

Tight Competitive and Variance Analyses of Matching Policies in Gig Platforms

Anonymous Author(s)

ABSTRACT

The gig economy features dynamic arriving agents and on-demand services provided. In this context, instant and irrevocable matching decisions are highly desirable due to the low patience of arriving requests. In this paper, we propose an online-matching-based model to tackle the two fundamental issues, matching, and pricing, existing in a wide range of real-world gig platforms, including ride-hailing (matching riders and drivers), crowdsourcing markets (pairing workers and tasks), and online recommendations (offering items to customers). Our model assumes the arriving distributions of dynamic agents (e.g., riders, workers, and buyers) are accessible in advance, and they can change over time, which is commonly referred to as *Known Adversary Distributions* (KAD).

In this paper, we initiate variance analysis for online matching algorithms under KAD. Unlike the popular competitive-ratio (CR) metric, the variance of online algorithms' performance is rarely studied due to inherent technical challenges, though it is well linked to robustness. We focus on two natural parameterized sampling policies, denoted by $\text{ATT}(\gamma)$ and $\text{SAMP}(\gamma)$, which appear as foundational bedrock in online algorithm design. We offer rigorous competitive ratio (CR) and variance analyses for both policies. Specifically, we show that $\text{ATT}(\gamma)$ with $\gamma \in [0, 1/2]$ achieves a CR of γ and a variance of $\gamma \cdot (1 - \gamma) \cdot B$ on the total number of matches with B being the total matching capacity. In contrast, $\text{SAMP}(\gamma)$ with $\gamma \in [0, 1]$ accomplishes a CR of $\gamma(1 - \gamma)$ and a variance of $\bar{\gamma}(1 - \bar{\gamma}) \cdot B$ with $\bar{\gamma} = \min(\gamma, 1/2)$. All CR and variance analyses are tight and unconditional of any benchmark. As a byproduct, we prove that $\text{ATT}(\gamma = 1/2)$ achieves an optimal CR of $1/2$.

CCS CONCEPTS

• **Theory of computation** → **Online algorithms**; • **Applied computing** → *Marketing*.

KEYWORDS

Online Matching, Competitive Analysis, Variance Analysis

ACM Reference Format:

Anonymous Author(s). 2018. Tight Competitive and Variance Analyses of Matching Policies in Gig Platforms. In *Proceedings of The Web Conference (TheWebConf)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/XXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

TheWebConf, May 13–May 17, 2024, Singapore

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXX.XXXXXXX>

1 INTRODUCTION

Markets in the gig economy feature dynamic arriving agents that are typically called *online* as opposed to static agents called *offline* and an instant decision-making requirement due to the low patience of online agents. Examples of gig markets include ride-hailing platforms, crowdsourcing markets, and online recommendations. There are two main categories of studies for gig platforms. The first focuses on the matching issue regarding how to pair up online and offline agents in the system, say, e.g., matching (online) riders and (offline) drivers in ride-hailing [25, 30, 37], pairing (online) workers and (offline) tasks in crowdsourcing markets [13, 14], and offering (offline) items to (online) customers in online recommendations [18]. These works aim to design an efficient matching policy to facilitate as many matches as possible. The second considers another one, called pricing, [4–6, 22, 26, 38], where they typically assume dynamic arriving agents have known or unknown value distributions and we need to design a pricing strategy to promote the total profit flowing into the system.

In this paper, we consider the two issues of matching and pricing simultaneously. Specifically, we assume that the arriving distributions of dynamic agents (e.g., riders in ride-hailing, workers in crowdsourcing, and buyers in online recommendations) and their value functions toward prices are all accessible in advance as part of the input. Our assumptions are inspired by the fact that we can typically learn and estimate these distributions by applying powerful machine learning techniques to massive historical data [9, 19, 34]. We present our model in detail as follows.

Matching and Pricing in Gig Economies. We use a bipartite graph $G = (I, J, E)$ to model the network between a set of offline (static) agent types I and a set of online (dynamic) agent types J , where an edge $e = (i, j)$ indicates the feasibility of matching agents between types i and j due to practical constraints. Examples include spatial and temporal constraints between a driver (of type) i and a rider (of type) j in ride-hailing, and the potential interest of a customer j toward an item i in online recommendations. We have a groundset of prices, denoted by $\mathcal{A} = \{a_k | k \in [K] = \{1, 2, \dots, K\}\}$. We refer to each tuple $f = (i, j, k)$ with $(i, j) \in E$ as an assignment, which represents to match an online agent j with an offline agent i and charge j at the price of a_k . Thus, our assignment consists of two parts: *matching and pricing*. For each offline agent $i \in I$, it has a capacity $b_i \in \mathbb{Z}_+$ representing an upper bound on the number of matches involving i . This captures matching caps imposed on i , which can be interpreted as the total number of drivers of type i featuring a specific working location in rideshare, the total amount of item i in stock in online recommendations, etc.

The online arrival process is as follows. We have a time horizon of T rounds and during each round (or time) $t \in [T] := \{1, 2, \dots, T\}$, one single online agent j will be sampled from J such that $\Pr[j =$

$j] = q_{j,t}$ with $\sum_{j \in J} q_{j,t} = 1$ (called j arrives at t).¹ Let $\mathcal{F} = \{f = (i, j, k) | (i, j) \in E\}$ be the collection of all feasible assignments. For each $f = (i, j, k) \in \mathcal{F}$ and $t \in [T]$, it is associated with a probability $p_{f,t} \in (0, 1]$ and a non-negative profit $w_{f,t}$, which denote the chance of online agent j accepts offline agent i with the price of a_k at time t and the total profit flowing into the system, respectively. Upon the arrival of online agent j at t , we (as the algorithm) should either reject j or select an assignment $f = (i, j, k)$ involving j . Our decision is instant and irrevocable, and it is then followed by a Bernoulli random choice of j toward the assignment $f = (i, j, k)$: with probability $p_{f,t}$, j accepts f (then i 's capacity gets reduced by one and we get a profit of $w_{f,t}$); and with probability $1 - p_{f,t}$, j rejects f and goes away.²

An input instance can be characterized as $\mathcal{I} = \{G = (I, J, E), \mathcal{A}, \{b_i\}, T, \{q_{j,t}\}, \{p_{f,t}, w_{f,t}\}\}$, whose information is all accessible to the algorithm. Our goal is to design a matching-and-pricing policy such that the total profit is maximized. Observe that we allow the arrival distributions of online agents to change over time, which is commonly referred to as *Known Adversarial Distributions* (KAD) [12, 40]. In light of this, we refer to our model as *Matching and Pricing under Known Adversarial Distributions* (MP-KAD).

Remarks on the Sources of Randomness in MP-KAD. There are three sources of randomness in total: the first is random arrivals of online agents (Q1); the second is potentially randomized choices of assignments made by a policy (Q2); the third is random acceptance/rejection decisions from online agents over assignments presented by a policy (Q3). We assume (Q1) and (Q3) are independent of each other. Also, the arriving distributions of online agents in (Q1) and their Bernoulli random choices in (Q3) are all independent over time.

1.1 Preliminaries

Competitive Ratio (CR). CR is a metric commonly used to evaluate the performance of online algorithms. Consider a given (online) policy (or algorithm) ALG and a clairvoyant optimal (OPT-OFF). Note that ALG is subject to the real-time decision-making requirement, *i.e.*, ALG has to make an irrevocable decision upon every arrival of online agents before observing future arrivals. In contrast, OPT-OFF is exempt from that requirement: it enjoys the privilege of observing the full arrival sequence of online agents *before* optimizing decisions. Consider a given instance \mathcal{I} of MP, and let $E[\text{ALG}]$ and OPT-OFF denote the *expected* total profit achieved by ALG and OPT-OFF, respectively. We say ALG achieves a CR of at least $\rho \in [0, 1]$ if $E[\text{ALG}] \geq \rho \cdot \text{OPT-OFF}$ for all possible instances \mathcal{I} . Essentially, CR captures the gap between a policy and a clairvoyant optimal due to the real-time decision-making requirement imposed on the former.

Variance Analysis in Online Algorithm Design. For online optimization problems like MP-KAD, we can run any policy *only once* on any given instance. This highlights the importance of variance analysis in online algorithm design, which offers valuable insights

¹We can always make it equal by creating a dummy node whose arrival simulates the case of no arrival at t .

²Generally, the profit $w_{f,t}$ collected by the system from $f = (i, j, k)$ at t can be estimated as the price a_k charged to agent j minus the payoff for agent i .

into the robustness of related policies. As pointed out by [41], “the CR-metric reflects only the gap between an online policy (ALG) and a clairvoyant optimal in terms of their *expected* performance: it has no guarantee on the variance or robustness of ALG.” Additionally, *we want to stress that variance analysis plays a critical role in risk evaluation, particularly for maximization problems as studied here compared with minimization versions.* For minimization problems (say, minimizing some cost), the information of expectation itself can help us upper bound the risk. Let $X \geq 0$ be the total cost incurred by a policy. By applying Markov’s inequality, we can upper bound the risk as $\Pr[X \geq N] \leq E[X]/N$ for any target $N > 0$. However, it is a totally different story for maximization as studied here. A lower bound on the expected profit offers no guarantee on the chance of a disastrous event occurring. Specifically, let X denote the random profit gained by a policy. It is possible that $\Pr[X \leq \epsilon] \geq 1 - \delta$ for a given threshold $0 < \epsilon \ll 1$ and any $0 < \delta < 1$ for whatever given value of $E[X]$ when the variance of X is missing. However, after adding the information of any upper bound on the variance, we can immediately estimate the risk by citing concentration inequalities, *e.g.*, $\Pr[X \leq \epsilon] \leq \Pr[|X - \mu| \geq \mu - \epsilon] \leq \text{Var}[X]/(\mu - \epsilon)^2$ for any $\epsilon < \mu = E[X]$.

