

Strategic Planning for Flexible Agent Availability in Large Taxi Fleets

Rajiv Ranjan Kumar
School of Computing and Information
Systems, Singapore Management
University
Singapore
rajivr.k.2017@phdcs.smu.edu.sg

Pradeep Varakantham
School of Computing and Information
Systems, Singapore Management
University
Singapore
pradeepv@smu.edu.sg

Shih-Fen Cheng
School of Computing and Information
Systems, Singapore Management
University
Singapore
sfcheng@smu.edu.sg

ABSTRACT

In large scale multi-agent systems like taxi fleets, individual agents (taxi drivers) are self interested (maximizing their own profits) and this can introduce inefficiencies in the system. One such inefficiency is with regards to the "required" availability of taxis at different time periods during the day. Since a taxi driver can work for limited number of hours in a day (e.g., 8-10 hours in a city like Singapore), there is a need to optimize the specific hours, so as to maximize individual as well as social welfare. Technically, this corresponds to solving a large scale multi-stage selfish routing game with transition uncertainty. Existing work in addressing this problem is either unable to handle "driver" constraints (e.g., breaks during work hours) or not scalable. To that end, we provide a novel mechanism that builds on replicator dynamics through ideas from behavior cloning. We demonstrate that our methods provide significantly better policies than the existing approach in terms of improving individual agent revenue and overall agent availability.

KEYWORDS

Large Taxi Fleet; Equilibrium Solution; Game Theory; Optimization; Replicator Dynamics; Data/Policy Completion

ACM Reference Format:

Rajiv Ranjan Kumar, Pradeep Varakantham, and Shih-Fen Cheng. 2023. Strategic Planning for Flexible Agent Availability in Large Taxi Fleets. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Aggregation systems offer time sensitive tasks/services (require workers to respond within a minute). Notable examples of such platforms include transport (e.g., Uber or Grab or Gojek), food delivery (e.g., Foodpanda or DoorDash), and grocery delivery (e.g., Amazon Fresh or Instacart). Workers on these platforms have to constantly make both spatial and temporal decisions: the "spatial" decision on "where" to work and the "temporal" decision on "when" to work. These decisions are challenging since the performance of workers depends not just on their understandings of the demands, but also on when and where other workers choose to work. So, in essence, there is a need to study strategic and operational decision problems together.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

In the literature, such decision problems are typically formulated using game theory and solved using equilibrium-seeking algorithms (since an individual has to best respond to all others' decisions, which are best responses to our best responses as well, *ad infinitum*). However, existing research [3, 27] consider either spatial decisions or temporal decisions but not both. Even in the one recent work where both spatial or temporal decisions are considered [9], there are rigid assumptions on workers' work arrangement (e.g., although a worker can choose when to start working, once work starts, this worker has to work consecutively for a fixed number of hours). Furthermore, the equilibrium computation approaches that are employed are computationally expensive (with 3 days to solve one instance) and do not scale easily to real problems of interest.

To that end, we develop scalable approaches that consider both spatial and temporal decisions for workers (henceforth referred to as agents), while allowing for flexibility in agent availability. Agents are profit-maximizing and their strategic decisions include: 1) pick the time period to enter the system and start working; 2) once in the system, choose the regions to move to at different times; and 3) while in the system, an agent can choose to take a temporary break of arbitrary length, subject to agents' total working hours.

Due to the scale of problems at hand with thousands of agents, existing methods have exploited homogeneity and anonymity (lack of identity) to compute symmetric equilibrium solutions [9, 27]. However, such approaches are not scalable and can take 3 days to run a single instance. To achieve scalability, we propose a simulation based mechanism based on the "replicator dynamics framework" [32] to compute approximate symmetric equilibrium solutions. Intuitively, at each iteration of our new method, the current symmetric policy is simulated for all the agents and the top "k" (where k is around 500) trajectories (in terms of revenue generated) are employed to compute a new symmetric policy. Since we consider a finite "k" it is feasible that some states are not covered by the trajectories. Therefore, we also explore multiple data imputation mechanisms to complete the policy.

Our paper makes the following key contributions:

- (1) We introduce a game-theoretic model that allows agents to choose their working schedule in a flexible manner. After picking a time to enter the system and begin working, agents can choose when to take breaks and where to work.
- (2) We propose a simulation-based best response method for computing symmetric equilibrium at scale.
- (3) To handle partial policy samples obtained with our simulation method, we propose various data imputation techniques

and show the utility of these methods in improving solution quality efficiently.

- (4) We demonstrate the scalability of our approach by solving problem instances that contain more than 20,000 agents in 10 minutes.

2 RELATED WORK

Past work has employed approaches for operational optimization of taxi fleets using heuristic approaches [13, 14], bipartite matching [8, 28], two stage optimization [12], multi-objective optimization models [16] and deep-learning based methods [4, 19, 29, 33]). There are multiple key differences with these line of works. First, above approaches assume a centralized optimization (by aggregating companies like Uber) and individual taxis are assumed to follow the inputs provided by the central entity. Second, the focus in the above approaches is typically on solving the operational decision problem of which zone to move to at different points of the day if there is no customer on board. In contrast, we focus on the combination of strategic (when to start shift) and operational (where to move to) problems with decentralized optimization (where individual taxis maximize their own revenue).

Saisubramanian et al. [23] used MIP to optimize the ambulance allocation in the emergency department setting. Lowalekar et al. [11] proposed online repositioning solution for bike sharing systems. Verma et al. [30] optimize a single agent’s sequential decision making process over a time horizon. For taxi fleet optimization there has been some recent work on decentralized approaches [17, 20, 27, 34, 35], however the emphasis there has been on the operational decision problem and not on the strategic decision problem (optimizing starting time of shift and breaks), which is a main focus in this paper.

The closest literature to our contributions in the paper are [3, 9]. While Gan and An [3] considers driver’s operating hours constraints explicitly to optimize constraints related to operation of taxi drivers, Kumar et al. [9] incorporates starting time period as part of agent’s strategies (not constraints) to provide advice on when to start the shift. Such approaches are not scalable and work with rigid assumptions on taxi’s work arrangement. In comparison, our work allows for flexibility in agent availability by optimizing and allowing breaks in operation. Furthermore, our approach is significantly more scalable (10 minutes vs 3 days) due to the simulation based best response framework employed to compute approximate symmetric equilibrium solutions.

