

Spatio-Temporal Pricing for Ridesharing Platforms*

Hongyao Ma[†] Fei Fang[‡] David C. Parkes[†]

October 4, 2018

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, so 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 work in a complete information, discrete time, multi-period, multi-location model, and 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. The mechanism is also welfare-optimal, envy-free, individually rational, budget balanced and core-selecting from any history onward. The proof of incentive alignment makes use of the M^{\natural} concavity of min-cost flow objectives. We also give an impossibility result, that there can be no dominant-strategy mechanism with the same economic properties. An empirical analysis conducted in simulation suggests that the STP mechanism can achieve significantly higher social welfare than a myopic pricing mechanism.

1 Introduction

Ridesharing platforms such as Uber and Lyft are rapidly disrupting traditional forms of transit, completing millions of trips each day. These platforms connect drivers and riders, with both sides in a customer relationship with the platforms. When a rider opens the app and enters an origin and destination, these platforms quote a price for the trip and an estimated wait time. If a rider requests the ride, the platform dispatches a nearby driver to pick up the rider, and if neither side cancels and the trip is completed, payment is made from the rider to the driver through the platform.¹

Comparing with traditional taxi systems, ridesharing platforms significantly increase the fraction of time when drivers have a rider in the back seat, reducing trip costs for riders [20]. The

*Thanks to Yakov Babichenko, John Byers, Peter Frazier, Keith Chen, Amos Fiat, Myrto Kalouptsidi, Scott Kominers, Robin Lee, Shengwu Li, Reshef Meir, Hamid Nazerzadeh, Michael Ostrovsky, Ariel Pakes, Garrett van Ryzin, Peng Shi, Glen Weyl, and the participants at MatchUp 2017, ACM EC'17, IFORS 2017, the Harvard IO seminar, USC CAIS seminar, Technion IE seminar, CAEC17, TSMO at WSDM'18, AI3 at AAMAS'18, and the Uber MODS workshop for helpful feedback. The authors also thank Kevin Leyton-Brown for discussions on core equivalence. Ma is supported by a Siebel scholarship. Fang was partially supported by Harvard Center for Research on Computation and Society fellowship.

[†]John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: {hma@seas,parkes@eecs}.harvard.edu.

[‡]School of Computer Science, Carnegie Mellon University, Hamerschlag Dr, 4126 Wean Hall, Pittsburgh, PA 15213, USA. Email: feifang@cmu.edu.

¹The actual practice is somewhat more complicated, in that platforms may operate multiple products, for example, high-end cars, sports utility vehicles, 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, safety of neighborhood). We ignore these effects in our model, and assume that, conditioned on the same trip, all drivers are equivalent from the perspective of riders and all riders equivalent from the perspective of drivers.

platforms also emphasize the importance of providing reliable transportation. For example, Uber’s mission is “to connect riders to reliable transportation, everywhere for everyone.”² This mission is also stated as “to make transportation as ubiquitous and reliable as running water.”³ Lyft’s mission is “to provide the best, most reliable service possible by making sure drivers are on the road when and where you need them most.”⁴ Whereas taxi systems have reliable pricing but often unreliable service, these platforms use dynamic “surge” pricing to guarantee rider wait times do not exceed a few minutes [33]. Moreover, the platforms provide the flexibility for drivers to drive on their own schedule. Uber advertises itself as “work that put you first— drive when you want, earn what you need,”⁵ and Lyft promises drivers “To drive or not to drive? It’s really up to you.”⁶

Despite their success, there remain a number of problems with the pricing and dispatching rules governing the ridesharing platforms, leading in turn to various kinds of market failure. A particular concern, is that trips may be mis-priced relative to each other, incentivizing drivers to strategize [19, 14].⁷ For example, many platforms hide trip destinations from drivers before the pick-up. However, experienced drivers will call riders to ask about trip details, and cancel those that are not worthwhile [19]. Drivers also strategize in the following scenarios, where there is spatial imbalance and temporal variation of rider trip flows:

- (spatial mis-pricing) when 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 spatial mis-pricing leads to drivers’ “chasing the surge”— turning off a ridesharing app while relocating to another location.⁸
- (temporal mis-pricing) when large events such a sports game will soon end, drivers can anticipate that prices will increase in order to balance supply and demand. In this case, drivers will decline trips and even go off-line in order to wait in place.
- (network externalities) the origin-based dynamic “surge pricing” used in standard practice does not correctly factor market conditions at the destination of a trip. As a result, drivers decline trips to destinations where the continuation payoffs are low, e.g. quiet suburbs with low prices and long wait times.⁹

These kinds of mis-pricing undercut the mission of reliable transport, with even high willingness-to-pay riders unable to get access to reliable service for certain trips, such as trips leaving the stadium before a game ends, and trips going to a quiet suburb.¹⁰ This can also lead to inequity, with demonstrated learning effects leading to differences in long-run earnings (e.g. a gender gap in

²<https://www.uber.com/legal/community-guidelines/us-en/>, visited September 1, 2017.

³<http://time.com/time-person-of-the-year-2015-runner-up-travis-kalanick/>, visited September 1, 2017.

⁴<https://help.lyft.com/hc/en-us/articles/115012926227>, visited September 1, 2017.

⁵<https://www.uber.com/drive/>, visited December 12, 2017.

⁶<https://blog.lyft.com/posts/reasons-to-drive-with-lyft>, visited December 12, 2017.

⁷There are also other incentive problems, including inconsistencies across classes of service, competition among platforms, drivers’ bonuses and off-platform incentives. In the interest of simplicity, we only model a single class of service and ignore cross-platform competition.

⁸<http://maximumridesharingprofits.com/advice-new-uber-drivers-dont-chase-surge/>, visited Jan 8, 2018; <https://www.youtube.com/watch?v=qI094xqMiPA>, visited Jan 8, 2018.

⁹<https://therideshareguy.com/uber-drops-destination-filters-back-to-2-trips-per-day/>, visited September 1, 2018; <https://www.marketwatch.com/story/new-york-citys-new-uber-rules-could-make-those-5-cancellation-fees-go-away-2018-08-16>, visited September 1, 2018.

¹⁰A related consideration is that these opportunities for strategic behavior make it very difficult for platforms to introduce additional mechanisms to accommodate idiosyncratic driver preferences, for example over location. Rather, these mechanisms are used as tools for strategic behavior by drivers. <https://www.uber.com/blog/180-days-of-change-more-flexibility-and-choice/>, visited September 1, 2018; <https://therideshareguy.com/uber-drops-destination-filters-back-to-2-trips-per-day/>, visited September 1, 2018.

driver hourly earnings [19]), with potential consequences around driver churn from the platform. The imposition of cancellation penalties on drivers does not solve these problems, since drivers can always go offline, or choose not to participate in the platform from certain locations or times.⁹

We conceptualize many of the problems with current platforms as arising from prices failing to be appropriately “smooth” in space and time— if prices for trips are higher in one location then they should be appropriately higher in adjacent locations; if demand would soon increase in a location then the current prices should already be appropriately higher; and if destinations differ in continuation payoffs then trip prices to these destinations need to reflect this. With appropriately smooth prices, drivers who retain the flexibility to choose when to work will choose to accept any trip to which they are dispatched (recall that drivers, too, are customers of the platform). Dispatching decisions are also suboptimal on current platforms; sending drivers away from a sports stadium five minutes before a game ends at low prices is inefficient, for example, and driver behavior may actually be moving a system back towards efficiency. Correctly designed, ridesharing platforms can succeed in optimally orchestrate trips in a centralized manner, without having the power to tell drivers what to do.

1.1 Our results

We propose the *Spatio-Temporal Pricing* (STP) mechanism for dispatching and pricing in the context of a ridesharing platform, addressing the problem of incentive alignment for drivers in the presence of spatial imbalance and temporal variation of rider trip flows. The STP mechanism has the following properties:

- welfare-optimality: maximizing total rider values minus driver costs
- incentive-alignment: drivers will always choose to accept any dispatched trips
- envy-freeness: drivers at the same location and time do not envy each other’s future payoff; riders requesting the same trips do not envy each other’s outcomes
- core-selecting: no coalition of riders and drivers can make a better plan among themselves
- robustness: the mechanism updates the plans after deviations from the original dispatches
- temporal-consistency: plans are computed and updated based on the current state but not past history, without using penalties or time-extended contracts

Welfare optimality and incentive alignment are standard desiderata. For the others, we consider envy-freeness and core-selecting to be of special importance for sharing economy systems such as ridesharing platforms. These properties both relate to fairness. An envy-free mechanism is fair, removing the variation in drivers’ income that depends on lucky dispatches or from learning by doing. A core outcome is fair, in that no group of riders and drivers could do better by breaking away and operating their own ridesharing economy. Even given incentive alignment, we consider robustness to be important in the face of unmodeled effects, erroneous predictions, or mistakes by participants. By being robust, the mechanism is able to ensure the other properties from any history onward. Without this, any solution would be necessarily brittle and poorly suited to practice. Finally, temporal-consistency is important, since using penalties and threatening to fire drivers or to shut down the system are incompatible with the spirit of the gig economy, and the the real-time flexibility of being able to choose when to work.

We work in a complete information, discrete time, multi-period and multi-location model. Thus, the challenge addressed in this paper is one of promoting desirable behavior in the absence of time-extended contracts, and not one of information asymmetry. At the beginning of each time period,

based on the history, current positioning of drivers, and current and future demand, the STP mechanism dispatches each available driver to a rider trip, or to relocate, or to exit the platform for the planning horizon. The mechanism also determines a payment to be made if the driver follows the dispatch. Each driver then decides whether to follow the suggested action, or to decline and stay or relocate to any location, or to exit the platform. After observing the driver actions in a period, the mechanism collects payments from the riders and makes payments to the drivers.

The main assumptions that we make are:

- (i) complete information about supply and demand over a finite planning horizon,
- (ii) impatient riders, who 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 with no intrinsic preference for locations or passengers, and no heterogeneity in costs.

We allow for heterogeneity in rider values and trip details (origin, destination, and time). For drivers, they can become available at different times and locations, and can also differ in whether they are already driving in the platform (for example, finishing a trip), or have not yet joined and thus need to make an entry decision (for example, the driver will be dropping off a child at school at a specific location and time, and is willing to drive afterwards). We also allow a driver whom is asked to exist the platform earlier than their intended exit time to incur a cost, modeling the forgone opportunity of outside options after driving for some time.

A rider’s value is modeled as her willingness-to-pay, over and above a base payment that covers a drivers’ extra costs for picking up a rider (extra wear and tear, inconvenience, etc), in comparison to the cost of driving the same trip as relocation without a rider (fuel, time, and so forth). This base payment is transferred directly to the driver, so that all drivers completing the same trip with or without a rider incur the same trip cost. The prices determined by the STP mechanism are the payments collected on top of the base payments.

We first prove the existence of anonymous, origin-destination, competitive equilibrium (CE) prices, allowing the unit price of a trip to depend on market conditions at both the origin and destination.¹¹ The STP mechanism uses *driver-pessimal CE prices*, recomputing a driver-pessimal CE plan in the event of any deviations from the current plan. The mechanism 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 main result is that the STP mechanism satisfies all the desiderata outlined above. Somewhat surprisingly, the use of driver-pessimal CE prices (vs., for example, driver-optimal CE prices as in Vickrey-Clarke-Groves mechanisms) is essential for achieving our main result, that accepting the mechanism’s dispatches at all times forms a subgame-perfect equilibrium among the drivers. The proof of incentive alignment makes use of the M^{\natural} concavity of min-cost flow objectives [27]. The same connection to min-cost flow leads to an efficient algorithm to compute an optimal dispatch plan and prices, and to operationalize the STP mechanism.

The rest of this paper is organized as follows. After a brief discussion on related work, we introduce the model in Section 2, and illustrate through an example that a myopic pricing mechanism, which naively clears the market for each location without considering future demand and supply,

¹¹The idea that the unit price for a trip might depend on the destination is already familiar in taxi systems, with trips to the suburbs billed with a surcharge to reflect that drivers may need to return to the city without a fare. Origin-destination pricing also seems practical. Indeed, ridesharing platforms are moving in this direction, and this is facilitated by the movement to quoting a total payment for a trip rather than an origin-based surge multiplier. <https://newsroom.uber.com/upfront-fares-no-math-and-no-surprises/>, visited September 1, 2017; <https://blog.lyft.com/posts/now-live-see-your-ride-total>, visited September 1, 2017; <https://www.bloomberg.com/news/articles/2017-05-19/uber-s-future-may-rely-on-predicting-how-much-you-re-willing-to-pay>, visited September 1, 2017.

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 2). We prove a welfare theorem, that a plan is welfare-optimal if and only if it is supported by some anonymous trip prices in competitive equilibrium (CE) (Lemma 3). Optimal CE plans always exist and are efficient to compute. We also prove that drivers’ total utilities among all CE plans form a lattice (Lemma 4), and that the CE outcomes and the core are utility-equivalent (Lemma 5). A class of static CE mechanisms are discussed, that implement an optimal CE plan in dominant strategy (Theorem 1), but may fail to be welfare-optimal or envy-free after driver deviations. Section 4 introduces the STP mechanism, and establishes that the STP mechanism is subgame-perfect incentive compatible, and is also welfare-optimal, envy-free, core-selecting, individually rational and budget balanced from any history onward (Theorem 2).¹² We also provide an impossibility result, that no dominant-strategy mechanism has the same economic properties (Theorem 3), followed by discussions on the effect of relaxing the model assumptions.

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 consider three stylized scenarios: the end of a sports event, the morning rush hour, and trips to and from the airport with unbalanced flows. We conclude in Section 6. Omitted proofs, additional examples, discussions, relations to the literature and simulation results are provided in the Appendix.

1.2 Related Work

To the best of our knowledge, this current paper is unique in that it considers both multiple locations and multiple time periods, along 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 space and time.

Earlier, Banerjee et al. [6] adopt a queuing-theoretic approach in analyzing the effect of dynamic pricing on the revenue and throughput of ridesharing platforms, assuming a single location and stationary system state. In this context, the optimal dynamic pricing strategy, where prices can depend on supply and demand conditions, does not achieve better performance than the optimal static pricing strategy when the platform correctly estimates supply and demand. However, dynamic pricing is more robust to fluctuations and to mis-estimation of system parameters. By analyzing a two-location, stationary state queueing model, Afèche et al. [1] study the impact of platform control on platform revenue and driver income.

By analyzing the equilibrium outcome under a continuum model (supply and demand), and with stationary demand and unlimited driver supply (at fixed opportunity costs), Bimpikis et al. [10] in independent, contemporaneous work study the steady-state and show that a ridesharing platform’s profit is maximized when the demand pattern across different locations is balanced. In simulation they show that, in comparison to setting a single price, pricing trips differently depending on trip origins improves the platform’s profit. The same simulation also shows that there is not a substantial, additional gain from using origin-destination based pricing in their model. Our model is quite distinct, in that it is not a continuum model, does not have unlimited driver supply, and is not stationary. These authors explain that the solution to optimal origin-destination prices in their model only has dimensionality linear in the number of locations, but do not offer a complete theory to explain their simulation results. Banerjee et al. [5], also independent and contemporaneous,

¹²The STP mechanism balances budget. Alternatively, we may think about the ridesharing platform taking a fixed fraction of the driver surplus. This does not affect the results presented in this paper, and the equilibrium outcome under STP still resides in some ϵ -core of the market, depending on the cut taken by the platform.

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.

Castillo et al. [12] study the impact of myopically dispatching the closest drivers to rider requests. In particular, they discuss a “wild goose chase” phenomena. When demand much exceeds supply, myopic dispatching means that drivers spend too much time driving to pick up riders instead of carrying riders on trips, which leads to decreased social welfare and revenue to the platform. The theoretical model assumes a stationary state, where driver supply is driven by hourly earnings, rider demand depends on trip prices and wait times, and the wait time increases as the number of idle cars decreases. They do not have an explicit model of location. Empirical evidence is provided by analyzing Uber data. The wild goose chase is an effect of the ridesharing platform always dispatching a driver as soon as any rider requests a ride. The authors establish the importance of dynamic pricing when using this kind of myopic dispatching scheme, in keeping enough cars idle to avoid the inefficiency of long pick-ups, and show this is superior to other possible solutions, for example limiting the pick-up radius.

There are various empirical studies of the Uber platform as a two-sided marketplace, analyzing Uber’s driver partners [23], the labor market equilibrium [22], consumer surplus [18] and the commission vs. medallion lease based compensation models [3]. By analyzing the hourly earnings of drivers on the Uber platform, Chen et al. [15] 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; Cook et al. [19] show that driving speed, preferred time and location to drive, driver experience and their ability to strategically canceling rides together contribute to a 7% gender gap in hourly earnings. In regard to dynamic pricing, Chen and Sheldon [16] 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 [21] 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.

The dynamic variations on the VCG mechanism [4, 7, 13] truthfully implement efficient decision policies, where agents receive private information over time. These mechanisms are not suitable for our problem, however, because some drivers may be paid negative payments for certain periods of time. The payment to an agent in a single period in the dynamic VCG mechanism is equal to the flow marginal externality imposed on the other agents by its presence in the current period [13]. The problem in our setting is that the existence of a driver for only one period may exert negative externality on the rest of the economy by inducing suboptimal positioning of the rest of the drivers in the subsequent time periods. See Appendix D.1 for examples and discussions.

The literature on *trading networks* studies economic models where agents in a network can trade via bilateral contracts [24, 25, 31]. Efficient, competitive equilibrium outcomes exist when agents’ valuation functions satisfy the “full substitution” property, and the utilities of agents on either end of an acyclic network form lattices. Under proper assumptions, the optimal dispatching problem in ridesharing can be reduced to a trading network problem, where drivers and riders trade the right to use a car for the rest of the planning horizon (see Appendix D.2). However, the underlying network in the ridesharing problem evolves as time progresses, and to our knowledge, the dynamics and incentives in trading networks with time-varying networks are not studied in the literature.

Principal-agent problems have been studied extensively in contract theory [11, 34], where problems with information asymmetry before the time of contracting are referred to as *adverse selection*, and problems where asymmetric information or hidden action arise after the time of contracting are referred to as *moral hazard*. Contracts specify how agents are going to be rewarded or penal-

ized based on observed performance measures. In the setting where contracts cannot be perfectly enforced, *relational incentive contracts* [26] have also been studied, which are self-enforcing by threatening to terminate an agent following poor performance. In our model, there are neither hidden actions nor asymmetric information. Instead, the challenge we address is one of incentive alignment in the absence of contracts. We insist on retaining flexibility for drivers, in the spirit of the customer relationship with platforms and the Gig economy, with drivers free to decide on their own actions without incurring penalties or termination threats.

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.¹³ Denote $[T] = \{0, 1, \dots, T\}$ and $[T - 1] = \{0, 1, \dots, T - 1\}$.

Let $\mathcal{L} = \{A, B, \dots\}$ 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 (in which case $a = b$). Let the distance $\delta : \mathcal{L} \times \mathcal{L} \rightarrow \mathbb{N}$ be the time to travel between locations, so that trip (a, b, t) ends at $t + \delta(a, b)$.¹⁴ 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 $a, b \in \mathcal{L}$. $\mathcal{T} \triangleq \{(a, b, t) \mid a \in \mathcal{L}, b \in \mathcal{L}, t \in \{0, 1, \dots, T - \delta(a, b)\}\}$ denotes the set of all feasible trips within the planning horizon.

Let \mathcal{D} denote the set of drivers, with $m \triangleq |\mathcal{D}|$. Each driver $i \in \mathcal{D}$ is characterized by *type* $\theta_i = (\beta_i, \ell_i, \underline{\tau}_i, \bar{\tau}_i)$ — driver i is able to enter the platform at location ℓ_i and time $\underline{\tau}_i$, and plans to exit the platform at time $\bar{\tau}_i$ (with $\underline{\tau}_i < \bar{\tau}_i$). β_i indicates driver i ’s entrance status. A driver with $\beta_i = 0$ has not yet entered the platform, and needs to make an entry decision. A driver with $\beta_i = 1$ has already entered the platform, and will become available to pick up again at $(\ell_i, \underline{\tau}_i)$ (she may be completing an earlier trip, or relocating to another location). 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 location, including where they finish their last trip in the planning horizon*.

Any driver who completes a trip $(a, b, t) \in \mathcal{T}$ without a rider incurs a trip cost $c_{a,b,t} \geq 0$, which models the cost of time, driving, fuel, etc. A driver who has already entered the platform may exit earlier than her intended exit time, in which case she will not be able to complete any trip in the remaining periods of the planning horizon. Exiting Δ periods earlier than time T incurs a cost of $\kappa_\Delta \geq 0$ (with $\kappa_0 = 0$), modeling the the forgone opportunity of outside employment options, after driving for the platform for some time. A driver with $\beta_i = 0$ who does not enter the platform at $(\ell_i, \underline{\tau}_i)$ does not incur any cost, and will not enter at a later time. Drivers have quasi-linear utilities, and seek to maximize the total payments received over the planning horizon minus the total costs.

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*,

¹³We may think about each time period as ~ 5 minutes, and with $T = 6$ the planning horizon would be half an hour.

¹⁴We can also allow the distance between a pair of locations to change over time, modeling the changes in traffic conditions, i.e. a trip from a to b starting at time t ends at time $t + \delta(a, b, t)$. This does not affect the results presented in this paper, and we keep $\delta(a, b)$ for simplicity of notation.

only value trips starting at τ_j , are not willing to relocate or walk from a drop-off point to their actual, intended destination, and do not have preference over drivers. The value v_j models the willingness-to-pay of the rider, over-and-above a base payment for a trip, which covers additional cost a driver incurs driving a rider from o_j to d_j at time τ_j (in comparison to relocating from o_j to d_j at time τ_j).¹⁵ 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 (travel times, trip costs, 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.¹⁶ Unless otherwise noted, we assume properties (S1), (S2), and complete, symmetric information throughout the paper, and discuss the effect of relaxing these assumptions in Section 4.3.

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 or is able to enter the platform, and has not yet exited, and is not *en route*. A driver who is available at time t and location a is able to complete a pick-up at this location and time. We allow a driver to drop-off a rider and pick-up another rider in the same location at the same time point (see Appendix A).

A *path* is a sequence of tuples (a, b, t) , representing driver entrance, exit, and the trips she takes over the planning horizon. 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. The path $Z_{i,0}$ includes no trip: for a driver with $\beta_i = 0$, $Z_{i,0}$ models the option to not enter the platform at all; for a driver s.t. $\beta_i = 1$, $Z_{i,0}$ models the option exit immediately at $(\ell_i, \underline{\tau}_i)$. For each $k = 1, \dots, |\mathcal{Z}_i|$, $Z_{i,k}$ is a path that starts $(\ell_i, \underline{\tau}_i)$, with the starting time and location of each successive trip equal to the ending time and location of the previous trip. Denote $(a, b, t) \in Z_{i,k}$ if path $Z_{i,k}$ includes (or *covers*) trip (a, b, t) , and let $\lambda_{i,k}$ be the total cost of the k^{th} path to driver i . We know that $\lambda_{i,0} = 0$ if $\beta_i = 0$, $\lambda_{i,0} = \kappa_{T-\underline{\tau}_i}$ if $\beta_i = 1$, and for $k > 1$, $\lambda_{i,k} = \sum_{(a,b,t) \in \mathcal{T}} c_{a,b,t} \mathbb{1}\{(a,b,t) \in Z_{i,k}\} + \kappa_{\Delta}$, if path $Z_{i,k}$ ends Δ periods earlier than T .

Driver i that takes the path $Z_{i,k}$ is able to pick up rider j if $(o_j, d_j, \tau_j) \in Z_{i,k}$, however, a path specifies only the movement in space and time, and does not specify whether a rider is picked up for each of the trips on the path.

Let an *action path* for driver i be a sequence of tuples, each of them can either be of the form (a, b, t) , representing a relocation trip from a to b at time t without a rider, or be of the form (a, b, t, j) , in which case the driver sends rider j from a to b at time t (thus requiring $(a, b, t) = (o_j, d_j, \tau_j)$). Let $\tilde{\mathcal{Z}}_i$ be the set of all feasible action paths of driver i . For an action path $\tilde{z}_i \in \tilde{\mathcal{Z}}_i$, denote $(a, b, t) \in \tilde{z}_i$ or $(a, b, t, j) \in \tilde{z}_i$ if the action path includes a relocation or rider trip from a to b at time t . A driver taking action path \tilde{z}_i that is *consistent* with path $Z_{i,k}$ (i.e. results in the same movement in space and time) incurs a total cost of $\lambda_{i,k}$.

Example 1. The planning horizon is $T = 2$ and there are two locations $\mathcal{L} = \{A, B\}$ with distance $\delta(A, A) = \delta(B, B) = 1$ and $\delta(A, B) = \delta(B, A) = 2$. See Figure 1. Trip costs are given by $c_{a,b,t} = 2\delta(a, b)$ for all $(a, b, t) \in \mathcal{T}$, and the opportunity cost of exiting early is $\kappa_{\Delta} = \Delta$. There is one driver, who has not yet entered the platform (i.e. $\beta_1 = 0$), but is able to enter at time $\underline{\tau}_1 = 0$ at location $\ell_1 = A$, and leaves at time $\bar{\tau}_1 = 2$. There are three riders:

¹⁵This base payment is made directly from a rider to the driver that picked her up, so that the cost any driver incurs for any trip $(a, b, t) \in \mathcal{T}$ is always $c_{a,b,t}$, regardless of whether a rider is in the back seat. The prices that we design, as a result, are the amounts that the riders pay on top of the base payments.

¹⁶More generally, it is sufficient that it be common knowledge amongst drivers that the platform has the correct information.




Figure 1: The economy in Example 1, with two locations A , B , two time periods and three riders.

- Rider 1: $o_1 = A$, $d_1 = A$, $\tau_1 = 0$, $v_1 = 5$,
- Rider 2: $o_2 = A$, $d_2 = A$, $\tau_2 = 1$, $v_1 = 6$,
- Rider 3: $o_3 = A$, $d_3 = B$, $\tau_3 = 0$, $v_3 = 8$.

In addition to not entering the platform at all— path $Z_{1,0}$ with cost $\lambda_{1,0} = 0$, there are three more feasible paths for driver 1: $Z_{1,1} = ((A, A, 0), (A, A, 1))$, $Z_{1,2} = ((A, B, 0))$, and $Z_{1,3} = ((A, A, 0))$. In $Z_{1,3}$, the driver exits one period before the end of planning horizon. The path costs are $\lambda_{1,1} = c_{A,A,0} + c_{A,A,1} = 2$, $\lambda_{1,2} = c_{A,B,0} = 4$, and $\lambda_{1,3} = c_{A,A,0} + \kappa_1 = 3$. Path $((A, A, 0), (A, B, 1))$ is infeasible, since the last trip ends later than the driver's leaving time. Similarly, paths $((A, B, 0), (B, B, 1))$ and $((A, A, 0), (B, B, 1))$ are infeasible.

In addition to not entering, there are eight feasible actions paths of rider 1. $((A, B, 0))$, relocating from A to B at time 0, and $((A, B, 0, 3))$, sending rider 3 from A to B at time 0, are both consistent with the path $Z_{1,2}$, and both have cost 4. Four action paths, $((A, A, 0), (A, A, 1))$, $((A, A, 0, 1), (A, A, 1))$, $((A, A, 0), (A, A, 1, 2))$, $((A, A, 0, 1), (A, A, 1, 2))$, are consistent with $Z_{1,1}$ and have cost 4. Both $((A, A, 0))$ and $((A, A, 0, 1))$ are both consistent with $Z_{1,3}$ and have cost 3. \square

We now provide an informal timeline of a *ridesharing mechanism* (see Section 4 for a formal definition). At each time point $t \in [T - 1]$, given the history of trips, current driver locations, driver availability status, and information about future driver supply and rider demand for trips:

1. The ridesharing mechanism determines for each rider with trip start time t , whether a driver will be dispatched to pick her up, and if so, the price of her trip.
2. The mechanism dispatches available drivers to pick up riders, to relocate, or to exit (for those already in the platform), or not to enter (for those have not entered, with $\tau_i = t$ and $\beta_i = 0$). The mechanism also determines the payments offered to drivers for accepting the dispatches.
3. Each available driver decides whether to accept the dispatch, or to deviate and either stay in the same location, or relocate, or exit/not enter. A driver may still decide to enter the platform even if asked not to do so. The mechanism collects and makes payments based on driver actions.

Any undispached, available driver makes their own choices of actions. We assume that any driver already *en route* will continue their current trip. A driver's payment in a period in which the driver declines a dispatch is zero, so that drivers are not charged penalties for deviation.

As a baseline, we define the following myopic pricing mechanism. For each rider $j \in \mathcal{R}$, denote the per-period surplus of her trip as $w_j \triangleq (v_j - c_{o_j, d_j, \tau_j}) / \delta(o_j, d_j)$.

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 with $(o_j, \tau_j) = (a, t)$ and $w_j \geq 0$, in decreasing order of w_j . The mechanism sets a market clearing rate $\rho_{a,t}$ (i.e. between highest unallocated w_j and lowest allocated w_j), and sets prices $p_{a,b,t} = \rho_{a,t} \delta(a, b) + c_{a,b,t}$ for each destination $b \in \mathcal{L}$, which is offered to all dispatched drivers and collected from all riders.




Figure 2: A Super Bowl game: time 0 plan under the myopic pricing mechanism.

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

Example 2 (Super Bowl example). Consider the economy illustrated in Figure 2, modeling the end of a sports event, with $T = 3$ time periods and three locations. Time $t = 0$ is 9:50pm, and 10 minutes before the Super Bowl ends, and each time period is 10 minutes. There are three locations A , B and C with symmetric distances $\delta(A, A) = \delta(B, B) = \delta(C, C) = \delta(A, B) = \delta(B, A) = \delta(B, C) = \delta(C, B) = 1$ 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, with exit times $\bar{\tau}_i = T$ for all $i \in \mathcal{D}$. At time 1, many riders with very high values show up at location C , where the game takes place. Riders' trips and their willingness to pay are as shown in the figure. Assume the trips cost drivers 10 per time period, i.e. $c_{a,b,t} = 10\delta(a, b)$ for all $(a, b, t) \in \mathcal{T}$, and exiting early costs 5 per period: $\kappa_\Delta = 5\Delta$.

