

Mix and Match: Markov Chains & Mixing Times for Matching in Rideshare

Michael Curry¹, John P. Dickerson¹, Karthik Abinav Sankararaman^{1,2}, Aravind Srinivasan¹, Yuhao Wan^{1,3}, and Pan Xu^{1,4}

¹ University of Maryland, College Park MD USA

² Facebook, Inc., Menlo Park CA USA

³ University of Washington, Seattle WA USA

⁴ New Jersey Institute of Technology, Newark NJ USA

{curry, john, srin}@cs.umd.edu

{karthikabinavs, yuhao.diane.wan, panxu0}@gmail.com

Abstract. Rideshare platforms such as Uber and Lyft dynamically dispatch drivers to match riders' requests. We model the dispatching process in rideshare as a Markov chain that takes into account the geographic mobility of both drivers and riders over time. Prior work explores dispatch policies in the limit of such Markov chains; we characterize when this limit assumption is valid, under a variety of natural dispatch policies. We give explicit bounds on convergence in general, and exact (including constants) convergence rates for special cases. Then, on simulated and real transit data, we show that our bounds characterize convergence rates—even when the necessary theoretical assumptions are relaxed. Additionally these policies compare well against a standard reinforcement learning algorithm which optimizes for profit without any convergence properties.

1 Introduction

Rideshare firms such as Uber, Lyft, and Didi Chuxing dynamically match riders to drivers via an online, digital platform. Riders request a driver through an online portal or mobile app; a driver is matched by the platform to a rider based on geographic proximity, driver preferences, pricing, and other factors. The rideshare driver then picks up the rider at her request location, transfers her to her destination, and reenters the platform to be matched again—albeit at a new geographic location. Part of the larger *sharing economy*, rideshare firms are increasingly competitive against traditional taxi services due to their ease of use, lower pricing, and immediacy of service [14].

Matching riders to rideshare drivers is nontrivial. While the core process is a form of the well-studied online matching problem [23], current models developed in the EconCS, AI, and Operations Research communities [28, 11, 5] do not completely capture the mobile aspects of both the drivers and riders. Drivers *and* riders are agents who move about a constrained space (e.g., city streets), becoming (in)active periodically due to the matching process. When a platform receives a request, it must make a near-real-time dispatch decision amongst nearby drivers who are available at the current time. The platform's goal is to maximize an objective (e.g., revenue or throughput) by servicing requests in an online fashion, subject to various real-world constraints and challenges like setting prices, predicting supply and demand, fairness considerations, competing with other firms, and so on [8, 17, 4].

In this paper, we study the dynamics of the nascent rideshare market under different *dispatch strategies*. Recent work uses Markov chains to model complex ride-sharing dynamics in a *closed-world* system—that is, a system with a fixed total supply of cars [28, 4, 5]. These assume the Markov chains reach their stationary distributions quickly, and thus all prior work optimizes for dispatch strategies in the limit. As a complement, the present paper characterizes—theoretically and empirically—when that limit assumption is valid under a variety of natural strategies.

Our contributions. The main contribution in this paper is *to show both theoretically and empirically the convergence rate of many natural policies to the stationary distribution*. First, we model the dispatching

problem in rideshare platforms as a Markov chain. The number of states in this Markov chain is exponential in the natural size of the problem; thus, unless the chain is *rapidly mixing*, the time taken to reach the stationary distribution is prohibitively large. Next, we consider two large natural classes of strategies and study the evolution of the driver distribution theoretically. We show that the Markov chains are rapidly mixing and give explicit bounds on the convergence rates. Then, we consider a special case of uniform arrival rates and compute the convergence rates *exactly*, including the constants. Finally, we conduct experiments on both simulated as well as a real-world large-scale dataset to corroborate our findings. In particular, even when the assumptions needed by the theory do not necessarily hold, simulations show that the convergence behavior does not change drastically, and experiments on real data show that the theory gives direct insight on the convergence properties in practice. Additionally, compared against a standard RL algorithm, these policies perform similarly. Hence, our policies are simpler and efficient to run with theoretical convergence guarantees while performing almost as good as more complicated algorithms without such properties.

2 Preliminaries

In this section, we define the formal model used throughout the paper. We begin with a brief primer on Markov chains, and then show how Markov chains can be used to model rideshare markets.

Definition 1 (Markov chain). A Markov chain \mathcal{M} is defined by a state space Ω and a transition matrix P . $P(\mathbf{x}, \mathbf{y})$ represents the probability of reaching state $\mathbf{y} \in \Omega$, in one step, from state $\mathbf{x} \in \Omega$. $P(\mathbf{x}, \mathbf{y})$ does not depend on states which the process was in prior to \mathbf{x} (i.e., it obeys the Markov property).

A central concept in Markov chain analysis is the notion of a *limiting* distribution and a *stationary* distribution. For two distributions μ and ν on the state space Ω , the *total variation distance* between μ and ν is defined as $\|\mu - \nu\|_{TV} = \frac{1}{2} \sum_{\mathbf{x} \in \Omega} |\mu(\mathbf{x}) - \nu(\mathbf{x})|$. A distribution π^* is said to be a *stationary distribution* if $\pi^* = \pi^* P$. A distribution π^* is said to be a *limiting distribution* if for every initial state \mathbf{x} , we have that $\lim_{t \rightarrow \infty} \|P^t(\mathbf{x}, \cdot) - \pi^*\|_{TV} = 0$, where $P^t(\mathbf{x}, \cdot)$ denotes that distribution of states after t steps, starting at \mathbf{x} .

To use Markov chains as an algorithmic tool, we need to understand the *rate* of convergence to the stationary distribution, commonly called its *mixing time*.

Mixing Time $\tau(\epsilon)$. Consider an irreducible and aperiodic Markov Chain \mathcal{M} with stationary distribution π^* . For a given t , let $d(t) = \max_{\mathbf{x} \in \Omega} \|P^t(\mathbf{x}, \cdot) - \pi^*\|_{TV}$. The mixing time of \mathcal{M} is defined as $\tau(\epsilon) = \min\{t : d(t') \leq \epsilon, \forall t' \geq t\}$. We say \mathcal{M} is *rapidly mixing* if $\tau(\epsilon) = O\left(\text{poly}\left(\log \frac{|\Omega|}{\epsilon}\right)\right)$.

In this paper, we consider a class of strategies (as motivated by, e.g., [11]) for the rideshare problem and cast it as a natural Markov chain. We then consider an objective function which depends on the limiting distribution of this chain and study convergence properties of that objective function, using the mixing properties of the Markov chain. We use Markov chains and relevant tools as the central concepts in this paper. Levin and Peres [18] details classical results in this space. The first *algorithmic* usage of Markov chain Monte Carlo (MCMC) methods can be traced back to the classical works of Hastings [15], Metropolis and Ulam [24], Metropolis et al. [25], commonly known as the Metropolis-Hastings algorithm. MCMC methods are a powerful tool in machine learning and we refer the reader to the survey [2].

A Markov chain model of rideshare. We now define our Markov-chain-based model of rideshare. Consider a two-dimensional grid \mathcal{U} consisting of n points (e.g., geographic locations). A *request type* $r = (u, u')$, represented by an *ordered* pair of points, is a set of requests that start and end at locations $u \in \mathcal{U}$ and $u' \in \mathcal{U}$, respectively. Let $\mathcal{R} = \{r = (u, u') | u \in \mathcal{U}, u' \in \mathcal{U}\}$ be the set of all request types. Note that we allow request types $r = (u, u)$ —that is, a request that both starts and ends at the same point $u \in \mathcal{U}$. This is just for notational convenience.

Given a time horizon T , at each time (or round) $t \in [T] \doteq \{1, 2, \dots, T\}$, a request of type r is sampled from \mathcal{R} with probability p_r .⁵ Our goal is to design a *matching* (or dispatching) scheme—assigning a driver (or car) to a request—that maximizes an overall objective after T rounds. In this paper, we assume that the sampling distribution $\{p_r\}$ in every round is *identical and independently distributed* (IID) but *unknown* to the algorithms. For notational simplicity, we also use r to denote a specific online request of type r when the context is clear.

Dispatching policy. Suppose the system has m identical drivers. We characterize the state of the system by a vector $\mathbf{x} \in \mathbb{Z}_+^n$, where x_u denotes the number of drivers in location u . Then, we can construct a Markov chain with state space $\Omega = \{\mathbf{x} : x_u \in \{0, 1, \dots, c\}, \sum_u x_u = m\}$. In this paper, we assume that $m \ll c \cdot n$ for some constant capacity c . A *dispatching policy* (or strategy) σ is a mapping from $\Omega \times \mathcal{R}$ to $\mathcal{U} \cup \{\emptyset\}$ such that at time t and a state $\mathbf{X}^{(t)}$, when a request r comes, σ assigns that request r to a potential driver at location $u_\sigma = \sigma(\mathbf{X}^{(t)}, r)$.⁶ Here $u_\sigma = \emptyset$ denotes that the policy σ rejects request r . We say σ successfully addresses the request $r = (u, v)$ at t if $X_{u_\sigma}^{(t)} \geq 1$ and $X_v^{(t)} < c$. If σ successfully addresses a request r , it receives a profit w_r .

Neighborhood of u . For each point $u \in \mathcal{U}$, let $\mathcal{N}(u)$ be the set of neighbors of u with Manhattan distance⁷ exactly 1 to u , i.e., $\mathcal{N}(u) = \{u' \in \mathcal{U} : |u - u'|_M = 1\}$, where $|u - u'|_M$ denotes the Manhattan distance between u and u' . We can assign each request with origin u only to a driver in the set $\{u\} \cup \mathcal{N}(u)$.

Objective functions. For a given policy σ , let $W(\sigma, t)$ be the expected profit obtained by the policy σ at time t . Denote $\mathbf{I}_\mathbf{X} := \mathbf{I}(X_{u_\sigma}^{(t)} \geq 1, X_v^{(t)} < c)$ which is the indicator for the event $X_{u_\sigma}^{(t)} \geq 1$ and $X_v^{(t)} < c$. Thus the expected performance of σ at time t and the expected average performance of σ over T rounds is,

$$W(\sigma, t) = \mathbb{E} \left[\sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \mathbf{I}_\mathbf{X} \right], \quad (1)$$

$$\widehat{W}(\sigma, T) = \frac{1}{T} \sum_{t=1}^T W(\sigma, t). \quad (2)$$

respectively. The randomness in the state $\mathbf{X}^{(t)}$ depends on two sources: the random arrival of requests from distribution $\{p_r\}$ and any internal randomness used in the execution of a randomized policy σ .

In this paper, our objective is to study a class of dispatching strategies that are both *effective* and *stable*. Specifically, suppose \mathbf{x}_0 is the initial state, $\mathbf{p} = \{p_r | r \in \mathcal{R}\}$ is the arrival distribution of the request types in each round, and $\mathbf{w} = \{w_r | r \in \mathcal{R}\}$ is the profit vector for the request types. We then wish to answer the following questions.

1. **Effectiveness.** For a given instance $\mathcal{I} = (\mathbf{x}_0, \mathbf{p}, \mathbf{w}, T)$, which non-adaptive strategy σ maximizes $W(\sigma, t)$ and $\widehat{W}(\sigma, T)$?
2. **Stability.** For a given non-adaptive strategy σ , do the limits $\lim_{t \rightarrow \infty} W(\sigma, t)$ and $\lim_{T \rightarrow \infty} \widehat{W}(\sigma, T)$ exist? If so, how fast do they converge to the respective limiting values?

3 Assumptions & Related Work

Rideshare is a recent and popular innovation; thus, the body of literature surrounding this paradigm is young and quickly growing. With that in mind, we now explicitly motivate and list the assumptions we make in our paper, and then place our model in the greater body of related work.

Model assumptions. First, we assume that the number of drivers in the system remains constant in the T rounds with no new driver either joining or leaving the system. This assumption is justified because (1) T

⁵ We have $\sum_{r \in \mathcal{R}} p_r \leq 1$. Thus, with probability $1 - \sum_{r \in \mathcal{R}} p_r$, there is no request in any given time.

⁶ The choice of u_σ can be random since σ can be a randomized policy.

⁷ It is not critical for our purposes, but the experiments use New York city and road-distance is measured in Manhattan distance.

rounds in the online phase is typically restricted to a few hours and (2) almost all trips are local (as described later in our experimental section on real data).

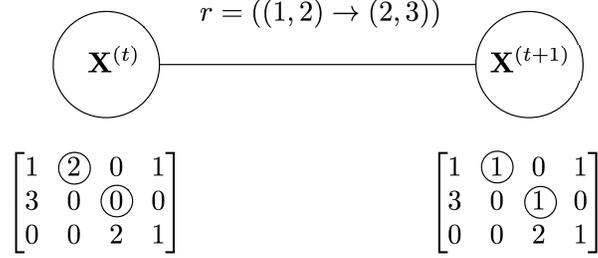


Fig. 1: A single state transition in the Markov chain; driver matched at (1, 2), trip ends at (2, 3).

Second, at time t when a request $r = (u, u')$ comes and is assigned to a driver at location v , we assume that (1) the system obtains a profit $w_r > 0$, which is proportional to the distance between u and u' ; (2) at time $t + 1$, the number of drivers at v reduces by 1 and the number at u' increases by 1 (*i.e.*, the request is completed instantaneously). From the real dataset, we observe that most trips in the Manhattan area are local and are completed within 15 mins (*i.e.*, a short time period). Hence things do not change drastically by making this simplifying assumption.

Third, we assume that all locations have the same capacity $c \in \mathbb{Z}_+$, an upper bound on the total number of drivers that can be present in each location at any time. Here c captures the maximum number of drivers allowed at any single location.

Fourth, we make the following global *hot-spot assumption* about the requests. There exists a location $u^* \in \mathcal{U}$ such that $p_{(u^*, u)} > 0$ and $p_{(u, u^*)} > 0$ for all $u \in \mathcal{U}, u \neq u^*$. We call this location u^* the “hot-spot”. This assumption naturally holds in many real scenarios (*e.g.*, cell (10, 5) in Figure 2), since most big cities have busy central locations (*e.g.*, the Empire State Building in New York City (NYC)), or have flow to and from an airport (*e.g.*, JFK or LGA, in the case of NYC).

