

Adwords with Unknown Budgets and Beyond

Rajan Udvani

University of California Berkeley, Industrial Engineering and Operations Research, rudwani@berkeley.edu

In the classic Adwords problem introduced by Mehta et al. (2007), we have a bipartite graph between advertisers and queries. Each advertiser has a maximum budget that is known a priori. Queries are unknown a priori and arrive sequentially. When a query arrives, advertisers make bids and we (immediately and irrevocably) decide which (if any) Ad to display based on the bids and advertiser budgets. The winning advertiser for each query pays their bid up to their remaining budget. Our goal is to maximize total budget utilized without any foreknowledge of the arrival sequence (which could be adversarial). We consider the setting where the online algorithm does not know the advertisers' budgets a priori and the budget of an advertiser is revealed to the algorithm only when it is exceeded. A naïve greedy algorithm is 0.5 competitive for this setting and finding an algorithm with better performance remained an open problem. We show that no deterministic algorithm has competitive ratio better than 0.5 and give the first (randomized) algorithm with strictly better performance guarantee. We show that the competitive ratio of our algorithm is at least 0.522 but also strictly less than $(1 - 1/e)$. We present novel applications of budget oblivious algorithms in search ads and beyond. In particular, we show that our algorithm achieves the best possible performance guarantee for deterministic online matching in the presence of multi-channel traffic (Manshadi et al. (2022)).

Key words: Adwords, Randomized Algorithms, Unknown Budgets, Competitive Ratio

1. Introduction

Online advertising has emerged as the dominant marketing channel in many parts of the world. According to some estimates, in the year 2019, more than 450 billion USD were spent on online ads, which accounts for over 60% of overall expenditure on ads¹. Internet search is a prominent channel for online advertisement². In this medium, also called *search ads*, advertisements are displayed alongside search results for key words relevant to the advertiser. Given a search request, advertisers make bids and then the platform decides which (if any) ads to show with the search results. The following model captures key elements of this problem.

The Adwords problem (Mehta et al. 2007): At the beginning of the planning period (typically a day), the platform has a set I of advertisers along with their budgets $(B_i)_{i \in I}$. Queries arrive sequentially on the platform and when a query t arrives, advertisers make bids $(b_{i,t})_{i \in I}$. Given the bids and advertiser budgets, the platform immediately picks at most one advertisers' ad to display

¹Digital advertising spending worldwide from 2019 to 2024. <https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/>

²In 2020 alone, Google made a revenue of over 104 billion USD from "search & other", which accounts for 70% of their total revenue from advertising and exceeds half the total revenue of parent company Alphabet (see Alphabet Year in Review 2020. <https://abc.xyz/investor/>).

alongside the search results for the query³. The chosen advertiser pays their bid but only up to their remaining budget, i.e., at the end of the planning period, the total payment of an advertiser does not exceed his/her budget. The objective of the platform is to maximize the total advertiser budgets utilized without any foreknowledge of the arrival sequence (which could be adversarial). A typical modeling assumption is that the maximum individual bid by any advertiser is much smaller than the advertiser's initial budget. This is called the *small bids* or *large budgets* assumption and it is in line with the practice of search ads.

In the face of uncertain future queries, which (if any) ad should the platform show for a given query? The greedy algorithm for this problem shows the ad with highest bid and (non-zero) available budget. In particular, if $I(t)$ is the set of advertisers which have non-zero remaining budget when query t arrives, the greedy algorithm shows ad,

$$\operatorname{argmax}_{i \in I(t)} b_{i,t}.$$

This algorithm does not distinguish between two advertisers that make identical bids, even if the budget of one advertiser is nearly used up and the other has plenty of budget remaining. Mehta et al. (2007) gave the state-of-the-art *bid pricing* algorithm that carefully accounts for the fraction of remaining budgets. Given remaining budget $(B_i(t))_{i \in I}$ on arrival of query t , the algorithm of Mehta et al. (2007) computes bid prices

$$b_{i,t} (1 - e^{-B_i(t)/B_i}) \quad \forall i \in I,$$

and shows the ad with the largest bid price. Given two advertisers with similar bids, this algorithm may select the advertiser with smaller bid and higher fraction of remaining budget. The trade off between bid and (fraction of) remaining budget is governed by the function $(1 - e^{-x})$ that arises quite naturally from the worst case performance analysis, a.k.a. *competitive ratio* analysis, of this algorithm. The competitive ratio of an algorithm is the worst case relative performance gap between the online algorithm and (optimal) offline algorithm, OPT, that knows all the queries and advertiser bids. Let G denote an instance of the problem (advertisers, arrivals, bids and budgets) and let $\mathcal{G}$ denote the set of all instances. Let $\text{ALG}(G)$ denote the expected total budget utilized by a (possibly) randomized online algorithm ALG on instance G . Similarly, let $\text{OPT}(G)$ denote the optimal offline value on instance G .

$$\textbf{Competitive ratio of ALG:} \quad \min_{G \in \mathcal{G}} \frac{\text{ALG}(G)}{\text{OPT}(G)}$$

³ For simplicity, the model considers at most one ad slot per search result. A generalization to multiple ad slots is described in Section 6 of Mehta et al. (2007).

The algorithm of Mehta et al. (2007) is $(1 - 1/e)$ competitive. In comparison, the greedy algorithm has a competitive ratio of 0.5. Remarkably, no (randomized) online algorithm has competitive ratio better than $(1 - 1/e)$ even for special cases of the Adwords problem (Kalyanasundaram and Pruhs 2000).

The knowledge of advertisers' initial budgets is essential for defining the bid pricing algorithm of Mehta et al. (2007). In fact, most (if not all) algorithms for Adwords in the literature (see Devanur and Hayes (2009), Alaei et al. (2012), Mirrokni et al. (2012), Mehta et al. (2013), Devanur et al. (2019), Balseiro et al. (2020)), rely critically on the knowledge of initial budgets for each advertiser. The naïve greedy algorithm is the one exception to this. For each query, greedy shows the ad with highest bid and (non-zero) available budget. Therefore, it is *budget oblivious*, i.e., does not require any advance information about budgets except the knowledge of which advertisers are still participating (have non-zero remaining budget). Clearly, a budget oblivious algorithm is more robust since it requires even less knowledge of the instance. Motivated by this, we consider the following question.

Is there a budget oblivious online algorithm for Adwords that outperforms greedy? What are the advantages of budget obliviousness?

We are not the first to ask these questions. In fact, the problem of Adwords with unknown budgets was first proposed by Mehta and Panigrahi (2012), who used a special case of the problem to analyze a related setting called online matching with stochastic rewards (described in more detail in Section 5.2). Prompted by this intriguing connection, Mehta et al. (2013) posed the question of finding a budget oblivious algorithm that is better than greedy (or proving that no such algorithm exists) as an open problem (see Open Question 20 in Mehta et al. (2013)). To the best of our knowledge, this question remained open prior to our work.

1.1. Our Contributions

When budgets are unknown, we show that greedy is the best possible deterministic algorithm even when the budgets are large, i.e., no deterministic algorithm has competitive ratio better than 0.5 when budgets large but unknown. In fact, we show that this is true even when every bid is either 0 or 1, in which case the problem reduces to online matching (with large budgets).

We show that one can do strictly better than greedy with randomized algorithms and give the first budget oblivious algorithm with competitive ratio better than 0.5. Our algorithm samples i.i.d. uniform random variables $x_i \in [0, 1] \forall i \in I$, and computes random bid prices,

$$b_{i,t}(1 - e^{-\beta x_i}) \quad \forall i \in I,$$

where $\beta > 0$ is a parameter that can be chosen by the platform. At each query, the algorithm shows the ad with the largest (random) bid price from the set of advertisers with non-zero remaining budget. The key difference from the algorithm of Mehta et al. (2007) is that bid prices are independent of budgets; the budget dependent factor, $(1 - e^{-B_i(t)/B_i})$, in the algorithm of Mehta et al. (2007), is replaced with the random factor $(1 - e^{-\beta x_i})$. In the special case of Adwords where $b_{i,t} \in \{0, B_i\} \forall i \in I, t \in T$, our algorithm (with $\beta = 1$) reduces to the Perturbed Greedy algorithm of Aggarwal et al. (2011). Due to this connection, we refer to our algorithm as *Generalized Perturbed Greedy* or GPG.

Under the standard assumption of small bids, we show that for $\beta = 1.15$, GPG is at least 0.522 competitive against the optimal offline algorithm that knows all bids and budgets. For $\beta = 1$, we show that the competitive ratio is least 0.508. Further, we prove that the competitive ratio of GPG is strictly less than 0.624 ($< (1 - 1/e)$), for all $\beta \geq 1$. In fact, we numerically observe a stronger upper bound of 0.604 for $\beta = 1^4$. To prove the competitive ratio lower bound we overcome a novel obstacle that arises from combinatorial interactions between time varying bids and the randomness inherent in the algorithm. In particular, a “dominance” property that is crucially used in most (if not all) previous analysis of (special cases) of GPG (Aggarwal et al. 2011, Devanur et al. 2013, Eden et al. 2021, Albers and Schubert 2021, Vazirani 2021), does not apply in the Adwords setting. We address this challenge by finding new structural insights into the problem. We also show that various parts of our analysis are individually tight but we believe that the overall analysis may not be tight.

It is Adwords folklore that for small bids (large budgets) deterministic algorithms are as powerful as randomized ones, i.e., if there is a randomized α competitive algorithm then there exists a deterministic algorithm with competitive ratio at least α . Our results for unknown budgets imply a perhaps surprising gap between the two classes of algorithms. At a high level, randomness allows an online algorithm to (with some probability) conserve the budget of advertisers with low initial budgets without actually knowing the budgets.

We demonstrate the usefulness of our budget oblivious algorithm in a variety of settings. First, we consider the setting of online matching (special case of Adwords) and multi-channel traffic. A stochastic generalization of this setting was introduced by Manshadi et al. (2022) to model online recommendations on platforms such as VolunteerMatch. We show that our budget oblivious algorithm achieves the best possible competitive ratio guarantee for this setting, closing the gap between the lower and upper bound shown in Manshadi et al. (2022). Second, we consider a previously well known application of budget oblivious algorithms in the setting of online matching

⁴Very recently, Liang et al. (2023) showed that for any function f , the algorithm that matches greedily using bid prices $b_{i,t}f(y_i)$ is strictly less than $(1 - 1/e)$ competitive.

with stochastic rewards (Mehta and Panigrahi 2012). Our result for Adwords with unknown budgets gives a new and more unified algorithmic result for this setting. Finally, we discuss potential applications of budget oblivious algorithms in auto-bidding and automated budget management, where a budget oblivious allocation algorithm may allow the platform to improve total budget utilization by re-optimizing advertisers’ budgets during the planning period.

1.2. Related Work

Karp et al. (1990) introduced the classic setting of online bipartite matching that corresponds to Adwords with binary bids and unit budget for all advertisers. A bid of 1 denotes an edge in the bipartite graph and bid of 0 denotes the absence of an edge. Every advertiser can be matched to at most one neighboring query. Karp et al. (1990) showed (among other results) that randomly ranking advertisers at the start and then matching every query to the best ranked unmatched advertiser is a $(1 - 1/e)$ competitive algorithm for this setting. In fact, this algorithm, called Ranking, achieves the best possible guarantee for the problem. The analysis of Ranking was clarified and considerably simplified by Birnbaum and Mathieu (2008) and Goel and Mehta (2008). Aggarwal et al. (2011) considered the more general vertex weighted version of this problem which corresponds to Adwords with bids $b_{i,t} \in \{0, B_i\} \forall i \in I, t \in T$. They gave the Perturbed Greedy algorithm (that we use in this paper), and showed that the algorithm is $(1 - 1/e)$ competitive.

Kalyanasundaram and Pruhs (2000) considered the problem of online b -matching, which is a special case of Adwords with binary bids and identical (large) budget b for every advertiser. They showed that as $b \rightarrow \infty$, the natural (deterministic) algorithm that balances the budget usage across advertisers is $(1 - 1/e)$ competitive. Generalizing this setting, Mehta et al. (2007) introduced the Adwords problem and gave the bid pricing based $(1 - 1/e)$ algorithm for Adwords under the small bid assumption. Buchbinder et al. (2007) gave a primal-dual analysis for this algorithm. Subsequently, Devanur et al. (2013) introduced the randomized primal-dual framework and used it to show all of the results mentioned above in a unified way.

The Adwords setting without the small bids assumption, generalizes each of the settings discussed above. The budget-aware greedy algorithm that matches each query $t \in T$ according to the following rule,

$$\operatorname{argmax}_{i \in I} (\min\{b_{i,t}, B_i(t)\}),$$

where $B_i(t)$ is the remaining budget of i on arrival of query t , is 0.5 competitive for Adwords without any assumption on the bids. Kapralov et al. (2013) showed that without the small bids assumption, no online algorithm has competitive ratio better than 0.612 for Adwords⁵. Recently,

⁵ This applies when the competitive ratio is evaluated against an LP upper bound on the offline problem.

Huang et al. (2020) gave the first algorithm with competitive ratio better than 0.5 for Adwords without the small bids assumption.

Concurrent with this paper: Albers and Schubert (2021) and Vazirani (2021, 2023) independently show that the Perturbed Greedy algorithm is $(1 - 1/e)$ competitive for Adwords with binary bids and arbitrary resource budgets. To the best of our knowledge, these approaches do not yield a performance guarantee better than 0.5 for Generalized Perturbed Greedy (GPG) in the Adwords setting and this is posed as an open problem in Vazirani (2023). Vazirani (2023) also identifies a key structural property, called *no surpassing*, the absence of which prevents a generalization of their result to the Adwords setting. Through various numerical experiments (based on synthetic data), they demonstrate that the performance of GPG is at par with the algorithm of Mehta et al. (2007). We include a more detailed comparison with these papers in Appendix A.

Very recently, Liang et al. (2023) showed that for any function f , the algorithm that matches greedily using bid prices $b_{i,t}f(y_i)$, is strictly less than $(1 - 1/e) - \delta$ competitive for some constant $\delta > 0$.

While the body of work discussed above considers an adversarial arrival sequence, there is also a long line of work on online matching and Adwords in stochastic and hybrid/mixed arrival models (for example, Goel and Mehta (2008), Feldman et al. (2009), Devanur and Hayes (2009), Karande et al. (2011), Manshadi et al. (2012), Alaei et al. (2012), Devanur et al. (2019), Mirrokni et al. (2012)). For a comprehensive review of these settings we refer to Mehta et al. (2013).

Outline for rest of the paper: Section 2 discusses the assumption of small bids and presents a resource allocation version of the Adwords setting that we use interchangeably with the original formulation. In Section 3, we show that one cannot obtain a deterministic algorithm with competitive ratio better than 0.5. In Section 4, we state and prove our main results for Adwords with unknown budgets. In Section 5, we show new results for our budget oblivious algorithm in settings beyond Adwords and discuss applications of budget obliviousness. Section 6 concludes our discussion.

2. Preliminaries

The Adwords problem generalizes a variety of settings in the literature on online resource allocation. We formalize this connection by discussing an online resource allocation problem that is known to be equivalent to the Adwords problem. Terminology from this setting and the Adwords problem will be used interchangeably throughout the paper.

Online Budgeted Allocation (OBA): Consider a complete bipartite graph G with vertex set (I, T) . Vertices $i \in I$, called *resources*, have capacities $(B_i)_{i \in I}$ and per-unit rewards $(r_i)_{i \in I}$.

Resources, their rewards, and capacities are known to us. Vertices $t \in T$, called *arrivals*, are unknown a priori and arrive sequentially. We use T to denote the set of arrivals as well as the total number of arrivals. When a vertex $t \in T$ arrives, we see *bids* $(b_{i,t})_{i \in I}$ that indicate t is interested in up to $b_{i,t}$ amount of resource $i \in I$. Given the bids for arrival t , we must immediately and irrevocably match the arrival to at most one resource. If t is matched to i , $b_{i,t}$ amount of i are consumed, subject to availability. If the remaining capacity of i is less than $b_{i,t}$, then matching t to i uses up all the remaining capacity. We receive a reward r_i per-unit of i 's consumed capacity. The total reward from matching i is capped at $r_i B_i$. The goal is to decide the allocation/matching for arrivals without any knowledge of future arrivals such that the total reward is maximized.

The Adwords problem is an instance of OBA where resources correspond to advertisers and arrivals correspond to queries. The per-unit rewards r_i are set to 1 in Adwords for every $i \in I$. On the other hand, an instance of OBA with bids $b_{i,t}$, capacities B_i , and per-unit rewards r_i , is equivalent to the Adwords setting with scaled bids $r_i b_{i,t}$ and budgets $r_i B_i$. Due to this observation, without loss of generality (w.l.o.g.), we let per-unit rewards

$$r_i = 1 \quad \forall i \in I.$$

A standard assumption in the Adwords setting is the that the maximum bid-to-budget ratio,

$$\gamma := \max_{i \in I, t \in T} \frac{b_{i,t}}{B_i}.$$

is small, i.e., $\gamma \rightarrow 0$. This is also called the *small bid* or the *large budget/capacity* assumption. This assumption is in line with the practice of search ads, where individual bids are typically much smaller than the overall budget. Recall that the $(1 - 1/e)$ guarantee of Mehta et al. (2007) holds only in the small bid regime. Even in this regime, no online algorithm has a competitive ratio better than $(1 - 1/e)$.

OBA with Unknown Capacities: In the setting we are interested in, the online algorithm has no prior knowledge of resource capacities. When the algorithm fully utilizes the capacity of a resource, the event is immediately revealed to the algorithm, i.e., capacity of a resource is revealed right after it is used up. We evaluate the competitive ratio of a capacity/budget oblivious online algorithm against the optimal offline matching with complete knowledge of the instance.

3. Upper Bound for Deterministic Algorithms

The folklore for Adwords and many other online resource allocation problems says that given an online randomized algorithm, one can construct an online deterministic algorithm that matches arrivals *fractionally* and emulates the expected performance of the randomized algorithm at every arrival. As we demonstrate below, this line of argument relies on prior knowledge of budgets.

Example 3.1 Consider an instance of online matching with a *big* resource that has capacity 2 and a *small* resource that has capacity 1. The first two arrivals have an edge to both resources (bids of 1). The Ranking algorithm starts with a random ranking of the two resources and matches every arrival to an available resource with the best rank. The first arrival is matched uniformly randomly. The second arrival is matched to the big resource with probability (w.p.) 1; either Ranking matches both arrivals to the *big* resource (w.p. 0.5) or matches the second arrival to the *big* resource after discovering (w.p. 0.5) that the budget of *small* resource has been used up. A deterministic algorithm that tries to emulate Ranking will match equal fractions (0.5) of the first arrival to both resources. Now, half of the *small* resource's budget is available at the second arrival and without knowing the budgets, the deterministic algorithm fails to distinguish between the two resources. Given a third arrival that can only be matched to the small resource, Ranking matches it w.p. 0.5 but the deterministic counterpart will fail to save the *small* resources' budget and cannot match any fraction of the final arrival.

Building on this example, we establish the following result.

THEOREM 1. *Every deterministic budget oblivious online algorithm for Adwords has competitive ratio at most 0.5, even on instances with binary bids and large budgets.*

Proof. Let ALG denote a deterministic budget oblivious online algorithm. For a deterministic algorithm, we can construct a worst case arrival sequence and assign resource capacities based on the decisions made by the algorithm. We construct an instance with n resources and binary bids. The arrival sequence has two phases. Given an arrival and a resource, an edge between the two corresponds to a bid of 1 and the absence of an edge corresponds to a bid of 0. Let $B > 0$ be an arbitrary integer value. We assign capacity nB to one resource and capacity B to the other $n - 1$ resources.

Phase one arrivals occur first and this phase has nB arrivals that have an edge to every resource. When ALG matches a resource B times, we reveal (to ALG) that the capacity of the resource is B , i.e., the resource has been used up. We do this up to $n - 1$ times in total, i.e., when a resource i is matched to B arrivals and i is not the sole resource with non-zero remaining capacity, we then inform the algorithm GPG that the capacity of i has been exhausted. There is at least one resource with non-zero remaining capacity at the end of phase one. Let j^* denote one such resource. Phase two has $(n - 1)B$ arrivals that occur after phase one arrivals and have an edge to every resource except j^* . Therefore, ALG can only match these arrivals to resources in $[n] \setminus \{j^*\}$ with available capacity. The capacity of resource j^* is nB and the capacity of all other resources is B .