Inspired by the work of [41], we focus on the variance on *the total number of successful assignments (scheduled and accepted)*, which is due to the randomness as outlined in (Q1, Q2 and Q3); see **Remarks on sources of randomness in MP-KAD**.³ A detailed discussion on the similarities and differences from the work [41] can be seen in Section 1.3.

1.2 Benchmark Linear Program (LP)

For ease of presentation, we assume WLOG that $b_i = 1$ for all $i \in I$ by creating b_i copies of offline agent i where each has a unit capacity. Thus, the total capacities $B = \sum_{i \in I} b_i = |I|$. For each assignment $f = (i, j, k) \in \mathcal{F}$, let $x_{f,t}$ be the probability that f is selected in a clairvoyant optimal OPT-OFF during round t , which includes the probability that j arrives at t (but excludes that f is accepted or rejected by j). For each given $j \in J$ and $i \in I$, let $\mathcal{F}_j = \{f = (i, j, k) | (i, j) \in E\}$ and $\mathcal{F}_i = \{f = (i, j, k) | (i, j) \in E\}$ be collections of assignments involving j and i , respectively.

$$\max \sum_{t \in [T]} \sum_{f \in \mathcal{F}} x_{f,t} \cdot p_{f,t} \cdot w_{f,t}, \quad (1)$$

$$\sum_{f=(i,j,k) \in \mathcal{F}_j} x_{f,t} \leq q_{j,t}, \forall j \in J, \quad \forall t \in [T] \quad (2)$$

$$\sum_{t \in [T]} \sum_{f=(i,j,k) \in \mathcal{F}_i} x_{f,t} \cdot p_{f,t} \leq b_i = 1, \forall i \in I \quad (3)$$

$$0 \leq x_{f,t}, \quad \forall f \in \mathcal{F}. \quad (4)$$

LEMMA 1. *The optimal value of LP (1) is a valid upper bound on the performance of a clairvoyant optimal of MP-KAD.*

PROOF. For ease of notation, we use OPT-OFF to denote both a clairvoyant optimal algorithm and its corresponding performance (*i.e.*, expected amount of profit obtained). For each assignment $f = (i, j, k)$, let $X_{f,t} = 1$ indicate that f is selected by OPT-OFF

³As noted by [41], variations in profits can lead to an unbounded variance in the total profit even for simple deterministic policies; thus, we study an unweighted version to make our problem technically tractable.

Table 1: Notations used throughout this paper.

$[n]$	Set of integers $\{1, 2, \dots, n\}$ for a generic integer n .
G	Input network graph $G = (I, J, E)$.
$I(J)$	Set of offline static agent types (online dynamic types).
$\mathcal{A} = (a_k)$	Groundset of prices.
$f = (i, j, k)$	An assignment of matching agents (of type) j and i and charging j a price of a_k .
T	Time horizon.
$q_t = (q_{j,t})$	Arriving distribution at t with $\sum_{j \in J} q_{j,t} = 1$ for all $t \in [T]$.
$p_{f=(i,j,k),t}$	Probability that agent j accepts i with price of a_k at time t .
$w_{f,t}$	Profit associated with assignment f if it is successfully made at t .
b_i	Matching capacity of agent $i \in I$.
OPT-LP	Optimal value of Benchmark LP (1).
OPT-OFF	A clairvoyant optimal and its corresponding performance.
OPT-ON	An online optimal and its performance (subject to the real-time decision restrictions).

at t (not necessarily accepted by j) with $E[X_{f,t}] = x_{f,t}$. Thus, the performance can be expressed as $\text{OPT-OFF} = \sum_{f \in \mathcal{F}} \sum_{t \in [T]} w_{f,t} \cdot p_{f,t} \cdot E[X_{f,t}] = \sum_{f \in \mathcal{F}} \sum_{t \in [T]} w_{f,t} \cdot p_{f,t} \cdot x_{f,t}$, which matches the objective of LP (1). To prove Lemma 1, it suffices to show that $\{x_{f,t}\}$ is feasible to LP (1).

Let $Y_{j,t} = 1$ indicate that j arrives at t with $E[Y_{j,t}] = q_{j,t}$. Observe that for any given j and t , $\sum_{f \in \mathcal{F}_j} X_{f,t} \leq Y_{j,t}$. Taking expectation on both sides, we get Constraint (2). For each $f = (i, j, k)$ and t , let $Z_{f,t} \sim \text{Ber}(p_{f,t})$ simulate the random choices of acceptance and rejection from j over f at t . Note that $\sum_{f=(i,j,k) \in \mathcal{F}_i} \sum_{t \in [T]} X_{f,t} \cdot Z_{f,t} \leq b_i$ due to the matching capacity of b_i on offline agent i . Taking expectation on both sides yields Constraint (3). The last constraint is trivial. Thus, we justify the feasibility of $\{x_{f,t}\}$ to LP (1) and establish Lemma 1. \square

1.3 Main Contributions and Related Work

In this paper, we introduce a stochastic optimization model designed to address two fundamental issues, matching and pricing, which are prevalent in a wide range of real-world matching markets. Our focus is on two *natural* LP-based sampling policies that serve as foundational elements in online algorithm design: One includes attenuations, while the other does not. It is important to emphasize that these two LP-based sampling policies and their variations appear in various contexts within online matching markets, as demonstrated by examples in [11, 12, 16, 27, 41]. *Therefore, our primary technical contributions do not lie in algorithm design but rather in providing rigorous and comprehensive competitive and variance analyses for these two representative policies, which we expect can be generalized to other similar settings.*

THEOREM 1. [Section 2] *There is an LP-based sampling policy **with** attenuations parameterized with $\gamma \in [0, 1/2]$, denoted by $\text{ATT}(\gamma)$, which achieves (1) a competitive ratio (CR) of at least γ for MP-KAD; (2) a variance of at most $\gamma \cdot (1 - \gamma) \cdot B$ on the total number of successful assignments, where $B = \sum_{i \in I} b_i$, and (3) $\text{ATT}(\gamma = 1/2)$ achieves an optimal CR of $1/2$ for MP-KAD. Both competitive and variance analyses are tight for any $\gamma \in [0, 1/2]$, which are irrespective of the benchmark LP.*

THEOREM 2. [Section 3] *There is an LP-based sampling policy **without** attenuations parameterized with $\gamma \in [0, 1]$, denoted by*

SAMP(γ), which achieves (1) a competitive ratio (CR) of at least $\gamma \cdot (1 - \gamma)$ for MP-KAD; (2) a variance of at most $\bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B$ on the total number of successful assignments, where $\bar{\gamma} = \min(1/2, \gamma)$ and $B = \sum_{i \in I} b_i$. Both competitive and variance analyses are tight for any $\gamma \in [0, 1]$, which are irrespective of the benchmark LP.

Finally, we implement both ATT and SAMP and compare them to several heuristics on a real dataset provided by DiDi, Inc., collected in Haikou, China. Detailed results can be seen in the Appendix.

Remarks on Theorems 1 and 2. (1) The tightness of the competitive analysis of $\text{ATT}(\gamma)$ unconditional of the benchmark LP means that we can identify an instance of MP-KAD on which $\text{ATT}(\gamma)$ achieves a CR *equal* to γ for any $\gamma \in [0, 1/2]$, where the ratio is computed directly against the clairvoyant optimal by definition instead of the LP value.⁴ Similarly, the tightness of the variance analysis suggests that we can identify an instance of MP-KAD on which $\text{ATT}(\gamma)$ achieves a variance *equal* to γ on the total number of successful matches for any $\gamma \in [0, 1/2]$. The same interpretation applies to $\text{SAMP}(\gamma)$. (2) Our analysis indicates that there exists an instance, where $\text{ATT}(\gamma)$ achieves the worst (smallest) CR of γ and the worst (largest) variance of $\gamma(1 - \gamma)B$ simultaneously for any $\gamma \in [0, 1/2]$; see the example shown in Figure 2. However, that is not necessarily true for SAMP. The CR worst-scenario instance for SAMP shown in Figure 3 has a very different structure from the variance one as illustrated in Figure 4. (3) The optimality of $\text{ATT}(\gamma = 1/2)$ is unconditional of the benchmark LP. In other words, no policy can achieve a CR strictly better than $1/2$ for MP-KAD even compared against the clairvoyant optimal directly; see the instance shown in Figure 1. (4) The value of $\gamma \cdot (1 - \gamma)$ strictly increases when $\gamma \in [0, 1/2]$. Thus, for $\text{ATT}(\gamma)$, the worst-scenario (WS) competitive ratio and variance both increase when $\gamma \in [0, 1/2]$, which suggests a larger profit could come along with a higher variance and vice versa. For $\text{SAMP}(\gamma)$, the same trend applies when $\gamma \in [0, 1/2]$, though the WS variance remains unchanged when $\gamma \in [1/2, 1]$. (5) As demonstrated in Theorems 1 and 2, $\text{ATT}(\gamma)$ attains a strictly superior competitive ratio (CR) while maintaining the same variance

⁴Note that the latter option yields only a lower bound on CR since the LP value is an upper bound on the clairvoyant optimal by Lemma 1.