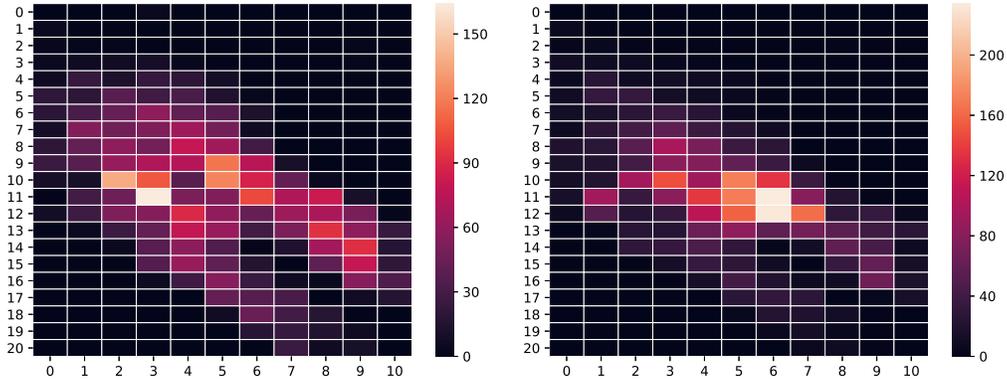


Fig. 2: Heatmap of starting (left) and ending (right) location of rides in NYC (7:00AM–11:00AM). X-axis, Y-axis: numbered from 0-19. Each cell represents a location. Heat index 0-black to 160-white.

Finally, we assume IID arrivals of online requests. This is necessary for convergence—even for very simple dispatching policies. Consider the following example with two locations u and v in the system and a single

driver. Define request types $r_a = (u, v)$ and $r_b = (v, u)$. Suppose the two online requests r_a and r_b arrive alternatively during odd and even rounds, and our dispatching policy is the simple greedy one: match a request r if there is one driver at the starting location of r , otherwise, reject it. Then the Markov chain will admit no stationary distribution, since it has a period of 2.

It is worth noting that the first, second and fifth assumptions are used in related work [4, 5]. The third assumption is a generalization of prior works which all considered $c = m$.

Related work. Research in rideshare 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 [28, 6, 4]. Apart from using Markov-chain methods, many allocation and scheduling problems have been studied in the rideshare context using methods from combinatorial optimization and machine learning (e.g., [11, 9, 16]). A large-scale mathematical and empirical study on the number of cars and the optimality of waiting times was recently analyzed by Alonso-Mora et al. [1]. In our work, we do not consider the waiting times; indeed, we assume that the match and the trips happen instantaneously. The role of pricing in the dynamics of drivers in rideshare platforms is also an active area of research in computational economics and AI/ML (e.g., [22, 30, 3]). Although rideshare platforms face challenges that are unique and different from *bike-sharing* ecosystems, there are some similarities between the two. Both of these deal with a matching market where the agents are constantly moving around and hence it is important to characterize the *flow* patterns of these agents. A number of works study the prediction aspects as well as the dynamics of flow patterns in both bike-sharing (e.g., [27, 12]) as well as in ridesharing (e.g., [35, 34, 17, 36, 29]) platforms.

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, [10, 21, 33, 31] have studied algorithms for matching in various dynamic markets such as kidney exchange, spatial crowdsourcing, labor markets, and so on. Online matching in *static* environments has been extensively studied in the literature; see Mehta et al. [23] for a detailed survey. We use a reinforcement learning algorithm as a baseline in our experiments. Apart from a recent work by Lin et al. [20], which explores the fleet management problem in ride-share, this tool has not been explored much in this domain. However, reinforcement learning (RL) has found success in other applications and we refer the reader to a recent survey [19].

Comparison with related concurrent work. In a recent work, Banerjee et al. [5] proposed an approximation framework for pricing in rideshare platforms. Our model shares many characteristics with theirs (i.e., the first, second and fifth assumptions). The main difference is the focus of the paper. In particular, they are interested in finding the optimal (online) dispatching policy that has a *good* approximation ratio with respect to expected revenue at the stationary distribution. Thus they consider the *known IID* arrival assumption for online requests while we consider the stronger *unknown IID* arrival assumption. Moreover they assume that the number of cars tends to infinity, and seek approximation ratios in this limit. The main focus of this paper is instead to characterize the *rate* of convergence to the stationary distribution. We seek to understand when using the expected reward in the stationary distribution is a good measure, especially if the rate of convergence is slow and/or non-existent.

From dispatching to pricing. The techniques in this paper can be extended to the scenario of dynamic pricing. Suppose that the system cannot directly match riders to drivers, but can set prices for every ride. Then, we can modify the schemes in this paper under certain conditions. Consider the pricing model as studied in Banerjee et al. [4]. Every rider arrives with a value, which can be viewed as a sample from a distribution associated with its request type. The system then sets a price (this can be arbitrarily correlated with the valuation distribution and may depend on both the time and request type but not the current state of the system) and the rider accepts the service if and only if their value is greater than the offered price. This process can be incorporated into the arrival rates of the requests and hence one can obtain similar convergence rates for this dynamic pricing problem. For brevity, we consider the problem of dispatching as the running example throughout this paper.

4 Dispatching Polices

In this section, we present our dispatching policies for matching riders to rideshare drivers. Specifically, we present two policies namely $\text{Samp}(\alpha)$ and $\text{Rand-Perm}(\phi)$. The main ideas of these are as follows.

When a request $r = (u, u')$ arrives, $\text{Samp}(\alpha)$ checks the availability of drivers in locations u and each of its neighbors in $\mathcal{N}(u)$ with respective probabilities α and $\frac{1-\alpha}{4}$, where $0 < \alpha < 1$ is a parameter. (Technically, nodes u on the boundary of the grid have less than 4 neighbors; we just assume that we do nothing if the realization is an invalid neighbor.)

When a request $r = (u, u')$ arrives, $\text{Rand-Perm}(\phi)$ will first check the availability of drivers at u and if there aren't any, then checks the neighbors in $\mathcal{N}(u)$ following a given order ϕ . Below, we present the theoretical results on their convergence rates.

Theorem 1 (Main Convergence Theorem). *For any policy $\text{Samp}(\alpha)$ with $\alpha > 0$ and $\text{Rand-Perm}(\phi)$ starting with any given initial state, the objective functions defined in (1) and (2) both converge to the same value with rates $\Theta(\beta^t)$ and $\Theta(T^{-1})$, respectively, for some $\beta \in (0, 1)$ independent of T (possibly dependent on the other input parameters such as m and n).*

Theorem 1 refers to the absolute error bound of the objectives in (1) and (2) from its limit value. In particular, Theorem 1 states that as the number of rounds increases, the expected revenue in every timestep converges to a stable value *exponentially* quickly.

Markov chains $\mathcal{M}(\text{Samp}(\alpha))$ and $\mathcal{M}(\text{Rand-Perm}(\phi))$. We make the simplifying assumption that for every $u \in \mathcal{U}$, the number of neighbors $|\mathcal{N}(u)| = 4$. Recall that $\Omega = \{\mathbf{x} : x_u \in \{0, 1, \dots, c\}, \sum_u x_u = m\}$. Define a Markov chain $\mathcal{M}(\text{Samp}(\alpha))$ for $\text{Samp}(\alpha)$ over Ω as follows. For any *ordered* pair (\mathbf{x}, \mathbf{y}) where $\mathbf{x}, \mathbf{y} \in \Omega$, we say (\mathbf{x}, \mathbf{y}) is a (u, u') -neighbor if $y_u = x_u - 1$, $y_{u'} = x_{u'} + 1$, and $x_v = y_v$ for every $v \notin \{u, u'\}$. The policy $\text{Samp}(\alpha)$ induces the following transition probability between all possible (u, u') -neighbors.

$$q(u, u') \doteq \alpha p_{(u, u')} + \frac{1-\alpha}{4} \sum_{v \in \mathcal{N}(u)} p_{(v, u')}.$$

Thus the transition matrix for the Markov chain $\mathcal{M}(\text{Samp}(\alpha))$ can be defined as follows.

$$P_\alpha(\mathbf{x}, \mathbf{y}) = q(u, u') \quad \text{iff } (\mathbf{x}, \mathbf{y}) \text{ is a } (u, u') \text{ neighbor} \quad (3)$$

Similarly, we can define a Markov chain $\mathcal{M}(\text{Rand-Perm}(\phi))$ over Ω for $\text{Rand-Perm}(\phi)$ as follows. For a given random permutation ϕ and $v \in \mathcal{N}(u)$, let $\phi(v)$ be the random order assigned by ϕ . Assume that $\text{Rand-Perm}(\phi)$ checks the neighbors $\mathcal{N}(u)$ in the order $(\phi^{-1}(1), \phi^{-1}(2), \phi^{-1}(3), \phi^{-1}(4))$. For a given state \mathbf{x} , location u and random order ϕ , the respective *supportive neighbor* of u is defined as $\mathcal{N}_{\mathbf{x}, \phi}(u) = \{v \in \mathcal{N}(u) : x_v = 0, x_{v'} = 0, \forall v' \in \mathcal{N}(v), \phi(v') < \phi(u)\}$. In other words, $\mathcal{N}_{\mathbf{x}, \phi}(u)$ includes those neighbors v of u such that when we are at state \mathbf{x} while a request r with origin v arrives, $\text{Rand-Perm}(\phi)$ will surely match r to a driver at u if present. The resultant transition matrix for the Markov chain $\mathcal{M}(\text{Rand-Perm}(\phi))$ is non-zero iff (\mathbf{x}, \mathbf{y}) is a (u, u') neighbor. The non-zero value is as follows.

$$P_\phi(\mathbf{x}, \mathbf{y}) = p_{(u, u')} + \sum_{r=(v, u'): v \in \mathcal{N}_{\mathbf{x}, \phi}(u)} p_r. \quad (4)$$

We use the following lemma about $\mathcal{M}(\text{Samp}(\alpha))$ and $\mathcal{M}(\text{Rand-Perm}(\phi))$. The proof is deferred to Appendix B.

Lemma 1. *Under the hot-spot assumption, Markov chains $\mathcal{M}(\text{Samp}(\alpha))$ and $\mathcal{M}(\text{Rand-Perm}(\phi))$ defined in (3) and (4) are both irreducible and aperiodic for any given $\alpha > 0$ and permutation function ϕ . Thus, from Theorem 5 in the primer on Markov chains, both $\mathcal{M}(\text{Samp}(\alpha))$ and $\mathcal{M}(\text{Rand-Perm}(\phi))$ admit a unique limiting distribution π^* , which coincides with the unique stationary distribution.*

4.1 Proof of Theorem 1

We prove this theorem by first showing the convergence rate for the objective (1) for both the policies. Finally we show how this can be adapted to give the convergence rates for the objective (2).

Convergence rate of Samp(α) for Objective (1). We first show the convergence rates for $\mathcal{M}(\text{Samp}(\alpha))$. From Lemma 1, we have that $\mathcal{M}(\text{Samp}(\alpha))$ is irreducible and aperiodic when $\alpha > 0$ and thus, it admits a unique limiting distribution (say π^*), regardless of the initial state. Suppose $\mathbf{X}^{(t)}$ is the random state after t steps (let the starting state be \mathbf{x}_0), which follows the distribution $P^t(\mathbf{x}^{(0)}, \cdot)$. Let $\mathbf{X}^{(\infty)}$ be the random state when $T \rightarrow \infty$, which follows the distribution π^* . For each $u, v \in \mathcal{U}$, define $\gamma_{u,v}$ as the probability that there is at least one driver at location u and there are less than c drivers at v in π^* . Thus by definition we have,

$$\gamma_{u,v} = \Pr[X_u^{(\infty)} \geq 1, X_v^{(\infty)} < c] = \sum_{\mathbf{x} \in \Omega: x_u \geq 1, x_v < c} \pi^*(\mathbf{x}).$$

For each u, v and t , define $\gamma_{u,v}^{(t)} := \Pr[X_u^{(t)} \geq 1, X_v^{(t)} < c]$. We prove the following Lemma 2 which we later use to prove the convergence rate in the main theorem.

Lemma 2. *There exists a scalar $C > 0$, independent of t and $\beta \in (0, 1)$ such that $|\gamma_{u,v}^{(t)} - \gamma_{u,v}| \leq 2C\beta^t$ for any $u, v \in \mathcal{U}$ and initial state \mathbf{x}_0 .*

Proof. Let $P^t(\mathbf{x}_0, \mathbf{y}) = \Pr[\mathbf{X}^{(t)} = \mathbf{y}]$ be the probability of reaching state \mathbf{y} after t steps with initial state \mathbf{x}_0 . Thus we have,

$$|\gamma_{u,v}^{(t)} - \gamma_{u,v}| = \left| \sum_{\mathbf{y}: y_u \geq 1, y_v < c} \Pr[\mathbf{X}^{(t)} = \mathbf{y}] - \gamma_{u,v} \right| \quad (5)$$

$$= \left| \sum_{\mathbf{y}: y_u \geq 1, y_v < c} P^t(\mathbf{x}_0, \mathbf{y}) - \sum_{\mathbf{z} \in \Omega: z_u \geq 1, z_v < c} \pi^*(\mathbf{z}) \right| \quad (6)$$

$$\leq \sum_{\mathbf{y}: y_u \geq 1, y_v < c} \left| P^t(\mathbf{x}_0, \mathbf{y}) - \pi^*(\mathbf{y}) \right| \quad (7)$$

$$\leq 2|P^t(\mathbf{x}_0, \cdot) - \pi^*|_{TV} \leq 2C\beta^t. \quad (8)$$

In equation (8), the first inequality is from the definition of total variation between two distributions while the second inequality is from the convergence theorem 4.9 in [18] (also Theorem 5 in primer on markov chains in Appendix A). \square

Convergence rate of $\mathcal{M}(\text{Rand-Perm}(\phi))$ for Objective (1). Recall that we proved in Lemma 1 that $\mathcal{M}(\text{Rand-Perm}(\phi))$ admits a unique limiting distribution. For notational convenience, we overload π^* and P to denote the limiting distribution and transition matrix of $\mathcal{M}(\text{Rand-Perm}(\phi))$. Set $\bar{\mathcal{N}}(u) = \{u\} \cup \mathcal{N}(u)$. Let $\mathbf{Y}^{(t)}$ be the random state in $\mathcal{M}(\text{Rand-Perm}(\phi))$ after t steps starting with \mathbf{y}_0 . For a given u and v let $\eta_{u,v}$ and $\eta_{u,v}^{(t)}$ be the respective probabilities, in the limiting distribution π^* and $\mathbf{Y}^{(t)}$, that there is at least one driver in $\bar{\mathcal{N}}(u)$ and less than c drivers at v . Therefore by definition we have that,

$$\eta_{u,v} = \Pr \left[\bigvee_{v \in \bar{\mathcal{N}}(u)} (Y_v^{(\infty)} \geq 1), Y_v^{(\infty)} < c \right] = \sum_{\mathbf{x}: \sum_{v \in \bar{\mathcal{N}}(u)} x_v \geq 1, x_v < c} \pi^*(\mathbf{x}),$$

$$\eta_{u,v}^{(t)} = \Pr \left[\bigvee_{v \in \bar{\mathcal{N}}(u)} (Y_v^{(t)} \geq 1), Y_v^{(t)} < c \right] = \sum_{\mathbf{x}: \sum_{v \in \bar{\mathcal{N}}(u)} x_v \geq 1, x_v < c} P^t(\mathbf{y}_0, \mathbf{x}).$$

Similar to Lemma 2, we have the following lemma for $\mathcal{M}(\text{Rand-Perm}(\phi))$.

