

Making Three out of Two: Three-Way Online Correlated Selection

Yongho Shin ✉

Department of Computer Science, Yonsei University, Seoul, South Korea

Hyung-Chan An¹ ✉ 

Department of Computer Science, Yonsei University, Seoul, South Korea

Abstract

Two-way online correlated selection (two-way OCS) is an online algorithm that, at each timestep, takes a pair of elements from the ground set and irrevocably chooses one of the two elements, while ensuring negative correlation in the algorithm's choices. Whilst OCS was initially invented by Fahrback, Huang, Tao, and Zadimoghaddam to break a natural long-standing barrier in the edge-weighted online bipartite matching problem, it is an interesting technique on its own due to its capability of introducing a powerful algorithmic tool, namely negative correlation, to online algorithms. As such, Fahrback et al. posed two tantalizing open questions in their paper, one of which was the following: Can we obtain *n-way OCS* for $n > 2$, in which the algorithm can be given $n > 2$ elements to choose from at each timestep?

In this paper, we affirmatively answer this open question by presenting a *three-way OCS*. Our algorithm uses two-way OCS as its building block and is simple to describe; however, as it internally runs two instances of two-way OCS, one of which is fed with the output of the other, the final output probability distribution becomes highly elusive. We tackle this difficulty by approximating the output distribution of OCS by a flat, less correlated function and using it as a safe “surrogate” of the real distribution. Our three-way OCS also yields a 0.5093-competitive algorithm for edge-weighted online matching, demonstrating its usefulness.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases online correlated selection, multi-way OCS, online algorithms, negative correlation, edge-weighted online bipartite matching

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2021.49

Funding This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF-2019R1C1C1008934).

Acknowledgements We thank the anonymous reviewers for their helpful comments.

1 Introduction

Online correlated selection (OCS) is an online algorithm that, at each timestep, takes a subset of the ground set as the input and irrevocably chooses a single element from the subset. When every input subset has cardinality n , we call it *n-way OCS* in particular. The aim of online correlated selection is to ensure a certain level of negative correlation in the choice made by the algorithm. For example, suppose we run a two-way OCS and afterwards specify m timesteps that contained some common element. The probability that this element was never chosen would be $(1/2)^m$ if the algorithm made independent and uniformly random choices; the goal of two-way OCS is to reduce this probability by introducing negative correlations. (See Definition 4 for a full definition that quantifies the desired amount of reduction.)

¹ Corresponding author. Department of Computer Science, Yonsei University, 50 Yonsei-ro, Seodaemun-gu, Seoul 03722, South Korea.



OCS was first invented by Fahrbach, Huang, Tao, and Zadimoghaddam [10] to attack the edge-weighted online bipartite matching problem. Negative correlation has proven to be a very powerful technique in approximation algorithms design (see, e.g., [34, 2, 7] for a limited list of examples); this suggests that OCS as well bears high potential as a general tool in online algorithms design rather than as a specialized technique to solve a particular problem. This opportunity was also observed by the breakthrough paper of Fahrbach et al. [10] and recently exemplified by Huang, Zhang, and Zhang [23], who devised a certain variant of OCS called *panoramic OCS* to solve the AdWords problem with general bids.

In light of such value of OCS as an algorithmic tool, Fahrbach et al. [10] raised in their paper two follow-up questions that arise quite naturally: Can we improve the performance of their two-way OCS? Can we obtain an n -way OCS for $n > 2$? This paper affirmatively answers the latter question. In this paper, we present a simple *three-way OCS* and analyze its performance. We will also show that our three-way OCS can be used to improve the previous competitive ratio of 0.5086 due to Fahrbach et al. to give a new 0.5093-competitive algorithm for edge-weighted online bipartite matching.

In fact, the construction itself of our three-way OCS is easy to describe. It internally executes two instances of two-way OCS. Upon arrival of a triple, we choose two of the three elements uniformly at random, and let the first two-way OCS choose one of them. We then pass its output, along with the element that was left out of the first OCS, to the second OCS. The second OCS chooses one of these two elements; this choice becomes the final output of this timestep.

In Section 3, we analyze the performance of our three-way OCS for a special case first: we bound the probability that a certain element, say u , is never chosen for k consecutive timesteps whose triples contain u . It is not that we require these k timesteps to be consecutive in the original input: they need to be consecutive in the subsequence of timesteps on which u appeared in the triple. By the definition of two-way γ -OCS, the probability that the second two-way OCS never chooses u is no greater than $(1/2)^j(1 - \gamma)^{\max(j-1, 0)}$ for some constant γ , where j is the number of times u was passed to the second OCS during those k timesteps. Since this bound depends only on j , the question really reduces to determining (the probability distribution of) j .

In order for u to be passed to the second OCS, it needs to be either left out of the first OCS or output by it. It is easy to count how many times u is left out of the first OCS: this follows a binomial distribution. Therefore, the challenge is in counting the number of times u is output by the first OCS. Unfortunately, its probability distribution highly depends on the actual *input* to the first OCS, rather than the number of times u is shown to the first OCS. Nonetheless, the following observation is crucial in coping with this difficulty: the probability distribution of the number of times Fahrbach et al.’s two-way OCS chooses u is a unimodal symmetric distribution. Recall that the probability that the second OCS never chooses u is bounded by $(1/2)^j(1 - \gamma)^{\max(j-1, 0)}$, which is “nearly” convex. Therefore, even though we cannot exactly calculate the probability distribution of j without the full knowledge of the input, the above observation implies that we can instead use a “flatter” unimodal symmetric distribution in lieu of the actual distribution of j . Thanks to the near-convexity, this would give a valid upper bound on the probability. We formalize what a “flatter” distribution is by defining the notion of *central dominance* as follows.

► **Definition 1 (Central Dominance).** *Given two discrete symmetric probability distributions D_1 and D_2 on $\{0, 1, \dots, x\}$ whose probability mass functions are p_1 and p_2 , respectively, we say D_1 centrally dominates D_2 if there exists $z \in [0, \frac{x}{2}]$ such that, for any integer $y \in [\frac{x}{2} - z, \frac{x}{2} + z]$, $p_1(y) \geq p_2(y)$, and for any integer $y \in [0, \frac{x}{2} - z) \cup (\frac{x}{2} + z, x]$, $p_1(y) \leq p_2(y)$.*

We then construct our “surrogate” distribution that is centrally dominated by any possible probability distribution of j . This distribution depends only on the number of times the first OCS is given u . It is therefore much more amenable and allows us to obtain a bound on the probability that our three-way OCS never chooses u from the given consecutive triples.

In Section 4, we generalize this bound to the non-consecutive case, i.e., a disjoint set of consecutive subsequences of timesteps containing u . To obtain this bound, we perform a set of surgical operations that modify the input to the first OCS, which are designed to reduce negative correlation. These operations are inspired by those of Fahrbach et al. [10] that they used to obtain a similar generalization. In our case, however, we face a new obstacle: previously, it sufficed to bound only the probability that the two-way OCS never chooses a given element, since the output of that OCS was the final output. Our final output on the other hand is determined by the second OCS, and if our modification changes the output distribution of the first OCS, this may affect the output of the second OCS in an obscure way. We prove that a set of careful surgical operations can remove all correlations while ensuring that the bound is not affected. Once the correlations are removed, the general-case bound can be simply given as the product of our bounds from Section 3 for single subsequences.

► **Theorem 2 (simplified).** *Consider a set of m disjoint consecutive subsequences of timesteps whose triples contain some element u of the ground set. Let k_1, \dots, k_m be the lengths of these subsequences. The probability that our three-way OCS never chooses u from these m subsequences is at most*

$$\prod_{i=1}^m \left[\left(\frac{2}{3} \right)^{k_i} (1 - \delta_1)^{\max(k_i - 1, 0)} (1 - \delta_2)^{\max(k_i - 2, 0)} \right],$$

where $\delta_1 = 0.0309587$ and $\delta_2 = 0.0165525$.