Under the myopic pricing mechanism, at time 0, drivers 1 and 2 are dispatched to pick up riders 1 and 2 respectively, and driver 3 is dispatched to pick up rider 4. At time 1, driver 1 picks up rider 5. Even assuming that drivers exit optimally (i.e. driver 2 exits at time 2 at a cost of $\kappa_1 = 5$, and drivers 2 and 3 exit at time 1 and each incur a cost of $\kappa_2 = 10$), the total social welfare would be no more than $v_1 + v_2 + v_4 + v_5 - 10 \times 4 - 5 - 10 - 10 = 25$.

The set of market clearing rates for location C at time 0 determined by a myopic pricing mechanism is $\rho_{C,0} \in [0, w_1] = [0, 10]$, thus the possible market clearing prices for the trip $(C, B, 0)$ is $p_{C,B,0} \in [10, 20]$. The price for the $(B, B, 1)$ trip is $p_{B,B,1} = c_{B,B,1} = 10$, since there is excessive supply. The highest possible total utility to driver 1 under any myopic pricing mechanism would in fact be $(20 - 10) + (10 - 10) - 5 = 5$. At time 1, since no driver is able to pick up the four riders at location C , the lowest market clearing rate is $\rho_{C,1} = w_6 = 90$, thus the prices must be at least $p_{C,B,1} \geq \rho_{C,1} + c_{C,B,1} = 100$ and $p_{C,A,1} \geq \delta(C, A)\rho_{C,1} + c_{C,A,1} = 200$.

Suppose driver 1 deviates from the dispatches and stays in location C until time 1. The mechanism would then dispatch driver 1 to pick up rider 6, and driver 1 would be paid the new market clearing price of $\delta(C, B)w_7 + c_{C,B,1} = 50$. This is a useful deviation, since exiting at time 2, the utility of driver 1 is $50 - 10 - 10 - 5 = 25$. Observe that by strategizing, driver 1 improves the social welfare by making use of information about the future demand and supply. Also observe that the outcome is not in the core: driver 1 at $(C, 1)$ has an incentive to make a side deal with rider 7, who is willing to pay more than 50, but is not picked up. \square

3 A Static CE Mechanism

In this section, we formulate the optimal planning problem and define competitive equilibrium (CE) prices. We also introduce the *static CE mechanism*, which computes an optimal plan at the start of the planning horizon and does not replan in the future. This makes the static CE mechanism fragile, for example it is not welfare optimal or envy-free forward from a period in which some driver has deviated from the suggested plan. The STP mechanism resolves this concern, combining driver-pessimal CE prices with replanning to provide robustness forward from any history.

3.1 Plans

A *plan* describes the paths taken by all drivers until the end of the planning horizon, rider pick-ups, as well as payments for riders and drivers for each trip associated with these paths.

Formally, a plan is the 4-tuple (x, \tilde{z}, q, r) , 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; \tilde{z} is a vector of action paths, where $\tilde{z}_i \in \tilde{\mathcal{Z}}_i$ is the dispatched action path for driver i ; q_j denotes the payment made by rider j , $r_{i,t}$ denotes the payment made to driver i at time t , and let $r_i \triangleq \sum_{t=0}^T r_{i,t}$ denote the total payment to driver i . If \tilde{z}_i is consistent with $Z_{i,k}$, the k^{th} feasible path of driver i , then the driver incurs a total cost of $\lambda_{i,k}$, and her utility is $\pi_i \triangleq r_i - \lambda_{i,k}$.

A plan (x, \tilde{z}, q, r) is *feasible* if for each rider $j \in \mathcal{R}$, $x_j = \sum_{i \in \mathcal{D}} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} \in \{0, 1\}$, where $\mathbb{1}\{\cdot\}$ is the indicator function. Unless stated otherwise, when we mention a plan in the rest of the paper, it is assumed to be feasible. For the *budget balance* (BB) of a plan, we need:

$$\sum_{j \in \mathcal{R}} q_j \geq \sum_{i \in \mathcal{D}} r_i, \quad (1)$$

with strict budget balance if (1) holds with equality. A plan is *individually rational for riders* if

$$x_j v_j \geq q_j, \quad \forall j \in \mathcal{R}.$$

A plan is *individually rational for drivers* if $\pi_i \geq 0$ for all $i \in \mathcal{D}$ s.t. $\beta_i = 0$, i.e. drivers that are not yet in the platform do not get negative utility from participating. A plan is *envy-free for riders* if no rider strictly prefers the outcome of another rider requesting the same trip, that is

$$x_j v_j - q_j \geq x_{j'} v_{j'} - q_{j'} \text{ for all } j, j' \in \mathcal{R} \text{ s.t. } o_j = o_{j'}, d_j = d_{j'}, \text{ and } \tau_j = \tau_{j'}. \quad (2)$$

A plan is *envy-free for drivers* if any pair of drivers with the same type have the same utility:

$$\pi_i = \pi_{i'} \text{ for all } i, i' \in \mathcal{D} \text{ s.t. } \tau_i = \tau_{i'}, \ell_i = \ell_{i'}, \text{ and } \beta_i = \beta_{i'}. \quad (3)$$

Definition 2 (Anonymous trip prices). A plan (x, \tilde{z}, q, r) uses *anonymous trip prices* if there exist $p = \{p_{a,b,t}\}_{(a,b,t) \in \mathcal{T}}$ such that for all $(a,b,t) \in \mathcal{T}$, we have:

- (i) all riders taking the same (a,b,t) trip are charged the same payment $p_{a,b,t}$, and there is no payment by riders who are not picked up, and
- (ii) all drivers that are dispatched on a rider trip from a to b at time t are paid the same amount $p_{a,b,t}$ for the trip at time t , and there is no other payment to or from any driver.

Given dispatches (x, \tilde{z}) and anonymous trips prices p , all payments are fully determined: the total payment to driver i is $r_i = \sum_{j \in \mathcal{R}} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} p_{o_j, d_j, \tau_j}$ and the payment made by rider j is $q_j = x_j p_{o_j, d_j, \tau_j}$. For this reason, we will represent plans with anonymous trip prices as (x, \tilde{z}, p) . By construction, plans with anonymous trip prices are strictly budget balanced.

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

- (i) (rider best response) 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) (driver best response) $\forall i \in \mathcal{D}$, $\pi_i = \max_{k=0, \dots, |\mathcal{Z}_i|} \left\{ \sum_{(a,b,t) \in \mathcal{Z}_{i,k}} \max\{p_{a,b,t}, 0\} - \lambda_{i,k} \right\}$, i.e. each driver achieves the highest possible utility given prices and the set of feasible paths.

Given any set of anonymous trip prices p , let anonymous trip prices p^+ be defined as $p_{a,b,t}^+ \triangleq \max\{p_{a,b,t}, 0\}$ for each $(a, b, t) \in \mathcal{T}$.

Lemma 1. Given any CE plan (x, \tilde{z}, p) , the plan with anonymous prices (x, \tilde{z}, p^+) also forms a CE, and has the same driver and rider payments and utilities as those under (x, \tilde{z}, p) .

The lemma implies that when studying the set of possible rider and driver payments and utilities among all CE outcomes, it is without loss to consider only anonymous trip prices that are non-negative. We leave the full proof to Appendix B.1. Prices must be non-negative for any trip that is requested by any rider, thus changing prices from p to p^+ does not affect the payments for any rider or driver, or the best response on the riders side. The driver best response property also continues to hold, since $\max\{p_{a,b,t}, 0\} = \max\{p_{a,b,t}^+, 0\}$ for all $(a, b, t) \in \mathcal{T}$.

Given any mechanism, complete information about supply and demand, and assuming that all drivers follow the dispatches of the mechanism at all times, the assignments of all riders, action paths taken by all drivers, and the corresponding payment schedule through the planning horizon can be computed at time 0, if all available drivers are dispatched at all times. We call this outcome the “time 0 plan” under the given mechanism.

3.2 Optimal Plans and CE Prices

The welfare-optimal planning problem can be formulated as an integer linear program (ILP) that determines optimal rider pick-ups and driver paths, followed by an assignment of riders to drivers whose paths cover the rider trips. 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:

$$\max_{x,y} \sum_{j \in \mathcal{R}} x_j v_j - \sum_{i \in \mathcal{D}} \sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} \lambda_{i,k} \quad (4)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{R}} x_j \mathbb{1}\{(o_j, d_j, \tau_j) = (a, b, t)\} \leq \sum_{i \in \mathcal{D}} \sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} \mathbb{1}\{(a, b, t) \in Z_{i,k}\}, \quad \forall (a, b, t) \in \mathcal{T} \quad (5)$$

$$\sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} = 1, \quad \forall i \in \mathcal{D} \quad (6)$$

$$x_j \in \{0, 1\}, \quad \forall j \in \mathcal{R} \quad (7)$$

$$y_{i,k} \in \{0, 1\}, \quad \forall i \in \mathcal{D}, k = 1, \dots, |\mathcal{Z}_i| \quad (8)$$

Constraint (6) requires that each driver takes exactly one path (which includes the path $Z_{i,0}$ representing not entering/exiting immediately). The feasibility constraint (5) requires that for all trips $(a, b, t) \in \mathcal{T}$, the number of riders who request this trip and are picked up is no more than the total number of drivers whose paths cover this trip. Once pick-ups x and paths y are computed, (5) guarantees that each rider with $x_j = 1$ can be assigned to a driver.

Relaxing the integrality constraints on variables x and y , we obtain the following linear program (LP) relaxation of the ILP:

$$\max_{x,y} \sum_{j \in \mathcal{R}} x_j v_j - \sum_{i \in \mathcal{D}} \sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} \lambda_{i,k} \quad (9)$$

$$\text{s.t. } \sum_{j \in \mathcal{R}} x_j \mathbb{1}\{(o_j, d_j, \tau_j) = (a, b, t)\} \leq \sum_{i \in \mathcal{D}} \sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} \mathbb{1}\{(a, b, t) \in Z_{i,k}\}, \quad \forall (a, b, t) \in \mathcal{T} \quad (10)$$

$$\sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} = 1, \quad \forall i \in \mathcal{D} \quad (11)$$

$$x_j \leq 1, \quad \forall j \in \mathcal{R} \quad (12)$$

$$x_j \geq 0, \quad \forall j \in \mathcal{R} \quad (13)$$

$$y_{i,k} \geq 0, \quad \forall i \in \mathcal{D}, k = 1, \dots, |\mathcal{Z}_i| \quad (14)$$

We refer to (9) as the primal LP. The constraint $y_{i,k} \leq 1$, that each path is taken by each driver at most once is guaranteed by imposing (11) and (14), and is omitted.

Lemma 2 (Integrality). There exists an integer optimal solution to the linear program (9).

We leave the proof to Appendix B.2, showing there a correspondence to a min cost flow (MCF) problem (that has integral optimal solutions), where drivers flow through a network with vertices corresponding to (location, time) pairs, edges corresponding to trips, and with edge costs equal to driver's costs minus riders' values. The reduction to MCF can also be used to solve the LP efficiently.

Let $p_{a,b,t}$, π_i and u_j denote the dual variables corresponding to the primal constraints (10), (11) and (12), respectively. The dual LP of (9) is:

$$\min \sum_{i \in \mathcal{D}} \pi_i + \sum_{j \in \mathcal{R}} u_j \quad (15)$$

$$\text{s.t. } \pi_i \geq \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} - \lambda_{i,k} \quad \forall k = 0, 1, \dots, |\mathcal{Z}_i|, \forall i \in \mathcal{D} \quad (16)$$

$$u_j \geq v_j - p_{o_j, d_j, \tau_j}, \quad \forall j \in \mathcal{R} \quad (17)$$

$$p_{a,b,t} \geq 0, \quad \forall (a, b, t) \in \mathcal{T} \quad (18)$$

$$u_j \geq 0, \quad \forall j \in \mathcal{R} \quad (19)$$

Lemma 3 (Welfare Theorem). A dispatching (x, \tilde{z}) is welfare-optimal if and only if there exists anonymous trip prices p s.t. the plan (x, \tilde{z}, p) forms a competitive equilibrium. Such optimal CE plans always exist and are efficient to compute. Moreover, these plans are strictly budget balanced, and are individually rational and envy-free for both riders and drivers.

Given optimal primal and dual solutions, we show that the dual variables π and u can be interpreted as utilities of drivers and riders, when the anonymous trip prices are given by p . We then make use of Lemma 1, and the standard observations about complementary slackness conditions and their connection with competitive equilibria [32, 8]. By integrality, optimal CE plans always exist, and can be efficiently computed by solving the primal and dual LPs of the MCF problem. See Appendix B.3 for the proof of the lemma.

For any two driver utility profiles $\pi = (\pi_1, \dots, \pi_m)$ and $\pi' = (\pi'_1, \dots, \pi'_m)$ that correspond to CE plans, let the join $\bar{\pi} = \pi \vee \pi'$ and the meet $\underline{\pi} = \pi \wedge \pi'$ be defined as $\bar{\pi}_i \triangleq \max\{\pi_i, \pi'_i\}$ and $\underline{\pi}_i \triangleq \min\{\pi_i, \pi'_i\}$ for all $i \in \mathcal{D}$. The following lemma shows that drivers' utilities among all CE outcomes form a lattice, meaning that there exist CE plans where the driver utilities are given by $\bar{\pi}$ or $\underline{\pi}$. The lemma also shows a connection between the top/bottom of the lattice and the welfare differences from losing/replicating a driver, which plays an important role in establishing the incentive properties of the STP mechanism.

Denote $W(\mathcal{D}, \mathcal{R})$ as the highest welfare achievable by drivers \mathcal{D} and riders \mathcal{R} (i.e. the optimal objective of (9)). For each driver $i \in \mathcal{D}$, define the *social welfare gain from replicating driver i* , and the *social welfare loss from losing driver i* , as:

$$\Phi_{D_i} \triangleq W(\mathcal{D} \cup \{i'\}, \mathcal{R}) - W(\mathcal{D}, \mathcal{R}), \quad (20)$$

$$\Psi_{D_i} \triangleq W(\mathcal{D}, \mathcal{R}) - W(\mathcal{D} \setminus \{i\}, \mathcal{R}), \quad (21)$$

where driver i' with $\theta_{i'} = \theta_i$ is a replica of driver i . A *driver-optimal plan* has a driver utility profile at the top of the lattice, and a *driver-pessimal plan* has a utility profile at the bottom of the lattice.

Lemma 4 (Lattice Structure). Drivers' utility profile π among all CE outcomes form a lattice. Moreover, for each driver $i \in \mathcal{D}$, Φ_{D_i} and Ψ_{D_i} are equal to utility of driver i in the driver-pessimal and driver-optimal CE plans, respectively.

We leave the proof of Lemma 4 to Appendix B.4. The lattice structure follows from the correspondences between driver utilities, optimal solutions to dual LP (15), and optimal solutions to the dual of the flow LP, and the fact that optimal dual solutions of MCF form a lattice. Standard arguments on shortest paths in the residual graph [2], and the connection between optimal dual solutions and subgradients (w.r.t. flow boundary conditions), then imply the correspondence between welfare gains/losses and driver pessimal/optimal utilities.

Unlike the classical unit-demand assignment problem, where the prices of items and the utilities of buyers both form lattices [35], it is drivers' utilities, and not trip prices or rider utilities that have a lattice structure. This is because although the driver-pessimal utilities are unique, the price for a trip under all driver-pessimal plans need not be unique. See Section 4.3 for further a development.

A plan is in the *core* of a ridesharing problem if no coalition of riders and drivers can break out of this plan and make an alternative plan among themselves, such that all drivers and riders in the coalition get at least their utilities from the original plan, and at least one of the drivers or riders is strictly better off. The following lemma proves that CE and core are utility-equivalent.

Lemma 5 (Core equivalence). All CE plans are in the core. Moreover, for any budget-balanced core outcome (x, \tilde{z}, q, r) , there exists prices p such that the plan with anonymous prices (x, \tilde{z}, p) forms a CE, and has the same driver and rider total utilities.

See Appendix B.5 for the full proof of this lemma. Intuitively, every CE plan is in the core since for any $\mathcal{D}' \subseteq \mathcal{D}$ and $\mathcal{R}' \subseteq \mathcal{R}$, the highest achievable coalitional welfare $W(\mathcal{D}', \mathcal{R}')$ is no greater than the sum of utilities of all driver and riders in this coalition under any CE plan. Given any core outcome, we can construct anonymous trip prices p that support the outcome in CE, and have the same driver and rider total payments: $q_j = x_j p_{o_j, d_j, \tau_j}$ and $r_i = \sum_{j \in \mathcal{R}} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} p_{o_j, d_j, \tau_j}$.

3.3 The Static CE Mechanism

A *static mechanism* announces a plan (x, \tilde{z}, q, r) at time $t = 0$, and never again updates the plan, even after deviations by drivers. Rather, each driver can choose to take any feasible path, but can




Figure 3: The Super Bowl example: the driver pessimal competitive equilibrium plan.

only pick up riders that are dispatched to her (and is only paid for the subset of these rider trips that are completed). Riders pay according to prices q in the event they are picked up by a driver.

Definition 4 (Static CE mechanism). A static CE mechanism announces an optimal CE plan (x, \tilde{z}, p) at the beginning of the planning horizon. Each driver $i \in \mathcal{D}$ then decides on the actual action path \tilde{z}'_i that she takes, and gets paid $\hat{r}_i = \sum_{j \in \mathcal{R}} p_{o_j, d_j, \tau_j} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i, (o_j, d_j, \tau_j, j) \in \tilde{z}'_i\}$. Each rider $j \in \mathcal{R}$ pays q_j only if she is picked up.

A static CE mechanism can be defined for any set of CE prices. Driver best response guarantees that no other path gives her a higher total utility, and it is a dominant strategy for each driver to follow the dispatched action path \tilde{z}_i .

Theorem 1. *A static CE mechanism implements an optimal CE plan in dominant strategy.*

In addition, if all riders and drivers follow the plan, the outcome under a static CE mechanism is strictly budget balanced and envy-free for both sides. The CE property also ensures that every rider that is picked up is happy to take the trip at the offered price, and that no rider who is not picked up has positive utility for the trip.¹⁷

The optimal static mechanism enjoys many good properties, but has a decisive design flaw—it is fragile to driver deviations since it does not react by replanning. Deviations could occur for many reasons: mistakes, unexpected contingencies, unexpected traffic, or unmodeled idiosyncratic preferences, etc. We show by revisiting the Super Bowl example that once a driver has deviated, the resulting outcome in the subsequent periods may no longer be welfare-optimal or envy-free.

Example 2 (Continued). For the Super Bowl game scenario, the static CE mechanism with a driver-pessimal plan adopts the plan illustrated in Figure 3. In this plan, all drivers stay at location C or re-position to location C at time 0, picking up riders with high values. Driver 3 exits at time 2 after dropping off rider 6. The total rider value is 300, and the total trip costs incurred is 80, and the total exit cost is 5, resulting in an optimal welfare of 205.

The anonymous trip prices are shown in italics, below the edges corresponding to the trips. For each feasible path of each driver, the total prices minus costs is 50, the welfare gain of the economy

¹⁷Still, Example 8 in Appendix C.3 shows that truthful reporting of a rider's value need not be a dominant strategy (and this can be the case whichever CE prices are selected).




Figure 4: The economy in Example 3 and the driver pessimal CE plan computed at time 0.

if the driver is replicated. The outcome forms a CE, that there is no other path with a higher total utility for any driver. All riders are happy with their dispatched trips given the prices, and there is no driver or rider envy.

Suppose now that driver 3 did not follow the plan at time 0 to pick-up rider 3 going from B to C , but stayed in location B until time 1. The effect of this deviation and not updating the plan is that driver 3 is no longer able to pick up rider 6, who strictly prefers to be picked up given the original price of $p_{C,B,1} = 75$. One of the drivers 1 and 2 who was supposed to pick up rider 8 with value 90 is actually able to pick up rider 6, and this would lead to a higher social welfare. Moreover, driver 3 is now able to pick up rider 5, however, she wouldn't be dispatched to do so. \square

One may think of a naive fix for this robustness issue of the static CE mechanisms, simply repeating the computation of the plan at all times. The following example shows that the mechanism that recomputes a driver pessimal plan at all times fails to be incentive compatible. Similarly, we show that the mechanism that repeatedly recomputes a *driver-optimal* plan is not envy-free and can have incentive issues (see Example 11 in Appendix C).

Example 3. Consider the economy as shown in Figure 4, where there are two locations $\mathcal{L} = \{A, B\}$ with distances $\delta(A, A) = \delta(A, B) = \delta(B, A) = \delta(B, B) = 1$. Assume for simplicity that all trip costs and opportunity costs are zero: $c_{a,b,t} = 0$ for all $(a, b, t) \in \mathcal{T}$, and $\kappa_\Delta = 0$ for all $\Delta = 0, 1, \dots, T$. In the driver-pessimal plan computed at time 0 as shown in the figure, the anonymous trip prices are $p_{B,B,1} = p_{A,A,1} = 5$. Assume that both drivers 1 and 2 follow the plan at time 0, and reach $(B, 1)$ and $(A, 1)$ respectively. If the mechanism re-computes the plan at time 1, the new driver-pessimal plan would set a new price of 0 for the trip $(B, B, 1)$ — the updated lowest market-clearing price for the trip. Therefore, if driver 1 follows the mechanism at all times, her total payment and utility would actually be 0. Now consider the scenario where driver 2 follows the mechanism at time 0, but driver 1 deviates and relocates to A , so that both drivers are at location A at time 1. At time 1, when the mechanism recomputes a driver-pessimal plan, both drivers would take the trip $(A, A, 1)$ and pick up riders 2 and 3 respectively. The updated price for the trip $(A, A, 1)$ would be 4, and this is a useful deviation for driver 1. \square

The challenge is to achieve robustness, but at the same time handle the new strategic considerations that can occur as a result of drivers being able to trigger re-planning through deviations.

4 The Spatial-Temporal Pricing Mechanism

The Spatio-Temporal Pricing (STP) mechanism computes a driver-pessimal CE plan at the beginning of the planning horizon, and recomputes a driver-pessimal plan upon any deviation. This leads to our main result: the STP mechanism achieves subgame-perfect incentive compatibility without the use of time-extended contracts, and is welfare-optimal, core-selecting, individually rational, budget balanced, and envy-free from any history onward.

4.1 A Dynamic Mechanism

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

Let $s_t = (s_{1,t}, s_{2,t}, \dots, s_{m,t})$ denote the *state* of the ridesharing platform at time t , where each $s_{i,t}$ describes the state of driver $i \in \mathcal{D}$. If driver i has entered the platform and is available at time t at location $a \in \mathcal{L}$, denote $s_{i,t} = (1, a, t)$. Otherwise, if driver i is *en route*, finishing the trip from a to b that she started at time $t' < t$ s.t. $t' + \delta(a, b) > t$, denote $s_{i,t} = (a, b, t')$ if she is relocating with no rider, or $s_{i,t} = (a, b, t', j)$ if she is taking a rider j from a to b at time t' . For drivers that had already exited or decided not to enter, denote $s_{i,t} = \phi$. For drivers with $\tau_i \geq t$, i.e. who enters or is able to enter now or in the future, $s_{i,t} = (\beta_i, \ell_i, \tau_i)$. The initial state of the platform is $s_0 = ((\beta_1, \ell_1, \tau_1), \dots, (\beta_m, \ell_m, \tau_m))$.

At each time t , each driver i takes an *action* $\alpha_{i,t}$. An available driver i with $s_{i,t} = (1, a, t)$ or $s_{i,t} = (0, a, t)$ may (enter and then) relocate to any location b within reach by the end of the planning horizon (i.e. $b \in \mathcal{L}$ s.t. $t + \delta(a, b) \leq T$), which we denote $\alpha_{i,t} = (a, b, t)$. She may pick up a rider $j \in \mathcal{R}$ with $\tau_j = t$ and $o_j = a$, in which case we write $\alpha_{i,t} = (a, d_j, t, j)$. She may also decide to exit (if $\beta_i = 1$) or not enter (if $\beta_i = 0$), for both cases we denote $\alpha_{a,t} = \phi$. For a driver i that is *en route* at time t , (i.e. $s_{i,t} = (a, b, t')$ or $s_{i,t} = (a, b, t', j)$ for some t' s.t. $t' + \delta(a, b) > t$), $\alpha_{i,t} = s_{i,t}$ —the only available action is to finish the current trip. For driver i with $\tau_i > t$, denote $\alpha_{i,t} = s_{i,t} = (\beta_i, \tau_i, \ell_i)$. A driver with $s_{i,t} = \phi$ takes no more actions: $\alpha_{i,t} = s_{i,t} = \phi$.

The action $\alpha_{i,t}$ taken by driver i at time t determines her state $s_{i,t+1}$ at time $t+1$:

- (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 $s_{i,t+1} = (1, b, t+1)$, meaning these drivers will become available at time $t+1$ at the destination of their trips,¹⁸
- (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) for $i \in \mathcal{D}$ s.t. $\alpha_{i,t} = (\beta_i, \tau_i, \ell_i)$, we have $s_{i,t+1} = (\beta_i, \tau_i, \ell_i)$,
- (already exited / never entered) if $\alpha_{i,t} = \phi$, then $s_{i,t+1} = \phi$.

Let $\alpha_t = (\alpha_{1,t}, \alpha_{2,t}, \dots, \alpha_{m,t})$ be the *action profile* of all drivers at time t , and let *history* $h_t \triangleq (s_0, \alpha_0, s_1, \alpha_1, \dots, s_{t-1}, \alpha_{t-1}, s_t)$, with $h_0 = (s_0)$. Finally, let $\mathcal{D}_t(h_t) = \{i \in \mathcal{D} \mid s_{i,t} = (1, a, t) \text{ or } s_{i,t} = (0, a, t) \text{ for some } a \in \mathcal{L}\}$ be the set of drivers available at time t .

Definition 5 (Dynamic ridesharing mechanism). A dynamic ridesharing mechanism is defined by its *dispatch rule* α^* , *driver payment rule* r^* and *rider payment rule* q^* . At each time t , given history h_t and rider information \mathcal{R} , the mechanism:

- uses its dispatch rule α^* to determine for each of a subset of available drivers, a *dispatch action* $\alpha_{i,t}^*(h_t)$ to either pick up a rider, or to relocate, or to exit/not enter.
- uses its driver payment rule r^* to determine, for each dispatched driver, a payment $r_{i,t}^*(h_t)$ in the event the driver takes the action ($r_{i,t}^*(h_t) = 0$ for available drivers that are not dispatched).
- dispatches each *en route* driver to keep driving (i.e. $\alpha_{i,t}^*(h_t) = s_{i,t}$), and does not make any payment to driver i in this period: $r_{i,t}^*(h_t) = 0$.
- determines for drivers entering in the future ($i \in \mathcal{D}$ s.t. $\tau_i > t$), and drivers who had already exited ($i \in \mathcal{D}$ s.t. $s_{i,t} = \phi$), $\alpha_{i,t}^*(h_t) = s_{i,t}$ and $r_{i,t}^*(h_t) = 0$.

¹⁸Here we assume that a driver that declines the mechanism's dispatch and decide to relocate from a to b also does so in time $\delta(a, b)$. We can also handle drivers who move more slowly when deviating, just as long as the mechanism knows when and where the driver will become available again.

- uses its rider payment rule q^* to determine, for each rider who receives a dispatch at time t , the payment $q_j^*(h_t)$ in the event that the rider is picked up.

Each driver then decides on which action $\alpha_{i,t} \in \mathcal{A}_{i,t}(h_t)$ to take, where $\mathcal{A}_{i,t}(h_t)$ is the set of actions available to agent i at time t given history h_t . For an available driver at (a, t) with dispatched action $\alpha_{i,t}^*(h_t)$, $\mathcal{A}_{i,t}(h_t) = \{\alpha_{i,t}^*(h_t)\} \cup \{(a, b, t) \mid b \in \mathcal{L} \text{ s.t. } t + \delta(a, b) \leq T\} \cup \{\phi\}$, i.e. the driver can either take the dispatched action, or to relocate to any location, or to exit or not enter; if an available driver at (a, t) is not dispatched, $\alpha_{i,t}^*(h_t)$, $\mathcal{A}_{i,t}(h_t) = \{(a, b, t) \mid b \in \mathcal{L} \text{ s.t. } t + \delta(a, b) \leq T\} \cup \{\phi\}$; for an *en route* driver, or a driver that enters in the future, or a driver that has already exited, $\mathcal{A}_{i,t}(h_t) = \{s_{i,t}\}$. After observing the action profile α_t , the mechanism pays each dispatched driver $\hat{r}_{i,t}(\alpha_{i,t}, h_t) = r_{i,t}^*(h_t) \mathbb{1}\{\alpha_{i,t} = \alpha_{i,t}^*\}$, and charges each rider $j \in \mathcal{R}$ with $\tau_j = t$ the amount $\hat{q}_j(\alpha_t) = q_j^*(h_t) \sum_{i \in \mathcal{D}_t} \mathbb{1}\{\alpha_{i,t} = (o_j, d_j, t, j)\}$.

A mechanism is *feasible* if (i) it is possible for each available driver to take the dispatched trip, i.e. $\forall t, \forall h_t, \forall i \in \mathcal{D}_t$, if $s_{i,t} = (1, a, t)$ or $s_{i,t} = (0, a, t)$ for some $a \in \mathcal{L}$, $\alpha_{i,t}^*(h_t) \in \{(a, b, t) \mid b \in \mathcal{L}, t + \delta(a, b) \leq T\} \cup \{(o_j, d_j, \tau_j, j) \mid j \in \mathcal{R}, \tau_j = t, o_j = a\}$, (ii) no rider is picked-up more than once, i.e. $\forall t, \forall h_t, \forall j \in \mathcal{R} \text{ s.t. } \tau_j = t, \sum_{i \in \mathcal{D}_t} \mathbb{1}\{\alpha_{i,t}^*(h_t) = (o_j, d_j, \tau_j, j)\} \leq 1$, and (iii) unavailable drivers are not dispatched. From Definition 5, there is no payment to or from unavailable or undispached drivers, or a dispatched driver i who deviated from $\alpha_{i,t}^*(h_t)$ at time t , or riders who are not picked up.

