Cardiff Economics Working Papers

Working Paper No. E2022/11

Incentivizing flexible workers in the gig economy: The case of ride-hailing

Cemil Selcuk and Bilal Gokpinar

June 2022

ISSN 1749-6010

Cardiff Business School
Cardiff University
Colum Drive
Cardiff CF10 3EU
United Kingdom
t: +44 (0)29 2087 4000
f: +44 (0)29 2087 4419
business.cardiff.ac.uk

Incentivizing flexible workers in the gig economy: The case of ride-hailing

Cemil Selcuk
Cardiff Business School, Cardiff University, Cardiff CF10 3EU, UK, selcukc@cardiff.ac.uk
Bilal Gokpinar
UCL School of Management, University College London, London E14 5AB, UK, b.gokpinar@ucl.ac.uk

Abstract: Creating the right incentives for a flexible workforce lies at the heart of the gig economy. For most companies, a key question is how to best connect a limited number of independent workers in their platforms with service-seeking consumers through the right pricing and matching mechanisms. We focus on ride-hailing where drivers have significant discretion over where and when to work across different locations. Building a spatial model, we study how a platform can create incentives for independent drivers via prices and commissions, and how such policies affect drivers' search behavior across a network of locations.

Contrary to common perception, we find that the flexibility of the commissions, and not the flexibility of prices, plays a dominant role in resolving local demand and supply mismatch. This is because location based price hikes at the bottlenecks negatively distort the local demand and generally do a poor job in incentivizing drivers towards such locations. Adjusting the commissions, on the other hand, does not interfere with the local demand; creates better incentives for the drivers, and therefore is more suitable to mitigate the effects of bottlenecks. Simulations based on actual ride patterns from New York City and Los Angeles confirm our insights.
Keywords: Ride-sharing, Gig workers' compensation, Flexible commission, Sharing economy

1 Introduction

Traditional taxi markets exhibit significant meeting frictions associated with taxi search and availability (Lagos, 2000; Frechette et al., 2016; Buchholz, 2021). Vacant cabs often spend a long time waiting at certain locations while passengers wait for cabs at other locations. Using little technology, cabs find passengers via slow dispatching from local companies or by random street hails. In recent years, making use of new technologies—mobile devices, location tracking, navigation—ride-hailing platforms such as Uber and Lyft disrupted the traditional taxi industry and revolutionized the transportation industry as a whole by improved matching efficiency. Founded in 2010, Uber has reported completion of 10 billion trips worldwide in just eight years (Uber, 2018) and is currently operating in over 83 countries and 10,000 cities (Uber, 2021).

A key challenge faced by the ride-hailing industry is the shortage of drivers. Finding and keeping drivers have always been a challenge for ride-hailing platforms as they experience significant driver turnover (Brown, 2019; Cook et al., 2020). Indeed, according to a report by Uber, 11% of new drivers stop driving within a month, and about half of them leave within a year (Huet, 2015c).

Driver shortage has been exacerbated recently due to a combination of factors including pandemic related health and safety concerns (Siddiqui, 2021), the emergence of alternative work options such as food and grocery delivery (Bursztynsky, 2021), and ongoing issues with drivers' employment rights, working and pay conditions (Paul, 2021). Although part of the shortage may be temporary as several platforms are reported to be "throwing money" to drivers to address severe shortage (D.Lee, 2021), limited supply of drivers will likely stay as a key feature of the ride-hailing platforms for the foreseeable future.

Given limited supply of drivers, how can a platform create the right incentives using alternative compensation schemes to generate as many matches as possible? Workforce compensation plans and incentives have been widely studied in the existing marketing literature e.g., salespeople compensation (Bhargava and Rubel, 2019; Kim et al., 2019). Such issues, however, require renewed attention within the context of an independent and flexible workforce commonly found in platform based industries. Indeed, the platform cannot instruct the drivers when/where to search as it does not employ them. It can, however, create incentives via prices and commissions, anticipating that the drivers will respond to such incentives.

Considering the spatial differentiation of supply and demand, if the incentives are not aligned properly, the platform inevitably experiences bottlenecks and shortages in some locations and excess supply at other locations. In order to study the strategic impact of such incentives, we consider alternative mechanisms in which prices and commission rates can be either fixed or flexible. With fixed-pricing, the platform charges the same per-mile price across the entire city, whereas with flexible-pricing it sets location-specific prices. Fixed and flexible commission rates are defined likewise. While the fixed commission model appears to be the norm among ride-sharing platforms (in most cities Uber gets a fixed 25% commission and drivers keep 75% of the revenue (Uber, 2020)), an alternative approach is to charge flexible commission rates based on ride locations. Indeed, Uber has already experimented with the possibility of charging different commission rates in different cities. For example, in 2015 it raised the commission rate from 20% to 25% in New York, Toronto, Indianapolis, Boston and Worcester (Huet, 2015a). Similarly, in a pilot program in San Francisco, it announced testing a tiered structure where UberX drivers pay a 30% commission on their first 20 rides in a week, 25% on their next 20 rides, and then 20% on any rides beyond that (Huet, 2015b). These anecdotes indicate that ridesharing platforms can indeed pursue more granular commission policies. In fact, various platforms operating two-sided markets are already charging flexible commissions tailored for each individual job (Taylor, 2014).

We should further mention that by flexible-pricing we do not mean "surge pricing" as practiced by Uber, which resolves temporary overflows in demand, due to, say, major sports events or bad weather. Instead we note that some locations have fundamentally different demand patterns as they consistently pull in and send away more traffic than others[1]. Such persistent long term differences call for flexible schemes in which prices and commission rates can be conditioned on location specific factors. If, despite such differences, the platform pursues a rigid, fixed-price-fixed-commission rule, then it is likely to generate fewer matches than it potentially could; thus lose out both on profits and consumer surplus.

With this in mind, we develop an analytically tractable model, in which the platform operates on a network of locations with differing distances and traffic flows. Being aware of the city's layout and traffic flows, and taking into account the prices and commission rates set by the platform, drivers self select themselves across these locations to search for customers. We analyze three alternative pricing and commission models and compare their performance in terms of the number of matches, profits and consumer surplus.

Our analysis reveals several important insights. First, the flexible commission system is more suitable to create the right incentives for drivers than the flexible price system, especially when the platform does not have many vehicles at its disposal. When operating with a limited number of cars, the platform's key task is to incentivize drivers towards bottleneck locations (undesirable locations with short trip lengths). To this end, price interventions not only distort the interior demand and hamper profits, but also do a poor job in incentivizing drivers to spread themselves evenly across the city. We show that even after price hikes at the bottlenecks, there still remains excess supply at other locations, which is a waste of precious resources. In contrast, with a flexible commission policy the number of cars does not become a constraint until the customer-to-cab ratio ("utilization rate") hits 100% everywhere, i.e. until no car idles at any location. Up to that point, by fine-tuning the commission rates—decreasing them at more desirable locations, increasing them at less desirable locations, or a combination—the platform can spread the cars evenly across the city and while doing so, it does not distort the interior demand with unnecessary price hikes. Thanks to these features, the flexible-commission policy outperforms the flexible-price policy by creating more matches, generating more profits and, depending on parameters, generating more consumer surplus.

[1] For instance, based on actual ride patterns, we observe that in Los Angeles there is a consistently high flow of traffic to and from Santa Monica, West Hollywood or the Los Angeles International Airport, but the same is not true for, say, Studio City or Pacific Palisades. Relatedly, locations have different expected trip lengths, e.g. if a ride originates from West Hollywood then it tends to last for 3.9 miles, but if it originates from Pacific Palisades then it lasts for more than three times, 12.3 miles.

A second advantage of the flexible commission policy is that while price interventions are observable to everyone (passengers and drivers), commission interventions are observable to drivers only. Customers do not seem to like fickle fares (Dholakia, 2015; Ariely, 2016), which means that the flexible-price model might alienate customers and reduce profits. In contrast, the flexible commission model achieves the demand-supply match in a more subtle way in that neither the commission rates nor the adjustments to them are observable to the customers.

To illustrate the results in a real-world setting, we calibrate the model for New York City and Los Angeles based on ride patterns we extracted from a publicly available connectome map on Uber's website. Our simulations, in line with the preceding analytical insights, underline the benefits of flexible commissions when there is a shortage of drivers and their ability to utilize every car in a ride. Once this aspect is controlled for, the gains from flexible pricing seem to be modest. Second, we document that the performance of pricing models depends on how homogenous a city's traffic structure is in terms of the trip lengths and traffic flows. If these parameters show significant variation across the city, then pursuing a non-flexible policy is more "costly" for the platform. Our data suggest that Los Angeles has a less homogenous traffic structure (both in terms of trip lengths and in traffic flows across locations) than New York; thus non-flexible rules tend to fare worse in Los Angeles than in New York. Our subsequent simulations based on randomly generated cities with varying distances and transition matrices further confirm these insights.

Related Literature. Our study builds upon the existing literature in the taxi industry. Lagos (2000) highlights endogenous search frictions in the taxi-cab market. Buchholz (2021) considers a non-stationary environment by employing data from New York City and analyzing the dynamic spatial equilibrium of taxi-cabs. The main difference between this stream of work and our model is that a platform sets prices and commission rates in our study, whereas there is no such platform in this line of work with exogenous prices and commissions.

Our work is related to the emerging literature on sharing economy and peer-to-peer matching platforms (Einav et al., 2016; Zervas et al., 2017; Eckhardt et al., 2019; Li and Srinivasan, 2019; Yao et al., 2022) with a focus on ridehailing (Cramer and Krueger, 2016; Wang et al., 2019; Zhang et al., 2022). While there has been increased attention on the design of on-demand ride-hailing platforms and corresponding incentive schemes with an ultimate goal of better matching demand with supply, most of these studies have focused on addressing short term demand fluctuations, i.e. instantaneous imbalances between demand and supply, with dynamic "surge pricing" for given locations (Chen and Sheldon, 2015; Banerjee et al., 2015; Castillo et al., 2017).

A main feature of our study is the investigation and direct comparison of the trade-offs associated with fixed vs. flexible price and commission models. Ride-sharing platforms' pricing, wage and compensation decisions have attracted attention (Cachon et al., 2017; Hu and Zhou, 2019); however, these studies do not explicitly take into account spatial features of the city in which the platform operates. Indeed, a key aspect of the process of matching demand with supply in ride-sharing is the spatial differentiation of consumer demand and the direct influence of pricing policies on strategic search behavior of drivers across various locations, which has received relatively little attention in the literature. Exceptions include Guda and Subramanian (2019) which study surge pricing and information sharing in a two-zone-two-period setup. Their main interest is to understand when and in which market zones the platform will use a surge price and explore its implications. Our focus, however, is stationary characteristics such as different demand patterns across various locations in a city (i.e., transition matrix that governs customers' moves), which are arguably as critical as short term demand fluctuations across locations. Bimpikis et al. (2019) explore spatial price discrimination for a ride-sharing platform; however, in their model, the platform has access to an infinite supply of potential drivers, thus the number of cars does not become a constraint. In contrast, in our model the number of drivers—especially if it is insufficient—plays a crucial role in explaining the performance of different compensation schemes.

Finally, our work has connections with the marketing literature on incentives and compensation plan design (Basu et al., 1985; Coughlan and Sen, 1989; Jain, 2012; Chan et al., 2014). Previous literature empirically analyzed implications of various compensation schemes including commission and bonus based plans (Misra and Nair, 2011; Kishore et al., 2013; Chung et al., 2014; Kim et al., 2019), and analytically explored how to best align salesperson's incentives with those of the firm's by considering moral hazard issues (Raju and Srinivasan, 1996; Kalra et al., 2003; Schöttner, 2017). More recent work focused on two-sided market platforms and examined compensation of salespeople employed by such platforms in the presence of network effects (Bhargava and Rubel, 2019). A common aim in this literature is to understand how different compensation schemes affect effort choices of the salespeople who have considerable autonomy and flexibility in their work. In a similar spirit, our study investigates how an on-demand platform designs incentives to effectively manage a highly independent and flexible workforce.

2 Model

Environment. Time is discrete and continues forever. We consider a city that consists of $n \geq 2$ locations and is populated by a continuum of people with size 1 and a continuum of cars with size θ. The number of people and cars at location i are denoted by $y_i > 0$ and $x_i > 0$ and they satisfy $\sum_{i=1}^{n} y_i = 1$ and $\sum_{i=1}^{n} x_i = \theta$. The physical distance between locations i and j is denoted by $\delta_{i,j}$ and people's moves across these locations are governed by a Markov process, characterized by the exogenous row stochastic transition matrix $T = (a_{i,j})_{n\times n}$ where $a_{i,j} > 0$ denotes the probability that a person at location i wants to go to location j.[2]

People and cabs are matched via an online platform that sets prices and commission rates. People's willingness to pay is uniformly distributed in $[0, 1]$, so, if the platform sets price p_i at location i then there are

$$r_i = y_i (1 - p_i)$$

passengers ("riders") willing to hire a cab at that location. Remaining people are assumed to use public transport or other means to travel, and they do not generate any revenue for the platform. The platform's software identifies cabs and passengers at location i and creates matches according to the following matching function:

$$m_i = \min \{r_i, x_i\}.$$

Cabs can drive only a single passenger per trip, and the assignments are random; thus, the probability that a driver who is searching at location i finds a passenger is equal to

$$\eta_i = \frac{m_i}{x_i} = \min \left\{ \frac{r_i}{x_i}, 1 \right\}.$$

Occasionally, we refer to η_i as the *utilization rate* at location i, because from the platform's point of view η_i represents the percentage of cabs utilized in a ride.

In addition to matching passengers to cabs, the platform sets prices and commission rates across the city. In terms of notation p_i refers to the per-mile price associated with rides originating from location i. Similarly the commission rate c_i refers to the percentage of the revenue that the driver keeps after completing a ride originating from location i. Drivers participate in this market if their expected earning is greater than or equal to their outside option, W, which is the exogenous wage they could earn in the labor market.

[2] The strict positivity of x_i, y_i and $a_{i,j}$ ensures that we avoid absorbing or null recurrent states in the steady state equilibrium.

The timing of events is as follows. Each period starts with a "matching session" in which vacant cars at each location are matched with passengers. All rides take one period to complete.[3] We ignore operating costs (petrol, insurance etc.) as one can redefine the outside option net of such costs. At the end of each period, passengers reach their destinations, matches are dissolved and the process starts again. Cars with no passengers must wait for the next matching session, however in the mean time they are free to relocate to another location. When deciding where to search, drivers not only take into account the probability of finding a customer η_i, but also the price p_i, the commission rate c_i and the average trip length originating from that location. Below we analyze the drivers' problem.

Drivers. Let V_i denote the value of searching at location i before the matching session starts, and U_i be the value of being unmatched once the session ends. A cab that is unable to get a passenger can either stay in the same location to search again in the next period, or move to another location if it is more advantegous to search there. It follows that the value of being unmatched at location i at the end of a period is equal to the discounted value of searching at the best location in the city at the beginning of the next period, i.e

$$U_i = \max\{\beta V_1, ..., \beta V_n\} \text{ for all } i.$$

Now turn to V_i. We have

$$V_i = \eta_i \sum_{j=1}^{n} a_{i,j} \max\{c_i p_i \delta_{i,j} + \beta V_j, U_i\} + (1-\eta_i) U_i.$$

With probability η_i the driver is assigned to a passenger, and with probability $a_{i,j}$ the passenger travels to location j. If the driver agrees to take this trip, then his payoff is equal to share of the revenue $c_i p_i \delta_{i,j}$ plus the discounted value of searching at location j, given by βV_j. If he refuses to travel to location j then he idles for that period and walks away with U_i. Drivers are allowed to refuse a match, but if they do so, they must wait until the next round, i.e. they cannot instantaneously re-enter the matching pool in an effort to draw a better ride. With probability $1-\eta_i$ he gets no passenger at all, in which case, again, he obtains U_i.

[3] In reality, trips may not finish exactly within the same time window; however, it is safe to assume that on a per mile basis long distance trips complete faster than short distance trips (e.g., expressways vs inner city roads); and therefore are more preferable to drivers. This preference creates an imbalance across the locations and affects drivers' search decision and the platform's pricing decision. The equal travel time assumption creates this imbalance in an analytically tractable way. It is also a common assumption in the literature; see for instance Besbes et al. (2021), Banerjee et al. (2018) or Lagos (2000).

