



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Surge Pricing and Its Spatial Supply Response

Omar Besbes, Francisco Castro, Ilan Lobel

To cite this article:

Omar Besbes, Francisco Castro, Ilan Lobel (2021) Surge Pricing and Its Spatial Supply Response. Management Science 67(3):1350-1367. <https://doi.org/10.1287/mnsc.2020.3622>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Surge Pricing and Its Spatial Supply Response

Omar Besbes,^a Francisco Castro,^b Ilan Lobel^c

^a Columbia Business School, Columbia University, New York, New York 10027; ^b Anderson School of Management, University of California, Los Angeles, Los Angeles, California 90095; ^c Stern School of Business, New York University, New York, New York 10012

Contact: ob2105@columbia.edu,  <https://orcid.org/0000-0002-2982-3794> (OB); francisco.castro@anderson.ucla.edu,  <https://orcid.org/0000-0003-0766-3491> (FC); ilobel@stern.nyu.edu,  <https://orcid.org/0000-0002-5396-8117> (IL)

Received: March 21, 2018

Revised: May 17, 2019; January 25, 2020

Accepted: February 8, 2020

Published Online in Articles in Advance:
October 8, 2020

<https://doi.org/10.1287/mnsc.2020.3622>

Copyright: © 2020 INFORMS

Abstract. We consider the pricing problem faced by a revenue-maximizing platform matching price-sensitive customers to flexible supply units within a geographic area. This can be interpreted as the problem faced in the short term by a ride-hailing platform. We propose a two-dimensional framework in which a platform selects prices for different locations and drivers respond by choosing where to relocate, in equilibrium, based on prices, travel costs, and driver congestion levels. The platform's problem is an infinite-dimensional optimization problem with equilibrium constraints. We elucidate structural properties of supply equilibria and the corresponding utilities that emerge and establish a form of spatial decomposition, which allows us to localize the analysis to regions of movement. In turn, uncovering an appropriate knapsack structure to the platform's problem, we establish a crisp local characterization of the optimal prices and the corresponding supply response. In the optimal solution, the platform applies different treatments to different locations. In some locations, prices are set so that supply and demand are perfectly matched; overcongestion is induced in other locations, and some less profitable locations are indirectly priced out. To obtain insights on the global structure of an optimal solution, we derive in quasi-closed form the optimal solution for a family of models characterized by a demand shock. The optimal solution, although better balancing supply and demand around the shock, quite interestingly also ends up inducing movement away from it.

History: Accepted by David Simchi-Levi, optimization.

Supplemental Material: The online appendices are available at <https://doi.org/10.1287/mnsc.2020.3622>.

Keywords: spatial pricing • revenue management • ride hailing • strategic supply • market design

1. Introduction

Pricing and revenue management have seen significant developments over the years in both practice and the literature. At a high level, the main focus has been to investigate tactical pricing decisions given the dynamic evolution of inventories, with prototypical examples coming from the airline, hospitality, and retail industries (Talluri and Van Ryzin 2004). With the emergence and multiplication of two-sided marketplaces, a new question has emerged: how to price when capacity/supply units are strategic and can decide when and where to participate. This is particularly relevant for ride-hailing platforms such as Uber and Lyft. In these platforms, drivers are independent contractors who have the ability to relocate strategically within their cities to boost their own profits. Although this operating model leads to a more flexible supply, it also restricts the platform from reallocating supply across locations at will. Instead, a platform must ensure that incentives are in place for drivers to select to reallocate themselves. Consider the spatial pricing problem within a city faced by a platform that shares its revenues with drivers. Suppose that there are different


demand and supply conditions across the city. The platform may want to increase prices at locations with high demand and low supply. Such an increase would have two effects. The first effect is a local demand response, which pushes the riders who are not willing to pay a higher price away from the system. The second effect is global in nature because drivers throughout the city may find the locations with high prices more attractive than the ones where they are currently located and may decide to relocate. In turn, this may create a deficit of drivers at some locations. In other words, prices set in *one region* of a city impact demand and supply in this region but also potentially impact supply in *other regions*. This brings to the foreground the question of how to price in space when supply units are strategic.

The central focus of this paper is to understand the interplay between spatial pricing and supply response. In particular, we aim to understand how to optimally set prices across locations in a city and what the impact of those prices is on the strategic repositioning of drivers. To that end, we consider a short-term model over a given timeframe where overall

supply is constant. That is, drivers respond to pricing and congestion by moving to other locations but not by entering or exiting the system. In our short-term framework, the platform's only tool for increasing the supply of drivers at a given location is to encourage drivers to relocate from other places. In turn, this time scale permits us to isolate the spatial implications on the different agents' strategic behavior. In this sense, our model can be thought of as a building block to better understand richer temporal-dynamic environments.

In more detail, we consider a revenue-maximizing platform that sets prices to match price-sensitive riders (demand) to strategic drivers (supply) who receive a fixed commission. In making their decisions, drivers take into account prices, supply levels across the city, and transportation costs. More formally, we consider a measure-theoretical Stackelberg game with three groups of players: a platform, drivers, and potential customers. Supply and demand are composed of nonatomic agents, who are initially arbitrarily positioned throughout the city. We use non-negative measures to model how these agents are distributed in the city. The players interact with each other in a two-dimensional (2D) city. Every location can admit different levels of supply and demand. For example, Figure 1 showcases an instance of a rectangular city in which the measure or distribution of supply has two peaks, whereas the distribution of demand has one peak in the center. The platform moves first, selecting prices for the different locations around the city. Once prices are set, the mass of customers willing to pay such prices is determined. Then drivers move in equilibrium in a simultaneous-move game, choosing where to reposition based on prices, supply levels, and driving costs. In fact, besides prices and transportation costs, supply levels across the city are a key consideration for drivers when optimizing their repositioning. If too many other drivers are at a given location, a driver relocating there will be less likely to be matched to a rider, negatively affecting that driver's utility. The platform's optimization problem consists of finding prices for all locations given that drivers move in equilibrium.

Figure 1. (Color online) Two-Dimensional Setting



Note. Distribution of initial supply and potential demand across locations.

1.1. Main Contributions

Our first set of contributions is at the modeling level. We propose a general measure-theoretical framework that encompasses a wide range of environments. Our setup can be used to study spatial interactions in both discrete- and continuous-location settings.

The platform's problem is an infinite-dimensional optimization problem with equilibrium constraints. This is a notably "hard" class of problems. In our second set of contributions, we develop a methodology to study the platform's problem. Our main result provides a structural characterization of the optimal prices and resulting drivers' equilibrium in regions of the city where drivers relocate. Our approach relies on a series of transformations, localization, and relaxations. In particular, we first establish that the platform's objective can be reformulated as a function of only the equilibrium utilities of drivers and their equilibrium postrelocation distribution. In turn, we establish structural properties of these two objects. We first characterize properties of the drivers' equilibrium utilities and prove that the city admits a form of spatial decomposition into regions where movement may emerge in equilibrium, "attraction regions," and the rest of the city. Furthermore, we establish that the equilibrium utility of drivers and the local equilibrium postrelocation supply are linked through a congestion bound. The former admits a fundamental upper bound parametrized by the latter. Based on these properties, we derive a relaxation of the platform's problem that takes the form of *coupled continuous bounded knapsack* problems. Notably, we establish that this relaxation is tight and, leveraging the knapsack structure, use it to obtain a crisp *local* structural characterization of an optimal pricing solution and its supply response.

Although this framework provides local structural properties of optimal pricing policies in arbitrary 2D regions of movement, in our third set of contributions, we shed light on the scope of prices as an incentive mechanism and provide insights into the *global* structure of an optimal policy. To that end, we focus on a family of one-dimensional (1D) instances (which could be interpreted as a cut of a symmetric 2D city) that are rich enough to capture the core interactions among supply agents but also confined enough to derive *quasi-closed-form solutions* that allow us to crisply identify some key features of an optimal solution. In particular, we study a family of cases in which a central location in the city, the origin, experiences a shock of demand.

Leveraging our earlier methodological results in conjunction with the derivation of new results, we characterize in quasi-closed form the optimal pricing policy and its corresponding supply response. Strikingly, the optimal pricing policy induces movement

toward the demand shock but potentially also *away* from the demand shock. The platform may create *damaged regions* through both prices and congestion to steer the flow of drivers toward more profitable regions. Compared with a heuristic that would only adjust the price at the shock location, the optimal solution incentivizes more drivers to travel toward the demand shock.

The optimal pricing policy splits the city into multiple regions around the origin (Figure 2). The mass of customers needing rides at the location of the shock is serviced by three subregions around it: the origin, the inner center, and the outer center. The origin is the most profitable location, so the platform surges its price, encouraging the movement of a mass of drivers to meet its high levels of demand. These drivers come from both the inner and outer center. In the former, locations are positively affected by the shock, and some drivers choose to stay in them, whereas others travel toward the origin. In the latter, drivers are too far from the demand shock, so the platform has to *deliberately damage* this region through prices to create incentives for drivers to relocate toward the origin. However, drivers in this region have an option: Instead of driving toward the demand shock at the origin, they could drive away from it. This gives rise to the next region, the inner periphery. Consider the marginal driver, that is, the furthest driver willing to travel to the origin. To incentivize the marginal driver to move to the origin, the platform is obligated to also damage conditions in the inner periphery. The optimal solution creates two subregions within the inner periphery. In the first, conditions are degraded through prices that make it unattractive for drivers. Drivers in this region leave toward the second region. That is, they drive in the direction opposite to the demand shock. The action of the platform in the second region is more subtle. Here the platform does not need to play with prices. The mere fact that drivers from the first region run away to this area creates congestion, and this is sufficient degradation to make the region unattractive for the marginal driver. The final region is the outer periphery, which is too far from the origin to be affected by its demand shock.


We complement our analysis with a set of numerics that highlights that the optimal policy can generate significantly more revenues than a heuristic that would simply respond locally to a shock in demand. In other words, anticipating the global supply response and taking advantage of the full flexibility of spatial pricing play a key role in revenue optimization.

2. Related Literature

Several recent papers examine the operations of ride-hailing platforms from diverse perspectives. We first review works that do not take spatial considerations into account. There is a recent but significant body of work on the impact of incentive schemes on agents' participation decisions. Gurvich et al. (2016) study the cost of self-scheduling capacity in a newsvendor-like model in which the firm chooses the number of agents it recruits and, in each period, selects a compensation level as well as a cap on the number of available workers. Cachon et al. (2017) analyze various compensation schemes in a setting in which the platform takes into account drivers' long- and short-term incentives. They establish that in high-demand periods, all stakeholders can benefit from dynamic pricing and fixed-commission contracts can be nearly optimal. The performance of such contracts in two-sided markets is analyzed by Hu and Zhou (2017), who derive performance guarantees. Taylor (2018) considers how uncertainty affects the price and wage decisions of on-demand platforms when facing delay-sensitive customers and autonomous capacity. Nikzad (2018) focuses on the effect of market thickness and competition on wages, prices, and welfare and shows that, in some circumstances, more supply could lead to higher wages and competition across platforms could lead to high prices and low consumer welfare.