Table 2: Comparison of the competitive ratio (CR) and variance achieved by the two LP-based sampling policies, ATT, and SAMP, in the two models proposed by [41] (results marked in blue) and in this paper (marked in red). For the validity of the assessment, the results in [41] listed below are obtained when $\Delta = 1$ (the number of resources incurred) since each assignment here could cost one unit of the corresponding offline agent’s matching capacity only. The terms “conditionally” and “unconditionally” are with respect to the benchmark LP. The unconditionally optimal CR of $1/2$ for MP-KAD contrasts with the conditionally tight CR of $1 - 1/e \sim 0.632$ for the model in [41], primarily due to the strict generalization in the arrival setting from KIID [41] to KAD, as studied here.

Algorithms	Range of γ	CR Bounds	Variance Bounds	Best CR
ATT(γ)	$\gamma \in [0, 1]$	$1 - e^{-\gamma}$	$(1 - e^{-2\gamma} - 2\gamma e^{-\gamma})T^2$	$1 - e^{-1}$ Conditionally Tight
	$\gamma \in [0, 1/2]$	γ	$\gamma(1 - \gamma) \cdot B$	$1/2$ Unconditionally Optimal
SAMP(γ)	$\gamma \in [0, 1]$	$1 - e^{-\gamma}$	$(1 - e^{-2\gamma} - 2\gamma e^{-\gamma})T^2$	$1 - e^{-1}$ Conditionally Tight
	$\gamma \in [0, 1]$	$\gamma(1 - \gamma)$	$\bar{\gamma}(1 - \bar{\gamma}) \cdot B, \bar{\gamma} = \min(1/2, \gamma)$	$1/4$ Unconditionally Tight

as SAMP(γ) for any $\gamma \in [0, 1/2]$. However, in practice, SAMP enjoys greater efficiency compared to ATT because it does not require the computation of attenuation factors, which is necessary for ATT.

Comparison Against the Work of [41]. We outline the distinctions between our work and that of [41] across three different dimensions. **Models.** Both studies focus on allocation policy design within an online-matching-based framework. However, while they concentrate solely on the matching aspect, our approach encompasses both matching and pricing. Additionally, both models accommodate stochastic arrivals of online agents under *known distributions*. In contrast to the assumption of *Known Identical Independent Distributions* (KIID) in [41], we operate within a more general setting known as *Known Adversary Distributions* (KAD), allowing for dynamic changes in arriving distributions over time.⁵ Moreover, the two models introduce randomness differently into the matching process. In [41], randomness arises from the cost realization: Each match $e = (i, j)$ incurs a *random* set of static resources. Any policy can match e only when a sufficient budget remains for every possibly needed resource. Importantly, once matched, e is guaranteed to be accepted by j (as no pricing is involved). In contrast, our model introduces randomness in the two possible outcomes—acceptance and rejection—of an arriving agent toward the price included in an assignment. **Algorithms.** Both papers introduce two parameterized LP-based sampling policies, denoted as ATT and SAMP. The key distinction lies in the presence of attenuations in ATT whereas SAMP does not incorporate them. However, it’s worth noting that the two papers propose different benchmark LPs due to variations in their respective models. Additionally, for the policy ATT, debilitation factors are computed precisely from the LP solution, whereas in [41], they are derived through Monte Carlo

simulations.⁶ **Results.** Both papers claim to conduct tight competitive ratio (CR) analyses for the two LP-based sampling policies. Notably, in [41], the CR tightness is relative to the corresponding benchmark LP, whereas in our analysis, it is unconditional. Furthermore, our variance bounds eliminate the dependence on the length of the time horizon T , which is typically assumed to be $T \gg 1$ and can be far larger than the budget B . More differences on the results can be seen in Table 2.

Other Related Work. There are quite a few previous works that have studied the matching [24, 30, 36] and pricing issues [3, 8, 17, 35] in matching markets. Unlike our setting here, most of them assume at least part of the input is unknown, say, *e.g.*, arriving distributions of online agents and/or acceptance and rejection probabilities are unknown. As a result, they proposed some machine-learning-based frameworks to manage the learning tasks during matching and pricing, among which reinforcement learning and multi-armed bandits are two of the most common paradigms. Another line of research studies has considered matching and pricing jointly in ride-hailing platforms but under an essentially static setting, where requests are assumed known in advance instead of arriving dynamically [23, 32]. In that case, authors typically utilize integer linear programming to resolve the matching issue.

Matching and pricing have also received significant attention in the Operations Research community. We list a few examples as follows. Özkan and Ward [31] have considered matching policy design for ride-hailing services, where they assume both drivers and riders join and depart the system stochastically, and thus, random sojourn time is allowed for every agent. Mahavir Varma et al. [28] have studied a similar setting featured by a two-sided arrival model, and they mainly utilize MDP-based techniques to address the matching and pricing issues. Vera et al. [39] have investigated online resource allocation but focus on evaluating the performance

⁵In practice, KAD (Known Adversary Distributions) is a more realistic arrival setting than KIID (Known Identical Independent Distributions) because the arriving distributions of online agents do exhibit variations over time. Please see Figure 8 in the Appendix for a justification on a real ride-hailing dataset.

⁶Monte-Carlo simulations are widely used to approximate attenuation factors [1, 7, 12, 20, 27]. Our policy ATT features that all attenuation factors can be pre-computed explicitly and exactly, which suggests superiority in efficiency in practice; see details of ATT in Algorithm 1.

in terms of regrets, the difference in the profit between a policy and a Prophet, which is very different from the CR metric here.

2 AN LP-BASED SAMPLING POLICY WITH ATTENUATIONS (ATT)

Slightly abusing the notation, we use $\{x_{f,t}\}$ to denote an optimal solution to LP (1). The overall picture of ATT(γ) with $\gamma \in (0, 1/2]$ works as follows: During each time t , we sample an assignment $f = (i, j, k) \in \mathcal{F}_j$ with probability $(x_{f,t}/q_{j,t}) \cdot (\gamma/\beta_{i,t})$ upon the arrival of online agent j , where γ represents the target competitive ratio we aim to achieve, and $\beta_{i,t}$ is a pre-calculated attenuation factor. The formal statement of ATT(γ) is as follows.

Algorithm 1: A sampling policy with attenuations for MP-KAD: ATT(γ), $\gamma \in [0, 1/2]$.

- 1 **Offline Phase:**
 - 2 Solve LP (1) and let $\{x_{f,t}\}$ be an optimal solution.
 - 3 Compute $\beta_{i,t} = 1 - \gamma \sum_{t' < t} \sum_{f \in \mathcal{F}_i} x_{f,t'} \cdot p_{f,t'}$ for all $i \in I$ and $t \in [T]$.
 - 4 **Online Phase:**
 - 5 **for** $t = 1, \dots, T$ **do**
 - 6 Let an online agent (of type) j arrive at time t .
 - 7 Sample an assignment $f = (i, j, k) \in \mathcal{F}_j$ with probability $\frac{x_{f,t}}{q_{j,t}} \frac{\gamma}{\beta_{i,t}}$.
 - 8 /* At most one assignment will be sampled in Step (7) since $\sum_{f \in \mathcal{F}_j} (x_{f,t}/q_{j,t}) \cdot (\gamma/\beta_{i,t}) \leq \sum_{f \in \mathcal{F}_j} x_{f,t}/q_{j,t} \leq 1$, which follows from $\beta_{i,t} \geq \gamma$ shown in Ineq. (5) and Const. (2) of LP (1). */
 - 9 **if** i is safe at t (the capacity of i remains), **then**
 - 10 Select the assignment f (i.e., assigning j to i and charging j a price of a_k);
 - 11 **else**
 - 12 Reject f .
-

2.1 Competitive Analysis of ATT(γ)

THEOREM 3. ATT(γ) with $\gamma \in [0, 1/2]$ achieves a competitive ratio of γ for MP-KAD.

PROOF. We first justify the validity of ATT by showing that the total sum of probabilities on Step (7) is no larger than one. Note that for any i and t ,

$$\beta_{i,t} = 1 - \gamma \cdot \sum_{t' < t} \sum_{f \in \mathcal{F}_i} x_{f,t'} \cdot p_{f,t'} \geq 1 - \gamma \geq \gamma, \quad (5)$$

where the first inequality is due to Constraint (3) with $b_i = 1$, and the second follows from $\gamma \leq 1/2$. Thus, $\sum_{f=(i,j,k) \in \mathcal{F}_j} \frac{x_{f,t}}{q_{j,t}} \frac{\gamma}{\beta_{i,t}} \leq \sum_{f=(i,j,k) \in \mathcal{F}_j} \frac{x_{f,t}}{q_{j,t}} \leq 1$, where the second inequality is due to Constraint (2) in LP (1).

