

Allocation Problems in Ride-Sharing Platforms: Online Matching with Offline Reusable Resources

JOHN P. DICKERSON, University of Maryland, College Park, USA

KARTHIK A. SANKARARAMAN, Facebook, USA

ARAVIND SRINIVASAN, University of Maryland, College Park, USA

PAN XU, New Jersey Institute of Technology, Newark, USA

Bipartite-matching markets pair agents on one side of a market with agents, items, or contracts on the opposing side. Prior work addresses online bipartite-matching markets, where agents arrive over time and are dynamically matched to a known set of disposable resources. In this paper, we propose a new model, *Online Matching with (offline) Reusable Resources under Known Adversarial Distributions* (OM-RR-KAD), in which resources on the offline side are *reusable* instead of disposable; that is, once matched, resources become available again at some point in the future. We show that our model is tractable by presenting an LP-based non-adaptive algorithm that achieves an online competitive ratio of $\frac{1}{2} - \epsilon$ for any given constant $\epsilon > 0$. We also show that no adaptive algorithm can achieve a ratio of $\frac{1}{2} + o(1)$ based on the same benchmark LP. Through a data-driven analysis on a massive openly-available dataset, we show our model is robust enough to capture the application of taxi dispatching services and ride-sharing systems. We also present heuristics that perform well in practice.

Additional Key Words and Phrases: Online-Matching, Ride-Sharing, Randomized Algorithms, Auction Design

ACM Reference Format:

John P. Dickerson, Karthik A. Sankararaman, Aravind Srinivasan, and Pan Xu. 2018. Allocation Problems in Ride-Sharing Platforms: Online Matching with Offline Reusable Resources. *ACM Transactions on Economics and Computation* 0, 0, Article 0 (January 2018), 17 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

In bipartite-matching problems, agents on one side of a market are paired with agents, contracts, or transactions on the other. Classical matching problems—assigning students to schools, papers to reviewers, or medical residents to hospitals—take place in a static setting, where all agents exist at the time of matching, are simultaneously matched, and then the market concludes. In contrast, many matching problems are dynamic, where one side of the market arrives in an *online* fashion and is matched sequentially to the other side.

Work done when KAS was a PhD student at University of Maryland, College Park. There, KAS supported in part by NSF Awards CNS 1010789 and CCF 1422569. PX was partially supported by NSF CRII Award IIS-1948157 and partially by NSF Awards CNS 1010789 and CCF 1422569. AS supported in part by NSF Awards CNS-1010789, CCF-1422569 and CCF-1749864, and by research awards from Adobe, Inc., Amazon, Inc., and Google Inc. JD was supported in part by NSF CAREER Award IIS-1846237, NIST MSE Award #20126334, DARPA GARD #HR00112020007, DARPA SI3-CMD #S4761, DoD WHS Award #HQ003420F0035, and a Google Faculty Research Award. A preliminary version of this work appeared in the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18) under the same title.

Authors' addresses: John P. Dickerson, University of Maryland, College Park, USA, john@cs.umd.edu; Karthik A. Sankararaman, Facebook, USA, karthikabinavs@gmail.com; Aravind Srinivasan, University of Maryland, College Park, USA, asriniv1@umd.edu; Pan Xu, New Jersey Institute of Technology, Newark, USA, pxu@njit.edu.

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 the author(s) 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.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

XXXX-XXXX/2018/1-ART0 \$15.00

<https://doi.org/0000001.0000001>

24 Online bipartite-matching problems are primarily motivated by Internet advertising. In the basic
25 version of the problem, we are given a bipartite graph $G = (U, V, E)$ where U represents the of-
26 fline vertices (advertisers) and V represents the online vertices (keywords or impressions). There
27 is an edge $e = (u, v)$ if advertiser u bids for a keyword v . When a keyword v arrives, a central
28 clearinghouse must make an instant and irrevocable decision to either reject v or assign v to one
29 of its “neighbors” u in G and obtain a profit w_e for the match $e = (u, v)$. When an advertiser u is
30 matched, it is no longer available for matches with other keywords (in the most basic case) or its
31 budget is reduced. The goal is to design an efficient online algorithm such that the expected total
32 weight (profit) of the matching obtained is maximized. Following the seminal work of Karp et al.
33 [29], there has been a large body of research on related variants (overviewed by Mehta [40]). One
34 particular flavor of problems is online-matching with known identical independent distributions
35 (OM-KIID) [14, 21, 26, 28, 36]. In this context, agents arrive over T rounds, and their arrival distri-
36 butions are assumed to be *identical and independent* over all T rounds; additionally, this distribution
37 is known to the algorithm beforehand.

38 Apart from the Internet-advertising application, online bipartite-matching models have been used
39 to capture a wide range of online resource allocation and scheduling problems. Typically we have
40 an offline and an online party representing, respectively, the service providers (SP) and online users;
41 once an online user arrives, we need to match it to an offline SP immediately. In many cases, the
42 service is *reusable* in the sense that once an SP is matched to a user, it will be gone for some time, but
43 will then rejoin the system afterwards. Besides that, in many real settings the arrival distributions of
44 online users do change from time to time (*i.e.*, they are not i.i.d.). Consider the following motivating
45 examples.

46 **Taxi Dispatching Services and Ride-Sharing Systems.** Traditional taxi services and ride-sharing
47 systems such as Uber and Didi Chuxing match drivers to would-be riders [30, 31, 50, 55]. Here, the
48 offline SPs are different vehicle drivers. Once an online request (potential rider) arrives, the system
49 matches it to a nearby driver instantly such that the rider’s waiting time is minimized. In most cases,
50 the driver will rejoin the system and can be matched again once she finishes the service. Addition-
51 ally, the arrival rates of requests changes dramatically across the day. Consider the online arrivals
52 during peak hours and off-peak hours for example: the arrival rates in the former case can be much
53 larger than the latter. Though our model is primarily motivated by taxi-dispatching services and
54 ride-sharing platforms, we acknowledge that it does not address the full generality of these applica-
55 tions. In particular, our work focuses on the temporal components of rideshare; spatial components
56 such as match-specific driver movement are not modeled [17, 18, 33]. Since we make the simpli-
57 fying assumption that the stochastic process that determines the parameters in the environment is
58 independent of the algorithm’s decisions, this may not fully capture the scenario where requests
59 are heterogenous with vastly-different ride times. One potential direction to extend this model is to
60 consider correlated reusable rates, which makes the theory much more challenging.

61 **Organ Allocation.** Chronic kidney disease affects tens of millions of people worldwide at great
62 societal and monetary cost [43, 49]. Organ donation—either via a deceased or living donor—is a
63 lifesaving alternative to organ failure. In the case of kidneys, a donor organ can last up to 15 years
64 in a patient before failing again. Various nationwide organ donation systems exist and operate under
65 different ethical and logistical constraints [12, 20, 37], but all share a common online structure: the
66 offline party is the set of patients (who reappear every 5 to 15 years based on donor organ longevity),
67 and the online party is the set of donors or donor organs, who arrive over time. Similarly, in some
68 blood or plasma donation settings, donors might reappear after some number of weeks.

69 Similar scenarios can be seen in other areas such as wireless network connection management
70 (SPs are different wireless access points) [59] and online cloud computing service scheduling [42,

60]. Inspired by the above applications, we generalize the model of OM-KIID in the following two ways.

Reusable Resources. Once we assign v to u , u will rejoin the system after C_e rounds with $e = (u, v)$, where $C_e \in \{0, 1, \dots, T\}$ is an integral random variable with known distribution. In this paper, we call C_e the *occupation time* of u w.r.t. e . In fact, we show that our setting can directly be extended to the case when C_e is time sensitive: when matching v to u at time t , u will rejoin the system after $C_{e,t}$ rounds. This extension makes our model adaptive to nuances in real-world settings. For example, consider the taxi dispatching or ride-sharing service: the occupation time of a driver u from a matching with an online user v does depend on both the user type of v (such as destination) and the time when the matching occurs (peak hours can differ significantly from off-peak hours).

Known Adversarial Distributions (KAD). Suppose we have T rounds and that for each round $t \in [T]^1$, a vertex v is sampled from V according to an arbitrary known distribution \mathcal{D} where the marginal for v is $\{p_{v,t}\}$ such that $\sum_{v \in V} p_{v,t} \leq 1$ for all t . Also, the arrivals at different times are independent (and according to these given distributions). The setting of KAD was introduced by [4, 5] and is known as Prophet Inequality matching.