In the context of matching in ride hailing without pricing, Feng et al. (2017) compare the waiting-time performance in a circular city of on-demand matching versus traditional street-hailing matching. Hu and Zhou (2016) analyze a dynamic matching problem as well as the structure of optimal policies. Relatedly, Ozkan and Ward (2020) develop a heuristic based on a

Figure 2. (Color online) The Optimal Solution Splits the City into Multiple Regions



continuous linear program to maximize the number of matches in a network. Afèche et al. (2017) study demand admission controls and drivers' repositioning in a two-location network, without pricing, and show that the value of the controls is large when both capacity is moderate and demand is imbalanced.

Most closely related to our work are papers that study pricing with spatial considerations. Castillo et al. (2017) take space into account but only in reduced form through the shape of the supply curve. This paper points out that surge pricing can help to avoid an inefficient situation termed the wild goose chase in which drivers' earnings are low because of long pickup times. Banerjee et al. (2015) consider a queuing network where drivers do not make decisions in the short term (no repositioning decisions), but they do care about their long-term earnings. They prove that a localized static policy is optimal as long as the system parameters are constant but that a dynamic pricing policy is more robust to changes in these parameters. Banerjee et al. (2016) find approximation methods to find source-destination prices in a network to maximize various long-run average metrics. In Banerjee et al. (2016), customers have a destination and react to prices, but supply units do not behave strategically. An important contribution to the field is that of Bimpikis et al. (2019). They study pricing under steady-state conditions in a network in which drivers behave in equilibrium and decide whether and when to provide service as well as where to reposition. They are able to isolate an interesting "balance" property of the network and establish its implications for prices, profits, and consumer surplus. Buchholz (2017) structurally estimates a spatial model to understand the welfare costs of taxi fare regulations. These papers investigate long-term implications of spatial pricing. In contrast, our work examines how the platform should respond to short-term supply-demand imbalances given that the supply units are strategic. Relatedly, Guda and Subramanian (2019) study a two-location setting and show the benefits of using strategic pricing in a short-term time scale. From an empirical perspective, short-term strategic repositioning decisions of drivers to changes in heat maps of prices have been demonstrated using Uber data in Lu et al. (2018). Our work complements such empirical and theoretical studies by providing a general multilocation framework for pricing while accounting for short-term strategic relocation decisions.

From a methodological point of view, our work borrows tools from the literature on nonatomic congestion games. Our equilibrium concept is similar to the one used by Roughgarden and Tardos (2002) and Cole et al. (2003) to analyze selfish routing under congestion in discrete settings: In equilibrium, drivers only depart for locations that yield the largest earnings.

We consider a more general measure-theoretical environment that can be traced back to Schmeidler (1973), Mas-Colell (1984), and, more recently, Blanchet and Carlier (2015) and Blanchet et al. (2018). The latter two papers build on the work of Mas-Colell by introducing an equilibrium notion that accounts for local congestion effects and relates it to the theory of optimal transport (see, e.g., Villani 2008). Moreover, Blanchet et al. (2018) exploit a regularization technique to compute equilibria in discrete settings. Our equilibrium concept relates to the one introduced by these authors in that it can be applied to general measure-theoretical frameworks and it captures local congestion effects. Once the platform sets prices, drivers must decide where to relocate. This creates a "flow" or a "transport plan" in the city from initial supply (initial measure) to postrelocation supply (final measure). A fundamental difference is that in our case the final measure is being optimized, and most of our work is focused on determining this measure with the goal of maximizing revenue. In particular, in our setting, there is a first-stage optimization of spatial prices, which is not present in the aforementioned paper.

Finally, some of our insights relate back to the damaged-goods literature. Deneckere and McAfee (1996) explain that a firm can strategically degrade a good in order to price discriminate. In our setting, the platform can damage some regions of the city through prices and congestion to steer drivers toward more profitable locations and thus increase revenues.

3. Problem Formulation

We will use measure-theoretic objects to represent supply, demand, and related concepts. This level of generality will enable us to capture the rich interactions that arise in the system through a spatial model that subsumes continuous and discrete settings. Our general framework allows us to focus on the central "physical" quantities that are not tailored to the nature of the model but also allows for quasi-closed-form solutions in special cases of interest. We now introduce some basic preliminaries to make the exposition rigorous.

3.1. Preliminaries

For an arbitrary metric set $\mathcal{X}$ equipped with a norm $\|\cdot\|$ and the Borel σ -algebra, we let $\mathcal{M}(\mathcal{X})$ denote the set of nonnegative finite measures on $\mathcal{X}$. The notation $\mathcal{T} \ll \mathcal{T}'$ represents measure $\mathcal{T}$ being absolutely continuous with respect to measure $\mathcal{T}'$. The notation $\text{ess sup}_{\mathcal{B}}$ corresponds to the essential supremum, which is the measure-theoretical version of a supremum that does not take into account sets of measure zero. The notation $\mathcal{T} - a.e.$ represents "almost everywhere with respect to measure $\mathcal{T}$." For any measure $\mathcal{T}$ in a product space $\mathcal{B} \times \mathcal{B}$, $\mathcal{T}_1$ and $\mathcal{T}_2$ denote, respectively,

the first and second marginals of $\mathcal{T}$. We use $\mathbf{1}_{\{\cdot\}}$ to denote the indicator function and S^c to represent the complement of a set S . We denote the closed and open line segments between any two points by $[x, y]$ and (x, y) , respectively.

3.2. Model Elements

Our model contains four fundamental elements: a city, a platform, drivers, and potential customers. For consistency, we use masculine pronouns to refer to drivers and feminine ones to refer to customers. We represent the city by a convex, compact subset $\mathcal{C}$ of $\mathbb{R}^2$ and a measure Γ in $\mathcal{M}(\mathcal{C})$. We refer to this measure as the “city measure,” and it characterizes the “size” of every location of the city. For example, if Γ has a point mass at some location, then that location is large enough to admit a point mass of supply and demand. In turn, Γ enables us to capture settings with a continuum of locations or discrete locations.

Demand (potential customers) and supply (drivers) are assumed to be infinitesimal—a single agent does not impact the outcome of the game—and initially distributed on $\mathcal{C}$. We denote the initial demand measure by $\Lambda(\cdot)$ and the supply measure by $\Theta(\cdot)$, with both measures belonging to $\mathcal{M}(\mathcal{C})$. For example, if Θ is the Lebesgue measure on $\mathcal{C}$, then drivers are uniformly distributed over the city. Both the demand and supply measures are assumed to be absolutely continuous with respect to the city measure, that is, $\Lambda, \Theta \ll \Gamma$. The proportion of customers at location $y \in \mathcal{C}$ with willingness to pay below q is given by $F_y(q)$.¹ We let $\bar{F}_y(q) = 1 - F_y(q)$. For all $y \in \mathcal{C}$, we assume that the revenue function $q \mapsto q \cdot \bar{F}_y(q)$ is continuous and unimodal and that F_y is strictly increasing over its support $[0, \bar{V}]$ for some finite positive $\bar{V}$.

We model the interactions between platform, customers, and supply as a game. The first player to act in this game is the platform. The platform selects fares across locations and facilitates the matching of drivers and customers. Specifically, the platform selects a measurable price mapping $p : \mathcal{C} \rightarrow [0, \bar{V}]$ so as to maximize its citywide revenues.

After prices are chosen, drivers select *whether* to relocate and *where* to do so. The relocation of drivers generates a flow/transport of mass from the initial measure of drivers Θ to some final endogenous measure of drivers. This final measure corresponds to the supply of drivers in the city after they have traveled to their chosen destinations. The movement of drivers across the city is modeled as a measure on $\mathcal{C} \times \mathcal{C}$, which we denote by $\mathcal{T}$. Any feasible flow has to preserve the initial mass of drivers in $\mathcal{C}$. That is, the first marginal of $\mathcal{T}$ should equal Θ . Moreover, $\mathcal{T}$ generates a new (after relocation) distribution of drivers in the city, which corresponds to the second

marginal of $\mathcal{T}$, $\mathcal{T}_2$. Formally, the set of feasible flows is defined as follows:

$$\mathcal{F}(\Theta) = \{\mathcal{T} \in \mathcal{M}(\mathcal{C} \times \mathcal{C}) : \mathcal{T}_1 = \Theta, \quad \mathcal{T}_2 \ll \Gamma\}.$$

The first condition ensures consistency with the initial positioning of drivers; the second condition ensures that there is no mass of relocated supply at locations where the city itself has measure zero and thus do not form part of the game. In particular, given the latter, the Radon–Nikodym derivatives of $\mathcal{T}_2$ and Λ with respect to Γ , $d\mathcal{T}_2(y)/d\Gamma$, and $d\Lambda(y)/d\Gamma$ are well defined; for ease of notation, we let, for any y in $\mathcal{C}$,

$$s^{\mathcal{T}}(y) \triangleq \frac{d\mathcal{T}_2}{d\Gamma}(y), \quad \text{and} \quad \lambda(y) \triangleq \frac{d\Lambda}{d\Gamma}(y).$$

Physically, $s^{\mathcal{T}}(y)$ represents the *postrelocation supply* at location y normalized by the size of location y , and $\lambda(y)$ corresponds to the potential demand at location y also normalized by the same size of such location.² Here and in what follows, we will refer to $s^{\mathcal{T}}(y)$ and $\lambda(y)$ as the *postrelocation supply* and potential demand at y , respectively. In order to lighten the exposition, and without loss of generality, we assume that $\lambda(y) > 0$ Γ -almost everywhere in $\mathcal{C}$.

