

Online Matching: A Brief Survey

ZHIYI HUANG

University of Hong Kong

and

ZHILIAO GAVIN TANG

Shanghai University of Finance and Economics

and

DAVID WAJC

Technion — Israel Institute of Technology

Matching, capturing allocation of items to unit-demand buyers, or tasks to workers, or pairs of collaborators, is a central problem in economics. Indeed, the growing prevalence of matching-based markets, many of which online in nature, has motivated much research in economics, operations research, computer science, and their intersection. This brief survey is meant as an introduction to the area of online matching, with an emphasis on recent trends, both technical and conceptual.

Categories and Subject Descriptors: F.2.2 [Nonnumerical Algorithms and Problems]: Matching Algorithms

General Terms: Algorithms

Additional Key Words and Phrases: Matching, Online Algorithms, Online Matching, Survey

1. INTRODUCTION

Matching theory lies at the heart of Economics, Computation and their intersection. Matching markets have played increasingly dominant roles in the world economy, both on the micro and macro scale. Such markets arise in domains as varied as Internet advertising, crowdsourcing of work and transportation, and organ transplantation. The repeated interactions and lack of certainty about future participants (buyers, sellers, etc.) result in these industries' dynamics being prime examples of *online* matching markets. See [Echenique et al. 2023] for detailed discussions. In this brief survey, we focus on three main aspects of recent developments in the study of such online matching and allocation problems.

Particularly prevalent are online *bipartite* matching markets. Often, agents on one side of the market arrive up front, while agents on the other side are revealed sequentially, to be matched (or not) immediately and irrevocably. For example, this dynamic abstracts the Internet advertising marketplace, with advertisers known up front and user queries (ad slots) revealed online. This motivates the study of online bipartite matching [Karp et al. 1990], and its generalization to weighted settings [Aggarwal et al. 2011; Feldman et al. 2009] and the AdWords problems [Mehta et al. 2007]. For more on the motivation from the Internet advertising application, see the influential survey on online bipartite matching and ad allocation by [Mehta 2013], and the more recent survey by [Devanur and Mehta 2022]. We outline recent

Authors' addresses: zhiyi@cs.hku.hk, tang.zhihao@mail.shufe.edu.cn, david.wajc@gmail.com

developments on these online bipartite matching and online ad allocation problems, in §2.

There are, of course, many aspects of modern online matching markets that are not bipartite, or that allow for agents on either side of bipartite matching markets to enter in an interleaved order. Similarly, while classic online matching models consider agents as having left the market after matching, in many crowdsourcing marketplaces (e.g., DoorDash, TaskRabbit, Uber/Lyft, etc.) freelance workers return to the market after being assigned a task (i.e., being matched) and completing their tasks. We discuss recent progress on modeling and addressing such problems, in §3.

The above-mentioned sections focus on the robust, but somewhat pessimistic, modeling choice of adversarial inputs and arrival orders. A less pessimistic model is that of random-order arrivals (“secretary models”), where the input is generated adversarially, but permuted by nature. Another modeling choice, motivated by the abundance of historical data from which to learn trends, is to posit a stochastic arrival model with parameters known to the algorithm. Here one can compare with either the best offline algorithm (computed by a “prophet” who knows the future) or the best online algorithm (computed by a “philosopher” who has enough time to think/compute). We discuss progress on online matching for such stochastic models, and their connection to mechanism design, in §4.

Finally, we give a brief glimpse of some overarching techniques that have played key recurring roles in the aforementioned recent developments, in §5. We illustrate some of the ideas with particularly short (and in our opinion, quite teachable) examples of some of the basic techniques in this space, in §A.

2. ONLINE MATCHING AND AD ALLOCATION

Researchers have made much progress in the past decade on the aforementioned online bipartite matching and ad allocation problems.

The most general problem along this line is *online submodular welfare maximization*. Consider a set of offline agents (advertisers), and a set of online items (impressions) that arrive one at a time. Each agent a has a submodular value function v_a over subsets of items; the algorithm can evaluate $v_a(S)$ for any subset S of the arrived items. On arrival of an online item, the algorithm must allocate it to an agent immediately and irrevocably. The basic benchmark is the greedy algorithm that allocates each impression to maximize the immediate increase in social welfare, which is $1/2$ -competitive. For the general problem, this ratio is optimal for polynomial-time algorithms under standard complexity-theoretic assumptions [Kapralov et al. 2013]. Most research has therefore focused on special cases of interest of this problem.

The (*unweighted*) *online bipartite matching* problem is the special case when each agent either likes or dislikes an item, and is willing to pay \$1 to get any one item they like: the value $v_a(S)$ equals 1 if agent a likes at least one item in S , and is 0 otherwise. Further, when a subset of items S is allocated to agent a , we can interpret this as matching a to any one item that a likes in S (e.g., the first one).

In online advertising, some advertisers may be able to pay more than others for an impression that they like. This can be captured by the *vertex-weighted*

generalization of online bipartite matching, where the value $v_a(S)$ is agent a 's weight w_a if agent a likes at least one item in S , and is 0 otherwise.

More generally, the same advertiser may have different values for different impressions, e.g., depending on the users' cookies and other information. This motivates considering edge weights instead of vertex weights; the advertiser only pays for one impression like in the unweighted and vertex-weighted case.¹ This is the *display ads* problem introduced by [Feldman et al. 2009], where $v_a(S)$ equals the largest edge weight w_{ai} among items $i \in S$.

Last but not least, some platforms let advertisers set a daily budget rather than a limit on the number of impressions. Given the allocation of impressions in a day, each advertiser a pays either the sum of its values for impressions it gets or its daily budget B_a , whichever is smaller; in other words, $v_a(S) = \min \{ \sum_{i \in S} w_{ai}, B_a \}$. Here, the weight w_{ai} is often referred to as agent a 's *bid* for impression i . This is the *AdWords* problem introduced by [Mehta et al. 2007].

Given the uncertainty over future items, when we decide how to allocate an online item, we do not want to put all our eggs in the same basket. It is easier to implement this old wisdom when the item is divisible (alternatively, if each agent has a large basket that can take many items). Imagine that each online item carries one litre of water (a divisible egg); each offline agent has a bucket (basket) of capacity one litre. The algorithm distributes an online item's water to its neighbors, where the amount of water going to each neighbor represents the fraction of the item allocated to the agent. The BALANCE algorithm (a.k.a., WATER LEVEL or WATER-FILLING [Azar and Litichevsky 2006]) lets the water flow to the least loaded bucket (the basket with the least amount of eggs); if there are multiple least loaded buckets, the water flows to them at an equal rate. This algorithm and its generalizations achieve the optimal $1 - 1/e$ competitive ratio in all the mentioned special cases of online submodular welfare maximization, including unweighted matching [Kalyanasundaram and Pruhs 2000], vertex-weighted matching [Buchbinder et al. 2007], display ads [Feldman et al. 2009], and AdWords [Mehta et al. 2007].

In the original problems where items are indivisible, we can distribute the risk through randomized decisions. However, making independent random decisions in each round is insufficient for getting a competitive ratio better than $1/2$, the baseline set by the greedy algorithm [Karp et al. 1990]. In the same paper that [Karp et al. 1990] introduced the online bipartite matching problem, they also gave an elegant RANKING algorithm achieving the optimal $1 - 1/e$ competitive ratio, later generalized by [Aggarwal et al. 2011] to vertex-weighted problem. The algorithm can be viewed as letting each offline vertex independently set a random price, and having each online vertex choose the lowest price offered by its unmatched neighbors. See §5 for a further discussion on this economic interpretation of RANKING and other online matching algorithms.

2.1 Breaking the $1/2$ Barrier in Longstanding Open Problems

Recall that for online (vertex-weighted) bipartite matching RANKING achieves an optimal competitive ratio, and in particular breaks the barrier of $1/2$. For display ads and AdWords, however, finding an online algorithm strictly better than the

¹Higher capacity can be simulated by creating multiple offline vertices per advertiser.

$1/2$ -competitive greedy algorithm had remained elusive for more than a decade. Fundamentally new ideas seemed necessary because a critical invariant in the analysis of RANKING fails to hold in these two problems.

In 2020, the $1/2$ barrier was broken for both problems using a new technique called Online Correlated Selection (OCS); see §5.4 for a further discussion on this technique. [Fahrback et al. 2022] introduced the concept of OCS and gave a 0.508-competitive algorithm for display ads. [Huang et al. 2020] modified the definition of OCS and applied it to the AdWords problem, and as a result, obtained a 0.501-competitive algorithm. The OCS technique has then been improved in a series of works by [Shin and An 2021], [Gao et al. 2021], and [Blanc and Charikar 2021]. The state-of-the-art competitive ratio for display ads is 0.536, given by a multi-way OCS algorithm by [Blanc and Charikar 2021].

Despite the aforementioned progress, we remark that there is no known evidence that the $(1 - 1/e)$ -competitive ratio cannot be achieved in display ads and AdWords. Hence, closing the gaps between the upper and lower bounds for these two problems remains an important open problem.