Our three-way OCS can be applied to edge-weighted and unweighted online bipartite matching [33].

1.1 Related Work

Introduced by Karp, Vazirani, and Vazirani [25], the unweighted online bipartite matching has been intensively and extensively studied with alternative proofs [16, 4, 8, 9] and under various settings including stochastic models [12, 29, 17, 32, 13, 20], fully online models [18, 19, 22], and general arrival models [14]. The study of edge-weighted online bipartite matching problem was initiated by Kalyanasundaram & Pruhs [24] and Khuller, Mitchell, & Vazirani [27], who independently considered this problem under the metric assumption. Feldman, Korula, Mirrokni, Muthukrishnan, and Pál [11] first investigated the edge-weighted version on arbitrary weights with free disposal, and more thorough understanding of this problem was achieved by subsequent work [28, 10]. Other variants and applications of online bipartite matching have been studied as well, including AdWords [31, 23], vertex-weighted version [1, 21, 13], stochastic or random arrival models [17, 26, 6, 20], and the windowed version [3]. We refer interested readers to the survey of Mehta [30].

Recent related works. Recently, after the authors independently obtained the present results but prior to announcing them, the authors learned that two closely related papers were announced on arXiv [15, 5].

The other open question raised by Fahrbach et al. [10] than the one answered by this paper was to improve two-way OCS. Gao, He, Huang, Nie, Yuan, and Zhong [15] addresses this question by giving a novel automata-based OCS, successfully departing from the previous

matching-based approach: their two-way OCS is a 0.167-OCS. In addition to this, they also give an improved primal-dual analysis and a variant of two-way OCS specifically adapted for edge-weighted online bipartite matching, yielding a 0.519-competitive algorithm. Finally, they consider a weaker relaxed notion of OCS called *semi-OCS*, where the probability bound holds only for consecutive prefixes. They provide a multi-way version of this semi-OCS that leads to a 0.593-competitive algorithm for unweighted/vertex-weighted online bipartite matching.

Blanc and Charikar [5] generalize Fahrbach et al.'s definition of OCS in two ways and give m -way OCS for any m . One of the two generalizations, called (F, m) -OCS, gives the probability bound specified as a discrete function rather than a fixed-form formula such as $(1/2)^j(1 - \gamma)^{\max(j-1, 0)}$. The other is called *continuous OCS*, which allows an element of the ground set to appear in a subset “to a fraction”, where the probability bound is now specified as a continuous function. Their continuous OCS along with their improved primal-dual analysis gives a 0.5368-competitive algorithm for edge-weighted online bipartite matching.

Considering these new results, an interesting question is whether the techniques from these papers and our independent result can together bring improvements in OCS or related problems such as online matching. In fact, since our framework treats the second two-way OCS as a black-box, any improved two-way OCS can be directly plugged into our three-way OCS. Combining the new two-way 0.167-OCS of Gao et al. [15] with our results, for example, immediately yields a 0.5132-competitive algorithm for edge-weighted online bipartite matching.

2 Preliminaries

In this section, we present some notation, definitions, and previous results to be used throughout this paper. Let us first introduce the definition of n -way OCS and two-way γ -OCS.

► **Definition 3** (n -way OCS). *Given a ground set, an n -way OCS is an online algorithm that, at each iteration, takes a subset of size n as the input and irrevocably chooses an element from the subset.*

For any element u of the ground set, we say a subsequence of subsets containing u is *consecutive* if every subset containing u that arrives between the first and last subsets of the subsequence is also in the subsequence. We also say that two or more subsequences of subsets containing u are *disjoint* if no two of these subsequences share a subset.

For example, suppose we are given $\{u, v\}$, $\{u, w\}$, $\{v, w\}$, $\{u, x\}$, and $\{u, y\}$ during five timesteps, in that order. Observe that $(\{u, v\}, \{u, w\}, \{u, x\})$ is a consecutive subsequence containing u , but $(\{u, v\}, \{u, x\}, \{u, y\})$ is not because $\{u, w\}$ appears between $\{u, v\}$ and $\{u, x\}$. Note also that $(\{u, v\}, \{u, w\})$ and $(\{u, y\})$ constitute a set of (two) disjoint consecutive subsequences containing u .

► **Definition 4** (Two-way γ -OCS). *A two-way γ -OCS is a two-way OCS such that, for any element u and a set of m disjoint consecutive subsequences of pairs containing u of lengths k_1, \dots, k_m , the probability that u never gets chosen by the OCS from any of the given subsequences is at most $\prod_{i=1}^m \left(\frac{1}{2}\right)^{k_i} (1 - \gamma)^{\max(k_i-1, 0)}$.*

Fahrbach et al. [10] presents two versions of two-way OCS with varying performance guarantees. We will use both in later sections. For the sake of completeness, we present a full description of the two-way 1/16-OCS of Fahrbach et al. below.

1/16-OCS. We will define what is called the *ex-ante graph* first. Although the algorithm does not need to explicitly construct this graph, it helps simplify the presentation. The vertices of the ex-ante graph correspond to the input pairs. For each pair, say $\{u, v\}$, we introduce an edge between this pair and the immediately following pair that contains u , and likewise for v . For example, if an element u appears at timesteps 2, 4, and 7, pairs 2 and 4 are adjacent, and so are 4 and 7. This implies that every vertex has degree of at most 4. We annotate each edge with the common element that caused this edge. For example, if two pairs $i = \{u, v\}$ and $j = \{v, w\}$ are adjacent due to v , let $(i, j)_v$ denote the edge (and its annotation). Note that the ex-ante graph may have parallel edges (with different annotations).

The algorithm samples exactly three random bits for each pair (or vertex). We will use them to define what is called the *ex-post graph* and further determine the output of the algorithm. The ex-post graph is on the same set of vertices. The first random bit is used to determine if the vertex becomes a “sender” or a “receiver”. If it becomes a sender, we use the second random bit to choose one of the two elements of the pair, and the sender will “want” to select the edge to the immediately following pair that shares the chosen element (if it exists). Similarly, if a vertex becomes a receiver, we use the second random bit to choose one element, and the receiver will “want” to select the edge to the immediately preceding pair that shares the element. Each edge of the ex-ante graph enters the ex-post graph if and only if both of its endpoints want to select the edge. Observe that the ex-post graph is a matching in the ex-ante graph.

Finally, the output of the algorithm is determined as follows. For every unmatched vertex, its output is determined solely by its third random bit. For each edge in the ex-post graph, we negatively correlate the choice of the two endpoints: we use the third random bit of the sender to determine the output of the sender, and the output of the receiver is determined so that the decision made for the shared element is the opposite. For example, if we have an edge annotated with u in the ex-post graph and the sender did not choose u , we choose u for the receiver; otherwise, we do not choose u for the receiver. The third random bit of the receiver is discarded.

► **Lemma 5** (Fahrbach et al. [10]). *This algorithm is a two-way $\frac{1}{16}$ -OCS.*

This algorithm also has the following useful property.

► **Lemma 6** (Fahrbach et al. [10]). *Given a consecutive subsequence of pairs containing an element of length k input to this algorithm, the probability that there does not exist any edges in the subgraph of the ex-post graph induced by the consecutive pairs is at most $(1 - 1/16)^{\max(k-1, 0)}$.*

The other version is a $\frac{13\sqrt{13}-35}{108}$ -OCS.