Note that, optimal offline (OPT) matches every arrival in phase one to resource j^* and matches phase two arrivals to the other resources, fully utilizing the total capacity.

$$\text{OPT} = (2n - 1)B.$$

At the end of both phases, ALG matches at most B arrivals to j^* and therefore, at least $(n - 1)B$ of the total resource capacity is unused in ALG. For $n \rightarrow +\infty$, we have

$$\frac{\text{ALG}}{\text{OPT}} \leq 1 - \lim_{n \rightarrow +\infty} \frac{n - 1}{2n - 1} = 0.5.$$

Since the value of $B > 0$ can be arbitrary, this upper bound also applies to instances with large budgets.

□

4. Randomized Budget Oblivious Algorithm and Analysis

In this section, we state and prove our main results for the Adwords/OBA problem with unknown budgets. Consider the following family of randomized algorithms with parameter $\beta > 0$.

ALGORITHM 1: Generalized Perturbed-Greedy (GPG)

Inputs: Set of advertisers I , parameter β ;

Let $g(x) = e^{\beta(x-1)}$;

For every $i \in I$ generate i.i.d. r.v. $y_i \sim U[0, 1]$;

for every new arrival t do

 Match t to $i^* = \arg \max_{i \in I} b_{i,t}(1 - g(y_i))$;

if i^* is out of budget **then** update $I = I \setminus \{i^*\}$;

Observe that Algorithm 1 (GPG) is budget oblivious and matches every arrival $t \in T$ greedily based on *randomized bid prices* $b_{i,t}(1 - g(y_i)) \forall i \in I$. The uniform random variables $(y_i)_{i \in I}$, called *seeds*, are sampled independently for each $i \in I$. On any given problem instance, ties between bid prices of any two (different) resources occur with a probability of 0.

The choice of function $g(\cdot)$ is influenced by the competitive ratio analysis of the algorithm. In particular, choosing $g(x) = e^{x-1}$ gives the optimal algorithm for closely related settings such as vertex weighted online matching (Aggarwal et al. 2011). Motivated by this, we will focus on analyzing the competitive ratio of GPG for the family of exponential functions $g(x) = e^{\beta(x-1)}$ for $\beta > 0$. Note that in Section 1.1, we presented GPG with seed values $x_i = 1 - y_i$. For consistency with prior work (Aggarwal et al. 2011, Devanur et al. 2013), we use seeds y_i in the subsequent discussion. As we will discuss later, it is challenging to show a non-trivial guarantee for GPG directly. We lower bound the competitive ratio of GPG by first analyzing Algorithm 2 (f -GPG), which is a fractional relaxation of GPG. Clearly, f -GPG is not budget oblivious. However, this is not of concern as the algorithm is only used as an intermediate step to analyze GPG. f -GPG

ALGORITHM 2: Fractional GPG (f -GPG)**Inputs:** Set of advertisers I , budgets $(B_i)_{i \in I}$, parameter β ;Initialize $I(0) = I$, $g(t) = e^{\beta(t-1)}$ and $B_i(0) = B_i$ for every $i \in I$;Generate i.i.d. sample $y_i \in U[0, 1]$ for every $i \in I$;**for** every new arrival t **do** Initialize total fractional of arrival matched $\delta = 0$; Initialize set of available advertisers $I(t) = I(t-1)$ and remaining budgets $B_i(t) = B_i(t-1) \forall i \in I$; **while** $\delta < 1$ **and** $I(t) \cap \{i \mid b_{i,t} > 0\} \neq \emptyset$ **do** Let $i^* = \arg \max_{i \in I(t)} b_{i,t} r_i (1 - g(y_i))$; Let $\delta_{i^*}(t) = \min \left\{ 1, \frac{B_{i^*}(t-1)}{b_{i^*,t}} \right\}$; **if** $\delta_{i^*}(t) > 0$ **then** update $B_{i^*}(t) = B_{i^*}(t-1) - b_{i^*,t} \delta_{i^*}(t)$ and $\delta = \delta + \delta_{i^*}(t)$; Update the set of available advertisers $I(t) = \{i \mid B_i(t) > 0\}$;

is also a randomized algorithm where random seeds $(y_i)_{i \in I}$ are chosen at the beginning. Similar to GPG, ties between bid prices occur with a probability of 0 in f -GPG. The key difference in f -GPG is that each arrival $t \in T$ is matched fractionally to (possibly) multiple resources. For $t \in T$ and some fixed values of the seeds $(y_i)_{i \in I}$, let $I(t)$ denote the set of resources with available budget when t arrives and let $B_i(t)$ denote the available budget of resource i right after arrival t departs. In GPG, we simply match an arrival t to resource i^* that has non-zero budget available and has the highest randomized bid price, whereas, in f -GPG, if the remaining budget of i^* is less than the bid $b_{i^*,t}$, i.e., $B_{i^*}(t-1) < b_{i^*,t}$, then we only match a fraction $\frac{B_{i^*}(t-1)}{b_{i^*,t}}$ of arrival t to i^* . For example, if $B_{i^*}(t-1) = 0.5b_{i^*,t}$, then we match 0.5 fraction of t to i^* and this uses up the remaining budget of i^* . Then, from the remaining set of resources we again select a resource with the highest (randomized) bid price and repeat the process, until t is fully (fractionally) matched or no resource is available. The total amount of resource i 's budget allocated to arrival t is given by $B_i(t) - B_i(t-1)$. Observe that an arrival is fractionally matched to at most $|I|$ different resources and at most $|I|$ arrivals are matched fractionally in f -GPG. We show the following lower bound on the competitive ratio of f -GPG against the optimal offline integral matching.

THEOREM 2. *With $\beta = 1.15$, f -GPG is at least 0.522 competitive against optimal (integer) offline allocation. When $\beta = 1$, f -GPG is at least 0.508 competitive.*

Given the similarities between f -GPG and GPG it might be tempting to conclude similar competitive ratio results for GPG in the small bids regime. Perhaps surprisingly, it turns out that for the same seed values, GPG and f -GPG can sometimes output very different matchings (see Example C.1 in Appendix C). Nonetheless, we show that the objective value of the two algorithms is always similar and using this we establish the following guarantee for GPG.

THEOREM 3. *Given a competitive ratio guarantee η for f -GPG (for some parameter value β), we have that GPG (with the same value of β) is $\frac{1}{1+\gamma}\eta$ competitive.*

COROLLARY 1. *For $\beta = 1.15$, GPG is at least $\frac{1}{1+\gamma}0.522$ competitive. When $\beta = 1$, GPG is at least $\frac{1}{1+\gamma}0.508$ competitive.*

The corollary above follows by combining Theorems 2 and 3. These results hold for any instance of OBA, i.e., for arbitrary value of γ . For small bids we have $\gamma \rightarrow 0$ and the competitive ratio of GPG is strictly better than 0.5. We also establish that the competitive ratio of GPG is strictly less than $(1 - 1/e)$.

THEOREM 4. *For every bid-to-budget ratio $\gamma \in (0, 1]$ and $\beta \geq 1$, the competitive ratio of GPG is at most 0.624.*

The proof of Theorem 4 is analytical but Monte Carlo simulations indicate that an even stronger upper bound of 0.604 may hold for $\beta = 1$. For special cases of Adwords, our analysis yields a stronger guarantee that surpasses these upper bounds, as described in the next lemma⁶.

THEOREM 5. *For $\beta = 1$, GPG is $\frac{1}{1+\gamma}(1 - 1/e)$ competitive for OBA when the bids are decomposable, i.e., $b_{i,t} \in \{0, b_i \times b_i\} \forall i \in I, t \in T$. For the special case of OBA with integer starting capacities and binary bids, GPG is $(1 - 1/e)$ competitive for arbitrary bid-to-budget ratio γ .*

In the following sections, we focus on proving Theorem 2 (Sections 4.1 and 4.2) and Theorem 3 (Section 4.3). In the course of proving these results, we discuss the insights behind the improved guarantees stated in Theorem 5 and defer a formal proof of Theorem 5 to Appendix G. We discuss the main insights behind Theorem 4 in Section 4 and defer the formal proof of this result to Appendix F.

4.1. Overview of the Analysis of f -GPG

We start with some notation. Let the number of resources $|I| = n$. Let OPT refer to the (integer) optimal offline algorithm. Overloading notation, we also use OPT to denote the total reward of the optimal offline algorithm. Since there are no unknowns in the offline problem, i.e., budgets and bids are all known, the optimal offline solution is a deterministic matching. Let OPT_i denote the set of arrivals matched to i as well as the total fraction of i 's budget that is matched in OPT. Note that $\text{OPT} = \sum_{i \in I} \text{OPT}_i$.

Because the discussion prior to Section 4.1.1 applies to both GPG and f -GPG, we use ALG as a unified reference in place of separate references to GPG and to f -GPG. Overloading notation,

⁶ The $(1 - 1/e)$ guarantee for OBA with integer starting capacities and binary bids was also shown concurrently in Albers and Schubert (2021) and Vazirani (2023). Further, the “no-surpassing” condition of Vazirani (2023) generalizes the notion of decomposable bids.

we also use ALG to refer to the expected reward of the online algorithms. Let Y denote the vector of seeds $(y_i)_{i \in I} \in [0, 1]^n$. For fixed seeds Y , ALG is deterministic, i.e., both GPG and $f\text{-GPG}$ are deterministic. For this reason, we also refer to the vector Y as a sample path of ALG . Let $\text{ALG}(Y)$ denote the (fractional) matching generated by ALG on sample path Y . Let $\text{ALG}_i(Y)$ denote the set of arrivals (fractionally) matched to i in the matching $\text{ALG}(Y)$. Overloading notation, we also use $\text{ALG}(Y)$ to denote the total reward of ALG with seed Y and $\text{ALG}_i(Y)$ to denote the total budget of i used in $\text{ALG}(Y)$. Let $\text{ALG}_t(Y)$ denote the set of resources (fractionally) matched to t in $\text{ALG}(Y)$ and let $\delta_{i,t}(Y)$ denote the fraction of t matched to i in $\text{ALG}(Y)$. Observe that in GPG , the set $\text{GPG}_t(Y)$ will never have more than one element.

The (randomized) primal-dual method of Devanur et al. (2013) is a standard technique for analyzing online algorithms for bipartite matching. Unfortunately, it is not obvious to us if one can use this method to obtain a non-trivial lower bound on the competitive ratio of ALG . We illustrate this in more detail in Appendix B. To prove Theorem 2, we use the flexible LP free analysis framework of Goyal et al. (2021). To prove a lower bound on the competitive ratio of an online algorithm ALG in this framework, it suffices to find a feasible solution to the following system of (linear) inequalities in variables λ_t and θ_i ,

$$\sum_{t \in T} \lambda_t + \sum_{i \in I} \theta_i \leq (1 + \varepsilon) \text{ALG} \quad (1)$$

$$\theta_i + \sum_{t \in \text{OPT}_i} \lambda_t \geq \alpha \text{OPT}_i \quad \forall i \in I, \quad (2)$$

$$\lambda_t \geq 0, \quad \theta_i \geq 0 \quad \forall t \in T, i \in I. \quad (3)$$

LEMMA 1. *Given a solution to the system defined by (1)–(3), we have that ALG is $\frac{\alpha}{1+\varepsilon}$ competitive against OPT .*

Proof. Summing up inequalities (2) over all $i \in I$, we have

$$\alpha \text{OPT} = \alpha \sum_{i \in I} \text{OPT}_i \leq \sum_{i \in I} \sum_{t \in \text{OPT}_i} \lambda_t + \sum_{i \in I} \theta_i \stackrel{(*)}{\leq} \sum_{t \in T} \lambda_t + \sum_{i \in I} \theta_i \leq (1 + \varepsilon) \text{ALG},$$

here inequality $(*)$ follows from the fact that the optimal offline solution is integral, i.e., OPT matches each arrival to at most one resource and $\text{OPT}_i \cap \text{OPT}_j = \emptyset$ for any two (distinct) resources $i, j \in I$. $\square$

Due to Lemma 1, our main goal is to find a feasible solution to the system (1)–(3) with a suitably large value of $\frac{\alpha}{1+\varepsilon}$. Recall that $\text{ALG}_t(Y)$ is the set of resources matched to t in $\text{ALG}(Y)$. To define the candidate solution, let

$$\lambda_t(Y) = \sum_{j \in \text{ALG}_t(Y)} b_{j,t} \delta_{j,t}(Y) (1 - g(y_j)) \quad \forall t \in T, Y \in [0, 1]^n, \quad (4)$$

$$\theta_i(Y) = \sum_{t \in \text{ALG}_i(Y)} b_{i,t} \delta_{i,t}(Y) g(y_i) = \text{ALG}_i(Y) g(y_i) \quad \forall i \in I, Y \in [0, 1]^n. \quad (5)$$

Our candidate solution is,

$$\lambda_t = E_Y [\lambda_\tau(Y)] \quad \forall t \in T \text{ and } \quad \theta_i = E_Y [\theta_i(Y)] \quad \forall i \in I, \quad (6)$$

here $E_Y[\cdot]$ denotes expectation over the random seeds in ALG. The main idea behind this candidate solution is as follows. Matching a $\delta_{i,t}(Y)$ fraction of t to i , generates reward $b_{i,t}\delta_{i,t}(Y)$. We split this reward into two parts; $(1 - g(y_i))$ fraction of the reward is added to $\lambda_t(Y)$ and the remaining $g(y_i)$ fraction of the reward is added to $\theta_i(Y)$. Then, we take expectation over the seeds Y . This reward splitting is a natural generalization of the one used Devanur et al. (2013) for vertex weighted online bipartite matching.

The main challenge is to prove that this candidate solution satisfies inequalities (2) for a large enough value of α . At a high level, we follow the same strategy as Devanur et al. (2013). Since it is challenging to analyze the expectation of multi-variate random variables such as $\lambda_t(Y)$ and $\theta_i(Y)$, we fix the seeds for all but one resource, say $i \in I$, and evaluate conditional expectations with respect to (w.r.t.) randomness in y_i . Let $Y_{-i} \in [0, 1]^{n-1}$ denote the vector of seeds of all resources except i . We fix Y_{-i} and establish non-trivial lower bounds on the conditional expectations $E_{y_i}[\lambda_t(y_i, Y_{-i}) \mid Y_{-i}]$ and $E_{y_i}[\theta_i(y_i, Y_{-i}) \mid Y_{-i}]$. Then, combining the lower bounds gives us (2).

To lower bound the conditional expectations, we examine the changes in the matching generated by ALG for different values of y_i . In particular, we compare the matching for any given value $y_i < 1$, with the matching for $y_i = 1$. The latter scenario serves as a base scenario where the bid price of i is 0 everywhere (since $1 - g(1) = 0$). To lower bound $E_{y_i}[\lambda_t(y_i, Y_{-i}) \mid Y_{-i}]$, we show that $\lambda_t(y_i, Y_{-i})$ takes its minimum value in the base scenario, i.e., when $y_i = 1$. This is similar to the “monotonicity” property shown in Devanur et al. (2013). To obtain a sufficiently strong lower bound on $E_{y_i}[\theta_i(y_i, Y_{-i}) \mid Y_{-i}]$, ideally, we hope that as y_i decreases from 1 to 0, the amount of i ’s budget used ($\text{ALG}_i(y_i, Y_{-i})$) increases monotonically at a “fast enough” rate. This corresponds to the “dominance” property in Devanur et al. (2013). All known analysis of (special cases) of the Perturbed Greedy algorithm reply on establishing these two properties (Aggarwal et al. 2011, Devanur et al. 2013, Goyal and Udwani 2023, Eden et al. 2021, Vazirani 2021, Albers and Schubert 2021, Delong et al. 2023).

4.1.1. Main Challenge. Unfortunately, as we illustrate in Example 4.1 below, the “dominance” property does not hold for GPG. In fact, $\text{GPG}_i(y_i, Y_{-i})$, which is the amount of i ’s budget used, may decrease sharply as y_i decreases from 1 to 0. This issue is somewhat mitigated in f -GPG, where $f\text{-GPG}_i(y_i, Y_{-i})$ increases as y_i decreases but the rate at which it increases may be quite small in comparison to what the “dominance” property requires. We overcome this challenge by finding new structural insights into the problem that we use to show a non-trivial lower bound on

$f\text{-GPG}_i(y_i, Y_{-i})$. To obtain this lower bound, our key new insight is to link $f\text{-GPG}_i(y_i, Y_{-i})$ with the total increase in bid prices from the base scenario i.e.,

$$\sum_{t \in T} (\lambda_t(y_i, Y_{-i}) - \lambda_t(1, Y_{-i})).$$

In particular, we show that any overall gain in bid prices originates from the budget utilization of i (Lemma 5). We also give an instance (Example 4.2) where only the budget utilization of i changes with y_i and this new lower bound on $f\text{-GPG}_i(y_i, Y_{-i})$ is tight. Importantly, this result is only true for $f\text{-GPG}$ and it is not true for GPG . The absence of this structural result for GPG is the main reason that we first analyze $f\text{-GPG}$ in order to show a guarantee for GPG . For more details, see Appendix E after reading Section 4.2.

In addition to the lower bound on $f\text{-GPG}_i(y_i, Y_{-i})$, we show that if i 's budget is unused for some $y_i < 1$, then the bid prices have to increase by a certain amount (Lemma 6). By combining these two properties, we establish (2) for $f\text{-GPG}$ with a sufficiently large value of α .

Now, we demonstrate that $\text{GPG}_i(Y)$ may decrease as y_i decreases with a simple example where individual bids are comparable to the overall budget. In Appendix C, we give a similar example in the small bids regime (see Example C.1) and also an example to show that the ‘‘dominance’’ property does not hold for $f\text{-GPG}$ (see Example C.2 after reading Section 4.2).

Example 4.1 Consider an instance with two resources and starting budgets $B_1 = 2$ and $B_2 = 1$. We have two arrivals, t and $t + 1$. Arrival t has bids $b_{1,t} = b_{2,t} = 1$. Arrival $t + 1$ has bids $b_{1,t+1} = 2$ and $b_{2,t+1} = 4$. Consider an execution of GPG with $\beta = 1$ and the random seed y_2 , for resource 2, fixed at 0.5. Observe that,

- (i) $b_{1,t+1}(1 - g(0)) < b_{2,t+1}(1 - g(0.5))$.
- (ii) $b_{1,t}(1 - g(0.5)) = b_{2,t}(1 - g(0.5))$.

Consider the matching generated by GPG as we decrease y_1 from 1 to 0. For $y_1 > 0.5$, GPG matches t to resource 2 and this uses up all of 2's budget. Consequently, $t + 1$ is matched to resource 1. For $y_1 < 0.5$, t is matched to 1 and $t + 1$ to 2. Thus, the amount of resource 1 matched in GPG decreases from 2 to 1 as y_1 decreases. This is somewhat surprising as the bid price of resource 1 increases when we decrease y_1 . At a high level, there are two main reasons behind this phenomenon.

The first (and main) reason is that bid prices vary with both time and seed values Y . In the example above, for $y_1 < 0.5$, the bid price of resource 1 exceeds that of resource 2 at arrival t but the time variation in bids ensures that the bid price of resource 2 dominates that of resource 1 at arrival $t + 1$. A second reason is that due to unknown budgets, GPG may overestimate the available budget of a resource. In the example above, recall that only 1 unit of resource 2 is available at

$t + 1$. For $y_1 < 0.5$, we have $b_{1t+1}r_1(1 - g(0.5)) > r_2(1 - g(y_2))$, i.e., if we account for the remaining budget of each resource, then resource 1 has a higher bid price at arrival $t + 1$ than resource 2. However, GPG, being ignorant of the budget, computes a higher bid price for resource 2 at $t + 1$. Observe that f -GPG is not budget oblivious and never overestimates the available budget of a resource. However, even in f -GPG, for $y_1 < 0.5$, the time variation in bids ensures that when $t + 1$ arrives, the bid price of resource 2 is higher than that of resource 1 and a fraction of arrival $t + 1$ is matched to resource 2.

