

What *Drives* the Efficiency in Ridesharing Markets?

Evidence from Austin, Texas

Motaz Al-Chanati*

Vinayak Iyer[†]

November 2021

[Click here for most recent version](#)

Abstract

In decentralized transportation markets, search and match frictions lead to inefficient outcomes. Ridesharing platforms, who act as intermediaries in traditional taxi markets, improve upon the status quo along two key dimensions: surge pricing and centralized matching. We study how and why these two features make the market more efficient; and explore how alternate pricing and matching rules can improve outcomes further. To this end, we develop a structural model of the ridesharing market with four components: (1) dynamically optimizing drivers who make entry, exit and search decisions; (2) stochastic demand; (3) surge pricing rule and (4) a matching technology. Relative to our benchmark model, surge pricing generates large gains for all agents; primarily during late nights. This is driven by the role surge plays in inducing drivers to enter the market. In contrast, centralized matching reduces match frictions and increases surplus for consumers, drivers, and the ridesharing platform, irrespective of the time of the day. We then show that a simple, more flexible pricing rule can generate even larger welfare gains for all agents. Our results highlight how and why centralized matching and surge pricing are able to make the market more efficient. We conclude by drawing policy implications for improving the competitiveness between taxis and ridesharing platforms.

Keywords: *ridesharing, surge pricing, centralized matching, dynamic games, search frictions*

JEL Classification: *L91, R41, D43, C73*

*Columbia University, Department of Economics. Email: mra2165@columbia.edu.

[†]Columbia University, Department of Economics. Email: vi2137@columbia.edu.

Acknowledgements: We are extremely grateful to Bernard Salanié, Don Davis, Suresh Naidu, Cailin Slattery, Tobias Salz and Michael Best for their valuable advice and support. We thank Yining Zhu for her initial contributions on this project. We are also grateful to Guy Aridor, Iain Bamford, Maggie Shi, Dilip Ravindran, Shoumitro Chatterjee, Anjali Priya Verma and participants at Columbia's Applied Micro Theory Colloquium, IO Colloquium, Trade Colloquium, Delhi School of Economics and UEA 2020 for their feedback. All errors are our own.

1 Introduction

Many decentralized markets are inefficient due to the presence of search and match frictions. Transport, labor and housing markets are well known examples of this. As a result, intermediaries exist in many of these markets and arguably improve efficiency. They primarily achieve this through two mechanisms: facilitating easier matching of buyers and sellers and by organizing the pricing structure in the market.¹ However there is relatively little quantitative evidence on the efficiency gains induced by these matching and price mechanisms as well as the specific role that they play in improving market outcomes.

This paper studies the case of Ride Austin, a ridesharing company in Austin, Texas to understand the role and relative contribution of matching and pricing in solving different market inefficiencies and improving market outcomes. In transportation markets, these inefficiencies broadly take two forms. Static inefficiencies arise when riders are unable to match with drivers at a particular point in time due to search frictions. We measure this as the proportion of rides requested which remained unfulfilled. Dynamic inefficiencies arise when a suboptimal spatial distribution of supply is induced by riders choosing to travel to relatively low demand areas. Similarly, the search, entry and exit decisions of drivers may also distort the spatial distribution of supply and induce further dynamic inefficiencies. We measure this dynamic inefficiency by the time drivers spend searching between rides.

In the particular context of ridesharing, the two instruments used by intermediaries such as Uber, Lyft and RideAustin are surge pricing and centralized matching. Each of these two features is important in overcoming the inefficiencies in the market albeit for different reasons. Surge (or dynamic) pricing allows fares to vary with market conditions, unlike the typical fixed rates of taxis.² As a result, it may play a crucial role in mitigating both static and dynamic inefficiencies by aiming to create a balanced pattern of demand and supply. Higher prices incentivize drivers to enter the market as well as to target their search towards specific locations while also allocating rides to consumers who value it the most. In two sided markets such as ridesharing, prices are crucial in ensuring participation by both sides of the market.³ Centralized matching connects riders and drivers through an app – without these two agents having to physically interact. Centralized matching potentially mitigates static inefficiencies by reducing the search frictions faced by riders and drivers. This likely induces efficiency gains over taxi markets, where a match occurs only if a rider and vacant driver are in the same location at the same time.

To quantify how much surge and matching overcome these market inefficiencies and to understand the specific role played by each feature, we estimate a structural model of the ridesharing market. We

¹Intermediaries like eBay, Airbnb and Uber are examples of intermediaries who organize the matching and price structure in existing markets. All these companies provide a platform for buyers and sellers to match with each other. Prices on eBay are determined through auctions; sellers set their own prices on Airbnb; Uber sets prices centrally. [Einav et al. \(2016\)](#) give an overview of these peer-to-peer markets.

²Surge pricing refers to a surge factor or multiplier (greater than 1) that is multiplied to a base fare.

³This channel has often been highlighted as the key factor in the efficiency gains of ridesharing markets. Its proponents emphasize that flexible pricing allows for improved allocations, resulting in higher consumer welfare ([Hall et al., 2015](#)). Opponents of surge pricing, however, argue that it acts as a form of price discrimination against riders ([Dholakia, 2015](#)), while being ineffective at incentivizing drivers ([Diakopoulos, 2015](#); [Lee et al., 2015](#)).

use data from a ridesharing company called RideAustin, which had the largest market share in Austin during our sample period of October 2016 to May 2017.⁴ Using this data we first establish key empirical facts about the role of surge and matching. We show that surge pricing influences drivers to target high demand areas by “chasing” the surge. Moreover, surge is concentrated in downtown Austin on weekends and is less common than one would expect. We also establish that the search frictions faced by drivers in RideAustin are significantly lower compared to the NYC taxi market.

Our model includes four components: (1) dynamically optimizing drivers choosing whether to enter, which location to search in, and whether to exit; (2) stochastic demand that is responsive to price and waiting time; (3) a matching technology that assigns riders to drivers across the city; and (4) a pricing algorithm which sets the surge factor for a location at a point in time. This builds upon [Buchholz \(2020\)](#) and [Fr chet te et al. \(2019\)](#) who model the intensive and extensive margin of labor supply respectively for taxi drivers in NYC. We don’t directly observe the pricing and matching algorithms used by RideAustin and therefore, before estimating the key parameters of our model, we approximate the pricing and matching algorithms used by Ride Austin. Our approximations match the patterns we observe in the data well. We then estimate the demand for rides using an IV approach which gives us estimates for the price and waiting time elasticity of demand. We finally estimate the key parameters governing the supply side of our model using Simulated Method of Moments. Our estimated model is able to match both targeted and non-targeted moments well.

To empirically quantify the welfare contribution of surge and matching as well as uncover the channels through which they solve market inefficiencies, we conduct policy counterfactuals and find three main results.

First, we show that both surge pricing and centralized matching are crucial in increasing consumer surplus, driver profits and reducing the proportion of unmet demand; our measure of static inefficiency. In particular we find that while centralized matching reduces unmet demand and induces welfare gains irrespective of the time of the day, surge pricing is most crucial late at night. Under a system without surge pricing, we find consumer surplus slightly increases (4%) during the afternoon, while it falls dramatically (65%) at night. Surge incentivizes driver entry, especially at night when they highly value their outside option. This in turn allows riders who requested rides to be matched to drivers. Thus while centralized matching reduces static inefficiencies by eliminating match frictions, surge pricing mitigates static inefficiencies through its effects on driver entry. We also find large spatial heterogeneity in consumer surplus in the absence of surge or centralized matching with the suburbs disproportionately worse off relative to downtown Austin. These results lead us to two policy recommendations. Policymakers in many cities who seek to implement seemingly pro-rider policies by banning or capping surge pricing may in fact harm consumers. In contrast, taxis who wish to compete with ridesharing should primarily focus on introducing a matching technology, rather than surge pricing.⁵

⁴In May 2016, the Austin city council passed a council ordinance requiring fingerprint-based background checks for ridesharing drivers. In response, Uber and Lyft ceased operations in the city. RideAustin entered the market and was one of the two major companies in the market till Uber and Lyft returned in May 2017.

⁵Cities like [New Delhi](#) and [Honolulu](#) aimed to regulate the surge pricing feature of ridesharing apps. On the other hand

Second, we assess whether the platform could further increase welfare by changing its pricing or matching rules and how these could mitigate dynamic inefficiencies. We find that higher surge prices reduces consumer surplus but increases driver profits. Moreover, there is somewhat limited scope for improving the matching technology. Match frictions can only be reduced so much, whereas there are many different pricing schemes the platform could use. We study one option, flexible surge, which allows for the surge factor to go below 1. Ridesharing companies tend to use one-sided surge:⁶ it only becomes activated in areas of relatively high demand. Fully flexible prices should also give discounts to places with relatively low demand. We find that this simple pricing strategy significantly increases welfare for all agents. These price discounts induce more riders into the market, which in turn allows existing drivers to be better matched to riders. As a result, flexible prices creates a more balanced pattern of demand and supply across locations thereby reducing the vacant time between rides; which is our measure of dynamic inefficiency.

Finally, we evaluate how our results would change under traditional driver compensation schemes as with Uber and Lyft. RideAustin’s compensation scheme involved taking \$0.99 for each completed ride with the entire fare going to the driver. While the qualitative patterns in a counterfactual with traditional driver compensation schemes are similar, we find lower total welfare relative to our benchmark model i.e. the model with drivers paying a \$1 fee per trip. The results imply that the introduction of such a platform creates a distortion which prevents the market from achieving the socially optimal outcome. This result is not surprising given that a traditional ridesharing platform does not internalize the effects of its compensation scheme on driver payoffs which in turn affects consumer surplus. These results however do not imply that there exist other compensation and pricing schemes that could be implemented which could benefit consumers and drivers relative to our benchmark model.

1.1 Related Literature

Our paper contributes to the empirical literature on transportation, dynamic pricing and intermediaries in markets with search and match frictions.

First, our paper contributes to the growing literature on rideshare/taxi markets. We quantify and decompose the efficiency gains arising from different features of ridesharing companies. We also uncover the margins that each of these features affects and the inefficiencies that they solve. Methodologically, our model builds on [Buchholz \(2020\)](#) and [Fréchette et al. \(2019\)](#) who study the New York City taxi market and conduct counterfactuals to see whether different pricing schemes and an “Uber-like” platform respectively can increase welfare. In particular, [Buchholz \(2020\)](#) develops a structural model with drivers searching the city, but without entry and exit decisions. In contrast, [Fréchette et al. \(2019\)](#) models drivers’ entry and exit decisions, but without strategic search. We extend the [Buchholz \(2020\)](#) model to incorporate both surge pricing and centralized matching, and drivers strategically choosing to enter, exit, and search. Our paper reinforces and extends the findings of [Buchholz \(2020\)](#) and [Fréchette et al.](#)

cities like New York are planning to [introduce surge pricing](#) for traditional taxis to increase competition with Uber and Lyft.

⁶The surge factor in Uber as well never falls below 1 making it flexible upwards only. [Castillo \(2019\)](#)

(2019) by not only by explicitly modeling the features of a ridesharing market but also documenting the role and joint importance of centralized matching and surge pricing in increasing welfare and solving the static and dynamic inefficiencies arising in ridesharing markets.

[Brancaccio et al. \(2020\)](#) study the decentralized dry bulk shipping industry and also document large gains from frictionless matching and optimal taxes. The notion of thick/thin and pooling externalities developed in their paper are similar to the static and dynamic inefficiencies we focus on. A key difference is that we allow entry and exit of drivers (carriers in their setting) which is an important aspect of the ridesharing market. Our paper complements their findings by highlighting the importance of centralized matching and surge pricing in mitigating static inefficiencies and increasing welfare. In particular, surge pricing is able to mitigate static inefficiencies by incentivizing driver entry. Further, we also show that fully flexible pricing is able to solve dynamic inefficiencies in transport markets by creating a balanced pattern of demand and supply across locations.

Another related paper is [Castillo \(2019\)](#), who studies Uber in Houston, Texas and focuses on distributional consequences of surge pricing. Like us, he also finds that consumer welfare falls when surge pricing is removed. We both find that higher surge factors, resulting in higher average prices, increase driver profits. Substantively, our paper differ along two key dimensions. First, his paper focuses on surge pricing and its distributional consequences for high and low income riders, while we focus on the relative contribution and role of centralized matching and surge pricing in solving market inefficiencies. Our results seem to indicate that matching plays an equal if not more important role in driving the efficiency gains of ridesharing platforms. Second, we model drivers as fully rational and forward-thinking, rather than taking a more heuristic approach. In particular, we explicitly account for how surge pricing and matching affects the drivers' value function. In contrast, [Castillo \(2019\)](#) is able to observe when agents log into the app, which allows estimation of short-run and long-run elasticities in response to changes in prices.

Other papers studying various features of ridesharing platforms such as [Ghili and Kumar \(2020\)](#) and [Shapiro \(2018\)](#) focus on the role of density in the success of ridesharing platforms. [Rosaia \(2020\)](#) studies the effect of mergers between competing platforms on traffic congestion in the city. [Liu et al. \(2019\)](#) study the role of centralized matching in increasing efficiency in a decentralized ridesharing market. Another set of papers estimate flexibility of labor supply amongst ridesharing drivers ([Chen et al., 2019](#); [Angrist et al., 2017](#)). Our paper complements these studies by finding large within-day variation in the flexibility of labor supply for drivers; thus highlighting the importance of surge pricing in inducing drivers to enter the market. Finally, our paper complements other studies of ridesharing platforms such as ([Cohen et al., 2016](#)), who estimate consumer surplus, ([Berger et al., 2018](#); [Hall et al., 2018](#)) who study how ridesharing changes traditional transport markets and [Buchholz et al. \(2020\)](#) who estimate consumer valuations of time. Our focus is on developing a structural model of a ridesharing platform, as well as understanding how they have made an existing market more efficient.

Second, our paper contributes to the literature on dynamic pricing in a variety of markets ([Joskow and Wolfram \(2012\)](#); [Wolak \(2011\)](#); [Lazarev \(2013\)](#); [Williams \(2020\)](#)). In addition there is a recent theo-

retical literature studying optimal pricing in ridesharing platforms (Bimpikis et al. (2019), Besbes et al. (2020), Castillo et al. (2017)). Our paper empirically verifies key theoretical properties of optimal pricing in ridesharing networks developed by Bimpikis et al. (2019) and Besbes et al. (2020). The key idea in these papers is that the demand and supply patterns induced by the optimal pricing rule should be “balanced” across locations. The surge factor used by ridesharing companies such as Uber, Lyft and RideAustin are capped below at one; making them only flexible upwards. By modifying the existing pricing rule to allow the surge factor to go below one for areas with relatively low demand, we find large welfare gains for all consumers and drivers. These results show that fully flexible surge pricing is able to solve the dynamic inefficiencies of transport markets by inducing balancing demand and supply across locations, thus ensuring drivers spend less time searching between rides.

Finally, our paper also contributes to the literature on the role and importance of intermediaries in markets with search and match frictions. Gavazza (2016), Hendel et al. (2009) and Salz (2020) study the role of intermediaries in improving market outcomes in the airline, housing and waste markets respectively. Our paper also is related to the literature on two-sided platforms (Rochet and Tirole (2003), Rochet and Tirole (2006)). Einav et al. (2016) provides a detailed overview of peer to peer markets. Farronato and Fradkin (2018) analyzes the welfare impact of AirBnb in the housing market.

The rest of the paper is organized as follows. Section 2 provides an overview of our setting and data, as well as motivating summary statistics. Section 4 details our model of the ridesharing market. We then outline how we estimate this model in Section 5. Finally, we present our results and counterfactuals in Sections 6 and 7.

2 Setting and Data

We study the case of RideAustin, a ridesharing company based in Austin, Texas, for the period of October 2016 to May 2017.

2.1 Background

In May 2016, the Austin city council passed a council ordinance requiring fingerprint-based background checks for ridesharing drivers. In response, Uber and Lyft ceased operations in the city. This led to many new companies entering to fill the gap, with RideAustin becoming one of the two major companies in the market.⁷ A statewide law eliminating the background check requirement passed on May 29, 2017. This superseded Austin’s local rules, after which Uber and Lyft shortly returned to the city.