Steady State Equilibrium. We focus on a steady state in which the distribution of people and cabs across locations remains stationary. To ensure this, we require the number of incoming rides to a location to be equal to the number of outgoing rides from that location, i.e.

$$m_i = a_{1,i}m_1 + ... + a_{n,i}m_n, \text{ for all } i.$$

The left hand side represents the outflow from i, whereas the right hand side is the inflow into i. Since each ride consists of one passenger and one car, the equation above ensures that in the steady state the number of cars *and* passengers at each location remain unchanged. The relationship holds across the entire city, so letting $\mathbf{m} = (m_1, ..., m_n)$ we write

$$\mathbf{m} = \mathbf{m}T, \tag{1}$$

where

$$T = \begin{bmatrix} 0 & a_{1,2} & .. & a_{1,n} \\ a_{2,1} & 0 & .. & a_{2,n} \\ : & : & : & : \\ a_{n,1} & a_{n,2} & .. & 0 \end{bmatrix}.$$

Lemma 1 *The number of rides in the steady state satisfies*

$$m_i = \sigma_i M, \tag{2}$$

where $\boldsymbol{\sigma} > \mathbf{0}$ is the unique steady state vector of the transition matrix T and $M = \sum_{j=1}^{n} m_j$.

In words, if there are a total M moves in the city, then a fraction $\sigma_i \in (0, 1)$ of those moves must be originating from location i. In the steady state, the incoming and outgoing traffic flows are equal to each other, so an alternative interpretation of (2) is that a fraction σ_i of the traffic must be directed towards i. Either way, the parameter σ_i is a proxy of how attractive the location is. If σ_i is high, then we infer that location i is highly attractive as it pulls in and sends away a lot of traffic.

Furthermore, we require that in the steady state no location should be more profitable than others; thus drivers should be indifferent across locations, i.e.

$$V_1 = ... = V_n = V. \tag{3}$$

The indifference condition has two implications. First, we have $U_i = \beta V$ for all i, which means that in equilirium drivers have no strict incentive to relocate and search at another location. Second, in equilibrium drivers would not turn down a match and go empty in search of a better opportunity. To see why, note that if a driver idles, he earns $U_i = \beta V$, whereas if he accepts a match, then he earns $p_i c_i \delta_{i,j} + \beta V$. Clearly the second expression is larger than the first; hence no driver idles voluntarily.

Our notion of the steady is characterized by (i) the stationarity of the distribution of people and cars across locations, captured by (1), and (ii) drivers' indifference across locations, captured by (3). Of course, temporary or even cyclical imbalances in the flow of traffic due to, say, rush hours, bad weather, football games etc. may actually violate these conditions. We ignore such fluctuations and take a rather long term view of the market, and posit that there cannot be a persistent difference between incoming and outgoing traffic at any location; thus (i) must hold. Similarly, there cannot be a persistent difference across locations in terms of expected profit; thus (ii) most hold. This notion of equilibrium is common in the literature, e.g. Lagos (2000), as it yields analytically tractable results. We show that such an equilibrium always exists and it is generally unique under operating models 2 and 3 (below we provide the descriptions of the operating models). Under model 1 we can prove existence in the interior case, but not in the corner case, though in the numerical simulations the equilibrium still materializes. Conditions (1) and (3), therefore, are necessary but not sufficient for the existence of a steady state equilibrium.

Simplifying V_i, we have

$$V(1-\beta) = \eta_i p_i c_i d_i \text{ for all } i, \tag{4}$$

where

$$d_i = \sum_{j=1}^{n} a_{i,j} \delta_{i,j} \tag{5}$$

is the average trip length of a ride originating from i. We label the locations from 1 to n in such a way that

$$d_1 < d_2 < ... < d_n$$

i.e. location 1 has the shortest expected trip length, whereas location n has the longest. (With no loss of generality we ignore equalities.) Note that d_i depends on both on $\delta_{i,j}$ and $a_{i,j}$; thus a location does not have to be physically the closest to other locations to have the smallest d_i.

Example 1. Consider the city in Figure 1 (left panel) and suppose that the transition probabilities are equal to each other (middle panel), i.e. $a_{i,j} = 1/4$ and $a_{i,i} = 0$. The expected trip lengths are

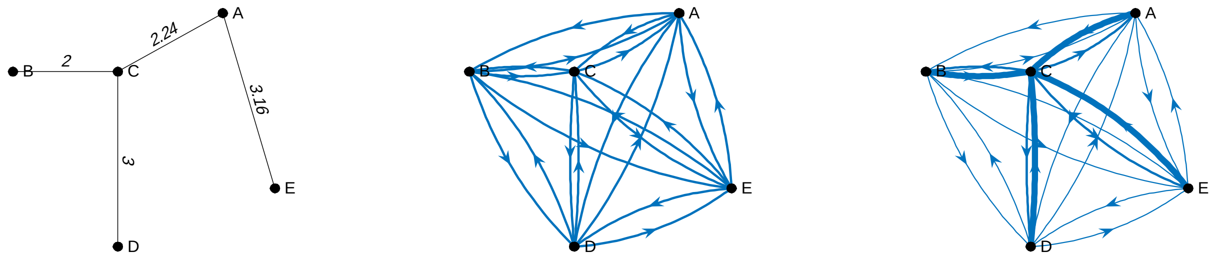

Figure 1: Layout and traffic flows

equal to[4]

$$d_A = 3.72, \quad d_B = 4.66 \quad d_C = 3.16, \quad d_D = 5.41, \quad d_E = 6.09.$$

Location C sits at the intersection of routes, and as expected, it has the shortest trip length; thus, it can be labelled as location 1. A ranks second; so it is 2, B is 3, D is 4 and E is 5. Now suppose C attracts more traffic than other locations, e.g. suppose that 70% of traffic out of any location is directed towards C, while the remaining 30% is shared equally between the other three locations (i.e. $a_{i,C} = 0.7$ and $a_{i,j} = 0.1$, where $i, j \neq C$). As for the traffic out of C, suppose it is still equally shared across the four destinations (i.e. $a_{C,i} = 0.25$). These flows are depicted in the right panel of Figure 1. The expected trip lengths are now equal to

$$d_A = 2.83, \quad d_B = 3.06, \quad d_C = 3.16, \quad d_D = 3.96, \quad d_E = 5.67.$$

Now A and B, despite being physically more remote, have shorter trip lengths than C. This is because when calculating, say, d_B the short distance $\delta_{B,C} = 2$ has a weight of 70% whereas the longer distances, say $\delta_{B,E} = 7.4$, have only 10% each. The imbalance in the traffic flow changes the ranking; so, now location A ought to be labelled as 1, B as 2, C as 3, D as 4 and finally E as 5.

The re-labelling is important for the following reason. From a driver's perspective the location with the minimum d_i is the least desirable location, and in model 1 this location turns into a bottleneck if there are not sufficiently many cars in the city. The bottleneck location is typically the most central one—the one with a short physical distance to every other location—however, as the example illustrates, this is not always the case. Moreover, the location of the bottleneck may change as the flow of traffic changes. Throughout the paper we use the numerical labels to refer to locations, but one should be wary that these labels are relative and may change as the transition matrix or the

[4] Consider location A and note that $\delta_{A,B} = 4.24$, $\delta_{A,C} = 2.24$, $\delta_{A,D} = 5.24$, and $\delta_{A,E} = 3.16$. Since $a_{i,j} = 1/4$, the expected length of a trip originating from A is equal to $d_A = (4.24 + 2.24 + 5.24 + 3.16)/4 = 3.72$. Other trip lengths can be calculated similarly.

layout of the city changes (e.g. road closures or new roads creating new links).

Drivers participate only if $V \geq W$, i.e. if their expected earnings is greater than or equal to their outside option W. The platform will not pay more than W, thus

$$\eta_i p_i c_i d_i = W(1-\beta) \equiv w, \tag{6}$$

where w can be thought as the per-period wage. After substituting for η_i the equality becomes

$$m_i p_i c_i d_i = x_i w, \text{ for all } i. \tag{7}$$

Combining (2) and (7) with the fact that $\sum_{i=1}^{n} x_i = \theta$ we obtain

$$x_i = \frac{\sigma_i p_i c_i d_i}{\sum_{j=1}^{n} \sigma_j p_j c_j d_j} \theta, \tag{8}$$

which pins down the number of cars at location i as a function of prices, commission rates, trip lengths and the attractiveness of each location. Drivers prefer locations that are more attractive (high σ_i[5]) and that have longer trip lengths (high d_i). In addition to these exogenous parameters, x_i depends on the price p_i and the commission rate c_i. The platform can encourage drivers to search at location i by raising p_i or c_i. Such decisions are part of the platform's problem, which we study next.

Platform's Problem and the Definition of the Equilibrium. The platform's per-period earnings at location i is equal to

$$\pi_i = \sum_{j=1}^{n} m_i p_i (1-c_i) a_{i,j} \delta_{i,j} = m_i p_i (1-c_i) d_i.$$

The second equality follows from (5). Adding different π_is across all locations yields the city-wide profit, π. We have

$$\pi = \sum_{i=1}^{n} \pi_i = \underbrace{\sum_{i=1}^{n} m_i p_i d_i}_{\text{Revenue}} - \underbrace{\sum_{i=1}^{n} m_i p_i c_i d_i}_{\text{Payout to Drivers}}.$$

Focus on the term relating to the payout to the drivers. Recall that drivers must be indifferent across

[5]In equilibrium y_i (the number of passengers at location i) is proportional to σ_i. Thus x_i tends to be high at locations where there are more customers.

locations, i.e. (7) must hold across all locations. Thus

$$\sum_{i=1}^{n} m_i p_i c_i d_i = \sum_{i=1}^{n} x_i w = \theta w. \tag{9}$$

There are θ drivers operating in the market, and via commissions, they each expect to earn w per period, so the total payout to drivers is equal to θw. Substituting this relationship into π yields

$$\pi = \sum_{i=1}^{n} m_i p_i d_i - \theta w. \tag{10}$$

The substitution eliminates commission rates from the platform's objective function. The platform picks prices to maximize the revenue while the commission rate(s) satisfy ensure drivers' participation and indifference via (7). Note that in the steady state payoffs are time invariant; thus the platform's lifetime profit is simply equal to $\Pi = \pi / (1 - \beta)$.

A steady state equilibrium is a time invariant tuple $\{(p_i, c_i, x_i, y_i)\}_{i=1}^{n}$ such that (i) the platform maximizes its lifetime profit; (ii) drivers participate and they are indifferent across locations; (iii) the inflow of moves equals to the outflow at each location; (iv) the total measures of cabs and passengers are equal to θ and 1, respectively. Within this setup, we will analyze three different operating models:

- Model 1 - Flexible Prices, Fixed Commission Rate: p_i is location specific, but $c_i = c$ for all i.
- Model 2 - Fixed Price, Flexible Commission Rates: $p_i = p$ for all i but c_i is location specific
- Model 3 - Flexible Prices, Flexible Commission Rates: Both p_i and c_i are location specific.

3 Operating Models

3.1 Model 1: Flexible Prices, Fixed Commission Rate

We start with the claim that in equilibrium no location exhibits excess demand.

Lemma 2 *There cannot be an equilibrium in which $x_i < r_i$ at any location.*

The proof is in the appendix; here we provide a brief sketch of the proof. Fix some $\mathbf{p}$ and let $\mathcal{S}_0$ denote the set of locations in which demand is greater than or equal to supply with at least one location exhibiting excess demand, and $\mathcal{S}_1$ the set of locations with excess supply. First suppose that $\mathcal{S}_1$ is non-empty. If the platform increases the price(s) in $\mathcal{S}_0$ infinitesimally and it leaves the prices in

$\mathcal{S}_1$ intact, then some of the drivers in $\mathcal{S}_1$ will move towards $\mathcal{S}_0$, creating more rides at (slightly) higher prices. Furthermore, if the price rise is kept at a minimum, then despite losing drivers, none of the locations in $\mathcal{S}_1$ will fall into excess demand, which means that the platform can still create the same number of rides at the same prices in $\mathcal{S}_1$. The intervention is profitable because the platform earns more in $\mathcal{S}_0$ while maintaining the same profits in $\mathcal{S}_1$. Now, suppose $\mathcal{S}_1$ is empty. This means that all cars in the city are being used in a ride and yet there is still excess demand at some location(s). But this implies that prices are too low; and, again there is a profitable intervention. Combining both arguments, we see that an equilibrium fails to exist if $x_i < r_i$.

Per Lemma 2, $r_i \leq x_i \Leftrightarrow m_i = r_i$ for all i. Since $m_i = \sigma_i M$ and $\sum_{i=1}^n y_i = 1$, we have

$$M = \frac{1}{h(\mathbf{p})} \quad \text{and} \quad m_i = r_i = \frac{\sigma_i}{h(\mathbf{p})},$$

where

$$h(\mathbf{p}) = \sum_{i=1}^n \frac{\sigma_i}{1-p_i}, \qquad g(\mathbf{p}) = \sum_{i=1}^n \sigma_i p_i d_i \quad \text{and} \quad \Omega(\mathbf{p}) = \frac{g(\mathbf{p})}{h(\mathbf{p})}. \tag{11}$$

Substituting for m_i, we have $\pi = \Omega(\mathbf{p}) - \theta w$. The platform solves

$$\max_{\mathbf{p}} \ \Omega(\mathbf{p}) - \theta w \ \ \text{s.t.} \ \ r_i \leq x_i \text{ for all } i.$$

Note that the constraints $r_i \leq x_i$ follow from Lemma 2 and they require the number of passengers to be less than or equal to the number of drivers at each location.

Lemma 3 *The objective function $\Omega(\mathbf{p})$ is strictly concave in $\mathbf{p}$.*

The proof is in the appendix. Now focus on the constraint $r_i \leq x_i$. Substituting for r_i and x_i we have

$$r_i \leq x_i \Leftrightarrow \Omega(\mathbf{p}) \leq p_i d_i \theta.$$

The Lagrangian of the platform's problem can be written as

$$\mathcal{L} = \Omega(\mathbf{p}) - \theta w + \sum \lambda_i \left[p_i d_i \theta - \Omega(\mathbf{p})\right],$$

where λ_i is the multiplier associated with the constraint at location i. The first order condition with respect to p_i is given by

$$\left[1 - \sum_{i=1}^n \lambda_i\right] \Omega_i'(\mathbf{p}) + \lambda_i d_i \theta = 0, \tag{12}$$

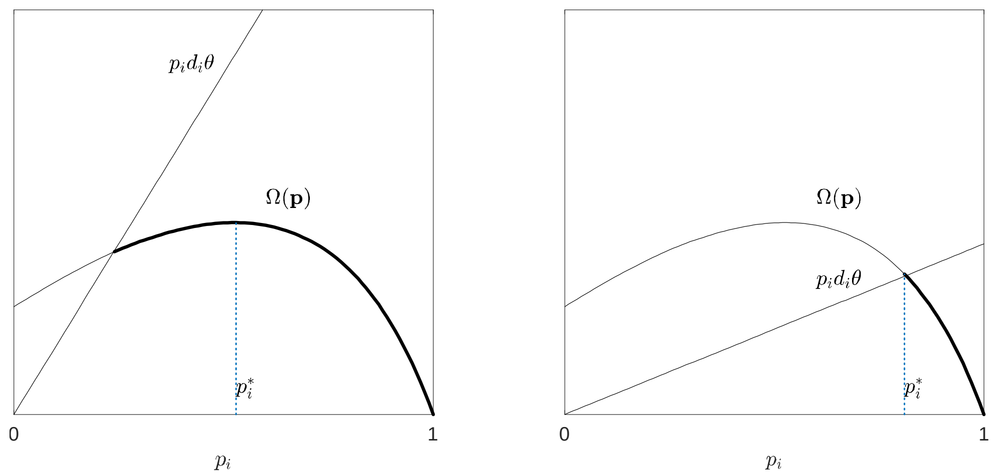

Figure 2: Left: nonbinding constraint (high θ), Right: binding constraint (low θ)

where

$$\Omega_i'(\mathbf{p}) = \frac{\sigma_i}{h(\mathbf{p})}\left[d_i - \frac{\Omega(\mathbf{p})}{(1-p_i)^2}\right].$$

Figure 2 illustrates two scenarios: one in which the constraint $p_i d_i \theta \geq \Omega$ associated with location i is slack (left panel) and the other in which the constraint is binding (right panel). The feasible portion of Ω satisfying $p_i d_i \theta \geq \Omega$ is highlighted in bold. Whether or not the constraint binds depends on how large θ is, i.e. how many cars there are in a city. If θ is high then the constraint is slack, and the platform can pick the interior price p_i^* satisfying the first order condition. If θ is low, then the constraint binds, pushing the platform to pick the corner price that equates demand and supply at that location ("price intervention"). In what follows we prove that the constraints bind in an orderly fashion, starting at the location with the shortest trip length (location 1), then at the location with the second shortest trip length (location 2) and so on.