Given the prices in place, the effective demand at a location y is given by $\lambda(y) \cdot \bar{F}_y(p(y))$, because at location y , only the fraction $\bar{F}_y(p(y))$ is willing to purchase at price $p(y)$. At the same time, the supply at y is given by $s^{\mathcal{T}}(y)$. Therefore, the ratio of effective (as opposed to potential) demand to supply at y is given by

$$\frac{\lambda(y) \cdot \bar{F}_y(p(y))}{s^{\mathcal{T}}(y)},$$

assuming $s^{\mathcal{T}}(y) > 0$. Because a driver can pick up at most one customer within the time frame of our game, a driver relocating to y will face a utilization rate of $\min\{1, \lambda(y) \cdot \bar{F}_y(p(y))/s^{\mathcal{T}}(y)\}$, again assuming $s^{\mathcal{T}}(y) > 0$. The effective utilization can be interpreted as the probability that a driver who relocated to y will be matched to a customer within the time frame of our game. In particular, if $s^{\mathcal{T}}(y) > \lambda(y) \cdot \bar{F}_y(p(y))$, there is driver congestion at location y , and not all drivers will be matched to a customer. If $s^{\mathcal{T}}(y) = 0$ at location y , we say that the utilization rate is one if the effective demand at y is positive and zero if the effective demand is zero. Formally, the utilization rate at location y is given by

$$R(y, p(y), s^{\mathcal{T}}(y)) \triangleq \begin{cases} \min\left\{1, \frac{\lambda(y) \cdot \bar{F}_y(p(y))}{s^{\mathcal{T}}(y)}\right\}, & \text{if } s^{\mathcal{T}}(y) > 0, \\ 1, & \text{if } s^{\mathcal{T}}(y) = 0, \lambda(y) \cdot \bar{F}_y(p(y)) > 0, \\ 0, & \text{if } \lambda(y) \cdot \bar{F}_y(p(y)) = 0. \end{cases}$$

When deciding whether to relocate, drivers take three effects into account: prices, travel distance, and congestion. The driver congestion effect (or utilization rate) is the one described in the preceding paragraph. We assume that the platform uses a commission model and transfers a fraction α in $(0, 1)$ of the fare to the driver. As a result, a driver who starts in location y and chooses to remain there earns utility equal to

$$U(y, p(y), s^{\mathcal{T}}(y)) \triangleq \alpha \cdot p(y) \cdot R(y, p(y), s^{\mathcal{T}}(y)). \quad (1)$$

That is, the utility is given by the compensation per ride times the probability of a match. We model the cost for drivers of repositioning from location x to location y through the distance between the locations $\|y - x\|$. Therefore, a driver originating in x who repositions to y earns utility

$$\Pi(x, y, p(y), s^{\mathcal{T}}(y)) \triangleq U(y, p(y), s^{\mathcal{T}}(y)) - \|y - x\|. \quad (2)$$

When clear from context, and with some abuse of notation, we omit the dependence on price and the supply-demand ratio, writing $U(y)$ and $\Pi(x, y)$. Thus far, we have introduced the objects to study a classic Cournot–Nash equilibrium (see, e.g., Blanchet and Carlier 2015). The space of players' types is $\mathcal{C}$ endowed with the type distribution $\Theta \in \mathcal{M}(\mathcal{C})$. The action space is the set of possible destinations $\mathcal{C}$, and there is an action distribution $\mathcal{V} \in \mathcal{M}(\mathcal{C})$ with density v . Given an action distribution $\mathcal{V}$, an agent with type x choosing action y receives a total utility given by $\Pi(x, y, p(y), v(y))$. We are now ready to define the notion of a supply equilibrium. An equilibrium specifying type–action pairs is defined through a flow $\mathcal{T} \in \mathcal{M}(\mathcal{C} \times \mathcal{C})$. This distribution should be consistent with the type and action distributions $\mathcal{T}_1 = \Theta$ and $\mathcal{T}_2 = \mathcal{V}$, and the actions of agents should be optimal given all other agents' actions. We next define precisely a supply equilibrium.

Definition 1 (Supply Equilibrium). A flow $\mathcal{T} \in \mathcal{F}(\Theta)$ is an equilibrium if it satisfies

$$\mathcal{T} \left(\left\{ (x, y) \in \mathcal{C} \times \mathcal{C} : \Pi(x, y, p(y), s^{\mathcal{T}}(y)) \right. \right. \\ \left. \left. = \operatorname{ess\,sup}_{\mathcal{C}} \Pi(x, \cdot, p(\cdot), s^{\mathcal{T}}(\cdot)) \right\} \right) = \Theta(\mathcal{C}),$$

where the essential supremum is taken with respect to the city measure Γ .

That is, an equilibrium flow of supply is a feasible flow such that essentially no driver wishes to unilaterally change his destination. As a result, the mass of drivers selecting the best location for themselves has to equal the original mass of drivers in the system.

The platform's objective is to maximize the revenues it garners across all locations in $\mathcal{C}$. With the

assumed commission structure in place, from a given location y , it earns $(1 - \alpha) \cdot p(y) \cdot \min\{s^{\mathcal{T}}(y), \lambda(y) \cdot \bar{F}_y(p(y))\}$. The term $(1 - \alpha) \cdot p(y)$ corresponds to the platform's share of each fare at location y , and the term $\min\{s^{\mathcal{T}}(y), \lambda(y) \cdot \bar{F}_y(p(y))\}$ denotes the quantity of matches of potential customers to drivers at location y . The platform's price-optimization problem can, in turn, be written as

$$\sup_{p(\cdot), \mathcal{T} \in \mathcal{F}(\Theta)} (1 - \alpha) \int_{\mathcal{C}} p(y) \cdot \min\{s^{\mathcal{T}}(y), \lambda(y) \cdot \bar{F}_y(p(y))\} d\Gamma(y)_1 \\ \text{(s.t.) } \mathcal{T} \text{ is a supply equilibrium,} \\ s^{\mathcal{T}} = \frac{d\mathcal{T}_2}{d\Gamma}. \quad (\mathcal{P}_1)$$

Remark 1. Our model may be interpreted as a basic model to understand the short-term operations of a ride-hailing company that needs to design a city-wide pricing policy to properly incentivize drivers. In particular, each driver completes at most one customer pickup within the time frame of our game, and there is not enough time for the entry of new drivers into the system. In the present model, we do not account explicitly for the destinations of the rides. We do so in order to isolate the interplay of supply incentives and pricing. In that regard, one could view our model as capturing origin-based pricing, a common practice in the ride-hailing industry. Note that our framework could be modified to partially account for origin–destination pricing by incorporating the fact that different locations lead to different average lengths of rides (such a factor would multiply the price on the right-hand side of Equation (1)). With such a modification, drivers, before repositioning to a location, would consider the corresponding average price for the total trip.

3.3. Solution Approach

Problem $(\mathcal{P}_1)$ can be classified as an infinite-dimensional mathematical program with equilibrium constraints. This is a notably difficult class of problems to solve in general, even numerically. For the particular problem we study, we will show that significant structure exists and that one can obtain a crisp local characterization of an optimal solution. We next lay out our approach.

In Section 4, we reformulate the platform's objective. The new objective is “well behaved” and showcases the fundamental structure of the problem. In particular, if we relaxed all equilibrium constraints, the problem would be a continuous knapsack problem in which the limited initial budget of drivers $\Theta(\mathcal{C})$ must be allocated across locations. However, the solution to the knapsack relaxation may not satisfy the equilibrium constraints. To address this, in

Sections 5 and 6, we derive equilibrium properties that we later add as constraints to the aforementioned relaxation. More precisely, in Section 5, we prove a fundamental upper bound on the amount of post-relocation supply that there could be at any location in equilibrium. In Section 6, we identify and localize the analysis to the regions in the city where drivers move. We study their properties and show that these regions, subject to appropriate constraints, can be optimized in isolation. In Section 7, we add these developments as constraints to problem $(\mathcal{P}_1)$ and relax the equilibrium constraints. The resulting problem is a relaxation of $(\mathcal{P}_1)$ in a region of potential movement that has the structure of a continuous bounded knapsack problem. We then show that this local relaxation is tight, thereby providing a characterization of the optimal solution within regions of potential movement. In Section 8, we combine these results, together with new results and a flow-mimicking technique (Online Appendix D.2), to solve in quasi-closed form for the optimal solution in a family of instances. Figure 3 provides a summary of our solution approach.

4. Structural Properties and Reformulation

A key challenge in solving the optimization problem presented in $(\mathcal{P}_1)$ is that the decision variables, the flow $\mathcal{T}$ and the price function $p(\cdot)$, are complicated objects. The flow $\mathcal{T}$, being a measure over movements on a 2D space, is obviously a complex object to manipulate. The price function will turn out to be a difficult object to manipulate as well. In order to analyze our

problem, we will need to introduce a better-behaved object. This object, which will be central to our analysis, is the (after movement) driver equilibrium utility.

4.1. Drivers' Utilities

For a given price function p and flow $\mathcal{T}$, we denote by $V_{\mathcal{B}}(x|p, \mathcal{T})$ the essential maximum utility that a driver departing from location x can garner by going anywhere within a measurable region $\mathcal{B} \subseteq \mathcal{C}$. In particular, the mapping $V_{\mathcal{B}}(\cdot|p, \mathcal{T}) : \mathcal{C} \rightarrow \mathbb{R}$ is defined as


$$V_{\mathcal{B}}(x|p, \mathcal{T}) \triangleq \operatorname{ess\,sup}_{\mathcal{B}} \Pi(x, \cdot, p(\cdot), s^{\mathcal{T}}(\cdot)). \quad (3)$$

When $\mathcal{B} = \mathcal{C}$, we use V instead of $V_{\mathcal{C}}$. Note that in the essential supremum, $\mathcal{B}$ enters Π when we evaluate the second argument of Π , $p(\cdot)$, and $s^{\mathcal{T}}(\cdot)$. By the definition of a supply equilibrium, essentially all drivers departing from location x earn $V(x|p, \mathcal{T})$ utility in equilibrium.

We now show that the equilibrium utility $V_{\mathcal{B}}(\cdot|p, \mathcal{T})$ must be 1-Lipschitz continuous. Intuitively, drivers from two different locations x and y who consider relocating see exactly the same potential destinations. So the largest utility drivers departing from x can garner must be greater or equal to that of the drivers departing from y minus the disutility stemming from relocating from x to y ; that is, $V_{\mathcal{B}}(x) \geq V_{\mathcal{B}}(y) - \|x - y\|$. Because this argument is symmetric, we deduce the 1-Lipschitz property.

Lemma 1 (Lipschitz). *Consider a measurable set $\mathcal{B} \subseteq \mathcal{C}$ such that $\Gamma(\mathcal{B}) > 0$. Let p be a measurable mapping $p : \mathcal{B} \rightarrow \mathbb{R}_+$, and let $\mathcal{T} \in \mathcal{F}(\Theta)$. Then the function $V_{\mathcal{B}}(\cdot|p, \mathcal{T})$ is 1-Lipschitz continuous.*

Figure 3. Solution Approach



Notes. (a) General solution approach to characterize the optimal solution within regions of movement. (b) Application of general results to obtain optimal solution in a family of instances with a shock of demand in the center of the city (Section 8).

4.2. Reformulating the Platform's Problem

We now introduce a reformulation of $(\mathcal{P}_1)$ that focuses on the equilibrium utility V and the postrelocation supply $s^\mathcal{T}$ as the central elements. The next result plays a key role in our solution approach and is what later motivates the development of the upper bound on $s^\mathcal{T}$ and the localization of the analysis to regions of movement (see Figure 3). In what follows, we define $\gamma \triangleq (1 - \alpha)/\alpha$.

Proposition 1 (Problem Reformulation). *The problem*

$$\begin{aligned} & \sup_{p(\cdot), \mathcal{T} \in \mathcal{T}(\Theta)} \gamma \cdot \int_{\mathcal{C}} V(x|p, \mathcal{T}) \cdot s^\mathcal{T}(x) d\Gamma(x) \\ & \text{s.t.} \quad \mathcal{T} \text{ is an equilibrium flow,} \\ & \quad V(x|p, \mathcal{T}) = \text{esssup}_{\mathcal{C}} \Pi(x, \cdot, p(\cdot), s^\mathcal{T}(\cdot)), \\ & \quad s^\mathcal{T} = \frac{d\mathcal{T}_2}{d\Gamma}, \end{aligned} \quad (\mathcal{P}_2)$$

admits the same value as the platform's optimization problem $(\mathcal{P}_1)$, and a pair $(p, \mathcal{T})$ that solves $(\mathcal{P}_2)$ also solves $(\mathcal{P}_1)$.