3 FLEXIBLE TAXI FLEETS

In many cities where taxis play an important role in providing point-to-point transport, it is of critical importance for drivers to decide how to position themselves when idling, so as to maximize their revenues. The driver’s decision on where to move to is strategic in nature, as drivers are competing for the same pool of jobs when they are close to each other.

In this paper, we aim to introduce another important dimension in driver’s strategy space, which is when to work. Similar to the spatial movement decisions, decisions on when to work are also strategic in nature for driver, as the number of drivers per time period also has a strong impact on individual’s revenue potential.

Compared to a recent work by Kumar et al. [9], which also allows drivers to pick their work starting time, our proposed model is much more flexible in that we allow drivers to freely leave and return to the service platform (as opposed to having to pick a starting time and work continuously for a fixed amount of time). Our proposed model is thus much more realistic and expressive.

We formally define our model as Flexible Selfish Routing under Uncertainty (FSRU), which allows workers to freely take break during their service. Our FSRU framework is a special case of the generic stochastic game model [15, 18, 24]: while the transition and reward functions for an agent in our model are dependent only on the aggregate distributions of other agent states, they can be dependent on specific state and action of every other agent in general stochastic games. FSRU extends on the Selfish Routing under Transition Uncertainty (SRT) model [9], by allowing for flexible operations of individual agents. FSRU represents problems with selfish agents and hence is different to cooperative models such as Decentralized POMDPs [1]. A FSRU instance is characterized by the tuple:

$$\langle \mathcal{P}, \mathcal{S}, \mathcal{A}, \mathcal{T}, \{\mathcal{R}_\tau\}_{\tau \in \Gamma}, B, \mathbf{d}^0, \delta, H \rangle$$

where \mathcal{P} is set of agents, i.e., the set of taxis. \mathcal{S} is the set of zones a taxi could move to, \mathcal{A} is the set of zones to which a taxi driver wishes to move. There is a special “Sink” state, s_{sink} , that is used to represent a taxi driver on a break and an action a_{sink} that represents taking a break and will move the state to s_{sink} state with probability 1. If a taxi in s_{sink} state takes any action, action will succeed with probability 1.

\mathcal{T} models the transition function of every agent and more specifically, $\mathcal{T}^t(s, a, s', \mathbf{d})$ represents the probability that an agent in state $s \in \mathcal{S}$ after taking action $a \in \mathcal{A}$ would transition to state s' , when the state distribution of all (active) agents is \mathbf{d} at time t . In case of taxi fleet, the transition function, \mathcal{T} , depends on the involuntary movements between zones. The involuntary movement between any two zones i and j at a time step t is determined by the number of customers (customer flow), $fl^t(i, j)$ moving between those zones. Thus, demand for taxis in a zone i at time step t is given by $\sum_j fl^t(i, j)$. We assume that the zones are small enough that if the number of taxis in a particular zone during a time period is fewer than the demand, then all taxis will be hired. If the number of taxis are more than the demand, a fraction of the taxis (equal to demand) will be hired and each taxi will have equal probability of getting hired. Equation (1) provides the expression for computing the transition probabilities between states.

$$\mathcal{T}^t(s, a, s', \mathbf{d}) = \begin{cases} \mathbf{1} & \text{if } (a = a_{sink} \ \& \ s' = s_{sink}) \text{ or } (s = s_{sink} \ \& \ a = s') \\ \mathbf{1} & \text{if } (a = a_{sink} \ \& \ s' \neq s_{sink}) \text{ or } (s = s_{sink} \ \& \ a \neq s') \\ \frac{fl^t(s, s')}{\sum_{\hat{s}} fl^t(s, \hat{s})} & \text{if } s \neq s_{sink}, a \neq a_{sink} \text{ and (C1)} \\ \frac{fl^t(s, s')}{d_s} & \text{if } s \neq s_{sink}, a \neq a_{sink} \text{ and (C2)} \\ 1 - \frac{\sum_{\hat{s} \neq s'} fl^t(s, \hat{s})}{d_s} & \text{if } s \neq s_{sink}, a \neq a_{sink} \text{ and (C3)} \end{cases} \quad (1)$$

The exact transition probabilities depend on normalized demands to other zones from the current zone (C1 represents this case; C1: $\sum_{\hat{s}} fl^t(s, \hat{s}) \geq d_s$). On the other hand, if the number of taxis is more than the number of customers in a zone, the transition

probabilities should depend on whether the action (intended zone) coincides with the destination zone (**C2** and **C3** represent these two cases). The condition **C2** is only possible when the taxi agent is hired by a customer heading towards any s' that is not a (**C2**: $\sum_{\hat{s}} fl^t(s, \hat{s}) < d_s$ and $a \neq s'$). For condition **C3**, the taxi agent can either be free or hired by a customer heading towards the agent's intended zone (**C3**: $\sum_{\hat{s}} fl^t(s, \hat{s}) < d_s$ and $a = s'$).

$\mathcal{R}_\tau^t(s, a, s', \mathbf{d})$ is the reward obtained by an agent of type τ at time t when in state s , taking action a and moving to state s' when the state distribution of other agents is \mathbf{d} at time t . Similar to the transition probabilities, the reward function for all active agents at time t . We will refer to the cost (of moving) and revenue (of moving a customer) from state s to s' at time t as $c^t(s, s')$ and $r^t(s, s')$ respectively. $\mathcal{R}^t(\cdot)$ is defined differently under these three conditions:

$$\mathcal{R}^t(s, a, \mathbf{d}) = \begin{cases} c^t(s, s_{sink}) & \text{if } a = a_{sink} \\ c^t(s_{sink}, s' = a) & \text{if } s = s_{sink} \\ \sum_{s'} \mathcal{T}^t(s, a, s', \mathbf{d}) \cdot (r^t(s, s') - c^t(s, s')) & \text{if } s \neq s_{sink}, a \neq a_{sink} \text{ and (C1)} \\ \sum_{s' \neq a} \mathcal{T}^t(s, a, s', \mathbf{d}) \cdot (r^t(s, s') - c^t(s, s')) & \text{if } s \neq s_{sink}, a \neq a_{sink} \text{ and (C2)} \\ \frac{fl^t(s, a)}{d_s} \cdot r^t(s, a) - \mathcal{T}^t(s, a, a) \cdot c^t(s, a) & \text{if } s \neq s_{sink}, a \neq a_{sink} \text{ and (C3)} \end{cases} \quad (2)$$

It should be noted that taxis are hired in conditions **C1** and **C2**; therefore, the expected rewards in these two cases are the sum of expected rewards to all feasible destinations. For **C3**, cost is incurred for sure, but revenue can only be earned if the taxi is hired. Our goal in solving taxi fleet optimization problem as a FSRU is to maximize expected revenue for individual taxi drivers who are perfectly rational and follow computed policies.