We call our new model Online Matching with (offline) Reusable Resources under Known Adversarial Distributions (OM-RR-KAD, henceforth). Note that the OM-KIID model can be viewed as a special case when C_e is a constant (with respect to T) and $\{p_{v,t} | v \in V\}$ are the same for all $t \in [T]$.

Competitive Ratio. Let $E[\text{ALG}(\mathcal{I}, \mathcal{D})]$ denote the expected value obtained by an algorithm ALG on an input \mathcal{I} and arrival distribution \mathcal{D} . Let $E[\text{OPT}(\mathcal{I})]$ denote the expected *offline optimal value*, which refers to the optimal solution when we are allowed to make choices after observing the entire sequence of online arrival vertices. Then, the competitive ratio is defined as $\min_{\mathcal{I}, \mathcal{D}} \frac{E[\text{ALG}(\mathcal{I}, \mathcal{D})]}{E[\text{OPT}(\mathcal{I})]}$. It is a common technique to use an LP optimal value to upper bound $E[\text{OPT}(\mathcal{I})]$ (called the benchmark LP) and hence get a valid lower bound on the resulting competitive ratio.

Adaptive vs. non-adaptive algorithms. An online algorithm ALG is called *non-adaptive* if the strategy for all time-steps t , denoted by $\text{ALG}(t)$, is pre-computed *before* the realizations of the online process. In contrast, adaptive algorithms can choose the strategy for time-step t *after* seeing the random realizations of all the random processes (*e.g.*, random seeds used in the algorithm, arrivals of online requests and occupation times of drivers) in steps $1, 2, \dots, t - 1$.

100 1.1 Our Contributions

101 First, we propose the new model of OM-RR-KAD to capture a wide range of real-world applica-
 102 tions related to online scheduling, organ allocation, rideshare dispatch, among others. We claim that
 103 this model is tractable enough to obtain good algorithms with theoretically provable guarantees and
 104 general enough to capture many real-life instances. Our model assumptions take a significant step
 105 forward from the usual assumptions in the online-matching literature where the offline side is as-
 106 sumed to be *single-use* or *disposable*. This leads to a larger range of potential applications which
 107 can be modeled by online-matching. The first part of the paper focusses on the abstract model and
 108 though the experiments are motivated using ride-share, we do not get into the application-specific
 109 assumptions/issues one may run into. For specific discussion on issues that arise when applied to
 110 ride-share, we defer to Section 5.

111 Second, we show how this model can be *cleanly* analyzed under a theoretical framework. We
 112 first construct a linear program (LP henceforth) LP (1) which we show is a valid upper bound on
 113 the expected offline optimal value (note that the latter is hard to characterize). Next, we propose

¹Throughout this paper, we use $[N]$ to denote the set $\{1, 2, \dots, N\}$, for any positive integer N .

114 an efficient *non-adaptive* algorithm that achieves a competitive ratio of $\frac{1}{2} - \epsilon$ for any given constant $\epsilon > 0$. This algorithm solves the LP and obtains an optimal fractional solution. It uses this
 115 optimal solution as a guide in the online phase. Using Monte-Carlo simulations (called simulations
 116 henceforth), and combining with this optimal solution, our algorithm makes the online decisions. In
 117 particular, Theorem 1 describes our first theoretical results formally.
 118

119 **THEOREM 1.** *LP (1) is a valid benchmark for OM-RR-KAD. There exists a non-adaptive on-*
 120 *line algorithm that achieves an online competitive ratio of $\frac{1}{2} - \epsilon$ for any given $\epsilon > 0$ against the*
 121 *benchmark LP (1).*

122 Third, we show that the online-competitive analysis of the non-adaptive algorithm is *tight* with
 123 respect to the choice of our benchmark LP. Specifically, we show that when restricted to LP (1), no
 124 adaptive algorithm can achieve a competitive ratio better than $\frac{1}{2}$ (Theorem 2), and no non-adaptive
 125 algorithm can beat $\frac{1}{2}$ even when all C_e are deterministic and equal (Theorem 3).

126 **THEOREM 2.** *No adaptive algorithm can achieve a competitive ratio better than $\frac{1}{2} + o(1)$ against*
 127 *the benchmark LP (1). Here, $o(1)$ is a vanishing term when T is sufficiently large.*

128 **THEOREM 3.** *No non-adaptive algorithm can achieve a competitive ratio better than $\frac{1}{2} + o(1)$*
 129 *against the benchmark LP (1) even when all C_e are deterministic and equal. Here, $o(1)$ is a vanishing*
 130 *term when both of C_e and T/C_e are sufficiently large.*

131 Finally, through a data-driven analysis on a massive openly-available dataset we show that our
 132 model is robust enough to capture the setting of taxi hailing/sharing at least. Additionally, we pro-
 133 vide certain *simpler* heuristics which also give good performance. Hence, we can combine these
 134 theoretically grounded algorithms with such heuristics to obtain further improved ratios in practice.
 135 Section 5 provides a detailed qualitative and quantitative discussion.

136 1.2 Other Related Work

137 In addition to the arrival assumptions of KIID and KAD, there are several other important, well-
 138 studied variants of online-matching problems. Under *adversarial ordering*, an adversary can arrange
 139 the arrival order of all items in an arbitrary way (e.g., online-matching [29, 53] and AdWords [16,
 140 41]). Under a *random arrival order*, all items arrive in a random permutation order (e.g., online-
 141 matching [35] and AdWords [24]). Finally, under *unknown distributions*, in each round, an item is
 142 sampled from a fixed but unknown distribution. (e.g., [19]). For each of the categories above, we
 143 list only a few examples considered under that setting. For a more complete list, please refer to the
 144 book by Mehta [40].

145 Despite the fact that our model is inspired by online bipartite-matching, it also overlaps with
 146 stochastic online scheduling problems (SOS) [38, 39, 52]. We first restate our model in the language
 147 of SOS: we have $|U|$ *nonidentical* parallel machines and $|V|$ jobs; at every time-step a single job v
 148 is sampled from V with probability $p_{v,t}$; the jobs have to be assigned immediately after its arrival
 149 (or rejected right away); additionally each job v can be processed *non-preemptively* on a specific
 150 subset of machines; once we assign v to u , we get a profit of w_e and u will be occupied for C_e
 151 rounds with $e = (u, v)$, where C_e is a random variable with known distribution. Observe that the key
 152 difference between our model and SOS is in the objective: the former is to maximize the expected
 153 profit from the completed jobs, while the latter is to minimize the total or the maximum completion
 154 time of all jobs. In a recent concurrent work [22], they consider the dynamic assortment of reusable
 155 resources. In their context, the offline side is the set of reusable resources, while the online side is
 156 the set of consumers. They consider the arrival setting of KAD (which they call the Bayesian model
 157 with non-identical distributions). The critical difference is that they assume that in each round, on

the arrival of an online customer, the algorithm should assign her a set of offline resources, where each offline resource has a given budget. The benchmark LP has an exponential number of variables, which is not solvable in polynomial time. They overcome this by assuming an offline oracle, that returns an optimal solution to the benchmark LP for any given arrival sequence of online customers. They design a similar simulation-based online policy that achieve a competitive ratio of $1/2$.

Research in ridesharing platforms and similar allocation problems is an active area of research within multiple fields, including computer science, operations research and transportation engineering. State-independent policies were studied previously using theory from control and queuing systems [11, 13, 46]. The role of pricing in the dynamics of drivers in ridesharing platforms is also an active area of research in computational economics and AI/ML (e.g., [2, 10, 17, 33, 44, 47, 58]). Our problem is a form of *online-matching in dynamic environments*, which is an active area of research within the AI/ML community. In particular, [20, 31, 54, 55] have studied algorithms for matching in various dynamic bipartite markets such as kidney exchange, spatial crowdsourcing, labor markets, and so on. A similar line of work on general graphs is also prominent in the literature (e.g., [3, 6–8]).