4.2. Lower Bound on Competitive Ratio of f -GPG (Theorem 2)

Recall that, to lower bound the competitive ratio of f -GPG, it suffices find a feasible solution to the system (1)–(3) with a suitably large value of $\frac{\alpha}{1+\varepsilon}$. Also recall the candidate solution (6) given by $\lambda_t = E_Y [\lambda_t(Y)]$ and $\theta_i = E_Y [\theta_i(Y)]$, where

$$\begin{aligned}\lambda_t(Y) &= \sum_{j \in f\text{-GPG}_t(Y)} b_{j,t} \delta_{j,t}(Y) (1 - g(y_j)) \quad \forall t \in T, Y \in [0, 1]^n, \\ \theta_i(Y) &= \sum_{t \in f\text{-GPG}_i(Y)} b_{i,t} \delta_{i,t}(Y) g(y_i) = f\text{-GPG}_i(Y) g(y_i) \quad \forall i \in I, Y \in [0, 1]^n.\end{aligned}$$

Observe that on every sample path Y , $\lambda_t(Y) > 0 \forall t \in T$ and $\theta_i(Y) > 0 \forall i \in I$. Thus, constraints (3) are satisfied.

LEMMA 2. *For the candidate solution given by (6), we have,*

$$\sum_{i \in I} \theta_i + \sum_{t \in T} \lambda_t = f\text{-GPG},$$

and constraint (1) is satisfied with $\varepsilon = 0$.

Proof. Fix an arbitrary seed Y and consider an arrival $t \in T$. The fraction of t matched to resource i in f -GPG(Y) is given by $\delta_{i,t}(Y)$. This fractional match contributes a reward of $b_{i,t} \delta_{i,t}(Y)$ to f -GPG(Y). By definition of $\theta_i(Y)$ and $\lambda_t(Y)$, the reward $(b_{i,t} \delta_{i,t}(Y))$ is split into two parts such that a fraction $g(y_i)$ of the reward goes to $\theta_i(Y)$ and the remaining $1 - g(y_i)$ fraction goes to $\lambda_t(Y)$. From this observation, it follows that,

$$\sum_{t \in T} \lambda_t(Y) + \sum_{i \in I} \theta_i(Y) = \sum_{i \in I, t \in T} b_{i,t} \delta_{i,t}(Y) = f\text{-GPG}(Y).$$

Taking expectation over Y on both sides completes the proof. $\square$

It remains to show that inequalities (2) are satisfied. In fact, we show a stronger statement as described in the following lemma.

LEMMA 3. Consider a resource $i \in I$ and let seed $Y_{-i} = (y_j)_{j \in I \setminus \{i\}}$ denote the random seed in f -GPG for all resources except $i \in I$. Suppose that for the candidate solution (6), we have,

$$E_{y_i} \left[\theta_i(Y) + \sum_{t \in \text{OPT}_i} \lambda_t(Y) \mid Y_{-i} \right] \geq \alpha \text{OPT}_i \quad \forall Y_{-i} \in [0, 1]^{n-1}, \quad (7)$$

for some value $\alpha > 0$. Then, inequality (2) is satisfied for resource i with the same α .

Proof. The lemma follows by taking expectation over Y_{-i} on both sides of (7). □

4.2.1. Notation and Definitions. We fix an arbitrary resource $i \in I$, seeds Y_{-i} , and establish (7) for a sufficiently large value α . For brevity, we suppress dependence on Y_{-i} from notation and highlight only the dependence on seed $y_i \in [0, 1]$. So $\lambda_t(y_i, Y_{-i})$ and $\theta_j(y_i, Y_{-i})$ are denoted as $\lambda_t(y_i)$ and $\theta_j(y_i)$ respectively. Further, $f\text{-GPG}(y_i)$ is the matching generated by f -GPG when it is executed with seed $Y = (y_i, Y_{-i})$. Similarly, $f\text{-GPG}_i(y_i)$ denotes the total amount of i 's budget allocated in $f\text{-GPG}(y_i)$. Let $I(t, y_i)$ denote the set of resources available in $f\text{-GPG}(y_i)$ when t arrives. Let $f\text{-GPG}_t(y_i)$ denote the set of resource that are fractionally matched to t in $f\text{-GPG}(y_i)$ and let $\delta_{i,t}(y_i)$ denote the fraction of t matched to i .

Overloading notation, we use T to also denote the total number of arrivals. For the analysis, it will be useful to view f -GPG as a continuous process during the interval $[1, T+1)$ where arrival t is matched continuously over the interval $[t, t+1)$. On arrival of t , we start with the set $I(t, y_i)$ of resources and let $I(\tau, y_i)$ denote the set of resources available at any moment $\tau \in [t, t+1)$. During the interval $[\tau, \tau + d\tau)$ we match an infinitesimal fraction $d\tau$ of arrival t to resource,

$$j_\tau = \operatorname{argmax}_{j \in I(\tau, y_i)} b_{j,t} (1 - g(y_j)),$$

consuming an infinitesimally small amount $b_{j_\tau, t} d\tau$ of the budget of resource j_τ . When the budget of resource j_τ is used up, we remove it from the set of available resources and the process continues till time $t+1$. Observe that this is an equivalent way to describe the fractional matching in f -GPG. In the continuous viewpoint, we refine the definition of OPT_i to be the set of all moments $\tau \in [t, t+1)$ such that arrival t is matched to i in OPT . We similarly refine the definition of $\lambda_t(\cdot)$ and let

$$\lambda_\tau(y_i) = \max_{j \in I(\tau, y_i)} b_{j,t} (1 - g(y_j)) \quad \forall \tau \in [t, t+1), t \in T.$$

Observe that $\lambda_t(y_i) = \int_{\tau=t}^{t+1} \lambda_\tau(y_i) d\tau$. To establish (7), we examine the changes in the matching generated by f -GPG for different values of y_i . In particular, we compare the matching for any given value $y_i < 1$, with the matching for $y_i = 1$. The latter scenario serves as a base scenario where

the bid price of i is 0 everywhere. To enable this comparison we define some new objects. For every $i \in I, t \in T$, and $\tau \in [t, t+1)$, we define the *critical threshold* $y_{i,\tau}^c \in [0, 1]$ as the value for which,

$$\textbf{Critical threshold } y_{i,\tau}^c: \quad b_{i,t}(1 - g(y_{i,\tau}^c)) = \max_{j \in I(\tau, 1)} b_{j,t}(1 - g(y_j)).$$

Set $y_{i,\tau}^c = 0$ if no such value exists and $y_{i,\tau}^c = 1$ if the set $I(\tau, 1)$ is empty. Due to the monotonicity of $g(x) = e^{\beta(x-1)}$, we have a unique value of $y_{i,\tau}^c$. Consider the set of distinct critical thresholds,

$$V := \{v \mid y_{i,\tau}^c = v \text{ for some } \tau \in \text{OPT}_i\}.$$

Let

$$b(v) := \sum_{t \in \text{OPT}_i} b_{i,t} \int_{\tau=t}^{t+1} \mathbb{1}(y_{i,\tau}^c = v) d\tau = \sum_{t \in \text{OPT}_i} b_{i,t} \delta_{i,t}(y_i) \quad \forall v \in V,$$

i.e., $b(v)$ is the cumulative bid on i from all moments that are in OPT_i and that have critical threshold v . We also define

$$B(y_i) = \sum_{v \geq y_i; v \in V} b(v) \quad \forall y_i \in [0, 1],$$

$$S(y_i) = \{\tau \mid y_i \leq y_{i,\tau}^c \text{ and } \tau \in \text{OPT}_i\}.$$

$B(y_i)$ is the cumulative bid on i from all moments in OPT_i that have critical threshold at least y_i and $S(y_i)$ is the set of all moments in OPT_i with critical threshold at least y_i . Recall that we use OPT_i to also denote the total fraction of i 's budget that is matched in OPT. Observe that, $B(0) = \sum_{v \in V} b(v) \text{OPT}_i$ and $S(0) = \text{OPT}_i$.

4.2.2. Establishing (2) for the Candidate Solution. At a high level, our approach is to separately lower bound $\lambda_\tau(y_i)$ and $\theta_i(y_i)$ and then combine the two lower bounds together. We start with $\lambda_\tau(y_i)$ and show that the maximum bid price at τ takes its minimum value in the base scenario where $y_i = 1$.

LEMMA 4. *Given $i \in I$ and seed Y_{-i} , for every $y_i \in [0, 1], t \in T$, and $\tau \in [t, t+1)$, we have*

$$\lambda_\tau(y_i) \geq \lambda_\tau(1) \geq b_{i,t}(1 - g(y_{i,\tau}^c)).$$

The proof is included in Appendix D. Next, to lower bound $\theta_i(y_i)$, we establish a lower bound on $f\text{-GPG}_i(y_i)$. Recall that, $f\text{-GPG}_i(y_i)$ is the set of arrivals matched to i in $f\text{-GPG}(y_i)$, as well as, the total budget of i used in $f\text{-GPG}_i(y_i)$. When bids $b_{j,t} \in \{0, b_j \times b_t\}$ for every $j \in I, t \in T$, we show in Appendix G that $f\text{-GPG}_i(y_i) \geq B(y_i)$, a lower bound that matches the ‘‘dominance’’ property. In general, $f\text{-GPG}_i(y_i)$ may be strictly smaller than $B(y_i)$ with non-zero probability (see Example C.2). In fact, there are instances where $\frac{f\text{-GPG}_i(y_i)}{B(y_i)} \rightarrow 0$ for some $y_i \in (0, 1)$ (see Example 4.2). We give a novel lower bound on $f\text{-GPG}_i(Y)$ by linking it with the increase in bid prices over all arrivals. In particular, we show that any overall gain in bid prices originates from the budget utilization of i .

LEMMA 5. Consider a resource $i \in I$ and seed vector $(y_i, Y_{-i}) \in [0, 1]^n$ such that $f\text{-GPG}_i(y_i) < B(y_i)$. Then, we have $f\text{-GPG}_i(y_i) \geq \sum_{v \in V, v \geq y_i} b(v) \frac{g(v) - g(y_i)}{1 - g(y_i)}$.

Proof. Given resource i and seed Y_{-i} , for every $y_i \in [0, 1]$, define

$$\lambda_{net}(y_i) = \int_{\tau=1}^{T+1} \lambda_{\tau}(y_i) d\tau = \sum_{t \in T} \lambda_t(y_i).$$

Using Lemma 4 we have, $\lambda_{net}(y_i) \geq \lambda_{net}(1) \forall y_i \in [0, 1]$. Now, fix a seed y_i such that $f\text{-GPG}_i(y_i) < B(y_i)$. Observe that the main claim follows from the following upper and lower bounds on $\lambda_{net}(y_i) - \lambda_{net}(1)$.

$$f\text{-GPG}_i(y_i) (1 - g(y_i)) \geq \lambda_{net}(y_i) - \lambda_{net}(1) \geq \sum_{v \in V, v \geq y_i} b(v) (g(v) - g(y_i)).$$

Proof of lower bound: Since $f\text{-GPG}_i(y_i) < B(y_i) \leq B_i$, we have that i is available at every moment in $f\text{-GPG}(y_i)$ and at every $\tau \in S(y_i) \cap [t, t+1)$, $f\text{-GPG}$ selects a resource $j \in I$ such that,

$$b_{j,t} (1 - g(y_j)) \geq b_{i,t} (1 - g(y_i)).$$

From this, we have for every $\tau \in S(y_i) \cap [t, t+1)$,

$$\lambda_{\tau}(y_i) - \lambda_{\tau}(1) \geq b_{i,t} (g(y_{i,\tau}^c) - g(y_i)),$$

where we used the following facts (i) For $y_{i,\tau}^c > 0$, we have $\lambda_{\tau}(1) = b_{i,t} (1 - g(y_{i,\tau}^c))$ and (ii) For $y_{i,\tau}^c = 0$ and $\tau \in S(y_i) \cap [t, t+1)$, we have $y_i = 0$ and therefore, $g(y_{i,\tau}^c) - g(y_i) = 0$.

Integrating over all moments in $S(y_i)$, we get

$$\int_{\tau \in S(y_i)} (\lambda_{\tau}(y_i) - \lambda_{\tau}(1)) d\tau \geq \sum_{v \in V, v \geq y_i} b(v) (g(v) - g(y_i)).$$

Finally, from Lemma 4 we have that $\lambda_{net}(y_i) - \lambda_{net}(1) \geq \int_{\tau \in S(y_i)} (\lambda_{\tau}(y_i) - \lambda_{\tau}(1)) d\tau$, completing the proof of the lower bound.

Proof of upper bound: We start by observing that for every seed $y \in [0, 1]$ of resource i ,

$$\lambda_{net}(y) = f\text{-GPG}_i(y) (1 - g(y)) + \sum_{j \in I \setminus \{i\}} f\text{-GPG}_j(y) (1 - g(y_j)).$$

Therefore,

$$\lambda_{net}(y_i) - \lambda_{net}(1) = f\text{-GPG}_i(y_i) (1 - g(y_i)) + \sum_{j \in I \setminus \{i\}} [f\text{-GPG}_j(y_i) - f\text{-GPG}_j(1)] (1 - g(y_j)),$$

where we used the fact that $g(1) = 1$ for every value of $\beta > 0$. Now, the desired upper bound on $\lambda_{net}(y_i) - \lambda_{net}(1)$ follows from the claim that,

$$f\text{-GPG}_j(y_i) \leq f\text{-GPG}_j(1) \quad \forall j \in I \setminus \{i\}.$$

Since $f\text{-GPG}_j(y_i) \leq B_i$, the claim is obviously true when $f\text{-GPG}_j(1) = B_j$, so let $f\text{-GPG}_j(1) < B_j$, i.e., resource j is available at every moment in $f\text{-GPG}(1)$. Therefore, in $f\text{-GPG}(1)$, every moment $\tau \in [t, t+1)$ that is not matched to j , must be matched to a resource with (strictly) higher bid price (since ties between bid prices do not occur except on a probability 0 set of seed values), i.e.,

$$\lambda_\tau(1) > b_{j,t}(1 - g(y_j)).$$

Since $\lambda_\tau(y_i) \geq \lambda_\tau(1)$ (from Lemma 4), we have that every moment that is not matched to j in $f\text{-GPG}(1)$, is not matched to j in $f\text{-GPG}(y_i)$ either. Therefore, $f\text{-GPG}_j(y_i) \leq f\text{-GPG}_j(1)$. $\square$

In Appendix E, we discuss why Lemma 5 does not hold for GPG when we consider the natural candidate solution given by (4)–(6). For $f\text{-GPG}$, the lower bound in Lemma 5 is tight and we give an example below.

Example 4.2 (Tightness of Lemma 5) Consider an instance with resources $i \in [n]$ and arrivals $t \in [2n - 1]$. Let

$$B_i = (1 - \epsilon)^{i-1} \quad \forall i \in [n - 1] \quad \text{and} \quad B_n = \frac{1 - (1 - \epsilon)^n}{\epsilon}.$$

Observe that $B_n = \sum_{i \in [n-1]} B_i$. For $i \in [n - 1]$, resource i has non-zero bids from arrivals $i, i + 1$, and $n + i$ such that,

$$b_{i,i} = b_{i,i+1} = b_{i,n+i} = (1 - \epsilon)^{i-1}.$$

Resource n has non-zero bids from the first n arrivals,

$$b_{n,t} = (1 - \epsilon)^{t-1} \quad \forall t \in [n].$$

OPT matches resource n to the first n arrivals and matches resource i to arrival $n + i$, $\forall i \in [n - 1]$. This fully utilizes the initial budget of every resource. It is not hard to see that $f\text{-GPG}$ outputs an integral matching for all seeds Y . Consider $f\text{-GPG}$ with seeds $y_i = 0.5 \forall i \in [n - 1]$. Notice that, $y_{n,t}^c = 0.5 \forall t \in [n]$. Now, $B(y_n) = B_n$ for $y_n < 0.5$ but for small enough $\delta > 0$, when $y_n = 0.5 - \delta$, $f\text{-GPG}$ matches arrival 1 to resource n and matches arrival t to resource $t - 1 \forall t \in [n] \setminus \{1\}$. Thus, $f\text{-GPG}_n(0.5 - \delta) = \frac{\epsilon}{1 - (1 - \epsilon)^n} B_n$. For $\epsilon = \frac{\beta\delta}{e^{0.5\beta} - 1}$, $\epsilon = \frac{1}{\sqrt{n}}$, and $n \rightarrow +\infty$, we have

$$\frac{\epsilon}{1 - (1 - \epsilon)^n} B_n \rightarrow \epsilon B_n \quad \text{and} \quad \frac{g(0.5) - g(0.5 - \delta)}{1 - g(0.5 - \delta)} B_n \rightarrow \epsilon B_n,$$

which shows that the lower bound in Lemma 5 is tight.

We now combine Lemma 4 and Lemma 5 to give a lower bound on $\theta_i(Y) + \sum_{t \in \text{OPT}_i} \lambda_t(Y)$.

LEMMA 6. *Given $i \in I$ and seed Y_{-i} , let $\mathbb{1}(y_i \leq v)$ indicate the event $y_i \leq v$. Then, for every $y_i \in [0, 1]$, we have*

$$\begin{aligned} \theta_i(y_i) + \sum_{t \in \text{OPT}_i} \lambda_t(y_i) \\ \geq \sum_{v \in V} b(v) \left[1 - g(v) + \mathbb{1}(y_i \leq v) \left(\min \left\{ g(y_i), \frac{g(v) - g(y_i)}{1 - g(y_i)} \right\} \right) \right]. \end{aligned}$$

Proof. Consider $i \in I$ and $Y_{-i} \in [0, 1]^{n-1}$ and fix an arbitrary seed $y_i \in [0, 1]$. As discussed earlier, $f\text{-GPG}_i(y_i)$ may be much smaller than $B(y_i)$. Keeping this in mind, we consider two cases. In the first case, all of i 's budget is used and $f\text{-GPG}_i(y_i)$ is sufficiently large. In the second case, some of i 's budget is unused and we show that the bid prices have to increase by a certain amount. Combining the two cases will give us the desired.

Case I: $f\text{-GPG}_i(y_i) \geq B(y_i)$. Thus,

$$\theta_i(y_i) \geq B(y_i)g(y_i) = \sum_{v \in V} \mathbb{1}(y_i \leq v) b(v) g(y_i).$$

Combining this with the lower bound from Lemma 4, we have

$$\sum_{t \in \text{OPT}_i} \lambda_t(y_i) + \theta_i(y_i) \geq \sum_{t \in \text{OPT}_i} \int_{\tau=t}^{t+1} \lambda_\tau(1) d\tau + \theta_i(y_i) \geq \sum_{v \in V} b(v) (1 - g(v) + \mathbb{1}(y_i \leq v) g(y_i)).$$

Case II: $f\text{-GPG}_i(y_i) < B(y_i)$ ($\leq B_i$), i.e., resource i is available at every moment in $f\text{-GPG}(y_i)$.

First, we establish a refined lower bound on $\sum_{t \in \text{OPT}_i} \lambda_t(y_i)$, followed by a lower bound on $\theta_i(y_i)$.