Lemma 3. *There exists a scalar $C > 0$ independent of t and $\beta \in (0, 1)$ such that $|\eta_{u,v}^{(t)} - \eta_{u,v}| \leq 2C\beta^t$ for any $u, v \in \mathcal{U}$ and initial state \mathbf{y}_0 .*

The proof of Lemma 3 is essentially the same as that of Lemma 2 since it does not use the properties of $\text{Samp}(\alpha)$.

Now we have all ingredients to prove main Theorem 1.

Proof. We focus on applying Lemma 2 to prove the results for $\mathcal{M}(\text{Samp}(\alpha))$ in Theorem 1. A similar analysis applies to the case $\mathcal{M}(\text{Rand-Perm}(\phi))$ by using Lemma 3.

Consider the case $\sigma = \text{Samp}(\alpha)$ and a given $r = (u, u')$. Policy $\text{Samp}(\alpha)$ implies that $u_\sigma = \sigma(\mathbf{X}^{(t)}, r)$ will be equal to u with probability α and any neighbor in $\mathcal{N}(u)$ with probability $\frac{1-\alpha}{4}$. Thus, we claim that

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E} \left[\mathbf{I}(X_{u_\sigma}^{(t)} \geq 1, X_v^{(t)} < c) \right] &= \alpha \Pr[X_u^{(\infty)} \geq 1, X_v^{(\infty)} < c] + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \Pr[X_k^{(\infty)} \geq 1, X_v^{(\infty)} < c] \\ &= \alpha \gamma_{u,v} + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \gamma_{k,v} \end{aligned}$$

Let $W(\alpha) = \lim_{t \rightarrow \infty} W(\text{Samp}(\alpha), t)$. Thus we claim that

$$\begin{aligned} W(\alpha) &= \lim_{t \rightarrow \infty} \mathbb{E} \left[\sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \mathbf{I}(X_{u_\sigma}^{(t)} \geq 1, X_v^{(t)} < c) \right] \\ &= \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha \gamma_{u,v} + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \gamma_{k,v} \right). \end{aligned} \quad (9)$$

Now we bound the convergence rate of $W(\text{Samp}(\alpha), t)$. Let $\Delta(t) \doteq |W(\text{Samp}(\alpha), t) - W(\alpha)|$ (we omit the subscript of α in $\Delta(t)$ here) and $w \doteq \max_{r \in \mathcal{R}} w_r$. We have

$$\Delta(t) = \left| \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha \Pr[X_u^{(t)} \geq 1, X_v^{(t)} < c] + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \Pr[X_k^{(t)} \geq 1, X_v^{(t)} < c] \right) \right| \quad (10)$$

$$- \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha \gamma_{u,v} + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \gamma_{k,v} \right) \Big| \quad (11)$$

$$\leq \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha |\gamma_{u,v}^{(t)} - \gamma_{u,v}| + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} |\gamma_{k,v}^{(t)} - \gamma_{k,v}| \right) \quad (12)$$

$$\leq \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot (2C\beta^t) \quad (13)$$

$$\leq 2wC\beta^t. \quad (14)$$

Inequality (13) is a direct application of Lemma 2 and the fact that $|\mathcal{N}(u)| \leq 4$ for each u ; Inequality (14) is due to the fact that $\sum_{r \in \mathcal{R}} p_r \leq 1$.

Recall that $\widehat{W}(\alpha) = \lim_{T \rightarrow \infty} \widehat{W}(\text{Samp}(\alpha), T)$. We can verify that $\widehat{W}(\alpha) = W(\alpha)$. Now we bound the convergence rate of $\widehat{W}(\text{Samp}(\alpha), T)$. Let $\widehat{\Delta}(T) \doteq |\widehat{W}(\text{Samp}(\alpha), T) - \widehat{W}(\alpha)|$. Notice that we have the following.

$$\widehat{\Delta}(T) \leq \frac{1}{T} \sum_{t=1}^T \Delta(t) = \frac{1}{T} \sum_{t=1}^T 2wC\beta^t \leq \frac{2wC}{T} \frac{\beta}{1-\beta}. \quad (15)$$

This completes the proof. \square

4.2 A Special Case: Uniform Arrivals of Online Requests

We now focus on the policy, $\text{Samp}(\alpha)$, and consider a special case when the requests arrive as a uniform sample in each round (*i.e.*, $p_r = p$ for all $r \in \mathcal{R}$). Then we can compute explicit values of C and β . Note that in this case, the transition matrix for $\mathcal{M}(\text{Samp}(\alpha))$ is $P_\alpha(\mathbf{x}, \mathbf{y}) = \alpha p_{(u, u')} + \frac{1-\alpha}{4} \sum_{v \in \mathcal{N}(u)} p_{(v, u')} = p$.

Consider a given instance of the rideshare setting $\mathcal{I} = (\mathbf{x}_0, \mathbf{p}, \mathbf{w}, T)$ and $\text{Samp}(\alpha)$.

Let $W(\alpha) \doteq \lim_{t \rightarrow \infty} W(\text{Samp}(\alpha), t)$, and $\widehat{W}(\alpha) \doteq \lim_{T \rightarrow \infty} \widehat{W}(\text{Samp}(\alpha), T)$, where $W(\text{Samp}(\alpha), t)$ and $\widehat{W}(\text{Samp}(\alpha), T)$ are defined in Equations (1) and (2) respectively. From Lemma 1, we have that the limits in Equations (1) and (2) both exist and are the same for any given $\alpha > 0$ with the limiting value as in Equation (16). Define $\Delta_\alpha(t) \doteq |W(\alpha) - W(\text{Samp}(\alpha), t)|$ and $\widehat{\Delta}_\alpha(T) \doteq |\widehat{W}(\alpha) - \widehat{W}(\text{Samp}(\alpha), T)|$. From Equations (9), (14) and (15) in the proof of Theorem 1 we have

$$W(\alpha) = \widehat{W}(\alpha) = \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha \gamma_{u,v} + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \gamma_{k,v} \right). \quad (16)$$

$$\Delta_\alpha(t) \leq 2wC\beta^t, \quad \widehat{\Delta}_\alpha(T) \leq \frac{1}{T} \sum_{t=1}^T \Delta_\alpha(t) \leq \frac{2wC}{T} \frac{1}{1-\beta}. \quad (17)$$

In Equation (16), $\gamma_{u,v}$ is the probability that location u has at least one driver and location v has fewer than c drivers in the limiting distribution π^* . In Inequalities (17), $w \doteq \max_{r \in \mathcal{R}} w_r$ and $C > 0$ and $\beta \in (0, 1)$ are two scalars which are independent of t and T , but may be related to parameters m and n . This yields the result for $\text{Samp}(\alpha)$; we can get similar results for the policy $\text{Rand-Perm}(\phi)$.

Observe that if (\mathbf{x}, \mathbf{y}) is a (u, u') -neighbor, then (\mathbf{y}, \mathbf{x}) is a (u', u) -neighbor and therefore we have $P_\alpha(\mathbf{y}, \mathbf{x}) = p$, implying that P_α is symmetric. This implies that $\mathcal{M}(\text{Samp}(\alpha))$ admits a unique uniform limiting distribution π^* over Ω (details in Appendix A). This enables us to get a closed-form expression for $W(\alpha)$ (Theorem 2), and derive explicit values of C and β (Theorem 3).

Theorem 2 (Closed-form expression for $W(\alpha)$). *Consider the uniform arrival distribution of the requests such that $p_r = p$ for all $r \in \mathcal{R}$ and $c = m$. Objectives defined in Equations (1) and (2) for $\text{Samp}(\alpha)$ both converge to $\frac{mp}{n+m-1} \sum_{r \in \mathcal{R}} w_r$ for any given $\alpha > 0$.*

The condition $c = m$ in Theorem 2 implies that there is no constraint on the maximum number of drivers in any location. The main idea of proof for Theorem 2 is to show that for any (u, v) , $\gamma_{u,v} = \frac{m}{n+m-1}$ for the special case. We defer the full proof of Theorem 2 to Appendix B.

Theorem 3 (Explicit values of C and β). *Consider the uniform arrival distribution of the requests such that $p_r = n^{-2}$ for all $r \in \mathcal{R}$ and $c \leq 2$. For any given $\alpha > 0$, Objectives in (1) and (2) for $\text{Samp}(\alpha)$ both converge with a rate upper bounded by (*i.e.*, as least as fast as) $\frac{4m \sum_{r \in \mathcal{R}} w_r}{n^2 \exp(t/n^2)}$ and $\frac{4m \sum_{r \in \mathcal{R}} w_r}{T}$, respectively.*

We first show the following useful lemma⁸. Consider the Markov Chain $\mathcal{M}(\text{Samp}(\alpha))$ as defined in (3) and let $\tau(\epsilon)$ be the mixing time of $\mathcal{M}(\text{Samp}(\alpha))$ as defined in Section 2.

Lemma 4. $\tau(\epsilon) \leq n^2 \ln(2m/\epsilon)$ when the arrival distribution of online requests is uniform such that $p_r = n^{-2}$ for all $r \in \mathcal{R}$ and $c \leq 2$.

We defer the full proof of Lemma 4 to Appendix B. Now we show the full proof of Theorem 3.

⁸ This can be viewed as a special case of Theorem 5 in Appendix A.

Proof. Setting $\tau(\epsilon) = t$ implies that $\epsilon \geq 2m \exp(-t/n^2)$ in Lemma 4. Let $\epsilon = 2m \exp(-t/n^2)$. Recall that the definition of $\tau(\epsilon) = \min\{t : d(t') \leq \epsilon, \forall t' \geq t\}$ where $d(t) = \max_{\mathbf{x} \in \Omega} \|P^t(\mathbf{x}, \cdot) - \pi^*\|_{TV}$. Thus, we have

$$\max_{\mathbf{x} \in \Omega} \|P^t(\mathbf{x}, \cdot) - \pi^*\|_{TV} \leq 2m (\exp(-1/n^2))^t.$$

Thus, we have that *when the arrival distribution of online requests is uniform such that $p_r = n^{-2}$ for all $r \in \mathcal{R}$ and $c \leq 2$, the scalars in Theorem 5 are explicit, namely, $C = 2m$ and $\beta = \exp(-1/n^2)$.* By Lemma 2, we have

$$|\gamma_{u,v}^{(t)} - \gamma_{u,v}| \leq 2C\beta^t = 4m \exp(-t/n^2).$$

Recall that $W(\alpha) = \lim_{t \rightarrow \infty} W(\text{Samp}(\alpha), t)$ and $\Delta_\alpha(t) = |W(\text{Samp}(\alpha), t) - W(\alpha)|$. Thus the RHS in Equation (14) can be upper-bounded by $\frac{4m \exp(-t/n^2)}{n^2} \sum_{r \in \mathcal{R}} w_r$. Similarly, we have that

$$\hat{\Delta}_\alpha(T) \leq \frac{1}{T} \sum_{t=1}^T \Delta_\alpha(t) \leq \frac{4m \sum_{r \in \mathcal{R}} w_r}{T \cdot n^2} \frac{1}{1 - \exp(-1/n^2)} \sim \frac{4m \sum_{r \in \mathcal{R}} w_r}{T}. \quad \square$$

4.3 A Lower Bound on Convergence Rates

We show that even in very special cases, the convergence rate shown in Theorem 3 is *almost* optimal. When $c = 1$, we have the following lower bound where the dependence on t and T nearly matches our upper bounds.