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 she takes $\alpha_{i,t} = \sigma_i(h_t) \in \mathcal{A}_{i,t}(h_t)$. For a mechanism that always dispatches all available drivers, σ_i^* denotes the *straightforward strategy* of always following the mechanism's dispatches at all times. Let $\sigma = (\sigma_1, \dots, \sigma_m)$ be the *strategy profile*, with $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_m)$. The strategy profile σ , together with the initial state s_0 and the rules of a mechanism, determine all actions and payments of all drivers through the planning horizon. Let $\sigma_i|_{h_t}$, $\sigma|_{h_t}$ and $\sigma_{-i}|_{h_t}$ denote the strategy profile from time t and history h_t onward for driver i , all drivers, and all drivers but i , respectively.

For each rider $j \in \mathcal{R}$, let $\hat{x}_j(\sigma) \in \{0, 1\}$ be the indicator that rider j is picked-up given strategy σ , and let $\hat{q}_j(\sigma) = \hat{x}_j(\sigma) q_j^*(h_{\tau_j})$ be her actual payment. For each driver $i \in \mathcal{D}$, $\hat{r}_i(\sigma) \triangleq \sum_{t=0}^{T-1} \hat{r}_{i,t}(\sigma_i(h_t), h_t)$ denotes 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 . Let $\hat{\pi}_{i,t}(\sigma_i(h_t), h_t)$ be the actual utility driver i gets at time t given history h_t and strategy σ_i , we know that if $\sigma_i(h_t) = (a, b, t)$ or $\sigma_i(h_t) = (a, b, t, j)$, then $\hat{\pi}_{i,t}(\sigma_i(h_t), h_t) = \hat{r}_{i,t}(\sigma_i(h_t), h_t) - c_{a,b,t}$; if $\sigma_i(h_t) = \phi$ and $s_{i,t} = (1, a, t)$ for some $a \in \mathcal{L}$, then $\hat{\pi}_{i,t}(\sigma_i(h_t), h_t) = \hat{r}_{i,t}(\sigma_i(h_t), h_t) - \kappa_{T-t}$. For every other scenario, $\hat{\pi}_{i,t}(\sigma_i(h_t), h_t) = \hat{r}_{i,t}(\sigma_i(h_t), h_t)$. Denote $\hat{\pi}_i(\sigma) \triangleq \sum_{t=0}^{T-1} \hat{\pi}_{i,t}(\sigma_i(h_t), h_t)$ as driver i 's total utility.

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 utility $\hat{\pi}_{i,t}(\alpha_{i,t}, h_t)$. The total utility $\hat{\pi}_i(\sigma)$ to each driver is determined by the rules of the mechanism.

We define the following properties.

Definition 6 (Budget balance). A ridesharing mechanism is *budget balanced* if for any set of riders and drivers, and any strategy profile σ taken by the drivers, we have

$$\sum_{j \in \mathcal{R}} \hat{q}_j(\sigma) \geq \sum_{i \in \mathcal{D}} \hat{r}_i(\sigma). \quad (22)$$

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) among the drivers, meaning for all $t \in [T - 1]$, for any history $h_t \in \mathcal{H}_t$,

$$\sum_{t'=t}^{T-1} \hat{\pi}_{i,t'}(\sigma_i^* | h_t, \sigma_{-i}^* | h_t) \geq \sum_{t'=t}^{T-1} \hat{\pi}_{i,t'}(\sigma_i | h_t, \sigma_{-i}^* | h_t), \quad \forall \sigma_i, \quad \forall i \in \mathcal{D}. \quad (23)$$

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 (Individual rationality (IR)). A ridesharing mechanism that always dispatches all available drivers is *individually rational in SPE for drivers* if for any set of riders and drivers, (i) the mechanism is SPIC for drivers, and (ii) assuming σ^* , drivers that have not yet entered do not get negative utility from participating, i.e.

$$\hat{\pi}_i(\sigma^*) \geq 0 \text{ for } i \in \mathcal{D} \text{ s.t. } \beta_i = 0.$$

A ridesharing mechanism is *individually rational for riders* if for any set of riders and drivers, and any strategy profile σ taken by the drivers,

$$\hat{x}_j(\sigma)v_j \geq \hat{q}_j(\sigma), \quad \forall i \in \mathcal{R}. \quad (24)$$

Definition 9 (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:

$$\sum_{t'=t}^{T-1} \hat{\pi}_{i,t'}(\sigma^* | h_t) = \sum_{t'=t}^{T-1} \hat{\pi}_{i',t'}(\sigma^* | h_t), \quad \forall i, i' \in \mathcal{D} \text{ s.t. } s_{i,t} = s_{i',t}. \quad (25)$$

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. $(o_j, d_j, \tau_j) = (o_{j'}, d_{j'}, \tau_{j'})$

$$\hat{x}_j(\sigma^*)v_j - \hat{q}_j(\sigma^*) \geq \hat{x}_{j'}(\sigma^*)v_j - \hat{q}_{j'}(\sigma^*). \quad (26)$$

Definition 10 (Core-selecting). A ridesharing mechanism that always dispatches all available drivers is *core-selecting* if for any set of riders and drivers, (i) the mechanism is SPIC, and (ii) for any time $t \in [T - 1]$ and any history $h_t \in \mathcal{H}_t$ onward, the outcome under the straightforward strategy σ^* is in the core.

Fix a ridesharing mechanism with dispatch rule α^* and payment rules q^* , r^* , where all available drivers are always dispatched. Recall that given complete information and the straightforward strategy σ^* , the outcome over the entire planning horizon can be computed at time 0, and is called the *time 0 plan* of the mechanism. If some driver deviated at time $t - 1$ for some $t > 0$, the downward outcomes given the dispatching and payment rules, assuming all drivers follow $\sigma^* | h_t$, can be thought of as an updated time t plan.

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 , and the remaining riders $\mathcal{R}^{(t)} = \{(o_j, d_j, \tau_j - t, v_j) \mid j \in \mathcal{R}, \tau_j \geq t\}$. For drivers, we have $\mathcal{D}^{(t)}(s_t) = \{\theta_i^{(t)} \mid i \in \mathcal{D}\}$, with types $\theta_i^{(t)} = (\beta_i^{(t)}, \underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)})$ determined as follows:

- for available drivers $i \in \mathcal{D}$ s.t. $s_{i,t} = (1, a, t)$ or $s_{i,t} = (0, a, t)$ for some $a \in \mathcal{L}$, let $\theta_i^{(t)} = (\beta_i^{(t)}, \underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) = (1, a, 0, \bar{\tau}_i - t)$ or $(1, a, 0, \bar{\tau}_i - t)$, respectively,
- for *en route* drivers $i \in \mathcal{D}$ s.t. $s_{i,t} = (a, b, t')$ or (a, b, t', j) where $t' + \delta(a, b) > t$, let $\theta_i^{(t)} = (\beta_i^{(t)}, \underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) = (1, b, t' + \delta(a, b) - t, \bar{\tau}_i - t)$,
- for each driver $i \in \mathcal{D}$ with $\underline{\tau}_i > t$, let $\theta_i^{(t)} = (\beta_i^{(t)}, \underline{\tau}_i^{(t)}, \bar{\tau}_i^{(t)}, \ell_i^{(t)}) = (\beta_i, \ell_i, \underline{\tau}_i - t, \bar{\tau}_i - t)$, and
- exclude drivers that have already exited or chose not to enter.

Definition 11 (Temporal consistency). A ridesharing mechanism is *temporally consistent* if after deviation at time $t - 1$, the updated plan is identical to that determined for economy $E^{(t)}(s_t)$.

Upon deviation(s) at time $t - 1$, a temporally consistent mechanism computes its updated plan from time t onward as if t is the beginning of the planning horizon, thus the mechanism does not make use of time-extended contracts, including penalties for previous actions. In fact, a temporally inconsistent mechanism would be able to trivially align incentives. For example, a mechanism that replans based on the history could fire any driver who has deviated in the past, while keeping the plans for the rest of the economy unchanged. A mechanism can also threaten to “shut down” and not make any further dispatches or payments to the drivers if any of them had deviated.

4.2 The Spatio-Temporal Pricing Mechanism

We define the STP mechanism by providing a method to plan or re-plan, this implicitly defining the dispatch and payment rules. For each $a \in \mathcal{L}$ and $t \in [T]$, denote the welfare gain from an additional driver at (a, t) that is already in the platform as,

$$\Phi_{a,t} \triangleq W(\mathcal{D} \cup \{(1, a, t, T)\}) - W(\mathcal{D}), \quad (27)$$

where $(1, a, t, T)$ represents the type of this driver that stays until the end of the planning horizon.

Definition 12 (Spatio-Temporal pricing mechanism). The *spatio-temporal pricing (STP) mechanism* is a dynamic ridesharing mechanism that always dispatches all available drivers. Given economy $E^{(0)}$ at the beginning of the planning horizon, or economy $E^{(t)}(s_t)$ immediately after a deviation by one or more drivers, the mechanism completes the following planning step:

- *Dispatch rule*: To determine the dispatches (α^*), compute an optimal solution (x, y) to the ILP (4), and dispatch each driver $i \in \mathcal{D}$ to take the path $Z_{i,k}$ for k s.t. $y_{i,k} = 1$, and pick up riders $j \in \mathcal{R}$ s.t. $x_j = 1$,
- *Payment rules*: To determine driver and rider payments (r^* and q^*), for each $(a, b, t) \in \mathcal{T}$, set anonymous trip prices to be $p_{a,b,t} = \Phi_{a,t} - \Phi_{b,t+\delta(a,b)} + c_{a,b,t}$:
 - 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 $i \in \mathcal{D}$, $r_{i,t}^* = \sum_{j \in \mathcal{R}, \tau_j = t} p_{o_j, d_j, t} \mathbb{1}\{\alpha_{i, \tau_j}^* = (o_j, d_j, t, j)\}$.

We now state the main result of the paper.

Theorem 2. *The spatio-temporal pricing mechanism is temporally consistent and subgame-perfect incentive compatible. It is also individually rational for riders and strictly budget balanced for any action profile taken by the drivers, and is welfare optimal, core-selecting, individually rational for drivers and envy-free in subgame-perfect equilibrium from any history onward.*




Figure 5: The Super Bowl example: replanning under the STP mechanism at time 1 after driver 3 deviated from the original plan and stayed in location B until time 1.

We provide the full proof of Theorem 2 in Appendix B.6. We first prove, via a telescoping sum, that the total utility of each driver under the STP mechanism is $\pi_i = \Phi_{D_i}$, the welfare gain from replicating driver i . Setting $u_j = \max\{v_j - p_{o_j, d_j, \tau_j}, 0\}$ for all riders $j \in \mathcal{R}$, we show that (p, π, u) forms an optimal solution to the dual LP (15) by observing (i) (Φ, u) forms an optimal solution to the dual of the corresponding MCF problem (the proof of Lemma 4), and (ii) a correspondence between the optimal solutions of the dual LP (15) and the optimal solutions of the dual of the MCF (Lemma 7 in Appendix B.3). This implies that the plan determined by the STP mechanism starting from any history onward forms a CE, and as a result is IR for riders and drivers, strictly budget balanced, envy-free, and resides in the core.

For incentive alignment, the single-deviation principle for finite horizon extensive-form games [30] means that we only need show that a single deviation from the mechanism's dispatching is not useful. For any driver in the platform, available at location a and time t , her total utility from time t onward, if all drivers follow the dispatches, is equal the welfare gain (at the time when the plan is computed) from adding a driver at location a and time t . We establish that this welfare gain is weakly higher than the welfare gain for the economy (at time $t + 1$) from replicating this driver at any location and time that the driver can deviate and relocate to.

For this, we use the M^\dagger concavity (and more specifically, the local exchange properties) of optimal objectives of MCF problems [27]. In particular, to maximize welfare, there is stronger substitution among drivers at the same location than among drivers at different locations. This shows that declining the mechanism's dispatch to stay/relocate is not useful. We also show that none of (i) exiting earlier than dispatched, (ii) not entering/exiting when asked to, and (iii) entering when dispatched not to, is a useful deviation. From incentive alignment and the CE property, we have the desired properties as stated.

Example 2 (Continued). Whereas the static CE mechanism fails to be welfare-optimal or envy-free for riders after driver 3 deviates from the dispatch and stays in location B until time 1, the plan recomputed under the STP mechanism at time 1 is as illustrated in Figure 5. Driver 3 is re-dispatched to pick up rider 5 and then exit from $(B, 2)$. Instead of picking up rider 8 whose value is 90, driver 2 now picks up rider 6 who was initially assigned to driver 3. If there existed an additional driver at $(C, 1)$, the driver will be dispatched to pick up rider 8, and contribute to a

welfare gain of $v_8 - c_{C,A,1} = 70$. An additional driver at $(A, 3)$ has no effect on welfare, thus the price for the $(C, A, 1)$ trip is updated to $p_{C,A,1} = \Phi_{C,1} - \Phi_{A,3} + c_{C,A,1} = 70 - 0 + 20 = 90$, and the utility of each driver from $(C, 1)$ onward is equal to $\Phi_{C,1} = 70$. Similarly, we have $p_{C,B,1} = \Phi_{C,1} - \Phi_{B,2} + c_{C,B,1} = 70 - (-5) + 10 = 85$ and $p_{B,B,1} = \Phi_{B,1} - \Phi_{B,2} + c_{B,B,1} = -10 - (-5) + 10 = 5$. The utility of driver 3 from $(B, 1)$ onward is $p_{B,B,1} - c_{B,B,1} - \kappa_1 = 5 - 10 - 5 = -10$. The outcome remains envy-free for riders, and welfare optimal from time 1 onward. \square

Under the STP mechanism, replanning 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 3. *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) IR for riders, (iii) budget balanced, and (iv) envy-free for riders and drivers.*

Proof. We show that for the economy in Example 3, as shown in Figure 4, under any mechanism that satisfies conditions (i)-(iv), following the dispatches at all times cannot be a DSE. We start by analyzing what must be the outcome at time 0 under such a mechanism. At time 0, optimal welfare is achieved by dispatching one of the two drivers to go to $(B, 1)$ so that at time 1 she can pick up rider 1, and the other driver to go to $(A, 1)$ to pick up rider 2. Assume w.l.o.g. that at time 0, driver 1 is dispatched to stay in B and driver 2 is dispatched to stay in A .

Now consider the scenario where driver 2 deviated, and took the trip $(A, B, 0)$ at time 0 instead. If driver 1 followed the mechanism's dispatch at time 0, both drivers are at $(B, 1)$ at time 1, and the welfare-optimal outcome is to pick up rider 1. Individual rationality requires that the highest amount of payment we can collect from rider 1 is 8. Budget balance and envy-freeness of drivers then imply that drivers 1 and 2 are each paid at most 4 at time 1. If driver 2 is going to deviate at time 0 and relocate to B , driver 1 may deviate from the mechanism's dispatch and relocate to $(A, 1)$ instead. In this case, at time 1 it is welfare optimal for driver 1 to pick up rider 2. Her payment for the trip $(A, A, 1)$ is at least 5, for otherwise rider 3 envies the outcome of rider 2. This is better than following the mechanism and get utility at most 4. \square

A natural variation of the STP mechanism is to consider the driver-optimal analog, which always computes a driver-optimal CE 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, Ψ_{D_i} , and corresponds to the reasoning of the VCG mechanism. The driver-optimal mechanism is, however, not incentive compatible. See Appendix C.1 for a detailed example and explanation.

4.3 Discussion

Throughout the paper, we assume (S1) that drivers stay until the end of the planning horizon and have no intrinsic preference over locations, and (S2), that riders are impatient. Examples 6 and 7 in Appendix C.2 show that if either assumption is relaxed, the LP relaxation (9) of the planning problem may no longer be integral, and thus, there may not exist anonymous, origin-destination CE prices. This is because the reduction to the minimum cost flow problem fails (see Appendix C.2).

The triangle inequalities in distance and trip costs, i.e. $\delta(a, a') \leq \delta(a, b) + \delta(b, a')$ and $c_{a,a',t} \leq c_{a,b,t} + c_{b,a',t+\delta(a,b)}$ for all $a, a', b \in \mathcal{L}$ and all $t \in [T]$, is not necessary for our results on drivers' incentives. On the rider's side, one concern in practice is that riders may try to break a long trip into a few shorter trips in order to get a lower total price, especially for platforms where riders

may set multiple stops for a single trip.¹⁹ With the triangle inequalities, we recover an arbitrage-proofness: the total price under the STP mechanism for the shorter trips is higher than for the single trip. This is because the shorter trips take a longer total time and incur a higher total cost, and the rider’s total payment, which equals the difference in the welfare gain from having an additional driver at the (origin, starting time) and the (destination, ending time) of the trip plus the total costs, is higher if the ending time is later.

The focus of this paper is to align incentives for drivers in a dynamic environment while maintaining driver flexibility, instead of that of information asymmetry. On the rider side, we show via Example 8 in Appendix C.3 that although the STP mechanism chooses a driver-pessimal CE plan, it may not be a dominant strategy for the riders to truthfully report their values. This is different from the classical, unit-demand assignment problem [35], where the set of CE prices form a lattice—the seller-pessimal CE prices correspond to the buyer-side VCG prices, and is incentive compatible for the buyers. In fact, we show in Theorem 4 that the rider-side VCG payment for a rider is equal to the minimum price for her trip among all CE outcomes, and such trip prices may not form a CE. This implies that no mechanism that computes optimal CE plans and balances budget incentivizes the riders to truthfully report their values. Moreover, we show that no BB, optimal and SPIC mechanism is incentive compatible on the rider’s side.

On the driver’s side, we may consider a scenario where the mechanism does not know about drivers’ entrance information, but asks the drivers to report their times and locations to enter the platform at the beginning of the planning horizon. We show in Appendix C.4 that under the STP mechanism, for a driver with entrance location and time (ℓ_i, τ_i) , there is no incentive for her to report feasible location and time $(\hat{\ell}_i, \hat{\tau}_i)$ where $\hat{\tau}_i \geq \tau_i + \delta(\ell_i, \hat{\ell}_i)$, and then actually enter the platform at $(\hat{\ell}_i, \hat{\tau}_i)$. If drivers are allowed to enter at any arbitrary time and location regardless of their reports, and if the mechanism simply replans without penalizing the drivers that deviated from their reports, then truthfully reporting entrance information is no longer a dominant strategy for drivers. We consider it less important to provide flexibility for drivers to enter the platform in a way that is different from their reports, since there is less uncertainty in driver’s ability to enter (e.g. many drivers start driving from home at a certain time).

5 Simulation Results

In this section, we compare, through numerical simulations, the performance of the STP mechanism against the myopic pricing mechanism, for three stylized scenarios: the end of a sporting event, the morning rush hour, and trips to and from the airport with unbalanced demand.

In addition to social welfare, we consider the *time-efficiency* of drivers, which is defined as the proportion of time where the drivers have a rider in the car, divided by the total time drivers spend on the platform. We also consider the *regret* to drivers for following the straightforward strategy in a non incentive-aligned mechanism: the highest additional amount a driver can gain by strategizing in comparison to following a mechanism’s dispatch, assuming that the rest of the drivers all follow the mechanism’s dispatches at all times. 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, prices are highly unstable, and drivers incur a high regret.

We define the myopic pricing mechanism to use the lowest market clearing prices (which market clearing prices are chosen is unimportant for the results). In addition, since this mechanism need not always dispatch all available drivers, we model any available driver who is not dispatched as

¹⁹<https://www.ridester.com/uber-adds-multiple-stops/>, visited January 8, 2018.




Figure 6: An example to illustrate the end of an event.




Figure 7: Comparison of welfare, driver time efficiency and average driver regret for the of an event.

randomly choosing a location that is within reach, and relocates there if this trip cost is no greater than the cost for exiting immediately (in which case she exits).

5.1 Scenario One: The End of a Sporting Event

We first consider the scenario in Figure 6, modeling the end of a sporting event. There are three locations $\mathcal{L} = \{A, B, C\}$ with unit distances $\delta(a, b) = 1$ for all $a, b \in \mathcal{L}$, and two time periods. Trip costs are 3 per period, and exiting early costs 1 per period: $c_{a,b,t} = 3\delta(a, b)$, and $\kappa_\Delta = \Delta$. In each economy, at time 0, there are 15 and 10 drivers that are already in the platform becoming available at locations C and B .²⁰ 20 riders request trip $(C, B, 0)$ and 10 riders request 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 from $(C, 1)$ to $(B, 2)$. The values of all riders are independently drawn from the exponential distribution with mean 10. As we vary the number of riders $N_{C,B,1}$ requesting the trip $(C, B, 1)$ from 0 to 100, we randomly generate 1,000 economies, and compare the average welfare, driver’s time efficiency and average regret as shown in Figure 7.

Figure 7a shows that the STP mechanism achieves a substantially higher social welfare than the myopic pricing mechanism, especially when there are a large number of drivers taking the trip $(C, B, 1)$. Figure 7b shows that the STP mechanism becomes less time-efficient as the number of $(C, B, 1)$ riders increases, as more of the 15 drivers starting at $(C, 0)$ stay in the same location until time 1. The high time-efficiency achieved by myopic is because of the fact that undispached drivers

²⁰We assume that all drivers are already in the platform, so that the results do not reflect the disadvantage of myopic for not making optimal entrance decisions.

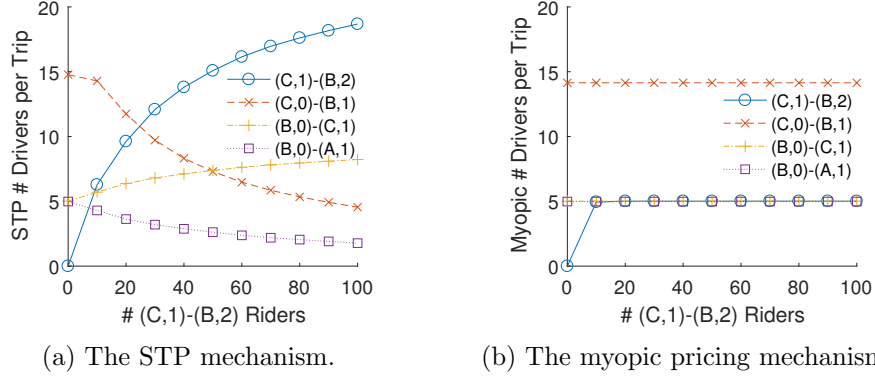


Figure 8: Comparison of the number of drivers per trip for the end of an event.

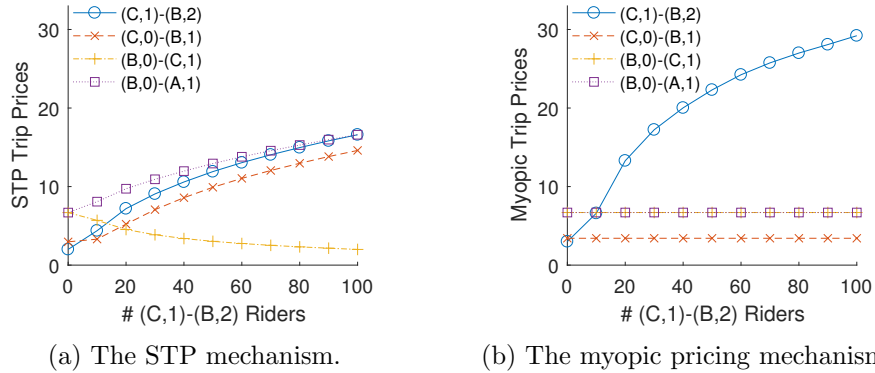


Figure 9: Comparison of trip prices for the end of an event.

decided to exit immediately. The effective use of driver’s time under myopic (total amount of time driver spend driving riders, divided by the total time a driver is willing to work) is in fact around 60 percent. We also see the extent to which the myopic pricing mechanism is not incentive aligned. 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. Figure 7c shows that the average regret of the 25 drivers increases significantly as $N_{C,B,1}$ increases.

The average number of drivers taking each of the four trips of interest under the two mechanisms are shown in Figure 8. As $N_{C,B,1}$ increases, the STP mechanism dispatches more drivers to $(C, 1)$ to pick-up higher-valued riders leaving C , while sending less drivers on trips $(C, B, 0)$ and $(B, A, 0)$. The 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 trip prices are plotted in Figure 9. First of all, prices are temporally “smooth”—trips leaving C at times 0 and 1 have very similar prices. On average, $p_{C,B,1}$ is higher than $p_{C,B,0}$, since drivers taking the $(C, 0)$ - $(B, 1)$ trip can exit at time 1 and incur a smaller total cost. The price for the trip $(B, A, 0)$ is the highest, so that a driver dispatched to A does not envy those dispatched to take the trip $(B, C, 0)$ and then $(C, B, 1)$. In contrast, the price for for the (C, B) trip drastically increases from time 0 to time 1 under the myopic pricing mechanism, and the drivers taking the trip $(B, A, 0)$ envy those that take $(B, C, 0)$ and subsequently $(C, B, 1)$. The “surge” for the trip $(C, B, 1)$ is significantly higher under the myopic pricing mechanism, implying that the platform is providing even less price reliability for the riders.

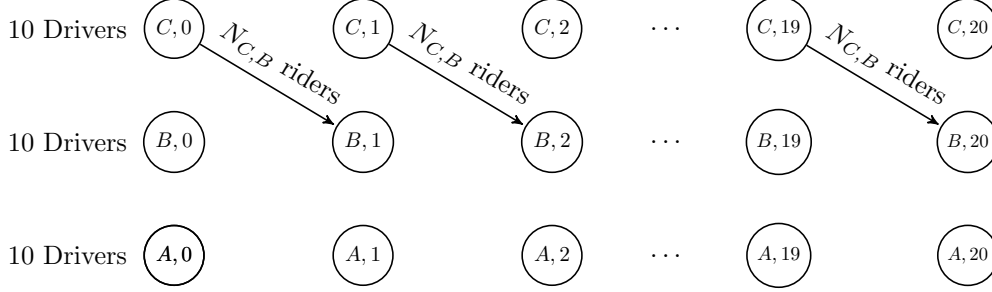


Figure 10: Morning rush hour.

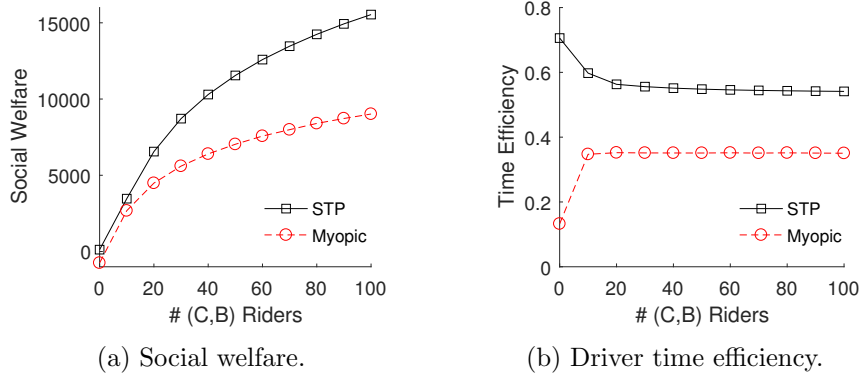


Figure 11: Comparison of social welfare and driver time efficiency for the morning rush hours.

5.2 Scenario Two: The Morning Rush Hour

We now compare the two mechanisms for the economy in Figure 10, modeling the demand pattern of the morning rush hour. There are $T = 20$ time periods and three locations $\mathcal{L} = \{A, B, C\}$ with $\delta(a, b) = 1$ for all $a, b \in \mathcal{L}$. Trip costs are 3 per period, and exiting early costs 1 per period: $c_{a,b,t} = 3\delta(a, b)$, and $\kappa_\Delta = \Delta$. C is a residential area, where there are a number of riders requesting rides to B , the downtown area, in every period. Location A models some other area in the city.

In each economy, at time $t = 0$, there are 10 drivers starting in each of the three locations A , B and C , who all stay until the end of the planning horizon. There are a total of 100 riders with trip origins and destinations independently drawn at random from \mathcal{L} and trip starting times randomly drawn from $[T - 1]$. In addition, in each period there are $N_{C,B}$ commuters traveling from C to B . We assume that the commuters' values for the rides are i.i.d. exponentially distributed with mean 20, whereas the random rides have values exponentially distributed with mean 10.

As we vary the $N_{C,B}$ from 0 to 100, the average social welfare achieved by the two mechanisms for 1,000 randomly generated economies is as shown in Figure 11a. The STP mechanism achieves higher social welfare than the myopic pricing mechanism. Figure 11b shows that the STP mechanism achieves much higher driver time efficiency. The time efficiency of STP mechanism actually decreases as the number of (C, B) riders per period increases above 10, since the mechanism sends more empty cars to C to pick up the higher value riders there.

For the four origin-destination pairs, (C, B) , (B, C) , (B, A) and (A, B) , Figures 12 and 13 plot the average number of drivers getting dispatched to take these trips in each time period (including both trips with a rider, and repositioning without a rider), and the average trip prices. For each economy, the number of drivers for each origin-destination (OD) pair and the trip prices for this OD

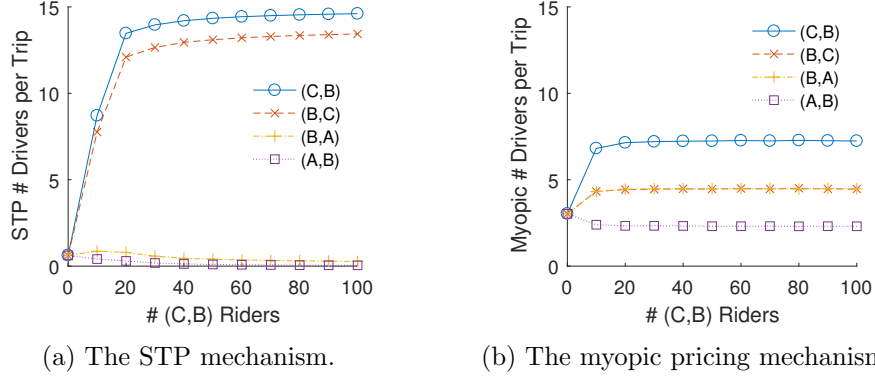


Figure 12: Comparison of the number of drivers per trip for the morning rush hour.