For each i and t , let $SF_{i,t} = 1$ indicate that offline agent (of type) i is safe at t , i.e., it has one unit capacity at (the beginning of) t before any online actions and $SF_{i,t} = 0$ otherwise. Let $\alpha_{i,t} = E[SF_{i,t}]$ be

the probability that i is safe at t . For each assignment f , let $\chi_{f,t} = 1$ indicate that f is successfully made at t (scheduled and accepted). We now show by induction on $t \in [T]$ that **(P1)** $\alpha_{i,t} = \beta_{i,t}$ and **(P2)** $E[\chi_{f,t}] = \gamma \cdot x_{f,t} \cdot p_{f,t}$ for all $i \in I, t \in [T]$, and $f \in \mathcal{F}$. Consider the base case $t = 1$. We see that $\alpha_{i,t} = \beta_{i,t} = 1$ for all $i \in I$. For each $f = (i, j, k)$, let $X_{j,t} = 1$ indicate that j arrives at t , $Y_{f,t} = 1$ indicate that f gets sampled at t , and $Z_{f,t} = 1$ indicate that j accepts f at t . Thus, for $t = 1$,

$$\begin{aligned} E[\chi_{f,t}] &= E[X_{j,t} \cdot Y_{f,t} \cdot SF_{i,t} \cdot Z_{f,t}] \\ &= q_{j,t} \cdot (x_{f,t}/q_{j,t}) \cdot (\gamma/\beta_{i,t}) \cdot \alpha_{i,t} \cdot p_{f,t} = \gamma \cdot x_{f,t} \cdot p_{f,t}. \end{aligned} \quad (6)$$


Now consider a given $t > 1$ and assume **(P1)** and **(P2)** are valid for all $t' < t$. Consider a given $i \in I$. Note that

$$\begin{aligned} \alpha_{i,t} &= E[SF_{i,t}] = 1 - E\left[\sum_{t' < t} \sum_{f \in \mathcal{F}_i} \chi_{f,t'}\right] \\ &= 1 - \sum_{t' < t} \sum_{f \in \mathcal{F}_i} \gamma \cdot x_{f,t'} \cdot p_{f,t'} = \beta_{i,t}, \end{aligned}$$

where the third equality is due to the inductive assumption **(P2)** and the last one follows from the definition of $\beta_{i,t}$. We can verify that Equation (6) remains valid for the given t and all f as long as $\alpha_{i,t} = \beta_{i,t}$ for t and all $i \in I$. Thus, we complete the inductive step on **(P1)** and **(P2)**. By linearity of expectation, the total expected profit of ATT(γ) is equal to

$$E[\text{ATT}(\gamma)] = \sum_{f \in \mathcal{F}} \sum_{t \in [T]} w_{f,t} \cdot E[\chi_{f,t}] = \sum_{f \in \mathcal{F}} \sum_{t \in [T]} w_{f,t} \cdot \gamma \cdot x_{f,t} \cdot p_{f,t},$$

which is a fraction of γ of the optimal value of LP (1). By Lemma 1, we claim that ATT achieves a competitive ratio at least γ . \square



$$\begin{aligned} p &= 1, T = 2, x_{f_1, t=1}^* = 1 - \epsilon, x_{f_3, t=2}^* = \epsilon, \\ \text{OPT-LP} &= \text{OPT-OFF} = 2 - \epsilon, \text{OPT-ON} = 1, \\ E[\text{ATT}(\gamma)] &= \gamma \cdot (2 - \epsilon), \frac{E[\text{ATT}(\gamma)]}{\text{OPT-OFF}} = \gamma. \end{aligned}$$

Figure 1: An example highlighting the tightness of competitive analysis of ATT(γ) for any $\gamma \in [0, 1/2]$ and the optimality of ATT($\gamma = 1/2$), where both claims are unconditional of the benchmark LP.

2.2 Tightness of the CR Analysis of $\text{ATT}(\gamma)$ with $\gamma \in [0, 1/2]$ and the Optimality at $\gamma = 1/2$

LEMMA 2. *There exists an instance of MP-KAD such that $\text{ATT}(\gamma)$ achieves a competitive ratio of γ for any $\gamma \in [0, 1/2]$. Meanwhile, $\text{ATT}(\gamma = 1/2)$ achieves an optimal competitive ratio of $1/2$ for MP-KAD. Both claims here are irrespective of the benchmark LP.*

EXAMPLE 1. *Consider such an instance as shown in Figure 1. We have one single offline agent with a unit budget that is connected to three online agents indexed by $j = 1, 2, 3$, respectively. Set $T = 2$, and the arriving distributions at $t = 1$ and $t = 2$ are $\mathbf{q}_1 = (1, 0, 0)$ and $\mathbf{q}_2 = (0, 1 - \epsilon, \epsilon)$ with $\epsilon > 0$, respectively. There is one single price, and thus, each edge itself represents an assignment. For ease of notation, we use j to index the three edges and the corresponding assignments, i.e., $f_j = (i = 1, j, *)$ for $j = 1, 2, 3$. Set $p_{f,t} = 1$ for all f and t , and the profit on assignments state as follows: $w_{f_1,t} = 1$, $w_{f_2,t} = 0$, and $w_{f_3,t} = 1/\epsilon$ for both $t = 1, 2$.*

PROOF. Consider the example shown in Figure 1. Let OPT-ON, OPT-OFF, and OPT-LP denote the (expected) performance of an online optimal policy, that of a clairvoyant optimal, and the optimal value of Benchmark LP (1) on the above example. We can verify the following facts. First, LP (1) has such an optimal solution that $x_{f_1,t=1}^* = 1 - \epsilon$, $x_{f_3,t=2}^* = \epsilon$ with all rest being zeros, and the corresponding optimal value is OPT-LP = $2 - \epsilon$. The performance of a clairvoyant optimal is OPT-OFF = $\epsilon \cdot (1/\epsilon) + (1 - \epsilon) = 2 - \epsilon$, while that of an online optimal is OPT-ON = 1, which is subject to the real-time decision-making requirements. Second, $\text{ATT}(\gamma)$ with $\gamma \in [0, 1/2]$ achieves a performance of $\gamma \cdot (1 - \epsilon) + \gamma \cdot \epsilon \cdot (1/\epsilon) = \gamma \cdot (2 - \epsilon)$. Thus, we claim that (1) $\text{ATT}(\gamma)$ achieves a CR of $\gamma(2 - \epsilon)/(2 - \epsilon) = \gamma$ using the benchmark of either LP (1) or the clairvoyant optimal for any $\gamma \in [0, 1/2]$; and (2) $\text{ATT}(\gamma = 1/2)$ achieves a CR of $1/2$ for MP-KAD, which is optimal since no algorithm can beat the CR of OPT-ON/OPT-OFF = $1/(2 - \epsilon)$, which approaches $1/2$ when $\epsilon \rightarrow 0_+$.

2.3 Variance Analysis of $\text{ATT}(\gamma)$

Recall that for each $f \in \mathcal{F}$ and $t \in [T]$, $\chi_{f,t} = 1$ indicates f is successfully made (scheduled and accepted) in $\text{ATT}(\gamma)$. Let $H_i = \sum_{f \in \mathcal{F}_i} \sum_{t \in [T]} \chi_{f,t}$ and $H = \sum_{i \in I} H_i$ denote the numbers of successful assignments involving i and in total, respectively. Note that $B = \sum_{i \in I} b_i = |I|$.

THEOREM 4. $\text{Var}[H] \leq \gamma \cdot (1 - \gamma) \cdot B$ with $\gamma \in [0, 1/2]$.


PROOF. For each assignment $f = (i, j, k)$, let $X_{j,t} = 1$, $Y_{f,t} = 1$, and $Z_{f,t} = 1$ indicate that j arrives at t , f gets sampled at t in $\text{ATT}(\gamma)$, and j accepts f at t , respectively. For each $f = (i, j, k)$, let $W_{f,t} = X_{j,t} \cdot Y_{f,t} \cdot Z_{f,t}$. Note for each given time t , we have $\sum_{f \in \mathcal{F}} W_{f,t} \leq \sum_{j \in J} X_{j,t} \leq 1$ since at most one single assignment will be sampled in Step (7) of $\text{ATT}(\gamma)$. Thus, by the Zero-One Principle [15], $\{W_{f,t} | f \in \mathcal{F}\}$ are negatively associated for each given t . Observe that $\mathcal{W}_t := \{W_{f,t} | f \in \mathcal{F}\}$ are independent from $\mathcal{W}_{t'}$ as long as $t \neq t'$. As a result, we claim $\{W_{f,t} | f \in \mathcal{F}, t \in [T]\}$ are negatively associated.

Let $\mathcal{W}_i = \{W_{f,t} | f \in \mathcal{F}_i, t \in [T]\}$ for each given $i \in I$. Recall that $H_i = \sum_{f \in \mathcal{F}_i} \sum_{t \in [T]} \chi_{f,t} = 1$ indicates that one assignment in \mathcal{F}_i is successfully made in $\text{ATT}(\gamma)$. Since each i has a unit capacity, $H_i = 1$

iff there exists at least one $W_{f,t} \in \mathcal{W}_i$ with $W_{f,t} = 1$. Consequently, $H_i = \min(1, \sum_{t \in [T]} \sum_{f \in \mathcal{F}_i} W_{f,t})$, which can be viewed as a non-decreasing function over \mathcal{W}_i . Therefore, $\{H_i\}$ can be regarded as a set of non-decreasing functions over disjoint subsets of negatively associated random variables of $\{\mathcal{W}_i | i \in I\}$, which suggests that $\{H_i | i \in I\}$ are also negatively associated [21]. By the result in [33], $\text{Var}[H] = \text{Var}[\sum_{i \in I} H_i] \leq \sum_{i \in I} \text{Var}[H_i]$.