2 MAIN MODEL

In this section, we present a formal statement of our main model. Suppose we have a bipartite graph $G = (U, V, E)$ where U and V represent the offline and online parties respectively. We have a finite time horizon T (known beforehand) and for each time $t \in [T]$, a vertex v will be sampled (we use the term *v arrives*) from a known probability distribution $\{p_{v,t}\}$ such that $\sum_{v \in V} p_{v,t} \leq 1^2$ (noting that such a choice is made independently for each round t). The expected number of times v arrives across the T rounds, $\sum_{t \in [T]} p_{v,t}$, is called the *arrival rate* for vertex v . Once a vertex v arrives, we need to make an *irrevocable decision immediately*: either to reject v or assign v to one of its neighbors in U . For each u , once it is assigned to some v , it becomes unavailable for C_e rounds with $e = (u, v)$, and subsequently rejoins the system. Here C_e is an integral random variable taking values from $\{0, 1, \dots, T\}$ and the distribution is known in advance. Each assignment e is associated with a weight w_e and our goal is to design an online assignment policy such that the total expected weights of all assignments made is maximized. Following prior work, we assume $|V| \gg |U|$ and $T \gg 1$. Throughout this paper, we use edge $e = (u, v)$ and assignment of v to u interchangeably.

For an assignment e , let $x_{e,t}$ be the probability that e is chosen at t in any offline optimal algorithm. For each u (likewise for v), let E_u (E_v) be the set of neighboring edges incident to u (v). We use the LP (1) as a benchmark to upper bound the offline optimal. We now interpret the constraints. For each round t , once an online vertex v arrives, we can assign it to at most one of its neighbors. Thus, we have: if v arrives at t , the total number of assignments for v at t is at most 1; if v does not arrive, the total is 0. The LHS of (2) is exactly the expected number of assignments made at t for v . It should be no more than the probability that v arrives at t , which is the RHS of (2). Constraint (3) is the *most* novel part of our problem formulation. Consider a given u and t . In the LHS, the first term (summation over $t' < t$ and $e \in E_u$) refers to the probability that u is not available at t while the second term (summation over $e \in E_u$) is the probability that u is assigned to some driver at t , which is no larger than probability u is available at t . Thus, the sum of the first term and second term on LHS is no larger than 1.³ This argument implies that the LP forms a valid upper-bound on the offline optimal solution and hence we have the first part of Lemma 4.

LEMMA 4. *The optimal value to LP (1) is a valid upper bound for the offline optimal. Moreover, suppose for some $\delta \geq 0$, we have an estimate $f(e, y)$ of $\Pr[C_e > y]$ for all edges e and $y \geq 0$, where*

²Thus, with probability $1 - \sum_{v \in V} p_{v,t}$, none of the vertices from V will arrive at t .

³We would like to point out that our LP constraint (3) on u is inspired by Ma [34]. The proof is similar to that by Alaei et al. [4] and Alaei et al. [5].

$$\text{maximize } \sum_{t \in [T]} \sum_{e \in E} w_e x_{e,t} \quad (1)$$

$$\text{subject to } \sum_{e \in E_v} x_{e,t} \leq p_{v,t} \quad \forall v \in V, t \in [T] \quad (2)$$

$$\sum_{t' < t} \sum_{e \in E_u} x_{e,t'} \Pr[C_e > t - t'] + \sum_{e \in E_u} x_{e,t} \leq 1 \quad \forall u \in U, t \in [T] \quad (3)$$

$$0 \leq x_{e,t} \leq 1 \quad \forall e \in E, t \in [T] \quad (4)$$

201 $f(e, y)/\Pr[C_e > y]$ always lies in $[1/(1 + \delta), 1 + \delta]$. Then, by using $f(e, t - t')$ in the LP instead
 202 of $\Pr[C_e > t - t']$ and scaling down the resultant vector \mathbf{x} by $(1 + \delta)$, we only get a further loss of
 203 $(1 + \delta)$ in the competitive ratio.

204 **PROOF.** Fix any offline optimal algorithm OPT. For each assignment $e = (u, v)$ and t , let $X_{e,t}$
 205 denote an indicator random variable for the event that e is matched at t in OPT, which includes the
 206 event that v arrives at time t . Let $\mathbb{E}[X_{e,t}] = x_{e,t}$. Therefore, by linearity of expectation, the expected
 207 performance of OPT is $\mathbb{E}[\text{OPT}] = \sum_t \sum_e w_e x_{e,t}$. Now we justify that the solution $\{x_{e,t}\}$ is feasible
 208 to all constraints in LP (1).

209 In constraint (2), the LHS denotes the probability that v is matched at t , which should be no larger
 210 than the probability that vertex v arrives at time t . In constraint (3), the first part
 211 $\sum_{t' < t} \sum_{e \in E_u} x_{e,t'} \Pr[C_e > t - t']$ denotes the probability that u is occupied due to assignments made
 212 prior to t while the second part $\sum_{e \in E_u} x_{e,t}$ is the probability that an assignment incident to u is made
 213 at time t . Thus, the sum of these two parts should be no larger than 1. Constraint (4) is satisfied since
 214 $\{x_{e,t}\}$ are all probability values. Therefore, we have shown that the values $\{x_{e,t}\}$ is feasible to all
 215 constraints in LP (1), which implies that the optimal value of LP (1) is a valid upper bound for the
 216 performance of any offline optimal.

217 The second part follows directly from the fact that \mathbf{x} is scaled down by a factor $(1 + \delta)$ and hence
 218 the objective is scaled down by a factor $(1 + \delta)$. \square

219 3 SIMULATION-BASED ALGORITHM

220 In this section, we present a simulation-based algorithm. We will first give a gentle introduction to
 221 simulation-based algorithms, that has been developed and used in prior works ([1] and [15]).

222 **Simulation-based algorithms.** We use the term *simulation* throughout this paper to refer to Monte
 223 Carlo simulation and the term *simulation-based attenuation* to refer to the simulation and attenua-
 224 tion techniques as shown in [1] and [15]⁴. At a high level, suppose we have a randomized algorithm
 225 such that for some event E (i.e., u is available at t) we have $\Pr[E] \geq c$, then we modify the algorithm
 226 as follows: (i) We first use simulation to estimate a value \hat{E} that lies in the range $[\Pr[E], (1 + \epsilon)\Pr[E]]$
 227 with probability at least $1 - \delta$. (ii) By “ignoring” E (i.e., *attenuation*, in a problem-specific manner)
 228 with probability $\sim 1 - c/\hat{E}$, we can ensure that the final effective value of $\Pr[E]$ is arbitrarily close
 229 to c , i.e., in the range $[c/(1 + \epsilon), c]$ with probability at least $1 - \delta$. This simple idea of attenuating the
 230 probability of an event to come down approximately to a certain value c is what we term simulation-
 231 based attenuation. The number of samples needed to obtain the estimate \hat{E} is $\Theta(\frac{1}{\epsilon^2} \cdot \log(\frac{1}{\delta}))$ via a
 232 standard Chernoff-bound argument. In our applications, we will take $\epsilon = 1/\text{poly}(N)$ where N is the
 233 problem-size, and the error ϵ will only impact lower-order terms in our approximations.

⁴This is called “dumping factor” in [1]. See Appendix B in [1] for a formal treatment.

234 We adapt the above simulation-based attenuation technique to our setting. Let \mathbf{x}^* denote an opti-
 235 mal solution to LP (1). Suppose we aim to develop an online algorithm achieving a ratio of $\gamma \in [0, 1]$.
 236 In particular, for every edge $e \in E$, we want to ensure that the $\Pr[e \text{ is matched}] \geq \gamma$. Consider an
 237 assignment $e = (u, v)$ when some v arrived at time t . Let $\text{SF}_{e,t}$ be the event that e is safe at t ,
 238 *i.e.*, u is available at t . Using Monte-Carlo simulations we obtain an approximate estimate $\beta_{e,t}$ of
 239 the quantity $\Pr[\text{SF}_{e,t}]$. Note that to obtain $\beta_{e,t}$, we need the estimates $\beta_{e,t'}$ for every $t' < t$. With the
 240 estimated quantity $\beta_{e,t}$, we match a safe edge e at time t with probability $\frac{x_{e,t}^* \gamma}{p_{v,t} \beta_{e,t}}$. This is a valid
 241 probability (*i.e.*, not exceeding 1), if and only if $\gamma \leq \beta_{e,t}$. We denote an algorithm $\text{ADAP}(\gamma)$ to be
 242 *valid* if and only if for every $e \in E$ and every $t \in [T]$ we have $\gamma \leq \beta_{e,t}$.