In terms of app interface and functionality, RideAustin is quite similar to Uber. Like Uber, RideAustin has surge pricing at discrete rates (in increments of 25 cents) and surge was mandatory.⁸ Unlike Uber,

⁷RideAustin claimed that they had approximately 50% market share in Austin’s ridesharing market during the period in which Uber and Lyft were not present. The other major competitor at the time was Fasten, which ceased operations in the U.S. in March 2018.

⁸In contrast, RideAustin’s competitor Fasten had optional surge which allowed a rider to move to the front of the waiting queue by accepting the surged price.

RideAustin did not offer pooled (multi-passenger) rides like UberPOOL or Lyft Shared. Another major difference is that Ride Austin is a 501(c)3 non-profit which only charged drivers a flat \$0.99 fee per ride. The rest of the trip fare went to the drivers.⁹ Uber, on the other hand, charges its a drivers a percentage commission, typically at a rate of 25% of the total fare. This is a major advantage in our setting because we can focus on rider and driver welfare without worrying about the profit maximization problem of the platform.

2.2 Data

We estimate our model using data from RideAustin. The company published their ride level data for June 2016 to April 2017 – a time when both Uber and Lyft were not operational in Austin. The publicly available data includes 1.2 million rides with over 4,000 distinct drivers. For each ride, we observe a unique driver identifier, the total fare, trip distance and surge factor for that trip, as well as the coordinates and times of the pick-up and drop-off. We supplement this with additional data obtained from RideAustin, which also included data on rides over a longer time period (July 2016 to October 2019). The additional data provides the time and driver’s location when they were dispatched to pick up a passenger, allowing us to track how far a driver had to travel to meet the rider. We also obtained the geographic definitions used by the company to set surge factors, as well as the value of surge in each area every 5 minutes for December 2016 to May 2017. A key limitation of our data is that we only observe *matched* rides. This means we do not observe requested trips that go unmatched or that are cancelled on. A larger concern is that we cannot observe drivers while they are not engaged in a trip (i.e. we cannot directly observe how they search while they do not have a passenger), the dispatch coordinates will help us in overcoming this challenge.

For our analysis, we focus on rides occurring between October 2016 and May 2017. This eliminates its start-up period of growth as well as its decline following the return of Uber and Lyft. As Figure (1) shows, this best captures when RideAustin is in equilibrium as a major market player. Figure (A1) shows further restrictions we make to ensure that our data represents a typical week of ridesharing.


3 Summary Statistics

In this section, we present key empirical patterns of demand, supply, prices and matching. We use these patterns in the data to motivate the features of the model. We also discuss how these patterns will help us identify the parameters in our model.

To analyze the data, it is useful to standardize the time and physical space that our data covers. We define a “day” as representing a 24 hour period starting at 4:00 a.m. We begin the day at 4 a.m. rather than midnight for two reasons. First, rides in the early morning are best thought of as late night activity continuing from the previous day’s evening rather than events separate from it. Second, 4:00 a.m. is

⁹The revenue per trip for the driver is $\pi = f(\delta, \tau) \times s - 0.99$, where $f(\delta, \tau)$ is the unsurged fare (a function of the trip’s distance δ and time taken τ) and s is the surge factor. $s = 1$ represents an unsurged ride.

Figure 1: RideAustin Trips Over Time



when the least number of rides within a day occur, which provides a reasonable point to divide the day. Furthermore, we divide the city of Austin into equally-sized hexagons that are approximately 0.9 kilometers wide using Uber’s H3 geospatial indexing system. We choose to divide the city in such a way since a main feature of our model is to characterize notions of *local* demand and supply.

3.1 Spatial and Temporal Distribution of Demand and Surge


We first determine where and when trips and surged rides most often occur in our data. We find that the demand for rides as well as the probability of rides being surged is concentrated in downtown Austin and on weekends. We also find that most rides are unsurged, with surge accounting for a quarter of rides in high activity areas and times.

In Figure (2), we plot the spatial distribution of where trips in our data begin. The outline superimposed on the map indicates the neighborhoods we use to aggregate the hexagons. This approximately matches the administrative boundaries of the city of Austin.¹⁰ Rides occurring outside these neighborhoods are infrequent; 94% of rides start *and* end within our defined neighborhoods, so we restrict all further analysis to those rides. Even within the city, there is still large spatial variation in the frequency of rides. We see that the overwhelmingly majority of trips begin in downtown Austin, the small central neighborhood depicted on the map. This is an area where many bars and restaurants are located. Rides in the outer, more suburban, neighborhoods are far more infrequent.

Next, we plot the spatial distribution of the probability of surge occurring. For each hexagon, we calculate the ratio of the total number of surged rides (rides that had a surge factor greater than 1)

¹⁰See Appendix B for an overview of the geographic units used in the analysis.

Figure 2: Distribution of Trips (by Start Location)




starting in that hexagon to the total number of rides starting from that hexagon. We also calculate this ratio for rides ending in each hexagon and then plot the spatial distributions of these ratios in Figure (3). We find that while most of the rides are unsurged, the probability of surge generally peaks for rides that start around the central downtown area. Rides beginning in the outer suburban areas are rarely surged. In contrast, there is no distinct pattern for where surged trips ends. This suggests that surge is only a function of conditions in your pick-up location, and not your destination.

Figure (3) also shows that surge is not a frequent occurrence. In the downtown neighborhood, where most rides (and most surged rides) occur, surge occurs in only 23% of rides. The relative infrequency of surge pricing is highlighted in Table (1) which plots the frequency distribution of the surge factor. We find that 86% of rides in our data have a surge factor of 1, i.e. they are unsurged. The remaining rides have a surge factor associated with them, most of these are the two lowest factors of 1.25 and 1.50. When focusing on Fridays and Saturdays we find that 78% of rides are unsurged. These statistics are similar to Castillo (2019), who finds that 77% of Uber rides in Houston are unsurged.¹¹ On Fridays and Saturdays, the highest activity days of the week, unsurged rides fall to 77% in our data. Therefore, while surge pricing is an often-discussed aspect of ridesharing, in reality, the majority of rides go unsurged.

We next look at the temporal distribution of rides and the surge factor. Panels (a) and (b) in Figure (4) show the distribution of rides and the probability of surged rides across a typical week. We find

¹¹In that context, Uber had a monopoly of the Houston ridesharing market. An earlier version of this paper used Uber data from San Francisco and he finds that 86% of Uber rides in San Francisco were unsurged.

Figure 3: Distribution of Surge




Notes: The left map shows the fraction of surged rides by trip starting location. The right map shows the fraction of surged rides by trip ending location. Hexagons that have less than 20 total rides over the sample period are not shaded (plotted in gray).

that most of our rides are concentrated on late nights during Fridays and Saturdays. In particular we find that rides are most likely to occur on Friday and Saturday nights with a sharp drop-off after 2 a.m., which is when bars in Austin are stipulated to close. Even though the closing time is a completely anticipated event, we find that rides are most likely to be surged at 2 a.m. on Friday nights and Saturday nights with nearly 60-70% of observed rides being surged in this time period. Rides are also likely to be surged around the evening, which corresponds to the evening post-work traffic rush on weekdays. These stylized facts imply that accounting for intra-day variations in demand and supply is important when modeling and estimating the ridesharing market.

3.2 Match Frictions in Ridesharing


Search frictions in a ridesharing market are relatively low. As a benchmark, we show that frictions in our setting are significantly lower compared to the New York City taxi market. Figure (5) shows that drivers' vacancy rate (time they spend searching for a ride as a percentage of total time spent driving) is relatively flat throughout the day, hovering around 30 to 35%.

Figure 4: Distribution of Rides and Surge Across and Within Days



Notes: In panel (a), each cell represents the proportion of rides occurring in that day-hour relative to all rides in the sample. In panel (b), for each cell we plot the proportion of observed rides in that day-hour that are surged.

Figure 5: Vacancy Rate Within Day



Notes: This plot shows the average ratio of search time to total driving time. Search time is equal to times between a previous ride’s drop-off and the subsequent ride’s dispatch. Driving time is time spent on a trip as well as driving to pick up the passenger. Data for NYC Taxis plot from [Fr chet te et al. \(2019\)](#).

This figure replicates a similar statistic calculated by [Fr chet te et al. \(2019\)](#) for New York City taxis, with the data from their analysis also plotted on the graph. We see that search is fundamentally different in ridesharing in two ways. First, the vacancy rate in ridesharing exhibits far less within-day variation. Second, the vacancy rate is lower at any time of the day. This suggests that search and match frictions

Table 1: Surge Frequency

Surge	Fraction of Rides		
	All	Fri.-Sat.	Sun.-Thurs.
1.00	86.02	77.31	93.6
1.25	5.32	7.77	3.18
1.50	4.17	6.76	1.91
1.75	1.5	2.59	0.56
2.00	0.96	1.65	0.35
2.25	0.31	0.59	0.07
2.50	0.47	0.88	0.12
2.75	0.22	0.43	0.04
≥ 3.00	1.03	2.03	0.16
<i>N</i>	1,403,783	653,025	750,758

are far lower in ridesharing markets, which is expected given that participants are matched using a centralized dispatcher. It also further highlights the importance of exit as ridesharing drivers can more easily exit the market rather than engage in costly search.

While match frictions are lower in ridesharing, these gains are not uniform across the city. In Figure (6), we show the mean search time for drivers based on the neighborhood of their last drop-off. We see that following rides that end in central areas, drivers spend an average of 11 minutes searching for their next ride. In contrast, rides that end in the non-central areas require drivers to search longer until their next ride. These highlight the presence of dynamic inefficiencies. Riders do not internalize that their destination choices can result in a spatial misallocation of drivers, who must then spend more time driving in order to find their next passenger.¹² Figure (C7) in the Appendix shows that these frictions exist from the viewpoint of riders as well. In particular, riders face a longer waiting time in non-central areas relative to central areas.

3.3 Evidence of Strategic Search, Entry, and Exit Decisions for Drivers


Finally we highlight the various factors that influence the strategic decisions of drivers. In particular we show that drivers strategically target high demand areas and that they tend to “chase” the surge.

3.3.1 Entry

Drivers in our data average only 3-6 rides per day over our sample period (this translates to working for about one hour). However, the actual workload varies each day with the modal number of trips being 1 ride per day. Figure (C5) shows these distributions. Unlike taxicabs, whose drivers tend to work longer and fixed shifts, the nature of ridesharing as being part of the ‘gig economy’ necessitates a focus on entry and exit.

¹²Brancaccio et al. (2020) refer to this inefficiency as “pooling externalities”.

Figure 6: Search Time By Drop-off Location



Notes: This plot shows the mean search time by neighborhood. Search time is equal to time (in minutes) between a previous ride's drop-off and the subsequent ride's dispatch. Neighborhood is the previous ride's drop-off area.


Figure (4a) shows that there is not only variation in the number of rides across days, but intra-day variation as well which is important to acknowledge. Figure (7) plots the distribution of hour of entry (i.e. what time do drivers begin working?) along with the average number of trips drivers complete given their entry hour.¹³ For example, the graph shows that 6.4% of drivers enter the market at 7pm and that drivers who enter at 7pm complete an average of 5 rides. It is clear from the figure that drivers' workload peak in the evening, and correspondingly, that is also when we observe the highest rates of entry during the day. This suggests that drivers are more likely to enter the market when they anticipate higher earnings.

3.3.2 Search and Exit

While centralized matching reduces search frictions, the search strategy of drivers still plays an important role in the rides they get and, consequently, their earnings. A driver's search strategy will be


¹³We cannot observe when the driver logs on or off the app (i.e. formally enters or exits the market). We assume that a driver enters the market 15 minutes before their first trip (as this is the mean time for drivers searching between trips, which we can observe). We also assume that a driver exits if we do not observe them on a trip for 60 minutes, and take the exit time as the drop-off time of their last trip. Figure (C11) shows the distribution of searching time between rides.

Figure 7: Entry and Rides Within Day (Fridays and Saturdays)



affected by not only where they expect higher earnings, but also the geography of the space they search. Ideally, if we could observe drivers at all times of the day, we could see how they choose to search. However, as mentioned before, our dataset limits us to only when drivers are involved in a trip. Figure (8) illustrates a driver’s shift and shows what we can observe. Note that we are able to see where a driver was at the time they were assigned (dispatched) to a passenger. By comparing this to where the driver’s *previous* ride ended, we know where the driver started searching from. The dispatch coordinates then gives us a snapshot of where the driver was on their search path at the time they were assigned their next ride. For example, in Figure (8), this strategy would correspond to comparing the coordinates of D_2 to the coordinates of E_1 .¹⁴

Figure 8: Example Driver’s Shift




Notes: The diagram shows a driver with two rides during their shift. D_i , S_i , and E_i are the dispatch, start, and end times for trip i . These three times (and the corresponding coordinates) are observable in our dataset.

To identify the drivers’ search strategy, we would ideally assign randomly a driver to a location on the map and observe how they search starting from their assigned spot. As drivers do not know the destination of the trip prior to accepting their assignment, the drop-off locations provide a close

¹⁴In Figure (8), the driver searches three times during their shift. However, only the search that occurs between two trips is used in this analysis, as it is the only one where we observe the start and end.

Figure 9: Search Path by Drop-off Location (Fridays and Saturdays)




analog to this experiment. For a given neighborhood where a drop-off occurs, we calculate the average coordinate of the driver's *subsequent* dispatch location. We do this for different bins of search time (the time between the driver's drop-off and their next dispatch). By piecing together these individual snapshots, we are able to trace out the average search path for drivers based on their start location. This is shown in Figure (9). We can see that as drivers have more time to search, they aim to get closer to the downtown area, which is both where most rides occur and is the most central location. This suggests that drivers are indeed strategic in their search and that there are some commonly held strategies.

Another implication of Figure (9) is that drivers perceive that non-central suburban areas have low earnings potential (since they do not spend time searching there). We should then expect higher rates of exits in these areas, i.e. when a driver completes a trip in a non-central area, we would expect that they would be less likely to continue working. As Figure (C8) shows, this is exactly what we find in the data. Table (C8) supplements this by showing exit rates of different regions broken down by the start location of the trip. Regardless of where the trip began, exit is more likely if the trip ends in a more non-central location. This highlights the spatial nature of drivers' labor supply decisions.

Finally, we also find that drivers target areas where they expect higher earnings, in particular, that

Figure 10: Pull Distance by Search Time



drivers appear to be “chasing the surge”. If drivers are targeting particular areas, then we should expect the distance that they are “pulled” to meet their assigned rider to be smaller.¹⁵ In Figure (10), we plot the average pull distance (how far a driver travels from their dispatch location to the trip’s start location) by how much time the driver has had to search since their previous trip was completed. This is important as drivers who get dispatched to their next ride soon after dropping off a passenger will not have enough time to strategically search.¹⁶ We calculate this under two scenarios: if the ride they are dispatched to is surged and if it unsurged. The patterns observed in this figure allow us to estimate how much of the driver behavior in our model is explained by continuation values of strategically searching in different locations.

We find that drivers are generally pulled from fairly short distances, with average distances ranging from 1 to 1.5 kilometers. However, this does show that it is difficult for drivers to perfectly predict where demand is going to occur (otherwise the pull distance would have been close to zero). We also see that the pull distance for rides soon after the dispatch (≤ 2 minutes) are generally longer than when the driver has more time to search. Lastly, when the subsequent ride is surged, drivers have much lower pull distances that fall rapidly in the first 10 minutes of search time. This evidence is consistent with drivers chasing the surge (if they have enough time). It also suggests that surge plays a role in helping drivers know where to target their search efforts.¹⁷

3.4 Takeaways

In our model of supply, the key parameters to be estimated will be the variances of idiosyncratic shocks relating to searching different locations, exiting, and entering, as well as the value of the driver’s outside

¹⁵In Figure (8), this is travel from D_i to E_i .

¹⁶Using the example in Figure (8), this would be plotting the distance from D_2 to S_2 on the y-axis and the difference in time from E_1 to D_2 on the x-axis.

¹⁷Like Uber, drivers for RideAustin are notified on the app which areas are currently experiencing surge pricing.

option. The patterns we find in the data shed insights into how we aim to identify the parameters.

The driver’s observed search strategy determines the relative importance of location fundamentals and idiosyncratic shocks. Figure (9) shows that, regardless of their starting locations, drivers exhibit a similar strategy of aiming to reach the same location: downtown Austin. As Figures (2) and (3) show, this is also the location of the highest demand and the highest rates of surged rides. Similarly, Figure (10) also shows that drivers try to position themselves closer to surged rides. These statistics suggest that drivers are incentivized by higher fares, and that their search is targeted at earning these higher fares.

