Spatio-Temporal Pricing for Ridesharing Platforms*

Hongyao Ma
Harvard SEAS
Cambridge, MA
hma@seas.harvard.edu

Fei Fang
Carnegie Mellon University
Pittsburge, PA
feifang@cmu.edu

David C. Parkes
Harvard SEAS
Cambridge, MA
parkes@eecs.harvard.edu

ABSTRACT

Ridesharing platforms match drivers and riders to trips, using dynamic prices to balance supply and demand. A challenge is to set prices that are appropriately smooth in space and time, in the sense that drivers will choose to accept their dispatched trips, rather than drive to another area or wait for higher prices or a better trip. We introduce the *Spatio-Temporal Pricing (STP) mechanism.* The mechanism is incentive-aligned, in that it is a subgame-perfect equilibrium for drivers to accept their dispatches, and the mechanism is welfare-optimal, envy-free, individually rational and budget balanced from any history onward. We work in a complete information, discrete time, multi-period, multi-location model, and prove the existence of anonymous, origin-destination, competitive equilibrium (CE) prices. The STP mechanism employs driver-pessimal CE prices, and the proof of incentive alignment makes use of the $M^{\natural}$ concavity of min-cost flow objectives. We also prove that there can be no dominant-strategy mechanism with the same economic properties. An empirical analysis conducted in simulation suggests that the STP mechanism achieves significantly higher social welfare than a myopic pricing mechanism, and highlights the failure of incentive alignment due to non-smooth prices in myopic mechanisms.

CCS CONCEPTS

• **Theory of computation → Algorithmic mechanism design**; **Market equilibria**; **Computational pricing and auctions**;

KEYWORDS

Ridesharing, mechanism design, dynamic pricing, matching

ACM Reference Format:

Hongyao Ma, Fei Fang, and David C. Parkes. 2018. Spatio-Temporal Pricing for Ridesharing Platforms. In *Proceedings of Workshop on Two-sided Marketplace Optimization:: Search, Pricing, Matching & Growth (TSMO'18).* ACM, Los Angeles, CA, USA, 9 pages.

1 INTRODUCTION

Ridesharing platforms such as Uber and Lyft are disrupting more traditional forms of transit. These are two-sided platforms, with both riders and drivers in a customer relationship with the platform. When a rider opens the app and enters an origin and destination, these platforms quote a price and an estimated wait time. If a rider requests the ride, the platform offers the pick-up opportunity to each of a sequence of nearby drivers until a driver accepts. At this point, if neither side cancels and the driver completes the pick-up then the trip begins. Once the trip is complete, payment is made from the rider to the driver through the platform.[1]

Both Uber and Lyft emphasize the importance of providing reliable transportation. Uber's mission is "to connect riders to reliable transportation, everywhere for everyone."[2] Lyft's mission is stated as "to provide the best, most reliable service possible by making sure drivers are on the road when and where you need them most."[3] Whereas taxi systems have reliable pricing but unreliable service [19], ridesharing platforms make use of dynamic pricing to achieve reliable service. These platforms also emphasize the flexibility for drivers, e.g. Uber advertises itself as "work that put you first— drive when you want, earn what you need".[4]

A known challenge with current ridesharing platforms is that incentives may fail to be aligned for drivers. A particular concern, is that trips may be mis-priced relative to other trip opportunities.[5] This leads to a loss in reliability— drivers may decline particular kinds of trips, or simply choose not to participate from certain locations at certain times, given that they have the flexibility to make ongoing decisions about participation. In this way, poorly designed pricing and dispatching systems undercut the ability of these platforms to fulfill their mission.

One kind of failure arises because of incorrect spatial pricing. Consider for example that if the price is substantially higher for trips that start in location A than an adjacent location B, drivers in location B that are close to the boundary will decline trips. This leads to drivers "chasing the surge"— turning off a ridesharing app while relocating to another location where prices are higher. This results in a loss in welfare, with even high willingness-to-pay riders in location B unable to access reliable transportation.

Problems with spatial pricing also arise because prices do not correctly factor market conditions at the destination of a trip. It has been standard practice to use *origin-based pricing*, with unit prices that depend on market conditions only at origin. Suppose that a driver in location A could be matched to a trip to a quiet suburb

*Thanks to Keith Chen, Michael Ostrovsky, and the participants at MatchUp 2017, ACM EC17, IFORS 2017, the Harvard IO workshop, the USC CAIS seminar, and the Uber MODS workshop for helpful feedback.

[1] The actual practice is somewhat more complicated, in that platforms may operate multiple products for example high-end cars, trips shared by multiple riders, etc. Moreover, drivers within even a single class are differentiated (e.g., cleanliness of car, skill of driving), as are riders (e.g., politeness, loud vs. quiet, pick-up and drop-off neighborhoods.) We ignore these effects, and assume that all riders are equivalent from the perspective of drivers and all drivers equivalent from the perspective of riders.
[2] https://www.uber.com/legal/community-guidelines/us-en/, visited September 1, 2017.
[3] https://help.lyft.com/hc/en-us/articles/115012926227, visited September 1, 2017.
[4] https://www.uber.com/drive/, visited December 12, 2017.
[5] There also exist other incentive problems, including inconsistencies across classes of service, competition among platforms, and drivers' off-platform incentives. We only model a single class of service, we ignore cross-platform competition, and we do not model location- or time-dependent opportunity cost.

B and or a trip to the busy downtown area C. We can expect the continuation payoff in location B to be smaller that in location C, because demand is greater in C and thus both a lower wait time and higher prices. Because of this, trips to C will be preferred to trips to B and there is a market failure, with even high willingness-to-pay (A, B) riders unable to access reliable transportation.

A second kind of market failure arises because of incorrect temporal pricing. Consider for example that a sports event will end soon, and drivers can anticipate that prices will increase in order to balance supply and demand. In this case, drivers will decline trips in anticipation of the surge. This results in a market failure, with even high willingness-to-pay riders unable to get matched to trips before the end of the sports event.

To avoid market failure, prices need to be appropriately "smooth" in both space and time in the sense that drivers who retain the flexibility to choose when to work will always choose to accept any trip to which they are dispatched. In this way, smooth prices are responsive to a central challenge of market design for ridesharing platforms, which is to optimally orchestrate trips without the power to tell drivers what to do. Because both drivers and riders are in a customer relationship with the platform, all parties must continually agree to accept any proposed action.[6]

We address the problem of incentive alignment for drivers in the absence of time-extended contracts. This recognizes the importance of real-time flexibility, allowing drivers to decide whether or not to provide rides at any moment. To this aim, we propose the *Spatio-Temporal Pricing (STP) mechanism*, under which accepting the mechanism's dispatches at all times forms a subgame-perfect equilibrium (SPE). From any history onward, the STP mechanism is individually rational, budget balanced, welfare optimal, and also envy-free, meaning that any pair of drivers in the same location at the same time, and not currently on a trip has the same continuation payoff. [7] We also give an impossibility result, that there can be no dominant-strategy mechanism with the same economic properties, and show via simulations that the STP mechanism achieves significantly higher social welfare than a myopic pricing mechanism.

We work in a complete information, discrete time, multi-period, multi-location model, and allow asymmetric, time-varying trip times, non-stationarity in demand, and riders with different values for completing a trip. At the beginning of each time period, based on the history, current driver positioning, and current and future demand, the mechanism dispatches a trip to each available driver (including the possibility of relocation), and determines a payment if the driver follows the dispatch. Each driver can decide whether to follow or to decline and stay or relocate to any location. After observing the driver actions in a period, the mechanism collects payment from the riders and makes payment to the drivers.