243 When the condition $\gamma \leq \beta_{e,t}$ is satisfied for all $e \in E$ and $t \in [T]$, we have that the probability an
 244 edge $e = (u, v)$ is matched conditioned on the event that v arrives and u is available is at least $\gamma x_{e,t}^*$.
 245 Thus, using linearity of expectation this immediately implies that the expected reward obtained by
 246 the algorithm is at least $\gamma \sum_{t \in [T]} \sum_{e \in E} w_e x_{e,t}^*$. Therefore, the competitive ratio is at least γ .

247 At the outset, this looks similar to the Inverse Propensity Scoring (IPS) used in the multi-armed
 248 bandit literature [9]. However, there is a key difference between IPS estimates and our estimates. In
 249 the bandit literature, one usually scales the value by the probability of playing an action, since this
 250 is the *cost* of observing only bandit feedback. However, here we scale by a quantity that depends
 251 on the probability of a certain event happening during the *run* of the algorithm, because of playing
 252 other actions. The linear program gives a distribution over the edges assuming that all the neighbors
 253 are available. Hence this scaling can be interpreted as the *cost* the algorithm needs to incur when
 254 some neighbors are already matched.

255 The simulation-based attenuation technique has been used previously for other problems, such as
 256 stochastic knapsack [34] and stochastic matching [1]. Throughout the analysis, we assume that we
 257 know the exact value of $\beta_{e,t} := \Pr[\text{SF}_{e,t}]$ for all t and e . (It is easy to see that the sampling error can
 258 be folded into a multiplicative factor of $(1 - \epsilon)$ in the competitive ratio by standard Chernoff bounds
 259 and hence, ignoring it leads to a cleaner presentation.). The formal statement of our algorithm,
 260 denoted by $\text{ADAP}(\gamma)$, is as follows. For each v and t , let $E_{v,t}$ be the set of *safe* assignments for v at
 261 t .

ALGORITHM 1: A simulation-based adaptive algorithm $\text{ADAP}(\gamma)$

For each time t , let v denote the request arriving at time t .

If $E_{v,t} = \emptyset$, then reject v ; otherwise choose $e \in E_{v,t}$ with probability $\frac{x_{e,t}^* \gamma}{p_{v,t} \beta_{e,t}}$ where $e = (u, v)$.

262 **Remarks.** Our simulation-based adaptive algorithm described above is *non-adaptive* according
 263 to our definition. Assume ADAP is valid with respect to a given $\gamma \in (0, 1)$. We can first solve
 264 the benchmark LP (1) and get an optimal solution $\{x_{e,t}^*\}_{e \in E, t \in [T]}$. Then by simulating the random
 265 arrivals of online agents according to known distributions $\{p_{v,t}\}_{v \in V, t \in [T]}$ and $\text{ADAP}(\gamma)$ itself se-
 266 quentially from $t = 1, 2, \dots, T$, we can get an arbitrarily accurate estimate of $\beta_{e,t}$ for each e and t .
 267 All these procedures can be done in an offline manner (*i.e.*, before the online process).

268 **LEMMA 5.** $\text{ADAP}(\gamma)$ is a valid algorithm (*i.e.*, $\gamma \leq \beta_{e,t}$ for every $e \in E$ and $t \in [T]$) when $\gamma = \frac{1}{2}$.

269 **PROOF.** Essentially we need to show that $\beta_{e,t} \geq \gamma = \frac{1}{2}$ for all e and t . We prove it by induction
 270 on t as follows.

271 When $t = 1$, $\beta_{e,t} = 1$ for all $e = (u, *)$. Therefore, we are done. Assume for all $t' < t$, $\beta_{e,t'} \geq 1/2$
 272 and $\text{ADAP}(\gamma)$ is valid for all rounds t' . In other words, we assume each e is assigned with probability
 273 *exactly* equal to $x_{e,t'}^* \cdot \frac{1}{2}$ for all $t' < t$. Now consider a given $e = (u, v)$. Observe that e is unsafe at t

274 iff u is assigned with some v' at $t' < t$ such that the assignment $e' = (u, v')$ makes u unavailable at
275 t . Therefore,

$$1 - \beta_{e,t} = 1 - \Pr[\text{SF}_{e,t}] = \sum_{t' < t} \sum_{e \in E_u} \frac{x_{e,t'}}{2} \Pr[C_e > t - t'] \leq \frac{1}{2}.$$

276 The last inequality above is due to Constraint (3) in the benchmark LP (1). Therefore, we have
277 $\beta_{e,t} \geq 1/2$ and we are done. \square

278 The main Theorem 1 follows directly from Lemmas 4 and 5.

279 **Extension from C_e to $C_{e,t}$.** Consider the case when the occupation time of u from e is sensitive to
280 t . In other words, each u will be unavailable for $C_{e,t}$ rounds from the assignment $e = (u, v)$ at t .
281 We can accommodate the extension by simply updating the constraints (3) on u in the benchmark
282 LP (1) to the following. We have that $\forall u \in U, t \in [T]$,

$$\sum_{t' < t} \sum_{e \in E_u} x_{e,t'} \Pr[C_{e,t'} > t - t'] + \sum_{e \in E_u} x_{e,t} \leq 1 \quad (5)$$

283 The rest of our algorithm remains the same as before. We can verify that (1) LP (1) with con-
284 straints (3) replaced by (5) is a valid benchmark; (2) ADAP achieves a competitive ratio of $\frac{1}{2} - \epsilon$ for
285 any given $\epsilon > 0$ for the new model based on the new valid benchmark LP. The modifications to the
286 analysis transfer through in a straightforward way and for brevity we omit the details here.

287 4 TIGHTNESS OF ONLINE ANALYSIS AGAINST THE BENCHMARK LP

288 In this section we prove the main result as stated in Theorem 2. Consider the following example.

289 **EXAMPLE 1.** Consider a star graph $G = (U, V, E)$, where U consists of one single node u and
290 $V = \{v_1, v_2\}$. Set $T = N + 1$. For $j = 1, 2$ and $t \in [T]$, let $p_{j,t}$ denote the arrival probability of the
291 vertex of type v_j at time t . For $t = 1$, $p_{11} = 1$ and $p_{21} = 0$. For $2 \leq t \leq T$, $p_{1,t} = 0$ and $p_{2,t} = 1/N$. In
292 other words, with probability $1 - 1/N$, no vertex will arrive (or we can assume a dummy node will
293 arrive) during round $t \geq 2$. Let $w_1 \doteq w_{(u,v_1)} = 1 - 1/N$ and $w_2 \doteq w_{(u,v_2)} = 1$. For $C_1 \doteq C_{(u,v_1)}$, it
294 takes value of T and 1 with respective probabilities $1 - 1/N$ and $1/N$. For $C_2 \doteq C_{(u,v_2)}$, it takes the
295 value 1 with probability 1.

296 **LEMMA 6.** The benchmark LP (1) has an optimal value of $2 - 1/N$ on Example 1.

297 **PROOF.** For ease of exposition, let x_t and y_t denote the probabilities that $e = (u, v_1)$ and $e =$
298 (u, v_2) get chosen at time t in any offline optimal, respectively. Thus, the updated benchmark LP is
299 as follows:

$$\text{maximize } \sum_{t \in [T]} (1 - 1/N)x_t + \sum_{t \in [T]} y_t \quad (6)$$

$$\text{s.t. } x_1 \leq 1, y_1 \leq 0 \quad (7)$$

$$x_t \leq 0, y_t \leq 1/N \quad \forall t \geq 2 \quad (8)$$

$$x_1(1 - 1/N) + y_t \leq 1 \quad \forall t \geq 2 \quad (9)$$

$$0 \leq x_t, y_t \leq 1 \quad \forall t \in [T] \quad (10)$$

300 We can verify that the optimal solution to the above LP is as follows: $x_1 = 1$ and $y_t = 1/N$ for all
301 $t \geq 2$, all the rest are zeros. Therefore, the optimal value is $(1 - 1/N) + (1/N) \cdot (T - 1) = 2 - 1/N$. \square

302 **LEMMA 7.** The optimal (adaptive) online algorithm has an expected performance of $1 + 1/N$ on
303 Example 1.