The first step in the proof of Proposition 1 is to rewrite the platform's objective in terms of the postrelocation supply $s^\mathcal{T}(x)$ and the premovement utility function $U(x, p(x), s^\mathcal{T}(x))$ (see Equation (1)). We note that the former step is driven by the commission structure of the contract offered by the platform, which aligns in some way the platform's objective with the drivers' utility. This transformation is not particularly useful per se because the function $U(x, p(x), s^\mathcal{T}(x))$ is not necessarily well behaved. The next step consists of establishing that $U(x, p(x), s^\mathcal{T}(x))$ coincides with $V(x|p, \mathcal{T})$ whenever a location has positive postmovement equilibrium supply (see Lemma A-2 in Online Appendix A). Indeed, whenever the equilibrium outcome is such that a location has positive supply, the utility generated by staying at that location has to be equal to the best utility one could obtain by traveling to any other location. In turn, one can effectively replace $U(x, p(x), s^\mathcal{T}(x))$ with $V(x|p, \mathcal{T})$ in the objective, which yields the alternative formulation.

The new formulation $(\mathcal{P}_2)$ offers a new perspective on the platform's problem. The key objects that drive the platform's revenue are $s^\mathcal{T}(\cdot)$ and $V(\cdot|p, \mathcal{T})$. Intuitively, for each unit of supply allocated to some x , the platform obtains $V(x|p, \mathcal{T})$. Observe that there is also a "budget" constraint that limits how much $s^\mathcal{T}(\cdot)$ can be allocated across locations because the total supply in the system is $\Theta(\mathcal{C})$. In turn, if we relaxed the equilibrium constraints in $(\mathcal{P}_2)$, the problem would correspond to a *knapsack* problem. In this case, we would set $s^\mathcal{T}(\cdot)$ as high as possible in the location where $V(\cdot|p, \mathcal{T})$ is the highest. However, doing so would violate the equilibrium constraints because that location would

experience high levels of congestion, which, as a result, would deter drivers from relocating to it in equilibrium. In order to deal with this, in the next section, we develop a congestion bound that controls for the effect on drivers' utilities introduced by congestion.

4.3. Connection to Optimal Transport

Our equilibrium concept is closely related to the notion of transport plan in the theory of optimal transport. For example, it is possible to establish that in any equilibrium $\mathcal{T}$, the total mass of drivers repositions in the most efficient way as to minimize the total transportation cost (see, e.g., Blanchet and Carlier 2015 for a related result). In contrast with optimal transport, in our case, the destination measure $\mathcal{T}_2$ is an endogenous object, and one of the central tasks is to find it via optimization by solving $(\mathcal{P}_2)$.

5. Congestion Bound

We first introduce some quantities that emerge from a classical capacitated monopoly pricing problem. Let us consider any location $x \in \mathcal{C}$ and ignore all other locations in the city. The problem that a monopolist faces when supply at x is s and demand is $\lambda(x)$ can be cast as

$$R_x^{\text{loc}}(s) \triangleq \max_{q \in [0, \bar{V}]} q \cdot \min\{s, \lambda(x) \cdot \bar{F}_x(q)\}, \quad (4)$$

with the price $\rho_x^{\text{loc}}(s)$ being defined as the argument that maximizes Equation (4). Because $q \cdot \bar{F}_x(q)$ is assumed to be unimodal in q , the optimal price $\rho_x^{\text{loc}}(s)$ is uniquely determined and is characterized as follows:

$$\begin{aligned} \rho_x^{\text{loc}}(s) &= \max\{\rho_x^{\text{bal}}(s), \rho_x^u\}, \text{ where } s = \lambda(x) \cdot \bar{F}_x(\rho_x^{\text{bal}}(s)), \\ \rho_x^u &\in \arg \max_{\rho \in [0, \bar{V}]} \{\rho \cdot \bar{F}_x(\rho)\}. \end{aligned} \quad (5)$$

That is, the optimal local price either balances supply and demand or maximizes the unconstrained local revenue. For a given local supply s , the maximum revenue that can be generated at location x is $R_x^{\text{loc}}(s)$, with a fraction α of that revenue being paid to the drivers. Therefore, $\alpha \cdot R_x^{\text{loc}}(s)/s$ is the maximum revenue a driver staying at this location can earn. To capture this notion, we introduce for every location x the supply congestion function $\psi_x : \mathbb{R}_+ \rightarrow [0, \alpha \cdot \bar{V}]$, which is defined as follows:

$$\psi_x(s) \triangleq \begin{cases} \alpha \cdot R_x^{\text{loc}}(s)/s & \text{if } s > 0, \\ \alpha \cdot \bar{V} & \text{if } s = 0. \end{cases}$$

In line with intuition, more drivers (in a single-location problem) imply lower revenues per driver; it is possible to show that the congestion function ψ_x is

decreasing (see Lemma A-3 in Online Appendix A). Crucially, the congestion function ψ_x yields an upper bound for the utility of drivers.

Proposition 2 (Congestion Bound). *Let $(p, \mathcal{T})$ be a feasible solution of $(\mathcal{P}_2)$. Then the equilibrium driver utility function is bounded as follows:*

$$V(x|p, \mathcal{T}) \leq \psi_x(s^{\mathcal{T}}(x)), \quad \Gamma - \text{a.e. } x \text{ in } \mathcal{C}.$$

When there is a single location, this inequality is an equality by the definition of ψ_x . For multiple locations, drivers may travel to any location, and there is no a priori connection between the utility that drivers originating from x can garner, $V(x|p, \mathcal{T})$ and $\psi_x(s^{\mathcal{T}}(x))$. The preceding result establishes that the latter upper bounds the former. The bound captures the structural property that as equilibrium supply increases at a location, and hence driver congestion increases, the drivers originating from that location will earn less utility.

As discussed in the preceding section, $(\mathcal{P}_2)$ admits a relaxation (not necessarily tight) that can be mapped to a knapsack problem. Proposition 2 provides a capacity constraint that can be added to such a relaxation. More precisely, because $\psi_x(\cdot)$ is strictly decreasing the congestion-bound delivers,

$$s^{\mathcal{T}}(x) \leq \psi_x^{-1}(V(x|p, \mathcal{T})), \quad \Gamma - \text{a.e. } x \text{ in } \mathcal{C}. \quad (6)$$

That is, given $V(\cdot|p, \mathcal{T})$, for almost every location in the city, the amount of postrelocation supply has a structural upper bound. Equation (6) can be plugged into the aforementioned relaxation to make it a bounded knapsack problem. Nevertheless, a solution to this new problem might still violate the equilibrium constraints and not be tight. The reason is that the resulting allocation may prescribe movement of drivers from faraway locations. To address this, in the next section, we localize the analysis to regions to which drivers are willing to travel and study their properties.

6. Localization to Regions of Potential Movement

A key feature of the problem at hand is that, in equilibrium, conditions at different locations are inherently linked because drivers select their destination among all locations. An important object that will help capture the link across various locations is the *indifference region* of a driver departing location x . The indifference region of x represents all the destinations to which drivers from x are willing to travel. Formally, the indifference region for a driver departing from $x \in \mathcal{C}$ under prices p and flow $\mathcal{T}$ is given by

$$\mathcal{IR}(x|p, \mathcal{T}) \triangleq \{y \in \mathcal{C} : V(y|p, \mathcal{T}) - \|y - x\| = V(x|p, \mathcal{T})\}.$$

Intuitively, the preceding definition says that if $y \in \mathcal{IR}(x|p, \mathcal{T})$, then drivers departing from x maximize their utility by relocating to y . The converse concept, which will turn out to be fundamental in our analysis, is the *attraction region* of a location z . The attraction region of z represents the set of all possible sources for which location z is their best option, as we formally define next.

Definition 2 (Attraction Region). *Let $(p, \mathcal{T})$ be a feasible solution of $(\mathcal{P}_2)$. For any $z \in \mathcal{C}$, its attraction region $A(z|p, \mathcal{T})$ is the set of locations from which drivers are willing to relocate to z :*

$$A(z|p, \mathcal{T}) \triangleq \{x \in \mathcal{C} : z \in \mathcal{IR}(x|p, \mathcal{T})\}.$$

We call a location $z \in \mathcal{C}$ a *sink* if its attraction region $A(z|p, \mathcal{T})$ is nonempty and $z \notin A(z'|p, \mathcal{T})$ for all $z' \neq z$.

Note that within an attraction region, the equilibrium utility of drivers $V(\cdot|p, \mathcal{T})$ is fully determined up to its value at the sink $V(z|p, \mathcal{T})$. Importantly, sinks and corresponding attraction regions emerge as soon as drivers move in the city, as formalized in the next proposition.

Proposition 3 (Existence of Attraction Regions). *Let $(p, \mathcal{T})$ be a feasible solution of $(\mathcal{P}_2)$.*


- Any flow in the city gives rise to an associated attraction region, that is, for any $\mathcal{L} \subseteq \mathcal{C} \times \mathcal{C}$ such that $\mathcal{T}(\mathcal{L}) > 0$, there exists $(x, y) \in \mathcal{L}$ such that $y \in \mathcal{IR}(x|p, \mathcal{T})$.*
- Moreover, if $y \in \mathcal{IR}(x|p, \mathcal{T})$ for some $x \neq y$, then there exists a sink $z \in \mathcal{C}$ such that $x, y \in A(z|p, \mathcal{T})$ and x, y, z are collinear points.*

6.1. Representation and Properties of an Attraction Region

We will anchor the coming discussion around Figure 4, which provides an illustration of an attraction region (shaded region) as well as its various structural properties.

In line with the literature on optimal transport (see, e.g., Ambrosio and Pratelli 2003), it will be useful in

Figure 4. Attraction Region



Notes. Illustration of structural properties of attraction regions. No flow crosses the boundaries of $A(z|p, \mathcal{T})$, and there is no flow traveling from one ray to another ray.

our analysis to study the behavior of drivers along rays around a particular location z . We use R_z to denote the set of all rays originating from z (excluding z) and index the elements of R_z by a . With such a representation, for a feasible solution $(p, \mathcal{T})$ and for any sink $z \in \mathcal{C}$, one can derive various structural properties.


