+(q_M) = (q_U, U);$$

$$\sigma^-(q_U) = (q_M, U), \quad \sigma^-(q_{U^2}) = (q_U, U), \quad \sigma^-(q_M) = \begin{cases} (q_U, M) & \text{w.p. } p \\ (q_{U^2}, M) & \text{w.p. } 1 - p \end{cases}.$$

For each pseudo-path, our forest constructor draws an initial state from the common stationary distribution of the automata. Then, when an arc arrives on the positive end, it calls σ^+ to update the state and to decide whether to include the arc into the pseudo-matching; similarly, when an arc arrives on the negative end, it calls σ^- . See Algorithm 5.

5.3.4 Properties of the Automata

Lemma 18. *The stationary distribution of the states of σ^+ and σ^- is:*

$$\pi = (\pi_U, \pi_{U^2}, \pi_M) = \left(\frac{1}{3-p}, \frac{1-p}{3-p}, \frac{1}{3-p} \right).$$

The next lemma formalizes the claim that σ^- is the inverse of σ^+ .

Lemma 19. *Consider two sequences of random variables:*

1. Sample q^0 from the stationary distribution π . Then recursively let:

$$(q^i, d^i) = \sigma^+(q^{i-1}), \quad 1 \leq i \leq \ell.$$

Algorithm 5 0.404-Forest constructor for general instances (when $p = 0.6616$)

State variables: (for each pseudo-path P)

- $q^+, q^- \in \{q_U, q_{U^2}, q_M\}$ of automata σ^+, σ^- respectively.
- Initialize (when the first arc in the pseudo-path arrives):

$$q^+ = q^- = \begin{cases} q_U & \text{w.p. } \frac{1}{3-p}; \\ q_{U^2} & \text{w.p. } \frac{1-p}{3-p}; \\ q_M & \text{w.p. } \frac{1}{3-p}. \end{cases}$$

For each arc: (of pseudo-path P)

1. Let $\tau = +$ if the arc is on the positive end, and $-$ otherwise.
 2. Let $(q^\tau, d) = \sigma^\tau(\text{state}^\tau)$.
 3. Include the arc into the pseudo-matching and thus the forest if $d = M$.
-

2. Sample \hat{q}^{i_0-1} from the stationary distribution π . Then recursively let:

$$\begin{aligned} (\hat{q}^i, \hat{d}^i) &= \sigma^+(\hat{q}^{i-1}), & i_0 \leq i \leq \ell; \\ (\hat{q}^{i-1}, \hat{d}^i) &= \sigma^-(\hat{q}^i), & 1 \leq i < i_0. \end{aligned}$$

The two sequences are identically distributed.

The second sequence corresponds to the states and decisions of automata σ^+, σ^- in Algorithm 5: \hat{q}^{i_0-1} is the common initial state of σ^+ and σ^- ; \hat{q}^i and \hat{d}^i for $i \geq i_0$ are the states and decisions on the positive end; \hat{q}^i and \hat{d}^i for $i < i_0$ are the states and decisions on the negative end. Lemma 19 allows us to analyze each pseudo-path as if the arcs' arrival order is from one end to the other.

Next we adopt the viewpoint of selecting arcs with only automaton σ^+ , and develop several properties of the corresponding sequence $(q^0, d^1, q^1, \dots, d^\ell, q^\ell)$. We start with three lemmas that follow by the definition of σ^+ .

Lemma 20. For any $1 \leq i \leq \ell$, $q^i = q_M$ if and only if $d^i = M$.

Lemma 21. For any $1 \leq i \leq \ell$:

$$\Pr[d^i = M] = \frac{1}{3-p}.$$

Lemma 22. For any $1 \leq i \leq \ell - 2$, at least one of d^i, d^{i+1} , and d^{i+2} equals M .

Lemma 20 asserts that every time σ^+ selects an arc into the pseudo-matching (and thus the forest), it resets to state q_M . Combining with Lemma 18, we get the marginal selection probability in Lemma 21. Lemma 22 further claims that it resets to state q_M at least once every three rounds. Hence, we focus on how the probabilistic automaton transitions starting from state q_M , and in particular how likely the i -th arc after that would be selected, which in turns characterizes the transition from the other two states. These probabilities are characterized by a recurrence:

$$f_i = \begin{cases} 1 & i = 0; \\ 0 & i = 1; \\ p & i = 2; \\ pf_{i-2} + (1-p)f_{i-3} & i \geq 3. \end{cases} \quad (3)$$

Lemma 23. For any $i \leq j$:

$$\Pr[d^j = M \mid q^i = q_M] = \Pr[d^j = M \mid d^i = M] = f_{j-i} .$$

Further, for any $i \leq j - 1$:

$$\begin{aligned} \Pr[d^j = M \mid q^i = q_{U^2}] &= f_{j-i-1} , \\ \Pr[d^j = M \mid q^i = q_U] &= f_{j-i+1} . \end{aligned}$$

Lemma 24. Suppose that $\frac{\sqrt{5}-1}{2} \leq p \leq \frac{2}{3}$. Then:

$$\begin{aligned} f_i &\geq 1 - p , & \forall i \geq 2 ; \\ f_i &\geq p^3 + (1 - p)^2 & \forall i \geq 4 . \end{aligned}$$

5.3.5 Analysis of Forest Constructor for General Instances: Proof of Lemma 11

This subsection proves Lemma 11 by showing that Algorithm 5 with $p = 0.6616$ is a 0.404-forest constructor. Below summarizes some properties that either follow by the definition of Algorithm 5, or have been established in the previous subsections:

- Constructing a good forest is the same as selecting a pseudo-matching from each pseudo-path. (Lemma 16)
- The selections of arcs in different pseudo-paths are independent. (Definition of Algorithm 5)
- The selections of arcs on a pseudo-path is equivalent to sampling a state from the stationary distribution π , and applying σ^+ to decide for each arc from one end to the other. (Lemma 19)
- Automaton σ^+ selects an arc and resets to state q_M at least once every three rounds. (Lemmas 20 and 22)
- The probability of selecting the i -th arc after any state is characterized by f_i , which can be lower bounded. (Lemmas 23 and 24)

Proving that Algorithm 5 Constructs a Good Forest. Consider an arbitrary pseudo-path. By the definition of σ^+ , it resets its state to q_M when the decision is M (Lemma 16), i.e., when an arc is selected, after which the next decision will be U. Therefore, the arcs selected from each pseudo-path are a pseudo-matching. By Lemma 16 this is a good forest.

Proof of Equation (2). To show the guarantee of a α -forest constructor, which we restate below:

$$\forall 0 \leq \beta \leq 1 : \quad \mathbf{E}(1 - \beta)^{|E_U^{\text{forest}}|} \leq (1 - \alpha\beta)^{|E_{U,e}^{\text{ex-ante}}|} , \quad (\text{Eqn. (2) restated})$$

it suffices to consider each pseudo-path P separately and to show that:

$$\forall 0 \leq \beta \leq 1 : \quad \mathbf{E}(1 - \beta)^{|E_U^{\text{forest}} \cap P|} \leq (1 - \alpha\beta)^{|E_{U,e}^{\text{ex-ante}} \cap P|} , \quad (4)$$

after which Eqn. (2) follows by taking the product of Eqn. (4) over all pseudo-paths, and by the independence of arc selections in different pseudo-paths.

The rest of the argument considers an arbitrary pseudo-path P and proves Eqn. (4). We start by establishing the last structural lemma about pseudo-paths, characterizing the subset of arcs that could contribute to the inequality by being counted in E_U^{forest} .

Lemma 25. *For any element e , and any subset of nodes $U \subseteq V$ involving e , there is a subset of arcs in P with both nodes inside U such that:*

1. *It is a superset of $E_{U,e}^{\text{ex-ante}} \cap P$;*
2. *It is the union of odd-length sub-pseudo-paths;*
3. *Any two of these sub-pseudo-paths are at least 3 arcs apart; and*
4. *Each sub-pseudo-path alternates between arcs with subscript e , i.e., arcs that also contribute to the right-hand-side of Eqn. (4), and arcs with other subscripts, i.e., arcs that only contribute to the left-hand-side. The arcs on the two ends have subscript e .*

Given Lemma 25, we may assume that the subset of arcs in P with both nodes involving e are sub-pseudo-paths starting from arc indices i_1, i_2, \dots, i_m and lengths $2k_1 + 1, 2k_2 + 1, \dots, 2k_m + 1$. Let \mathcal{I}_j denote the set of indices of the j -th sub-pseudo-path:

$$\mathcal{I}_j = \{i_j, i_j + 1, \dots, i_j + 2k_j\} .$$

Let $\mathcal{I} = \cup_{j=1}^m \mathcal{I}_j$ denote the set of arc indices in these pseudo-paths. For each arc $i \in \mathcal{I}$, consider the indicator of if the i -th arc on the pseudo-path is selected into the pseudo-matching:

$$X_i \stackrel{\text{def}}{=} \mathbf{1}(d^i = \mathbf{M}) .$$

Since any arc in \mathcal{I} have both nodes in U by definition, we have:

$$|E_U^{\text{forest}} \cap P| \geq \sum_{i \in \mathcal{I}} X_i .$$

Further, there are $k_j + 1$ arcs with subscripts e in the j -th sub-pseudo-path (i.e., contributing to $E_{U,e}^{\text{ex-ante}} \cap P$) by Lemma 25. Hence, to prove Eqn. (4) it is sufficient to show:

$$\mathbf{E}(1 - \beta)^{\sum_{i \in \mathcal{I}} X_i} \leq (1 - \alpha\beta)^{\sum_{j=1}^m (k_j + 1)} .$$

We next argue that it suffices to consider the case when all sub-pseudo-paths have unit lengths, because we can reduce the general case to it. Suppose that the j -th pseudo-path has length $2k_j + 1 > 1$. When $2k_j + 1 = 3$ or $2k_j + 1 \geq 7$, by Lemma 22 we have:

$$\sum_{i \in \mathcal{I}_j} X_i \geq \left\lfloor \frac{2k_j + 1}{3} \right\rfloor \geq \frac{k_j + 1}{2} .$$

Hence, regardless of the realization of randomness in the forest constructor (Algorithm 5), for $\alpha = 0.404$ and for any $\beta \in [0, 1]$ we always have:

$$(1 - \beta)^{\sum_{i \in \mathcal{I}_j} X_i} \leq (1 - \beta)^{\frac{k_j + 1}{2}} \leq \left(1 - \frac{\beta}{2}\right)^{k_j + 1} \leq (1 - \alpha\beta)^{k_j + 1} .$$

In other words, we can without loss of generality remove the j -th sub-pseudo-path and prove Eqn. (4) for the remaining instance.

When $2k_j + 1 = 5$, i.e., $k_j = 2$, also by Lemma 22 we have:

$$\sum_{i \in \mathcal{I}_j, i \neq i_j} X_i \geq 1 .$$

Hence, regardless of the realization of randomness in the forest constructor (Algorithm 5), for $\alpha = 0.404$ and for any $\beta \in [0, 1]$ we always have:

$$(1 - \beta)^{\sum_{i \in \mathcal{I}_j, i \neq i_j} X_i} \leq (1 - \beta) \leq \left(1 - \frac{\beta}{2}\right)^2 \leq (1 - \alpha\beta)^2 .$$

That is, we can without loss of generality remove the arcs *other than* i_j from j -th sub-pseudo-path and prove Eqn. (4) for the remaining instance.

Finally, consider Eqn. (4) when all sub-pseudo-paths have unit lengths. In other words, for a subset of indices $\mathcal{I} = \{i_1, i_2, \dots, i_m\}$ such that any two indices differ by at least 4 (Lemma 25), we shall prove that:

$$\forall 0 \leq \beta \leq 1 : \quad \mathbf{E}(1 - \beta)^{\sum_{i \in \mathcal{I}} X_i} \leq (1 - \alpha\beta)^m .$$

Our proof is an induction on m . The base case when $m = 1$ follows by:

$$\begin{aligned} \mathbf{E}(1 - \beta)^{X_{i_1}} &\leq \frac{1}{3 - p}(1 - \beta) + \frac{2 - p}{3 - p} && \text{(Lemma 21)} \\ &= 1 - \frac{1}{3 - p}\beta \\ &\leq 1 - \alpha\beta . && (\alpha = 0.404, p = 0.6616) \end{aligned}$$

Suppose that the inequality holds for up to $m - 1$ indices. We next prove it for m indices. Suppose without loss of generality that $i_1 < i_2 < \dots < i_m$. By the inductive hypothesis:

$$\mathbf{E}(1 - \beta)^{\sum_{i \in \mathcal{I}, i \neq i_m} X_i} \leq (1 - \alpha\beta)^{m-1} .$$

It suffices to prove that *for any realized* $X_{i_1}, X_{i_2}, \dots, X_{i_{m-1}}$:

$$\mathbf{E}[X_{i_m} \mid X_{i_1}, \dots, X_{i_{m-1}}] \geq \alpha ,$$

as it would imply that:

$$\mathbf{E}\left[(1 - \beta)^{X_{i_m}} \mid X_{i_1}, \dots, X_{i_{m-1}} = 1\right] \leq (1 - \beta)\alpha + (1 - \alpha) = 1 - \alpha\beta .$$

For any realization such that $X_{i_{m-1}} = 1$:

$$\begin{aligned} \mathbf{E}[X_{i_m} \mid X_{i_1}, \dots, X_{i_{m-1}} = 1] &= f_{i_m - i_{m-1}} && \text{(Lemmas 20, 23)} \\ &\geq p^3 + (1 - p)^2 && (i_m - i_{m-1} \geq 4 \text{ and Lemma 24)} \\ &\geq \alpha . && (\alpha = 0.404, p = 0.6616) \end{aligned}$$

For a realization such that $X_{i_{m-1}} = 0$, the state after processing arc i_{m-1} could be q_U or q_{U^2} . We first argue that it is the former at least a $\frac{1}{1+p}$ fraction of the time.

Lemma 26. *For any realization such that $X_{i_{m-1}} = 0$:*

$$\mathbf{Pr}\left[q^{i_{m-1}} = q_U \mid X_{i_1}, \dots, X_{i_{m-1}} = 0\right] \geq \frac{1}{1 + p} .$$

Given Lemma 26, we can lower bound $\mathbf{E}[X_{i_m} \mid X_{i_1}, \dots, X_{i_{m-1}} = 0]$ by:

$$\frac{1}{1 + p}\mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_U] + \frac{p}{1 + p} \min \left\{ \mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_U], \mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_{U^2}] \right\} .$$

Algorithm 6 $(\sqrt{2} - 1)$ -Forest OCS

Input:

- An online selection instance represented by its ex-ante graph $G^{\text{ex-ante}} = (V, E^{\text{ex-ante}})$.
- A good forest $G^{\text{forest}} = (V, E^{\text{forest}})$.

State variables:

- Label $\ell_t(e) \in \{\text{H}, \text{T}\}$ for every round $1 \leq t \leq T$ and every element $e \in \mathcal{E}^t$.
- State q^t of automaton σ^* for every node $1 \leq t \leq T$.

For each round t : (suppose that $\mathcal{E}^t = \{e_1, e_2\}$)

1. If t is the root of a directed binary tree in G^{forest} , let the labels be $\ell_t(e_1) = \text{H}$, $\ell_t(e_2) = \text{T}$.
 2. Otherwise, suppose without loss of generality that $(t', t)_{e_1}$ is the in-arc of node t in G^{forest} , let $\ell_t(e_1) = \ell_{t'}(e_1)$ and let $\ell_t(e_2)$ be the other label from $\{\text{H}, \text{T}\}$.
 3. Let $(q^t, \ell) = \sigma^*(q^{t'})$ (artificially let $q^{t'} = q_{\text{O}}$ if t is the root of a directed binary tree).
 4. Select the element with label ℓ .
-

Suppose that $i_m - i_{m-1} = 4$. By Lemma 24:

$$\mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_{\text{U}}] = f_5 = 2p(1-p) ;$$

$$\mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_{\text{U}^2}] = f_3 = 1-p .$$

Observe that $2p(1-p) > 1-p$ for $p = 0.6616$. We get that:

$$\begin{aligned} \mathbf{E}[X_{i_m} \mid X_{i_1}, \dots, X_{i_{m-1}} = 0] &\geq \frac{1}{1+p} \cdot 2p(1-p) + \frac{p}{1+p} \cdot (1-p) \\ &= \frac{3p(1-p)}{1+p} \geq \alpha . \end{aligned} \quad (\alpha = 0.404, p = 0.6616)$$

Otherwise, we have that $i_m - i_{m-1} \geq 5$. By Lemma 24:

$$\mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_{\text{U}}] = f_{i_m - i_{m-1} + 1} \geq p^3 + (1-p)^2 ;$$

$$\mathbf{E}[X_{i_m} \mid q^{i_{m-1}} = q_{\text{U}^2}] = f_{i_m - i_{m-1} - 1} \geq p^3 + (1-p)^2 .$$

Hence:

$$\mathbf{E}[X_{i_m} \mid X_{i_1}, \dots, X_{i_{m-1}} = 0] \geq p^3 + (1-p)^2 \geq \alpha . \quad (\alpha = 0.404, p = 0.6616)$$

5.4 Forest OCS

5.4.1 Algorithm

On observing the elements \mathcal{E}^t of each round t , the forest OCS labels the elements by head (H) and tail (T). Then, it calls automaton σ^* from Subsection 5.2.1 with the state of t 's parent in the good forest as input to select a label and to get the state of t . Finally, it selects the element whose label is selected by automaton σ^* . See Algorithm 6.