Since i is available at every moment in $f\text{-GPG}(y_i)$, we have,

$$\lambda_\tau(y_i) \geq b_{i,t} (1 - g(y_i)) = b_{i,t} (1 - g(y_{i,\tau}^c) + g(y_{i,\tau}^c) - g(y_i)) \quad \forall t \in T, \tau \in S(y_i) \cap [t, t+1).$$

Using this we have,

$$\sum_{t \in \text{OPT}_i} \lambda_t(y_i) \geq \sum_{v \in V} b(v) [1 - g(v) + \mathbb{1}(y_i \leq v) (g(v) - g(y_i))]. \quad (8)$$

Now, $\theta_i(y_i) = f\text{-GPG}_i(y_i) g(y_i)$. From Lemma 5,

$$f\text{-GPG}_i(y_i) g(y_i) \geq \sum_{v \in V; v \geq y_i} b(v) \frac{g(v) - g(y_i)}{1 - g(y_i)} g(y_i). \quad (9)$$

Combining (8) and (9) we get,

$$\begin{aligned}
& \theta_i(y_i) + \sum_{t \in \text{OPT}_i} \lambda_t(y_i) \\
& \geq \sum_{v \in V} b(v) \left[1 - g(v) + \mathbb{1}(y_i \leq v) \left(g(v) - g(y_i) + \frac{g(v) - g(y_i)}{1 - g(y_i)} g(y_i) \right) \right] \\
& \geq \sum_{v \in V} b(v) \left[1 - g(v) + \mathbb{1}(y_i \leq v) \frac{g(v) - g(y_i)}{1 - g(y_i)} \right]
\end{aligned}$$

For any given y_i , we could be in the worse of the two cases above, resulting in the following combined lower bound,

$$\theta_i(y_i) + \sum_{t \in \text{OPT}_i} \lambda_t(y_i) \geq \sum_{v \in V} b(v) \left[1 - g(v) + \mathbb{1}(y_i \leq v) \min \left\{ g(y_i), \frac{g(v) - g(y_i)}{1 - g(y_i)} \right\} \right].$$

□

Proof of Theorem 2. Let $\alpha = \min_{v \in [0,1]} \left[1 - g(v) + \int_0^v \left(\min \left\{ g(y), \frac{g(v) - g(y)}{1 - g(y)} \right\} \right) dy \right]$. Taking expectation over the randomness in seed y_i on both sides of the inequality in Lemma 6, we have

$$\begin{aligned}
& E_{y_i} \left[\theta_i(y_i) + \sum_{t \in \text{OPT}_i} \lambda_t(y_i) \mid Y_{-i} \right] \\
& \geq \sum_{v \in V} b(v) \left[1 - g(v) + \int_0^v \left(\min \left\{ g(y), \frac{g(v) - g(y)}{1 - g(y)} \right\} \right) dy \right], \\
& \geq \alpha \sum_{v \in V} b(v) = \alpha r_i \text{OPT}_i.
\end{aligned}$$

It remains to lower bound α . We do this step numerically. For $g(x) = e^{x-1}$, we obtain $\alpha > 0.508$ (minimum at $x = 0.586$) and for $g(x) = e^{1.15(x-1)}$ we obtain $\alpha > 0.522$ (minimum at $x = 0.789$, see Figure 1).

□

Remark: As stated in Theorem 5, GPG with $\beta = 1$ is $(1 - 1/e)$ competitive for special cases of Adwords. We present the proof in Appendix G.

4.3. From Fractional to Integral Algorithm (Theorem 3)

Since GPG is oblivious to budgets but f -GPG is not, there can be a substantial difference between the output of these algorithm for the same seed Y (see Example C.1 in Appendix C). Therefore, we compare GPG not with f -GPG on the same instance but with the performance of f -GPG on a modified instance where the budgets of resources are increased as described next.

Budget augmentation: Consider an instance where the budget of every item is augmented as follows,

$$B_i^a = B_i + \max_{t \in T} b_{i,t} \quad \forall i \in I.$$




Figure 1 Plot of $\alpha(x) = \left[1 - g(x) + \int_0^x \left(\min \left\{g(y), \frac{g(x) - g(y)}{1 - g(y)}\right\}\right) dy\right]$ for $g(x) = e^{1.15(x-1)}$.

For instances with bid-to-budget ratio γ , observe that,

$$\frac{B_i^a}{B_i} \leq 1 + \gamma \quad \forall i \in I.$$

Proof of Theorem 3. Let OPT^a denote the total reward of optimal offline on an instance with augmented budgets B_i^a . Let $f\text{-GPG}^a$ denote the expected total reward of $f\text{-GPG}$ on the augmented instance. For any seed Y , let $\text{GPG}(Y)$ denote the total reward of GPG as well as the matching output by GPG on the original instance with seed Y . Similarly, let $f\text{-GPG}^a(Y)$ denote the total reward of $f\text{-GPG}$ as well as the matching output by $f\text{-GPG}$ on the augmented instance with seed Y . For any resource $i \in I$, let $f\text{-GPG}_i^a(Y)$ denote the total revenue from resource i in $f\text{-GPG}$ on the augmented instance with seed Y . Similarly, let $\text{GPG}_i(Y)$ denote the total revenue from i in GPG on the original instance with seed Y . Finally, let $x_j^{a,f}(t, Y)$ denote the remaining budget of resource $j \in I$ when t arrives in $f\text{-GPG}^a(Y)$. Similarly, let $x_j(t, Y)$ denote the remaining budget of j after arrival $t - 1$ in $\text{GPG}(Y)$.

Now, fix an arbitrary seed Y and a resource $i \in I$. We claim that,

$$\frac{B_i^a}{B_i} \text{GPG}_i(Y) \geq f\text{-GPG}_i^a(Y). \quad (10)$$

Before proving (10), we show that (10) suffices to prove the theorem. Taking expectation over seeds Y and also a summation over $i \in I$ on both sides of (10) gives us,

$$\text{GPG} \geq \frac{1}{1 + \gamma} f\text{-GPG}^a, \quad (11)$$

here we use the fact that $\frac{B_i^a}{B_i} \leq 1 + \gamma \quad \forall i \in I$. Since an allocation that is feasible in the original instance is also feasible after augmenting the budgets, we have

$$\text{OPT}^a \geq \text{OPT}.$$

Given a lower bound of η on the competitive ratio guarantee of f -GPG we have,

$$f\text{-GPG}^a \geq \eta \text{OPT}^a \geq \eta \text{OPT}. \quad (12)$$

Combining (11) and (12) proves the theorem. It remains to show (10) and we establish it in cases.

Case I: All B_i of i 's budget is allocated in $\text{GPG}(Y)$ i.e., $\text{GPG}_i(Y) = B_i$. Then, (10) follows by observing that $f\text{-GPG}_i^a(Y) \leq B_i^a = B_i^a \frac{\text{GPG}_i(Y)}{B_i}$.

Case II: Some of i 's budget is not allocated to any arrival in $\text{GPG}(Y)$. In this case, we show more strongly that $f\text{-GPG}_i^a(Y) \leq \text{GPG}_i(Y)$. Define

$$\begin{aligned} z_j^{a,f}(t, Y) &= \max\{0, B_j - x_j^{a,f}(t, Y)\} = (B_j - x_j^{a,f}(t, Y))^+ \quad \forall j \in I, \\ z_j(t, Y) &= (B_j - x_j(t, Y))^+ \quad \forall j \in I. \end{aligned}$$

Notice that we use the original budget B_i for defining $z_j^{a,f}(\cdot, \cdot)$ as well as $z_j(\cdot, \cdot)$. As some of i 's budget is unused in $\text{GPG}(Y)$ after the final arrival T , we have $z_i(T+1, Y) > 0$.

To prove that $f\text{-GPG}_i^a(Y) \leq \text{GPG}_i(Y)$, it suffices to show that $z_i^{a,f}(T+1, Y) \geq z_i(T+1, Y)$. In fact, we show more strongly that

$$z_j^{a,f}(t, Y) \geq z_j(t, Y) \quad \forall j \in I, t \in T. \quad (13)$$

This is true at $t = 1$. Suppose this is true just prior to arrival of t . We will show that the inequality also holds just prior to arrival of $t + 1$. Then by induction we have the desired.

Given $z_j^{a,f}(t, Y) \geq z_j(t, Y)$, we have

$$\{j \mid z_j(t, Y) > 0\} \subseteq \{j \mid z_j^{a,f}(t, Y) > 0\}, \quad (14)$$

where the LHS is the set of resources available in $\text{GPG}(Y)$ at t . Let e be the resource matched to t in $\text{GPG}(Y)$. Recall that when f -GPG matches an arrival to multiple resources, it matches fractions of the arrival to each resource in descending order of bid prices. Consider resources in the set $\{j \mid z_j(t, Y) > 0\}$ and let w denote the first resource (if any) in this set that is fractionally matched to t in $f\text{-GPG}^a(Y)$. We consider two cases based on whether w exists.

Case II.A: Resource w does not exist, i.e., no resource in the set $\{j \mid z_j(t, Y) > 0\}$ is matched to t in $f\text{-GPG}^a(Y)$. Then,

$$\begin{aligned} z_q^{a,f}(t+1, Y) &= z_q^{a,f}(t, Y) \geq z_q(t, Y) \geq z_q(t+1, Y) \quad \forall q \in \{j \mid z_j(t, Y) > 0\}, \\ z_q^{a,f}(t+1, Y) &\geq 0 = z_q(t, Y) = z_q(t+1, Y) \quad \forall q \in \{j \mid z_j(t, Y) = 0\}. \end{aligned}$$

Case II.B: Resource w exists. Since both GPG and f -GPG match greedily according to bid prices, we have

$$b_{w,t}(1 - g(y_w)) = \max_{\{j | z_j(t, Y) > 0\}} b_{j,t}(1 - g(y_j)) = b_{e,t}(1 - g(y_e)),$$

here the first equality follows from (14) and the fact that w is the first resource in $\{j | z_j(t, Y) > 0\}$ that is matched to t in f -GPG^a(Y). The second equality follows by definition of e . Since the bid prices of two different resources are unequal w.p. 1, we have that,

$$w = e.$$

Now, using $z_e^{a,f}(t, Y) \geq z_e(t, Y)$ and the fact that t is fractionally matched in f -GPG(Y), we have,

$$z_e^{a,f}(t + 1, Y) \geq z_e(t + 1, Y).$$

Further, since $z_e^{a,f}(t, Y) > 0$, at least $\max_{t \in T} b_{e,t}$ of e 's augmented budget is available in f -GPG^a(Y) when t arrives. Hence, in f -GPG^a(Y), no resource from the set $\{j | j \neq e, z_j(t, Y) > 0\}$ is matched (fractionally) to t , i.e.,

$$z_q^{a,f}(t + 1, Y) = z_q^{a,f}(t, Y) \geq z_q(t, Y) = z_q(t + 1, Y) \quad \forall q \in \{j | j \neq e, z_j(t, Y) > 0\}.$$

Similar to Case II.A, we also have,

$$z_q^{a,f}(t + 1, Y) \geq 0 = z_q(t, Y) = z_q(t + 1, Y) \quad \forall q \in \{j | z_j(t, Y) = 0\}.$$

This completes the proof. $\square$

Remark: When every bid is either 0 or 1 (b -matching case), it is easy to see that f -GPG and GPG are identical on every sample path. Since b -matching is a special case of Adwords with decomposable bids ($b_{i,t} \in \{0, b_i \times b_t\} \forall i \in I, t \in T$), and f -GPG is $(1 - 1/e)$ competitive for Adwords with decomposable bids (Appendix G), we have that, GPG is $(1 - 1/e)$ competitive for b -matching (with arbitrary budgets).

4.4. Upper Bound on Competitive Ratio of GPG

The upper bound of $(1 - 1/e)$ for online bipartite matching and its many generalizations (including Adwords), follows from a family of instances where arrivals get pickier over time; the first arrival has an edge to every resource and for every subsequent arrival, one edge is dropped uniformly randomly, i.e., the t -th arrival has exactly $n - t + 1$ edges and the $t + 1$ -th arrival has an edge to all but one (randomly chosen) neighbor of arrival t (Karp et al. 1990). This family of instances has identical bids and budgets; features that are also sufficient for a lower bound of $(1 - 1/e)$ on the competitive ratio of GPG (Theorem 5). To obtain a tighter upper bound on GPG, we start

by constructing an instance where bids are not decomposable and GPG is at most $(1 - 1/e)$ OPT. Then, by a suitable modification of this instance, we obtain a family of instances such that for every $\beta \geq 1$, there exists an instance in the family where GPG is strictly less than $(1 - 1/e)$ OPT.

To describe the main ideas we let $\beta = 1$ and allow arbitrary bid to budget ratio. Consider an instance with $n + 1$ resources and $2n$ arrivals. Each arrival bids 1 on every resource in $\{1, 2, \dots, n\}$. Arrival $t \in [n]$ bids

$$w(t) = \frac{1 - e^{\frac{t}{n+1}-1}}{1 - e^{-1}},$$

on resource $n + 1$. Arrivals $t \geq n + 1$ do not have an edge to resource $n + 1$. Resources $\{1, 2, \dots, n\}$ have unit budget and resource $n + 1$ has a budget of $\sum_{t \in [n]} w(t)$. The optimal matching allocates resource $n + 1$ to every $t \in [n]$ and matches arrival t to resource $t - n$ for $t \geq n + 1$. We have,

$$\text{OPT} = n + \sum_{t \in [n]} w(t),$$

here $\lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{t \in [n]} w(t) \rightarrow \frac{1}{e-1}$. Therefore, $\lim_{n \rightarrow +\infty} \frac{1}{n} \text{OPT} \rightarrow \frac{e}{e-1}$.

Let y_i denote the random seed of resource $i \in [n + 1]$ in GPG. For $k \in [n]$, let z_k denote the k -th order statistic of the random variables $\{y_i\}_{i \in [n]}$, i.e., the k -th smallest seed value. For $k \in [n]$, random variable z_k has mean $\frac{k}{n+1}$. Let $I(t)$ denote the (random) set of resources available at t in GPG. Resource $n + 1$ is matched to arrival $t \in [m]$ if

$$w(t)(1 - e^{y_{n+1}-1}) > \max_{i \in I(t)} (1 - e^{y_i-1}) \geq (1 - g(z_t)),$$

here the last inequality follows from the fact that the t -th lowest seeded resource in $[n]$ will not be matched to any of the arrivals preceding t . For the sake of intuition, suppose that for $n \rightarrow +\infty$, we have $z_t \approx E[z_t] = \frac{t}{n+1} \forall t \in [n]$, with high probability (w.h.p.). Then, at a high level, we argue as follows. Resource $n + 1$ is matched to t only if $y_{n+1} \approx 0$, which occurs with a vanishingly small probability. By linearity of expectation, the expected total consumed capacity of resource $n + 1$ is ≈ 0 . Therefore, the expected value of $\frac{1}{n} \text{GPG}$ is ≈ 1 , and we have $\frac{\text{GPG}}{\text{OPT}} \approx (1 - 1/e)$.

Now, consider a family of instances parameterized by $\epsilon \geq 0$. Every instance in the family is similar to the instance above except that the bid from arrival $t \in [n]$ to resource $n + 1$ is,

$$w(t, \epsilon) = \frac{1 - e^{\frac{t}{n+1}-1}}{1 - e^{-1+\epsilon}},$$

and the capacity of resource $n + 1$ is $\sum_{t \in [n]} w(t, \epsilon)$. All other bids and parameters are left unchanged. Observe that for a fixed t , $w(t, \epsilon)$ is a strictly increasing function of ϵ . For non-zero ϵ , the value of offline (OPT) increases by

$$\sum_{t \in [n]} (w(t, \epsilon) - w(t, 0)) = \sum_{t \in [n]} w(t, 0) \frac{e^{-1}(e^\epsilon - 1)}{1 - e^{-1+\epsilon}} = O(\epsilon n).$$

To obtain an upper bound that is strictly less than $(1 - 1/e)$, we show that the expected total value of GPG increases by an amount that is significantly lower than $O(\epsilon n)$. The proof is presented in Appendix F. We would like to note that for GPG with $\beta = 1$, performing Monte Carlo simulations with $n = 5000$ indicates a stronger upper bound of 0.604 when bids from arrivals $t \in [n]$ to resource $n + 1$ are set as follows,

$$\hat{w}(t, \epsilon) = \frac{1 - e^{\frac{t}{n+1} - 1 - \epsilon}}{1 - e^{-1}},$$

with $\epsilon = 0.2$.

5. Extensions and Applications

In this section, we demonstrate the flexibility and usefulness of budget oblivious algorithms through different settings. Our first two results are in settings that are different from Adwords. Then, we discuss potential applications of budget oblivious allocation in automated budget management for Adwords.

5.1. Online Matching with Multi-channel Traffic

Inspired by applications in two-sided matching platforms such as VolunteerMatch – where volunteers are matched to nonprofits – Manshadi et al. (2022) introduce a setting where volunteer traffic comes from two types of sources. *Internal* traffic consists of volunteers who arrive on the platform with an interest in (possibly) many nonprofits and the platform selects a set of opportunities to recommend. *External* traffic models volunteers who directly arrive at a nonprofit’s volunteering page via an external link that bypasses the platform’s matching/recommendation system. Manshadi et al. (2022) model this as an instance of online matching where nonprofit opportunities with a finite need for volunteers correspond to offline resources with finite capacity and each volunteer corresponds to an online arrival. Upon arrival, a volunteer is probabilistically matched to at most one opportunity. External arrivals have an edge to exactly one resource. They show that when the number of external arrivals is a small fraction of the total capacity of all resources, a natural generalization of the algorithm of Mehta et al. (2007), which is oblivious to the source of arrivals, achieves the optimal competitive ratio for the problem. However, when external arrivals are a significant fraction of the total capacity of all resources, one can achieve strictly better performance with an algorithm that reacts to external and internal arrivals in different ways. For a deterministic version of their (stochastic) problem, we show that the budget (and arrival source) oblivious GPG algorithm achieves the best possible guarantee for every possible composition of the arrival sequence. This closes the (non-zero) gap between the best known guarantee and the best known upper bound shown for the deterministic version of the setting in Manshadi et al. (2022).

Formally, we consider a setting with I resources that have arbitrary capacities $(B_i)_{i \in I}$. Every arrival $t \in T$ comes with a set of edges to resources $(i, t) \in E$. Matching an arrival to a resource uses

one unit of the resource's capacity and an arrival can only be matched to a neighboring resource with available capacity. The goal is to match the maximum number of arrivals. Every arrival is either internal or external. Let T^{In} and T^{Ex} denote the set of internal arrivals and external arrivals, respectively. Let T_i^{Ex} denote the set of external arrivals with an edge to resource $i \in I$. Manshadi et al. (2022) define the *fraction of external traffic* (of a given instance) as the fraction of total capacity that can be matched to external arrivals.

$$\text{Effective Fraction of External Traffic: } \delta = \frac{\sum_{i \in I} \min\{|T_i^{\text{Ex}}|, B_i\}}{\sum_{i \in I} B_i}.$$

For any given δ , this setting corresponds to OBA with binary bids and without the large capacity assumption. Therefore, GPG with $\beta = 1$ is $(1 - 1/e)$ competitive for this setting (Theorem 5). We analyze the competitive ratio of GPG with $\beta = 1$ as a function of δ . Since $\beta = 1$ leads to the best possible guarantee, we hereby fix β at 1 and refer to GPG with $\beta = 1$ as simply GPG. Note that, Manshadi et al. (2022) consider a model where matches can fail with some probability and the starting capacities are large. We do not make the large capacity assumption but consider a deterministic setting where every match succeeds with probability one.

Our main goal is to find how the performance of GPG changes with δ . While GPG is at least $(1 - 1/e)$ competitive for every $\delta \in [0, 1]$, at a high level, as $\delta \rightarrow 1$, the competitive ratio of GPG approaches 1. To see this, consider an arbitrary resource $i \in I$ and sample path Y in GPG. On this sample path, GPG matches every arrival in T_i^{Ex} provided there is some budget remaining. Thus, GPG uses at least $\min\{|T_i^{\text{Ex}}|, B_i\}$ of resource i 's budget on every sample path. Now, by definition of δ , the expected value of GPG is at least $\delta \sum_{i \in I} B_i \geq \delta \text{OPT}$. Overall, we have

$$\text{GPG} \geq \max\{\delta, (1 - 1/e)\} \text{OPT}.$$

This gives a lower bound on the competitive ratio that improves as δ increases but the dependence on δ is not tight. For example, we show that GPG is $1 + (1 - 1/e) \ln(1 - 1/e)$ (> 0.71) competitive for $\delta = (1 - 1/e)$. From the upper bound shown in Manshadi et al. (2022), this is the best guarantee achievable by any online algorithm.

THEOREM 1 IN MANSHADI ET AL. (2022). *For any value of $\delta \in [0, 1]$ and any resource capacities, no (randomized) online algorithm can achieve a competitive ratio better than $\max\{(1 - 1/e), 1 + \delta \ln \delta\}$.*

We show that the competitive ratio of GPG matches this upper bound for all values of $\delta \in [0, 1]$. This improves the competitive ratio achieved by the novel Adaptive Capacity (AC) algorithm proposed in Manshadi et al. (2022)⁷. It is perhaps surprising that, unlike the AC algorithm, GPG

⁷ Unfortunately, there is no closed form expression for the guarantee of AC. See Theorem 2 in Manshadi et al. (2022) for the factor revealing LP that gives a lower bound on the competitive ratio of AC. Figure 2 in Manshadi et al. (2022) illustrates that there is a non-zero gap between the known competitive ratio of AC and the upper bound.

achieves the optimal guarantee without requiring any information about the channel that each arrival belongs to. We would like to note that this result does not require the new structural properties (such as Lemma 5) that we show for the Adwords problem.