► **Lemma 7** (Fahrbach et al. [10]). *There exists a two-way $\frac{13\sqrt{13}-35}{108}$ -OCS.*

In the *unweighted online bipartite matching problem*, we are given a bipartite graph $G = (L \cup R, E)$ where we only know L in advance. Each vertex $v \in R$ arrives one by one, and the edges adjacent with v are only then revealed. Upon each arrival of v , we irrevocably decide whether we match v , and if so, to which exposed adjacent vertex in L we match. The objective is to find a matching of the maximum size in G .

In the *edge-weighted online bipartite matching problem with free disposal*, we are given a bipartite graph $G = (L \cup R, E)$ as well as an edge weight $w_{uv} \geq 0$ for each $(u, v) \in E$, where we only know L in advance, again. Each vertex $v \in R$ arrives one at a timestep, and the

edges adjacent with v and their edge weights are disclosed at that time. Upon each arrival of v , we decide whether we match v , and if so, to which adjacent vertex in L we match. We remark that v can be matched to a vertex u that is already matched, after disposing of the edge incident with u in the current matching (*free disposal*). The objective is to find a maximum-weight matching in G .

3 Three-Way Online Correlated Selection

In this section, we present our three-way OCS and analyze it by considering the special case where we are interested in a single consecutive subsequence of triples containing a common element. We will extend this to the general case in Section 4.

3.1 Algorithm

Let \mathcal{A} and \mathcal{B} be two-way OCS algorithms whose random choices are independent. In particular, we choose Fahrback et al.'s 1/16-OCS as \mathcal{A} and $\frac{13\sqrt{13}-35}{108}$ -OCS as \mathcal{B} . Let $\gamma_{\mathcal{A}} := 1/16$ and $\gamma_{\mathcal{B}} := \frac{13\sqrt{13}-35}{108}$. We use the 1/16-OCS as \mathcal{A} due to its property that, in its ex-post graph, each edge's existence is determined by (the random bits of) its endpoints only.

Upon arrival of a triple, say $\{u, v, w\}$, we choose a pair uniformly at random out of $\{\{u, v\}, \{u, w\}, \{v, w\}\}$. We refer to this step as the *random pair choice phase*. Without loss of generality, suppose that $\{u, v\}$ is chosen. We then input $\{u, v\}$ to \mathcal{A} and let it choose one element from the pair. Let us say (without loss of generality again) that u is returned by \mathcal{A} . Now we let \mathcal{B} choose one element from $\{u, w\}$. The element chosen by \mathcal{B} is the final output of this iteration.

3.2 Overview of the Analysis

The goal of Section 3 is to prove the following theorem.

► **Theorem 8.** *Consider a consecutive subsequence of timesteps whose triples contain some element u of the ground set. Let k be the length of the subsequence. The probability that our three-way OCS never chooses u during these k timesteps is at most*

$$\eta(k) := c_1 t_1^k + c_2 t_2^k - c_3 t_3^k - c_4 t_4^k$$

for some $c_1 \approx 0.957795$, $c_2 \approx 0.176756$, $c_3 \approx 0.011047$, $c_4 \approx 0.131738$, $t_1 \approx 0.630024$, $t_2 \approx 0.599919$, $t_3 \approx 0.148345$, and $t_4 = 0.3125$, which is bounded from above by

$$\left(\frac{2}{3}\right)^k (1 - \delta_1)^{\max(k-1, 0)} (1 - \delta_2)^{\max(k-2, 0)},$$

where $\delta_1 = 0.0309587$ and $\delta_2 = 0.0165525$.

In what follows, we will fix an arbitrary set of random choices made during the random pair choice phase, and condition on them. Let x be the number of pairs containing u from the subsequence that are passed to \mathcal{A} . Observe that these pairs also form a consecutive subsequence to \mathcal{A} , and that the number of u 's that are left out of the first OCS is $k - x$.

For $y = 0, \dots, x$, let $p(x, y)$ be the (conditional) probability that \mathcal{A} returns y number of u 's from these pairs. Then, we can see that the (conditional) probability that u never gets chosen from the given consecutive triples is bounded from above by

$$\sum_{y=0}^x p(x, y) \left(\frac{1}{2}\right)^{k-x+y} (1 - \gamma_{\mathcal{B}})^{\max(k-x+y-1, 0)} \quad (1)$$

by Lemma 7 and the fact that the pairs containing u passed to \mathcal{B} are also consecutive (see also Definition 4). To simplify the presentation, we introduce two shorthands: for a probability mass function p' defined on $\{0, \dots, x\}$, let

$$\theta(x, p') := \sum_{y=0}^x p'(y) \cdot \left(\frac{1}{2}\right)^y (1 - \gamma_{\mathcal{B}})^{y-1}; \text{ and} \tag{2}$$

$$\theta'(x, p') := \sum_{y=0}^x p'(y) \cdot \left(\frac{1}{2}\right)^y (1 - \gamma_{\mathcal{B}})^{\max(y-1, 0)}. \tag{3}$$

Observe that (1) is equal to $\left(\frac{1-\gamma_{\mathcal{B}}}{2}\right)^{k-x} \theta(x, p(x, \cdot))$ if $k > x$ and $\theta'(k, p(k, \cdot))$ otherwise (if $k = x$).

We would like to bound (1), but unfortunately, $p(x, y)$ depends on the actual input to \mathcal{A} , not just x . Yet, we can circumvent this problem by exploiting the behavior of \mathcal{A} . Recall that \mathcal{A} constructs the ex-post graph (which is a matching in the ex-ante graph) and negatively correlate the two endpoints of each edge in the ex-post graph. Therefore, when we consider the subgraph of the ex-post graph induced by the given consecutive pairs, we can observe that

- for each edge in this subgraph, exactly one u is chosen from its two endpoints, and
- if a vertex is isolated in this subgraph, u is chosen with probability $1/2$, independently from the other vertices.

This observation implies that the probability distribution $p(x, \cdot)$ is a unimodal symmetric distribution. Moreover, the more likely \mathcal{A} puts an edge in the ex-post graph, the pointier the distribution would be. Using this property, we will construct an imaginary probability distribution $\{p^*(x, y)\}_{y=0, \dots, x}$ such that

- $p^*(x, y)$ is flatter than (or, formally, centrally dominated by) any distribution p that results from \mathcal{A} (see the proof of Lemma 10); and
- $p^*(x, y)$ only depends on x , not the input itself.

We will further demonstrate that, in (2) and (3), substituting p' with a centrally dominated distribution would only overestimate (2) and (3) (Lemma 9). Intuitively speaking, this is because $(1/2)^y(1 - \gamma_{\mathcal{B}})^{y-1}$ and $(1/2)^y(1 - \gamma_{\mathcal{B}})^{\max(y-1, 0)}$ are (nearly) convex functions.

Thus, the probability that our three-way OCS never chooses u from the given consecutive triples can finally be bounded from above by

$$\eta(k) := \sum_{x=0}^k \text{binom}\left(k, x, \frac{2}{3}\right) \left[\sum_{y=0}^x p^*(x, y) \left(\frac{1}{2}\right)^{k-x+y} (1 - \gamma_{\mathcal{B}})^{\max(k-x+y-1, 0)} \right], \tag{4}$$

where $\text{binom}(k, x, r)$ represents the probability of the binomial distribution for x successes out of k trials with probability r (Lemma 11).

3.3 Using a Centrally Dominated Distribution

Recall that k denotes the length of the subsequence, i.e., the number of triples, and x denotes the number of pairs containing u passed to \mathcal{A} . Let y be the number of u 's returned by \mathcal{A} ; then, \mathcal{B} receives a consecutive subsequence of pairs containing u of length $k - x + y$.