5.4.2 Analysis: Proof of Lemma 12

The lemmas in this subsection assume $\beta = \sqrt{2} - 1$, which we shall not restate repeatedly. We first establish a structural lemma about the subset of arcs with both nodes involving an element e in any good forest, and about the labels in Algorithm 6. Recall that E_U^{forest} denotes the subset of arcs in forest G^{forest} with both nodes in U .

Lemma 27. *For any good forest $G^{\text{forest}} = (V, E^{\text{forest}})$, any element e , and any subset of nodes $U \subseteq V$ involving e , E_U^{forest} consists of a collection of tree-paths such that:*

1. *There is no arc between any two nodes in distinct tree-paths; and*
2. *Element e has the same label in each tree-path.*

We next prove Lemma 12 by an induction on the number of tree-paths in E_U^{forest} . The base case with zero tree-path holds vacuously.

Next for some $m \geq 1$ suppose that the lemma holds with at most $m - 1$ tree-paths. Consider an arbitrary instance for which E_U^{forest} consists of m tree-paths satisfying the properties of Lemma 27. Let t_1, t_2, \dots, t_m denote the first node of these tree-paths. Let k_1, k_2, \dots, k_m be the lengths. We assume without loss of generality that t_m 's height in its directed binary tree is greater than or equal to the height of any other t_i from the same tree.

Case 1: Suppose that t_m is the root of a directed binary tree in G^{forest} . First by the inductive hypothesis on the sub-instance that removes the tree rooted at t_m , the probability of not selecting element e in these tree-paths is at most $2^{-\sum_{i=1}^{m-1} k_i} (1 - \beta)^{\sum_{i=1}^{m-1} (k_i - 1)}$. Then, *conditioned on any realized randomness on the other tree-paths*, the probability of never selecting e on the tree-path starting with t_m is at most $2^{-k_m} (1 - \beta)^{k_m - 1}$, because of the second part of Lemma 27 and a special case of the next lemma when $i = 1$.

Lemma 28. *For any k consecutive positive integers $i, i + 1, \dots, i + k - 1$, and any label $\ell \in \{\text{H}, \text{T}\}$, the probability that automaton σ^* , starting from the original state q_0 , does not select label ℓ in its i -th to $(i + k - 1)$ -th selections is at most $2^{-k} (1 - \beta)^{k - 1}$.*

Case 2: Suppose that t_m is not the root of any directed binary tree in G^{forest} , but its sibling in G^{forest} (if any) is not one of t_1, t_2, \dots, t_{m-1} . By the latter assumption, the fact that t_m 's parent is not on the other tree-paths (Lemma 27, first part), and the assumption that t_m has the largest height compared to other t_i from the same directed binary tree, we get that the nodes rooted from t_m 's parent are not on the other tree-paths. Then, by the inductive hypothesis on the sub-instance that removes nodes rooted from t_m 's parent, the probability of not selecting element e in these tree-paths is at most $2^{-\sum_{i=1}^{m-1} k_i} (1 - \beta)^{\sum_{i=1}^{m-1} (k_i - 1)}$. Then, *conditioned on any realized randomness on the other tree-paths*, the probability of never selecting e on the tree-path starting with t_m is at most $2^{-k_m} (1 - \beta)^{k_m - 1}$, because of the second part of Lemma 27 and the next lemma.

Lemma 29. *For any k consecutive positive integers $i, i + 1, \dots, i + k - 1$ starting from $i \geq 2$, and any label $\ell \in \{\text{H}, \text{T}\}$, the probability that automaton σ^* , starting from an arbitrary state, does not select label ℓ in its i -th to $(i + k - 1)$ -th selections is at most $2^{-k} (1 - \beta)^{k - 1}$.*

Case 3: Suppose that t_m is not the root of any directed binary tree in G^{forest} , and further its sibling in G^{forest} is t_j , the starting node of another tree-path in Lemma 27. Since t_m 's parent is not on the other tree-paths (Lemma 27, first part), and further by the assumption that t_m has the

largest height compared to other t_i from the same directed binary tree, the nodes rooted from t_m and t_j 's parent are not on the other tree-paths. By the inductive hypothesis on the sub-instance that removes the nodes rooted from t_m 's parent, the probability of not selecting element e in these tree-paths is at most $2^{-\sum_{i \neq j, m} k_i} (1 - \beta)^{\sum_{i \neq j, m} (k_i - 1)}$. Then, *conditioned on any realized randomness on the other tree-paths*, the probability of never selecting e on the tree-paths starting with t_j and t_m is at most $2^{-k_j - k_m} (1 - \beta)^{(k_j - 1) + (k_m - 1)}$, because of the second part of Lemma 27, the observation that e must get different labels on these two tree-paths,⁷ and the next lemma.

Lemma 30. *Consider two independent copies of automaton σ^* from an arbitrary but identical initial state. Then, for any $k, \hat{k} \geq 1$, the probability that the first copy never selects \top in the first k rounds, and the second copy never selects H in the first \hat{k} rounds is at most $2^{-k - \hat{k}} (1 - \beta)^{(k-1) + (\hat{k}-1)}$.*

5.5 Hardness

Fahrback et al. [18] rule out the possibility of 1-OCS. This subsection improves the upper bound to $\frac{1}{4}$. This holds even for algorithms with unlimited computational power, and even if the algorithms know the instance beforehand.

Theorem 31. *There is no $(\frac{1}{4} + \epsilon)$ -OCS for any constant $\epsilon > 0$.*

Proof. Consider any γ -OCS. Consider an online selection instance with three elements $\{0, 1, 2\}$ and $T = 2i + 1$ rounds. The elements in odd rounds are $\{0, 1\}$; the elements in even rounds are $\{0, 2\}$. We shall prove that $\gamma \leq \frac{i}{4i-1}$ even if we omit the properties of γ -OCS concerning three or more rounds. Theorem 31 then follows by choosing a sufficiently large i .

For any even $1 \leq t \leq T$ (so that the elements are $\{0, 2\}$), let A^t be the event that the algorithm selects element 0 in rounds t . We have:

$$\forall 1 \leq t \leq T, t \equiv 0 \pmod{2} : \quad \Pr[A^t] = \frac{1}{2}.$$

For any $1 \leq t < t' \leq T$ with distinct parities (so that 0 is the only common element), let $B^{t, t'}$ be the event that the algorithms does not select element 0 in both rounds t and t' . We have:

$$\forall 1 \leq t < t' \leq T, t + t' \equiv 1 \pmod{2} : \quad \Pr[B^{t, t'}] \leq \begin{cases} \frac{1}{4} & t' \geq t + 3 ; \\ \frac{1-\gamma}{4} & t' = t + 1 . \end{cases}$$

For any $1 \leq t < t' \leq T$ with the same parity (so that they have the same two elements), let $C^{t, t'}$ be the event that the algorithms selects element 0 in both rounds t and t' (i.e., it does not select the other element in both rounds). We have:

$$\forall 1 \leq t < t' \leq T, t + t' \equiv 0 \pmod{2} : \quad \Pr[C^{t, t'}] \leq \begin{cases} \frac{1}{4} & t' \geq t + 4 ; \\ \frac{1-\gamma}{4} & t' = t + 2 . \end{cases}$$

⁷The parent p of t_j, t_m must not contain e by the second requirement of good forest. Suppose that p 's common elements with t_j, t_m are $e_j, e_m \neq e$ respectively. Then, e 's labels in t_j, t_m are the labels of e_m, e_j in p respectively.

Summing together:

$$\begin{aligned}
& \sum_{t \equiv 0 \pmod 2} \Pr[A^t] + \sum_{t < t': t+t' \equiv 1 \pmod 2} \Pr[B^{t,t'}] + \sum_{t < t': t+t' \equiv 1 \pmod 2} \Pr[C^{t,t'}] \\
& \leq \underbrace{\frac{1}{2} \cdot i}_{\text{events } A^t} + \underbrace{\frac{1}{4} \cdot i(i-1) + \frac{1-\gamma}{4} \cdot 2i}_{\text{events } B^{t,t'}} + \underbrace{\frac{1}{4} \cdot (i-1)^2 + \frac{1-\gamma}{4} \cdot (2i-1)}_{\text{events } C^{t,t'}} \\
& = \frac{2i^2 + 3i}{4} - \frac{4i-1}{4} \gamma. \tag{5}
\end{aligned}$$

Next consider any selections (s^1, s^2, \dots, s^T) by the algorithm. Let j_o and j_e be the numbers of odd and even rounds that select $s^t = 0$ respectively. The number of events that it satisfies equals:

$$\begin{aligned}
& \underbrace{j_e}_{\text{events } A^t} + \underbrace{(i+1-j_o)(i-j_e)}_{\text{events } B^{t,t'}} + \underbrace{\binom{j_o}{2} + \binom{j_e}{2}}_{\text{events } C^{t,t'}} = \frac{1}{2}(j_o + j_e - i - \frac{1}{2})^2 + \frac{i(i+1)}{2} - \frac{1}{8} \\
& \geq \frac{i(i+1)}{2}. \quad (i, j_o, j_e \text{ are integers})
\end{aligned}$$

As a result we get that:

$$\sum_{t \equiv 0 \pmod 2} \Pr[A^t] + \sum_{t < t': t+t' \equiv 1 \pmod 2} \Pr[B^{t,t'}] + \sum_{t < t': t+t' \equiv 1 \pmod 2} \Pr[C^{t,t'}] \geq \frac{i(i+1)}{2}. \tag{6}$$

Combining Equations (5) and (6) gives $\gamma \leq \frac{i}{4i-1}$ as desired. \square

6 Applications in Online Bipartite Matching

6.1 Online Bipartite Matching Preliminaries

Consider an undirected bipartite graph $G = (L, R, E)$, where L and R are the sets of left-hand-side and right-hand-side vertices respectively, and E is the set of edges. Each edge $(u, v) \in E$ has a positive edge-weight $w_{uv} > 0$. The problem is *unweighted* if $w_{uv} = 1$ for all $(u, v) \in E$, is *vertex-weighted* if $w_{uv} = w_u$ for some positive vertex-weights $(w_u)_{u \in L}$ of the left-hand-side vertices, and is *edge-weighted* if the edge-weights could be arbitrary.

In online bipartite matching problems, we refer to the left-hand-side and right-hand-side vertices as *offline* and *online* vertices respectively. Initially, the algorithm only knows the offline vertices, and the vertex-weights in the vertex-weighted case. Then, the online vertices arrive one at a time. When an online vertex $v \in R$ arrives, the algorithm sees its incident edges, and the edge-weights in the edge-weighted case. The algorithm then immediately and irrevocably matches v to an offline neighbor $u \in L$.

The objective is to maximize the sum of the maximal edge-weight matched to each offline vertex. In the unweighted and vertex-weighted problems, matching an offline vertex more than once does not further increase the objective. Therefore, we may assume without loss of generality that the algorithm matches each offline vertex at most once and the matched edges indeed form matching. The objectives in these two cases are equivalent to maximizing the cardinality of the matching, and maximizing the sum of the vertex-weights of matched offline vertices, respectively.

In edge-weighted online bipartite matching, we may alternatively view the above objective as allowing disposals of previously matched edges so that a matched offline vertex could be rematched

to a new edge with a larger edge-weight. In other words, we may think of the matching as being comprised of the heaviest edge matched to each offline vertex, and seek to maximize the total edge-weight of the matching. Further in online advertising, it corresponds to displaying an advertiser's ad multiple times but only charges for the most valuable one. Feldman et al. [19] introduce this *free disposal* model which has then become the standard model of edge-weighted online bipartite matching under worst-case competitive analysis.

We compare the expected objective of the matching by the algorithm, and the optimal matching that maximizes the objective in hindsight given full information of the bipartite graph $G = (L, R, E)$ and the edge-weights $(w_{uv})_{(u,v) \in E}$. The competitive ratio of an online algorithm is the infimum of this ratio over all possible instances.

6.2 Semi-OCS and Unweighted and Vertex-weighted Online Bipartite Matching

Fahrback et al. [18] give a two-choice greedy algorithm for unweighted online bipartite matching, using a semi-OCS as a sub-routine. Their original theorem is only for the unweighted problem and only for the guarantee of γ -semi OCS, i.e., $p(k) = 2^{-k}(1 - \gamma)^{k-1}$. Nonetheless, the algorithm and analysis generalize to the vertex-weighted case and for general $p(k)$ by standard techniques in the online matching literature. We state the more general theorem below.

Theorem 32 (c.f., Fahrback et al. [18]). *Given a semi-OCS such that the probability of never selecting an element e that appears k times is at most $p(k)$, there is a Γ -competitive two-choice greedy algorithm for unweighted and vertex-weighted online bipartite matching, where the competitive ratio Γ is the optimal value of the following linear program (LP):*

$$\text{maximize } \Gamma \quad (\text{MATCHING LP})$$

$$\text{subject to } a(k) + b(k) \leq p(k) - p(k+1) \quad \forall k \geq 0 \quad (7)$$

$$\sum_{i=0}^{k-1} a(i) + 2b(k) \geq \Gamma \quad \forall k \geq 0 \quad (8)$$

$$b(k+1) \leq b(k) \quad \forall k \geq 0 \quad (9)$$

$$a(k), b(k) \geq 0 \quad \forall k \geq 0$$

We will not present the generalized algorithm and the proof of Theorem 32 because they will be subsumed by the algorithm and theorem in the next subsection. Instead, the main result of this subsection is an explicit optimal solution to the LP. By contrast, Fahrback et al. [18] rely on solving a finite approximation of the LP numerically using LP solvers.

Theorem 33. *Suppose that $p(0) = 1$ and $p(k+1) \leq \frac{2}{3}p(k)$ for any $k \geq 0$. Then, the MATCHING LP admits an optimal solution as follows:*

$$\Gamma = 1 - \frac{1}{3} \sum_{i=0}^{\infty} \left(\frac{2}{3}\right)^i p(i) ;$$

$$b(k) = \frac{1}{3} \sum_{i=k}^{\infty} \left(\frac{2}{3}\right)^{i-k} (p(i) - p(i+1)) \quad \forall k \geq 0 ;$$

$$a(k) = p(k) - p(k+1) - b(k) \quad \forall k \geq 0 .$$

The assumption of $p(k+1) \leq \frac{2}{3}p(k)$ is essentially without loss of generality since any natural online selection algorithm shall at least halve the unselected probability after each round involving the element. Indeed, even the trivial independent sampling satisfies the stronger $p(k+1) \leq \frac{1}{2}p(k)$. The proof of this theorem is deferred to Appendix D.1.

For γ -semi-OCS, it recovers a result by Huang et al. [24] as a corollary.

Corollary 34 (c.f., Huang et al. [24]). *Suppose that $p(k) = 2^{-k}(1-\gamma)^{k-1}$. Then the optimal value of the MATCHING LP is:*

$$\frac{3+2\gamma}{6+3\gamma}.$$

Since the optimal semi-OCS in Section 3 gives $p(k) = 2^{-2^k+1}$, we have the next corollary through a numerical calculation.