The main assumptions are (i) complete information about supply and demand over a planning horizon, (ii) impatient riders that need to be picked-up at a particular time and location (and without preferences over drivers), and (iii) drivers who are willing to take trips until the end of the planning horizon and have no intrinsic preference for driving in one location over another. We do allow drivers to first become available to dispatch from diverse locations, with different drivers coming online at different times. We model a rider's value as her willingness-to-pay over-and-above a base payment that covers driver costs. In this way, the prices that are computed through the STP mechanism correspond to the "surge" price, over-and-above the base payment for a trip.

We prove the existence of anonymous, origin-destination, competitive equilibrium (CE) prices. By adopting origin-destination prices, the unit price of a trip can depend on market conditions at both the origin and destination. The STP mechanism employs driver-pessimal CE prices, computing a driver-pessimal CE plan at the beginning of the planning horizon and after driver deviations. This induces an extensive-form game among the drivers, where the total payoff to each driver is determined by the mechanism's dispatch and payment rules. The use of driver pessimal CE prices is an essential part of achieving smooth prices, and the proof of incentive alignment also makes use of the $M^{\natural}$ concavity of min-cost flow objectives [17]. The same connection to min-cost flow leads to an efficient algorithm to compute an optimal dispatching and prices, and operationalize the STP mechanism.

We introduce the model in Section 2, and illustrate through an example that a baseline myopic pricing mechanism fails to be welfare-optimal or incentive aligned. In Section 3, we formulate the optimal planning problem and establish integrality properties of a linear-programming relaxation (Lemma 3.1). We show that a plan with anonymous trip prices is welfare-optimal if and only if it forms a competitive equilibrium (CE) (Lemma 3.2), and that optimal CE plans exist and are efficient to compute. We also prove that drivers' total payments among all CE plans form a lattice (Lemma 3.3).

We prove our main result in Section 4, that the STP mechanism is subgame-perfect incentive compatible, and is also individually rational, budget balanced, envy-free, and welfare optimal from any history onward (Theorem 4.1). We also prove an impossibility result, that no dominant-strategy mechanism has the same economic properties (Theorem 4.2). An empirical analysis conducted through simulation (Section 5) suggests that the STP mechanism can achieve significantly higher social welfare than a myopic pricing mechanism, and highlights the failure of incentive alignment due to non-smooth prices in myopic mechanisms.

We compare in Section 5 through simulation, the performance of the STP mechanism against the myopic pricing mechanism for an economy modeling the end of an event. The STP mechanism achieves significantly higher social welfare, whereas under the myopic pricing mechanism, drivers incur a high regret. We conclude in Section 6 with discussions on the effect of relaxing the model assumptions and directions of future work. Additional examples, discussions, omitted proofs, simulation results and relations to the literature are provided in the full version of this paper [16].

1.1 Related Work

Banerjee et al. [3] take a queuing-theoretic approach to analyze the effect of dynamic pricing on revenue and throughput of ridesharing platforms, assuming a single location and stationary system state.

[6] There is an echo here to *market-oriented programming* [20] and *agoric systems* [11]. There, markets with virtual prices were suggested as a means for achieving optimization in decentralized systems. But whereas this earlier work adopted market-based methods for their ability to optimize, and prices were virtual, there is an additional need in the present setting to align incentives.

[7] The mechanism that we design in this paper balances budget. Alternatively, we many think about the ridesharing platform taking a fixed percentage of the prices. This does not affect the results presented in the paper.

The optimal dynamic pricing strategy does not outperform the optimal static pricing strategy when parameters on supply and demand are correctly estimated, however, dynamic pricing is more robust to fluctuations and mis-estimation of system parameters.

By analyzing the equilibrium outcome under a continuum, stationary model with unlimited driver supply at fixed costs, Bimpikis et al. [6] show that a ridesharing platform's profits and consumer surplus are maximized when the demand pattern across different locations is balanced, and that the platform can benefit significantly from pricing trips differently depending on trip origins. Relatd to this, Banerjee et al. [2] model a shared vehicle system as a continuous-time Markov chain, and establish approximation guarantees for a static, state-independent pricing policy (i.e. fixed prices that do not depend on the spatial distribution of cars), w.r.t. the optimal, state-dependent policy.

The "wild goose chase" phenomena is analyzed by Castillo et al. [7]: when demand significantly exceeds supply, drivers spend too much time driving to pick up the riders instead of having riders in the backseat, leading to decreased revenue and social welfare. This is an effect of the ridesharing platform always dispatching a driver as soon as any rider request a ride. In comparison to setting fixed prices, social welfare is improved if prices can be set based on market conditions, and this solution is shown to be superior to other solutions e.g. limiting the pick-up radius.

Our model differs from these previous works, in that we consider both multiple locations and multiple time periods, with rider demand, rider willingness to pay, and driver supply that can vary across both space and time. This leads to the focus of the present paper on the design of a ridesharing mechanism with prices that are smooth in both time and space.

There are various empirical studies of the Uber platform as a two-sided marketplace [12, 14, 15], analyzing Uber's driver partners, labor market equilibrium and consumer surplus. By analyzing the hourly earnings of drivers on the Uber platform, Chen et al. [9] show that drivers' reservation wages vary significantly over time, and that the real-time flexibility of being able to choose when to work increases both driver surplus and the supply of drivers. In regard to dynamic pricing, Chen and Sheldon [10] show by analyzing the trips provided by a subset of driver partners in several US cities from 2014-2015 that surge pricing increases the supply of drivers on the Uber platform at times when the surge pricing is high. A case study [13] into an outage of Uber's surge pricing during the 2014-2015 New Year's Eve in New York City found a large increase in riders' waiting time after requesting a ride, and a large decrease in the percentage of requests completed.

2 PRELIMINARIES

Let T be the length of the planning horizon, starting at time $t = 0$ and ending at time $t = T$. We adopt a discrete time model, and refer to each time point t as "time t", and call the duration between time t and time $t + 1$ a *time period*. Trips start and end at time points. Denote $[T] = \{0, 1, \ldots, T\}$ and $[T - 1] = \{0, 1, \ldots, T - 1\}$.

Let $\mathcal{L} = \{A, B, \ldots, \}$ be a set of $|\mathcal{L}|$ discrete locations, and we adopt a and b to denote generic locations. For all $a, b \in \mathcal{L}$ and $t \in [T]$, the triple (a, b, t) denotes a *trip* with origin a, destination b, starting at time t. Each trip can represent (i) taking a rider from a to b at time t, (ii) relocating without a rider from a to b at time t, and (iii) staying in the same location for one period of time. Let $\delta : \mathcal{L}\times\mathcal{L} \to \mathbb{N}$ denote the time to travel between locations, so that trip (a, b, t) ends at $t + \delta(a, b)$.[8] We allow $\delta(a, b) \neq \delta(b, a)$ for locations $a \neq b$, modeling asymmetric traffic flows. We assume $\delta(a, b) \geq 1$ and $\delta(a, a) = 1$ for all $a, b \in \mathcal{L}$, and that the triangle inequality holds.[9] Let $\mathcal{T} = \{(a, b, t) \mid a \in \mathcal{L},\ b \in \mathcal{L},\ t \in \{0, 1, \ldots, T - \delta(a, b)\}\}$ denote the set of feasible trips within the planning horizon.

Let $\mathcal{D}$ denote the set of drivers, with $m = |\mathcal{D}|$. Each driver $i \in \mathcal{D}$ is characterized by *type* $\theta_i = (\underline{\tau}_i, \bar{\tau}_i, \ell_i)$— driver i enters the market at time $\underline{\tau}_i$ and location ℓ_i, and plans to exit the market at time $\bar{\tau}_i$ (with $\underline{\tau}_i < \bar{\tau}_i$). Here we make the assumption (S1) that *driver types are known to the mechanism and that all drivers stay until at least the end of planning horizon, and do not have an intrinsic preference over locations.* Each driver seeks to maximize the total payment received over the planning horizon.