As we illustrate in Figure (8), a data challenge we face is that we do not observe what drivers do outside of a trip. In our estimation, we leverage that the dispatch location and time provides a snapshot to the driver’s search activity. First, the pull distance (difference between dispatch and start locations) tells us how close the driver was to the rider. This logic, which we use to interpret Figure (10), highlights that the pull distance can be used as a measure of how strategic drivers are. Second, in Figure (9), we observe what drivers do when they have more time to search. If drivers have longer search times, this implies that their search is not targeted towards finding passengers.

Finally, we observe that drivers tend to complete only a few rides per day. We therefore develop a model that incorporates entry and exit of drivers. The entry and exit rates will be informative in identifying the outside option parameter, as this will show whether drivers prefer to engage in ridesharing over their alternate activity. Figure (7) shows that drivers are more likely to enter at higher activity times of the day, when they expect higher earnings. Figure (C8) shows that there is spatial variation in the exit rates of drivers; in particular, we see higher exits if they are forced to take riders to low-value areas. In our model, the geographical variation in the entry of drivers will also be used to understand how much of the entry behavior we observe is driven by model fundamentals versus driven by unobservables.

4 Model

Our model incorporates four key features of ridesharing markets: (1) demand for rides that depends on price and waiting time, (2) supply of drivers where drivers strategically choose to search, enter or exit, (3) a matching algorithm, and (4) a surge pricing algorithm. We extend [Buchholz \(2020\)](#), which models the taxi market; and incorporate the features of the ridesharing market.

In the demand side of our model, at every point in time, riders (r) in every location i open the app and decide whether or not to request a ride after observing the price and expected waiting time. On the supply side, drivers who have already decided to enter are assigned to riders to drop them off to their desired locations; if they are unassigned, they can then choose to search the city (by driving to a different location) or exit the market. Drivers currently not in the market choose whether or not to enter the market. Additionally, our model has a matching technology that assigns drivers to riders. This technology approximates a “first-dispatch protocol” that RideAustin uses, where the rider is assigned to the nearest available driver. Finally, our model includes an endogenous surge multiplier that is observable

by all agents. In the following sections, we outline each of these parts of our model in more detail.

We assume that the city is a network of L locations given by $\mathcal{J} = \{1, 2, \dots, L\}$. As we are analyzing the ridesharing market within a day, we assume that time is discrete and that the time horizon is finite i.e. $t \in \{1, 2, \dots, T\}$. The distance between any two locations i and j is denoted by δ_{ij} and the time taken to travel between locations is given by τ_{ij} . There are two types of agents in the market: riders (r) and drivers (v). At a location i in time t , there are r_i^t riders who request a ride and v_i^t drivers. We denote the vector of riders and drivers in locations across the city at time t as r^t and v^t , respectively.

4.1 Demand for Rides

In our model, an exogenous pool of potential riders O_i^t open the app and observe the price (surge factor) and expected waiting time (time for a driver to arrive at the pick-up location). Based on these factors, some customers will choose to request a ride. Among those who request, if a rider is matched to a driver, they will be picked up and have their trip completed. If a rider is unmatched, then with some probability they will re-request in the next period, or otherwise exit the market.

To implement this, we assume a reduced form stochastic demand model in each location-time cell. This results in a realized demand after observing the base price for a given origin-destination pair. We assume demand is of the form:

$$\log(\tilde{r}_i^t) = \alpha_i^t + \theta_1 \log s_i^t + \theta_2 \log w_i^t + \theta_3 \log O_i^t + \varepsilon_i^t \quad (1)$$

where \tilde{r}_i^t is the number of new riders requesting in location i at time t , s_i^t is the surge factor and w_i^t is the average waiting time. O_i^t is the number of people opening the app, and represents the (exogenous) potential pool of riders.

From the \tilde{r}_i^t people who request rides, the app will match as many of them as possible to vacant drivers (this process will be described in the next section). Those who do not have their requests fulfilled will either exit the market (with probability ω) or request again in $t + 1$ (with probability $1 - \omega$). If they request again in $t + 1$ and still find themselves without a match, they will either exit the market or request again in $t + 2$, with the same probabilities as before. This process repeats until a maximum waiting period of G . This means that at time t there will be riders who initially requested their ride in periods $t - G$ to $t - 1$. We will denote the number of re-requesters as \bar{r}_i^t . Therefore, the total number of requested rides in location i at time t , denoted as r_i^t , is the sum of \tilde{r}_i^t (new requests) and \bar{r}_i^t (re-requesters).

4.2 Matching

Ridesharing platforms act as a centralized dispatcher who assigns available drivers to customers requesting a ride. In a traditional taxi market, e.g. as studied by [Buchholz \(2020\)](#), drivers and riders must be in the same location to be able to find each other and create a match. However, ridesharing apps allow for matching of drivers and riders who are in completely different locations.

In this section, we describe the matching algorithm to assign riders to drivers. We develop an algorithm that approximates a “first-dispatch protocol” wherein riders are matched to nearest available driver (Yan et al., 2019). We show in Appendix D.1 that the patterns we observe in the data justify this approach.

As denoted above, the vector of riders and drivers locations across the city as r^t and v^t . The matching algorithm $m(\cdot, \cdot)$ takes in as inputs the vector of riders and drivers (r^t, v^t) and outputs a square matrix P of *pull probabilities* from each location to every other location. Each entry p_{ij} is the probability that a driver in location i is matched to a rider in location j (we refer to this process as “pulling”).

$$P^t = m(\{r_1^t, \dots, r_L^t\}, \{v_1^t, \dots, v_L^t\}) \quad (2)$$


The matching algorithm focuses on pulling drivers to the locations where they are needed. Within a location, we randomly assign drivers to riders. However, we give priority to riders who have waited longer (i.e. we first assign drivers to riders who have waited G periods, then $G - 1$, and so on).

Next, we show how we recover the P matrix via simulation and properties of the matching algorithm.

4.2.1 Implementing the Algorithm

Recall that we have divided the city of Austin into equally sized hexagons (Figure 1). Each of these hexagons correspond to a location i in our model. For each hexagon i , we define its local neighborhoods as follows. We denote the six adjacent hexagons as Ring 1 of that hexagon i . The 12 hexagons adjacent to the Ring 1 hexagons are known as Ring 2 and so on. This can be seen in Figure (11) where we show the different rings of a single hexagon. While this figure denotes this for just one hexagon in the city, we create these “rings” for every hexagon in our grid. This means that for any two hexagons $i, j \in \mathcal{J}$, we can calculate their ring relationship $Ring(i, j) = Ring(j, i) \in \mathbb{Z}_{\geq 0}$.

Figure 11: Hexagon System



Next, to obtain the matrix P , we use a matching algorithm that emulates the first-dispatch protocol. There are two aspects to the matching algorithm. First, it is frictionless: if a driver is matched to a rider by the central dispatcher, the match will occur and the trip will be completed with certainty. This can

be represented as a Leontief matching function. Second, it involves stages of pulling in drivers. To give more detail, in each location i , the algorithm initially matches as many drivers to riders who are both in that location ($\min\{r_i^t, v_i^t\}$). After this step, locations will either have excess demand (unmatched riders) or excess supply (vacant drivers) or neither. Next, the algorithm takes all locations i that still have excess demand and matches riders in i randomly to vacant drivers in j , where j is a location in Ring 1 of i . After this, excess demand and supply counts update and the process repeats with Ring 2. We continue the process until excess demand is satisfied everywhere or supply runs out. As the pulling in involves randomness, the matching is a stochastic process. For a given distribution of riders and drivers (r^t, v^t) , we average observed matches from simulations of the matching algorithm to calculate the pull matrix \mathbf{p} .¹⁸

The first three stages of the matching procedure is illustrated in Figure (12) for two hexagons highlighted in yellow. First, all demand within the yellow hexagons are met using a Leontief matching function. If there is still unmet demand in those hexagons, then we match drivers in Ring 1 to the remaining riders in the yellow hexes. This can be seen in panel (b) of Figure (12) as pulling from the blue hexagons. If there is still unmet demand in the yellow hexagons, we move on to Ring 2 as denoted in panel (c) of Figure (12) and match available drivers to remaining riders.¹⁹

Figure 12: Matching Procedure




Figure 3 (a)

Figure 3 (b)

Figure 3 (c)

4.2.2 Properties of the Matching Algorithm

In this section, we examine some properties of our simulated matching function. First we show that as the market thickness increases, matching becomes more efficient.²⁰ Second, we show that drivers in


¹⁸This matching function differs from that in [Castillo \(2019\)](#) in several ways. His matching function *sequentially* matches riders to the closest driver according a random order, while ours function matches *simultaneously* at a city-wide level. This means that our matching reduces the possibility of mismatch (i.e. that a driver could have instead been assigned to a closer rider). In addition, we are able to generate explicit pull probabilities, which can explicitly enter the driver's value function.

¹⁹For illustrative purposes, this figure only shows the matching for two hexagons. The full algorithm is done simultaneously for all hexagons in every time period t .

²⁰Market thickness is captured by the number of agents in the market (i.e. given the fixed geography, an increase in the number of agents results in a denser market).

our data seem to be strategically choosing where to locate themselves. To do this, we simulate how the outcome of the matching algorithm changes as the market becomes thicker. In each of our simulations, we keep the ratio of riders to drivers the same, i.e. we scale up demand and supply proportionately.

Figure 13: Matching Function Properties



Notes: The graph shows three simulations where the distribution of riders (demand) and drivers (supply) are varied. The red line distributes demand according to the empirical probabilities for each hexagon, while supply is distributed uniformly across the hexagons. The blue line uses the empirical distribution for both demand and supply. The green line distributes both demand and supply uniformly.

In Figure (13) on the x-axis, we plot total demand across the city while simultaneously increasing supply so that the resulting is ratio fixed at 1/3. On the y-axis, we plot the average pull distance as a result of the matching. We consider three different distributions of total demand and total supply. First, we distribute demand and supply across locations as we observe in the data. Second, we consider a uniform distribution of demand and supply over locations. Finally, we consider a scenario where demand is distributed as observed in the data, but supply is distributed uniformly.

We see that as market thickness increases, the matches that occur are more likely to involve lower pulling distances. This is expected as a thicker market increases the chance to find a driver closer to the rider’s location. However, the spatial distribution of riders and drivers is also an important determinant in the efficiency of the matching algorithm, i.e. the average pull distance of drivers. When demand and supply are equally distributed across space, the matching algorithm is most efficient. As the demand-to-supply ratio is below 1, all demand will (in expectation) be satisfied by supply in the same location. Matching is least efficient when demand is distributed as in the data, but supply is distributed uniformly across space. In this situation, even as the market thickness increases, the matching is never fully efficient as the pull distance remains above zero. This is driven by areas of high demand that continue to have

excess demand even as the market gets thicker. The case when drivers and riders are distributed as in the data falls in between the two prior scenarios. This is consistent with drivers strategically searching in areas with high demand. It also shows how difficult it is to perfectly predict demand since the matching is not as efficient as in the case of uniformly distributed demand and supply. However, as the market becomes thicker, the difference between the empirically and uniformly distributed cases vanishes. In a sufficiently thick market, even if an individual driver mis-strategizes where to search, we should not expect this to have any aggregate efficiency implications.

4.3 Supply of Rides

We model the supply side of this market with drivers dynamically deciding entry, search and exit decisions. Drivers decide whether to enter the market and thus make themselves available for trips. Once in the market, they can search around the city. If a driver gets matched to a rider, they first go to pick-up the rider, then complete the trip and receive the trip's fare. If they are unmatched, they can continue searching or exit the market. We first define the payoffs from completing a single trip from location k to j . After that, we describe our model of search and exit followed by our model of entry. To contrast it to the previous literature, our supply side model extends [Buchholz \(2020\)](#) to account for surge pricing, centralized matching as well as allowing for drivers to enter and exit on top allowing them to search across locations.

4.3.1 Payoff per Trip

In this section, we outline how the trip fare is calculated. A trip at time t is defined as a rider in location k wanting to go to j who has been matched to a driver currently in i . We set the trip fare using the same formula used by RideAustin. The fare that a driver gets depends on: the flag-drop fare b , the distance-based fare π_δ , the distance between the pick-up and drop-off locations denoted by δ_{kj} , the time-based fare π_τ , the travel time between the pick-up and drop-off locations denoted by τ_{kj} , as well as the surge factor s_k^t applicable for that trip (which varies with the origin location k and request time t). The total revenue that a driver earns from the trip as defined above is $(b + \pi_\delta \delta_{kj} + \pi_\tau \tau_{kj}) s_k^t$.²¹ The costs incurred by the driver involves fuel costs in driving from i to pick up the rider in k , denoted by c_{ik} as well as the fuel costs involved in dropping off the rider in j (c_{jk}).²² The net payoff per ride for a driver is therefore:

$$\Pi_{ikj} = (b + \pi_\delta \delta_{kj} + \pi_\tau \tau_{kj}) s_k^t - c_{ik} - c_{jk} \quad (3)$$

²¹RideAustin additionally ensures that the minimum fare is \$5, which we also impose in the estimation. Per the company's policy, the driver keeps the entire fare from the trip net of a flat \$0.99 fee which the company keeps from each trip to cover operating costs.

²²This payoff follows the structure of [Buchholz \(2020\)](#), who sets the fuel cost at 12.4 cents per mile for New York City taxis. [Castillo \(2019\)](#) sets driving costs for Uber drivers at 26 cents per mile, though this includes fuel as well as maintenance, repairs, and depreciation.

4.3.2 Search and Exit

In this section, we focus on drivers who are already in the market. Drivers who have already entered the market can either be *vacant* (available for matching) or *in transit* (engaged in a trip). We denote the state variables for a driver w who has already chosen to enter in location i at time t as $x_w^t = (l_w^t, e_w^t, \mathcal{S}^t)$, where l_w^t is the location of driver w at time t , e_w^t is an indicator for whether driver w at time t is on a trip, and \mathcal{S}^t is the spatial distribution of demand and vacant drivers at t . This is defined as $\mathcal{S}^t = \{v_i^t, r_i^t\}_{i \in \mathcal{J}}$, where v_i^t is the number of vacant drivers in location i and r_i^t is the total demand at location i . Drivers who are in transit ($e_w^t = 1$) have no decisions to make until their current trip is complete. Therefore, to model the decisions of drivers, we only need to focus on vacant drivers ($e_w^t = 0$).

Vacant drivers must decide which location to drive to (search) or to stop working (exit). The set of locations that they can choose to search from is denoted by $\mathcal{A}(i)$. This decision is based on their current location (l_w^t) as well the distribution of demand and other vacant drivers (\mathcal{S}^t). Drivers will search by driving to the location that provides them the highest value. The value of a location for a vacant driver can be broken up into two sources.

1. With some probability, the driver gets dispatched from that location to a trip. The driver then drives to pick-up the passenger and drops them off in their destination, thus receiving the trip fare. Then they continue as a vacant driver from the drop-off location.
2. Alternatively, the driver remains unmatched and can choose a new location to search (the set of available locations denoted by $\mathcal{A}(i)$) or they can choose to exit.

Given the spatial distribution of drivers and riders, \mathcal{S}^t , the *ex-ante* value for each vacant driver in location i at time t is given by:²³


$$V_i^t(\mathcal{S}^t) = \mathbb{E}_{\mathcal{S}|\mathcal{S}^t} \left[\sum_{k \in \mathcal{J}} p_{ik}^t \sum_{j \in \mathcal{J}} M_{kj}^t \left(\Pi_{ikj} + V_j^{t+\tau_{ikj}}(\mathcal{S}^{t+\tau_{ikj}}) \right) + \left(1 - \sum_{k \in \mathcal{J}} p_{ik}^t \right) \mathbb{E}_\varepsilon \left[\max_{j \in \mathcal{A}(i) \cup e} \left\{ V_j^{t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}}) - c_{ij} + \varepsilon_{ij} \right\} \right] \right] \quad (4)$$

The first part of the value function captures the value for a driver located in i if they are matched by the platform to a rider traveling from k to j . p_{ik}^t is the probability that a driver in location i will be pulled in to location k which is derived from the matching function described in the previous section.²⁴ M_{kj}^t is the probability that a rider in k wants to go to location j at time t . Π_{ikj} is the net payoff for the driver, as shown in Equation (3), of starting in location i , picking up the passenger in k , and dropping them off in j . This driver will complete their assignment at time $t + \tau_{ikj}$, where τ_{ikj} is the number of periods to go from i to k and then k to j , i.e. $\tau_{ikj} = \tau_{ik} + \tau_{kj}$. $V_j^{t+\tau_{ikj}}(\mathcal{S}^{t+\tau_{ikj}})$ is the continuation value of the vacant

²³The expectation is taken over the distribution of riders and drivers in future periods.