We remark that the probability that u is never chosen by \mathcal{B} , conditioned on the event that \mathcal{A} takes x pairs containing u and outputs y number of u 's, can be bounded by $\left(\frac{1}{2}\right)^{k-x+y} \cdot (1 - \gamma_{\mathcal{B}})^{\max(k-x+y-1, 0)}$ from Lemma 7. Note that this bound depends only on the lengths k , x , and y , not the actual input to \mathcal{B} .

For \mathcal{A} , on the other hand, we need to calculate the probability that exactly y number of u 's are chosen by \mathcal{A} from the given subsequence of pairs. However, this probability distribution highly depends on the input to \mathcal{A} ; therefore, we cannot conveniently fix a distribution parameterized by length x anymore. Nonetheless, we will show that there exists a bounding distribution, depending only on x , such that it yields a valid upper bound on the probability that u is never chosen by our entire algorithm from the given subsequence of triples.

3.3.1 Central Dominance

To construct the bounding distribution, we need the notion of central dominance. Recall the definition.

► **Definition 1** (Central Dominance). *Given two discrete symmetric probability distributions D_1 and D_2 on $\{0, 1, \dots, x\}$ whose probability mass functions are p_1 and p_2 , respectively, we say D_1 centrally dominates D_2 if there exists $z \in [0, \frac{x}{2}]$ such that, for any integer $y \in [\frac{x}{2} - z, \frac{x}{2} + z]$, $p_1(y) \geq p_2(y)$, and for any integer $y \in [0, \frac{x}{2} - z) \cup (\frac{x}{2} + z, x]$, $p_1(y) \leq p_2(y)$.*

Intuitively speaking, this definition indicates how flat a symmetric probability distribution is. Following is the lemma that formalizes why a flat distribution helps bound the probability that u never gets chosen from our three-way OCS. The definitions of $\theta(\cdot, \cdot)$ and $\theta'(\cdot, \cdot)$ can be found in (2) and (3), respectively. It is easy to intuitively see that the first half of the lemma should hold; yet, θ' is only *nearly* convex and requires some work.

► **Lemma 9.** *Suppose we are given two discrete symmetric distributions D_1 and D_2 on $\{0, 1, \dots, x\}$ whose probability mass functions are p_1 and p_2 , respectively. If D_1 centrally dominates D_2 , we have $\theta(x, p_1) \leq \theta(x, p_2)$ and $\theta'(x, p_1) \leq \theta'(x, p_2)$.*

Proof. If $x = 0, 1$, it is trivial since there is only one possible symmetric distribution. Now we assume that $x \geq 2$. Given a discrete symmetric distribution with a probability mass function p , let us consider the following operation: For some $w, z \in [0, \frac{x}{2}]$ such that $w > z$, we decrease $p(\frac{x}{2} - z)$ and $p(\frac{x}{2} + z)$ by some $\epsilon > 0$, and increase $p(\frac{x}{2} - w)$ and $p(\frac{x}{2} + w)$ by ϵ , instead. We can observe that p still forms a valid probability distribution, and changes $\theta(x, p)$ by

$$\Delta\theta(x, p) = \epsilon \left(\frac{1}{2}\right)^{x/2} (1 - \gamma_{\mathcal{B}})^{x/2-1} \left[\left(\frac{1 - \gamma_{\mathcal{B}}}{2}\right)^w + \left(\frac{1 - \gamma_{\mathcal{B}}}{2}\right)^{-w} - \left(\frac{1 - \gamma_{\mathcal{B}}}{2}\right)^z - \left(\frac{1 - \gamma_{\mathcal{B}}}{2}\right)^{-z} \right].$$

Let $\alpha := (1 - \gamma_{\mathcal{B}})/2$ and $f(t) := \alpha^t + \alpha^{-t}$. Since $f'(t) = \ln \alpha \cdot (\alpha^t - \alpha^{-t})$, we can verify that $f(t)$ is increasing on $t \geq 0$. Hence, we can see that

$$\Delta\theta(x, p) = \epsilon \left(\frac{1}{2}\right)^{x/2} (1 - \gamma_{\mathcal{B}})^{x/2-1} (f(w) - f(z)) \geq 0.$$

Since D_1 centrally dominates D_2 , we can transform p_1 into p_2 by transferring the mass through a few times of the above operation with proper w 's, z 's, and ϵ 's. For each application, $\theta(x, p)$ only increases, yielding that $\theta(x, p_1) \leq \theta(x, p_2)$.

Proving the other statement is more subtle. Suppose we transform p_1 into p_2 through the same method. If we increase the probabilities of $\frac{x}{2} - w$ and $\frac{x}{2} + w$ where $w < \frac{x}{2}$, we can apply to the above analysis since the change of $\theta'(x, p)$ is exactly the same as $\Delta\theta(x, p)$.

It remains to consider when the probabilities of 0 and x get increased by ϵ . Assume that we instead decrease the probabilities of $\frac{x}{2} - z$ and $\frac{x}{2} + z$ where $0 \leq z \leq \frac{x}{2} - 1$. The change of $\theta'(x, p)$ would be

$$\Delta\theta'(x, p) := \epsilon \left[1 + \left(\frac{1}{2}\right)^x (1 - \gamma_{\mathcal{B}})^{x-1} - \left(\frac{1}{2}\right)^{x/2-z} (1 - \gamma_{\mathcal{B}})^{x/2-z-1} - \left(\frac{1}{2}\right)^{x/2+z} (1 - \gamma_{\mathcal{B}})^{x/2+z-1} \right].$$

Let $\beta := (1/2)^{x/2}(1 - \gamma_{\mathcal{B}})^{x/2-1}$ and $g(z) := (2/(1 - \gamma_{\mathcal{B}}))^z$. We then rewrite $\Delta\theta'(x, p)$ as

$$\Delta\theta'(x, p) = \epsilon [1 + \beta^2(1 - \gamma_{\mathcal{B}}) - \beta(g(z) + 1/g(z))].$$

The partial derivative of $\Delta\theta'(x, p)$ with respect to z is

$$\frac{\partial\Delta\theta'(x, p)}{\partial z} = -\epsilon\beta \ln\left(\frac{2}{1 - \gamma_{\mathcal{B}}}\right) \left[\left(\frac{2}{1 - \gamma_{\mathcal{B}}}\right)^z - \left(\frac{1 - \gamma_{\mathcal{B}}}{2}\right)^z \right].$$

Note that $\frac{\partial\Delta\theta'(x, p)}{\partial z} \leq 0$ for $z \geq 0$. Thus, we can bound $\Delta\theta'(x, p)$ from below by plugging $z = \frac{x}{2} - 1$, i.e.,

$$\Delta\theta'(x, p) \geq \epsilon \left[\frac{1}{2} - \frac{1 + \gamma_{\mathcal{B}}}{2} \left(\frac{1}{2}\right)^{x-1} (1 - \gamma_{\mathcal{B}})^{x-2} \right].$$

Since $x \geq 2$, we have

$$\Delta\theta'(x, p) \geq \epsilon \left[\frac{1}{2} - \frac{1 + \gamma_{\mathcal{B}}}{4} \right] \geq 0,$$

where the last inequality comes from the fact that $\gamma_{\mathcal{B}} < 1$. ◀

3.3.2 Bounding Distribution

In this section, we construct the bounding distribution. As was noted earlier, \mathcal{A} negatively correlates the endpoints of every edge of the ex-post graph. Therefore, we can write the probability $p(x, y)$ that the two-way OCS chooses precisely y number of u 's out of the given x consecutive pairs if we are given $q \in \mathbb{R}_+^{\lfloor x/2 \rfloor + 1}$ defined as follows: for $i = 0, \dots, \lfloor \frac{x}{2} \rfloor$, q_i is the probability that exactly i edges of the ex-post graph have both endpoints in the given subsequence. Note that $\sum_{i=0}^{\lfloor x/2 \rfloor} q_i = 1$. We have

$$p(x, y) = q_0 \cdot \binom{x}{y} \left(\frac{1}{2}\right)^x + q_1 \cdot \binom{x-2}{y-1} \left(\frac{1}{2}\right)^{x-2} + \dots + q_{\lfloor x/2 \rfloor} \cdot \binom{x-2\lfloor x/2 \rfloor}{y-\lfloor x/2 \rfloor} \left(\frac{1}{2}\right)^{x-2\lfloor x/2 \rfloor},$$

where $\binom{b}{c} := 0$ if $b < c$ or $c < 0$.