Lemma 4 *If $\lambda_k = 0$ then $\lambda_{k+1} = 0$. Similarly if $\lambda_{k+1} > 0$ then $\lambda_k > 0$.*

Letting $k = 0, 1, ..$ we refer to *regime-k* as the outcome in which the first k constraints are active. Prices in regime-k satisfy

$$p_i^k d_i \theta = \Omega(\mathbf{p}^k) \text{ for } i = 1, ..., k \quad \text{and} \quad (1 - p_i^k)^2 d_i = \Omega(\mathbf{p}^k) \text{ for } i = k+1, ..., n \tag{13}$$

The first set of equations follow from the fact that the constraints bind at locations 1 through k, whereas the second set of equations are due to the first order conditions at locations $k+1$ through n. The (unconstrained) case, $k = 0$, can be solved analytically, which we report in the following proposition. For $k \geq 1$ we need numerical simulations.

Proposition 1 *If $\theta > \bar{\theta}_1$, where $\bar{\theta}_1$ is given by (19), then the platform sets*

$$p_i^{interior} = 1 - \frac{\mathbb{E}_\sigma(d)}{2\sqrt{d_i}\mathbb{E}_\sigma(\sqrt{d})} \quad \text{and} \quad c^{interior} = \frac{4\theta w \mathbb{E}_\sigma^2(\sqrt{d})}{\mathbb{E}_\sigma^2(d)},$$

where

$$\mathbb{E}_\sigma(d) = \sum_{i=1}^{n} \sigma_i d_i \qquad \text{and} \quad \mathbb{E}_\sigma(\sqrt{d}) = \sum_{i=1}^{n} \sigma_i \sqrt{d_i}. \tag{14}$$

The interior equilibrium emerges if there are sufficiently many cars in the city ($\theta > \bar{\theta}_1$). The abundance of cars allows the platform to set interior prices without worrying about incentivizing drivers towards undesirable locations. In such an equilibrium, prices satisfy $p_i^{\text{interior}} < p_{i+1}^{\text{interior}}$, i.e. the platform sets higher prices at locations with longer trip lengths. This relationship can also be seen in Figure 3 (left panel), where prices in the interior region $\theta > \bar{\theta}_1$ satisfy $p_1 < ... < p_5$.[6] The platform faces a standard trade-off between extensive and intensive margin effects. On the extensive margin, it generates more matches by lowering prices, whereas on the intensive margin it raises more money from each ride by increasing prices. The intensive margin effect is stronger at locations with longer trip lengths: raising the price may cause a drop in the local demand, but if the location has a high d_i then the long trip length more than covers this loss. The opposite is true at locations with a low d_i. Taking these considerations into account, the platform sets prices satisfying the above relationship.

Drivers prefer higher prices and longer trip lengths, so one might ask how they can be indifferent across locations if in equilibrium p_i and d_i move in the same direction. The answer is the probability of finding a customer. Location n has the longest expected trip length and the highest price; so, it attracts more cars (per customer) than any other location. As a result, it has the lowest customer-to-cab ratio, η_i, making it the most difficult location for a driver to find a customer. In contrast, location 1 has the highest η_i, which makes up for the short trip length and the low price. In general, the customer-to-cab ratios satisfy $\eta_n < ... < \eta_1$. Notice, however, since θ is sufficiently large, even at location 1 there is excess supply, i.e. $\eta_1 < 1$.

If θ falls below $\bar{\theta}_1$ (0.74 in Figure 3) then the interior demand cannot be sustained with the available cars in the city and the constraints $r_i \leq x_i$ start to bind, starting at location 1. In response,

[6] In the simulations, for the purpose of exposition, we consider a city with only five locations. The layout of the city and the transition matrix are as in Figure 1, right panel. In Section 4.5 we calibrate the model for New York City and Los Angeles using real world ride patterns from Uber. The results from both sets of simulations are qualitatively very similar.

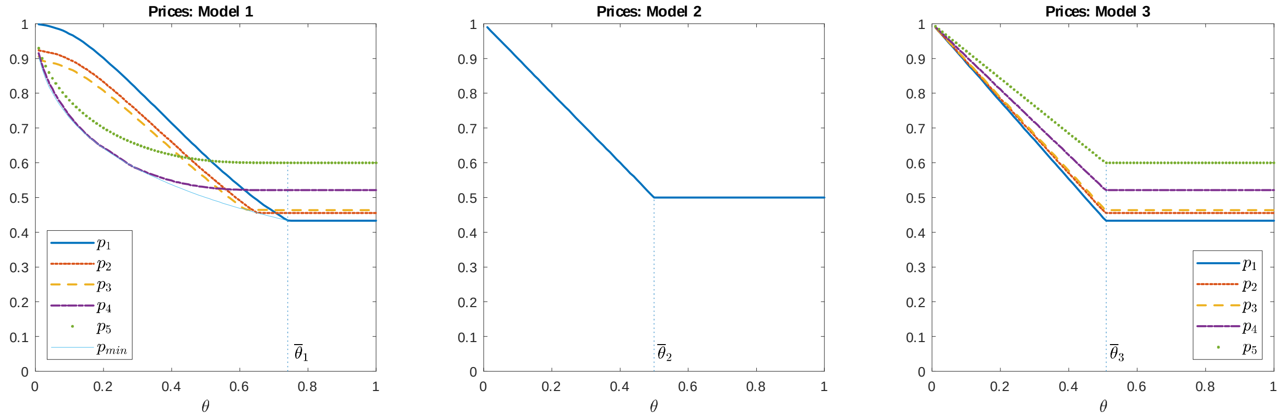

Figure 3: Equilibrium Prices

the platform increases p_1, which diminishes the local demand r_1 and increases the local supply x_1 (by encouraging more drivers towards that location).[7] The price intervention matches the demand and supply, so $\eta_1 = 1$ at location 1, but at every other location η_i is still less than 1.[8] In other words, the insufficient number of cars in the city forces the platform to conduct a price intervention at location 1—a diversion from the interior equilibrium—yet there is still excess supply at other locations. If the platform could somehow divert drivers from excess supply locations towards location 1 without disturbing the interior price there, then it could earn more. However, within the confines of the current model (single commission rate) it cannot do that. Location 1, from the platform's perspective, turns into a *bottleneck*.

As θ falls further down, the constraints $r_i \leq x_i$ start to bind at more locations, forcing the platform to intervene and raise the price at those locations: if θ falls below 0.65 then p_2 starts to rise, and if it falls below 0.62 then p_3 starts to rise to match the demand and supply at those locations. Notice that even though the model cannot avoid bottlenecks, it still deals with them locally. The platform uses location specific prices to intervene, and as a result, it does not disturb the demand at other locations too much. For instance, in Figure 3 when θ falls below $\bar{\theta}_1$ the price at location 1 surges up, but prices at remaining (unconstrained) locations stay rather unchanged.

The above observations seem to resonate with the surge pricing strategy employed by Uber. The surge pricing scheme kicks in when the number of passengers asking for a ride at a location exceeds the number of available drivers at that location, which in our model is equivalent to the constraint $r_i \leq x_i$ becoming active. Uber executives defend the surge pricing practice saying it serves their

[7] It is straightforward to verify that $r_1 = \frac{\sigma_1}{h(\mathbf{p})}$ decreases while $x_1 = \frac{\sigma_1 p_1 d_1 \theta}{g(\mathbf{p})}$ increases in p_1.

[8] Per Lemma 4 the constraints bind in an orderly fashion. If θ falls below $\bar{\theta}_1$ then $\eta_1 = 1$ and $\eta_i < 1$ for $i \geq 2$. If it falls further down, then $\eta_1 = \eta_2 = 1$ and $\eta_i < 1$ for $i \geq 3$, and so on.

goal of "relentless reliability to manage the marketplace math so that supply and demand match as perfectly as possible in the face of ever-shifting, highly unpredictable circumstances"(Wohlsen, 2013). Our results seem to confirm a similar insight. Indeed when p_1 surges up at location 1, the local demand r_1 goes down while the local supply x_1 goes up. Furthermore, Uber's practice is location specific, i.e. while the price may surge at an excess demand location, it remains unchanged at other locations. This outcome is also similar to what we observe in our simulations.

In highlighting our model's insights that may be relevant to real practice such as surge pricing, we should note the following caveat. Our model is based on a steady state setting; as such Figure 3 depicts steady state equilibrium prices associated with different values of θ. Uber's surge pricing practice, on the other hand, appears to be a temporary and transitional solution; thus, it may not be directly comparable to our steady state results.

Proposition 2 *Prices in regime-k are bounded below by* $p_{\min} = 1 + \theta/2 - \sqrt{\theta^2/4 + \theta}$ *and satisfy* $p_1^k > ... > p_k^k > p_{\min}$ *and* $p_{\min} < p_{k+1}^k < ... < p_n^k$.

Even though we are unable to analytically characterize prices in regime-k when $k \geq 1$, we can still pin down their lower bound $p_{\min}$. It is easy to verify that as θ drops $p_{\min}$ rises, which indicates that the platform responds to a decreasing number of cars by increasing prices. The second part of the proposition establishes that at locations $1, ..., k$ prices are inversely related to the trip length, which is the opposite of what we have seen in Proposition 1. The reason is that at locations $1, ..., k$ the constraints $r_i \leq x_i$ bind, thus drivers' indifference condition boils down to $p_i d_i = p_j d_j$. If d_i is less than d_j then p_i must exceed p_j; else drivers cannot be indifferent.

3.2 Model 2: Fixed Price, Flexible Commission Rates

We start by arguing that Lemma 2 is still valid, i.e. there cannot be an equilibrium in which a location exhibits excess demand. Save for some minor differences (instead of prices, the platform uses commission rates to incentivize drivers towards excess demand locations) the proof remains the same; thus we only provide a sketch of the proof in here. To start, suppose that $\mathcal{S}_1$ is non-empty. If the platform leaves the price p as well as the commission rates in $\mathcal{S}_0$ intact, but reduces the commission rates in $\mathcal{S}_1$, then some drivers in $\mathcal{S}_1$ would flow towards locations in $\mathcal{S}_0$, creating more rides there.[9] If the reduction is infinitesimally small (call it ε_i), then despite losing drivers to $\mathcal{S}_0$, none of the locations in $\mathcal{S}_1$ would fall into excess demand, and the remaining drivers would still be

[9] The number of drivers at location i, $x_i = \frac{\sigma_i c_i d_i \theta}{\sum \sigma_j c_j d_j}$, increases in c_i and decreases in c_js.

able to serve the initial demand.[10] Overall, the platform would not lose any profits in $\mathcal{S}_1$, yet it would create more rides and more profits in $\mathcal{S}_0$ rendering the intervention profitable. Now suppose $\mathcal{S}_1$ is empty, i.e. at every location we have $x_i \leq r_i$ with at least one inequality strict. Since all cars are being used in a ride, and there is still some excess demand, the platform can earn more by increasing p to the point where demand equals to supply at every location, rendering the conjectured outcome a non-equilibrium. In conclusion, so long as there is excess demand in the city, the platform has a profitable intervention; thus, there cannot be an equilibrium in which $x_i < r_i$ at any i.

The claim establishes that $r_i \leq x_i$, and therefore, $m_i = r_i = y_i (1 - p)$ for all i. In addition, since $m_i = \sigma_i M$ and $\sum_{i=1}^n y_i = 1$ we have

$$m_i = r_i = \sigma_i (1 - p).$$

Substituting this relationship, and $p_i = p$ into (10) yields

$$\pi = (1 - p) p \mathbb{E}_\sigma (d) - \theta w. \tag{15}$$

The platform solves $\max_p \pi$ s.t. $r_i \leq x_i$ for all i. Substituting for r_i and x_i, we have

$$r_i \leq x_i \Leftrightarrow (1 - p) \sum_{i=1}^n \sigma_i c_i d_i \leq c_i d_i \theta.$$

A commission vector $\mathbf{c}$ is *incentive compatible* if it satisfies drivers' indifference across locations, i.e. it satisfies equation (7). After substituting for m_i and x_i this is equivalent to

$$(1 - p) p \sum_{i=1}^n \sigma_i d_i c_i = \theta w. \tag{16}$$

It follows that

$$r_i \leq x_i \Leftrightarrow w \leq p c_i d_i.$$

Recall that in model 1 the constraints became active in an orderly fashion, starting at location 1 and then at location 2 and so on. Here this is no longer the case.

Lemma 5 *Fix p. Suppose there exists an incentive compatible $\mathbf{c}$ under which $r_i < x_i$ for $i \leq k$ and $r_i = x_i$ for $i > k$. Then there exists another incentive compatible $\hat{\mathbf{c}}$ under which $r_i < x_i$ for all i.*

[10] Such ε_i exists because locations in $\mathcal{S}_1$ exhibit excess supply. Since $x_i > r_i \Leftrightarrow c_i p d_i > w$ for all $i \in \mathcal{S}_1$, and since the inequality is strict, there exists $\varepsilon_i > 0$ satisfying $(c_i - \varepsilon_i) p d_i > w$.

The proof is in the appendix. The idea is that we can generate a new $\hat{\mathbf{c}}$ by marginally shaving off the rates of $\mathbf{c}$ at locations where the constraint is slack (but without rendering any of these constraints binding) and marginally increasing the rates at locations where the constraint is binding. Thus, by construction all constraints will be slack under the new $\hat{\mathbf{c}}$.

The Lemma rules out the possibility that $r_i < x_i$ for some locations and $r_i = x_i$ at other locations. Either the constraints are slack at all locations or they bind at all locations. We can now characterize the equilibrium.

Proposition 3 *If $\theta > \bar{\theta}_2 = \frac{1}{2}$ then all locations exhibit excess supply, i.e. $r_i < x_i$ for all i. The platform sets $p^{interior} = \frac{1}{2}$, however the commission rates are indeterminate: there exists a continuum of $\mathbf{c}$ satisfying incentive compatibility while ensuring $r_i < x_i$ for all i. If, however, $\theta \leq \bar{\theta}_2$, then $r_i = x_i$ for all i, i.e. no cab idles at any location, and the platform sets*

$$p^{corner} = 1 - \theta \quad and \quad c_i^{corner} = \frac{w}{(1-\theta)\, d_i}.$$

If θ is sufficiently large, then the platform can experiment with a wide range of commission schemes and still keep the customer-to-cab ratio η_i below 1 at all locations. Below we present an example for such an outcome.

Example 2. Consider the city in Figure 1 and let $\theta = 1$ and $w = 0.6$. Furthermore recall that $d_5 > d_4 > ... > d_1$, i.e. location 5 has the longest expected travel distance, location 4 has the second longest, and so on. Since $\theta > \bar{\theta}_2$ we deal with the interior equilibrium, which exhibits a continuum of commission rates, including $\mathbf{c}$ and $\hat{\mathbf{c}}$ below.

Location	c_i	η_i	$\hat{c}_i$	$\hat{\eta}_i$
1	0.85	0.5	0.67	0.63
2	0.78	0.5	0.67	0.59
3	0.76	0.5	0.67	0.57
4	0.61	0.5	0.67	0.45
5	0.42	0.5	0.67	0.32

Table 1 – Equilibrium Commissions

It is straightforward to check that $\mathbf{c}$ and $\hat{\mathbf{c}}$ are incentive compatible (they satisfy 16) and that the constraints $w \leq pc_i d_i$ are all slack under both commission vectors; hence they both can be posted in equilibrium.[11] With $\mathbf{c}$ the rates satisfy $c_5 < ... < c_1$, i.e. the longer a location's expected travel distance d_i, the lower its commission rate. Drivers want both c_i and d_i to be large, so under $\mathbf{c}$ no location is more popular for the drivers than the other (note that the customer-to-cab ratio η_i is the same across locations), because the commission rates counterbalance the travel distances. In contrast, $\hat{\mathbf{c}}$ is a fixed-commission system where $\hat{c}_1 = ... = \hat{c}_5$. As a result, location 5 becomes the most popular among the drivers—it has the smallest $\hat{\eta}_i$—location 4 becomes the second most popular and so on.