²⁴We assume that the matching function does not take into account drivers who are about to end trips in that region. Additionally, we also assume that drivers cannot be matched en-route to a location.

Figure 14: Vacant Driver Valuation



driver who is at location j at time $t + \tau_{ikj}$. By modeling it this way, drivers in our model will not only want to locate themselves in areas which have high demand, but also in areas which have high demand in neighboring areas as well.

The second part of the equation represents the value of being unmatched, which happens with probability $(1 - \sum_{k \in \mathcal{J}} p_{ik}^t)$. If this happens, a driver can choose to search in any of the locations accessible to them or can choose to exit. $A(i)$ is the set of all locations in \mathcal{J} accessible from i and e is the option to exit. Driving to location j from location i incurs a fuel cost c_{ij} and takes time τ_{ij} . $V_j^{t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}})$ is the continuation value of the vacant driver who is at location j at time $t + \tau_{ij}$. ϵ_{ij} is drawn from a Type-1 Extreme Value (EV) distribution with scale parameter σ_ϵ . This shock captures preferences of drivers for a particular location. We assume that exiting is permanent and that the value of exiting is $V_e^t(\mathcal{S}^t) = \mu(T - t) = \mu_t$, where μ is an outside option that accumulates each period. The fuel cost and travel time for exiting are both zero, $c_{ie} = \tau_{ie} = 0$. The model for drivers is summarized in Figure (14).

The driver's problem highlights the two sources of inefficiencies present in the market. At the first step, there is a probability that a driver is matched to a rider. In Equation (4), this is represented by $\sum_{k \in \mathcal{J}} p_{ik}^t$. This represents the static inefficiency: the presence of frictions preventing riders and drivers matching from one another. As the probability of being pulled increases, this is associated with a decrease in the static inefficiency. Note that under centralized matching, $p_{ik}^t > 0$ for some $k \neq i$. Without centralized matching, it would be equal to $p_{ik}^t = 0$ for all $k \neq i$ as the driver and rider are not in the same location to be able to find each other. If the driver does get matched, then the rider chooses the destination. This second step is what creates the dynamic inefficiency: riders do not internalize their destination choice, which may send drivers to low-value locations. In Equation (4), this is dictated by the term M_{kj}^t , which represents drivers' trip preferences. While a vacant driver is able to choose their next location, a matched driver has no choice in where the rider will take them and will be forced to continue their search at a potentially sub-optimal location.

At the end of each period, vacant drivers decide where to search or whether to exit the market. This

means that a driver in location i solves the following problem:

$$j^* = \operatorname{argmax}_{j \in A(i) \cup e} \left\{ V_j^{t+\tau_{ij}} (\mathcal{S}^{t+\tau_{ij}}) - c_{ij} + \varepsilon_{ij} \right\} \quad (5)$$

Define the *ex-ante* choice specific value function, conditional on taking action j , as:

$$W_i^t(j, \mathcal{S}^t) = \mathbb{E}_{\mathcal{S} | \mathcal{S}^t} \left[V_j^{t+\tau_{ij}} (\mathcal{S}^{t+\tau_{ij}}) - c_{ij} \right] \quad (6)$$

Given our assumption that ε_{ij} is i.i.d. Type 1 EV with scale parameter σ_ε , the probability of choosing $j \in A(i) \cup e$, before observing the draw of ε , takes a simple analytical form given by:

$$Pr_i^t(j | \mathcal{S}^t) = \frac{\exp(W_i^t(j, \mathcal{S}^t) / \sigma_\varepsilon)}{\sum_{k \in A(i) \cup e} \exp(W_i^t(k, \mathcal{S}^t) / \sigma_\varepsilon)} \quad (7)$$

This captures the policy functions of the drivers, which we denote as $\sigma_i^t = \left\{ Pr_i^t(j | \mathcal{S}^t)_{j \in \mathcal{J}} \right\}$. Time ends at T and if a driver has chosen not to exit till then, the continuation value is set as $V_i^t = 0$ for all $i \in \mathcal{J}$ and $t > T$.

4.3.3 Entry

Next we consider the problem of a driver deciding whether or not to enter. We assume that in time t , at each location i , an exogenous number of drivers consider entering the market, denoted as \bar{V}_i^t . Each potential driver decides whether to enter or not. If they enter, they immediately become available to be matched to riders. If they do not enter, they receive μ_t for the rest of day.²⁵ The potential driver obtains draws η_i^t from a Type 1 EV distribution with scale parameter σ_η for both options in their choice set: entering the market in location i at time t or staying out for the remainder of the day. Given this setup, the utility from entering is given by:

$$V_{i1}^t = V_i^t(\mathcal{S}^t) + \eta_{i1}^t$$

The utility of the driver's outside option (staying out of the market) is given by :

$$V_{i0}^t = \mu_t + \eta_{i0}^t$$

Therefore a driver in location i solves the following problem:

$$j^* = \operatorname{argmax}_{j \in \{0,1\}} \left\{ V_{ij}^t \right\} \quad (8)$$

²⁵This setup restricts drivers to only one opportunity to enter the market. When estimating the model, we divide up the day into 3-hour shifts, which means that the driver is restricted to making one entry decision within a 3 hour period (however they could still enter in different shifts).

Given the distributional assumption on the error term, this implies that the probability with which a driver enters in location i at time t can be expressed as:

$$q_i^t(\mathcal{S}^t) = \frac{\exp(V_i^t(\mathcal{S}^t)/\sigma_\eta)}{\exp(V_i^t(\mathcal{S}^t)/\sigma_\eta) + \exp(\mu_t/\sigma_\eta)} \quad (9)$$

4.4 Surge Pricing Algorithm

The final piece of the model is surge pricing. Surge pricing is a vector of prices for each location-time, that are a function of local market conditions. To implement this, we need to specify how surge prices are set. While RideAustin’s exact pricing algorithm is unknown to us, the company provided us with all the data which serves as an input into their surge pricing algorithm as well as the spatial and temporal level at which surge is set. Moreover, company representatives told us that RideAustin uses some function of demand and supply at each time and location to decide on the surge factor.

We use this information to approximate the company’s pricing algorithm by a function $f(\cdot)$. Surge is calculated by the company at a broader geographic unit than the hexagons we have used so far. We call these units “surge areas”; they are shown in Figure (B2b). This means our algorithm will produce a surge factor s_a^t , which is the surge factor in area a at 5-minute interval t .²⁶ As these surge areas are larger than our hexagons, we set the surge factor in the hexagons as equal to its area’s surge factor, i.e. the surge factor of hexagon i is s_a^t for all $i \in a$. Details are provided in Appendix Section D.2.

4.5 Intraday Timing

The timing of events both across periods and within periods is outlined in detail below.

1. At $t = 0$, there is an initial distribution of drivers and riders denoted by \mathcal{S}^0 . This is assumed to be common knowledge among all the drivers.
2. At every time t , the following process occurs:
 - (a) Vacant drivers arrive in locations, either through completed trips, searches, or entry of drivers. This determines the vector of drivers:

$$v^t = [v_1^t, v_2^t, \dots, v_L^t]$$

- (b) Riders appear in each location as a function of surge in the previous period (\bar{r}_i^t) along with unmatched riders from previous periods who re-request (\bar{r}_i^t). This determines the vector of riders:

$$r^t(s^{t-1}) = [r_1^t(s_1^{t-1}), r_2^t(s_2^{t-1}), \dots, r_L^t(s_L^{t-1})]$$

²⁶As shown in Table (1), the actual surge factor s_{at} goes up in increments of 0.25 starting at 1. The surge factor never goes below 1.

- (c) Based on the distribution of drivers and riders, the matching algorithm assigns drivers to riders using the matching function:

$$P^t = m(r^t, v^t)$$

- (d) Matched drivers go to their pick-up locations to begin their trips
(e) Unmatched drivers choose a location to search or exit according to their policy functions σ^t
(f) Unmatched riders exit the market with probability ω . Those who have waited G periods exit with probability 1.

3. Drop-offs, search strategies, and driver entry decisions determine where drivers will be in the next period. This, along with the re-requesting riders, will determine next period's distribution S^1 .

4. This process continues until the end of the day S^T is reached.

We present a more detailed discussion of the intraday state transition and equilibrium in Appendix D.

5 Estimation

The key parameters to estimate in this model are σ_ε (S.D. of the idiosyncratic shock in driver's search/exit problem; Equation 4), σ_η (S.D. of the idiosyncratic shock in driver's entry problem; Equation 9), and μ (the value of the outside option).

This section is divided into three parts. First, we discuss how we prepare our data for estimation. Second, we discuss the objects that are directly identified from the data. Finally, we discuss how we identify the remaining parameters. To compute the equilibrium, we follow a two-step procedure, where we first non-parametrically estimate many parameters directly from the data. After this, we solve the driver's dynamic optimization problems up to the unknown parameters $\sigma_\varepsilon, \sigma_\eta$ and μ . In an outer loop, we then find the parameters values that best match model moments to data moments.

5.1 Preparing the Data

We discretize the data in space and time. We define a "day" as finite set of discrete time intervals representing a 24 hour period starting at 4:00 a.m. We choose to divide the day into 5 minute intervals, which is a natural choice given that RideAustin prices their surge every 5 minutes. As referenced before, the city has been divided into hexagons. However, we make further simplifications to make our estimation computationally feasible.

We focus only on Fridays and Saturdays, as these days have the most number of rides each week and exhibit similar intra-day trip patterns. Additionally, we divide the day into four three-hour shifts

from 3 p.m. to 3 a.m. the following day and estimate our model separately on each shift.²⁷ Using the model notation, time is given by $t \in \{1, 2, \dots, T\}$, where $t = 1$ is the start of the shift (e.g. 3 p.m.), $t = T$ is three hours after the shift start time (e.g. 6 p.m.), and each t is a 5 minute interval.

In terms of space, we also take steps to simplify this further. Given that suburban areas are not dense and feature a small fraction of rides in individual hexagons, we aggregate the non-central areas into a representative hexagon to reduce the dimensionality of our state space (Figure B3a). In the end, we are left with 83 locations in the city. In Appendix B, we describe the process of creating $A(i)$, the accessible areas from each location, which is visualized in Figure (B3b).

5.2 Objects Identified from Data

We can directly identify a number of objects directly by taking averages from the data.

- We can identify the consumer transition probabilities M_{ij}^t by taking the mean of the probabilities of passengers going from i to j over all data in that shift. This means that M_{ij}^t differs across shifts, but is constant within a shift.
- We set the distance δ_{ij} for a trip as equal to the ring relationship between the hexagons i and j : $\delta_{ij} = Ring(i, j)$.²⁸ We find in the data that drivers can travel on average 2 rings in a 5 minute interval, and so we set $\tau_{ij} = \lfloor \frac{\delta_{ij} + 1}{2} \rfloor$ for $\delta_{ij} > 0$ and $\tau_{ij} = 1$ for $\delta_{ij} = 0$ (i.e. in one time interval, a driver can travel 0, 1, or 2 rings; in two intervals, they can travel 3 or 4 rings etc.).
- The company policy gives us the flag drop fare $b = 5$, the distance-based fare $\pi_\delta = 0.99$ per mile, and the time-based fare $\pi_\tau = 0.25$ per minute.²⁹
- We parametrize the fuel cost c_{ij} as $c_{ij} = \frac{g\delta_{ij}}{MPG}$, where g is the average per-gallon price of gas in Austin (\$2.21 per gallon) and MPG is the average fuel efficiency of cars in the data. Using information about the cars' models, we calculate MPG to be 22.7. This gives us the fuel cost c as equivalent to 9.8 cents per mile.
- From the trip data, we set the time of each driver's first dispatch within a shift less 15 minutes as the driver's entry time.³⁰ From this, we can calculate the number of entering drivers at each t for all days in our sample. We set \bar{V}^t as 90th percentile of this distribution, which captures the pool of potential drivers. \bar{V}_i^t is calculated by distributing the total entrants \bar{V}^t using the empirical probabilities of first dispatch locations within a shift.

²⁷While ridesharing drivers do not work in defined shifts, splitting the day in this way means that when a driver exits, we assume they will not return the remainder of the three-hour shift. This does leave the possibility that they could return for work in a different shift, however, our model does not consider any links between shifts.

²⁸If $i = j$, then $\delta_{ii} = 0$ unless i is in the non-central areas, then $\delta_{ii} > 0$ (see Appendix B for more details).

²⁹We first convert δ_{ij} into miles and τ_{ij} into minutes based on regressions from the data so that the fare components better match the data.

³⁰We shift it back by 15 minutes as this is the average searching time for drivers between rides

- We set $G = 2$ and $\omega = 0.5$, which corresponds to a 50% probability of unmatched riders waiting an additional 5 minutes, and riders waiting a maximum of 10 minutes.³¹
- We set the maximum ring distance from which a driver can be pulled in from to be 10 (Figure C15).

After estimating these objects from the data, we next solve the drivers' dynamic programming problem up to the parameters $(\sigma_\varepsilon, \sigma_\eta, \mu)$.

5.3 Model Estimation and Identification

In this section, we explain how we solve our model and estimate the remaining parameters after having estimated the above objects directly from the data. In the first step of this we solve the drivers' problem and then in an outer loop find the parameters $(\sigma_\varepsilon, \sigma_\eta, \mu)$ to match data moments to model moments.

5.3.1 Solving for Equilibrium

The dynamic programming problem of the drivers in Equation (4) suffers from the curse of dimensionality and thus makes it impossible to solve. Hence, we follow Buchholz (2020), who builds on insights in the firm dynamics literature pioneered by Hopenhayn (1992). We leverage the fact that given the large number of drivers, in aggregate, the driver's problem boils down to a single agent problem with deterministic state transitions. Hence we only need to compute the value functions along the equilibrium path of the states S^t for all t . We adapt the Taxi Equilibrium Algorithm of Buchholz (2020) to solve for the equilibrium value functions. See Algorithm (1) in the Appendix for more details.

The algorithm begins with an initial guess of the state S_0^t and then given the state, updates the value functions by backward induction and then uses forward simulation to generate the transition path of the states given the value functions. We iterate until the implied state transitions and the associated value functions are mutually consistent with each other.

Figure (D17) plots the computed value functions for all the locations over time. We can see that equilibrium value functions are closely bunched together and are downward sloping as expected. One might expect perfect spatial arbitrage across locations, but that is prevented by travel costs as well as travel time between locations which leads to heterogeneity in the value functions across space. Vacant drivers prefer to locate themselves in regions with higher and more profitable demand, but locating themselves in those regions after dropping off consumers in non-central locations is costly in terms of higher fuel costs and higher travel time. In Figure (D18), we show the model's convergence for the estimated parameters. For this, we show for iterations $k = 1, \dots, 100$ following Algorithm 1 how the mean difference in driver counts in each location $\left(\frac{1}{|\mathcal{J}|} \sum_{i \in \mathcal{J}} |v_{i,k}^t - v_{i,k-1}^t|\right)$ changes. On average, we find that this count varies by 2 in successive iterations, which is unlikely to result in changes in equilibrium driver behavior.

³¹In the data, the median time a matched rider waits for their driver to arrive is 5.4 minutes. 10 minutes is the 87th percentile in the waiting distribution (Figure C10).

5.3.2 Estimation of Scale Parameters and the Outside Option

With the equilibrium value functions in hand, we can then use the Simulated Method of Moments to construct model moments that we then match with data moments. The unknown parameters of the model that remain to be estimated are:

- (a) the variance of the Type 1 errors of the search and exit problem (σ_ε);
- (b) the variance of the Type 1 errors of the entry problem (σ_η);
- (c) the per-period outside option (μ).