304 **PROOF.** Let ALG be an optimal online algorithm and suppose that it matches the edge $e = (u, v_1)$
 305 with probability $\alpha \in [0, 1]$ at $t = 1$. Let $\mathbf{E}[\text{ALG}]$ be the expected total weight of all matches achieved
 306 by ALG. Thus,

$$\mathbf{E}[\text{ALG}] = \alpha \left(1 + (1/N) \cdot (1/N) \cdot N \right) + (1 - \alpha) \cdot (1/N) \cdot N = 1 + \alpha/N.$$

307 Thus, the optimal online algorithm will choose $\alpha = 1$ and the resultant expected performance is
 308 $1 + 1/N$. \square

309 **Proof of Theorem 2.**

310 **PROOF.** Combining the two lemmas above, we see that the optimal algorithm can achieve a
 311 competitive ratio of $(1 + 1/N)/(2 - 1/N)$ on Example 1 w.r.t. to benchmark LP (1). This completes
 312 the proof. \square

313 **4.1 The special case of deterministic C_e being a constant**

314 Consider a complete bipartite graph $G = (U, V, E)$ where $|U| = K$, $|V| = n^2$. Suppose we have
 315 $T = n$ rounds and $p_{v,t} = \frac{1}{n^2}$ for each v and t . In other words, in each round t , each v is sampled
 316 uniformly from V . For each e , let C_e be deterministically equal to K , which implies that each u
 317 will be unavailable for a constant K rounds after each assignment. Assume all assignments have a
 318 uniform weight (*i.e.*, $w_e = 1$ for all e). Split the whole online process of n rounds into $n - K + 1$
 319 consecutive windows $\mathcal{W} = \{W_\ell\}$ such that $W_\ell = \{\ell, \ell + 1, \dots, \ell + K - 1\}$ for each $1 \leq \ell \leq n - K + 1$.
 320 The benchmark LP (1) then reduces to the following.

$$\max \sum_{t \in [T]} \sum_{e \in E} x_{e,t} \quad (11)$$

$$\text{s.t.} \sum_{e \in E_v} x_{e,t} \leq \frac{1}{n^2} \quad \forall v \in V, t \in [T] \quad (12)$$

$$\sum_{t \in W_\ell} \sum_{e \in E_u} x_{e,t} \leq 1 \quad \forall u \in U, 1 \leq \ell \leq n - K + 1 \quad (13)$$

$$0 \leq x_{e,t} \leq 1 \quad \forall e \in E, t \in [T] \quad (14)$$

321 We can verify that an optimal solution to the above LP is as follows: $x_{e,t}^* = 1/(n^2 K)$ for all e and
 322 t with the optimal objective value of n . We investigate the performance of any optimal non-adaptive
 323 algorithm. Notice that the expected arrivals of any v in the full sequence of online arrivals is $1/n$.
 324 Thus for any non-adaptive algorithm NADAP, it needs to specify the allocation distribution \mathcal{D}_v for
 325 each v during the first arrival. Consider a given NADAP parameterized by $\{\alpha_{u,v} \in [0, 1]\}$ for each
 326 v and $u \in E_v$ such that $\sum_{u \in E_v} \alpha_{u,v} \leq 1$ for each v . In other words, NADAP will assign v to u with
 327 probability $\alpha_{u,v}$ when v comes for the first time and u is available.

Let $\beta_u = \sum_{v \in E_u} \alpha_{u,v} * \frac{1}{n^2}$, which is the probability that u is matched in each round if it is safe at
 the beginning of that round, when running NADAP. Hence,

$$\sum_{u \in U} \beta_u = \sum_{u \in U} \sum_{v \in E_u} \alpha_{u,v} \cdot \frac{1}{n^2} = \sum_{v \in V} \sum_{u \in E_v} \alpha_{u,v} \cdot \frac{1}{n^2} \leq 1$$

328 Consider a given u with β_u and let $\gamma_{u,t}$ be the probability that u is available at t . Then the expected
 329 number of matches of u after the n rounds is $\sum_t \beta_u \gamma_{u,t}$. We have the recursive inequalities on $\gamma_{u,t}$
 330 as in Lemma 8, with $\gamma_{u,t} = 1, t = 1$.

331 LEMMA 8. $\forall 1 < t \leq n$, we have

$$\gamma_{u,t} + \beta_u \sum_{t-K+1 \leq t' < t} \gamma_{u,t'} = 1$$

PROOF. The inequality for $t = 1$ is due to the fact that u is safe at $t = 1$. For each time $t > 1$, Let $SF_{u,t}$ be the event that u is safe at t and $A_{u,t}$ be the event that u is matched at t . Observe that for each window of K time slots, $\{SF_{u,t}, A_{u,t'}, t - K + 1 \leq t' < t\}$ are mutually exclusive and collectively exhaustive events. Therefore,

$$\begin{aligned} 1 &= \Pr[SF_{u,t}] + \sum_{t-K+1 \leq t' < t} \Pr[A_{u,t'}] \\ &= \gamma_{u,t} + \beta_u \sum_{t-K+1 \leq t' < t} \gamma_{u,t'} \end{aligned}$$

332

□

333 Note that the OPT of our benchmark LP is n while the performance of NADAP is $\sum_u \sum_t \beta_u \gamma_{u,t}$.
334 The resulting competitive ratio achieved by an optimal NADAP is captured by the following maxi-
335 mization problem.

$$\max \quad \frac{\sum_u \sum_t \beta_u \gamma_{u,t}}{n} \quad (15)$$

$$\text{s.t.} \quad \sum_{u \in U} \beta_u \leq 1 \quad (16)$$

$$\gamma_{u,t} + \beta_u \sum_{t-K+1 \leq t' < t} \gamma_{u,t'} = 1 \quad \forall 1 < t \leq n, u \in U \quad (17)$$

$$\beta_u \geq 0, \gamma_{u,1} = 1 \quad \forall u \in U \quad (18)$$

336 We prove the following Lemma which implies Theorem 3.

337 LEMMA 9. *The optimal value of the program (15) is at most $\frac{1}{2-1/K} + K/n$.*

338 PROOF. Focus on a given vertex $u \in U$. Notice that $\gamma_{u,t} + \beta_u \sum_{t-K+1 \leq t' < t} \gamma_{u,t'} = 1$ for all
339 $1 \leq t \leq n$. Summing both sides over $t \in [n]$, we have the following.

$$\begin{aligned} \left(1 + \beta_u(K-1)\right) \sum_{t \in [n]} \gamma_{u,t} &= n + \beta_u(K-1)\gamma_{u,n} + \beta_u(K-2)\gamma_{u,n-1} + \cdots + \beta_u\gamma_{u,n-K+2} \\ &\leq n + K - 1 \end{aligned}$$

340 Therefore we have,

$$\sum_{t \in [n]} \gamma_{u,t} \leq \frac{n}{1 + \beta_u(K-1)} + \frac{K-1}{1 + \beta_u(K-1)} \leq \frac{n}{1 + \beta_u(K-1)} + \frac{1}{\beta_u}$$

341 Define $H_u \doteq \sum_t \beta_u \gamma_{u,t}$. From the above analysis, we have that $H_u \leq \frac{n\beta_u}{1 + \beta_u(K-1)} + 1$. Thus the
342 objective value in the program (15) can be upper-bounded as follows.

$$\frac{\sum_u \sum_t \beta_u \gamma_{u,t}}{n} = \sum_{u \in U} \frac{H_u}{n} \leq \sum_{u \in U} \frac{\beta_u}{1 + \beta_u(K-1)} + \frac{K}{n}$$

343 We claim that the optimal value to the program (15) can be upper bounded by the following
 344 maximization program.

$$\left\{ \max \sum_{u \in [U]} \frac{\beta_u}{1 + \beta_u(K - 1)} + \frac{K}{n} : \sum_{u \in U} \beta_u = 1, \beta_u \geq 0, \forall u \in U \right\}$$

345 According to our assumption $K = o(n)$, the second term can be ignored. Let $g(x) = x/(1 +$
 346 $x(K - 1))$. For any $K \geq 2$, it is a concave function, which implies that maximization of g subject to
 347 $\sum_u \beta_u = 1$ will be achieved when all $\beta_u = 1/K$. The resultant value is $\frac{1}{2-1/K} + o(1)$. Thus we are
 348 done. \square

349 **Hardness against the ex-post optimal solution.** Here we show a hardness result against the ex-
 350 post optimal solution. Manshadi et al. [36] prove that for the online-matching problem under known
 351 IID distributions (but disposable offline vertices), no algorithm can achieve a ratio better than 0.823.
 352 Since our setting generalizes this, their hardness result directly applies to our problem as well.
 353 It is worth noting that the gap between the LP relaxation (1) and the ex-post optimal solution is
 354 currently unknown. Thus, either the hardness result can be improved or the algorithm can use a
 355 tighter relaxation of the offline optimal solution.