B represents maximum number of allowed breaks in operation. δ is the maximum number of time steps any agent can be active. Finally H represents the time horizon of the decision making process.

The good thing about the FSRU model is that all the elements of the tuple can be computed from any real world taxi dataset, where information (i.e., source, destination, fare, cost) about customer requests is available.

3.1 Example

We show how transition probabilities and the rewards are computed for a small problem extended from the one introduced in Kumar et al. [9]:

EXAMPLE 1. Consider a map with three zones, $\mathcal{S} = \{s_0, s_1, s_2\} \cup s_{sink}$. For simplicity we set flow values to 1 across all time periods; i.e., one passenger goes from each zone to an adjacent zone in all time periods. We also set rewards and costs to be fixed at $r^t(s, s') = 1$ and $c^t(s, s') = 0$ for all time periods t and all zones $s, s' \in \mathcal{S}$. If the distribution of taxis at a given time period t is $\mathbf{d}^t = (1, 1, 4, 0)$, then:

- The transition function $\mathcal{T}^t(s, a, s', \mathbf{d}^t)$ is specified by matrix $m(s)$, in which the row label represents action a , and the column label represents destination zone s' . Transition function for s_0 and s_2

are: $m(s_0)$ & $m(s_2) =$

$$\begin{pmatrix} 0.0 & 0.5 & 0.5 & 0.0 \\ 0.0 & 0.5 & 0.5 & 0.0 \\ 0.0 & 0.5 & 0.5 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{pmatrix} \& \begin{pmatrix} 0.75 & 0.25 & 0.0 & 0.0 \\ 0.25 & 0.75 & 0.0 & 0.0 \\ 0.25 & 0.25 & 0.5 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{pmatrix}$$

- Similarly, the reward function $\mathcal{R}^t(s, a, \mathbf{d})$ is specified as a matrix, in which the row label represents current zone s , and the column label represents action a :

$$\begin{pmatrix} 1.0 & 1.0 & 1.0 & 0.0 \\ 1.0 & 1.0 & 1.0 & 0.0 \\ 0.5 & 0.5 & 0.5 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 \end{pmatrix}$$

Consider state s_2 : $fl^t(s_2, s') = 1$ for $s' \in \{s_0, s_1\}$ and $d^t(s_2) = 4$. The transition and reward functions are computed using (1) and (2). For $a = s_2$ and $s' \in \{s_0, s_1\}$, $\mathcal{T}^t(s_2, a, s', \mathbf{d}^t) = 1/4$ by **C2**. For $a = s_2$ and $s' = s_2$, $\mathcal{T}^t(s_2, a, s', \mathbf{d}^t) = 1 - 2/4 = 1/2$ by **C3**. For $a = s_2$ and $s' = s_{sink}$, $\mathcal{T}^t(s_2, a, s', \mathbf{d}^t) = 0$. Rest of the transition and reward function values can be computed similarly.

4 BRUTE-FORCE APPROACH: FP

Algorithm 1 Finding ϵ -Equilibrium

```

1:  $\pi_0 = \text{GetRandomPolicy}()$ ,  $x_0 = \text{GetAgentFlow}(\pi_0, N)$ 
2:  $i = 1$ , converged = false
3: while converged = false do
4:    $probability = \text{GetAgentCountProbability}(\pi_0, N-1)$ 
5:    $x_1 = \text{GetBestResponse}(probability)$ 
6:    $x_1^t(s, n, b, a) = (x_0^t(s, n, b, a) \cdot i +$ 
7:      $x_1^t(s, n, b, a)) / (i + 1)$ ,  $\forall t, s, a, n$ 
8:    $\pi_1^t(s, n, b, a) = x_1^t(s, n, b, a) / \sum_a x_1^t(s, n, b, a)$ ,  $\forall t, s, n, b, a$ 
9:   if  $|\pi_0 - \pi_1| \leq \epsilon_\phi$  then
10:     converged = TRUE
11:   else
12:      $x_0 = x_1$ ,  $\pi_0 = \pi_1$ 
13:   end if
14:    $i + 1$ 
15: end while

```

In this section, we describe the Fictitious Play (FP [2]) approach for computing approximate symmetric Nash Equilibrium (NE) solutions in FSRU.

Fictitious play is an iterative approach, where an agent plays best response against aggregate policy (aggregated over all the previous iterations) of all other agents at each iteration. Algorithm 1 provides this fictitious play algorithm for FSRU building on work by Varakantham et al. [27], where the best response for an agent (line 7) is computed by using Algorithm 2.

Towards achieving a flexible shift optimization, we embed the policy with shift information of agents, i.e., policy for any agent is defined as, $\pi^t(s, n, b, a)$, where, t is time horizon, s is state, a is action, n is number of hours served by the agents and b is breaks taken by the agent.

$$\sum_a \pi^t(s, n, b, a) = 1 \quad \forall t, s, n, b$$

Algorithm 2 GetBestResponse(p)

s_S : sink state, s_{-S} : not sink state

a_S : sink action, a_{-S} : not sink action

M : maximum allowed operating hours

B : maximum allowed breaks in operation

$$\max \sum_{t,s,n,b,a,i} x^t(s, n, b, a) \cdot p_i^t(s) \cdot R^t(s, a, i + 1) \quad (3)$$

$$\mathbf{s.t.} \sum_a x^0(s_S, 0, 0, a) = 1 \quad (4)$$

$$\sum_{n \in (1, M), b \in (1, B), a} x^0(s_S, n, b, a) = 0 \quad (5)$$

$$\begin{aligned} & \sum_a x^t(s_{-S}, n, b, a) \\ &= \sum_{s', a'} x^{t-1}(s'_{-S}, n-1, b, a'_{-S}) \cdot p_i^{t-1}(s') \cdot \phi_{i+1}^{t-1}(s', a, s) \\ & \quad + \sum_{s', a_S} x^{t-1}(s'_{-S}, n-1, b-1, a_S) \cdot p_i^{t-1}(s') \\ & \quad \cdot \phi_{i+1}^{t-1}(s', a, s) + x^t(s_S, n, b, a = s), \quad \forall t, n, s_{-S} \end{aligned} \quad (6)$$