Denote $\mathcal{R}$ as the set of riders, each intending to take a single trip during the planning horizon. The *type* of rider $j \in \mathcal{R}$ is (o_j, d_j, τ_j, v_j), where o_j and d_j are the trip origin and destination, τ_j the requested start time, and $v_j \geq 0$ the *value* for the trip. We assume (S2) *that riders are impatient, only value trips starting at τ_j, are not willing to relocate or walk.* The value v_j models the willingness to pay of the rider, over-and-above a base payment for a trip. This base payment is an amount that is collected by the platform, so that in aggregate drivers' costs are covered. Accordingly, the prices we derive in this paper correspond to "surge" prices, over-and-above these base payments. Rider utility is quasi-linear, with utility $v_j - p$ to rider j for a completed trip at (incremental to base) price p.

We assume the platform has complete information about supply and demand over the planning horizon (driver and rider types, including driver entry during the planning horizon). We assume drivers have the same information, and that this is common knowledge amongst drivers. More generally, it would be sufficient that it be common knowledge amongst drivers that the platform has the correct information. Unless otherwise noted, we assume (S1), (S2), and complete, symmetric information throughout the paper.

At each time t, a driver is *en route* if she started her last trip from a to b at time t' (with or without a rider), and $t < t' + \delta(a, b)$. A driver is *available* if she has entered the platform ($t \geq \underline{\tau}_i$) and is not en route. An available driver at time t and location a is able to complete a pick-up at (a, t). We allow a driver to drop-off a rider and pick-up another rider in the same location and time point.

A *path* is a sequence of tuples (a, b, t), representing the trips taken by a driver over the planning horizon. A *feasible path* for driver i starts at $(\ell_i, \underline{\tau}_i)$, with the starting time and location of each successive trip equal to the ending time and location of the previous trip. Let $\mathcal{Z}_i$ denote the set of all feasible paths of driver i, with $Z_{i,k} \in \mathcal{Z}_i$ to denote the k^{th} feasible path. Denote $(a, b, t) \in Z_{i,k}$ if path $Z_{i,k}$ includes (or *covers*) trip (a, b, t). Driver i with path $Z_{i,k}$ is able to pick up rider j if $(o_j, d_j, \tau_j) \in Z_{i,k}$.

As a baseline, we define the following mechanism:

[8] We can allow the distance between a pair of locations to change over time, modeling changes in traffic conditions. This does not affect the results presented in this paper.

[9] With the triangle inequality, $\forall a, b, c \in \mathcal{L}$, $\delta(a, c) \leq \delta(a, b) + \delta(b, c)$, then we have an additional property— riders would not have incentives in the STP mechanism to break a long trip into several shorter trips in order to get a lower price. See Section **??** for discussions on rider incentives.

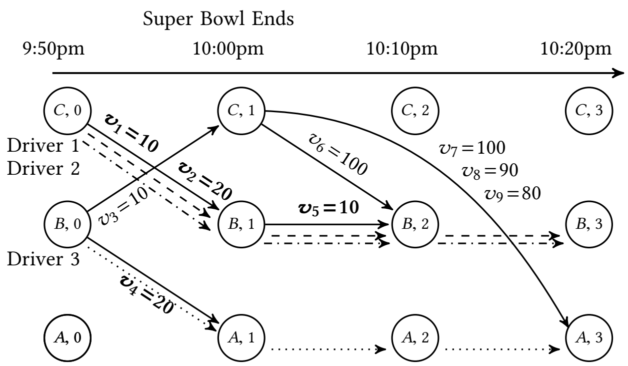

Figure 1: A Super Bowl game: time 0 plan under the myopic pricing mechanism.

Definition 1 (Myopic pricing mechanism)**.** At each time point $t \in [T]$, for each location $a \in \mathcal{L}$, the *myopic pricing mechanism* dispatches available drivers at (a, t) to riders requesting rides from (a, t) in decreasing order of riders' values, and sets a market clearing price $p_{a,t}$ that is offered to all dispatched drivers (and will be collected from riders).

The market clearing price may not be unique, and a fully defined myopic mechanism must provide a rule for picking a particular clearing price. This mechanism has anonymous, origin-based pricing, and is very simple in ignoring the need for smooth pricing, or future supply and demand.

Example 1 (Super Bowl example)**.** Consider the economy illustrated in Figure 1, modeling the end of a sports event. Time $t = 0$ is 9:50pm, and 10 minutes before the Super Bowl ends. There are three locations A, B and C with symmetric distances $\delta(A, B) = \delta(B, A) = \delta(B, C) = \delta(C, B) = 1$ and $\delta(A, C) = \delta(C, A) = 2$. Drivers 1 and 2 enter at location C at time 0, while driver 3 enters at location B at time 0. At time 1, many riders with very high value show up at location C, where the game takes place.

Under the myopic pricing mechanism, at time 0, drivers 1, 2 and 3 are dispatched to pick up riders 1, 2 and 4, respectively. At time 1, one of drivers 1 and 2 picks up rider 5, and the total social welfare is $v_1 + v_2 + v_4 + v_5 = 60$. The paths taken by drivers 1 and 2 are $((C, B, 0), (B, B, 1), (B, B, 2))$ (the dash and dash-dotted) and the path taken by driver 3 is $((B, A, 0), (A, A, 1), (A, A, 2))$ (the dotted). The set of all possible market clearing price for trip $(C, B, 0)$ is $p_{C,B,0} \in [0, 10]$, and the price for $(B, B, 1)$ is $p_{B,B,1} = 0$ since there is excessive supply. The highest total payment to driver 1 under any myopic pricing mechanism would be 10. At time 1, since no driver is able to pick up the four riders at location C, the lowest market clearing prices are $p_{C,B,1} = p_{C,A,1} = 100$. Suppose driver 1 deviates and stays in location C until time 1. The mechanism would then dispatch her to pick up rider 6, and she would be paid the new market clearing price of 100. This is a useful deviation. □

3 OPTIMAL PLANNING

A *plan* describes the paths taken by all drivers until the end of the planning horizon, as well as payments, for riders and drivers, for each trip associated with these paths. Formally, a plan is the 4-tuple (x, z, q, π), where: x is the indicator of rider pick-ups, where for all $j \in \mathcal{R}$, $x_j = 1$ if rider j is picked-up according to the plan, and $x_j = 0$ otherwise; z is a vector of paths, where $z_i \in \mathcal{Z}_i$ is the dispatched path for driver i; q_j denotes the payment made by rider j, and $\pi_{i,t}$ denotes the payment made to driver i at time t; and $\pi_i = \sum_{t=0}^{T} \pi_{i,t}$ denotes the total payment to driver i.

A plan (x, z, q, π) is *feasible* if $\forall (a, b, t) \in \mathcal{T}$, $\sum_{j \in \mathcal{R}} x_j \mathbb{1}\{(o_j, d_j, \tau_j) = (a, b, t)\} \leq \sum_{i \in \mathcal{D}} \mathbb{1}\{(a, b, t) \in z_i\}$, where $\mathbb{1}\{\cdot\}$ is the indicator function. Unless otherwise indicated, when we mention a plan in the rest of the paper, it is assumed to be feasible.

For *budget balance* (BB), we need: $\sum_{j \in \mathcal{R}} q_j \geq \sum_{i \in \mathcal{D}} \pi_i$, with strict budget balance if the inequality holds with equality. A plan is *individually rational* (IR) for riders if $\forall j \in \mathcal{R} : x_j v_j \geq q_j$. A plan is *envy-free for drivers* if $\forall i,\ i' \in \mathcal{D}$, s.t. $\underline{\tau}_i = \underline{\tau}_{i'},\ \ell_i = \ell_{i'} : \pi_i = \pi_{i'}$. A plan is *envy-free for riders* if $\forall j,\ j' \in \mathcal{R}$, s.t. $o_j = o_{j'},\ d_j = d_{j'},\ \tau_j = \tau_{j'} : x_j v_j - q_j \geq x_{j'} v_j - q_{j'}$.