5.1.1. Overview of Analysis. In the same vein as Manshadi et al. (2022), we lower bound the competitive ratio by combining the LP free analysis technique with a factor revealing program (FRP). At a high level, every possible instance of the problem can be projected to a feasible solution for FRP with objective value equal to the ratio $\frac{\text{GPG}}{\text{OPT}}$ on that instance. Computing the optimal value of the FRP gives a lower bound on the worst case ratio, i.e., the competitive ratio of GPG. In general, such a program may be intractable to analyze but given the right set of constraints one can find a strong lower bound on the optimal value of the program. We derive the constraints of our FRP using the LP free framework and structural properties of GPG and OPT. We start with some useful structural properties of worst case instances that simplify the analysis.

LEMMA 7. *For any given $\delta \in [0, 1]$, to lower bound the competitive ratio of GPG it suffices to consider instances where*

- (i) *OPT matches every arrival.*
- (ii) *Internal traffic arrives before external traffic, i.e., there is no internal traffic after the first external arrival.*

We present the proof in Appendix H.3. To state the FRP, let $G(x) = \int_0^x g(u)du = e^{x-1} - e^{-1}$ and notice that $G(1) = (1 - 1/e)$. The FRP has decision variables $\{o_i^{\text{In}}, o_i^{\text{Ex}}, x_i^{\text{In}}, x_i^{\text{Ex}}\}_{i \in I}$ and ζ .

$$\text{FRP :} \quad \inf_{\zeta, \{o_i^{\text{In}}, o_i^{\text{Ex}}, x_i^{\text{In}}, x_i^{\text{Ex}}\}_{i \in I}} \quad \zeta$$

$$\text{s.t.} \quad o_i^{\text{In}} + o_i^{\text{Ex}} \leq B_i \quad \forall i \in I, \quad (15)$$

$$x_i^{\text{Ex}} \leq o_i^{\text{Ex}} \quad \forall i \in I, \quad (16)$$

$$\sum_{i \in I} o_i^{\text{Ex}} = \delta \sum_{i \in I} B_i, \quad (17)$$

$$\sum_{i \in I} x_i^{\text{Ex}} \geq (\zeta - 1 + \delta) \sum_{i \in I} B_i, \quad (18)$$

$$\zeta \sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}}) \geq \sum_{i \in I} \left(G(1) o_i^{\text{In}} + x_i^{\text{Ex}} + o_i^{\text{Ex}} G \left(1 - \frac{\sum_{j \in I} x_j^{\text{Ex}}}{\sum_{j \in I} o_j^{\text{Ex}}} \right) \right). \quad (19)$$

THEOREM 6. *For any given $\delta \in (0, 1]$, the optimal value of FRP is a lower bound on the competitive ratio of f -GPG. For $\delta \geq \frac{1}{e}$, the optimal value of FRP is at least $1 + \delta \ln \delta$.*

We give an overview of the first part of the Theorem 6 here and present the full proof of the theorem in Appendix H. The first part states that the optimal value of FRP is a lower bound

on the competitive ratio. The main idea is as follows. Given an instance with fraction of external traffic δ , we show that there exists a feasible solution to FRP with objective value equal to the performance ratio $\frac{\text{GPG}}{\text{OPT}}$ on that instance. Consider the following candidate solution for the FRP.

Candidate solution for FRP: For resource $i \in I$, let o_i^{Ex} be the total budget of i allocated to external arrivals in OPT. Let x_i^{Ex} be the expected budget of i allocated to external arrivals in f -GPG. Let o_i^{In} and x_i^{In} be the (expected) budget of i allocated to internal arrivals in OPT and in f -GPG, respectively. Finally, let

$$\zeta = \frac{\sum_{i \in I} (x_i^{\text{In}} + x_i^{\text{Ex}})}{\sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}})}.$$

Observe that ζ is equal to the performance ratio $\frac{\text{GPG}}{\text{OPT}}$. We start by noting that for $\delta > 0$, all constraints (including (19)) are well defined. In particular, for $\delta > 0$, using Lemma 7(i) we have $\sum_{i \in I} o_i^{\text{Ex}} > 0$ on every non-trivial instance of the problem. By using the structure of the instance and borrowing ideas from Manshadi et al. (2022), it can be shown that the candidate solution satisfies constraints (15)–(18). The crux of our analysis is to show that inequality (19) is satisfied and we give a detailed proof in Appendix H.1.

5.2. Online Matching with Stochastic Rewards

Introduced by Mehta and Panigrahi (2012), the problem of online matching with stochastic rewards generalizes online bipartite matching by associating a probability of success $p_{i,t}$ with every pair $i \in I, t \in T$. When a match is made, i.e., edge is chosen, it succeeds independently with this probability. If the match succeeds, the resource cannot be matched to any other arrival. If the match fails, the arrival departs but the resource is available for future rematch. The edge probabilities are revealed sequentially with each arrival. When edge probabilities are binary, i.e., $p_{i,t} \in \{0, 1\} \forall i \in I, t \in T$, this setting reduces to online bipartite matching. The natural greedy algorithm that matches each arrival to an available resource with the highest success probability is 0.5 competitive for this problem (Golrezaei et al. 2014). In general, this is the best known result but better algorithms are known for several well studied special cases (Mehta and Panigrahi 2012, Mehta et al. 2015, Goyal and Udwani 2023, Huang and Zhang 2020). We discuss the state-of-the-art below and refer to Goyal and Udwani (2023) and Huang and Zhang (2020) for a more detailed review.

- (i) **Decomposable probabilities:** Goyal and Udwani (2023) introduced the special case where edge probabilities are decomposable i.e., $p_{i,t} \in \{0, p_i \times p_t\} \forall i \in I, t \in T$. They showed that Algorithm 3 with $\beta = 1$ is $(1 - 1/e)$ competitive. Udwani (2024) recently gave a simpler and more general proof of this result. Note that Algorithm 3 corresponds to the GPG algorithm with bids replaced by edge probabilities. Since online bipartite matching is a special case of the setting of decomposable probabilities, $(1 - 1/e)$ is also the best possible competitive ratio guarantee.

- (ii) **Vanishing probabilities:** Another well studied special case is when the edge probabilities are vanishingly small, i.e., $p_{i,t} \rightarrow 0 \forall i \in I, t \in T$. Goyal and Udwani (2023) showed that a *deterministic* algorithm, originally proposed by Mehta et al. (2015) and distinct from Perturbed Greedy, has competitive ratio guarantee of at least 0.596 in this setting.

ALGORITHM 3: Generalized Perturbed-Greedy for Stochastic Rewards

Inputs: Set of resources I , parameter β ;

Let $g(x) = e^{\beta(x-1)}$;

For every $i \in I$ generate i.i.d. r.v. $y_i \sim U[0, 1]$;

for every new arrival t **do**

 Match t to $i^* = \arg \max_{i \in I} p_{i,t}(1 - g(y_i))$;

if match succeeds **update** $I = I \setminus \{i^*\}$;

Inspired by these positive results, a perhaps natural question is if the algorithms developed for special cases point to a candidate algorithm that may outperform greedy for the general problem setting. Since the setting of stochastic rewards includes online bipartite matching as a special case, no deterministic algorithm can beat greedy in general (Karp et al. 1990). This rules out all the algorithms that are known to beat greedy in the special case of vanishing probabilities. That leaves Algorithm 3, which gives the optimal competitive ratio guarantee for decomposable probabilities (for $\beta = 1$). Previously, it was not known if this algorithm outperforms greedy for vanishing probabilities. Goyal and Udwani (2023) highlight some challenges with proving such a result and leave the analysis of Algorithm 3 for vanishing probabilities as an open problem. We show that Algorithm 3 with $\beta = 1$ is at least 0.508 for vanishing probabilities. While this does not beat the best known competitive ratio result for vanishing probabilities, it gives a single algorithm (namely Algorithm 3) that outperforms greedy in *all* previously studied special case of the problem. Recently, Udwani (2024) showed that in general, the competitive ratio of Algorithm 3 is strictly less than $(1 - 1/e)$.

THEOREM 7. *For online matching with stochastic rewards and vanishing probabilities, i.e., $\max_{i \in I, t \in T} p_{i,t} \rightarrow 0$, Algorithm 3 is 0.522 competitive for $\beta = 1.15$ and 0.508 competitive for $\beta = 1$.*

We present the proof in Appendix I. The result follows directly from an equivalence between the setting of vanishing probabilities and Adwords with unknown (stochastic) budgets. This equivalence was first observed by Mehta and Panigrahi (2012) and later generalized by Huang and Zhang (2020) and Goyal and Udwani (2023). The result states an instance of online matching with vanishing probabilities is equivalent to an instance of Adwords with unknown budgets (on the same graph) with bids equal to the edge probabilities, i.e., $b_{i,t} = p_{i,t} \forall i \in I, t \in T$, and unknown budgets sampled independently for each resource from the exponential probability distribution with unit mean, i.e., $B_i \sim \text{Exp}(1) \forall i \in I$.

5.3. Advantage of Budget Obliviousness in Automated Budget Management

The high level goal for a typical advertiser is to target their customers through multiple ad campaigns and marketing channels. Starting with an overall budget, an advertiser must determine a good distribution of their budget to individual ad campaigns. For search ads, advertisers must also determine the bids for relevant key words. While these decisions play a crucial role in the success of their ad campaigns, determining a good distribution of budgets between diverse options and deciding optimal bids for specific campaigns can be an incredibly challenging task for any advertiser. As a result, many advertisers rely on automated bidding and budget management tools. The motivation behind these tools is to improve performance for advertisers while simplifying the usage of ad platforms (Aggarwal et al. 2019). Budget oblivious allocation algorithms, such as GPG, do not require fixed initial budget as input. This increases the degrees of freedom available to automated budget management tools and may allow significantly better outcomes for both the advertisers and the platform. We illustrate via a stylized example in Appendix J.

6. Conclusion

We considered the classic Adwords setting with unknown budgets and showed that a natural generalization of Perturbed Greedy algorithm, that computes random bid prices for resources without using budget information, is 0.522 competitive. This is the first result that improves on the guarantee of 0.5 obtained by the greedy algorithm. To show the result, first, we analyze the fractional version of the algorithm using recent innovations in analysis of online matching algorithms, alongside various novel structural insights. We also showed that no deterministic algorithm can do better than greedy and also gave an upper bound of 0.624 on our randomized algorithm (and observed a stronger upper bound of 0.604 in simulations). Finally, we demonstrated the usefulness of budget oblivious algorithms for online resource allocation in a variety of applications. An immediate open question from our work is to find out the best possible competitive ratio guarantee for the Adwords problem with unknown budgets.

Acknowledgments

We thank the anonymous referees for their valuable and helpful feedback. We are also grateful to Vijay Vazirani for many insightful discussions on the topic of this paper. Finally, we thank Will Ma for pointing us to a simpler proof of Theorem 1 and thank Sebastian Schubert for pointing out an error in an earlier version of this paper.

References

Aggarwal G, Badanidiyuru A, Mehta A (2019) Autobidding with constraints. *International Conference on Web and Internet Economics*, 17–30 (Springer).

- Aggarwal G, Goel G, Karande C, Mehta A (2011) Online vertex-weighted bipartite matching and single-bid budgeted allocations. *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, 1253–1264 (SIAM).
- Alaei S, Hajiaghayi M, Liaghat V (2012) Online prophet-inequality matching with applications to ad allocation. *Proceedings of the 13th ACM Conference on Electronic Commerce*, 18–35.
- Albers S, Schubert S (2021) Optimal Algorithms for Online b-Matching with Variable Vertex Capacities. Wootters M, Sanità L, eds., *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, volume 207 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 2:1–2:18 (Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik).
- Arnold BC, Balakrishnan N, Nagaraja HN (2008) *A First Course in Order Statistics* (Society for Industrial and Applied Mathematics), URL <http://dx.doi.org/10.1137/1.9780898719062>.
- Balseiro S, Lu H, Mirrokni V (2020) Dual mirror descent for online allocation problems. III HD, Singh A, eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 613–628 (PMLR).
- Birnbaum B, Mathieu C (2008) On-line bipartite matching made simple. *ACM SIGACT News* 39(1):80–87.
- Buchbinder N, Jain K, Naor JS (2007) Online primal-dual algorithms for maximizing ad-auctions revenue. *European Symposium on Algorithms*, 253–264 (Springer).
- Delong S, Farhadi A, Niazadeh R, Sivan B, Udmani R (2023) Online bipartite matching with reusable resources. *Mathematics of Operations Research* .
- Devanur NR, Hayes TP (2009) The adwords problem: online keyword matching with budgeted bidders under random permutations. *Proceedings of the 10th ACM conference on Electronic commerce*, 71–78 (ACM).
- Devanur NR, Jain K, Kleinberg RD (2013) Randomized primal-dual analysis of ranking for online bipartite matching. *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, 101–107 (Society for Industrial and Applied Mathematics).
- Devanur NR, Jain K, Sivan B, Wilkens CA (2019) Near optimal online algorithms and fast approximation algorithms for resource allocation problems. *J. ACM* 66(1), ISSN 0004-5411, URL <http://dx.doi.org/10.1145/3284177>.
- Eden A, Feldman M, Fiat A, Segal K (2021) An economics-based analysis of ranking for online bipartite matching. *Symposium on Simplicity in Algorithms (SOSA)*, 107–110 (SIAM).
- Feldman J, Mehta A, Mirrokni V, Muthukrishnan S (2009) Online stochastic matching: Beating $1-1/e$. *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, 117–126 (IEEE).
- Goel G, Mehta A (2008) Online budgeted matching in random input models with applications to adwords. *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, 982–991 (Society for Industrial and Applied Mathematics).

-
- Golrezaei N, Nazerzadeh H, Rusmevichientong P (2014) Real-time optimization of personalized assortments. *Management Science* 60(6):1532–1551.
- Goyal V, Iyengar G, Udwani R (2021) Asymptotically optimal competitive ratio for online allocation of reusable resources. Feldman M, Fu H, Talgam-Cohen I, eds., *Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings*, volume 13112 of *Lecture Notes in Computer Science*, 543 (Springer), URL <https://link.springer.com/content/pdf/bbm%3A978-3-030-94676-0%2F1>.
- Goyal V, Udwani R (2023) Online matching with stochastic rewards: Optimal competitive ratio via path-based formulation. *Oper. Res.* 71(2):563–580, URL <http://dx.doi.org/10.1287/OPRE.2022.2345>.
- Huang Z, Zhang Q (2020) Online primal dual meets online matching with stochastic rewards: configuration lp to the rescue. *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 1153–1164.
- Huang Z, Zhang Q, Zhang Y (2020) Adwords in a panorama. *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, 1416–1426 (IEEE).
- Kalyanasundaram B, Pruhs KR (2000) An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science* 233(1-2):319–325.
- Kapralov M, Post I, Vondrák J (2013) Online submodular welfare maximization: Greedy is optimal. *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, 1216–1225 (SIAM).
- Karande C, Mehta A, Tripathi P (2011) Online bipartite matching with unknown distributions. *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 587–596 (ACM).
- Karp RM, Vazirani UV, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, 352–358 (ACM).
- Liang J, Tang ZG, Xu YE, Zhang Y, Zhou R (2023) On the perturbation function of ranking and balance for weighted online bipartite matching 274:80:1–80:15, URL <http://dx.doi.org/10.4230/LIPICS.ESA.2023.80>.
- Manshadi V, Rodilitz S, Saban D, Suresh A (2022) Online algorithms for matching platforms with multi-channel traffic. *Proceedings of the 23rd ACM Conference on Economics and Computation*, 986–987, EC '22 (New York, NY, USA: Association for Computing Machinery), ISBN 9781450391504, URL <http://dx.doi.org/10.1145/3490486.3538326>.
- Manshadi VH, Gharan SO, Saberi A (2012) Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37(4):559–573.
- Mehta A, Panigrahi D (2012) Online matching with stochastic rewards. *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, 728–737 (IEEE).
- Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *J. ACM* 54(5):22–es, ISSN 0004-5411.

- Mehta A, Waggoner B, Zadimoghaddam M (2015) Online stochastic matching with unequal probabilities. *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, 1388–1404 (Society for Industrial and Applied Mathematics).
- Mehta A, et al. (2013) Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8(4):265–368.
- Mirroknii VS, Gharan SO, Zadimoghaddam M (2012) Simultaneous approximations for adversarial and stochastic online budgeted allocation. *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, 1690–1701 (Society for Industrial and Applied Mathematics).
- Udwani R (2024) When stochastic rewards reduce to deterministic rewards in online bipartite matching. *2024 Symposium on Simplicity in Algorithms (SOSA)*, 321–330 (SIAM).
- Vazirani VV (2021) Randomized online algorithms for adwords. *arXiv preprint arXiv:2107.10777* .
- Vazirani VV (2023) Towards a practical, budget-oblivious algorithm for the adwords problem under small bids. Bouyer P, Srinivasan S, eds., *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2023, December 18-20, 2023, IIIT Hyderabad, Telangana, India*, volume 284 of *LIPICs*, 21:1–21:14 (Schloss Dagstuhl - Leibniz-Zentrum für Informatik), URL <http://dx.doi.org/10.4230/LIPICS.FSTTCS.2023.21>.

Appendix A: Further Discussion of Concurrent Work

Albers and Schubert (2021) and Vazirani (2021, 2023) concurrently show that the Perturbed Greedy algorithm is $(1 - 1/e)$ competitive for Adwords with binary bids and arbitrary resource budgets, similar to Theorem 5. The overall proof structure in these results is essentially the same; each proof is based on finding a feasible solution to a linear system. Further, the candidate solutions proposed in the papers are also essentially the same and the proof of feasibility relies on establishing “monotonicity” and “dominance” properties similar to Devanur et al. (2013). The minor differences arise from the linear system used in each paper. Albers and Schubert (2021) perform a primal-dual analysis based on a configuration linear program (LP) introduced for a related problem by Huang and Zhang (2020). Vazirani (2023) generalize the ‘economic viewpoint’ of Eden et al. (2021) to obtain a linear system that closely resembles that of Albers and Schubert (2021) (which is stated below).

$$\begin{aligned} \sum_{t \in T} \lambda_t + \sum_{i \in I} \theta_i &\leq \text{GPG}, \\ \sum_{t \in \Delta} \lambda_t + \theta_i &\geq \alpha \min\{|\Delta|, B_i\} \quad \forall i \in I, \Delta \subseteq T, \\ \lambda_t, \theta_i &\geq 0. \end{aligned}$$

Comparing this with the constraints (1)–(3) that we introduce in Section 4.1, the main difference is that we only impose the second constraint for a specific set OPT_i , which is the set of arrivals matched to i in the optimal offline solution. This difference does not play an important role here but can make a significant difference in settings with stochastic elements (see, for example, (Goyal and Udwani 2023, Goyal et al. 2021)).

These approaches do not yield a performance guarantee better than 0.5 in the general Adwords setting. In particular, Vazirani (2023) identifies a key structural property, called *no surpassing*, the absence of which prevents a generalization of their result to the Adwords setting. For the general setting, we replace the “dominance” with a new lower bound (shown in Lemma 5) and analyze a fractional algorithm (f -GPG) in order to analyze GPG.