Theorem 4 (Lower bound for $\Theta(\beta^t)$). *There is an instance with uniform arrival distribution where $p_r = n^{-2}$ for all $r \in \mathcal{R}$ and $c = 1$ such that Objective (1) for $\text{Samp}(1)$ has an asymptotic convergence rate of $\frac{2m \sum_{r \in \mathcal{R}} w_r}{n^3 \exp(t/n)}$.*

Proof. Consider the following instance. Let $p_r = 1/n^2$ for all $r \in \mathcal{R}$ with $|\mathcal{R}| = n^2$ and $c = 1$. Thus the state space $\Omega = \{\mathbf{x} = \{0, 1\}^n : \sum_u x_u = m\}$. Suppose we index the n locations of \mathcal{U} as $u = 1, 2, \dots, n$ and we set the initial state $\mathbf{x}^{(0)}$ as $x_u^{(0)} = 1$ for $1 \leq u \leq m$ and $x_u^{(0)} = 0$ for $u > m$. In other words, we have that at $t = 0$, the first m locations has one driver each and the rest have no driver. Recall that from assumptions we have that $m \ll cn$ which simplifies to $m \ll n$. Set $w_r = 1$ if $\bar{r} = (m+1, 1)$ (with origin of $m+1$ and destination of 1) and $w_r = 0$ otherwise. Note that by definition of $\text{Samp}(\alpha = 1)$, we will check the availability of driver at location u with probability 1 when a request $r = (u, *)$ arrives. Since all request brings a profit of 0 except $\bar{r} = (m+1, 1)$, the only source of profits is when $\bar{r} = (m+1, 1)$ arrives, the system has one drive at location $m+1$ but no drive at location 1. Thus, the limiting performance in (1) is $W(\text{Samp}(1)) = \frac{1}{n^2} \gamma_{m+1,1}$, where $\gamma_{m+1,1}$ refers to the probability that the system has one drive at location $m+1$ while no drive at location 1 in the limiting distribution.

As before, from Lemma 1 we have that $\mathcal{M}(1)$ admits a unique uniform distribution π^* over Ω . Thus

$$\gamma_{m+1,1} = \sum_{\mathbf{x}: x_{m+1} \geq 1, x_1 < 1} \pi^*(\mathbf{x}) = \frac{\binom{n-2}{m-1}}{\binom{n}{m}} = \frac{m}{n} \frac{n-m}{n-1}.$$

Consider the locations $m+1$ and 1 and let $\gamma_{m+1,1,t} = \Pr[X_{m+1}^{(t)} \geq 1, X_1^{(t)} < 1]$, where $\mathbf{X}^{(t)}$ is the random state in Markov Chain $\mathcal{M}(\text{Samp}(\alpha = 1))$ as defined in (3) with initial state $\mathbf{x}^{(0)}$. Now we compute an explicit formula for $\gamma_{m+1,1,t}$.

Let $u^* = 1$ and $v^* = m+1$. There are 4 possible states for the number of drivers in u and v : $\mathbf{s}_1 = (1, 0)$, $\mathbf{s}_2 = (1, 1)$, $\mathbf{s}_3 = (0, 0)$ and $\mathbf{s}_4 = (0, 1)$, where \mathbf{s}_1 denotes that there is one driver at $u = 1$ and no driver at $m+1$.

We define a Markov chain $\mathcal{M}(u^*, v^*)$ over these 4 states as follows

$$P_{(u^*, v^*)} = \begin{pmatrix} \frac{n^2-n+1}{n^2} & \frac{m-1}{n^2} & \frac{n-m-1}{n^2} & \frac{1}{n^2} \\ \frac{n-m}{n^2} & \frac{n^2-2(n-m)}{n^2} & 0 & \frac{n-m}{n^2} \\ \frac{m}{n^2} & 0 & \frac{n^2-2m}{n^2} & \frac{m}{n^2} \\ \frac{1}{n^2} & \frac{m-1}{n^2} & \frac{n-m-1}{n^2} & \frac{n^2-n+1}{n^2} \end{pmatrix}$$

Consider the value $P(\mathbf{s}_1, \mathbf{s}_2)$. From state \mathbf{s}_1 to \mathbf{s}_2 , we have that this event occurs iff a request ends at $m+1$ but starts at some location with one driver. Since the system has m drivers in total, we have that there are $m-1$ request types contributing to the transition from \mathbf{s}_1 to \mathbf{s}_2 . Thus $P(\mathbf{s}_1, \mathbf{s}_2) = (m-1)/n^2$. Similarly we can get the other values.

Note that for $\mathcal{M}(u^*, v^*)$, we start with state \mathbf{s}_1 and therefore $\gamma_{m+1,1,t} = P_{(u^*, v^*)}^t(\mathbf{s}_1, \mathbf{s}_4)$. We have that

$$|\gamma_{m+1,1,t} - \gamma_{m+1,1}| = \left(1 - \frac{1}{n}\right)^t \cdot \frac{2mn - n - 2m^2}{n(n-2)} + \left(1 - \frac{2}{n} + \frac{2}{n^2}\right)^t \cdot \frac{(m-1)(n-m-1)}{(n-1)(n-2)}. \quad (18)$$

$$\geq \left(1 - \frac{1}{n}\right)^t \cdot \left(\frac{2m}{n} - \frac{1}{n} - \frac{2m^2}{n^2}\right) + \left(1 - \frac{2}{n} + \frac{2}{n^2}\right)^t \cdot \left(\frac{m}{n} - \frac{1}{n}\right)\left(1 - \frac{m}{n} - \frac{1}{n}\right). \quad (19)$$

When $1 \ll m \ll n$ and $t \geq 10n$ the above expression can be approximated by

$$|\gamma_{m+1,1,t} - \gamma_{m+1,1}| \sim \exp(-t/n) \frac{2m}{n} + \exp(-2t/n) \frac{m}{n} \sim \exp(-t/n) \frac{2m}{n}.$$

Therefore we have,

$$\Delta_t = |W(\text{Samp}(1)) - W(\text{Samp}(1), t)| = \frac{1}{n^2} |\gamma_{m+1,1} - \gamma_{m+1,1,t}| \sim \frac{2m \sum_{r \in \mathcal{R}} w_r}{n^3 \exp(t/n)}. \quad \square$$

5 Experiments

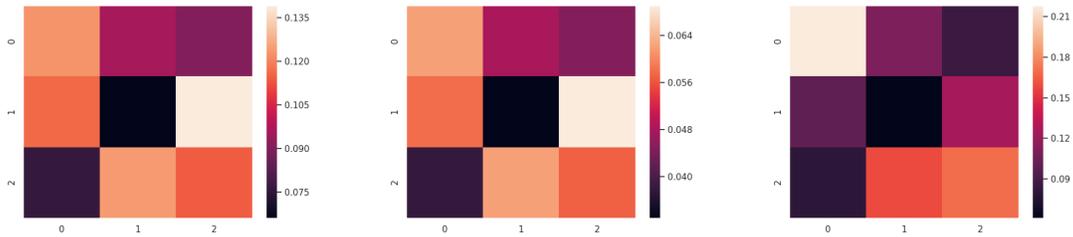


Fig. 3: 3×3 grids; black to white represent low to high values. Plots 1, 2, 3 represent the distribution of requests with either starting or ending at a cell, requests starting at a cell and percentage of time spent by cars in a cell respectively by the optimal policy.

In this section, we present experimental results on simulated synthetic data and a real-world dataset⁹: the trip records for Jan 2013 of New York yellow cabs. We aim to experimentally study the convergence properties of objectives in Equations (1) and (2) for various algorithms. We consider three kinds of algorithms;

⁹ <http://www.andresmh.com/nyctaxitrips/>

Samp($\alpha = 0.8$) and Rand-Perm(ϕ) described in the main section of the paper, where ϕ is chosen as the clockwise permutation and the Greedy algorithm whose action is defined as follows. When a request r arrives at time t with the system in state $\mathbf{X}^{(t)}$, we choose the neighbors such that the number of cars in state $\mathbf{X}^{(t)}$ are in decreasing order. Where appropriate, we also compare against a standard RL [32] algorithm, namely, Deep-Q-Learning [26], with agent and environment adapted from OpenAI Gym [7] for 5 episodes of 2500 steps and experience replay.¹⁰ Fig. 3 represents exploratory analysis using optimal policy obtained via value iteration.¹¹

Preprocessing. The records are pre-processed by restricting to trips whose starting and ending coordinates lie in the range of (40.7014, 40.8024) for latitudes and (-74.0041, -73.9552) for longitudes. We then split the records into 31 parts with one part per date and further split them into Morning, Afternoon, Evening segments, corresponding to periods 7:00AM-11:00AM, 11:00AM-3:00PM, and 3:00PM-7:00PM, respectively. For brevity we only show results on the Morning segment; Appendix C describes the other results. The Manhattan area is divided into a grid of 21×11 . Using latitude-longitude coordinates, we map every record to a *request* in this grid.

Simulated experiments. In the first experiment, we consider a grid of 5×5 cells and a total of 10 cars with no accumulation limit on the number of cars in each location. We consider¹² $T = 10^4$ and run 10^3 independent experiments to compute the Obj. (1) and (2). Each algorithm starts at the same *adversarial* initial state (all 10 cars are allocated in one single location) for every experiment. At each time step t , a request r is sampled *uniformly* at random. Figs. 4 and 5 show the convergence of Obj. (1) and (2) for the three algorithms. In particular, we observe that the Obj. (1) and (2) of Samp(0.8) both converge to the exact value as predicted in Theorem 2 (shown as the black line in the figure).

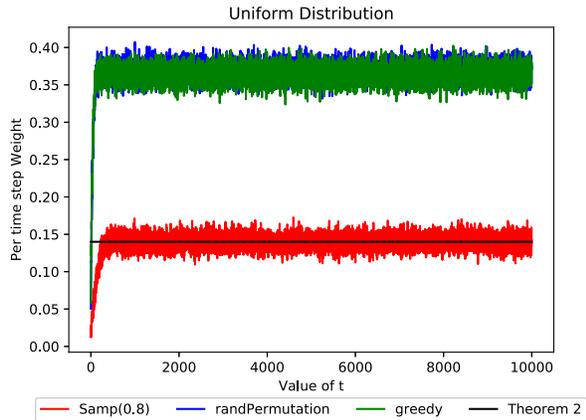


Fig. 4: Obj. (1) on simulation with Uniform Arrivals

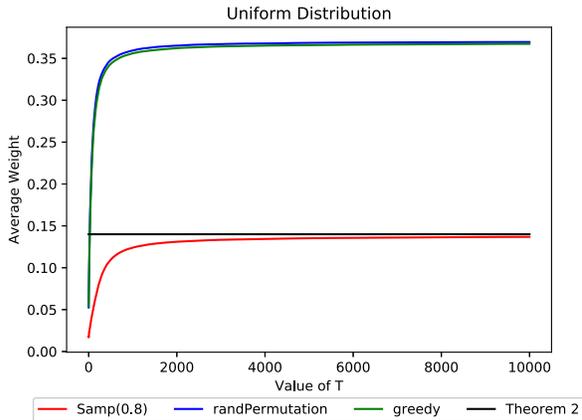


Fig. 5: Obj. (2) on simulation with Uniform Arrivals

Synthetic experiments using learnt parameters. Here, we “learn” the parameters from the processed taxi dataset, but *simulate* the requests from these learnt parameters. In particular, we first learn the arrival rates of each of the requests using the data from Jan 2013 using a standard unbiased estimator. We then set the *weight* for each request as the distance between the starting and ending location. At each time step t , a request r is drawn from this learnt distribution. The grid size and the number of cars are identical to the real-world experiments (described below). Figs. 6 and 7 show the convergence of Obj. (1) and (2) while Figs. 9 and 10 show the corresponding absolute error bound. In Fig. 9 we fit a statistical best-fit exponential curve using least squares and in Fig. 10 we fit a statistical best-fit linear function. This corroborates our

¹⁰ <https://github.com/keon/deep-q-learning>

¹¹ https://www.ics.uci.edu/~dechter/publications/r42a-mdp_report.pdf

¹² We consider $T = 10^4$ arrivals within a 3-hour window. Additionally, the trips take multiple time-steps to complete.

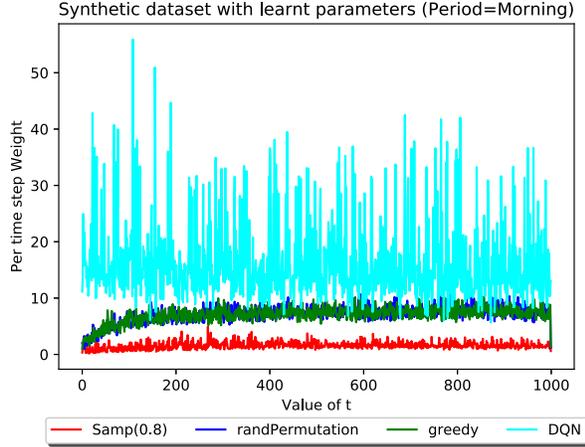


Fig. 6: Obj. (1), synthetic data, parameters learnt from real data

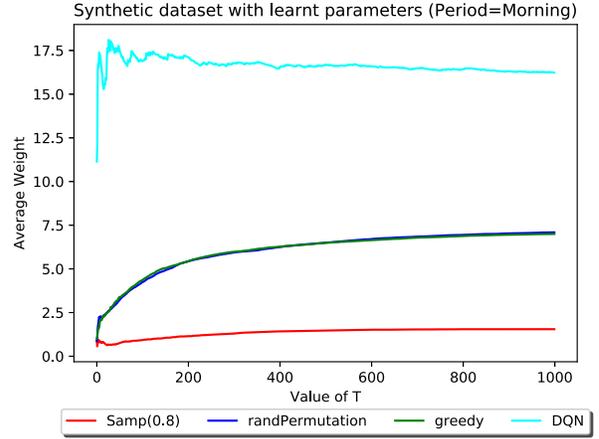


Fig. 7: Obj. (2), synthetic data, parameters learnt from real-data

Main Theorem 1 as we see the same asymptotic behavior. Note that the parameters in Figs. 8 and 7 are the same except the arrival assumption, thus the sharp difference does shed light on the importance of arrival assumption for convergence.