Definition 2 (Anonymous trip prices)**.** A plan uses *anonymous trip prices* if there exist $p = \{p_{a,b,t}\}_{(a,b,t) \in \mathcal{T}}$ s.t. for all $\forall (a, b, t) \in \mathcal{T}$,

(i) all riders taking trip (a, b, t) are charged $p_{a,b,t}$, and there is no payment by riders who are not picked up, and
(ii) all drivers that are dispatched on a trip from a to b at time t (with or without a rider) are paid $p_{a,b,t}$ for the trip.

Given dispatches (x, z) and anonymous trips prices p, all payments are fully determined: the total payment to driver i is $\pi_i = \sum_{(a,b,t) \in z_i} p_{a,b,t}$ and the payment made by rider j is $q_j = x_j p_{o_j,d_j,\tau_j}$. We represent plans with anonymous trip prices as (x, z, p).

Definition 3 (Competitive Equilibrium)**.** A plan with anonymous trip prices (x, z, p) forms a *competitive equilibrium* (CE) if:

(i) all riders $j \in \mathcal{R}$ that can afford the ride are picked up, i.e. $v_j > p_{o_j,d_j,\tau_j} \Rightarrow x_j = 1$, and all riders that are picked up can afford the price $x_j = 1 \Rightarrow v_j \geq p_{o_j,d_j,\tau_j}$,
(ii) $\forall i \in \mathcal{D}$, $z_i \in \arg\max_{z_i' \in \mathcal{Z}_i} \left\{ \sum_{(a,b,t) \in z_i'} p_{a,b,t} \right\}$, i.e. each driver takes one of her feasible paths with the highest total payment, and
(iii) $\forall (a, b, t) \in \mathcal{T}$, $\sum_{j \in \mathcal{R},\ (o_j,d_j,\tau_j)=(a,b,t)} x_j < \sum_{i \in \mathcal{D}} \mathbb{1}\{(a, b, t) \in z_i\} \Rightarrow p_{a,b,t} = 0$, meaning that any trip with excess supply has zero price.

The welfare-optimal planning problem can be formulated as an integer linear program (ILP). Let x_j be the indicator that rider $j \in \mathcal{R}$ is picked up, and $y_{i,k}$ be the indicator that driver i takes $Z_{i,k}$, her k^{th} feasible path in $\mathcal{Z}_i$. We have:

$$\max_{x,y} \sum_{j \in \mathcal{R}} x_j v_j \tag{1}$$

$$\text{s.t.} \sum_{j \in \mathcal{R},\ (o_j,d_j,\tau_j)=(a,b,t)} x_j \leq \sum_{i \in \mathcal{D}} \sum_{k=1}^{|Z_i|} y_{i,k} \mathbb{1}\{(a, b, t) \in Z_{i,k}\}, \quad \forall (a, b, t) \in \mathcal{T} \tag{2}$$

$$\sum_{k=1}^{|\mathcal{Z}_i|} y_{i,k} \leq 1,\ \forall i \in \mathcal{D} \tag{3}$$

$$x_j \in \{0, 1\},\ \forall j \in \mathcal{R},\ y_{i,k} \in \{0, 1\},\ \forall i \in \mathcal{D},\ k = 1, \ldots, |\mathcal{Z}_i|$$

(2) requires that for all $(a, b, t) \in \mathcal{T}$, the number of riders taking this trip is no more than the number of drivers whose paths cover this trip. (3) requires that each driver takes at most one path.

Relaxing the integrality constraints, we obtain the following linear program (LP) relaxation, which we refer to as the primal LP:

$$\max_{x,y} \sum_{j \in \mathcal{R}} x_j v_j \tag{4}$$

$$\text{s.t.} \sum_{j \in \mathcal{R},\ (o_j, d_j, \tau_j)=(a,b,t)} x_j \leq \sum_{i \in \mathcal{D}} \sum_{k=1}^{|Z_i|} y_{i,k} \cdot \mathbb{1}\{(a,b,t) \in Z_{i,k}\}, \quad \forall (a,b,t) \in \mathcal{T} \tag{5}$$

$$\sum_{k=1}^{|\mathcal{Z}_i|} y_{i,k} \leq 1,\ \forall i \in \mathcal{D} \tag{6}$$

$$x_j \leq 1,\ \forall j \in \mathcal{R} \tag{7}$$

$$x_j \geq 0,\ \forall j \in \mathcal{R},\ y_{i,k} \geq 0,\ \forall i \in \mathcal{D},\ k = 1, \ldots, |\mathcal{Z}_i|$$

Lemma 3.1 (Integrality). *There exists an integer optimal solution to the linear program* (4).

We leave the proof to the full version of the paper [16]. For every instance optimal planning problem, we can construct a *minimum cost flow (MCF) problem*, where drivers flow through a network with vertices corresponding to (location, time) pairs, edges corresponding to trips, and edge costs corresponding to rider values. We show a correspondence between optimal solutions of (4) and that of the MCF, and appeal to the integrality of the optimal solution of MCF [17]. This reduction can also be used to solve (4) efficiently.

Let $p_{a,b,t}$, π_i and u_j be the dual variables corresponding to constraints (5), (6) and (7), respectively. The dual LP of LP (4) is:

$$\min \sum_{i \in \mathcal{D}} \pi_i + \sum_{j \in \mathcal{R}} u_j \tag{8}$$

$$\text{s.t. } \pi_i \geq \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t}, \quad \forall k = 1, \ldots, |\mathcal{Z}_i|,\ \forall i \in \mathcal{D} \tag{9}$$

$$u_j \geq v_j - p_{o_j, d_j, \tau_j}, \quad \forall j \in \mathcal{R} \tag{10}$$

$$p_{a,b,t} \geq 0, \quad \forall (a,b,t) \in \mathcal{T} \tag{11}$$

$$\pi_i \geq 0, \quad \forall i \in \mathcal{D} \tag{12}$$

$$u_j \geq 0, \quad \forall j \in \mathcal{R} \tag{13}$$

Lemma 3.2 (Welfare Theorem). *A plan with anonymous trip prices (x, y, p) is welfare-optimal if and only if it forms a CE. Such optimal CE plans always exist, are efficient to compute, and are individually rational for riders, strictly budget balanced, and envy-free.*

See [16] for the proof of this lemma. Given an optimal primal solution and optimal dual solution, the dual variables π and u can be interpreted as the total payment to drivers and utility of riders, when the anonymous trip prices are given by p. We then make use of standard observations about complementary slackness conditions and their connection with CE [5, 18]. By integrality, optimal CE plans always exist, and can be efficiently computed by solving the primal LP and the dual LP of the corresponding MCF problem.

There is a lattice structure on drivers' total payments among all CE outcomes, and a connection between drivers' total payments among CE outcomes and the welfare differences from replicating/losing a driver. For any two sets of total payments to drivers, $\pi = (\pi_1, \pi_2, \ldots, \pi_m)$ and $\pi' = (\pi'_1, \pi'_2, \ldots, \pi'_m)$, let the join $\bar{\pi} = \pi \vee \pi'$ and the meet $\underline{\pi} = \pi \wedge \pi'$ be defined as, for all $i \in \mathcal{D}$,

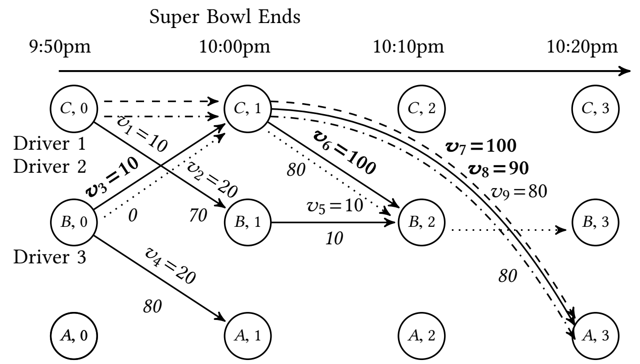

Figure 2: The Super Bowl example: the driver-pessimal plan, illustrating competitive equilibrium and thus smooth prices.