Observe that H_i is a Bernoulli random variable with mean $E[H_i] = \sum_{t \in [T]} \sum_{f \in \mathcal{F}_i} E[\chi_{f,t}] = \sum_{t \in [T]} \sum_{f \in \mathcal{F}_i} \gamma \cdot x_{f,t} \cdot p_{f,t} \leq \gamma$, where the second equality is due to Equation (6), while the last inequality is due to Constraint (3) in LP (1). Thus, $\text{Var}[H_i] \leq \gamma \cdot (1 - \gamma)$ since $\gamma \in [0, 1/2]$. Therefore, we claim that

$$\text{Var}[H] = \text{Var}\left[\sum_{i \in I} H_i\right] \leq \sum_{i \in I} \text{Var}[H_i] \leq \gamma \cdot (1 - \gamma) \cdot B. \quad \square$$



$$|I| = |J| = T = m, p = 1, w = 1, q_{j=t,t} = 1, q_{j \neq t,t} = 0, \forall t \in [T],$$

$$x_{i=t,j=t,t}^* = 1, \forall t \in [T], H = \sum_{i \in I} H_i, H_i = \text{Ber}(\gamma), \forall i \in I,$$

$$\text{Var}[H] = \gamma \cdot (1 - \gamma) \cdot m = \gamma \cdot (1 - \gamma) \cdot B,$$

$$\text{OPT-LP} = \text{OPT-OFF} = m, E[\text{ATT}(\gamma)]/\text{OPT-OFF} = \gamma.$$

Figure 2: An example where $\text{ATT}(\gamma)$ achieves the tight CR bound of γ and the tight variance bound of $\gamma(1 - \gamma) \cdot B$ simultaneously for any $\gamma \in [0, 1/2]$.

2.4 Tightness of the Variance Analysis of $\text{ATT}(\gamma)$ for any $\gamma \in [0, 1/2]$

LEMMA 3. *There exists an instance of MP-KAD such that $\text{ATT}(\gamma)$ achieves a variance of $\gamma \cdot (1 - \gamma) \cdot B$ on the (random) number of successful assignments for any $\gamma \in [0, 1/2]$.*

PROOF. Consider the instance as shown in Figure 2. We have a graph $G = (I, J, E)$ with $|I| = |J| = |E| = m$, where $E = \{(i, j) | i = j \in [m]\}$ that consists of m edges. There is one single price, and thus, each edge one-one corresponds to an assignment. For ease of notation, we use e to represent the corresponding assignment. Each offline agent type i has a unit capacity $b = 1$. Let $T = m$ and during each time $t \in [T]$, $j = t$ arrives with probability one and no others will arrive. For every assignment $e \in E$ and time $t \in [T]$, $w_{e,t} = p_{e,t} = 1$. We can verify that (1) an optimal solution to LP (1) is as follows: $x_{(i,j),t}^* = 1$ if $i = j = t \in [m]$ and 0 otherwise; (2) the optimal LP value and the performance of a clairvoyant optimal are both m . Note that $\text{ATT}(\gamma)$ works as follows: during each round $t \in [T]$ when $j = t$ arrives, $\text{ATT}(\gamma)$ selects the assignment $e = (i = t, j = t)$ with probability γ since $\beta_{i=t,t} = 1$.

Let H_i be the number of successful assignments on $i \in I$. We claim that $H_i \sim \text{Ber}(\gamma)$, and thus, $\text{Var}[H_i] = \gamma \cdot (1 - \gamma)$. Observe that $\{H_i\}$ are all independent. Therefore,

$$\text{Var}[H] = \sum_{i \in I} \text{Var}[H_i] = \gamma \cdot (1 - \gamma) \cdot m = \gamma \cdot (1 - \gamma) \cdot B,$$

$$\mathbb{E}[\text{ATT}(\gamma)] = \sum_{i \in I} \mathbb{E}[H_i] = \gamma \cdot m.$$

□

3 AN LP-BASED SAMPLING POLICY WITHOUT ATTENUATIONS (SAMP)

In this section, we present another LP-based sampling policy but without attenuations, which is formally stated in Algorithm 2.

Algorithm 2: An LP-based sampling policy SAMP(γ) for MP-KAD with $\gamma \in [0, 1]$.

- 1 **Offline Phase:**
 - 2 Solve LP (1) and let $\{x_{f,t}\}$ be an optimal solution.
 - 3 **Online Phase:**
 - 4 **for** $t = 1, \dots, T$ **do**
 - 5 Let an online agent (of type) j arrive at time t .
 - 6 Sample an assignment $f = (i, j, k) \in \mathcal{F}_j$ with probability $\gamma \cdot x_{f,t}/q_{j,t}$.
 - 7 /* At most one assignment will be sampled in Step (6) since $\sum_{f \in \mathcal{F}_j} \gamma \cdot x_{f,t}/q_{j,t} \leq \sum_{f \in \mathcal{F}_j} x_{f,t}/q_{j,t} \leq 1$ by Const. (2) of LP (1). */
 - 8 **if** i is safe at t (i.e., the unit capacity remains), **then**
 - 9 Select the assignment f (i.e., assigning j to i and charging j a price of a_k);
 - 10 **else**
 - 11 Reject f .
-

3.1 Competitive Analysis of SAMP(γ)

THEOREM 5. SAMP(γ) with $\gamma \in [0, 1]$ achieves a competitive ratio of $\gamma \cdot (1 - \gamma)$ for MP-KAD.

PROOF. Consider a given \bar{i} and a given time $\bar{t} \in [T]$. We show that \bar{i} is safe at (the beginning of) \bar{t} with a probability of at least $1 - \gamma$. Observe that since \bar{i} has a unit capacity, \bar{i} is safe at \bar{t} iff no successful assignment $f \in \mathcal{F}_{\bar{i}}$ has ever been made before \bar{t} .

For each $t < \bar{t}$, let $W_t = \sum_{f=(\bar{i},j,*) \in \mathcal{F}_{\bar{i}}} X_{j,t} \cdot Y_{f,t} \cdot Z_{f,t}$, where $X_{j,t} \sim \text{Ber}(q_{j,t})$, $Y_{f,t} \sim \text{Ber}(\gamma \cdot x_{f,t}/q_{j,t})$, and $Z_{f,t} \sim \text{Ber}(p_{f,t})$ are all Bernoulli random variables simulating the three independent random events, which are j arrives at t ($X_{j,t} = 1$), f gets sampled at t in SAMP ($Y_{f,t} = 1$), and j accepts f at t ($Z_{f,t} = 1$), respectively. Observe that (1) $W_t \in \{0, 1\}$ since at any time t , there is at most one single online agent arriving, denoted by j , and at most one assignment $f \in \mathcal{F}_j$ gets sampled in Step (6) of SAMP, and (2) $\{W_t | 1 \leq t < \bar{t}\}$ are all independent. Let $\text{SF}_{\bar{i}, \bar{t}} = 1$ indicate that \bar{i} is

safe at \bar{t} . We see that \bar{i} is safe at \bar{t} iff $W_t = 0$ for all $t < \bar{t}$. Thus,

$$\begin{aligned} \mathbb{E}[\text{SF}_{\bar{i}, \bar{t}}] &= \Pr \left[\bigwedge_{t < \bar{t}} (W_t = 0) \right] = \prod_{t < \bar{t}} \Pr[W_t = 0] \\ &\quad \text{(by independence of } \{W_t | t < \bar{t}\} \text{)} \\ &= \prod_{t < \bar{t}} (1 - \mathbb{E}[W_t = 1]) = \prod_{t < \bar{t}} \left(1 - \mathbb{E} \left[\sum_{f=(\bar{i},j,*) \in \mathcal{F}_{\bar{i}}} X_{j,t} \cdot Y_{f,t} \cdot Z_{f,t} \right] \right) \\ &= \prod_{t < \bar{t}} \left(1 - \sum_{f=(\bar{i},j,*) \in \mathcal{F}_{\bar{i}}} \mathbb{E}[X_{j,t}] \cdot \mathbb{E}[Y_{f,t}] \cdot \mathbb{E}[Z_{f,t}] \right) \\ &\quad \text{(by independence of } \{X_{j,t}, Y_{f,t}, Z_{f,t}\} \text{)} \\ &= \prod_{t < \bar{t}} \left(1 - \sum_{f=(\bar{i},j,*) \in \mathcal{F}_{\bar{i}}} q_{j,t} \cdot (\gamma \cdot x_{f,t}/q_{j,t}) \cdot p_{f,t} \right) \\ &= \prod_{t < \bar{t}} \left(1 - \sum_{f=(\bar{i},j,*) \in \mathcal{F}_{\bar{i}}} \gamma \cdot x_{f,t} \cdot p_{f,t} \right) \\ &\geq 1 - \sum_{t < \bar{t}} \sum_{f=(\bar{i},j,*) \in \mathcal{F}_{\bar{i}}} \gamma \cdot x_{f,t} \cdot p_{f,t} \geq 1 - \gamma. \end{aligned} \quad (7)$$

(by Constraint (3) of LP (1))

The analysis above suggests that for any given assignment $\bar{f} = (\bar{i}, \bar{j}, *) \in \mathcal{F}_{\bar{i}}$, it is successfully made at \bar{t} by SAMP(γ), denoted by $\chi_{\bar{f}, \bar{t}} = 1$, with probability of at least

$$\mathbb{E}[\chi_{\bar{f}, \bar{t}}] = \mathbb{E} \left[X_{\bar{j}, \bar{t}} \cdot Y_{\bar{f}, \bar{t}} \cdot \text{SF}_{\bar{i}, \bar{t}} \cdot Z_{\bar{f}, \bar{t}} \right] \geq (1 - \gamma) \cdot \gamma \cdot x_{\bar{f}, \bar{t}} \cdot p_{\bar{f}, \bar{t}}.$$