The selection of commissions is not confined to $\mathbf{c}$ and $\hat{\mathbf{c}}$; if $\theta > \bar{\theta}_2$ then the platform can pick a continuum of other rates. But as θ gets closer to the threshold $\bar{\theta}_2$ the platform must increase the rates at locations with a small d_i and decrease the rates at locations with a large d_i to counterbalance the distance effect.

If θ falls below $\bar{\theta}_2$, then the platform can no longer avoid the constraints, so η_i hits 1 at all locations and the platform operates at a 100% utilization rate everywhere in the city. Since the number of cars in the city is not sufficient to address the interior demand, the price, inevitably, starts to rise (Figure 3, the middle panel). The equilibrium price materializes at the point where the aggregate demand can be addressed with the number of available cars.

Model 2 gives the platform the ability to avoid bottlenecks. Thanks to the flexible commission structure, the platform does not resort to a price intervention until the passenger-to-cab ratio is equal to 100% at *every* location. Up to that point, by adjusting the commission rates—increasing them at less desirable locations, decreasing them at more desirable locations, or a combination—the platform manages to spread the cars evenly and serve the (unconstrained) demand associated with the interior solution. Thus, in contrast to the previous model, no location turns into a bottleneck.

3.3 Model 3: Flexible Prices, Flexible Commission Rates

Finally we turn to the most flexible scheme. As before, there cannot be an equilibrium in which $x_i < r_i$ at any i. Since $r_i \leq x_i$ we can write $m_i = r_i = y_i (1 - p_i)$ and since $m_i = \sigma_i M$, and $\sum_{i=1}^{n} y_i = 1$, we have

$$m_i = r_i = \frac{\sigma_i}{h(\mathbf{p})}.$$

[11] Recall that in our benchmark (the right hand side panel of Figure 1) the trip lengths are $d_1 = 2.83$, $d_2 = 3.06$, $d_3 = 3.16$, $d_4 = 3.96$, $d_5 = 5.67$. The steady state vector is equal to $\boldsymbol{\sigma} = (0.15, 0.15, 0.4, 0.15, 0.15)$. Using these numbers one can verify that (i) $(1-p)\, p \sum_{i=1}^{n} \sigma_i d_i c_i = \theta w$ and (ii) $w < pc_i d_i$ under both $\mathbf{c}$ and $\hat{\mathbf{c}}$.

Substituting for m_i, the platform's profit is equal to

$$\pi = \Omega(\mathbf{p}) - \theta w,$$

where $\Omega(p)$ is given by (11). The constraint $r_i \leq x_i$, after substituting for r_i and x_i, can be written as

$$r_i \leq x_i \Leftrightarrow \frac{\sum_{i=1}^{n} \sigma_i p_i c_i d_i}{h(\mathbf{p})} \leq p_i c_i d_i \theta \text{ for all } i. \tag{17}$$

The commission vector $\mathbf{c}$ must be incentive compatible, i.e. it must satisfy (7). Substituting for m_i and x_i, the condition becomes

$$\frac{\sum_{i=1}^{n} \sigma_i p_i d_i c_i}{h(\mathbf{p})} = \theta w. \tag{18}$$

Combining (17) and (18), the constraint associated with location i can be rewritten as

$$r_i \leq x_i \Leftrightarrow w \leq p_i c_i d_i \text{ for all } i.$$

Since the commission rates are flexible, either the constraints are slack at all locations (i.e. $p_i c_i d_i > w$ for all i) or they bind at all locations (i.e. $p_i c_i d_i = w$ for all i). In other words, there cannot be a scenario where $p_i c_i d_i > w$ for some locations and $p_i c_i d_i = w$ at other locations. When faced with such an outcome, the platform can simply shave off the the commission rates at excess supply locations and increase the rates at constrained locations to slacken those constraints. The proof of this claim is practically the same as the proof of Lemma 5; thus it is skipped in here.

Proposition 4 *If $\theta > \bar{\theta}_3 \equiv \frac{\mathbb{E}_\sigma(d)}{2\mathbb{E}_\sigma^2(\sqrt{d})}$ then all locations exhibit excess supply, i.e. $r_i < x_i$ for all i. The platform sets*

$$p_i^{interior} = 1 - \frac{\mathbb{E}_\sigma(d)}{2\sqrt{d_i}\mathbb{E}_\sigma(\sqrt{d})},$$

but the interior commission rates are indeterminate. If, however, $\theta \leq \bar{\theta}_3$, then no location exhibits excess supply. Along such an outcome the platform sets

$$p_i^{corner} = 1 - \frac{\mathbb{E}_\sigma(\sqrt{d})\theta}{\sqrt{d_i}} \quad \text{and} \quad c_i^{corner} = \frac{w}{d_i - \sqrt{d_i}\mathbb{E}_\sigma(\sqrt{d})\theta}.$$

With sufficiently many cars in the city ($\theta > \bar{\theta}_3$) the equilibrium is interior and no constraint $r_i \leq x_i$ is active. Thanks to the flexible nature of the commissions, the platform avoids the constraints until $\theta = \bar{\theta}_3$. Up to that point, by fine tuning the location specific rates, the platform incentivizes the drivers to spread themselves across the city in an even way, and thereby, it avoids bottlenecks. If, however, θ falls below $\bar{\theta}_3$, then the interior demand cannot be addressed with the number of available cars, so prices start to rise (Figure 3, right panel).

Equilibrium prices satisfy $p_i < p_{i+1}$, i.e. the platform sets higher prices at locations with a high d_i. This relationship is similar to what we saw in model 1 and it has the same intuition: at locations with a high d_i the intensive margin effect (raising the price to generate more revenue) is more dominant when compared to the extensive margin effect (lowering the price to create more matches). Raising the price may cause a drop in demand, but if the location has a high d_i then the trip length more than covers the loss in demand, which induces the platform to set marginally higher prices at such locations. In model 1 this relationship broke down in the corner region because in that model prices at bottleneck locations had to surge up to incentivize drivers. In here commission rates are employed for this purpose; thus the relationship remains valid both in the interior and the corner equilibria.

Note that the interior commission rates are indeterminate. If there are sufficiently many cars in the city then there exists a continuum of payoff-equivalent equilibria in which the platform may select from alternative commission structures. A corollary is that, if θ is sufficiently large then the platform can implement a fixed commission system ($c_i = c$ for all i) without loss of optimality.

Model 3 will serve as a benchmark in our comparisons as the other models are special cases of Model 3.

4 Comparisons

We now compare the schemes in terms of the number of matches, the amount of profits, and the amount of consumer surplus they generate. The comparisons, inevitably, depend on whether equilibria are interior or corner, so we start with the relevant $\bar{\theta}_i$s.

Proposition 5 *The thresholds satisfy* $\bar{\theta}_1 > \bar{\theta}_3 > \bar{\theta}_2$.

To see why $\bar{\theta}_1 > \bar{\theta}_3$, note that in model 3, the parameter θ does not become a constraint until $\eta_i = 1$ at every location. In contrast, in model 1 the constraint becomes active as soon as $\eta_1 = 1$ at location 1 while there is still excess supply at other locations. This explains why model 3 needs fewer cars to maintain its unconstrained outcome. As for $\bar{\theta}_3 > \bar{\theta}_2$, below we show that model 3 generates more matches than model 2 in its interior equilibrium. Naturally, it needs more cars in order to sustain those matches.

When comparing the performance of different models, one has to pay attention to these thresholds. For instance, model 1 allows the platform to tailor prices according to location specific factors; so, one would expect it to generate more profits than the single-price model (model 2). Notice, however, model 1 requires a large θ to sustain its interior equilibrium whereas model 2 does not. If θ is less than $\bar{\theta}_1$ but more than $\bar{\theta}_2$ then the comparison involves the corner outcome of model 1 and the interior outcome of model 2, which is not clear.

4.1 Number of Matches

Under models 1 and 3 we have $m_i = \sigma_i / h(\mathbf{p})$, and therefore $M = 1/h(\mathbf{p})$. Similarly, under model 2, $m_i = \sigma_i(1-p)$, thus $M = (1-p)$. After substituting for equilibrium prices we have

$$M_1 = \begin{cases} \frac{\mathbb{E}_\sigma(d)}{2\mathbb{E}_\sigma^2(\sqrt{d})} & \text{if } \theta > \bar{\theta}_1 \\ \text{simulations} & \text{if } \theta \le \bar{\theta}_1 \end{cases} \qquad M_2 = \begin{cases} 1/2 & \text{if } \theta > \bar{\theta}_2 \\ \theta & \text{if } \theta \le \bar{\theta}_2 \end{cases} \qquad M_3 = \begin{cases} \frac{\mathbb{E}_\sigma(d)}{2\mathbb{E}_\sigma^2(\sqrt{d})} & \text{if } \theta > \bar{\theta}_3 \\ \theta & \text{if } \theta \le \bar{\theta}_3 \end{cases}$$

Proposition 6 *In terms of generating matches model 3 outperforms model 2: $M_3 > M_2$ if $\theta > \theta_2$ and $M_3 = M_2$ if $\theta \le \theta_2$. Model 1's performance varies: if $\theta > \bar{\theta}_1$, then it creates as many matches as model 3 does, but if $\theta \le \bar{\theta}_2$, then it creates fewer matches than model 2 does.*

In the interior equilibrium of a model $(\theta > \bar{\theta}_i)$ the number of cars does not enter into the platform's problem as a constraint; thus, the resulting M_i is independent of θ. This is why M_i associated with model i is flat when $\theta > \bar{\theta}_i$ in Figure 4. The parameter θ filters into M_i only when there is an insufficient number of cars in the city.

Models 1 and 3 produce the same number of matches, and indeed the same prices and profits, if θ is large enough for both models to exhibit their interior outcomes $(\theta > \bar{\theta}_1)$. This is because both models are based on flexible-pricing, as such, the platform maximizes the same objective function. The models differ in terms of their flexibility in commissions, but if θ is large enough this difference becomes irrelevant because with sufficiently many cars at its disposal, the platform can satisfy those constraints with either commission format.

As θ starts to fall, however, the number of cars becomes relevant and M_i depends mainly on the platform's ability to spread a limited number of cars evenly without running into bottlenecks. With few cars in the city (generally if $\theta < \bar{\theta}_1$) we see that models 2 and 3 generate more matches than model 1. This is because the utilization rate–equilibrium passenger-to-cab ratio–in those models

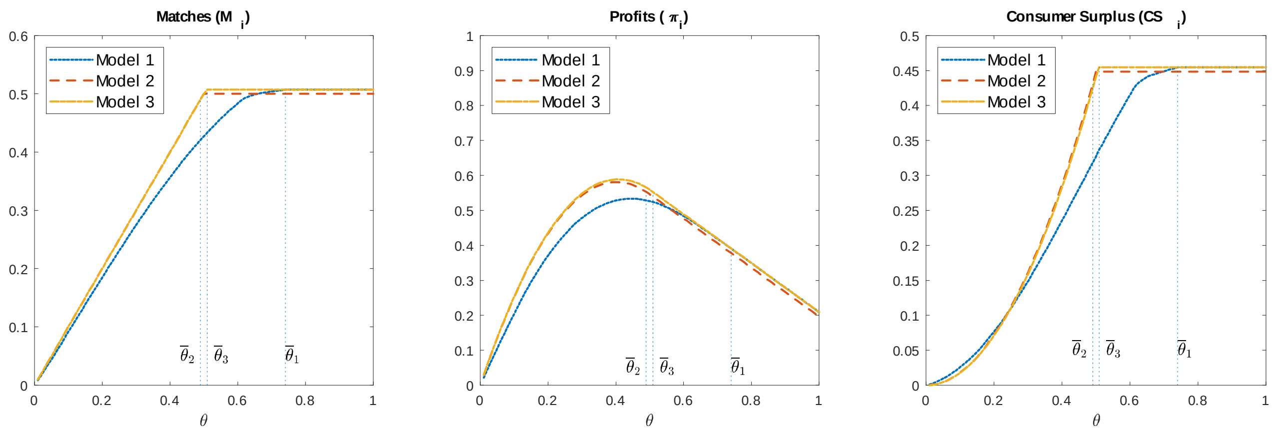

Figure 4: Number of Matches (left), Profits (middle), Consumer Surplus (right)

hits 100% at every location. In contrast, the utilization rate in model 1 is equal to 100% only at locations with a small d_i while remaining locations exhibit excess supply. This, of course, is a waste of resources, causing the platform to create fewer matches than it potentially could.

4.2 Profits

In models 1 and 3 we have $\pi = \Omega(\mathbf{p}) - \theta w$. Substituting for equilibrium prices we have

$$\pi_1 = \begin{cases} \frac{\mathbb{E}_\sigma^2(d)}{4\mathbb{E}_\sigma^2(\sqrt{d})} - \theta w & \text{if } \theta > \bar{\theta}_1 \\ \text{simulations} & \text{if } \theta \leq \bar{\theta}_1 \end{cases}, \text{ and } \pi_3 = \begin{cases} \frac{\mathbb{E}_\sigma^2(d)}{4\mathbb{E}_\sigma^2(\sqrt{d})} - \theta w & \text{if } \theta > \bar{\theta}_3 \\ \theta[\mathbb{E}_\sigma(d) - \theta \mathbb{E}_\sigma^2(\sqrt{d})] - \theta w & \text{if } \theta \leq \bar{\theta}_3 \end{cases}$$

Going through similar steps, π_2 can be written as

$$\pi_2 = \begin{cases} \frac{\mathbb{E}_\sigma(d)}{4} - \theta w & \text{if } \theta > \bar{\theta}_2 \\ \theta(1-\theta)\mathbb{E}_\sigma(d) - \theta w & \text{if } \theta \leq \bar{\theta}_2 \end{cases}.$$

Figure 4 (middle panel) shows that π_i has a concave trajectory, peaking around $\bar{\theta}_i$. In the region $\theta > \bar{\theta}_i$ the number of cars does not enter into the platform's problem as a constraint; thus neither m_i nor p_i, and therefore, nor the platform's revenue depend on θ. The cost θw, however, increases in θ. Since the revenue is flat while the cost increases, the profit falls when $\theta > \bar{\theta}_i$. From the platform's point of view having a θ more than $\bar{\theta}_i$ is a waste of resources, because the additional drivers do not generate any new matches (they constitute excess supply), yet they ought to be paid w each, which is a drain on the profits.

In the region $\theta \leq \bar{\theta}_i$ the constraint(s) become active. As θ decreases, fewer matches are created and the revenue falls. The cost component θw, too, falls with θ, so the direction of change in the profit depends on the strength of these movements. The simulation suggests that when θ is sufficiently close to $\bar{\theta}_i$, the fall in revenue is only minimal, thus the profit initially increases and reaches its maximum slightly below $\bar{\theta}_i$. This suggests that when θ is endogenized, the platform would want to create a small degree of shortage in the number of cabs.

Proposition 7 *In terms of profits model 3 strictly outperforms model 2, i.e. $\pi_3 > \pi_2$ for any θ. Model 1's performance varies: if $\theta > \bar{\theta}_1$, then it generates as much profit as model 3 does, thus outperforms model 2, but if $\theta \leq \theta_2$, then it performs worse than model 2.*

The relative performance of models depends on how large θ is. If $\theta > \bar{\theta}_1$ then, as noted above, the distinguishing feature for a model is whether its prices are flexible, because the constraints $r_i \leq x_i$ can be satisfied with either commission structure. The proposition establishes that if θ is large then flexible price models (1 and 3) generate more profits than the single price model (2).

However, with an insufficient θ, the platform's ability to avoid bottlenecks and spread the cars around evenly becomes a more important feature than fine tuning the prices. Indeed, when $\theta < \bar{\theta}_2$ we see that $\pi_3 \approx \pi_2 > \pi_1$. With a limited number of cars, flexible commission models (2 and 3) achieve a 100% passenger-to-cab ratio at every location, whereas in the single commission model (1) this ratio is equal to 100% only at some location(s) while other locations exhibit excess supply. It is no surprise that with such a waste of precious resources the latter model is not as profitable.