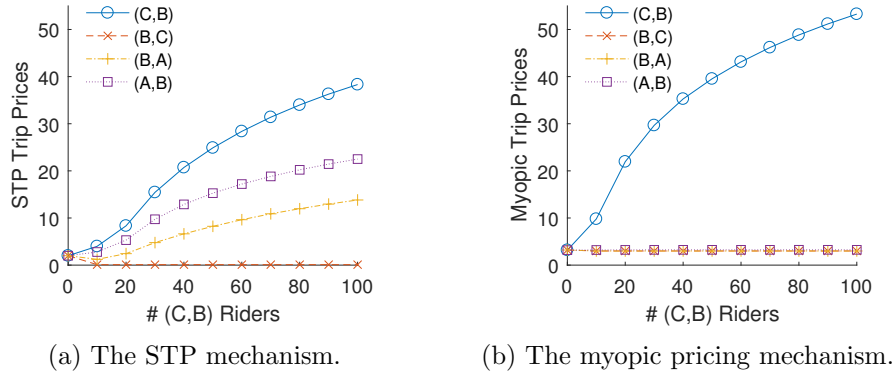


Figure 13: Comparison of trip prices for the morning rush hour.

pair are averaged over the entire planning horizon. Results on the other five trips, (A, A) , (A, C) , (B, B) , (C, A) and (C, C) can be interpreted similarly, therefore are deferred to Appendix E.

Under the STP mechanism, given the large demand for trips from C to B in each time period, there is a large number of drivers taking the trip (C, B) , and also a large number of drivers relocating from B to C in order to pick up future riders from C (see Figure 12a). A small number of drivers are dispatched from B to A due to the lack of future demand at A . Because of the abundance of supply at B that are brought in by the (C, B) trips, very few drivers are sent from A to B . Under the myopic pricing mechanism, the number of drivers dispatched to take each trip, in contrast, does not contribute to the repositioning of drivers. See Figure 12b. There are an equal number of drivers traveling from B to A and C despite the significant difference in the demand conditions at the two destinations. Moreover, a non-trivial number of drivers are traveling from A to B despite the fact that there are already too much of supply at location B .

Regarding the average prices under the STP mechanism plotted in Figure 13a, the morning commute route (C, B) has a higher average price due to the large demand for this trip. The (B, A) trip is less costly than the (A, B) trip since there is plenty of supply of drivers that are brought to B by the (C, B) trips, so that the marginal value of supply at B is low. The (A, B) price is high so that not too many drivers are dispatched from A to B . The (B, C) trips are priced almost always at zero, despite the fact that the trip cost for the drivers is 3, since it is beneficial for the economy for drivers move to C to pick up the commuters. Finally, Figure 13b shows that the (C, B) trip has a much higher average price under the myopic pricing mechanism than the STP mechanism,




Figure 14: Imbalance in trips to and from the airport.

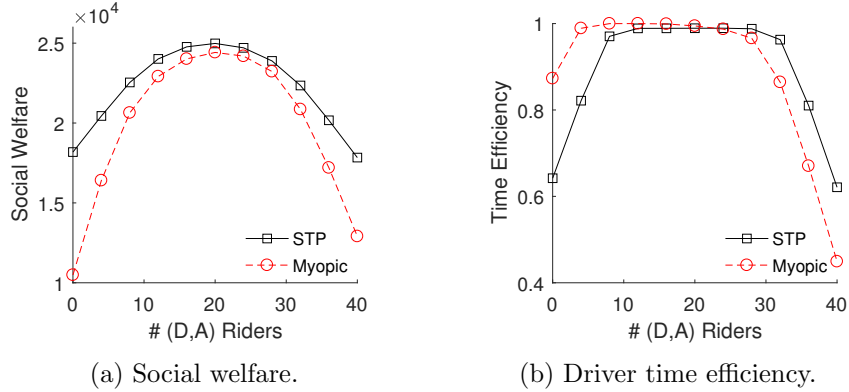


Figure 15: Comparison of social welfare and driver time efficiency for the morning rush hours.


whereas the rest three trips are priced at the trip costs— this is because without optimizing for the supply of drivers at each location, locations A and B almost always have plenty of supply to pick up all riders starting from these locations, whereas there is far from enough drivers to satisfy the large demand at location C .

5.3 Scenario Three: Unbalanced Airport Trips


In this scenario, we consider the imbalance between trips to and from the airport from the downtown area, as illustrated in Figure 14. There are a total of $T = 20$ time periods and two locations $\mathcal{L} = \{A, D\}$, modeling the airport and the downtown area respectively. $\delta(A, A) = \delta(D, D) = 1$, whereas trips in between downtown and the airport are longer: $\delta(A, D) = \delta(D, A) = 2$. Trip costs are 3 per period, and exiting early costs 1 per period: $c_{a,b,t} = 3\delta(a, b)$, and $\kappa_\Delta = \Delta$.

In each economy, there are 20 available drivers at each of A and D at time 0. Within the downtown area, there are 40 riders requesting rides in each period, whereas in between the downtown area and the airport, there are a total of 40 riders heading toward or leaving the airport, which may be unevenly distributed on the two directions. The value of each of the downtown riders are drawn i.i.d from the exponential distribution with mean 10, and the value for each trip to or from the airport is drawn i.i.d from the exponential distribution with mean 40. Since the airport trips are twice as long, we are modeling the scenario where the airport travelers are less price sensitive, and are willing to pay twice as much, in comparison to the downtown riders.

As we vary $N_{D,A}$ from 0 to 40 (thus at the same time varying the number of (A, D) riders from 40 to 0), the average social welfare and driver-time efficiency achieved by the two mechanisms over 1,000 randomly generated economies are as shown in Figure 15. We first observe that the more balanced the trip flow to and from the airport is (i.e. when $N_{D,A}$ is closer to 20), the higher the social welfare and driver time efficiency achieved by STP. This is because when the trip flow is more




(a) The STP mechanism.




(b) The myopic pricing mechanism.

Figure 16: Comparison of driver numbers for the unbalanced trips to and from the airport.



(a) The STP mechanism.



(b) The myopic pricing mechanism.

Figure 17: Comparison of trip prices for the unbalanced trips to and from the airport.

balanced, it is more likely for a driver to pick up riders with high values for the trips both to and from the airport, whereas when the flow is unbalanced, drivers may relocate with an empty car or pick up riders with low values for one of the two directions. The myopic pricing mechanism achieves comparatively good welfare and driver time-efficiency when trip flows are reasonably balanced, however, the performance downgrades quickly as the flow becomes more unbalanced.

Figure 16a shows that regardless of the flow imbalance, the number of drivers doing downtown or airport trips stay reasonably stable under the STP mechanism. However, Figure 16b shows that when $N_{D,A}$ is too small, too many drivers staying in downtown, forgoing the opportunity to pick up the large number of drivers hoping to return to downtown from the airport. Similarly, when $N_{D,A}$ is too large, the myopic pricing mechanism sends too many drivers to the airport, since they have higher per period surplus. As a result, a large number of drivers line up at the airport, and too few rider traveling within downtown are picked up, resulting in much lower efficiency.

The average prices for trips under the two mechanisms are shown in Figure 17. Comparing the two mechanisms, we can see that the price for the (D,D) downtown trip is almost constant, regardless of how unbalanced the airport trip flows are, however, the price for the downtown trip is seriously affected by the conditions at the airport. For the trips to and from the airport, we can see that both mechanisms increase the prices for the direction that is over-demanded, and lowers the price for other direction. The differences are (i) the price surges are much lower under STP than under myopic, providing riders more price stability, and (ii) the price for the under-demanded direction is close to zero under STP, reflecting the need to relocate cars even when there is no

demand, however, myopic insists on setting a prices to at least cover the trip cost.

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 prices are highly unstable and drivers incur a large regret.

Throughout the paper, we assumed complete information and a finite planning horizon. Our on-going work include generalizing the model to settings where there is uncertainty in demand/supply prediction, and where the planning horizon rolls forward as the uncertainty gets resolved over time. One challenge is that of balancing budget— strict balance does not hold for arbitrary driver action profiles, but may still hold in subgame perfect equilibrium. Another challenge is that the class of M^\natural concave functions is not closed under addition, thus the continuation value of a distribution of drivers at the end of the planning horizon (which is the expectation of its continuation value for each realized future demand pattern) may no longer be M^\natural concave, and this affects the integrality of the optimal planning problem and the existence of CE.

Other interesting directions include: (i) empirical analysis on the ridesharing platforms, especially on the loss of welfare due to myopia and the strategic behavior of drivers, (ii) the fairness in regard to effect of welfare-optimal planning on trip prices to and from different neighborhoods with different demand patterns, (iii) truthful elicitation of information in scenarios where drivers poses information on local demand/supply in the near future, and where drivers have preferences with respect to locations or heterogeneity in costs, and (iv) pricing and dispatching with multiple classes of services, and with the existence of competition between platforms.

References

- [1] Philipp Afèche, Zhe Liu, and Costis Maglaras. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Technical report, Columbia Business School. Available at SSRN 3120544, 2018.
- [2] Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. *Network flows: theory, algorithms, and applications*. Prentice hall, 1993.
- [3] Joshua D Angrist, Sydnee Caldwell, and Jonathan V Hall. Uber vs. taxi: A drivers eye view. *NBER Working Paper No. 23891*, 2017.
- [4] Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.
- [5] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 517–517. ACM, 2017.

- [6] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 639–639. ACM, 2015.
- [7] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.
- [8] Dimitri P Bertsekas. The auction algorithm for assignment and other network flow problems: A tutorial. *Interfaces*, 20(4):133–149, 1990.
- [9] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.
- [10] Kostas Bimpikis, Ozan Candogan, and Saban Daniela. Spatial Pricing in Ride-Sharing Networks. *Working paper, available at SSRN 2868080*, 2016.
- [11] Patrick Bolton and Mathias Dewatripont. *Contract theory*. MIT press, 2005.
- [12] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 241–242. ACM, 2017.
- [13] Ruggiero Cavallo, David C Parkes, and Satinder Singh. Efficient mechanisms with dynamic populations and dynamic types. Technical report, Harvard University, 2009.
- [14] Harshal A Chaudhari, John W Byers, and Evimaria Terzi. Putting data in the driver’s seat: Optimizing earnings for on-demand ride-hailing. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 90–98. ACM, 2018.
- [15] M Keith Chen, Judith A Chevalier, Peter E Rossi, and Emily Oehlsen. The value of flexible work: Evidence from uber drivers. *NBER Working Paper No. 23296*, 2017.
- [16] M Keith Chen and Michael Sheldon. 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.
- [17] Scott H Clearwater. *Market-based control: A paradigm for distributed resource allocation*. World Scientific, 1996.
- [18] Peter Cohen, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe. Using big data to estimate consumer surplus: The case of uber. *NBER Working Paper No. 22627*, 2016.
- [19] Cody Cook, Rebecca Diamond, Jonathan Hall, John List, Paul Oyer, et al. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. Technical report, Stanford University, 2018.
- [20] Judd Cramer and Alan B Krueger. Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–82, 2016.
- [21] Jonathan Hall, Cory Kendrick, and Chris Nosko. The effects of uber’s surge pricing: A case study. Technical report, The University of Chicago Booth School of Business, 2015.
- [22] Jonathan V Hall, John J Horton, and Daniel T Knoepfle. Labor market equilibration: Evidence from uber. Technical report, New York University Stern School of Business, 2017.

- [23] Jonathan V Hall and Alan B Krueger. An analysis of the labor market for ubers driver-partners in the united states. *NBER Working Paper No. 22843*, 2016.
- [24] John William Hatfield, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp. Stability and competitive equilibrium in trading networks. *Journal of Political Economy*, 121(5):966–1005, 2013.
- [25] John William Hatfield, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp. Chain stability in trading networks. Technical report, 2015.
- [26] Jonathan Levin. Relational incentive contracts. *The American Economic Review*, 93(3):835, 2003.
- [27] Kazuo Murota. *Discrete convex analysis*. SIAM, 2003.
- [28] Kazuo Murota. Discrete convex analysis: A tool for economics and game theory. *Journal of Mechanism and Institution Design*, 1(1):151–273, 2016.
- [29] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- [30] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.
- [31] Michael Ostrovsky et al. Stability in supply chain networks. *American Economic Review*, 98(3):897, 2008.
- [32] David C Parkes and Lyle H Ungar. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI'00)*, pages 74–81. AAAI, 2000.
- [33] Lisa Rayle, Susan Shaheen, Nelson Chan, Danielle Dai, and Robert Cervero. App-based, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in san francisco. Technical report, University of California Transportation Center (UCTC), 2014.
- [34] Bernard Salanié. *The economics of contracts: a primer*. MIT press, 2005.
- [35] Lloyd S Shapley and Martin Shubik. The assignment game i: The core. *International Journal of game theory*, 1(1):111–130, 1971.
- [36] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.
- [37] Laurence A Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, Incorporated, Somerset, 1999.

Appendix

Appendix A provides a continuous-time interpretation of the discrete time model that we adopted. Appendix B includes proofs that are omitted from the body of the paper. Appendix C provides examples and discussions on the driver-optimal mechanism, integrality of the LP relaxation and existence of CE, incentives of riders, entrance as drivers' private information, and the naive recomputation of optimal CE plans. Appendix D discusses the relationship with the literature on trading networks and the dynamic VCG mechanism, and why they do not solve the ridesharing problem. Finally, additional simulation results are presented in Appendix E.

A Continuous Time Interpretation

Under the discrete time model introduced in Section 2, trips within the same location takes $\delta(a, a) = 1$ unit of time for all $a \in \mathcal{L}$, however, we also assume that a driver can drop-off a rider and pick-up a new rider in the same location at the same time point.




Figure 18: Time-line for a within-location trip $A \rightarrow A$ which takes $\delta(A, A) = 1$ period of time, and a between-location trip $A \rightarrow B$ which takes $\delta(A, B) = 2$ periods of time.

Figure 18 illustrates the continuous-time interpretation of this discrete-time model. There are two trips: an A to A trip at time 0 which takes $\delta(A, A) = 1$ unit of time, and an A to B trip at time 1, which takes $\delta(A, B) = 2$ units of time and ends at time $t = 3$. The time after the drop-off of the first rider at A and the time for the second pick-up at A is the time the driver takes to travel within location A to pick up the second rider. In the discrete time model, both the drop-off and the pick-up are both considered to happen at time $t = 1$.

B Proofs

B.1 Proof of Lemma 1

Lemma 1. Given any CE plan (x, \tilde{z}, p) , the plan with anonymous prices (x, \tilde{z}, p^+) also forms a CE, and has the same driver and rider payments and utilities as those under (x, \tilde{z}, p) .

Proof. Assume that there exists $(a, b, t) \in \mathcal{T}$ s.t. $\exists j \in \mathcal{R}$ s.t. $(o_j, d_j, \tau_j) = (a, b, t)$ and $p_{a,b,t} < 0$. Rider best response implies $x_j = 1$, thus there exists $i \in \mathcal{D}$ s.t. $(o_j, d_j, \tau_j, j) \in \tilde{z}_i$, and this is paid $p_{a,b,t} < 0$ at time t . This violates driver best response, since keeping the rest of the action path unchanged, but choosing not to be paid for this trip, the driver would get a higher total payment.

This implies that $p_{o_j, d_j, \tau_j} = p_{o_j, d_j, \tau_j}^+$ holds for all $j \in \mathcal{R}$, therefore given the dispatching (x, \tilde{z}) , rider best response under prices p implies rider best response under prices p^+ . Denote the total payment to each driver and the utility of each driver given (x, \tilde{z}, p) as r_i^+ and π_i^+ . We know that for each driver $i \in \mathcal{D}$, $r_i^+ = \sum_{j \in \mathcal{R}} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} p_{o_j, d_j, \tau_j}^+ = \sum_{j \in \mathcal{R}} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} p_{o_j, d_j, \tau_j} = r_i$, thus $\pi_i^+ = \pi_i$. This implies driver best response: $\pi_i^+ = \pi_i = \max_{k=0, \dots, |Z_i|} \{\sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}, 0\} - \lambda_{i,k}\} = \max_{k=0, \dots, |Z_i|} \{\sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}^+, 0\} - \lambda_{i,k}\}$,

This completes the proof that (x, \tilde{z}, p) also forms a CE, and that the driver and rider payments and utilities under the two plans are identical. \square

B.2 Proof of Lemma 2

B.2.1 Minimum Cost Flow Problems

We first provide the formulation of the minimum cost flow (MCF) problem, and the reduction of an optimal dispatching problem to an MCF problem where drivers flow through a network with nodes corresponding to the initial states of drivers and the (location, time) pairs, and the edge costs equal to the trip costs for drivers minus the rider values.

Let $G = (\mathcal{N}, \mathcal{E})$ be a directed graph with a node set \mathcal{N} and an edge set \mathcal{E} . Let $\underline{\zeta} : \mathcal{E} \rightarrow \mathbb{Z} \cup \{-\infty\}$ be the lower capacity function, $\bar{\zeta} : \mathcal{E} \rightarrow \mathbb{Z} \cup \{+\infty\}$ be the upper capacity function, and let $\gamma : \mathcal{E} \rightarrow \mathbb{R}$ be the cost function. For each edge $e \in \mathcal{E}$, denote $\partial^+ e \in \mathcal{N}$ as the initial (tail) node of e and $\partial^- e \in \mathcal{N}$ as the terminal (head) node of e . That is, $\partial^+ e = n_1$ and $\partial^- e = n_2$ for the edge $e = (n_1, n_2)$.

A feasible flow f is a function $f : \mathcal{E} \rightarrow \mathbb{R}$ such that $\underline{\zeta}(e) \leq f(e) \leq \bar{\zeta}(e)$ for each $e \in \mathcal{E}$. Its boundary $\partial f : \mathcal{N} \rightarrow \mathbb{R}$ is defined as

$$\partial f(n) = \sum \{f(e) \mid e \in \mathcal{E}, \partial^+ e = n\} - \sum \{f(e) \mid e \in \mathcal{E}, \partial^- e = n\}.$$

A node n for which $\partial f(n) > 0$ is a source of the flow, and a node n is a sink if $\partial f(n) < 0$. Let ξ be a vector in $\mathbb{R}^{|\mathcal{N}|}$, the minimum cost for any flow with boundary condition ξ is:

$$\omega(\xi) = \inf_f \left\{ \sum_{e \in \mathcal{E}} c(e) f(e) \mid f : \text{feasible flow with } \partial f = \xi \right\}. \quad (28)$$

B.2.2 Reducing Optimal Dispatching to MCF

Given an instance of the optimal dispatching problem with planning horizon T , locations \mathcal{L} , distances δ , costs $\{c_{a,b,t}\}_{(a,b,t) \in \mathcal{T}}$ and $\{\kappa_\Delta\}_{\Delta=1,\dots,T}$, riders \mathcal{R} and drivers \mathcal{D} , we construct a corresponding MCF problem. Let $G = (\mathcal{N}, \mathcal{E})$ be the graph, where the nodes \mathcal{N} consists of (location, time) pairs, the nodes $\{D_i\}_{i \in \mathcal{D}}$ modeling the initial states of drivers, and an additional ‘‘sink’’ node S representing the end of time:

$$\mathcal{N} = \{(a, t) \mid a \in \mathcal{L}, t \in [T]\} \cup \{D_i \mid i \in \mathcal{D}\} \cup \{S\}.$$

The set of edges $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3 \cup \mathcal{E}_4$ consists of the following parts:

- $\mathcal{E}_1 = \{R_j \mid j \in \mathcal{R}\}$ corresponds to rider trips, where the edge $R_j = ((o_j, \tau_j), (d_j, \tau_j + \delta(o_j, d_j)))$ corresponds to the trip requested by rider j and has minimum capacity $\underline{\zeta}(R_j) = 0$, maximum capacity $\bar{\zeta}(R_j) = 1$ and cost $\gamma(R_j) = -v_j + c_{o_j, d_j, \tau_j}$. Intuitively, if a unit of driver flows through the edge corresponding to rider j (i.e. rider j is picked up by a driver), we incur a cost of c_{o_j, d_j, τ_j} , and gain value v_j .
- \mathcal{E}_2 consists of edges that are feasible relocating trips without riders:

$$\mathcal{E}_2 = \{((a, t), (b, t + \delta(a, b))) \mid (a, b, t) \in \mathcal{T}\}.$$

Recall that $\mathcal{T} = \{(a, b, t) \mid a, b \in \mathcal{L}, t + \delta(a, b) \leq T\}$ denotes the set of feasible trips within the planning horizon. There is no upper bound on capacities of these edges: $\forall e \in \mathcal{E}_2, \underline{\zeta}(e) = 0, \bar{\zeta}(e) = +\infty$. The edge costs are $\gamma(e) = c_{a,b,t}$, if $e = ((a, t), (b, t + \delta(a, b)))$.

- \mathcal{E}_3 consists of edges that connect all nodes (a, t) to the sink S , representing the exit of a driver from location a and time t :

$$\mathcal{E}_3 = \{((a, t), S) \mid a \in \mathcal{L}, t \in [T]\}.$$

Similar to \mathcal{E}_2 , there is no upper capacity constraint: $\forall e \in \mathcal{E}_3, \underline{\zeta}(e) = 0, \bar{\zeta}(e) = +\infty$. A driver exiting at time t incurs an early exiting opportunity cost of κ_{T-t} , therefore $\gamma(e) = \kappa_{T-t}$ for all $e = ((a, t), S) \in \mathcal{E}_3$.

- \mathcal{E}_4 consists of two sets of edges. The first set of edges allow each driver to enter the platform at their entrance location and time (ℓ_i, τ_i) , and the second set of edges allow drivers with $\beta_i = 0$ to not enter the platform at all:

$$\mathcal{E}_4 = \{(D_i, (\ell_i, \tau_i)) \mid i \in \mathcal{D}\} \cup \{(D_i, S) \mid i \in \mathcal{D} \text{ s.t. } \beta_i = 0\}.$$

There is no upper capacity constraint, and no additional cost for these edges: $\forall e \in \mathcal{E}_4, \underline{\zeta}(e) = 0, \bar{\zeta}(e) = +\infty, \text{ and } \gamma(e) = 0$.

The boundary condition of the MCF problem is given by:

$$\begin{aligned} \xi_{D_i} &= 1, \quad \forall i \in \mathcal{D}, \\ \xi_n &= 0, \quad \text{if } n = (a, t) \text{ for some } a \in \mathcal{L} \text{ and } t \in [T], \\ \xi_S &= -m. \end{aligned}$$

Flow LP Given this construction, there are non-zero edge costs and upper flow capacity constraints only for edges in \mathcal{E}_1 . The minimum cost flow problem (28) can therefore be simplified and rewritten in the following form:

$$\begin{aligned} \min_f \quad & \sum_{j \in \mathcal{R}} (c_{o_j, d_j, \tau_j} - v_j) f(R_j) + \sum_{(a, b, t) \in \mathcal{T}} c_{a, b, t} f(((a, t), (b, t + \delta(a, b)))) \\ & + \sum_{a \in \mathcal{L}, t \in [T]} \kappa_{T-t} f(((a, t), S)) \end{aligned} \quad (29)$$

$$\text{s.t.} \quad \sum_{e \in \mathcal{E}, \partial^+ e = (a, t)} f(e) - \sum_{e \in \mathcal{E}, \partial^- e = (a, t)} f(e) = 0, \quad \forall a \in \mathcal{L}, \forall t \in [T] \quad (30)$$

$$\sum_{e \in \mathcal{E}, \partial^+ e = D_i} f(e) = 1, \quad \forall i \in \mathcal{D} \quad (31)$$

$$f(R_j) \leq 1, \quad \forall j \in \mathcal{R} \quad (32)$$

$$f(e) \geq 0, \quad \forall e \in \mathcal{E} \quad (33)$$

Note that given (30) and (31), the flow balance constraint at the sink, $\sum_{e \in \mathcal{E}, \partial^- e = S} f(e) = m$, is redundant, and therefore omitted from the above formulation in order to achieve better interpretability of the dual variables. Observing that minimizing the negation of the total value of riders that are picked up is equivalent to maximizing the total value of riders that are picked up,

we can rewrite (29) in the following form:

$$\begin{aligned} \max_f \quad & \sum_{j \in \mathcal{R}} (v_j - c_{o_j, d_j, \tau_j}) f(R_j) - \sum_{(a, b, t) \in \mathcal{T}} c_{a, b, t} f((a, t), (b, t + \delta(a, b))) \\ & - \sum_{a \in \mathcal{L}, t \in [T]} \kappa_{T-t} f((a, t), S) \end{aligned} \quad (34)$$

$$\text{s.t.} \quad \sum_{e \in \mathcal{E}, \partial^+ e = (a, t)} f(e) - \sum_{e \in \mathcal{E}, \partial^- e = (a, t)} f(e) = 0, \quad \forall a \in \mathcal{L}, \forall t \in [T] \quad (35)$$

$$\sum_{e \in \mathcal{E}, \partial^+ e = D_i} f(e) = 1, \quad \forall i \in \mathcal{D} \quad (36)$$

$$f(R_j) \leq 1, \quad \forall j \in \mathcal{R} \quad (37)$$

$$f(e) \geq 0, \quad \forall e \in \mathcal{E} \quad (38)$$

We refer to (34) as the flow LP.

B.2.3 Proof of Lemma 2

Lemma 2 (Integrality). There exists an integer optimal solution to the linear program (9).

Proof. It is known that the MCF problems with certain structure have integral optimal solutions [27]. The flow LP (34) is integral since (I) the flow balance constraints (35) and (36) can be written in matrix form $Ff = \xi$ where F is total unimodular and ξ has only integer entries and (II) the edge capacity constraints (37) are all integral. See Section III.1.2 in [37] for details on total unimodularity and the integrality of polyhedron.

To prove the integrality of the original LP (9), we show that

- (i) for each feasible solution to the LP (9), there exists a feasible solution to the flow LP (34) with the same objective, and
- (ii) for each integral feasible solution to the flow LP (34), there a corresponding integral feasible solution to the LP (9) with the same objective.

The integrality of MCF then implies that there exists an integral optimal solution of (9), since the optimal objective of the LP (9) cannot exceed the optimal objective of the MCF, which is achieved at some integral feasible solution of the MCF, and therefore also at some integral feasible solution of (9). We now prove (i) and (ii).

Part (i). Let (x, y) be a feasible solution to the LP (9). A solution f to the flow LP (34) can be constructed as follows:

- For each $j \in \mathcal{R}$, let $f(R_j) = x_j$. We know that $0 \leq f(R_j) \leq 1$ for all $j \in \mathcal{R}$.
- For each $e = ((a, t), (b, t + \delta(a, b))) \in \mathcal{E}_2$ corresponding to the relocation trip $(a, b, t) \in \mathcal{T}$, let $f(e) = \sum_{i \in \mathcal{D}} \sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k} \mathbb{1}\{(a, b, t) \in \mathcal{Z}_{i,k}\} - \sum_{j \in \mathcal{R}} x_j \mathbb{1}\{(o_j, d_j, \tau_j) = (a, b, t)\}$. Constraint (10) guarantees that $f(e) \geq 0$.
- For each driver $i \in \mathcal{D}$ s.t. $\beta_i = 1$, let $f((D_i, (\ell_i, \tau_i))) = \sum_{k=0}^{|\mathcal{Z}_i|} y_{i,k}$. For each $i \in \mathcal{D}$ s.t. $\beta_i = 0$, let $f((D_i, (\ell_i, \tau_i))) = \sum_{k=1}^{|\mathcal{Z}_i|} y_{i,k}$, and $f((D_i, S)) = y_{i,0}$.

- For each $e = ((a, T), S) \in \mathcal{E}_3$, let $f(e) = \sum_{e' \in \mathcal{E}, \partial^- e' = (a, T)} f(e')$ to balance the flow in and out of (a, T) . This is the total number of drivers that existed the platform from (a, t) .

The edge capacity constraints (37) and (38) are satisfied by construction. Given constraint (11) and the fact that each $Z_{i,k}$ is a feasible path, constraints (35) and (36) are satisfied. Moreover, it is obvious that the objective of the two linear programs coincide, thus f is a feasible solution to the flow LP (9) with the same objective.

Part (ii). Given a feasible, integral solution f to the flow LP (34), we construct an integral feasible solution to the original LP. For the riders, let $x_j = f(R_j)$ for all $j \in \mathcal{R}$. For the drivers, from the standard flow decomposition arguments [9], the m units of flow in f that all converge in S can be decomposed into m paths of single units of flow, that correspond to each driver's feasible path taken over the entire planning horizon. This gives us a feasible solution to the original LP, and it is easy to see that the objectives are the same. This completes the proof of the lemma. \square

The reduction to MCF can also be used to solve the original LP efficiently. In the optimal dispatching problem, the number of feasible paths for each driver is exponential in $|\mathcal{L}|$ and T , thus there are exponentially many decision variables in the LP (9). The numbers of decision variables and constraints of the flow LP are, in contrast, polynomial in $|\mathcal{R}|$, $|\mathcal{L}|$ and T , and there are efficient algorithms for solving network flow problems (see [2]).

B.3 Proof of Lemma 3

Before proving the lemma, we first state the complementary slackness (CS) conditions [9]. Given a feasible solution (x, y) to the primal LP (9), and a feasible solution (p, π, u) to the dual LP (15), both solutions are optimal if and only if the following conditions hold:

(CS-1) for all $j \in \mathcal{R}$, $x_j > 0 \Rightarrow u_j = v_j - p_{o_j, d_j, \tau_j}$,

(CS-2) for all $j \in \mathcal{R}$, $u_j > 0 \Rightarrow x_j = 1$,

(CS-3) for all $i \in \mathcal{D}$ and all $k = 1, \dots, |Z_i|$, $y_{i,k} > 0 \Rightarrow \pi_i = \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} - \lambda_{i,k}$,

(CS-4) for all $(a, b, t) \in \mathcal{T}$,

$$p_{a,b,t} > 0 \Rightarrow \sum_{j \in \mathcal{R}} x_j \mathbb{1}\{(o_j, d_j, \tau_j) = (a, b, t)\} = \sum_{i \in \mathcal{D}} \sum_{k=0}^{|Z_i|} y_{i,k} \mathbb{1}\{(a, b, t) \in Z_{i,k}\}.$$

We also provide this following lemma, showing that given any CE outcome, any trip with excessive driver supply has a non-positive price.