Appendix B: Obstacles with Using Classic Primal-Dual Analysis

The primal-dual framework of Buchbinder et al. (2007) and Devanur et al. (2013) is a general technique for proving guarantees for online matching and related problems. To describe the framework, consider the following primal and dual problems that upper bound the optimal offline solution for the Adwords problem.

$$\begin{aligned} \textbf{Primal:} \quad & \min \sum_{i \in I, t \in T} b_{i,t} x_{it} \\ & s.t. \quad \sum_{t \in T} b_{i,t} x_{it} \leq B_i \quad \forall i \in I, \\ & \quad \sum_{i \in I} x_{it} \leq 1 \quad \forall t \in T, \\ & \quad x_{it} \geq 0 \quad \forall i \in I, t \in T. \\ \textbf{Dual:} \quad & \min \sum_{t \in T} \lambda_t + \sum_{i \in I} B_i \theta_i \\ & s.t. \quad \lambda_t + b_{i,t} \theta_i \geq b_{i,t} \quad \forall i \in I, t \in T \\ & \quad \lambda_t, \theta_i \geq 0 \quad \forall t \in T, i \in I. \end{aligned}$$

Primal-dual certificate Devanur et al. (2013): To prove α competitiveness for GPG, it suffices to find a set of non-negative values λ_t, θ_i such that,

$$(i) \quad \lambda_t + b_{i,t} \theta_i \geq \alpha b_{i,t} \quad \forall i \in I, t \in T,$$

$$(ii) \quad \varepsilon \text{ GPG} \geq \sum_{t \in T} \lambda_t + \sum_{i \in I} B_i \theta_i,$$

To understand the obstacles in using this framework to analyze GPG for OBA, we start by defining a natural candidate solution for the system based on the decisions of GPG. Let GPG denote Algorithm 1 as well as its expected reward. In GPG, let $g(t) = e^{t-1}$, i.e., $\beta = 1$. Given that we have external randomness in GPG through $Y = (y_i)_{i \in I}$, we shall define variables λ_t^Y, θ_i^Y and subsequently set $\lambda_t = E_Y[\lambda_t^Y]$ and $\theta_i = E_Y[\theta_i^Y]$. Inspired by Devanur et al. (2013), we set λ_t^Y and θ_i^Y as follows. Initialize all dual variables to 0. Conditioned on Y , for any match (i, t) in GPG set,

$$\lambda_t^Y = b_{i,t}(1 - g(y_i)) \text{ and increment } \theta_i^Y \text{ by } \frac{b_{i,t}}{B_i} g(y_i). \quad (20)$$

Clearly, λ_t^Y is set uniquely since GPG offers at most one resource i to arrival t , and θ_i^Y takes a non-zero value only if it is also accepted by some t , and if this occurs θ_i^Y is never re-set. The following lemma (stated without proof) declares that this candidate solution satisfies constraint (ii) with ε close to 1 in the small bid regime.

LEMMA 8. *For the candidate solution given by (20), constraint (ii) in the primal-dual certificate is satisfied with $\varepsilon = 1 + \gamma$.*

Unfortunately, there exist instances such that constraints (i) do not hold for any value of $\alpha > 0.5$. For example, suppose that the first arrival has a non-zero bid exclusively from resource i . In fact, let $b_{i,1} = 1$. Consequently, the first arrival is always matched to i and we have, $\lambda_1^Y = (1 - g(y_i)) \forall Y \in [0, 1]^n$. Subsequent arrivals have higher bids from many resources other than i such that they are matched to i with very small probability. Thus, $\theta_i \approx \frac{1}{B_i} \int_0^1 g(x) dx$, and for $B_i \rightarrow +\infty$, we have

$$E_Y[\lambda_1^Y] + \theta_i \approx \left(\int_0^1 (1 - g(x)) dx + \frac{1}{B_i} \int_0^1 g(x) dx \right) \rightarrow \int_0^1 (1 - g(x)) dx.$$

Notice that when $g(x) = e^{x-1}$, we have $\int_0^1 (1 - g(x)) dx = g(0) = 1/e < 0.5$.

Now, let OPT denote the offline solution and let OPT_i denote the set of arrivals matched to i in OPT. In contrast to the primal-dual scheme, the LP free scheme imposes the following linear combination of the LP constraints,

$$\sum_{t \in \text{OPT}_i} \lambda_t + B_i \theta_i \geq \alpha \text{OPT}_i,$$

where we summed constraints (i) over all arrivals in OPT_i and scaled θ_i in accordance with the LP free system (the counterpart of constraint (ii) in the LP free system replaces $B_i \theta_i$ with θ_i). Since i is never matched to any arrival after the first one, we have that, $\lambda_t \geq b_{i,t}(1 - g(0)) \forall t \geq 2$. Thus,

$$\sum_{t \in \text{OPT}_i} \lambda_t + \theta_i \geq (1 - 1/e) \left(\sum_{t \in \text{OPT}_i, t \geq 2} b_{i,t} \right) + \int_0^1 g(x) dx \geq (1 - 1/e) \text{OPT}_i.$$

Appendix C: Additional Examples to Demonstrate Main Analytical Challenge

The first example shows that for GPG, even in the small bids regime, the total used budget of a resource i can decrease as y_i decreases and the bid prices of resource i increase. The example also demonstrates that GPG and f -GPG can output a very different matching on the same instance.

Example C.1 Consider an instance with n resources $\{1, \dots, n\}$, with budget $n - 1$ for resources $j \in [n - 1]$ and budget $(n - 1)^{1.98}$ for resource n . We focus on a snippet of this instance by considering arrivals $\{t + 1, \dots, t + 2n - 2\} \subset T$. We execute Algorithm 1 (with $\beta = 1$) on this instance with seed Y_n for resources $j \in [n - 1]$ fixed, and observe the change in output as seed y_n varies. Suppose that exactly 1 unit of budget is available at arrival $t + 1$ for every resource $j \in [n - 1]$ and every value of $y_n \in [0, 1]$. Further, $\forall y_n \in [0, 1]$, suppose that all of resource n 's budget is available at $t + 1$. The bids are as follows.

- (i) Arrival $\tau \in \{t + 1, \dots, t + n - 1\}$ bids 1 for resource $j = \tau - t$, bids 2 for resource n , and bids 0 for all other resources.
- (ii) Arrival $\tau \in \{t + n, \dots, t + 2n - 2\}$ bids $(n - 1)^{0.99}$ for resource $j = \tau - (t + n - 1)$, bids $(n - 1)^{0.98}$ for resource n , and bids 0 for all other resources.

Let $n \rightarrow +\infty$ so that we are in the small bid regime. Now, observe that there exists some $c \in (0, 1)$ such that when $y_n \in (c, 1]$, GPG matches resource n to the last $n - 1$ arrivals and all of n 's budget is utilized. On the other hand, for $y_n \in (0, c)$, GPG matches resource n only to the first n arrivals and almost none of n 's budget is used. Thus, as y_n decreases from 1 to 0 and the bid price of resource n increases, less of n 's budget is used in GPG. This is despite the fact that GPG makes allocation greedily w.r.t. the bid prices.

f -GPG, which is not budget oblivious, does not exhibit the above behavior when started at arrival $t + 1$ with the same initial conditions. Since at $t + 1$ only 1 unit of budget is available for resources $j \in [n - 1]$, for all values of y_n , f -GPG will match at most a tiny fraction of the last n arrivals to resources $j \in [n - 1]$. In other words, for all values of y_n , f -GPG (fractionally) matches resource n to the last n arrivals and uses up almost the entire budget of n . Thus, the matchings output by the two algorithms differ substantially for $y_n \in (0, c)$.

Next, we consider an example that illustrates the challenge with analyzing f -GPG. As we discussed in Section 4.1.1, the main difficulty stems from the fact that the “dominance” property does not hold. Consider a resource i and fix all seed values except y_i . To define the example we use the notion of critical thresholds and the quantity $B(y_i)$ defined in Section 4.2.1. Recall that $B(y_i)$ is the cumulative bid on i from all arrivals matched to i in OPT that have critical threshold at least y_i . The “dominance” property states that the amount of i 's budget used is always at least $B(y_i)$.

Example C.2 Consider a snippet of an instance with 3 resources $\{1, \dots, 3\}$, budget n for resources 1 and 2 and $1.5n$ for resource 3. Consider a sequence of $[2n]$ arrivals. For every arrival $t \in [2n]$, we have $b_{i,t} = 1$ from $i \in \{1, 2\}$. Bids from resource 3 are as follows, $b_{3t} = 1$ for $t \in [n]$ and $b_{3t} = 0.5$ otherwise. Notice that matching all $2n$ arrivals to resource 3 exactly utilizes the budget of the resource.

Let $n \rightarrow +\infty$. We execute f -GPG with $\beta = 1$ on this instance with seeds (y_1, y_2) for resources $\{1, 2\}$ fixed, and observe the change in output as seed y_3 varies. In particular, let $y_1 \in (1/8, 1/4)$ and $y_2 \in (3/4, 1)$.

- (i) For $y_3 = 1$, the first n arrivals are matched to resource 1 and the next n to resource 2.
- (ii) For $y_3 \in (1/4, 0.4)$, on the last n arrivals, the bid price of resource 3 is higher than the bid price of resource 2. Therefore, the first n arrivals are matched to resource 1 and the next n to resource 3.

- (iii) For $y_3 \leq 1/8$, the first n arrivals are matched to resource 3. However, the last n arrivals are now matched to resource 1.

Notice that the critical threshold of every arrival w.r.t. resource 3 is at least $1/8$. Thus,

$$B(y_3) = 1.5n \quad \forall y_3 < 1/8.$$

Despite this, only the first n arrivals are matched to resource 3 when $y_3 \leq 1/8$ and the total budget of resource 3 used in f -GPG does not exceed n for $y_3 < 1/8$.

Appendix D: Proof of Lemma 4

LEMMA 4. *Given $i \in I$ and seed Y_{-i} , for every $y_i \in [0, 1]$, $t \in T$, and $\tau \in [t, t+1]$, we have*

$$\lambda_\tau(y_i) \geq \lambda_\tau(1) \geq b_{i,t}(1 - g(y_{i,\tau}^c)).$$

Proof. Given i and Y_{-i} , consider an arbitrary seed $y_i \in [0, 1]$ and arrival $t \in T$. By definition of $y_{i,\tau}^c$, we have $\lambda_\tau(1) \geq b_{i,t}(1 - g(y_{i,\tau}^c))$. It remains to show that $\lambda_\tau(y_i) \geq \lambda_\tau(1)$. We claim that this follows from,

$$I(\tau, 1) \setminus \{i\} \subseteq I(\tau, y_i).$$

To see this, observe that given the above nesting of sets, we have

$$\lambda_\tau(y_i) = \max_{j \in I(\tau, y_i)} b_{j,t}(1 - g(y_j)) \geq \max_{j \in I(\tau, 1) \setminus \{i\}} b_{j,t}(1 - g(y_j)) = \max_{j \in I(\tau, 1)} b_{j,t}(1 - g(y_j)) = \lambda_\tau(1).$$

We prove that $I(\tau, 1) \setminus \{i\} \subseteq I(\tau, y_i)$ by contradiction. Let y_i be such that $I(\tau, 1) \setminus \{i\} \not\subseteq I(\tau, y_i)$. Let $\tau_1 \leq \tau$ be the earliest moment such that there exists a resource $i_1 \in I(\tau_1, 1) \setminus (I(\tau_1, y_i) \cup \{i\})$, i.e., $i_1 (\neq i)$ is available at τ_1 in f -GPG(1) but unavailable at τ_1 in f -GPG(y_i). This occurs only if i_1 is matched at some moment $\tau_0 < \tau_1$ in f -GPG(y_i) but not matched at τ_0 in f -GPG(1). Now, the following statements are true.

(i) $I(\tau_0, 1) \setminus \{i\} \subseteq I(\tau_0, y_i)$. This follows from the definition of τ_1 and the fact that $\tau_0 < \tau_1$.

(ii) $i_1 \in I(\tau_0, 1) \setminus \{i\}$. Follows from $i_1 \in I(\tau_1, 1) \setminus \{i\}$ and the fact that $I(\tau_1, 1) \subseteq I(\tau_0, 1)$.

Overall, we have $i_1 \in I(\tau_0, 1) \setminus \{i\} \subseteq I(\tau_0, y_i)$. At every arrival, f -GPG picks an available resource with highest bid price. So if τ_0 is matched to i_1 in f -GPG(y_i), then from (i) and (ii) it must also be matched to i_1 in f -GPG(1) (the bid price of i_1 at τ_0 does not change). This contradicts the definition of τ_0 . $\square$

Appendix E: Obstacle to Analyzing GPG Directly

Let $I(t, Y)$ denote the set of resources available at t in GPG. Consider the candidate solution,

$$\lambda_t = E_Y [\lambda_\tau(Y)] \quad \forall t \in T \text{ and } \theta_i = E_Y [\theta_i(Y)] \quad \forall i \in I, \quad (21)$$

where,

$$\lambda_t(Y) = \max_{j \in I(t, Y)} b_{j,t}(1 - g(y_j)) \quad \forall t \in T, Y \in [0, 1]^n, \quad (22)$$

$$\theta_i(Y) = \text{GPG}_i(Y) g(y_i) \quad \forall i \in I, Y \in [0, 1]^n. \quad (23)$$

This definition is a natural generalization of the candidate solution used by Devanur et al. (2013) to analyze GPG for vertex weighted online bipartite matching.

On every sample path Y , $\lambda_t(Y) > 0 \quad \forall t \in T$ and $\theta_i(Y) > 0 \quad \forall i \in I$. Thus, constraints (3) are satisfied. We claim (without proof) that Lemmas 2 and 4 from the analysis of f -GPG, can be generalized for GPG as stated below.

LEMMA 9. For the candidate solution given by (6), we have,

$$\sum_{i \in I} \theta_i + \sum_{t \in T} \lambda_t \leq \sum_{i \in I} \left(1 + \max_{t \in T} \frac{b_{i,t}}{B_i} \right) \text{GPG}_i,$$

and constraint (1) is satisfied with $\varepsilon = 1 + \gamma$.

LEMMA 10. Given $i \in I$ and seed Y_{-i} , for every $y_i \in [0, 1], t \in T$, we have

$$\lambda_t(y_i) \geq \lambda_t(1) \geq b_{i,t}(1 - g(y_i^c(t))).$$

Unfortunately, Lemma 5 need not hold for GPG and this is the main reason that we first analyze f -GPG and then link the performance of GPG with f -GPG. In the proof of Lemma 5, the main step that fails is the one that links $\theta_i(y_i)$ to the difference $\lambda_{net}(y_i) - \lambda_{net}(1)$. In GPG, we have that for every seed value $y \in [0, 1]$ of resource i ,

$$\lambda_{net}(y) \geq \text{GPG}_i(y)(1 - g(y)) + \sum_{j \in I \setminus \{i\}} \text{GPG}_j(y)(1 - g(y_j)).$$

Notice that this is an inequality whereas in case of f -GPG the left hand side equals the right hand side. In GPG, the value of $\lambda_{net}(y)$ may be strictly larger than the right hand side because GPG may overestimate the remaining budget of resources. In particular, when i is matched to arrival t , the remaining budget of i may be smaller than the bid $b_{i,t}$. However, we set $\lambda_t(y) = b_{i,t}(1 - g(y))$, which may be larger than the reward from matching i to t . Now, consider the difference

$$\lambda_{net}(y_i) - \lambda_{net}(1) = (1 - g(y_i)) \sum_{t \in f\text{-GPG}_i(y_i)} b_{i,t} + \sum_{j \in I \setminus \{i\}} \left[\sum_{t \in f\text{-GPG}_j(y_i)} b_{j,t} - \sum_{t \in f\text{-GPG}_j(1)} b_{j,t} \right] (1 - g(y_j)),$$

where we used the fact that $g(1) = 1$ for every value of $\beta > 0$. In the analysis of f -GPG, we show that the second term is non-positive. However, in GPG the second term may be strictly positive. In particular, it is possible that for $j \in I \setminus \{i\}$,

$$\text{GPG}_j(1) = \sum_{t \in f\text{-GPG}_j(1)} b_{j,t} = B_j$$

but

$$\sum_{t \in f\text{-GPG}_j(y_i)} b_{j,t} > B_j.$$

In general, the increase in $\lambda_{net}(y_i) - \lambda_{net}(1)$ is not due to $\text{GPG}_i(y_i)$ alone and Lemma 5 does not hold for GPG.

Appendix F: Proof of Theorem 4

THEOREM 4. For every bid-to-budget ratio $\gamma \in (0, 1]$ and $\beta \geq 1$, the competitive ratio of GPG is at most 0.624.

Proof. Consider a family of instances parameterized by $\epsilon \geq 0$, integer budget to bid ratio $m = \frac{1}{\gamma} \geq 1$, and $\beta \geq 1$. Each instance in the family has $n + 1$ resources. There are $2mn$ arrivals and each arrival bids $\frac{1}{m}$ on every resource in $\{1, 2, \dots, n\}$. Arrival $t \in [mn]$ bids

$$w(t, \epsilon) = \frac{1 - e^{\beta(\frac{t}{n+1} - 1)}}{m(1 - e^{-\beta(1 - \epsilon)})}$$

on resource $n + 1$. Arrivals $t \geq mn + 1$ do not have an edge to resource $n + 1$. Resources $\{1, 2, \dots, n\}$ have unit budget and resource $n + 1$ has a budget of $\sum_{t \in [mn]} w(t, \epsilon)$. The optimal matching allocates resource $n + 1$ to every $t \in [mn]$ and matches arrival t to resource $\lceil \frac{t}{m} \rceil - n$ for $t \geq mn + 1$. We have,

$$\text{OPT} = n + \sum_{t \in [mn]} w(t, \epsilon),$$

here $\lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{t \in [mn]} w(t, 0) \rightarrow \frac{e^{\beta(1-\beta^{-1})+\beta^{-1}}}{e^{\beta}-1}$. Therefore, for $\epsilon = 0$, we have, $\lim_{n \rightarrow +\infty} \frac{1}{n} \text{OPT} \rightarrow \frac{e^{\beta(2-\beta^{-1})-1+\beta^{-1}}}{e^{\beta}-1}$. Let

$$H(\beta) = \frac{e^{\beta(2-\beta^{-1})-1+\beta^{-1}}}{e^{\beta}-1}.$$

Observe that $H(1) = \frac{e}{e-1}$ and $H(\beta) > \frac{e}{e-1}$ for $\beta \geq 1$.

To understand the high level idea, consider $\epsilon = 0$ and $\beta = 1$. Let y_i denote the random seed of resource $i \in [n + 1]$ in GPG. For $k \in [n]$, let z_k denote the k -th order statistic of the random variables $\{y_i\}_{i \in [n]}$, i.e., the k -th smallest seed value. For $k \in [n]$, random variable z_k has mean $\frac{k}{n+1}$. Let $I(t)$ denote the (random) set of resources available at t in GPG. Resource $n + 1$ is matched to arrival $t \in [nm]$ if

$$w(t, 0)(1 - e^{y_{n+1}-1}) > \max_{i \in I(t)} \frac{1}{m} (1 - e^{y_i-1}) \geq \frac{1}{m} (1 - g(z_t)),$$

here the last inequality follows from the fact that the t -th lowest seeded resource in $[n]$ will not be matched to any of the arrivals preceding t . For the sake of intuition, suppose that for $n \rightarrow +\infty$, we have $z_{\lceil \frac{t}{m} \rceil} \approx E[z_{\lceil \frac{t}{m} \rceil}] = \frac{\lceil \frac{t}{m} \rceil}{n+1} \forall t \in [mn]$, with high probability (w.h.p.). Then, resource $n + 1$ is matched to t only if $y_{n+1} \approx 0$, which occurs with a vanishingly small probability. By linearity of expectation, the expected total consumed capacity of resource $n + 1$ is ≈ 0 . Therefore, the expected value of $\frac{1}{n} \text{GPG}$ is ≈ 1 , and we have $\frac{\text{GPG}}{\text{OPT}} \approx (1 - 1/e)$.