Corollary 35. *The two-choice greedy algorithm using the optimal semi-OCS as a sub-routine is at least 0.536-competitive for unweighted and vertex-weighted online bipartite matching.*

6.3 OCS and Edge-weighted Online Bipartite Matching

6.3.1 Online Primal-Dual Algorithm

This subsection gives a variant of the online primal-dual algorithm of Fahrbach et al. [18] for edge-weighted online bipartite matching, using an OCS as a sub-routine. This variant simplifies the analysis in the next subsection. To simplify exposition, we assume that for every online vertex v there is a unique offline dummy vertex such that the edge between them has weight 0. Then, every online vertex will be matched, although being matched to the dummy vertex is the same as being left unmatched.

For each online vertex v , the algorithm shortlists two candidates u_1, u_2 from v 's neighbors. If the shortlisted candidates are the same, the algorithm matches v to it. Otherwise, the algorithm lets the OCS selects one of them and matches v to the selected one. To explain how the algorithm makes the shortlists, let $k_u(w)$ be the number of times that u is shortlisted thus far due to online vertices with edge-weight $w_{uv} \geq w$. In a round in which $u_1 = u_2 = u$, the corresponding $k_u(w)$'s increase by 2. We remark that $k_u(w) = 0$ for any dummy offline vertex u and for any $w > 0$. The algorithm is parameterized by the optimal solution to the MATCHING LP in Theorem 33. Given the optimal solution, define the “value” of matching an online vertex v to an offline vertex u as:

$$\Delta_u \beta_v \stackrel{\text{def}}{=} \int_0^{w_{uv}} b(k_u(w)) dw - \frac{1}{2} \int_{w_{uv}}^{\infty} \sum_{i=0}^{k_u(w)-1} a(i) dw. \quad (10)$$

For each online vertex v , the algorithm first finds u_1 with the maximum $\Delta_u \beta_v$, and then finds u_2 with the maximum *updated* $\Delta_u \beta_v$. If $u_1 = u_2$, match v to it. Otherwise, match v to the one that the OCS selects. Following the terminology of Fahrbach et al. [18], we call the former a deterministic round, and the latter a randomized round. Their algorithm computes the “values” of deterministic and randomized rounds using different equations. By contrast, our variant computes the “values” using the same Eqn. (10). See Algorithm 7.

6.3.2 Improved Online Primal-Dual Analysis

This subsection improves the analysis of Fahrbach et al. [18] in twofold. First, our edge-weighted result uses the LP in Theorem 32 and its optimal solution in Theorem 33, same as the unweighted

Algorithm 7 Online primal-dual edge-weighted bipartite matching algorithm

State variables: (for each offline vertex u)

- $k_u(w)$: the number of times $u_1 = u$ or $u_2 = u$ and further its edge weight is at least w .

On the arrival of an online vertex $v \in R$:

1. For $\ell \in \{1, 2\}$:
 - (a) Find u_ℓ with maximum $\Delta_u \beta_v$ given by Eqn. (10).
 - (b) Increase $k_{u_\ell}(w)$ by 1 for $0 \leq w \leq w_{u_\ell v}$.
 2. If $u_1 \neq u_2$, let the OCS select one of them, and match v to it. **(Randomized round)**
 3. Otherwise, match v to $u_1 = u_2$. **(Deterministic round)**
-

and vertex-weighted cases. By contrast, the analysis of Fahrbach et al. [18] for edge-weighted online bipartite matching needs to consider a LP with additional constraints. Second, our analysis indicates that the online selection algorithm only needs to guarantee a condition strictly weaker than the property of γ -OCS. It enables us to further explore a variant of OCS in the next subsection to further improve the competitive ratio in edge-weighted online bipartite matching.

Theorem 36. *Suppose that $(p(k))_{k \geq 0}$ is non-increasing and satisfies $p(0) = 1$, and Γ , $(a(k))_{k \geq 0}$, and $(b(k))_{k \geq 0}$ form a solution to the MATCHING LP. Algorithm 7 is Γ -competitive for edge-weighted online bipartite matching if the OCS ensures that for any online selection instance, any element e , and any consecutive subsequences of the rounds involving the element with lengths k_1, k_2, \dots, k_m , element e is unselected in these rounds with probability at most:*

$$p\left(\sum_{i=1}^m k_i\right) + \frac{1}{2} \sum_{i=2}^m \sum_{j=0}^{k_1 + \dots + k_{i-1} - 1} a(j). \quad (11)$$

We make three remarks before presenting the proof of the theorem. First, for unweighted and vertex-weighted online bipartite matching, the online selection algorithm only needs to ensure the above property for the subset of all k rounds involving an element. Then, it degenerates to the guarantee of semi-OCS because $m = 1$ and thus the second term involving the $a(j)$'s disappears. The proof below shall make this explicit.

Further, the guarantee in Eqn. (11) holds almost trivially for natural online selection algorithms when $m \geq 3$. On the one hand, any natural algorithm would at least halve the unselected probability for every round involving the element. Hence, after $\sum_{i=1}^m k_i \geq 3$ rounds, the unselected probability is at most $\frac{1}{8}$. On the other hand, the optimal LP solution from Theorem 33 satisfies that $a(0) \geq \frac{2}{9}$ for all online selection algorithms in the literature and in this paper, and even for the overly idealized algorithm that ensure selecting an element when it appears more than once. Hence, the $a(0)$'s in the second term of Eqn. (11) sum to at least $\frac{2}{9} > \frac{1}{8}$. A similar argument shows that the guarantee holds almost trivially for $m = 2$ if $k_1 + k_2 \geq 3$. Hence, it suffices to slightly enhance the semi-OCS guarantee to further handle either a single consecutive subsequence (but not necessarily starting from the earliest round involving the element as in semi-OCS), or two very short consecutive subsequences. This motivates the variant of OCS in the next subsection.

Finally, Theorem 36 subsumes the analysis of Fahrbach et al. [18] because the original guarantee of γ -OCS satisfies Eqn. (11), as we will prove in the next lemma.

Lemma 37. *Suppose that $\gamma \in [0, \frac{1}{4}]$,⁸ and $p(k) = 2^{-k}(1 - \gamma)^{k-1}$ for $k \geq 0$. Let $(a(k))_{k \geq 0}$ take*

⁸Theorem 31 shows that there is no γ -OCS for $\gamma > \frac{1}{4}$.

values as in the optimal LP solution in Theorem 33. Then, for any positive integers k_1, k_2, \dots, k_m :

$$\prod_{i=1}^m 2^{-k_i} (1-\gamma)^{k_i-1} \leq 2^{-\sum_{i=1}^m k_i} (1-\gamma)^{\sum_{i=1}^m k_i-1} + \frac{1}{2} \sum_{i=2}^m \sum_{j=0}^{k_1+\dots+k_{i-1}-1} a(j) .$$

Proof. In fact we will prove it even dropping all $a(j)$'s for $j \geq 1$. If $m = 1$ the left-hand-side equals the first term on the right-hand-side. If $m \geq 2$, the difference between the left-hand-side and the first term on the right-hand-side is:

$$\begin{aligned} (1 - (1-\gamma)^{m-1}) \prod_{i=1}^m 2^{-k_i} (1-\gamma)^{k_i-1} &\leq (m-1)\gamma \prod_{i=1}^m 2^{-k_i} (1-\gamma)^{k_i-1} \\ &\leq \frac{(m-1)\gamma}{4} . \end{aligned} \quad (\sum_{i=1}^m k_i \geq 2)$$

On the other hand, Theorem 33 indicates that for $p(k) = 2^{-k}(1-\gamma)^{k-1}$:

$$a(0) = \frac{3+\gamma}{12+6\gamma} \geq \frac{\gamma}{2} ,$$

for any $\gamma \leq \frac{1}{4}$.⁹ Hence, the $a(0)$'s in the second term on the right sum to at least $\frac{(m-1)\gamma}{4}$. \square

Proof of Theorem 36. For any offline vertex u , consider the subset of rounds in which u is shortlisted as u_1 or u_2 by Algorithm 7. Further for any weight level $w > 0$, suppose that the subset of rounds in which u is shortlisted by an online vertex v with edge-weight $w_{uv} \geq w$ form consecutive subsequences of lengths k_1, k_2, \dots, k_m . We remark that if $u_1 = u_2 = u$ in the round of some online vertex v with $w_{uv} \geq w$, this deterministic round contributes 2 to the corresponding k_i . In such cases the sequel probability bounds hold trivially because u is matched to an edge weight at least w with certainty. The binding case of our analysis is when there are only randomized rounds. The OCS guarantee in the theorem statement ensures that the probability of matching u to one of them is at least:

$$y_u(w) \stackrel{\text{def}}{=} 1 - p\left(\sum_{i=1}^m k_i\right) - \frac{1}{2} \sum_{i=2}^m \sum_{j=0}^{k_1+\dots+k_{i-1}-1} a(j) . \quad (12)$$

Therefore, the expected maximum edge-weight matched to vertex u is at least $\int_0^\infty y_u(w) dw$. The expected total weight of the matching by Algorithm 7 is at least:

$$\text{ALG} \stackrel{\text{def}}{=} \sum_{u \in L} \int_0^\infty y_u(w) dw .$$

The competitive analysis is a charging argument. For every online vertex $v \in R$, we split the changes of ALG among the shortlisted offline vertices u_1, u_2 and the online vertex v . Formally, let $\alpha_u = \int_0^\infty \alpha_u(w) dw$ be the gain of each offline vertex $u \in L$, where $\alpha_u(w)$ is the contribution from weight-level w . Let β_v denote the gain of each online vertex v . Both are initially zero. Then, as an online vertex v arrives and when $u \in \{u_1, u_2\}$ is shortlisted, suppose that $y_u(w)$ changes by $\Delta y_u(w)$ for any $0 \leq w \leq w_{uv}$:

- Increase β_v by $\Delta_u \beta_v$ according to Eqn. (10), which we restate below:

$$\Delta_u \beta_v \stackrel{\text{def}}{=} \int_0^{w_{uv}} b(k_u(w)) dw - \frac{1}{2} \int_{w_{uv}}^\infty \sum_{i=0}^{k_u(w)-1} a(i) dw .$$

⁹In fact, this holds for any $0 \leq \gamma \leq \frac{\sqrt{61}-5}{6} \approx 0.468$

- Increase $\alpha_u(w)$ by:

$$\Delta\alpha_u(w) \stackrel{\text{def}}{=} \begin{cases} \Delta y_u(w) - b(k_u(w)) & \text{if } w_{uv} \geq w \\ \frac{1}{2} \sum_{i=0}^{k_u(w)-1} a(i) & \text{if } 0 \leq w_{uv} < w \end{cases}. \quad (13)$$

We remark that the values of $k_u(w)$'s in the above charging rules are at the moment when u is shortlisted by the algorithm for online vertex v .

Feasibility of the Charging Rule. We first verify that the total change in α_u and β_v equals the change of ALG due to online vertex v . By Equations (10) and (13), the total change in the vertices' gains equals:

$$\underbrace{\int_0^{w_{uv}} (\Delta y_u(w) - b(k_u(w))) dw + \frac{1}{2} \int_{w_{uv}}^{\infty} \sum_{i=0}^{k_u(w)-1} a(i) dw}_{\text{change of } \alpha_u} + \underbrace{\int_0^{w_{uv}} b(k_u(w)) dw - \frac{1}{2} \int_{w_{uv}}^{\infty} \sum_{i=0}^{k_u(w)-1} a(i) dw}_{\text{change of } \beta_v} = \int_0^{w_{uv}} \Delta y_u(w) dw .$$

Invariant of Offline Gain. Next we show that for any offline vertex u , and any positive weight-level $w > 0$:

$$\alpha_u(w) \geq \sum_{i=0}^{k_u(w)-1} a(i) . \quad (14)$$

Consider the rounds in which u is shortlisted and u_1 or u_2 (or both) and the edge-weight is at least w . Partition them into consecutive subsequences of the rounds that shortlist u , regardless the edge-weights. Let k_1, k_2, \dots, k_m be the lengths of the consecutive subsequences. By considering the changes to $\alpha_i(w)$ due to the rounds in the subsequences, and any $m-1$ rounds involving u between the subsequences, one for each pair of neighboring subsequences, we get that:

$$\begin{aligned} \alpha_u(w) &\geq \underbrace{y_u(w) - \sum_{i=0}^{k_u(w)-1} b(i)}_{\text{rounds in subsequences, 1st case of Eqn. (13)}} + \underbrace{\frac{1}{2} \sum_{i=2}^m \sum_{j=0}^{k_1+\dots+k_{i-1}-1} a(i)}_{\text{rounds in between, 2nd case of Eqn. (13)}} \\ &= 1 - p(k_u(w)) - \sum_{i=0}^{k_u(w)-1} b(i) \quad (\text{Eqn. (12), and } k_u(w) = k_1 + \dots + k_m) \\ &= \sum_{i=0}^{k_u(w)-1} (p(i) - p(i+1) - b(i)) \quad (p(0) = 1) \\ &\geq \sum_{i=0}^{k_u(w)-1} a(i) . \quad (\text{Eqn. (7)}) \end{aligned}$$

This is the only place in our argument that uses Eqn. (11) about the online selection algorithm, indirectly through Eqn. (12). We remark that in unweighted and vertex-weighted online bipartite

matching, there is only one weight level $w = 1$ or $w = w_u$ of concern for any offline vertex u . Hence, there is only a single subsequence with all rounds that shortlist u in the above argument. It suffices to replace Eqn. (11) by the weaker property of semi-OCS.

Non-negativity of Gains. The non-negativity of offline gains follows from the above invariant. The non-negativity of online gains follows by that $\Delta_u\beta_v = 0$ for the dummy vertex u . Hence, the offline neighbors u_1, u_2 shortlisted by Algorithm 7 have non-negative $\Delta_{u_1}\beta_v, \Delta_{u_2}\beta_v$.

Γ -Approximate Equilibrium. The gains cumulated by the online and offline vertices satisfy an approximate equilibrium condition in the sense that for any edge $(u, v) \in E$ the total gain of u and v is at least Γ times the edge weight w_{uv} . By the definition of Algorithm 7, and by that $\Delta_u\beta_v$ in Eqn. (10) is non-increasing in $k_u(w)$'s, we have $\beta_v \geq 2\Delta_u\beta_v$ even when we compute $\Delta_u\beta_v$ using the final values of $k_u(w)$'s. Hence:

$$\alpha_u + \beta_v \geq \int_0^\infty \sum_{i=0}^{k_u(w)-1} a(i)dw + 2\Delta_u\beta_v \quad (\text{Eqn. (14)})$$

$$= \int_0^{w_{uv}} \sum_{i=0}^{k_u(w)-1} a(i)dw + 2 \int_0^{w_{uv}} b(k_u(w))dw \quad (\text{Eqn. (10)})$$

$$\geq \Gamma w_{uv} . \quad (\text{Eqn. (8)})$$

Then, consider an optimal matching $M \subseteq E$. Algorithm 7 is Γ -competitive because:

$$\begin{aligned} \text{ALG} &= \sum_{u \in L} \alpha_u + \sum_{v \in R} \beta_v \\ &\geq \sum_{(u,v) \in M} (\alpha_u + \beta_v) \\ &\geq \Gamma \sum_{(u,v) \in M} w_{uv} . \end{aligned}$$

Finally, we remark that the above analysis is mathematically equivalent an online primal dual analysis under the framework of Devanur et al. [11] for online bipartite matching, and also Devanur et al. [12] and Fahrback et al. [18] for the edge-weighted case. We choose the above exposition to avoid having to introduce the more general framework. \square

Combining Algorithm 7 with the improved 0.167-OCS from Theorem 10 in Section 5 surpasses the state-of-the-art 0.508-competitive algorithm for edge-weighted online bipartite matching by Fahrback et al. [18].

Corollary 38. *There is an two-choice greedy algorithm for edge-weighted online bipartite matching that is at least 0.512-competitive.*

6.4 A Variant of OCS and Edge-weighted Online Bipartite Matching

This subsection considers another online selection algorithm tailored for the relaxed condition in Eqn. (11). Each element is associated with a flag 1 or 0, initialized uniformly at random. In each round t , the algorithm samples an element e from \mathcal{E}^t uniformly at random to probe its flag. If its

Algorithm 8 A variant of OCS designed for edge-weighted online bipartite matching

State variables: (for each element e)

- $\tau_e \in \{0, 1\}$; its initial value τ_e^0 is independently and uniformly at random.