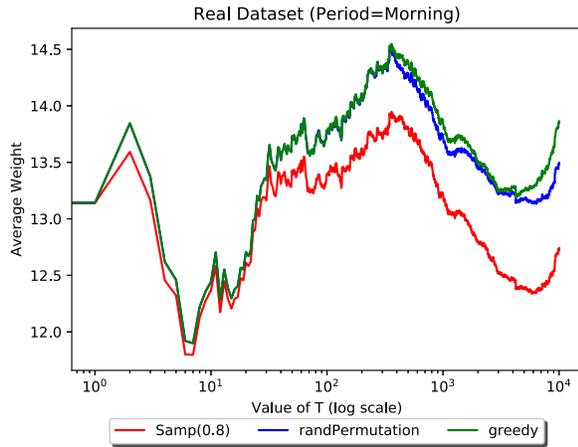


Fig. 8: Obj. (2) on real dataset. x -axis in log-scale

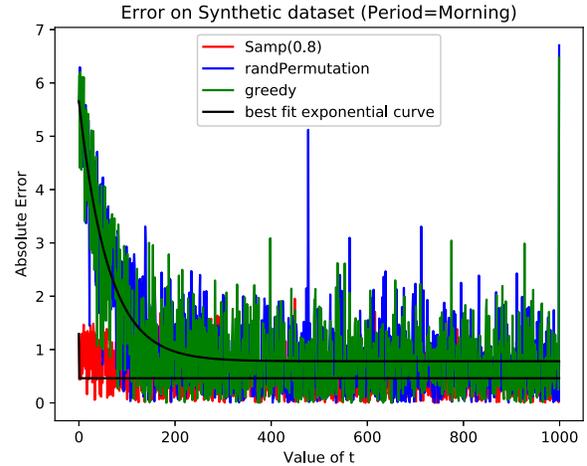


Fig. 9: Error bound of Obj. (1) on synthetic data

Real-world experiments. First, a random subset of 5000 cars from the dataset is chosen and records corresponding to trips of these cars are sub-sampled. We use the real-time per-second arrival sequence of this sub-sampled dataset with 5000 cars and a maximum congestion capacity of 50 cars per cell as the experimental parameters. We study the convergence properties of Obj. (2) on this real-world dataset, for each period Morning, Afternoon and Evening. We use $T = 10^4$ and use the 31 days as “independent” runs to average our results over. Figs. 8 and 11 show the evolution of Obj. (2) for the Morning period with x -axis having log-scale spacing and the corresponding error bound respectively. Fig. 11 uses the average of all T rounds as the target convergence value.

Further discussion. Our simulations show that even when starting in a pathologically bad state, all algorithms show the “correct” convergence behavior on both objectives (Figs. 4, 5, 6 and 7). In Fig. 6, our

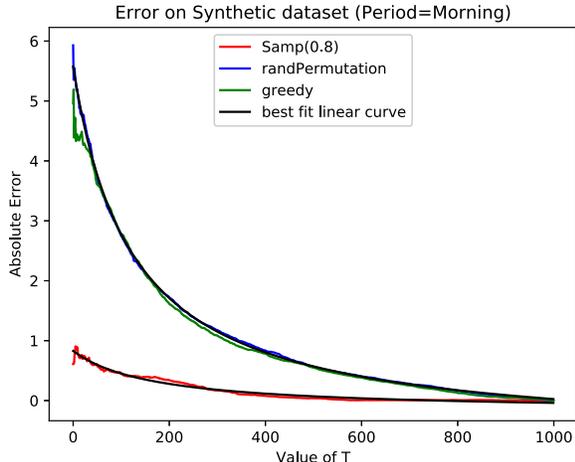


Fig. 10: Error bound of Obj. (2) on synthetic data

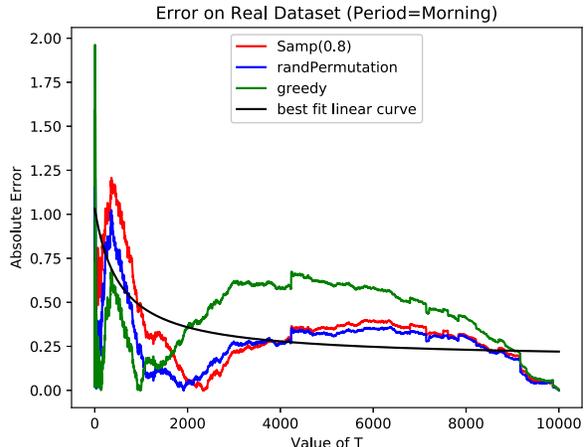


Fig. 11: Error bound of Obj. (2) on real dataset

state-dependent policy does almost as well as the standard RL algorithm while having *good* convergence properties. DQN, on the other hand, wildly oscillates even after running for a large number of time-steps, thus supporting our main message on the importance of convergence. Additionally, algorithms proposed in this paper are *simple* and extremely fast compared to DQN. On the real dataset, there are huge variations in the beginning steps and eventually the values do not vary a lot. However, they still do not fully converge like in the IID setting (Figs. 8 and 11). This is mainly due to the fact that in practice the distribution is not *truly* identical over all time steps and that full convergence is unlikely (*e.g.*, we provided such an example of non-i.i.d. arrivals in preliminaries). Nonetheless, the theory proved with the appropriate assumptions gives good insight into the performance in practice. In supplementary C, we show results for the Afternoon and Evening periods; the qualitative behaviors are the same. The fluctuations in the beginning also justify the importance of reaching a stationary distribution since platforms are *risk-averse* and would prefer schemes with guaranteed throughput.

6 Conclusions & Future Research

In this paper, we considered the rideshare dispatching problem and the corresponding Markov chain for the evolution of the matching process. In turn, we characterized the driver distribution over time. In particular, we showed that the mixing time of this Markov chain is small, and gave explicit bounds on the convergence rate of various algorithms. Under practical assumptions, our theory shows that the convergence rates are fast, and in practice one would reach the stationary distribution in a short amount of time. To complement the theory we ran extensive experiments on both simulated and real datasets. Our theory gives accurate insight into rideshare dynamics in practice, and also complements a growing body of related research (*e.g.*, [28, 4, 5]) that relies on assumptions that were, until now, underexplored and poorly understood. Still, relaxing any of the five commonly-made assumptions discussed in the preliminaries (§2) would be practically useful.

Acknowledgments. John Dickerson and Michael Curry were supported in part by NSF CAREER Award IIS-1846237 and DARPA SI3-CMD Award S4761. Aravind Srinivasan’s research was supported in part by NSF CNS-1010789, CCF-1422569 and CCF-1749864, and by research awards from Adobe, Inc. The research of Karthik Sankararaman and Pan Xu was supported in part by NSF Awards CNS 1010789 and CCF 1422569. Yuhao Wan was supported via an REU grant, NSF CCF-1852352, and was advised by John Dickerson. John Dickerson and Aravind Srinivasan were both supported by a gift from Google.

Appendix

A Primers on Markov Chains

In this section we provide a brief background on important concepts and results from Markov chains. For more details, please refer to these excellent monographs ([18, 13]). Throughout, we denote a Markov Chain by \mathcal{M} with state space Ω and transition matrix P . Additionally, we assume that the number of states in \mathcal{M} is finite (i.e., $|\Omega| < \infty$).

Irreducible Markov Chains. \mathcal{M} is said irreducible if for any two states $\mathbf{x}, \mathbf{y} \in \Omega$, there exists an integer t such that $P^t(\mathbf{x}, \mathbf{y}) > 0$. In other words, it is possible to start at any state and reach any other state with a positive probability after finite number of transitions.

Aperiodic Markov Chains. For any state \mathbf{x} , let $\mathcal{T}(\mathbf{x}) \doteq \{t \geq 1, P^t(\mathbf{x}, \mathbf{x}) > 0\}$. The period of state \mathbf{x} is defined to be the greatest common divisor of the entries in the set $\mathcal{T}(\mathbf{x})$. A Markov Chain \mathcal{M} is said to be aperiodic if all states have period 1.

Lemma 1.6 in [18] shows that all states in an irreducible chain have the same period. Thus to prove an irreducible chain \mathcal{M} is aperiodic, it suffices to show that there exists a state whose period is 1.

For two distributions μ and ν on the state space Ω , the *total variation distance* between μ and ν is defined as $\|\mu - \nu\|_{TV} = \frac{1}{2} \sum_{\mathbf{x} \in \Omega} |\mu(\mathbf{x}) - \nu(\mathbf{x})|$.

Definition 2 (Stationary Distribution). A distribution π^* is said to be a stationary distribution if $\pi^* = \pi^* P$.

Definition 3 (Limiting Distribution). A distribution π^* is said to be a limiting distribution if for every initial state \mathbf{x} , we have $\lim_{t \rightarrow \infty} \|P^t(\mathbf{x}, \cdot) - \pi^*\|_{TV} = 0$. Here $P^t(\mathbf{x}, \cdot)$ denotes that distribution over when after starting at \mathbf{x} and running the process for t steps.

Stationary and Limiting Distributions of Irreducible and Aperiodic Markov Chains. Any (finite-state) irreducible and aperiodic Markov Chain admits a unique stationary distribution. Additionally this is also the unique limiting distribution. More formally, we can prove the following convergence theorem.

Theorem 5. (Convergence Theorem 4.9 in [18]) Suppose a (finite-state) Markov Chain \mathcal{M} with transition matrix P is irreducible and aperiodic with stationary distribution π^* . Then there exists a constant $C > 0$ and $\beta \in (0, 1)$ such that

$$\max_{\mathbf{x} \in \Omega} \|P^t(\mathbf{x}, \cdot) - \pi^*\|_{TV} \leq C\beta^t.$$

Theorem 5 implies that the total variation distance between the distribution after t steps and the unique stationary distribution exponentially decreases, irrespective of the initial state. Note that the terms C and β may depend on the parameters of the Markov Chain (e.g., the size of $|\Omega|$) and is independent of t .

When is the Distribution π^* Uniform? For any irreducible and aperiodic Markov Chain \mathcal{M} , its unique stationary distribution π^* is uniform over Ω iff in the transition matrix P , the sum of every column is equal to 1. In particular, any symmetric matrix P has all columns sum to 1, since the sum of every row in P is equal to 1.

Rapid Mixing via coupling arguments. Coupling and path coupling are common techniques used to prove that a Markov Chain is rapidly mixing. A (causal) coupling of the Markov Chain \mathcal{M} is a joint process $(\mathcal{X}, \mathcal{Y}) = (\mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$ on $\Omega \times \Omega$, such that each of the process \mathcal{X} and \mathcal{Y} is a faithful copy of \mathcal{M} , if considered marginally. In other words, we require that, for all possible $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}' \in \Omega$, we have that

$$\Pr[\mathbf{X}^{(t+1)} = \mathbf{x}' | \mathbf{X}^{(t)} = \mathbf{x}] = P(\mathbf{x}, \mathbf{x}'), \quad \Pr[\mathbf{Y}^{(t+1)} = \mathbf{y}' | \mathbf{Y}^{(t)} = \mathbf{y}] = P(\mathbf{y}, \mathbf{y}').$$

Lemma 5 (Path Coupling Lemma 6.3 in [13]). Let \mathbf{d} be an integer-valued metric on $\Omega \times \Omega$ which has range over $\{0, 1, \dots, D\}$. Let \mathcal{S} be a subset of $\Omega \times \Omega$ such that for each $(\mathbf{x}, \mathbf{y}) \in \Omega \times \Omega$, there exists a path $\mathbf{x} = \mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_t = \mathbf{y}$ connecting \mathbf{x} and \mathbf{y} where $(\mathbf{z}_\ell, \mathbf{z}_{\ell+1}) \in \mathcal{S}$ for all $0 \leq \ell < t$ and $\mathbf{d}(\mathbf{x}, \mathbf{y}) = \sum_{\ell=0}^{t-1} \mathbf{d}(\mathbf{z}_\ell, \mathbf{z}_{\ell+1})$. In other words, we just need to specify a subgraph \mathcal{H} with vertex set Ω and edge set \mathcal{S} and define a metric over \mathcal{S} . Then $\mathbf{d}(\mathbf{x}, \mathbf{y})$ is simply the shortest path between \mathbf{x} and \mathbf{y} over \mathcal{H} . Suppose we design a coupling $(\mathbf{X}, \mathbf{Y}) \rightarrow (\mathbf{X}', \mathbf{Y}')$ of the original Markov Chain \mathcal{M} , which is defined over all pairs $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$ such that there exists a $\beta \in [0, 1)$ such that $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq \beta \mathbf{d}(\mathbf{X}, \mathbf{Y})$. Then the mixing time $\tau(\epsilon)$ of \mathcal{M} satisfies $\tau(\epsilon) \leq \frac{\ln(D/\epsilon)}{1-\beta}$.

B Missing Proofs in Main Sections

Proof of Lemma 1.

Lemma 6. (Restated Lemma 1). Under the hot-spot assumption, Markov chains $\mathcal{M}(\alpha)$ and $\mathcal{M}(\phi)$ defined in (3) and (4) are both irreducible and aperiodic for any given $\alpha > 0$ and ϕ .

It is well-known (see Theorem 5 in the primer) that a Markov chain that is both irreducible and aperiodic admits a unique limiting distribution π^* , which coincides with the unique stationary distribution.

For this proof we refer to the transition matrix P_α of $\mathcal{M}(\alpha)$ defined in (3) as P for notation simplicity (i.e., we drop the subscript P_α).

Proof. $\mathcal{M}(\alpha)$ is irreducible. From the definition of $P(\cdot)$ in (3), we have that for any (u, u') -neighbor (\mathbf{x}, \mathbf{y}) , the inequality $P(\mathbf{x}, \mathbf{y}) \geq \alpha p_{(u, u')}$ holds¹³. Thus from the condition stated in Lemma 1, we get that $P(\mathbf{x}, \mathbf{y}) > 0$ for any (u, u') -neighbor (\mathbf{x}, \mathbf{y}) with either $u = u^*$ or $u' = u^*$.

In what follows, we assume that $u \neq u^*$ and $u' \neq u^*$. Consider any two given states \mathbf{x}, \mathbf{y} in Ω . We show that there exists a positive integer $t > 0$ such that $P^t(\mathbf{x}, \mathbf{y}) > 0$. This can be proved using the following cases.

- **Case $\mathbf{x} \neq \mathbf{y}$.** W.l.o.g. assume that (\mathbf{x}, \mathbf{y}) is a (u, u') -neighbor such that $u \neq u^*$ and $u' \neq u^*$. Let $x_u = a, x_{u'} = b, x_{u^*} = d$. From the definition of (u, u') -neighbor, we have that $y_u = a - 1, y_{u'} = b + 1$ and $y_{u^*} = d$. The values on the ordered pairs of locations (u, u', u^*) can be enumerated for various values of d .