4.3 Consumer Surplus

The consumer surplus at location i is given by

$$cs_i = \sum_{j=1}^{n} y_i a_{i,j} \delta_{i,j} \int_{p_i}^{1} (v - p_i)\, dF(v),$$

where $F(v)$ is the cdf governing a passenger's willingness to pay and y_i is the mass of passengers at location i. Noting that $\sum_{j=1}^{n} a_{i,j}\delta_{i,j} = d_i$ and $F(v) = v$ (uniform distribution) we have

$$cs_i = \frac{1}{2} y_i d_i (1 - p_i)^2 = \frac{1}{2} m_i d_i (1 - p_i).$$

The second equation follows from the fact that in equilibrium $m_i = y_i (1 - p_i)$. The consumer surplus for the entire city is equal to

$$CS = \sum_{i=1}^{n} cs_i = \frac{1}{2} \sum_{i=1}^{n} m_i d_i (1 - p_i).$$

Similar to profits, consumer surplus is linked to the number of matches and, generally, it follows the same pattern as M_i (Figure 4, right panel). For instance, if θ is large, then we see that the flexible price models (1 and 3) generate more matches, and therefore, more consumer surplus than the single-price model (2). As θ falls, the flexibility of the commission structure starts to become more important. In this region (roughly $\theta < \bar{\theta}_1$) flexible commission models (2 and 3) create more matches, and more consumer surplus, than the single commission model (1).

Combining these observations with the ones in the previous section, we note that if θ falls below $\bar{\theta}_1$ then models 2 and 3 create not only more profits but also more consumer surplus than model 1. The observation is somewhat counter-intuitive, as one may think that if profits are high, then consumer surplus should be low and vice versa. This would be the case if the size of the pie (total surplus) was fixed; however in here, models 2 and 3 generate more matches, and therefore a higher total surplus than model 1—or a bigger pie to continue with the analogy—so, they may indeed create more profits *and* more consumer surplus.

4.4 Comparative Statics

Here we explore comparative statics pertaining to the transition matrix T. A change in some $a_{i,j}$ changes the steady state vector, expected trip lengths, and thereby all equilibrium objects. For analytic tractability, suppose that we are interested in some location i and that passengers at any other location j wish to move to i with the same probability $q \in (0, 1)$, that is $a_{j,i} = q$. Furthermore, passengers at j visit remaining locations with identical probabilities, i.e. $a_{j,k} = \frac{1-q}{n-2}$ for $j, k \neq i$. Finally, passengers at i visit other locations also with identical probabilities, i.e. $a_{i,j} = \frac{1}{n-1}$. Figure 5 (right panel) provides an illustration of these flows. With this specification, the steady state distribution can be computed as:

$$\sigma_i = \frac{q}{1+q} \quad \text{and} \quad \sigma_j = \frac{1}{(n-1)(1+q)} \text{ for } j \neq i.$$

A rise in q increases the flow of traffic towards i. In the steady state, such an increase must be met with a drop in the flows towards other locations. This is why an increase in q causes σ_i to rise but

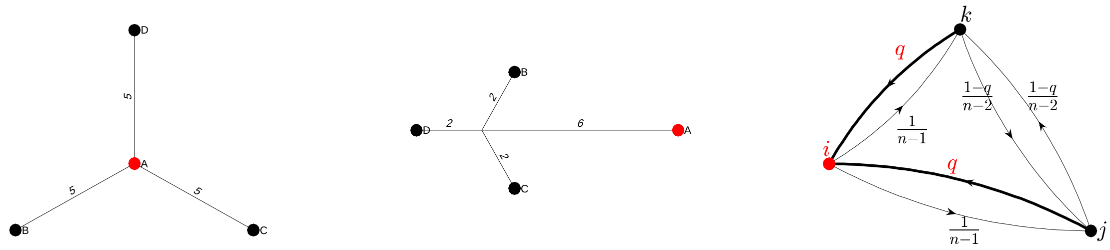

Figure 5: Comparative Statics

σ_js to fall. Recall that a high value of σ_i indicates that the location is attractive as it pulls in and sends away a lot of traffic.

A change in q affects expected trip lengths as well. For tractability suppose that the distance between the location of interest and every other location is equal to δ and the distance between any other location is equal to $\tilde{\delta}$, i.e.

$$\delta_{j,i} = \delta \quad \text{and} \quad \delta_{j,k} = \tilde{\delta} \text{ for } j, k \neq i.$$

Clearly, if $\delta < \tilde{\delta}$ then i is a central location and if $\delta > \tilde{\delta}$ then it is peripheral. Figure 5 provides two examples for such layouts.[12] With this simplification the expected trip lengths can be computed as:

$$d_i = \delta \quad \text{and} \quad d_j = q\delta + (1-q)\,\tilde{\delta}.$$

In what follows we study the impact of q on prices and profits, starting with model 3.

Model 3: Multiple Prices, Multiple Commission Rates

Proposition 8 *Suppose q increases and location i becomes more attractive. If $\delta < \tilde{\delta}$, i.e. if location i is a short distance away from other locations, then p_i increases while the platform's profit decreases. If $\delta > \tilde{\delta}$ then these relationships are reversed.*

In the proof we show that dp_i/dq is inversely related to the rate of change in $\mathbb{E}_\sigma(\sqrt{d})$, which represents the average trip length in the city. If $\delta < \tilde{\delta}$ then a rise in q decreases $\mathbb{E}_\sigma(\sqrt{d})$. Indeed if a short-distance location becomes more attractive, then the average trip length in the entire city falls. In response, the platform increases the price at that location.

[12] We are interested in location A; thus $i = A$. In the left panel the distance between i and any other location is $\delta = 5$ whereas the distance between two other locations is $\tilde{\delta} = 10$. In this layout i is physically central as it is only a short distance away from other locations. In contrast, in the middle panel the distance between i and any other location is $\delta = 8$ whereas the distance between two other locations is $\tilde{\delta} = 4$. In this layout i is peripheral.

Even though the price goes up, the profits fall. There are two reasons for this. First, the rising q increases σ_i but decreases σ_js; thus we end up with more short-distance trips (associated with location i) and fewer long-distance trips (associated with other locations). Second, the increasing q causes d_js to fall, which means that even the previously longer trips get shorter. Profits fall as a result. If $\delta > \tilde{\delta}$, i.e. if location i sits far away from other locations, then these relationships are reversed.

Model 2: Single Price, Multiple Commission Rates

Remark 1 *In model 2 the price is independent of the transition matrix T. In response to a change in the flow of traffic the platform does not update the price; instead it updates the commissions. Customers, crucially, do not observe such adjustments. The direction of change in the platform's profit is same as above: if a location with a relatively small d_i becomes more attractive then the platform earns less (and vice versa).*

In model 2 both the interior price and the corner price are independent of d_i and σ_i. This means that fluctuations in the transition matrix do not affect the equilibrium price at all—the platform charges the same per-mile price irrespective of T. This outcome, incidentally, is not specific to the simplified environment presented in this section—prices are independent of d_i and σ_i under any physical layout and transition matrix. Instead of the price, the platform works with the commission rates to respond to the fluctuations in the flow of traffic. If, for instance, a change in T causes d_i to increase then the platform responds by decreasing c_i at that location, because in equilibrium c_i and d_i are inversely related (Proposition 3).

In this model the platform conducts its fine-tuning via commissions, and crucially these changes are not seen by the customers (they are observed only by the drivers); so from the customers' point of view, the per-mile price not only is the same everywhere, but also is stable. We know that model 3 outperforms model 2 (in terms of profits) due to its additional flexibility in pricing. However, if the transition matrix T changes then so do all prices and commissions in model 3, which means that the platform not only has to set location-specific price and commission rates, but also must be actively updating those objects based on fluctuations in T. So much variation and adjustment in prices might turn off some customers and cause them to turn to alternative means of travel. If one factors in the presence of such (behavioral) customers then the profits under model 3 would diminish as the customer base would shrink. The simulations suggest that the difference between π_2 and π_3 is modest. Taking the above consideration into account, model 2 might, in fact, outperform model 3 if a non-negligible portion of the customers were to behave as above.

Model 1: Flexible Prices, Single Commission

Comparative statics pertaining to model 1 in its interior region is the same as model 3, because the unconstrained versions of both models produce identical equilibrium objects (prices, matches, profits). In its corner region, in addition to the factors listed in Proposition 8, the direction of change in p_i depends on whether location i is a bottleneck or not. Since the net effect is somewhat cumbersome to disentangle, we skip the details in here.

More importantly, model 1 is based on flexible pricing, as such it is subject to the same criticism above: the platform posts different prices at different locations and it must update those prices as the transition matrix changes. So much variation and adjustment in prices might turn off some customers.

4.5 Calibration

In what follows we calibrate the model for New York City and Los Angeles based on real world ride patterns we extracted from a publicly available connectome map of rides on Uber's website (Uber, 2019; Bimpikis et al., 2019). Much like a classical connectome map showing point-to-point spatial connectivity of neural pathways in the brain, the connectome map that was available on Uber's website included a visual map of the ride patterns during a single month among the neighborhoods of these cities. Actual ride frequencies, however, were not readily available in this visual map. Using the open html code of the website, we were able to extract data including borders defining various neighborhoods in both cities (much like the n locations in our model), the latitude/longitude coordinates defining the center of each neighborhood and the relative likelihood of a ride going from one coordinate to the other during July 2014. We then created the transition matrix T using the ride patterns between the nodes, and the distance matrix via Google Maps API using the coordinates of the nodes (see the Appendix for the transition and distance matrices for NYC).[13]

Based on the transition matrix T which comprises of actual aggregate ride patterns in both cities, we next calculate the platform's profit under each scheme and cross-compare their performances. This helps us illustrate our results in a real-world setting. Figure 6 depicts equilibrium profits against the

[13] In both cities some diagonal elements of the transition matrices are non-zero, which means that a number of rides started and ended within the same neighborhood, e.g. Upper East Side in NYC. Similarly some non-diagonal elements were zero, indicating no rides took place between those locations. These facts violate our assumptions that $a_{i,j} > 0$ and $a_{i,i} = 0$, but they do not affect the inner workings of the model. Our assumptions are sufficient, but not necessary, to ensure that the transition matrix has a unique steady state vector σ. Our calculations show that the transition matrices associated with both cities are ergodic; thus they both have unique steady state vectors.

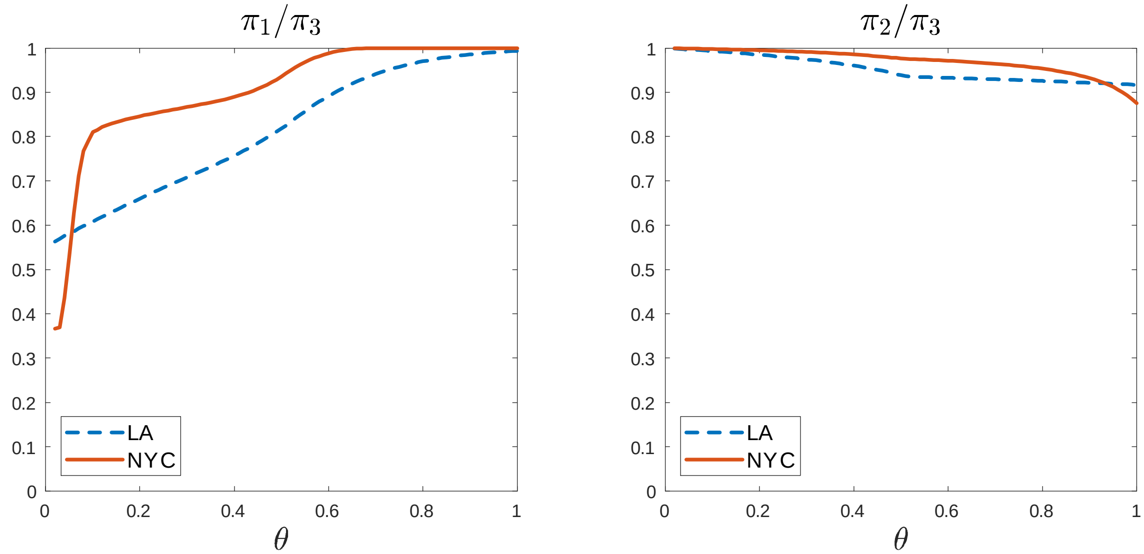

Figure 6: Profit Ratios

the number of available cars, θ. Model 3 (flexible prices, flexible commissions) is the most versatile operating system and encompasses the other models as special cases; as such π_3 serves as a benchmark in the simulations. The panel on the left hand side illustrates π_1 as a fraction of π_3; likewise the panel on the right hand side illustrates π_2/π_3.

First, the fixed-commission model (model 1) performs well when θ is large but it suffers when θ is small. For instance in New York if $\theta > 0.6$ (left panel) then π_1/π_3 exceeds 90%; however, if $\theta < 0.2$ then the ratio drops below 70%. The corresponding numbers in Los Angeles are even lower, because its traffic pattern is more varied than New York (more on this below). To see why, notice that with a large fleet (high θ) the platform does not run into bottlenecks, so, the flexibility of the commissions becomes rather unimportant. With a small fleet, however, bottlenecks are unavoidable and in model 1, due to the fixed commission structure, it is impossible to achieve a high utilization rate. This is why we observe a sharp drop in relative performance when θ falls.

Interestingly, model 2 exhibits the opposite pattern: if θ is large then, due to its fixed-price nature, it underperforms the benchmark, but if θ is small then the difference vanishes. The underperformance is not too severe either: in New York it is generally less than 5% and in Los Angeles less than 10%. The result underlines the importance of flexible commissions and its ability to utilize every car in a ride. Once this aspect is controlled for, the gains from location-specific pricing seem to be modest in comparison, especially when the fleet size is small.

While flexible pricing may appear as a first option to addess local demand and supply mismatch,

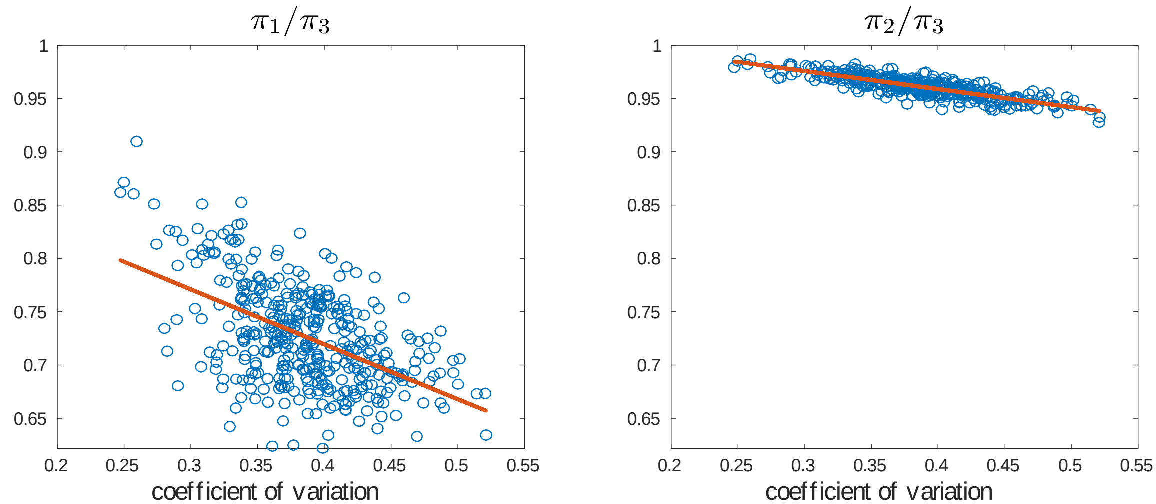

Figure 7: Traffic Imbalance and Profits in Randomly Generated Cities

our simulations highlight the importance of a flexible commission structure when there is a limited number of drivers. We find that when faced with a shortage of drivers it is the flexibility of the commissions, and not the flexibility of prices, that enables the platform to create more matches and generate more profits.

Finally, the performance of a model depends on how uniform the traffic pattern is. If the components of $\boldsymbol{\sigma}$ vary significantly, then we say the traffic in that city is rather non-uniform (relative to another city) as some locations are significantly more popular than others. A similar argument applies if the components of the distance vector $\mathbf{d}$ vary too much. In our model the equilibrium objects (prices and commissions) depend on σ_i and d_i; if they vary too much across locations, then so do prices and commissions. A non-uniform traffic structure, therefore, calls for significantly varied prices and commission rates across locations. The implication is that in a city where these variables differ significantly, pursuing a non-flexible policy is more costly. In New York the coefficient of variation for $\boldsymbol{\sigma}$ is 0.98, for $\mathbf{d}$ it is 0.32. The corresponding numbers in Los Angeles are 1.52 and 0.35, implying that Los Angeles indeed has a more varied traffic structure than New York. This explains why in the simulations the inflexible rules—price-wise or commission-wise—fare worse in Los Angeles than they do in New York.