For each round t :

1. Draw $e^t \in \mathcal{E}^t$ uniformly at random.
 2. If $\tau_{e^t} = 1$, select e^t and let $\tau_{e^t} = 0$.
 3. Otherwise, select the other element in \mathcal{E}^t and set $\tau_{e^t} = 1$.
-

flag is 1, the algorithm selects e and sets its flag to 0. Otherwise, the algorithm selects the other element and sets e 's flag to 1. In other words, the algorithm randomly samples an element, lets its flag decides the selection, and flips the flag. See Algorithm 8.

Theorem 39. *Algorithm 8 ensures the selection probability in Eqn. (11) for:*

$$p(k) = 2^{-k - \min\{k, \lceil \frac{k+2}{2} \rceil\}} + k \cdot 2^{-k - \min\{k, \lceil \frac{k+3}{2} \rceil\}} .$$

Combining with Theorem 36 further improves the competitive ratio of edge-weighted online bipartite matching.

Corollary 40. *Algorithm 7, using Algorithm 8 for online selections, is at least 0.519-competitive for edge-weighted online bipartite matching.*

Preliminaries on Boolean Formula with Uniform Input. Algorithm 8 uses two kinds of random bits that are sampled independently and uniformly: the initial flags $(\tau_e^0)_{e \in \mathcal{E}}$, and the sampled elements $(e^t)_{1 \leq t \leq T}$. Viewing these random bits as boolean variables, we will represent each selection event by an XOR clause, i.e., an XOR of a subset these boolean variables, their negates, and the constant 1, such as $X_1 \oplus \neg X_2 \oplus X_3 \oplus 1$. Next we introduce two properties related to XOR clauses with uniform input.

Lemma 41. *For any uniform and independent boolean variables and m XOR clauses such that each variable is in at most one clause, the probability of satisfying all clauses equals 2^{-m} .*

Proof. This is because each clause independently holds with probability half. □

Lemma 42. *For any uniform and independent boolean variables and m XOR clauses such that each variable is in at most two clauses, the probability of satisfying all clauses is at most $2^{-\lceil \frac{m}{2} \rceil}$.*

Proof. Consider an undirected graph $G = (V, E)$ in which the vertices $V = \{1, 2, \dots, m\}$ correspond to clauses, and the edges correspond the boolean variables that appears in two clauses:

$$E = \{(i, j)_k : \text{clauses } i \text{ and } j \text{ both involve variable } k\} .$$

Next, consider any maximal matching M of G . Let $\ell = |M|$ be the size of the matching. Let $U \subseteq V$ denote the set of $m - 2\ell$ unmatched vertices (i.e., clauses). Since the matching is maximal, the unmatched clauses do not share any variables. Hence, over the randomness of the variables not in the matching M , the probability of satisfying all unmatched clauses equals $2^{-m+2\ell}$ by Lemma 41.

Further, *conditioned on any realization of the variables not in the matching M* and over the randomness of the variables in M , each pair of matched clauses hold with probability at most $\frac{1}{2}$.

Therefore, the probability of satisfying all clauses is at most $2^{-m+2\ell} \cdot 2^{-\ell} = 2^{-m+\ell}$. The lemma then follows by $\ell \leq \lfloor \frac{m}{2} \rfloor$. □

Selection Probabilities. We next develop a lemma about probability of not selecting an element e conditioned on the sampled elements e^{t_i} 's. The proof of Theorem 39 will repeatedly use the lemma.

Lemma 43. *For any element e and any k rounds $t_1 < t_2 < \dots < t_k$ involving e , conditioned on any realization of e^{t_1}, \dots, e^{t_k} , the probability that e is never selected in these k rounds is at most:*

$$2^{-\min\{k, \lceil \frac{k+2}{2} \rceil\}} .$$

Proof. We will prove a stronger result. If there are d distinct elements in the realized e^{t_1}, \dots, e^{t_k} , then e is unselected in these rounds with probability at most:

$$\begin{cases} 2^{-k} & \text{if } d = 1; \\ 2^{-\lceil \frac{k+d}{2} \rceil} & \text{if } d \geq 2. \end{cases}$$

We next introduce an XOR clause for each t_i , $1 \leq i \leq k$, so that not selecting e in these rounds is equivalent to satisfying all k clauses. If round t_i is the earliest among these k rounds that samples element e^{t_i} , i.e., $e^{t_i} \neq e^{t_j}$ for any $j < i$, consider a clause that represents the value of flag $\tau_{e^{t_i}}$ at the beginning of round t_i :

$$\begin{cases} \tau_{e^{t_i}}^0 \oplus \left(\bigoplus_{t < t_i: e^{t_i} \in \mathcal{E}^t} \mathbf{1}(e^t = e^{t_i}) \right) & e^{t_i} = e ; \\ 1 \oplus \tau_{e^{t_i}}^0 \oplus \left(\bigoplus_{t < t_i: e^{t_i} \in \mathcal{E}^t} \mathbf{1}(e^t = e^{t_i}) \right) & e^{t_i} \neq e . \end{cases}$$

We shall refer to such clauses as type-A clauses.

Otherwise, suppose that element e^{t_i} was most recently sampled in t_j , i.e., $e^{t_i} = e^{t_j}$ and $e^{t_i} \neq e^{t_\ell}$ for $j < \ell < i$. Consider a clause that represents the parity of the number of times that flag $\tau_{e^{t_i}}$ flips between the two rounds, *including the flip due to round t_j* , i.e.:

$$1 \oplus \left(\bigoplus_{t_j < t < t_i: e^{t_i} \in \mathcal{E}^t} \mathbf{1}(e^t = e^{t_i}) \right) .$$

We shall refer to such clauses as type-B clauses. It captures if the value of τ_e at the beginning of t_i is the same as that at the beginning of round t_j , and thus still leads to not selecting element e .

If $d = 1$, there are 1 type-A clause and $k - 1$ type-B clauses. Further, each variable appears in at most one clause. It then follows by Lemma 41.

If $d \geq 2$, there are d type-A clauses and $k - d$ type-B clauses. First consider the type-B clauses and the random variables corresponding to the sampled elements $(e^t)_{1 \leq t \leq T}$. Each of these variables appears in at most two clauses. By Lemma 42 the probability of satisfying all these clauses is at most $2^{-\lceil \frac{k-d}{2} \rceil}$. Further, each type-A clause has a unique variable $\tau_{e^{t_i}}^0$. Hence, over any realization of the sampled elements, the probability of satisfying all d type-A clauses is 2^{-d} . Combining the two bounds proves the lemma. \square

Proof of Theorem 39. First recall the requirement of Eqn. (11). For any element and any consecutive subsequences of the rounds involving the element with lengths k_1, k_2, \dots, k_m , we will upper bound the probability that Algorithm 8 never selects e in these rounds by:

$$p\left(\sum_{i=1}^m k_i\right) + \frac{1}{2} \sum_{i=2}^m \sum_{j=0}^{k_1 + \dots + k_{i-1} - 1} a(j) ,$$

where:

$$p(k) = 2^{-k - \min\{k, \lceil \frac{k+2}{2} \rceil\}} + k 2^{-k - \min\{k, \lceil \frac{k+3}{2} \rceil\}} ,$$

and $(a(j))_{j \geq 0}$ take values as in the optimal solution given by Theorem 33. Importantly:

$$a(0) \approx 0.2403 > 0.24 .$$

Further recall a remark after Theorem 36 that the guarantee in Eqn. (11) holds almost trivially for three or more consecutive subsequences and for two subsequences whose total lengths are more than three. Hence, the proof will first handle the remaining cases before substantiating the remark.

One Subsequence. In this case the second term in Eqn. (11) disappears, so we will upper bound the probability by $p(k)$ alone. Suppose that t_1, t_2, \dots, t_k are the rounds in the consecutive subsequence. Consider the number of them that sample $e^{t_i} = e$. If there are at least two, the flag τ_e must be 1 in at least one of them, and by definition Algorithm 8 selects e there.

If exactly one out of the k rounds samples $e^{t_i} = e$, which happens with probability $k2^{-k}$ over the randomness of e^{t_i} 's, Lemma 43 indicates that the probability of never selecting e in the $k - 1$ rounds with $e^{t_i} \neq e$, *conditioned on the realized e^{t_i} 's*, is at most $2^{-\min\{k-1, \lceil \frac{k+1}{2} \rceil\}}$. Further, over the randomness of the initial flag τ_e , the round with $e^{t_i} = e$ selects e with probability half, independent to the realization of the other $k - 1$ rounds. In sum, this case happens with probability at most:

$$k2^{-k-\min\{k, \lceil \frac{k+3}{2} \rceil\}} .$$

Finally, if none of the k rounds samples $e^{t_i} = e$, which happens with probability 2^{-k} over the randomness of e^{t_i} 's, Lemma 43 indicates that the probability of never selecting e in these round is at most $2^{-\min\{k, \lceil \frac{k+2}{2} \rceil\}}$. In sum, this case happens with probability at most:

$$2^{-k-\min\{k, \lceil \frac{k+2}{2} \rceil\}} .$$

Summing the probability bounds in the last two cases gives exactly $p(k)$.

Two Subsequences, Two Rounds. By Lemma 43, the probability of not selecting element e in the two rounds is at most $\frac{1}{4}$. It then follows by $p(2) = \frac{3}{16}$ and by $a(0) > 0.24$.

Two Subsequences, Three or More Rounds. At least one subsequence must have at least two rounds. If two neighboring rounds in the same subsequence both sample $e^t = e$, which happens with probability $\frac{1}{4}$, flag τ_e must be 1 in one of them and Algorithm 8 selects e in that round. Otherwise, Lemma 43 indicates that the probability of never selecting e in these three or more rounds is at most $\frac{1}{8}$. In total, the probability is at most $(1 - \frac{1}{4})\frac{1}{8} = \frac{3}{32}$. It then follows by $a(0) > 0.24$.

Three or More Subsequences. Lemma 43 indicates that the probability of never selecting e in these three or more rounds is at most $\frac{1}{8}$. It then follows by $a(0) > 0.24$, since the second term in Eqn. (11) sum to at least $a(0)$. \square

6.5 Multi-way Semi-OCS and Unweighted and Vertex-weighted Online Bipartite Matching

This subsection introduces an algorithm BALANCE-OCS that combines an unbounded variant of the BALANCE algorithm [25, 28] and a multi-way semi-OCS. The former assigns one unit of masses to the offline neighbors of each online vertex. The latter then selects one of them to which the online vertex will match. We shall analyze its competitive ratio in the unweighted and vertex-weighted online bipartite matching problems.

Algorithm 9 BALANCE-OCS

State variables: (for each offline vertex u)

- Total mass y_u allocated to offline vertex u far; initially, $y_u = 0$.

On the arrival of an online vertex $v \in R$:

1. Find threshold $\theta \in [0, \infty)$ that satisfies Eqn. (16).
 2. For each neighbor u , let $x_u^v = (y_u(\theta) - y_u)^+$ be its mass in this round.
 3. Match v to the neighbor that the multi-way semi-OCS selects, with \mathbf{x}^v as the mass vector in this round.
-

BALANCE is parameterized by a non-increasing discounting function $b : [0, +\infty) \rightarrow [0, 1]$. For each online vertex v , it continuously assigns one unit of masses to v 's neighbors,¹⁰ prioritizing the ones with the largest discounted weight $w_u b(y_u)$ where y_u denotes the total mass assigned to u so far. To describe it as an algorithm instead of a continuous process, for any offline vertex u and any threshold marginal utility $\theta \geq 0$, define:

$$y_u(\theta) = b^{-1}\left(\frac{\theta}{w_u}\right). \quad (15)$$

We will explain shortly how to interpret the algorithm if b is not continuous or is not strictly increasing, i.e., when the inverse function is not well-defined. In any case, the discount function b used by our algorithm is continuous and strictly monotone.

Define $y_u(\theta) = 0$ if $w_u b(0) < \theta$, e.g., when $\theta > w_u$. Let z^+ denote $\max\{z, 0\}$. Then, we may equivalently interpret the BALANCE algorithm as choosing a threshold θ such that:

$$\sum_{u:(u,v) \in E} (y_u(\theta) - y_u)^+ = 1, \quad (16)$$

and then assigning mass $(y_u(\theta) - y_u)^+$ to each vertex u .

For a discount function b whose inverse is not well-defined, let $y_u^-(\theta) = \sup\{y \geq 0 : w_u b(y) < \theta\}$ and $y_u^+(\theta) = \inf\{y \geq 0 : w_u b(y) > \theta\}$. The BALANCE algorithm chooses an appropriate threshold θ and choose $y_u(\theta) \in [y_u^-(\theta), y_u^+(\theta)]$ to satisfy Eqn. (16).

The original BALANCE algorithm cannot assign more than one unit of total mass to any offline vertex, which introduces boundary considerations that complicate the above description. In our setting, however, the masses are merely input of the multi-way semi-OCS, and therefore the total mass of an offline vertex could be arbitrarily large.

Theorem 44. *Suppose that $p : [0, \infty) \rightarrow [0, 1]$ is decreasing and differentiable, and $p(0) = 1$. Then, unbounded BALANCE with a p -multi-way semi-OCS (Algorithm 9) is Γ -competitive for unweighted and vertex-weighted online bipartite matching, where the competitive ratio Γ and the corresponding*

¹⁰In the context of fractional online matching, BALANCE fractionally matches v to its neighbors.

discount function b are from an optimal solution of the following continuous LP:

$$\text{maximize } \Gamma \quad (\text{BALANCE LP})$$

$$\text{subject to } a(y) + b(y) \leq -p'(y) \quad \forall y \geq 0 \quad (17)$$

$$\int_0^y a(z)dz + b(y) \geq \Gamma \quad \forall y \geq 0 \quad (18)$$

$$b(y') \leq b(y) \quad \forall y' \geq y \quad (19)$$

$$a(y), b(y) \geq 0 \quad \forall y \geq 0$$

Proof. By the guarantee of p -multi-way semi-OCS, each offline vertex u is matched by unbounded BALANCE with probability at least:

$$1 - p(y_u) .$$

Therefore, the expected total weight of the matched vertices is at least:

$$\text{ALG} \stackrel{\text{def}}{=} \sum_{u \in L} w_u (1 - p(y_u)) .$$

Similar to the competitive analysis of the two-choice algorithm (Algorithm 7), for every online vertex v we will distribute the increase of ALG between vertex v and its offline neighbors. Let α_u and β_v be the distributed gain of any offline vertex u and any online vertex v . They are initially zero. In the round of online vertex v , for each offline neighbor u , increase α_u by:

$$w_u \int_{y_u}^{y_u + x_u^v} a(z)dz ,$$

where y_u is the value right before v arrives.

Further, let β_v be:

$$\beta_v \stackrel{\text{def}}{=} \sum_{u: (u,v) \in E} w_u \int_{y_u}^{y_u + x_u^v} b(z)dz .$$

Feasibility of the Charging Rule. The total gain distributed above is upper bounded by the increase in ALG because:

$$\begin{aligned} \sum_{u: (u,v) \in E} w_u \int_{y_u}^{y_u + x_u^v} (a(z) + b(z))dz &\leq \sum_{u: (u,v) \in E} w_u \int_{y_u}^{y_u + x_u^v} -p'(z)dz && (\text{Eqn. (17)}) \\ &= \sum_{u: (u,v) \in E} w_u \left((1 - p(y_u + x_u^v)) - (1 - p(y_u)) \right) . \end{aligned}$$

Invariant of Offline Gain. By definition, for any offline vertex u :

$$\alpha_u = w_u \int_0^{y_u} a(z)dz .$$

Invariant of Online Gain. For any online vertex v , since the algorithm prefers neighbors with larger $w_u b(y_u)$ and assigns one unit of mass, we get that for any v 's neighbor u :

$$\beta_v \geq w_u b(y_u) .$$

This holds for the final value of y_u because y_u only increase over time and the discount function b is non-increasing.

Γ -Approximate Equilibrium. The gains satisfy an approximate equilibrium condition in the sense that for any edge (u, v) , the total gain of u and v is at least Γ times the vertex-weight w_u :

$$\begin{aligned} \alpha_u + \beta_v &\geq w_u \int_0^{y_u} a(z) dz + w_u b(y_u) && \text{(Invariants)} \\ &\geq \Gamma w_u . && \text{(Eqn. (18))} \end{aligned}$$

Then, for any optimal matching $M \subseteq E$, unbounded BALANCE is Γ -competitive because:

$$\begin{aligned} \text{ALG} &\geq \sum_{u \in L} \alpha_u + \sum_{v \in R} \beta_v \\ &\geq \sum_{(u,v) \in M} (\alpha_u + \beta_v) \\ &\geq \Gamma \sum_{(u,v) \in M} w_u . \end{aligned}$$

□

The LP in Theorem 44 has continuously many variables and constraints. Fortunately, we have an explicit optimal solution for most natural p functions.