1. First case is when $d < c$, where c is the capacity of the number of drivers in each location. Observe that there exists a feasible path from \mathbf{x} and \mathbf{y} , namely,

$$\mathbf{x} = (a, b, d) \rightarrow \mathbf{x}' = (a - 1, b, d + 1) \rightarrow \mathbf{y} = (a - 1, b + 1, d).$$

Therefore we have that $P^2(\mathbf{x}, \mathbf{y}) > P(\mathbf{x}, \mathbf{x}') \cdot P(\mathbf{x}', \mathbf{y}) > 0$ since $(\mathbf{x}, \mathbf{x}')$ is a (u, u^*) -neighbor while $(\mathbf{x}', \mathbf{y})$ is a (u^*, u') -neighbor.

2. The other case is when $d = c$. Then there exists another feasible path connecting \mathbf{x} and \mathbf{y} , namely,

$$\mathbf{x} = (a, b, c) \rightarrow \mathbf{x}' = (a, b + 1, c - 1) \rightarrow \mathbf{y} = (a - 1, b + 1, c).$$

Once again we have that $P^2(\mathbf{x}, \mathbf{y}) > 0$.

Note that in both these cases we critically used the *hot-spot* assumption.

- **Case $\mathbf{x} = \mathbf{y}$.** W.l.o.g. assume that $x_u = a$ with $a < c$ for some $u \neq u^*$ and $x_{u^*} = b$. In this case we enumerate the various values for the ordered pair (u, u^*) .

¹³ A similar corresponding property holds for $\mathcal{M}(\phi)$.

1. ($b = 0$). We consider a few different scenarios in this case. The easiest case is when $a \neq 0$. In this case we have the following non-zero path from \mathbf{x} to \mathbf{x} namely,

$$\mathbf{x} = (a, b) \rightarrow \mathbf{x}' = (a - 1, b + 1) \rightarrow \mathbf{x} = (a, b).$$

Using the *hot-spot* assumption, we have that $P^2(\mathbf{x}, \mathbf{x}) > P(\mathbf{x}, \mathbf{x}') \cdot P(\mathbf{x}', \mathbf{x}) > 0$.

The second case is when $a = 0$ while at least one of $v \in \mathcal{N}(u)$ has a driver. We have a non-zero path in this case, with the changes in the values of the ordered pair (u, v, u^*) as follows.

$$\mathbf{x} = (0, h \neq 0, b) \rightarrow \mathbf{x}' = (0, h - 1, b + 1) \rightarrow \mathbf{x} = (0, h, b).$$

Using the *hot-spot* assumption, we have that $P^2(\mathbf{x}, \mathbf{x}) > P(\mathbf{x}, \mathbf{x}') \cdot P(\mathbf{x}', \mathbf{x}) > 0$.

The last case is when no drivers exist in u and its neighborhood $\mathcal{N}(u)$. In this case, any algorithm that dispatches using the drivers in this location will not find a driver and hence $P(\mathbf{x}, \mathbf{x}) = 1 > 0$ for all algorithms.

2. ($0 < b$). Then observe that there exists a feasible loop on \mathbf{x} as follows (using the *hot-spot* assumption).

$$\mathbf{x} = (a, b) \rightarrow \mathbf{x}' = (a + 1, b - 1) \rightarrow \mathbf{x} = (a, b).$$

Therefore we have $P^2(\mathbf{x}, \mathbf{x}) > P(\mathbf{x}, \mathbf{x}') \cdot P(\mathbf{x}', \mathbf{x}) > 0$.

$\mathcal{M}(\alpha)$ is **aperiodic**. From Lemma 1.6 in [18], it suffices to show that there exists a state with period 1. We do so by considering two cases on the relationship between m and c .

- **Case $m \geq c$.** Consider a given state \mathbf{x} with $x_{u^*} = c$. We have that $P(\mathbf{x}, \mathbf{x}) \geq \sum_{u \neq u^*} p_{(u, u^*)} > 0$. The above inequality is valid for the following reason. When we are at state \mathbf{x} with $x_{u^*} = c$, then for any online request type r with ending location u^* , any algorithm has to reject r , and hence we remain in state \mathbf{x} for the next round.
- **Case $m < c$.** In this case, we consider a state \mathbf{x} with $x_{u^*} = m$ while $x_u = 0$ for all $u \neq u^*$. Like before we have that $P(\mathbf{x}, \mathbf{x}) \geq \sum_{u \neq u^*} \alpha p_{(u, u^*)} > 0$.

Thus for each of the two cases we can always identify a state \mathbf{x} with $P(\mathbf{x}, \mathbf{x}) > 0$, implying that \mathbf{x} has a period of 1. \square

Proof of Theorem 2.

Proof. When $c = m$, we have that $\Omega = \{\mathbf{x} \in \mathbb{Z}_+^n : \sum_u x_u = m, x_u \geq 0, \forall u\}$. Note that

$$|\Omega| = \binom{m+n-1}{n-1} = \binom{m+n-1}{m}.$$

Let γ_u be the probability that location u has at least one driver in the limiting distribution. We have that

$$\begin{aligned} \gamma_u &= \Pr[X_u^{(\infty)} \geq 1] = \sum_{\mathbf{x}: x_u \geq 1} \pi^*(\mathbf{x}). \\ &= 1 - \frac{|\{\mathbf{x} \in \Omega : x_u = 0\}|}{|\Omega|} = 1 - \frac{\binom{m+n-2}{m}}{\binom{m+n-1}{m}} = \frac{m}{n+m-1}. \end{aligned}$$

Thus we claim that $\gamma_u = \frac{m}{n+m-1}$ for all u . Notice that in our case $c = m$, $\gamma_{u,v} = \gamma_u$ for all $u, v \in \mathcal{U}$ since if there is a driver at some location u , it surely implies that all other locations have less than c drivers. Therefore,

$$\widehat{W}(\text{Samp}(\alpha)) = \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha \gamma_{u,v} + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \gamma_{k,v} \right) \quad (20)$$

$$= \sum_{r=(u,v) \in \mathcal{R}} p_r \cdot w_r \cdot \left(\alpha \gamma_u + \frac{1-\alpha}{4} \sum_{k \in \mathcal{N}(u)} \gamma_k \right) \quad (21)$$

$$= \sum_{r \in \mathcal{R}} p_r \cdot w_r \cdot \frac{m}{n+m-1} = \frac{m \cdot p}{n+m-1} \sum_{r \in \mathcal{R}} w_r. \quad (22)$$

Note that Equality (22) assumes that each u has 4 neighbors with Manhattan distance 1. \square

Proof of Lemma 4. We use Lemma 5 to prove Lemma 4. When $p_r = 1/n^2$ for all $r \in \mathcal{R}$, we have that the transition matrix P of $\mathcal{M}(\alpha)$ is reduced to the following form: $P(\mathbf{x}, \mathbf{y}) = 1/n^2$ iff (\mathbf{x}, \mathbf{y}) is a (u, u') -neighbor for all possible (u, u') with $u \neq u'$ and $P(\mathbf{x}, \mathbf{y}) = 0$ if $\mathbf{x} \neq \mathbf{y}$ and (\mathbf{x}, \mathbf{y}) is not a (u, u') -neighbor. For each given state \mathbf{x} , let $\mathcal{N}(\mathbf{x})$ be set of states \mathbf{y} such that (\mathbf{x}, \mathbf{y}) is a (u, u') -neighbor for some (u, u') . The Markov Chain $\mathcal{M}(\alpha)$ thus can be viewed as follows: we start with an initial state $\mathbf{x}_0 \in \Omega$; suppose at time t we are at state $\mathbf{X}^{(t)}$; then $\mathbf{X}^{(t+1)} = \mathbf{y}$ with probability $1/n^2$ for each $\mathbf{y} \in \mathcal{N}(\mathbf{X}^{(t)})$ and $\mathbf{X}^{(t+1)} = \mathbf{X}^{(t)}$ with the remaining probability.

Define a natural metric \mathbf{d} over $\Omega \times \Omega$ as follows: $\mathbf{d}(\mathbf{x}, \mathbf{y}) = \sum_u |x_u - y_u|$. Let \mathcal{S} be the set of all possible (u, u') -neighbors on Ω . Observe that $\mathbf{d}(\mathbf{x}, \mathbf{y}) \geq 2$ for any $\mathbf{x} \neq \mathbf{y}$ and $\mathbf{d}(\mathbf{x}, \mathbf{y}) = 2$ for all $(\mathbf{x}, \mathbf{y}) \in \mathcal{S}$. We can verify that $\mathbf{d}(\mathbf{x}, \mathbf{y}) = \sum_u |x_u - y_u| \leq D = 2m$ for all possible (\mathbf{x}, \mathbf{y}) . To apply Lemma 5, we design a natural coupling $(\mathbf{X}, \mathbf{Y}) \rightarrow (\mathbf{X}', \mathbf{Y}')$ over all possible $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$ such that \mathbf{X}' and \mathbf{Y}' each is a one-step transition from \mathbf{X} and \mathbf{Y} respectively according the transition matrix P . In other words, \mathbf{X}' and \mathbf{Y}' are random states each follows the respective distribution $P(\mathbf{X}, \cdot)$ and $P(\mathbf{Y}, \cdot)$. Therefore we claim that \mathbf{X}' and \mathbf{Y}' is a faithful copy of $\mathcal{M}(\alpha)$ with respective P . In the following, we prove that $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq (1 - \frac{1}{n^2})\mathbf{d}(\mathbf{X}, \mathbf{Y})$.

Proof. Consider the case $c = 2$. Let $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$ is a (a, b) -neighbor such that $Y_a = X_a - 1$ and $Y_b = X_b + 1$. Consider the following alternate view of the transition of $\mathcal{M}(\alpha)$. Suppose we are at state \mathbf{X} ; each round we sample a request $r = (u, u')$ with probability $1/n^2$; if $X_u \geq 1$ and $X_{u'} \leq c - 1$, then set $X'_u = X_u - 1$, $X'_{u'} = X_{u'} + 1$ and $X'_v = X_v$ for all $v \notin \{u, u'\}$, such that $(\mathbf{X}, \mathbf{X}')$ forms a (u, u') -neighbor; otherwise set $\mathbf{X}' = \mathbf{X}$. Consider the following cases.

Case 1: $X_a = 1, X_b = 0$. Consider an arbitrary request $r = (u, u')$. If $u = a$ and $u' = b$, then we have that $\mathbf{X}' = \mathbf{Y}$ and $\mathbf{Y}' = \mathbf{Y}$ and thus $\mathbf{d}(\mathbf{X}', \mathbf{Y}') = 0$, which occurs with probability $1/n^2$. Similarly, if $u = b$ and $u' = a$, we have $\mathbf{X}' = \mathbf{X}$ and $\mathbf{Y}' = \mathbf{X}$ and $\mathbf{d}(\mathbf{X}', \mathbf{Y}') = 0$, which occurs with probability $1/n^2$. We can verify that for all other request types, $\mathbf{d}(\mathbf{X}', \mathbf{Y}') \leq 2 = \mathbf{d}(\mathbf{X}, \mathbf{Y})$. Thus, we claim that $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq (1 - \frac{2}{n^2})\mathbb{E}[\mathbf{d}(\mathbf{X}, \mathbf{Y})]$. The details are summarized in the table below.

Table 1: $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$: $X_a = 1, X_b = 0$ and $Y_a = 0, Y_b = 1$.

$r = (u, u')$	$\mathbf{d}(\mathbf{X}', \mathbf{Y}')$
$u \notin \{a, b\}$	2
$u = a, u' \neq b$	2
$u = a, u' = b$	0
$u = b, u' \neq a$	2
$u = b, u' = a$	0

Case 2: $X_a = 1, X_b = 1$. Applying a similar analysis, we claim that $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq (1 - \frac{1}{n^2})\mathbb{E}[\mathbf{d}(\mathbf{X}, \mathbf{Y})]$. The details are as follows.

Table 2: $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$: $X_a = 1, X_b = 1$ and $Y_a = 0, Y_b = 2$.

$r = (u, u')$	$\mathbf{d}(\mathbf{X}', \mathbf{Y}')$
$u \notin \{a, b\}$	2
$u = a, u' \neq b$	2
$u = a, u' = b$	0
$u = b$	2

Case 3: $X_a = 2, Y_a = 1$ and $X_b = 1, Y_b = 2$.

Consider the following cases summarized in the table below. Thus we have $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq (1 - \frac{2}{n^2})\mathbb{E}[\mathbf{d}(\mathbf{X}, \mathbf{Y})]$.

Table 3: $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$: $X_a = 2, X_b = 1$ and $Y_a = 1, Y_b = 2$

$r = (u, u')$	$\mathbf{d}(\mathbf{X}', \mathbf{Y}')$
$u \notin \{a, b\}$	2
$u = a, u' \neq b$	2
$u = a, u' = b$	0
$u = b, u' \neq a$	2
$u = b, u' = a$	0

Case 4: $X_a = 2, X_b = 0$ and $Y_a = 1, Y_b = 1$. Thus from the table below, we have $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq (1 - \frac{1}{n^2})\mathbb{E}[\mathbf{d}(\mathbf{X}, \mathbf{Y})]$.

Table 4: $(\mathbf{X}, \mathbf{Y}) \in \mathcal{S}$: $X_a = 2, X_b = 0$ and $Y_a = 1, Y_b = 1$

$r = (u, u')$	$\mathbf{d}(\mathbf{X}', \mathbf{Y}')$
$u \notin \{a, b\}$	2
$u = a$	2
$u = b, u' \neq a$	2
$u = b, u' = a$	0

From the case based analysis above, we have $\mathbb{E}[\mathbf{d}(\mathbf{X}', \mathbf{Y}')] \leq \beta\mathbb{E}[\mathbf{d}(\mathbf{X}, \mathbf{Y})]$ with $\beta = 1 - \frac{1}{n^2}$. By applying a similar analysis, we get the result for the case when $c = 1$. Thus, by Lemma 5, we have $\tau(\epsilon) \leq n^2 \ln(2m/\epsilon)$. \square

C Additional Experiments

In this section we describe the results from the additional experiments. In particular, we conduct both simulated and real-world experiments similar to the main section for the periods Morning and Evening. Figures 12, 13 describes the request pattern for Afternoon while Figures 14, 15 show the patterns for Evening. Figures 16, 17 denote the results from using the Afternoon data for simulation while Figures 20, 21 denote the results when the Evening data is used for simulation. Figures 24, 25 denotes the performance on the

real-time requests for the Afternoon and Evening periods respectively. In general, all results follow the same pattern as those reported in the main section.