For the general online submodular welfare maximization problem, we recall that the simple greedy algorithm is $1/2$ competitive, and (barring surprising developments in complexity theory) no polynomial-time online algorithms can do better [Kapralov et al. 2013]. That being said, this impossibility relied on the computational hardness of maximizing a submodular function. It would be interesting to explore online algorithms with unlimited computational capacity, because practical heuristics can often solve these optimization problems better than the worst-case approximation ratio promises, and positive results along this line may point to other special submodular functions that are computationally tractable.

	Fractional/Divisible Relaxation	Original Problem
Unweighted	$1 - 1/e$ [Kalyanasundaram and Pruhs 2000]	$1 - 1/e$ [Karp et al. 1990]
Vertex-Weighted	$1 - 1/e$ [Buchbinder et al. 2007, Section 5]	$1 - 1/e$ [Aggarwal et al. 2011]
Display Ads	$1 - 1/e$ [Feldman et al. 2009]	0.536 [Blanc and Charikar 2021]
AdWords	$1 - 1/e$ [Mehta et al. 2007]	0.501 [Huang et al. 2020]

Table I. State-of-the-art for central online bipartite matching & allocation problems.

2.2 Stochastic Rewards and Oblivious Budget

Many online advertising platforms adopt the *pay-per-click* model. In this model, an advertiser pays each time a user clicks on its advertisement. Since these platforms cannot control the user’s behavior, they resort to the next best option: modeling a user’s behavior stochastically, and estimating the probability that the user clicks the advertisement, known as the *click-through-rate* (CTR). As a result of the users’ stochastic behavior, the platform’s revenue from assigning an impression to an advertiser is also stochastic.

[Mehta and Panigrahi 2012] introduced online matching with stochastic rewards. They analyzed both the RANKING algorithm and a variant of the BALANCE algorithm, and showed competitive ratios better than $1/2$ for uniform CTRs. That is, the CTR is either p or 0 (if the advertiser is not interested in this impression).

Progress made in the past decade on stochastic rewards is threefold. First, [Mehta et al. 2014] gave the first algorithm that breaks the $1/2$ barrier for non-uniform but sufficiently small CTRs. On one hand, small CTRs are arguably the most relevant case in practice because most keywords' CTRs are less than 10%. On the other hand, 10% or even 1% is larger than the assumption made by [Mehta et al. 2014] and its follow-up works. Hence, it remains an important open problem to design better online algorithms for less restrictive CTRs.

The second line of improvements comes from combining the online primal-dual framework and a more expressive linear program for the problem. This new analysis method gives a better understanding of classical algorithms BALANCE [Huang and Zhang 2020] and RANKING [Huang et al. 2023] in the presence of stochastic rewards.

Researchers have also tried to gain new insight by considering a weaker clairvoyant benchmark. [Goyal and Udhwani 2023] showed that against the weaker benchmark, RANKING achieves the optimal $1 - 1/e$ competitive ratio for uniform CTRs. They also analyzed BALANCE for small CTRs, and obtained a ratio better than the aforementioned state-of-the-art against the offline optimum benchmark. The latter result was later improved by [Huang et al. 2023] to 0.611.

Last but not least, stochastic rewards are closely related to budget-oblivious algorithms for AdWords, i.e., algorithms which do not know an agent's budget until the moment it is depleted. By a reduction by [Mehta et al. 2014], a competitive online algorithm for the latter model would yield the same competitive ratio in the former model (but not vice versa). Again, the greedy algorithm is a $1/2$ -competitive budget-oblivious algorithm for this problem. The survey by [Mehta 2013] listed finding a better budget-oblivious algorithm as an open problem. [Vazirani 2023] suggested a variant of RANKING as a candidate algorithm. [Liang et al. 2023] showed that no variant of this algorithm is $(1 - 1/e)$ -competitive. Finally, [Udhwani 2023] proved that the candidate is at least 0.508-competitive, and a variant of this algorithm is 0.522-competitive, both under the small-bids assumption, whereby the bids w_{ai} 's are small compared to the agent's budget B_a .

3. BEYOND ONLINE BIPARTITE MATCHING AND AD ALLOCATION

The preceding online matching models, largely motivated by online advertising, crucially rely on the assumption that one side of the bipartite graph is fixed and known upfront. This prevents the theory of online matching being applied to other modern applications, including ride-hailing, ride-sharing, rental services, etc.

In this section, we discuss generalizations of classic online bipartite matching. The first two generalizations are motivated by ride-hailing and ride-sharing, that allow all vertices to arrive online and allow general (non-bipartite) graphs. The third is somewhat theoretical in nature, but is the most general problem. The last generalization is motivated by rental services and freelance labor markets, and so captures the reusability of resources.

3.1 Fully Online Model: Vertices with Arrivals and Deadlines

In online ride-hailing platforms (e.g., Uber, Lyft, DiDi), ride requests are submitted to the platform in an online fashion and are active in the system for a few minutes. The platform assigns each request to a currently available taxi (or self-employed driver). Requests and taxis can be modeled as vertices in a bipartite

	Fractional Relaxation	Original Problem
Fully Online	0.6 [Tang and Zhang 2022]	0.569 [Huang et al. 2020]
General Vertex Arrival	0.526 [Wang and Wong 2015]	$1/2 + \Omega(1)$ [Gamlath et al. 2019]
Edge Arrival	$1/2$ [Gamlath et al. 2019]	$1/2$ [Gamlath et al. 2019]
Reusable Resources	$1 - 1/e$ [Goyal et al. 2021]	0.589 [DeLong et al. 2023]
	$1 - 1/e$ [Feng et al. 2021]	

Table II. State-of-the-art for online matching problems beyond bipartite matching & ad allocation.

graph with edges between compatible ride requests and taxis. This is an online bipartite matching problem but does not fit into the classic model, since all vertices (both the requests and the taxis) arrive online. Similarly, ride-sharing platforms, which match ride requests (pairing up riders) are naturally modeled as an online matching problem on *general (non-bipartite) graphs*.

[Huang et al. 2020] introduced the fully online matching model to capture the above scenarios, though the same model was studied earlier by [Blum et al. 2006] in the context of liquidity in clearing markets. Let $G = (V, E)$ be the underlying graph, initially *completely unknown*. Each time step is either an arrival or a deadline of a vertex. Upon the arrival of a vertex, its incident edges to their previously-arrived neighbors are revealed. A vertex can be matched at any point until its *deadline*, with this time revealed on its arrival. Naturally, we assume the deadline of a vertex is after its arrival, and all edges incident to a vertex are revealed before its deadline. This model generalizes the classic one-sided online bipartite matching model, where all offline vertices arrive at the beginning and have deadlines at the end, and every online vertex has its deadline right after its arrival.

For the fully online matching problem, [Huang et al. 2020] proved that RANKING achieves a tight $\Omega \approx 0.567$ (the unique solution to $\Omega \cdot e^\Omega = 1$) competitive ratio for bipartite graphs, and a competitive ratio of 0.521 for general graphs. For the fractional variant of the problem, [Huang et al. 2019] established a tight $2 - \sqrt{2} \approx 0.585$ competitive ratio of BALANCE. Remarkably, RANKING and BALANCE are known to be optimal in the classic model, but the claimed tightness here only applies to the two algorithms themselves. Indeed, [Huang et al. 2020] introduced the BALANCED-RANKING algorithm that achieves a competitive ratio of 0.569 for bipartite graphs, and the EAGER WATER-FILLING algorithm that achieves a competitive ratio of 0.592 for the fractional variant. The later result was further improved by [Tang and Zhang 2022] to 0.6. On the negative side, the state-of-the-art upper bound (i.e., hardness) is 0.613 [Huang et al. 2020; Eckl et al. 2021; Tang and Zhang 2022], separating the fully online model from the classic online bipartite matching model.

This setting is also known as the windowed online matching problem. [Ashlagi et al. 2023] assumed a first-in-first-out structure on the active windows (i.e., arrivals and deadlines) of vertices, and achieved a $1/4$ competitive ratio for edge-weighted graphs through a reduction to the Display Ads problem by suffering an extra factor of 2. Combined with the state-of-the-art algorithm for the Display Ads problem by [Blanc and Charikar 2021], their competitive ratio can be improved to 0.268.

3.2 General Vertex Arrival

Generalizing fully online matching is the online matching with general vertex arrivals problem, introduced by [Wang and Wong 2015]. Again, the input is a graph $G = (V, E)$, initially unknown, with vertices arriving online. Upon the arrival of a vertex v , its incident edges to its previously-arrived neighbors are revealed. The algorithm either matches v to an unmatched neighbor immediately or leaves v unmatched, possibly matching it to a later-arriving neighbor u upon u 's arrival. The inability to match vertices at any point before their departure (and lack of this information) makes this model more restrictive than the fully online model, and so algorithms for general vertex arrivals are also algorithms in the fully online model, with the same competitive ratio.