Theorem 45. *Suppose that function $p : [0, \infty) \rightarrow [0, 1]$ is decreasing, convex, and differentiable, and $p(0) = 1$. Then, an optimal solution to the BALANCE LP is:*

$$\begin{aligned} \Gamma &= \int_0^\infty e^{-z} (1 - p(z)) dz ; \\ b(y) &= -e^y \int_y^\infty p'(z) e^{-z} dz && \forall y \geq 0 ; \\ a(y) &= -p'(y) - b(y) && \forall y \geq 0 . \end{aligned}$$

The proof of this theorem is deferred to Appendix D.2.

Corollary 46. *Unbounded BALANCE with the multi-way semi-OCS from Theorem 6 in Section 4 is at least 0.593-competitive for unweighted and vertex-weighted online bipartite matching.*

Acknowledgment

We thank Zhihao Gavin Tang and Hu Fu for helpful discussions on online contention resolution schemes.

References

- [1] Marek Adamczyk and Michał Włodarczyk. Random order contention resolution schemes. In *59th Annual IEEE Symposium on Foundations of Computer Science*, pages 790–801. IEEE, 2018.
- [2] Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1253–1264. SIAM, 2011.

- [3] Kenneth S Alexander. A counterexample to a correlation inequality in finite sampling. *The Annals of Statistics*, pages 436–439, 1989.
- [4] Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph (Seffi) Naor. A polylogarithmic-competitive algorithm for the k-server problem. *Journal of the ACM*, 62(5):1–49, 2015.
- [5] Guy Blanc and Moses Charikar. Multiway online correlated selection. In *62nd Annual IEEE Symposium on Foundations of Computer Science*, 2021.
- [6] Niv Buchbinder and Joseph (Seffi) Naor. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research*, 34(2):270–286, 2009.
- [7] Niv Buchbinder, Joseph, Naor, and David Wajc. A randomness threshold for online bipartite matching, via lossless online rounding. *arXiv preprint arXiv:2106.04863*, 2021.
- [8] T-H Hubert Chan, Zhiyi Huang, Shaofeng H-C Jiang, Ning Kang, and Zhihao Gavin Tang. Online submodular maximization with free disposal. *ACM Transactions on Algorithms*, 14(4): 1–29, 2018.
- [9] Ilan Reuven Cohen and David Wajc. Randomized online matching in regular graphs. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 960–979. SIAM, 2018.
- [10] Ilan Reuven Cohen, Binghui Peng, and David Wajc. Tight bounds for online edge coloring. In *60th Annual IEEE Symposium on Foundations of Computer Science*, pages 1–25. IEEE, 2019.
- [11] Nikhil R Devanur, Kamal Jain, and Robert D Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 101–107. SIAM, 2013.
- [12] Nikhil R Devanur, Zhiyi Huang, Nitish Korula, Vahab S Mirrokni, and Qiqi Yan. Whole-page optimization and submodular welfare maximization with online bidders. *ACM Transactions on Economics and Computation*, 4(3):1–20, 2016.
- [13] Shaddin Dughmi. A truthful randomized mechanism for combinatorial public projects via convex optimization. In *12th ACM conference on Electronic Commerce*, pages 263–272, 2011.
- [14] Shaddin Dughmi. The outer limits of contention resolution on matroids and connections to the secretary problem. In *47th International Colloquium on Automata, Languages, and Programming*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [15] Shaddin Dughmi. Matroid secretary is equivalent to contention resolution. *arXiv preprint arXiv:2103.04205*, 2021.
- [16] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. Optimal mechanisms for combinatorial auctions and combinatorial public projects via convex rounding. *Journal of the ACM*, 63(4): 1–33, 2016.
- [17] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Online stochastic max-weight matching: prophet inequality for vertex and edge arrival models. In *21st ACM Conference on Economics and Computation*, pages 769–787, 2020.
- [18] Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. Edge-weighted online bipartite matching. In *61st Annual IEEE Symposium on Foundations of Computer Science*, pages 412–423, 2020.

- [19] Jon Feldman, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pál. Online ad assignment with free disposal. In *5th International Workshop on Internet and Network Economics*, pages 374–385. Springer, 2009.
- [20] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1014–1033. SIAM, 2016.
- [21] Hu Fu, Zhihao Gavin Tang, Hongxun Wu, Jinzhao Wu, and Qianfan Zhang. Random order vertex arrival contention resolution schemes for matching, with applications. In *48th International Colloquium on Automata, Languages, and Programming*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [22] Buddhima Gamlath, Sagar Kale, and Ola Svensson. Beating greedy for stochastic bipartite matching. In *30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2841–2854. SIAM, 2019.
- [23] Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online matching with general arrivals. In *60th Annual IEEE Symposium on Foundations of Computer Science*, pages 26–37. IEEE, 2019.
- [24] Zhiyi Huang, Qiankun Zhang, and Yuhao Zhang. Adwords in a panorama. In *61st Annual IEEE Symposium on Foundations of Computer Science*, pages 1416–1426, 2020.
- [25] Bala Kalyanasundaram and Kirk R Pruhs. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science*, 233(1-2):319–325, 2000.
- [26] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for online bipartite matching. In *22nd Annual ACM Symposium on Theory of Computing*, pages 352–358, 1990.
- [27] Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2018.
- [28] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):22–es, 2007.
- [29] Aranyak Mehta et al. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.
- [30] Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *22nd ACM conference on Electronic Commerce*, 2021.
- [31] Amin Saberi and David Wajc. The Greedy Algorithm Is not Optimal for On-Line Edge Coloring. In *48th International Colloquium on Automata, Languages, and Programming*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [32] Yongho Shin and Hyung-Chan An. Making three out of two: Three-way online correlated selection. *arXiv preprint arXiv:2107.02605*, 2021.

A Missing Proofs in Section 3

A.1 Positive Correlation in 3-Way Sampling without Replacement

Consider the following counter-example which shows that there could be positive correlation in 3-way (unweighted) sampling without replacement. The elements are integers from 1 to 9. It has 7 rounds:

$$\mathcal{E}^1 = \mathcal{E}^2 = \{1, 4, 5\}, \mathcal{E}^3 = \mathcal{E}^4 = \{2, 6, 7\}, \mathcal{E}^5 = \mathcal{E}^6 = \{3, 8, 9\}, \mathcal{E}^7 = \{1, 2, 3\}.$$

Recall that \mathcal{U}^t denotes the subset of unselected elements after round t , and thus $e \in \mathcal{U}^t$ denotes the event that element e remains unselected after round t . Further $\mathcal{U} = \mathcal{U}^7$ denotes the subset of unselected elements at the end. On the one hand:

$$\Pr[1, 2 \in \mathcal{U}] = \Pr[s^1, s^2 \neq 1] \Pr[s^3, s^4 \neq 2] \Pr[s^5, s^6 \neq 3] \Pr[s^7 = 3 \mid 1, 2, 3 \in \mathcal{U}^6] = \left(\frac{1}{3}\right)^4 = \frac{1}{81}.$$

On the other hand:

$$\begin{aligned} \Pr[1 \in \mathcal{U}] &= \Pr[s^1, s^2 \neq 1] \left(\Pr[s^3, s^4 \neq 2] \Pr[s^5, s^6 \neq 3] \Pr[s^7 \neq 1 \mid 1, 2, 3 \in \mathcal{U}^6] \right. \\ &\quad + \Pr[s^3, s^4 \neq 2] \Pr[3 \in \{s^5, s^6\}] \Pr[s^7 \neq 1 \mid 1, 2 \in \mathcal{U}^6, 3 \notin \mathcal{U}^6] \\ &\quad \left. + \Pr[2 \in \{s^3, s^4\}] \Pr[s^5, s^6 \neq 3] \Pr[s^7 \neq 1 \mid 1, 3 \in \mathcal{U}^6, 2 \notin \mathcal{U}^6] \right) \\ &= \frac{1}{3} \left(\frac{1}{3} \cdot \frac{1}{3} \cdot \frac{2}{3} + \frac{1}{3} \cdot \frac{2}{3} \cdot \frac{1}{2} + \frac{2}{3} \cdot \frac{1}{3} \cdot \frac{1}{2} \right) = \frac{8}{81}. \end{aligned}$$

Further by symmetry:

$$\Pr[2 \in \mathcal{U}] = \frac{8}{81}.$$

Therefore:

$$\frac{\Pr[1, 2 \in \mathcal{U}]}{\Pr[1 \in \mathcal{U}] \Pr[2 \in \mathcal{U}]} = \frac{81}{64} > 1.$$

A.2 Proof of Theorem 4

We shall construct a distribution of instances and prove the desired probability bound holds on average. By considering an appropriate distribution of instance, we ensure that the randomness of the instance dictates the selection result.

Consider elements $\mathcal{E} = \{e_1^0 = 1, e_2^0 = 2, \dots, e_{2^k}^0 = 2^k\}$. Further for i from 1 to k , recursively define e_j^i to be either e_{2j-1}^{i-1} or e_{2j}^{i-1} uniformly at random. The instance has pairs $\{e_{2j-1}^i, e_{2j}^i\}$ for all $0 \leq i \leq k-1$ and all $1 \leq j \leq 2^{k-i}$, in ascending order of i ; the order with respect to j for any fixed i is unimportant yet for concreteness we define it to be in ascending order as well.

In other words, this is a knockout-tournament-like instance. First partition the 2^k elements into 2^{k-1} pairs in lexicographical order. We shall refer to these pairs as the first stage of the instance. Then, a randomly chosen “winner” from each pair advances to the next stage. Repeat this process until we have the final “winner”, denoted as e_1^k by the construction above. We shall refer to the pairs defined with respect to elements e_j^{i-1} , $1 \leq j \leq 2^{k-i+1}$, as the i -th stage of the instance. For example, consider 2^3 elements $\{1, 2, \dots, 8\}$. A possible realization of the random instance proceeds as $\{1, 2\}, \{3, 4\}, \{5, 6\}, \{7, 8\}, \{1, 3\}, \{5, 8\}, \{3, 5\}$. See Figure 1 for an illustration.

Next we show that with probability at least 2^{-2^k+1} , the final “winner” e_1^k is never selected by the algorithm, despite its k appearances. Here the probability space is over both the randomness of the algorithm and that of the instance. In fact, we shall inductively prove the following stronger invariant; the above claim is the special case when $i = k$.




Figure 5: Knockout tournament-like input sequence.

Invariant: For any $1 \leq i \leq k$, after processing the pairs involving elements e_j^{i-1} 's, i.e., after the first i stages of the tournament, all elements e_j^i , $1 \leq j \leq 2^{k-i}$, in the next stage remain unselected with probability at least:

$$2^{-2^k+2^{k-i}}.$$

The base case when $i = 0$ is vacuously true.

Next suppose that the invariant holds for $i - 1$, and consider the case of i . Below are a set of sufficient conditions under which all elements e_j^i are unselected after the first i stages:

1. for any $1 \leq j \leq 2^{k-i+1}$, e_j^{i-1} is unselected after the first $i - 1$ stages;
2. for any $1 \leq j \leq 2^{k-i}$, $e_j^i \in \{e_{2j-1}^{i-1}, e_{2j}^{i-1}\}$ is *not* the one selected by the algorithm for the pair.

By the construction of the random instance, the two events are independent. The first holds with probability at least $2^{-2^k+2^{k-i+1}}$ by the inductively hypothesis. The second holds with probability $2^{-2^{k-i}}$ by the construction of the random instance. Hence, the probability that all elements e_j^i remain unselected after the first i stages is at least:

$$2^{-2^k+2^{k-i+1}} \cdot 2^{-2^{k-i}} = 2^{-2^k+2^{k-i}}.$$

B Missing Proofs in Section 4

B.1 Proof of Lemma 7

Both sides equal 1 when $x = 0$. The rest of the proof considers $0 < x < 1$. Let $t = \frac{y}{x}$. We have:

$$\begin{aligned} \frac{x}{1-x}w(y) + 1 - \frac{w(y+x)}{w(y)} &= \frac{x}{1-x}w(tx) + 1 - \frac{w(tx+x)}{w(tx)} \\ &= \int_0^x \left(B(x', t)w(tx') - A(x', t)\frac{w(tx'+x')}{w(tx')} \right) dx' \\ &= \int_0^x B(x', t)w(tx') \left(1 - \frac{A(x', t)w(tx'+x')}{B(x', t)w(tx')^2} \right) dx' \\ &= \int_0^x B(x', t)w(tx') \int_0^{x'} \frac{C(x'', t)w(tx''+x'')}{B(x'', t)^2w(tx'')^2} dx'' dx', \end{aligned}$$

where:

$$Q(x) = \frac{1}{w(x)} \frac{dw(x)}{dx} = 1 + x + 3cx^2 ,$$

$$A(x, t) = (t + 1)Q(tx + x) - tQ(tx) ,$$

$$B(x, t) = \frac{1}{(1 - x)^2} + \frac{tx}{1 - x} Q(tx) ,$$

$$C(x, t) = A(x, t) \frac{dB}{dx}(x, t) - B(x, t) \frac{dA}{dx}(x, t) - A(x, t)B(x, t) ((t + 1)Q(tx + x) - 2tQ(tx)) .$$

To prove the lemma, i.e.:

$$\frac{w(y + x)}{w(y)} \leq \frac{x}{1 - x} w(y) + 1 ,$$

for any $0 < x < 1$ and $y \geq 0$, we only need to prove $C(x, t) \geq 0$ for any $0 < x < 1$ and any $t \geq 0$. It is equivalent to show that for any $x, t \geq 0$:

$$\frac{(x + 1)^6}{x} C \left(\frac{x}{1 + x}, t \right) \geq 0 .$$

By computation:

$$\frac{(x + 1)^6}{x} C \left(\frac{x}{1 + x}, t \right) = \sum_{i=0}^8 x^i P_i(t) ,$$

where

$$\begin{aligned}
P_0(t) &= -18ct^2 - 18ct - 6c \\
&\quad + 6t^2 + 2, \\
P_1(t) &= 12ct^3 - 126ct^2 - 126ct - 42c \\
&\quad + 8t^3 + 36t^2 + 17, \\
P_2(t) &= 42ct^4 + 54ct^3 - 408ct^2 - 390ct - 126c \\
&\quad + 5t^4 + 38t^3 + 88t^2 + 5t + 64, \\
P_3(t) &= 54c^2t^5 - 27c^2t^4 - 144c^2t^3 - 135c^2t^2 - 54c^2t - 9c^2 \\
&\quad + 36ct^5 + 156ct^4 + 36ct^3 - 780ct^2 - 672ct - 204c \\
&\quad + 2t^5 + 17t^4 + 72t^3 + 115t^2 + 30t + 139, \\
P_4(t) &= 63c^2t^6 + 135c^2t^5 - 306c^2t^4 - 729c^2t^3 - 594c^2t^2 - 225c^2t - 36c^2 \\
&\quad + 21ct^6 + 75ct^5 + 183ct^4 - 117ct^3 - 882ct^2 - 654ct - 180c \\
&\quad + 6t^5 + 21t^4 + 72t^3 + 91t^2 + 74t + 190, \\
P_5(t) &= 72c^2t^7 - 144c^2t^5 - 783c^2t^4 - 1242c^2t^3 - 927c^2t^2 - 342c^2t - 54c^2 \\
&\quad + 42ct^6 + 42ct^5 + 96ct^4 - 174ct^3 - 510ct^2 - 300ct - 66c \\
&\quad + 6t^5 + 11t^4 + 44t^3 + 49t^2 + 96t + 167, \\
P_6(t) &= 81c^3t^8 - 162c^3t^7 - 459c^3t^6 - 405c^3t^5 - 162c^3t^4 - 27c^3t^3 \\
&\quad + 72c^2t^7 - 63c^2t^6 - 171c^2t^5 - 531c^2t^4 - 801c^2t^3 - 603c^2t^2 - 225c^2t - 36c^2 \\
&\quad + 21ct^6 + 3ct^5 + 57ct^4 - 51ct^3 - 66ct^2 + 12ct + 18c \\
&\quad + 2t^5 + 2t^4 + 18t^3 + 19t^2 + 69t + 92, \\
P_7(t) &= 54c^2t^5 - 27c^2t^4 - 144c^2t^3 - 135c^2t^2 - 54c^2t - 9c^2 \\
&\quad + 30ct^4 + 12ct^3 + 60ct^2 + 66ct + 24c \\
&\quad + 4t^3 + 4t^2 + 26t + 29, \\
P_8(t) &= 18ct^2 + 18ct + 6c \\
&\quad + 4t + 4.
\end{aligned}$$

When $c = \frac{4-2\sqrt{3}}{3}$, we can verify, using numerical computation software, that $P_i(t)$ has positive leading coefficient and no non-negative real roots for $i = 1, 2, \dots, 8$, and:

$$P_0(t) = (4\sqrt{3} - 6) (1 - \sqrt{3}t)^2 \geq 0.$$

Therefore, for any $x, t \geq 0$:

$$\frac{(x+1)^6}{x} C\left(\frac{x}{1+x}, t\right) \geq 0.$$

So the lemma holds.