$$\sum_a x^t(s_S, n, b+1, a) = \sum_{s'} x^{t-1}(s', n, b, a_S), \quad \forall t \in (1, T), n \quad (7)$$

$$\sum_{t,s,a_{-S}} x^t(s, n = M, b, a_{-S}) = 0 \quad (8)$$

$$\sum_{t,s,a_{-S}} x^t(s, n, b = B, a_{-S}) = 0 \quad (9)$$

By embedding shift information inside policy we allow for agents operating hours to be split into multiple parts. We provide a FP-based mechanism (Algorithm 1) to compute an approximate equilibrium solution if the approach converges. This FP process at each iteration computes best response (line 5) against aggregate policy of all other agents (line 4). It uses Algorithm 2 for computing best response and is based on the dual LP formulation for solving MDPs. It is more complicated than a basic MDP, because of the flexible operating hours constraint (breaks).

There are two key challenges. First, policy space has two new dimensions [$\pi^t(s, a) \rightarrow \pi^t(s, n, b, a)$], which increases the complexity of fictitious play process significantly. Second, in best response computation, initial distribution is not known and agents can leave (take a break) and join the simulation multiple times (we do not know when and where agent starts its operation and agent can start multiple times after taking a break).

To handle issue on the initial distribution, we introduce a dummy node for agents. At the beginning of the process, i.e., at time step 0, all agents start their operation from a node decided by policy, to avoid initial distribution issue (where initial distribution is unknown) all agents are assumed to be at dummy node and policy will decide to move the taxi to any other node in the next time step. Agents can move from and to the dummy node. This ensures that the decision problem is no longer the initial distribution, but the action to take in each state. These changes are incorporated into Algorithm 2 for best response computation.

5 BEST RESPONSE VIA SIMULATION

For one instance, the FP process requires many iterations to converge and in each iteration it has to go through a time consuming best response computation. Due to this, it can take weeks to compute a solution. More importantly, the solutions computed are not effective if the process is stopped early. In addition, we also have to run multiple instances with different input data values to get the best strategy.

Therefore, for efficient and effective computation, we replace the Algorithm 2 with a simulation-based best response method in Algorithm 4. Simulation-based best response takes policy as input, as the policy is simulated by all agents to identify top trajectories. This is unlike in the FP approach of the previous section, where agent count probability (line 6 of Algorithm 2) is the input.

In, Algorithm 4, we first compute revenue weighted occupancy measure over k-best response paths (line 1) by using Algorithm 3 and form the path matrix (line 2) which contains no information (marked as NAN or NULL) for non-visited states-actions and the weighted occupancy measure for visited states-actions. Algorithm 3 computes the weighted best response path for top-k agents. We first simulate agent policies (in line 2) and identify the paths (or trajectories) taken by top-k agents (lines 3-4) in terms of revenue. We then compute one path from top-k paths, weighted according to their revenues (line 5).

Algorithm 3 GetBestResponsePath-Sim-Topk(π)

- 1: P, R: set of agent paths, R: set of agent revenue
- 2: P, R = SimulateAgents(π)
- 3: P, R = SortPathAndRevenue_RevenueWise(P, R)
- 4: $TopP_k, TopR_k = GetTopKPathsAndRevenue(P, R)$
- 5: weightedPath = getWeightedPath($TopP_k, TopR_k$)
- 6: **return** weightedPath

Algorithm 4 GetBestResponsePolicy-Sim-Topk(π)

- 1: weightedPath = GetBestResponsePath-Sim-Topk(π)
- 2: pathMatrix = getPathMatrix(weightedPath)
- 3: pathMatrix = Impute(pathMatrix)
- 4: $\pi = GetPolicy(pathMatrix)$
- 5: **return** π

Simulation-based best response computation developed here is built on the idea of “Replicator Dynamics” [32] from the “Evolutionary Game Theory”. Replicator dynamics has two key properties. First, if a trajectory/path is absent from the population, then it must always have been absent and will always be absent at any time in the future. Second, if a trajectory/path is present in the population, then it must always have been present and will always remain present at any time in the future.

When we compute best response using simulation, the best response is bounded by the input policy, e.g., if in a state s , the probability for taking an action a is 0, then the best response agent will never be able to take action a in state s , and this may lead to sub optimal best response because agent is not able to explore all possibilities. To ensure best response computation is able to explore

all possibilities before making its decision, instead of using a random starting policy (as done in line 1 of Algorithm 1), we design a uniform starting policy that assigns equal importance along both spatial as well as temporal dimensions (given the constraints on maximum hours to serve and maximum number of breaks) of state space.

If we only consider the best path, only a small part of policy space will be updated in one iteration of FP process, and it will result in extremely slow policy update. To remedy this, instead of only considering “the best path”, we compute top k-best paths at each step, and policy update is done using a (weighted) average over k-best response paths. In experiments we used k as 500 (there are around 20,000 agents in the system). Even with k-best responses, best response paths do not provide enough information to extract a complete policy out of it. Thus, we explore policy completion approaches.

5.1 Policy Completion

Policy completion can be treated as missing data imputation problem. The “right” methods for handling missing data depend strongly on why data is missed. Therefore, it is important to describe the different types of missing data and more specifically the type of missing data in the policy completion problem.

Rubin [22] classified missing data problems into three categories: i) MCAR: where data is missing due to unrelated reason, e.g., due to not collecting required information or collecting incomplete or false information. ii) MAR: where the missing data may only depend on the observable variables. iii) MNAR: where missing data depends on the values of other variables and its own value. Formally, if $X = (X_{obs}, X_{mis})$ is a data matrix where X_{obs}, X_{mis} represents observed and missing values respectively. M represents missing data pattern. Then,

$$\text{MCAR} : P(M|X) = P(M), \quad \forall X$$

$$\text{MAR} : P(M|X) = P(M|X_{obs}), \quad \forall X_{mis}$$

$$\text{MNAR} : P(M|X) = P(M|X_{obs}, Y_{mis}), \quad \exists Y_{mis} \in X_{mis}$$

In best response policy, policy value depends on time(t) and state(s) information along with number of hours served(m) and breaks taken(b), since

$$\sum_a \pi^t(s, m, b, a) = 1 \quad \forall t, s, m, b$$

Therefore missing value for an action is also affected by other actions, which may be missing as well. Therefore in this case, missing values in policy belong to the MNAR category. It is shown that non-parametric “missForest” that is based on the unsupervised imputation is robust for MNAR category of missing data [5].