The two scale parameters tell us how much of the drivers' entry, exit and search behavior is driven by model unobservables relative to observables. Large values of σ_ε imply that drivers are more likely to search uniformly across all locations, whereas a low σ_ε equilibrium implies that drivers concentrate their search in areas where the value functions are highest, i.e. in downtown Austin. Similarly, a high σ_η equilibrium implies that drivers are not more likely to enter (conditional on being distributed across space) in downtown Austin relative to suburban Austin, while a low σ_η equilibrium should imply higher entry in the high-activity downtown areas. The outside option μ can be thought of as a reduced form of the drivers' options to earn money driving for other ridesharing platforms or working their main job if driving is a part-time job. The moments we use from the data are:

1. Average pull distance of drivers picking up riders
2. Average length of time a driver is vacant between rides
3. Number of drivers who have trips
4. Exit probabilities in different regions of the city
5. Entry probabilities in different regions of the city

(1) is informative about σ_ε since a higher σ_ε implies the behavior of drivers is determined more by unobservables. This implies that the data will show less targeted searching, which in turn implies larger pull distances. For similar reasons, (2) is also informative about σ_ε since it implies more spread out drivers and longer vacant times for drivers. (3) helps in identifying the outside option μ since the number of drivers who have trips is determined by how many choose to enter and exit. (4) is also informative about the outside option μ since the exit rates in differentially valuable regions tell us about the outside option of drivers. (5) is informative about σ_η and μ since a higher σ_η implies more spread out entry of drivers and a higher μ implies relatively higher entry in downtown Austin relative to the suburbs.

6 Results

6.1 Demand Elasticities

We estimate a version of the reduced form demand function given in Equation (1). As our surge (s_i^t) and potential riders (O_i^t) data are only available at the surge area level, we estimate the following equation:

$$\log(\tilde{r}_{ad}^t) = \alpha_a^t + \theta_1 \log s_{ad}^t + \theta_2 \log w_{ad}^t + \theta_3 \log O_{ad}^t + \varepsilon_{ad}^t \quad (10)$$

where a is a surge area, t is a 5 minute interval, and d is a day in our sample (the unit of observation is the same as in Equation (11)). We take waiting time (w_{ad}^t) to be the average waiting time in minutes for the given a, d, t cell.

Estimating Equation (10) by OLS suffers from the classic endogeneity issue. To circumvent this, we need an exogenous supply shifter to causally estimate the parameters of the demand equation. We use the number of rides ending in location a as an instrument for the price (s_{ad}^t).³² An increase in the number of rides ending in location a imply a positive shift in the supply of drivers in that area. This means there will be a greater number of drivers available for meeting demand in location a , which should push down the price in that location. Note that the number of trips ending are not an endogenous choice of drivers, but rather, are a result of riders' decisions from previous periods. This is important for the exclusion restriction to hold. It is plausible that the decision of a rider choosing to travel from another location b to location a and arriving at time t is orthogonal to the decision of a rider in location a requesting a ride at time t .

The results from the estimation are shown in Table (2). The first column shows the result from the first stage of the model. An increase in the number of rides which are completed in location a lead to a decrease in the surge factor in a . In the second column, the absolute value of the price and waiting time elasticity of demand are 1.62 and 0.014 respectively. The elasticity estimates are similar to those found in Buchholz (2020) and Buchholz et al. (2020), which study taxis in NYC and ridesharing in Prague respectively.

Our estimates however are much higher than those found in Cohen et al. (2016) and Buchholz (2020). We believe that these results reflect a young market after the exit of Uber and Lyft. Table (E9) in the Appendix shows how these estimates vary within the day.³³ While the price elasticity of demand estimates are much higher during the day, waiting time elasticities are higher at night. This is unsurprising given that we are focusing on Fridays and Saturdays when a lot of people frequent restaurants and bars at night. We use the estimates from Table (E9) and use them as an input into the estimation of the remain-

³²While one may expect that waiting time is endogenous as well, we do not instrument for waiting time because we do not expect the endogeneity to be very large given the rich set of fixed effects we include. Cohen et al. (2016) follows a similar approach. Moreover, our estimates of waiting time elasticity are in line with those found by Buchholz et al. (2020). However in Appendix Table E10, we estimate the demand equation using 2 instrumental variables i.e. with lagged ending rides as an additional instrument for waiting time. The results are similar and hence we proceed with our baseline specification.

³³The estimates for 3 p.m.–9 p.m. are similar to the price elasticity estimates Buchholz (2020) finds for short trips (< 4 miles). Almost 60% of our rides are short trips.

der of the structural model.³⁴

Table 2: IV Demand Results

	First Stage	Second Stage
	Surge (Log)	(Log) Rides Requested
Surge (Log)		-1.62*** (0.23)
# Completed Rides (Log)	-0.012*** (0.00)	
Wait time (Log)	0.011*** (0.00)	-0.014*** (0.00)
# People Opening App (Log)	0.011*** (0.00)	0.014*** (0.00)
Area \times Time Interval FE	Yes	Yes
F - stat	403.09	.
Observations	80388	80388

6.2 Structural Estimation Results

The unknown parameters σ_ε , σ_η and μ are estimated by matching the model moments to the data moments described in Section (5.3.2). The results of the matching are discussed below in Section (7.1). The estimated parameters for the four shifts are shown in Table (3). The estimates for the variance of the Type 1 error, σ_ε and σ_η , are similar, but typically lower than those found in Buchholz (2020) and Fréchet et al. (2019). The estimate of the outside option μ generally increases through the day, from an hourly outside option of approximately \$20/hour in the afternoon shift to \$40/hour late at night. This implies that, unsurprisingly, the threshold for drivers logging on to the app and continue driving for the app is higher at night than earlier in the day.³⁵

6.3 Model Fit

We show that our model is able to fit both targeted and non-targeted moments well. The different shifts capture high-activity periods that also have fundamentally different spatial patterns (e.g. the late night shift is primarily trips of people returning to residential areas from the downtown neighborhood). In the results below, we group the city’s neighborhoods into three regions in terms of their centrality, with Region 1 being the downtown area and Region 3 being the outermost suburbs (see Figure B4).

³⁴The parameters help us generate demand at the surge area level. We then distribute the area-level demand to the hexagon-level according to the empirical distributions we observe in our trips data.

³⁵These numbers must not be interpreted as the hourly wage that they would have earned if not driving for RideAustin. This reflects the *value* that drivers have by *not* driving for RideAustin for the hour.

Table 3: Parameter Estimates

Shift	Time	σ_ε	μ	σ_η
Afternoon	3 p.m. – 6 p.m.	2.79 (0.211)	1.64 (0.110)	2.93 (0.179)
Evening	6 p.m. – 9 p.m.	1.64 (0.307)	2.64 (0.208)	1.50 (0.257)
Night	9 p.m. – 12 p.m.	3.21 (0.445)	3.21 (0.150)	2.50 (0.214)
Late Night	12 a.m. – 3 a.m.	1.79 (0.269)	3.36 (0.307)	3.36 (0.315)


Notes: Parameters for each shift are estimated separately. Standard errors in parenthesis.

Table 4: Targeted Moments

	<i>Afternoon</i>		<i>Evening</i>		<i>Night</i>		<i>Late Night</i>	
	Data	Model	Data	Model	Data	Model	Data	Model
Average Pull Distance (# of hexagons)	1.92	2.07	1.83	1.97	1.51	1.90	1.26	1.70
Vacant Searching Times (5 min intervals)	5.32	7.00	3.88	5.63	3.46	7.44	3.03	5.87
Drivers With Trips	481	581	644	684	695	721	559	598
P(Entry in Region 1)	0.24	0.17	0.25	0.18	0.27	0.22	0.45	0.40
P(Entry in Region 2)	0.38	0.46	0.41	0.46	0.44	0.50	0.40	0.46
P(Exit in Region 1)	0.27	0.35	0.42	0.47	0.39	0.49	0.30	0.32
P(Exit in Region 2)	0.39	0.42	0.42	0.44	0.42	0.40	0.45	0.43

We first show how our model matches the targeted moments for both the shifts and then present results for non-targeted moments. As described in the previous section, the moments we use for estimation are: (1) Average pull distance of drivers picking up riders, (2) Average length of time a driver is vacant, (3) Number of drivers with trips, (4) Exit probabilities in different regions of the city, and (5) Entry probabilities in different regions of the city. Table (4) shows how well our model is able to match targeted moments while Table (F12) in the Appendix shows how our model is able to match non-targeted moments. Notably, drivers in our model spend longer searching than what is observed in the data. Our benchmark model also matches non-targeted moments in the data well as shown in Table (F12), with a similar fraction of surged rides but a lower number of completed trips (which helps explain the lower average earnings and longer search times). Finally, Figure (15) shows that our baseline model matches the distribution of surge prices observed in the data closely.

Figure 15: Distribution of Surge in the Data and Model



Notes: In panel (a), we show how the distribution of surge factors from the model compares to the distribution from the data. In panel (b), we show this for the late night shift. The results for the other two shifts are similar.

7 Counterfactuals

Next, we present the results from different policy counterfactuals that we conduct where we try and disentangle the relative contributions of surge pricing and centralized matching in the efficiency gains from ridesharing. This section is organized as follows: (1) First, we show how data moments change under three counterfactual scenarios, where we remove matching and surge pricing separately and at the same time. (2) Next, we look at welfare and show how varying surge pricing and centralized matching affects consumer surplus, driver and platform revenue. (3) Given our findings, we propose a simple change to the pricing algorithm to increase welfare and also understand the implications of alternate driver compensation schemes.

7.1 Role of Surge and Matching

We evaluate the effect of surge pricing and matching in the ridesharing market by running counterfactuals where we remove each feature. Through this exercise we are able to isolate the mechanisms by which surge and matching drive the efficiency gains of ridesharing platforms. We present the results from three counterfactuals in this subsection. First, we shut off surge pricing by fixing the surge factor to be 1 in all times and locations. Therefore the fare structure is now independent of local demand and supply conditions in all locations at all times. Second, we make supply local by only allowing matches when drivers and riders are in the same location. More precisely, we allow for Leontief matching only *within* a hexagon, but do not allow drivers to be pulled from other locations to be matched to a rider. This effectively removes the centralized matching that ridesharing provides. Finally, we shut off both surge pricing and centralized matching. We call this counterfactual the “gig taxi” equilibrium as this

represents a traditional taxi market that lacks surge and matching but where drivers can still enter and exit.

7.1.1 Response of Moments

When we turn off surge pricing, the number of drivers drops dramatically in the late night shift. The drivers who do enter are in high demand, which results in them having no vacant search times as they quickly get matched to their next ride after dropping off their previous passenger. Given that 47% of rides in the data are surged during the late night shift, this implies that surge is crucial in inducing driver entry. This aligns with our parameter estimates as we find that the value of the outside option (μ) is highest during this shift (Table 3). In the benchmark model, the surged rides primarily occur in the downtown area (Region 1). Consistent with this, when surge is removed, we find that drivers are more likely to enter (and less likely to exit) in the downtown area as these are the highest value locations. Despite the total number of rides decreasing, note that the drivers who do enter earn 50% more than the benchmark as they complete more trips per driver. However, with so few drivers participating in the market, more ride requests are expected to go unmet (69% versus 3% in the benchmark). Figure (16) shows that the unmet demand occurs throughout the city, even in the central locations. In contrast to the late night results, the no surge counterfactual for the afternoon (3 p.m. – 6 p.m.) shift has starkly different outcomes (Appendix F.1). In particular, removing surge has little effect on the equilibrium in the afternoon. This is expected as only 18% of rides are surged in the benchmark during that part of the day and the estimate of the outside option is lowest during that shift. This exercise shows that surge works by inducing drivers to enter, which is especially important during times when drivers highly value their outside option.

Table 5: 12 a.m. to 3 a.m. Shift - Targeted Moments

	Data	Benchmark	No Surge	No Match	Gig Taxi
Average Pull Distance	1.26	1.70	1.38	0.72	0.86
Vacant Searching Times	3.03	5.87	0.00	6.30	1.83
Drivers With Trips	559	598	86	499	62
$P(\text{Entry in Region 1})$	0.45	0.40	0.50	0.47	0.58
$P(\text{Entry in Region 2})$	0.40	0.46	0.47	0.38	0.34
$P(\text{Exit in Region 1})$	0.30	0.32	0.23	0.37	0.27
$P(\text{Exit in Region 2})$	0.45	0.43	0.36	0.49	0.61

Notes: Average Pull distance is the average ring distance between the rider and driver.

Vacant Searching time is the number of time periods drivers search between rides.

$P(\text{Exit/Entry in Region } i)$ is the fraction of rides ending/starting in a region which were the last/first ride.

In contrast, when we shut off matching, drivers spend longer searching for a ride and also exit at much higher rates than our benchmark model. While drivers generally search in the right areas,

the fact that they cannot exactly predict the location of demand means that unmet demand is high (36%). This results in a much lower number of completed trips and lower average earnings for drivers. This highlights the important role that matching plays in balancing demand and supply across space and eliminating match frictions. In Figure (16), we see that the increase in unmet demand for this counterfactual comes almost entirely from the non-central areas. Drivers are unlikely to travel to those areas: central areas provide higher rides, rides that are more likely to be surged, and have lower travel costs to reach them. Moreover drivers tend to drive short shifts and exit sooner. These results speaks to how ridesharing platforms, through centralized matching, can reduce the spatial inequality in access to transport. In summary, matching solves the *static inefficiency* i.e. riders and drivers being unable to match with each other due to imbalances in demand and supply in localized regions.

Table 6: 12 a.m. to 3 a.m. Shift - Non-targeted Moments

	Data	Benchmark	No Surge	No Match	Gig Taxi
Average Earnings	54.56	39.83	59.85	28.40	29.47
Surged Rides (%)	46.88	48.10	0.00	38.15	0.00
Unmet Demand (%)	–	2.86	69.25	35.91	89.04
Last Rides (%)	26.38	36.64	16.54	47.48	34.25
Average Search Distance	–	20.86	–	25.36	2.83
Number of Trips	2,075	1,632	520	1,051	181
$P(\text{Search in Region 1})$	–	0.32	–	0.29	0.35
$P(\text{Search in Region 2})$	–	0.65	–	0.67	0.62
$P(\text{Search in Region 3})$	–	0.03	–	0.04	0.03

Notes: Average Earnings is the average revenue earned by a driver during the shift. Surged Rides is the proportion rides which are surged. Unmet Demand is the proportion of requested rides which were unmet. Last Rides is the proportion of total rides which were last rides for a driver. Average Search Distance is the ring distance between the drop-off location for trip k and dispatch location of trip $k + 1$. $P(\text{Search in Region } i)$ refers to the proportion of total searches in Region i

In addition, we see the combined effect of surge and matching from the “gig taxi” counterfactual. Here, we remove surge pricing and centralized matching as before but maintain the “gig economy” aspect of ridesharing by still allowing drivers to enter and exit. By removing matching, this makes it more difficult for drivers to find rides as compared to the no surge counterfactual. Consequently, this magnifies the issues of removing surge: the number of drivers and total trips falls even further, while the proportion of unmet demand increases. Unlike the no surge counterfactual, drivers do have to search for their next ride, but on average they do not have to search for a long time. In the afternoon shift, there is little difference between the no match and gig taxi counterfactuals – as before, surge plays a minimal role and so removing it does not have a large impact on the equilibrium outcome.

Figure 16: Spatial Distribution of Proportion of Unmet Demand (12 a.m. – 3 a.m.)






Figure 17: Comparison of Different Simulations to the Taxi Equilibrium

(a) 3 p.m. – 6 p.m.



(b) 12 a.m. – 3 a.m.



Notes: In panel (a), we show how each of the moments for our counterfactuals compares to the taxi equilibrium for the afternoon shift. In panel (b), we show this for the late night shift.

We summarize the key comparisons across the counterfactuals in Figure (17) where the black dotted line represents the gig taxi equilibrium. The red bars represent our benchmark model, the blue bars represent the counterfactual without surge pricing, while the gray bars represent the counterfactual without centralized matching. All values have been normalized to be relative to the gig taxi equilibrium. While the effects of shutting of centralized matching are consistent across the shifts, the effects of the surge counterfactual have dramatically different implications across the shifts. The results also highlight the importance of not only spatial, but also temporal variation when evaluating counterfactual outcomes.