B.2 Proof of Lemma 8

We will use the following lemma which follows by the definition of the weight function w in Eqn. (1).

Lemma 47. *Function $\log \circ w$ is convex in $[0, \infty)$.*

If $\sum_{i=1}^k x_i \geq 1$, the left-hand-side is zero so the inequality holds trivially. If $\sum_{i=1}^k x_i = 0$, both sides are equal to 1 so the inequality also holds. The rest of the proof consider $0 < \sum_{i=1}^k x_i < 1$.

Fix any non-negative $y_1, y_2, \dots, y_k \geq 0$. Define function $f : [0, +\infty)^k \rightarrow (-\infty, +\infty)$ over $\mathbf{x} = (x_1, x_2, \dots, x_k)$ be the difference between the logarithms of the two sides, i.e.:

$$f(\mathbf{x}) = \log \frac{1 - \sum_{i=1}^k x_i}{\sum_{i=1}^k x_i w(y_i) + 1 - \sum_{i=1}^k x_i} - \sum_{i=1}^k \log \frac{w(y_i)}{w(y_i + x_i)} .$$

We first argue that f is convex over a simplex, using the log-convexity of the weight function w (Lemma 47). Consider any $\mathbf{x} = (x_1, \dots, x_k), \mathbf{x}' = (x'_1, \dots, x'_k) \in [0, +\infty)^k$ such that:

$$\sum_{i=1}^k x_i = \sum_{i=1}^k x'_i .$$

Then, for any $t \in [0, 1]$ and the linear $\mathbf{x}'' = t\mathbf{x} + (1-t)\mathbf{x}'$, we have:

$$\begin{aligned} tf(\mathbf{x}) + (1-t)f(\mathbf{x}') &= \log \left(1 - \sum_{i=1}^k x_i \right) - t \log \left(\sum_{i=1}^k x_i w(y_i) + 1 - \sum_{i=1}^k x_i \right) \\ &\quad - (1-t) \log \left(\sum_{i=1}^k x'_i w(y_i) + 1 - \sum_{i=1}^k x_i \right) \\ &\quad - \sum_{i=1}^k \log w(y_i) + \sum_{i=1}^k (t \log w(y_i + x_i) + (1-t) \log w(y_i + x'_i)) . \end{aligned}$$

By the log-convexity of w , and by Jensen's inequality on the second and third terms and on the last two terms, this is at least:

$$\log \left(1 - \sum_{i=1}^k x_i \right) - \log \left(\sum_{i=1}^k x''_i w(y_i) + 1 - \sum_{i=1}^k x_i \right) - \sum_{i=1}^k \log w(y_i) + \sum_{i=1}^k \log w(y_i + x''_i) = f(\mathbf{x}'') .$$

Therefore, we have:

$$\begin{aligned} f(\mathbf{x}) &\leq \frac{1}{\sum_{i=1}^k x_i} \sum_{i=1}^k x_i f \left(0, \dots, 0, \underbrace{\sum_{j=1}^k x_j}_{i\text{-th entry}}, 0, \dots, 0 \right) \quad (\text{convexity of } f \text{ on simplex}) \\ &= \frac{1}{\sum_{i=1}^k x_i} \sum_{i=1}^k x_i \left(\log \frac{w(y_i + \sum_{j=1}^k x_j)}{w(y_i)} - \log \left(\frac{\sum_{j=1}^k x_j}{1 - \sum_{j=1}^k x_j} w(y_i) + 1 \right) \right) \\ &\leq 0 . \end{aligned} \tag{Lemma 7}$$

B.3 A Weaker Version of Lemma 7

The following lemma is a weak version of Lemma 7. We provide it and a proof that does not involve computer-aided numerical verification.

Lemma 48. For any $y \geq 0$ and any $0 < \delta < 1$:

$$\exp\left((y + \delta) + \frac{(y + \delta)^2}{2} + \frac{(y + \delta)^3}{6}\right) \leq \exp\left(y + \frac{y^2}{2} + \frac{y^3}{6}\right) + \frac{\delta}{1 - \delta} \exp\left(2y + y^2 + \frac{y^3}{3}\right).$$

Proof. The inequality holds with equality when $\delta = 0$. Hence, it suffices to consider the partials of both sides with respect to δ , and to show that the partial of the right-hand-side is larger, i.e.:

$$\left(1 + (y + \delta) + \frac{(y + \delta)^2}{2}\right) \exp\left((y + \delta) + \frac{(y + \delta)^2}{2} + \frac{(y + \delta)^3}{6}\right) \leq \frac{1}{(1 - \delta)^2} \exp\left(2y + y^2 + \frac{y^3}{3}\right).$$

Taking logarithm of both sides and rearrange terms, it is equivalent to:

$$\ln\left(1 + (y + \delta) + \frac{(y + \delta)^2}{2}\right) + 2 \ln(1 - \delta) \leq (y - \delta) + \frac{(y - \delta)^2}{2} + \frac{(y - \delta)^3}{6} - \delta^2 - y\delta^2.$$

Since $\ln(1 - \delta) \leq -\delta - \frac{\delta^2}{2} - \frac{\delta^3}{3} - \frac{\delta^4}{4}$ for all $0 < \delta < 1$, after an rearrangement of terms it suffices to show that:

$$\ln\left(1 + (y + \delta) + \frac{(y + \delta)^2}{2}\right) \leq (y + \delta) + \frac{(y - \delta)^2}{2} + \frac{y^3}{3} + \frac{2\delta^3}{3} - \frac{(y + \delta)^3}{6} + \frac{\delta^4}{2}.$$

By the inequality of arithmetic and geometric means:

$$\frac{(y - \delta)^2}{2} + \frac{\delta^4}{2} \geq |y - \delta|\delta^2.$$

Hence, we arrive at the final inequality that is sufficient for establishing the inequality regarding the partial derivatives with respect to y and thus, the correct of the lemma:

$$\ln\left(1 + (y + \delta) + \frac{(y + \delta)^2}{2}\right) \leq (y + \delta) + |y - \delta|\delta^2 + \frac{y^3}{3} + \frac{2\delta^3}{3} - \frac{(y + \delta)^3}{6}. \quad (20)$$

Roadmap of Proving Eqn. (20). The natural next step is to upper bound $\ln\left(1 + x + \frac{x^2}{2}\right)$ by a polynomial of x to transform the left-hand-side of the inequality a polynomial over y and δ , just like the right-hand-side. The Taylor series $\ln\left(1 + x + \frac{x^2}{2}\right) = x - \frac{x^3}{6} + \frac{x^4}{8} - \frac{x^5}{20} + O(x^7)$ suggests natural upper bounds such as $\ln\left(1 + x + \frac{x^2}{2}\right) \leq x$ and $\ln\left(1 + x + \frac{x^2}{2}\right) \leq x - \frac{x^3}{6} + \frac{x^4}{8}$. Unfortunately, neither of these bounds proves Eqn. (20) for all $y \geq 0$ and all $0 < \delta < 1$. In particular, the former fails when $\delta = y$, and the latter leaves a degree-4 term that cannot be bounded by the right-hand-side of Eqn. (20) for sufficiently large y . Instead, we shall consider three polynomial upper bounds of $\ln\left(1 + x + \frac{x^2}{2}\right)$ depending on the range of x and together they cover all the cases. These upper bounds are from the next lemma, whose proof is deferred to the end of the subsection.

Lemma 49. The function:

$$\frac{x - \ln\left(1 + x + \frac{x^2}{2}\right)}{x^3}.$$

is decreasing for $x > 0$

Case 1: Use $\ln(1+x+\frac{x^2}{2}) \leq x$ for $x \geq 0$, when $y \geq (1+\sqrt[3]{2})\delta$. In other words, by inequality $\ln(1+(y+\delta)+\frac{(y+\delta)^2}{2}) \leq y+\delta$ and $y > \delta$, Eqn. (20) reduces to:

$$\frac{y^3}{6} - \frac{y^2\delta}{2} + \frac{y\delta^2}{2} - \frac{\delta^3}{2} \geq 0,$$

or equivalently:

$$\left(\frac{y}{\delta}\right)^3 - 3\left(\frac{y}{\delta}\right)^2 + 3\frac{y}{\delta} - 3 = \left(\frac{y}{\delta} - 1\right)^3 - 2 \geq 0.$$

We remark that this approach can also prove Eqn. (20) for $y \leq c\delta$ for $c \approx 0.83$. We skip it since it is covered by the other cases.

Case 2: Use $\ln(1+x+\frac{x^2}{2}) \leq x - \frac{x^3}{24}$ for $0 \leq x \leq \frac{7}{3}$, if $y+\delta \leq \frac{7}{3}$. The stated inequality follows by Lemma 49 and that for $x = \frac{7}{3}$:

$$\frac{x - \ln(1+x+\frac{x^2}{2})}{x^3} \approx 0.0419 > \frac{1}{24}.$$

By $\ln(1+(y+\delta)+\frac{(y+\delta)^2}{2}) \leq y+\delta - \frac{(y+\delta)^3}{24}$, Eqn. (20) reduces to:

$$|y-\delta|\delta^2 + \frac{y^3}{3} + \frac{2\delta^3}{3} - \frac{(y+\delta)^3}{8} \geq 0.$$

If $y \geq \delta$, the left-hand-side equals:

$$\frac{5}{24}y^3 - \frac{3}{8}y^2\delta + \frac{5}{8}y\delta^2 - \frac{11}{24}\delta^3 = \frac{1}{24}(y-\delta)(5y^2 - 4y\delta + 11\delta^2) \geq 0.$$

If $y < \delta$, the left-hand-side equals:

$$\frac{5}{24}y^3 - \frac{3}{8}y^2\delta - \frac{11}{8}y\delta^2 + \frac{37}{24}\delta^3 = \frac{1}{24}(\delta-y)(37\delta^2 + 4\delta y - 5y^2) \geq 0.$$

Case 3: Use $\ln(1+x+\frac{x^2}{2}) \leq x - \frac{x^3}{36}$ for $0 \leq x \leq 2 + \sqrt[3]{2}$, if $\frac{7}{3} \leq y+\delta \leq 2 + \sqrt[3]{2}$. The stated inequality follows by Lemma 49 and that for $x = 2 + \sqrt[3]{2}$:

$$\frac{x - \ln(1+x+\frac{x^2}{2})}{x^3} \approx 0.028 > \frac{1}{36}.$$

The argument for this case combines the assumption of $y+\delta > \frac{7}{3}$ and that $0 < \delta < 1$ to derive:

$$y > \frac{4}{3}\delta.$$

By $\ln(1+(y+\delta)+\frac{(y+\delta)^2}{2}) \leq y+\delta - \frac{(y+\delta)^3}{36}$ and by $y > \delta$, Eqn. (20) reduces to:

$$(y-\delta)\delta^2 + \frac{y^3}{3} + \frac{2\delta^3}{3} - \frac{5(y+\delta)^3}{36} \geq 0.$$

Rearranging terms, the left-hand-side equals:

$$\frac{7}{36}y^3 - \frac{5}{12}y^2\delta + \frac{7}{12}y\delta^2 - \frac{17}{36}\delta^3 = \frac{1}{144}(5y-6\delta)^2y + \frac{1}{48}(y^2-\delta^2)y + \frac{17}{48}\delta^2(y-\frac{4}{3}\delta) \geq 0.$$

The three cases cover all $y \geq 0$ and $0 < \delta < 1$ since the last two cases prove the lemma for any y, δ that satisfies $y+\delta \leq 2 + \sqrt[3]{2}$. For any y, δ with $y+\delta > 2 + \sqrt[3]{2}$ (and $0 < \delta < 1$) must satisfy $y > (1 + \sqrt[3]{2})\delta$ and therefore is covered by the first case. \square

Proof of Lemma 49. We shall prove that the derivative is non-positive, i.e.:

$$\frac{d}{dx} \frac{x - \ln(1 + x + \frac{x^2}{2})}{x^3} = \frac{3}{x^4} \left(\ln(1 + x + \frac{x^2}{2}) - x + \frac{x^3}{6(1 + x + \frac{x^2}{2})} \right) \leq 0.$$

Equivalently, we need to show that:

$$\ln(1 + x + \frac{x^2}{2}) - x + \frac{x^3}{6(1 + x + \frac{x^2}{2})} \leq 0.$$

Since this holds with equality at $x = 0$, it suffices to prove that its derivative is non-positive. This follows by:

$$\frac{d}{dx} \left(\ln(1 + x + \frac{x^2}{2}) - x + \frac{x^3}{6(1 + x + \frac{x^2}{2})} \right) = -\frac{x^3(1+x)}{6(1+x+\frac{x^2}{2})^2} \leq 0.$$

□

C Missing Proofs in Section 5

C.1 Proof of Lemma 15

It suffices to show that the pseudo-paths are pairwise disjoint. Consider an arc (p, c) and a pseudo-path P that involves it. It is clear that (p, c) can be adjacent with at most two arcs: the other in-arc of c (if exists) and the other out-arc of p (if exists). As P is maximal, if c has another in-arc in $G^{\text{ex-ante}}$, it should be adjacent to (p, c) in P . Similarly, if p has another out-arc (p, c') in $G^{\text{ex-ante}}$ such that rounds p, c, c' have a common element, it should also be adjacent to (p, c) in P . Therefore, for any arc (p, c) , the set of its adjacent arcs in any P is fixed according to the ex-ante dependence graph. Hence, for any arc (p, c) , there is a unique pseudo-path in the collection that involves it.

C.2 Proof of Lemma 16

By the definition of good forests, pseudo-paths and pseudo-matchings, the following statements are equivalent:

1. A subgraph of the ex-ante dependence graph is a good forest.
2. A subgraph satisfies: each node has at most one in-arc; and there is no node p with two out-arcs (p, c) and (p, c') such that rounds p, c, c' have a common element.
3. A subgraph is a union of pseudo-matchings, one for each pseudo-path.

In particular, the equivalence between the first two follows by the definition of good forests. The equivalence between the last two follows by the definitions of pseudo-paths and pseudo-matchings.

C.3 Proof of Lemma 17

Lemma 50. *For any three different rounds $1 \leq p < c < c' \leq T$ such that $(p, c), (p, c') \in E^{\text{ex-ante}}$ and all rounds have a common element, rounds p, c' have the same set of elements, i.e., $\mathcal{E}^p = \mathcal{E}^{c'}$.*

Proof. Suppose $\mathcal{E}^p = \{a, b\}$ and the subscript of (p, c) is a . Then the subscript of (p, c') is b . Since c' is first round involving element b after round p , we have $b \notin \mathcal{E}^c$. Further by that p, c, c' have a common element, the common element can only be a . In sum, $\mathcal{E}^{c'} = \{a, b\} = \mathcal{E}^p$. □

We shall prove the lemma by contradiction. Suppose for contrary that there is a pseudo-path P and $1 \leq t \leq T$ such that the subset of arcs in P that arrive in the first t rounds is not a sub-pseudo-path. Since all arcs in P arrive after round T and form a single pseudo-path, there must be a round $t' > t$ such that in-arcs of t' concatenate two sub-pseudo-paths.