Lemma 6. Given any plan with anonymous trip prices (x, \tilde{z}, p) that forms a CE, for any $(a, b, t) \in \mathcal{T}$, if there exists a driver $i \in \mathcal{D}$ s.t. $(a, b, t) \in \tilde{z}_i$, then $p_{a,b,t} \leq 0$.

The proof is straightforward. If there exists any trip (a, b, t) with a positive price, and a driver that takes this trip as relocation without a rider, the driver is not getting paid for this trip. This violates driver best response, since getting paid for this trip improves total payment to the driver.

We are now ready to prove Lemma 3.

Lemma 3 (Welfare Theorem). A dispatching (x, \tilde{z}) is welfare-optimal if and only if there exists anonymous trip prices p s.t. the plan (x, \tilde{z}, p) forms a competitive equilibrium. Such optimal CE plans always exist and are efficient to compute. Moreover, these plans are strictly budget balanced, and are individually rational and envy-free for both riders and drivers.

Proof. Given any welfare optimal dispatching (x, \tilde{z}) , we can construct an integral optimal solution (x, y) to the LP (9), where for all $i \in \mathcal{D}$, $y_{i,k} = 1$ if the movement of driver i in space and time according to the action path \tilde{z}_i is consistent with the path $Z_{i,k}$. Given any CE plan (x, \tilde{z}, p) , Lemma 1 implies that (x, \tilde{z}, p^+) also forms a CE, where p^+ is defined s.t. $p_{a,b,t}^+ = \max\{p_{a,b,t}, 0\}$. We prove the lemma in two steps:

- Step 1. Given any optimal dispatching (x, \tilde{z}) , and any optimal solution (p, π, u) to the dual LP (15), the CS conditions imply that π and u can be interpreted as drivers' and rider's utilities, if the anonymous trip prices is given by p . Optimal dual conditions guarantee driver and rider best responses, thus the plan (x, \tilde{z}, p) forms a CE.
- Step 2. Given a CE plan (x, \tilde{z}, p) , let (x, y) be the corresponding primal solution, and construct a dual solution (p^+, π, u) , where π and u are the corresponding driver and rider utilities. CS conditions are satisfied between (x, y) and (p^+, π, u) , thus (x, \tilde{z}, p) is welfare optimal.

This proves the correspondence between CE and optimal plans and the existence of CE. Lemma 7 in Appendix B.4 implies that the CE prices can be efficiently computed from solving the dual of the flow LP. Regarding the properties: rider IR and envy-freeness is guaranteed by anonymous trip prices and CE; strict budget balance is guaranteed by the definition of anonymous trip prices; for driver envy-freeness, given any two drivers with the same initial (starting location/time, whether the driver had entered the platform or not), they have the same set of feasible paths, therefore both get the same highest total utility among those paths.

We now prove the above two steps.

Step 1: Optimal primal and dual solutions \Rightarrow CE.

Given an optimal dispatch (x, \tilde{z}, p) , let (x, y) be the corresponding optimal integral solution to the primal LP (9), and let (p, π, u) be any optimal solution to the dual LP (15). We first show that if the anonymous trip prices are given by p , then the dual variables π and u correspond to drivers' and riders' utilities, respectively:

1. $x_j > 0 \Rightarrow u_j = v_j - p_{o_j, d_j, \tau_j}$ from (CS-1), thus for riders that are picked up, u_j represent the utilities of the rider, which is her value minus the price for her trip.
2. $u_j > 0 \Rightarrow x_j = 1$ from (CS-2), i.e. in order for a rider to have positive utility, the rider must be picked up. This implies that $x_j = 0 \Rightarrow u_j = 0$, i.e. riders that are not picked up have zero utilities. Combining 1. and 2., we know that u_j correspond to the the rider's utilities.
3. $y_{i,k} > 0 \Rightarrow \pi_i = \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} - \lambda_{i,k}$ from (CS-3), i.e. if driver i takes her k^{th} feasible path, then π_i equals the sum of the prices of each trip covered by this path minus the total cost of this path. $\sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t}$ is equal to the driver's total payment since (I) for any rider trip, i.e. $(a, b, t) \in Z_{i,k}$ s.t. $\exists j \in \mathcal{R}$ s.t. $(a, b, t, j) \in \tilde{z}_i$, the driver is paid $p_{a,b,t}$, and (II), for (a, b, t) where the driver relocates without a rider, (CS-4) implies that $p_{a,b,t} = 0$, therefore $p_{a,b,t}$ is also the driver's payment. As a result, π_i coincides with the total utility of driver i .

We now show that this outcome forms a CE. For rider best response: constraint $u_j \geq 0$ guarantees IR for riders, thus riders that are picked up can afford the price; $v_j - p_{o_j, d_j, \tau_j} > 0 \Rightarrow u_j > 0 \Rightarrow x_j = 1$ implies that all riders that strictly prefer getting pickup up are dispatched to some driver. For driver best response, the dual constraints (16) and (18) guarantee that for all $i \in \mathcal{D}$, $\pi_i \geq \max_{k=0, \dots, |Z_i|} \left\{ \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} - \lambda_{i,k} \right\} = \max_{k=0, \dots, |Z_i|} \left\{ \sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}, 0\} - \lambda_{i,k} \right\}$.

Step 2: CE \Rightarrow Optimal primal and dual solutions.

Let (x, \tilde{z}, p) be a CE plan with anonymous trip prices, and let (u, π) be riders' and drivers' utilities under this plan. The plan being feasible implies that corresponding (x, y) is a feasible and integral primal solution. We show that (p^+, π, u) is a feasible solution to the dual LP (15): (18) holds by definition of p^+ ; Lemma 1 and rider best response implies (17) and (19); driver best response implies (16).

We now prove that (x, y) and (p^+, π, u) must both be optimal, by checking the CS conditions:

1. For (CS-1): given any $j \in \mathcal{R}$ s.t. $x_j > 0$, we know that the rider is picked up, pays p_{o_j, d_j, τ_j} and gets utility $u_j = v_j - p_{o_j, d_j, \tau_j}$.
2. For (CS-2): for any rider $j \in \mathcal{R}$ that gets utility $u_j > 0$, she must be picked up since otherwise her utility would be zero, therefore $x_j = 1$.
3. For (CS-3), for each driver $i \in \mathcal{D}$, $y_{i,k} > 0$ implies that driver i takes her k^{th} feasible path. For each trip (a, b, t) on the path, the driver gets paid $p_{a,b,t}$ regardless of whether she picks up a rider (see Lemma 6), thus her total utility is the sum of the prices minus her cost, thus $\pi_i = \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} - \lambda_{i,k}$.
4. For (CS-4), $p_{a,b,t} > 0 \Rightarrow \sum_{j \in \mathcal{R}} x_j \mathbb{1}\{(o_j, d_j, \tau_j) = (a, b, t)\} = \sum_{i \in \mathcal{D}} \sum_{k=0}^{|Z_i|} y_{i,k} \mathbb{1}\{(a, b, t) \in Z_{i,k}\}$ is implied by Lemma 6— otherwise, there is excessive supply for trip (a, b, t) , implying $p_{a,b,t} = 0$.

This completes the proof of the lemma. \square

B.4 Proof of Lemma 4

B.4.1 Dual of the Flow LP

Before proving Lemma 4, we first discuss the dual of the flow LP, and its correspondence to the dual LP (15). Let $\varphi_{a,t}$, φ_{D_i} , and μ_j be the dual variables corresponding to constraints (35), (36) and (37), respectively. The dual LP of (34) can be written as:

$$\min \sum_{i \in \mathcal{D}} \varphi_{D_i} + \sum_{j \in \mathcal{R}} \mu_j \quad (39)$$

$$\text{s.t. } \varphi_{o_j, \tau_j} - \varphi_{d_j, \tau_j + \delta(o_j, d_j)} + \mu_j \geq v_j - c_{o_j, d_j, \tau_j}, \quad \forall j \in \mathcal{R} \quad (40)$$

$$\varphi_{a,t} - \varphi_{b,t + \delta(a,b)} \geq -c_{a,b,t}, \quad \forall (a, b, t) \in \mathcal{T} \quad (41)$$

$$\varphi_{a,t} \geq -\kappa_{T-t}, \quad \forall a \in \mathcal{L}, \forall t \in [T] \quad (42)$$

$$\varphi_{D_i} \geq \varphi_{\ell_i, \tau_i}, \quad \forall i \in \mathcal{D} \quad (43)$$

$$\varphi_{D_i} \geq 0, \quad \forall i \in \mathcal{D} \text{ s.t. } \beta_i = 0, \quad (44)$$

$$\mu_j \geq 0, \quad \forall j \in \mathcal{R} \quad (45)$$

Given a solution (φ, μ) of (39), the φ variables corresponding to the flow balance constraints are usually referred to as the *potential* of the nodes, and we call φ an *optimal potential* of the MCF problem if there exist $\mu \in \mathbb{R}^{|\mathcal{R}|}$ s.t. (φ, μ) is an optimal solution of (39). The potential for each node can be interpreted as how “useful” it is to have an additional unit of flow originating from this node, and μ_j for each j can be interpreted as the utility of the rider j .

Complementary Slackness Conditions Given a feasible solution f to the flow primal LP (34) and a feasible solution (φ, μ) to the flow dual LP (39), both solutions are optimal if and only if the following complementary slackness conditions [8] are satisfied.

- (CS_F-1) for all $j \in \mathcal{R}$, $f(R_j) > 0 \Rightarrow \varphi_{o_j, \tau_j} - \varphi_{d_j, \tau_j + \delta(o_j, d_j)} + \mu_j = v_j - c_{o_j, d_j, \tau_j}$,
- (CS_F-2) for all $j \in \mathcal{R}$, $\mu_j > 0 \Rightarrow f(R_j) = 1$,
- (CS_F-3) for all $(a, b, t) \in \mathcal{T}$, $f(((a, t), (b, t + \delta(a, b)))) > 0 \Rightarrow \varphi_{a, t} - \varphi_{b, t + \delta(a, b)} = -c_{a, b, t}$,
- (CS_F-4) for all $a \in \mathcal{L}$ and $t \in [T]$, $f(((a, t), S)) > 0 \Rightarrow \varphi_{a, t} = -\kappa_{T-t}$.
- (CS_F-5) for all $i \in \mathcal{D}$, $f((D_i, (\ell_i, \tau_i))) > 0 \Rightarrow \varphi_{D_i} = \varphi_{\ell_i, \tau_i}$.
- (CS_F-6) for all $i \in \mathcal{D}$ s.t. $\beta_i = 0$, $f((D_i, S)) > 0 \Rightarrow \varphi_{D_i} = 0$.

The following lemma establishes a one-to-one correspondence between the π_i variables in optimal solutions to the dual LP (15), and the φ_{D_i} variables in optimal solutions to (39).

Lemma 7. Given any ridesharing problem satisfying assumptions (S1) and (S2), for any optimal solution (p, π, u) to the dual LP (15), there exists an optimal solution (φ, μ) to the dual of the flow LP (39) such that $\varphi_{D_i} = \pi_i$ for all $i \in \mathcal{D}$, $u_j = \mu_j$ for all $j \in \mathcal{R}$, and vice versa.

Proof. We prove the following two directions by construction:

- (i) Given an optimal solution (p, π, u) to (15), there exists an optimal solution (φ, μ) to (39) s.t. $\varphi_{D_i} = \pi_i$ for all $i \in \mathcal{D}$ and that $\mu_j = u_j$ for all $j \in \mathcal{R}$.
- (ii) Given an optimal solution (φ, μ) to (39), there exists an optimal solution (p, π, u) to (15) s.t. $\pi_i = \varphi_{D_i}$ for all $i \in \mathcal{D}$ and that $u_j = \mu_j$ for all $j \in \mathcal{R}$.

Part (i). Given any optimal solution (p, π, u) to (15), we construct a solution (φ, μ) to (39) from the prices p as follows, where $\varphi_{a, t}$ represents the highest continuation payoff for any driver from location a and time t onward, φ_{D_i} represents the highest achievable payoff of driver i , and μ_j represents the highest achievable utility of rider j :

- For all $a \in \mathcal{L}$, let $\varphi_{a, T} = -\kappa_{T-T} = 0$.
- For all $a \in \mathcal{L}$ and all $t = T - 1, T - 2, \dots, 0$, let

$$\varphi_{a, t} = \max \left\{ \max_{b \in \mathcal{L} \text{ s.t. } t + \delta(a, b) \leq T} \{ \varphi_{b, t + \delta(a, b)} + p_{a, b, t} - c_{a, b, t} \}, -\kappa_{T-t} \right\}. \quad (46)$$

- For all $i \in \mathcal{D}$, s.t. $\beta_i = 1$, let $\varphi_{D_i} = \varphi_{\ell_i, \tau_i}$; for $i \in \mathcal{D}$ s.t. $\beta_i = 0$, let $\varphi_{D_i} = \max\{\varphi_{\ell_i, \tau_i}, 0\}$.
- For all $j \in \mathcal{R}$, let $\mu_j = \max\{v_j - p'_{o_j, d_j, \tau_j}, 0\}$, where

$$p'_{a, b, t} \triangleq \varphi_{a, t} - \varphi_{b, t + \delta(a, b)} + c_{a, b, t}, \quad \forall (a, b, t) \in \mathcal{T}. \quad (47)$$

Note that $p'_{a, b, t} \geq p_{a, b, t}$ holds for all $(a, b, t) \in \mathcal{T}$, since $\varphi_{a, t} \geq \varphi_{b, t + \delta(a, b)} + p_{a, b, t} - c_{a, b, t}$ from (46). Moreover, we claim that for any optimal solution (x, y) to the LP (9),

$$p_{a, b, t} = p'_{a, b, t} \text{ for all } (a, b, t) \in \mathcal{T} \text{ s.t. } \sum_{i \in \mathcal{D}} \sum_{k=1}^{|Z_i|} y_{i, k} \mathbb{1}\{(a, b, t) \in Z_{i, k}\} > 0, \quad (48)$$

i.e. the prices must coincide for any trips that is taken by at least one driver.

To prove (48), first observe that $\forall (a, t) \in \mathcal{L} \times [T]$, $\varphi_{a, t}$ as in (46) is equal to the highest total utility among all possible paths starting from (a, t) to the end of time, given the prices p — this is obvious for $t = T$ (since $\varphi_{a, T} = 0$ for all a) and also for $t < T$ by induction. Now consider any

trip (a, b, t) taken by some driver, say driver 1. Since the outcome forms a CE (since (x, y) and (p, π, u) are optimal primal and dual solutions), we know that the total utility to driver 1 from time t onward must be $\varphi_{a,t}$, the highest total utility among all possible paths starting from (a, t) . Similarly, the total utility to driver 1 from location b and time $t + \delta(a, b)$ onward is $\varphi_{b,t+\delta(a,b)}$. Since the (a, b, t) trip pays the driver $p_{a,b,t}$ and costs $c_{a,b,t}$, we know that $\varphi_{a,t} = p_{a,b,t} + \varphi_{b,t+\delta(a,b)} - c_{a,b,t}$ must hold, which gives us $p'_{a,b,t} = \varphi_{a,t} - \varphi_{b,t+\delta(a,b)} = p_{a,b,t}$.

We now show that (φ, μ) forms an optimal solution to (39). Given the non-negativity of *price*, constraints (40) to (45) are satisfied by construction, thus what is left to prove is optimality. Let (x, y) be some optimal integral solution to (9). We know that (x, y) and (p, π, u) satisfy the CS conditions (CS-1)-(CS-4), and form a CE. We construct an optimal integral solution f to (34) from (x, y) in the same way as in the proof of Lemma 2, and it is sufficient for the optimality to prove that (CS_F-1)-(CS_F-6) hold between f and (φ, μ) :

1. To show (CS_F-1), first observe that for all $j \in \mathcal{R}$, $f(e_j) > 0 \Rightarrow x_j > 0$ implies that rider j is picked up, thus the trip (o_j, d_j, τ_j) is taken by some driver, thus $p_{o_j, d_j, \tau_j} = p'_{o_j, d_j, \tau_j}$ by (48). Moreover, $u_j = v_j - p_{o_j, d_j, \tau_j}$ from (CS-1), implying $v_j - p'_{o_j, d_j, \tau_j} \geq 0$. This gives us: $f(e_j) > 0 \Rightarrow \mu_j = v_j - p'_{o_j, d_j, \tau_j} \Rightarrow \varphi_{o_j, \tau_j} - \varphi_{d_j, \tau_j + \delta(o_j, d_j)} + \mu_j = v_j - c_{o_j, d_j, \tau_j}$.
2. To show (CS_F-2), recall that $p'_{a,b,t} \geq p_{a,b,t}$ for all $(a, b, t) \in \mathcal{T}$. Therefore, given (17) and (CS-2), we know that for all $j \in \mathcal{R}$, $\mu_j > 0 \Rightarrow v_j - p'_{o_j, d_j, \tau_j} > 0 \Rightarrow v_j - p_{o_j, d_j, \tau_j} > 0 \Rightarrow u_j > 0 \Rightarrow x_j = 1 \Rightarrow f(R_j) = 1$.
3. (CS_F-3) holds since $f(e) > 0$ only when there is excessive supply in the dispatching (x, y) for the trip (a, b, t) , therefore $p_{a,b,t} = 0$ given (CS-4). (48) then implies that $p'_{a,b,t} = 0$, thus $\varphi_{a,t} - \varphi_{b,t+\delta(a,b)} = p'_{a,b,t} - c_{a,b,t} = -c_{a,b,t}$.
4. (CS_F-4) holds, since $f((a, t), S) > 0$ implies that given dispatch (x, y) , there exists at least one driver that exited the platform from (a, t) , therefore gets utility $-\kappa_{T-t}$ from time t onward. As a result, the highest utility for all paths from (a, t) onward must be $\varphi_{a,t} = -\kappa_{T-t}$ without violating CE.
5. (CS_F-5) holds by construction for i s.t. $\beta_i = 1$. For i s.t. $\beta_0 = 0$, since when $f((D_i, (\ell_i, \underline{\tau}_i))) > 0$, driver i entered the platform (instead of not entering and getting zero utility), thus her utility $\varphi_{\ell_i, \underline{\tau}_i}$ from $(\ell_i, \underline{\tau}_i)$ onward must not be negative. and as a result, $\varphi_{D_i} = \max\{\varphi_{\ell_i, \underline{\tau}_i}, 0\} = \varphi_{\ell_i, \underline{\tau}_i}$.
6. (CS_F-6) holds, since when $f((D_i, S)) > 0$, driver i did not enter the platform at all according to the dispatch (x, y) . As a result, CE implies that entering must not give her positive utility, therefore $\varphi_{\ell_i, \underline{\tau}_i} \leq 0$, and her utility $\varphi_{D_i} = \max\{\varphi_{\ell_i, \underline{\tau}_i}, 0\} = 0$.

What is left to show is that $\varphi_{D_i} = \pi_i$ for all $i \in \mathcal{D}$ and that $\mu_j = u_j$ for all $j \in \mathcal{D}$. $\varphi_{D_i} = \pi_i$ holds, since φ_{D_i} as constructed is the highest achievable utility for driver i among all feasible paths, and π_i must take this value given CE. For $j \in \mathcal{R}$ s.t. $x_j = 0$, we know that $f(R_j) = 0$, and $u_j = \mu_j = 0$ must hold. For $j \in \mathcal{R}$ s.t. $x_j = 1$, $p_{o_j, d_j, \tau_j} = p'_{o_j, d_j, \tau_j}$ holds given (48), therefore $f(R_j) \Rightarrow \mu_j = v_j - p'_{o_j, d_j, \tau_j} = v_j - p_{o_j, d_j, \tau_j} = u_j$.

Part (ii). Let (φ, μ) be an optimal solution to (39). We now construct a solution (p, π, u) to (15), where the price $p_{a,b,t}$ is the loss of potential between the origin node (a, t) and the destination node $(b, t + \delta(a, b))$ plus the trip cost $c_{a,b,t}$, and the driver and rider utilities are given by φ_{D_i} and μ_i :

$$\begin{aligned} u_j &= \mu_j, & \forall j \in \mathcal{R} \\ \pi_i &= \varphi_{D_i}, & \forall i \in \mathcal{D} \\ p_{a,b,t} &= \varphi_{a,t} - \varphi_{b,t+\delta(a,b)} + c_{a,b,t}, & \forall (a, b, t) \in \mathcal{T} \end{aligned}$$

We first show that (p, π, u) is a feasible solution to (15).

1. From telescoping sum, for any feasible path $Z_{i,k}$ of driver i , starting at (ℓ_i, τ_i) and ending at (a', t) for some $a' \in \mathcal{L}$ and some $t' \in [T]$, the total utility from taking the path is $\sum_{(a,b,t) \in Z_{i,k}} (p_{a,b,t} - c_{a,b,t}) - \kappa_{T-t'} = \varphi_{\ell_i, \tau_i} - \varphi_{a', t'} - \kappa_{T-t'}$, which is at most φ_{ℓ_i, τ_i} (since $\varphi_{a', t'} \geq -\kappa_{T-t'}$ for all $a' \in \mathcal{L}$ and $t' \in [T]$, guaranteed by (42)). This implies that the utility $\pi_i = \varphi_{D_i} \geq \varphi_{\ell_i, \tau_i} \geq \sum_{(a,b,t) \in Z_{i,k}} (p_{a,b,t} - c_{a,b,t}) - \kappa_{T-t'} = \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} - \lambda_{i,k}$ for any $k \in \{1, \dots, |\mathcal{Z}_i|\}$, therefore (16) holds.
2. (40) implies $u_j = \mu_j \geq v_j - (\varphi_{o_j, \tau_j} - \varphi_{d_j, \tau_j + \delta(o_j, d_j)} + c_{o_j, d_j, \tau_j}) = v_j - p_{o_j, d_j, \tau_j}$ thus (17) holds.
3. (41) implies that $p_{a,b,t} \geq 0$ thus (18) holds.
4. Lastly, (45) implies $u_j = \mu_j \geq 0$, which is (19).

Therefore, (p, π, u) is a feasible solution to (15). Regarding the optimality of (p, π, u) , we know by construction that the objective of (15) is equal to that of (39). Recall the correspondence of optimal solutions that we established in Appendix B.2, that the optimal objective of the flow LP (34) is equal to that of the original LP (9). This implies that the optimal objective of the dual (15) is equal to the optimal objective of the LP (9), therefore (p, π, u) is an optimal solution such that $\pi_i = \varphi_{D_i}$ for all $i \in \mathcal{D}$ and that $u_j = \mu_j$ for all $j \in \mathcal{R}$.

This completes the proof of the lemma. \square

B.4.2 Proof of Lemma 4

Lemma 4 (Lattice Structure). Drivers' utility profile π among all CE outcomes form a lattice. Moreover, for each driver $i \in \mathcal{D}$, Φ_{D_i} and Ψ_{D_i} are equal to utility of driver i in the driver-pessimal and driver-optimal CE plans, respectively.

Proof. Step 2 of the proof of Lemma 3 established that the set of possible driver utilities among all CE outcomes correspond to the π variables among the set of optimal solutions (p, π, u) to the dual LP (15). Since Lemma 7 established the correspondence between the π variables and the φ_{D_i} variables in optimal solutions to (15) and (39), what we need to show is the lattice structure of φ in optimal solutions of (39), and that Φ and Ψ reside on the bottom and the top of the lattice.

Step 1. Proof of the Lattice Structure

We first prove the lattice structure. Let (φ, μ) and (φ', μ') be two optimal solutions of (39). We prove that the join and the meet of φ and φ' are both optimal potentials. Let the join and the meet be defined as: For all $(a, t) \in \mathcal{L} \times [T]$, let the join and the meet of the potentials be

$$\begin{aligned} \bar{\varphi}_{a,t} &\triangleq \max\{\varphi_{a,t}, \varphi'_{a,t}\}, \quad \forall (a, t) \in \mathcal{L} \times [T], \quad \bar{\varphi}_{D_i} \triangleq \max\{\varphi_{D_i}, \varphi'_{D_i}\}, \quad \forall i \in \mathcal{D}, \\ \underline{\varphi}_{a,t} &\triangleq \min\{\varphi_{a,t}, \varphi'_{a,t}\}, \quad \forall (a, t) \in \mathcal{L} \times [T], \quad \underline{\varphi}_{D_i} \triangleq \min\{\varphi_{D_i}, \varphi'_{D_i}\}, \quad \forall i \in \mathcal{D}, \end{aligned}$$

For convenience of notation, for all $(a, b, t) \in \mathcal{T}$, denote

$$\begin{aligned} p_{a,b,t} &\triangleq \varphi_{a,t} - \varphi_{b,t+\delta(a,b)} + c_{a,b,t}, \\ p'_{a,b,t} &\triangleq \varphi'_{a,t} - \varphi'_{b,t+\delta(a,b)} + c_{a,b,t}, \end{aligned}$$

and let \bar{p} and \underline{p} be the prices constructed from the join and the meet of the potentials:

$$\begin{aligned} \bar{p}_{a,b,t} &\triangleq \bar{\varphi}_{a,t} - \bar{\varphi}_{b,t+\delta(a,b)} + c_{a,b,t}, \\ \underline{p}_{a,b,t} &\triangleq \underline{\varphi}_{a,t} - \underline{\varphi}_{b,t+\delta(a,b)} + c_{a,b,t}. \end{aligned}$$

Finally, for all $j \in \mathcal{R}$, let

$$\begin{aligned}\bar{\mu}_j &\triangleq \max\{v_j - \bar{p}_{o_j, d_j, \tau_j}, 0\}, \\ \underline{\mu}_j &\triangleq \max\{v_j - \underline{p}_{o_j, d_j, \tau_j}, 0\}.\end{aligned}$$

We first prove that both $(\bar{\varphi}, \bar{\mu})$ and $(\underline{\varphi}, \underline{\mu})$ are feasible solutions to (39). Constraints (40), (42), (44) and (45) hold by construction. For constraint (41), we first show that for all $(a, b, t) \in \mathcal{T}$,

$$\bar{p}_{a,b,t} \in [\min\{p_{a,b,t}, p'_{a,b,t}\}, \max\{p_{a,b,t}, p'_{a,b,t}\}], \quad (49)$$

$$\underline{p}_{a,b,t} \in [\min\{p_{a,b,t}, p'_{a,b,t}\}, \max\{p_{a,b,t}, p'_{a,b,t}\}]. \quad (50)$$

We only prove $\bar{p}_{a,b,t} \geq \min\{p_{a,b,t}, p'_{a,b,t}\}$ here. The proof for the other three inequalities are very similar. Assume w.l.o.g. that $\varphi_{a,t} \geq \varphi'_{a,t}$. This implies $\bar{\varphi}_{a,t} = \varphi_{a,t}$. Consider two scenarios: (I) $\varphi_{b,t+\delta(a,b)} \geq \varphi'_{b,t+\delta(a,b)}$ and (II) $\varphi_{b,t+\delta(a,b)} < \varphi'_{b,t+\delta(a,b)}$. For (I), $\bar{\varphi}_{b,t+\delta(a,b)} = \varphi_{b,t+\delta(a,b)}$ thus $\bar{p}_{a,b,t} = \bar{\varphi}_{a,t} - \bar{\varphi}_{b,t+\delta(a,b)} + c_{a,b,t} = \varphi_{a,t} - \varphi_{b,t+\delta(a,b)} + c_{a,b,t} = p_{a,b,t} \geq \min\{p_{a,b,t}, p'_{a,b,t}\}$. For (II), we know that $\bar{\varphi}_{b,t+\delta(a,b)} = \max\{\varphi_{b,t+\delta(a,b)}, \varphi'_{b,t+\delta(a,b)}\} = \varphi'_{b,t+\delta(a,b)}$, thus $\bar{p}_{a,b,t} = \bar{\varphi}_{a,t} - \bar{\varphi}_{b,t+\delta(a,b)} + c_{a,b,t} = \varphi_{a,t} - \varphi'_{b,t+\delta(a,b)} + c_{a,b,t} \geq \varphi'_{a,t} - \varphi'_{b,t+\delta(a,b)} + c_{a,b,t} = p'_{a,b,t} \geq \min\{p_{a,b,t}, p'_{a,b,t}\}$.

φ and φ' satisfying (41) implies $p_{a,b,t} \geq 0$ and $p'_{a,b,t} \geq 0$, which means that $\min\{p_{a,b,t}, p'_{a,b,t}\} \geq 0$ for all $(a, b, t) \in \mathcal{T}$. Therefore $\bar{p}_{a,b,t} \geq 0$ and $\underline{p}_{a,b,t} \geq 0$ both hold. This proves $\bar{\varphi}$ and $\underline{\varphi}$ satisfy (41). Constraint (43) also holds, since for all $i \in \mathcal{D}$, $\bar{\varphi}_{D_i} = \max\{\varphi_{D_i}, \varphi'_{D_i}\} \geq \max\{\varphi_{\ell_i, \mathcal{I}_i}, \varphi'_{\ell_i, \mathcal{I}_i}\} = \bar{\varphi}'_{\ell_i, \mathcal{I}_i}$, and $\underline{\varphi}_{D_i} = \min\{\varphi_{D_i}, \varphi'_{D_i}\} \geq \min\{\varphi_{\ell_i, \mathcal{I}_i}, \varphi'_{\ell_i, \mathcal{I}_i}\} = \underline{\varphi}'_{\ell_i, \mathcal{I}_i}$. Thus $(\bar{\varphi}, \bar{\mu})$ and $(\underline{\varphi}, \underline{\mu})$ are both feasible.