$\bar{\pi}_i = \max\{\pi_i, \pi'_i\}$ and $\underline{\pi}_i = \min\{\pi_i, \pi'_i\}$, respectively. For each driver $i \in \mathcal{D}$, denote the *social welfare gain from replicating driver i*, and the *social welfare loss from eliminating driver i*, as:

$$\Phi_{\ell_i, \underline{\tau}_i} \triangleq W(\mathcal{D} \cup \{(\underline{\tau}_i, \bar{\tau}_i, \ell_i)\}) - W(\mathcal{D}), \tag{14}$$

$$\Psi_{\ell_i, \underline{\tau}_i} \triangleq W(\mathcal{D}) - W(\mathcal{D} \backslash \{(\underline{\tau}_i, \bar{\tau}_i, \ell_i)\}), \tag{15}$$

where $W(\mathcal{D})$ is the highest social welfare achievable by a set of drivers $\mathcal{D}$. A *driver-optimal plan* has a payment profile at the top of the lattice, and a *driver-pessimal plan* has a payment profile at the bottom of the lattice.

Lemma 3.3 (Lattice Structure). *Drivers' total payments π among all CE outcomes form a lattice. Moreover, for each driver $i \in \mathcal{D}$, $\Phi_{\ell_i, \underline{\tau}_i}$ and $\Psi_{\ell_i, \underline{\tau}_i}$ are equal to the total payments to driver i in the driver-pessimal and driver-optimal CE plans, respectively.*

We leave the proof to [16]. For the lattice structure, we prove a correspondence between the optimal solutions to the dual LP (8) and the optimal solutions to the dual LP of the MCF problem, and establish a lattice structure of the MCF optimal dual solutions. For the driver optimal and pessimal payments, we show that the welfare gains and losses form optimal dual solutions to the MCF by standard arguments on shortest paths in the residual graph [1]. The fact that the optimal dual solutions must be a subgradient of the objective of the dual LP for MCF (as a function of the flow boundary conditions) implies that the welfare gains and losses correspond to the bottom and the top of the lattice of the dual variables.

Example 1 (Continued)**.** For the Super Bowl scenario, the myopic pricing mechanism with the lowest market-clearing prices sets $p_{B,C,0} = p_{B,A,0} = 0$ and $p_{C,B,1} = p_{C,A,1} = 100$. The prices for trips leaving C increase significantly at time 1, and the plan looking forward from time 0 does not form a CE, since the alternative path $((C, C, 0),\ (C, A, 1))$ has total payment of 100, higher than that of driver 1's path $((C, B, 0),\ (B, B, 1),\ (B, B, 2))$, which pays 0.

By contrast, the driver-pessimal CE plan is illustrated in Figure 2. The anonymous trip prices are shown in italics, below the edge corresponding to each trip. In this plan, all drivers stay in or reposition to location C at time 0, pick up riders with high values, and achieve the optimal welfare of 300. The outcome forms a CE, that there is no other path with a higher total payment for any driver.

All riders are happy with their dispatched trips given the prices, and there is no driver or rider envy. □

4 SPATIAL-TEMPORAL PRICING

We first formally define a dynamic mechanism, that can use the history of actions to update the plan forward from the current state.

Let $s_t = (s_{1,t}, s_{2,t}, \ldots, s_{m,t})$ be the *state* of the platform at time t, where each $s_{i,t}$ describes the state of driver $i \in \mathcal{D}$. If a driver is available at time t at location a, we denote $s_{i,t} = (a, t)$. Otherwise, if driver i is en route, finishing the trip from a to b that she started at time $t' < t$ s.t. $t' + \delta(a, b) > t$, we denote $s_{i,t} = (a, b, t')$ if the driver is relocating, or $s_{i,t} = (a, b, t', j)$ if the driver is taking a rider j from a to b. For drivers that have not yet entered, we write $s_{i,t} = (\underline{\tau}_i, \ell_i)$. The initial state of the system is determined by the types of drivers: $s_0 = ((\underline{\tau}_1, \ell_1), \ldots, (\underline{\tau}_m, \ell_m))$.

At each time t, each driver i takes an *action* $\alpha_{i,t}$. Assuming $s_{i,t} = (a, t)$, the action may be to relocate to location b, which we denote $\alpha_{i,t} = (a, b, t)$, or to pick up a rider j, in which case $\alpha_{i,t} = (a, d_j, t, j)$. For a driver i that is en route at time t, $\alpha_{i,t} = s_{i,t}$— the only available action is to finish the current trip. For driver i that has not yet entered, $\alpha_{i,t} = (\underline{\tau}_i, \ell_i)$. The action $\alpha_{i,t}$ taken by driver i at time t determines her state $s_{i,t+1}$:

- (will complete trips at $t+1$) if $\alpha_{i,t} = (a, b, t')$ or $\alpha_{i,t} = (a, b, t', j)$ s.t. $t' + \delta(a, b) = t + 1$, then $s_{i,t+1} = (b, t + 1)$, meaning these drivers will become available at time $t + 1$ at their destinations,
- (still en route) if $\alpha_{i,t} = (a, b, t')$ or $\alpha_{i,t} = (a, b, t', j)$ s.t. $t' + \delta(a, b) > t + 1$, then $s_{i,t+1} = \alpha_{i,t}$,
- (not yet entered) if $i \in \mathcal{D}$ s.t. $\alpha_{i,t} = (\underline{\tau}_i, \ell_i)$, then $s_{i,t+1} = (\underline{\tau}_i, \ell_i)$.

Let $\alpha_t = (\alpha_{1,t}, \alpha_{2,t}, \ldots, \alpha_{m,t})$ be the *action profile* of all drivers at time t, and let *history* $h_t = (s_0, \alpha_0, s_1, \alpha_1, \ldots, s_{t-1}, \alpha_{t-1}, s_t)$, with $h_0 = (s_0)$. Finally, let $\mathcal{D}_t(h_t) = \{i \in \mathcal{D} \mid s_{i,t} = (a, t) \text{ for some } a \in \mathcal{L}\}$ be the set of drivers available at time t.

Definition 4 (Dynamic Ridesharing Mechanism). A dynamic ridesharing mechanism is defined by its dispatch and payment rules (α^*, π^*, q^*). At each time t, given history h_t and rider information $\mathcal{R}$, the mechanism announces:

- for a subset of available drivers $\mathcal{D}_t(h_t)$, a dispatch action $\alpha^*_{i,t}(h_t)$ to either pick up a rider or relocate, and a payment $\pi^*_{i,t}(h_t)$ if the driver follows the dispatch and takes this action ($\pi^*_{i,t}(h_t) = 0$ for available drivers that are not dispatched).
- for each en route driver $i \in \mathcal{D} \backslash \mathcal{D}_t$, we have $\alpha^*_{i,t}(h_t) = s_{i,t}$ (keep driving) and $\pi^*_{i,t}(h_t) = 0$ (no more payment).
- for each rider who receives a dispatch at time t, the payment $q^*_j(h_t)$, in the event that she is picked up.