[Wang and Wong 2015] presented a fractional 0.526-competitive algorithm for the fractional version of the problem. [Gamlath et al. 2019] designed a rounding of Wang and Wong's fractional algorithm and established a $1/2 + \Omega(1)$ competitive ratio for the integral matching problem. This result stands as the only non-trivial integral algorithm so far. On the negative side, [Wang and Wong 2015; Buchbinder et al. 2019; Tang et al. 2022] established an upper bound of 0.583, separating the general vertex arrival model from the fully online model.

3.3 Edge Arrivals

Finally, we remark that the most general online matching setting is the edge arrival model. That is, edges of an underlying graph arrive in a sequence and the algorithm decides whether to select an edge immediately on its arrival. Here edges correspond to fleeting collaboration opportunities between agents. A competitive ratio of $1/2$ can be trivially achieved by a greedy algorithm that matches each edge on arrival if both its endpoints are free, and this is optimal for deterministic algorithms. Unfortunately, [Gamlath et al. 2019] proved that no online algorithm achieves a better than $1/2 + 1/2^n$ competitive ratio, even for the fractional version of the problem. Positive results are known assuming structure, including low-degree graphs and trees [Buchbinder et al. 2019], batching [Lee and Singla 2017], random-order arrivals [Guruganesh and Singla 2017], or stochastic arrivals [Gravin et al. 2021] (see §4 for more on the latter models).

3.4 Reusable Resources

In sponsored search, the advertisers' budgets, viewed as resources, are non-reusable. In contrast, in such markets as cloud computing (e.g., AWS, Azure), short-term rentals (e.g., Airbnb), and freelancer labor (e.g., TaskRabbit), the allocated resources (be it compute, housing or labor) are reusable, and can be reallocated after being used.

The above motivates online bipartite matching with reusable resources, where after an offline vertex (a rental service) is matched, it becomes available again after d time steps, where d is a known parameter that corresponds to the usage duration of the vertices. The classic online bipartite matching problem is a special case of the reusable resources model when $d = \infty$.

This model was first introduced by [Gong et al. 2022] in a more general setting of online assortment optimization. [Goyal et al. 2021; Feng et al. 2021] generalized

the BALANCE algorithm and achieves an optimal $1 - 1/e$ competitive ratio for the fractional version of the problem.² For the integral version of the problem, [DeLong et al. 2023] proposed the PERIODIC RERANKING algorithm (a variant of RANKING that reranks the offline vertices every d time steps), and show that it achieves a competitive ratio of 0.589, and an online correlated selection-based algorithm achieves a competitive ratio of 0.505. All these results extend to the vertex-weighted setting. We remark that the results of [DeLong et al. 2023] heavily rely on the assumption that all vertices have identical usage durations d . The case of heterogeneous usage durations (i.e., each vertex v has an individual duration time d_v) remains open.

[Feng et al. 2022] further studied online assortment of reusable resources in the stochastic setting. Reusable resources due to additional production have also been considered in infinite-horizon *stochastic* settings [Aouad and Saritaç 2020; Collina et al. 2020; Kessel et al. 2022; Patel and Wajc 2024]. We discuss stochastic settings (without reusable resources) in more detail in the following section.

4. STOCHASTIC MODELS: SECRETARIES, PROPHETS, AND PHILOSOPHERS

The preceding sections focused on adversarial models, where both input graph and arrival order are chosen by an adversary. This modeling choice, while robust, is quite pessimistic, and naturally results in worse guarantees than possibly achievable for real-world applications of interest. A natural way to obtain improved provable guarantees is to either consider random arrival orders (but adversarial input), or to posit a stochastic generative model, possibly learnt from historical data. Such models hearken back to classic results in optimal stopping theory concerning online Bayesian selection problems.

In the most basic setting, a buyer has a single item to sell, and impatient buyers arrive one after another and make take-it-or-leave-it bids for this single item. The buyer must select which bid to accept, immediately and irrevocably when the bid is made. Under adversarial models, a buyer cannot be competitive with the hindsight-optimal solution. In contrast, if the bids arrive in random order (referred to as the *secretary problem*), then a competitive ratio of $1/e$ is optimal [Dynkin 1963]. Similarly, if the successive bids v_i are drawn independently from known distributions \mathcal{D}_i , then the optimal competitive ratio is $1/2$ [Krengel and Sucheston 1978], i.e., the buyer can guarantee an expected gain at least half of that obtained by a “prophet” who knows the realization of the randomness,

$$\mathbb{E}[\text{Gain}] \geq \frac{1}{2} \cdot \mathbb{E}[\max_i v_i].$$

Such guarantees contrasting with the offline optimal, or prophet, are referred to as *prophet inequalities*. One may also contrast with the (computationally-unbounded) optimal online algorithm for such problems, which for reasons elaborated below we refer to as *philosopher inequalities*.

These models have been generalized and extended significantly over the years. In this section, we focus on recent developments for generalizations of the above to bipartite matching markets, where the buyer wishes to sell multiple heterogeneous items, and each arriving buyer proposes a different bid for each item. Put other-

²Equivalently, they assumed that resources (offline vertices) have large capacities.

wise, we focus on online bipartite matching models. We note that the buyer and seller terminology are not accidental, and these models have tight connections to questions in *mechanism design*, which we also discuss in this section.

	impossibility	algorithmic
Secretary Matching	$1/e$ [Dynkin 1963]	$1/e$ [Kesselheim et al. 2013]
Prophet Matching	$1/2$ [Krengel and Sucheston 1978]	$1/2$ [Feldman et al. 2015]
Philosopher Matching	0.99999 [Papadimitriou et al. 2021]	0.652 [Naor et al. 2023]

Table III. State-of-the-art for online bipartite weighted matching in stochastic settings.

4.1 Secretary Problems

For edge-weighted online bipartite matching with online vertices arriving in random order, [Korula and Pál 2009] were the first to obtain a constant-competitive ratio, specifically a $1/8$ -competitive algorithm. This was later improved to the optimal $1/e$ ratio by [Kesselheim et al. 2013], generalizing the classic single-item result of [Dynkin 1963], which we recall is the special case of a single offline vertex. For k heterogeneous offline vertices, [Kleinberg 2005] showed $(1 - 1/\sqrt{k})$ -competitive and truthful mechanism. [Ezra et al. 2022a] study secretary matching in general graphs (with vertices arriving with edges to their previously-arrived neighbors, as in §3.1 and §3.2). They show that the optimal competitive ratio in this level of generality is $5/12$, notably strictly greater than achievable for bipartite graphs.³

The random-order model similarly allows for improved guarantees for the special cases of vertex-weighted and unweighted online bipartite matching. For unweighted matching, [Guruganesh and Singla 2017] show that for *edges* revealed in random order a better than $1/2$ competitive ratio is possible, notably beyond the worst-case optimal for adversarial arrivals [Gamlath et al. 2019]. Similarly, random-order vertex arrivals in bipartite graphs (with arrivals on only one side of the graph) allow one to surpass the worst-case optimal $1 - 1/e$: a generalization of RANKING achieves a competitive ratio of 0.662 [Huang et al. 2019; Jin and Williamson 2021]. This generalizes results of [Karande et al. 2011; Mahdian and Yan 2011], who showed that for the unweighted problem RANKING (unchanged) achieves competitiveness beyond $1 - 1/e$, with the best known bound being 0.696 [Mahdian and Yan 2011]. As noted by these last two works, these results for random-order arrivals imply the same competitive ratios for stochastic matching problems with unknown i.i.d. distributions over arrival types. This remains the best known result for unknown distributions. In the following sections, we discuss the types of guarantees achievable under *known* distributions.

³For some intuition as to why this is not a contradiction, note that in the star example (i.e., the single-item problem), the center of the star arrives after the highest-bidding neighbor with probability $1/2 > 1/e$, and so greedily matching the center when it arrives is $1/2$ -competitive, and the lower bound of [Dynkin 1963] for bipartite graphs does not apply.

4.2 Prophet Inequalities

The optimal competitive ratio of $1/2$ for the single-item problem due to [Krengel and Sucheston 1978] was also obtained several years later using a single-threshold (i.e., posted-price) algorithm by [Samuel-Cahn 1984]. This results in truthful welfare-approximating mechanisms for single-item auctions. This connection between (pricing-based) prophet inequalities and mechanism design was later elaborated upon by researchers at the intersection of Economics and Computation [Hajjaghayy et al. 2007; Chawla et al. 2010; Kleinberg and Weinberg 2019]. Interestingly, very recently [Banihashem et al. 2024] show that *any* guarantee achieved by an online Bayesian selection algorithm can be achieved by a (dynamic) posted-price policy, implying that studying the non-strategic setting results in truthful mechanisms which achieve the same approximation of the social welfare as the algorithm in strategic settings.