For the rest of this proof, we are interested in the asymptotic case where $n \rightarrow +\infty$. We use $O(\cdot)$ notation to hide constants independent of n . Now, observe that for a fixed t , $w(t, \epsilon)$ is a strictly increasing function of ϵ . When we increase ϵ from 0, the value of offline (OPT) increases by

$$\sum_{t \in [mn]} (w(t, \epsilon) - w(t, 0)) = \sum_{t \in [mn]} w(t, 0) \frac{e^{-\beta}(e^{\epsilon\beta} - 1)}{1 - e^{\beta(-1+\epsilon)}} = n(H(\beta) - 1) \frac{e^{-\beta}(e^{\epsilon\beta} - 1)}{1 - e^{\beta(-1+\epsilon)}} = O(\epsilon n). \quad (24)$$

We show that the expected total value of GPG increases by an amount that is significantly lower than $O(\epsilon n)$. Let $z_0 = 0$ and $z_{n+1} = 1$. The minimum of all ranks (first order statistic), z_1 , is a random variable with mean $\frac{1}{n+1}$ and variance $\frac{n}{(n+1)^2(n+2)} = \frac{O(1)}{n^2}$ (Arnold et al. 2008). Using Chebyshev's inequality,

$$\mathbb{P}\left(\left|z_1 - z_0 - \frac{1}{n+1}\right| \geq \frac{1}{n^{4/3}}\right) \leq \frac{O(1)}{n^{4/3}}.$$

In general, for every $k \in \{0, 1, \dots, n\}$, the random variable $z_{k+1} - z_k$ has mean $\frac{1}{n+1}$ and variance $\frac{n}{(n+1)^2(n+2)} = \frac{O(1)}{n^2}$ (Arnold et al. 2008). Again, using Chebyshev's inequality,

$$\mathbb{P}\left(\left|z_{k+1} - z_k - \frac{1}{n+1}\right| \geq \frac{1}{n^{4/3}}\right) \leq \frac{O(1)}{n^{4/3}} \quad \forall k \in [n] \cup \{0\}.$$

Applying the union bound, we have,

$$\mathbb{P}\left(\bigcup_{k \in [n] \cup \{0\}} \left\{\left|z_{k+1} - z_k - \frac{1}{n+1}\right| \geq \frac{1}{n^{4/3}}\right\}\right) \leq \frac{O(1)}{n^{1/3}}.$$

We have that for large n , the following event occurs w.h.p.,

$$\left| z_{k+1} - z_k - \frac{1}{n+1} \right| < \frac{1}{n^{4/3}} \quad \forall k \in [n] \cup \{0\}. \quad (25)$$

In the rest of this proof, we condition on this high probability event. Consequently, we have

$$\left| z_{\lceil \frac{t}{m} \rceil} - \frac{\lceil \frac{t}{m} \rceil}{n+1} \right| \leq \frac{\lceil \frac{t}{m} \rceil}{n^{4/3}} \quad \forall t \in [mn].$$

We will show that in expectation, GPG uses a very small fraction of the budget of resource $n+1$. For any given $t \leq mn - mn^{3/4}$, resource $n+1$ is matched to arrival t if,

$$\begin{aligned} mw(t, \epsilon)(1 - g(y_{n+1})) &> 1 - g(z_{\lceil \frac{t}{m} \rceil}), \\ \frac{1 - e^{\beta \left(\frac{\lceil \frac{t}{m} \rceil}{n+1} - 1 \right)}}{1 - e^{\beta(-1+\epsilon)}} (1 - e^{\beta(y_{n+1}-1)}) &> 1 - e^{\beta \left(\frac{\lceil \frac{t}{m} \rceil}{n+1} + \frac{t}{n^{4/3}} - 1 \right)}, \\ \Rightarrow y_{n+1} &< \epsilon + \frac{O(1)}{n^{1/3}}. \end{aligned}$$

For large enough n , resource $n+1$ is not matched to any arrival in GPG when $y_{n+1} > \epsilon + \frac{1}{n^{1/4}}$. In the following discussion, we condition on the event,

$$y_{n+1} \leq \epsilon + \frac{1}{n^{1/4}}, \quad (26)$$

which occurs w.p. $\epsilon + \frac{1}{n^{1/4}}$ (independent of event (25)). Finally, we claim that resource $n+1$ is matched to at most $\epsilon m(n+1) + mn^{3/4}$, i.e., $O(\epsilon mn)$, of the first mn arrivals. To prove this, it suffices to show that for every $t \geq m(1 + \epsilon(n+1) + n^{3/4})$, resource $\tau(t) = \lceil \frac{t}{m} \rceil - n^{3/4} - \epsilon(n+1)$ is matched to one of the first mn arrivals w.p. 1. Consider an arbitrary $t \geq m(1 + \epsilon(n+1) + n^{3/4})$. Resource $\tau(t)$ is not matched to any arrival prior to t if,

$$\begin{aligned} \frac{1 - e^{\beta \left(\frac{\lceil \frac{t}{m} \rceil}{n+1} - 1 \right)}}{1 - e^{\beta(-1+\epsilon)}} (1 - e^{\beta(y_{n+1}-1)}) &> 1 - e^{\beta \left(\frac{\tau(t)}{n+1} + \frac{\lceil \frac{t}{m} \rceil}{n^{4/3}} - 1 \right)}, \\ &> 1 - e^{\beta \left(\frac{\lceil \frac{t}{m} \rceil}{n+1} - 1 - \epsilon \right)}, \end{aligned}$$

which does not occur for any value of y_{n+1} . Overall, conditioned on events (25) and (26), the best that GPG can do is to match the first $O(\epsilon mn)$ arrivals to resource $n+1$, using up $O(\epsilon n)$ of the resource's budget. More precisely, w.p. 1, GPG uses at most,

$$n \frac{\int_0^\epsilon (1 - e^{\beta(u-1)}) du}{1 - e^{\beta(-1+\epsilon)}} + o(n),$$

of resource $n+1$'s budget. Unconditioning on events (25) and (26), the expected budget of resource $n+1$ used in GPG is $O(\epsilon^2 n)$. Thus,

$$\frac{\text{GPG}}{n} \leq 1 + O(\epsilon^2) + o(1).$$

From (24), we have

$$\frac{\text{OPT}}{n} \geq H(\beta) + O(\epsilon) + o(1).$$

Combining these inequalities, for $n \rightarrow +\infty$, we have $\frac{\text{GPG}}{\text{OPT}} \leq \frac{1}{H(\beta)} < (1 - 1/e)$ for some $\epsilon > 0$ and any $\beta \geq 1$. As m is arbitrary, the upper bound holds for small bids as well. We can obtain a stronger upper bound by considering a value of ϵ that minimizes the ratio in the limit $n \rightarrow +\infty$, i.e.,

$$\min_{\epsilon \geq 0} \lim_{n \rightarrow +\infty} \frac{\text{GPG}}{\text{OPT}} \leq \min_{\epsilon \geq 0} \frac{1 + \epsilon \frac{\int_0^\epsilon (1 - e^{\beta(u-1)}) du}{1 - e^{\beta(-1+\epsilon)}}}{H(\beta) + (H(\beta) - 1) \frac{e^{-\beta}(\epsilon\beta - 1)}{(1 - e^{\beta(-1+\epsilon)})}} < 0.624 \quad \forall \beta \geq 1,$$

where the final inequality follows by setting $\epsilon \approx 0.1336$.

□

Appendix G: Proof of Theorem 5

To prove Theorem 5, we first need the following intermediate result for f -GPG.

LEMMA 11. *Given $i \in I$, seed Y_{-i} in f -GPG, and decomposable bids $b_{i,t} \in \{0, b_i b_t\} \forall i \in I, t \in T$, we have,*

$$E_{y_i} [\min\{\theta_i(y_i), g(y_i) \text{OPT}_i\} | Y_{-i}] + \sum_{t \in \text{OPT}_i} E_{y_i} [\lambda_t(y_i) | Y_{-i}] \geq \sum_{v \in V} b(v) \left[1 - g(v) + \int_0^v g(x) dx \right].$$

Proof. It suffices to show that,

$$\theta_i(y_i) \geq B(y_i) g(y_i) \quad \forall y_i \in [0, 1]. \quad (27)$$

Given this inequality, using the fact that $\text{OPT}_i \geq B(y_i) \forall y_i \in [0, 1]$ and taking expectation over $y_i \sim U[0, 1]$, we have,

$$E_{y_i} [\min\{\theta_i(y_i), g(y_i) \text{OPT}_i\} | Y_{-i}] \geq \sum_{v \in V} \left(b(v) \int_0^v g(x) dx \right).$$

Combining the inequality above with the lower bound in Lemma 4 gives us the desired.

Notice that, if $b_i = 0$, then (27) is trivially true. Further, all arrivals $t \in T$ where $b_t = 0$ can be ignored. W.l.o.g., let $b_i > 0$ and $b_t > 0 \forall t \in T$. Now, it suffices to show that

$$f\text{-GPG}_i(y_i) \geq B(y_i) \forall y_i \in [0, 1].$$

Given these inequalities, (27) follows by definition of $\theta_i(y_i)$ (see (5)). For the sake of contradiction, consider a value $y_i = y_0 \in [0, 1)$ such that $f\text{-GPG}_i(y_0) < B(y_0) \leq B_i$. Then, there exists an arrival $\tau \in \text{OPT}_i$ such that, $y_i^c(\tau) \geq y_0$, but τ is not matched to i in $f\text{-GPG}(y_0)$. Since, $f\text{-GPG}(y_0) < B_i$, resource i is available at τ . Therefore, in $f\text{-GPG}(y_0)$, τ is matched to a resource i_1 such that,

$$b_{i_1, \tau}(1 - g(y_{i_1})) \geq b_{i, \tau}(1 - g(y_0)) \geq b_{i, \tau}(1 - g(y_i^c(\tau))),$$

here the first inequality is strict w.p. 1 (ties occur w.p. 0). Using the decomposability of bids, we have (w.p. 1),

$$b_{i_1}(1 - g(y_{i_1})) > b_i(1 - g(y_0)).$$

Therefore, i_1 is preferred over i at all arrivals. We say that i_1 is *better* than i , or $i_1 \succ i$ (in $f\text{-GPG}(y_0)$). Now, notice that i_1 must be unavailable at τ in $f\text{-GPG}(1)$. Therefore, there exists an arrival τ_{i_1} prior to τ such that, τ_{i_1} is matched to i_1 in $f\text{-GPG}(1)$ but in $f\text{-GPG}(y_0)$, τ_{i_1} it is matched to a resource i_2 that is better than i_1 . Repeating this argument a number of times, we get a sequence of arrivals $\tau > \tau_{i_1} > \dots > \tau_{i_k}$ and resources $i \prec i_1 \prec \dots \prec i_k$. The number of resources is finite, so the sequence must terminate and, w.l.o.g., $i_k = i_\ell$ for some $\ell < k \leq n$, contradiction.

□

THEOREM 5. *GPG with $\beta = 1$ is $\frac{1}{1+\gamma}(1-1/e)$ competitive for OBA when the bids are decomposable, i.e., $b_{i,t} \in \{0, b_i \times b_t\} \forall i \in I, t \in T$. For the special case of OBA with integer starting capacities and binary bids, GPG is $(1-1/e)$ competitive for arbitrary bid-to-budget ratio γ .*

Proof. For f -GPG, when $g(x) = e^{x-1}$ we have $\forall i \in I, Y_{-i} \in [0, 1]^{n-1}$,

$$\begin{aligned} \sum_{t \in \text{OPT}_i} E_{y_i}[\lambda_t(y_i) | Y_{-i}] + E_{y_i}[\theta_i(y_i) | Y_{-i}] &\geq \sum_{v \in V} b(v) (1 - e^{-v} + e^{-v} - e^{-1}), \\ &= (1 - e^{-1}) \sum_{v \in V} b(v), \\ &= (1 - e^{-1}) \text{OPT}_i. \end{aligned}$$

Now, combining this with Lemma 2 we have that f -GPG with $\beta = 1$ is $(1-1/e)$ competitive for decomposable bids. Using Theorem 2, we have that GPG is $\frac{1}{1+\gamma}(1-1/e)$ competitive for decomposable bids. The case of binary bids is a special case of decomposable bids where GPG and f -GPG are identical. Thus, GPG with $\beta = 1$ is unconditionally $(1-1/e)$ competitive for binary bids. $\square$

Appendix H: Missing Proofs for Online Matching with Multi-Channel Traffic

Candidate solution for FRP: Recall the candidate solution for FRP. For resource $i \in I$, let o_i^{Ex} be the total budget of i allocated to external arrivals in OPT. Let x_i^{Ex} be the expected budget of i allocated to external arrivals in f -GPG. Let o_i^{In} and x_i^{In} be the (expected) budget of i allocated to internal arrivals in OPT and in f -GPG, respectively. Finally, let

$$\zeta = \frac{\sum_{i \in I} (x_i^{\text{In}} + x_i^{\text{Ex}})}{\sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}})}.$$

Observe that ζ is equal to the performance ratio $\frac{\text{GPG}}{\text{OPT}}$. We start by noting that for $\delta > 0$, all constraints (including (19)) are well defined. In particular, for $\delta > 0$, using Lemma 7(i) we have $\sum_{i \in I} o_i^{\text{Ex}} > 0$ on every non-trivial instance of the problem. By using the structure of the instance and borrowing ideas from Manshadi et al. (2022), it can be shown that the candidate solution satisfies constraints (15)–(18). The crux of our analysis is to show that inequality (19) is satisfied and we show this in the next section. The rest of the proof of Theorem 6 is included in Appendix H.2.

H.1. Proof of Inequality (19)

We start with some notation. Recall that T^{In} and T^{Ex} denote the set of internal arrivals and external arrivals, respectively. On sample path Y in f -GPG, let $x_i^{\text{Ex}}(Y)$ and $x_i^{\text{In}}(Y)$ denote the budget of i allocated to external and internal arrivals, respectively. Also, on sample path Y in GPG, let $I(t, Y)$ denote the set of neighboring resources available at t . GPG may match more than o_i^{In} internal arrivals to i on some sample paths. To capture these “excess” matches we define,

$$xs_i(Y) = \max\{x_i^{\text{In}}(Y) - o_i^{\text{In}}, 0\} \quad \forall i \in I.$$

Excess matches play a crucial role in proving (19). In fact, using the LP free framework, we first establish a lower bound on $\zeta \sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}})$ in terms of the expectation $E_Y[\sum_{i \in I} xs_i(Y) g(y_i)]$. Given a candidate

solution to the FRP, the random variables $\{xs_i(Y)\}_{i \in I}$ and their expected values are not uniquely defined. We address this by selecting the “worst” possible way to define these random variables for any given candidate solution. To obtain a lower bound on $E_Y[\sum_{i \in I} xs_i(Y)g(y_i)]$ we now state two useful lemmas. The first lemma gives a lower bound on $xs_i(Y)$ in terms of $x_i^{\text{Ex}}(Y)$, which also a random variable that is not uniquely defined (given a candidate solution). The second lemma gives a result for real valued functions that is reminiscent of the calculus of variations.

LEMMA 12. *On every sample path Y , $xs_i(Y) \geq o_i^{\text{Ex}} - x_i^{\text{Ex}}(Y) \quad \forall i \in I$.*

Proof. Suppose that every external arrival with an edge to i is matched in $f\text{-GPG}(Y)$, then from Lemma 7(i) we have, $x_i^{\text{Ex}}(Y) = o_i^{\text{Ex}}$, and the claim follows. If an external arrival with edge to i is unmatched in $f\text{-GPG}(Y)$, then all of i 's budget is consumed in $f\text{-GPG}(Y)$, i.e., $x_i^{\text{In}}(Y) + x_i^{\text{Ex}}(Y) = B_i$. Using $B_i \geq o_i^{\text{In}} + o_i^{\text{Ex}}$ and the fact that $xs_i(Y) \geq x_i^{\text{In}}(Y) - o_i^{\text{In}}$, completes the proof. $\square$

LEMMA 13. *Given $g(x) = e^{x-1}$, $G(x) = \int_0^x g(u)du$, real values $h_0, H_0 > 0$ and an integrable function $h : [0, 1] \rightarrow [0, H_0]$ with $\int_0^1 h(u)du = h_0$, we have*

$$\int_0^1 h(u)g(u)du \leq H_0 \left(G(1) - G\left(1 - \frac{h_0}{H_0}\right) \right).$$

Proof. Consider the function

$$f(u) = \begin{cases} 0 & u \in [0, 1 - \frac{h_0}{H_0}), \\ H_0 & u \in [1 - \frac{h_0}{H_0}, 1]. \end{cases}$$

Observe that $\int_0^1 f(u)du = h_0$ and $\int_0^1 f(u)g(u)du = H_0 \left(G(1) - G(1 - \frac{h_0}{H_0}) \right)$. It suffices to show that, $\int_0^1 h(u)g(u)du \leq \int_0^1 f(u)g(u)du$. Let $\mathcal{H}$ denote the set of all integrable functions $h : [0, 1] \rightarrow [0, H_0]$ with $\int_0^1 h(u)du = h_0$. Let

$$h^* = \sup_{h \in \mathcal{H}} \int_0^1 h(u)g(u)du.$$

Now, if $\int_0^1 f(u)g(u)du < \int_0^1 h^*(u)g(u)du$, then there exists a closed interval $I_1 \subseteq [0, 1 - \frac{h_0}{H_0})$ with non-zero measure and a closed interval $I_2 \subseteq [1 - \frac{h_0}{H_0}, 1]$ with non-zero measure such that,

$$h^*(u) > 0 \quad \forall u \in I_1 \quad \text{and} \quad h^*(u) < H_0 \quad \forall u \in I_2.$$

Let $\epsilon = \min \{ \min_{u \in I_1} h^*(u), \min_{u \in I_2} H_0 - h^*(u) \}$. Let $|I|$ denote the length/measure of an interval I . Consider function h' such that $h'(u) = h^*(u) - \epsilon|I_1|^{-1} \quad \forall u \in I_1$, $h'(u) = h^*(u) + \epsilon|I_2|^{-1} \quad \forall u \in I_2$ and $h'(u) = h^*(u)$ for all other values of $u \in [0, 1]$. We have, $\int_0^1 h'(u)du = h_0$ and $\int_0^1 h'(u)g(u)du > \int_0^1 h^*(u)g(u)du$, contradiction. $\square$

Next, we state two intermediate inequalities that we use to establish inequality (19).

LEMMA 14. *On every problem instance under consideration, the candidate solution to FRP satisfies,*

$$(i) \quad \zeta \sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}}) \geq \sum_{i \in I} (G(1) o_i^{\text{In}} + x_i^{\text{Ex}} + E_Y[xs_i(Y)g(y_i)]).$$

$$(ii) \quad \sum_{i \in I} E_Y[xs_i(Y)g(y_i)] \geq \sum_{i \in I} o_i^{\text{Ex}} G\left(1 - \frac{\sum_{j \in I} x_j^{\text{Ex}}}{\sum_{j \in I} o_j^{\text{Ex}}}\right).$$

Before we prove Lemma 14, observe that a straightforward combination of inequalities (i) and (ii) in the lemma proves that the candidate solution satisfies (19).

Proof of Lemma 14. The proof of part (i) is based on the LP free framework. In particular, we find a feasible solution to the following system of inequalities,

$$\sum_{t \in T} \lambda_t + \sum_{i \in I} \theta_i \leq \zeta \sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}}) \quad (28)$$

$$\theta_i + \sum_{t \in \text{OPT}_i} \lambda_t \geq G(1)o_i^{\text{In}} + x_i^{\text{Ex}} + E_Y[xs_i(Y)g(y_i)] \quad \forall i \in I, \quad (29)$$

$$\lambda_t \geq 0, \quad \theta_i \geq 0 \quad \forall t \in T, i \in I. \quad (30)$$

Given a feasible solution the system, summing up inequalities (29) over all $i \in I$ gives us the desired. Our candidate feasible solution is based on the random variables,

$$\theta_i^{\text{In}}(Y) = x_i^{\text{In}}(Y)g(y_i) \quad \forall i \in I, \quad (31)$$

$$\theta_i^{\text{Ex}}(Y) = x_i^{\text{Ex}}(Y) \quad \forall i \in I, \quad (32)$$

$$\lambda_t(Y) = \begin{cases} \max_{j \in I(t, Y)} (1 - g(y_j)) & t \in T^{\text{In}}, \\ 0 & t \in T^{\text{Ex}}. \end{cases} \quad (33)$$

We note that random variables $\{\lambda_t(Y)\}_{t \in T^{\text{In}}}$ and $\{\theta_i^{\text{In}}(Y)\}_{i \in I}$ are the same as their counterparts ((4) and (5)) in Section 4.2. For external arrivals, the fact that there is exactly one edge incident on each arrival eliminates the need to set a non-zero value for λ_t and we set $\lambda_t(Y) = 0 \forall t \in T^{\text{Ex}}$. Consider the (non-negative) candidate solution,

$$\lambda_t = E_Y[\lambda_t(Y)] \quad \forall t \in T, \quad \theta_i = E_Y[\theta_i^{\text{In}}(Y) + \theta_i^{\text{Ex}}(Y)] \quad \forall i \in I.$$

Observe that,

$$E_Y \left[\sum_{i \in I} (\theta_i^{\text{In}}(Y) + \theta_i^{\text{Ex}}(Y)) + \sum_{t \in T} \lambda_t(Y) \right] = \sum_{i \in I} (x_i^{\text{In}} + x_i^{\text{Ex}}),$$

Thus, our candidate solution satisfies inequality (28). It remains to show that inequality (29) holds for all $i \in I$. We prove this via the same high level approach as the analysis of f -GPG in Section 4. Fix an arbitrary resource $i \in I$ and seeds Y_{-i} for all resources except i . To prove (29), it suffices to show that,

$$E_{y_i} \left[\theta_i^{\text{In}}(Y) + \theta_i^{\text{Ex}}(Y) + \sum_{t \in \text{OPT}_i} \lambda_t(Y) \mid Y_{-i} \right] \geq G(1)o_i^{\text{In}} + x_i^{\text{Ex}} + E_{y_i}[xs_i(Y)g(y_i) \mid Y_{-i}]. \quad (34)$$

To show inequality (34), we first consider a “truncated” instance of the problem without the external arrivals, i.e., the arrival sequence is T^{In} and all budgets and bids are unchanged. Let OPT^{In} denote the optimal offline solution of the truncated instance. W.l.o.g., OPT^{In} is identical to the internal traffic portion of OPT . This follows from the fact that OPT matches every arrival in the original instance (Lemma 7(i)) and OPT^{In} can do no better than match all arrivals. For every $i \in I$, exactly o_i^{In} of resource i ’s budget is used in OPT^{In} and the set of arrivals matched to i is given by $\text{OPT}_i^{\text{In}} = T^{\text{In}} \cap \text{OPT}_i$. Now, using Lemma 11 for decomposable bids (see Appendix G), we have

$$E_{y_i} \left[\min\{x_i^{\text{In}}(Y), o_i^{\text{In}}\}g(y_i) + \sum_{t \in \text{OPT}_i^{\text{In}}} \lambda_t(Y) \mid Y_{-i} \right] \geq G(1)o_i^{\text{In}}, \quad (35)$$

where, $x_i^{\text{In}}(Y)g(y_i) = \theta_i^{\text{In}}(Y)$, by definition (see (31)). Now, we claim that,

$$\theta_i^{\text{In}}(y_i) + \theta_i^{\text{Ex}}(y_i) = (\min\{x_i^{\text{In}}(y_i), o_i^{\text{In}}\} + xs_i(y_i))g(y_i) + x_i^{\text{Ex}}(y_i) \quad \forall y_i \in [0, 1].$$

Combining this equality with (35) gives us (34) and completes the proof of part (i) of the lemma. To show the equality above, fix an arbitrary $y_i \in [0, 1]$ and consider the following cases.