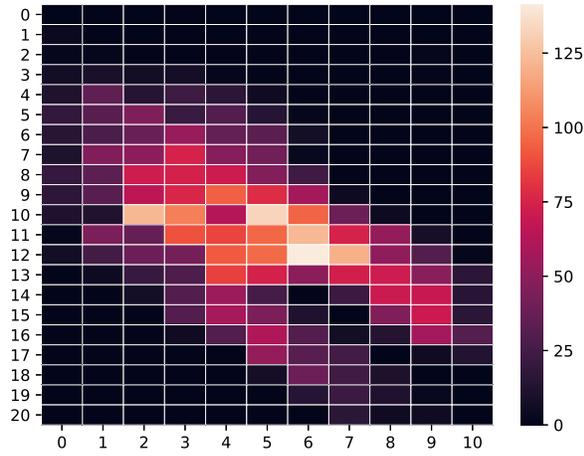


Fig. 12: Heat map of Starting Location of rides in NYC (11:00AM - 3:00PM)

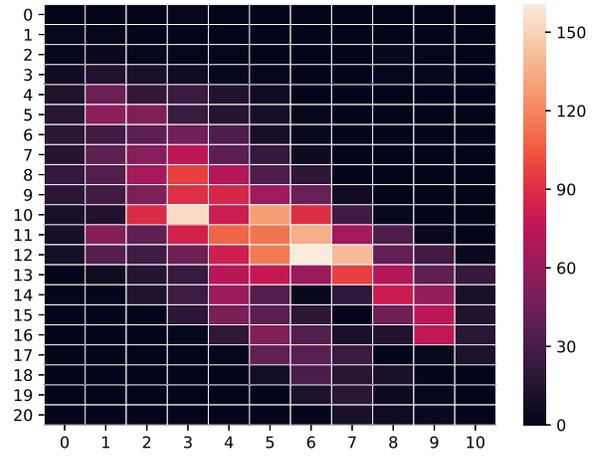


Fig. 13: Heat map of Ending Location of rides in NYC (11:00AM - 3:00PM)

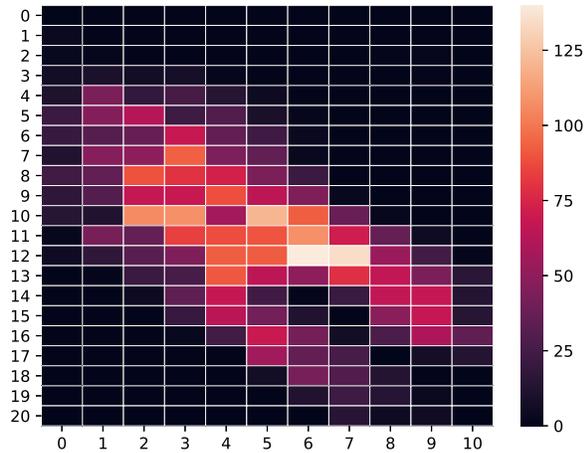


Fig. 14: Heat map of Starting Location of rides in NYC (3:00PM - 7:00PM)

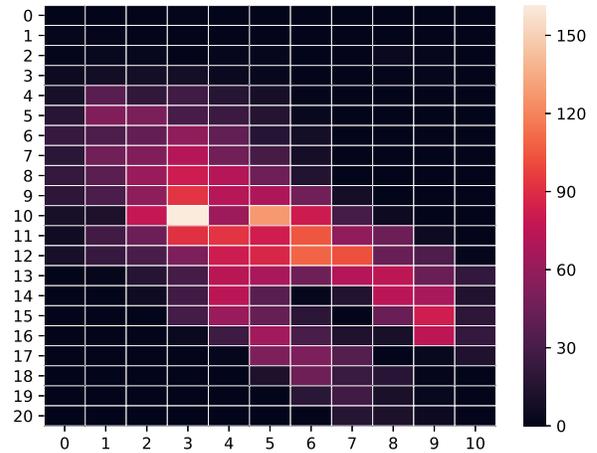


Fig. 15: Heat map of Ending Location of rides in NYC (3:00PM - 7:00PM)

D Details on Deep-Q-Learning

We used deep Q-learning with experience replay, trained in a custom OpenAI Gym environment. Memory size was 2000 events, with a batch size of 64 for experience replay. The discount factor was 0.9. The initial

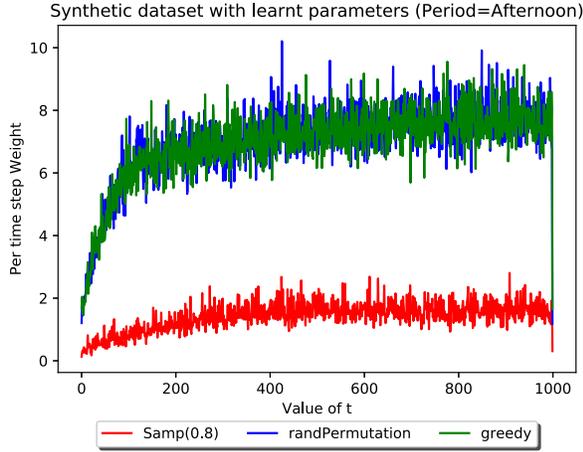


Fig. 16: Obj. (1) on simulation with real-data parameters (11:00AM-3:00PM)

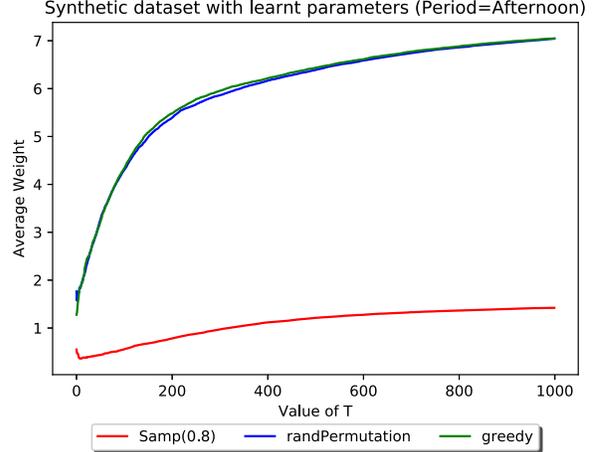


Fig. 17: Obj. (2) on simulation with real-data parameters (11:00AM-3:00PM)

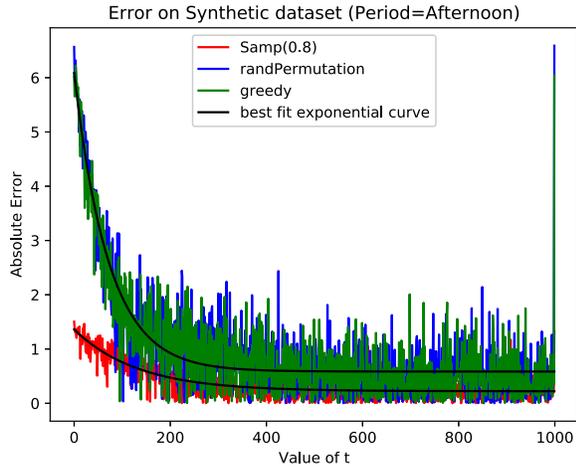


Fig. 18: Error in Obj. (1) on simulation with real-data parameters (11:00AM-3:00PM)

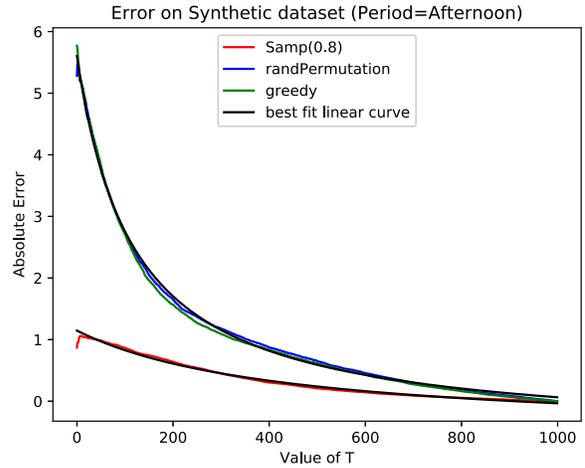


Fig. 19: Error in Obj. (2) on simulation with real-data parameters (11:00AM-3:00PM)

exploration rate was 1.0, decaying with a decay factor of .995 to a minimum of 0.01. The learning rate for the Q-learning update was 0.001. Training consisted of 5 episodes, each randomly initialized and continuing for 2500 steps. The network itself had 2 dense hidden layers with 24 units and ReLU activation, and a final layer with 6 outputs and a linear activation; the loss was quadratic. As with the $\text{Samp}(\alpha)$ policy, any attempt to make a match that would be illegal was simply ignored and given reward 0.

E Details on Value Iteration

We used standard value iteration algorithm to solve for an exact optimal policy on small grids under uniform IID requests. To illustrate the optimal policy, we present the distribution of different measures of activity

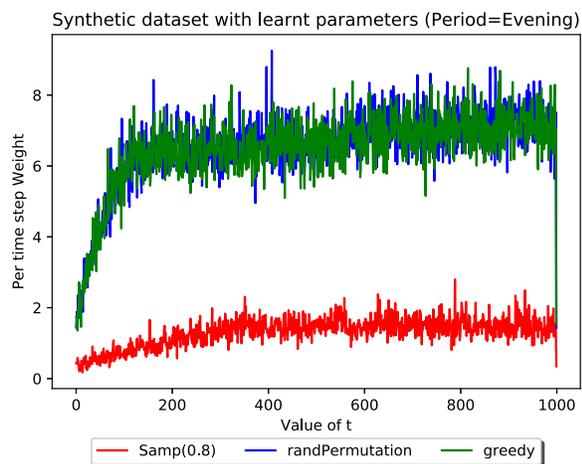


Fig. 20: Obj. (1) on simulation with real-data parameters (3:00PM-7:00PM)

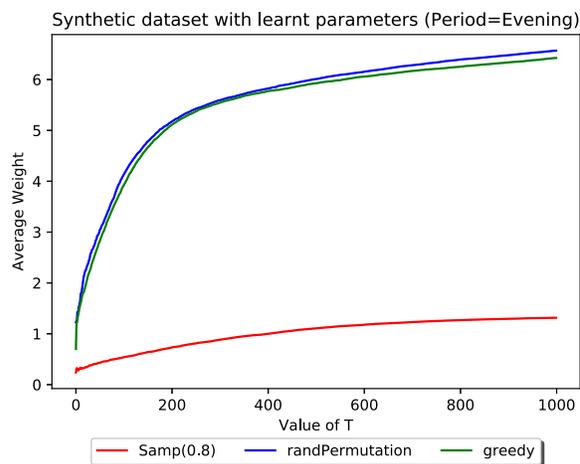


Fig. 21: Obj. (2) on simulation with real-data parameters (3:00PM-7:00PM)

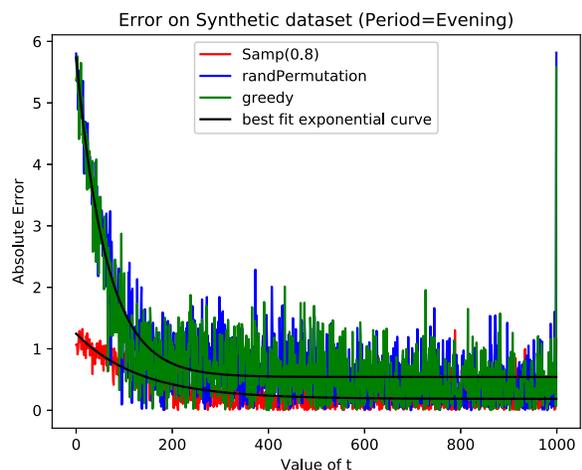


Fig. 22: Error in Obj. (1) on simulation with real-data parameters (3:00PM-7:00PM)

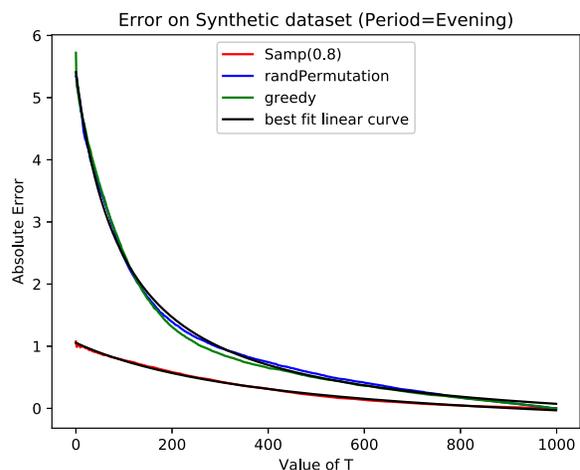


Fig. 23: Error in Obj. (2) on simulation with real-data parameters (3:00PM-7:00PM)

as heatmaps on the grid. We initialize the state with 1 car at left-top location and zero car at remaining locations. We then simulate an episode of 1000 periods under the optimal policy by drawing uniform IID requests. The measure “time covered” is the percentage of periods in which the location is occupied by at least one car. The measure “drop rate” is the percentage of periods in which the location is either a starting or ending location. The measure “start location” is the percentage of periods in which the location is a starting location. The reason for considering small grids is because of the huge number of states. In particular, 2×2 grid with capacity 2, capacity 3 and 3×3 grid with capacity 1 have 324, 1024 and 41472 states respectively.