The connection between (combinatorial) prophet inequalities and mechanisms design continues to motivate a flurry of results on prophet inequalities for increasingly involved markets, with more and more sophisticated combinatorial constraints on the sets of buyers that may be serviced, or items sold. See the excellent surveys [Hill and Kertz 1992; Correa et al. 2018] and [Hartline 2012; Lucier 2017] for more on prophet inequalities and their connection to mechanism design, respectively. In what follows, we focus on prophet inequalities subject to matching constraints.

For unweighted online bipartite matching, [Feldman et al. 2009] were the first to show that stochastic inputs allow for competitive ratio beyond the worst-case optimal $1 - 1/e \approx 0.632$. Specifically, they show that if online vertices' neighborhoods are drawn i.i.d. from a single known distribution, then a competitive ratio of 0.67 is achievable. There has been a long line of work studying this question, most recently [Jaillet and Lu 2013; Brubach et al. 2021; Huang and Shu 2021; Huang et al. 2022; Tang et al. 2022], with the current best competitive ratios being 0.7299 and 0.716 assuming integral and arbitrary arrival rates [Brubach et al. 2021; Huang et al. 2022].⁴ For edge-weighted matching, a number of results were obtained under integral arrival rates [Haeupler et al. 2011; Brubach et al. 2021], with the best ratio standing at 0.704, while for arbitrary arrivals the ratio of $1 - 1/e$ was only recently beaten [Yan 2024; Qiu et al. 2023]. In contrast, by a work of [Manshadi et al. 2012], no competitive ratios greater than $1 - 1/e^2 \approx 0.864$ and 0.823 are possible in the same settings.

For unweighted and vertex-weighted bipartite matching under (much more general) time-varying independent distributions, [Tang et al. 2022] recently provided the first algorithm surpassing the competitive ratio of $1 - 1/e$, presenting a 0.666-competitive algorithm. For edge-weighted matching a competitive ratio of $1/2$ is best possible, as this generalizes the single-item problem of [Krengel and Sucheston 1978]. This ratio is known to be achievable via numerous approaches [Feldman et al. 2015; Dütting et al. 2020; Ezra et al. 2022b]. This ratio of $1/2$ is even achievable under vertex arrivals in general graphs [Ezra et al. 2022b], or correlated arrivals [Aouad and Ma 2023]. In contrast, the problem is strictly harder under edge arrivals, where the best known competitive ratio is in the range $[0.344, 0.4]$ [MacRury et al. 2023]

⁴The arrival rate is the expected number of arrivals of a particular online type.

and $[0.349, 3/7]$ for bipartite graphs [MacRury et al. 2023; Correa et al. 2023]. For unweighted matching a competitive ratio of 0.502 is possible [Gravin et al. 2021].

4.3 Philosopher Inequalities

While a competitive ratio of $1/2$ is worst-case optimal for online Bayesian selection subject to bipartite matching constraints, this is still a pessimistic worst-case guarantee, as the lower bounds focus on worst-case distributions. The optimal algorithms for distributions of interest may allow for better competitive ratios. This optimal algorithm, which is the solution of a Markov Decision Problem (MDP), is computable in polynomial space via standard techniques. As shown by [Papadimitriou et al. 2021], this is the right characterization, and even *approximating* the optimal policy beyond some 0.999 ratio is PSPACE-complete (i.e., is as hard as the hardest problems requiring polynomial space). Hence, under standard complexity-theoretic assumptions, this optimal policy is not computable in polynomial time. Put otherwise, it is likely computable only by a character with sufficient time to “think” (i.e., compute), therefore naturally referred to as a “philosopher”. This motivates the study of polynomial-time approximation of the optimal online algorithm, which, in analogy with prophet inequalities (approximation of the optimal offline algorithm), we term *philosopher inequalities*.

[Anari et al. 2019] were the first to consider the approximation of the optimal online algorithm for online Bayesian selection. They considered bounded-depth and production-constrained laminar matroids, for which they provided $(1 + \epsilon)$ -approximate philosopher inequalities for any constant $\epsilon > 0$. [Dütting et al. 2023] obtained the same bounds for random-order (secretary) philosopher inequalities for a single item. For online bipartite matching, by [Papadimitriou et al. 2021] such an approximation would result in surprising developments in complexity theory. On the positive side, a successive line of work [Papadimitriou et al. 2021; Saberi and Wajc 2021; Braverman et al. 2022; Naor et al. 2023] showed that (increasingly) better than $1/2$ -approximate philosopher inequalities are possible. The current best bound stands at 0.652 (notably, above the natural bound of $1 - 1/e$ for online matching algorithms). We note that the recent work of [Banihashem et al. 2024] also translates the above (polynomial-time) policies approximating the optimal policy into pricing-based (and hence truthful) mechanisms providing the same approximation of the optimal mechanism.

5. OVERARCHING TECHNIQUES

5.1 Primal-Dual Algorithms

The primal-dual method has found wide applications in the area of online algorithms. Refer to [Buchbinder and Naor 2009] for a comprehensive survey. For online matching and related problems, the primal-dual schema was first adapted by [Buchbinder et al. 2007] to analyze BALANCE for the AdWords problem. We illustrate this idea in the special case of (unweighted) online bipartite matching.

We start with an economic interpretation of BALANCE. Consider the offline vertices as divisible items, and the online vertices as 0-1 unit-demand buyers. At any moment, each offline vertex v *prices* itself at $g(x_v)$ per (fractional) unit based on the current water level x_v (matched fraction of v) and thus its neighbor receives a *utility*

of $1 - g(x_v)$ per unit of v (fractionally) assigned to it, where $g(\cdot)$ is an increasing function. (Note that the price of v increases over time.) Upon its arrival, an online vertex u continuously chooses the unmatched neighbors giving u the largest utility. Recall that the dual of the maximum matching problem is the minimum vertex cover problem. The economic interpretation suggests a natural way to set the dual variables: for each offline vertex, let its dual variable be the total collected price, and for each online vertex, let its dual variable be the utility. The primal-dual framework asserts that in order to establish a Γ competitive ratio, it suffices to prove that the total gain of each item-buyer pair is Γ . To illustrate this approach, we provide a formal yet brief analysis in Appendix A.

Almost all fractional online matching algorithms were analyzed within the primal-dual framework. This includes algorithms for AdWords [Buchbinder et al. 2007], Display Ads [Devanur et al. 2016], fully online matching [Huang et al. 2019; Huang et al. 2020; Tang and Zhang 2022], general vertex arrivals [Wang and Wong 2015], stochastic matching [Tang et al. 2022], etc.

The primal-dual method for fractional matching crucially requires the dual constraints to be satisfied *always*. In contrast, [Devanur et al. 2013] noticed that it suffices to have the dual constraints hold *in expectation* for randomized algorithms, and used this observation to provide a simplified competitive analysis of RANKING for online (vertex-weighted) bipartite matching. Their approach is now referred to as the *randomized* primal-dual schema.

Their proof relied on an intuitive economic interpretation of RANKING that is similar to the economic interpretation of BALANCE. Instead of maintaining a dynamic price that depends on the water level, each offline vertex sets a randomized fixed price (according to the random permutation generated by RANKING) at the beginning. Then on the arrival of each online vertex, it buys the cheapest remaining neighbor. Again, we split the gain of each matched edge between its two endpoints (i.e., set the corresponding dual variables), according to the price of the offline vertex and the utility of the online vertex. For our EC readers, refer to [Eden et al. 2021] for a proof that is written explicitly in the language of price and utility and avoids duality.

A remarkable property of the randomized primal-dual schema is its intrinsic robustness for vertex-weighted graphs for all variants of online bipartite matching. Indeed, this schema often (if not always) provides a “free lunch”, allowing one to extend a result on unweighted graphs to vertex-weighted graphs while preserving the same competitive ratio. E.g., [Devanur et al. 2013; Huang and Zhang 2020; Huang and Shu 2021; Huang et al. 2022; Tang et al. 2022]. Going beyond the online bipartite matching model, in the fully online matching model [Huang et al. 2020] further developed the randomized primal-dual schema, by introducing a novel charging mechanic that allows a vertex other than the two endpoints of a matched edge to share the gain. [Levin and Wajc 2021] further found an application of the randomized primal dual framework for submodular maximization.

5.2 Randomized Rounding and Contention Resolution Schemes

The relax-and-round framework considers fractional relaxations as guides for randomized algorithms’ probabilistic choices. This section discusses the prevalence of this approach for online matching problems.

For bipartite matching, the standard relaxation allows us to assign a fractional value $x_e \in [0, 1]$ to each edge so that any node v has at most one unit assigned to its edges, $\sum_{e \ni v} x_e \leq 1$. Intuitively, x_e can be thought of as the marginal matching probability of edge e by some randomized algorithm, and use these fractions to obtain randomized algorithms. Indeed, since every fractional bipartite matching is the convex combination of integral matchings M_1, \dots, M_k , this intuition can be made formal, by randomly picking one such matching M_i with probability equal to its coefficient in the convex combination. This results in each edge e being matched with probability x_e , and thus preserves any linear objectives, $\sum_e w_e \cdot x_e$. We refer to such rounding schemes matching each edge with probability x_e as *lossless* rounding schemes.