Property 1 (Closed Attraction Region). *The attraction region $A(z|p, \mathcal{T})$ is closed and can be represented as a collection of segments of the form $[z, X_a(z|p, \mathcal{T})]$, where $X_a(z|p, \mathcal{T})$ is the last point in the attraction region on ray a (see Figure 4). We provide a formal statement in Lemma B-1 in Online Appendix B.*

Property 2 (Flow Separation). *An attraction region does not receive external supply, and supply units within such a region do not travel outside. Furthermore, any movement takes place along the rays originating from the sink z . We provide a formal statement in Proposition B-1 in Online Appendix B.*

Property 3 (Pasting). *Informally, this result states the following pasting property. Suppose that we start from a price-equilibrium pair $(p, \mathcal{T})$ and a sink z and its attraction region $A(z|p, \mathcal{T})$. Then we can replace the price flow within $A(z|p, \mathcal{T})$ with any other local price equilibrium, say $(\tilde{p}, \tilde{\mathcal{T}})$, within that attraction region as long as we do not change $V(x|p, \mathcal{T})$ for $x \in A(z|p, \mathcal{T})$ that have initial supply. One can obtain a new feasible solution $(\hat{p}, \hat{\mathcal{T}})$ in $\mathcal{C}$ by merging the old solution $(p, \mathcal{T})$ in the complement of $A(z|p, \mathcal{T})$ with the modified solution $(\tilde{p}, \tilde{\mathcal{T}})$ in the attraction region $A(z|p, \mathcal{T})$. We illustrate this property in Figure 5. We provide a formal statement in Proposition B-2 in Online Appendix B.*

These properties have important implications. Attraction regions lead to a natural decoupling of the platform's problem, providing a way of segmenting the city. These regions can be optimized in isolation,

Figure 5. (Color online) Illustration of the Pasting Property (Property 3)



Notes. In $A(z|p, \mathcal{T})$, we can modify flows, postrelocation supply, and prices in a feasible manner. We show the pre- and postmodification supplies along a line segment. The initial supply corresponds to $s^{\mathcal{T}}(x)$ (solid shaded). The supply after modification is $s^{\tilde{\mathcal{T}}}(x)$ inside $A(z|p, \mathcal{T})$ (striped) and coincides with $s^{\mathcal{T}}(x)$ outside $A(z|p, \mathcal{T})$.

and the new solution within these regions can then be pasted to the old solution outside these regions to obtain a new global feasible solution to the platform's problem. In the next section, we introduce a local relaxation of $(\mathcal{P}_2)$, derive its solution, and prove its tightness. We construct this relaxation by leveraging our reformulation, congestion bound, and localization results.

7. Structure of Optimal Solution Within an Attraction Region

In this section, we leverage our previous results to derive a structural characterization of the optimal prices and postrelocation supply of drivers within any attraction region.

Consider an arbitrary feasible solution $(p, \mathcal{T})$ of $(\mathcal{P}_2)$. Let $z \in \mathcal{C}$ be a sink and $A(z|p, \mathcal{T})$ its corresponding attraction region. The next result establishes that one can construct a second feasible solution of $(\mathcal{P}_2)$ with (weakly) greater revenue and, in turn, uncover the structure of prices and supply in an optimal solution.

Theorem 1 (Optimal Supply and Prices Within an Attraction Region). *Consider a feasible solution $(p, \mathcal{T})$ of $(\mathcal{P}_2)$, and let $z \in \mathcal{C}$ be a sink. Then there exists another feasible solution $(\hat{p}, \hat{\mathcal{T}})$ that weakly revenue dominates $(p, \mathcal{T})$ and is such that its supply $s^{\hat{\mathcal{T}}}$ in $A(z|p, \mathcal{T})$ is given by*

$$s^{\hat{\mathcal{T}}}(x) = \begin{cases} \psi_x^{-1} \left(V(z|p, \mathcal{T}) - \|x - z\| \right) & \text{if } x \in \bigcup_{a \in R_z} [z, r_a), \\ s_a & \text{if } x = r_a, a \in R_z, \\ 0 & \text{otherwise,} \end{cases}$$

for a set of values $\{r_a\}$ such that $r_a \in [z, X_a(z|p, \mathcal{T})]$ and $s_a \geq 0, a \in R_z$. Furthermore,


$$\hat{p}(x) = \begin{cases} \rho_x^{\text{loc}}(s^{\hat{\mathcal{T}}}(x)) & \text{if } x \in A(z|p, \mathcal{T}) \setminus \bigcup_{a \in R_z} \{r_a\}, \\ p_a & \text{if } x = r_a, a \in R_z, \end{cases}$$

where p_a is such that $U(r_a, p_a, s_a) = V(r_a|p, \mathcal{T})$ for $a \in R_z$.

Theorem 1 characterizes the structure of an optimal solution, including both prices and flows, within an attraction region.

7.1. Illustration of Theorem 1

It is useful to consider a prototypical example to illustrate the structure of an optimal solution. In Figure 6, we consider a disk-shaped attraction region with a sink at its center z . We assume that the demand density is a cone centered at z that flattens out after a certain distance from z (see Figure 6(a)). Given that the region plotted is an attraction region, the equilibrium utility must be a cone centered at z (see Figure 6(b)).

Figure 6. (Color online) Illustration of Inputs for Theorem 1

Notes. We consider a disk-shaped attraction region with sink z (light-gray region). In panel (b), we plot the potential demand for rides $\lambda(x)$. It increases toward the sink, and it is positive but constant in the region between the dashed and solid lines. In panel (a), we depict the equilibrium utility of drivers $V(x|p, \tau) = V(z|p, \tau) - \|z - x\|$. We consider $\alpha = 0.8$, $V(z|p, \tau) = 0.56$, $X_a(z|p, \tau) = 0.35$ for all $a \in \mathbb{R}_z$, and $\lambda(x) = \max\{0.9 - 4\|z - x\|, 0.3\}$. The norm is the Euclidean distance.

In Figure 7, we display the structure of an optimal solution within this attraction region. Figure 7 depicts (a) the optimal prices, (b) the effective demand, and (c) the postrelocation equilibrium supply. Three subregions emerge. In the central subregion (subregion (iii)), the optimal prices correspond to the prices that balance induced supply and demand; that is, $s^{\hat{\mathcal{T}}}(x) = \lambda(x) \cdot (1 - \hat{p}(x))$. In this subregion, there is positive supply; thus, from the equilibrium condition, we must have that $\Pi(x, x) = V(x)$. In turn, because the utilization is one, $\alpha \hat{p}(x) = V(x)$, which implies prices grow linearly toward the sink. In this subregion, drivers in locations that are further from the sink earn lower profits as prices decrease. In the middle subregion (subregion (ii)), prices all equal the unconstrained optimal price. Here there is more supply than effective demand. In this subregion, drivers' utility decreases as we move away from the sink because of congestion rather than price changes. Finally, in the outer subregion (subregion (i)), there is no postrelocation supply. In this subregion, prices are set


in such a way that drivers are better off repositioning toward the sink rather than staying in subregion (i).

Theorem 1 establishes the general structure of an optimal solution within regions of movement. It also showcases the rich behavior that emerges, in the form of a menu of subregions, within an attraction region. In some locations, supply and demand are perfectly matched, other locations are overcongested, and some less profitable locations are priced out. We will exploit Theorem 1 in Section 8 to study a prototypical family of instances, where we will characterize all attraction regions and, in turn, the optimal solution across the city in quasi-closed form. Before doing so, we discuss the main ideas that lead to the characterization of the optimal solution within attraction regions.

7.2. Key Ideas for Theorem 1

The key idea underlying the proof of the result is based on optimizing the contribution of the attraction region $A(z|p, \mathcal{T})$ to the overall objective by reallocating the supply around the sink and then showing that this reallocation of supply constitutes an equilibrium flow in the original problem. In order to optimize the supply around the sink, we consider the following optimization problem, which, as explained later, is a relaxation of $(\mathcal{P}_2)$ within $A(z|p, \mathcal{T})$:

$$\begin{aligned} \max_{\tilde{s}(\cdot) \geq 0} \quad & \int_{A(z|p, \mathcal{T})} V(x|p, \mathcal{T}) \cdot \tilde{s}(x) d\Gamma(x) \quad (\mathcal{P}_{KP}(z)) \\ \text{s.t.} \quad & \tilde{s}(x) \leq \psi_x^{-1}(V(x)), \quad \Gamma - a.e. \ x \text{ in } \mathcal{C}, \\ & \quad \quad \quad \text{(Congestion Bound)} \\ & \int_{A(z|p, \mathcal{T})} \tilde{s}(x) d\Gamma(x) = \mathcal{T}_c, \quad \text{(Flow Conservation)} \\ & \int_{(z, X_a(z|p, \mathcal{T}))} \tilde{s}(x) d\Gamma_a(x) \leq \mathcal{T}_a, \quad \Gamma - a.e. \ a \in \mathbb{R}_z, \\ & \quad \quad \quad \text{(No Flow Crossing Rays)} \end{aligned}$$

Figure 7. (Color online) Illustration of Theorem 1's Solution

Notes. Given $A(z|p, \tau)$, the figure depicts the structure of an optimal solution for the conditions associated with Figure 6. There are three subregions that emerge. (a) The optimal prices. In (iii), optimal prices perfectly match supply and demand ρ_x^{bal} . In (ii), optimal prices correspond to the unconstrained optimal price ρ_x^u . In (i), prices are such that demand is shutdown. (b) The effective demand $\lambda(x) \cdot \tilde{F}_x(p(x))$. (c) The postrelocation supply. There is no supply in (i); in (ii), there is more supply than effective demand; and in (iii), supply and demand are matched. We consider $r_a = 0.25$ for all $a \in \mathbb{R}_z$ and $F_x(q) = q$ for all x .

where $\mathcal{T}_c$ corresponds to the total flow that $\mathcal{T}$ transports from $A(z|p, \mathcal{T})$ to $A(z|p, \mathcal{T})$, and $\mathcal{T}_a$ correspond to the total flow in $A(z|p, \mathcal{T})$ that is transported along ray a , excluding z . The measures Γ_a correspond to the contribution of rays indexed by a on the total city measure, and Γ^p is a measure over rays that allows us to integrate the contribution of the measures Γ_a to Γ (for more details, see the proof of Theorem 1 in Online Appendix C). Recall that given the postrelocation supply $\tilde{s}$, the quantity $\int_{\mathcal{B}} \tilde{s}(x) d\Gamma(x)$ represents the postrelocation supply induced by $\tilde{s}$ in $\mathcal{B}$. Thus, the last two constraints in $(\mathcal{P}_{KP}(z))$ stand for consistency of the total postrelocation supply in each one of the relevant subregions of $A(z|p, \mathcal{T})$. The key is to observe that this is a relaxation of the original problem in the attraction region. In particular, the equilibrium constraint implies the conservation constraint (see Proposition B-1(i) in Online Appendix B) and the no-flow-crossing constraints (see Proposition B-1(ii) in Online Appendix B). The congestion bound is also a consequence of the equilibrium constraint (see Proposition 2). In words, in this formulation, we relax the equilibrium constraint but impose the implications of it. We constrain the amount of mass that we can allocate in each direction around z , but we fix the total amount of mass in $A(z|p, \mathcal{T})$.