Therefore, we claim that the expected amount of profit gained by SAMP(γ) should be at least

$$\begin{aligned} \mathbb{E}[\text{SAMP}(\gamma)] &= \sum_{f \in \mathcal{F}} \sum_{t \in [T]} w_{f,t} \cdot \mathbb{E}[\chi_{f,t}] \geq (1 - \gamma) \cdot \gamma \cdot x_{f,t} \cdot p_{f,t} \cdot w_{f,t} \\ &= (1 - \gamma) \cdot \gamma \cdot \text{OPT-LP} \geq (1 - \gamma) \cdot \gamma \cdot \text{OPT-OFF}, \end{aligned}$$

where OPT-LP refers to the optimal value of LP-(1) and OPT-OFF the performance of a clairvoyant optimal, and the last inequality is due to Lemma 1. We establish the claim that SAMP(γ) achieves a competitive ratio of at least $\gamma \cdot (1 - \gamma)$. □

3.2 Tightness of the Competitive Analysis of SAMP(γ) for any $\gamma \in [0, 1]$

LEMMA 4. There exists an instance of MP-KAD such that SAMP(γ) achieves a competitive ratio of $\gamma \cdot (1 - \gamma)$ for any $\gamma \in [0, 1]$ regardless of the benchmark LP.

PROOF. Consider the instance as shown in Figure 3, which is almost the same as that in Figure 1 except that the profit on f_3 is $w_{f_3,t} = 1/\epsilon^2$ instead of $1/\epsilon$ for both $t = 1, 2$. We can verify the following facts. First, LP (1) has such an optimal solution that $x_{f_1,t=1}^* = 1 - \epsilon$, $x_{f_3,t=2}^* = \epsilon$ with all rest being zeros, and the corresponding optimal value is $\text{OPT-LP} = 1/\epsilon + 1 - \epsilon$. The performance of a clairvoyant optimal is $\text{OPT-OFF} = 1/\epsilon + 1 - \epsilon$. Second, SAMP(γ) with $\gamma \in [0, 1]$ samples f_1 with probability $\gamma(1 - \epsilon)$ when $j = 1$ arrives at $t = 1$, and it samples f_3 with probability γ if $j = 3$ arrives at $t = 2$. As a result, the probability that f_1 is successfully made is $\gamma(1 - \epsilon)$ and that of f_3 is $\epsilon \cdot \gamma \cdot \mathbb{E}[\text{SF}_{\bar{i}, \bar{t}}] = \epsilon \cdot \gamma \cdot (1 - \gamma \cdot (1 - \epsilon)) =$

$\epsilon \cdot \gamma \cdot (1 - \gamma + \gamma \cdot \epsilon)$. This suggests SAMP(γ) gains an expected amount of profit of $\gamma(1 - \epsilon) + (1/\epsilon) \cdot \gamma \cdot (1 - \gamma + \gamma \cdot \epsilon)$. Thus, we claim that SAMP(γ) achieves a CR of $\gamma(1 - \gamma)$ when $\epsilon \rightarrow 0_+$ based on the benchmark of either LP (1) or the clairvoyant optimal for any $\gamma \in [0, 1]$. \square




Figure 3: An example highlighting the unconditional tightness of competitive analysis of SAMP(γ) for any $\gamma \in [0, 1]$, which is irrespective of the benchmark LP.

3.3 Variance Analysis of SAMP(γ)

Recall that for each $f \in \mathcal{F}$ and $t \in [T]$, $\chi_{f,t} = 1$ indicates f is successfully made (scheduled and accepted) in SAMP(γ) with $\gamma \in [0, 1]$. Let $H_i = \sum_{f \in \mathcal{F}_i} \sum_{t \in [T]} \chi_{f,t}$ and $H = \sum_{i \in I} H_i$ denote the numbers of successful assignments involving i and in total, respectively.

THEOREM 6. $\text{Var}[H] \leq \bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B$, where $\bar{\gamma} = \min(1/2, \gamma)$ and $B = \sum_{i \in I} b_i$ with $\gamma \in [0, 1]$.

PROOF. For each $f = (i, j, k)$, let $X_{j,t} = 1$, $Y_{f,t} = 1$, and $Z_{f,t} = 1$ indicate that j arrives at t , f gets sampled at t in SAMP(γ), and j accepts f at t , respectively. For each f and t , let $W_{f,t} = X_{j,t} \cdot Y_{f,t} \cdot Z_{f,t}$. Observe that $\mathbb{E}[W_{f,t}] = \gamma \cdot x_{f,t} \cdot p_{f,t}$. For each i and t , let $W_{i,t} = \sum_{f \in \mathcal{F}_i} W_{f,t}$ and $W_i = \sum_{t \in [T]} W_{i,t}$. Following analyses in the proof of Theorem 5, we have (1) $\{W_{i,t} | t \in [T]\}$ are independent for each given $i \in I$, and each $W_{i,t} \in \{0, 1\}$ with $\mathbb{E}[W_{i,t}] = \sum_{f \in \mathcal{F}_i} \gamma \cdot x_{f,t} \cdot p_{f,t}$; and (2) i is not matched in the end, i.e., $H_i = 0$, iff $W_i = 0$. Thus, we see that $H_i = \min(W_i, b_i) = \min(W_i, 1)$. Following similar analyses in the proof of Theorem 4, we can show that $\{W_{f,t} | f \in \mathcal{F}, t \in [T]\}$ are negatively associated, and so are $\{H_i | i \in I\}$. Therefore, $\text{Var}[H] = \text{Var}[\sum_{i \in I} H_i] \leq \sum_{i \in I} \text{Var}[H_i]$.

Consider a given $i \in I$. Observe that i is not matched is equivalent to that i stays safe to the end of T , denoted by $\text{SF}_{i, T+1}$. By Inequality (7),

$$\mathbb{E}[H_i] = 1 - \Pr[H_i = 0] = 1 - \mathbb{E}[\text{SF}_{i, T+1}] \leq \gamma,$$

which suggests that $\text{Var}[H_i] \leq \bar{\gamma} \cdot (1 - \bar{\gamma})$ with $\bar{\gamma} = \min(1/2, \gamma)$. Consequently, we have $\text{Var}[H] = \text{Var}[\sum_{i \in I} H_i] \leq \sum_{i \in I} \text{Var}[H_i] \leq \bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B$. \square

3.4 Tightness of the Variance Analysis of SAMP(γ) for any $\gamma \in [0, 1]$


LEMMA 5. *There exists an instance of MP-KAD such that SAMP(γ) achieves a variance of $\bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B$ on the (random) number of successful assignments for any $\gamma \in [0, 1]$, where $\bar{\gamma} = \min(1/2, \gamma)$.*

PROOF. Consider the instance as shown in Figure 4, which is almost the same as that in Figure 2 except that $p_{f,t} = \min(1, (1/2)/\gamma)$ instead of 1 for all f and t . SAMP(γ) works as follows: during each round $t \in [T]$ when $j = t$ arrives, SAMP(γ) selects the assignment $f = (i = t, j = t)$ with probability γ , which suggests f is successfully made with probability $\gamma \cdot p_{f,t} = \bar{\gamma}$.

Let H_i be the number of successful assignments on $i \in I$. We claim that $H_i \sim \text{Ber}(\bar{\gamma})$, and thus, $\text{Var}[H_i] = \bar{\gamma} \cdot (1 - \bar{\gamma})$. Observe that $\{H_i\}$ are all independent, therefore,

$$\begin{aligned} \text{Var}[H] &= \text{Var}\left[\sum_{i \in I} H_i\right] = \sum_{i \in I} \text{Var}[H_i] \\ &= \bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot m = \bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B. \end{aligned}$$

\square



$$|I| = |J| = T = m, p = \min(1, (1/2)/\gamma),$$

$$q_{j=t,t} = 1, q_{j \neq t,t} = 0, \forall t \in [T],$$

$$x_{i=t,j=t,t}^* = 1, \forall t \in [T],$$

$$H = \sum_{i \in I} H_i, H_i = \text{Ber}(\bar{\gamma}), \forall i \in I,$$

$$\text{Var}[H] = \bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot m = \bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B.$$

Figure 4: An example highlighting the tightness of variance analysis of SAMP(γ) for any $\gamma \in [0, 1]$.

4 CONCLUSION

In this paper, we study matching and pricing simultaneously emerging in various gig platforms. We focus on two fundamental LP-based sampling algorithms and provide tight competitive and variance analyses for each of them. Our research opens a few directions. The immediate one is to extrapolate the current variance-analysis techniques to more general settings such as a weighted objective, i.e., the variance of the total profit instead of the total number of assignments, and less strict arriving assumptions, e.g., random arrival order and adversary. We expect more technical challenges there, which perhaps require us to add extra assumptions to make the problem tractable.