Perhaps surprisingly, and as hinted at by Table I, for *online* edge-weighted bipartite matching and ad allocation problems, there exists a gap between our understanding of fractional algorithms and indivisible randomized algorithms. Indeed, as pointed out by [Devanur et al. 2013, Footnote 3], the above-mentioned integrality does not carry over to the online setting: for every randomized algorithm, there exist graphs on which the (optimally) $(1 - 1/e)$ -competitive fractional algorithm BALANCE achieves value $8/7$ times higher than any randomized algorithms. Therefore, rounding fractional algorithms seems to require losing a large multiplicative factor (in the worst case). At face value, this large gap seems to rule out the use of the relax-and-round approach to obtain good randomized (integral) online matching algorithms.

Despite the above, a large number of results in online matching in recent years are obtained by (or can be interpreted as employing) online randomized rounding of fractional solutions, often obtained using the primal-dual schema, §5.1. See [the FOCS23 workshop on the topic](#). There are three flavors of results in this vein.

5.2.1 Lossless rounding. The Online Correlated Selection (OCS) technique, elaborated upon in §5.4 [Fahrbach et al. 2022; Gao et al. 2021; Blanc and Charikar 2021] can be seen as losslessly rounding *structured* fractional bipartite matchings in online settings. [Buchbinder et al. 2023] were the first to explicitly ask what structure is *necessary* for lossless online rounding of bipartite matching, i.e., allowing one to match each edge with probability *exactly* x_e . They considered other structured fractional online matching algorithms and provided lossless online rounding schemes for these, which they used to obtain generalizations of OCS and sharp randomness thresholds for beating deterministic algorithms for online bipartite matching. Similarly, “spread out” fractional matchings, e.g., ones assigning value $1/\Delta$ in graphs of maximum degree Δ , can be rounded *nearly* losslessly, i.e., one can match each edge e with probability $x_e \cdot (1 - \epsilon)$, and this is key to numerous results for online edge coloring, e.g., [Cohen and Wajc 2018; Wajc 2020; Blikstad et al. 2024].

5.2.2 Approximate Rounding. In the other extreme, [Naor et al. 2023] ask how well *arbitrary* fractional matchings \vec{x} can be rounded online, and provide approximate rounding schemes that match each edge e with probability $0.652 \cdot x_e$, notably breaking the barrier of $1 - 1/e$ for this problem. They then use this scheme to obtain improved results for online edge coloring of multigraphs and philosopher inequalities, among others. Some prior results for philosopher inequalities [Papadim-

itriou et al. 2021; Saberi and Wajc 2021] are also obtained by such approximate online rounding schemes applied to LP relaxations incorporating constraints only applicable to online algorithms [Torrìco and Toriello 2022; Buchbinder et al. 2014]. Similarly, the multiway OCS of [Gao et al. 2021] can be seen (and used) as an approximate rounding scheme that provides guarantees per offline vertex, as opposed to per edge, lending itself to results for vertex-weighted matching. The result of [Gamlath et al. 2019] for online matching under general vertex arrivals likewise follows a lossy rounding approach (applied to the fractional algorithm of [Wang and Wong 2015]), though here the approximation guarantees are more global than per-vertex or per-edge. More approximate rounding schemes are obtained by Online Contention Resolution Schemes (OCRS), whose guarantees are weak in the context of adversarial settings, but are central to prophet inequalities, as we now discuss.

5.3 Online Contention Resolution Schemes

Contention resolution schemes (CRS) have their origins in the (offline) submodular optimization literature [Chekuri et al. 2014], and follow a natural rounding approach: *activate* each element independently with probability x_e , and then select a high-valued feasible subset (in our case, a matching) among the active elements. This is obtained by guaranteeing each active element be selected with as high a probability as possible. This probability $\Pr[e \text{ selected} \mid e \text{ active}]$ is referred to as the *balance ratio* of the CRS.

The above approach can be generalized to online settings [Feldman et al. 2016], where inclusion of active element (in our case, edge) e must be made immediately upon its activation. By an approach due to [Yan 2011], using an appropriate convex *ex-ante relaxation*, an OCRS with balance ratio c provides c -approximate prophet inequalities for the same setting, by considering an element (buyer) active if their bid is in their x_e -th percentile. As [Lee and Singla 2018] show, the opposite is true: prophet inequalities that are c -competitive with respect to this relaxation yield OCRS that are c -balanced. By preceding discussions on pricing-based prophet inequalities being derivable from arbitrary prophet inequalities, we find that OCRS yield welfare-approximating truthful mechanisms [Banihashem et al. 2024]. Generally, OCRS have found widespread applications since their introduction. See [Patel and Wajc 2024] for a discussion.

When specified to matchings, for batched OCRS (capturing vertex arrivals), balance ratios of $1/2$ and $(1 + 1/e^2)/2 \approx 0.567$ are optimal for adversarial vertex arrivals in general graphs [Ezra et al. 2022b] and random-order vertex arrivals in bipartite graphs [MacRury and Ma 2024], respectively. For other settings, despite much progress in recent years, the optimal balance ratio is still unknown for adversarial edge arrivals [Gravin and Wang 2019; Correa et al. 2023; MacRury et al. 2023], random-order edge arrivals [Brubach et al. 2021; Pollner et al. 2022; MacRury et al. 2023] and random-order vertex arrivals [Fu et al. 2021; MacRury and Ma 2024].

5.4 Online Rounding: Online Correlated Selection

In this section we elaborate on one particular general rounding scheme for online bipartite matching and its generalizations, termed Online Correlated Selection (OCS). In each round, the OCS observes the online item *and a fractional allocation of this item* to the offline agents. It must then allocate this item in whole to one

of the agents. We measure the quality of an OCS by how an agent’s value for the allocated subset of items depends on the fractional allocation to the agent.

Consider unweighted matching as a running example. We expect that an offline agent gets matched with a higher chance as the total fractional allocation to the agent increases. Consider the following natural baseline algorithm: match each item i to an unmatched agent a interested in the item with probability proportional to x_{ai} , the fractional allocation of item i to agent a . Let x_a denote the total fractional allocation to agent a . This rounding algorithm guarantees that agent a is matched with probability at least $1 - e^{-x_a}$, and this bound is tight (c.f., [Gao et al. 2021]). Unfortunately, combining the Balance algorithm and this baseline bound only gives the trivial 0.5 competitive ratio. See [Gao et al. 2021] for the state-of-the-art techniques for designing OCS for unweighted and vertex-weighted matching, which improved the above bound for the probability of matching agent a , and the resulting competitive ratios for these two problems. [Hosseini et al. 2023] further applied the OCS technique to an unweighted online matching problem with class fairness as its main objective.

The definitions of OCS for display ads and AdWords are too technical to be covered concisely in this brief survey. We refer readers to [Fahrbach et al. 2022] for the original definition of OCS for display ads and a proof of concept that non-trivial OCS exists, and [Blanc and Charikar 2021] for the best existing result along this line. See [Huang et al. 2020] for the definition of OCS for AdWords; the technique is similar to the proof of concept for display ads by [Fahrbach et al. 2022].

Although the concept of OCS was first introduced for problems in non-stochastic models, the technique has found applications in stochastic models as well. In hindsight, this is perhaps unsurprising. The general recipe for online matching algorithms in stochastic models is to first solve an LP relaxation and then make online matching decisions taking the LP solution as a guide. By design, an OCS can treat the LP solution as a fractional allocation and convert it into online matching decisions. [Tang et al. 2022] used this approach to obtain the first non-trivial algorithm for the non-IID model of unweighted and vertex-weighted online stochastic matching. [Huang et al. 2022] further gave an OCS tailored for the stochastic model to get the best competitive ratios to date for IID unweighted and vertex-weighted online stochastic matching.

Acknowledgements

Zhiyi Huang is supported by an NSFC grant 6212290003. Zhihao Gavin Tang is supported by NSFC grant 61932002. David Wajc is supported in part by a Taub Family Foundation “Leader in Science and Technology” fellowship. We thank Niv Buchbinder for his short analysis of the BALANCE algorithm (in the appendix), and Sam Taggart for valuable feedback.

A. ADDENDUM: A TEACHABLE MOMENT

Following a quote attributed to Feynman, namely “If you want to master something, teach it,” we present short and self-contained (and in our opinion, quite teachable) proofs of two basic results in online bipartite matching: a competitive ratio of $1 - 1/e$ for BALANCE extended to vertex-weighted matching [Buchbinder et al. 2007], and an even shorter proof that no fractional algorithm can do better.

Recall that for the online vertex-weighted bipartite matching problem, each offline vertex i and its positive weight $w_i \geq 0$ are known up front. At each time t , online vertex t arrives, together with its edges $(i, t) \in E$ to its neighbors $i \in N(t)$, and we must decide to what extent $x_{i,t}$ to assign t to its neighbor i , from which the algorithm accrues a value of $w_i \cdot x_{i,t}$. Both offline vertices and online vertices must be assigned to a total extent of at most one.