7.2 Welfare

Given the underlying mechanisms that surge and matching play in the determining equilibrium outcomes in the model, in this section, we turn to evaluating how different counterfactual scenarios translate to changes in consumer surplus, driver revenue as well as platform profits. We then draw policy implications for policy makers wanting to regulate ridesharing platforms such as Uber and Lyft as well as aiming to improve the competitiveness of traditional taxis. We first define how we calculate each of these quantities. Driver revenues are defined as:

$$\text{Driver Revenues} = \sum_v \left[\sum_{x \in \text{Trips}} [(b + \pi_\delta \delta_x + \pi_\tau \tau_x) s_x - c_x] - \sum_{y \in \text{Pulls}} c_y - \sum_{z \in \text{Search}} c_z \right]$$

This is simply the sum of payoffs (fare less fuel cost and dollar per ride) earned on each trip x , the fuel cost for every pull y , and the fuel cost for every search z . We sum this for every driver v .

To compute consumer surplus, we use the demand function given in Equation (10) and compute the area under the curve. More precisely:

$$\text{Consumer Surplus} = \sum_{it} \left(\frac{m_{it}}{r_{it}} \right) \int_{s_{it}^*}^{\infty} \int_{w_{it}^*}^{\infty} e^{\alpha_{it}} s_{it}^{\theta_1} w_{it}^{\theta_2} O_{it}^{\theta_3} dw_{it} dp_{it}$$


where $\frac{m_{it}}{r_{it}}$ is the proportion of rides requested which got a ride. For implementation purposes, since these elasticities are local, evaluating consumer surplus at high values of surge and waiting time requires extrapolation (Buchholz, 2020). Hence we set the upper bound on surge and waiting time to be 7 and 30 respectively. Computing consumer surplus this way also allows us to understand the spatial heterogeneity in consumer surplus across the different counterfactuals.

Computing platform revenues is straightforward in our setting since Ride Austin charges a flat \$0.99 fee for each trip. Hence platform revenues equal the number of trips. We estimate an annual consumer surplus of \$5.5 million for the afternoon shift and \$44.5 million for the late night shift in our benchmark estimation.³⁶ Driver profits are estimated to be \$4.3 and \$6.9 million for the two shifts, while platform revenues are \$457,000 and \$601,000, respectively. We also compute these quantities in our two main counterfactuals and compare them to the taxi equilibrium. The results are shown in Figure (18).

³⁶This is an upper bound given that these numbers are calculated based on data for Fridays and Saturdays.

In the 3 p.m.–6 p.m. shift, barring average driver revenue for the counterfactual without matching, all the models seem to generate substantial gains across the board relative to the taxi equilibrium. If surge pricing is shut off, consumer surplus increases by 4.1% while driver revenues fall by 20% relative to the benchmark model. Moreover, platform revenue is unchanged.³⁷ However if matching is shut off, the fall in consumer surplus, driver revenues and platform revenues are substantially larger relative to the benchmark. These results however are extremely different if we analyze the late night 12 a.m.–3 a.m. shift. While consumers and drivers are typically better off in any simulation relative to the taxi equilibrium, consumers, drivers and the platform are all worse relative to the benchmark in all counterfactual scenarios.

Figure 18: Comparison of Different Simulations to the Taxi Equilibrium




Notes: In panel (a), we show how consumer surplus, driver revenue and platform revenue for our counterfactuals compares to the taxi equilibrium for Shift 3. In panel (b), we show how consumer surplus, driver revenue and platform revenue for our counterfactuals compares to the taxi equilibrium for Shift 6.

Interestingly, there is heterogeneity in the spatial distribution of consumer surplus. Figure (19) shows that consumer surplus skews towards the central areas when we shut off matching in the afternoon shift. Centralized matching is key to ensuring an equal distribution of consumer surplus across space. We see a similar spatial distribution of consumer surplus in the taxi market. These results for the late night are shown in Appendix Figure (F21). While the results there are less stark, the same general pattern appears.

Overall our results indicate that centralized matching plays a key role in the efficiency of ride-sharing through the course of the day. Surge pricing has small negative effects on consumer surplus in the day (relative to our benchmark model), however surge pricing plays a key role in increasing welfare late at night. In particular, during late night shifts, surge pricing and centralized matching seem to act on different dimensions and thus are likely complementary. As discussed in the previous section, surge

³⁷While economic theory predicts that flexible pricing must improve consumer surplus, this may not be as straightforward in ridesharing markets. Since surge factors are never set below one, it is not inconceivable that the downward rigidity of surge pricing may in fact not be welfare improving for consumers.

Figure 19: Spatial Distribution of Consumer Surplus (3 p.m. – 6 p.m.)



solves the driver entry problem while centralized matching minimizes match frictions. These findings which hint at the complementarity between surge pricing and centralized matching are analyzed further in the next subsection.

7.3 Complementarity

Our model incorporates three key elements of ridesharing: surge pricing, matching technology, and flexible entry/exit by drivers. We can evaluate complementarities between these features by comparing counterfactuals where we activate each feature. This gives us 8 possible counterfactuals, whose welfare implications for the 12 a.m.–3a.m. shift are shown in Figure (F22) in the Appendix. We compare all outcomes relative to the taxi equilibrium, which has none of three ridesharing features. Consumer surplus is highest with a traditional taxi system endowed with a centralized matching technology. The

intuition behind this result follows from the fact that in a traditional taxi market, supply is held fixed and therefore surge doesn't play a critical role in increasing consumer surplus. Drivers on the other hand benefit the most in a ridesharing world due to the majority of rides being surged late at night.³⁸

These figures also provide evidence for complementarities between surge and matching. Moving from the flexible taxi counterfactual to the no surge counterfactual adds matching technology. Likewise, moving from the flexible taxi counterfactual to the no matching counterfactual adds surge pricing. In either case, for the 12 a.m. – 3 a.m. shift, there are gains to consumer surplus, driver and platform revenues. However, moving from flexible taxis to ridesharing simultaneously adds surge and matching. This change increases welfare for all agents by *more* than the sum of the gains from adding each feature separately. These results are captured in Table (7). We interpret this as evidence for the complementarity between surge and matching. To be more precise, surge pricing induces more drivers to enter and leads to a thicker market which ensures that matching works well i.e. drivers are not pulled in from too far away. Appendix Section (F3) includes further discussion of results from these counterfactuals.

Our results allow us to draw two policy implications. First, banning or capping surge pricing, as many cities are considering may in fact harm both riders and drivers on the whole. This can be seen in Figure (F22). Second, for taxicabs who wish to compete with ridesharing platforms, the results show that surge pricing is *not* the answer. Instead, by introducing a matching technology, taxicabs would be able to potentially increase consumer and driver welfare. The different policy recommendations stem from the nature of supply for each of the two systems.

Table 7: 12 a.m. to 3 a.m. Shift - Complementarity

Welfare	No Surge	No Match	Benchmark	Complementarity
Consumer Surplus	159.42	414.89	677.73	103.42
Driver Revenue	181.39	514.73	1067.66	371.55
Platform Revenue	187.29	480.66	801.66	133.70

Notes: No Surge is the counterfactual without surge pricing. No Match is the counterfactual without matching technology. Benchmark is the ride sharing model (surge and matching). All figures represent welfare gains relative to the Flexible Taxi counterfactual (no surge or matching, but entry/exit as in all previous models). The fourth column represents the complementarity gain = Benchmark – (No Surge + No Match)


To further assess how these two features interact, we simulate the model by varying both at a more granular level. For matching, we vary the maximum ring (in other words the maximum distance) from which a driver could be pulled from. The no match counterfactual only pulls from rings less than 1, while the benchmark model pulls from rings less than 10.³⁹ For surge, we amplify (or dampen) surge by a factor $\theta \geq 0$ such that a surge factor s becomes amplified to $s_\theta = 1 + \theta(s - 1)$.⁴⁰ The no surge

³⁸A similar graph for the 3 p.m.–6 p.m. shift is shown in Figure (F24) in the Appendix.

³⁹A maximum ring distance of 10 implies that the maximum distance from which drivers could be pulled in is approximately 9 kms.

⁴⁰For example, a surge of 1.25 and a factor $\theta = 2$ would be amplified to $s_\theta = 1.5$.

Figure 20: Welfare Effects of Varying Matching and Surge



Notes: Results are simulated for each shift and then aggregated over the day. Total welfare is the sum of consumer surplus, driver profits, and platform profits.


counterfactual has $\theta = 0$ and the benchmark has $\theta = 1$.

In Figure (20), we show the effects of varying each of these feature separately and compare the welfare effects to that of the benchmark model. In the left panel, we see that increasing the maximum pull ring (i.e. reducing match frictions) has unambiguously positive welfare effects for all agents. Unsurprisingly, there are diminishing returns to the gains from reducing match frictions, evidenced by the concavity of the benefits for all agents. In contrast, amplifying surge reduces consumer and total welfare, while only increasing drivers’ profits. In particular, given the other benchmark settings, the current surge is optimal for consumer welfare, but not for drivers (a result consistent with [Castillo \(2019\)](#)). This is because the higher surge prices out low value customers, resulting in a lower number of rides but higher overall fares.⁴¹

Next, we vary both features together, as shown in the benchmark plots i.e. the first column of plots in Figure (21). These graphs plot how consumer surplus, driver profits and total welfare compare relative our benchmark model in the yellow square. From this, we can see that while both riders and drivers view reducing match frictions as welfare increasing, there is a conflicting role for surge. Higher surge generally reduces consumer welfare but increases driver profits. By increasing both surge and the maximum pull distance, the platform would be able to increase total welfare. Moreover, the plots suggest that drivers have convex iso-profit lines for surge and matching, which supports the complementarity between the two features.

⁴¹Figure (F25) shows the effects of average prices and pull distances under these simulations. In particular, a surge amplification of 5 corresponds to approximately an average price (surge factor) of 1.8.

Figure 21: Welfare Effects of Varying Matching and Surge Simultaneously



Notes: Results are simulated for each shift and then aggregated over the day. Total welfare is the sum of consumer surplus, driver profits, and platform profits. All values are relative to the benchmark model with max pull ring at 10 and surge amplification at 1 (indicated by the yellow rectangle). Benchmark indicates the baseline model that was estimated. Platform commission is a counterfactual where the platform receives 25% of the fare. Flexible surge is a counterfactual where surge is allowed to go below 1.

7.4 Impact of Alternate Driver Compensation and Pricing Schemes

Our results so far suggest that while matching primarily drives the efficiency gains of ridesharing, the platform’s pricing structure is still key to understanding the full welfare effects. Moreover, given the unique nature of the ridesharing platform in our setting, it is not obvious how our results would apply to a scenario in the presence of more traditional driver compensation schemes as implemented by ridesharing platforms like Uber or Lyft. Therefore, in this section, we look to answer two questions. First, would these results generalize to other ridesharing companies such as Uber and Lyft? Second, are there alternative pricing structures that could increase welfare? Figure (21) shows these results. The first column of graphs presents results about how consumer surplus, driver profit and total welfare⁴² change relative to our benchmark model; denoted by the yellow square, by varying the match frictions as well as amplifying surge in the model we have developed so far. The second column of graphs performs these same counterfactuals in a world with a traditional ridesharing platform which charges a commission. Finally the last column of graphs modifies the surge pricing algorithm by allowing it to be fully flexible.

7.4.1 Alternate Driver Compensation Schemes

The company we study, RideAustin, is a non-profit, which only takes 99 cents from every completed trip (with the rest going to the driver). A key difference between RideAustin and a platform like Uber is that Uber takes a 25% commission on rides.⁴³ To see whether this would affect our results, we simulate a counterfactual where the platform receives 25% of the total fare. This is shown in the second column of graphs under the header “Platform Commission” of Figure (21). We see that under this fare structure, consumers welfare and total welfare are lower in all possible combinations, as compared to the benchmark model. The platform taking a proportional cut creates a distortion which prevents it from achieving the socially optimal allocation. In particular, consumers are strictly better off in a world with a platform like RideAustin relative to the presence of traditional platforms like Uber and Lyft who follow the same pricing rule as Ride Austin. This is not surprising given that traditional ride-sharing platforms do not necessarily internalize the effects of their compensation schedules on driver payoffs which in turn affects consumer surplus. These results however do not imply that there exist other compensation and pricing schemes that could be implemented which could benefit consumers and drivers relative to our benchmark model. However, there is scope to improve drivers’ profits, though less so than under the flat rate benchmark i.e. in the model where drivers keep the entire fare of the ride. Importantly, the general welfare differentials of surge and matching continue to hold under this counterfactual. We interpret this as our results continuing to hold, even under a pay structure such as that of Uber.

To make this point clearer we consider another driver compensation scheme where drivers keep a proportion α of the trip revenue instead of paying a \$1 fee so as to keep their total earnings the same.

⁴²Total welfare is computed as the sum of consumer surplus, driver profit and platform revenue.

⁴³Uber’s website states “Uber charges partners 25% fee on all fares”, which is also consistent with Castillo (2019).

This α is calculated to be 92% of the trip revenue. We rerun our model with drivers keeping 92% of the trip revenue instead of paying a \$1 fee and find no difference from the benchmark results. This is seen in Figure (F23) in the Appendix. This is driven by the fact that the lumpsum fee affects the extensive margin (entry and exit) which is highly responsive while the proportional compensation scheme affects the intensive margin as well (search) which is not as responsive.

7.4.2 Alternate Pricing Schemes

While matching appears to be preferred by both sides of the market, there is a limitation: match frictions can only be reduced so much. In our context, the gains from pulling from a further distance diminish as eventually the entire city is covered and there are no additional areas from which to pull drivers. Selecting a different pricing structure provides many more options for the platform (or social planner) to increase welfare. Our results so far show that surge plays a more limited role in improving consumer welfare. This may be surprising given that we would expect prices that are more responsive to market conditions would be welfare-improving. However, one key limitation of surge, as implemented by companies like RideAustin and Uber, is that it is one-sided. The surge factor is never set below 1, which means that prices *increase* in relatively *high* demand areas, but relatively *low* demand areas do not receive a discount. Fully flexible prices should be responsive to both types of market fluctuations. This simple idea is echoed in [Bimpikis et al. \(2019\)](#) and [Besbes et al. \(2020\)](#) which study the optimal pricing in ridesharing platforms. The key property characterizing optimal prices is that demand and supply are “balanced” in different locations. In essence, once drivers drop off riders in a non central location i with low demand, there is a mismatch in relative demand and supply at that point in time. Instead of drivers searching their way back to high demand areas, it is better to give a “discount” to location i to induce higher demand so that supply is optimally utilized across space. To see how such a pricing rule impacts welfare, we run a counterfactual allowing for flexible surge (a surge factor that can go below 1, with a lower bound of 0.25). Our surge pricing algorithm allows us to approximate the “balancedness” property since the inputs into the algorithm are the demand and supply at a given location. The results from this exercise are shown in the third column of graphs (“Flexible Surge”) of Figure (21). We see that this simple change in the pricing rule greatly increases welfare for all agents as compared to the benchmark. We find evidence consistent with the market becoming more “balanced” leading to better matching as the time drivers spend searching between rides falls by 16.5%. Under the flexible surge structure, the trade-offs between surge and matching still persist, though the drivers’ iso-profit curves appear to be more convex. This suggests a greater complementarity between the two features.


7.5 Static and Dynamic Inefficiencies

As argued before, transportation markets exhibit two types of inefficiencies. The first is a static inefficiency: at a given point in time, riders and drivers may be unable to match with each due to search frictions. The second is a dynamic inefficiency: drivers’ location is affected by riders’ destination choices,

which may result in a suboptimal spatial distribution of supply. It is important to then see what role surge and matching play in reducing these inefficiencies.

We measure the static inefficiency as the proportion of rides requested which remained unfulfilled. Panel (a) of Figure (22) plots the proportion of unmet demand under different configurations. The benchmark model exhibits a small proportion of unmet demand. Drawing on the results shown before, both surge and matching play a role in clearing the market. Surge pricing solves the driver entry problem i.e. creates a thick market while centralized matching helps in clearing the market. In the counterfactuals without surge or matching, the proportion of unmet demand goes up to nearly 30%. The last bar in the panel shows a scenario where we fix supply to be the median number of drivers and don't allow them to exit and also shut off surge pricing. This shows that given a large supply of available drivers, matching is able to solve most of the *static inefficiency* in the market. Thus both surge and matching play independent and crucial roles in mitigating the static inefficiency.