356 5 EXPERIMENTS

357 To validate the approaches presented in this paper, we use the New York City Yellow Cabs dataset,⁵
 358 which contains the trip records for trips in Manhattan, Brooklyn, and Queens for the year 2013. The
 359 dataset is split into 12 months. For each month we have numerous records each corresponding to a
 360 single trip. Each record has the following structure. We have an anonymized license number which
 361 is the primary key corresponding to a car. For privacy purposes a long string is used as opposed to
 362 the actual license number. We then have the time at which the trip was initiated, the time at which
 363 the trip ended, and the total time of the trip in seconds. This is followed by the starting coordinates
 364 (*i.e.*, latitude and longitude) of the trip and the destination coordinates of the trip.

365 **Assumptions.** We make two assumptions specific to our experimental setup. Firstly, we assume that
 366 every car starts and ends at the same location, for *all* trips that it makes. Initially, we assign every car
 367 a location (potentially the same) which corresponds to its *docking* position. On receiving a request,
 368 the car leaves from this docking position to the point of pick-up, executes the trip and returns to this
 369 docking position. Secondly, we assume that occupation time distributions (OTD) associated with
 370 all matches are identically (and independently) distributed, *i.e.*, $\{C_e\}$ follow the same distribution.
 371 Note that this is a much stronger assumption than what we made in the model, and is completely
 372 inspired by the dataset (see Section 5.2). We test our model on two specific distributions, namely a
 373 *normal* distribution and a *power-law* distribution (see Figure 5). The docking position of each car
 374 and parameters associated with each distribution are all learned from the training dataset (described
 375 below in the **Training** discussion).

376 5.1 Experimental Setup

377 For our experimental setup, we randomly select 30 cabs (each cab is denoted by u). We discretize the
 378 Manhattan map into cells such that each cell is approximately 4 miles (increments of 0.15 degrees
 379 in latitude and longitude). For each pair of locations, say (a, b) , we create a request *type* v , which
 380 represents all trips with starting and ending locations falling into a and b respectively. In our model,
 381 we have $|U| = 30$ and $|V| \approx 550$ (variations depending on day to day requests with low variance).

⁵<http://www.andresmh.com/nyctaxitrips/>

382 We focus on the month of January 2013. We split the records into 31 parts, each corresponding to a
 383 day of January. We choose a random set of 12 parts for *training* purposes and use the remaining for
 384 *testing* purposes.

385 The edge weight w_e on $e = (u, v)$ (*i.e.*, edge from a car u to type v) is set as a function of two
 386 distances in our setup. The first is the trip distance (*i.e.*, the distance from the starting location to the
 387 ending location of v , denoted L_1) while the second is the docking distance (*i.e.*, the distance from the
 388 docking position of u to the starting/ending location of v , denoted L_2). We set $w_e = \max(L_1 - \alpha L_2, 0)$,
 389 where α is a parameter capturing the subtle balance between the positive contribution from the trip
 390 distance and negative contribution from the docking distance to the final profit. We set $\alpha = 0.5$ for
 391 the experiments. We consider each single day as the time horizon and set the total number of rounds
 392 $T = \frac{24 \cdot 60}{5} = 288$ by discretizing the 24-hour period into a time-step of 5 minutes. Throughout this
 393 section, we use time-step and round interchangeably.

394 **Training.** We use the training dataset of 12 days to learn various parameters. As for the arrival rates
 395 $\{p_{v,t}\}$, we count the total number of appearances of each request type v at time-step t in the 12
 396 parts (denote it by $c_{v,t}$) and set $p_{v,t} = c_{v,t}/12$ under KAD (Note that $c_{v,t}$ is at most 12 and hence
 397 this value is always less than 1). When assuming KIID, we set $p_v = p_{v,t} = (c_v/12)/T$ where we
 398 have $c_v = \sum_{t \in [T]} c_{v,t}$ (*i.e.*, the arrival distributions are assumed the same across all the time-steps
 399 for each v). The estimation of parameters for the two different occupation time distributions are
 400 processed as follows. We first compute the average number of seconds between two *requests* in the
 401 dataset (note this was 5 minutes in the experimental setup). We then assume that each *time-step* of
 402 our online process corresponds to a time-difference of this average in seconds. We then compute the
 403 sample mean and sample variance of the trip lengths (as number of seconds taken by the trip divided
 404 by five minutes) in the 12 parts. Hence we use the normal distribution obtained by this sample mean
 405 and standard deviation as the distribution with which a car is unavailable. We assign the docking
 406 position of each car to the location (in the discretized space) in which the majority of the requests
 407 were initiated (*i.e.*, starting location of a request) and matched to this car.

408 5.2 Justifying The Two Important Model Assumptions

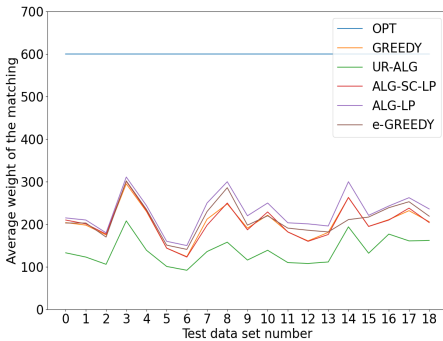


Fig. 1. OTD is normal distribution under KIID

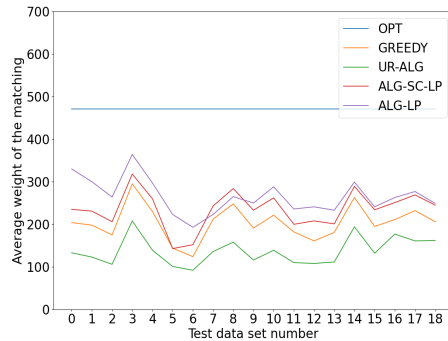


Fig. 2. OTD is normal distribution under KAD

409 **Known Adversarial Distributions.** Figure 4 plots the number of arrivals of a particular type at var-
 410 ious times during the day. Notice the significant increase in the number of requests in the middle of
 411 the day as opposed to the mornings and nights. This justified our arrival assumption of KAD which

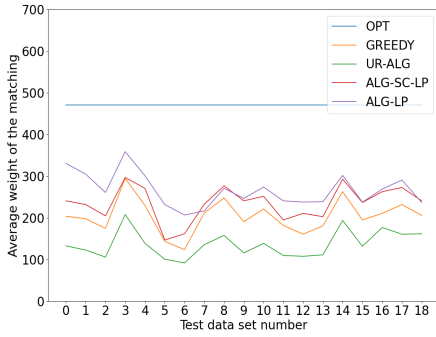


Fig. 3. OTD is power law distribution under KAD




Fig. 4. The number of requests of a given type at various time-steps. x-axis: time-step, y-axis: number of requests

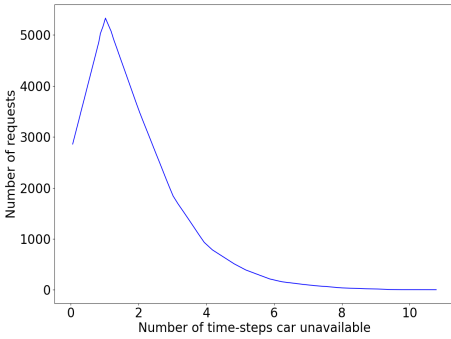


Fig. 5. Occupation time distribution of all cars. x-axis: number of time-steps, y-axis: number of requests

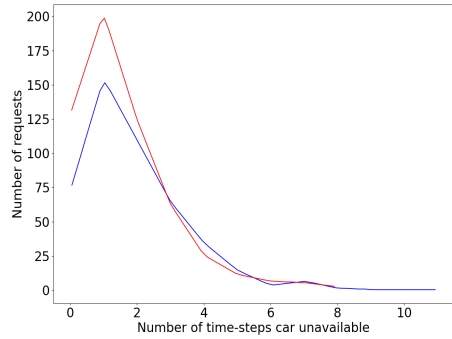


Fig. 6. Occupation time distribution of two different cars. x-axis: number of time-steps, y-axis: number of requests

412 assumes different arrival distributions at different time-steps. Hence the LP (and the correspond-
 413 ing algorithm) can exploit this vast difference in the arrival rates and potentially obtain improved
 414 results compared to the assumption of Known Identical Independent Distributions (KIID). This is
 415 confirmed by our experimental results shown in Figures 1 and 2.