Because of the definition of pseudo-paths and Lemma 15, for each round t , the in-arcs of t are simultaneously added into a pseudo-path such that they are adjacent in the pseudo-path. As a result, there are only two possibilities of the concatenation:

1. If t' has only one in-arc, this in-arc is next to only at most one other arc in the pseudo-path by definition. Hence, it cannot concatenate two sub-pseudo-paths.
2. If t' has two in-arcs (p, t') , (p', t') and it concatenates two sub-pseudo-paths together in round t' , there must be $c, c' < t'$ such that $(p, c), (p', c') \in P$, and further rounds p, c, t' must have a common element and rounds p', c', t' must have a common element. By Lemma 50, we have $\mathcal{E}^p = \mathcal{E}^t = \mathcal{E}^{p'}$. If $p \neq p'$, the in-arc from the earlier one, e.g., (p, t') shall not exist in the ex-ante dependence graph by definition. If $p = p'$, on the other hand, arcs (p, c) (p', c) do not exist in the ex-ante dependence graph. In fact, they shall be two parallel arcs (p, t') that form a pseudo-path on their own, like the right-most pseudo-path in Figure 3a.

In sum, there is always a contraction in all cases.

C.4 Proof of Lemma 18

With the rows and columns in the order of q_U, q_{U^2}, q_M , the transition matrices P^+ of σ^+ and P^- of σ^- are as follows:

$$P^+ = \begin{bmatrix} 0 & 1-p & p \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}, \quad P^- = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ p & 1-p & 0 \end{bmatrix}.$$

Then, the proof of the stationary distribution follows from the next two equations:

$$\begin{aligned} \pi P^+ &= \left(\frac{1}{3-p}, (1-p) \cdot \frac{1}{3-p}, p \cdot \frac{1}{3-p} + \frac{1-p}{3-p} \right) = \left(\frac{1}{3-p}, \frac{1-p}{3-p}, \frac{1}{3-p} \right) = \pi, \\ \pi P^- &= \left(\frac{1-p}{3-p} + p \cdot \frac{1}{3-p}, (1-p) \cdot \frac{1}{3-p}, \frac{1}{3-p} \right) = \left(\frac{1}{3-p}, \frac{1-p}{3-p}, \frac{1}{3-p} \right) = \pi. \end{aligned}$$

C.5 Proof of Lemma 19

By the definitions of σ^+ and σ^- , the state sequences determine the corresponding choice sequence, i.e. $q^i = q_M$ if and only if $d_i = M$ and $\hat{q}^i = q_M$ if and only if $\hat{d}_i = M$. Therefore, it suffices to show that the distributions of the state sequences $(q^i)_{0 \leq i \leq \ell}$ and $(\hat{q}^i)_{0 \leq i \leq \ell}$ are the same.

With the chain rule:

$$\begin{aligned} \Pr[q^0, \dots, q^\ell] &= \prod_{i=0}^{\ell} \Pr[q^i | q^0, \dots, q^{i-1}], \\ \Pr[\hat{q}^0, \dots, \hat{q}^\ell] &= \prod_{i=0}^{\ell} \Pr[\hat{q}^i | \hat{q}^0, \dots, \hat{q}^{i-1}]. \end{aligned}$$

It suffices to show that for any $1 \leq i \leq \ell$:

$$\Pr[q^i | q^0, \dots, q^{i-1}] = \Pr[\hat{q}^i | \hat{q}^0, \dots, \hat{q}^{i-1}] .$$

The case when $i = 0$ follows by Lemma 18, i.e., by that states q^0, \hat{q}^0 both follow the common stationary distribution of σ^+, σ^- . Next, we consider the other conditional probabilities.

For $i_0 \leq i \leq \ell$, it follows directly from the memoryless property of probabilistic automata:

$$\Pr[q^i | q^0, \dots, q^{i-1}] = \Pr[q^i | q^{i-1}] = \Pr[\hat{q}^i | \hat{q}^{i-1}] = \Pr[\hat{q}^i | \hat{q}^0, \dots, \hat{q}^{i-1}] .$$

For $1 \leq i < i_0$, the memoryless property still holds for q^i . On the other hand, it can be deduced that the memoryless property also holds for \hat{q}^i :

$$\begin{aligned} \Pr[\hat{q}^i | \hat{q}^0, \dots, \hat{q}^{i-1}] &= \frac{\Pr[\hat{q}^0, \dots, \hat{q}^i]}{\Pr[\hat{q}^0, \dots, \hat{q}^{i-1}]} \\ &= \frac{\Pr[\hat{q}^0, \dots, \hat{q}^{i-2} | \hat{q}^{i-1}, \hat{q}^i] \cdot \Pr[\hat{q}^{i-1}, \hat{q}^i]}{\Pr[\hat{q}^0, \dots, \hat{q}^{i-2} | \hat{q}^{i-1}] \cdot \Pr[\hat{q}^{i-1}]} \\ &= \Pr[\hat{q}^i | \hat{q}^{i-1}] . \end{aligned}$$

It remains to verify that the joint distributions of state pairs (q^i, q^{i-1}) and $(\hat{q}^i, \hat{q}^{i-1})$ are identical, which follows by:

$$\text{diag}(\boldsymbol{\pi})P^+ = \begin{bmatrix} 0 & 1-p & p \\ 0 & 0 & 1-p \\ 1 & 0 & 0 \end{bmatrix} = \left(\text{diag}(\boldsymbol{\pi})P^- \right)^T .$$

C.6 Proof of Lemma 23

By Lemma 20, to show the first equation it suffices to show for any $i \leq j$:

$$\Pr[q^j = q_M | q^i = q_M] = f_{i-j} .$$

The proof is an induction that corresponds to the recurrence. First consider the base cases. The case when $j = i$ is trivial. The case when $j = i + 1$ holds because $q^i = q_M$ implies $q^{i+1} = q_U$ by the definition of σ^+ . The case when $j = i + 2$ holds because $q^i = q_M$ implies $q^{i+1} = q_U$, from which σ^+ transition to $q^{i+2} = q_M$ with probability p by definition.

Finally consider the case when $j \geq i + 3$. There are only two possibilities in the two rounds after i . The first case is $q^{i+1} = q_U$ and $q^{i+2} = q_M$, which happens with with probability p , and after which $q^j = q_M$ with probability f_{j-i-2} by the inductive hypothesis. The other case is $q^{i+1} = q_U$, $q^{i+2} = q_{U^2}$, and $q^{i+3} = q_M$, which happens with with probability $1 - p$, and after which $q^j = q_M$ with probability f_{j-i-3} by the inductive hypothesis. Putting together:

$$\begin{aligned} \Pr[q^j = q_M | q^i = q_M] &= p \cdot f_{j-i-2} + (1-p) \cdot f_{j-i-3} \\ &= f_{j-i} . \end{aligned}$$

Further, according to the definition of σ^+ , $q^i = q_{U^2}$ if and only if $q^{i+1} = q_M$ while $q^i = q_U$ if and only if $q^{i-1} = q_M$. Thus, it follows that:

$$\begin{aligned} \Pr[d_j = M | q^i = q_{U^2}] &= \Pr[d_j = M | q^{i+1} = q_M] = f_{j-i-1} , \\ \Pr[d_j = M | q^i = q_U] &= \Pr[d_j = M | q^{i-1} = q_M] = f_{j-i+1} . \end{aligned}$$

C.7 Proof of Lemma 24

By the recurrence, the first seven terms of the sequence $\{f_k\}$ are:

$$\begin{aligned} f_0 &= 1, & f_1 &= 0, & f_2 &= p, & f_3 &= 1-p, \\ f_4 &= p^2, & f_5 &= 2p(1-p), & f_6 &= p^3 + (1-p)^2. \end{aligned}$$

As a result of $p \geq \frac{\sqrt{5}-1}{2}$, $f_2 = p, f_4 = p^2 \geq f_3 = 1-p$. For any $k \geq 5$, as $f_k = pf_{k-2} + (1-p)f_{k-3}$, it is easy to see that $f_k \geq f_3 = 1-p$ by induction.

On the other hand, as $f_6 = pf_4 + (1-p)f_3$ and $f_4 \geq f_3, f_6 \leq f_4$. Further, as $p \leq \frac{2}{3}$, $f_5 = 2p(1-p) \geq p^2 = f_4 \geq f_6$. For any $k \geq 7$, as $f_k = pf_{k-2} + (1-p)f_{k-3}$, it is easy to see that $f_k \geq f_6 = p^3 + (1-p)^2$ by induction.

C.8 Proof of Lemma 25

Consider any element e , any subset of nodes $U \subseteq V$ involving e and any pseudo-path $P = ((t_i, t'_i)_{e_i})_{1 \leq i \leq \ell}$. Consider the subset that is the union of $E_{U,e}^{\text{ex-ante}} \cap P$ and the subset of arcs in P that are adjacent to two distinct arcs of $E_{U,e}^{\text{ex-ante}} \cap P$. Simply by definition, the first statement holds for this subset. Note that any two adjacent arcs in P are either two in-arcs or two out-arcs of a same node, any two adjacent arcs in P have different subscripts. Therefore any two arcs in $E_{U,e}^{\text{ex-ante}} \cap P$ are not adjacent in P and thus each maximal sub-pseudo-path satisfies the fourth statement, i.e. it alternates between arcs with subscript e and arcs with other subscripts, and the second statement, i.e. it is odd-length. Next, we shall show the third statement for this subset.

By definition of the subset and the fourth statement, it suffices to show that any two arcs with subscript e cannot be two arcs apart in P . Since in-arcs of a node are added simultaneously into the same pseudo-path in its corresponding round, the first arcs added into the pseudo-path should be the out-arcs of a node, say the initiative node of P . Let i_0 be the minimum index among the first arcs of P . The initiative node is then t'_{i_0} . Let $\{a, b\}$ be the set of elements of initiative node, i.e. $\mathcal{E}^{t'_{i_0}} = \{a, b\}$, and a be the subscript of the i_0 -th arc, i.e. $e_{i_0} = a$. Then, the subscripts of arcs in P on the pseudo-path can be characterized by the following lemma.

Lemma 51. *The subscripts of arcs in P satisfies:*

1. For any $1 \leq i \leq i_0$, the subscript of the i -th arc is a if and only if $i_0 - i$ is even; and
2. For any $i_0 < i \leq \ell$, the subscript of the i -th arc is b if and only if $i - i_0$ is odd.

Proof. For $1 \leq i \leq i_0$, we shall prove some stronger results:

1. For any $1 \leq i \leq i_0$, the subscript of the i -th arc is a if and only if $i_0 - i$ is even;
2. For any $1 \leq i \leq i_0$, t'_i involves element a ;
3. For any $2 \leq i \leq i_0$, if $i - i_0$ is odd, $t_{i-1} = t_i$ and otherwise $t'_{i-1} = t'_i$.

We shall prove it by induction. The base case is that the subscript of the i_0 -th arc is a , t'_{i_0} involves a and that $t_{i_0-1} = t_{i_0}$. For any $1 \leq i \leq i_0 - 1$, given that the subscript of the i -th arc is a , t'_i involves a and $t_{i-1} = t_i$, the subscript of the $(i-1)$ -th arc can't be a , as adjacent arcs have different subscripts, and $t_{i-2} \neq t_{i-1}$ (if $i \geq 3$), as each node has at most 2 out-arcs. Then it is clear $t'_i = t'_{i-1}$ and round $t_i = t_{i-1}, t'_i, t'_{i-1}$ have a common element. By Lemma 50, rounds t_{i-1}, t'_{i-1} have the same set of elements. Therefore t'_{i-1} involves a . On the other hand, given that the subscript of the i -th

arc is not a , t'_i involves a and $t'_{i-1} = t'_i$, the subscript of the $(i-1)$ -th arc is a , as t'_i has two out-arcs and one of them has subscript a , and $t'_{i-2} \neq t'_{i-1}$ (if $i \geq 3$). Therefore, it is clear that t'_{i-1} involves a and $t_{i-2} = t_{i-1}$.

If $i_0 < \ell$, i.e. there is another in-arc of the initiative node, for any $i_0 < i \leq \ell$, it follows the symmetry with $1 \leq i \leq i_0$. \square

With the characterization of the subscripts, it is clear that for any $i \neq i_0 - 1$, the subscripts of i -th arc and $(i+3)$ -th arc are different. The last possible violation of the third statement is that the $(i_0 - 1)$ -th arc of P and the $(i_0 + 2)$ -th arc of P have a common element. Note that these two arcs are out-arcs of the origins of in-arcs of the initiative node because of the characterization, this violation can be ruled out by the following lemma.

Lemma 52. *If the two origins of in-arcs of the initiative node have a same element, i.e. $\mathcal{E}^{t_{i_0}} \cap \mathcal{E}^{t_{i_0+1}} \neq \emptyset$, at least one of them doesn't extend another out-arc in later rounds.*

Proof. If the two origins are the same, it is clear that the origin has already extended two out-arcs and it can't further extend. Otherwise, suppose the element is c . One of the origins have already extended an out-arc with subscript c before round t'_{i_0} and it can't further extend. \square

C.9 Proof of Lemma 26

We shall prove a stronger claim that lets the left-hand-side probability be further conditioned on an arbitrary realization of $q^{i_{m-1}-4}$. This probability equals:

$$\Pr[q^{i_{m-1}} = q_U \mid X_{i_1}, \dots, X_{i_{m-1}} = 0, q^{i_{m-1}-4}] = \Pr[q^{i_{m-1}} = q_U \mid X_{i_{m-1}} = 0, q^{i_{m-1}-4}].$$

By Bayes' rule, it is further equal to:

$$\frac{\Pr[q^{i_{m-1}} = q_U \mid q^{i_{m-1}-4}]}{\Pr[q^{i_{m-1}} = q_U \mid q^{i_{m-1}-4}] + \Pr[q^{i_{m-1}} = q_{U^2} \mid q^{i_{m-1}-4}]}.$$

It remains to show that:

$$p \cdot \Pr[q^{i_{m-1}} = q_U \mid q^{i_{m-1}-4}] \geq \Pr[q^{i_{m-1}} = q_{U^2} \mid q^{i_{m-1}-4}]. \quad (21)$$

Consider the transition matrix P^+ of the automaton σ^+ , with columns and rows in the order of q_U , q_{U^2} , and q_M :

$$P^+ = \begin{bmatrix} 0 & 1-p & p \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}.$$

The transition after four steps (from $i_{m-1} - 4$ to i_{m-1}) is:

$$(P^+)^4 = \begin{bmatrix} p^2 & (1-p)^2 & 2p(1-p) \\ p & 0 & 1-p \\ 1-p & p(1-p) & p^2 \end{bmatrix}.$$

Since the first column multiplied by p dominates the second column in every entry, we prove Eqn. (21), and thus the lemma.

C.10 Proof of Lemma 27

Consider the collection of all maximal tree-paths consisting of nodes in U .

Note that in the good forest, for any node p with two children c and c' , the corresponding rounds have no common element. Since every node in U involves element e , for any node $p \in U$, at most one of its children is in U . Since the collection consists of maximal tree-paths consisting of nodes in U , for any $p, c \in U$ such that p is the parent of c , p is on one tree-path in the collection if and only if c is on the tree-path. Therefore, for any node p and any tree-path in the collection, the neighbors (i.e., parent or children) of p in the path is fixed. It is clear that any two distinct tree-paths are disjoint. Moreover, if there is an arc (p, c) between nodes in distinct tree-paths, there should be two tree-paths involving node p , which contradicts to the fact the tree-paths are pairwise disjoint.

Consider one path in the collection consisting of nodes $t_1, t_2, \dots, t_k \in U$. For each $2 \leq i \leq k$, if arc (t_{i-1}, t_i) have subscript e , it is clear that the algorithm sets $\ell_{t_i}(e) = \ell_{t_{i-1}}(e)$. Otherwise, as both node t_{i-1} and t_i involve e and the element corresponding to subscript of the arc, the label of the other element in these nodes clearly imply $\ell_{t_i}(e) = \ell_{t_{i-1}}(e)$. Therefore, element e has the same label in each tree-path.

C.11 Proof of Lemma 28

By symmetry, consider label $\ell = \text{H}$ without loss of generality. The lemma holds vacuously for $k \geq 4$ since by design the automaton never selects the same label four times in a roll.¹¹ Next we prove the cases of $k = 1, 2, 3$. The remaining argument lets q^j denote the state after round j , and lets s^j denote the selected label in round j .