Note that q (and therefore in turn p) depends on the actual input to \mathcal{A} . Given any probability vector $q \in \mathbb{R}_+^{\lfloor x/2 \rfloor + 1}$, let $D(q)$ (or $D(q_0, \dots, q_{\lfloor x/2 \rfloor})$) denote the probability distribution $\{p(x, y)\}_{y=0, \dots, x}$. It is easy to verify that, for all q , $D(q)$ is a valid probability distribution that is symmetric on $y = \frac{x}{2}$.

Now we construct the bounding distribution. For $x = 0, \dots, k$, let $\alpha_x := (1 - \gamma_{\mathcal{A}})^{\max(x-1, 0)}$ for $\gamma_{\mathcal{A}} = 1/16$. We choose $D(\alpha_x, 1 - \alpha_x, 0, \dots, 0)$ as our bounding distribution. Let $p^*(x, \cdot)$ be the probability mass function of the bounding distribution, i.e., for each $y = 0, \dots, x$,

$$p^*(x, y) := \begin{cases} \alpha_x \left(\frac{1}{2}\right)^x, & \text{if } y = 0 \text{ or } y = x, \\ \alpha_x \binom{x}{y} \left(\frac{1}{2}\right)^x + (1 - \alpha_x) \binom{x-2}{y-1} \left(\frac{1}{2}\right)^{x-2}, & \text{otherwise.} \end{cases}$$

Following is a key lemma to show that $D(\alpha_x, 1 - \alpha_x, 0, \dots, 0)$ is a good choice of the centrally dominated distribution.

► **Lemma 10.** *Given a consecutive subsequence of pairs containing u of length x ($0 \leq x \leq k$), let $q \in \mathbb{R}_+^{\lfloor x/2 \rfloor + 1}$ be the probability vector realized by \mathcal{A} : i.e., q_i is the probability that exactly i edges of the ex-post graph have both endpoints in the given subsequence. We then have $\theta(x, p(x, \cdot)) \leq \theta(x, p^*(x, \cdot))$ and $\theta'(x, p(x, \cdot)) \leq \theta'(x, p^*(x, \cdot))$, where $p(x, \cdot)$ follows the distribution $D(q)$.*

Proof. For $x = 0, 1$, it is easy to see that $p(x, y) = p^*(x, y)$ for every possible y since $q_0 = 1 = \alpha_x$. For $x \geq 2$, we construct intermediate distributions as follows:

$$\begin{aligned} D_0 &:= D(q_0, 1 - q_0, 0, \dots, 0), \\ D_1 &:= D(q_0, q_1, 1 - q_0 - q_1, \dots, 0), \\ &\dots, \\ D_{\lfloor x/2 \rfloor - 1} &:= D(q_0, q_1, q_2, \dots, 1 - \sum_{i=0}^{\lfloor x/2 \rfloor - 1} q_i) = D(q). \end{aligned}$$

We claim that D_0 centrally dominates $D(\alpha_x, 1 - \alpha_x, 0, \dots, 0)$ and, for each $j = 1, \dots, \lfloor \frac{x}{2} \rfloor - 1$, D_j centrally dominates D_{j-1} . Then, by repeatedly applying Lemma 9, we can prove the lemma.

Let us first prove that D_0 centrally dominates $D(\alpha_x, 1 - \alpha_x, 0, \dots, 0)$. Since both distributions are symmetric, it is sufficient for us to consider $y \in [0, \frac{x}{2}]$. By Lemma 6, we have $q_0 \leq \alpha_x$. If $q_0 = \alpha_x$, D_0 is equivalent to $D(\alpha_x, 1 - \alpha_x, 0, \dots, 0)$, completing the proof. We now assume $q_0 < \alpha_x$. If $y = 0$, D_0 has the smaller probability. This implies that there must exist a point in $[1, \frac{x}{2}]$ where D_0 has the greater probability since both distributions are valid.

Let y^* be the smallest such point. We show that for any integer $y \in [y^*, \frac{x}{2}]$, D_0 has the greater probability. Let $g(y) := (q_0 - \alpha_x) \binom{x}{y} \left(\frac{1}{2}\right)^x + (\alpha_x - q_0) \binom{x-2}{y-1} \left(\frac{1}{2}\right)^{x-2}$ be the subtraction of the probability of D_0 from that of $D(\alpha_x, 1 - \alpha_x, 0, \dots, 0)$ at point y . It suffices to prove that $g(y)$ is non-decreasing on $[y^*, \frac{x}{2}] \subseteq [1, \frac{x}{2}]$. By rewriting the formula, we have $g(y) = (\alpha_x - q_0) \binom{x}{y} \left(\frac{1}{2}\right)^x \left(\frac{y(x-y)}{4x(x-1)} - 1\right)$, and it is not hard to see that the function is increasing on $[1, \frac{x}{2}]$.

A similar argument can be applied to show that, for each $j = 1, \dots, \lfloor \frac{x}{2} \rfloor - 1$, D_j centrally dominates D_{j-1} . Indeed, if $q_j = 1 - \sum_{i=0}^{j-1} q_i$, D_j is equivalent to D_{j-1} . We thus assume that $q_j < 1 - \sum_{i=0}^{j-1} q_i$. Observe that D_j has the same probabilities as D_{j-1} for $y = 0, \dots, j-1$, and has the smaller probability for $y = j$; hence, there exists a point in $[j+1, \frac{x}{2}]$ where D_j has the greater probability, and let y^* be the smallest such value. Again, let $g(y)$ be the difference obtained by subtracting the probability of D_{j-1} from that of D_j at point y . We can write

$$\begin{aligned} g(y) &= \left(1 - \sum_{i=0}^j q_i\right) \left[\binom{x-2j-2}{y-j-1} \left(\frac{1}{2}\right)^{x-2j-2} - \binom{x-2j}{y-j} \left(\frac{1}{2}\right)^{x-2j} \right] \\ &= \left(1 - \sum_{i=0}^j q_i\right) \binom{x-2j}{y-j} \left(\frac{1}{2}\right)^{x-2j} \left[\frac{(y-j)(x-y-j)}{4(x-2j)(x-2j-1)} - 1 \right], \end{aligned}$$

implying that $g(y)$ is increasing on $[y^*, \frac{x}{2}] \subseteq [j+1, \frac{x}{2}]$. ◀

We are now ready to bound the probability that our three-way OCS never chooses the element u from the consecutive subsequence of triples of length k . Let $\text{binom}(k, x, r)$ be the probability mass function of the binomial distribution for x successes out of k trials with probability r , i.e., $\text{binom}(k, x, r) := \binom{k}{x} r^x (1-r)^{k-x}$. Finally, we define $\eta(k)$ as follows and it will be the desired bound.

$$\eta(k) := \sum_{x=0}^k \text{binom} \left(k, x, \frac{2}{3} \right) \left[\sum_{y=0}^x p^*(x, y) \left(\frac{1}{2} \right)^{k-x+y} (1 - \gamma_{\mathcal{B}})^{\max(k-x+y-1, 0)} \right].$$

► **Lemma 11.** *Given a consecutive subsequence of triples containing u of length k , the probability that our three-way OCS never chooses u from the subsequence is at most $\eta(k)$.*

Proof. Conditioned on the pairs selected by the random pair choice phase, let x be the number of pairs containing u inserted into \mathcal{A} . Recall that these pairs form a consecutive subsequence. Let $p(x, y)$ denote the probability that \mathcal{A} returns y number of u 's from this subsequence. Note that p follows $D(q)$ for some q .