416 **Identical-Occupation-Time Distribution.** We assume each car will be available again via an in-
 417 dependent and identical random process regardless of the matches it received. The validity of our
 418 assumptions can be seen in Figures 5 and 6, where the x -axis represents the different occupation
 419 time and the y -axis represents the corresponding number of requests in the dataset responsible for
 420 each occupation time. It is clear that for most requests the occupation time is around 2-3 time-steps
 421 and dropping drastically beyond that with a long tail. Figure 6 displays occupation times for two
 422 representative (we chose two out of the many cars we plotted, at random) cars in the dataset; we
 423 see that the distributions roughly coincide with each other, suggesting that such distributions can be
 424 learned from historical data and used as a guide for future matches.

425 5.3 Results

426 Inspired by the experimental setup by [55, 56], we run five different algorithms on our dataset. The
 427 first algorithm is the ALG-LP. In this algorithm, when a request v arrives, we choose a neighbor u
 428 with probability $x_{e,t}^*/p_{v,t}$ with $e = (u, v)$ if u is available. Here $x_{e,t}^*$ is an optimal solution to our
 429 benchmark LP and $p_{v,t}$ is the arrival rate of type v at time-step t . The second algorithm is called
 430 ALG-SC-LP. Recall that $E_{v,t}$ is the set of “safe” or available assignments with respect to v when the
 431 type v arrives at t . Let $x_{v,t} = \sum_{e \in E_{v,t}} x_{e,t}^*$. In ALG-SC-LP, we sample a safe assignment for v with
 432 probability $x_{e,t}^*/x_{v,t}$. The next two algorithms are heuristics oblivious to the underlying LP. Our
 433 third algorithm is called GREEDY which is as follows. When a request v comes, match it to the safe
 434 neighbor u with the highest edge weight. Our fourth algorithm is called UR-ALG which chooses one
 435 of the safe neighbors uniformly at random. Finally, we use a combination of LP-oblivious algorithm
 436 and LP-based algorithm called ϵ -GREEDY. In this algorithm when a type v comes, with probability ϵ
 437 we use the greedy choice and with probability $1 - \epsilon$ we use the optimal LP choice. In our algorithm,
 438 we optimized the value of ϵ and set it to $\epsilon = 0.1$. We summarize our results in the following
 439 plots. Figures 1, 2, and 3 show the performance of the five algorithms and OPT (optimal value of
 440 the benchmark LP) under the different assumptions of the OTD (normal or power law) and online
 441 arrives (KIID or KAD). In all three figures the x-axis represents test data-set number and the y-axis
 442 represents average weight of matching.

443 **Discussion.** From the figures, it is clear that both the LP-based solutions, namely ALG-LP and
 444 ALG-SC-LP, do better than choosing a free neighbor uniformly at random. Additionally, with dis-
 445 tributional assumptions the LP-based solutions outperform greedy algorithm as well. We would like
 446 to draw attention to a few interesting details in these results. Firstly, compared to the LP optimal
 447 solution, our LP-based algorithms have a competitive ratio in the range of 0.5 to 0.7. We believe this
 448 is because of our experimental setup. In particular, we have that the rates are high (> 0.1) only in a
 449 few time-steps while in all other time-steps the rates are very close to 0. This means that it resembles
 450 the structure of the *theoretical* worst case example we showed in Section 4. In future experiments,
 451 running our algorithms during *peak* periods (where the request rates are significantly larger than 0)
 452 may show that competitive ratios in those cases approach 1. Secondly, it is surprising that our algo-
 453 rithm is fairly robust to the *actual* distributional assumption we made. In particular, from Figures 2
 454 and 3 it is clear that the difference between the assumption of normal distribution versus power-law
 455 distribution for the unavailability of cars is *negligible*. This is important since it might not be easy
 456 to learn the *exact* distribution in many cases (*e.g.*, cases where the sample complexity is high) and
 457 this shows that a close approximation will still be as good.

458 **Simulation based algorithm.** We omit the results of the simulation based algorithm, since the per-
 459 formance was similar to the algorithm without the scaling (*i.e.*, ALG-LP). Here we briefly describe
 460 the implementation details on performing the simulations efficiently in practice. The estimates are
 461 computed even before the start of the algorithm. We first simulate the entire sequence of T requests,
 462 δ times. Using these δ samples we first compute the estimates for the first time-step. We now re-use
 463 the same δ samples and the computed estimates in the first time-step to obtain the estimates for the
 464 second time-step. Hence in a sequential manner, we compute estimates at time t using the samples
 465 from time-steps $1, 2, \dots, t - 1$. The overall run-time of this implementation is $O(\delta T + \delta T \kappa)$, where
 466 κ denotes the running time of ADAP in every time-step. Hence during the online phase, the running
 467 time of ADAP is same as that of ALG-LP.

6 CONCLUSION AND FUTURE DIRECTIONS

In this work, we provide a model that captures the application of assignment in ride-sharing platforms. One key aspect in our model is to consider the *reusable* aspect of the offline resources. This helps in modeling many other important applications where agents enter and leave the system *multiple* times (e.g., organ allocation, crowdsourcing markets [27], etc.). Our work opens several important research directions. The first direction is to generalize the online model to the *batch* setting (e.g., this subsequent work). In other words, in each round we assume multiple arrivals from V . This assumption is useful in crowdsourcing markets (for example) where multiple tasks—but not all—become available at some time. The second direction is to consider a Markov model on the driver starting position. In this work, we assumed that each driver returns to her docking position. However, in many ride-sharing systems, drivers start a new trip from the position of the last-drop off. This leads to a Markovian system on the offline types, as opposed to the assumed static types in the present work. Finally, pairing our current work with more-applied stochastic-optimization and reinforcement-learning approaches would be of practical interest to policymakers running taxi and bikeshare services [23, 31, 45, 51, 57]. Following the initial conference publication of this paper, subsequent work has appeared that addresses some of the aforementioned directions. The multi-capacitated version of the problem in the *batch* setting was studied in [32] where the authors devise an algorithm that has a competitive ratio of 0.317. Driver-rider matching in rideshare was modeled as a Markovian system [18]; they show that under a practical condition this system converges to the stationary distribution very fast. The unweighted multi-capacity version of the problem considered in this paper under the adversarial arrival model was studied in [25]. They proved an optimal competitive ratio $1 - 1/e$ improving over the ratio of $1/2$ proved in [22, 48].

Acknowledgements. The authors would like to thank the anonymous reviewers of AAAI-2018 and the TEAC reviewers for valuable comments on the presentation of this paper.

REFERENCES

- [1] Marek Adamczyk, Fabrizio Grandoni, and Joydeep Mukherjee. Improved approximation algorithms for stochastic matching. In *ESA-15*. 2015.
- [2] Daniel Adelman. Price-directed control of a closed logistics queueing network. *Operations Research*, 2007.
- [3] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Dynamic matching market design. *EC*, 2014.
- [4] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *EC-12*, 2012.
- [5] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. The online stochastic generalized assignment problem. In *APPROX-RANDOM-13*. 2013.
- [6] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.
- [7] Itai Ashlagi, Patrick Jaillet, and Vahideh H Manshadi. Kidney exchange in dynamic sparse heterogeneous pools. *EC*, 2013.
- [8] Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. 2017.
- [9] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [10] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Dynamic pricing in ridesharing platforms. *ACM SIGecom Exchanges*, 2016.
- [11] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. In *EC*, 2017.
- [12] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1), 2013.
- [13] Anton Braverman, Jim G Dai, Xin Liu, and Lei Ying. Empty-car routing in ridesharing systems. *arXiv preprint arXiv:1609.07219*, 2016.