Case I: $x_i^{\text{In}}(y_i) \geq o_i^{\text{In}}$. We have, $o_i^{\text{In}} + xs_i(y_i) = x_i^{\text{In}}(y_i)$. From (31) and (32),

$$\theta_i^{\text{In}}(y_i) + \theta_i^{\text{Ex}}(y_i) = (o_i^{\text{In}} + xs_i(y_i))g(y_i) + x_i^{\text{Ex}}(y_i),$$

as desired.

Case II: $x_i^{\text{In}}(y_i) < o_i^{\text{In}} (\leq B_i)$, i.e., resource i is available at every internal arrival in $f\text{-GPG}(y_i)$. In this case, $xs_i(y_i) = 0$ and

$$\theta_i^{\text{In}}(y_i) + \theta_i^{\text{Ex}}(y_i) = (x_i^{\text{In}}(y_i) + xs_i(y_i))g(y_i) + x_i^{\text{Ex}}(y_i),$$

as desired.

Part (ii) of the lemma gives a lower bound on $\sum_{i \in I} E[g(y_i) xs_i(y_i)]$ in terms of x_i^{Ex} and o_i^{Ex} . Since $xs_i(y_i) < o_i^{\text{Ex}}$, we can ignore any resource $i \in I$ for which $o_i^{\text{Ex}} = 0$ (for this part). Let $o_i^{\text{Ex}} > 0 \forall i \in I$ and recall that $G(x) = \int_0^x g(u)du = e^{x-1} - e^{-1}$. Now, to prove (ii) we claim that it suffices to show,

$$E_{y_i}[xs_i(y_i)g(y_i)] \geq G\left(1 - \frac{x_i^{\text{Ex}}}{o_i^{\text{Ex}}}\right) o_i^{\text{Ex}} \quad \forall i \in I, Y_{-i} \in [0, 1]^{n-1}. \quad (36)$$

First, let us show why (36) is sufficient.

$$\begin{aligned} \sum_{i \in I} G\left(1 - \frac{x_i^{\text{Ex}}}{o_i^{\text{Ex}}}\right) o_i^{\text{Ex}} &= \left(\sum_{i \in I} o_i^{\text{Ex}}\right) \left[\sum_{i \in I} G\left(1 - \frac{x_i^{\text{Ex}}}{o_i^{\text{Ex}}}\right) \frac{o_i^{\text{Ex}}}{\sum_{j \in I} o_j^{\text{Ex}}}\right], \\ &\geq \left(\sum_{i \in I} o_i^{\text{Ex}}\right) \left[G\left(1 - \frac{\sum_{j \in I} x_j^{\text{Ex}}}{\sum_{j \in I} o_j^{\text{Ex}}}\right)\right], \end{aligned} \quad (37)$$

here inequality (37) follows by using the fact that $G(x)$ is a convex function of x and applying Jensen's inequality. It remains to show (36). First, by a direct application of Lemma 12,

$$E_{y_i}[xs_i(y_i)g(y_i)] \geq E_{y_i}[(o_i^{\text{Ex}} - x_i^{\text{Ex}}(y_i))g(y_i)] = G(1) o_i^{\text{Ex}} - \int_0^1 x_i^{\text{Ex}}(u)g(u)du.$$

Using the fact that there are exactly o_i^{Ex} external arrivals with an edge to i (Lemma 7(i)), we have, $x_i^{\text{Ex}}(y_i) \leq o_i^{\text{Ex}}$. Further, $x_i^{\text{Ex}}(u)$ is an integrable function of u with $\int_0^1 x_i^{\text{Ex}}(u)du = x_i^{\text{Ex}}$. Applying Lemma 13 with $h(u) = x_i^{\text{Ex}}(u)$, $h_0 = x_i^{\text{Ex}}$, and $H_0 = o_i^{\text{Ex}}$, we have,

$$\int_0^1 x_i^{\text{Ex}}(u)g(u)du \leq o_i^{\text{Ex}} \left(G(1) - G\left(1 - \frac{x_i^{\text{Ex}}}{o_i^{\text{Ex}}}\right)\right)$$

Therefore,

$$E_{y_i}[xs_i(y_i)g(y_i)] \geq G(1) o_i^{\text{Ex}} - o_i^{\text{Ex}} \left(G(1) - G\left(1 - \frac{x_i^{\text{Ex}}}{o_i^{\text{Ex}}}\right)\right) \geq G\left(1 - \frac{x_i^{\text{Ex}}}{o_i^{\text{Ex}}}\right) o_i^{\text{Ex}}.$$

□

H.2. Proof of Theorem 6

THEOREM 6. *For any given $\delta \in (0, 1]$, the optimal value of FRP is a lower bound on the competitive ratio of f -GPG. For $\delta \geq \frac{1}{e}$, the optimal value of FRP is at least $1 + \delta \ln \delta$.*

Proof. Given an instance with fraction of external traffic δ , we show that there exists a feasible solution to FRP with objective value equal to the performance ratio $\frac{\text{GPG}}{\text{OPT}}$ on that instance. Consider the following candidate solution for the FRP. For resource $i \in I$, let o_i^{Ex} be the total budget of i allocated to external arrivals in OPT. Let x_i^{Ex} be the expected budget of i allocated to external arrivals in f -GPG. Let o_i^{In} and x_i^{In} be the (expected) budget of i allocated to internal arrivals in OPT and in f -GPG, respectively. Finally, let

$$\zeta = \frac{\sum_{i \in I} (x_i^{\text{In}} + x_i^{\text{Ex}})}{\sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}})}.$$

Observe that ζ is equal to the performance ratio $\frac{\text{GPG}}{\text{OPT}}$. We start by noting that for $\delta > 0$, all constraints (including (19)) are well defined. In particular, for $\delta > 0$, using Lemma 7(i) we have $\sum_{i \in I} o_i^{\text{Ex}} > 0$ on every non-trivial instance of the problem. Now, constraints (15) follow from the fact that OPT cannot match a resource more times than the resource's initial budget. Since OPT matches every arrival (Lemma 7(i)), no algorithm can match more than o_i^{Ex} external arrivals to resource $i \in I$. Thus, the candidate solution satisfies (16). Constraint (17) follows directly from the definition of δ . Constraint (18) is inspired by a constraint in the factor revealing program of Manshadi et al. (2022). To show that the candidate solution satisfies (18), we start by giving a lower bound on the expected number of internal arrivals matched by GPG. Using the fact that OPT matches all arrivals we have,

$$\sum_{i \in I} x_i^{\text{In}} \leq \sum_{i \in I} o_i^{\text{In}} \leq (1 - \delta) \sum_{i \in I} B_i.$$

Applying the inequalities above in the definition of ζ , we get,

$$\begin{aligned} \sum_{i \in I} x_i^{\text{Ex}} &= \zeta \sum_i (o_i^{\text{In}} + o_i^{\text{Ex}}) - \sum_{i \in I} x_i^{\text{In}}, \\ &\geq \zeta \delta \sum_{i \in I} B_i - (1 - \zeta) \sum_{i \in I} o_i^{\text{In}}, \\ &\geq (\zeta \delta - (1 - \zeta)(1 - \delta)) \sum_{i \in I} B_i, \\ &= (\zeta - 1 + \delta) \sum_{i \in I} B_i. \end{aligned}$$

Recall that we showed constraint (19) is satisfied by the candidate solution in Section H.1. It remains to show that the optimal value of FRP is at least $1 + \delta \ln \delta$ for $\delta \geq 1/e$.

Let $W^{\text{Ex}} = \sum_{i \in I} o_i^{\text{Ex}}$ and $X^{\text{Ex}} = \sum_{i \in I} x_i^{\text{Ex}}$. From (18), we have,

$$X^{\text{Ex}} \geq (\zeta - 1 + \delta) \left(\sum_{i \in I} B_i \right) \geq 0, \quad (38)$$

here the non-negativity of the expression follows from the fact that $\zeta \geq (1 - 1/e) \geq 1 - \delta$. From (19), we have

$$\zeta \sum_{i \in I} (o_i^{\text{In}} + o_i^{\text{Ex}}) \geq \eta \sum_{i \in I} o_i^{\text{In}} + W^{\text{Ex}} \left(\frac{X^{\text{Ex}}}{W^{\text{Ex}}} + G \left(1 - \frac{X^{\text{Ex}}}{W^{\text{Ex}}} \right) \right),$$

$$\begin{aligned}
& \zeta \sum_{i \in I} o_i^{\text{In}} + \zeta \delta \sum_{i \in I} B_i \stackrel{(a)}{\geq} \eta \sum_{i \in I} o_i^{\text{In}} + \delta \left(\sum_{i \in I} B_i \right) \left(\frac{\zeta - 1 + \delta}{\delta} + G \left(1 - \frac{\zeta - 1 + \delta}{\delta} \right) \right), \\
& (\zeta - \eta) \frac{\sum_{i \in I} o_i^{\text{In}}}{\sum_{i \in I} B_i} + \zeta \delta \geq \delta \left(\frac{\zeta - 1 + \delta}{\delta} + G \left(\frac{1 - \zeta}{\delta} \right) \right), \\
& (\zeta - \eta)(1 - \delta) + \zeta \delta \stackrel{(b)}{\geq} \zeta - 1 + \delta + \delta G \left(\frac{1 - \zeta}{\delta} \right), \\
& (1 - \eta) \frac{1 - \delta}{\delta} \stackrel{(c)}{\geq} e^{\frac{1 - \zeta - \delta}{\delta}} - e^{-1}, \\
& \zeta \geq 1 - \delta \left(1 + \ln \left(\frac{1}{e\delta} \right) \right), \\
& = 1 + \delta \ln \delta.
\end{aligned}$$

Inequality (a) follows from (17), (38), and the fact that $J(u) = u + G(1 - u)$ is an increasing function for $u \geq 0$. Inequality (b) follows from the upper bound $\sum_{i \in I} o_i^{\text{In}} \leq (1 - \delta) \sum_{i \in I} B_i$ that is implied by (15) and (17). Finally, inequality (c) follows by plugging in $G(u) = e^{u-1} - e^{-1}$. $\square$

H.3. Other Missing Proofs for Multi-Channel Traffic

LEMMA 7. *For any given $\delta \in [0, 1]$, to lower bound the competitive ratio of GPG it suffices to consider instances where*

(i) *OPT matches every arrival.*

(ii) *Internal traffic arrives before external traffic, i.e., there is no internal traffic after the first external arrival.*

Proof. To prove part (i), it suffices to show that for any graph G and capacities $(B_i)_{i \in I}$, removing unmatched arrivals from G does not decrease the ratio $\frac{\text{GPG}}{\text{OPT}}$. Clearly, removing arrivals that are not matched in OPT has no effect on OPT. It suffices to show that on every sample path Y , $\text{GPG}(Y)$ does not increase if we remove an arrival $t \in T$. Consider two cases based on the type of resource matched to t in $\text{GPG}(Y)$. First, suppose that t is either unmatched or matched to a resource with non-zero remaining budget at the end of the arrival sequence. Then, removing t does not change the matching. Second, suppose that t is matched to resource i such that $\text{GPG}_i(Y) = B_i$. Removing t may change the matching but does not increase the total budget consumption of any resource.

Consider an external arrival $t \in T$ in an arbitrary arrival sequence. First, notice that changing the sequence of arrivals does not change OPT. To prove part (ii), it suffices to show that moving arrival t to the end of the arrival sequence (keeping all other arrivals fixed) can not increase GPG. Consider an arbitrary sample path Y in GPG. If t is unmatched in $\text{GPG}(Y)$, moving t to the end of the sequence does not change the matching. However, if t is matched in $\text{GPG}(Y)$ then moving t to the end of the arrival sequence does not increase the total budget consumption of any resource. $\square$

Appendix I: Missing Proofs for Online Matching with Stochastic Rewards

One common way to state asymptotic competitive ratio results for vanishing probabilities is as follows,

$$\text{ALG} \geq \alpha \text{OPT} - o(1), \tag{39}$$

where α is the competitive ratio and $o(1)$ is a term that goes to 0 as $\max_{i \in I, t \in T} p_{i,t} \rightarrow 0$. This is the form used (explicitly) in Huang and Zhang (2020) and (implicitly) in Mehta et al. (2015). An alternative approach is to define an entirely multiplicative ratio as follows,

$$\text{ALG} \geq (1 - o(1)) \alpha \text{OPT},$$

where $o(1)$ is a term that goes to zero as $\max_{i \in I, t \in T} p_{i,t} \rightarrow 0$. Guarantees of this form were shown in Mehta and Panigrahi (2012) and Goyal and Udwani (2023). For simplicity, we state and prove our result in the sense of (39).

From stochastic rewards to Adwords: Given an instance of online matching with stochastic rewards where the edge probabilities are given by $p_{i,t}$ and $p = \max_{i \in I, t \in T} p_{i,t}$, consider the Adwords instance on the same graph with bids

$$b_{i,t} = p_{i,t} \quad \forall i \in I, t \in T.$$

and unknown stochastic budgets,

$$B_i \sim \text{Exp}(1) \quad \forall i \in I,$$

here $\text{Exp}(1)$ is the exponential distribution with mean 1. Taking expectation over the budgets, we have the following performance guarantee for this stochastic instance of Adwords.

LEMMA 15. *For the stochastic instance of Adwords described above, GPG is at least $0.522 - o(1)$ competitive for $\beta = 1.15$ and at least $0.508 - o(1)$ competitive for $\beta = 1$ (in expectation over the random budgets and seed values).*

Proof. Given starting budgets $\mathcal{B} = \{B_i\}_{i \in I}$, let

$$\gamma_i(B_i) = \max_{t \in T} \frac{b_{i,t}}{B_i} \leq \frac{p}{B_i} \quad \forall i \in I.$$

Let $\text{OPT}(\mathcal{B})$ and $\text{GPG}(\mathcal{B})$ denote the (expected) total reward in OPT and GPG respectively. For any fixed value of β , let α denote the asymptotic performance guarantee of GPG in the limit $\gamma \rightarrow +\infty$. We have,

$$\alpha(\mathcal{B}) = \frac{\text{GPG}(\mathcal{B})}{\text{OPT}(\mathcal{B})} \geq \frac{1}{1 + \max_{i \in I} \gamma_i(B_i)} \alpha \geq \frac{\min_{i \in I} B_i}{\min_{i \in I} B_i + p} \alpha.$$

Thus, $E_{\mathcal{B}}[\alpha(\mathcal{B})] \geq E_{\mathcal{B}}\left[\frac{\min_{i \in I} B_i}{\min_{i \in I} B_i + p}\right] \alpha$. Using the inequality $\frac{1}{1+x} \geq 1-x$, we have,

$$E_{\mathcal{B}}\left[\frac{\min_{i \in I} B_i}{\min_{i \in I} B_i + p}\right] \geq 1 - E_{\mathcal{B}}\left[\frac{p}{\min_{i \in I} B_i}\right].$$

For any finite set I , let E denote the event that $\min_{i \in I} B_i \geq \sqrt{p}$. Notice that the probability that E occurs equals 1 in the limit of $p \rightarrow 0$. Thus, $\alpha > E_{\mathcal{B}}[\alpha(\mathcal{B})] \geq \alpha - o(1)$, as desired. $\square$

LEMMA 3. *Lemma 14 in Huang and Zhang (2020). An online algorithm that is $\alpha - o(1)$ competitive for the stochastic instances of Adwords defined above, is also $\alpha - o(1)$ competitive for online matching with stochastic rewards.*

Proof of Theorem 7. Using Lemma 14 in Huang and Zhang (2020) (stated above) and Lemma 15, gives us the desired. $\square$

Appendix J: Potential Application in Automated Budget Management

In using a automated tools to manage their portfolio, the advertiser decides the overall budget, creates a portfolio of ad campaigns, and specifies a high level performance goal for the portfolio. Using these specifications, the tool aims to automatically determine (over time) a good budget distribution for the portfolio, as well as, daily budgets and bids for key words in each campaign. As the example below illustrates, fixing a (daily) budget at the start of the day may be highly sub-optimal.

Example J.1 Consider an advertiser with a portfolio composed of two search ad campaigns labeled $\{1, 2\}$. In the absence of a budget constraint, let the total expenditure (per day) in campaign 1 be a Bernoulli random variable $X_1 \in \{0, 1\}$. Similarly, let $X_2 \in \{0, 1\}$ represent the expenditure (per day) in campaign 2. We know that X_1 and X_2 are identically distributed with mean 0.5 and $X_1 = 1 - X_2$, i.e., the campaign expenditures are in perfect negative correlation. Now, suppose we have a total daily budget of 1 that needs to be distributed between the two campaigns at the start of each day. If we distribute the overall budget evenly i.e., allocate a daily budget of 0.5 to each campaign, then the total budget utilization in campaign $i \in \{1, 2\}$ is at most $E[\min\{0.5, X_i\}] = 0.25$, i.e., only 50% of the total budget is used. Splitting the budget across the campaigns in a different proportion does not improve the total utilization.

Now, suppose that instead of fixing the daily budget for each campaign at the beginning, we start by giving a budget of 0.5 to each campaign and use live cross-campaign data to adjust the distribution during the day. Since the campaigns are negatively correlated, exactly one of the two campaigns will have a non-zero expenditure on each day. Let i^* denote this campaign on a given day. Before the initial budget of i^* runs out, we can redirect the rest of the budget to i^* and achieve a budget utilization of 100%.

Platforms with in-house tools for automated budget optimization (such as Search Ads 360 by Google) may, in fact, have the cross-campaign data necessary to perform adjustments to budget distribution in real-time. At a conceptual level, armed with a budget oblivious allocation algorithm and live cross-campaign data, such a platform can start the day with some distribution of budget to each campaign and adjust the distribution during the day to improve overall utilization. When a campaign is about to run out of its tentatively assigned budget, the platform can either transfer unused budget from another active campaign to this one, or, let the campaign expire for the day. In this way, campaigns that do particularly well during the day would receive more budget and this improves the overall utilization of budget for the portfolio. A budget oblivious allocation algorithm is agnostic to such adjustments and only needs to be notified when the budget of a campaign expires for the day.