We now describe the different imputation methods that complete an incomplete policy matrix as input, where incomplete values are usually marked with NAN or NULL.

5.1.1 Matrix Factorization. Matrix factorization (MF) [10] is a strategy for imputing missing values in a matrix. MF takes a partially observed $n \times m$ matrix X as input. O denotes the set of indices in X whose values are known. The goal of MF is to impute missing values in X (i.e., $(i, j) \notin O$). MF assumes that if X , were it fully-known, would be effectively low-rank. That is to say, X can be well

approximated by a matrix $Y = U^T V$ of low-rank r (r is a hyper-parameter), where $r \ll \min(m, n)$, U and V are matrices with r rows, and u_i and v_j are latent factors that represents i^{th} row and j^{th} column of matrix X . Then U and V can be estimated using following optimization:

$$U, V = \operatorname{argmin}_{U^*, V^*} \sum_{i, j \in O} (x_{ij} - u_i^{T*} v_j^*)^2 + \lambda (\|U^*\|_F^2 + \|V^*\|_F^2)$$

where, the second term is the regularization term used to reduce over-fitting. To impute the missing values in X , it just needs to compute the inner-product, $u_i^T v_j$ of the learned latent factors.

5.1.2 MissForest. MissForest is one of the best imputation techniques, and is considered as state of art technique for missing data imputation. It is based on random forest imputation algorithm for missing data. This imputation technique works as follows:

- (1) Initialization: Impute all missing data using the mean (for continuous variables) or mode (for categorical variables) value in each data column.
- (2) Multiple Imputation Step:
 - Impute: For each variable with missing values, fit a random forest on the observed part and predicts the missing part.
 - Stop: Repeat this process of training and predicting the missing values until a stopping criterion is met, or a maximum number of iterations is reached.

MissForest can be applied to mixed data types without any need for pre-processing (no standardization, normalization, scaling, data splitting, etc) on the dataset. It has been proven robust against noises in dataset, due to effective build-in feature selection of random forests. In MissForest, Imputation time, increases with the number of observations, predictors and number of predictors containing missing values. It is based on random forest therefore it inherits the lack of interpretability usually seen in random forests techniques.

MissForest has provided better imputation results in several studies as compared to standard imputation methods [7, 21, 25]. In study conducted by [31] missForest was found to be consistently produce the lowest imputation error when compared against other state of art imputation methods such as k-nearest neighbors (k-NN) imputation and “MICE” [26]. Apart from missForest imputation method, we also experimented with other state of art imputation methods for policy completion.

5.1.3 Supervised Imputation (Algorithm 5). Here we provide a supervised imputation method similar to missForest (Algorithm5), where we replace the time consuming random forest regression model with the ground-truth policy. The ground truth policy evolves as planning process goes through the iterations of the FP process. In our experiments we used the average policy learned through fictitious play process as ground-truth policy. As we do not need to train a regression model here, this method is many fold faster than missForest imputation method, while providing comparable results.

5.1.4 Multivariate imputation by chained equations (MICE). MICE operates under the assumption that given the variables used in the imputation procedure, the missing data are Missing At Random (MAR). Implementing MICE when data are not MAR could result in biased estimates. MICE runs a series of regression models to

Algorithm 5 Supervised Imputation**Require:** X as $n \times p$ matrix

- 1: $k \leftarrow$ vector of sorted indices of columns in X w.r.t. increasing amount of missing values;
- 2: $X_{old}^{imp} \leftarrow$ store previously imputed matrix;
- 3: **for** s in k **do** **do**
- 4: extract $y_{mis}^{(s)}$ from ground-truth policy using $x_{mis}^{(s)}$;
- 5: $X_{new}^{imp} \leftarrow$ update imputed matrix, using predicted $y_{mis}^{(s)}$;
- 6: **end for**
- 7: **return** the imputed matrix X^{imp}

model each variable with missing data conditioned upon the other variables in the data. This means each variable can be modeled according to its distribution. It works as follows:

- (1) Perform a simple imputation for every missing value, e.g., imputing the mean, treat is as “place holders”.
- (2) Set the “place holder” imputed value for one variable (“ v_{miss} ”) back to missing.
- (3) Treat “ v_{miss} ” as dependent variable and all the other variables as independent variable in a regression model. Regress the observed values from the variable “ v_{miss} ” in previous step based on the other variables in the imputation model, which may or may not contain all variables in the dataset.
- (4) Replace the missing values for “ v_{miss} ” with predictions from the regression model. In subsequent steps, when “ v_{miss} ” is used as an independent variable in the regression models for other variables, both observed as well as imputed values will be used.
- (5) Repeat steps 2–4 for each variable with a missing value. Once every missing value is imputed, it is considered as one cycle. At the end of a cycle, replace all of the missing values with the predictions from regression model. Repeat it for multiple cycles.

5.1.5 Generative Adversarial Imputation Nets (GAIN). It is relatively new approach to use generative adversarial networks (GANs) for handling missing data. GANs can learn the hidden data distribution very well and through feedback loop between Generator and Discriminator it can result into high accuracy outputs, and therefore it can provide good results for missing data imputation.

Generative Adversarial Imputation Nets (GAIN) is a very popular GAN architecture for missing data imputation. The idea behind GAIN is very simple: 1) Generator takes the real data as input and imputes the missing values in it. 2) Discriminator takes the imputed data, and tries to figure out which data is imputed.

Generator in GAIN usage same parametric function as any generator (in a GAN network), but it takes more inputs, so that it can fill the missing values close to the true distribution of the data. It generates and uses three different matrices, these matrices are generated from original incomplete dataset. a) The first matrix (X) is the incomplete data matrix that contains all observed data values and 0 for the missing values; b) Second matrix is a indicator matrix for missing values. i.e., this matrix indicates which of the components of X are being observed (indicated by 1) and which are missing (indicated by 0); c) The last matrix is a random noise matrix that

has 0 at the positions of the observed values and random values at the position of missing values.