Note that, if \mathcal{A} returns y number of u 's, \mathcal{B} eventually takes $k - x + y$ consecutive pairs that originate from the given consecutive triples. By Lemma 7, we can see that the probability that u is never selected from the given consecutive subsequences of triples is bounded from above by

$$\sum_{y=0}^x p(x, y) \left(\frac{1}{2} \right)^{k-x+y} (1 - \gamma_{\mathcal{B}})^{\max(k-x+y-1, 0)},$$

which can be rewritten as $\left(\frac{1-\gamma_{\mathcal{B}}}{2} \right)^{k-x} \theta(x, p(x, \cdot))$ if $x < k$ or $\theta'(k, p(x, \cdot))$ if $x = k$. In any case, by Lemma 10, it can be further bounded by

$$\sum_{y=0}^x p^*(x, y) \left(\frac{1}{2} \right)^{k-x+y} (1 - \gamma_{\mathcal{B}})^{\max(k-x+y-1, 0)}.$$

It is noteworthy that this value depends only on the number x of u 's selected by the random pair choice phase. Observe that the probability that x pairs containing u are chosen from k consecutive triples by the random pair choice phase is exactly $\text{binom}(k, x, 2/3)$, completing our proof. ◀

It still remains to calculate the bounds on $\eta(k)$ in the same form as stated in Theorem 8. We refer the interested readers to [33].

4 General Bound

We now extend our discussion to the general case and provide a bound for a set of disjoint consecutive subsequences. The following theorem states this bound.

► **Theorem 2 (restated).** *Consider a set of m disjoint consecutive subsequences of triples containing an element u . Let k_1, \dots, k_m be the lengths of these subsequences. The probability that our three-way OCS never chooses u from these m subsequences is at most $\prod_{i=1}^m \eta(k_i)$.*

Let $\mathbf{k} = (k_1, \dots, k_m)$ be a vector whose i -th entry is the length of the i -th subsequence. In what follows, we fix (and condition upon) the choices made by the *random pair choice phase*. Let $\mathbf{x} := (x_1, \dots, x_m)$ be the (now constant) vector whose i -th entry represents the number of pairs containing u from the i -th subsequence that are passed to \mathcal{A} . For $\mathbf{y} := (y_1, \dots, y_m) \leq \mathbf{x}$, let $p(\mathbf{x}, \mathbf{y})$ be the conditional probability that \mathcal{A} chooses y_i number of u 's from the i -th subsequence of pairs for all i . Finally, let p_0 be the probability that our three-way OCS never chooses u . By Lemma 7, we can observe that p_0 is no greater than

$$\sum_{\mathbf{y} \leq \mathbf{x}} p(\mathbf{x}, \mathbf{y}) \prod_{i=1}^m \left(\frac{1}{2} \right)^{k_i - x_i + y_i} (1 - \gamma_{\mathcal{B}})^{\max(k_i - x_i + y_i - 1, 0)}. \tag{5}$$

To prove the desired bound, we will first present how we can modify the input to \mathcal{A} without ever decreasing (5). After finitely many modifications, we will be able to “decouple” the random choices across different subsequences.

Recall that each pair inserted into \mathcal{A} takes three independent random bits. The first random bit determines if the pair is a sender or a receiver. The second random bit decides which adjacent vertex to correlate. The third random bit selects an element to be output for this pair, unless the pair becomes matched as a receiver. Note that, with the first two random bits, the ex-post graph can be determined.

For each vertex j containing u in the ex-ante graph, let $\text{pred}(j, u)$ and $\text{pred}(j, -u)$ be the immediate predecessor of j linked by u and by the other element than u , respectively. Without loss of generality, let v be the other element of j than u . If we have j' and j , residing in different subsequences and j' appearing before j , such that

- (A) $j' = \text{pred}(j, u)$ or $j' = \text{pred}(j, -u)$ (or both); or
- (B) there exists \hat{j} such that
 - $\hat{j} = \text{pred}(j', u)$ or $\hat{j} = \text{pred}(j', -u)$,
 - $\hat{j} = \text{pred}(j, -u)$, and
 - \hat{j} is not contained in any subsequence,

we call them *violations*. In particular, the first type of violations are called Type A violations and the other Type B. Our goal is to modify the input so that there are no violations while (5) never decreased.

For notational simplicity, we let $\phi(y) := \left(\frac{1}{2}\right)^{k-x+y} (1 - \gamma_B)^{\max(k-x+y-1, 0)}$ for all $y = 0, \dots, x$. Observe that $\phi(y)$ is decreasing over y . We can now consider (5) as the expected value of $\prod_{i=1}^m \phi(y_i)$.

Removing Type A Violations. Let us first consider the case when $j' = \text{pred}(j, u) = \text{pred}(j, -u)$. We build the new input to \mathcal{A} by inserting $\bar{j}_1 = (u, \star)$ and $\bar{j}_2 = (v, \diamond)$ right before j is input where \star and \diamond are elements appearing nowhere else and $\star \neq \diamond$.

Let us fix some random bits of \mathcal{A} : one can think of this as further conditioning on those bits. In particular, we will fix the first two random bits of every pair except for \bar{j}_1 and \bar{j}_2 .

Suppose for now that the random bits dictate that neither $(j', j)_u$ nor $(j', j)_v$ is present in the ex-post graph of the original input. We claim that the probability distribution p conditioned on this event stays the same even after the modification and therefore the “contribution” to the expectation (5) is not affected either (or in other words, the conditional expectation remains equal).

We will show a stronger statement that the marginal probability distribution of the pairs in the subsequences stays the same. (Then the probability distribution p will automatically be the same.) Note that the ex-post graph of the two inputs will look almost identical, except for possible addition of some edges incident with \bar{j}_1 or \bar{j}_2 . Therefore, the output choice of every pair in the subsequences will be determined by the same random bit in both inputs, except that, it may be the case that j was not adjacent with any pair in the ex-post graph of the original input but becomes a receiver of \bar{j}_1 (or \bar{j}_2) after the modification. In this case, the output choice of j was determined by the third random bit of j before but by that of \bar{j}_1 (or \bar{j}_2) now. However, the only pair in the subsequence whose output choice is determined by the third random bit of \bar{j}_1 (or \bar{j}_2) is j . Moreover, the third random bits of the pairs are i.i.d.; hence, the marginal distribution of the pairs in the subsequences stays the same, as was claimed.

Now consider the remaining case where j' and j are adjacent in the ex-post graph of the original input. In this case, we will show that the conditional expectation cannot decrease (as opposed to staying the same). To ease the argument, we will fix more random bits (and

argue that the conditional expectation cannot decrease in all cases): we fix the third random bits of all pairs, except for j' , j , \bar{j}_1 , and \bar{j}_2 . Let z' be the number of u 's chosen from the subsequence containing j' , except for j' itself, in the original input. (Since we did not fix the third random bit of j' , we cannot determine what the output choice of j' is.) Similarly, let z be the number of u 's chosen from the subsequence containing j , except for j itself.