Each dispatched driver decides whether to follow the dispatch and take action $\alpha_{i,t} = \alpha^*_{i,t}(h_t)$, or to deviate and take an action $\alpha_{i,t} \in \mathcal{A}_{i,t}(h_t) \backslash \{\alpha^*_{i,t}(h_t)\}$. For a driver in location a, the set $\mathcal{A}_{i,t}(h_t)$ of available actions at time t is defined as $\mathcal{A}_{i,t}(h_t) = \{\alpha^*_{i,t}(h_t)\} \cup \{(a, b, t) \mid b \in \mathcal{L} \text{ s.t. } t + \delta(a, b) \le T\}$. For an available driver in location a who is not dispatched, the set of available actions in period t is $\mathcal{A}_{i,t}(h_t) = \{(a, b, t) \mid b \in \mathcal{L} \text{ s.t. } t + \delta(a, b) \le T\}$. An en route driver must continue the trip, and $\mathcal{A}_{i,t}(h_t) = \{\alpha^*_{i,t}(h_t)\}$.

After observing the action profile α_t at time t, the mechanism pays each dispatched driver $\hat{\pi}_{i,t}(\alpha_{i,t}, h_t) = \pi^*_{i,t}(h_t) \mathbb{1}\{\alpha_{i,t} = \alpha^*_{i,t}\}$ and charges $\hat{q}_j(\alpha_t) = q^*_j(h_t) \sum_{i \in \mathcal{D}_t} \mathbb{1}\{\alpha_{i,t} = (o_j, d_j, t, j)\}$ from each rider that requests a ride at time t.

A mechanism is *feasible* if $\forall t \in [T]$, $\forall h_t$, (i) it is possible for each available driver to take the trip dispatched to her, i.e. if $s_{i,t} = (a, t)$ for some $a \in \mathcal{L}$, $\alpha^*_{i,t}(h_t) \in \{(a, b, t) \mid b \in \mathcal{L},\ t + \delta(a, b) \le T\} \cup \left\{(o_j, d_j, \tau_j, j) \;\middle|\; j \in \mathcal{R},\ \tau_j = t,\ o_j = a\right\}$, and (ii) no rider is picked-up more than once: $\forall j \in \mathcal{R}_t$, $\sum_{i \in \mathcal{D}_t} \mathbb{1}\{\alpha^*_{i,t}(h_t) = (o_j, d_j, \tau_j, j)\} \le 1$.

Let $\mathcal{H}_t$ be the set of all possible *histories* up to time t. A *strategy* σ_i of driver i defines for all times $t \in [T-1]$ and all histories $h_t \in \mathcal{H}_t$, the action $\alpha_{i,t} = \sigma_i(h_t) \in \mathcal{A}_{i,t}(h_t)$. For a mechanism that always dispatches all available drivers, let σ^*_i denote the *straightforward strategy* of always following the mechanism's dispatches.

Let $\sigma = (\sigma_1, \ldots, \sigma_m)$ be the *strategy profile* of all drivers, with $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_m)$. The strategy profile σ, together with the initial state s_0, and the rules of a mechanism, determines all actions and payments of all drivers through the entire planning horizon. Let $\sigma_i|_{h_t}$, $\sigma|_{h_t}$ and $\sigma_{-i}|_{h_t}$ denote the strategy profile from time t onward given history h_t for driver i, all drivers, and all drivers but i, respectively.

For each driver $i \in \mathcal{D}$, let $\hat{\pi}_i(\sigma) = \sum_{t=0}^{T-1} \hat{\pi}_{i,t}(\sigma_i(h_t), h_t)$ denote the total actual payments made to driver i, where drivers follow σ and the history h_t is induced by the initial state and strategy σ. For the riders, let $\hat{x}_j(\sigma) \in \{0, 1\}$ be the indicator that rider j is picked-up at strategy profile σ. Let $\hat{q}_j(\sigma) = \hat{x}_j(\sigma) q^*_j(h_{\tau_j})$ be the payment made by rider j at time τ_j.

Fixing driver and rider types, a ridesharing mechanism induces an extensive form game. At each time point t, each driver decides on an action $\alpha_{i,t} = \sigma_i(h_t) \in \mathcal{A}_{i,t}(h_t)$ to take based on strategy σ_i and the history h_t, and receives payment $\hat{\pi}_{i,t}(\alpha_{i,t}, h_t)$. The total payment $\hat{\pi}_i(\sigma)$ to each driver is determined by the rules of the mechanism. We define the following properties.

Definition 5 (Budget Balance (BB)). A ridesharing mechanism is budget balanced if for any set of riders and drivers, and any strategy profile σ taken by the drivers, $\sum_{j \in \mathcal{R}} \hat{q}_j(\sigma) \ge \sum_{i \in \mathcal{D}} \hat{\pi}_i(\sigma)$.

Definition 6 (Individual Rationality (IR)). A ridesharing mechanism is individually rational for riders if for any set of riders and drivers and for any driver strategy σ, $\hat{x}_j(\sigma) v_j \ge \hat{q}_j(\sigma)$, $\forall i \in \mathcal{R}$.

Definition 7 (Subgame-Perfect Incentive Compatibility). A ridesharing mechanism that always dispatches all available drivers is *subgame-perfect incentive compatible* (SPIC) for drivers if given any set of riders and drivers, following the mechanism's dispatches at all times forms a subgame-perfect equilibrium (SPE) for the drivers, meaning for all times $t \in [T-1]$, for any history $h_t \in \mathcal{H}_t$,

$$\sum_{t'=t}^{T-1} \hat{\pi}_{i,t'}(\sigma^*_i|_{h_t}, \sigma^*_{-i}|_{h_t}) \ge \sum_{t'=t}^{T-1} \hat{\pi}_{i,t'}(\sigma_i|_{h_t}, \sigma^*_{-i}|_{h_t}),\ \forall \sigma_i,\ \forall i \in \mathcal{D}.$$

A ridesharing mechanism is *dominant strategy incentive compatible (DSIC)* if for any driver, following the mechanism's dispatches at all time points that the driver is dispatched maximizes her total payment, regardless of the actions taken by the rest of the drivers.

Definition 8 (Envy-freeness in SPE). A ridesharing mechanism that always dispatches all available drivers is *envy-free in SPE for drivers* if for any set of riders and drivers, (i) the mechanism is SPIC for drivers, and (ii) for any time $t \in [T-1]$, for all history $h_t \in \mathcal{H}_t$, all drivers with the same state at time t are paid the same total amount in the subsequent periods, assuming all drivers follow the mechanism's dispatches:

$$\sum_{t'=t}^{T-1} \hat{\pi}_{i,t'}(\sigma^*|_{h_t}) = \sum_{t'=t}^{T-1} \hat{\pi}_{i',t'}(\sigma^*|_{h_t}),\ \forall i, i' \in \mathcal{D} \text{ s.t. } s_{i,t} = s_{i',t}.$$

A ridesharing mechanism is *envy-free in SPE for riders* if (i) the mechanism is SPIC for drivers, and (ii) for all $j \in \mathcal{R}$, for all possible $h_{\tau_j} \in \mathcal{H}_{\tau_j}$, and all $j' \in \mathcal{R}$ s.t. $(o_j, d_j, \tau_j) = (o_{j'}, d_{j'}, \tau_{j'})$

$$\hat{x}_j(\sigma^*)v_j - \hat{q}_j(\sigma^*) \geq \hat{x}_{j'}(\sigma^*)v_j - \hat{q}_{j'}(\sigma^*).$$

4.1 The Spatio-Temporal Pricing Mechanism

In defining the STP mechanism, we conceptualize the dispatch and payment rule as defining a time 0 plan, and then defining an updated "time t plan" in the event of a deviation by any driver (from the current plan) at time $t-1$. In each case, the plans are induced by the dispatch rule assuming that drivers followed the proposed dispatch actions. The dispatch and payment rules of a mechanism can also be considered to be determined by the way in which a "planning rule" is used to compute and update the plan. We adopt this viewpoint in defining the STP mechanism.