REFERENCES

- [1] Marek Adamczyk, Fabrizio Grandoni, and Joydeep Mukherjee. 2015. Improved approximation algorithms for stochastic matching. In *Algorithms-ESA 2015*. Springer, 1–12.
- [2] J. Angrist, Sydney Caldwell, and Jonathan V. Hall. 2017. Uber vs. Taxi: A Driver’s Eye View. *Randomized Social Experiments eJournal* (2017).
- [3] Eric Balkanski and Jason D. Hartline. 2016. Bayesian Budget Feasibility with Posted Pricing. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*. ACM, 189–203.
- [4] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. 2017. Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation* (Cambridge, Massachusetts, USA) (EC ’17). ACM, New York, NY, USA, 517–517. <https://doi.org/10.1145/3033274.3085099>
- [5] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. 2016. Dynamic Pricing in Ridesharing Platforms. *SIGecom Exch.* 15, 1 (Sept. 2016), 65–70. <https://doi.org/10.1145/2994501.2994505>
- [6] Kostas Bimpikis, Ozan Candogan, and Daniela Saban. 2019. Spatial pricing in ride-sharing networks. *Operations Research* (2019).
- [7] Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2020. Attenuate locally, win globally: An attenuation-based framework for online stochastic matching with timeouts. *Algorithmica* 82, 1 (2020), 64–87.
- [8] Haipeng Chen, Yan Jiao, Zhiwei Qin, Xiaocheng Tang, Hao Li, Bo An, Hongtu Zhu, and Jieping Ye. 2019. Inbede: integrating contextual bandit with td learning for joint pricing and dispatch of ride-hailing platforms. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 61–70.
- [9] Wang Chi Cheung, Will Ma, David Simchi-Levi, and Xinshang Wang. 2018. Inventory balancing with online learning. *arXiv preprint arXiv:1810.05640* (2018).
- [10] Michael B. Cohen, Y. Lee, and Zhao Song. 2019. Solving linear programs in the current matrix multiplication time. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019).
- [11] John Dickerson, Karthik Sankararaman, Kanthi Sarpatwar, Aravind Srinivasan, Kung-Lu Wu, and Pan Xu. 2019. Online resource allocation with matching constraints. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [12] John Dickerson, Karthik Sankararaman, Aravind Srinivasan, and Pan Xu. 2020. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. *ACM Transactions on Economics and Computation* (2020).
- [13] John P. Dickerson, Karthik Abinav Sankararaman, Kanthi Kiran Sarpatwar, Aravind Srinivasan, Kun-Lung Wu, and Pan Xu. 2019. Online Resource Allocation with Matching Constraints. In *AAMAS’19*. 1681–1689.
- [14] John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2018. Assigning Tasks to Workers based on Historical Data: Online Task Assignment with Two-sided Arrivals. In *AAMAS ’18*. 318–326.
- [15] Devdatt P Dubhashi and Desh Ranjan. 1996. Balls and bins: A study in negative dependence. *BRICS Report Series* 3, 25 (1996).
- [16] Seyed Esmaeili, Sharmila Dupplala, Davidson Cheng, Vedant Nanda, Aravind Srinivasan, and John P. Dickerson. 2023. Rawlsian fairness in online bipartite matching: Two-sided, group, and individual. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 5624–5632.
- [17] Zhixuan Fang, Longbo Huang, and Adam Wierman. 2017. Prices and Subsidies in the Sharing Economy. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*. ACM, 53–62.
- [18] Elaheh Fata, Will Ma, and David Simchi-Levi. 2019. Multi-stage and multi-customer assortment optimization with inventory constraints. *Available at SSRN 3443109* (2019).
- [19] Jiekun Feng, Mark Gluzman, and Jim G Dai. 2021. Scalable deep reinforcement learning for ride-hailing. In *2021 American Control Conference (ACC)*. IEEE, 3743–3748.
- [20] Yiding Feng, Rad Niazadeh, and Amin Saberi. 2019. Linear programming based online policies for real-time assortment of reusable resources. *Available at SSRN 3421227* (2019).
- [21] Kumar Joag-Dev and Frank Proschan. 1983. Negative association of random variables with applications. *The Annals of Statistics* (1983), 286–295.
- [22] Yash Kanoria and Pengyu Qian. 2019. Near Optimal Control of a Ride-Hailing Platform via Mirror Backpressure. *arXiv preprint arXiv:1903.02764* (2019).
- [23] Nixie S Lesmana, Xuan Zhang, and Xiaohui Bei. 2019. Balancing efficiency and fairness in on-demand ridesourcing. In *Advances in Neural Information Processing Systems*. 5310–5320.
- [24] Minne Li, Zhiwei (Tony) Qin, Yan Jiao, Yaodong Yang, Jun Wang, Chenxi Wang, Guobin Wu, and Jieping Ye. 2019. Efficient Ridesharing Order Dispatching with Mean Field Multi-Agent Reinforcement Learning. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*. ACM, 983–994.
- [25] Meghna Lowalekar, Pradeep Varakantham, and Patrick Jaillet. 2018. Online spatio-temporal matching in stochastic and dynamic domains. *Artificial Intelligence* 261 (2018), 71 – 112. <https://doi.org/10.1016/j.artint.2018.04.005>
- [26] Hongyao Ma, Fei Fang, and David C Parkes. 2019. Spatio-temporal pricing for ridesharing platforms. In *ACM EC ’19*.
- [27] Will Ma. 2018. Improvements and generalizations of stochastic knapsack and markovian bandits approximation algorithms. *Mathematics of Operations Research* 43, 3 (2018), 789–812.
- [28] Sushil Mahavir Varma, Pornpawee Bumpensanti, Siva Theja Maguluri, and He Wang. 2020. Dynamic pricing and matching for two-sided queues. *ACM SIGMETRICS Performance Evaluation Review* 48, 1 (2020), 105–106.
- [29] Vedant Nanda, Pan Xu, Karthik Abhinav Sankararaman, John Dickerson, and Aravind Srinivasan. 2020. Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 2210–2217.
- [30] Takuma Oda. 2021. Equilibrium Inverse Reinforcement Learning for Ride-hailing Vehicle Network. In *WWW ’21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*. ACM / IW3C2, 2281–2290.
- [31] Erhun Ozkan and Amy R Ward. 2020. Dynamic matching for real-time ride sharing. *Stochastic Systems* 10, 1 (2020), 29–70.
- [32] Sanket Shah, Meghna Lowalekar, and Pradeep Varakantham. 2022. Joint Pricing and Matching for City-Scale Ride-Pooling. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 32. 499–507.
- [33] Qi-Man Shao. 2000. A comparison theorem on moment inequalities between negatively associated and independent random variables. *Journal of Theoretical Probability* 13, 2 (2000), 343–356.
- [34] Dingyuan Shi, Yongxin Tong, Zimu Zhou, Bingchen Song, Weifeng Lv, and Qiang Yang. 2021. Learning to Assign: Towards Fair Task Assignment in Large-Scale Ride Hailing. In *ACM SIGKDD ’21*. 3549–3557.
- [35] Kirill Solovev and Nicolas Pröllochs. 2021. Integrating Floor Plans into Hedonic Models for Rent Price Appraisal. In *WWW ’21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*. ACM / IW3C2, 2838–2847.
- [36] Yi Su, Magd Bayoumi, and Thorsten Joachims. 2022. Optimizing Rankings for Recommendation in Matching Markets. In *WWW ’22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*. ACM, 328–338.
- [37] Yongxin Tong, Jieying She, Bolin Ding, Libin Wang, and Lei Chen. 2016. Online mobile micro-task allocation in spatial crowdsourcing. In *32nd IEEE International Conference on Data Engineering (ICDE ’16)*. 49–60.
- [38] Yongxin Tong, Libin Wang, Zimu Zhou, Lei Chen, Bowen Du, and Jieping Ye. 2018. Dynamic pricing in spatial crowdsourcing: A matching-based approach. In *Proceedings of the 2018 International Conference on Management of Data*. 773–788.
- [39] Alberto Vera, Siddhartha Banerjee, and Itai Gurvich. 2021. Online allocation and pricing: Constant regret via bellman inequalities. *Operations Research* 69, 3 (2021), 821–840.
- [40] Ruoyu Wu, Wei Bao, and Liming Ge. 2023. Online Task Assignment with Controllable Processing Time. *arXiv preprint arXiv:2305.04453* (2023).
- [41] Pan Xu. 2022. Exploring the Tradeoff between Competitive Ratio and Variance in Online-Matching Markets. *arXiv preprint arXiv:2209.07580* (2022).
- [42] Haoran Yu, Ermin Wei, and R. Berry. 2020. Learning to price vehicle service with unknown demand. *Proceedings of the Twenty-First International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing* (2020).

A EXPERIMENTAL RESULTS

Table 3: Parameter settings, where $|I|$ and $|J|$ denotes the numbers of driver and rider types, $[b, \bar{b}]$ denotes the range of capacities among all driver types, B is the total capacities of all driver types (*i.e.*, the total number of drivers over all types), T is the total number of online rounds, \mathcal{A} is set of basic prices per kilometer, and γ is the parameter in ATT and SAMP.

	Duration	$ I $	$ J $	$[b, \bar{b}]$	B	T	\mathcal{A}	γ
General Case	8:00-20:00	50	80	$\{[1,3],[1,7],[1,11],$ $[1,15],[1,19],[1,23]\}$	$\{100,200,\dots,600\}$	4200	$\{2.2,2.4,$ $\dots,3.2\}$	$\{0.1,0.2,\dots,0.5\}$
Special Case with Large Capacities				$\{[10,20],[20,30],[30,40],$ $[40,50],[50,60]\}$	$\{750,1250,\dots,2750\}$			$10 \times B$




Figure 5: Results for the general case on a real dataset offered by DiDi, Inc., collected in Haikou, China.