A linear programming (LP) relaxation of the problem (allowing us to match each edge (i, t) to an extent of $x_{i,t}$), together with this LP's dual, are as follows.

$$\begin{array}{ll}
 \text{(P) max} & \sum_{(i,t) \in E} w_i \cdot x_{i,t} \\
 \text{s.t.} & \sum_t x_{i,t} \leq 1 \quad \forall i \\
 & \sum_i x_{i,t} \leq 1 \quad \forall t \\
 & x \geq 0 \\
 \text{(D) min} & \sum_i y_i + \sum_t z_t \\
 \text{s.t.} & y_i + z_t \geq w_i \quad \forall (i, t) \in E \\
 & y, z \geq 0
 \end{array}$$

A.1 Short analysis of BALANCE

The BALANCE algorithm: Initialize zero primal and dual solutions, \vec{x} and \vec{y}, \vec{z} . Let $g(x) := \frac{e^x - 1}{e - 1}$. For every online vertex t on arrival, letting $x_i := \sum_{t' < t} x_{i,t'}$ be the fractional degree of neighboring vertex $i \in N(t)$ before arrival of t , we increase $x_{i,t}$ for all $i \in A$, where

$$A := \arg \max_{i \in N(t)} \{w_i \cdot (1 - g(x_i + x_{i,t}))\},$$

so that this set A grows monotonically, until $\sum_i x_{i,t} = 1$ or $\min_{i \in N(t)} (x_i + x_{i,t}) = 1$.

Note that for the unweighted case ($w_i = 1$ for all i), this algorithm is precisely the WATER LEVEL algorithm described in §2. The intuition behind this algorithm is clear: since each offline vertex is equally likely to have no future neighbors, we wish to maximize the value assigned to edges of the least (fractionally) matched offline vertex, in case it (and it alone) has no future edges. That generalizing this approach allows to get a competitive ratio better than $1/2$ (also for vertex-weighted matching) is perhaps less immediate. We present a proof of this fact using LP duality, and specifically dual fitting.

Dual fitting: For our analysis (only), for each offline vertex i and online vertex t , we set dual values

$$\begin{aligned}
 y_i &\leftarrow w_i \cdot g\left(\sum_{t'} x_{i,t'}\right), \\
 z_t &\leftarrow \max_{i \in N(t)} \left(w_i \cdot \left(1 - g\left(\sum_{t' \leq t} x_{i,t'}\right)\right) \right).
 \end{aligned}$$

The first step of any primal-dual-based proof involves showing that the constructed dual is feasible, and hence its cost upper bounds the maximum gain (in hindsight).

LEMMA A.1. *Vectors \vec{y}, \vec{z} are dual feasible, i.e., they are positive and $y_i + z_t \geq w_i$ for all edges $(i, t) \in E$. Consequently, $\sum_i y_i + \sum_t z_t \geq OPT$, where OPT is the weight of a maximum vertex-weighted matching.*

PROOF. As $g : [0, 1] \rightarrow [0, 1]$, we have that $y_i, z_t \geq 0$. On the other hand, by monotonicity of g (and positivity of \vec{x}), for every edge $(i, t) \in E$ we have that $y_i = w_i \cdot g(\sum_{t'} x_{i,t'}) \geq w_i \cdot g(\sum_{t' \leq t} x_{i,t'})$, and so

$$y_i + z_t \geq w_i \cdot g\left(\sum_{t' \leq t} x_{i,t'}\right) + w_i \cdot \left(1 - g\left(\sum_{t' \leq t} x_{i,t'}\right)\right) = w_i.$$

The lower bound $\sum_i y_i + \sum_t z_t \geq OPT$ then follows from weak LP duality, together with the primal LP being a fractional relaxation of the problem. \square

The second step in a primal-dual proof involves bounding the ratio of the primal and dual solutions' gain/cost. For online algorithms this typically boils down to bounding the ratio of the change in these values per cost, as in the following.

LEMMA A.2. *For each time t , the increase in primal value, $\Delta P_t := \sum_i w_i \cdot x_{i,t}$, and the increase in the dual cost, $\Delta D_t := z_t + \sum_i (w_i \cdot g(x_i + x_{i,t}) - w_i \cdot g(x_i))$, satisfy*

$$\Delta P_t / \Delta D_t \geq 1 - 1/e.$$

PROOF. The increase in dual cost satisfies

$$\begin{aligned} \Delta D_t &= \sum_i w_i \cdot (g(x_i + x_{i,t}) - g(x_i)) + z_t \\ &\leq \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + z_t \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + \max_i w_i \cdot (1 - g(x_i + x_{i,t})) \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + \sum_i x_{i,t} \cdot \max_j w_j \cdot (1 - g(x_j + x_{j,t})) \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(g(x_i + x_{i,t}) + \frac{1}{e-1}\right) + \sum_i x_{i,t} \cdot w_i \cdot (1 - g(x_i + x_{i,t})) \\ &= \sum_i w_i \cdot x_{i,t} \cdot \left(1 + \frac{1}{e-1}\right) \\ &= \Delta P_t \cdot \frac{e}{e-1}. \end{aligned}$$

Above, the single inequality relied on the definition of $g(x) = \frac{e^x - 1}{e - 1}$ implying that $g'(x) = g(x) + \frac{1}{e-1}$ is monotone increasing in x . The second and third equalities follow by definition of z_t and either $\sum_i x_{i,t} = 1$ or $\min_{i \in N(t)} (x_i + x_{i,t}) = 1$, where the latter implies that $g(x_i + x_{i,t}) = 1$ for all $i \in N(t)$. The fourth equality follows from $x_{i,t} = 0$ unless $i \in \arg \max_{i \in N(t)} \{w_i \cdot (1 - g(x_i + x_{i,t}))\}$. \square

Combining both lemmas, the algorithm's competitive ratio follows.

THEOREM A.3. *Algorithm BALANCE is $(1 - 1/e)$ -competitive for vertex-weighted online bipartite matching.*

PROOF. Summing over all times t , we have by Lemmas A.2 and A.1 that

$$\sum_{i,t} w_i \cdot x_{i,t} = \sum_t \Delta P_t \geq (1 - 1/e) \cdot \sum_t \Delta D_t = \sum_i w_i \cdot g \left(\sum_{t'} x_{i,t'} \right) + \sum_t z_t \geq OPT,$$

where the last equality uses the telescoping sum

$$\sum_t \left(g \left(\sum_{t' \leq t} x_{i,t'} \right) - g \left(\sum_{t' < t} x_{i,t'} \right) \right) = g \left(\sum_{t'} x_{i,t'} \right) - g(0) = g \left(\sum_{t'} x_{i,t'} \right),$$

where the last equality follows from $g(0) = 0$. \square

A.2 Matching Impossibility Result

We now show that BALANCE is optimal, up to vanishingly small lower-order terms, even for the special case of *unweighted* online bipartite matching. The underlying input family is the same as that used by [Karp et al. 1990] to prove that randomized algorithms are at best $(1 - 1/e)$ -competitive. The following proof shows in a more direct manner that this ratio is best possible even for (potentially more powerful) fractional algorithms. The intuition behind the proof is precisely that guiding WATER LEVEL described above.

THEOREM A.4. *Any fractional online bipartite (unweighted) matching algorithm \mathcal{A} has competitive ratio at most $1 - 1/e + o(1)$.*

PROOF. We consider inputs consisting of n offline and online vertices. The first online vertex neighbors all offline vertices; each following online vertex has the same neighborhood as its predecessor, barring one vertex. (Under appropriate labeling the “bipartite” adjacency matrix of this graph is upper triangular.) This graph clearly has a perfect matching, and therefore optimum value of n . To prove our theorem we show that for any fractional algorithm \mathcal{A} , a judicious choice of input forces \mathcal{A} to achieve value at most $n(1 - 1/e) + 1$.

For every online vertex t , the neighbor i with no future edges is chosen to be a neighbor which was matched below the average, e.g., which minimizes $\sum_{t' \leq t} x_{i,t'}$. As the t -th online vertex to arrive neighbors $n - t + 1$ offline vertices, a proof by induction shows that offline vertices that do neighbor online vertex $t+1$ are matched before time $t + 1$ to an average of at least $\sum_{t' \leq t} x_{i,t'} \geq \min\{1, \sum_{t'=1}^t \frac{1}{n-t'+1}\}$.⁵ Consequently, for $t = n(1 - 1/e) + 1$, every offline vertex i that neighbors $t + 1$ is fractionally matched to an extent of at least one, since

$$\sum_{t'=1}^{n(1-1/e)+1} \frac{1}{n-t'+1} = \sum_{i=n/e}^n \frac{1}{i} \geq \int_{n/e}^{n+1} \frac{1}{x} dx = \ln(n+1) - \ln(n/e) \geq 1.$$

⁵This holds only if the algorithm \mathcal{A} is greedy, in the sense that it exhausts every online vertex t , i.e., sets $\sum_i x_{i,t} = 1$, when feasible. However, it is easy to modify any algorithm with a greedy algorithm which does no worse.