In order to confirm the last insight, we randomly generated 400 cities, each consisting of 25 locations with distances varying from 2 to 10 miles. Accompanying transition matrices, too, were randomly generated. In each map, we computed the profits under each scheme (π_1, π_2 and π_3), as well as the coefficient of variation of $\boldsymbol{\sigma}$: the higher the coefficient of variation, the more varied the traffic pattern in that city. Figure 7 plots π_1/π_3 and π_2/π_3 in each map against the corresponding coefficient of variation. The downward trend in both panels confirms the above insight: pursuing a fixed rule—price-wise, commission-wise or both—becomes more costly for the platform as traffic patterns become more varied (as the coefficient of variation rises).

5 Concluding Remarks

On-demand platforms exhibit unique characteristics in terms of the nature of the work and supply of the workers. With limited number of independent workers having significant discretion over when and where to work, platforms need to rely on economic instruments and market forces to effectively serve consumers. In such an incentive-based, fluid economy, it is imperative to understand market implications and dynamics associated with various, often competing, incentive schemes. Focusing on ride-sharing platforms with spatially differentiated supply and demand, we analyze the selection and performance of prices and commissions— two common economic instruments—in creating the right incentives for independent drivers.

Our results suggest that a location-specific flexible commission policy is more effective in matching demand and supply than flexible pricing, especially if the platform does not have sufficiently many drivers at its disposal. In the absence of flexible commissions, the platform resorts to price hikes to prevent excess demand; however such interventions not only distort the unconstrained demand—and hamper profits—but also do a relatively poor job in incentivizing drivers to spread themselves evenly across the city. We show that, even after such price hikes, there is still excess supply at other more desirable locations, which, from the platform's point of view, is a waste of precious resources.

In contrast, with a flexible commission policy, by fine tuning the per-mile rates—decreasing them at more desirable locations, increasing them at less desirable locations, or a combination—the platform can spread the cars evenly across the city and avoid bottlenecks. With such a policy, the number of cars does not become a constraint until the customer-to-cab ratio hits 100% everywhere, i.e. until no car idles at any location. In our comparisons, we find that avoiding bottlenecks and utilizing every car in a ride is a significantly more important factor for profits than tailoring prices according to location-specific elements.

A second potential advantage of using commissions, rather than prices, in aligning supply with demand is the fact that commission rates are not observable by the customers. Whether the platform charges different commission rates at different locations, or whether it changes them, or keeps them intact, are only observed by the drivers. Customers experience the impact of such interventions (e.g.

an increase in the number of drivers at a location), but they do not observe the actual commission rates. In contrast, price modifications are observed by everyone, including the customers. It is conceivable that such flexible price policy might alienate certain customers and reduce profitability.

The recent rise of the 'gig' economy brings about new and unique challenges to marketplaces in that it is the norm rather than the exception for on-demand platforms to rely on a limited number of independent contractors as opposed to permanent workers to serve customers. On-demand providers from rides (Uber, Lyft) to deliveries (Grubhub, Deliveroo) to household tasks (Handy, Task Rabbit) employ gig workers who work on a fluid and flexible basis, yet the legal status of such workers is a major area of contention. Unlike regular employers, these providers cannot dictate such flexible workforce on what to do or when to do with regards to their work. Instead, they must rely on right incentives and market forces to provide effective service to consumers. Our work provides an important contribution by offering critical insights into the use and implications of two main economic instruments—prices and commissions—in effective management of gig workers in emerging two-sided market platforms.

References

Ariely, D. (2016). Why we hate uber's surge pricing. *The Wall Street Journal.* (June 20, 2016), https://www.wsj.com/articles/why-we-hate-ubers-surge-pricing-1453311373/.

Banerjee, S., Johari, R., and Riquelme, C. (2015). Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 639–639. ACM.

Banerjee, S., Kanoria, Y., and Qian, P. (2018). Dynamic assignment control of a closed queueing network under complete resource pooling. *arXiv*.

Basu, A. K., Lal, R., Srinivasan, V., and Staelin, R. (1985). Salesforce compensation plans: An agency theoretic perspective. *Marketing science*, 4(4):267–291.

Besbes, O., Castro, F., and Lobel, I. (2021). Surge pricing and its spatial supply response. *Management Science*, 67(3):1350–1367.

Bhargava, H. K. and Rubel, O. (2019). Sales force compensation design for two-sided market platforms. *Journal of Marketing Research*, 56(4):666–678.

Bimpikis, K., Candogan, O., and Saban, D. (2019). Spatial pricing in ride-sharing networks. *Operations Research*, 67(3):744–769.

Boyd, S., Boyd, S. P., and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.

Brown, E. (2019). Uber and lyft face hurdle of finding and keeping drivers. *The Wall Street Journal*. (May 12), https://www.wsj.com/articles/uber-and-lyft-face-tough-test-of-finding-and-keeping-drivers-11557673863.

Buchholz, N. (2021). Spatial equilibrium, search frictions and efficient regulation in the taxi industry. *Forthcoming, Review of Economic Studies*.

Bursztynsky, J. (2021). Why many uber and lyft drivers arenâĂŹt coming back. *CNBC*. (Jul 4), https://www.cnbc.com/2021/07/04/why-many-uber-and-lyft-drivers-arent-coming-back.html.

Cachon, G. P., Daniels, K. M., and Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3):368–384.

Castillo, J. C., Knoepfle, D., and Weyl, G. (2017). Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 241–242. ACM.

Chan, T. Y., Li, J., and Pierce, L. (2014). Compensation and peer effects in competing sales teams. *Management Science*, 60(8):1965–1984.

Chen, M. K. and Sheldon, M. (2015). Dynamic pricing in a labor market: Surge pricing and the supply of uber driver-partners. *Working Paper. (UCLA Anderson School)*.

Chung, D. J., Steenburgh, T., and Sudhir, K. (2014). Do bonuses enhance sales productivity? a dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2):165–187.

Cook, C., Diamond, R., Hall, J. V., List, J. A., and Oyer, P. (2020). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies*. (27 November 2020), https://doi.org/10.1093/restud/rdaa081.

Coughlan, A. T. and Sen, S. K. (1989). Salesforce compensation: Theory and managerial implications. *Marketing Science*, 8(4):324–342.

Cramer, J. and Krueger, A. B. (2016). Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–82.

Dholakia, U. (2015). Everyone hates uber's surge pricing here's how to fix it. *Harvard Business Review*. (December 21, 2015), https://hbr.org/2015/12/everyone-hates-ubers-surge-pricing-heres-how-to-fix-it/.

D.Lee (2021). Uber and lyft âĂŸthrowing moneyâĂŹ at us drivers to ease shortage. *The Financial Times*. (April 7), https://www.ft.com/content/2b71e009-bb96-46ee-b024-4c574ef41988.

Eckhardt, G. M., Houston, M. B., Jiang, B., Lamberton, C., Rindfleisch, A., and Zervas, G. (2019). Marketing in the sharing economy. *Journal of Marketing*, 83(5):5–27.

Einav, L., Farronato, C., and Levin, J. (2016). Peer-to-peer markets. *Annual Review of Economics*, 8:615–635.

Frechette, G. R., Lizzeri, A., and Salz, T. (2016). Frictions in a competitive, regulated market evidence from taxis. *CEPR Discussion Paper No. DP11626*. https://ssrn.com/abstract=2869080.

Grinstead, C. M. and Snell, J. L. (1998). *Introduction to probability*. American Mathematical Soc.

Guda, H. and Subramanian, U. (2019). Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science*, 65(5):1995–2014.

Hu, M. and Zhou, Y. (2019). Price, wage and fixed commission in on-demand matching. *Working Paper*. https://ssrn.com/abstract=2949513.

Huet, E. (2015a). Uber raises uberx commission to 25 percent in five more markets. *Forbes*. (September 11), https://www.forbes.com/sites/ellenhuet/2015/09/11/uber-raises-uberx-commission-to-25-percent-in-five-more-markets/3e4213723214.

Huet, E. (2015b). Uber tests taking even more from its drivers with 30 *Forbes Magazine*. (May 18), https://www.forbes.com/sites/ellenhuet/2015/05/18/uber-new-uberx-tiered-commission-30-percent/5928545e43f6.

Huet, E. (2015c). Uber's ever-renewing workforce: One-fourth of its current u.s. drivers joined last month. *Forbes*. (Jan 22), https://www.forbes.com/sites/ellenhuet/2015/01/22/uber-study-workforce/?sh=1bc99b7e367a.

Jain, S. (2012). Self-control and incentives: An analysis of multiperiod quota plans. *Marketing Science*, 31(5):855–869.

Kalra, A., Shi, M., and Srinivasan, K. (2003). Salesforce compensation scheme and consumer inferences. *Management Science*, 49(5):655–672.

Kim, M., Sudhir, K., Uetake, K., and Canales, R. (2019). When salespeople manage customer relationships: Multidimensional incentives and private information. *Journal of Marketing Research*, 56(5):749–766.

Kishore, S., Rao, R. S., Narasimhan, O., and John, G. (2013). Bonuses versus commissions: A field study. *Journal of Marketing Research*, 50(3):317–333.

Lagos, R. (2000). An alternative approach to search frictions. *Journal of Political Economy*, 108(5):851–873.

Li, H. and Srinivasan, K. (2019). Competitive dynamics in the sharing economy: an analysis in the context of airbnb and hotels. *Marketing Science*, 38(3):365–391.

Misra, S. and Nair, H. S. (2011). A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3):211–257.

Paul, K. (2021). Uber and lyft drivers join day-long strike over working conditions. *The Guardian*. (Jul 22), https://www.theguardian.com/technology/2021/jul/21/uber-lyft-drivers-strike-app-based-work-gig-economy.

Raju, J. S. and Srinivasan, V. (1996). Quota-based compensation plans for multiterritory heterogeneous salesforces. *Management Science*, 42(10):1454–1462.

Schöttner, A. (2017). Optimal sales force compensation in dynamic settings: Commissions vs. bonuses. *Management Science*, 63(5):1529–1544.

Siddiqui, F. (2021). You may be paying more for uber, but drivers arenâĂŹt getting their cut of the fare hike. *The Washington Post*. (June 10), https://www.washingtonpost.com/technology/2021/06/09/uber-lyft-drivers-price-hike/.

Taylor, C. (2014). Lawn love wants to bring the whole software eating the world thing to your own backyard. *TechCrunch*. (July 24), https://techcrunch.com/2014/07/24/lawn-love-brings-the-whole-software-eating-the-world-thing-to-your-own-backyard/.

Uber (2018). Uber newsroom. *Uber Offical Website*. (December 19), https://www.uber.com/newsroom/10-billion/.

Uber (2019). Uber newsroom. *Uber Offical Website*. (April 14), https://newsroom.uber.com/wp-content/uploads/2014/07/uber$_{l}a_{c}onnectome_{.}html$.

Uber (2020). Uber service fee. *Uber Offical Website*. (March 15), https://www.uber.com/gh/en/drive/basics/tracking-your-earnings/.

Uber (2021). Uber cities. *Uber Offical Website*. (September 1), https://www.uber.com/en-GB/cities/.

Wang, Y., Wu, C., and Zhu, T. (2019). Mobile hailing technology and taxi driving behaviors. *Marketing Science*, 38(5):734–755.

Wohlsen, M. (2013). Uber boss says surging prices rescue people from the snow. *Wired Magazine*. (December 12), https://www.wired.com/2013/12/uber-surge-pricing/.

Yao, D., Tang, C., and Chu, J. (2022). A dynamic model of owner acceptance in peer-to-peer sharing markets. *Marketing Science, Forthcoming*.

Zervas, G., Proserpio, D., and Byers, J. W. (2017). The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Journal of marketing research*, 54(5):687–705.

Zhang, S., Lee, D., Singh, P., and Mukhopadhyay, T. (2022). Demand interactions in sharing economies: Evidence from a natural experiment involving airbnb and uber/lyft. *Journal of Marketing Research*, 59(2):374–391.

6 Appendix A– Proofs

Proof of Lemma 1. A Markov chain with a finite state space is said to be *regular* if a power of its transition matrix has only positive entries. In our model n is finite and since $a_{i,j} > 0$ it is easy to verify that T^2 has only positive entries. It follows that the Markov chain associated with T is regular, and thus ergodic. For an ergodic Markov chain, there is a unique steady state vector

$$\boldsymbol{\sigma} = (\sigma_1, .., \sigma_n) \text{ with } \sigma_i > 0 \text{ and } \sum_{i=1}^{n} \sigma_i = 1$$

satisfying $\boldsymbol{\sigma} = \boldsymbol{\sigma} T$. Furthermore, any vector $\mathbf{v} > \mathbf{0}$ such that $\mathbf{v} = \mathbf{v}T$ must be a multiple of $\boldsymbol{\sigma}$ (see Grinstead and Snell (1998), Theorem 11.10). The steady state condition $\mathbf{m} = \mathbf{m}T$, therefore, implies that $m_i = \zeta \sigma_i$, where $\zeta > 0$ is a positive scalar. Since $\sum_{i=1}^{n} \sigma_i = 1$ we have $\zeta = \sum_{i=1}^{n} m_i \equiv M$. ■

Proof of Lemma 2. Let $\mathcal{S}_0$ denote the set of locations in which demand is less than or equal to supply with at least one location with excess demand, i.e. $\mathcal{S}_0 = \{i \in \mathbb{N}_+ : r_i \leq x_i\}$ with at least one inequality strict. Similarly let $\mathcal{S}_1$ denote the set of excess supply locations, i.e. $\mathcal{S}_1 = \{i \in \mathbb{N}_+ : r_i > x_i\}$.

Case 1 - $\mathcal{S}_1 \neq \emptyset$: Since $x_i \leq r_i$ we have $\eta_i = 1$ for all $i \in \mathcal{S}_0$. Similarly $x_i > r_i \Leftrightarrow \eta_i < 1$ for all $i \in \mathcal{S}_1$. The indifference condition (6) implies

$$p_i d_i = p_j d_j, \text{ for all } i, j \in \mathcal{S}_0 \qquad \text{and} \qquad p_i d_i < p_j d_j, \text{ for all } i \in \mathcal{S}_0 \text{ and } j \in \mathcal{S}_1.$$

Suppose that the platform leaves prices in $\mathcal{S}_1$ intact but increases prices in $\mathcal{S}_0$ to $p_i' = p_i + \varepsilon_i$, where ε_i is positive but infinitesimally small and satisfies

$$\varepsilon_i d_i = \varepsilon_j d_j, \text{ for all } i, j \in \mathcal{S}_0.$$

Note that $p_i' d_i = p_j' d_j$, which means $\eta_i' = \eta_j'$ for all $i, j \in \mathcal{S}_0$. It follows that either $\eta_i' = 1$ or $\eta_i' < 1$ for all $i \in \mathcal{S}_0$. If ε_i is small enough then one can ensure that

$$\eta_i' = 1, \text{ for all } i \in \mathcal{S}_0 \qquad \text{and} \qquad p_i' d_i < p_j d_j, \text{ for all } i \in \mathcal{S}_0 \text{ and } j \in \mathcal{S}_1,$$

i.e. no location in $\mathcal{S}_0$ exhibits excess supply and no location in $\mathcal{S}_1$ exhibits excess demand. Recall that $x_i < r_i$ at least at one location in $\mathcal{S}_0$, whereas $x_i > r_i$ at every location in $\mathcal{S}_1$. Such an ε_i exists because the inequalities are strict and ε_i can be infinitesimally small. Now we can compare profits.

Locations in $\mathcal{S}_0$: Prices are higher after the intervention. As for the number of rides, recall that prior to the intervention $m_i = x_i$. The fact that $\eta_i' = 1$ implies that after the intervention we still have $m_i' = x_i'$; however, note that $x_i' > x_i$. To see why, recall that

$$x_i = \frac{\sigma_i p_i d_i \theta}{\sum_{j=1}^n \sigma_j p_j d_j}.$$

After the intervention, prices in $\mathcal{S}_0$ go up while prices in $\mathcal{S}_1$ remain the same. Since $\varepsilon_i d_i = \varepsilon_j d_j$ we have $p_i' d_i = p_j' d_j$ for all $i, j \in \mathcal{S}_0$. Taking location i as a reference point we have

$$x_i' = \frac{\sigma_i p_i' d_i \theta}{p_i' d_i \sum_{j \in \mathcal{S}_0} \sigma_j + \sum_{j \in \mathcal{S}_1} \sigma_j p_j d_j} \text{ for all } i \in \mathcal{S}_0.$$

It is easy to verify that $x_i' > x_i$ because $p_i' > p_i$. Since both the prices and the number of rides go up, the platform earns more in $\mathcal{S}_0$ than it did before.