In $(\mathcal{P}_{KP}(z))$, we fix the driver utilities and ask what should be the optimal allocation of drivers while satisfying flow balance in the regions $\{[z, X_a]\}_{z \in R_z}$ and imposing the congestion bound. Clearly, selecting $\tilde{s} = s^{\mathcal{T}}$ is feasible for the preceding problem, and hence the optimal value upper bounds the value generated by the initial price-equilibrium pair $(p, \mathcal{T})$ in the region $A(z|p, \mathcal{T})$. In the proof, we show that this relaxation is tight. That is, it is possible to construct prices and equilibrium flows achieving the value of Problem $(\mathcal{P}_{KP}(z))$. The proof consists of two main steps: (1) verifying the structure of the optimal $\tilde{s}$ and (2) showing that the postrelocation supply that solves the relaxation can actually be obtained from appropriate prices and flows. For step 1, the main idea relies on recognizing that Problem $(\mathcal{P}_{KP}(z))$ is a measure-theoretical instance of a coupled collection of *continuous bounded knapsack problems*. The solution to $(\mathcal{P}_{KP}(z))$ is obtained by allocating as much as possible at locations where we can make the most revenue per unit of volume; that is, we would like to make $\tilde{s}(x)$ as large as possible at locations where $V(x|p, \mathcal{T})$ is the largest. For step 2, we explicitly construct prices and flows that yield the same postrelocation supply and objective as the solution to $(\mathcal{P}_{KP}(z))$. The optimal prices correspond to the optimal local prices $\rho_x^{loc}(s^{\hat{\mathcal{T}}}(x))$. We obtain the optimal flows by constructing a transport plan with two components. The first transports drivers from all locations in $A(z|p, \mathcal{T})$ to z ; the second transports

drivers within rays, excluding z . To obtain the latter flows along each segment $(z, X_a]$, we solve an optimal transport problem with cost function equal to the distance between any two points, initial measure equal to the reminder mass that was not sent to z , and final measure equal to the restriction of the solution of Problem $(\mathcal{P}_{KP}(z))$ in $(z, X_a]$. Finally, we apply the pasting result (Proposition B-2 in Online Appendix B) to obtain a feasible price equilibrium in the whole city $\mathcal{C}$.

8. A 1D Family of Instances: Global Optimal Solution and Insights

The results derived in previous sections characterize the structure of an optimal pricing policy and the corresponding supply response *locally in attraction regions* for general demand and supply conditions in two dimensions. In this section, to further understand the interplay of spatial supply incentives and spatial pricing, we aim to develop insights into the global optimal policy and the interactions across attraction regions. To that end, we focus on a family of instances that will be rich enough to capture spatial supply demand imbalances while isolating the interplay above. Leveraging previous results for the local form of optimal solutions but also developing new results to characterize attraction regions, we determine the global optimal solution for this family of instances in quasi-closed form. We will see how the optimal solution creates a spatial menu of prices to incentivize but also to discourage supply movement.

In particular, we focus on a 1D city and a family of models that captures a potential local surge in demand. The restriction to a 1D setting is for simplicity and exposition. The results in this section can be viewed as the 1D cut of the solution to a disk-shaped city in which the conditions along diameters are symmetric. (That is, exhibiting the 1D solution is enough to characterize the 2D solution in this special case.) More precisely, we specialize the model to the case where the city measure is supported on the interval $\mathcal{C} = [-H, H]$ and is given by³

$$\Gamma(\mathcal{B}) = \mathbf{1}_{\{0 \in \mathcal{B}\}} + \int_{\mathcal{B}} dx, \quad \text{for any measurable set } \mathcal{B} \subseteq \mathcal{C};$$

that is, the origin may admit point masses of supply and demand, whereas the rest of the locations in $\mathcal{C}$ only admit infinitesimal amounts of supply and demand. We fix the city measure throughout, but we parametrize the supply and demand measures. Supply is initially evenly distributed throughout the city, with a density of drivers equal to θ_1 everywhere. Potential demand will also be assumed to have a uniform density on the line interval, except potentially at the origin.

We analyze what happens when a potential demand shock at the origin (the potential high-demand location) materializes, and in particular, we investigate the optimal pricing policy in response to such a shock. We represent the demand shock by a Dirac delta at this location. Therefore, for any measurable set $B \subseteq \mathcal{C}$, the potential demand measure (after the shock) is given by

$$\Lambda(\mathcal{B}) = \lambda_0 \cdot \mathbf{1}_{\{0 \in \mathcal{B}\}} + \int_{\mathcal{B}} \lambda_1 dx,$$

where $\lambda_0 \geq 0$ and $\lambda_1 > 0$. In particular, we refer to the case $\lambda_0 = 0$ as the *predemand shock environment* and the case $\lambda_0 > 0$ as the *demand shock environment*. For this family of models, we assume that customer willingness-to-pay distribution is location independent and denoted by $F(\cdot)$. This special structure will enable us to elucidate the spatial supply response induced by surge pricing and the structural insights on the optimal policies that emerge.

Throughout this section, we will use short-hand notation to present the optimal solution in a streamlined fashion. Let $(p, \mathcal{T})$ be a price-equilibrium pair; we use $A(0)$, X_l and X_r to denote $A(0|p, \mathcal{T})$ and the end points of the left and right rays around z , respectively. Moreover, when clear from context, we write $V(\cdot)$ instead of $V(\cdot|p, \mathcal{T})$.

8.1. The Predemand Shock Environment

We start by analyzing the preshock environment. Both demand and supply are uniformly distributed along the city, with respective densities λ_1 and θ_1 . In line with intuition, in this highly symmetric setting, one can show that the optimal price policy does not induce any movement of supply, and the optimal price at each location is the same and simply that of a single-location-capacitated pricing problem. We denote the predemand shock optimal price by $\rho_1 = \rho_x^{loc}(\theta_1)$ (not location dependent) and use ψ_1 to denote $\psi_x(\theta_1)$ (also not location dependent). For completeness, this is formalized in Proposition D-1 in Online Appendix D.

8.2. Benchmark: Myopic Price Response to a Demand Shock

We next start our analysis of the demand shock environment. Before turning our attention to an optimal policy in Section 8.2, we first focus on a simple type of pricing heuristic that responds to changes in demand conditions through changes in prices *only* where these changes occur. In particular, in the context of the demand shock model, this corresponds to responding to a shock in demand at the origin by only adjusting the price at the origin; we call this policy the *myopic price response*. This provides a benchmark to better understand the structure and performance of an

optimal policy. We next characterize the optimal myopic price response.

Proposition 4 (Myopic Price Response to a Demand Shock). *Fix $\lambda_0 > 0$. Suppose that $p(x) = \rho_1$ for all x in $\mathcal{C} \setminus \{0\}$ and that the firm optimizes the price $p(0)$. Then*

a. (Prices) *The optimal price at the origin is given by $p(0) = \rho_0^{loc}(s^{\mathcal{T}}(0))$ and $p(0) \geq \rho_1$.*

b. (Movement) *There exists two thresholds $X_r \geq X_r^0 \geq 0$ such that $X_r > 0$ and*


• *For all x in $[-X_r^0, X_r^0]$, all the supply units move to the origin.*

• *For all x in $[-X_r, -X_r^0]$ and all x in $[X_r^0, X_r]$, a fraction of the supply units moves to the origin, and the other fraction does not move.*

• *For all x in $\mathcal{C} \setminus [-X_r, X_r]$, no supply unit moves. Furthermore, the platform's revenue is strictly larger than in the predemand shock environment.*

This result characterizes the structure of a myopic price response as well as the structure of the supply movement it induces. Figure 8 depicts the structure of the supply response. In particular, the myopic price response leads to a higher price at the origin to respond to the surge in demand at that location. In turn, this higher price attracts drivers from a symmetric region around the origin. In that region, for locations close to the origin, all supply units move to the origin. After a given threshold X_r^0 , only a fraction of the drivers will move to the origin. Intuitively, as one gets further from the origin, traveling to it becomes a less attractive option compared with staying put or traveling elsewhere, and an increasingly smaller fraction of units travels to the origin. We also establish that supply units have no incentive to travel anywhere else in the city. As a result, units that do not travel to the origin stay put and serve local demand. Beyond the threshold X_r , no supply units move in the equilibrium induced by the myopic price response. In a supply-constrained regime $\theta_1 \leq \lambda_1 \cdot \bar{F}(\rho^u)$, all drivers within $[-X_r, X_r]$ drive to the origin; that is, $X_r^0 = X_r$. However, in a supply-unconstrained regime $\theta_1 > \lambda_1 \cdot \bar{F}(\rho^u)$, the two thresholds are different, $X_r^0 < X_r$, as depicted in Figure 8. This occurs because in locations further from the origin but still within $[-X_r, X_r]$ because underutilized drivers drive toward the origin, conditions at the departing point improve, and in equilibrium, staying put becomes competitive with driving to the origin.

Figure 8. (Color online) Myopic Price Response: Induced Supply Response for a Case with $\theta_1 > \lambda_1 \cdot \bar{F}(\rho^u)$



8.3. Optimal Solution

In this section, we focus on the optimal *global* price response across all locations in the city. To that end, we will first develop results to identify the attraction regions in the city and then leverage the results developed for the general model to ultimately obtain a quasi-closed-form solution to the platform's problem in this specialized setting. Our first result demonstrates that we can focus on price-equilibrium pairs such that the high-demand location is a sink.

Lemma 2 (Origin Is a Sink). *Without loss of optimality, one can restrict attention to price-equilibrium pairs $(p, \mathcal{T})$ such that the origin is a sink such that $X_l < 0 < X_r$.*

The intuition behind this lemma harks back to the fact that the performance of the preshock environment is dominated by that of the myopic price response solution. Solutions for which the origin is not a sink have revenues capped by that of the predemand shock environment. At a high level, in those solutions, there is no positive mass of drivers willing to travel to the demand-shock location, and thus the city resembles a city without a demand shock. However, the myopic price response solution incentivizes drivers from both sides to travel to the demand shock and has a strictly larger revenue. This implies that at optimality we must have drivers coming from both sides to the origin; that is, $X_l < 0 < X_r$. An important consequence of Lemma 2 is that the attraction region $A(0)$ is a well-defined nonempty set. We thus could apply Theorem 1 to obtain a local characterization of the optimal solution within $A(0)$. However, our goal in this section is to obtain the full global optimal solution as opposed to just a solution in $A(0)$. Hence, before we use Theorem 1, in what follows we characterize all attraction regions in $\mathcal{C}$. To make our exposition clear and highlight the solution's spatial aspects, we call the interval $[X_l, X_r]$ the *center* region; the region outside of it will be referred to as the *periphery*.

8.3.1. Equilibrium Utilities and Attraction Regions. In this section, we characterize $V(\cdot)$ throughout $\mathcal{C}$. This characterization is key because it will enable us to identify all the attraction regions in $\mathcal{C}$.

Theorem 2 (Equilibrium Utilities). *Under an optimal price-equilibrium pair $(p, \mathcal{T})$, the equilibrium utility function $V(\cdot)$ is fully parametrized by the three values $V(0)$ and X_l, X_r as follows:*


$$V(x) = \begin{cases} V(0) - |x|, & \text{if } x \in [X_l, X_r], \\ \min\{V(0) - 2X_r + x, \psi_1\}, & \text{if } x > X_r, \\ \min\{V(0) - 2|X_l| + |x|, \psi_1\}, & \text{if } x < X_l. \end{cases}$$

Moreover, $V(0) > \psi_1$ and $V(X_l), V(X_r) \leq \psi_1$.