In the original input, note that the output choice of j' and j are both determined by the third random bit of j' . That is, with probability $1/2$, z' and $z+1$ respectively are the number of u 's chosen from the subsequences containing j' and j , and with probability $1/2$, $z'+1$ and z are. In the modified input, however, while we do not know whether j' is adjacent with \bar{j}_1 or \bar{j}_2 in the ex-post graph, \bar{j}_1 and \bar{j}_2 can each be adjacent with at most one pair in the ex-post graph. Therefore, in the marginal distribution of the pairs in the subsequences, the output choice of j' is independent from all other pairs, including j . This shows that the number of u 's chosen from the subsequence containing j' in the modified input is z' with probability $1/2$ and $z'+1$ with probability $1/2$, and independently from that, the number of u 's from the subsequence containing j is z with probability $1/2$ and $z+1$ with probability $1/2$.

Let us now calculate the increase of the conditional expectation of $\prod_{i=1}^m \phi(y_i)$, but since we are only interested in its sign, we will ignore the common terms, i.e., $\phi(y_i)$'s for the subsequences other than ones containing j and j' . Note that

$$\begin{aligned} & \frac{1}{4} [\phi(z')\phi(z) + \phi(z')\phi(z+1) + \phi(z'+1)\phi(z) + \phi(z'+1)\phi(z+1)] \\ & - \frac{1}{2} [\phi(z')\phi(z+1) + \phi(z'+1)\phi(z)] \\ = & \frac{1}{4} (\phi(z') - \phi(z'+1))(\phi(z) - \phi(z+1)) \geq 0, \end{aligned}$$

where the inequality follows from the fact that $\phi(y)$ is decreasing over y . This shows our claim.

Let us now consider the case where $j' = \text{pred}(j, u) \neq \text{pred}(j, -u)$ or vice versa. In this case, we insert $\bar{j} = (u, \star)$ or $\bar{j} = (v, \star)$ right before j , where \star again is a unique element. A similar argument shows that we can show that (5) can only increase in the new input.

Removing Type B Violations. We build the new input to \mathcal{A} by inserting $\bar{j} = (v, \star)$ right before j , where \star is a new unique element. Let us fix the first two random bits of every pair (including \bar{j}). We claim that the marginal distributions of the pairs in the subsequences are identical in both inputs.

Consider the ex-post graph of the two inputs. They will be almost identical only with the following possible differences:

- (\hat{j}, j) may be only in the original ex-post graph;
- (\hat{j}, \bar{j}) or (\bar{j}, j) may be only in the new ex-post graph.

For each pair x in the subsequences, the ex-post graph determines, among the third random bits of all pairs, which one determines the output choice of x .

Suppose (\hat{j}, j) is in the original ex-post graph. In this case, j is the only pair in the subsequences whose output choice is determined by the third random bit of \hat{j} . (Recall that \hat{j} is not in any subsequence.) In the modified input, if (\bar{j}, j) is in the ex-post graph, j will be the only pair in the subsequences whose output choice is determined by the third random bit of \bar{j} . If (\bar{j}, j) is not in the ex-post graph, j will be the only pair whose output choice is determined by j . Note that the output choice of every other pair in the subsequences is determined by the same random bit in both inputs. Again, since the third random bits of the pairs are i.i.d., this shows that the marginal distribution stays the same.

49:14 Making Three out of Two: Three-Way Online Correlated Selection

Suppose now (\hat{j}, j) is not in the original ex-post graph. Let us focus on j , since it is the only pair whose output choice may be determined by a different random bit in the modified input. If j was adjacent with another pair in the original ex-post graph, the edge will remain in the new ex-post graph, too, in which case there is nothing to prove. Otherwise, it may be the case that (\bar{j}, j) may be newly introduced as an edge in the ex-post graph, but this only means that j will be the only pair in the subsequences whose output choice is determined by the third random bit of \bar{j} (instead of j). The marginal distribution therefore stays the same.

Further Modification. We can remove all Type A and Type B violations after finitely many modifications of the input. We can interpret \mathcal{A} as an algorithm that first constructs the ex-ante graph using the first two random bits and then determines the output using only the third random bits. Our next modification will delete some edges directly from the ex-ante graph rather than modifying the input pairs. (You could alternatively think of this edge being “disabled,” which can never appear in the ex-post graph even if it is chosen by the first two random bits of the pairs.) We will now show how we can modify this input graph without affecting (5). Suppose we have some pair j' and j such that

- j' appears before j ,
- j' and j are adjacent in the ex-ante graph, and
- j is not in any subsequence.

If we delete this edge from the ex-ante graph, this can change the output choice of only j in \mathcal{A} 's output. (Imagine we fix all random bits, and we can easily observe this fact.) Therefore, we can safely delete all such edges without affecting the marginal distribution of the pairs in the subsequences.

Once we remove all such edges, the following argument shows that no two pairs in different subsequences can belong to the same connected component of the ex-ante graph: since two pairs that are not in any subsequences cannot be adjacent, the only way of having two pairs from different subsequences in the same connected component is by having a direct edge in-between (which would be a Type A violation) or a length-2 path whose “midpoint” is a predecessor of the two pairs (which would be a Type B violation: note that, in the definition of Type B violation, \hat{j} cannot be $\text{pred}(j, u)$ because j' appears before j).

Conclusion. Let $p'(\mathbf{x}, \mathbf{y})$ be the probability that \mathcal{A} chooses y_i number of u 's from the i -th subsequence of pairs for all i in the *modified* input. We have so far shown that

$$p_0 \leq \sum_{\mathbf{y} \leq \mathbf{x}} p'(\mathbf{x}, \mathbf{y}) \prod_{i=1}^m \left(\frac{1}{2}\right)^{k_i - x_i + y_i} (1 - \gamma_{\mathcal{B}})^{\max(k_i - x_i + y_i - 1, 0)}.$$

For each i , let $p'_i(x_i, y_i)$ be the probability that \mathcal{A} chooses y_i number of u 's from the i -th subsequence of pairs. Since \mathcal{A} may introduce negative correlation only on those pair of vertices that are adjacent in the ex-ante graph, output choices made across different connected components of the ex-ante graph will be independent. This implies that $p'(\mathbf{x}, \mathbf{y}) = \prod_{i=1}^m p'_i(x_i, y_i)$, yielding

$$\begin{aligned} p_0 &\leq \sum_{\mathbf{y} \leq \mathbf{x}} p'(\mathbf{x}, \mathbf{y}) \prod_{i=1}^m \left(\frac{1}{2}\right)^{k_i - x_i + y_i} (1 - \gamma_{\mathcal{B}})^{\max(k_i - x_i + y_i - 1, 0)} \\ &= \sum_{\mathbf{y} \leq \mathbf{x}} \prod_{i=1}^m p'_i(x_i, y_i) \left(\frac{1}{2}\right)^{k_i - x_i + y_i} (1 - \gamma_{\mathcal{B}})^{\max(k_i - x_i + y_i - 1, 0)} \end{aligned}$$

$$\begin{aligned}
&= \prod_{i=1}^m \sum_{y_i=0}^{x_i} p'_i(x_i, y_i) \left(\frac{1}{2}\right)^{k_i-x_i+y_i} (1-\gamma_{\mathcal{B}})^{\max(k_i-x_i+y_i-1,0)} \\
&\leq \prod_{i=1}^m \sum_{y_i=0}^{x_i} p^*(x_i, y_i) \left(\frac{1}{2}\right)^{k_i-x_i+y_i} (1-\gamma_{\mathcal{B}})^{\max(k_i-x_i+y_i-1,0)},
\end{aligned}$$

where the last inequality follows from the proof of Lemma 11. Since $p^*(x, \cdot)$ depends only on x , rather than the actual input, the bound on p_0 we obtain depends only on \mathbf{x} . Let $\text{binom}(\mathbf{k}, \mathbf{x}, r) := \prod_{i=1}^m \text{binom}(k_i, x_i, r)$. Now the probability that our three-way OCS never chooses u from the given subsequences is no greater than