Locations in $\mathcal{S}_1$: Prices remain intact. The number of rides is also the same as before, because after the intervention we still have $\eta_i' < 1$; thus the number of rides is still equal to $y_i (1 - p_i)$ (recall that p_i remains unchanged). It follows that the platform earns the same in $\mathcal{S}_1$ as it did before. The intervention allows the platform to move idle drivers in $\mathcal{S}_1$ to $\mathcal{S}_0$ and earn more; thus the initially conjectured outcome cannot be an equilibrium.

Case 2 - $\mathcal{S}_1 = \emptyset$: Along this outcome $x_i \leq r_i$ for all $i = 1, ..., n$ with at least one inequality strict; thus $\eta_i = 1$ for all $i = 1, .., n$. Pick location j as a reference point, and note that since $\eta_i = 1$ the indifference condition (6) becomes $p_i d_i = p_j d_j$, for all i. Substituting this relationship into (8) we have $x_i = \sigma_i \theta$ for all i. Recall that $r_i = y_i (1 - p_i)$; thus $x_i \leq r_i \Leftrightarrow y_i \geq \sigma_i \theta / (1 - p_i)$. It follows that

$$\sum_{i=1}^n y_i > \sum_{i=1}^n \frac{\sigma_i \theta}{1 - p_i} \Leftrightarrow \Delta(p_j) < 0,$$

where

$$\Delta(p_j) = \theta - \left[\sum_{i=1}^n \frac{\sigma_i d_i}{d_i - p_j d_j} \right]^{-1}.$$

The second step obtains because $\sum_{i=1}^n y_i = 1$ and $p_i d_i = p_j d_j$. The inequality $\Delta(p_j) < 0$ is strict because at least one location has $x_i < r_i$. Note that Δ increases in p_j and $\Delta(1) > 0$. Since $\Delta(p_j) < 0$, there exists some $p_j' \in (p_j, 1)$ satisfying $\Delta(p_j') = 0$. So, if the platform increases p_j to p_j' at location j, while also ensuring that $p_i' d_i = p_j' d_j$ at other locations, then $x_i' = r_i'$ for all i, i.e. no location exhibits excess demand. Prior to the intervention we had $x_i \leq r_i$, with at least one strict inequality; thus the number of rides was equal to $m_i = x_i = \sigma_i \theta$ for all i. After the intervention we have $x_i' = r_i'$; thus, the number of rides is still equal to $m_i' = x_i' = \sigma_i \theta$ for all i. Prices, on the other hand, are now higher, which means that the platform earns more than before. It follows that the initially conjectured outcome cannot be an equilibrium. ■

Proof of Lemma 3. We start by showing that $h^{-1}(\mathbf{p})$is strictly concave. The strategy is to establish that $h^{-1}(\mathbf{p})$ lies underneath its linearization at some $\mathbf{p}^0$, which is given by

$$\hat{h}^{-1}(\mathbf{p}) = h^{-1}(\mathbf{p}^0) + \nabla h^{-1}\left(\mathbf{p}^0\right)\left(\mathbf{p} - \mathbf{p}^0\right) = \frac{\sum_i \frac{\sigma_i(1-p_i)}{(1-p_i^0)^2}}{\left[\sum_{i=1} \frac{\sigma_i}{1-p_i^0}\right]^2}.$$

The function is concave if $h^{-1}(\mathbf{p}) < \hat{h}^{-1}(\mathbf{p})$, i.e. if

$$\left[\sum_{i=1}^{n} \frac{\sigma_i}{1-p_i^0}\right]^2 < \sum_{i=1}^{n} \frac{\sigma_i\left(1-p_i\right)}{(1-p_i^0)^2} \sum_{i=1}^{n} \frac{\sigma_i}{1-p_i}.$$

Letting $t_i \equiv \sqrt{\frac{\sigma_i(1-p_i)}{(1-p_i^0)^2}}$ and $s_i \equiv \sqrt{\frac{\sigma_i}{1-p_i}}$, the inequality becomes

$$\left[\sum_{i=1}^{n} t_i s_i\right]^2 < \sum_{i=1}^{n} t_i^2 \sum_{i=1}^{n} s_i^2.$$

The result follows from Cauchy-Scwharz. Note that the inequality is strict; thus $h^{-1}(\mathbf{p})$is strictly concave. Observe that $\Omega = g(\mathbf{p})h^{-1}(\mathbf{p})$, where g is linear and increasing; whereas h^{-1} is strictly concave and decreasing in p. Thus Ω is strictly concave (see Boyd et al. (2004), pg. 119).■

Proof of Lemma 4. Start with the first claim. If $\lambda_k = 0$ then the constraint is slack at location k, thus (i) $p_k d_k \theta > \Omega(\mathbf{p})$. Furthermore the first order condition implies $\Omega_k' = 0$; thus (ii) $d_k\left(1-p_k\right)^2 = \Omega(\mathbf{p})$. (In the first order condition, the term $1 - \sum_{i=1}^{n} \lambda_i$ cannot be equal to zero, because otherwise (12) fails to hold at any location with a positive λ_i.)

Now by contradiction suppose $\lambda_{k+1} > 0$. Since the constraint is assumed to bind at location $k+1$ we have (iii) $p_{k+1} d_{k+1} \theta = \Omega(\mathbf{p})$. Furthermore, since Ω is strictly concave and the constraint is assumed to bind we have $\Omega_{k+1}' < 0$, thus (iv) $d_{k+1}\left(1-p_{k+1}\right)^2 < \Omega(\mathbf{p})$. Since $d_{k+1} > d_k$, equations (ii) and (iv) together imply that

$$d_k\left(1-p_k\right)^2 > d_{k+1}\left(1-p_{k+1}\right)^2 \Rightarrow p_{k+1} > p_k.$$

Notice, however, (i) and (iii) together imply that $p_k > p_{k+1}$; a contradiction. Thus λ_{k+1} cannot be positive; so, it must be zero. The second part of the Lemma is proved similarly. ■

Proof of Proposition 1. In the unconstrained case ($k = 0$) prices satisfy $(1 - p_i)^2 d_i = \Omega$ for all $i = 1, ..., n$. (We omit the superscript 0 when understood.) It follows that $p_j = 1 - (1 - p_i)\sqrt{d_i/d_j}$; hence

$$\Omega = \frac{(1 - p_i)\sqrt{d_i}}{\mathbb{E}_\sigma(\sqrt{d})}[1 - (1 - p_i)\sqrt{d_i}\mathbb{E}_\sigma(\sqrt{d})].$$

Solving $(1 - p_i)^2 d_i = \Omega$ for p_i yields the expression in the body of the proposition. The equilibrium commission rate can be obtained by substituting p_i into (9). For this (interior) equilibrium to emerge we need $p_1 d_1 \theta > \Omega$, i.e. the constraint at location 1 ought to be slack. After substituting for prices, the condition is equivalent to

$$\theta > \bar{\theta}_1 \equiv \frac{\mathbb{E}_\sigma^2(d)}{2\mathbb{E}_\sigma(\sqrt{d})[2d_1\mathbb{E}_\sigma(\sqrt{d}) - \sqrt{d_1}\mathbb{E}_\sigma(d)]}. \tag{19}$$

If the constraint is slack at location 1 then it is slack at every other location (Lemma 4); thus $\theta > \bar{\theta}_1$ is sufficient. ■

Proof of Proposition 2. In regime-k we have

$$(1 - p_k^k)^2 d_k < \Omega(\mathbf{p}^k) \quad \text{and} \quad p_{k+1}^k d_{k+1}\theta > \Omega(\mathbf{p}^k)$$

The first inequality is due to the fact that the constraint binds at location k, whereas the second one obtains because the constraint is slack at location $k + 1$. Since constraints kick in in an orderly fashion we do not need to worry about constraints prior to k or after $k + 1$. Furthermore, per (13), p_k^k and p_{k+1}^k satisfy

$$p_k^k d_k \theta = d_{k+1}(1 - p_{k+1}^k)^2 = \Omega(\mathbf{p}^k).$$

Substituting these relationships into the inequalities above reveals that both p_k^k and p_{k+1}^k exceed $p_{\min}$. The inequalities pertaining to the other prices follow from the fact that $d_i < d_{i+1}$.■

Proof of Lemma 5. The commission vector $\mathbf{c}$ is incentive compatible if it satisfies (7), which, after substituting for m_i and x_i is equivalent to

$$(1 - p)\, p \sum_{i=1}^{n} \sigma_i d_i c_i = \theta w.$$

Furthermore, recall that $r_i \leq x_i \Leftrightarrow w \leq pc_i d_i$. Per our conjecture, under $\mathbf{c}$ we have $pc_i d_i > w$ for $i \leq k$ and $pc_i d_i = w$ for $i \geq k+1$.[14] We will construct a new $\hat{\mathbf{c}}$ by marginally shaving off the rates of $\mathbf{c}$ at locations where the constraint is slack (but without rendering any of these constraints binding) and marginally increasing the rates at locations where the constraint is binding. Let

$$\hat{c}_i = c_i - \varepsilon_i \text{ for } i \leq k \quad \text{and} \quad \hat{c}_i = c_i + \varepsilon_i \text{ for } i \geq k+1,$$

where $(\varepsilon_1, ..., \varepsilon_n) \in (0,1)^n$ is an arbitrarily small tuple satisfying

$$\sum_{i=1}^{k} \sigma_i d_i \varepsilon_i = \sum_{i=k+1}^{n} \sigma_i d_i \varepsilon_i.$$

Note that

$$\sum_{i=1}^{n} \sigma_i d_i \hat{c}_i = \sum_{i=1}^{k} \sigma_i d_i (c_i - \varepsilon_i) + \sum_{i=k+1}^{n} \sigma_i d_i (c_i + \varepsilon_i) = \sum \sigma_i d_i c_i,$$

thus $\hat{\mathbf{c}}$ is incentive compatible. In addition, since $(\varepsilon_1, ..., \varepsilon_n)$ can be picked arbitrarily small, the inequality $p_i \hat{c}_i d_i > w$ can be satisfied for all i. ■

Proof of Proposition 3. Ignoring the constraints, it is straightforward to show that the platform sets $p = 1/2$. The constraints are slack if $w < pc_i d_i$, i.e. if $2w < c_i d_i$ for all i. The commission rates must satisfy (16), which, after substituting for $p = 1/2$ becomes

$$\sum_{i=1}^{n} \sigma_i d_i c_i = 4\theta w. \tag{20}$$

There are n commission rates, a single equality constraint and n inequality constraints. Substituting $2w = c_i d_i$ into (20) yields the lower bound $\bar{\theta}_2 = 0.5$; thus if $\theta > \bar{\theta}_2$ then there exists a continuum of commission rates satisfying the system above, i.e. we have a continuum of equilibria, each with the same (interior) price $p = 1/2$, but different commission rates. If $\theta \leq 0.5$ then all constraints must bind, i.e. $w = pc_i d_i$ for all i. These equalities, together with (16), uniquely pin down the solution as $p = 1 - \theta$ and $c_i = w/(1-\theta) d_i$.■

Proof of Proposition 4. Suppose the constraints are slack. Then the platform's problem is the same as the unconstrained problem in model 1; thus p_i^{interior} is the same as the interior price there.

[14] For ease of exposition, we assume that the constraints are slack at locations $1, ..., k$ and bind at $k+1, .., n$; however, this is without loss in generality. The proof can be recast when the constraints are slack/binding at some randomly selected locations.

Now suppose the constraints are active, i.e. suppose $p_i c_i d_i = w$, for all i. Substituting these equalities into (18) yields $h(\mathbf{p}) = 1/\theta$. The platform, therefore, solves

$$\max_{\mathbf{p}} \ \Omega(\mathbf{p}) - \theta w \ \text{ s.t. } \ h(\mathbf{p}) = 1/\theta,$$

while the commission rates are uniquely pinned down via $p_i c_i d_i = w$, for all i. Letting λ denote the Lagrange multiplier, the first order condition with respect to p_i is given by (recall that $\Omega(\mathbf{p})$ is strictly concave and $h(\mathbf{p})$ is strictly convex)

$$d_i (1 - p_i)^2 = \Omega + \lambda h(\mathbf{p}), \text{ for all } i.$$

Since the right hand side is not indexed by i, we have

$$d_i (1 - p_i)^2 = d_j (1 - p_j)^2 \Rightarrow p_i = 1 - (1 - p_j) \sqrt{d_j / d_i}.$$

Combining this relationship with the constraint $h(\mathbf{p}) = 1/\theta$ yields the expression for the corner price in the body of the proposition. The commission rates can be recovered from $p_i c_i d_i = w$, whereas the threshold $\bar{\theta}_3$ can be obtained via $\theta = 1/h\left(\mathbf{p}^{interior}\right)$.■

Proof of Proposition 5. First we show that $\mathbb{E}_\sigma(d) > \mathbb{E}_\sigma^2(\sqrt{d})$. The inequality can be written as

$$\sum \sigma_i d_i > \left[\sum_{i=1}^{n} \sigma_i \sqrt{d_i}\right]^2.$$

Letting $t_i \equiv \sqrt{\sigma_i d_i}$ and $s_i \equiv \sqrt{\sigma_i}$ and noting that $\sum \sigma_i = 1$, the inequality becomes

$$\sum_{i=1}^{n} t_i^2 \sum_{i=1}^{n} s_i^2 > \left[\sum_{i=1}^{n} t_i s_i\right]^2.$$

The result follows from Cauchy-Schwarz. It is straightforward to show that the inequalities (i) $\bar{\theta}_1 > \bar{\theta}_3$ and (ii) $\bar{\theta}_3 > \bar{\theta}_2$ boil down to, respectively, (i) $\mathbb{E}_\sigma^2(\sqrt{d}) > d_1$ and (ii) $\mathbb{E}_\sigma(d) > \mathbb{E}_\sigma^2(\sqrt{d})$. The first inequality holds because $d_1 < d_i$ whereas the second inequality is already established above. ■

Proof of Proposition 6. Focus on M_2 and M_3 and suppose that $\theta \leq \bar{\theta}_2$. Since $\bar{\theta}_3 > \bar{\theta}_2$, we have $M_2 = M_3 = \theta$. Now suppose $\bar{\theta}_2 < \theta \leq \bar{\theta}_3$. In this region $M_3 = \theta$, whereas $M_2 = 1/2$. Since $\theta > \bar{\theta}_2 = 0.5$ we have $M_3 > M_2$. Finally, if $\theta > \bar{\theta}_3$ then $M_3 = \mathbb{E}_\sigma(d) / 2\mathbb{E}_\sigma^2(\sqrt{d})$ and $M_2 = 1/2$. The inequality $M_3 > M_2$ obtains because $\mathbb{E}_\sigma(d) > \mathbb{E}_\sigma^2(\sqrt{d})$. ■

Proof of Proposition 7. We will show that $\pi_3 > \pi_2$. There are three cases to consider:

(i) Suppose $\theta > \bar{\theta}_3$. In this region $\pi_3 > \pi_2$ obtains because $\mathbb{E}_\sigma(d) > \mathbb{E}^2_\sigma(\sqrt{d})$.

(ii) Suppose $\bar{\theta}_2 < \theta \leq \bar{\theta}_3$. The inequality $\pi_3 > \pi_2$ holds if

$$\theta[\mathbb{E}_\sigma(d) - \theta\mathbb{E}^2_\sigma(\sqrt{d})] > \mathbb{E}_\sigma(d)/4.$$

Since $\theta > \bar{\theta}_3$ the left hand side is increasing in θ; thus for a sufficient condition substitute $\theta = \bar{\theta}_2$ and note that inequality holds if $\mathbb{E}_\sigma(d) > \mathbb{E}^2_\sigma(\sqrt{d})$, which is true (proof of Proposition 5)

(iii) Suppose $\bar{\theta}_2 \geq \theta$. The inequality $\pi_3 > \pi_2$ holds if $\mathbb{E}_\sigma(d) > \mathbb{E}^2_\sigma(\sqrt{d})$, which is true.