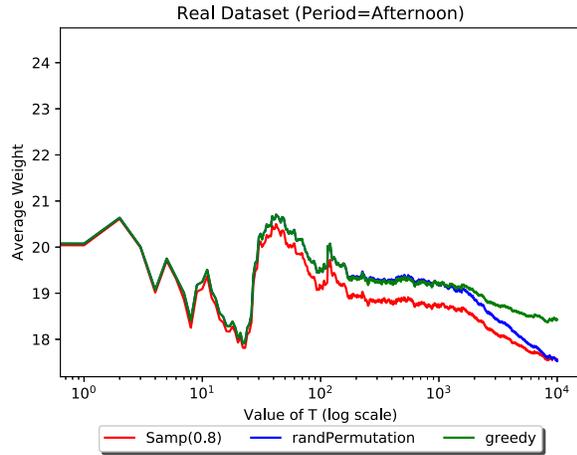


Fig. 24: Obj. (1) on real-dataset (11:00AM-3:00AM). x -axis in log-scale

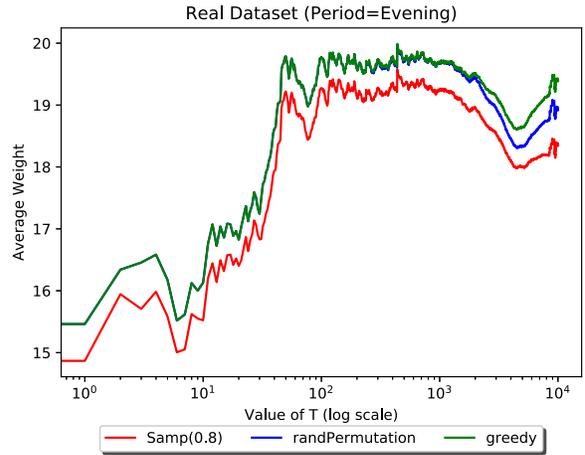


Fig. 25: Obj. (2) on real-dataset (3:00PM-7:00PM). x -axis in log-scale

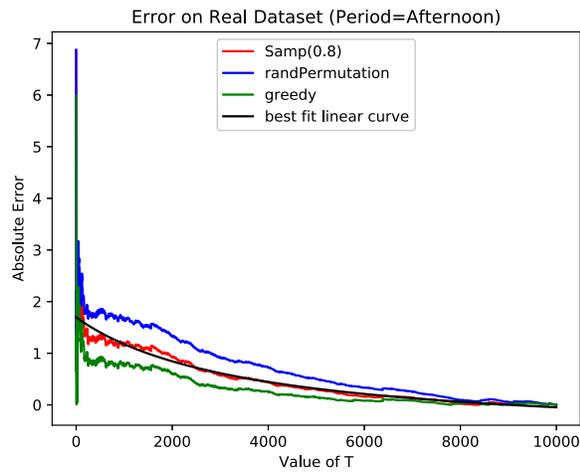


Fig. 26: Error in Obj. (1) on real-dataset (11:00AM-3:00PM). x -axis in log-scale

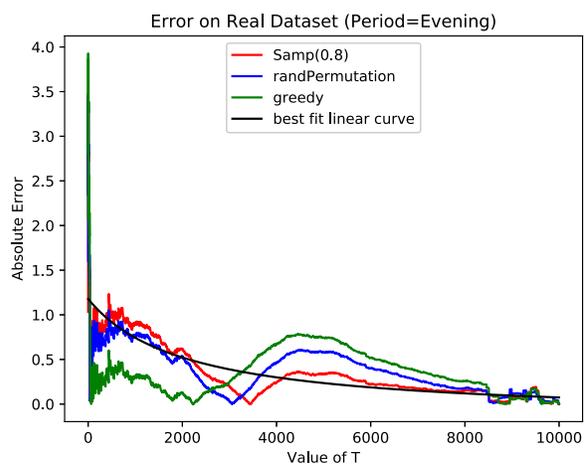


Fig. 27: Error in Obj. (2) on real-dataset (3:00PM-7:00PM). x -axis in log-scale

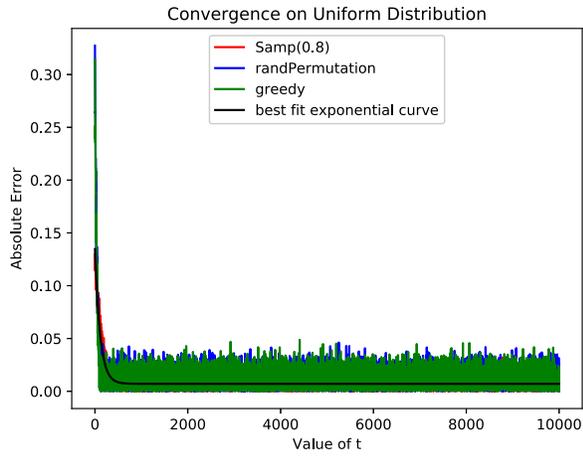


Fig. 28: Error in Obj. (1) on simulated data

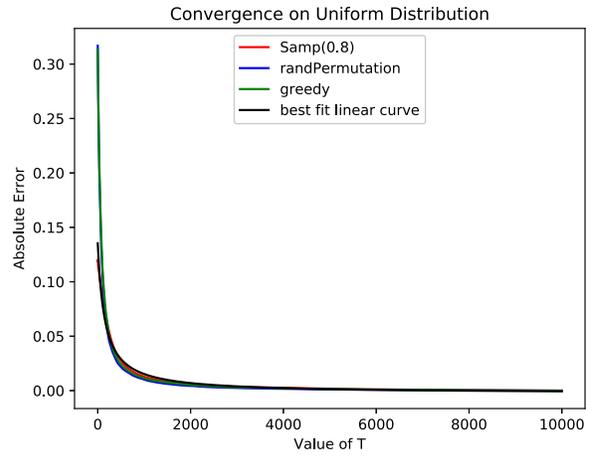


Fig. 29: Error in Obj. (2) on simulated data

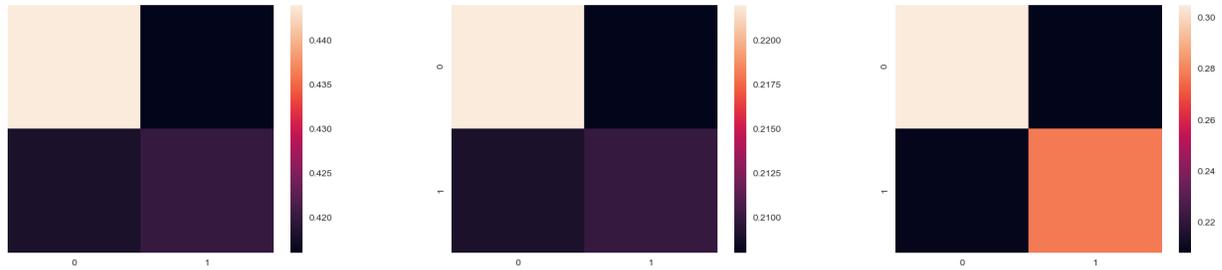


Fig. 30: 2×2 grids with capacity 3; black to white represent low to high values. Plots 1, 2, 3 represent the distribution of requests with either starting or ending at a cell, requests starting at a cell and percentage of time spent by cars in a cell respectively by the optimal policy.

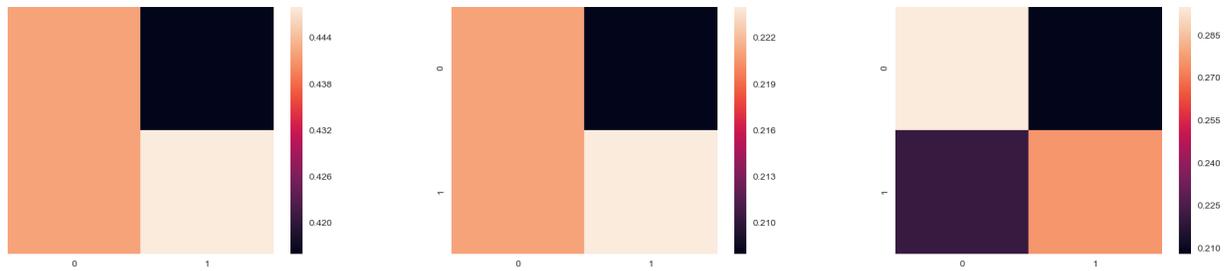


Fig. 31: 2×2 grids with capacity 2; black to white represent low to high values. Plots 1, 2, 3 represent the distribution of requests with either starting or ending at a cell, requests starting at a cell and percentage of time spent by cars in a cell respectively by the optimal policy.

Bibliography

- [1] Javier Alonso-Mora, Samitha Samaranyake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, 114(3):462–467, 2017.
- [2] Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I Jordan. An introduction to mcmc for machine learning. *Machine learning*, 50(1-2):5–43, 2003.
- [3] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Dynamic pricing in ridesharing platforms. *SIGecom Exch.*, 15(1):65–70, September 2016. ISSN 1551-9031. <https://doi.org/10.1145/2994501.2994505>. URL <http://doi.acm.org/10.1145/2994501.2994505>.
- [4] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 517–517, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4527-9. <https://doi.org/10.1145/3033274.3085099>. URL <http://doi.acm.org/10.1145/3033274.3085099>.
- [5] Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. State dependent control of closed queueing networks with application to ride-hailing. In *Proceedings of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '18, 2018. Complete working paper available at arXiv:1803.04959.
- [6] Anton Braverman, Jim G Dai, Xin Liu, and Lei Ying. Empty-car routing in ridesharing systems. *arXiv preprint arXiv:1609.07219*, 2016.
- [7] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [8] Gerard P Cachon, Kaitlin M Daniels, and Ruben Lobel. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3):368–384, 2017.
- [9] Abhimanyu Das, Sreenivas Gollapudi, Anthony Kim, Debmalya Panigrahi, and Chaitanya Swamy. Minimizing latency in online ride and delivery services. In *Proceedings of the 2018 World Wide Web Conference*, pages 379–388. International World Wide Web Conferences Steering Committee, 2018.
- [10] John P Dickerson and Tuomas Sandholm. Futurematch: Combining human value judgments and machine learning to match in dynamic environments. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [11] John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, AAAI '18, pages 1007–1014, 2018.
- [12] Supriyo Ghosh, Pradeep Varakantham, Yossiri Adulyasak, and Patrick Jaillet. Dynamic repositioning to reduce lost demand in bike sharing systems. *Journal of Artificial Intelligence Research*, 58:387–430, 2017.
- [13] Venkatesan Guruswami. Rapidly mixing markov chains: a comparison of techniques (a survey). *arXiv:1603.01512*, 2016.
- [14] Robert Hahn and Robert Metcalfe. The ridesharing revolution: Economic survey and synthesis. Technical report, 2017.
- [15] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 1970.
- [16] Yongzheng Jia, Wei Xu, and Xue Liu. An optimization framework for online ride-sharing markets. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, pages 826–835. IEEE, 2017.
- [17] Nikolay Laptev, Jason Yosinski, Li Erran Li, and Slawek Smyl. Time-series extreme event forecasting with neural networks at uber. In *International Conference on Machine Learning*, volume 34, pages 1–5, 2017.
- [18] David A Levin and Yuval Peres. *Markov chains and mixing times*. American Mathematical Soc., 2017.
- [19] Yuxi Li. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*, 2017.

- [20] Kaixiang Lin, Renyu Zhao, Zhe Xu, and Jiayu Zhou. Efficient large-scale fleet management via multi-agent deep reinforcement learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1774–1783. ACM, 2018.
- [21] Meghna Lowalekar, Pradeep Varakantham, and Patrick Jaillet. Online spatio-temporal matching in stochastic and dynamic domains. *Artificial Intelligence*, 261:71 – 112, 2018. ISSN 0004-3702. <https://doi.org/https://doi.org/10.1016/j.artint.2018.04.005>. URL <http://www.sciencedirect.com/science/article/pii/S0004370218302030>.
- [22] Hongyao Ma, Fei Fang, and David C Parkes. Spatio-temporal pricing for ridesharing platforms. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, 2019.
- [23] Aranyak Mehta et al. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368, 2013.
- [24] Nicholas Metropolis and Stanislaw Ulam. The monte carlo method. *Journal of the American statistical association*, 44(247):335–341, 1949.
- [25] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- [26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [27] Eoin O’Mahony and David B Shmoys. Data analysis and optimization for (citi) bike sharing. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
- [28] Erhun Ozkan and Amy Ward. Dynamic matching for real-time ridesharing. *Available at SSRN 2844451*, 2017.
- [29] Meng Qu, Hengshu Zhu, Junming Liu, Guannan Liu, and Hui Xiong. A cost-effective recommender system for taxi drivers. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 45–54. ACM, 2014.
- [30] Jiang Rong, Tao Qin, and Bo An. Dynamic pricing for reusable resources in competitive market with stochastic demand. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [31] Yong Sun, Jun Wang, and Wenan Tan. Online algorithms of task allocation in spatial crowdsourcing. In *Proceedings of the 12th Chinese Conference on Computer Supported Cooperative Work and Social Computing*, pages 205–208. ACM, 2017.
- [32] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [33] Yongxin Tong, Jieying She, Bolin Ding, Lei Chen, Tianyu Wo, and Ke Xu. Online minimum matching in real-time spatial data: experiments and analysis. *Proceedings of the VLDB Endowment*, 9(12):1053–1064, 2016.
- [34] Tanvi Verma, Pradeep Varakantham, Sarit Kraus, and Hoong Chuin Lau. Augmenting decisions of taxi drivers through reinforcement learning for improving revenues. In *Twenty-Seventh International Conference on Automated Planning and Scheduling*, 2017.
- [35] Jian Wen, Jinhua Zhao, and Patrick Jaillet. Rebalancing shared mobility-on-demand systems: A reinforcement learning approach. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 220–225. IEEE, 2017.
- [36] Huaxiu Yao, Fei Wu, Jintao Ke, Xianfeng Tang, Yitian Jia, Siyu Lu, Pinghua Gong, Jieping Ye, and Zhenhui Li. Deep multi-view spatial-temporal network for taxi demand prediction. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, AAAI '18, pages 2588–2595, 2018.