Let f be an integral optimal solution to the flow LP (34). We prove that $(\bar{\varphi}, \bar{\mu})$ and $(\underline{\varphi}, \underline{\mu})$ are both optimal solutions to (39) by showing that the CS conditions (CS_{F-1})-(CS_{F-6}) hold between f and $(\bar{\varphi}, \bar{\mu})$, and also between f and $(\underline{\varphi}, \underline{\mu})$. First note that (CS_{F-1})-(CS_{F-6}) hold in between f and (φ, μ) and between f and (φ', μ') .

1. To show (CS_{F-1}), note that (CS_{F-1}) between f and (φ, μ) , (φ', μ') imply that if $f(R_j) > 0$, $\mu_j = v_j - p_{o_j, d_j, \tau_j} \geq 0$ and $\mu'_j = v_j - p'_{o_j, d_j, \tau_j} \geq 0$. Applying (49) and (50), we get $v_j - \bar{p}_{o_j, d_j, \tau_j} \geq v_j - \max\{p_{o_j, d_j, \tau_j}, p'_{o_j, d_j, \tau_j}\} = \min\{v_j - p_{o_j, d_j, \tau_j}, v_j - p'_{o_j, d_j, \tau_j}\} \geq 0$ and $v_j - \underline{p}_{o_j, d_j, \tau_j} \geq v_j - \max\{p_{o_j, d_j, \tau_j}, p'_{o_j, d_j, \tau_j}\} = \min\{v_j - p_{o_j, d_j, \tau_j}, v_j - p'_{o_j, d_j, \tau_j}\} \geq 0$. The definitions of $\bar{\mu}_j$ and $\underline{\mu}_j$ then imply $\bar{\mu}_j = v_j - \bar{p}_{o_j, d_j, \tau_j}$ and $\underline{\mu}_j = v_j - \underline{p}_{o_j, d_j, \tau_j}$ hold.
2. To show (CS_{F-2}), observe that for all $j \in \mathcal{R}$, $\bar{\mu}_j > 0 \Rightarrow v_j - \bar{p}_{o_j, d_j, \tau_j} > 0 \Rightarrow v_j - \min\{p_{o_j, d_j, \tau_j}, p'_{o_j, d_j, \tau_j}\} > 0 \Rightarrow \max\{v_j - p_{o_j, d_j, \tau_j}, v_j - p'_{o_j, d_j, \tau_j}\} > 0 \Rightarrow \max\{\mu_j, \mu'_j\} > 0 \Rightarrow f(R_j) = 1$. Similarly, $\underline{\mu}_j > 0 \Rightarrow f(R_j) = 1$.
3. We now consider (CS_{F-3}). For any $e = ((a, t), (b, t+\delta(a, b))) \in \mathcal{E}_2$, $f(e) > 0 \Rightarrow \varphi_{a,t} - \varphi_{b,t+\delta(a,b)} = \varphi'_{a,t} - \varphi'_{b,t+\delta(a,b)} = -c_{a,b,t}$. This implies that $p_{a,b,t} = p'_{a,b,t} = 0$, and as a result, $0 \leq \bar{p}_{a,b,t}, \underline{p}_{a,b,t} \leq 0$, therefore $\bar{\varphi}_{a,t} - \bar{\varphi}_{b,t+\delta(a,b)} = -c_{a,b,t}$ and $\underline{\varphi}_{a,t} - \underline{\varphi}_{b,t+\delta(a,b)} = -c_{a,b,t}$ both hold.
4. (CS_{F-4}) holds since for each $e = ((a, t), S) \in \mathcal{E}$ for some $a \in \mathcal{L}$ and $t \in [T]$, $f(e) > 0 \Rightarrow \varphi_{a,T} = -\kappa_{T-t}$ and $\varphi'_{a,T} = -\kappa_{T-t}$. Thus, $\bar{\varphi}_{a,T} = \underline{\varphi}_{a,T} = -\kappa_{T-t}$.
5. For (CS_{F-5}): $f((D_i, (\ell_i, \mathcal{I}_i))) > 0$ implies $\varphi_{D_i} = \varphi_{\ell_i, \mathcal{I}_i}$ and $\varphi'_{D_i} = \varphi'_{\ell_i, \mathcal{I}_i}$, thus $\bar{\varphi}_{D_i} = \max\{\varphi_{D_i}, \varphi'_{D_i}\} = \max\{\varphi_{\ell_i, \mathcal{I}_i}, \varphi'_{\ell_i, \mathcal{I}_i}\} = \bar{\varphi}'_{\ell_i, \mathcal{I}_i}$ and $\underline{\varphi}_{D_i} = \min\{\varphi_{D_i}, \varphi'_{D_i}\} = \min\{\varphi_{\ell_i, \mathcal{I}_i}, \varphi'_{\ell_i, \mathcal{I}_i}\} = \underline{\varphi}'_{\ell_i, \mathcal{I}_i}$.
6. For (CS_{F-6}), for $i \in \mathcal{D}$ s.t. $\beta_i = 0$ and $f((D_i, S)) > 0$, $\varphi_{D_i} = \varphi'_{D_i} = 0$, thus $\bar{\varphi}_{D_i} = \underline{\varphi}_{D_i} = 0$.

This completes the proof of the lattice structure of drivers' total utilities.

Step 2. Driver Optimal and Pessimial Plans.

We now prove the correspondence between the welfare changes and the top and bottom of the lattice. Recall that Φ_{D_i} and Ψ_{D_i} are the welfare gain/loss from replicating/losing driver i , respectively, and $\Phi_{a,t}$ is the welfare gain from adding another driver (that has entered) to location a and time t . Here we define the welfare loss from losing one driver from location a at time t as:

$$\Psi_{a,t} \triangleq W(\mathcal{D}, \mathcal{R}) - W(\mathcal{D} \setminus \{(1, a, t, T)\}, \mathcal{R}), \quad (51)$$

where $W(\mathcal{D} \setminus \{(1, a, t, T)\})$ is the highest achievable social welfare, if one of the drivers in \mathcal{D} who was supposed to exit the platform at time T now needs to exit the platform at location a at time t . Note that this does not specify which particular driver exits, and can be considered as the objective of the flow LP where we simply subtract 1 from the boundary condition $\xi_{a,t}$ at the node (a, t) .

We first show via standard arguments with the residual graph that Φ and Ψ as we defined in (27) and (51) are optimal potentials for the flow LP. We then show via subgradient arguments that Φ and Ψ are the bottom and the top of the lattice of the potentials, respectively. Given Lemma 7, and the fact that driver payments among CE outcomes correspond to the optimal solutions of the dual LP (15), we know Φ_{D_i} and Ψ_{D_i} correspond to the bottom and the top of the lattice of driver's total payments among all CE outcomes, hence Lemma 4.

Step 2.1. Φ and Ψ are Optimal Potentials:

Given the MCF problem (29) with graph $G = (\mathcal{N}, \mathcal{E})$ and an optimal integral solution f (which is guaranteed to exist), we first construct the standard residual graph $\tilde{G} = (\mathcal{N}, \tilde{\mathcal{E}})$ where the set of nodes remains the same, and the set of edges $\tilde{\mathcal{E}} = \tilde{\mathcal{E}}_1 \cup \tilde{\mathcal{E}}_2 \cup \tilde{\mathcal{E}}_3 \cup \tilde{\mathcal{E}}_4$ consists of:

- $\tilde{\mathcal{E}}_1 = \{R_j \mid j \in \mathcal{R}, f(R_j) = 0\} \cup \{\tilde{R}_j \mid j \in \mathcal{R}, f(R_j) = 1\}$, where $R_j = ((o_j, \tau_j), (d_j, \tau_j + \delta(o_j, d_j)))$ is the edge corresponding to rider j with $\gamma(R_j) = -v_j + c_{o_j, d_j, \tau_j}$, $\underline{\zeta}(R_j) = 0$, and $\bar{\zeta}(\tilde{R}_j) = 1$; $\tilde{e}_j = ((d_j, \tau_j + \delta(o_j, d_j), (o_j, \tau_j)))$ is the reversed edge corresponding to rider j s.t. $f(R_j) = 1$, with $\gamma(\tilde{R}_j) = v_j - c_{o_j, d_j, \tau_j}$, $\underline{\zeta}(\tilde{R}_j) = 0$ and $\bar{\zeta}(\tilde{R}_j) = 1$.
- $\tilde{\mathcal{E}}_2 = \mathcal{E}_2 \cup \{\tilde{e} \mid e \in \mathcal{E}_2, f(e) > 0\}$, where for each $e = ((a, t), (b, t + \delta(a, b))) \in \mathcal{E}_2$ with $f(e) > 0$, $\tilde{e} = ((b, t + \delta(a, b)), (a, t))$, and has $\gamma(\tilde{e}) = -c_{a, b, t}$, $\underline{\zeta}(\tilde{e}) = 0$ and $\bar{\zeta}(\tilde{e}) = f(e)$.
- $\tilde{\mathcal{E}}_3 = \mathcal{E}_3 \cup \{\tilde{e} \mid e \in \mathcal{E}_3, f(e) > 0\}$ where for each $e = ((a, t), S) \in \mathcal{E}_3$, $\tilde{e} = (S, (a, t))$ with $\gamma(\tilde{e}) = -\kappa_{T-t}$, $\underline{\zeta}(\tilde{e}) = 0$ and $\bar{\zeta}(\tilde{e}) = f(e)$.
- $\tilde{\mathcal{E}}_4 = \{(D_i, (\ell_i, \tau_i)) \mid i \in \mathcal{D}, f((D_i, (\ell_i, \tau_i))) = 0\} \cup \{((\ell_i, \tau_i), D_i) \mid i \in \mathcal{D}, f((D_i, (\ell_i, \tau_i))) = 1\} \cup \{(D_i, S) \mid i \in \mathcal{D}, \beta_i = 0, f((D_i, S)) = 0\} \cup \{(S, D_i) \mid i \in \mathcal{D}, \beta_i = 0, f((D_i, S)) = 1\}$. For the forward edges, i.e. $e \in \tilde{\mathcal{E}}_4$ s.t. $e = (D_i, (\ell_i, \tau_i))$ or $e = (D_i, S)$, we have $\gamma(e) = 0$, $\underline{\zeta}(e) = 0$. and $\bar{\zeta}(e) = +\infty$. For each $e = ((\ell_i, \tau_i), D_i) \in \tilde{\mathcal{E}}_4$, we have $\gamma(e) = 0$, $\underline{\zeta}(e) = 0$, and $\bar{\zeta}(e) = f((D_i, (\ell_i, \tau_i)))$, and for each $e = (S, D_i) \in \tilde{\mathcal{E}}_4$, we have $\gamma(e) = 0$, $\underline{\zeta}(e) = 0$, and $\bar{\zeta}(e) = f((D_i, S))$.

From the standard argument on the residual graphs [2], we know that the cost of a feasible flow in the residual graph is equal to the incremental cost of the same flow in the original graph. For any node $n = \mathcal{N}$, the ‘‘shortest distance’’ from this node to the sink S refers to the smallest total cost among all paths from n to S in the residual graph. Since the edge costs are equal to driver costs minus rider values, the shortest distance corresponds to the negation of the maximum

incremental welfare created by an additional unit of driver flow starting from n , i.e. $-\Phi_{a,t}$ at node (a, t) , or $-\Phi_{D_i}$ at node D_i . Given $-\Phi$ as the (negation of the) shortest distances, define:

$$p_{a,b,t} \triangleq \Phi_{a,t} - \Phi_{b,t+\delta(a,b)} + c_{a,b,t}, \quad \forall (a, b, t) \in \mathcal{T},$$

$$\mu_j \triangleq \max\{v_j - p_{o_j,d_j,\tau_j}, 0\}, \quad \forall j \in \mathcal{R},$$

we show that (Φ, μ) forms an optimal solution to (39). The argument is very similar to that of the reduced cost optimality, however, we include the proof here for completeness. We first show the feasibility of (Φ, μ) :

1. Constraint (40) holds by definition of μ .
2. For (41), observe that for all $(a, b, t) \in \mathcal{T}$, there exists an edge $((a, t), (b, t + \delta(a, b))) \in \tilde{\mathcal{E}}$ with cost $c_{a,b,t}$, thus the shortest distance from (a, t) to S is at most $c_{a,b,t}$ plus the shortest distance from $(b, t + \delta(a, b))$ to S , implying $-\Phi_{a,t} \leq -\Phi_{b,t+\delta(a,b)} + c_{a,b,t} \Rightarrow \Phi_{a,t} - \Phi_{b,t+\delta(a,b)} \geq -c_{a,b,t}$.
3. For (42), note that $\forall a \in \mathcal{L}$ and $\forall t \in [T]$, there exists $((a, t), S) \in \tilde{\mathcal{E}}$ with cost $\gamma(e) = \kappa_{T-t}$. Therefore, the shortest distance $-\Phi_{a,T}$ between (a, t) and S is at most κ_{T-t} , i.e. $-\Phi_{a,T} \leq \kappa_{T-t} \Rightarrow \Phi_{a,T} \geq -\kappa_{T-t}$.
4. For (43), we know that for each $i \in \mathcal{D}$, there exists an edge $(D_i, (\ell_i, \tau_i)) \in \tilde{\mathcal{E}}$ with unlimited capacity and zero cost, therefore the shortest path from D_i to the sink S satisfies $-\Phi_{D_i} \leq -\Phi_{\ell_i, \tau_i} \Rightarrow \Phi_{D_i} \geq \Phi_{\ell_i, \tau_i}$.
5. For (44), since for each $i \in \mathcal{D}$ s.t. $\beta_i = 0$, there exists $e = (D, S) \in \tilde{\mathcal{E}}$ with unlimited capacity and zero cost, thus $-\Phi_{D_i} \leq 0 \Rightarrow \Phi_{D_i} \geq 0$.
6. (45) holds by definition of μ .

We now show the optimality by examining that the CS conditions (CS_{F-1})-(CS_{F-6}) hold between the optimal integral flow f and (Φ, μ) :

1. To show (CS_{F-1}), given how μ_j is defined, we only need to show that when $f(R_j) > 0$, $v_j - (\Phi_{o_j, \tau_j} - \Phi_{d_j, \tau_j + \delta(o_j, d_j)} + c_{o_j, d_j, \tau_j}) \geq 0$. This holds, since in \tilde{G} , there exists edge \tilde{R}_j from $(d_j, \tau_j + \delta(o_j, d_j))$ to (o_j, τ_j) with cost $v_j - c_{o_j, d_j, \tau_j}$, thus the shortest distances must satisfy: $-\Phi_{d_j, \tau_j + \delta(o_j, d_j)} \leq v_j - c_{o_j, d_j, \tau_j} - \Phi_{o_j, \tau_j} \Rightarrow v_j - (\Phi_{o_j, \tau_j} - \Phi_{d_j, \tau_j + \delta(o_j, d_j)} + c_{o_j, d_j, \tau_j}) \geq 0$.
2. Now consider (CS_{F-2}). Observe that when $\mu_j > 0$, we must have $v_j - (\Phi_{o_j, \tau_j} - \Phi_{d_j, \tau_j + \delta(o_j, d_j)} + c_{o_j, d_j, \tau_j}) > 0$, implying $-\Phi_{o_j, \tau_j} > -\Phi_{d_j, \tau_j + \delta(o_j, d_j)} - v_j + c_{o_j, d_j, \tau_j}$, i.e. in the residual graph, the shortest distance from (o_j, τ_j) to the sink is longer than $-v_j + c_{o_j, d_j, \tau_j}$ plus the shortest distance from $(d_j, \tau_j + \delta(o_j, d_j))$ to the sink. This means that the edge $R_j = ((o_j, \tau_j), (d_j, \tau_j + \delta(o_j, d_j)))$ with capacity 1 and cost $-v_j + c_{o_j, d_j, \tau_j}$ cannot be present in the residual graph, which is the case only if $f(R_j) = 1$.
3. For (CS_{F-3}): we proved $\Phi_{a,t} - \Phi_{b,t+\delta(a,b)} \geq -c_{a,b,t}$ above for feasibility, thus we only need to show the other direction of the inequality. Observing that with $f((a, t), (b, t + \delta(a, b))) > 0$, there exists an edge from $(b, t + \delta(a, b))$ to (a, t) in the residual graph with cost $-c_{a,b,t}$ and non-zero capacity, thus the shortest distance from $(b, t + \delta(a, b))$ to the sink is at most $-\Phi_{a,t} - c_{a,b,t}$, implying $\Phi_{a,t} - \Phi_{b,t+\delta(a,b)} \leq -c_{a,b,t}$.
4. Assume (CS_{F-4}) does not hold and given feasibility, we know that there exists $a \in \mathcal{L}$ and $t \in [T]$ s.t. $f(((a, t), S)) > 0$ and $-\Phi_{a,T} < \kappa_{T-t}$. This implies that the minimum cost for an extra unit of flow from (a, t) to the sink is lower than κ_{T-t} , and the objective of the flow LP can be improved by routing one unit of flow that goes from (a, t) directly to the S through this alternative shortest path. This contradicts the optimality of f .

5. For (CS_F-5): given $i \in \mathcal{D}$ s.t. $f((D_i, (\ell_i, \mathcal{I}_i))) > 0$, there exists an edge from (ℓ_i, \mathcal{I}_i) to D_i with zero cost, thus $-\Phi_{\ell_i, \mathcal{I}_i} \leq -\Phi_{D_i} + 0 \Rightarrow \Phi_{D_i} \leq \Phi_{\ell_i, \mathcal{I}_i}$. Together with (43), we know $\varphi_{D_i} = \Phi_{\ell_i, \mathcal{I}_i}$.
6. For (CS_F-6), given $f((D_i, S)) > 0$, we know that there's a unit of flow from D_i to S generating a total cost of zero. If $-\Phi_{D_i} < 0$, there exists a path from D_i to S for which routing a unit of driver flow improves the objective (in comparison to going directly from D_i to S). This contradicts the optimality, thus $-\Phi_{D_i} \geq 0 \Rightarrow \Phi_{D_i} \leq 0$. Given feasibility, we know $\Phi_{D_i} = 0$.

This completes the argument that (Φ, μ) form an optimal solution to (34), thus the unit replica welfare gain $\{\Phi_{D_i}\}_{i \in \mathcal{D}}$ is indeed a CE driver utility profile. Similarly, we can show that $-\Psi_{a,t}$ and $-\Psi_{D_i}$ corresponds to the shortest distance from the sink S to the node (a, t) and the node D_i , respectively, and that there exists $\mu' \in \mathbb{R}^{|\mathcal{R}|}$ (can be constructed in similar ways as the above μ) s.t. (Ψ, μ') forms an optimal solution to (34).

Step 2.2. Φ and Ψ are the Bottom and Top of the Potential Lattice:

What is left to show is that Φ and Ψ must be the bottom and top of the lattice formed by all optimal potentials of (34). For convenience of notation, we now work with the dual of the original flow LP (29) where the objective is to minimize the negation of the total social welfare: let $\psi_{a,t}$, ψ_{D_i} and η_j be the dual variables corresponding to the constraints (30)-(32), respectively, the dual of (29) can be written in the following form:

$$\begin{aligned}
\max \quad & \sum_{i \in \mathcal{D}} \psi_{D_i} + \sum_{j \in \mathcal{R}} \eta_j & (52) \\
\text{s.t.} \quad & \psi_{o_j, \tau_j} - \psi_{d_j, \tau_j + \delta(o_j, d_j)} + \eta_j \leq -v_j + c_{o_j, d_j, \tau_j}, & \forall j \in \mathcal{R} \\
& \psi_{a,t} - \psi_{b,t + \delta(a,b)} \leq c_{a,b,t}, & \forall (a,b,t) \in \mathcal{T} \\
& \psi_{a,t} \leq \kappa_{T,t}, & \forall a \in \mathcal{L}, t \in [T] \\
& \psi_{D_i} \leq \psi_{\ell_i, \mathcal{I}_i}, & \forall i \in \mathcal{D} \\
& \psi_{D_i} \leq 0, & \forall i \in \mathcal{D} \text{ s.t. } \beta_i = 0 \\
& \eta_j \leq 0, & \forall j \in \mathcal{R}
\end{aligned}$$

For any optimal solution (φ, μ) to (39), (ψ, η) where $\psi = -\varphi$ and $\eta = -\mu$ is an optimal solution to (52), and vice versa. Thus we know the ψ variables among optimal solutions of (52) also form a lattice, and what is left to show is that $-\Phi$ and $-\Psi$ must be the top and bottom of the lattice formed by all optimal potentials of (52).

Recall that for a MCF problem, ξ denotes the boundary condition, so that for each node $n \in \mathcal{N}$, ξ_n is the number of the units of flow that enters (or exits, if negative) the network from node n . For our problem, ξ is a $|\mathcal{D}| + |\mathcal{L}|(T+1)$ dimensional vector, where $\xi_{D_i} = 1$ and $\xi_{(a,t)} = 0$ (recall that the condition $\xi_S = -|\mathcal{D}|$ is redundant given the flow balance constraints, therefore is omitted). Keeping everything else the same, the optimal objective of (29) can be thought of as a function of the boundary condition ξ , which we denote as $\omega(\xi)$. It is known that any potential from the set of all optimal solutions of (52) must be a subgradient of the function $\omega(\xi)$ (see the proof of Theorem 5.2 in [9]), but we still include the proof here for completeness. First, ω is a convex function of ξ (Theorem 5.1 in [9] can easily be generalized to incorporate inequality constraints). Recall that a vector ψ is a *subgradient* of a convex function ω at ξ if for all ξ' ,

$$\omega(\xi) + \psi \cdot (\xi' - \xi) \leq \omega(\xi').$$

Let (ψ, η) be an optimal solution to (52). The strong duality implies $\psi \cdot \xi + \eta \cdot \vec{1} = \omega(\xi)$. Now consider any arbitrary ξ' . For any feasible flow f given the boundary condition ξ' , weak duality implies

$\psi \cdot \xi' + \eta \cdot \vec{1} \leq f \cdot \gamma$ where γ is the vector of all edge costs. Taking the minimum over all feasible flow f , we obtain $\psi \cdot \xi' + \eta \cdot \vec{1} \leq \omega(\xi')$. Hence $\psi \cdot \xi' + \eta \cdot \vec{1} - (\psi \cdot \xi + \eta \cdot \vec{1}) \leq \omega(\xi') - \omega(\xi) \Leftrightarrow \omega(\xi) + \psi(\xi' - \xi) \leq \omega(\xi')$, i.e. ψ is a subgradient of ω at ξ .

Now we show that for any subgradient ψ of ω at ξ , the entries ψ_{D_i} is bounded by $-\Psi_{D_i} \leq \psi_{D_i} \leq -\Phi_{D_i}$.²¹ Let χ_{D_i} be $|\mathcal{D}| + |\mathcal{L}|(T+1)$ by 1 vector which takes value 0 except for the D_i entry, and $\chi_{D_i} = 1$. We know that for any subgradient ψ , $\omega(\xi) + \psi \cdot \chi_{D_i} \leq \omega(\xi + \chi_{D_i}) \Rightarrow \omega(\xi) + \psi_{D_i} \leq \omega(\xi + \chi_{D_i}) \Rightarrow \psi_{D_i} \leq \omega(\xi + \chi_{D_i}) - \omega(\xi) = -\Phi_{D_i}$. The last equality holds since the objective ω is the negation of the optimal total welfare achievable by the vector ξ of driver inflow. Similarly, $\omega(\xi) + \psi \cdot (-\chi_{D_i}) \leq \omega(\xi - \chi_{D_i}) \Rightarrow \omega(\xi) - \psi_{D_i} \leq \omega(\xi - \chi_{D_i}) \Rightarrow \psi_{D_i} \geq \omega(\xi) - \omega(\xi - \chi_{D_i}) = -\Psi_{D_i}$. We can similarly prove $-\Psi_{a,t} \leq \psi_{a,t} \leq -\Phi_{a,t}$. This implies that $-\Phi$ and $-\Psi$ are the top and the bottom of the lattice formed by the optimal potentials of (29), respectively, and therefore completes the proof of the lemma. \square

B.5 Proof of Lemma 5

Lemma 5 (Core equivalence). All CE plans are in the core. Moreover, for any budget-balanced core outcome (x, \tilde{z}, q, r) , there exists prices p such that the plan with anonymous prices (x, \tilde{z}, p) forms a CE, and has the same driver and rider total utilities.

Proof. We first prove that every CE plan is in the core. Let (x, \tilde{z}, p) be a CE plan with anonymous trip price p , where the rider and driver utilities are given by u and π . Fix any coalition $(\mathcal{D}', \mathcal{R}')$ of riders and drivers for some $\mathcal{D}' \subseteq \mathcal{D}$ and $\mathcal{R}' \subseteq \mathcal{R}$. We prove that

$$\sum_{i \in \mathcal{D}'} \pi_i + \sum_{j \in \mathcal{R}'} u_j \geq W(\mathcal{D}', \mathcal{R}'), \quad (53)$$

meaning the total utilities for all drivers and riders in the coalition, under the CE plan, is weakly higher than the highest achievable welfare among themselves. This implies that there is no way for the coalition to make an alternative plan, so that everyone has weakly higher utilities, and at least one driver or rider is strictly better off.

We now prove (53). Let (x', \tilde{z}') be an optimal dispatch that achieves the highest coalitional welfare $W(\mathcal{D}', \mathcal{R}')$. For all $j \in \mathcal{R}'$ s.t. $x'_j = 1$, let her payment be $q'_j = p_{o_j, d_j, \tau_j}$, the anonymous trip price for the trip according to the original CE plan (x, \tilde{z}, p) . Accordingly, let the payment to each driver $i \in \mathcal{D}'$ be $r'_i = \sum_{j \in \mathcal{R}'} \mathbf{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}'_i\} p_{o_j, d_j, \tau_j}$. Under this new plan (x', \tilde{z}', q', r') , the utility of each rider $j \in \mathcal{R}'$ is therefore $u'_j = x'_j(v_j - q'_j)$, and the utility of driver $i \in \mathcal{D}'$ is $\pi'_i = r'_i - \lambda_{i,k}$, if the dispatched action path \tilde{z}'_i is consistent with the k^{th} feasible path of driver i and has total cost of $\lambda_{i,k}$.

Note that the plan (x', \tilde{z}', q', r') is strictly budget balanced, therefore the utility of drivers and riders under this plan add up to the welfare: $\sum_{i \in \mathcal{D}'} \pi'_i + \sum_{j \in \mathcal{R}'} u'_j = W(\mathcal{D}', \mathcal{R}')$. What is left to show is that $u'_j \leq u_j$ for all $j \in \mathcal{R}'$ and $\pi'_i \leq \pi_i$ for all $i \in \mathcal{D}'$ both hold. This is a consequence of the original plan forming a CE. For the riders, if $x'_j = 0$, then $u'_j = 0 \leq u_j$; if $x'_j = 1$, then $u'_j = v_j - p_{o_j, d_j, \tau_j} \leq u_j$. For each $i \in \mathcal{D}'$ with \tilde{z}'_i consistent with the k^{th} feasible path of driver i ,

$$\pi_i \geq \sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}, 0\} - \lambda_{i,k} \geq \sum_{j \in \mathcal{R}'} \mathbf{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}'_i\} p_{o_j, d_j, \tau_j} - \lambda_{i,k} = \pi'_i.$$

This completes the proof of (53), thus all CE plans are in the core.

²¹This is a result of the convexity of ω and the relationship between directional derivatives and subgradients (see Theorem 3.1.14 in [29]). We include a simple proof here for completeness.

We now prove the second part of this lemma, that every core outcome that balances budget can be “priced” in CE. Let (x, \tilde{z}, q, r) be a budget balanced core outcome. The following are immediate implications of a core outcome:

1. The outcome (x, \tilde{z}, q, r) must be strictly budget balanced, i.e. $\sum_{j \in \mathcal{R}} q_j = \sum_{i \in \mathcal{D}} r_i$, since otherwise, the entire economy $(\mathcal{D}, \mathcal{R})$ will be a blocking coalition.
2. The outcome (x, \tilde{z}, q, r) must be welfare-optimal and achieve $W(\mathcal{D}, \mathcal{R})$, otherwise the entire economy $(\mathcal{D}, \mathcal{R})$ is blocking since an improvement of total utilities is possible.
3. The plan is individually rational for riders, i.e. $x_j = 0 \Rightarrow q_j \leq 0$ and $x_j = 1 \Rightarrow q_j \leq v_j$ for each $j \in \mathcal{R}$, otherwise dropping out improves their utilities.

We now claim that $\forall i \in \mathcal{D}$, $\sum_{j \in \mathcal{R}} q_j \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} = r_i$, i.e. we must have strict budget balance among driver i and the riders that she picked up. First, $\sum_{j \in \mathcal{R}} q_j \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} \leq r_i$ holds since otherwise this set of driver and riders will be blocking. Since this inequality holds for all $i \in \mathcal{D}$, the plan is budget balanced, and $q_j \leq 0$ for j s.t. $x_j = 0$, we have

$$\begin{aligned} \sum_{i \in \mathcal{D}} r_i &= \sum_{j \in \mathcal{R}} q_j = \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{R}} q_j \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} + \sum_{j \in \mathcal{R}, x_j=0} q_j \\ &\leq \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{R}} q_j \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} \leq \sum_{i \in \mathcal{D}} r_i, \end{aligned}$$

which requires that all inequalities hold with equality. This implies that $\sum_{j \in \mathcal{R}} q_j \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} = r_i$ for all $i \in \mathcal{D}$, and moreover, $q_j = 0$ for all $j \in \mathcal{R}$ s.t. $x_j = 0$.

We now show that prices must be anonymous for the riders, i.e. for any two riders $j \neq j'$ s.t. $(o_j, d_j, \tau_j) = (o_{j'}, d_{j'}, \tau_{j'})$ and $x_j = x_{j'} = 1$, we must have $q_j = q_{j'}$. Otherwise, assume w.l.o.g. that $q_j < q_{j'}$, and that riders j and j' are picked up by drivers i and i' respectively, we know that rider j' , driver i , and all of the riders picked up by driver i except for rider j , would form a blocking coalition. We now construct a set of anonymous trip prices p . For any trip $(a, b, t) \in \mathcal{T}$,

- (i) if no rider requests this trip, i.e. $(o_j, d_j, \tau_j) \neq (a, b, t)$ for all $j \in \mathcal{R}$, then let $p_{a,b,t} = 0$.
- (ii) if some rider requests this trip, but no rider is picked up, then let the price be the highest value for this trip: $p_{a,b,t} = \max_{j \in \mathcal{R}, (o_j, d_j, \tau_j) = (a, b, t)} v_j$.
- (iii) if some rider is picked up, i.e. if $\exists j \in \mathcal{R}$ s.t. $(o_j, d_j, \tau_j) = (a, b, t)$ and $x_j = 1$, let $p_{a,b,t} = q_j$.