- 518 [14] Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. New algorithms, better bounds, and a
519 novel model for online stochastic matching. In *ESA-16*, 2016.
- 520 [15] Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Attenuate locally, win globally: An
521 attenuation-based framework for online stochastic matching with timeouts. In *Sixteenth International Conference on*
522 *Autonomous Agents and Multiagent Systems (AAMAS 2017)*. 2017.
- 523 [16] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. Online primal-dual algorithms for maximizing ad-auctions rev-
524 enue. In *ESA-07*, 2007.
- 525 [17] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. Surge pricing solves the wild goose chase. In *EC*, 2017.
- 526 [18] Michael Curry, John P Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, Yuhao Wan, and Pan Xu. Mix
527 and match: Markov chains and mixing times for matching in rideshare. In *International Conference on Web and Internet*
528 *Economics (WINE)*, pages 129–141. Springer, 2019.
- 529 [19] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms
530 and fast approximation algorithms for resource allocation problems. In *EC-11*, 2011.
- 531 [20] John P. Dickerson and Tuomas Sandholm. FutureMatch: Combining human value judgments and machine learning to
532 match in dynamic environments. In *AAAI-15*, 2015.
- 533 [21] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S Muthukrishnan. Online stochastic matching: Beating $1-1/e$. In
534 *FOCS-09*, 2009.
- 535 [22] Yiding Feng, Rad Niazadeh, and Amin Saberi. Linear programming based online policies for real-time assortment of
536 reusable resources. Available at SSRN 3421227, 2019.
- 537 [23] Supriyo Ghosh, Pradeep Varakantham, Yossiri Adulyasak, and Patrick Jaillet. Dynamic repositioning to reduce lost
538 demand in bike sharing systems. *Journal of Artificial Intelligence Research (JAIR)*, 58:387–430, 2017.
- 539 [24] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In
540 *SODA-08*, 2008.
- 541 [25] Vineet Goyal, Garud Iyengar, and Rajan Udhwani. Online allocation of reusable resources: Achieving optimal compet-
542 itive ratio. *arXiv preprint arXiv:2002.02430*, 2020.
- 543 [26] Bernhard Haeupler, Vahab S Mirrokni, and Morteza Zadimoghaddam. Online stochastic weighted matching: Improved
544 approximation algorithms. In *WINE-11*, 2011.
- 545 [27] Chien-Ju Ho and Jennifer Wortman Vaughan. Online task assignment in crowdsourcing markets. In *AAAI-12*, 2012.
- 546 [28] Patrick Jaillet and Xin Lu. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations*
547 *Research*, 39(3), 2013.
- 548 [29] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In
549 *STOC-90*, 1990.
- 550 [30] Der-Hong Lee, Hao Wang, Ruey Cheu, and Siew Teo. Taxi dispatch system based on current demands and real-time
551 traffic conditions. *Transportation Research Record: Journal of the Transportation Research Board*, (1882), 2004.
- 552 [31] Meghna Lowalekar, Pradeep Varakantham, and Patrick Jaillet. Online spatio-temporal matching in stochastic and
553 dynamic domains. In *AAAI-16*, 2016.
- 554 [32] Meghna Lowalekar, Pradeep Varakantham, and Patrick Jaillet. Competitive ratios for online multi-capacity ridesharing.
555 In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 771–779,
556 2020.
- 557 [33] Hongyao Ma, Fei Fang, and David C Parkes. Spatio-temporal pricing for ridesharing platforms. In *ACM Conference*
558 *on Economics and Computation (EC)*, pages 583–583, 2019.
- 559 [34] Will Ma. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms.
560 In *SODA-14*, 2014.
- 561 [35] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: an approach based on strongly
562 factor-revealing lps. In *STOC-11*, 2011.
- 563 [36] Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on
564 offline statistics. *Mathematics of Operations Research*, 37(4), 2012.
- 565 [37] Nicholas Mattei, Abdallah Saffidine, and Toby Walsh. Mechanisms for online organ matching. In *IJCAI-17*, 2017.
- 566 [38] Nicole Megow, Marc Uetz, and Tjark Vredeveld. Stochastic online scheduling on parallel machines. In *WAOA-04*,
567 2004.
- 568 [39] Nicole Megow, Marc Uetz, and Tjark Vredeveld. Models and algorithms for stochastic online scheduling. *Mathematics*
569 *of Operations Research*, 31(3), 2006.
- 570 [40] Aranyak Mehta. Online matching and ad allocation. *Theoretical Computer Science*, 8(4), 2012.
- 571 [41] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal*
572 *of the ACM (JACM)*, 54(5), 2007.
- 573 [42] Michael Miller. *Cloud computing: Web-based applications that change the way you work and collaborate online*. Que
574 publishing, 2008.

- 575 [43] Brendon L Neuen, Georgina E Taylor, Alessandro R Demaio, and Vlado Perkovic. Global kidney disease. *The Lancet*,
576 382(9900), 2013.
- 577 [44] Afshin Nikzad. Thickness and competition in ride-sharing markets. Technical report.
- 578 [45] Eoin O’Mahony and David B Shmoys. Data analysis and optimization for (citi) bike sharing. In *AAAI-15*, 2015.
- 579 [46] Erhun Ozkan and Amy Ward. Dynamic matching for real-time ridesharing. *Working Paper*, 2017.
- 580 [47] Jiang Rong, Tao Qin, and Bo An. Dynamic pricing for reusable resources in competitive market with stochastic demand.
581 In *AAAI*, 2018.
- 582 [48] Paat Rusmevichientong, Mika Sumida, and Huseyin Topaloglu. Dynamic assortment optimization for reusable products
583 with random usage durations. *Management Science*, 2020.
- 584 [49] Rajiv Saran, Yi Li, Bruce Robinson, John Ayanian, Rajesh Balkrishnan, Jennifer Bragg-Gresham, JT Chen, Elizabeth
585 Cope, Debbie Gipson, Kevin He, et al. US renal data system 2014 annual data report: Epidemiology of kidney disease
586 in the United States. *American Journal of Kidney Diseases*, 65(6 Suppl 1), 2015.
- 587 [50] Kiam Tian Seow, Nam Hai Dang, and Der-Horng Lee. A collaborative multiagent taxi-dispatch system. *IEEE Trans-*
588 *actions on Automation Science and Engineering*, 7(3), 2010.
- 589 [51] Divya Singhvi, Somya Singhvi, Peter I Frazier, Shane G Henderson, Eoin O’Mahony, David B Shmoys, and Dawn B
590 Woodard. Predicting bike usage for new york city’s bike sharing system. In *AAAI-15 Workshop on Computational*
591 *Sustainability*, 2015.
- 592 [52] Martin Skutella, Maxim Sviridenko, and Marc Uetz. Unrelated machine scheduling with stochastic processing times.
593 *Mathematics of operations research*, 41(3), 2016.
- 594 [53] Xiaoming Sun, Jia Zhang, and Jialin Zhang. Near optimal algorithms for online weighted bipartite matching in adver-
595 sary model. *Journal of Combinatorial Optimization*, 2016.
- 596 [54] Yong Sun, Jun Wang, and Wenan Tan. Online algorithms of task allocation in spatial crowdsourcing. In *ICDE*, 2017.
- 597 [55] Yongxin Tong, Jieying She, Bolin Ding, Lei Chen, Tianyu Wo, and Ke Xu. Online minimum matching in real-time
598 spatial data: experiments and analysis. *Proceedings of the VLDB Endowment*, 9(12), 2016.
- 599 [56] Yongxin Tong, Jieying She, Bolin Ding, Libin Wang, and Lei Chen. Online mobile micro-task allocation in spatial
600 crowdsourcing. In *ICDE-16*, 2016.
- 601 [57] Tanvi Verma, Pradeep Varakantham, Sarit Kraus, and Hoong Chuin Lau. Augmenting decisions of taxi drivers through
602 reinforcement learning for improving revenues. 2017.
- 603 [58] Ariel Wasserhole and Vincent Jost. Pricing in vehicle sharing systems: Optimization in queuing networks with product
604 forms. *EURO Journal on Transportation and Logistics*, 2016.
- 605 [59] Man Lung Yiu, Kyriakos Mouratidis, Nikos Mamoulis, et al. Capacity constrained assignment in spatial databases. In
606 *SIGMOD-08*, 2008.
- 607 [60] Andrew J Younge, Gregor Von Laszewski, Lizhe Wang, Sonia Lopez-Alarcon, and Warren Carithers. Efficient resource
608 management for cloud computing environments. In *International Green Computing Conference*, 2010.