Discriminator in GAIN tries to distinguish between observed values and imputed value for every data point in \hat{X} (imputed X). It takes \hat{X} as well as a hint matrix as input. The output of Discriminator is a mask matrix that gives the probability of each data point in \hat{X} being observed. The hint matrix provides Discriminator some of the imputed values and some observed values, and therefore it provides support to Discriminator in distinguishing whether a value is imputed or observed.

6 EXPERIMENTAL RESULTS

In this section we provide comparison/analysis of our simulation based best response (SBD) approach equipped with the five imputation methods against an existing approach for strategic planning [9]. We first provide the details of data set and the evaluation method used in this work. Then we provide comparison results for solution quality and run-time of our approaches against the benchmark approach. Next, we analyzed the results to identify best policy for different scenarios. At the end we provide robustness analysis for our different approaches. The comparison/analysis is provided on a real world taxi data set of a company that operates in Singapore.

6.1 Dataset Description

In this work, we use a data set for taxis operating in Singapore. We use 3 months of trip data from 2017. The data are generated from approximately more than 20,000 taxis operating everyday. The data set contains raw information about: a) the pickup/drop-off date and timestamp; b) the taxi trip duration; c) the pickup/drop-off coordinates; d) and the start and end zone ids (entire region is divided in 100 zones). Each pair of a pickup and a drop-off point is defined as a taxi trip. After some data filtering, where we removed taxis operating for more 15 hours in a day, total trips correspond to, a total of more than 0.3 million taxi trips every day.

6.2 Evaluation method

To evaluate our method against the benchmark approach [9], we generate policy for each day of training data set and evaluate it on individual days of test data set.

Generate policy for each day: Using a single day of data, we generated policies using different methods: a) Benchmark method [9]; b) Using our SBR method, where SBR is coupled with different imputation techniques studied in this work. An agent can take up to 2 breaks. Policy convergence criteria and maximum number of hours an agent can operate is same across all methods including benchmark approach. We generated such policies for each day in data set.

Training/Testing data set: For all evaluations/analysis done in this work, 1 week of data is treated as training data set and remaining 3 weeks of the same month is treated as testing data set. All results included in the paper are for week 1 as training set and individual days in weeks 2-4 as test set for the month of March and April (unless stated otherwise). We obtained similar results for other combinations of training and test data.

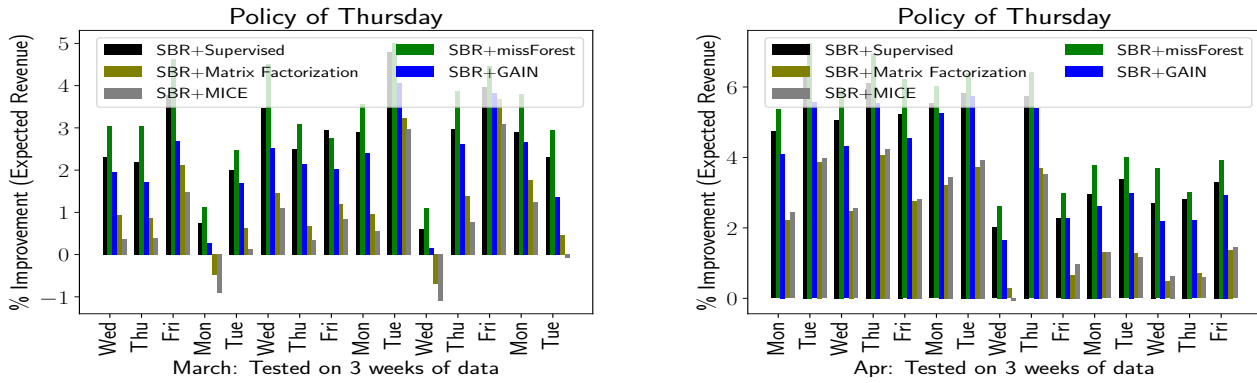


Figure 1: Solution Quality: percentage improvement over benchmark, when tested over 3 test weeks for the policy generated by our simulation based best response method using demand data for Thursday in week 1

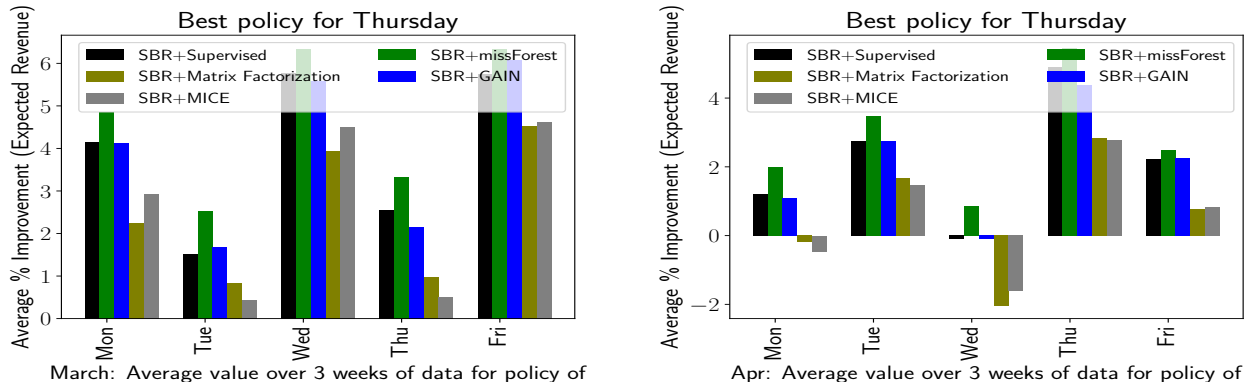


Figure 2: Average Solution Quality (% improvement over baseline) when different weekday policies are tested on Thursday of 3 test weeks on the month of March and Apr

6.3 Quality comparison

Solution quality comparison: We compare the solution quality (expected revenue) of different methods on test data. Weekday policies are evaluated on weekdays, and weekend policies are evaluated on weekends of the 3 testing weeks.

Run-Time comparison: Here we provide an estimate of (on average) run-time taken by different method to generate a converged policy on training data.

Results of Quality Comparison: We generated 7 polices for each day of training data (week 1) with 5 imputation methods over simulation-based best response and the benchmark approach. Key results are as follows:

- Our simulation based best response method enabled with missForest Imputation method provides the best results with up to $\approx 10.5\%$ improvement over the benchmark policy. This is followed closely by GAIN and supervised imputation methods.
- Our simulation method enabled with matrix factorization and MICE imputation do not fare as well, but are mostly better than the benchmark approach.