Given the anonymity that we proved above, for any rider that is picked up, she pays p_{o_j, d_j, τ_j} . We claim that $p_{a,b,t} \geq 0$ for all $(a, b, t) \in \mathcal{T}$. This is obvious for cases (i) and (ii) above. For case (iii), we only need to show that payments made by riders that are picked up must be non-negative, i.e. $q_j \geq 0$. This holds, since otherwise the driver who picks up this rider, together with the rest of the riders that this driver picks up, will form a blocking coalition.

We also claim that $p_{a,b,t} = 0$ for trips with excessive supply, i.e. if $(a, b, t) \in \tilde{z}_i$ for some $i \in \mathcal{D}$. Consider some trip with $p_{a,b,t} > 0$. We know that either case (ii) holds, where there is some rider j willing to pay up to $p_{a,b,t}$ but is not picked up, or case (iii) holds, where some rider j is paying $p_{a,b,t}$ to be picked up. In both cases, rider j , driver i (who takes the trip (a, b, t) without a rider), and the rest of the riders picked up by driver i will form a blocking coalition—the rest of the riders can pay the same amounts, driver i can get a higher payment, whereas rider j either gets picked up (case (ii)) or pays less (case (iii)).

We now prove that under plan with anonymous trip prices (x, \tilde{z}, p) , the rider and driver total payments (and therefore utilities) coincide with the original plan. This is obvious for the riders. For each driver, the total payment under plan (x, \tilde{z}, p) is equal to $\sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}, 0\}$, if \tilde{z}

is consistent with the k^{th} path of driver i . Given the non-negativity of p , and the fact that trips with excessive supply has zero prices, we know that driver i is paid $\sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}, 0\} = \sum_{(a,b,t) \in Z_{i,k}} p_{a,b,t} = \sum_{j \in \mathcal{R}} p_{o_j, d_j, \tau_j} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} = \sum_{j \in \mathcal{R}} q_j \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} = r_i$.

We complete the proof of this lemma by showing that (x, \tilde{z}, p) forms a CE. Rider best-response is implied by IR for riders requesting trips where no rider is picked up (case (ii)), and for riders that are already picked up (case (iii)). For rider j s.t. $x_j = 0$ but there exists j' s.t. $(o_j, d_j, \tau_j) = (o_{j'}, d_{j'}, \tau_{j'})$ and $x_{j'} = 1$, we claim $v_j \leq p_{o_j, d_j, \tau_j}$. Otherwise, assume that rider j' is picked up by driver i , we know rider j , driver i and the rest of the riders picked up by driver i will form a blocking coalition.

What is left to show is driver best response. Assume that driver best response doesn't hold for driver i , we know that if \tilde{z}_i is consistent with $Z_{i,k}$, there exists an alternative path $Z_{i,k'} \neq Z_{i,k}$ s.t. $\sum_{(a,b,t) \in Z_{i,k'}} \max\{p_{a,b,t}, 0\} - \lambda_{i,k'} > \sum_{(a,b,t) \in Z_{i,k}} \max\{p_{a,b,t}, 0\} - \lambda_{i,k}$. For each $(a, b, t) \in Z_{i,k'}$ s.t. $p_{a,b,t} > 0$, there exists a rider that is either paying $p_{a,b,t}$ to be picked up, or is not picked up but willing to pay $p_{a,b,t}$. Driver i , together with all of these riders, will form a blocking coalition. This proves that (x, \tilde{z}, p) forms a CE, and therefore completes the proof of this theorem. \square

Note that a core outcome does not necessarily use anonymous trip prices. The following example shows that the CE plan with anonymous trip prices (x, \tilde{z}, p) constructed from a core outcome (x, \tilde{z}, q, r) may not pay $r_{i,t}$ to driver i at time t , and we can only guarantee utility equivalence, i.e. the total payment to each driver is equal to r_i .

Example 4. Consider an economy with a single location A , two time periods, one driver and four riders with $o_j = A$, $d_j = A$, and $v_j = 4$ for all $j \in \mathcal{R}$. Moreover, $\tau_1 = \tau_2 = 0$ and $\tau_3 = \tau_4 = 1$. Assume all costs are zero. Consider the plan (x, \tilde{z}, q, r) , where riders 1 and 3 are picked up and each pays 4: $x_1 = x_3 = 1$ and $q_1 = q_3 = 4$. The driver takes action path $\tilde{z}_1 = ((A, A, 0, 1), (A, A, 1, 3))$, however $r_{1,0} = 2$ and $r_{1,1} = 6$, i.e. the driver is paid 2 at time 0 and 6 at time 1. It is easy to see that the outcome is in the core, however, given any CE plan with anonymous trip prices, $p_{A,A,0} = p_{A,A,1} = 4$, so the driver needs to be paid 4 at each of time 0 and time 1. \square

B.6 Proof of Theorem 2

Theorem 2. *The spatio-temporal pricing mechanism is temporally consistent and subgame-perfect incentive compatible. It is also individually rational for riders and strictly budget balanced for any action profile taken by the drivers, and is welfare optimal, core-selecting, individually rational for drivers and envy-free in subgame-perfect equilibrium from any history onward.*

Proof. As is outlined in the body of the paper, what is left to show is incentive alignment. We first show a correspondence of drivers' continuation utilities and the unit replica welfare gains (which implies that the plan determined by the STP mechanism at any time forms a competitive equilibrium), then we show that there is no useful single deviation, implying that always accepting the mechanism's dispatches forms an SPE.

Step 1. Let (x, \tilde{z}) be the optimal dispatch determined by the STP mechanism, and let f be a corresponding optimal solution to the flow LP (34), constructed in the same way as in the proof of Lemma 2. Setting $u_j = \max\{v_j - p_{o_j, d_j, \tau_j}, 0\}$, we know from the proof of Lemma 4 that (Φ, u) forms an optimal solution to the dual of the flow LP (39), and satisfies the CS conditions with f .

Consider any driver $i \in \mathcal{D}$, who is in the platform, and is available at some location a and time t . Assume that the dispatched action path \tilde{z}_i is consistent with $Z_{i,k}$, the k^{th} feasible path of driver i , and assume that path $Z_{i,k}$ ends at location a' and time t' (i.e. the driver is dispatched to exit the

platform at (a', t')). Assuming that all drivers follow the dispatches of the platform at all times, the total payment to driver i from time t onward is:

$$\sum_{j \in \mathcal{R}, \tau_j \geq t} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} p_{o_j, d_j, \tau_j} = \sum_{t'' \geq t} \mathbb{1}\{(a'', b'', t'') \in Z_{i,k}\} p_{a'', b'', t''},$$

since when the driver takes a relocation trip (a'', b'', t'') without a rider, $f(((a'', t''), (b'', t + \delta(a'', b'')))) > 0$, and the complementary slackness condition (CS_{F-3}) implies that the trip price $p_{a'', b'', t''} = 0$. Moreover, when there exists a driver exiting from (a', t') , $f((a', t'), S) > 0$, (CS_{F-4}) implies $\Phi_{a', t'} = -\kappa_{T-t'}$. As a result, the utility of driver i from time t onward is

$$\begin{aligned} & \sum_{j \in \mathcal{R}, \tau_j \geq t} \mathbb{1}\{(o_j, d_j, \tau_j, j) \in \tilde{z}_i\} p_{o_j, d_j, \tau_j} - \sum_{t'' \geq t} \mathbb{1}\{(a'', b'', t'') \in Z_{i,k}\} c_{a'', b'', t''} - \kappa_{T-t'} \\ &= \sum_{t'' \geq t} \mathbb{1}\{(a'', b'', t'') \in Z_{i,k}\} (p_{a'', b'', t''} - c_{a'', b'', t''}) - \kappa_{T-t'} \\ &= \sum_{t'' \geq t} \mathbb{1}\{(a'', b'', t'') \in Z_{i,k}\} (\Phi_{a'', t''} - \Phi_{b'', t'' + \delta(a'', b'')} + c_{a'', b'', t''} - c_{a'', b'', t''}) - \kappa_{T-t'} \\ &= \Phi_{a, t} - \Phi_{a', t'} - \kappa_{T-t'} \\ &= \Phi_{a, t}. \end{aligned}$$

A first implication is that for a driver with $\beta_i = 1$, her total utility over the planning horizon is $\pi_i = \Phi_{\ell_i, \underline{\tau}_i}$. This is equal to Φ_{D_i} give (CS_{F-5}), since $f((D_i, (\ell_i, \underline{\tau}_i))) = 1 \rightarrow \Phi_{D_i} = \Phi_{\ell_i, \underline{\tau}_i}$. For a driver with $\beta_i = 0$ but was dispatched to enter the platform, $\pi_i = \Phi_{D_i} = \Phi_{\ell_i, \underline{\tau}_i}$ holds for the same reason. Drivers with $\beta_i = 0$ and was not dispatched to enter the platform get $\pi_i = 0$, which is also equal to Φ_{D_i} , since $f((D_i, S)) > 0 \Rightarrow \Phi_{D_i} = 0$. Therefore, $\pi_i = \Phi_{D_i}$ holds for all $i \in \mathcal{D}$. Lemma 3 and Lemma 7 then imply that the plan determined by the STP mechanism forms a CE.

Step 2. We now prove that a single deviation from the mechanism's dispatches by any driver at any time is not useful. For drivers who are (at time t given state s_t) *en route*, or have already exited, or has not yet entered, there is effectively only one actions that is available to them, so there is no useful deviation. Therefore we only need to consider a driver that is at time t available.

Given any time t' and state $s_{t'}$, let $\Phi_{a, t}^{(t')}(s_{t'})$ be the welfare gain from adding an additional driver (who is already in the platform) at time $t \geq t'$ in location a , to the economy starting at time t' and state $s_{t'}$:

$$\Phi_{a, t}^{(t')}(s_{t'}) \triangleq W(\mathcal{D}^{(t')}(s_{t'}) \cup \{(1, a, t - t', T - t')\}) - W(\mathcal{D}^{(t')}(s_{t'})). \quad (54)$$

Here, $(1, a, t - t', T - t')$ is the type of the additional driver that starts at a at time t in the original economy, and therefore at time $t - t'$ in the time-shifted economy $E^{(t')}(s_{t'})$ (where the length of the planning horizon is $T - t'$).

Assume that the current plan $(x^{(t')}(s_{t'}), z^{(t')}(s_{t'}), q^{(t')}(s_{t'}), \pi^{(t')}(s_{t'}))$ is computed at time t' given state $s_{t'}$, and that no driver had deviated from the plan since time t' . Fix any time $t \geq t'$ and let s_t be the state of the platform at time t , if all drivers followed the plan up to time t . Consider a driver, say driver i , that is available at time t at location a , i.e. $s_{i, t} = (1, a, t)$ or $s_{i, t} = (0, a, t)$. We first argue that deviating from the dispatch to exit (when dispatched to stay) or not enter (when dispatched to enter) is not a useful deviation. This is because exiting or not entering is equivalent to the driver's choosing a path different than the one determined by the plan, and the plan forming a CE implies that no alternative path is more profitable.

What is left to consider is the case where a driver deviated from the dispatches (regardless of which action she is dispatched to take), and did not exit the platform. The only possible deviation action that the driver can take in this case is to relocate to some location $b \in \mathcal{L}$ that is within reach (i.e. $b \in \mathcal{L}$ s.t. $t + \delta(a, b) \leq T$). From Step 1, we know that if all drivers follow the plan until the end of the planning horizon, then a driver with $s_{i,t} = (1, a, t)$ gets utility $\Phi_{a,t}^{(t')}(s_{t'})$ in the remaining time periods, and a driver with $s_{i,t} = (0, a, t)$ gets utility $\max\{\Phi_{a,t}^{(t')}(s_{t'}), 0\}$. The rest of the proof of this theorem shows that by deviating to to drive to b , the utility of the driver from time t onward is upper bounded by $\Phi_{a,t}^{(t')}(s_{t'})$, thus this is not a useful deviation.

If all drivers followed the plan at time t , denote the state of the platform at time $t + 1$ as s_{t+1} . Now, at state s_t , consider the scenario where the rest of the drivers all follow the plan at time t , but driver i deviates and relocates to some location $b \in \mathcal{L}$. Denote the state of the platform at time $t + 1$ as $\tilde{s}_{t+1} \triangleq (s_{-i,t+1}, (1, b, t + \delta(a, b)))$ —the states of the rest of the drivers are the same as the case if all drivers follow the plan, and driver i will be available at location b at time $t + \delta(a, b)$. Driver i is not paid at time t , but incurs cost $c_{a,b,t}$ from driving toward b . The mechanism replans at time $t + 1$, and from time $t + 1$ onward, driver i 's total utility under σ^* would be $\Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1})$, the welfare gain from replicating the driver at $(b, t + \delta(a, b))$, computed at time $t + 1$ given state \tilde{s}_{t+1} . We prove $\Phi_{a,t}^{(t')}(s_{t'}) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1}) - c_{a,b,t}$ by showing:

$$(i) \quad \Phi_{a,t}^{(t')}(s_{t'}) \geq \Phi_{a,t}^{(t)}(s_t),$$

$$(ii) \quad \Phi_{a,t}^{(t)}(s_t) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) - c_{a,b,t} \text{ for all } b \in \mathcal{L} \text{ s.t. } t + \delta(a, b) \leq T, \text{ and}$$

$$(iii) \quad \Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1}).$$

Part (i): $\Phi_{a,t}^{(t')}(s_{t'}) \geq \Phi_{a,t}^{(t)}(s_t)$. The inequality trivially holds if $t = t'$. Consider $t > t'$. The highest achievable welfare at state $s_{t'}$ with an additional driver at (a, t) is weakly higher than the welfare of the scenario where all drivers follow the original plan until time t , and then optimize at time t with all the existing drivers (whose states are now s_t) and the additional driver at (a, t) :

$$W(\mathcal{D}^{(t')}(s_{t'}) \cup \{(1, a, t - t', T - t')\}) \geq \left[W(\mathcal{D}^{(t')}(s_{t'})) - W(\mathcal{D}^{(t)}(s_t)) \right] + W(\mathcal{D}^{(t)}(s_t) \cup \{(1, a, 0, T - t)\}).$$

Here, $(1, a, 0, T - t) = (1, a, t - t, T - t)$ is the type of the additional driver entering at (a, t) , in the time-shifted economy starting from s_t . The unit-replica welfare gain is therefore

$$\begin{aligned} \Phi_{a,t}^{(t')}(s_{t'}) &= W(\mathcal{D}^{(t')}(s_{t'}) \cup \{(1, a, t - t', T - t')\}) - W(\mathcal{D}^{(t')}(s_{t'})) \\ &\geq W(\mathcal{D}^{(t)}(s_t) \cup \{(1, a, 0, T - t)\}) - W(\mathcal{D}^{(t)}(s_t)) \\ &= \Phi_{a,t}^{(t)}(s_t). \end{aligned}$$

Part (ii): $\Phi_{a,t}^{(t)}(s_t) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) - c_{a,b,t}$ for all $b \in \mathcal{L}$ s.t. $t + \delta(a, b) \leq T$. This is similar to part (i), observing that at state s_t , the additional driver at (a, t) can relocate to b at a cost of $c_{a,b,t}$ while the rest of the drivers follow the original plan at time t , and then optimize at time $t + 1$:

$$\begin{aligned} &W(\mathcal{D}^{(t)}(s_t) \cup \{(1, a, 0, T - t)\}) \\ &\geq \left[W(\mathcal{D}^{(t)}(s_t)) - W(\mathcal{D}^{(t+1)}(s_{t+1})) \right] - c_{a,b,t} + W(\mathcal{D}^{(t+1)}(s_{t+1}) \cup \{(1, b, \delta(a, b) - 1, T - (t + 1))\}). \end{aligned}$$

Here, $(1, b, \delta(a, b) - 1, T - (t + 1)) = (1, b, t + \delta(a, b) - (t + 1), T - (t + 1))$ is the type of the additional driver at $(b, t + \delta(a, b))$, time shifted by $t + 1$. This gives us:

$$\begin{aligned}\Phi_{a,t}^{(t)}(s_t) &= W(\mathcal{D}^{(t)}(s_t) \cup \{(1, a, 0, T - t)\}) - W(\mathcal{D}^{(t)}(s_t)) \\ &\geq W(\mathcal{D}^{(t+1)}(s_{t+1}) \cup \{(1, b, \delta(a, b) - 1, T - (t + 1))\}) - W(\mathcal{D}^{(t+1)}(s_{t+1})) - c_{a,b,t} \\ &= \Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) - c_{a,b,t}.\end{aligned}$$

Part (iii): $\Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1})$. First, observe that the only possible difference between s_{t+1} and \tilde{s}_{t+1} is the state of driver i . Fixing the state of the rest of the riders as $s_{-i,t+1}$, $\Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1})$ is the welfare gain from adding an additional driver at $(b, t + \delta(a, b))$ where driver i is at $s_{i,t+1}$ (the state of driver i if she followed the dispatch at time t), whereas $\Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1})$ is the welfare gain from adding an additional driver at $(b, t + \delta(a, b))$ where driver i is at $(b, t + \delta(a, b))$ (the state of driver i that had deviated while the replan happens at time $t + 1$).

When $s_{i,t+1} = (b, t + \delta(a, b))$ (i.e. the driver's deviation resulted in the same future state at time $t + 1$ as in the scenario that she didn't deviate, e.g. instead of picking up rider j who travels to d_j , the driver relocates with an empty car to d_j instead), the inequality trivially holds. When $s_{i,t+1} = \phi$, i.e. when the driver is asked to exit (or not enter) at time t , then $\Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1})$ is implied by the fact that drivers are substitutes, i.e. the more drivers there are, the smaller the marginal welfare contribution of each driver. When $s_{i,t+1} \neq (b, t + \delta(a, b))$ and when $s_{i,t+1} \neq \phi$, intuitively, the marginal value of an available driver when there is another available driver at the same location is smaller than the marginal value of an available driver when the existing available driver at some other location, i.e. there is stronger substitution among drivers at the same locations, in comparison to that among drivers at different locations.

More formally, let $\xi^* \in \mathbb{Z}^{|\mathcal{D}|+|\mathcal{L}|(T+1)}$ be the vector of sources of driver flow given state $s_{-i,t+1}$, s.t. for all $a' \in \mathcal{L}$, for all $t' \in [T]$, let

$$\begin{aligned}\xi_{a',t'}^* &= \sum_{i' \neq i} \mathbb{1}\{s_{i',t+1} = (1, a', t')\} + \sum_{i' \neq i} \mathbb{1}\{s_{i',t+1} = (a'', a', t''), t'' + \delta(a'', a') = t'\} \\ &\quad + \sum_{i' \neq i} \mathbb{1}\{s_{i',t+1} = (o_j, d_j, \tau_j, j), d_j = a', \tau_j + \delta(o_j, d_j) = t'\},\end{aligned}$$

and for each driver $i' \neq i$, let $\xi_{D_{i'}}^* = \mathbb{1}\{\tau_i \geq t + 1\}$. Intuitively, $\xi_{a',t'}^*$ is the number of drivers in $\mathcal{D} \setminus \{i\}$ who are in the platform and available at (a', t') , plus the number of drivers who are en-route relocating to (a', t') , plus the number of drivers who are driving a rider to (a', t') .

Let $\omega(\xi)$ be the objective of the flow LP (34) where the flow boundary condition is given by ξ , and χ_n be the vector of all zeros but a single 1 for the entry corresponding to node n . If driver i is dispatched to exit (or not enter) the platform at time t , i.e. $s_{i,t+1} = \phi$, the desired property $\Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1})$ is equivalent to

$$\omega(\xi^* + \chi_{(b,t+\delta(a,b))}) - \omega(\xi^*) \geq \omega(\xi^* + 2\chi_{(b,t+\delta(a,b))}) - \omega(\xi^* + \chi_{b,t+\delta(a,b)}).$$

This identity corresponds to the first local exchange property of M^{\natural} concave functions (equation (4.5) of Theorem 4.1 in [28]), and that the objective function of MCF problems is M^{\natural} concave (see Example 5 in Section 4.6 of [28]).²² The objective of the flow problem there is defined as a function

²²See [28] for a general introduction of M^{\natural} concavity, and also Chapter 9 of [27] for the related properties of the objectives of network flow problems.

of the “sink” nodes in the flow graph, however, the roles of sinks and sources are symmetric: our MCF problem can also be formulated as having a source node at the end of time, where edges go back in time, and the node corresponding to the entering location/time of each driver sinks at most one unit of flow.

Finally, for the case where driver i is not dispatched to exit the platform at time t , let (a', t') be the location and time where driver i will become available again if she followed the dispatch.

The identity that we need to prove $\Phi_{b,t+\delta(a,b)}^{(t+1)}(s_{t+1}) \geq \Phi_{b,t+\delta(a,b)}^{(t+1)}(\tilde{s}_{t+1})$ can now be written as

$$\omega(\xi^* + \chi(a', t') + \chi(b, t+\delta(a,b))) - \omega(\xi^* + \chi(a', t')) \geq \omega(\xi^* + 2\chi(b, t+\delta(a,b))) - \omega(\xi^* + \chi(b, t+\delta(a,b))).$$

This corresponds to the third local exchange property of M^\natural concave functions (equation (4.7) of Theorem 4.1 in [28]), which intuitively means that there is stronger substitution among drivers at the same location and time, in comparison to drivers that are at different locations and time. This completes the proof of the theorem. \square

C Additional Discussions and Examples

We provide in this section additional examples and discussions omitted from the body of the paper.

C.1 The Driver-Optimal Mechanism

A natural variation on the STP mechanism is to consider the driver-optimal analog, 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 following example shows, however, that the driver-optimal mechanism is not incentive compatible. This is because as time progresses, the set of paths that are available to the drivers shrinks, thus the welfare loss from losing some driver may increase. Because of this, it may sometimes be profitable for such drivers to deviate, trigger the replanning and get higher total payments in subsequent time periods.

Example 5 (Driver optimal mechanism is not IC). Consider the economy illustrated in Figure 19 with three locations, three time periods and symmetric distances $\delta(A, A) = \delta(B, B) = \delta(C, C) = \delta(B, C) = 1$ and $\delta(A, B) = \delta(A, C) = 2$. Assume all trip costs and early exit costs are zero. Two drivers enter the platform at time 0 at location B , and three riders have types:

- Rider 1: $o_1 = C$, $d_1 = C$, $\tau_1 = 1$, $v_1 = 1$,
- Rider 2: $o_2 = C$, $d_2 = C$, $\tau_2 = 2$, $v_2 = 5$,
- Rider 3: $o_3 = A$, $d_3 = A$, $\tau_3 = 1$, $v_3 = 1$.

In a welfare-optimal dispatching as shown in Figure 19, driver 1 is dispatched to take the path $z_1 = ((B, C, 0), (C, C, 1), (C, C, 2))$ and to pick up riders 1 and 2. Driver 2 takes the path $z_2 = ((B, A, 0), (A, A, 2))$ and picks up rider 3. One driver-optimal CE plan sets anonymous trip prices $p_{C,C,1} = 0$ and $p_{C,C,2} = p_{A,A,2} = 1$, so that the total utility of each driver equal to the welfare-loss of 1, if she was eliminated from the economy.

Assume that driver 2 follows the mechanism and starts to drive toward A at time 0, we show a useful deviation of driver 1 by rejecting the dispatched relocation to C and staying in location B to trigger a replanning at time 1. At time 1, driver 2 is already *en route* to A thus is only able to pick up rider 3 at time 2. Driver 1 would be asked to relocate to C and pick up rider 2. The price $p_{C,C,2}$ in the updated driver-optimal CE plan is 5, the welfare loss if the economy at time 1 loses driver 1. This is higher than driver 1’s payment from following the dispatches at all times. \square




Figure 19: Illustration of the toy economy II in Example 5 with three locations A, B, C , three time periods, two drivers starting at $(B, 0)$ and three riders with values 1, 5 and 1, respectively. Under a welfare optimal plan, driver 1 picks up riders 1 and 2 and driver 2 picks up rider 3.

A variation on the driver-optimal mechanism where drivers' payments are shifted in time is equivalent to the *dynamic VCG mechanism* [7, 13]. The dynamic VCG mechanism is not incentive compatible because the existence of a driver at some certain time may exert negative externality on the economy, in which case the payment to the driver would be negative for that time period. The driver would have incentives to decline the dispatch and avoid such payment. See the discussions and examples in Appendix D.1.

C.2 LP Integrality and Existence of CE

We show via the following two examples that when either of the assumptions (S1) and (S2) is violated, the LP relaxation (9) of the ILP (4) may no longer be integral, and that welfare-optimal competitive equilibrium outcomes as defined in Definition 3 may not exist. We first examine the case where drivers may have different times of exiting the platform.

Example 6 (Different driver exit times). Consider the economy as shown in Figure 20 with three locations $\mathcal{L} = \{A, B, C\}$ and three time periods. The distances are symmetric and given by $\delta(A, A) = \delta(B, B) = \delta(C, C) = \delta(A, B) = \delta(B, C) = 1$, and $\delta(A, C) = 2$, and assume all trip costs and exit costs are zero. There are three drivers, entering and exiting at:

- $\ell_1 = A, \underline{\tau}_1 = 0, \bar{\tau}_1 = 3,$
- $\ell_2 = B, \underline{\tau}_2 = 0, \bar{\tau}_2 = 2,$
- $\ell_3 = B, \underline{\tau}_3 = 1, \bar{\tau}_3 = 3,$

and there are six riders with types:

- $o_1 = A, d_1 = C, \tau_1 = 0, v_1 = 5,$
- $o_2 = A, d_2 = B, \tau_2 = 1, v_2 = 7,$
- $o_3 = A, d_3 = B, \tau_3 = 1, v_3 = 1,$
- $o_4 = B, d_4 = A, \tau_4 = 1, v_4 = 2,$
- $o_5 = B, d_5 = A, \tau_5 = 1, v_5 = 5,$
- $o_6 = B, d_6 = A, \tau_6 = 2, v_6 = 4.$

In the unique optimal integral solution, Driver 1 takes the path $z_1^* = ((A, A, 0), (A, B, 1), (B, A, 2))$ and picks up riders 2 and 6. Driver 2 takes the path $z_2^* = ((B, B, 0), (B, A, 1))$ and picks up rider 4. Driver 3 takes the path $z_3^* = ((B, A, 1), (A, A, 2))$ and picks up rider 5. The total social welfare is $v_2 + v_6 + v_4 + v_5 = 18$. The optimal solution of the LP, however, is not integral. Each driver i takes each of their two paths z_i and z'_i with probably 0.5:

- $z_1 = ((A, C, 0), (C, C, 2)), z'_1 = ((A, A, 0), (A, B, 1), (B, A, 2)),$
- $z_2 = ((B, B, 0), (B, A, 1)), z'_2 = ((B, A, 0), (A, B, 1)),$




Figure 20: The economy in Example 6 with three locations A , B , C , three time periods, 6 riders, three drivers starting at $(A, 0)$, $(B, 0)$ and $(B, 1)$, where driver 2 exits the platform at time 2.

- $z_3 = ((B, A, 1), (A, A, 2))$, $z'_3 = ((B, B, 1), (B, A, 2))$.

The riders 2, 5 and 6 are picked up with probability 1, whereas rider 1 is picked up with probability 0.5. The total social welfare is $0.5v_1 + v_2 + v_5 + v_6 = 18.5 > 18$. There is a unique solution to the dual LP (15), which implies anonymous trip prices of $p_{A,C,0} = 5$ and $p_{A,B,1} = p_{B,A,1} = p_{B,A,2} = 2.5$. These prices do not support the optimal integral solution, since rider 4 is willing to pay only $v_2 = 2$ but is picked up and charged 2.5.

Moreover, we show that no anonymous origin-destination prices support the optimal integral dispatch in competitive equilibrium.²³ First, rider 1 with value 5 is not picked up, therefore the price for the $(A, C, 0)$ trip needs to be at least $p_{A,C,0} \geq 5$. In order for driver 1 to not regret not taking this trip, the prices for her trips need to be at least $p_{A,B,1} + p_{B,A,2} \geq 5$.

Since rider 4 with value 2 is picked up, the price for trip $(B, A, 1)$ can be at most 2. As a consequence, the price for the trip $(A, B, 1)$ cannot exceed 2 either, since otherwise, driver 2 would have incentive to take the path $((B, A, 0), (A, B, 1))$ instead. This implies that the price for the trip $(B, A, 2)$ needs to be at least 3. Note that driver 3 now prefers taking the path $((B, B, 1), (B, A, 2))$ and get paid at least 3, in comparison to the dispatched trip $((B, A, 1), (A, A, 2))$ and gets paid at most 2. This is a contradiction, and shows that no anonymous OD price supports the welfare-optimal outcome in competitive equilibrium. \square

The reason integrality fails is that the ridesharing problem can no longer be reduced to an MCF problem in the way that we discuss in Appendix B.2.2 without loss of generality. In the standard MCF problem, there is a single type of flow flowing through the network, and the optimal flow is guaranteed to be integral. When drivers have different exiting times, if all units of flow are still treated as homogeneous, the resulting decomposed flow may not send the correct drivers to leave at the correct times. As an example, the optimal homogeneous flow with the same boundary condition in this example can be decomposed into the following three paths: $((A, C, 0))$, $((B, A, 0), (A, B, 1), (B, A, 2))$, $((B, A, 1), (A, A, 2))$ with a total social welfare of 21. However, it cannot be implemented since it is driver 2 who enters at $(B, 0)$ and needs to exit at time 2, but in this decomposition, the flow that corresponds to driver 1 exits at time 2.

When drivers have different exiting times, the MCF problem has heterogeneous flows. Similarly, we can construct examples showing that the optimal solutions to the ILP and LP do not coincide if drivers have preference over which location to they end up with at the end of the planning horizon, unless all drivers start at the same time and location.