Figure 6: Results for the special case of large capacities on a real dataset offered by DiDi, Inc., collected in Haikou, China.

Data Preprocessing. We conduct our experiments on a real-world ride-hailing dataset offered by DiDi Inc., which is collected in Haikou, China. Each trip record consists of 23 columns, including pick-up and drop-off locations, the timestamp when the trip started and ended, and the estimated total fare for the order, *etc.* Following the setting of work [29], we set drivers offline agents while riders are online agents that arrive dynamically. We focus on the case where longitude and latitude range in $(110.18, 110.48)$ and $(19.90, 20.10)$, respectively, and we partition the area into $15 \times 10 = 150$ grids with equal size.

We construct the input instance as follows. Focus on the time window from 8:00 to 20:00 on May 5, 2017. We choose $T = 4200$ and split the total 12 hours in the window into T intervals such that each spans 30 seconds. For each grid, we create a driver type i , and sample a serving capacity b_i from $[b, \bar{b}]$ uniformly at random, which denotes the number of drivers of type i in the system. For each pair of starting

and ending grids, we create a rider type j , and set its arriving probability at time t (i.e., $q_{j,t}$) as the proportion of its record number among all rider types during t . We select the top 50 driver types and 80 rider types in terms of their number of records. For each pair of driver-rider types, we add an edge with probability 1 if the two share the same starting grid; with probability $1/4$ if they are adjacent to each other.

In our context of ride-hailing services, the groundset of prices applying to all ride types can be huge due to the large variance in trip length. Instead, we introduce the groundset of basic prices per kilometer (denoted by \mathcal{A} as well) and for each $a_k \in \mathcal{A}$, the assignment $f = (i, j, k)$ should be re-interpreted as assigning j to i and charge j at a price of a_k per kilometer (and thus, the total price charged to j should be equal to a_k multiplied by its trip length). For each given rider type j and time t , let $F_{j,t}$ be the CDF (Cumulative Distribution Function) of private value distribution of j over basic prices at t , i.e., $F_{j,t}(a_k)$ denotes the probability that the expected price (per kilometer) from j is no larger than a_k at time t . Commonly, the (random) private value is set to follow either a normal distribution [38, 42] or an exponential distribution [6, 26]. In our case, we consider a truncated normal distribution for each $F_{j,t}$, where the mean and variance are estimated based on samples of records relevant to rider type j at t . In this way, for each assignment $f = (i, j, k)$, we set the acceptance probability of j accepting f as $p_{f,t} = 1 - F_{j,t}(a_k)$, i.e., the probability that ride type j holds a private value larger or equal to the given price a_k . For each $f = (i, j, k)$, we set $w_{f,t} = \lambda \cdot a_k \cdot \kappa_j$, where λ refers to the royalty rate and κ_j is the trip length of rider type j . We set λ to 25% by default [2].

Algorithms. In addition to the two LP-based algorithms ATT and SAMP proposed in this paper, we implement several baselines as follows. Consider a rider of type j arrives at time t , and let $\mathcal{F}_{j,t}$ be the collection of *safe* assignments with respect to j at t (note that $\mathcal{F}_{j,t}$ may get reduced as time). Recall that $\{x_{f,t}\}$ is an optimal solution to the benchmark LP (1). If $\mathcal{F}_{j,t}$ is empty, then we have to reject j . Otherwise, (a) ATT-B (a boosted version of ATT): Samples an assignment $f = (i, j, k) \in \mathcal{F}_{j,t}$ with probability $\frac{x_{f,t}}{q_{j,t} \cdot \beta_{i,t}} / \sum_{f \in \mathcal{F}_{j,t}} \frac{x_{f,t}}{q_{j,t} \cdot \beta_{i,t}}$; (b) SAMP-B (a boosted version of SAMP): Sample an assignment $f = (i, j, k) \in \mathcal{F}_{j,t}$ with probability $\frac{x_{f,t}}{q_{j,t}} / \sum_{f \in \mathcal{F}_{j,t}} \frac{x_{f,t}}{q_{j,t}}$; (c) GRY: Select an assignment $f \in \mathcal{F}_{j,t}$ that maximizes $p_{f,t} \cdot w_{f,t}$ (and break ties arbitrarily); (d) UNM: Sample an assignment in $\mathcal{F}_{j,t}$ uniformly at random.

Computational Complexity of ATT and SAMP. Both ATT and SAMP consist of two parts: **Offline Phase** and **Online Phase**. As for **Offline Phase**, both ATT and SAMP need to solve the benchmark LP (1) that has $N := |E| \cdot K \cdot T$ variables. Thus, theoretically the running time on the part of solving LP (1) can be as low as $O^*(N^{2+1/6} \log(N/\delta))$ [10], where δ is the relative accuracy and $N = |E| \cdot K \cdot T$ with K and T being the total number of prices and online rounds, respectively. For SAMP, **Offline Phase** involves one extra procedure of computing all attenuation factors $\{\beta_{i,t}\}$ that takes another $O(T \cdot |E| \cdot K)$. As for **Online Phase**, both ATT and SAMP just need to sample an assignment from a one-dimensional vector with a size no larger than $|I| \cdot K$, which takes $O(|I| \cdot K)$ time. Thus, the dominant part of the running time will be solving the benchmark LP (1) in **Offline Phase**. Fortunately, all computations in **Offline Phase** can be done well before the online process starts.

Results and Discussions. For each instance, we run all algorithms for 100 times and take the average as the final performance (the total expected profits obtained). We compute the ratio of the performance of each algorithm to the optimal value of LP (1) as the final competitive ratio achieved. Additionally, we output the total number of successful assignments (that are scheduled and accepted) and the resulting variance.

Figure 5a shows the performance of $\text{ATT}(\gamma)$ is quite stable: its competitive ratios always stay slightly above and almost match the theoretical lower bounds of γ . This confirms our theoretical prediction in Theorem 1 and highlights the tightness of our competitive-ratio analysis. Though $\text{ATT}(\gamma)$ with $\gamma = 1/2$ proves optimal in the worst-case, its practical performance seems not as good as the boosted version ATT-B. This is mainly due to the fact that real-world instances deviate largely from the worst-case version. Figure 5a suggests that when the total number of arrivals of riders is fixed ($T = 4200$), heuristics like GRY and UNM are dominated by ATT-B for any given B , showing that ATT-B can work well in practice. In addition, $\text{ATT}(0.5)$ has a prominent advantage over the two heuristics when B is moderate and still slightly outperforms them even when B is extremely small or large. In contrast, Figure 5c suggests that the variance increases almost linearly as B for each fixed γ . These results align perfectly with our theoretical upper bounds in Theorem 1.

As for the special case with large capacities, Figure 6a shows that as γ increases, the competitive ratios of $\text{SAMP}(\gamma)$ will increase as well. This confirms our theoretical prediction in the first part of Theorem 2. Moreover, $\text{SAMP}(0.8)$ and $\text{SAMP}(1.0)$ always outperform GRY and UNM over different choices of B . Although GRY and UNM perform well in the metric of total successful assignments, they fail to achieve a good competitive ratio finally. This is mainly due to the myopia of heuristic-based strategies: they will run out the supplies very quickly and lose the opportunities to serve those profitable demands in the future. On the other hand, the LP-based algorithms will optimize the usage of supplies globally, which results in an extremely high competitive ratio overall. Figure 6c shows that the variance for all instances is upper bounded by $\bar{\gamma} \cdot (1 - \bar{\gamma}) \cdot B$, as suggested in the second part of Theorem 2.

B GRID PARTITION IN HAIKOU, CHINA.

In our experiments, we focus on the area where longitude and latitude range in $(110.18, 110.48)$ and $(19.90, 20.10)$, respectively, and we partition the area into $15 \times 10 = 150$ grids with equal size, as shown in Figure 7.




Figure 7: Grid partition in Haikou, China.

C JUSTIFICATION OF ASSUMPTIONS IN THE MODEL

First, we validate our KAD (Known Adversary Distributions) arrival assumption by plotting the distribution of the total number of arrivals for all rider types at various times; see Figure 8. The results demonstrate that the arrival patterns of each rider type do indeed change significantly over time. Notably, there is a peak-hour trend between 16:00 and 18:00 when most rider types experience their highest arrival rates.

Second, we validate the dependency of the acceptance probability on rider types and time by plotting the distribution of the number of trip records at different times for each given rider type and by plotting the distribution of the number of trip records for different rider types within each predefined time interval, as illustrated in Figure 9 and Figure 10. In Figure 10, it's evident that for a given time slot between 16:00 and 18:00, acceptance prices are partially concentrated around 3.0 per kilometer, but the distributions still exhibit significant variation across rider types. Figure 9 reveals that for rider type 4, the charged prices during the afternoon peak hours (from 16:00 to 18:00) are significantly higher than those for other time slots. This observation suggests an over-demand tendency during the afternoon peak hours.




Figure 8: The distribution of the total number of arrivals for all rider types at various times on May 5, 2017, which suggests a peak-hour trend between 16:00 and 18:00 when most rider types experience their highest arrival rates.

1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566

1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624





Figure 9: The number of trip records for rider type 4 from 8:00 to 20:00. The charged prices during the afternoon peak hours (from 16:00 to 18:00) are significantly higher than those for other time slots.

Figure 10: The number of trip records for all rider types between 16:00 and 18:00.