- Our simulation method enabled with supervised imputation has the least run-time with policies being generated in ≈ 10 minutes. In contrast, benchmark approach takes up to 3 days. MICE, missForest, matrix factorization, and GAIN generate solutions in 3-8 hours, with GAIN being the slowest. We used 100 training epochs for training the GAIN.

6.4 Analysis

Analyze best policy for different days: Next, we analyzed the results to identify best policy for different scenarios.

- Best policy for each day (on average): Based on average performance on test set, we can identify best policy for each day, Table 1 shows the best policy for each day of the week with training data as Week1 (March). Similar results are observed with different settings of training and test data.
- Best policy for each weekdays/weekend: With week1 (March) as training data, policy generated for Wednesday works well on all weekdays and policy generated for Sunday works well on all weekends. Similar results are observed with different data (different month and week) where one policy works well on all weekdays and one policy works well on all weekends.

Day	Mon	Tues	Wed	Thurs	Fri	Sat	Sun
BestPolicy	Fri	Wed	Wed	Fri	Wed	Sun	Sun

Table 1: Best policy for each day (based on avg. performance on 3 test weeks), when trained w/ week 1 data in March.

Robustness analysis: Here we provide robustness analysis of our approaches. We ran the worst case analysis of different polices on 3 test weeks. Detailed result for worst case, best case and average case scenarios with all imputation methods is shown in Table 2 for training week 1 (March) when tested on weekdays of 3 test weeks, Table 3 for training week 1 (March) when tested on weekends of 3 test weeks, Table 5 for training week 1 (Apr) when tested on weekdays of 3 test weeks and Table 4 for training week 1 (Apr) when tested on weekends of 3 test weeks. In summary, worst case performance of missForest imputation with best policy (one policy for all weekdays/weekend of test days) is as follows, It is providing average improvement of 4.47% and 3.48% for weekdays and weekends in March, and 2.48% and 3.62% for weekdays and weekends in April.

Method	Average Case	Best Case	Worst Case
Supervised	3.86	10.67	-0.22
Matrix	2.5	9.25	-1.57
MICE	2.59	9.81	-1.13
missForest	4.47	10.57	0.58
GAIN	3.84	10.93	-0.01

Table 2: [Training data: Week 1 (March)]: Average, Best and Worst Case Performance (% improvement over baseline) when tested on weekdays of 3 test weeks

Method	Average Case	Best Case	Worst Case
Supervised	2.49	3.99	0.15
Matrix	0.04	1.47	-3.71
MICE	0.0	1.4	-3.63
missForest	3.48	5.04	1.7
GAIN	1.76	3.04	-0.85

Table 3: [Training data: Week 1 (March)]: Average, Best and Worst Case Performance (% improvement over baseline) when tested on weekend of 3 test weeks

Method	Average Case	Best Case	Worst Case
Supervised	1.8	6.6	-2.72
Matrix	0.16	5.06	-4.55
MICE	0.15	4.73	-4.14
missForest	2.48	7.24	-1.57
GAIN	1.72	6.78	-2.34

Table 4: [Training data: Week 1 (Apr)]: Average, Best and Worst Case Performance (% improvement over baseline) when tested on weekdays of 3 test weeks

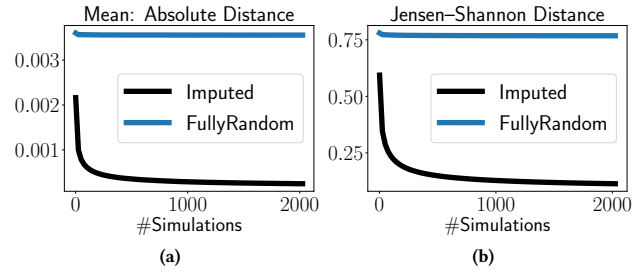


Figure 3: Distance of occupancy matrix from best response (without imputation) occupancy matrix

Method	Average Case	Best Case	Worst Case
Supervised	2.25	6.41	-2.21
Matrix	-0.76	3.08	-5.67
MICE	-0.4	3.22	-4.91
missForest	3.62	7.77	-1.38
GAIN	1.11	5.76	-4.73

Table 5: [Training data: Week 1 (Apr)]: Average, Best and Worst Case Performance (% improvement over baseline) when tested on weekend of 3 test weeks

7 DISCUSSION

When generating policy from a given set of trajectories [6], a theoretical sound way of evaluating correctness is the distance between occupancy measures (visitation frequencies of different state, action pairs) of the policy and that of the trajectories. In Figure 3, we plot the distance as number of simulations is increased for two measures: mean absolute distance and Jensen-Shannon distance.

We compare the distance of best response policy from (a) our imputation method; and (b) a fully random policy. For this, we simulated each policy multiple times and compared distance between occupancy matrices (visitation frequencies of different state, action pairs). Occupancy matrix generated by our imputation method policy is much closer to the best response trajectories than the fully random policy. As number of simulations is increased, distance between occupancy matrix generated by imputed policy and the best response policy decreases.

8 CONCLUSION

In this work we introduce mechanism to provide better flexibility for selfish agent while improving the performance of the entire system. We introduced a simulation based (faster) equilibrium computation method. We studied and analyzed different imputation methods and show that a good imputation method coupled with a well designed simulation based best response computation can help in achieving better symmetric equilibrium for large scale systems, in a time efficient manner. In our experiments we also show that we can reuse policy for multiple days. We analyzed and provided, (on average) best policy for each day of the week on data set used in this experiment.

9 ACKNOWLEDGMENT

This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-016).