²³In general, the non-existence of CE does not imply that there do not exist dynamic ridesharing mechanisms that are SPIC, since a mechanism determining a CE plan is not necessary for the mechanism to be incentive compatible.




Figure 21: The economy in Example 7 with three locations A, B, C , three time periods, 2 riders, and one driver entering the platform at $(B, 0)$.

The following example examines the case where riders are patient and may be willing to wait.

Example 7 (Patient riders). Consider the economy as illustrated in Figure 21 with three locations $\mathcal{L} = \{A, B, C\}$ and three time periods. The pairwise distances are symmetric and given by $\delta(A, A) = \delta(B, B) = \delta(C, C) = \delta(B, C) = 1$ whereas $\delta(A, B) = \delta(A, C) = 2$. Assume all trip costs and early exit costs are zero. There is a single driver entering at location B at time 0 who would stay until the end of the planning horizon. There are two riders. Rider 1 is impatient, requesting a trip at time 0 from B to A , and is willing to pay $v_1 = 9$. Rider 2 is willing to pay $v_2 = 5$ for a trip from B to C at time 0, but is willing to wait for at most two time periods.

In the optimal integral solution, driver 1 takes the path $((B, A, 0), (A, A, 2))$, picks up rider 1, and achieves total social welfare of $v_1 = 9$. In the optimal solution to the LP (9), however, the driver takes the paths $((B, A, 0), (A, A, 2))$ and $((B, C, 0), (C, B, 1), (B, C, 2))$, each with probability 0.5. The path $((B, C, 0), (C, B, 1), (B, C, 2))$ seemingly have a total value of 10, therefore the objective of the LP would be $10 \times 0.5 + 9 \times 0.5 = 9.5 > 9$. The optimal integral solution is not supported by anonymous OD prices in CE either— since rider 2 is not picked up, the prices for the (BC) trips starting at times 0, 1 and 2 need to be at least 5. Thus the total payment for the path $((B, C, 0), (C, B, 1), (B, C, 2))$ is at least 10, however, the driver’s dispatched trip $(B, A, 0)$ pays at most $v_2 = 9$. \square

Similar to the case when drivers have different exit times, integrality fails with patient riders also because there is no direct way of reducing the problem to an integral MCF problem without loss. In the MCF problem, each rider corresponds to a single edge in the flow graph with edge cost equal to the trip cost minus the rider’s value. If the rider is patient, there may be multiple edges that correspond to the same rider, and there is no easy way expressing the constraint that a rider cannot be picked up more than once without breaking the integrality of the MCF problem.

C.3 Rider Incentives

The following example illustrates (i) the trip-prices and rider-utilities under all CE outcomes may not have lattice structure and (ii) the rider-side VCG prices do not coincide with the prices in the driver-pessimal CE plan, and (iii) no welfare-optimal CE mechanism, including the STP mechanism, incentivizes riders to truthfully report their values.

Example 8. Consider the economy in Figure 22, where all trip costs and early exit costs are assumed to be zero. Driver 1 enters at location A and time 0, and stays until the end of the planning horizon. Under the welfare-optimal dispatching, the driver takes the path $((A, A, 0), (A, A, 1))$ and




Figure 22: The economy in Example 8, with two locations A , B , two time periods and three riders.

picks up riders 1 and 2, achieving social welfare $v_1 + v_2 = 11$. In the driver pessimal CE plan, the prices for the trips are be $p_{A,B,0} = 8$, and $p_{A,A,0} + p_{A,A,1} = 8$, therefore $p_{A,A,0} = 5$, $p_{A,A,1} = 3$ and $p_{A,A,0} = 2$ and $p_{A,A,1} = 6$ would both support the driver-pessimal CE outcome.

Lattice Structure: It is easy to check that the lowest prices for the trips $(A, A, 0)$ and $(A, A, 1)$ under all CE outcomes are be 2 and 3 respectively. However, setting $p_{A,A,0} = 2$ and $p_{A,A,1} = 3$ would not form a CE, since rider 3 is willing to pay 8, thus $p_{A,B,0} \geq 8$ and this violates driver best-response. This implies that trip prices under all CE outcomes do not form a lattice. As a consequence, riders' utilities under all CE outcomes do not form a lattice either.

Rider-side VCG Prices: Moreover, we can check that $p_{A,A,0} = 2$ is what rider 1 should be charged under the rider-side VCG payment rule: if rider 1 is not present, rider 3 gets picked up, thus the total welfare for the rest of the economy increases from $v_2 = 6$ to $v_3 = 8$. Similarly, rider 2's VCG payment would be $p_{A,A,1} = 3$. This shows that the VCG payment on the rider side may not support a welfare-optimal outcome in CE.

Rider-side IC: This example also implies that the STP mechanism is not incentive compatible on the rider's side. Under any driver-pessimal outcome, we know that one of the riders 1 and 2 would be charged a payment that is higher than their VCG price. A simple analysis would show that if the rider who is charged higher than the VCG price reports the VCG price as her value, then her payment under the STP mechanism would be exactly her VCG price. This is a useful deviation. More generally, this shows that no welfare-optimal CE outcome would be incentive compatible on the rider's side, since $p_{A,A,0} + p_{A,A,1} \geq 8$ under any CE outcome. \square

It is not a coincidence that the lowest possible prices for each rider under all CE outcomes is equal to their rider-side VCG prices. The following theorem shows that the minimum CE prices and the rider-side VCG prices always coincide.

Theorem 4 (Minimum CE = rider-side VCG). *For any rider that is picked up in some welfare-optimal dispatching, her rider-side VCG price is equal to the minimum price for her trip among all CE outcomes.*

Proof. For simplicity of notation, assume that driver $j \in \mathcal{R}$ requests the trip (a, b, t) , has value v_j , and is picked up under some welfare-optimal dispatching. We are going to prove:

- (i) the price $p_{a,b,t}$ under any CE outcome is at least the rider-side VCG payment for rider j , and
- (ii) there exists an CE outcome where rider j 's trip price is at most her rider-side VCG payment.

Combining (i) and (ii), we know that the rider-side VCG payment has to be the lowest CE price for the trip among all CE outcomes.

Let $\tilde{W}(\mathcal{D}, \mathcal{R})$ be the optimal welfare achieved by the set of drivers \mathcal{D} and the set of riders \mathcal{R} , i.e. the optimal objective of (4). Moreover, we denote $\tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R})$ as the optimal objective of (4) if the trip capacity constraint (5) for this specific trip is relaxed by 1, i.e. where there is an additional driver that is only able to complete an (a, b, t) trip.

Similarly, denote $W(\mathcal{D}, \mathcal{R})$ as the optimal objective of (9) and $W(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R})$ as the optimal objective of (9) with an extra (a, b, t) trip capacity. From the integrality of the LP (9) under (S1) and (S2), we know that $\tilde{W}(\mathcal{D}, \mathcal{R}) = W(\mathcal{D}, \mathcal{R})$, however, we only know $\tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) \leq W(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R})$ since with the additional capacity of 1 for the (a, b, t) trip, it is not obvious whether the LP relaxation would still be integral.

Part (i): To prove (i), first observe that with the same argument on subgradients as in the proof of Lemma 4, we can show that under any CE outcome, the price $p_{a,b,t}$ as the subgradient w.r.t. the RHS of capacity constraint (10) in the LP (9), is lower bounded by the welfare gain from relaxing the capacity constraint by 1, i.e. $p_{a,b,t} \geq W(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) - W(\mathcal{D}, \mathcal{R})$. This implies that $p_{a,b,t} \geq \tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) - \tilde{W}(\mathcal{D}, \mathcal{R})$, i.e. any CE price must be at least the welfare contribution of an additional (a, b, t) trip to the original economy at no cost.

What is left to prove is that $\tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) - \tilde{W}(\mathcal{D}, \mathcal{R}) \geq p_{a,b,t}^{vcg}$, where $p_{a,b,t}^{vcg} = \tilde{W}(\mathcal{D}, \mathcal{R} \setminus \{j\}) - (\tilde{W}(\mathcal{D}, \mathcal{R}) - v_j)$, i.e. the optimal welfare in the economy without rider j minus the welfare of the rest of the riders in the economy with rider j . This holds since:

$$\begin{aligned} & \tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) - \tilde{W}(\mathcal{D}, \mathcal{R}) - p_{a,b,t}^{vcg} \\ &= \tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) - \tilde{W}(\mathcal{D}, \mathcal{R}) - (\tilde{W}(\mathcal{D}, \mathcal{R} \setminus \{j\}) - (\tilde{W}(\mathcal{D}, \mathcal{R}) - v_j)) \\ &= \tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R}) - (\tilde{W}(\mathcal{D}, \mathcal{R} \setminus \{j\}) + v_j) \\ &\geq 0. \end{aligned}$$

The last inequality holds because $\tilde{W}(\mathcal{D} \cup \{(a, b, t)\}, \mathcal{R})$, the optimal welfare from adding both a trip (a, b, t) (at no cost) and a rider j to the economy $(\mathcal{D}, \mathcal{R} \setminus \{j\})$, is weakly higher than assigning the trip (a, b, t) to rider j and keeping the plan for the rest of the economy unchanged.

This completes the proof of part (i), that any CE price is weakly higher than the VCG payment.

Part (ii): Given \mathcal{D} and \mathcal{R} , we construct an alternative economy $E' = (\mathcal{D}, \mathcal{R}')$ where \mathcal{R}' and \mathcal{R} coincide, except for the value of rider j : instead of having value v_j , we change her value to her VCG payment in the original economy, i.e.

$$v'_j = p_{a,b,t}^{vcg} = \tilde{W}(\mathcal{D}, \mathcal{R} \setminus \{j\}) - (\tilde{W}(\mathcal{D}, \mathcal{R}) - v_j).$$

Now consider the optimal dispatching in the economy E' . If rider j' is not picked up, the optimal welfare is equal to $\tilde{W}(\mathcal{D}, \mathcal{R} \setminus \{j\})$, the highest welfare achievable for the rest of the economy. If rider j' is picked up, the highest achievable welfare for the rest of the economy is equal to $(\tilde{W}(\mathcal{D}, \mathcal{R}) - v_j)$, therefore the total welfare is $v'_j + (\tilde{W}(\mathcal{D}, \mathcal{R}) - v_j) = \tilde{W}(\mathcal{D}, \mathcal{R} \setminus \{j\})$. This implies that in at least one of the optimal plans in economy E' , rider j' is picked up. Let (x', y') be an optimal dispatching in economy E' where rider j' is picked up, and let p' be any CE prices. First, $p'_{a,b,t} \leq v'_j = p_{a,b,t}^{vcg}$ holds, since the outcome forms a CE and rider j' is picked up and must have non-negative utility.

We also claim that the plan with anonymous prices (x', y', p') also forms a CE in the original economy. Since (x', y', p') is a CE in economy E' , we know that for drivers, the dispatched paths under y' gives them the highest total payment given prices p' . We also know that trips with excessive supply have zero prices. For any rider other than j , her values in E' and E are the same thus rider best-response holds. For rider j , her value in E is $v_j \geq p_{a,b,t}^{vcg} = v'_j \geq p'_{a,b,t}$ thus she weakly prefers getting picked up $x'_j = 1$ and is also best-responding. This shows that there exists a




Figure 23: The economy in Example 9, with two locations A, B , two time periods and four riders.

CE outcome in E where the price for the (a, b, t) trip is at most v'_j , rider j 's VCG payment. This completes the proof of part (ii), and also the theorem. \square

Finally, we show via the following example that a mechanism that always computes an welfare-optimal dispatching together with rider-side VCG prices (at the beginning of the planning horizon, and also after any driver deviation), is not SPIC for the drivers. This is not implied by Theorem 4 since a mechanism's plan forming a CE is not necessary for a mechanism being SPIC for drivers.

Example 9. Consider the economy in Figure 23, with $T = 2$ and two locations $\mathcal{L} = \{A, B\}$ with unit distances: $\delta(a, b) = 1, \forall a, b \in \mathcal{L}$. Assume all trip costs and exit costs are zero. There is one driver entering at time $\tau_1 = 0$ at location $\ell_1 = A$ and leaves at time $\bar{\tau}_1 = 2$. There are four riders:

- Rider 1: $o_1 = A, d_1 = A, \tau_1 = 0, v_1 = 5,$
- Rider 2: $o_2 = A, d_2 = A, \tau_2 = 1, v_2 = 6,$
- Rider 3: $o_3 = B, d_3 = B, \tau_3 = 1, v_3 = 8,$
- Rider 4: $o_4 = B, d_4 = B, \tau_4 = 1, v_4 = 8.$

The optimal plan computed at time 0 has driver 1 taking the path $((A, A, 0), (A, A, 1))$ and picking up riders 1 and 2. The rider-side VCG prices for riders 1 and 2 would be 2 and 3 respectively, thus the driver's total payment, if she follows the dispatches at all times, would be 5. Now consider the scenario where the driver relocates to location B at time 0 instead. When time 1 comes, the updated plan would dispatch driver 1 to pick up one of riders 3 or 4, and the updated VCG payment for this trip would be 8. This is a useful deviation, thus the mechanism is not SPIC. \square

C.4 Truthful Reporting of Driver Entrance

Throughout the paper, we assumed a complete information model, where the mechanism knows about the entering location and time for all the drivers. In this section, we discuss the scenario where the location and time where a driver first becomes available to pick up are drivers' private information, and the mechanism needs to ask the drivers to report their entrance information. Here, we still assume that all drivers stay until at least the end of the planning horizon.

Theorem 5. *Under the STP mechanism, for driver i who is available to pick up at location ℓ_i starting at time τ_i , it is not useful for her to report some entrance location and time $(\hat{\ell}_i, \hat{\tau}_i) \in \mathcal{L} \times [T]$ where $\hat{\tau}_i \geq \tau_i + \delta(\ell_i, \hat{\ell}_i)$, and then enter the platform at $(\hat{\ell}_i, \hat{\tau}_i)$.*

Proof. First, observe that for driver i whose true entering location and time is (ℓ_i, τ_i) , the driver is only able to enter at (ℓ_i, τ_i) , or at some $(\hat{\ell}_i, \hat{\tau}_i) \in \mathcal{L} \times [T]$ where $\hat{\tau}_i \geq \tau_i + \delta(\ell_i, \hat{\ell}_i)$. Assume that drivers all follow the SPE once they entered the platform, and always accepts the dispatches of the mechanisms. If driver i reports truthfully, her total payment would be $\Phi_{D_i} = \max\{\Phi_{\ell_i, \tau_i}, 0\}$, the welfare gain of the economy from replicating this driver.

Following the same notation as in the proof of Theorem 2, we use ξ^* to denote the boundary condition (given the initial state) for the flow problem of the economy except for driver i , i.e.

$\xi_{D_{i'}}^* = 1$ for all $i' \neq i$, and $\xi_n^* = 0$ for all other node $n \in \mathcal{N}$. Let $\omega(\cdot)$ be the optimal objective of the corresponding flow problem, Φ_{ℓ_i, τ_i} can be written as:

$$\Phi_{\ell_i, \tau_i} = \omega(\xi^* + 2\chi_{(\ell_i, \tau_i)}) - W(\xi^* + \chi_{(\ell_i, \tau_i)}).$$

Here, $\chi_{(\ell_i, \tau_i)}$ is a $|\mathcal{L}|(T+1) + |\mathcal{D}|$ by 1 vector with all zero entries, except a single 1 at the (ℓ_i, τ_i) entry. If the driver reports $(\hat{\ell}_i, \hat{\tau}_i)$ as her entering location and time, and actually enters at $(\hat{\ell}_i, \hat{\tau}_i)$, her equilibrium payoff for the rest of the planning horizon can be written as

$$\omega(\xi^* + 2\chi_{(\hat{\ell}_i, \hat{\tau}_i)}) - W(\xi^* + \chi_{(\hat{\ell}_i, \hat{\tau}_i)}).$$

Let g be the lowest cost that the driver has to incur, while moving from (ℓ_i, τ_i) to $(\hat{\ell}_i, \hat{\tau}_i)$. By reporting and entering at $(\hat{\ell}_i, \hat{\tau}_i)$, the agent's total utility is at most $\omega(\xi^* + 2\chi_{(\hat{\ell}_i, \hat{\tau}_i)}) - W(\xi^* + \chi_{(\hat{\ell}_i, \hat{\tau}_i)}) - g$. We show that this is not a useful deviation, since

$$\begin{aligned} & W(\xi + 2\chi_{\ell_i, \tau_i}) - W(\xi + \chi_{\ell_i, \tau_i}) - (W(\xi + 2\chi_{\hat{\ell}_i, \hat{\tau}_i}) - W(\xi^* + \chi_{(\hat{\ell}_i, \hat{\tau}_i)}) - g) \\ & \geq W(\xi + 2\chi_{\ell_i, \tau_i}) - W(\xi + \chi_{\ell_i, \tau_i}) - (W(\xi + \chi_{\ell_i, \tau_i} + \chi_{\hat{\ell}_i, \hat{\tau}_i}) - W(\xi + \chi_{\ell_i, \tau_i}) - g) \\ & = W(\xi + 2\chi_{\ell_i, \tau_i}) - W(\xi + \chi_{\ell_i, \tau_i} + \chi_{\hat{\ell}_i, \hat{\tau}_i}) - g \\ & \geq 0. \end{aligned}$$

The first inequality holds due to the local exchange property of the $M^{\mathfrak{h}}$ concave functions (Equation (4.7) in [28]), and the last inequality holds since the highest achievable welfare achievable from two additional drivers at (ℓ_i, τ_i) is (weakly) higher than the scenario where one of these drivers has to move to $(\hat{\ell}_i, \hat{\tau}_i)$ (at the lowest possible cost g). \square

This result on the truthfulness of driver entrance reports, however, considers only the scenario that the driver actually enters at the location and time as she reported. We may also consider a mechanism that takes the drivers' reports of entering location and time, plans accordingly at the beginning of the planning horizon, but replans if any driver's entrance action turns out to be different from expected/reported, without penalizing any driver that had deviated. The following example shows that when allowing arbitrary driver entrance regardless of their report, the STP mechanism does not incentivize the drivers to truthfully report their entering location and time.

Example 10. Consider the economy as shown in Figure 24. The planning horizon is $T = 3$ and there are two locations $\mathcal{L} = \{A, B\}$ with unit distances $\delta(a, b) = 1$ for all $a, b \in \mathcal{L}$. Assume that all trip costs and early exiting costs are zero. There is one driver entering at time $\tau_1 = 0$ at location $\ell_1 = A$ and leaves at time $\bar{\tau}_1 = 3$. There is another driver, whose true entering time and location is $\tau_2 = 2$ and $\ell_2 = B$. There are four riders with type:

- Rider 1: $o_1 = B, d_1 = B, \tau_1 = 1, v_1 = 10,$
- Rider 2: $o_2 = A, d_2 = A, \tau_2 = 1, v_2 = 8,$
- Rider 3: $o_3 = B, d_3 = B, \tau_3 = 2, v_3 = 5,$
- Rider 4: $o_4 = B, d_4 = B, \tau_4 = 2, v_4 = 4.$

Under the STP mechanism, if both drivers report their entrance location and time truthfully, the welfare-optimal plan dispatches driver 1 to take the path $((A, B, 0), (B, B, 1), (B, B, 2))$ and pick up riders 1 and 3. Driver 2 is dispatched to pick up rider 4, and her payment would be 0, the welfare gain from an additional driver entering at $(B, 2)$.

However, if driver 2 reports $(B, 1)$ as her entering time and location, then under the STP mechanism, driver 1 would be dispatched to go to $(A, 1)$ to pick up rider 2 at time 1. When time 1




Figure 24: The economy in Example 10, with two locations, three time periods and four riders.

comes, driver 2 fails to enter at $(B, 1)$, and regardless of any future entrance of driver 2, the optimal plan at time 1 is for driver 1 to pick up rider 2. When time 2 comes, driver 2 can then decide to actually enter at location B . The mechanism would replan again, dispatching driver 2 to pick up rider 3. The welfare gain from an additional driver at $(B, 2)$ would now be 4, thus the driver's new payment would be 4, and this is a useful deviation. \square

C.5 Naive Update of Static Plans

The last example in this section shows that a mechanism that always re-computes a driver-optimal plan at all times is not envy-free for the drivers, and may not be incentive compatible for drivers, depending on how the mechanism breaks ties among different driver-optimal plans.

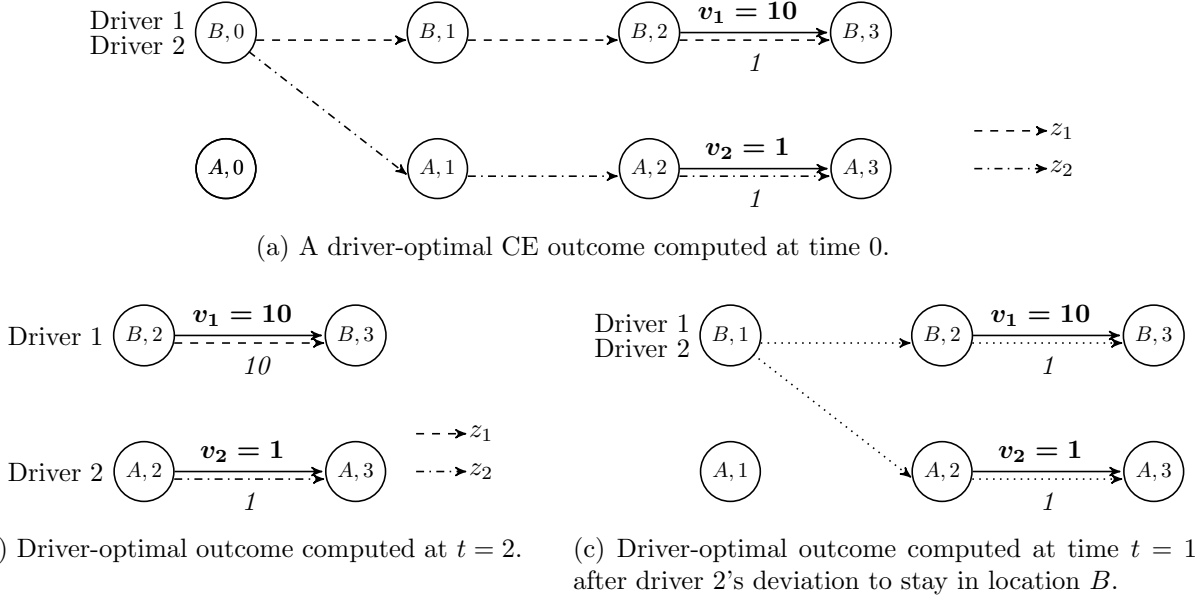


Figure 25: The economy in Example 11, and driver-optimal CE plans computed at different states.

Example 11 (Repeated driver-optimal static CE mechanism). Consider the economy as illustrated in Figure 25 where all costs are zero, and a mechanism that repeatedly computes a driver-optimal CE outcome at all times, regardless of whether any deviation happened. Under the CE outcome computed at time 0 as in Figure 25a, both riders 1 and 2 are picked up, and the prices for the trips are both $p_{B,B,2} = p_{A,A,2} = 1$.

Assume that both drivers follow the mechanism until time 2, the new driver-optimal outcome computed at time 2 is as illustrated in Figure 25b, where the price for trip $(B, B, 2)$ becomes

10, the highest market-clearing price at this state. This shows that the “time 0” plan under the mechanism, i.e. the actual outcome where all drivers follow the mechanism’s dispatch at all times, is not envy-free for drivers, since the total payment to driver 1 is 10, higher than that of driver 2.

Now consider the scenario where driver 2 stayed in location B at time 0 instead of following the dispatch and relocate to A . Under the optimal CE outcome from time 1 onward as in Figure 25c, the two drivers take the paths $((B, B, 1), (B, B, 2))$ and $((B, A, 1), (A, A, 2))$ respectively, and pick up both riders. The IC property of the mechanism now depends on how the mechanism breaks ties among driver-optimal CE plans, but as long as the mechanism dispatches driver 2 to take the path $((B, B, 1), (B, B, 2))$ with non-zero probability (which would be the case if ties are broken uniformly at random), this would be a useful deviation for driver 2. Once driver 2 followed the plan and reach $(B, 2)$ as driver 1 arrived at $(A, 2)$, the newly updated price for the trip $(B, B, 2)$ would again become 10, higher than driver 2’s original payment. \square

D Relation to the Literature

D.1 Dynamic VCG

The dynamic VCG mechanisms [7, 13] truthfully implement efficient decision policies, where agent receive private information over time. Under the dynamic VCG mechanisms, the payment to agent i in each period is equal to the flow marginal externality imposed on the other agents by its presence in this period only [13].

The dynamic VCG mechanism can be adapted for the ridesharing problem, where there is no uncertainty in the transition of states (the actions taken by all drivers at time t fully determines the state of the platform at time $t + 1$) and no private information from the drivers’ side (the location of the driver is known to the mechanism and we assume homogeneous driver costs and no location preferences). We actually show that a variation of the driver-optimal dynamic mechanism that we discussed in Section 4, where driver payments are “shifted” over time, is equivalent to the dynamic VCG mechanism.

The dynamic VCG mechanism for ridesharing, however, fails to be incentive compatible, since some drivers may be paid a negative payment for certain periods of time, and the drivers would have incentive to decline the dispatch at such times to avoid making the payments. This is because the existence of some driver for only one period of time may exert negative externality on the rest of the economy by inducing seemingly efficient actions that result in suboptimal positioning of the rest of the drivers in the subsequent time periods.

We illustrate this via analyzing the economy introduced in Example 5, as shown in Figure 19.

Example 6 (Continued). Without driver 1, driver 2 would be dispatched to pick up riders 1 and 2 and achieve a total welfare of 6. With driver 1, one welfare-optimal dispatching plan sends driver 1 to take the path $((B, C, 0), (C, C, 1), (C, C, 2))$ and sends driver 2 to take the path $((B, A, 0), (A, A, 2))$.

At time 0, driver 1 takes trip $(B, C, 0)$ and driver 2 takes trip $(B, A, 0)$. Driver 1 contributes 0 to welfare at time 0 since she did not pick up any driver. When time 1 comes, if driver 1 appears for only one period of time, the optimal welfare achieved by the rest of the economy would only be 1— driver 1 disappears and driver 2 picked up rider 3. Therefore, the payment to driver 1 at time 0 would be -5 , since exerted a negative externality of 5 on the rest of the economy by appearing only at time 0. Similarly, we can compute that the payment to driver 1 at times 1 and 2 would be 1 and 5 respectively, giving her a total payment of $-5 + 1 + 5 = 1$.

Now consider the scenario where driver 1 declines the dispatch, refuses to make the payment and stays in location B , and assume that driver 2 still followed the mechanism and drove to location A . When time 1 comes, driver 1 would again be dispatched to drive to C at time 1 and pick up rider 2 at time 2. We can check that the payment to driver 1 at time 1 would be 0, and the payment to driver 1 at time 2 would be 5— the amount the rest of the economy gains from the existence of driver 2 at that time. This is a useful deviation, thus the dynamic VCG mechanism where drivers are allowed to freely decline trips is not IC. \square

D.2 Trading Networks

The literature on trading networks studies economic models where agents in a network can trade via bilateral contracts [24, 25, 31]. Efficient, competitive equilibrium outcomes exist when agents’ valuation functions satisfy the “full substitution” property, and the utilities of agents on either end of the trading network form lattices.

Assume that all trip costs and early exit costs are zero, the optimal dispatching problem of ridesharing platforms can be formulated as a trading network problem in the following way:

- For each driver or rider, there is a node in the network.
- For each driver $i \in \mathcal{D}$ and each rider $j \in \mathcal{R}$, there is an edge from i to j if $\tau_j \geq \tau_i + \delta(\ell_i, o_j)$, i.e. driver i is able to pick up rider j if she drives directly to o_j after entering.
- For any two riders j and j' in \mathcal{D} , there is an edge from j to j' if (i) $\tau_j + \delta(o_j, d_j) + \delta(d_j, o_{j'}) \leq \tau_{j'}$ assuming $d_j \neq o_{j'}$ or (ii) $\tau_j + \delta(o_j, d_j) \leq \tau_{j'}$ if $d_j = o_{j'}$. Intuitively, riders j can trade to rider j' if a driver is still able to pick up rider j' after dropping off rider j .

What is being traded in the network is the right to use the car over the rest of the planning horizon. Each driver is able to trade with at most one rider. A driver’s utilities is zero if she does not trade, and her utility is equal to the her payment if she did trade. Each rider values buying the right to use at most one car, and values it at v_j . If she did buy the right to use one car, her utility is v_j minus the price she paid to buy the right to use the car, plus the payment she collected from the rider who bought the right to use the car from her. Riders that did not buy a car cannot sell (values such contracts at $-\infty$).


With existing results in the trading network literature, we show the existence of welfare-optimal, competitive equilibrium outcomes, and the lattice structure of drivers’ total payments under all CE outcomes. This does not solve our problem, since there isn’t language in the trading network literature that describes the temporal evolvement of the network structure and the corresponding incentive issues— as time progresses, the set of reachable riders for each driver decreases, thus the network becomes sparser.

E Additional Simulation Results


We present in this section the additional simulation results that are omitted from the body of the paper.

E.1 Morning Rush Hour

Figures 26 and 27 show the average number of drivers and average prices for each of the five trips that are not analyzed in Section 5.2 for the morning rush hour scenario.




(a) The STP mechanism.




(b) The myopic pricing mechanism.

Figure 26: Additional comparison of the number of drivers per trip for the morning rush hour.

The STP mechanism dispatches a reasonably high number of drivers to the (A, C) trip since there is a high demand for drivers at C (see Figure 26). In contrast, though the myopic pricing mechanism is not sending too many drivers from C to C or A , many drivers linger around B due to the excessive supply, and the mechanism did not relocate more driver from A to C than from A to A , despite the imbalance in demand in these locations. Prices as shown in Figure 27 are also intuitive and easy to interpret.



(a) The STP mechanism.



(b) The myopic pricing mechanism.

Figure 27: Additional comparison trip prices for the morning rush hour.