By the symmetry of automaton σ^* , for any round j :

$$\Pr[q^j = q_{\text{H}}] = \Pr[q^j = q_{\text{T}}] , \quad \Pr[q^j = q_{\text{H}^2}] = \Pr[q^j = q_{\text{T}^2}] . \quad (22)$$

Further, by the above symmetry and by:

$$\begin{aligned} \Pr[q^j = q_{\text{H}^2}] &= \Pr[q^{j-1} = q_{\text{H}}] \Pr[s^j = \text{H} \mid q^{j-1} = q_{\text{H}}] = \Pr[q^{j-1} = q_{\text{H}}] \frac{1-\beta}{2} , \\ \Pr[q^j = q_{\text{H}}] &\geq \Pr[q^{j-1} = q_{\text{T}}] \Pr[s^j = \text{H} \mid q^{j-1} = q_{\text{T}}] = \Pr[q^{j-1} = q_{\text{T}}] \frac{1+\beta}{2} , \end{aligned}$$

we have:

$$\frac{\Pr[q^j = q_{\text{H}}]}{\Pr[q^j = q_{\text{H}^2}]} \geq \frac{1+\beta}{1-\beta} . \quad (23)$$

If $k = 1$, the symmetry implies that the marginal probability of selecting each label in round i equals $\frac{1}{2}$.

¹¹The longest identical sections are selecting H in three consecutive rounds from state T^2 .

If $k = 2$, the probabilities of selecting \top two consecutive times from each of the states are:

$$\begin{aligned}\Pr[s^i = s^{i+1} = \top \mid q^{i-1} = q_{\mathcal{O}}] &= \Pr[s^i = \top \mid q^{i-1} = q_{\mathcal{O}}] \Pr[s^{i+1} = \top \mid q^i = q_{\top}] = \frac{1-\beta}{4}, \\ \Pr[s^i = s^{i+1} = \top \mid q^{i-1} = q_{\mathcal{H}}] &= \Pr[s^i = \top \mid q^{i-1} = q_{\mathcal{H}}] \Pr[s^{i+1} = \top \mid q^i = q_{\top}] = \frac{1-\beta^2}{4}, \\ \Pr[s^i = s^{i+1} = \top \mid q^{i-1} = q_{\top}] &= \Pr[s^i = \top \mid q^{i-1} = q_{\top}] \Pr[s^{i+1} = \top \mid q^i = q_{\top^2}] = 0, \\ \Pr[s^i = s^{i+1} = \top \mid q^{i-1} = q_{\mathcal{H}^2}] &= \Pr[s^i = \top \mid q^{i-1} = q_{\mathcal{H}^2}] \Pr[s^{i+1} = \top \mid q^i = q_{\mathcal{O}}] = \frac{1}{2}, \\ \Pr[s^i = s^{i+1} = \top \mid q^{i-1} = q_{\top^2}] &= 0.\end{aligned}$$

Hence:

$$\Pr[s^i = s^{i+1} = \top] = \Pr[q^{i-1} = q_{\mathcal{O}}] \frac{1-\beta}{4} + \Pr[q^{i-1} = q_{\mathcal{H}}] \frac{1-\beta^2}{4} + \Pr[q^{i-1} = q_{\mathcal{H}^2}] \frac{1}{2}.$$

By Eqn. (23), we further get that:

$$\begin{aligned}\Pr[s^i = s^{i+1} = \top] &\geq \Pr[q^{i-1} = q_{\mathcal{O}}] \frac{1-\beta}{4} + \Pr[q^{i-1} = q_{\mathcal{H}} \text{ or } q_{\mathcal{H}^2}] \left(\frac{1+\beta}{2} \frac{1-\beta^2}{4} + \frac{1-\beta}{2} \frac{1}{2} \right) \\ &= \Pr[q^{i-1} = q_{\mathcal{O}}] \frac{1-\beta}{4} + \Pr[q^{i-1} = q_{\mathcal{H}} \text{ or } q_{\mathcal{H}^2}] \left(\frac{(1+\beta)^2}{2} + 1 \right) \frac{1-\beta}{4} \\ &= \left(\Pr[q^{i-1} = q_{\mathcal{O}}] + 2\Pr[q^{i-1} = q_{\mathcal{H}} \text{ or } q_{\mathcal{H}^2}] \right) \frac{1-\beta}{4} \quad (\beta = \sqrt{2} - 1) \\ &= \frac{1-\beta}{4}.\end{aligned}\tag{Eqn. (22)}$$

If $k = 3$, the automaton must start from $q_{\mathcal{H}^2}$ in order to select \top in three consecutive rounds. The probability equals:

$$\Pr[q^{i-1} = q_{\mathcal{H}^2}] \Pr[s^i = \top \mid q^{i-1} = q_{\mathcal{H}^2}] \Pr[s^{i+1} = \top \mid q^i = q_{\mathcal{O}}] \Pr[s^{i+2} = \top \mid q^{i+1} = q_{\top}].$$

The first term is at most $\frac{1-\beta}{4}$ by Equations (22) and (23). The last three equal 1, $\frac{1}{2}$, and $\frac{1-\beta}{2}$ respectively. Hence:

$$\Pr[s^i = s^{i+1} = s^{i+2} = \top] = \frac{(1-\beta)^2}{16} < \frac{(1-\beta)^2}{8}.$$

C.12 Proof of Lemma 29

It suffices to prove it for $i = 2$, since otherwise it reduces to the case of $i = 2$ by conditioning on the state after round $i - 2$. Further, if the automaton starts from the original state $q_{\mathcal{O}}$, it follows from Lemma 28. If the automaton starts from $q_{\mathcal{H}^2}$ or q_{\top^2} , it resets back to the original state $q_{\mathcal{O}}$ after the first round and once again the lemma reduces to Lemma 28. Finally the lemma holds vacuously for $k \geq 4$ since automaton σ^* never selects the same label in four consecutive rounds.

The remaining proof consider starting from $q_{\mathcal{H}}$ and q_{\top} and $k \in \{1, 2, 3\}$. By symmetry, we

consider label $\ell = \text{H}$ without loss of generality. If the automaton starts from q_{H} :

$$\begin{aligned}
\Pr[s^2 = \text{T} \mid q^0 = q_{\text{H}}] &= \Pr[s^1 = \text{H} \mid q^0 = q_{\text{H}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{H}^2}] \\
&\quad + \Pr[s^1 = \text{T} \mid q^0 = q_{\text{H}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{T}}] \\
&= \frac{1-\beta}{2} \cdot 1 + \frac{1+\beta}{2} \cdot \frac{1-\beta}{2} \\
&= \frac{1}{2}; \qquad (\beta = \sqrt{2} - 1)
\end{aligned}$$

$$\begin{aligned}
\Pr[s^2 = s^3 = \text{T} \mid q^0 = q_{\text{H}}] &= \Pr[s^1 = \text{H} \mid q^0 = q_{\text{H}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{H}^2}] \Pr[s^3 = \text{T} \mid q^2 = q_{\text{O}}] \\
&\quad + \Pr[s^1 = \text{T} \mid q^0 = q_{\text{H}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{T}}] \Pr[s^3 = \text{T} \mid q^2 = q_{\text{T}^2}] \\
&= \frac{1-\beta}{2} \cdot 1 \cdot \frac{1}{2} + \frac{1+\beta}{2} \cdot \frac{1-\beta}{2} \cdot 0 \\
&= \frac{1-\beta}{4};
\end{aligned}$$

$$\begin{aligned}
\Pr[s^2 = s^3 = s^4 = \text{T} \mid q^0 = q_{\text{H}}] &= \Pr[s^1 = \text{H} \mid q^0 = q_{\text{H}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{H}^2}] \\
&\quad \Pr[s^3 = \text{T} \mid q^2 = q_{\text{O}}] \Pr[s^4 = \text{T} \mid q^3 = q_{\text{H}}] \\
&= \frac{1-\beta}{2} \cdot 1 \cdot \frac{1}{2} \cdot \frac{1-\beta}{2} \\
&= \frac{(1-\beta)^2}{8}.
\end{aligned}$$

The last case omits the $s^1 = \text{T}$ option because the automaton by design never selects T in four consecutive rounds.

If the automaton starts from q_{T} :

$$\begin{aligned}
\Pr[s^2 = \text{T} \mid q^0 = q_{\text{T}}] &= \Pr[s^1 = \text{H} \mid q^0 = q_{\text{T}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{H}}] \\
&\quad + \Pr[s^1 = \text{T} \mid q^0 = q_{\text{T}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{T}^2}] \\
&= \frac{1+\beta}{2} \cdot \frac{1+\beta}{2} + \frac{1-\beta}{2} \cdot 0 \\
&= \frac{1}{2}; \qquad (\beta = \sqrt{2} - 1)
\end{aligned}$$

$$\begin{aligned}
\Pr[s^2 = s^3 = \text{T} \mid q^0 = q_{\text{T}}] &= \Pr[s^1 = \text{H} \mid q^0 = q_{\text{T}}] \Pr[s^2 = \text{T} \mid q^1 = q_{\text{H}}] \Pr[s^3 = \text{T} \mid q^2 = q_{\text{T}}] \\
&= \frac{1+\beta}{2} \cdot \frac{1+\beta}{2} \cdot \frac{1-\beta}{2} \\
&= \frac{1-\beta}{4}. \qquad (\beta = \sqrt{2} - 1)
\end{aligned}$$

The second case omits the $s^1 = \text{T}$ option since automaton σ^* cannot select T in three consecutive rounds starting from state q_{T} .

Finally, it is impossible to have $s^2 = s^3 = s^4 = \text{T}$ starting from state $q^0 = q_{\text{T}}$ because from here we cannot have $q^1 = q_{\text{H}^2}$, the only state of automaton σ^* that could lead to selecting T in the next three rounds.

C.13 Proof of Lemma 30

Let q^j, s^j denote the states and selected labels of the first copy, and let \hat{q}^j, \hat{s}^j denote those of the second copy. If the initial state is the original state $q_{\mathcal{O}}$, it follows by Lemma 28. If the initial state is $q_{\mathcal{H}^2}$ or $q_{\mathcal{T}^2}$, the lemma holds because the first selections in the two copies are the same.

By symmetry, we next without loss of generality that the initial state is $q^0 = \hat{q}^0 = q_{\mathcal{H}}$. From $q_{\mathcal{H}}$ it is impossible to select \mathcal{H} in the next two rounds, or to select \mathcal{T} in the next three rounds. Hence, the lemma follows if $k \geq 2$ or $\hat{k} \geq 3$. It remains to consider $k = 1$, and $\hat{k} = 1, 2$.

If $k = \hat{k} = 1$:

$$\begin{aligned} \Pr[s^1 = \mathcal{H}, \hat{s}^1 = \mathcal{T} \mid q^0 = \hat{q}^0 = q_{\mathcal{H}}] &= \Pr[s^1 = \mathcal{H} \mid q^0 = q_{\mathcal{H}}] \Pr[\hat{s}^1 = \mathcal{T} \mid \hat{q}^0 = q_{\mathcal{H}}] \\ &= \frac{1 - \beta}{2} \frac{1 + \beta}{2} \\ &< \frac{1}{4}. \end{aligned}$$

If $k = 1$ and $\hat{k} = 2$:

$$\begin{aligned} \Pr[s^1 = \mathcal{H}, \hat{s}^1 = \hat{s}^2 = \mathcal{T} \mid q^0 = \hat{q}^0 = q_{\mathcal{H}}] &= \Pr[s^1 = \mathcal{H} \mid q^0 = q_{\mathcal{H}}] \Pr[\hat{s}^1 = \hat{s}^2 = \mathcal{T} \mid \hat{q}^0 = q_{\mathcal{H}}] \\ &= \frac{1 - \beta}{2} \frac{1 - \beta^2}{4} \\ &< \frac{1 - \beta}{8}. \end{aligned}$$

D Missing Proofs in Section 6

D.1 Proof of Theorem 33

Proof. We first verify the *feasibility* of the stated solution. Constraint (7) holds with equality by the definitions of $a(k)$ and $b(k)$.

Constraint (8) also holds with equality. When $k = 0$, it follows by:

$$\begin{aligned} \Gamma &= p(0) - \sum_{i=0}^{\infty} \left(\left(\frac{2}{3}\right)^i - \left(\frac{2}{3}\right)^{i+1} \right) p(i) && \text{(Definition of } \Gamma, p(0) = 1) \\ &= \sum_{i=0}^{\infty} \left(\frac{2}{3}\right)^{i+1} p(i) - \sum_{i=1}^{\infty} \left(\frac{2}{3}\right)^i p(i) \\ &= \sum_{i=0}^{\infty} \left(\frac{2}{3}\right)^{i+1} (p(i) - p(i+1)). \end{aligned} \tag{24}$$

This equals $2b(0)$ by definition. Then, it further holds inductively for $k \geq 1$ because:

$$\begin{aligned} b(k) &= \frac{3}{2}b(k-1) - \frac{1}{2}(p(k-1) - p(k)) && \text{(Definition of } b(k), b(k-1)) \\ &= b(k-1) - \frac{1}{2}a(k-1). && \text{(Definition of } a(k-1)) \end{aligned}$$

That is, the left-hand-side of Constraint (8) stays the same from $k-1$ to k . Since the above equation $b(k) = b(k-1) - \frac{1}{2}a(k-1)$ would also imply Constraint (9) provided that $a(k-1) \geq 0$,

it remains to verify that $a(k)$ and $b(k)$ are nonnegative. By its definition and by $p(k+1) \leq \frac{2}{3}p(k)$, we get that $b(k) \geq 0$. The non-negativity of $a(k)$ follows by its definition and by:

$$\begin{aligned} b(k) &= \frac{1}{3}p(k) - \sum_{i=k+1}^{\infty} \left(\left(\frac{2}{3}\right)^{i-k-1} - \left(\frac{2}{3}\right)^{i-k} \right) p(i) && \text{(Definition of } b(k)) \\ &\leq \frac{1}{3}p(k) \\ &\leq p(k) - p(k+1) . && (p(k+1) \leq \frac{2}{3}p(k)) \end{aligned}$$

Next we establish its *optimality*. Multiplying Constraint (8) by $\left(\frac{2}{3}\right)^k$ and summing over $k \geq 0$:

$$\sum_{k=0}^{\infty} \left(\frac{2}{3}\right)^k \left(\sum_{i=0}^{k-1} a(i) + 2b(k) \right) \geq \sum_{k=0}^{\infty} \left(\frac{2}{3}\right)^k \Gamma .$$

Grouping terms on the left and dividing both sides by 3, this is:

$$\sum_{i=0}^{\infty} \left(\frac{2}{3}\right)^{i+1} a(i) + \sum_{i=0}^{\infty} \left(\frac{2}{3}\right)^{i+1} b(i) \geq \Gamma .$$

Further by Constraint (7), the left-hand-side is at most:

$$\sum_{i=0}^{\infty} \left(\frac{2}{3}\right)^{i+1} (p(i) - p(i+1)) .$$

This equals the optimal Γ in the theorem by Eqn. (24). □

D.2 Proof of Theorem 45

Proof. We first verify its *feasibility*. Constraint Eqn. (17) holds with equality by definition, i.e.:

$$a(y) + b(y) = -p'(y) , \quad \forall y \geq 0 . \quad (25)$$

Constraint Eqn. (18) holds for $y = 0$ from integration by parts:

$$b(0) = - \int_0^{\infty} p'(z) e^{-z} dz = \int_0^{\infty} e^{-z} (1 - p(y)) dz = \Gamma .$$

It further holds for $y > 0$ since its left-hand-side is a constant for all y . Indeed, the derivative of the left-hand-side is:

$$\begin{aligned} a(y) + b'(y) &= a(y) - e^y \int_y^{\infty} p'(z) e^{-z} dz + p'(y) && \text{(Definition of } b(y)) \\ &= a(y) + b(y) + p'(y) && \text{(Definition of } b(y)) \\ &= 0 . && \text{(Eqn. (25))} \end{aligned}$$

Constraint 19 holds, i.e., $b(y)$ is decreasing because:

$$\begin{aligned} b'(y) &= -e^y \int_y^{\infty} p'(z) e^{-z} dz + p'(y) && \text{(Definition of } b(y)) \\ &= -e^y \int_y^{\infty} (p'(y) - p'(z)) e^{-z} dz \\ &\leq 0 . && \text{(Convexity of } p) \end{aligned}$$

Finally, $b(y)$ is non-negative by definition and by that p is decreasing. The non-negativity of $a(y)$ follows by:

$$a(y) = -p'(y) + e^y \int_y^\infty p'(z)e^{-z} dz \quad (\text{Definitions of } a(y), b(y))$$

$$= e^y \int_y^\infty e^{-z} p''(z) dz \quad (\text{Integration by parts})$$

$$\geq 0 . \quad (\text{Convexity of } p)$$

□