REFERENCES

- [1] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research* 27, 4 (2002), 819–840.
- [2] George W Brown. 1951. Iterative solution of games by fictitious play. *Activity analysis of production and allocation* 13, 1 (1951), 374–376.
- [3] Jiarui Gan and Bo An. 2017. Game-theoretic considerations for optimizing taxi system efficiency. *IEEE Intelligent Systems* 32, 3 (2017), 46–52.
- [4] Yong Gao, Dan Jiang, and Yan Xu. 2018. Optimize taxi driving strategies based on reinforcement learning. *International Journal of Geographical Information Science* 32, 8 (2018), 1677–1696.
- [5] Chao-Yu Guo, Ying-Chen Yang, and Yi-Hau Chen. 2021. The Optimal Machine Learning-Based Missing Data Imputation for the Cox Proportional Hazard Model. *Frontiers in Public Health* 9 (2021), 881. <https://doi.org/10.3389/fpubh.2021.680054>
- [6] Jonathan Ho and Stefano Ermon. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems* 29 (2016), 4565–4573.
- [7] Shangzhi Hong and Henry S Lynn. 2020. Accuracy of random-forest-based imputation of missing data in the presence of non-normality, non-linearity, and interaction. *BMC medical research methodology* 20, 1 (2020), 1–12.
- [8] Sebastian Hörl, Claudio Ruch, Felix Becker, Emilio Frazzoli, and Kay W Axhausen. 2019. Fleet operational policies for automated mobility: A simulation assessment for Zurich. *Transportation Research Part C: Emerging Technologies* 102 (2019), 20–31.
- [9] Rajiv Ranjan Kumar, Pradeep Varakantham, and Shih-Fen Cheng. 2021. Adaptive Operating Hours for Improved Performance of Taxi Fleets. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 728–736.
- [10] Daniel D Lee and H Sebastian Seung. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 6755 (1999), 788–791.
- [11] Meghna Lowalekar, Pradeep Varakantham, Supriyo Ghosh, Sanjay Dominik Jena, and Patrick Jaillet. 2017. Online repositioning in bike sharing systems. In *Twenty-seventh international conference on automated planning and scheduling*.
- [12] Meghna Lowalekar, Pradeep Varakantham, and Patrick Jaillet. 2018. Online spatio-temporal matching in stochastic and dynamic domains. *Artificial Intelligence* 261 (2018), 71–112.
- [13] Michal Maciejewski and Joschka Bischoff. 2015. Large-scale microscopic simulation of taxi services. *Procedia Computer Science* 52 (2015), 358–364.
- [14] Michal Maciejewski, Joschka Bischoff, and Kai Nagel. 2016. An assignment-based approach to efficient real-time city-scale taxi dispatching. (2016).
- [15] J-F Mertens and Abraham Neyman. 1981. Stochastic games. *International Journal of Game Theory* 10, 2 (1981), 53–66.
- [16] Fei Miao, Shuo Han, Shan Lin, John A Stankovic, Desheng Zhang, Sirajum Munir, Hua Huang, Tian He, and George J Pappas. 2016. Taxi dispatch with real-time sensing data in metropolitan areas: A receding horizon control approach. *IEEE Transactions on Automation Science and Engineering* 13, 2 (2016), 463–478.
- [17] Luis Moreira-Matias, João Gama, Michel Ferreira, and Luis Damas. 2012. A predictive model for the passenger demand on a taxi network. In *2012 15th International IEEE Conference on Intelligent Transportation Systems*. IEEE, 1014–1019.
- [18] Abraham Neyman, Sylvain Sorin, and S Sorin. 2003. *Stochastic games and applications*. Vol. 570. Springer Science & Business Media.
- [19] Takuma Oda and Carlee Joe-Wong. 2018. Movi: A model-free approach to dynamic fleet management. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2708–2716.
- [20] Meng Qu, Hengshu Zhu, Junming Liu, Guannan Liu, and Hui Xiong. 2014. A cost-effective recommender system for taxi drivers. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 45–54.
- [21] Burim Ramosaj and Markus Pauly. 2019. Predicting missing values: a comparative study on non-parametric approaches for imputation. *Computational Statistics* (2019), 1–24.
- [22] Donald B Rubin. 1976. Inference and missing data. *Biometrika* 63, 3 (1976), 581–592.
- [23] Sandhya Saisubramanian, Pradeep Varakantham, and Hoong Chuin Lau. 2015. Risk based optimization for improving emergency medical systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.
- [24] Lloyd S Shapley. 1953. Stochastic games. *Proceedings of the national academy of sciences* 39, 10 (1953), 1095–1100.
- [25] Fei Tang and Hemant Ishwaran. 2017. Random forest missing data algorithms. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 10, 6 (2017), 363–377.
- [26] Stef Van Buuren and Karin Groothuis-Oudshoorn. 2011. mice: Multivariate imputation by chained equations in R. *Journal of statistical software* 45, 1 (2011), 1–67.
- [27] Pradeep Varakantham, Shih-Fen Cheng, Geoff Gordon, and Asrar Ahmed. 2012. Decision support for agent populations in uncertain and congested environments. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.
- [28] Mohammad M Vazifeh, Paolo Santi, Giovanni Resta, Steven H Strogatz, and Carlo Ratti. 2018. Addressing the minimum fleet problem in on-demand urban mobility. *Nature* 557, 7706 (2018), 534–538.
- [29] Tanvi Verma and Pradeep Varakantham. 2020. Correlated learning for aggregation systems. In *Uncertainty in Artificial Intelligence*. PMLR, 60–70.
- [30] Tanvi Verma, Pradeep Varakantham, Sarit Kraus, and Hoong Chuin Lau. 2017. Augmenting decisions of taxi drivers through reinforcement learning for improving revenues. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 27. 409–417.
- [31] Akbar K Waljee, Ashin Mukherjee, Amit G Singal, Yiwei Zhang, Jeffrey Warren, Ulysses Balis, Jorge Marrero, Ji Zhu, and Peter DR Higgins. 2013. Comparison of imputation methods for missing laboratory data in medicine. *BMJ open* 3, 8 (2013), e002847.
- [32] Jörgen W Weibull. 1997. *Evolutionary game theory*. MIT press.
- [33] Zhe Xu, Zhixin Li, Qingwen Guan, Dingshui Zhang, Qiang Li, Junxiao Nan, Chunyang Liu, Wei Bian, and Jieping Ye. 2018. Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 905–913.
- [34] Nicholas Jing Yuan, Yu Zheng, Lihang Zhang, and Xing Xie. 2012. T-finder: A recommender system for finding passengers and vacant taxis. *IEEE Transactions on knowledge and data engineering* 25, 10 (2012), 2390–2403.
- [35] Kai Zhang, Zhiyong Feng, Shizhan Chen, Keman Huang, and Guiling Wang. 2016. A framework for passengers demand prediction and recommendation. In *2016 IEEE International Conference on Services Computing (SCC)*. IEEE, 340–347.