Figure 22: Static and Dynamic Inefficiencies



Notes: In panel (a), we show how surge and matching reduce unmet demand; our measure of static inefficiency. No Match and No Surge refer to counterfactuals without matching and pricing respectively. No Surge (Fixed SS) refers to a counterfactual without surge pricing and fixed supply of drivers. In panel (b), we show the relative importance of surge over matching in reducing the vacant time between rides for drivers. Flexi Surge refers to a counterfactual where the surge factor is allowed to go below 1. No Match (Flexi) refers to a counterfactual without matching but with flexible pricing.

The dynamic inefficiency is shown in Panel (b) of Figure (22) and is measured as the vacant time periods (measured in 5 minute intervals) between successive rides. The benchmark model shows that drivers wait 30 minutes between rides in the benchmark. While there exists an option value of waiting and searching for a ride, the lower the time drivers wait between rides, the lower the inefficiency. The inefficiency here can be understood through an example : if a driver drops off a rider in the suburbs and then has to drive all the way back to downtown Austin to be matched again, this is strictly worse off

compared to a situation where the driver is matched to a rider traveling from the suburbs to the center. The benchmark surge and matching technologies are unable to address this problem. As mentioned above, the flexible pricing counterfactual conducted above aims to solve exactly that. In the counterfactual where prices are allowed to be fully flexible i.e. the surge factor goes below 1, we find that the vacant search time is almost halved. Thus the ability to induce demand in the suburbs and suppress demand in downtown leads to a more balanced pattern of demand and supply leading to lower wait-times between rides. The last bar shows that matching is still an important factor over and above the flexible surge since the market still needs to be cleared even after the flexible prices leads to a balanced distribution of demand and supply.

8 Conclusion

Modeling the ridesharing platform, RideAustin allowed to understand the relative importance of surge pricing and centralized matching in overcoming existing market inefficiencies and increasing welfare relative to the status quo. We conduct policy counterfactuals and find that both centralized matching and surge pricing are crucial in increasing consumer surplus as well as driver profits. The mechanisms through which they work however are different. Surge improves outcomes through its effect on incentivizing driver entry while centralized matching improves outcomes by reducing match frictions, especially in non-central areas. We also find complementarities between surge and matching i.e. matching improves efficiency more given surge, and vice-versa. We implement a simple change in the pricing rule by making it more flexible which generates large welfare gains for all agents due to its effect on reducing spatial misallocation. Our results indicate that policymakers should think before banning surge pricing as this may reduce consumer and driver surplus. In contrast, taxis who wish to compete with ridesharing should focus on introducing a matching technology, rather than implementing surge pricing.

Our data does not allow us to model platform competition and its impact on riders, drivers and overall market efficiency. This is an important avenue for further research. Since for-profit platforms have been pursuing market share aggressively, it would not only require modeling their objectives in a dynamic framework but also collect and utilize long term data from multiple such platforms. Finally, it would be interesting to study how different pricing and compensation schemes may enhance welfare even beyond the simple changes we simulated in this paper.

References

- Angrist, Joshua D, Sydney Caldwell, and Jonathan V Hall**, “Uber vs. Taxi: A Driver’s Eye View,” *NBER Working Paper*, 2017.
- Berger, Thor, Chinchih Chen, and Carl Benedikt Frey**, “Drivers of disruption? Estimating the Uber effect,” *European Economic Review*, 2018, 110, 197–210.
- Besbes, Omar, Francisco Castro, and Ilan Lobel**, “Surge pricing and its spatial supply response,” *Management Science*, 2020.
- Bimpikis, Kostas, Ozan Candogan, and Daniela Saban**, “Spatial pricing in ride-sharing networks,” *Operations Research*, 2019, 67 (3), 744–769.
- Brancaccio, Giulia, Myrto Kalouptsi, Theodore Papageorgiou, and Nicola Rosaia**, “Search Frictions and Efficiency in Decentralized Transportation Markets,” Technical Report, National Bureau of Economic Research 2020.
- Buchholz, Nicholas**, “Spatial Equilibrium, Search Frictions and Dynamic Efficiency in the Taxi Industry,” *Working Paper*, 2020.
- , **Laura Doval, Jakub Kastl, Filip Matějka, and Tobias Salz**, “The Value of Time: Evidence From Auctioned Cab Rides,” Technical Report, National Bureau of Economic Research 2020.
- Castillo, Juan Camillo**, “Who Benefits from Surge Pricing?,” *Working Paper*, 2019.
- Castillo, Juan Camilo, Dan Knoepfle, and Glen Weyl**, “Surge pricing solves the wild goose chase,” in “Proceedings of the 2017 ACM Conference on Economics and Computation” 2017, pp. 241–242.
- Chen, M Keith, Judith A Chevalier, Peter E Rossi, and Emily Oehlsen**, “The Value of Flexible Work: Evidence from Uber Drivers,” *NBER Working Paper*, 2019.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe**, “Using Big Data to Estimate Consumer Surplus: The Case of Uber,” *NBER Working Paper*, 2016.
- Dholakia, Utpal M**, “Everyone Hates Uber’s Surge Pricing - Here’s How to Fix It,” *Harvard Business Review*, dec 2015.
- Diakopoulos, Nicholas**, “How Uber surge pricing really works,” 2015.
- Einav, Liran, Chiara Farronato, and Jonathan Levin**, “Peer-to-peer markets,” *Annual Review of Economics*, 2016, 8, 615–635.
- Farronato, Chiara and Andrey Fradkin**, “The welfare effects of peer entry in the accommodation market: The case of airbnb,” Technical Report, National Bureau of Economic Research 2018.

- Fréchette, Guillaume R, Alessandro Lizzeri, and Tobias Salz**, “Frictions in a Competitive, Regulated Market: Evidence from Taxis,” *American Economic Review* (forthcoming), 2019.
- Gavazza, Alessandro**, “An empirical equilibrium model of a decentralized asset market,” *Econometrica*, 2016, 84 (5), 1755–1798.
- Ghili, Soheil and Vineet Kumar**, “Spatial Distribution of Supply and the Role of Market Thickness: Theory and Evidence from Ride Sharing,” 2020.
- Hall, Jonathan, Cory Kendrick, and Chris Nosko**, “The effects of Uber’s surge pricing: A case study,” *The University of Chicago Booth School of Business*, 2015.
- Hall, Jonathan D, Craig Palsson, and Joseph Price**, “Is Uber a substitute or complement for public transit?,” *Journal of Urban Economics*, 2018, 108, 36–50.
- Hendel, Igal, Aviv Nevo, and François Ortalo-Magné**, “The relative performance of real estate marketing platforms: MLS versus FSBO Madison.com,” *American Economic Review*, 2009, 99 (5), 1878–98.
- Hopenhayn, Hugo A**, “Entry, Exit, and firm Dynamics in Long Run Equilibrium,” *Econometrica*, 1992, 60 (5), 1127–1150.
- Joskow, Paul L and Catherine D Wolfram**, “Dynamic pricing of electricity,” *American Economic Review*, 2012, 102 (3), 381–85.
- Lazarev, John**, “The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in US monopoly markets,” *New York University*, 2013.
- Lee, Min Kyung, Daniel Kusbit, Evan Metsky, and Laura Dabbish**, “Working with machines: The impact of algorithmic and data-driven management on human workers,” in “Proceedings of the 33rd annual ACM conference on human factors in computing systems” 2015, pp. 1603–1612.
- Liu, Tracy, Zhixi Wan, and Chenyu Yang**, “The efficiency of a dynamic decentralized two-sided matching market,” *Available at SSRN 3339394*, 2019.
- Rochet, Jean-Charles and Jean Tirole**, “Platform competition in two-sided markets,” *Journal of the european economic association*, 2003, 1 (4), 990–1029.
- and —, “Two-sided markets: a progress report,” *The RAND journal of economics*, 2006, 37 (3), 645–667.
- Rosaia, Nicola**, “Competing platforms and transport equilibrium: Evidence from New York City,” Technical Report, mimeo, Harvard University 2020.
- Salz, Tobias**, “Intermediation and competition in search markets: An empirical case study,” Technical Report, National Bureau of Economic Research 2020.
- Shapiro, Matthew H**, “Density of Demand and the Benefit of Uber,” 2018.

Williams, Kevin, "Dynamic airline pricing and seat availability," 2020.

Wolak, Frank A, "Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment," *American Economic Review*, 2011, 101 (3), 83–87.


Yan, Chiwei, Helin Zhu, Nikita Korolko, and Dawn Woodard, "Dynamic Pricing and Matching in Ride-Hailing Platforms," *Naval Research Logistics*, 2019.

Appendix

A Data Preparation

We restrict to rides between October 1, 2016 and May 31, 2017. Within this, we eliminate other rides to ensure that our data comes from a typical week. As Figure (A1) shows, within our sample period, there are two periods where this is unlikely to be true. First, is South by Southwest (SXSW), an annual festival hosted in Austin. Second, is the Christmas/New Year's holiday period. We exclude these two periods from our data. We also remove rides that are categorized as premium (a more expensive luxury option but only 5% of rides) and rides that appear to have data entry issues (3% of rides). After this, we also mostly focus on rides that begin and end inside Austin city (see Figure B2a), which further removes another 6% from the sample. In the end, we have over 1.3 million rides in the sample.

Figure A1: RideAustin Trips Over Time (October 2016–May 2017)



To prepare the data for the model, we focus more specifically on rides that occur on Fridays and Saturdays. The plurality of rides (46.5% of our sample) occur on these two days and they exhibit fairly similar patterns during the day.⁴⁴


B Geographic Units

We focus on rides in Austin city. Since the official city limits are highly irregular (gray area in Figure B2), we use a collection of census tracts to approximate the boundaries of the city in a more compact manner. We then divide up this 'city region' into neighborhoods as shown in Figure (B2a). Neighborhoods are

⁴⁴More precisely, since our days begin at 4 a.m., we restrict to rides between 4 a.m. on Friday and 4 a.m. on Sunday.

based on collections of RideAustin’s surge areas (Figure B2b). These surge areas are the level at which the company sets surge. While the process of creating neighborhoods is ad hoc, we try best to mimic known neighborhoods in Austin (e.g. Downtown, South Central Austin, North Central Austin) and use reasonable boundaries (e.g. highways) as guides.

Figure B2: Austin Map




Notes: The gray area is the official municipal definition of Austin. The orange lines are highways. The colored regions are neighborhoods created by aggregating RideAustin surge areas intersected with a group of census tracts that approximate Austin city. Four points of interest are plotted on the map for illustrative purposes.

To have better notions of local demand and supply, it is useful to discretize the space into smaller units. For this we use Uber’s H3 geospatial indexing system (h3geo.org) and use hexagons of resolution 8 (approximately 0.93 kilometers wide). A useful feature of this system is that it allows for a convenient way to describe the relationship (distance) between any two hexagons. In Figure (11), the red central hexagon is the reference hexagon. We can refer to the ring itself as being in “ring 0”. The hexagons in blue that are adjacent to the reference hexagon are said to be in “ring 1”. This idea continues for rings 2, 3, and so on. As all hexagons are identical, we use “rings” as a measure of distance. A ring distance is equal to the hexagon’s width, i.e. 1 ring is approximately 1 kilometer.

Given that suburban areas are not dense and feature a very small fraction of our rides, we aggregate hexagons in the suburbs to reduce the dimensionality of our state space. This greatly reduces the computational burden in the model estimation. To do this, we find the weighted centroid of ride locations (start and end coordinates) within each surge area. We then map this centroid to a hexagon and take that to be the representative hexagon of the surge area. For the areas in the central part of the city, we do not do this aggregation as these are high activity areas in our data. In the end, we are left with 83 locations

in the city as shown in Figure (B3). 54 of these are in the central part and are not aggregated; these are plotted in yellow in the figure. 86% of rides start or end in these hexagons, so this level of aggregation does not create too much data loss. This aggregation has an implication for distances in the model. We can partition the city into the central areas (\mathcal{J}_C) and non-central areas (\mathcal{J}_N), i.e. $\mathcal{J} = \mathcal{J}_C \cup \mathcal{J}_N$. To calculate distances, we set $\delta_{ij} = \text{Ring}(i, j)$ if $i \neq j$ or $i = j \in \mathcal{J}_C$. For $i = j \in \mathcal{J}_N$, we set δ_{ii} as the average ring distance between all the hexagons that i is a representative for.

Figure B3: Model Map




Notes: This represents the map used in the model, where each hexagon is a location. The yellow central hexagons are unaggregated. The blue outer hexagons are representative of a broader area. In (a), the black borders capture the area which the representative hexagon covers. In (b), a line connecting two hexagons represents a route (i.e. connection) between the two hexagons.

In the model, we restrict which hexagons can be accessed from each hexagon via the correspondence $A(i)$. This is graphically represented in Figure (B3). This can be thought of as capturing main streets and highways. We create these routes by assigning each hexagon access to all routes within ring 2 (areas it can access within one time period). For hexagons that have fewer than 6 connections from this (i.e. the areas in the periphery), we assign them access to the three nearest outer hexagons (in blue) and the three nearest central hexagons (in gold). All access routes are symmetrical (i.e. if area i can access j , then j can also access i).

Finally, our estimation relies on moments capturing average entry and exit. We divide up our neighborhoods into three regions, as shown in Figure (B4). The region numbers capture the centrality of the area, where Region 1 is the central downtown neighborhood.


Figure B4: Regions



C Additional Summary Statistics

In Figure (C5), we plot the distribution of the average number of rides per day for a driver. We calculate this in two ways. The first method (driver-day count) treats each driver in each day as a separate unit of observation. Here, we see that distribution is heavily right-skewed, with a mode of only 1 ride. The second method (within-driver average) plots the average number of trips that a driver makes within a day across all the days that we observe said driver having a ride. In contrast to the first approach, this distribution is less skewed. This difference has an important – and intuitive – interpretation. On average, a driver may expect to have 3-6 rides per day; this translates to working for about an hour. However, on an individual day, a driver’s actual rides could vary widely from days with only one ride to days with over 15 rides. This is what one would expect given that ridesharing is part of the ‘gig economy’ where drivers supplement their income by driving for short shifts in the day, but that this income source is quite volatile. This also implies that entry and exit decisions of drivers are crucial in any model of the ridesharing market, in sharp contrast to the traditional taxi market as modeled by [Buchholz \(2020\)](#).

Figure C5: Distribution of Rides per Day



Notes: Let N_{it} be the number of rides taken by driver i on day t . Driver-day plots the non-zero distribution of $\{N_{it}\}_{i,t}$. Within-driver average plots the distribution of $\left\{ \frac{1}{\sum_t \mathbb{1}\{N_{it}>0\}} \sum_t N_{it} \right\}_i$.


Table C8: Probability of Exit

Start Region	End Region		
	Region 1	Region 2	Region 3
Region 1	14.2	16.5	24.3
Region 2	14.1	15.9	22.8
Region 3	13.2	14.9	19.5

$N = 1,053,069$

Notes: Each cell represents the proportion of trips that that start in Region i (row) and end in Region j (column) that are driver's last ride. Last ride is defined as a ride where the driver has no subsequent trips for the next 60 minutes. See Figure (B4) for regions.

Figure C7: Wait Time By Pick-Up Location



Notes: This plot shows the mean waiting time by neighborhood for riders requesting trips.

Figure C6: Distribution of Trips (by End Location)






Figure C8: Exit Rates by Neighborhood



Notes: For each neighborhood, the value represents the proportion of trips that end in the neighborhood which are a driver's last ride. Last ride is defined as a ride where the driver has no subsequent trips for the next 60 minutes.

Figure C9: Active Drivers Within Day



Notes: This plots the average number of active drivers (searching or in a trip) by hour of the day for the full sample (All Days) and the model sample (Fri.-Sat.).

Figure C10: Rider Wait Time Distribution




Figure C11: Time Searching Between Rides





Figure C12: Shift Length by Starting Hour



Notes: This plots the average shift length by the hour in which the shift starts. This is done for both the full sample (All Days) and the model sample (Fri.-Sat.). A shift begins when a driver is dispatched for the first time after 60 minutes of having no trips in a given day. A shift ends when the driver completes a trip with no subsequent trips in the next 60 minutes within that same day.