The first main implication of this result is that we know, up to $V(0)$, X_l , and X_r , how much utility each supply unit garners under optimal prices throughout the entire city. Quite strikingly, the characterization of $V(\cdot)$ is “independent” of the flows. That is, in order to characterize the equilibrium utility, we did not need to pin down the distribution of postmovement supply.

The second implication is that the city has at most three types of regions. Figure 9 depicts the equilibrium utility function. The center $[X_l, X_r]$ is by definition an attraction region. Let W_r and Y_r be defined as the points to the left and right of X_r , where the driver's equilibrium utility function equals the preshock utility level ψ_1 . To the right of the origin (and similarly to the left), we can observe three main regions. We first have the interval $[0, W_r]$, where drivers' utilities are above the preshock level. Drivers in this region are positively impacted by the shock of demand at the origin (and the global optimal prices). The second region $[W_r, Y_r]$ is notable. Here drivers garner strictly less utility compared with the preshock environment. In $[W_r, X_r]$, drivers are “too far” from the origin, so their utilities are negatively affected by the cost of driving to the origin. Drivers in $[X_r, Y_r]$ are outside the origin's attraction region and thus do not relocate to the origin. This interval forms an attraction region with sink Y_r ; that is, Y_r belongs to the indifference region of any location in the interval, and Y_r does not belong to the indifference region of any other location. In turn, besides $A(0)$, there are two other attraction regions, $A(Y_l)$ and $A(Y_r)$, in $\mathcal{C}$. Interestingly, drivers in $[X_r, Y_r]$ suffer because the platform has to make sure that the drivers in $[0, X_r]$ stay within the attraction region of the origin. For the marginal drivers at X_r to be willing to travel to the origin, the conditions to the right of X_r should not be too attractive. The final region corresponds to $[Y_r, H]$; this region is not affected by the shock of demand because it is far from the origin.

Figure 9. Drivers' Equilibrium Utility Under an Optimal Pricing Policy



Note. The equilibrium utility is fully characterized up to $V(0)$, X_l , and X_r ; the intervals $[Y_l, X_l]$, $[X_r, Y_r]$ and $[X_l, X_r]$ are attraction regions.

8.3.1.1. Key Ideas for the Proof of Theorem 2. The proof relies on leveraging structural properties of the equilibrium utility function, the congestion bound, and a novel flow-mimicking technique. At a high level, we focus on and solve for $V(\cdot)$ in each region separately, center and periphery.

We start by considering the center region, which is easy to analyze. Lemma 2 establishes that we can focus on solutions such that $A(0) = [X_l, X_r]$ is a nonempty interval that strictly contains the origin. Then, by definition of $A(0)$, $V(x) = V(0) - |x|$, for all $x \in [X_l, X_r]$. Importantly, the characterization of $V(\cdot)$ in this region only depends on three parameters, namely $V(0)$, X_l , and X_r .

Switching attention to the periphery, consider the right periphery $(X_r, H]$ (the treatment for the left periphery is analogous). We first argue that in this region, the drivers' equilibrium utility has a non-trivial upper bound

$$V(x) \leq \min\{V(X_r) + x - X_r, \psi_1\}, \quad \text{for all } x \in (X_r, H]. \quad (7)$$

This upper bound follows from two bounds. A first upper bound can be derived using the 1-Lipschitz property of V (Lemma 1), which limits the growth rate of V . Thus, $V(x)$ is bounded by $V(X_r) + x - X_r$. A second bound may be obtained by leveraging the congestion bound (Proposition 2). One may show that drivers from almost any location who do not have an incentive to travel to the origin have their utilities capped by the predemand shock utility level ψ_1 .

The core of the argument toward characterizing the equilibrium utilities in the periphery resides in establishing that the upper bound in Equation (7) is always binding. We show this in two steps. We first establish that the value function has to be nondecreasing in $[X_r, H]$ (see Proposition D-2 in Online Appendix D); this implies that drivers only move right (or do not move) in the right peripheral region. Then, exploiting the monotonicity, we use a flow-mimicking argument to establish that the upper bound is achieved under an optimal pricing policy (Proposition D-3 in Online Appendix D).

8.3.2. From Equilibrium Utilities to Supply Distribution and Optimal Prices. Given that we pinned down the equilibrium utility function across the city and all attraction regions, we next solve for prices and supply through the problem reformulation in Proposition 1. Leveraging Theorem 1 and a symmetry argument, we can solve for the optimal $s^\mathcal{T}$ and the corresponding prices in each attraction region. The solution for the no-movement regions reduces to the preshock environment. Then we use the pasting property (compare Property 3 in Section 6) to paste the solution from each

region and, in turn, obtain a quasi-closed-form characterization of the optimal solution to the platform's problem as presented in Theorem 3.

Theorem 3 (Optimal Prices and Flows). *An optimal price-equilibrium pair $(p, \mathcal{T})$ is such that $V(\cdot)$ is as in Theorem 2, $X_r = -X_l$, and prices and flows are characterized as follows:*

1. (Prices) *The optimal prices are given by $p(x) = \rho_x^{\text{loc}}(s^\mathcal{T}(x))$, where $s^\mathcal{T}(x)$ is as below.*

2. (Postrelocation supply) *There exists unique $\beta_c \in [0, W_r]$ and $\beta_p \in [X_r, Y_r]$ such that*

$$\begin{aligned} \int_{-\beta_c}^{\beta_c} \psi_x^{-1}(V(x)) d\Gamma(x) &= \theta_1 \cdot 2 \cdot X_r \quad \text{and} \\ \int_{\beta_p}^{Y_r} \psi_x^{-1}(V(x)) d\Gamma(x) &= \theta_1 \cdot (Y_r - X_r), \end{aligned}$$

and the optimal postrelocation supply is given by

$$s^\mathcal{T}(x) = \begin{cases} 0, & \text{if } x \in (\beta_c, \beta_p) \cup (-\beta_p, \beta_c), \\ \psi_x^{-1}(V(x)), & \text{otherwise.} \end{cases}$$

3. (Movement)


- For all x in $[-\beta_c, \beta_c]$, drivers move in the direction of the origin.
- For all x in $[-X_r, -\beta_c) \cup (\beta_c, X_r]$, all drivers move to $[-\beta_c, \beta_c]$.
- For all x in $[X_r, \beta_p]$ (respectively $(-\beta_p, -X_r]$), all drivers move to $[\beta_p, Y_r]$ (respectively $[-Y_r, -\beta_p]$).
- For all x in $[\beta_p, Y_r]$ (respectively $[-Y_r, -\beta_p]$), drivers move in the direction of Y_r (respectively $-Y_r$).
- For all x in $[-H, -Y_r) \cup (Y_r, H]$, drivers do not relocate.

In other words, we have fully characterized the optimal solution across the city, and it is fully parametrized by only on two values, $V(0)$ and X_r .

8.3.2.1. Discussion. We depict in Figure 10 the structure of the solution obtained in Theorem 3. The main feature of the optimal solution is that it separates each side of the city with respect to the origin into multiple regions. For clarity, we focus our discussion on the right side of the city.

The origin receives a mass of supply equal to $\psi_0^{-1}(V(0))$. This mass of drivers comes from two regions, the inner and the outer center, which we now define. The first corresponds to the interval $(0, \beta_c]$. Some drivers in this region choose to stay put, whereas others, attracted by the favorable conditions at the center of the city, choose to drive to the origin. In equilibrium, drivers staying or traveling to the origin garner the same utility. The outer center is the interval $(\beta_c, X_r]$. Here the platform sets prices to $\bar{V}$ (or 0), and therefore, supply is equal to zero. That is, the platform chooses prices to shut down demand, giving no incentive for drivers to stay there (or, alternatively, sets prices at zero to again

Figure 10. (Color online) Supply Response (Solid Line) Induced by Optimal Prices (Dashed Line)



give no incentive for drivers to stay there). In turn, this incentivizes all drivers in this region to move somewhere else. In order to incentivize these drivers to move toward the origin, the platform creates one more region: the inner periphery.

The inner periphery corresponds to the interval $(X_r, Y_r]$. The platform “artificially” degrades the conditions for drivers in this interval in two different ways, leading to the two subregions, (i) and (ii) in Figure 10. In region (i), the platform sets prices equal to $\bar{V}$ (or 0) in $(X_r, \beta_p]$, shutting down demand, so no drivers want to either travel to or stay in this region. As a result, the interval $(\beta_p, Y_r]$ receives all drivers from $(X_r, \beta_p]$. This creates driver congestion and thus endogenously worsens driver conditions in the interval $(\beta_p, Y_r]$. The reason the platform selects these inner periphery prices is to discourage drivers in the outer center from driving toward the periphery. Quite strikingly, the optimal global price response to a demand shock at the origin induces supply movement *away* from the origin in the inner periphery. The final region is the outer periphery. All drivers in this region stay put, leading to $s^T(x) = \theta_1$. Here drivers collect the same utility they would if there was no demand shock at the origin.

In sum, the optimal global price response to a demand shock, while correcting the supply demand imbalance at the origin, also creates significant imbalances across the city. This is driven by the self-interested nature of capacity units and the need to incentivize them through spatial pricing. See Proposition 4 for how the optimal policy differs from the myopic best response.

8.4. Myopic Price Response vs. the Optimal Solution

In this section, we will use the myopic price response solution as a benchmark for comparison and put the optimal solution into perspective. The objective is to illustrate through several metrics the different features of the optimal solution as well as its performance in terms of revenue maximization (we complement the comparison, including welfare performance, in Online Appendix E). Throughout this section, we use

superscripts to label relevant quantities associated with the myopic price response and optimal solution, respectively (except when obvious from context).

We first observe that the attraction region around the origin of the demand-shock location is always wider under the optimal solution than under the local best response. That is, $A^{\text{my}}(0) \subset A^{\text{opt}}(0)$. In particular, this means that more locations are affected by a demand shock in the optimal solution than under the myopic price response. Hence, the largest interval in which both solutions differ corresponds to $[-Y_r^{\text{opt}}, Y_r^{\text{opt}}]$. We denote this interval by $\mathcal{C}_{\text{diff}}$.

Next, we illustrate and discuss through a set of numerics the differences between the two policies. In order to obtain numerical solutions for the global optimal policy, we rely on Theorems 2 and 3. From Theorem 2, we know that $V(x)$ is characterized by three values: $V(0), X_r, X_l$. In Theorem 3, given $V(\cdot)$, we provide a full characterization of the optimal solution. Also, we establish that $X_l = -X_r$. In turn, to find the optimal solution, we perform a grid search over $[0, H] \times [0, \bar{V}]$. For the myopic price response, we proceed in a similar fashion by making use of the closed-form expressions developed in the proof, in Online Appendix D, of Proposition 4. We consider a range of instances that includes various levels of supply availability. We fix the city to be characterized by $H = 1$ and assume that the demand is uniformly distributed across locations with $\lambda_1 = 4$. The origin experiences a shock of demand ranging from low to high: $\lambda_0 \in \{3, 6, 9\}$. We vary the initial supply $\theta_1 \in \{1, 1.5, 2, \dots, 4.5, 5\}$ so that when it is low, the city (excluding the origin) is supply constrained, and when it is high, the city is supply unconstrained. Consumer valuation is uniformly distributed in the unit interval. Note that the city (excluding the origin) is supply constrained whenever $\theta_1 < \lambda_1 \cdot \bar{F}(p^u) = 2$. To eliminate any strong dependence on the choice of H , for each instance, we compare the myopic price response performance and optimal solution performances within the subregion of the city corresponding to the largest interval in which both solutions differ $\mathcal{C}_{\text{diff}}$.