$$\begin{aligned}
&\sum_{\mathbf{x} \leq \mathbf{k}} \text{binom}\left(\mathbf{k}, \mathbf{x}, \frac{2}{3}\right) \left[\prod_{i=1}^m \sum_{y_i=0}^{x_i} p^*(x_i, y_i) \left(\frac{1}{2}\right)^{k_i-x_i+y_i} (1-\gamma_{\mathcal{B}})^{\max(k_i-x_i+y_i-1,0)} \right] \\
&= \sum_{\mathbf{x} \leq \mathbf{k}} \left[\prod_{i=1}^m \text{binom}\left(k_i, x_i, \frac{2}{3}\right) \right] \left[\prod_{i=1}^m \sum_{y_i=0}^{x_i} p^*(x_i, y_i) \left(\frac{1}{2}\right)^{k_i-x_i+y_i} (1-\gamma_{\mathcal{B}})^{\max(k_i-x_i+y_i-1,0)} \right] \\
&= \prod_{i=1}^m \sum_{x_i=0}^{k_i} \text{binom}\left(k_i, x_i, \frac{2}{3}\right) \left[\sum_{y_i=0}^{x_i} p^*(x_i, y_i) \left(\frac{1}{2}\right)^{k_i-x_i+y_i} (1-\gamma_{\mathcal{B}})^{\max(k_i-x_i+y_i-1,0)} \right] \\
&= \prod_{i=1}^m \eta(k_i).
\end{aligned}$$

This completes the proof of Theorem 2.

5 Application to Online Bipartite Matching

Our three-way OCS can be applied to online bipartite matching problems. While the algorithm and the factor-revealing LP had to be generalized to a higher dimension, our algorithm and analysis are still analogous to Fahrback et al. [10]. We refer the interested readers to [33].

► **Theorem 12.** *There exists a 0.5096-competitive algorithm for unweighted online bipartite matching.*

► **Theorem 13.** *There exists a 0.5093-competitive algorithm for edge-weighted online bipartite matching with free disposal.*

References

- 1 Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Symposium on Discrete Algorithms (SODA)*, pages 1253–1264. SIAM, 2011.
- 2 Arash Asadpour, Michel X Goemans, Aleksander Mądry, Shayan Oveis Gharan, and Amin Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. *Operations Research*, 65(4):1043–1061, 2017.
- 3 Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. Edge weighted online windowed matching. In *Conference on Economics and Computation (EC)*, pages 729–742, 2019.
- 4 Benjamin E. Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1):80–87, 2008.
- 5 Guy Blanc and Moses Charikar. Multiway online correlated selection. *arXiv preprint*, 2021. To appear in Symposium on Foundations of Computer Science (FOCS 2021). [arXiv:2106.05579](https://arxiv.org/abs/2106.05579).

- 6 Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. New algorithms, better bounds, and a novel model for online stochastic matching. In *European Symposium on Algorithms (ESA)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- 7 Chandra Chekuri and Jan Vondrák. Randomized pipage rounding for matroid polytopes and applications. *CoRR*, abs/0909.4348, 2009. [arXiv:0909.4348](#).
- 8 Nikhil R. Devanur, Kamal Jain, and Robert D. Kleinberg. Randomized primal-dual analysis of RANKING for online bipartite matching. In *Symposium on Discrete Algorithms (SODA)*, pages 101–107. SIAM, 2013.
- 9 Alon Eden, Michal Feldman, Amos Fiat, and Kineret Segal. An economics-based analysis of ranking for online bipartite matching. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 107–110. SIAM, 2021.
- 10 Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. Edge-weighted online bipartite matching. In *Symposium on Foundations of Computer Science (FOCS)*, pages 412–423. IEEE, 2020.
- 11 Jon Feldman, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pál. Online ad assignment with free disposal. In *International Workshop on Internet and Network Economics (WINE)*, pages 374–385. Springer, 2009.
- 12 Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and Shan Muthukrishnan. Online stochastic matching: Beating $1-1/e$. In *Symposium on Foundations of Computer Science (FOCS)*, pages 117–126. IEEE, 2009.
- 13 Buddhima Gamlath, Sagar Kale, and Ola Svensson. Beating greedy for stochastic bipartite matching. In *Symposium on Discrete Algorithms (SODA)*, pages 2841–2854. SIAM, 2019.
- 14 Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online matching with general arrivals. In *Symposium on Foundations of Computer Science (FOCS)*, pages 26–37. IEEE, 2019.
- 15 Ruiquan Gao, Zhongtian He, Zhiyi Huang, Zipei Nie, Bijun Yuan, and Yan Zhong. Improved online correlated selection. *arXiv preprint*, 2021. To appear in Symposium on Foundations of Computer Science (FOCS 2021). [arXiv:2106.04224](#).
- 16 Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *Symposium on Discrete Algorithms (SODA)*, pages 982–991. SIAM, 2008.
- 17 Bernhard Haeupler, Vahab S Mirrokni, and Morteza Zadimoghaddam. Online stochastic weighted matching: Improved approximation algorithms. In *International Workshop on Internet and Network Economics (WINE)*, pages 170–181. Springer, 2011.
- 18 Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. Fully online matching. *Journal of the ACM*, 67(3):1–25, 2020.
- 19 Zhiyi Huang, Binghui Peng, Zhihao Gavin Tang, Runzhou Tao, Xiaowei Wu, and Yuhao Zhang. Tight competitive ratios of classic matching algorithms in the fully online model. In *Symposium on Discrete Algorithms (SODA)*, pages 2875–2886. SIAM, 2019.
- 20 Zhiyi Huang and Xinkai Shu. Online stochastic matching, poisson arrivals, and the natural linear program. In *Symposium on Theory of Computing (STOC)*, pages 682–693. ACM, 2021.
- 21 Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. *ACM Transactions on Algorithms*, 15(3):1–15, 2019.
- 22 Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Fully online matching ii: Beating ranking and water-filling. In *Symposium on Foundations of Computer Science (FOCS)*, pages 1380–1391. IEEE, 2020.
- 23 Zhiyi Huang, Qiankun Zhang, and Yuhao Zhang. Adwords in a panorama. In *Symposium on Foundations of Computer Science (FOCS)*, pages 1416–1426. IEEE, 2020.
- 24 Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *Journal of Algorithms*, 14(3):478–488, 1993.

- 25 Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Symposium on Theory of Computing (STOC)*, pages 352–358. ACM, 1990.
- 26 Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *European Symposium on Algorithms (ESA)*, pages 589–600. Springer, 2013.
- 27 Samir Khuller, Stephen G Mitchell, and Vijay V Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theoretical Computer Science*, 127(2):255–267, 1994.
- 28 Nitish Korula, Vahab S Mirrokni, and Morteza Zadimoghaddam. Bicriteria online matching: Maximizing weight and cardinality. In *International Conference on Web and Internet Economics (WINE)*, pages 305–318. Springer, 2013.
- 29 Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4):559–573, 2012.
- 30 Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.
- 31 Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):22, 2007.
- 32 Aranyak Mehta, Bo Waggoner, and Morteza Zadimoghaddam. Online stochastic matching with unequal probabilities. In *Symposium on Discrete Algorithms (SODA)*, pages 1388–1404. SIAM, 2015.
- 33 Yongho Shin and Hyung-Chan An. Making three out of two: Three-way online correlated selection. *arXiv preprint*, 2021. [arXiv:2107.02605](https://arxiv.org/abs/2107.02605).
- 34 Aravind Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *Symposium on Foundations of Computer Science (FOCS)*, pages 588–597. IEEE, 2001.