For any time $t \in [T]$, given any state s_t of the platform, let $E^{(t)}(s_t)$ represent the *time-shifted economy* starting at state s_t, with planning horizon $T^{(t)} = T - t$, the same set of locations $\mathcal{L}$ and distances δ, the remaining riders $\mathcal{R}^{(t)} = \{(o_j, d_j, \tau_j - t, v_j) \mid j \in \mathcal{R}, \tau_j \geq t\}$ and a set of drivers $\mathcal{D}^{(t)}(s_t) = \{(\underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) \mid i \in \mathcal{D}\}$ with types determined by s_t as follows:

(i) for available drivers $i \in \mathcal{D}$ s.t. $s_{i,t} = (a, t)$ for some $a \in \mathcal{L}$, $(\underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) = (0, \bar{\tau}_i - t, a)$,
(ii) for en route drivers $i \in \mathcal{D}$ s.t. $s_{i,t} = (a, b, t')$ or (a, b, t', j), $(\underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) = (t' + \delta(a,b) - t, \bar{\tau}_i - t, b)$, and
(iii) for driver $i \in \mathcal{D}$ that had not entered, $(\underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) = (\underline{\tau}_i - t, \bar{\tau}_i - t, \ell_i)$.

For each location $a \in \mathcal{L}$ and time $t \in [T]$, define the welfare gain in the economy from adding a driver at (a, t) as,

$$\Phi_{a,t} \triangleq W(\mathcal{D} \cup \{(t, T, a)\}) - W(\mathcal{D}), \tag{16}$$

where (t, T, a) represents the type of this additional driver. We now define the STP mechanism.

Definition 9 (Spatio-Temporal Pricing)**.** The *spatio-temporal pricing (STP) mechanism* is a dynamic ridesharing mechanism that always dispatches all available drivers. Given $E^{(0)}$ at the beginning of the planning horizon, or $E^{(t)}(s_t)$ immediately after some driver's deviation, the mechanism computes a plan as follows:

- To determine the dispatches (α^*), compute any optimal solution (x, y) to (1), and dispatch each driver $i \in \mathcal{D}$ to take the path $z_i = Z_{i,k}$ for k s.t. $y_{i,k} = 1$ and pick up riders $j \in \mathcal{R}$ s.t. $x_j = 1$,
- To determine driver and rider payments (π^* and q^*), for each trip $(a, b, t) \in \mathcal{T}$, set anonymous trip prices $p_{a,b,t} = \Phi_{a,t} - \Phi_{b,t+\delta(a,b)}$ where $\Phi_{a,t}$ is the welfare gain as defined in (16):
  - for each rider $j \in \mathcal{R}$, $q_j^* = p_{o_j, d_j, \tau_j} \sum_{i \in \mathcal{D}} \mathbb{1}\{\alpha^*_{i,\tau_j} = (o_j, d_j, \tau_j, j)\}$,
  - for each driver $i \in \mathcal{D}$, $\pi^*_{i,t} = \sum_{a,b \in \mathcal{L}} p_{a,b,t} \mathbb{1}\{(a,b,t) \in z_i\}$.

The STP mechanism uses a history independent and time-invariant planning rule: at any point of time, when computing a new forward-looking plan, this is computed as if the time an updated plan was computed was the beginning of the planning horizon. This is the sense in which the STP mechanism does not make use of time-extended contracts, including penalties for previous actions.[10] We now state the main result of the paper.

THEOREM 4.1. *The spatio-temporal pricing mechanism is subgame-perfect incentive compatible. It is also individually rational and strictly budget balanced for any action profile taken by the drivers, and is welfare optimal and envy-free in subgame-perfect equilibrium.*

We leave the proof to [16]. Observe that $\Phi_{a,T} = 0$ for all $a \in \mathcal{L}$ since a driver that enters at time T cannot pick up any rider thus does not improve welfare. Any feasible path of each driver $i \in \mathcal{D}$ over the planning horizon starts at $(\ell_i, \underline{\tau}_i)$ and ends at (a, T) for some $a \in \mathcal{L}$. By telescoping sum, the total payment to driver i under the planning rule of the STP mechanism is

$$\pi_i = \sum_{(a,b,t) \in z_i} p_{a,b,t} = \sum_{(a,b,t) \in z_i} \Phi_{a,t} - \Phi_{b,t+\delta(a,b)} = \Phi_{\ell_i, \underline{\tau}_i},$$

which is the welfare gain of the system from replicating driver i.

Lemma 3.3 implies that the STP mechanism can therefore be interpreted as a dynamic mechanism that always announces a driver-pessimal CE plan at the beginning of the planning horizon and also after driver deviation. We may also consider the following driver-optimal mechanism, which always computes and updates the driver-optimal CE plan. A natural variation on the STP mechanism is to consider the driver-optimal analogue, which always computes a driver-optimal competitive equilibrium plan at the beginning of the planning horizon, or upon deviation of any driver. This mechanism pays each driver the externality she brings to the economy, and corresponds to the reasoning of the VCG mechanism. The driver-optimal mechanism is, however, not incentive compatible. See Ma et al. [16] for a detailed example and explanation.

Under the STP mechanism, recomputation of the plan can be triggered by the deviation of any driver, and for this reason the total payment of a driver is affected by the actions of others and the straightforward behavior is not a dominant strategy. We show that no mechanism can implement the desired properties in a dominant-strategy equilibrium.

THEOREM 4.2. *Following the mechanism's dispatch at all times does not form a dominant strategy equilibrium under any dynamic ridesharing mechanism that is, from any history onward, (i) welfare-optimal (ii) individually rational, (iii) budget balanced and (iv) envy-free for riders and drivers.*

5 SIMULATION RESULTS

In this section we compare through simulation, the performance of the STP mechanism against the myopic pricing mechanism with the lowest market clearing prices (which market clearing prices are

[10] A mechanism that is not history-independent or time-invariant can trivially align incentives. For example, a mechanism that is not history-independent could "fire" any driver that had deviated in the past according to the history, while keeping the plans for the rest of the economy unchanged. A mechanism that is history-independent but not time-invariant can threaten to "shut down" and not make any further dispatches or payments to the drivers if any of them had deviated.

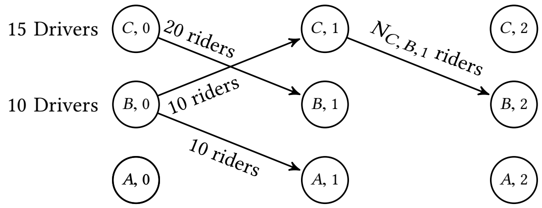

Figure 3: The end of an event.

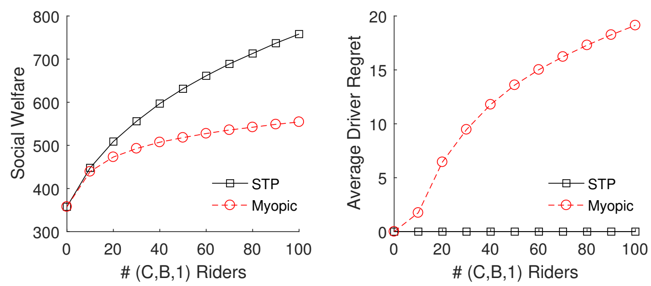

(a) Social welfare. **(b) Average driver regret.**

Figure 4: Comparison of welfare and average driver regret.

chosen is unimportant for the results) for the economy as illustrated in Figure 3, modeling the end of an event. The analysis suggests that the STP mechanism achieves substantially higher social welfare, as well as time-efficiency for drivers, whereas, under the myopic pricing mechanism, drivers incur a high regret.

There are three locations $\mathcal{L} = \{A, B, C\}$ with distance $\delta(a, b) = 1$ for all $a, b \in \mathcal{L}$, and two time periods. The event ends at location C at time 1, where there would be a large number of riders requesting rides. In each economy, there are 15 and 10 drivers entering at locations C and B at time 0. 20 riders request the trip $(C, B, 0)$, and 10 riders request the trips $(B, C, 0)$ and $(B, A, 0)$ respectively. When the event ends, there are $N_{C,B,1}$ riders hoping to take a ride to $(B, 2)$. For each economy, the values of all riders drawn i.i.e. from an exponential distribution with mean 10.