Table 1. Revenue Improvement (in %) of Optimal Solution over Myopic Price Response in $\mathcal{C}_{\text{diff}}$

θ_1	1	1.5	2	2.5	3	3.5	4	4.5	5
$\lambda_0 = 3$	2.05	4.64	9.59	13.02	13.87	12.92	11.00	8.60	5.91
$\lambda_0 = 6$	2.17	3.11	4.99	8.73	9.96	10.01	9.56	8.92	8.21
$\lambda_0 = 9$	2.69	3.51	4.69	8.75	10.16	10.30	9.81	9.10	8.29

8.4.1. Revenue Improvement. The revenue performance of the optimal solution with respect to our benchmark in $\mathcal{C}_{\text{diff}}$ is shown in Table 1. For any level of demand shock, we observe that the revenue improvement reaches its maximum value for medium to high levels of supply and can be significant ($>10\%$). In order to appreciate where the revenue gains stem from, consider Table 2, which summarizes some key quantities for the case $\theta_1 = 3, \lambda_0 = 9$ (so that $\psi_1 = 0.27$).

Let us analyze the various contributions to revenues under both policies. We start by noticing that the drivers' equilibrium utility at the shock location is lower under the optimal solution than under the myopic price response, $V^{\text{opt}}(0) = 0.62$ and $V^{\text{my}}(0) = 0.65$. However, because $X_r^{\text{opt}} = 0.46$ and $X_r^{\text{my}} = 0.38$, the optimal solution is able to incentivize the movement of a larger mass of drivers toward the demand shock, leading to a mass $s^{\text{opt}}(0) = 1.97$ versus $s^{\text{my}}(0) = 1.66$. Focusing on the objective reformulation in Proposition 1, this extra mass of drivers delivers 0.14 units ($0.62 \times 1.97 - 0.65 \times 1.66$) of extra revenue to the platform. The revenue difference is further increased by the fact that the remaining 0.79 units of drivers in the attraction region of zero ($2 \times 3 \times 0.46 - 1.97$) in the optimal solution travel to locations nearby the demand shock, where $V(\cdot)$ is close to 0.62. In contrast, the benchmark solution has the remainder 0.62 drivers ($2 \times 3 \times 0.38 - 1.66$) staying within $[X_r^{0,\text{my}}, X_r^{\text{my}}]$ where $V(\cdot)$ is below 0.37 ($V^{\text{my}}(0) - X_r^{0,\text{my}}$). Through these two mechanisms, the optimal policy garners more revenue than the benchmark solution in the region $[-X_r^{\text{opt}}, X_r^{\text{opt}}]$.

However, the benefits come at a cost. To induce the “right” incentives in the shock's attraction region, the platform has to alter conditions to the right of the attraction region. In order to incentivize the movement of drivers in $[-X_r^{\text{opt}}, X_r^{\text{opt}}]$ toward the demand shock, the region $[X_r^{\text{opt}}, Y_r^{\text{opt}}]$ is damaged by having the 0.22 units of drivers in it ($2 \times (0.57 - 0.46)$) contributing values strictly below $\psi_1 = 0.27$ to the platform's objective. The same units of drivers in the benchmark solution contribute exactly 0.27 per unit to the platform's revenue. This cost is offset by the

proceeds that incentivizing the movement of a larger amount of drivers toward the demand shock generates.

9. Conclusion

This study analyzes the short-term pricing problem faced by a platform matching spatially distributed demand to strategic supply units. Given supply and demand conditions across a 2D region, the platform sets prices at every location, and supply units select where to reposition in equilibrium. To analyze this problem, we employ a measure-theoretic framework that subsumes both discrete and continuous settings. The resulting problem is a mathematical program with equilibrium constraints for which no standard solution approach is readily available. We provide two main contributions. First, we establish a characterization of the optimal solution, prices and flows, within regions of potential movement (attraction regions). Our approach consists of relaxing some of the equilibrium constraints and identifying that our relaxation, localized to attraction regions, is tight and leads to coupled continuous bounded knapsack problems, which we solve to optimality. Our second contribution is in terms of managerial insights. Our results highlight how the optimal solution may employ both positive and negative incentives to induce the right movement of drivers, raising prices in profitable undersupplied regions but also potentially damaging (using prices) less profitable regions, thus incentivizing the relocation of drivers toward regions that are more beneficial for the platform.

There are several potential directions for future work. Our results may be used to study possible heuristics that approximate the global optimal solution in arbitrary instances. A direction would be to develop two-stage heuristics that first set candidate locations and shapes for attraction regions and then leverage the results developed to optimize within those. One approach is to first parametrize their shape, for example, using circles or hexagons, and then do a search for regions of large supply-demand imbalances in order to identify sink locations. Each region then would be parametrized by the shape of its border and the utility at the sink location. We can then solve for the optimal solution within each of them and paste to obtain a candidate solution. A master global optimization would follow to tune the shape and the utilities at the sink locations across the city. That is, the general methodology we developed may be leveraged to compute parametric global solutions to the

Table 2. Metrics for the Local Response and Optimal Solution for the Case $\theta_1 = 3, \lambda_0 = 9$

$V^{\text{opt}}(0)$	$s^{\text{opt}}(0)$	$p^{\text{opt}}(0)$	X_r^{opt}	Y_r^{opt}	$V^{\text{my}}(0)$	$s^{\text{my}}(0)$	$p^{\text{my}}(0)$	X_r^{my}	$X_r^{0,\text{my}}$
0.62	1.97	0.78	0.46	0.57	0.65	1.66	0.81	0.38	0.25

platform's problem. Another important extension is the incorporation of dynamics. The model studied in this paper can be regarded as a two-stage model. In the first stage, drivers are initially positioned. In the second stage, they reposition in equilibrium given prices and demand conditions. This does not consider that drivers who are closer to a given location might be more likely to be matched to riders in such a location (they can get to that location faster than other drivers). It also does not consider the continuation game that emerges after drivers are matched. Studying these different settings and their variations constitute interesting avenues of future inquiry.

Finally, the framework and ideas developed here could be leveraged in settings beyond ride hailing. Given the generality of our framework, all results before Section 8 can be extended to higher-dimensional settings. One could use our framework to study problems in which agents have different types characterized by high-dimensional vectors. Agents can modify their types by exerting some costly effort, but their potential earning when doing so will depend on how many other agents end up being of the same type.

Endnotes

¹ In precise terms, the mass of customers with willingness to pay below q in a measurable region $\mathcal{B} \subseteq \mathcal{C}$ is given by $\int_{\mathcal{B}} F_y(q) d\Lambda(y)$. We assume that $y \mapsto F_y(q)$ is a measurable mapping for any q .

² The Radon–Nikodym derivative can also be interpreted as a measure of the units of demand or supply per unit of area (e.g., square mile).

³ Observe that, thanks to the generality of our measure-theoretic framework, all the structural results developed thus far apply to this 1D setting.

References

- Afèche P, Liu Z, Maglaras C (2017) Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Working paper, University of Toronto, Toronto.
- Ambrosio L, Pratelli A (2003) Existence and stability results in the L^1 theory of optimal transportation. Caffarelli LA, Salsa S, eds. *Optimal Transportation and Applications* (Springer, Berlin), 123–160.
- Banerjee S, Freund D, Lykouris T (2016) Multi-objective pricing for shared vehicle systems. Preprint, submitted August 24, <https://arxiv.org/abs/1608.06819>.
- Banerjee S, Riquelme C, Johari R (2015) Pricing in ride-share platforms: A queueing-theoretic approach. Working paper, Cornell University, Ithaca, NY.
- Bimpikis K, Candogan O, Saban D (2019) Spatial pricing in ride-sharing networks. *Oper. Res.* 67(3):744–769.
- Blanchet A, Carlier G (2015) Optimal transport and Cournot–Nash equilibria. *Math. Oper. Res.* 41(1):125–145.
- Blanchet A, Carlier G, Nenna L (2018) Computation of Cournot–Nash equilibria by entropic regularization. *Vietnam J. Math.* 46(1):15–31.
- Buchholz N (2017) Spatial equilibrium, search frictions and efficient regulation in the taxi industry. Working paper, Princeton University, Princeton, NJ.
- Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing Service Oper. Management* 19(3):368–384.
- Castillo JC, Knoepfle D, Weyl G (2017) Surge pricing solves the wild goose chase. *Proc. 2017 Conf. Econom. Comput.* (ACM, New York), 241–242.
- Cole R, Dodis Y, Roughgarden T (2003) Pricing network edges for heterogeneous selfish users. *Proc. 35th Annual ACM Sympos. Theory Comput.* (ACM, New York), 521–530.
- Deneckere RJ, McAfee PR (1996) Damaged goods. *J. Econom. Management Strategy* 5(2):149–174.
- Feng G, Kong G, Wang Z (2017) We are on the way: Analysis of on-demand ride-hailing systems. Working paper, University of Minnesota, Minneapolis.
- Guda H, Subramanian U (2019) Your Uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Sci.* 65(5):1995–2014.
- Gurvich I, Larivière M, Moreno A (2016) Operations in the on-demand economy: Staffing services with self-scheduling capacity. Working paper, Northwestern University, Evanston, IL.
- Hu M, Zhou Y (2016) Dynamic type matching. Working paper, University of Toronto, Toronto.
- Hu M, Zhou Y (2017) Price, wage and fixed commission in on-demand matching. Working paper, University of Toronto, Toronto.
- Lu A, Frazier PI, Kislev O (2018) Surge pricing moves Uber's driver-partners. *Proc. 2018 ACM Conf. Econom. Comput.* (ACM, New York), 3.
- Mas-Colell A (1984) On a theorem of Schmeidler. *J. Math. Econom.* 13(3):201–206.
- Nikzad A (2018) Thickness and competition in ride-sharing markets. Working paper, Stanford University, Stanford, CA.
- Özkan E, Ward AR (2020) Dynamic matching for real-time ride-sharing. *Stochastic Systems* 10(1):29–70.
- Roughgarden T, Tardos É (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.
- Schmeidler D (1973) Equilibrium points of nonatomic games. *J. Statist. Phys.* 7(4):295–300.
- Talluri KT, Van Ryzin GJ (2004) *The Theory and Practice of Revenue Management*, International Series in Operations Research and Management Science, vol. 68 (Springer US, New York).
- Taylor T (2018) On-demand service platforms. *Manufacturing Service Oper. Management* 20(4):704–720.
- Villani C (2008) *Optimal Transport: Old and New*, vol. 338 (Springer Science & Business Media, Berlin).