But since $\sum_{t'} x_{i,t'} \leq 1$, this implies that every offline vertex i neighboring online vertices after time t is already fully (fractionally) matched by time t , and so the algorithm gains no further profit from such i . Consequently, \mathcal{A} achieves a value of at most $\sum_{i,t'} x_{i,t'} \leq n(1 - 1/e) + 1$, and so has competitive ratio at most $1 - 1/e + 1/n$. \square

See [Feige 2018] for tight bounds on the $o(1)$ term in the above upper bound.

REFERENCES

ReferencesRef

- AGGARWAL, G., GOEL, G., KARANDE, C., AND MEHTA, A. 2011. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1253–1264.
- ANARI, N., NIAZADEH, R., SABERI, A., AND SHAMELI, A. 2019. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 91–92.
- AOUAD, A. AND MA, W. 2023. A nonparametric framework for online stochastic matching with correlated arrivals. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 114–114.
- AOUAD, A. AND SARITAÇ, Ö. 2020. Dynamic stochastic matching under limited time. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 789–790.
- ASHLAGI, I., BURQ, M., DUTTA, C., JAILLET, P., SABERI, A., AND SHOLLEY, C. 2023. Edge-weighted online windowed matching. *Mathematics of Operations Research* 48, 2, 999–1016.
- AZAR, Y. AND LITICHEVSKY, A. 2006. Maximizing throughput in multi-queue switches. *Algorithmica* 45, 1, 69–90.
- BANIHASHEM, K., HAJIAGHAYI, M., KOWALSKI, D. R., KRISTA, P., AND OLKOWSKI, J. 2024. Power of posted-price mechanisms for prophet inequalities. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 4580–4604.
- BLANC, G. AND CHARIKAR, M. 2021. Multiway online correlated selection. In *Proceedings of the 62nd Symposium on Foundations of Computer Science (FOCS)*. 1277–1284.
- BLIKSTAD, J., SVENSSON, O., VINTAN, R., AND WAJC, D. 2024. Online edge coloring is (nearly) as easy as offline. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*. To Appear.
- BLUM, A., SANDHOLM, T., AND ZINKEVICH, M. 2006. Online algorithms for market clearing. *Journal of the ACM (JACM)* 53, 5, 845–879.
- BRAVERMAN, M., DERAKHSHAN, M., AND LOVETT, A. M. 2022. Max-weight online stochastic matching: Improved approximations against the online benchmark. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 967–985.
- BRUBACH, B., GRAMMEL, N., MA, W., AND SRINIVASAN, A. 2021. Improved guarantees for offline stochastic matching via new ordered contention resolution schemes. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- BUCHBINDER, N., JAIN, K., AND NAOR, J. S. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Proceedings of the 15th Annual European Symposium on Algorithms (ESA)*. 253–264.
- BUCHBINDER, N., JAIN, K., AND SINGH, M. 2014. Secretary problems via linear programming. *Mathematics of Operations Research* 39, 1, 190–206.
- BUCHBINDER, N., NAOR, J., AND WAJC, D. 2023. Lossless online rounding for online bipartite matching (despite its impossibility). In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2030–2068.
- BUCHBINDER, N. AND NAOR, J. S. 2009. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends® in Theoretical Computer Science (TCS)* 3, 2–3, 93–263.
- BUCHBINDER, N., SEGEV, D., AND TKACH, Y. 2019. Online algorithms for maximum cardinality matching with edge arrivals. *Algorithmica* 81, 5, 1781–1799.

- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*. 311–320.
- CHEKURI, C., VONDRÁK, J., AND ZENKLUSEN, R. 2014. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing (SICOMP)* 43, 6, 1831–1879.
- COHEN, I. R. AND WAJC, D. 2018. Randomized online matching in regular graphs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 960–979.
- COLLINA, N., IMMORLICA, N., LEYTON-BROWN, K., LUCIER, B., AND NEWMAN, N. 2020. Dynamic weighted matching with heterogeneous arrival and departure rates. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*. 17–30.
- CORREA, J., CRISTI, A., FIELBAUM, A., POLLNER, T., AND WEINBERG, S. M. 2023. Optimal item pricing in online combinatorial auctions. *Mathematical Programming*, 1–32.
- CORREA, J., FONCEA, P., HOEKSMAN, R., OOSTERWIJK, T., AND VREDEVELD, T. 2018. Recent developments in prophet inequalities. *SIGecom Exchanges* 17, 1, 61–70.
- DELONG, S., FARHADI, A., NIAZADEH, R., SIVAN, B., AND UDWANI, R. 2023. Online bipartite matching with reusable resources. *Mathematics of Operations Research*.
- DEVANUR, N. AND MEHTA, A. 2022. Online matching in advertisement auctions. In *Online and Matching-Based Market Design*. Cambridge University Press.
- DEVANUR, N. R., HUANG, Z., KORULA, N., MIRROKNI, V. S., AND YAN, Q. 2016. Whole-page optimization and submodular welfare maximization with online bidders. *ACM Trans. Economics and Comput.* 4, 3, 14:1–14:20.
- DEVANUR, N. R., JAIN, K., AND KLEINBERG, R. D. 2013. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 101–107.
- DÜTTING, P., FELDMAN, M., KESSELHEIM, T., AND LUCIER, B. 2020. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing (SICOMP)* 49, 3, 540–582.
- DÜTTING, P., GERGATSOULI, E., REZVAN, R., TENG, Y., AND TSIGONIAS-DIMITRIADIS, A. 2023. Prophet secretary against the online optimal. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 561–581.
- DYNKIN, E. B. 1963. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics* 4, 627–629.
- ECHENIQUE, F., IMMORLICA, N., AND VAZIRANI, V. V. 2023. *Online and matching-based market design*. Cambridge University Press.
- ECKL, A., KIRSCHBAUM, A., LEICHTER, M., AND SCHEWIOR, K. 2021. A stronger impossibility for fully online matching. *Operations Research Letters* 49, 5, 802–808.
- EDEN, A., FELDMAN, M., FIAT, A., AND SEGAL, K. 2021. An economics-based analysis of ranking for online bipartite matching. In *Proceedings of the 4th Symposium on Simplicity in Algorithms (SOSA)*. 107–110.
- EZRA, T., FELDMAN, M., GRAVIN, N., AND TANG, Z. G. 2022a. General graphs are easier than bipartite graphs: Tight bounds for secretary matching. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 1148–1177.
- EZRA, T., FELDMAN, M., GRAVIN, N., AND TANG, Z. G. 2022b. Prophet matching with general arrivals. *Mathematics of Operations Research* 47, 2, 878–898.
- FAHRBACH, M., HUANG, Z., TAO, R., AND ZADIMOGHADDAM, M. 2022. Edge-weighted online bipartite matching. *Journal of the ACM (JACM)* 69, 6, 1–35.
- FEIGE, U. 2018. Tighter bounds for online bipartite matching. *arXiv preprint arXiv:1812.11774*.
- FELDMAN, J., KORULA, N., MIRROKNI, V., MUTHUKRISHNAN, S., AND PÁL, M. 2009. Online ad assignment with free disposal. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*. 374–385.
- FELDMAN, J., MEHTA, A., MIRROKNI, V., AND MUTHUKRISHNAN, S. 2009. Online stochastic matching: Beating $1-1/e$. In *Proceedings of the 50th Symposium on Foundations of Computer Science (FOCS)*. 117–126.