Figure C13: Surge versus Demand/Supply Ratio




Figure C14: Recommended versus Actual Surge





Figure C15: Distribution of Pull Distances




D Model Details

D.1 Matching Algorithm

We want to show that our matching algorithm is empirically justified, in particular, that RideAustin aims to match riders to the nearest available driver. Figure (D16) shows that as the time between successive rides in the same location is small, then riders are matched to drivers who are on average further away. The reasoning behind this pattern is simple. Consider two riders who request rides from the same location within minutes of each other. Then, once the first rider has been matched to a driver, the supply of drivers is temporarily reduced in that location and its neighboring areas. So one would expect that the

second rider is matched to a driver who is further away than the first driver if RideAustin matches riders to the closest riders. This is precisely the pattern we see in Figure (D16), where we plot the difference in time between successive rides on the x-axis against the difference in distance to the riders for the two drivers. As the time between successive ride requests gets smaller, then the difference in distance to the riders gets larger. On the other hand, as the time between successive rides increases, the market gets more time to adjust and thus the drivers are approximately the same distance away from the riders.

Figure D16: First Dispatch Protocol Observed in the Data



Notes: The x-axis captures the difference in dispatch time for two successive rides in the same location. The y-axis is the average difference in distance the drivers are from their rider. The x-axis values are binned at 1 minute intervals.

D.2 Surge Pricing Algorithm

We use the elements of the data to motivate a two-step approach. In the data, we observe for every area a and 5-minute interval t , the number of people requesting trips (demand), the number of vacant drivers (supply), a recommended surge factor, and an actual surge factor.⁴⁵ In the first step, the algorithm computes a deterministic recommended surge factor \tilde{s}_a^t for every area a and 5-minute interval t based on local market conditions. In the second step, it converts the recommended surge \tilde{s}_a^t into the actual surge factor s_a^t (the value that agents observe) using a probabilistic approach. We estimate both of these steps based on the empirical patterns seen in the data.

To determine the recommended surge, we fit a function to predict the recommended surge based on demand and supply across locations and time periods. In Figure (C13), we see that there is a strong relationship between surge and the ratio of the demand and supply of rides, thus providing support for

⁴⁵This data is not observable in our trips dataset. However, as mentioned in 2.2, we have access to a separate dataset which has this information.

this approach. We estimate the following equation:

$$\tilde{s}_{ad}^t = \gamma_1 \text{Demand}_{ad}^t + \gamma_2 \text{Supply}_{ad}^t + \gamma_3 s_{ad}^{t-1} \times \mathbb{1} \left\{ \tilde{s}_{ad}^{t-1} > 1 \right\} + \alpha_a + \lambda^t + \varepsilon_{ad}^t \quad (11)$$

In the above equation, a indicates a surge area (rather than a hexagon), t is a 5-minute interval, and d is a day in our sample. $\mathbb{1} \{ \tilde{s}_{ad,t-1} > 1 \}$ is an indicator which turns on if the recommended surge in the previous period was bigger than 1. This captures the persistence of surge in a location once it is activated. We use the estimated parameters from Equation (11), to calculate the recommended surge \tilde{s}_a^t in the first step of our algorithm.

For the second step, we convert the recommended surge into actual surge using a probabilistic approach.⁴⁶ Each recommended surge factor is converted to an actual surge factor with some probability. We determine these probabilities using the empirical distribution in the data. Denote F as the set of actual surge factors which we observe in our data. Given a fixed recommended surge \tilde{s} , we can compute from our data the probability of actual surge being $s \in F$. The model converts \tilde{s}_a^t to s_a^t using draws according to the estimated probabilities.

D.3 State Transition

Here, we specify the state transition process as a function of the key variables. We consider the transition from $\mathcal{S}^t = \{v_i^t, r_i^t\}_{i \in \mathcal{J}}$ to $\mathcal{S}^{t+1} = \{v_i^{t+1}, r_i^{t+1}\}_{i \in \mathcal{J}}$.

The total number of drivers available at location i at time $t + 1$ depends on:

- (i) the number of drivers who reach location i at time $t + 1$ after choosing to search there
- (ii) the number of drivers who drop off passengers there
- (iii) the number of drivers who choose to enter in that location

Vacant drivers choose to search or exit according to the policy function σ_i^t (Equation 7) and potential drivers choose to enter with the policy functions q_i^t (Equation 9). Drivers drop off passengers according to the probabilities denoted by M_{ij}^t (Equation 4). Call the history of spatial distribution, destination probabilities, and surge as $H^t = \{S^z, M^z, s^z\}_{z=1}^t$. The state transition kernel of driver location is therefore defined as:

$$Q \left(v^{t+1} | H^t \right) = \sum_{z=1}^t \left(\underbrace{Q_S \left(v_S^{t+1} | S^z, \sigma^z \right)}_{\text{Drivers arriving from search}} + \underbrace{Q_T \left(v_{IT}^{t+1} | H^z \right)}_{\text{Drivers completing trips}} \right) + \underbrace{Q_E \left(\bar{v}^{t+1} | S^t, q^t \right)}_{\text{New drivers entering}} \quad (12)$$

where v is the vector of total vacant drivers, v_S is the vector of searching drivers, v_{IT} is the vector of in-transit drivers who are engaged in trip, and \bar{v} is the vector of entering drivers. The transition kernel for searching drivers, matched drivers finishing trips, and entering drivers are denoted by Q_S , Q_T , and Q_E ,

⁴⁶Figure (C14) in the Appendix shows that on average the actual surge is lower than the recommended surge, though there can be large variation.

respectively. Note that the transitions rely on the history of states and policy functions, rather than just those from t , as drivers can take multiple periods to arrive from their search or assigned trip (however, new drivers simply enter based on the current period's conditions).

The transition kernel for demand is $Q_R(r^{t+1} | H^t, \omega)$, where next period's demand is a function of the current period's surge and the exogenous drop-out probability ω . Together, the transition kernel of driver location Q and the transition kernel of demand Q_R determine the state transition of \mathcal{S}^t to \mathcal{S}^{t+1} .

D.4 Equilibrium

The equilibrium in our setting is determined by the driver's beliefs about the entire distribution of riders and drivers across space as well as the beliefs over the preferences of consumers to be dropped off in different locations. Denote this belief of a driver in location i as \tilde{Q}_i^t . Given that the drivers' optimal policies only depend on the current state and their beliefs about the evolution of the state, we can define the following Markovian equilibrium in our setting similar to [Buchholz \(2020\)](#).

Definition. Equilibrium is a sequence of state vectors $\{\mathcal{S}^t\}$, transition beliefs $\{\tilde{Q}_i^t\}$, and policy functions $\{\sigma_i^t, q_i^t\}$ over each location $i \in \mathcal{J}$ and an initial state $\{\mathcal{S}_i^0\}$ such that:

1. In each location $i \in \{1, 2, \dots, L\}$, at the end of each period, unmatched drivers search or exit according to their policy functions $\{\sigma_i^t\}$, which solves Equation (5), and derive expectations under beliefs \tilde{Q}_i^t . This determines the transition kernel of unmatched drivers given by $Q_S(v_S^{t+1} | \mathcal{S}^z, \sigma^z)$, $\forall z \in [0, 1, \dots, t]$.
2. In each location $i \in \{1, 2, \dots, L\}$, at the start of each period, potential drivers enter according to their policy functions $\{q_i^t\}$, which solves Equation (8). This determines the transition kernel of entering drivers given by $Q_E(\bar{v}^{t+1} | \mathcal{S}^t, q^t)$.
3. In each location $i \in \{1, 2, \dots, L\}$, at the start of each period, in transit drivers arrive at drop-off locations following the transition matrix M^t . This determines the transition kernel for matched drivers $Q_T(v_{IT}^{t+1} | H^t)$, $\forall z \in [0, 1, \dots, t]$.
4. The state transitions are given by the combination of matched, unmatched, and entering drivers which is given by Equation (12).
5. Driver's expectations are rational which implies that the transition beliefs are self-fulfilling given the optimizing behavior i.e. $\tilde{Q}_i^t = Q_i^t$ for all i and t .

Algorithm 1 Algorithm Describing the Equilibrium Computing Algorithm

Input the demand and empirical consumer match probabilities M^t

Fix parameter σ_ε and guess initial S_0^T and compute $V^T(S_0^T)$

while diff > tol **do**

 Set counter $k = 0$

for $t = T - 1$ to 1 **do**

 Guess S_0^t and compute $V^t(S_0^t)$

 Derive choice specific value functions $W_i^t(j, S^t)$ for all i, j, t

for $t = 1$ to $T - 1$ **do**

 Derive policy functions σ from choice specific value

σ , M , and q imply transition from S^t to S^{t+1}

 Update S_k^t and V_k^t to S_{k+1}^t and V_{k+1}^t

$k = k + 1$

end for

end for

 Compute diff = $\max |V_{k+1} - V_k|$

end while

Figure D17: 12 am. to 3 a.m. Value Functions





Figure D18: 12 am. to 3 a.m. Convergence



Notes: This graph plots the mean absolute difference between successive supply matrices over 100 iterations of the equilibrium algorithm that we use. On average there is a change of 2 drivers in a single location across successive iterations. Given that most suburban areas rarely have any drivers and that drivers are concentrated in central Austin, this implies that a change of less than 2 drivers is unlikely to change equilibrium driver behavior.

E Demand Elasticities

Table (E9) shows how the price and waiting time elasticity vary through the day. The first two columns present results from 3 p.m. to 9 p.m. while the last two columns present results from 9 p.m. to 3 a.m. The price elasticity of demand is much higher during the day/evening than late at night while the waiting time elasticity is higher at night.

Table E9: IV Demand Results

	1st Stage (3 pm - 9 pm)	2nd Stage (3 pm - 9 pm)	1st Stage (9 pm - 3 am)	2nd Stage (9 pm - 3 pm)
	Surge (Log)	(Log) Rides Requested	Surge (Log)	(Log) Rides Requested
Surge (Log)		-2.16** (0.98)		-0.18 (0.20)
# Completed Rides (Log)	-0.0050*** (0.00)		-0.022*** (0.00)	
Wait time (Log)	0.0094*** (0.00)	-0.015 (0.01)	0.018*** (0.00)	-0.034*** (0.01)
# People Opening App (Log)	0.011*** (0.00)	0.022** (0.01)	0.017*** (0.00)	0.0027 (0.00)
Area \times Time Interval FE	Yes	Yes	Yes	Yes
<i>F</i> - stat	22.90	.	270.04	.
Observations	28139	28139	21896	21896

This table presents the results from using both number of rides ending in location i at time t as well as the lagged number of ending rides to instrument for price and waiting time.

Table E10: IV Demand Results with Two IVs

	First Stage		Second Stage
	Surge (Log)	Wait time (Log)	(Log) Rides Requested
# Completed Rides (Log)	-0.0084*** (0.00)	0.098*** (0.00)	
Lagged Completed Rides (Log)	-0.0032*** (0.00)	-0.064*** (0.00)	
# People Opening App (Log)	0.010*** (0.00)	-0.0078*** (0.00)	0.021*** (0.00)
Surge (Log)			-2.44*** (0.38)
Wait time (Log)			-0.23*** (0.06)
Area \times Time Interval FE	Yes	Yes	Yes
F - stat	309.14	259.76	.
Observations	134368	134503	80387

F Model Fit and Counterfactuals

F.1 Model Fit

Table F11: 3 p.m. to 6 p.m. Shift - Non-targeted Moments

	Data	Benchmark	No Surge	No Match	Gig Taxi
Average Earnings	31.51	27.63	25.76	21.67	19.16
Surged Rides (%)	18.96	17.08	0.00	18.15	0.00
Unmet Demand (%)	–	4.44	5.77	29.29	30.79
Last Rides (%)	40.16	46.59	45.17	55.46	56.19
Average Search Distance	–	16.61	16.42	19.08	20.07
Number of Trips	1,197	1,247	1,242	898	897
$P(\text{Search in Region 1})$	–	0.39	0.40	0.37	0.36
$P(\text{Search in Region 2})$	–	0.57	0.56	0.59	0.60
$P(\text{Search in Region 3})$	–	0.04	0.04	0.04	0.04


Notes: Average Earnings is the average revenue earned by a driver. Surged Rides is the proportion rides which are surged. Unmet Demand is the proportion of requested rides which were unmet. Last Rides is the proportion of total rides which were last rides for a driver. Average Search Distance is the ring distance between the dropoff location for trip k and dispatch location of trip $k + 1$. $P(\text{Search in Region } i)$ refers to the proportion of total searches in Region i

Table F12: Non Targeted Moments

	<i>Afternoon</i>		<i>Evening</i>		<i>Night</i>		<i>Late Night</i>	
	Data	Model	Data	Model	Data	Model	Data	Model
Average Earnings	31.51	27.63	34.97	25.50	42.62	26.52	54.56	39.83
Surged Rides (%)	18.96	17.08	21.44	19.80	22.23	20.45	46.88	48.10
Last Rides (%)	40.16	46.59	32.67	47.70	27.64	46.52	26.38	36.64
Number of Trips	1,197	1,247	1,977	1,434	2,494	1,550	2075	1632

Notes: Average Earnings is the average revenue earned by a driver. Surged Rides is the proportion of rides which are surged. Last Rides is the proportion of total rides which were last rides for a driver.


Figure F19: Spatial Distribution of Proportion of Unmet Demand (3 p.m. – 6 p.m.)



Notes: This figure shows the spatial distribution of the proportion of rides which are requested go unmatched.


F.2 Additional Counterfactual Results

Figure F20: Average Pull versus End-Start Difference



Notes: This is a binned scatter plot, where each unit of observation is a location-time cell as shown in Figure (?). The x-axis represents the difference in the count of rides ending in the location-time and the count of rides starting in the location-time. We restrict to x-axis values between the 1st and 99th percentile of the distribution. The y-axis represents the average pull for rides in the corresponding bin. The line represents a LOWESS function line. The results are shown from the benchmark model.

Figure F21: Spatial Distribution of Consumer Surplus (12 a.m. – 3 a.m.)




F.3 Complementarities

To expand on the complementarity, while surge pricing induces more drivers to enter, it can only indirectly allocate drivers to where they are needed. Matching not only corrects the spatial misallocation issue, but allows the platform to rely less on surge. With a lower surge, more consumers will request rides, and the matching technology almost certainly ensures they will be picked up. However, matching pulls drivers farther, especially if there are not enough drivers to satisfy all of demand. When surge pricing is added – and more drivers enter – this will lead to lower pull distances than without it.

The complementarities between surge and matching also interact with the third feature of ridesharing: that drivers can choose to enter and exit throughout the day. Indeed, surge pricing is effective in large part because it induces drivers to enter the market. However, under a taxi system where supply is fixed (i.e. drivers cannot enter and exit), there is smaller scope for surge pricing to increase welfare. This is seen in Figure (F22) and (F24) by comparing the surged taxi counterfactual to the dotted line. This counterfactual performs worse than the taxi equilibrium irrespective of the time of day. Matching is also highly effective because it can re-allocate drivers without the risk of incentivizing drivers to exit due to pulling them to locations they do not want to be.

Figure F22: Welfare Comparison (12 a.m. – 3 a.m.)



Surge	True	False	True	True	True	False	False
Matching	True	True	False	True	False	True	False
Entry/Exit	True	True	True	False	False	False	True

Figure F23: Comparison of Driver Compensation Scheme with $\alpha = 0.92$





Figure F24: Welfare Comparison (3 p.m. – 6 p.m.)



Surge	True	False	True	True	True	False	False
Matching	True	True	False	True	False	True	False
Entry/Exit	True	True	True	False	False	False	True

Figure F25: Effects of Varying Matching and Surge




Table F13: Varying Matching and Surge Simultaneously

	Benchmark	Platform Commission	Flexible Surge
Consumer Surplus	286,105.81	266,294.52	321,905.75
Driver Profit	49,345.83	33,128.75	64,197.76
Platform Profit	5,848.00	5,588.00	10,702.00
Total Welfare	341,299.65	305,011.27	396,805.51
Total Rides	5,848	5,588	10,702
Driver Count	2,596	2,232	2,842
Average Surge	1.13	1.19	0.67
Fraction Surge	0.23	0.28	0.13
Vacant Time	6.88	7.01	4.24
Average Pull	1.91	2.03	1.55

Notes: Benchmark indicates the baseline model that was estimated. Platform commission is a counterfactual where the platform receives 25% of the fare. Flexible surge is a counterfactual where surge is allowed to go below 1.