As $N_{C,B,1}$ varies from 0 to 100, we randomly generate $1,000$ economies, and compare the average welfare under the two mechanisms, as shown in Figure 4a. The STP mechanism achieves a significantly higher social welfare than the myopic pricing mechanism, especially when there there are a large number of drivers taking the trip $(C, B, 1)$.

The STP mechanism is incentive compatible, however, the myopic pricing mechanism is not. Drivers that are dispatched to the trips $(C, B, 1)$ and $(B, A, 1)$ may regret having not relocated to C instead and get paid a large amount at time 1. Define the *regret* of a driver as the highest additional amount a driver can gain by strategizing, in comparison to following a mechanism's dispatch, assuming the rest of the drivers all follow the mechanism's dispatch. Figure 4b shows that the average regret of the 25 drivers increase significantly as more riders request the trip $(C, B, 1)$.

The number of drivers completing each of the four trips of interest under the two mechanism are as shown in Figure 5. As $N_{C,B,1}$ increases, the STP mechanism dispatches more drivers to $(C, 1)$ to pick-up the higher-valued riders leaving C, while sending less drivers away to pick up riders for trips $(C, B, 0)$ and $(B, A, 0)$. The

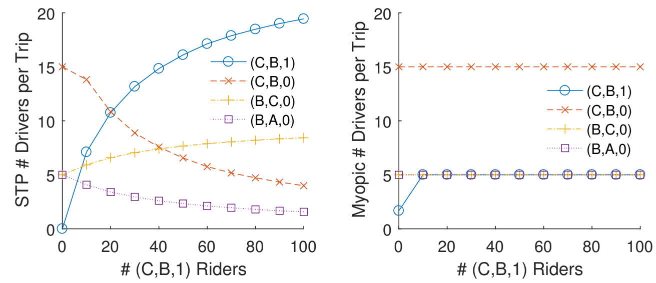

(a) The STP mechanism. **(b) The myopic mechanism.**

Figure 5: Comparison of the number of drivers per trip.

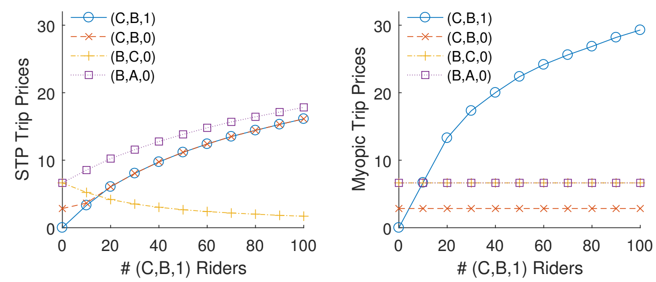

(a) The STP mechanism. **(b) The myopic mechanism.**

Figure 6: Comparison of trip prices.

myopic pricing mechanism, being oblivious to future demand, sends all drivers starting at $(C, 0)$ to location B, and an average of only 5 drivers to $(C, 1)$ from $(B, 1)$.

The average prices for the four trips under the two mechanisms are as shown in Figure 6. First of all, prices are temporally "smooth" under STP— trips leaving C at times 0 and 1 have very similar prices. Moreover, the prices for trips $(B, C, 0)$ and $(C, B, 1)$ add up to the price of the trip $(B, A, 0)$, so that drivers starting at $(B, 1)$ would not envy each other. In contrast, prices for trips leaving C increase significantly under the myopic pricing mechanism, and the drivers that are dispatched the trip $(B, A, 0)$ envy those that are dispatched $(B, C, 0)$ and subsequently $(C, B, 1)$. The "surge" for the trip $(C, B, 1)$ is significantly higher under the myopic pricing mechanism than under the STP mechanism, implying that the platform is providing less price reliability for the riders.

6 CONCLUSION

We study the problem of optimal dispatching and pricing in two-sided ridesharing platforms in a way that drivers would choose to accept the platform's dispatches instead of driving to another area or waiting for a higher price. Under a complete information, discrete time, multi-period and multi-location model, we show that always following the mechanism's dispatches forms a subgame-perfect equilibrium among the drivers under the spatio-temporal pricing mechanism, which always computes a driver-pessimal competitive equilibrium plan at the beginning of the planning horizon as well as after any deviations. Our empirical study suggests that the STP mechanism achieves substantially higher social welfare and drivers' time efficiency in comparison to the myopic pricing mechanism, where in addition drivers incur a high regret.

REFERENCES

[1] Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. 1993. Network flows: theory, algorithms, and applications. (1993).

[2] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. 2016. Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework. *CoRR abs/1608.06819* (2016).

[3] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. 2015. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 639–639.

[4] Dirk Bergemann and Juuso Välimäki. 2010. The dynamic pivot mechanism. *Econometrica* 78, 2 (2010), 771–789.

[5] Dimitri P Bertsekas. 1990. The auction algorithm for assignment and other network flow problems: A tutorial. *Interfaces* 20, 4 (1990), 133–149.

[6] Kostas Bimpikis, Ozan Candogan, and Saban Daniela. 2016. Spatial Pricing in Ride-Sharing Networks. *SSRN Electronic Journal* (2016), 1–51. https://doi.org/10.2139/ssrn.2868080

[7] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. 2017. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 241–242.

[8] Ruggiero Cavallo, David C Parkes, and Satinder Singh. 2009. Efficient mechanisms with dynamic populations and dynamic types. (2009).

[9] M Keith Chen, Judith A Chevalier, Peter E Rossi, and Emily Oehlsen. 2017. *The Value of Flexible Work: Evidence from Uber Drivers*. Technical Report. National Bureau of Economic Research.

[10] M Keith Chen and Michael Sheldon. 2015. Dynamic Pricing in a Labor Market: Surge Pricing and the Supply of Uber Driver-Partners. *University of California (Los Angeles) Working Paper URL http://citeseerx. ist. psu. edu/viewdoc/download* (2015).

[11] Scott H Clearwater. 1996. *Market-based control: A paradigm for distributed resource allocation*. World Scientific.

[12] Peter Cohen, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe. 2016. *Using big data to estimate consumer surplus: The case of uber*. Technical Report. National Bureau of Economic Research.

[13] Jonathan Hall, Cory Kendrick, and Chris Nosko. 2015. The effects of UberâĂŹs surge pricing: A case study. *The University of Chicago Booth School of Business* (2015).

[14] Jonathan V Hall, John J Horton, and Daniel T Knoepfle. 2017. Labor Market Equilibration: Evidence from Uber. (2017).

[15] Jonathan V Hall and Alan B Krueger. 2016. *An analysis of the labor market for UberâĂŹs driver-partners in the United States*. Technical Report. National Bureau of Economic Research.

[16] Hongyao Ma, Fei Fang, and David C. Parkes. 2018. Spatio-Temporal Pricing for Ridesharing Platforms. *arXiv preprint arXiv:1801.04015* (2018).

[17] Kazuo Murota. 2003. *Discrete convex analysis*. SIAM.

[18] David C Parkes and Lyle H Ungar. 2000. Iterative combinatorial auctions: Theory and practice. *AAAI/IAAI* 7481 (2000).

[19] Lisa Rayle, Susan Shaheen, Nelson Chan, Danielle Dai, and Robert Cervero. 2014. *App-based, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in san francisco university of california transportation center (uctc)*. Technical Report. UCTC-FR-2014-08.

[20] Michael P. Wellman. 1993. A market-oriented programming environment and its application to distributed multicommodity flow problems. *JAIR* 1 (1993), 1–23.