- FELDMAN, M., GRAVIN, N., AND LUCIER, B. 2015. Combinatorial auctions via posted prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 123–135.
- FELDMAN, M., SVENSSON, O., AND ZENKLUSEN, R. 2016. Online contention resolution schemes. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1014–1033.
- FENG, Y., NIAZADEH, R., AND SABERI, A. 2021. Robustness of online inventory balancing algorithm to inventory shocks. Available at SSRN: <https://ssrn.com/abstract=3795056>.
- FENG, Y., NIAZADEH, R., AND SABERI, A. 2022. Near-optimal bayesian online assortment of reusable resources. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 964–965.
- FU, H., TANG, Z. G., WU, H., WU, J., AND ZHANG, Q. 2021. Random order vertex arrival contention resolution schemes for matching, with applications. In *Proceedings of the 48th International Colloquium on Automata, Languages and Programming (ICALP)*.
- GAMLATH, B., KAPRALOV, M., MAGGIORI, A., SVENSSON, O., AND WAJC, D. 2019. Online matching with general arrivals. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*. 26–38.
- GAO, R., HE, Z., HUANG, Z., NIE, Z., YUAN, B., AND ZHONG, Y. 2021. Improved online correlated selection. In *Proceedings of the 62nd Symposium on Foundations of Computer Science (FOCS)*. 1265–1276.
- GONG, X., GOYAL, V., IYENGAR, G. N., SIMCHI-LEVI, D., UDWANI, R., AND WANG, S. 2022. Online assortment optimization with reusable resources. *Management Science* 68, 7, 4772–4785.
- GOYAL, V., IYENGAR, G., AND UDWANI, R. 2021. Asymptotically optimal competitive ratio for online allocation of reusable resources. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. 543.
- GOYAL, V. AND UDWANI, R. 2023. Online matching with stochastic rewards: optimal competitive ratio via path-based formulation. *Operations Research* 71, 2, 563–580.
- GRAVIN, N., TANG, Z. G., AND WANG, K. 2021. Online stochastic matching with edge arrivals. In *Proceedings of the 48th International Colloquium on Automata, Languages and Programming (ICALP)*. 74:1–74:20.
- GRAVIN, N. AND WANG, H. 2019. Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 93–109.
- GURUGANESH, G. P. AND SINGLA, S. 2017. Online matroid intersection: Beating half for random arrival. In *Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 241–253.
- HAEUPLER, B., MIRROKNI, V. S., AND ZADIMOGHADDAM, M. 2011. Online stochastic weighted matching: Improved approximation algorithms. In *Proceedings of the 7th Conference on Web and Internet Economics (WINE)*. 170–181.
- HAJIAGHAYI, M. T., KLEINBERG, R., AND SANDHOLM, T. 2007. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*. 58–65.
- HARTLINE, J. D. 2012. Approximation in mechanism design. *American Economic Review* 102, 3, 330–36.
- HILL, T. P. AND KERTZ, R. P. 1992. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics* 125, 191–207.
- HOSSEINI, H., HUANG, Z., IGARASHI, A., AND SHAH, N. 2023. Class fairness in online matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 5673–5680.
- HUANG, Z., JIANG, H., SHEN, A., SONG, J., WU, Z., AND ZHANG, Q. 2023. Online matching with stochastic rewards: Advanced analyses using configuration linear programs. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE)*. 384–401.
- HUANG, Z., KANG, N., TANG, Z. G., WU, X., ZHANG, Y., AND ZHU, X. 2020. Fully online matching. *Journal of the ACM (JACM)* 67, 3, 1–25.

- HUANG, Z., PENG, B., TANG, Z. G., TAO, R., WU, X., AND ZHANG, Y. 2019. Tight competitive ratios of classic matching algorithms in the fully online model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2875–2886.
- HUANG, Z. AND SHU, X. 2021. Online stochastic matching, poisson arrivals, and the natural linear program. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*. 682–693.
- HUANG, Z., SHU, X., AND YAN, S. 2022. The power of multiple choices in online stochastic matching. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 91–103.
- HUANG, Z., TANG, Z. G., WU, X., AND ZHANG, Y. 2019. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. *ACM Transactions on Algorithms (TALG)* 15, 3, 1–15.
- HUANG, Z., TANG, Z. G., WU, X., AND ZHANG, Y. 2020. Fully online matching ii: Beating ranking and water-filling. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*. 1380–1391.
- HUANG, Z. AND ZHANG, Q. 2020. Online primal dual meets online matching with stochastic rewards: Configuration lp to the rescue. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*. 1153–1164.
- HUANG, Z., ZHANG, Q., AND ZHANG, Y. 2020. Adwords in a panorama. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*. 1416–1426.
- JAILLET, P. AND LU, X. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*.
- JIN, B. AND WILLIAMSON, D. P. 2021. Improved analysis of ranking for online vertex-weighted bipartite matching in the random order model. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. 207–225.
- KALYANASUNDARAM, B. AND PRUHS, K. R. 2000. An optimal deterministic algorithm for online b -matching. *Theoretical Computer Science (TCS)* 233, 1, 319–325.
- KAPRALOV, M., POST, I., AND VONDRÁK, J. 2013. Online submodular welfare maximization: Greedy is optimal. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1216–1225.
- KARANDE, C., MEHTA, A., AND TRIPATHI, P. 2011. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 587–596.
- KARP, R. M., VAZIRANI, U. V., AND VAZIRANI, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*. 352–358.
- KESSEL, K., SHAMELI, A., SABERI, A., AND WAJC, D. 2022. The stationary prophet inequality problem. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 243–244.
- KESSELHEIM, T., RADKE, K., TÖNNIS, A., AND VÖCKING, B. 2013. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Proceedings of the 21st Annual European Symposium on Algorithms (ESA)*. 589–600.
- KLEINBERG, R. AND WEINBERG, S. M. 2019. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior* 113, 97–115.
- KLEINBERG, R. D. 2005. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 630–631.
- KORULA, N. AND PÁL, M. 2009. Algorithms for secretary problems on graphs and hypergraphs. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*. 508–520.
- KRENGEL, U. AND SUCHESTON, L. 1978. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces* 4, 197–266.

- LEE, E. AND SINGLA, S. 2017. Maximum matching in the online batch-arrival model. In *Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 355–367.
- LEE, E. AND SINGLA, S. 2018. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *Proceedings of the 26th Annual European Symposium on Algorithms (ESA)*. 57:1–57:14.
- LEVIN, R. AND WAJC, D. 2021. Streaming submodular matching meets the primal-dual method. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1914–1933.
- LIANG, J., TANG, Z. G., XU, Y. E., ZHANG, Y., AND ZHOU, R. 2023. On the perturbation function of ranking and balance for weighted online bipartite matching. In *Proceedings of the 31st Annual European Symposium on Algorithms (ESA)*. 80:1–80:15.
- LUCIER, B. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16, 1, 24–47.
- MACRURY, C. AND MA, W. 2024. Random-order contention resolution via continuous induction: Tightness for bipartite matching under vertex arrivals. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*. To Appear.
- MACRURY, C., MA, W., AND GRAMMEL, N. 2023. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs. In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1995–2014.
- MAHDIAN, M. AND YAN, Q. 2011. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 597–606.
- MANSHADI, V. H., GHARAN, S. O., AND SABERI, A. 2012. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37, 4, 559–573.
- MEHTA, A. 2013. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8, 4, 265–368.
- MEHTA, A. AND PANIGRAHI, D. 2012. Online matching with stochastic rewards. In *Proceedings of the 53rd Symposium on Foundations of Computer Science (FOCS)*. 728–737.
- MEHTA, A., SABERI, A., VAZIRANI, U., AND VAZIRANI, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5, 22.
- MEHTA, A., WAGGONER, B., AND ZADIMOGHADDAM, M. 2014. Online stochastic matching with unequal probabilities. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1388–1404.
- NAOR, J. S., SRINIVASAN, A., AND WAJC, D. 2023. Online dependent rounding schemes. *arXiv preprint arXiv:2301.08680*.
- PAPADIMITRIOU, C., POLLNER, T., SABERI, A., AND WAJC, D. 2021. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 763–764.
- PATEL, N. AND WAJC, D. 2024. Combinatorial stationary prophet inequalities. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 4605–4630.
- POLLNER, T., ROGHANI, M., SABERI, A., AND WAJC, D. 2022. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*.
- QIU, G., FENG, Y., ZHOU, S., AND WU, X. 2023. Improved competitive ratio for edge-weighted online stochastic matching. In *Proceedings of the 20th International Conference on Web and Internet Economics (WINE)*. Springer, 527–544.
- SABERI, A. AND WAJC, D. 2021. The greedy algorithm is not optimal for online edge coloring. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*. 109:1–109:18.
- SAMUEL-CAHN, E. 1984. Comparison of threshold stop rules and maximum for independent non-negative random variables. *the Annals of Probability* 12, 4, 1213–1216.

- SHIN, Y. AND AN, H. C. 2021. Making three out of two: three-way online correlated selection. In *Proceedings of the 32nd Annual International Symposium on Algorithms and Computation (ISAAC)*. 49.
- TANG, Z. G., WU, J., AND WU, H. 2022. (fractional) online stochastic matching via fine-grained offline statistics. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 77–90.
- TANG, Z. G. AND ZHANG, Y. 2022. Improved bounds for fractional online matching problems. *arXiv preprint arXiv:2202.02948*.
- TORRICO, A. AND TORIELLO, A. 2022. Dynamic relaxations for online bipartite matching. *INFORMS Journal on Computing* 34, 4, 1871–1884.
- UDWANI, R. 2023. Adwords with unknown budgets and beyond. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 1128–1128.
- VAZIRANI, V. V. 2023. Towards a practical, budget-oblivious algorithm for the adwords problem under small bids. In *Proceedings of the 43rd Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*.
- WAJC, D. 2020. Matching theory under uncertainty. Ph.D. thesis, Carnegie Mellon University.
- WANG, Y. AND WONG, S. C.-W. 2015. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP)*. 1070–1081.
- YAN, Q. 2011. Mechanism design via correlation gap. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 710–719.
- YAN, S. 2024. Edge-weighted online stochastic matching: Beating. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 4631–4640.