As for π_1 note that if $\theta \geq \bar{\theta}_1$ then $\pi_1 = \pi_3$; thus, per the arguments above $\pi_1 > \pi_2$. If, however, $\theta < \bar{\theta}_1$ then we need numerical simulations to establish the magnitude of π_1.■

Proof of Proposition 8. Since d_i is independent of q, we have

$$\frac{dp_i^{interior}}{dq} \propto -\frac{d}{dq}\left[\frac{\mathbb{E}_\sigma(d)}{\mathbb{E}_\sigma(\sqrt{d})}\right] \quad \text{and} \quad \frac{dp_i^{corner}}{dq} \propto -\frac{d}{dq}\mathbb{E}_\sigma(\sqrt{d}),$$

where

$$\mathbb{E}_\sigma(d) = \frac{2q\delta + (1-q)\tilde{\delta}}{1+q} \qquad \text{and} \qquad \mathbb{E}_\sigma(\sqrt{d}) = \frac{q\sqrt{\delta} + \sqrt{q\delta + (1-q)\tilde{\delta}}}{1+q}.$$

It is straightforward to show that if $\delta < \tilde{\delta}$ then both $\mathbb{E}'_\sigma(\sqrt{d}) < 0$ and $[\mathbb{E}_\sigma(d)/\mathbb{E}_\sigma(\sqrt{d})]' < 0$; thus p_i increases in q. Turning to profits,

$$\frac{d\pi_3^{interior}}{dq} \propto \frac{d}{dq}\left[\frac{\mathbb{E}_\sigma(d)}{\mathbb{E}_\sigma(\sqrt{d})}\right] \quad \text{and} \quad \frac{d\pi_3^{corner}}{dq} \propto \frac{d}{dq}\left[\mathbb{E}_\sigma(d) - \theta\mathbb{E}^2_\sigma(\sqrt{d})\right].$$

From the previous step it is clear that if $\delta < \tilde{\delta}$ then $\pi_3^{interior}$ falls in q. As for π_3^{corner} observe that $d\pi_3^{corner}/dq$ decreases in θ. If $\theta = \bar{\theta}_3$ then $\pi_3^{interior} = \pi_3^{corner}$, thus $d\pi_3^{corner}/dq < 0$. In the corner equilibrium $\theta \leq \bar{\theta}_3$, thus $d\pi_3^{corner}/dq < 0$ whenever $\theta < \bar{\theta}_3$. ■

7 Appendix B– Transition and Distance Matrices

Table 1: Transition Matrix for NYC

		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	σ_i
Battery Park	1	0	0	0	0	0	0	0	0	10.4	10.3	8.3	6.5	17.2	9.9	3.8	5.2	5	0	5	3.5	4.4	3.3	0	0	7	*1.11*
Carnegie Hill	2	0	0	0	0	0	0	0	0	10.4	12.1	0	0	30.9	10.4	12.4	0	0	0	0	0	0	0	0	0	23.4	*0.64*
Central Park	3	0	0	0	0	0	0	0	0	8.6	13.7	7.9	0	24.4	15.4	9.9	0	0	0	0	0	0	0	0	0	20.1	*0.73*
Harlem	4	0	0	0	0	0	0	0	0	0	20	0	0	36.9	18.7	0	0	0	0	0	0	0	0	0	0	24.2	*0.25*
N. Sutton Area	5	0	0	0	0	0	0	0	0	12.2	8.8	6	6.9	25.5	11.6	6.8	6.8	0	0	0	0	0	0	0	0	15.3	*0.54*
Chinatown	6	0	0	0	0	0	0	0	0	16.7	7.9	8.7	10.1	18	14.6	0	0	0	0	9	0	7.5	0	0	0	7.3	*0.24*
Yorkville	7	0	0	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0	0	0	0	*0.12*
East Harlem	8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	100	*0.04*
West Village	9	1.4	0.8	0.8	0	0.7	0.6	0	0	6.6	8.2	6.3	6.2	16.3	10.5	4.3	3.6	3.9	2.5	4.6	3.7	3.6	3.2	2.4	1.8	8.3	*7.71*
Gramercy	10	1.6	0.7	1.1	0.7	0.5	0.5	0	0	7.5	7.2	5.2	6.6	16.3	9.1	5.2	4.5	4.6	2.5	3.5	3.3	3.9	2.8	1.9	2.8	8.1	*8.17*
Soho	11	1.2	0.4	0.7	0.4	0.6	0.4	0	0	10.7	8.1	6.3	6	14.5	11.7	3.6	4.2	3.4	2.7	5.2	3.4	3.5	2.8	1.1	2.6	6.5	*5.22*
Greenwich Vil.	12	1.7	0.8	0.8	0.4	0.7	0.5	0.3	0	7.7	7.9	4.9	6.4	14.8	10.1	4.9	4	3.7	3.5	5.1	3.3	3.8	2.6	2.4	1.9	8	*6.38*
Midtown	13	1.3	0.9	1	0.5	0.7	0.3	0.3	0.2	7.6	8.8	4.7	6.5	12.7	9.2	8.1	4.6	3.1	2.6	4	2.7	2.1	3.2	1.6	1	12.1	*16.1*
Chelsea	14	1.1	0.7	0.7	0.5	0.6	0.4	0.3	0.2	8	7	5.6	6.8	15.9	10.3	5.1	3.5	4.4	3.1	3.7	3.2	3.4	4.1	1.9	1.6	7.8	*9.95*
U. West Side	15	1	1.3	1.2	0	0	0	0	0	5.1	7.5	4.7	4.7	23.7	9.8	9.3	4.6	2.5	2.8	2.5	2.4	1.8	2.8	1	1.2	10.1	*5.51*
Garment Dist.	16	1.4	1	0.8	0	0.9	0	0	0	6.5	8.4	4.7	6.2	17.9	9.1	5.3	2.9	4.1	2.6	3.7	2.8	2.1	2.4	1.8	1.8	13.2	*4.26*
East Village	17	0.9	0	0.6	0	0	0	0	0	6.4	10.7	7.1	6.8	12.9	11.5	3.6	6	4.5	2.6	3.1	4.5	4.2	2.6	2.3	3.1	6.4	*3.82*
Murray Hill	18	0.9	0	0.9	0	0.8	0	0	0	8.2	7.3	5.6	7.9	15	11.5	5.7	4.9	5.2	3.7	3.6	2.8	2.7	2.5	1.7	1.3	7.9	*2.55*
Tribeca	19	1.3	0	0	0	0	0.6	0	0	10.6	7.6	6.2	8.1	12.6	10.5	3.3	3.9	4	1.9	5.4	3.8	3.8	2.4	3.6	2.6	7.7	*3.93*
Financial Dist.	20	1	0	0	0	1.1	0	0	0	11.5	7.3	5.4	7.6	12.9	9.9	4.1	3.7	5.4	2.6	5.2	5.1	4.3	2.8	1.8	2.2	6	*3.09*
L. East Side	21	1.5	0	0	0	0	0	0	0	9	11.2	5	7.3	11.3	10.6	3.5	4	5	3.4	3.5	4.1	4.9	2.9	1.8	3.7	7.1	*3.09*
Clinton	22	0.8	0	0	0	0	0	0	0	7.6	7.9	5.4	6.2	18.3	12.1	5.1	4.5	4	2.7	4	2.5	3.1	5.7	1.7	1.8	6.5	*2.91*
Little Italy	23	0	0	0	0	0	0	0	0	9.4	8.1	3.4	5.7	17.4	11.6	3.6	5.4	5	2.5	4.7	3.9	3.1	2.5	2.3	3.9	7.5	*1.78*
Williamsburg	24	0	0	0	0	0	0	0	0	7.6	9.4	6.3	8.2	9.7	8.6	3.5	2.4	10	0	4.7	5	6.7	2.5	3.5	7.2	4.5	*1.86*
U. East Side	25	0.7	1.2	1.2	0.3	1.1	0	0.4	0	5.9	7	3.8	6	20.5	7.6	5.6	5.7	2.5	2.3	3.3	2	2	2.3	1.2	0.7	16.7	*9.79*

The transition matrix T provides the likelihood of rides across locations. Consider, for instance, location 1 (Battery Park): 10.4% of rides originating from this location are directed towards West Village, 10.3% are directed towards Gramercy, 8.3% towards Soho and so on. Given T, we calculate its steady state distribution σ_i (last column), which reveals the "centrality" of each location. Midtown appears to be the most active location as it sends (and receives) 16.1% of the total ride-sharing traffic. It is followed by Chelsea (9.95%), Upper East Side (9.79%), Gramercy (8.17%) and West Village (7.71%).

Table 2: Distance Matrix for NYC

		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	d_i
Battery Park	1	0.5	9.32	10	11.6	6.44	2.58	9.18	9.99	1.97	5.22	1.57	2.39	7.01	3.01	7.35	3.98	4.32	5.81	0.93	1.66	2.81	3.73	2.03	7.79	7.73	*4.13*
Carnegie Hill	2	9.06	0.5	1.03	2.67	3.2	8.07	0.94	1.75	6.99	4.67	7.53	4.32	2.33	6.26	1.95	3.35	6.45	4.11	9.79	8.06	6.92	5.15	7.08	8.56	1.29	*3.22*
Central Park	3	9.01	1	0.5	3.39	3.14	8.01	1.66	2.47	6.76	4.62	7.47	4.14	2.15	6.04	0.87	3.17	6.4	4.05	9.73	8.01	6.86	4.92	7.03	8.51	1.21	*3.59*
Harlem	4	11.1	2.54	3.4	0.5	5.3	10.1	3.26	1.39	8.61	6.77	9.63	8.1	6.25	7.42	3.66	7.02	8.56	6.21	11.8	10.1	9.02	6.31	9.19	11	3.87	*5.98*
N. Sutton Area	5	7.2	3.15	2.68	5.5	0.5	6.21	3.01	3.82	4.64	2.81	5.67	4.14	1.34	3.25	3.28	2.35	4.59	1.73	7.93	6.2	5.06	2.53	5.22	6.7	1.47	*2.77*
Chinatown	6	2.13	8.23	8.98	10.5	5.35	0.5	8.09	8.9	1.7	4.13	0.86	1.71	3.61	2.87	7.22	3.21	1.47	4.72	0.9	1.13	0.88	3.6	0.5	2.74	6.64	*2.56*
Yorkville	7	9.55	1.04	1.78	2.31	3.68	8.55	0.5	0.8	6.99	5.16	8.02	6.49	4.64	6.61	2.64	5.41	6.94	4.59	10.2	8.55	7.41	5.68	7.57	9.05	1.77	*4.64*
East Harlem	8	10	1.81	2.56	1.97	4.2	9.07	1.36	0.5	7.51	5.68	8.53	7.01	5.15	7.13	3.31	5.93	7.46	5.11	10.7	9.07	7.92	12.2	8.09	10.6	2.17	*2.17*
West Village	9	1.94	6.57	6.31	7.91	8.2	1.82	6.83	7.64	0.5	2.55	1.35	0.92	3.17	1.35	5.7	2.33	2.4	3.11	1.47	3.43	2.31	2.08	1.53	4.18	9.5	*2.97*
Gramercy	10	5.04	5.04	5.79	7.39	1.83	1.92	4.9	5.71	1.82	0.5	2.14	1.32	2.15	1.44	5.26	1.73	1.29	0.92	2.5	4.04	2.9	2.41	1.71	3.72	3.46	*2.35*
Soho	11	1.97	7.68	7.2	10	4.8	1.06	7.55	8.36	0.97	2.13	0.5	0.85	2.75	2.24	6.59	2.34	1.66	2.68	0.87	2.01	1.32	2.97	0.51	3.03	6.1	*2.38*
Greenwich Vil.	12	2.3	6.29	4.82	8.63	3.07	1.42	6.15	6.96	0.92	1.34	1.3	0.5	2.43	1.73	6.56	2.02	1.17	1.9	1.63	2.42	1.8	2.7	0.99	3.48	4.7	*2.29*
Midtown	13	7.09	2.53	2.39	4.92	1.21	3.81	3.19	4	2.49	2.15	3.38	2.44	0.5	1.9	2.99	1.01	4.48	1.32	3.63	6.09	4.95	1.19	3.41	7.44	1.92	*2.42*
Chelsea	14	3.44	6.25	5.91	7.51	2.93	2.59	6.11	6.92	1.27	1.3	2.16	1.49	1.82	0.5	5.3	0.98	2.5	1.86	2.41	4.93	4.11	1.41	2.34	4.83	4.66	*2.37*
U. West Side	15	7.89	1.77	0.82	4.06	3.75	8.62	2.34	3.15	6.29	5.22	7.49	4.74	2.75	5.57	0.5	5.35	7	4.66	7.42	8.61	7.47	4.45	7.64	9.11	1.81	*4.21*
Garment Dist.	16	6.34	3.37	3.23	7.36	1.95	3.06	4.88	5.69	1.74	1.41	2.63	1.69	0.84	1.15	5	0.5	3.73	0.78	2.88	5.34	4.2	1.11	2.67	6.72	2.8	*2.25*
East Village	17	4.3	5.85	6.6	8.2	2.64	1.39	5.72	6.53	2.14	1.29	1.67	1.04	3.21	2.54	6.08	2.83	0.5	2.02	2.03	3.3	0.96	4.13	1.2	2.77	4.27	*2.46*
Murray Hill	18	6.16	2.89	3.03	6.58	1.17	5.16	4.1	4.91	2.48	1.02	2.87	1.98	1.19	1.9	3.63	1	3.55	0.5	6.88	5.16	4.01	1.49	2.56	6.31	1.99	*2.36*
Tribeca	19	1	8.74	7.58	11	5.86	0.77	8.6	9.41	1.45	2.79	0.66	1.51	3.41	2.62	6.97	3.01	2.33	3.35	0.5	1.4	1.4	3.35	0.92	3.37	7.16	*2.64*
Financial Dist.	20	1.12	9.1	9.85	11.4	6.22	2.37	8.96	9.77	2.83	5	2.01	2.86	6.79	3.87	8.21	6.41	4.1	5.6	1.79	0.5	2.59	4.59	2.69	4.67	7.51	*4.26*
L. East Side	21	3.58	7.21	7.96	9.55	4.33	0.8	7.07	7.88	2.27	3.11	1.29	1.88	4.9	3.57	7.43	4.52	1.09	3.71	1.4	1.93	0.5	5.21	0.88	2.06	5.63	*3.16*
Clinton	22	3.8	4.75	4.49	6.09	2.4	3.69	5.01	5.82	2.29	2.35	3.4	2.88	1.19	1.39	3.88	0.94	4.68	1.52	3.34	5.29	5.14	0.5	3.55	7.64	3.13	*2.46*
Little Italy	23	2.01	7.21	7.96	9.56	4.33	0.4	7.07	7.88	1.64	1.69	0.74	1.18	3.39	2.9	7.25	2.99	1.19	2.24	1.24	1.62	0.81	3.63	0.5	2.52	5.63	*2.61*
Williamsburg	24	7.54	8.88	7.67	12.6	8.4	3.98	8.74	11.7	5.58	7.92	4.74	3.92	8.55	6.75	8.27	8.15	2.86	7.33	4.69	4.72	2.88	8.82	2.93	0.5	6.64	*5.29*
U. East Side	25	8.11	1.07	1.2	3.74	1.19	7.11	1.43	2.24	5.55	3.71	6.57	5.04	1.97	5.17	1.8	3.96	5.5	3.15	8.83	7.11	5.96	3.42	6.13	7.61	0.5	*3.43*

The distance matrix D is obtained by querying Google Maps API with the lat-long coordinates of the above locations. The API's result reflects the distance along the fastest route at the time of querying, which depends on the traffic at that time, as such $\delta_{i,j}$ is not necessarily equal to $\delta_{j,i}$. A second issue is the trip lengths for rides starting and ending in the same neighborhood. We assigned 0.5 mile for such rides; however, the results remain the same with other (sensible) trip lengths. Given T and D, we can calculate the expected distances d_i associated with each location (last column). In this dimension, Harlem is the leading location: trips originating from Harlem are expected to take 5.98 miles. For the majority of other locations this number is less than 3 miles. Given $\mathbf{d}$ and $\boldsymbol{\sigma}$ we can calculate the equilibrium prices, commissions and profits under each operating model, which we discuss in the main text. The corresponding tables for Los Angeles are available upon request.